

Economics needs to treat the economy as a complex system

J. Doyne Farmer^{1,2}

¹*Department of Mathematics, the University of Oxford,*

Institute for New Economic Thinking at the Oxford Martin School

²*Santa Fe Institute, 1399 Hyde Park road, Santa Fe, NM 87501, USA*

May 3, 2012

Abstract

The path to better understanding the economy requires treating the economy as the complex system that it really is. We need more realistic behavioral models, but even more important, we need to capture the most important components of the economy and their most important interactions, and make realistic models of institutions. The complex systems approach is intermediate between traditional economic theory and econometrics. Traditional economic theory is top-down, modeling decision making from first principles, and then testing against data later. By “first principles” I mean the requirement that theories have “economic content”, i.e. that they derive behaviors from preferences. Econometrics, in contrast, takes a bottom up, data-driven, but fundamentally *ad hoc* approach. The complex systems approach sits in the middle, taking a bottom up data-driven approach that differs from traditional econometrics by explicitly representing agents and institutions and modeling their interactions, without the requirement that everything be derived from fundamental principles. It has the potential to free behavioralism from the straightjacket of equilibrium modeling, and to bring the computer revolution fully into economics. This path will at times involve abandoning economic content in favor of economic realism.

1 Linear mathematics and non-complex economics

At the 2012 INET conference “Paradigm Lost: Rethinking Economics and Politics”, I was asked to speak in a session titled “Taking Stock of Complexity Economics: Which Problems Does It Illuminate?”. This title made me laugh because it reminded me of a lecture by the great mathematician Stan Ulam that

I heard as a physics graduate student at UC Santa Cruz in 1977. The lecture was titled “nonlinear mathematics”. Ulam began by saying that he was embarrassed by the inappropriateness of title: Since almost all of mathematics is nonlinear, he said, it is like calling an animal a “non-elephant animal”. He argued that almost all interesting mathematics, and almost all interesting physical phenomena, are inherently nonlinear. This is not just a technical distinction: In most cases the essence of physical, natural and social phenomena depends on nonlinearity. The importance of nonlinear interactions is perhaps even more important in biology and the social sciences than it is in physics, even if the underlying models are less well developed.

Just as almost all mathematics is nonlinear, almost all economic phenomena are complex, and a discussion of which problems are illuminated by complexity economics is silly. A more tractable topic would be whether there are any problems it does not illuminate, or should not illuminate. The fact that such a session needs to occur and that its title is couched in such defensive terms is symptomatic of a gaping hole in the subject. The importance of complexity to economics should be taken for granted. The fact that INET has chosen the moniker “new economic thinking” makes the title of this session particularly ironic. If there were a similar conference in a parallel universe where the field of economics is at it should be, the complex systems approach would be well represented. (But then again in that universe this conference would be supported by the ISET – the Institute of Sound Economic Thinking).

The study of complex systems is built on top of nonlinear mathematics. Complex systems is the study of how interesting emergent phenomena arise from the interactions of low-level building blocks. By definition, if the interaction is linear, the whole is just the sum of the parts, but if the interactions of the parts are nonlinear, the whole can be more than the sum of its parts. In the most interesting cases the whole is qualitatively different from the sum of its parts, in which case we say it is an *emergent phenomenon*.

Complex systems extends the study of nonlinear systems to focus on collective phenomena among simple or not-so-simple building blocks. While the underlying cause of interesting emergent phenomena is nonlinearity in the interaction of the building blocks, this is a weak statement: Nonlinear interactions can take many forms and can have many different results. The study of complex systems focuses on the synthetic properties of such interactions. The goal is to characterize emergent phenomena, and to understand which types of low level interactions lead to which kind of phenomena.

In view of its origins, the fact that economics can be singled out as a field of science where the complex systems revolution has had the least impact is an odd twist of history. Within the study of complex systems Adam Smith is widely regarded as the first to clearly articulate the concept of an emergent phenomenon. In his studies of the distribution of income, Pareto was the first to describe a power law distribution, a concept that has played an important role in

the study of nonequilibrium phenomena and complex systems in particular. In my view, the failure to embrace complex systems thinking is probably the main reason that economics has not experienced the rapid progress seen in so many other sciences during the last fifty years. This is closely connected to the lack of serious computer simulation, which goes hand-in-hand with the complex systems approach.

In this essay I intend to be provocative. I will give my views on the underlying reasons why the complex systems and computer revolutions have so far passed economics by, to motivate why it is urgent to change this. I do not have space here to present a comprehensive review of what has already accomplished in complex economics, nor to thoroughly discuss how the alternatives should be implemented, but rather will just provide references.

2 What is the complex systems approach?

Complex systems, and its close cousin and near synonym “complexity”, are terms that refer to a movement in science that has blossomed and begun to bear fruit during the last thirty years or so. It is focused on the question of how emergent phenomena emerge from the interaction of low level building blocks. In a certain sense one might argue that this is nothing new – in a narrow sense physics and other disciplines have been doing this for a long time. One might argue that in some narrow sense of the word, the movement of the planets is an emergent phenomenon that comes about from the very simple low-level postulates of Newton’s laws of gravitation and mechanics. A better example would be Adam Smith’s invisible hand: How do the collective wealth and shared benefits of society come about from the striving of selfish individuals? Unlike the solar system, we are still far from understanding the mysteries of the invisible hand.

The debate here is not whether this is an interesting thing to understand – we can all agree that it is – but how we should go about understanding it. An advocate of the complex systems approach would argue that we will never get there with simple models in which we have a few representative agents. Experience from the study of complex systems in other fields of science suggests that to understand where the richness of the economy comes from, we need to instead study the collective interaction of a large number of heterogeneous participants.

Complex systems seeks to develop concepts, methods and tools that transcend specific applications and disciplines. It also seeks to apply this approach to problems that are well beyond the traditional scope of physics, such as adaptive systems, like those commonly encountered in biology and social science. While of course manifestations of complex systems are different in every field, the commonalities can be striking. It is possible, for example, to show that models of adaptive systems including neural networks, classifier systems, autocatalytic networks, immune networks, and evolutionary game theory can all be mapped into

a common mathematical framework [1]. For a general review of the literature in complex systems see [2, 3, 4, 5, 6, 7]

Complex systems has antecedents that have gone under a variety of names, including cybernetics, self-organization, and synergetics. One of the visionaries of complex systems, John von Neumann, correctly predicted the future in 1950 when he said:

Science and technology will shift from a past emphasis on motion, force, and energy to communication, organization, programming, and control.

To this I would add form, function, structure, information, computation, emergence and evolution. The change in language that von Neumann anticipated reflects both a change in focus of the problems that are considered, and a change in the way of thinking about them and the methods that are brought to bear to understand them. Complex systems is a highly interdisciplinary field that has incorporated ideas and methods from statistical mechanics, dynamical systems, game theory, evolutionary biology, ecology, machine learning, information theory, theory of computation, and graph theory. The knowledge involved has been imported from many different disciplines.

It is of course very ambitious to try to find general principles spanning so many fields. This program is far from being realized, and many question whether it can ever be fully realized, but there are some modest successes. These include network theory and agent-based modeling, which are both highly relevant to economics and will be discussed here.

The complex systems approach is strongly data driven. One begins by studying a system to identify its components and their interactions. Then one typically constructs either a simulation or a theoretical model that takes these interactions into account, and in particular focuses on understanding the emergent collective behavior that they generate. For an economy, for example, one would naturally identify the key institutions, such as households, firms, banks, and financial markets and construct rules for their interactions. Because the economy is a rather complicated complex system, a natural starting point is a computer simulation, such as an agent-based model. Of course, because the key decision makers are human beings, the basic building blocks of such a model are extremely complicated, and this is a challenging task. Nonetheless, the fact that this approach has never been tried in economics is striking [8]. See [9] for a vision of how this approach can be used in economics.

Even if simulation models are a cornerstone of the complex systems approach, they are by no means its only component. Analytic models can also be very useful. But under the complex systems approach analytic analysis almost always goes hand in hand with simulation.

3 Economic content vs. economic realism

The failure for complex systems methods to be widely used in economics is closely associated with the evolution within the economics profession of the ironclad requirement that theories must have “economic content” to be considered valuable. This phrase does not mean what one would naively think, that a theory must explain an economic phenomenon. Instead, a theory is said to have “economic content” if its starting point is the assumption that selfish individuals maximize their preferences. If the theory doesn’t start there, it has no economic content, and in fact, it doesn’t even deserve the word “theory”.

Constructing theories with economic content is a laudable goal. There are many situations where this approach is valuable, and there are good reasons why this requirement has evolved, the most famous being the Lucas critique [10]. Many important problems have economic content, there are many situations in which maximizing preferences plays an important role, and this is in no way incompatible with complex systems theory. As I already said, the invisible hand of Adam Smith is the first clear articulation of an emergent phenomenon. Nonetheless, making this an absolute requirement needlessly restricts what can be accomplished and slows down progress in economics. It has choked off other approaches and become a straight jacket that limits the usefulness of economics.

There are three problems with the strict requirement that theories have “economic content”:

1. Even though many problems may ultimately depend on some version of selfishly maximizing preferences, the requirement that this be the starting point for all theories is excessively stringent. Many phenomena depend on factors that cannot be addressed using the heavy mathematical baggage that “economic content” requires.
2. There are many important economic phenomena that either have no economic content or that require minimal behavioral assumptions to explain.
3. There are many circumstances in which preferences are difficult to characterize, are intrinsically plastic and context dependent, or are dictated by other requirements, so that this is simply the wrong approach.
4. Such an a priori restriction diminishes the diversity of effort and impoverishes the resulting science.

The first problem comes from the fact that science is a resource-limited enterprise. One cannot do everything well; it is necessary to allocate effort efficiently and focus on a few things at a time. The derivation of economic equilibria is a mathematically challenging activity; as one adds structure to a problem the complexity of the mathematical machinery increases, and there is inevitably a

critical point where a problem becomes intractable. As a result, any economic theory that has economic content has to have a simple framework. One must keep the number of variables to a minimum and it is only possible to treat a few interactions at a time. This often means that important aspects of a problem must be neglected. All too often the baby goes out with the bathwater [11, 8].

The mandate that all theories have economic content is like requiring that all theories in physics start at the level of quarks. In fact, physicists often construct *phenomenological theories*, that simply connect different phenomena. This is sometimes done using approximations that can be justified based on more fundamental theories, but this is not always the case; sometimes physicists begin with an *ansatz*, i.e. a seemingly arbitrary assumption, and showing that such an assumption leads to interesting conclusions. (A good example is Newton's postulate that gravity obeys a force that decreases as the square of the radius, which was severely criticized by the Cartesians, and still lacks a truly fundamental explanation). The main agenda of science is not to derive everything from first principles, but rather to relate diverse phenomena to each other and thereby simplify our description of the world.

The second problem can be viewed as failure in cost-benefit analysis. One of the cardinal rules of science is to pick low hanging fruit first. If a problem doesn't require a given assumption, why use it? Examples of problems in which it is possible to learn a lot while assuming little or no economic content include the size distribution of mutual funds [12], variations in the growth of firms [13], the relationship between order flow and the bid-ask spread in the continuous double auction [14], and indeed one of the most important problems in economics, namely technological progress and economic growth [15, ?, 16, 17, 18]. Since papers that do not have "economic content" are very difficult to publish in top economics journals, there is a great deal of low hanging fruit waiting to be picked for those adventurous enough to do so.

The third problem is that in most cases utility does not provide a good description of preferences. Prospect theory might give a small improvement, but it is still wide of the mark. Studies of the factors that influence subjective well-being show that the real determinants of happiness are much more complicated than the standard utility function. But there is an even deeper problem: In an evolutionary setting preferences are not arbitrary – they are dictated by the imperative of survival. From the theory of gambling it is well known that only individuals with log utility survive in the long run [19, 20, 21]. This matches up with evolutionary principles – the bottom line is survival, not utility. In these circumstances utility is no longer fundamental since it is not a matter of choice, and maximizing preferences is misleading, as it gives the impression that preferences are arbitrary.

4 How does the complex systems approach differ from traditional economics?

We can now return to ask how the complex systems approach differs from the approach used in mainstream economics. The skeptic might argue that the approach outlined in Section 3 is precisely that taken by economists. The key difference is in where the effort is focused, how the components are identified, which kind of interactions are considered, the method used to understand their effect, and what is considered an interesting result.

As currently practiced, economics is a field polarized between two extremes: Economic theory sits on one hand and econometrics sits on the other. The basic structure of an economic theory is always the same: One postulates preferences and beliefs and then derives an equilibrium. This top down approach is justified by the argument that economic phenomena are noisy and complicated, and unless one imposes strong priors, it is easy to get lost [22].

Econometrics, in contrast, is a bottom-up, strongly data-driven activity. Econometricians use functional forms that are typically fairly arbitrary and use statistical methods to estimate their parameters from data. This can be extremely useful; for example, it is the basis for the principle models used to make economic forecasts. But it is intrinsically limited in the understanding it can provide.

At least in a certain sense, a typical complex systems model can be viewed as sitting in between these two poles. Complex systems models are not constrained to have economic content (though there is no prohibition against it, and in recent years some economists have begun to consider properties such as network interactions that are very much in the spirit of complex systems [23]). The difference is that the focus of a complex systems model is on representing the interactions among the elements, which in an economic model involves agents (individuals) and institutions (firms, regulatory agencies, etc.).

A good complex systems model both begins and ends with data: Low level data is used to formulate the assumptions about the building blocks of the model, and both high and low level data is also used to test whether the resulting emergent phenomena properly correspond to those observed in the real world. Economic content is often sacrificed in favor of economic realism.

5 Economic phenomena are often far from equilibrium

Current economic theory is almost entirely based on the notion of equilibrium. A typical model postulates utility functions for the relevant economic agents, makes assumptions about their strategies for maximizing it, and then computes equilibria. Under the right circumstances, i.e. when all the underlying assumptions

apply, this approach can be very useful. However, in many situations there is no unique equilibrium. When there are multiple equilibria it may be difficult to predict which agents will converge to; in other circumstances they may fail to converge to any equilibrium at all.

How frequent are these “bad” situations where convergence to equilibrium fails? Is there any way we can know in advance when such problems are likely to occur? Tobias Galla and I have investigated questions within the context of game theory. We do this by making up games at random. That is, for each possible combination of moves by the players we assign them random payoffs. We then assume the players try to learn better strategies for making their moves using a common approach called reinforcement learning, which has been shown in economic experiments to provide a good characterization of how real people learn in games. The games we study are “complicated” in the sense that each player has many possible actions. Think of the stock market, where a player has a choice of buying any of thousands of possible stocks. For the details please see [24].

For simplicity we began by studying two player games. We find that there are two key parameters that characterize when the system will converge to an equilibrium. The *competition parameter* Γ is the correlation between the payoffs of the players: If $\Gamma = -1$ the game is zero sum, meaning that if one player wins the other loses. Similarly if $\Gamma = 0$ the payoffs of the two players are uncorrelated, and if $\Gamma > 0$ the payoffs are positively correlated, i.e. the payoffs are likely to be either win-win or lose-lose, depending on the players’ moves. Thus the smaller Γ , the stronger the competition. The *timescale parameter* α controls how much attention the players pay to the distant past. If $\alpha = 0$ then actions in the distant past are weighted the same as actions in the recent past, and if $\alpha > 0$ more recent actions are given higher weight, so the bigger α the faster the player forgets the past.

One might think that, since the payoffs are made up at random, the results would have no particular pattern. This is not the case at all. For a given choice of α and Γ we see fairly consistent behavior as shown in in Figure 1: The regimes are:

1. The strategies converge to a unique equilibrium. This occurs when α is large and Γ is small, i.e. when the strategies weight the future more strongly than the past and when the games are competitive.
2. The strategies converge to one of many possible equilibria. The number of equilibria is often very large, e.g. many hundreds, beyond our ability to find all of them. This occurs when α is small and Γ is positive. i.e. when players do not forget moves from the distant past and the payoffs of the players are aligned.
3. The strategies usually fail to converge, instead wandering around on a limit




Figure 1: Schematic illustration of the regimes of behavior for complicated games when the players use reinforcement learning. α gives the timescale for the importance that players give to past moves: $\alpha = 0$ implies the players weight the importance of all past moves equally, whereas a large value of α implies that moves in the recent past have more influence in formulating the strategy. Γ is a competition parameter: $\Gamma = -1$ implies a zero sum game, $\Gamma > 0$ a positive sum game.

cycle or chaotic attractor. The resulting chaotic attractors can be of very high dimension, indicating that for all practical purposes the behavior is effectively random. This occurs when α is small and Γ is negative, i.e. when the players remember the past and the games are competitive.

We also find the game dynamics have many interesting and unexpected properties. One is that when the strategies fail to converge to an equilibrium the games exhibit “fads”, in which the two players spend extended periods of time favoring certain moves over others, and then suddenly switch to favoring a different set of moves. This comes about through an interaction between the moves the players make, which determine the outcomes, which in turn affect the moves – what Soros has termed market reflexivity [25]. This gives rise to fluctuations in payoffs that look like the “clustered volatility” of financial markets: There are periods in which the payoffs change rapidly, and other periods in which they are

relatively constant, as shown in Figure 2.




Figure 2: The change in total payoff to the two players is plotted as a function of time for a randomly selected game that displays high dimensional chaos in the learning dynamics of the two players. There are periods in which the payoffs change very little and other in which they change rapidly, reminiscent of clustered volatility in financial markets.

The lesson is that whether or not the usual equilibrium approach is useful depends on the situation. Situations where competition is strong and learning is myopic are well-described in terms of a unique equilibrium. If the situation is not competitive, in the sense that one player's gain makes another player's gain more likely, and if learning has long memory, then there are an enormous number of possible equilibria and it becomes difficult to know which one the system will coordinate on. Finally, in situations where competition is strong and learning has long-memory, the dynamics are complicated and random, and equilibria are irrelevant.

Together with James Sanders we have now extended this analysis to games with many players [26]. In this situation we find that the chaotic region grows larger and larger as more players are added. In other words, for competitive games like the stock market where there are many players who can make many possible moves, chaos is likely (and equilibrium is not).

If the economic setting falls into regimes 2 or 3 then the equilibrium theories that most of economics is based on are out the window. For complicated competitive games, such as financial markets, we predict that high dimensional behavior is to be expected. One must then use alternative approaches, such as agent-based modeling.

6 Bringing the computer revolution to economics

During the last 50 years the computer has caused a revolution in the physical and natural sciences. Computers have affected science in many ways, including data gathering, data processing, and the ability to perform regressions, not to mention email and word processing. But perhaps the most profound impact has been the ability to simulate complex systems. Until the advent of the computer nonlinear systems could not be solved, and the only way to make simulations was by constructing crude mechanical analogs, such as Irving Fisher's famous hydraulic model of the economy. The impact on science has been profound.

There is a famous joke about the mathematician searching for his glasses under the lamp post, who complains that although he is sure he dropped his glasses somewhere else, he is forced to search under the lamp post because it is the only place that he can see. Similarly, until computers became widely available almost all mathematical science had to be based on linearized equations, and because of this was restricted to studying simplified non-complex problems in a reductionistic mode. Computers have changed all this, providing the flashlight that lets the mathematician find his keys.

Computer simulations are now essential to model a broad range of phenomena, including weather, traffic, epidemics, fluid turbulence, general relativity, earthquakes, and neural systems. Simulation is essential because it makes it possible to use reductionism to study complexity. Reductionism is a program in science that seeks to understand basic interactions at a low level, typically by studying the parts of a system in isolation. Complexity is the study of phenomena that emerge from the interaction of low level building blocks. Using the tools of simulation, complexity works hand-in-hand with reductionism. When the interaction rules for the low level building blocks are well understood, as they sometimes are in physical science, the task of complex systems is relatively (but only relatively!) easy; one constructs simulations, reproduces emergent phenomena, and then seeks to quantify their regularities and discover more elegant mathematics to characterize them.

In the more common case that the interaction rules are not well understood, or the proper characterization of lower level building blocks is not clear, the problem is much harder. This is almost always the case in biology and social science. (In fact even for many physical problems, such as modeling the climate, the nature of the low level interactions is far from well-understood). When the low level building blocks and their interaction rules are uncertain, one is forced to work from both ends, experimenting with the low level rules and using the emergence of the correct higher level phenomena as a test for whether the low level rules are correct.

Economists use computers in many ways, but surprisingly, there is very little effort on simulation outside of the small field of agent-based modeling [8]. The reason that most economists would give for this is that, unlike physics, the

low level rules of interaction are not well understood. Since the Lucas critique economists have felt that any model of the economy needed to incorporate the ability of agents to reason about their world. This has driven a push toward models that are “micro-founded”, which in practice generally means the solution of equilibrium conditions of a rational representative agent.

Econometric, DSGE and CGE models are solved on computers. But when I say that these are not simulations I mean that they neglect most of the structural and institutional features of real economies. Econometric models are statistical time series models, that are entirely data driven, without a fundamental model; DSGE models attempt a fundamental model, but in order to accommodate the underlying equilibrium assumptions they are of necessity highly stylized. This is in contrast to the models used elsewhere in science to simulate phenomena such as weather, climate, traffic or epidemics, which attempt to capture the structural features of the underlying systems they are targeted at understanding, building in as much realism as needed.

To see the impact such a simulation model could have on economics it is useful to compare to the role of large scale simulations in other fields. In meteorology, for example, large comprehensive computer models sit at center stage. The numerical weather prediction models that the National Center for Atmospheric Research uses incorporate the current best science. A great deal of effort goes into diagnosing the failures of the models, and these failures which provide feedback to the entire field of meteorology and climate. For example, the failure of global circulation models to yield good predictions for cloud formation has driven better observation technology, new methods in artificial intelligence for identifying clouds from satellite measurements, fundamental theory about cloud formation, and the development of faster computers and better numerical methods. There is no corresponding role for the simulation of large scale models in directing the efforts of economists. The last statement needs to be qualified: The failures of DSGE and econometric models do indeed inform the activities of many theoretical macro-economists and econometricians, but as already mentioned, these are not really large scale simulations of a real economy.

7 Behavioralism and the straight jacket of equilibrium

Behavioral economics is one of the recent success stories in economics. Behavioral models used to be unacceptable to the mainstream, but over the last twenty years or so they have made their way in, as indicated by the Nobel prize to Daniel Kahneman and Vernon Smith in 2002. So far, however, behavioral economics is mainly about showing how the assumptions of rationality are violated, and has made very little progress toward building a quantitative, positive theory that can

be used to analyze policies or make conditional forecasts.

In my view, there are two principal reasons for this. The first is that most studies in behavioral economics operate at too general a level. Behavioralists have done a good job of showing how basic psychological biases cause major deviations from rational behavior. For example, people (and men in particular) are highly overconfident, they make characteristic errors in estimating probabilities, and they are prone toward seeing patterns even when none exist. While it is essential to understand these things, and why they make it clear why rational models fail in many cases, they do not tell us what we need to know to understand market phenomena such as a credit crisis. For this we need behavioral observations in very specific real world contexts. For example, how will lenders behave when they observe a previously trusted counterparty fail – is the resulting tightening of credit rational, or is it exaggerated due to emotion?

The second problem is that the whole style of modeling that has built up in economics based around rational expectations equilibrium is ill suited to incorporate behavioralism. Behavioral views are often inconsistent by their very nature, and in many cases forcing them into an equilibrium context is inappropriate.

Agent-based models have the potential to allow behavioralism and experimental economics to realize their potential, and to take center stage as the foundations for quantitative simulation models. The exercise of building an agent-based model forces one to think through the key decisions that economic agents must make. By performing experiments in controlled laboratory conditions one can understand how real agents will behave, formulate decision rules that incorporate the key aspects of the decision making, and build these into agent-based models that simulate the role of each agent in the economy. In a fairly simplified setting such methods have been used by the group of Cars Hommes to build models of speculative behavior [27].

8 Concluding remark

Bringing the complex systems approach into economics should be the central agenda of New Economic Thinking.

Acknowledgments

This work was supported by the National Science Foundation under Grant No. 0965673, the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement CRISIS-ICT-2011-288501, and by the Sloan Foundation. The opinions are purely those of the author and do not necessarily reflect the views of any of the funding institutions.

References

- [1] J. D. Farmer, “A rosetta stone for connectionism,” *Physica D*, vol. 42, pp. 153–187, 1990.
- [2] M. E. J. Newman, “The structure and function of complex networks,” *Siam Journal on Applied Mathematics*, vol. 45, no. 2, pp. 167–256, 2003.
- [3] M. Mitchell, “Complex systems: Network thinking,” *Artificial Intelligence*, vol. 170, no. 18, pp. 1194–1212, 2006.
- [4] E. Beinhocker, *The Origin of Wealth: Evolution, Complexity and the Radical Remaking of Economics*. Harvard Business School Press.
- [5] *Complexity: A Guided Tour*. Oxford University Press, 2009.
- [6] M. E. J. Newman, “Complex systems: A survey,” *American Journal of Physics*, vol. 79, pp. 800–810, 2011.
- [7] A. Kirman, *Complex Economics: Individual and Collective Rationality*. Routledge, 2010.
- [8] J. D. Farmer and D. Foley, “The economy needs agent-based modelling,” *Nature*, vol. 460, pp. 685–686, 2009.
- [9] J. D. Farmer, M. Gallegati, C. Hommes, A. Kirman, P. Omerod, S. Cincotti, A. Sanchez, and D. Helbing, “A complex systems approach to constructing better models for managing financial markets and the economy,” tech. rep., FuturICT, 2012.
- [10] R. E. Lucas, “Econometric policy evaluation: A critique,” *Carnegie Rochester Conference Series on Public Policy*, pp. 19–46, 1976.
- [11] J. D. Farmer and J. Geanakoplos, “The virtues and vices of equilibrium and the future of financial economics,” *Complexity*, vol. 14, pp. 11–38, 2009. to appear in *Complexity*.
- [12] Y. Schwarzkopf and J. D. Farmer, “Time evolution of the mutual fund size distribution,” tech. rep., Santa Fe Institute working paper, 2008.
- [13] Y. Schwarzkopf and J. Farmer, “The cause of universality in growth fluctuations,” tech. rep., Santa Fe Institute working paper, 2010.
- [14] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy, “Statistical theory of the continuous double auction,” *Quantitative Finance*, vol. 3, no. 6, pp. 481–514, 2003.

- [15] W. Arthur, *The Nature of Technology: What it is and how it evolves*. Simon and Shuster, 2009.
- [16] J. McNerney, J. D. Farmer, S. Redner, and J. Trancik, “Role of design complexity in technology improvement.,” *Proceedings of the National Academy of Science*, vol. 108, no. 22, pp. 9008–9013., 2011.
- [17] C. Hidalgo, B. Klinger, A. Barabasi, and R. Hausmann, “The product space conditions the development of nations,” *Science*, vol. 27, no. 2837, pp. 482–487, 2007.
- [18] C. Hidalgo and R. Hausmann, “The building blocks of economic complexity,” *Proceedings of the National Academy of Science*, vol. 106, no. 26, pp. 10570–10575, 2009.
- [19] J. L. Kelly, “A new interpretation of information rate,” *The Bell System Technical Journal*, vol. 35, no. 4, pp. 917–926, 1956.
- [20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley Series in Telecommunications, 1991.
- [21] D. Cherkashin, *Perception Game*. PhD thesis, University of Chicago, 2004.
- [22] C. A. Sims, “Macroeconomics and reality,” *Econometrica*, vol. 48, no. 1, pp. 1–48, 1980.
- [23] M. Jackson, *Social and Economics Networks*. Princeton: Princeton University Press, 2008.
- [24] T. Galla and J. Farmer, “Complex dynamics in learning complicated games,” tech. rep., Santa Fe Institute working paper, 2011.
- [25] G. Soros, *The Alchemy of Finance*. John Wiley and Sons, 1987.
- [26] J. Sanders, T. Galla, and J. Farmer, “Complex dynamics in learning multi-player games,” tech. rep., Santa Fe Institute, 2012.
- [27] C. Hommes, “The heterogeneous expectations hypothesis: Some evidence from the lab,” *Journal of Economic Dynamics and Control*, vol. 35, pp. 1–24, 2011.