
Is the Study of Complex Adaptive Systems Going to Solve the Mystery of Adam Smith's "Invisible Hand"?

————— ◆ —————

SREEKALA KOCHUGOVINDAN

AND

NICOLAAS J. VRIEND

Among the main tasks of economic theory are explaining the outcomes and understanding the fundamental mechanisms of trade in a decentralized economy. Perhaps the well-known difficulties of explaining the fluctuations in aggregate economic activity known as the business cycle, for example, merely attest that we have not yet achieved a precise understanding of how decentralized economies operate. However, recent empirical evidence pertaining to Eastern Europe suggests that decentralized trade presents deeper questions that contemporary economic theory cannot answer (see also Leijonhufvud 1993; Sargent 1993). For many of these countries some of the enigmas are related to political and institutional reforms. But in the case of the German reunification, the main change was that a large number of East German people with their preferences and endowments were added to the West German economy. Nevertheless, economists reached no consensus about the immediate German economic prospects, and forecasts, which could no longer rely on a simple extrapolation and error correction of the latest West German figures, differed widely.

Sreekala Kochugovindan is a doctoral student in economics, and **Nicolaas J. Vriend** is a lecturer in economics, both at Queen Mary and Westfield College, University of London.

The Independent Review, v.III, n.1, Summer 1998, ISSN 1086-1653, Copyright © 1998, pp. 53–66

In the former centrally planned economies, a very visible and conspicuous, although not very successful, ordering hand had produced a long period of “frozen” order. As soon as the heavy hand of the planners was removed, after glasnost had made everything transparent and perestroika had dismantled most of their institutions, chaos ensued. At the moment it seems that no hand is available, whether visible or not, and that all faith concerning the coordination of economic activities is being placed in the working of a free price system. New patterns of trade may well emerge. But many questions remain. In seeking answers, one would like the help of economic theory.

In contrast to the former centrally planned economies, traditionally decentralized economies appear to have organized themselves into relatively orderly networks of economic relations without the hand of a planner. However, the apparent order is not absolute—witness the macroeconomic fluctuations, the high volatility of prices in stock markets, and the job mobility and labor turnover. Moreover, new markets emerge and others collapse. Decentralized economies seem to be moving constantly in a region between absolute order and complete chaos.

The key problems of economic theory have remained the same since Adam Smith articulated them in 1776. Smith’s main accomplishment was to put forward as the central theme of economics the systematic analysis of the behavior of individuals pursuing their self-interest under conditions of competition. The most eloquent quotations in this respect are presumably these: “It is not from the benevolence of the butcher, the brewer, or the baker, that we expect our dinner, but from their regard to their own interest” (Smith [1776] 1976, 26–27), and “he intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention” (456). Smith viewed each individual as acting out of simple self-interest, but only by appealing to a transcendental “invisible hand”—a concept embracing the entire economy but itself standing above the level of the individual agents—was he able to conclude that the self-interested behavior of individuals results in a coordinated overall outcome.

Nowadays economists typically appeal to the metaphor of the invisible hand to explain the emergence of regularities, often loosely invoking some kind of “as if” argument when referring to the invisible hand and thereby leaving it basically a black-box concept. This contrasts with the Smithian conception. For Smith, the hand is transcendental: individuals are literally led by an invisible hand and do not behave simply as if such a guiding force exists (see also Vaughn 1989).

Some economists have claimed that General Equilibrium Theory, as formalized, for example, by Gerard Debreu (1959), “has finally proved mathematically what Smith argued two centuries ago.” For example, the First Fundamental Theorem of Welfare Economics is considered to be “a formal and very general confirmation of Adam Smith’s asserted ‘invisible hand’ property of the market” (Mas-Colell, Whinston, and Green 1995, 549).¹

As one of us has argued elsewhere (Vriend 1994), however, General Equilibrium models cannot be considered ideal representations of decentralized economies because their main contribution has been to make Smith’s transcendental hand visible by imposing various centralized concepts. These models base their analysis of a decentralized economy on the actions of autonomous agents, pretending to discard any kind of external determination of the behavior of the individual agents. But it turns out that in order to explain anything, such models must introduce many concepts and structures that transcend the level of the individual agents, much as Smith brought his invisible hand into play: the “auctioneer,” the division of time into real time and meta-time, and the rules of the game in these models—none of which arises from the behavior of the individual agents. Moreover, the rules and structure of the model predetermine the set of possible actions of the individual agents, and each individual takes the structure of the model into account in calculating his choices, *anticipating* the equilibrium character of the overall outcome. Instead, the outcome, equilibrium or not, should be explained by their actions. Unless he is God, the man with the invisible hand, or the auctioneer, how can an individual agent in one of these models understand that the economy will turn out to be in equilibrium?

More than two centuries after the publication of Smith’s masterpiece, economists still do not know how, why, and when the invisible hand works in a truly decentralized system of interacting individual agents. In other words, even more basic questions underlie such questions as when should one expect stability, booms, or slumps in a decentralized economy. Most fundamentally, we would like to know how it is possible that many individual agents, each pursuing self-interest, produce order rather than chaos, and under what conditions would they do so. We shall argue that the study of *complex adaptive systems* might help us to understand the underlying interactive processes governing the behavior of decentralized economies and leading to the emergence of regularities.

The Study of Complex Adaptive Systems and Economics

Recently a framework has been developed that may help to answer the sorts of questions raised above. This is the study of complex adaptive systems. A complex system is one consisting of a large number of relatively independent parts that are interconnected and interactive. Such a system is adaptive if the parts are agents that change their actions as a result of events occurring in the process of interaction. Examples of complex systems include biological systems, immunology systems, brains, weather systems, ecologies, and societies.² A decentralized economy, consisting as it does of a

1. This theorem states that if a set of prices and a corresponding allocation of commodities constitute a competitive equilibrium then this allocation is Pareto optimal, which means that no agent can be made better off without making at least one other agent worse off.

large number of locally interconnected and interacting rational agents who are all continually pursuing advantageous opportunities, also exemplifies a complex adaptive system. Persuasive arguments that this is the case appear in the work of F. A. Hayek (1948), written long before the term *complex adaptive system* was known.³ An essential characteristic of such systems is that their global properties cannot be derived simply from an examination of the individual components. Even when each individual agent is inherently simple, the behavior of both the system as a whole and the individual agents may become complex (see Holland 1992, Langton 1989, and Kauffman 1993).

A complex system is not the same as a chaotic system. In general, complex systems tend to evolve away from the extremes of, on the one side, absolute order and, on the other side, what appears to be complete randomness. Currently the key theoretical concepts are self-organization (the formation of regularities in the patterns of interaction) and selection (through system constraints). Selection seems to act in many self-organizing systems to constantly push the system back to some boundary between order and chaos. Around this edge, these systems appear to carry out the most complex behavior and adapt most readily to changing environments.

Because the interactions between the individual agents are in general nonlinear, from a mathematical point of view such systems are often intractable. In economic applications the analytical apparatus borrowed from graph theory, statistical mechanics, and the theory of interacting particle systems diminishes the economic content of the models. It seems especially difficult to incorporate the essential fact that the interactions taking place between economic agents in reality are not determined by their given position in a grid, graph, or lattice or by some kind of anonymous matching device.

In a decentralized economy most interactions are determined by economic agents who are themselves actively pursuing the interactions most advantageous to them. Transactions do not take place in Walrasian central markets or through anonymous random matching devices. Instead, market interactions depend in a crucial way on local knowledge of the identity of some potential trading partners. A market in a decentralized economy, then, is not a central place where a certain good is exchanged; nor is it the aggregate supply and demand of a good. In general, markets emerge as the result of locally interacting individual agents' pursuit of advantageous contacts; that is, they are self-organized.

Given the limitations of the existing formal methods, analysts often must rely on computational methods. One approach is to model each individual agent and the

2. For an introduction to the topic of complexity in a broad sense, see Waldrop 1992, Lewin 1993, Casti 1994, and Gell-Mann 1994.

3. These arguments did not influence many formal economic models, in which direct interaction between the agents is usually absent. See Kirman 1994 for a recent survey of economic models with interacting agents. As one of the referees suggested, it would be interesting to develop the striking similarities between Hayekian theory and the complex-system approach, but doing so is beyond the scope of this paper.

interactions explicitly; in other words, to provide the agents (real robots or simulated machines) with artificial "flesh and blood." That "bottom-up" approach is also known as the artificial life, or *alife*, approach (Langton 1989). Modeling *homo oeconomicus* as a "machine" does not seem to pose any particular conceptual difficulties for economic theorists. After all, as Robert Lucas puts it, doing economics means "programming robot imitations of people" (quoted in Klammer 1984, 49). Although Lucas was speaking metaphorically, current computational capabilities suggest taking the metaphor literally and considering its usefulness for economic theory. Lucas was referring, of course, to the idea that *homo oeconomicus* is a rather mechanical representation of actual human beings. The fundamental characteristic of *homo oeconomicus* is simply that he chooses the most preferred option in his perceived opportunity set. In fact, *homo oeconomicus* is an "opportunist," always doing the best he can (see also Vriend 1996).

In this light, the modeling of the perceived opportunity sets of the individual agents becomes a central concern. During the process of interaction between the individual agents in a decentralized economy, perceived opportunities evolve, either because of a change in the underlying situation or because of a change in the agents' perception of the situation, that is, learning. A general characteristic of agents living in the complexity of a "large world" is that they do not have a true, well-specified model to work with: the agents' problem situation is ill defined (see Arthur 1992, 1994). Hence, instead of basing their actions on deductive reasoning from universal truths, they are forced to use inductive reasoning, proceeding from the actual situation they face. In the terminology of Leonard J. Savage (1954), they follow the "cross that bridge when you meet it" principle. In a large world, individual agents constantly search, learn, and adapt to their environment.

In a decentralized economy each individual's activities are "involved" in a certain way in the activities and decisions of some other agents. Hence, each agent has a different relevant "environment" for different kinds of activities, and each environment is influenced by the actions of other individual agents. In other words, while an individual agent is adapting to his environment, parts of the environment are adapting to him. In biology this phenomenon is known as coevolution. In principle, such a process might go on forever (see Sigmund 1995). Therefore, rather than analyzing whether an equilibrium exists for an economy with some given structure, in this approach one analyzes how structures and patterns emerge as regularities in the process of interaction of the individual agents. Not nineteenth-century physics but modern biology or meteorology provides the relevant metaphors for this approach to the study of decentralized economies.

Having simulated the operation of an artificial economy, one has to characterize the history of actions and outcomes by looking for the emergence of regularities in the generated data set. In particular, one searches for regularities that cannot be deduced directly from the built-in properties of the individual agents or some other

microeconomic aspect of the model—or at least cannot be deduced by any argument substantially shorter than that of producing the regularity by running the simulation itself (see Lane 1992).

The complexity approach differs in a fundamental sense from the microeconomic-foundations-of-macroeconomics approach (see Weintraub 1979, Benassy 1982). The latter is essentially a reductionist, top-down approach, based on the idea that one can come to understand a given system by analyzing more and more detailed parts of that system. The complexity approach follows an alternative route: taking simple microeconomic properties as given, one tries to understand how the macroeconomic dynamics (both short-term and long-term) of a decentralized economy can emerge from the process of interactions among many individual agents.

Examples of Economies as Complex Adaptive Systems

This approach has been followed in studies of the emergence of a medium of exchange (Marimon, McGrattan, and Sargent 1990), emergent patterns in macroeconomic dynamics (Bak and others 1993), the emergence of social structures and group behavior (Epstein and Axtell 1996), the emergence of trading rules and patterns in prices and volumes on a stock market (Arthur and others 1996), the emergence of patterns of interaction in an iterated prisoner's dilemma game (Stanley, Ashlock, and Tesfatsion 1994), and the emergence of self-organized markets (Vriend 1995).⁴

R. Marimon, E. McGrattan, and T. J. Sargent (1990) study the emergence of fiat money as a medium of exchange. Fiat money is an intrinsically worthless commodity because it does not appear in any utility or production function. People accept such a commodity in exchange only when they believe that others will accept it from them. Theoretical work by N. Kiyotaki and R. Wright (1989) showed that equilibria—that is, situations from which no individual agent can improve his own condition by deviating unilaterally—may exist in which fiat money takes on value for exactly that reason. In principle, fully rational agents could reason what such an equilibrium would be, and behave accordingly without further delay, but there might be some practical difficulties in coordinating on such an equilibrium. One way to achieve such an equilibrium would be through the hand of a planner who sets up a complex system of exchange institutions, designating one commodity as mandatory medium of exchange. The question analyzed by Marimon, McGrattan, and Sargent, however, is whether simple adaptive agents can “set up” such complex social arrangements spontaneously, without the hand of a planner.

Their basic model is one in which they impose the need for indirect trade. When two agents meet, direct trade is possible only if agent 1 has what agent 2 wants, while

4. Arthur, Durlauf, and Lane 1997 is a good starting point for additional recent papers following a similar approach.

at the same time agent 1 wants what agent 2 has. This condition is called a “double coincidence of wants.” Lacking a double coincidence of wants, agents must resort to indirect trade. If agent 1 has a commodity that agent 2 wants, he might accept from agent 2 in exchange a commodity that he does not want but for which he expects to find some agent 3 who does want it, an agent 3 who possesses a commodity that agent 1 himself wants. Such a situation with indirect trade is called a “Wicksell triangle.” In their basic model, Marimon, McGrattan, and Sargent assume that there are just three types of agents, with the production and utility functions of each type defined such that type 1 agents want to consume commodity 1 but can produce only commodity 2, type 2 agents want to consume good 2 but can produce only good 3, and type 3 agents want to consume good 3 but can produce only good 1. Hence indirect trade is necessary. After an analysis of this basic model, the analysts consider the following important variant. Suppose there exists one additional good, commodity 0, which cannot be produced or consumed by any of these agents. Will the agents learn to use that commodity as a general medium of exchange?

Marimon, McGrattan, and Sargent drop the assumption of rational agents in favor of “artificially intelligent” agents who modify their behavior over time on the basis of their experience. Each agent begins a time period carrying a particular good. A random process then matches the agents, who decide whether or not to trade goods. Once the decision is made, the agents must decide to consume the good or carry it into the next period. Consumption of the first good leads to the production of a new good that will be carried into the next period instead. To make their trading and consumption decisions, agents use two Classifier Systems. The first uses the pre-trade information of the goods held by each partner to determine the trading decision. Trade takes place only when both agents agree to do so. The second Classifier System examines the post-trade goods to determine the consumption decision. The Classifier Systems are a collection of decision rules with an accounting system that selects rules to be followed. The basic learning idea is that agents experiment using different decision rules through trial and error and are more likely to use rules that led to better payoffs in the past. The aim is to select “co-adaptive rules” that will work successfully over the range of situations the agent may face. Computer simulations of two sorts of Classifier Systems are implemented for the Kiyotaki-Wright model. The first carries a complete enumeration of all possible rules. The second does not carry a complete enumeration, but if an unforeseen situation arises, the Classifier System will create and experiment with new rules using a Genetic Algorithm. The simulation results show that commodity 0 emerges as a generally accepted medium of exchange. Hence, Marimon, McGrattan, and Sargent demonstrate how simple adaptive agents can “set up” complex arrangements such as fiat money without the hand of a planner.

The study by W. B. Arthur and others (1996) is similar in spirit but deals with financial markets. The analysts ask, first, what kind of regularities exist in such markets, and second, how can we explain such regularities as resulting from the behavior of a

large number of interacting individual agents. Traditionally, academic theorists and market traders have held two distinct and opposing views of financial markets. Standard theory assumes identical investors. These agents share rational expectations of future asset prices and use all market information to adjust their price expectations. No one can consistently earn speculative profits in such an environment. Technical trading, where traders use patterns in past price movements to predict future patterns, never arises; nor do price overreactions such as temporary bubbles and crashes. Trading volume is very low or zero, and volatility in volume and in prices are not serially correlated, which means volatility in the current period is not linked to volatility in past periods. Thus, according to standard theory, the market must be rational and efficient. Traders, however, believe that the market *does* offer speculative opportunities; technical trading can be profitable; a “market psychology” exists; bubbles and crashes occur and are often unrelated to market news. The traders regard the standard academic theory as unrealistic. Evidence has forced academics to reconsider the standard economic view. Studies of the 1987 crash of the financial markets have found no significant correlation between the crash and market information of the time. Volume and price volatility have been found to be large, technical trading rules have been shown to provide statistically significant long-term profits, and significant serial correlation has been proven to exist.

Arthur and others analyze a simple model of an artificial stock market in which there is a risky stock paying a stochastic dividend and a risk-free bond. Each individual agent uses a set of predictors to form forecasts of next period’s prices and dividend. These predictors are rules based on pattern recognition. Some rules summarize market fundamentals such as price-dividend ratios, whereas others are technical trading indicators such as price-trend movements. Because the price depends on the behavior of other agents, people need to forecast the forecasts of other peoples’ forecasts of other peoples’ forecasts, and so on. Given their forecasts, they compute their optimal action, and the market price is determined. Using a Classifier System they update their confidence in the rules they have been using, and with a Genetic Algorithm they generate new forecasting rules to be tested. In this way, individuals form their own theories of the market, test them, and use those with the most predictive power to make decisions to buy or sell. Expectations therefore evolve and compete in a market shaped by other agents’ expectations. The crucial question is whether these heterogeneous expectations coevolve into homogeneous rational-expectations beliefs, supporting the efficient-market theory held by academics, or instead a richer pattern of individual and collective behavior emerges, upholding the traders’ view.

The results confirm both views, under different regimes. The difference lies in the speed at which agents adapt their forecasts to new market information. If agents explore alternative expectation rules at a low rate, the market price converges rapidly to the homogeneous rational-expectations equilibrium value as theorized by the academics. If some agents forecast different values, then the fact that most of their coun-

terparts have forecasts close to the homogeneous rational-expectations equilibrium value returns a market-clearing price that eliminates most of the deviation. With a perhaps more realistic rate of exploration of alternative forecast rules by the individual agents, the traders' view of a complex pattern of beliefs seems to be confirmed. The appearance of bubbles and crashes reveals that technical trading has emerged in the market as agents buy and sell into trends. Increased volatility in prices and volume as well as an increase in volume are apparent. This new set of behavior is referred to as a complex regime. It should be stressed, however, that in both regimes the regularity that emerges at the market level grows out of the interaction of many individual agents, all continuously trying to learn to do the best they can.

In their study, J. M. Epstein and R. Axtell (1996) consider the art of growing artificial societies. The basic purpose of their models is to study the emergence of social structure and group behavior from the interactions of the individual agents. They do so in a "Sugarscape," which consists of two parts. First, there is the Sugarscape proper—the physical environment or landscape—which is divided into cells. Each cell holds a general resource, for example, sugar, that the agents wish to consume. Some regions in the landscape are rich in sugar, some poor. Second, there are the agents, or the "people," in these artificial societies. Each agent has an internal state and a set of behavioral rules. The states can be fixed, such as sex and metabolism, or they can be altered through interactions with other agents. Examples of the latter include economic preferences, wealth, and health. Changes of states depend on rules of behavior for the agents and on the space or environment in which they live. The rules govern the agents' interactions with other agents and the agents' interactions with the environment; they also determine the computing ability of the agents and the amount of information available to each individual agent. Populations expand through sexual reproduction, and cultural characteristics are transmitted, producing local "tribes."

In the basic model, agents are born with a certain metabolism and burn sugar with each move. The behavioral rule is to look around as far as the individual's vision permits, to find the spot with the most sugar, then to go there and eat it. The following phenomena emerge from this "Sugarscape" model. It is evident that the environment can support only a finite population. A wealth distribution (measured in sugar) exists, and the analyst can examine the variety of situations that lead to skewed distributions indicating income inequality. If seasons are introduced into the model, then migration occurs; some regions face greater competition for limited resources, and the host agents confront national security implications. If the landscape contains two commodities (sugar and spice), then trade patterns emerge. The agents have different sugar and spice metabolisms, which determine the relative preferences for the two goods. Agents must have positive amounts of each in order to survive. Pairs of agents bargain to a local price and then exchange if both parties are made better off. When credit relationships are allowed, complex networks arise. Lenders assess prospective borrowers' ability to repay the loan, using information about past income. Granted

loans are repaid on the due date or renegotiated.

Epstein and Axtell show that these artificial-society models serve as a convenient means of studying the dynamics of models in which heterogeneous agents have genetic and culturally inherited rules and characteristics. Over thousands of time periods, the analyst can examine the emergence of sudden cultural fads and other temporary phenomena. Agent-based models track the individuals over time, revealing the emergence of macro-level social structures, and because the agents interact with the institutions they have jointly created, one can also study how these institutions in turn affect the agents' behavior.

E. A. Stanley, D. Ashlock, and L. Tesfatsion (1994) analyze the well-known iterated prisoner's dilemma game (see, for example, Axelrod 1984). This game is a popular metaphor for a common social dilemma, which has the following general feature. Each player chooses one of two possible actions, say, Cooperate and Defect. Whatever the actions chosen by the other players, the immediate payoff to an individual player is higher when choosing Defect than when choosing Cooperate. However, if all players choose Defect, their payoff is lower than when all choose Cooperate. Hence, the fundamental question is whether in a population of egoists we might expect the emergence of Cooperation.

In the literature, the mere repetition of the game has emerged as a factor enhancing cooperation. Repetition allows a player to reward cooperative opponents by continuing to play Cooperate himself while punishing defecting opponents by switching to Defect himself. Hence, a defecting player gains in the short run, but in the longer run he might very well lose much more. It turns out that agents with very limited strategic reasoning capabilities, playing strategies like "tit-for-tat" (start with Cooperate, and in each succeeding round choose Cooperate whenever your opponent chose Cooperate in the preceding round, Defect whenever your opponent chose Defect in the previous round) and similar strategies do relatively well from the individual players' point of view. But obviously, in real life, economic agents do not interact blindly. A simple alternative way to punish a defecting partner is to refuse to play additional games with him. Stanley, Ashlock, and Tesfatsion consider this alternative.

They analyze an iterated prisoner's dilemma (IPD) game with a fixed pool of agents that use simple rules of thumb, where each rule is a function of what the opponent did in the previous couple of rounds. The agents modify their behavior every now and then: the probability of adopting a certain rule of thumb depends on the payoffs generated by other players in the population using that and similar rules in the past.⁵

The special feature of this model is that players may choose or refuse partners. This endogenous determination of playing partners is implemented as follows. Each player keeps track of the payoff he realizes with each and every individual. If it falls

5. Following Miller (1996), the individual players are modeled as finite automata, whereas the modification takes place with help of a Genetic Algorithm.

below some minimum tolerance level, the player will no longer accept game offers from that opponent. Simulation results were obtained for this evolutionary IPD game both before and after the introduction of the choice-and-refusal feature. The change fundamentally modifies how the players interact. For example, refusal gives players a way to protect themselves from defections without having to defect themselves, and ostracism of defectors occurs endogenously. Choice and refusal speed up the emergence of cooperation, but also the emergence of stable uncooperative player interaction patterns. For example, high values for the minimum tolerance level or the wallflower payoff (when an agent neither proposes nor accepts a game) lead to the emergence of ecologies with primarily antisocial agents. Parasitic relations also emerge, as choice and refusal allow clever "rip-off" players to exploit vulnerable partners. The underlying player interaction patterns induced by choice and refusal can be complex and can vary over time, with agents clustering into distinct subpopulations, even when expressed play behavior is largely cooperative.

N. J. Vriend (1995) studies a model of decentralized trade with firms that produce a given commodity and consumers who wish to purchase repeatedly one unit of that commodity. Consumers "shop around," while firms attract the attention of potential customers by sending information signals and offering good service. The main objective of the study is to present an example of a computational approach to the following question: How do self-organized markets emerge in the economy, and what are their characteristics? Trading opportunities perceived by the locally interacting individual agents change endogenously from period to period, and a combination of Genetic Algorithms and Classifier Systems is used to model the adaptive behavior of the individual agents. Vriend performs a simulation of the model with fifty firms and five thousand consumers for two thousand periods. After an initial phase of "overall" learning, the macroeconomic situation is characterized by comparatively steady aggregates. Competition appears to lead to coordination of economic activities; communication by firms and patronage by consumers play important roles; and high communication expenditures are the main source of macroeconomic inefficiency. The microeconomic distributions underlying the aggregates show strong differences between the market shares of firms and between the shopping behavior of consumers. However, all firms offer an identical service rate and incur the same costs per unit sold, and all consumers have the same rate of success.

To illustrate the significance of the emergence of self-organized markets—that is, of the fact that all actions and outcomes in the simulated economy emerge purely *endogenously* as the result of locally interacting and learning individual agents who are all continuously looking for advantageous opportunities—Vriend analyzes a model in which one aspect of the behavior of the individual agents is fixed exogenously. The exogenous restriction is that if a consumer has been satisfied by a firm, then he will patronize that firm the next day. This apparently small and intuitively reasonable modification of the choice menu of the consumers leads to very different market outcomes.

In this variant with “fixed” patronage, firms realize higher sales with a lower average production level, implying that less unsold stock perishes, and they spend much less on signaling. As a result, the firms obtain a larger profit. Because the consumers are on average slightly better off as well, the overall efficiency of the economy is much higher. The main source of this improvement is the substantial saving of communication expenditures. Note that the only difference from the standard version of the model is a *restriction* of the consumers’ behavior—all the actions and favorable outcomes of the variant with fixed patronage are also feasible in the unrestricted standard version. In that version a satisfied individual consumer does not always repatronize because it would not be rational to do so. Repatronizing by satisfied consumers makes consumers better off on average only when patronage is fixed so that both the firms and all the other consumers change their behavior. Furthermore, in the standard version without fixed patronage, firms could decide to signal only very little. Then consumers would be more or less forced to patronize, making everybody better off, especially the firms. Again, the point is that for an individual firm that decision would not be rational. This case illustrates an important difference between self-organized markets and centrally planned markets.

Conclusion

The examples presented show that many economic regularities traditionally related to the metaphor of the “invisible hand” can be analyzed explicitly as emergent properties of a process of interacting individual agents. Hence, for economic analysis the approach followed in the study of complex adaptive systems allows us to look inside the black box of the metaphor of the invisible hand and to understand the underlying interactive processes that govern the behavior of decentralized economies.

The ultimate objective of this kind of analysis is, of course, not to become wise with respect to artificial worlds but to understand what is going on in real decentralized economies. Can analysts actually recover known real-world regularities in simple, simulated models and identify how those regularities depend on parameter choices or modeled mechanisms? Preliminary results suggest an affirmative answer; see, for example, the study by A. P. Kirman and N. J. Vriend (1997), which uses a model of the market with simple adaptive agents to analyze the stylized facts of loyalty and price dispersion as observed in the wholesale fish market in Marseille. Simulations of artificial economies serve the same purpose as any formal, mathematical model that abstracts from many aspects of reality. They may suggest how one might *understand* what is going on in a decentralized economy. And because emergent phenomena are inherently unpredictable to some extent, the study of complex adaptive systems confirms that the true objective of science is not prediction but understanding.

References

- Arthur, W. B. 1992. *On Learning and Adaptation in the Economy*. Working Paper 92-07-038, Santa Fe Institute.
- . 1994. Inductive Reasoning and Bounded Rationality. *American Economic Review* 84: 406–11.
- Arthur, W. B., S. N. Durlauf, and D. A. Lane 1997. *The Economy as an Evolving Complex System II*. Vol. 27. Santa Fe Institute Studies in the Sciences of Complexity. Redwood City, Calif.: Addison-Wesley.
- Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler. 1996. *Asset Pricing under Endogenous Expectations in an Artificial Stock Market*. Working Paper 96-12-093, Santa Fe Institute.
- Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Bak, P., K. Chen, J. A. Scheinkman, and M. Woodford. 1993. *Aggregate Fluctuations from Independent Sectoral Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics*. Working Paper 93-01-004, Santa Fe Institute.
- Benassy, J. P. 1982. *The Economics of Market Disequilibrium*. New York: Academic Press.
- Casti, J. L. 1994. *Complexification: Explaining a Paradoxical World through the Science of Surprise*. New York: HarperCollins.
- Debreu, G. 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New York: Wiley.
- Epstein, J. M., and R. Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, D.C.: Brookings/MIT Press.
- Gell-Mann, M. 1994. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. New York: Freeman.
- Hayek, F. A. 1948. *Individualism and Economic Order*. Chicago: University of Chicago Press.
- Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. 2nd ed. Cambridge, Mass.: MIT Press.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.
- Kiyotaki, N., and R. Wright. 1989. On Money as a Medium of Exchange. *Journal of Political Economy* 97: 927–54.
- Kirman, A. P. 1994. *Economies with Interacting Agents*. Working Paper 94-05-030, Santa Fe Institute.
- Kirman, A. P., and N. J. Vriend. 1997. Evolving Market Structure: A Model of Price Dispersion and Loyalty. Manuscript.
- Klamer, A. 1984. *The New Classical Macroeconomics: Conversations with New Classical Economists and Their Opponents*. Brighton: Wheatsheaf.
- Lane, D. A. 1992. *Artificial Worlds and Economics*. Working Paper 92-09-048, Santa Fe Institute.

- Langton, C. G., ed. 1989. *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, 1988*. Vol. 6, Santa Fe Institute Studies in the Sciences of Complexity. Redwood City, Calif.: Addison-Wesley.
- Leijonhufvud, A. 1993. Towards a Not-Too-Rational Macroeconomics. *Southern Economic Journal* 60: 1–13.
- Lewin, R. 1993. *Complexity: Life at the Edge of Chaos*. London: Phoenix.
- Mas-Colell, A., M. D. Whinston, and J. R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.
- Marimon, R., E. McGrattan, and T. J. Sargent. 1990. Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents. *Journal of Economic Dynamics and Control* 14: 329–73.
- Miller, J. H. 1996. The Coevolution of Automata in the Repeated Prisoner’s Dilemma. *Journal of Economic Behavior and Organization* 29: 87–112.
- Sargent, T. J. 1993. *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: Wiley.
- Sigmund, K. 1995. Darwin’s “Circles of Complexity”: Assembling Ecological Communities. *Complexity* 1: 40–44.
- Smith, A. [1776] 1976. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Glasgow edition. Oxford: Clarendon.
- Stanley, E. A., D. Ashlock, and L. Tesfatsion. 1994. Iterated Prisoner’s Dilemma with Choice and Refusal of Partners. In *Artificial Life III*, edited by C. G. Langton. Vol. 17, Santa Fe Institute Studies in the Sciences of Complexity. Redwood City, Calif.: Addison-Wesley.
- Vaughn, K. I. 1989. Invisible Hand. In *The Invisible Hand (The New Palgrave)*, edited by J. Eatwell, M. Milgate, and P. Newman. London: Macmillan.
- Vriend, N. J. 1994. A New Perspective on Decentralized Trade. *Economie Appliquée* 47: 5–22.
- . 1995. Self-Organization of Markets: An Example of a Computational Approach. *Computational Economics* 8: 205–31.
- . 1996. Rational Behavior and Economic Theory. *Journal of Economic Behavior and Organization* 29: 263–85.
- Waldrop, M. M. 1992. *Complexity: The Emerging Science at the Edge of Order and Chaos*. London: Viking.
- Weintraub, E.R. 1979. *Microfoundations: The Compatibility of Microeconomics and Macroeconomics*. Cambridge: Cambridge University Press.

Acknowledgments: We wish to thank John Casti, Joop Hartog, and Leigh Tesfatsion for critical comments and helpful discussions. The comments by two anonymous referees led to considerable improvements of the paper. All errors and responsibilities are ours.