Special Issue article on Complexity

Politics, Philosophy & Economics
11(1) 5–25

DOI: 10.1177/1470594X11433745
ppe.sagepub.com

A complexity perspective on institutional design

Scott E. Page
University of Michigan, USA

Abstract

The task of designing effective economic and political institutions requires substantial foresight. The designer must anticipate not only the behavior of individual actors, but also how that behavior will aggregate. Rising complexity brought about by increases in speeds of adaptation, diversity, connectedness, and interdependence make institutional design all the more challenging. Given the focus on equilibria, the extant literature on mechanism design might appear incapable of coping with this complexity. Yet, I suggest that a deeper engagement with the origins of the mechanism-design framework reveals insights that may help us design robust, adaptive institutions that can harness complexity.

Keywords

mechanism design, complexity, robustness, learning

1. Introduction

Our world increases in complexity with each passing day. Whether one studies national economies, global politics, or civic cultures, one notices that actions have become more interdependent, economies and people more connected, the relevant actors more diverse, and the rate of adaptation greater. On just this last point, when you walk into a modern grocery store and buy a carton of milk, your actions tip off a web of intelligent software that updates inventories and, if necessary, phones the cow. The rate of change in the

Corresponding author:
Scott E. Page, Center for the Study of Complex Systems, 317 West Hall, University of Michigan, Ann Arbor, MI 48106, USA
Email: spage@umich.edu

economy alone has increased to such a point that leading consultants quantify not just the state and direction of the economy, but also its acceleration (Hagel et al., 2009).

In this article, I contemplate the implications of increasing complexity for the problem of institutional and mechanism design. In particular, I ask whether and how we should take complexity into account when constructing economic, political, and social institutions. At its core, design requires foresight. The designer must anticipate responses and outcomes. Thus, the existence of complexity creates a problem for design: by definition, complex systems are difficult to predict (Page, 2010). Thus, as the world becomes more complex, even the best efforts of the sharpest minds cannot make accurate long-term forecasts (Jervis, 1998; Orrell, 2007; Tetlock, 2006; Watts, 2011).[1] Nevertheless, I will argue, we can still predict some characteristics of outputs, and some institutions may create more predictable outcomes than others, provided they are designed with an understanding of the causes of complexity.

As an organizing framework, I rely on an extensive literature devoted to mechanism design (Maskin, 2007). The mechanism-design literature characterizes an economic or political institution as consisting of six parts: an *environment*, a *message space*, a space of *outcomes*, a *response function* (or behavioral rule) for individuals, an *outcome function* that maps behaviors into the space of outcomes, and a *social choice correspondence*: a set of idealized outcomes given the environment. This analytic framework proves sufficiently general to encompass most institutional settings, including exchange economies, networks of banks, and legislative bodies. It can also help to organize our thinking about how complexity arises, why complexity matters, and what we might do to harness complexity for our betterment (Axelrod and Cohen, 2000; Beinhocker, 2006).

Mechanism design and complexity theory might at first appear incommensurable. The former focuses on the equilibria of systems. The standard mechanism-design perspective on institutions can be summed up as follows: institutions produce equilibria; better institutions produce better equilibria. A complexity perspective, while not denying equilibria, admits other classes of phenomena, such as cycles, randomness, and complex dynamics, that can produce large events such as stock-market crashes and the collapse of markets.

One of my challenges, then, will be to describe how the mechanism-design approach can be modified to embrace complexity. This will not require 'round holing' a square peg. In fact, it is more of a return to the origins of mechanism design, or what Hurwicz called *adjustment processes*. Adjustment processes can produce both equilibria and complexity (Hurwicz, 1994). To give away a bit of the plot, mechanism design, as currently practiced, assumes that agents optimize. Adjustment-process models assume that individuals follow rules. They consider how and whether those rules lead to equilibria. For example, in Fisher's seminal model of price adjustment without an auctioneer (1972), buyers sample a fixed, random sample of sellers. Such behavior could not possibly be optimal, but it suffices to get to equilibrium. Complex systems models can lie in between, with agents relying on simple rules, but adapting those rules over time. This adaptation will be one of the building blocks needed for a mechanism to allow for complexity.

The remainder of this article consists of four parts. In the first, I provide a brief overview of the components of a complex adaptive system and describe the phenomena produced by such systems (Miller and Page, 2008; Mitchell, 2009). In the second part, I provide a primer on mechanism design. By juxtaposing these two theoretical frameworks, I intend not to make any normative comparison, but to highlight core assumptions of each approach to thinking about the world. In part three, I describe how to extend mechanism design to include the core assumptions of complexity theory. Finally, in part four, I discuss what we might learn about how to design better institutions by integrating ideas from complexity into mechanism design.

2. Complex adaptive systems: attributes and outputs

I begin with a brief overview of complexity and what many call 'complex adaptive systems'. I distinguish between the attributes of a complex system and the phenomena, or outputs, that such systems produce. The attributes of a complex system include *diverse agents*, that are *connected* either virtually or geographically, who follow *adaptive, rule-based behaviors*, and whose choices are *interdependent* in meaningful and often non-linear ways (Holland, 1995, 1998).

Two points are in order. First, a system that has these attributes can, but need not, produce complex outcomes. Therefore, we might more accurately refer to them as 'complex (capable) systems'. Clearly, an economy possesses all of the attributes of a system capable of producing complexity, but not all of the outcomes in an economy are complex. Second, systems that lack these attributes tend not to produce complexity. Thus, the toy models used to teach economics do not produce complexity because they lack the necessary attributes.

I do not intend the previous sentence as a criticism, but as a statement of fact. By creating models that lack, say, diversity or adaptation, economists obtain well-behaved outcomes allowing them to formulate clean hypotheses and make coherent forecasts and policy recommendations.[2] They can also design optimal institutions. But that optimality applies only in a limited context.

When I say that a system produces complexity, that still leaves open the question of what exactly complexity is. Complexity has an abundance of definitions – at least 20. Surveys of the literature distinguish between classes of definitions of complexity (Mitchell, 2009). In a recent book (Page, 2010), I suggest that these many definitions fall into the following two broad categories.

BOAR. Complexity lies between order and randomness.
DEEP. Complexity cannot be easily described, evolved, engineered, or predicted.

BOAR and DEEP are more similar than they might appear. BOAR places complexity between ordered and random. Ordered systems are not DEEP. They can be described, engineered, and predicted. Often, they can be evolved as well. Random processes are also easily described and engineered, and if stationary, easily predicted (at least at the distribution level). In contrast, a complex process cannot be described in just a few words, nor be predicted accurately. Thus, what lies between order and randomness will tend to be DEEP.

Although these definitions do not provide a clean test for complexity, they do roughly agree on what systems they call complex. But they also differ in their particulars. One process may be more ordered than another, but harder to evolve if that order results from elaborate feedback structures. Moreover, that difference may be a strength. Complexity, by definition, implies high information content. Thus, for many of the same reasons that anthropologists have created hundreds of definitions of culture, we also want an abundance of definitions of complexity.

At this point, it is worth commenting on whether the empirical evidence for complexity can be documented using these measures.[3] Time series of GDP, housing prices, or aggregate stock returns suggest complexity exists. That said, the patterns that we see in those data are not the elaborate structured patterns of the kind that result from the fixed rules shown in Wolfram (2002). That lack of structure emerges because any predictable pattern could be exploited for gain. Some patterns though, such as the clustered volatility of stock-market prices, cannot be exploited, or at least cannot be exploited easily. As a result, complexity scholars view the clustered volatility as evidence of complexity. Stock-market returns also exhibit fat tails (Mandelbrot, 1963). These large movements in prices suggest interdependent behavior, another characteristic of complex systems.

Some would call these fluctuations evidence of market uncertainty and not complexity. I do not see this as mere semantics. Uncertainty is not the same thing as complexity.[4] Uncertainty refers to a lack of information about the state of the world. Complexity arises from interactions between diverse actors with interdependent behaviors who adapt to one another (Miller and Page, 2008).

2.1 Complexity, but not always

I next reiterate a distinction between the attributes that define systems capable of producing complexity and the complex outcomes themselves. As mentioned earlier, just because a system has diverse, interacting agents does not mean that it produces complex outcomes. It could be that the system can be explained by a few macro-level variables. For example, many complex systems models of the economy take agents as the primitive and then explore how their actions aggregate (Tefatsion, 2006). These models need not produce complex outcomes. They can produce efficient equilibria. Economics has a long history of demonstrating how interactions can produce simplicity (Hayek, 1945). More recent experimental and computational work unpacks exactly how market mechanisms produce equilibria despite complex underpinnings (Dickhaut et al., 2010; Gode and Sunder, 1993).

In all these cases, aggregation reduces complexity. Thus, as a prelude to my discussion of mechanisms, I want to spend a moment on aggregation. To organize this discussion, I construct an *aggregation operator A* that maps the vector of the agents' states at time t into an aggregate variable Y^t or vector of variables $\vec{Y}^t$. I also define an *aggregate-level mapping H* that transforms the aggregate variable at time t to its state at time $t+1$ (see Figure 1).[5]

Consider a macroeconomic model of the economy. It focuses on the aggregates (the top row of Figure 1), ignoring the lower level. Even though, in reality, the economy consists of many individual agents all interacting, macro models ignore them. In other words, the model assumes that the particulars do not matter. Stated mathematically, this

$$
\begin{tikzcd}
\begin{array}{c}\text{Aggregate State at t}\\ \vec{Y}^t = A(\vec{x}_1^t, \vec{x}_2^t, ..\vec{x}_n^t)\end{array} \arrow[r, "H"'] & \begin{array}{c}\text{Aggregate State at t+1}\\ \vec{Y}^{t+1} = H(Y^t)\end{array} \\
\begin{array}{c}\text{Agents' State at t}\\ (\vec{x}_1^t, \vec{x}_2^t, ..\vec{x}_n^t)\end{array} \arrow[u, "A"] \arrow[r, "F"'] & \begin{array}{c}\text{Agents' State at t+1}\\ (\vec{x}_1^{t+1}, ..\vec{x}_n^{t+1}) = F(\vec{x}_1^t, ...\vec{x}_n^t)\end{array} \arrow[u, "A"']
\end{tikzcd}
$$

Figure 1. The aggregation diagram

means that the diagram *commutes*, that is, $H[A(x_1^t, x_2^t, \ldots x_n^t)] = A[F(x_1^t, x_2^t, \ldots x_n^t)]$. Isawa et al. (1987) provide conditions for that to be the case.[6]

One way to make the diagram commute is to make all agents identical. In that setting, all of the x_i's will take on the same value and Figure 1 is more likely to commute. However, we also need that the system has limited interdependency among the agents. The centrality of interdependence to complexity will be a consistent theme in this article, so it is worth unpacking.

I will contrast two examples, one formal and one informal. The first will be a standard model of consumer demand for a single good. In that model, the primitives will be the agents' preferences and incomes. Assume that agents' preferences are such that they spend a fixed proportion of their income on the good. Let x_i^t denote the income of agent i at time t and let the demand of agent i' for the product equal $a_i + bx_i^t$. The function that describes aggregate states can then be written as $F(x_1^t, x_2^t, \ldots x_n^t) = (x_1^t + 1, x_2^t + 1, \ldots x_n^t + 1)$. Let Y^t denote aggregate income, that is, the sum of the agents' incomes.

A straightforward calculation shows that Figure 1 commutes. Let

$$A(\vec{x}_1^{\,t}, \vec{x}_2^{\,t}, ..\vec{x}_n^{\,t}) = \sum\nolimits_{i=1}^{n} \vec{x}_i^{\,t}$$

and

$$H(Y^t) = Y^t + n$$

It follows that $H(Y^t) = A[F(\vec{x}_1^{\,t}, \vec{x}_2^{\,t}, ..\vec{x}_n^{\,t})]$. To calculate the sum of the agents' demands, we can use either the agent-based model or the aggregate model. In the former case, we just sum the individual demands of the agents. In the latter case, aggregate demand in time t equals $A + bY^t$, where

$$A = \sum\nolimits_{i=1}^{n} a_i$$

In this example, agents spend a fixed proportion of their income on the product. Therefore, the entire population of agents also spends a fixed proportion of its income on the good, and, as a result, the aggregate model works perfectly. This condition (that is, identical proportional behavior across agents) implies almost no interactions, and therefore aggregation works. Implicit in this construction was the idea that agents optimize effortlessly, and that reduces the potential for complexity. Microeconomic models of exchange economies with well-behaved preferences such as described here will tend to produce some complexity as agents learn effective trading rules over time (Gintis, 2007). From this example, we learn that exchange economies have few interdependencies, and, as a result, they are not likely to be complex.

For my second example, I consider financial markets. In a financial market, a person's value of a good depends on what other people think about the value of that good as well (LeBaron, 2001, 2006). A bundle of mortgages could be worth a lot of money if everyone else thinks the mortgages will be paid off and could be worth nothing if people think that they will not be paid off. In a summary of financial folly spanning centuries, Reinhart and Rogoff (2009) show, among other results, that leverage ratios need not be a good predictor of impending doom. Markets can be highly leveraged yet remain stable. They also can be highly leveraged and on the brink of collapse. Market prices depend on people's opinions, and those can move around in response to information and the actions of others. The result is that financial markets produce complex outcomes, including sequences of temporary patterns and occasional large events (Bak, 1996) such as the 1987 crash or the more recent home mortgage crises. In sum then, systems with lots of interdependencies and diversity may not aggregate as easily as those that do not. This makes constructing accurate simple models difficult, if not impossible. I now turn to what this implies for mechanism design.

3. Mechanism design

Recall from the introduction that a mechanism consists of six parts: an *environment*, a *message space*, a *response function* (g), an *outcome function* (h), a *space of outcomes*, and a *social choice function correspondence* (F). Figure 2, known as the Mount-Reiter diagram, shows the relationships between the six parts.

The environment resides in the upper left of the Mount-Reiter diagram. The environment consists of all relevant information for the mechanism. For the purposes of this article, I will assume a set of agents each having a type θ_i that captures everything relevant about the agent. The type will encompass preferences, information sets, and beliefs. If the agent represents a firm, then its type includes the production technologies available to the firm. The entire profile of types $\vec{\theta} = (\theta_1, \theta_2, \ldots \theta_n)$ constitutes what will be called the 'environment'.

The message space lies at the bottom of the diagram. The message space characterizes the language that the agents use to communicate with one another. In a market, messages could be quantity demands for products or they could be price demands to sell a unit of a good. In a political system, messages could be votes. Designing the message space is one of the tasks confronting the designer of a mechanism.

The other task for the designer is to understand the agents' response function, denoted by g. This function describes how agents map their types into messages. For example, in an exchange economy types might represent initial endowments and messages might represent requested trades. In that setting, the response function could be a requested trade, say, four apples for three bananas.

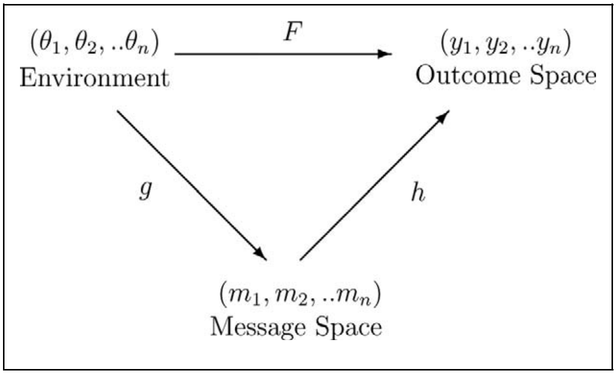

Figure 2. The Mount-Reiter diagram

The messages sent depend on incentives. Those incentives are determined by the outcome function, *h*. The art of mechanism design lies in constructing outcome functions that induce agents to send the right messages. The function maps those messages into outcomes: allocations of goods, electoral outcomes, and so on. So far, my description of a mechanism ignores normative aspects of the outcomes produced. The final piece of a mechanism, the social choice function or correspondence (denoted by *F*), adds a normative dimension.

This distinction between functions and correspondences has important consequences. A *social choice function* is a one-to-one mapping from type profiles to outcomes. For each environment, it prescribes a unique outcome. A *social choice correspondence* does not prescribe a single outcome, but many. Mathematically, a correspondence is a one-to-many mapping from type profiles to outcomes.

Some examples show why this distinction plays such a central role in mechanism design. If types represent preferences over a ballot proposition and outcomes must be either yea or nay, then we might want to assign, for each profile, the outcome preferred by a majority. If so, we would be constructing a social choice function. For each preference profile, we are choosing a single outcome. On the other hand, if types represent preferences over goods and initial allocations of goods and if outcomes represent allocations, then we might only desire Pareto-efficient allocations. If so, for each profile of types, there would exist multiple desired outcomes, and we would assume a social choice correspondence.

If we take a complexity perspective, the social choice correspondence can be seen even more expansively. We might want outcomes in a single period to belong to a set, but we might also want the path of outcomes to belong to a set of paths that represent a flourishing economy. For example, in designing financial markets, we might want outcome paths that do not produce large bubbles and crashes.

Of course, the space between desires and reality can be a chasm. Therein lies the strength of the Mount-Reiter diagram. It graphically captures the fundamental problem of design: the tension between the clean, crisp normative ideal and the messier, sometimes grim reality. Across the top, the social choice correspondence *F* describes the heights to which we aspire. Down below, individuals respond to their perceptions of reality through g and the outcome function h adjudicates the competing claims of messages captured in m. In a perfect world, the diagram would commute. Reality would produce the ideal. Often though, that proves impossible, and no mechanism proves capable of implementing the correspondence F.

3.1. Realization versus implementation of outcomes

A breakthrough in mechanism design occurred when social scientists discovered that given an outcome function, an environment, and a message space, a mechanism created a *game* in which the messages can be thought of as actions and the outcome function renamed a 'payoff function'. Given that formulation, the outcome of a mechanism can be captured as an equilibrium of the game. If an equilibrium of the game lies within the correspondence F, then the mechanism *implements* the social choice correspondence F (Reiter, 1977).

During the past 25 years, implementation has been the standard approach in mechanism design – to solve for equilibria. Thus, given a mechanism, one can restrict attention to the equilibria the mechanism implements. This leaves open the question of what can be implemented. Given the vast space of possible mechanisms, this question would seem unanswerable. Fortunately, progress has been made on this question. Maskin (2007) showed the necessity of the monotonicity of the social choice function. Monotonicity implies that if outcome A is chosen in one environment and A does not become less attractive for anyone in a new environment, then A must still be chosen.

A second breakthrough was the development of the *Revelation Principle* (Myerson, 1992). The Revelation Principle states that any equilibrium implemented by a mechanism can also be implemented by a *direct revelation mechanism*, that is, a mechanism in which the message space equals the type space. A direct revelation mechanism is said to be *incentive compatible* if the truthful revelation of types is a Nash equilibrium. This means that if all other agents truthfully reveal their types, then agent i should also truthfully reveal her type. Note that this does not preclude other equilibria in which agents misrepresent their types. The mechanism could induce many such equilibria.[7]

Now, most of mechanism design considers Nash equilibria, but that was not always true. Hurwicz (1994) originally described what he called *adjustment processes*. Under an adjustment process, the response functions are *behavioral rules* characterized as algorithms. As an example, consider a two-agent pure exchange economy. The behavioral rule for one agent could be to use a randomizing device to produce potential trades and to propose the first trade that he prefers to the status quo, and the behavioral rule for the other agent could be to accept any proposed trade that she prefers to the status quo. These behavioral rules produce a sequence of allocations in which both agents' utilities weakly increase. However, these rules are not optimal because the second agent might not want to accept a trade that she only prefers over the status quo by a small amount, especially if she believes that this trade would greatly benefit the first agent. Accepting such a trade might harm the second agent's future position.

The outcome that would result from the repeated application of these behavioral rules is said to be *realized* by the mechanism. This distinction between *implemented* (equilibrium) outcomes and *realized* outcomes will be central to the analysis that follows. Within economics and political science, implementation has become the standard approach. In simpler language, this means that the rationality assumption has reigned. Why? There are many reasons. Rationality makes a good benchmark (Myerson, 1992). Furthermore, any other behavioral assumption can be criticized as both ad hoc *and* inconsistent with the idea that individuals keep improving their actions. If an agent can improve her lot by deviating from a fixed behavioral rule, why would she not do so?

Two changes in perspective have begun to call into question rationality as a benchmark. Experimental and behavioral economics have raised empirical challenges to whether people optimize. Many studies suggest, instead, that people exhibit common biases (Camerer, 2003). Complementing the increased attention to actual, not idealized behavior has been a recognition of the difficulty of some choice contexts and the complexity of some environments. Difficult problems consist of many variables that interact in nonlinear ways (Page, 2008). In difficult and complex environments, the optimization assumption becomes problematic because of computational constraints. Some of the most interesting work in experimental mechanism design has explored how people perform when the environments are more complex (Brunner et al., 2010; Ledyard et al., 1997).

3.2. Informational and computational requirements

Notice that direct revelation mechanisms remove much of the actual mechanism from the analysis. To see why, consider an exchange economy in which an agent's type would be her endowment and her preferences. In the austere world of blackboard economics, preferences might be a simple vector of parameters, but in the real world preferences might be rather high-dimensional objects. In a direct revelation mechanism, these objects would have to be communicated to an outcome function which would then describe the trades between the agents. Framed this way, the direct revelation mechanism can be thought of as centrally planned exchange. People send all of the relevant information to a central source, which then informs them of their trades.

Contrast this with a price mechanism. A price mechanism would create one agent, a Walrasian auctioneer, who announces a set of prices. The other agents could then communicate their demands for the various goods by sending their ideal vectors of goods. These vectors must cost the same as the agents' endowments given the price vector. Note that revelation of an ideal vector of goods would not be equilibrium behavior – agents could benefit by being more strategic (more on that in a moment). For now, focus only the informational efficiency of the market – How long are the messages that agents send? In a market, the agents send much less information than they do in a centrally planned economy.

This question (that is, the derivation of minimal dimensional message spaces) animated much of the early research in mechanism design. Jordan (1982) showed that a market-like mechanism realized (not implemented) the Walrasian correspondence with the fewest possible dimensions – the idea being that the dimension of the message space is a proxy for informational cost. Reichelstein and Reiter (1988) proved that a slight modification of a market mechanism implemented the Walrasian correspondence in minimal dimensions. Both of these results precluded the use of dimension-reducing encodings. These results imply that market-like mechanisms improve on central planning not because of the outcomes they generate, but because of the informational efficiency with which they generate those outcomes.

The focus on the dimensionality of messages leaves out the amount of computation required of the agents.[8] If one mechanism uses lower dimensional messages than another, but requires more onerous computations, is it really more efficient? It is exactly this question that Mount and Reiter (2002) consider by constructing a model that assumes a set of primitive computations. Loosely speaking, they calculate the computational costs of a mechanism as the number of primitive computations required by the response functions. This construction enables them to derive an efficient frontier of mechanisms. A mechanism lies on this frontier if no other mechanism both has a lower dimensional message space and requires less computation.

3.3. Dominated and unbounded strategies and robust mechanisms

The computational approach of Mount and Reiter captures the difficulty of calculating the equilibrium message, but it does not address the difficulty of coordinating on or selecting the equilibrium. It restricts attention to the direct computation of the equilibrium and ignores the cognitive costs involved in choosing which equilibrium to compute. This second aspect of computation proves challenging to quantify. As a rule, mechanism designers prefer implementation in dominant strategies which obviates the problem of equilibrium selection. However, implementation in dominant strategies often fails.

Designers also prefer not to use mechanisms that include *undominated strategies* (Jackson, 1992). In the early days of mechanism design, a common trick was to append a 'name the largest integer game' to the mechanism. Payoffs could then be written as follows: 'If the equilibrium messages are sent, apply the desired outcome function.' If a player sends a nonequilibrium message, then an enormous payoff goes to the agent who lists the largest integer – in a 'name the largest integer game'. Since the largest integer game has no equilibrium (it is unbounded), appending it to the mechanism wipes out all undesirable equilibria.[9] Mechanisms that include undominated strategies are less predictable. They contain no guarantee that the agents will avoid playing 'name the largest integer' games.

A related and separate question concerns the beliefs that agents have about the types of other agents. Everything discussed so far assumes common knowledge of beliefs. But what if the agents' higher-order beliefs are not common knowledge, which they probably are not? Bergemann and Morris (2005) refer to mechanisms that can implement a social choice correspondence with relaxed assumptions about higher-order beliefs as *robust*.[10] In brief, robust implementation is only possible in what they call *separable* environments. These environments limit the interdependence of preferences, that is, an agent can only care about its own outcome and a global variable.

4. A complexity perspective

Now that I have provided skeletal descriptions of both complex systems and mechanism design, I present an argument that the latter provides a useful framework for thinking about complexity in economic and political systems. In what follows, I introduce core features of complex systems and interpret them through a mechanism-design perspective. I do so by reconsidering the distinction between realization and implementation in five ways. I first consider classes of outcomes. Systems need not attain equilibria. They can also cycle, be chaotic, or produce complexity. I then consider behavioral diversity. What if agents differ in the rules they follow? Third, I discuss the effects of adding learning and adaptation of behavioral rules. Fourth, I introduce networks (that is, agents situated in space) and, finally, I incorporate the idea of institutional ensembles.

After this reconsideration of implementation and realization, I then embed individual mechanisms within a larger context. This construction builds on earlier work on ensembles of games (Bednar and Page, 2008). In that research, we find that multiple games can produce emergent cultures – similar behavioral patterns across games. Here, I relate the idea of a cognitive culture back to the issue of common knowledge of priors discussed by Bergemann and Morris (2005).

4.1. Realizing more than equilibria

Recall the distinction between implementation and realization. Implementation assumes that agents optimize, while realization assumes that agents follow rules. Both approaches restrict attention to equilibria. Implemented equilibria are also Nash equilibria, but realized equilibria need not be. They need only be fixed points of the behavioral rules. Economic models, for the most part, restrict attention to equilibria. Yet, as mentioned, systems of interacting agents can also produce patterns, randomness, or complexity (Wolfram, 2002).

In order for a mechanism to produce these other types of outcome, it must be dynamic. A single-shot game cannot produce a cycle, a random sequence of outcomes, or complexity. Therefore, to adopt a complexity perspective, we have to abandon the idea of implementing equilibria and instead adopt a view that agents follow rules – these rules could be optimal given other agents' actions, but they must be rules. A mechanism would then produce a sequence of outcomes.

As a first step, assume that the rules that the agents follow are fixed. These could be selections from best-response correspondences, in which agents choose a best response to the existing set of messages, or they could be idiosyncratic. Choose a random set of initial messages and then allow the process to iterate in discrete time steps. In some cases, the result will be an equilibrium. If the process does reach an equilibrium, then, following Epstein (2005), we can say that the equilibrium has been *generated*. Note that generativeness is stronger than existence. An equilibrium can exist, but be difficult or impossible to generate.

In one-shot settings with high stakes, we may have no better approach than to assume something close to optimizing behavior on the part of agents *and* to assume that the agents believe that the actions of others will be at or near an equilibrium (Camerer, 2003). Such might be the case on bidding on an oil lease or a military contract. In settings with lower stakes, agents are more likely to follow rules. Those rules may not be optimal. They could be habitual. They could be 'rational' given an incorrect model. They could be rules that suffer from a behavioral bias. Regardless, we can probably assume that whatever rules are in play have survived some sort of winnowing process, and that, on average, they function reasonably well (Gigerenzer and Selten, 2001).

Online auction mechanisms such as EBay provide a nice example of low-stake, one-shot interactions. Some bidders ‘nibble’. They raise their bids by small increments. Their primary goal is to win, but they do not want to pay too much. Other bidders swoop in. They submit electronic bids that enter at the last instant at a fixed price. These bidders’ primary concern is finding great deals (Mittermeier, 2010). From a complex systems perspective, we can think of the community of EBay bidders as an ecology of strategies. The outcomes produced by this ecology may not be exactly what mechanism design would predict, but the theory probably would not miss by far. That is because we can show theoretically that auctions are quite robust to small behavioral errors. Other research, such as Gode and Sunder’s model of zero-intelligence traders (1993), demonstrates the robustness of exchange market mechanisms to nonrational behaviors.

None of this is to deny that the equilibria are not valuable as benchmarks. But one way to evaluate mechanisms might be to consider a variety of initial conditions and a variety of possible behavioral rules and to examine what arises given those combinations. If under a wide variety of assumptions the system goes to equilibrium, then we can have some measure of confidence that comparing equilibria is sensible. If, though, it is extremely difficult to produce equilibria, then equilibria may not be the appropriate solution concept.

4.2. Diverse rules

I next consider the inclusion of diversity. Behavioral diversity often requires that some agents behave suboptimally – if we were all behaving differently, we might not all be able to be optimizing. If, and this is often the case, optimal behavior is unique, then to be different implies not taking the optimal action.[11] In complex systems models, the standard assumption is that agents follow rules, and rather simple ones at that. There exist an enormous number of plausible behavioral rules, and, on top of that, a variety of ways to learn how to adjust those rules.

Just because diversity exists does not imply that it has a meaningful influence on outcomes. Mistakes could cancel out due to large numbers of agents. In addition, as mentioned above, exchange markets do in fact prove robust to errors. In more elaborate mechanisms, diversity does matter. The behavior of a representative agent will not align with what occurs when diverse agents interact, as Kirman (1992) shows in his analysis of spatial markets.

One way to investigate the effects of diversity involves revisiting the insights from the robust implementation literature relating to separability. If an agent does not interact in meaningful ways with the actions of the other agents, then that agent has a clear path of action – what game theorists call a ‘dominant strategy’. The lack of common knowledge in beliefs about types does not have any effect. In a rule-based complex system, individuals do not necessarily optimize with respect to the actions of the other agents, but they do follow coherent rules that tend to produce good payoffs. If the actions are separable, then variations in actions by other agents should not change the behavior specified by a good rule any more than it changes the behavior specified by optimizing behavior. Thus, a central-limit-type logic applies, and variation will just be so much noise about the mean.

If, instead, agents' choices interact in meaningful ways, then diversity can influence results. To put some flesh on this, consider the race-to-the-bottom game (Nagel, 1995). In this game, each individual chooses a real number in the interval [0,100]. The winner is the person who guesses the number closest to two-thirds of the mean guess. The unique Nash equilibrium to this game is to bid zero, but in experiments few people bid that amount in their first play of the game.

To see the effects of diverse rules of behavior, imagine that this game is played over and over with 10 agents. Consider first a case in which every agent initially guesses 10 and thereafter guesses approximately two-thirds of the mean in the previous period. The result will be that the means converge to zero according to the following continuing sequence:

$$10, \frac{20}{3}, \frac{40}{9}, \frac{80}{27}$$

Now, replace one of the 10 players with someone who always bids 70 and replace another player with someone who always bids zero. A quick calculation shows that the result will be a constant mean of 15, a very different outcome than when there was no diversity. In this case, the payoff to an action depends critically on the actions of other agents – so diversity matters.

4.3. Learning and adaptation

The examples of behavior rules described so far have all been fixed. It is also possible that agents can *learn* new rules. That learning could be individualized or social (Vriend, 2000). So, for example, an agent could learn which percentage to shade her bid. Now, of course, the learning rule will itself be fixed, so at the deepest level the agents follow fixed rules (Miller and Page, 2008). Therefore, I will distinguish between *direct rules* (rules that describe actions) and *meta-rules* (rules that modify direct rules). People possess both types of rules. We follow direct rules when we take actions. We employ meta-rules when we decide what rules to use to make actions. Direct rules capture actions. Meta-rules capture learning and adaptation.

If we plop learning agents into a mechanism, we create an alternative criteria for comparing mechanisms. One mechanism might be better than another if the former's good outcome is more easily learned. An explicit example helps to clarify this point. One of the most famous results in mechanism design is the *Revenue Equivalence Theorem* (Myerson, 1992). The theorem states that (given certain assumptions) any auction mechanism will yield the same expected revenue. Suppose that we want to compare a sealed-bid, first-price auction (the good goes to the highest bidder at the highest bid price) to a second-price auction (the good goes to the highest bidder at the second highest bid). Suppose that our social choice correspondence, our goal, calls for the good to go to the buyer with the highest valuation and to yield the most revenue to the seller. By the Revenue Equivalence Theorem both auctions are identical in expectation.

To decide between the two, we then might look at the dimension of the message space. Here again, we have no way to adjudicate between the two mechanisms, as they both use only a single dimension. We might next take the Mount-Reiter approach. In the first-price auction, individuals must do sophisticated calculations based on their beliefs about the distribution of the bid values of the other agents. In the second-price auction, individuals need only bid their true values. So, here, all signs would point to the second-price auction as being simpler. However, as Andreoni and Miller (1995) find in a model using artificial adaptive agents, the second-price auction creates a starker learning environment. Given a fixed set of bids by other players, the payoff function for the other agent has large flat regions. Thus, ease of learnability of an equilibrium is not equivalent to the ease of computability of an equilibrium.

4.4. Networks

Standard mechanism-design models assume that all interactions take place at a single location in an instant in time. That assumption makes sense in many contexts. When you buy a book on the Internet or participate in an online auction, your physical location does not matter. In other contexts, physical or virtual space may play a significant role in how events transpire (Barabási, 2002; Jackson, 2008), even permitting cooperation where it did not exist previously (Nowak and May, 1993).

To see the importance of network structures on outcomes, consider a coordination game in which agents want to match the actions of other agents. Assume that there are only two possible actions A and B. Consider the following behavioral rule: *in the first period, play randomly, then take whichever action was more prevalent in the previous period*. If all agents interact in a single location and if the number of agents is odd, then this rule will produce a random collection of A's and B's in the first period and then either all A's or all B's in the second period, depending on which was more prevalent initially.

We can add a network structure to this interaction by placing the agents in a line and assuming that an agent can only see the actions of neighbors within some distance k in each direction. If we assume the same behavioral rule, with the caveat that the agents now only see the actions of $2\ k$ neighbors plus their own action, then we no longer need get full coordination. This system most often produces clusters of A's and B's. Glaser et al. (1996) construct a variant of this model and interpret A as criminal behavior and B as law-abiding behavior. They then argue that the clusters of criminal behavior produced by the model may help explain local variations in crime.

The addition of networks to a mechanism means that agents' behavioral rules do not apply to the full vector of messages of the other agents. Instead, they only depend on the actions of the agents' neighbors. As seen in the example of the coordination game, the restriction to local interactions can result in multiple equilibria, and, as a consequence, less predictable outcomes and the potential inefficacy of some policy prescriptions. For example, Brock and Durlauf (2002) construct a model of social influence that produces two stable equilibria, one of which can be interpreted as a poverty trap.[12]

4.5. Ensembles of mechanisms, culture, and priors

For my last complexification of mechanism design, I consider multiple mechanisms. This extension aligns with one of the core principles of complexity research: that everything is interconnected. Complex adaptive-systems models often include disease, war, politics, and economic activity simultaneously (Epstein and Axtell, 1996). Though this makes the models harder to interpret, it captures unanticipated interactions between systems. This idea from complex systems has implications for mechanism design in that it suggests that the behavior that emerges in an institutional setting may depend on how that institution is situated. Put differently, complexity theory suggests that by thinking of mechanisms in isolation we are making a substantial error. Behaviors may depend significantly on the entire collection of institutions facing an agent.

The idea of systems-level effects is central in economics. A seminal contribution of neoclassical economics was the development of general equilibrium theory (Debreu, 1959). General equilibrium theory takes into account the full effect of changes in an economy. If a technological breakthrough lowers the labor required to produce a unit of a commodity, then partial equilibrium analysis tells us that the price of that commodity will fall and the amount sold will rise. General equilibrium theory tells us that the prices of all other goods as well as the amounts sold will also change, as will the price of labor, as demand for labor has fallen. These secondary effects can ripple through the economy and aggregate into something larger.

These linkages across systems can be substantial. In *Plagues and Peoples*, McNeil (1976) shows that diseases such as smallpox played as large a role in armed conflict as military might. More recently, Diamond (1997) has sketched out how disparities in wealth and power also depend on disease, climate, and even the sizes of seeds. A complex systems approach to institutions offers the potential for similar types of insights by allowing for the consideration of ensembles of institutions (Bednar and Page, 2008). At present, the results do not exhibit the closed-form grandeur of general equilibrium theory (nor may they ever). Nevertheless, the idea of ensembles of institutions merits discussion and analysis.

Consider an individual in a developed economy. This individual interacts in multiple institutional settings. In the language of mechanism design, she sends messages and receives payoffs in more than one mechanism. These mechanisms may be public, such as markets or online auctions, or they may be within the boundaries of a firm or partnership. We can call the set of all mechanisms that an individual interacts within her *ensemble*. Modeling this individual not as an optimizer, but as a collection of behavioral rules highlights the inefficiency of an individual developing distinct rules for each setting. If the reasons that people follow rules rather than optimize relates to the costs of computation, as Samuelson (2001) assumes, then agents should choose an ensemble-relevant set of behaviors.

How an agent acts in one setting can often be influenced by the other mechanisms in her ensemble. Studies of individual mechanisms preclude ensemble-level effects on cognitive processes. In a world with optimizing agents, we need not worry about cognitive spillovers. Everyone is optimizing. If we instead assume that people follow rules, then we might naturally assume that these rules would flow across contexts. This might occur directly through case-based applications of rules across domains (Gilboa and Schmeidler, 1995) or it might occur indirectly through searches for new rules that combine parts of existing rules. In experimental work, Bednar et al. (2011) show that individuals who learn to alternate actions in one game often transfer that heuristic to other games.

Ensemble dependency provides an entry into the study of cultural behavior. Loosely defined, cultural behaviors can be thought of as behaviors shared by a community of interacting individuals that exhibits some consistency across domains. Thus, cultures can be categorized as trusting, individualistic, risk taking, and so on (Inglehart, 1997). Note the resonance between ensemble effects (Bednar and Page, 2008; Samuelson, 2001) and the cost-of-computation approach of Mount and Reiter (2002). Both assume that thinking takes resources. Mount and Reiter adhere to the optimizing paradigm and seek out mechanisms that balance message-space size and computation. They include the costs of communicating and thinking, but they do not sacrifice any deviations from allocative optimality for potentially large decreases in computation. Samuelson (2001) and Bednar and Page (2008) are willing to trade off a little allocative optimality in exchange for less cognitive effort.

5. Conclusion

In this article, I have provided an introduction to some basic concepts from complex systems and an overview of mechanism design. I have then discussed how a complexity perspective might move mechanism design into new and interesting directions. The first step in that process may be a 'doozy'. It requires abandoning two fundamental assumptions from mechanism design. First, agents must be assumed to adapt rules and not to optimize. Second, systems must be allowed to do something other than attain equilibria.

By making these two changes, economists can include more interdependencies. The emphasis on implementation and, recently, on robust implementation has concentrated attention on separable environments. These are precisely the environments least likely to produce complex outcomes. The two aforementioned changes (that is, allowing adaptive rules and considering broader classes of outcomes) would have the effect of raising the proverbial streetlight under which economists spend their time. The higher light may not shine as brightly, but it will cover more ground.

These emendations to mechanism design do not require abandonment of its core principles and assumptions. If anything, the adoption of rule-based behavior can be seen as a return to the origins of mechanism design. Prior to the rise of game theory, most economists assumed that people follow rules. If the agents apply meta-rules that allow them to adapt better rules, then the assumption of rule-based behavior morphs into learning agents, which lies on the frontier of economics. I would argue that learning models from behavioral economics have more in common with rule-based models than they do with optimization models of behavior.

The second change requires a willingness to consider outcome phenomena other than equilibria. Here again, the change to standard practice will be minor. Many economic models, including learning models and dynamic, stochastic general equilibrium models already consider time to equilibrium. In addition, Markov-perfect equilibria allow for equilibrium distributions as opposed to equilibrium points (Young, 2001). Modern economics admits outcome phenomena other than equilibria. Economic models produce cycles (of the entire economy and of fashions) and they produce bubbles and crashes. Mechanism design, though, has not been as expansive in the outcomes it considers. I believe this to be primarily because of the interests of researchers and not the limits of the approach.

The complexity-based mechanism design that I advocate would rely on the Mount-Reiter diagram as a point of departure, but rather than invoking the Revelation Principle, it would emphasize real mechanisms – the message spaces and the outcome functions. It would take seriously how agents communicate with one another and how those communications get translated into outcomes. I do not advocate abandoning the Revelation Principle. Direct mechanisms are useful as benchmarks, but they should not be an end in themselves. More emphasis should be placed on how the individual behavior aggregates. I am not advocating a return to adjustment processes: agents should be adaptive. Their response functions should change in response to the reactions of others. In an adjustment process, the rules that the agents follow spring from the minds of economists. In what I propose, those rules would emerge, not be built into the model.

If we follow such a path, the important earlier work on message-space dimensions and computational requirements could then be extended to account for the complexity of the learning environment (Mount and Reiter, 2002). That complexity would depend on dimensionality and computational requirements to be sure, and it would also include interdependencies between actions and the network structure. In principle, new measures of the complexity of mechanisms might be constructed that are better suited to the issues that concern economists.

The positive application of a complexity-based mechanism design would be straightforward. Given an environment, two mechanisms (the behaviors, the messages, the outcome function, and the outcomes themselves) can be compared by the paths of outcomes that they produce. These paths can be evaluated by multiple criteria, including average payoff, equality, and robustness. By 'robustness', I do not mean robustness to different prior beliefs, as mentioned earlier; I mean the robustness of the outcomes to the internal dynamics of the agents and to changes in the environment (Bednar, 2009). One need only consider the magnitude of the 2008 bailout of the Troubled Asset Relief Program in the USA to recognize the importance of including the robustness of the outcomes that a mechanism produces.

The normative application of a complexity-based approach would be more challenging. Given rule-based behavior, the social choice correspondence F needs to be reformulated. Currently, F describes the set of desirable equilibrium outcomes given the environment. However, F might describe a set acceptable outcomes or might characterize trajectories that the mechanism would produce. In the later case, the normative concern would be whether a mechanism produces paths of outcomes that lie within that set of trajectories under a broad class of assumptions about behavior.

A final (and speculative) effect of a complexity perspective also relates to normative goals. A case could be made that the current emphasis on efficiency is an artifact of an earlier, mechanistic view of the economy. Just as now we compute the static efficiency of mechanisms, we might one day compute their dynamic complexity. The complexity of the world in which we live is mainly of our own creation. We will not understand that complexity by pulling out the parts and examining them one by one. That is a necessary task, but not a sufficient one. If we want to understand not just outcomes, but emergent phenomena such as culture and large events, we will need more expansive models of institutions. We have to ask ourselves: How much complexity do we want? Clearly, we do not want to rid economic systems of all complexity. Systems that continue to churn, including ecosystems and economies, are more robust. This is partly because they maintain sufficient diversity to continue to explore their surroundings (Page, 2010). Yet, it is also clear that we would like some bound on complexity, lest we produce unpredictable, fragile systems and frequent, large negative events (Crutchfield, 2009). Thus, to marry the insights of Bednar (2009) and Axelrod and Cohen (2000), our goal should be to develop principles of design that harness complexity. To help us on that path, we may have no better starting-off point than a renewed engagement with the deep and varied contributions of mechanism design.

Notes

The author would like to thank the NSF, the Army Office of Research, and the Air Force for funding that supported research related to this project. An earlier draft of this article was presented at a conference on complexity in New Orleans sponsored by this journal. This version includes corrections, improvements, and changes in direction and focus suggested by participants at that conference.

1. For example, the 2002 Congressional Budget Office projections anticipated a government surplus exceeding US$100 billion in 2009. Fiscal year 2009 produced a deficit in excess of US$1.4 trillion (Congressional Budget Office, 2002).
2. Advocates of complexity thinking would argue that this adherence to equilibrium thinking results in unexpected large events, such as the collapse of financial markets. Such large events become likely when we allow a system to become complex (Crutchfield, 2009; Tesfatsion, 1997, 2006).
3. See Durlauf (2012) in this same issue on this point.
4. In Page (2008), I provide an elaboration of the differences between the two concepts.
5. A variant of Figure 1 can also be found in Iwasa et al. (1987).
6. Those conditions limit either the diversity of the agents or the interdependencies between them, and thus rule out much complexity.
7. The proof of the Revelation Principle is straightforward. If the response function g maps a person's type into some message m, and h determines a payoff as a function of m, we can collapse h and g into a single function based on the agent's type.
8. There exists a separate literature on the computation of equilibria by a central authority. See, for example, Scarf (1977).
9. Not all mechanisms that rely on undominated strategies include name-the-largest-integer games. In some cases, such as the Groves-Ledyard solution to the free-rider problem (Groves and Ledyard, 1976), the undominated messages went unrecognized for decades (Page and Tassier, 2010).
10. Bergemann and Morris (2005) distinguish between full and partial robust implementation. The former requires that all equilibria lie in the correspondence. The latter requires that there exists an equilibrium consistent with the social choice correspondence.
11. Many games have multiple best responses, and, in particular, mixed-strategy equilibria allow for diverse optimal behaviors.
12. See Durlauf (2012) in this volume for an elaboration of these ideas.

References

Andreoni J and Miller J (1995) Auctions with artificial adaptive agents. *Games and Economic Behavior* 10(1): 39–64.

Axelrod R and Cohen MD (2000) *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. New York: Free Press.

Bak P (1996) *How Nature Works: The Science of Self-organized Criticality*. New York: Copernicus.

Barabási AL (2002) *Linked: The New Science of Networks*. New York: Basic Books.

Bednar J (2009) *The Robust Federation*. Cambridge: Cambridge University Press.

Bednar J and Page SE (2008) Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Rationality and Society* 19(1): 65–97.

Bednar J, Chen Y, Liu X et al. (2012) Behavioral spillovers in multiple games: An experimental study. *Games and Economic Behavior* 74(1): 12–31.

Beinhocker E (2006) *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics*. Cambridge, MA: Harvard University Press.

Bergemann D and Morris S (2005) Robust mechanism design. *Econometrica* 73: 1771–813.

Brock WA and Durlauf SN (2002) A multinomial choice model with neighborhood effects. *American Economic Review* 92: 298–303.

Brunner C, Goeree J, Holt C et al. (2010) An experimental test of flexible combinatorial spectrum auction formats. *American Economic Journal: Micro-Economics* 2(1): 39–57.

Camerer C (2003) *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton, NJ: Princeton University Press.

Congressional Budget Office (2002) *The Budget and Economic Outlook: Fiscal Years 2003–2012*. Washington, DC: CBO.

Crutchfield JP (2009) The hidden fragility of complex systems? Consequences of change, changing consequences. In: Alsina P and Perello J (eds) *Cultures of Change: Changing Cultures*. Barcelona: ACTAR Publishers.

Debreu G (1959) *The Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New Haven, CT: Cowles Foundation, Yale University.

Diamond J (1997) *Guns, Germs, and Steel: The Fates of Human Societies*. New York: Norton.

Dickhaut J, Lin S, Porter D et al. (2010) *Durability, re-trading, and market performance*. Technical report. Orange, CA: Economic Science Institute, Chapman University.

Durlauf SN (2012) Complexity, economics, and public policy. *Politics, Philosophy and Economics* 11(1): 45-75.

Epstein J (2005) *Generative Social Science*. Princeton, NJ: Princeton University Press.

Epstein J and Axtell R (1996) *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA: MIT Press.

Fisher F (1972) On price adjustment without an auctioneer. *Review of Economic Studies* 39: 1–15.

Gigerenzer G and Selten R (eds) (2001) *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.

Gilboa I and Schmeidler D (1995) Case-based decision theory. *Quarterly Journal of Economics* 110(3): 605–39.

Gintis H (2007) The dynamics of general equilibrium. *Economic Journal* 117: 1289–309.

Glaser E, Sacerdote B and Scheinkman J (1996) Crime and social interactions. *Quarterly Journal of Economics* 111: 507–48.

Gode DK and Sunder S (1993) Allocative Efficiency of Markets with Zero Intelligence Traders. *Journal of Political Economy* 101: 119–37.

Groves T and Ledyard J (1976) Optimal allocation of public goods: A solution to the free rider problem. *Econometrica* 45(4): 783–809.

Hagel J III, Browne JS and Davison L (2009) *Measuring the Forces of Long-term Change: The 2009 Shift Index*. Deloitte Touche.

Hayek FA (1945) The use of knowledge in society. *American Economic Review* 35(4): 519–30.

Holland J (1995) *Hidden Order: How Adaptation Builds Complexity*. Helix Books.

Holland J (1998) *Emergence: From Chaos to Order*. Oxford: Oxford University Press.

Hurwicz L (1994) Economic design, adjustment processes, mechanisms, and institutions. *Economic Design* 1(1): 1–14.

Inglehart R (1997) *Modernization and Postmodernization: Cultural, Economic, and Political Change in 43 Societies*. Princeton, NJ: Princeton University Press.

Iwasa Y, Andreasen V and Levin S (1987) Aggregation in model ecosystems I: Perfect aggregation. *Ecological Modeling* 37: 287–302.

Jackson MO (1992) Implementation in undominated strategies: A look at bounded mechanisms. *Review of Economic Studies* 59(4): 757–75.

Jackson MO (2008) *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Jervis R (1998) *System Effects: Complexity in Political and Social Life*. Princeton, NJ: Princeton University Press.

Jordan J (1982) The competitive allocation process is informationally efficient uniquely. *Journal of Economic Theory* 28: 1–18.

Kirman A (1992) Whom or what does the representative individual represent? *Journal of Economic Perspectives* 6: 117–36.

LeBaron B (2001) Stochastic volatility as a simple generator of apparent financial power laws and long memory. *Quantitative Finance* 1: 621–31.

LeBaron B (2006) Agent-based financial markets: Matching stylized facts with style. In: Colander D (ed.) *Post Walrasian Macroeconomics: Beyond the DSGE Mode*. Cambridge: Cambridge University Press, pp. 221–35.

Ledyard J, Porter D and Rangle A (1997) Experiments testing multiobject allocation mechanisms. *Journal of Economics and Management Strategy* 6(3): 639–75.

McNeill WH (1976) *Plagues and Peoples*. Garden City, NY: Doubleday.

Mandelbrot B (1963) The variation of certain speculative prices. *Journal of Business* 36: 394–419.

Maskin E (2007) Mechanism design: How to implement social goals. *Nobel Memorial Lecture*, 8 December.

Miller J and Page SE (2008) *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton, NJ: Princeton University Press.

Mitchell M (2009) *Complexity: A Guided Tour*. Oxford: Oxford University Press.

Mittermeier Stefan (2010) *Heuristics of guile*. Research proposal.

Mount K and Reiter S (2002) *Computation and Complexity in Economic Behavior and Organization*. Cambridge: Cambridge University Press.

Myerson R (1992) Revelation Principle. In: Eatwell J, Milgate M and Newman P (eds) *New Palgrave Dictionary of Money and Finance*. New York: Stockton Press.

Nagel R (1995) Unraveling in guessing games: An experimental study. *American Economic Review* 85(5): 1313–26.

Nowak MA and May RM (1993) The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos* 3(1): 35–78.

Orrell D (2007) *The Future of Everything: The Science of Prediction*. New York: Basic Books.

Page SE (2008) Uncertainty, difficulty, and complexity. *Journal of Theoretical Politics* 20: 115–49.

Page SE (2010) *Diversity and Complexity*. Princeton, NJ: Princeton University Press.
Page SE and Tassier T (2010) A characterization of equilibria in the Groves-Ledyard mechanism. *Journal of Mathematical Economics* 46(6): 1229–42.
Reichelstein S and Reiter S (1988) Game forms with minimal strategy spaces. *Econometrica* 56: 661–92.
Reinhart CM and Rogoff K (2009) *This Time is Different: Eight Centuries of Financial Folly*. Princeton, NJ: Princeton University Press.
Reiter S (1977) Information, incentives, and performance in the (new) economid. *American Economic Review* 67(1): 226–34.
Samuelson L (2001) Analogies, adaptations, and anomalies. *Journal of Economic Theory* 97: 320–67.
Scarf H. (1977) *The computation of equilibrium prices: An exposition*. Cowles Foundation Working Paper Series, no. 473. New Haven, CT: Cowles Foundation.
Tesfatsion L (1997) How economists can get a-life. In: Arthur WB, Durlauf SN and Lane D (eds) *The Economy as a Complex Evolving System II*. Reading, MA: Addison Wesley, pp. 533–65.
Tesfatsion L (2006) Agent-based computational economics: A constructive approach to economic theory. In: Tesfatsion L and Judd KL (eds) *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Amsterdam: Elsevier, pp. 831–80.
Tetlock P (2006) *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.
Vriend N (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control* 24(1): 1–19.
Watts D (2011) *Everything is Obvious once you Know the Answer*. New York: Crown.
Wolfram S (2002) *A New Kind of Science*. Champaign, IL: Wolfram Media.
Young P (2001) *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.

About the author

Scott E. Page is the Leonid Hurwicz Collegiate Professor of Complex Systems, Political Science, and Economics at the University of Michigan-Ann Arbor. His most recent book, *Diversity and Complexity*, explores the various roles that diversity plays in complex systems.