Agent-Based Models

Scott de Marchi[1] and Scott E. Page[2]

[1]Department of Political Science, Duke University, Durham, North Carolina 27708; email: demarchi@duke.edu

[2]Center for the Study of Complex Systems, Departments of Political Science and Economics, University of Michigan, Ann Arbor, Michigan 48106; email: spage@umich.edu

Annu. Rev. Polit. Sci. 2014. 17:1–20

First published online as a Review in Advance on March 17, 2014

The *Annual Review of Political Science* is online at polisci.annualreviews.org

This article's doi: 10.1146/annurev-polisci-080812-191558

Keywords

computational, complexity, game theory, social science

Abstract

Agent-based models (ABMs) provide a methodology to explore systems of interacting, adaptive, diverse, spatially situated actors. Outcomes in ABMs can be equilibrium points, equilibrium distributions, cycles, randomness, or complex patterns; these outcomes are not directly determined by assumptions but instead emerge from the interactions of actors in the model. These behaviors may range from rational and payoff-maximizing strategies to rules that mimic heuristics identified by cognitive science. Agent-based techniques can be applied in isolation to create high-fidelity models and to explore new questions using simple constructions. They can also be used as a complement to deductive techniques. Overall, ABMs offer the potential to advance social sciences and to help us better understand our complex world.

INTRODUCTION

Agent-based models (ABMs) consist of autonomous, interacting computational objects, called agents, often situated in space and time (Holland & Miller 1991). These agents can be identical, or each can be endowed with unique attributes (Bankes 2002). The agents can be as few as one or number in the millions. Because these models rely on computations, the agents' behavior is rule based (Miller & Page 2007). Rules can be simple (e.g., backward induction on an extensive game form) or more complex (e.g., based on heuristics derived from cognitive psychology or neuroscience). The rules can even be created by social scientists competing for cash prizes (Fowler & Laver 2008).

In an ABM, agents receive inputs from their environment and take actions in response to those inputs; as such, many ABMs implement social networks or spatial relationships that factor into decision making. The aggregation of these individual behaviors instantiated in software produces system-level outcomes such as cooperation and price fluctuations that are referred to as emergent, generative (Cederman 2005), or bottom-up (Richards 2000, Epstein 2007).

The outcomes in ABMs can be temporal, ABMs of two-party electoral competition produce time series of vote totals for the parties (Kollman et al. 1992, 1997; de Marchi 1999; Laver 2005; Fowler & Smirnov 2005; Smirnov & Fowler 2007); they can be stochastic, as in the case in Bednar et al.'s (2010) model of coordination and consistency; and they can be static equilibria, such as in Schelling's racial segregation model (1978) or Axelrod's culture model (1997).

Unlike much of the modeling work done in the social sciences, ABMs often focus on the dynamics of the behavior in question—equilibria may be difficult to find or may not exist at all, and outcomes may be complex. But, even in the latter case, patterns of behavior may still exist, and agents may follow optimization/search rules that can be tested empirically (Laver & Sergenti 2011).

In the minds of many, ABMs are synonymous with complexity. This is an unfortunate conflation. ABMs are a methodology, a tool that scientists from many disciplines use to construct models (Kollman et al. 2003). Complexity refers to systems and processes that are difficult to explain, predict, evolve, or engineer (Gell-Mann 1994, Mitchell 2009, Page 2011).

By drawing this distinction, we do not mean to downplay the importance of complexity or the usefulness of ABMs in understanding it. Complex systems are of interest for many reasons, just one being their potential to produce large unexpected events and tipping points (Lamberson & Page 2012). Many scholars believe that the political world is becoming increasingly complex and that for research to advance, political science will need a deeper understanding of complexity (Cederman 1997, Jervis 1997). That may well be true. And if it is, ABMs, which are well suited to make sense of complex systems, may well grow in popularity, with the caveat that standards for good ABMs continue to be developed (Axelrod et al. 1996, Miller 1998, de Marchi & Page 2009, Jones-Rooy & Page 2012).

The uptake of agent-based modeling within political science has been steady and growing (Kollman & Page 2005). Yet, knowledge of ABMs among political scientists varies widely and misperceptions are common. As previously mentioned, many people incorrectly conflate ABMs with the field of complexity. This is understandable. Complex phenomena lie between ordered equilibrium regimes and pure randomness. Thus, complex phenomena, such as the patterns of flu outbreaks (Epstein et al. 2006), which cannot be easily explained, described, or predicted (Mitchell 2009, Page 2011), become a natural area of research to apply ABMs.

There exists a range of different approaches to modeling under the umbrella of agent-based modeling. Like game-theoretic models, agents in computational models often assume an explicit payoff function. Thus, in many cases, the distinction between ABMs and game-theoretic models is not as great as some maintain: It often comes down to how explicit one is about the assumptions implicit in one's model (Kollman 2012). There are differences, though. In game-theoretic models, agents optimize and most often the researchers settle on a set of assumptions (e.g., equilibrium refinements, stationarity) designed to produce plausible equilibria. Game theorists also place a primary emphasis on equilibrium properties rather than dynamics, and they rarely expose their models to behavioral robustness checks (i.e., testing whether or not results hold under different behavioral assumptions).[1] Game theorists do devote substantial attention to the robustness of the equilibrium to changes in parameters by characterizing the domains under which results hold.

In contrast, for agent-based modelers, equilibrium holds no special place in the mind of the modeler. Thus, agents are not constrained to use backward induction as an algorithm to select strategies. Instead, agents may have cognitively plausible limitations on memory or cognition, and algorithms may be based on research on human decision making (Kim et al. 2010). Though some see rationality assumptions as stronger than heuristic assumptions, this intuition ignores the fact that models that assume rational behavior must, by definition, be solvable. Therefore, rationality can only be assumed when the model's author has the capacity to derive what optimal behavior is.

For problems where optimal behavior has yet to be determined, neither mathematics nor computation appears to have much of an advantage. It may therefore come as some surprise to most social scientists that research on complex problems such as poker—which has long been a topic of study by noncooperative-game theorists—has most often been fruitful when deductive approaches have been combined with computational approaches (Billings et al. 1998).[2]

Advocates of agent-based approaches, among whom we count ourselves, emphasize the resonance between ABMs—their attributes and the phenomena they were capable of producing—and the real world. Many of the topics studied by political scientists—political participation, social movements, terrorism, economic regulation, cooperation, and alliances—involve interactions among diverse entities that adapt their behavior over time. Thus, political processes produce a mix of equilibria, patterns, and seeming randomness.

ABMs offer the possibility of increasing the reach of existing models and making them more general by relaxing assumptions. The potential for agent-based modeling rests on some specific capabilities, including their capacity to unpack aggregative properties (Anderson 1972), their ability to include space and networks (Newman 2003), their potential to link multiple domains (Epstein & Axtell 1996), their potential for greater agent heterogeneity (Page 2007, 2011), their production of outcome contingency, their emphasis on outcome robustness (Jen 2005), and their plasticity, which allows the construction and analysis of neighborhoods of models (de Marchi 2005).

This flexibility comes with two costs: a lack of analytic tractability and an expansion of the parameter space of the model (de Marchi 2005). Although these are serious concerns, they result

[1]There is no natural way and therefore multiple ways to define an equivalence class for games, which complicates robustness checks. Thus, changing the assumptions of a game simply produces a different game with different results. The rightly famous Arrow result is important because the assumptions are of substantive interest and the result is very general. This is rare—typically, the assumptions are not of interest and lack verisimilitude, and the conclusions are overly responsive to these assumptions.

[2]In games with the order of complexity of poker or chess, one cannot solve the extensive form directly for a Nash equilibrium. In chess, researchers have derived "idiosyncratic utility functions" that allow for the evaluation of the game prior to the terminal nodes (see de Marchi 2005 for an overview). In poker, researchers limit the size of the extensive form in two ways. First, they add constraints (e.g., they limit the number of betting opportunities or the range of bet size), and second, they compute approximations to the Nash equilibrium [e.g., see the work on regret minimization in Zinkevich et al. (2007)]. In the social sciences, bargaining theory has also benefitted from a combination of noncooperative-game theory and computational approaches (see Ansolabehere et al. 2005). In this work on bargaining, a computational model of the power of political actors was necessary to make progress in applying bargaining models to the n-player case of proportional representation government formation.

from the fact that ABMs must be explicit about the choices that are made regarding timing, interactions, and adaptation. To be concrete, let us take bargaining models as an example. As Laver et al. (2011) point out, noncooperative models of bargaining in the Rubinstein and Baron & Ferejohn (1989) tradition assume that agents adhere to stationarity and that the reversion point to bargaining is for all players to receive a payoff of zero; for a strict sequence to be observed in offers rather than an open-outcry assumption[3]; and for an exogenous, uniform, and independent selection mechanism to exist that chooses the proposer in each time period. Whereas the result of models in the noncooperative work on bargaining is purely deductive, it is not the case that these are trivial assumptions or even general assumptions that are of themselves substantively interesting.[4] Nor is it the case that the equilibrium results are robust to changes in these assumptions.

We have organized this review article as follows. We begin with a brief unpacking of the specific capabilities of ABMs. We then provide a basic primer on how to construct an ABM. More complete sources exist (Epstein & Axtell 1996, de Marchi 2005, Miller & Page 2007, Mitchell 2009). We intend for this to be a self-contained introduction that provides an entry to the literature. To show the relative strengths and weaknesses of ABMs, we then compare and contrast them with traditional deductive techniques. We claim that political science would improve if modelers combined the two techniques but that ABMs also have stand-alone value. We conclude with a discussion of empirical testing of ABMs and some general comments on the use of models in social science.

CAPABILITIES OF AGENT-BASED MODELS

In this section we highlight specific capabilities of ABMs and their relevance to social scientific inquiry. Given the diversity of computational and agent-based models, we focus on general tendencies. At one end of the spectrum computational work resembles Monte Carlo experiments in applied statistics; on the other end, researchers create large numbers of sophisticated, adaptive agents and examine the resultant dynamics of the system.

First, and at the most basic level, ABMs allow researchers to link behavioral rules to aggregate patterns, be they equilibria or complex patterns. For this reason, they are commonly referred to as bottom-up. For example, Laver (2005) explores the dynamics produced by electoral competition and shows how those dynamics differ depending on the behavioral rules of the parties. The idea that different behaviors should produce different patterns makes intuitive sense, especially if one accounts for the evidence that in many contexts, agents are involved in optimization problems (i.e., search) rather than simply solving for equilibrium. That a particular set of learning rules produces results that match actual electoral outcomes in Ireland is a strong argument for why one should use computational approaches. ABMs allow researchers to construct this micro-to-macro mapping and to test the empirical validity of the results produced (Fowler & Laver 2008).

In many cases, computational models that assume psychologically based behavioral rules may better fit experimental and empirical data than do rational-actor models (Camerer 2003, Lodge 2004). ABMs can encode these behaviors and help us understand how they aggregate. This capacity to encode any behavioral rule leads some to see ABMs as opposed to the notion of rational-choice theory and as wedded to the behavioral approach. This is a misrepresentation. ABMs rely on behavioral rules. Those rules could be best-response functions given some utility function and

[3] In open-outcry models, every actor may speak at any point and there is no strict sequence of play that distinguishes between the proposer and the passive respondents to offers. There is also no need for an exogenous selection mechanism for a proposer in each round.

[4] This is in contrast to the Arrow result, which rests on assumptions that are of vital substantive interest and that are extremely general.

the agents' information sets. Though often the rules are heuristics, nothing precludes rational behavior, and the problem itself generally pushes the modeler to adopt a particular set of behavioral rules. Verisimilitude is most often not the goal in and of itself; rather, researchers resort to computational models when purely deductive models produce poor outcomes that are useful in generating results on substantive problems (Tesfatsion 1997, Leombruni & Richiardi 2005).

The literature on turnout is one prominent example. Bendor et al. (2003) built a computational model where they demonstrate how a learning rule with more verisimilitude than purely deductive models accounts for positive turnout; Fowler (2006b) improves on this with a learning rule that allows for habituation If the goal is to study turnout as it exists in real electoral systems, the resort to computational models makes perfect sense—it allows researchers to examine different (and more interesting) learning rules, inspect the resultant dynamics, and use these models to explain real-world data or to create similar patterns (Palmer et al. 1994).

Second, ABMs enable the inclusion of geographic and social space at multiple levels of resolution. In politics, work in applied statistical models has shown that space matters (Hoff & Ward 2004). People segregate by race, income, religion, and other characteristics, and (evidence suggests) individuals' behavior depends partly on that of people around them (Glaeser et al. 1996). Further, political representation is often geographically based. These three stylized facts combine with evidence from rich empirical case studies (Huckfeldt et al. 2004) to suggest that even if it is not true that all politics are local, what goes on spatially probably does matter sufficiently for it to be included in our models.

Relatedly, ABMs can also include networks (Newman 2003). These could be policy networks (Carpenter et al. 2004), interaction networks (Lazer 2001), or friendship networks. Given the growth of network measures and statistical techniques, testing network effects has become much easier (Fowler 2006a,c).

Third, ABMs have the capacity to include multiple processes simultaneously. To understand the nature of this contribution, we start with the more familiar distinction between partial and general equilibrium analysis. A partial equilibrium analysis might consider only the health-care market following the passage of a mandatory health-care bill. A general equilibrium analysis would consider how health-care payments affect markets for houses, cars, and exercise equipment. In an ABM, general equilibrium spillover effects can extend to the behavioral realm and cross domains. ABMs with these features have been used extensively for research into policy-resistant problems by the National Institutes of Health (under the Systems Science request for applications) and the Department of Defense.

Given that agents learn from one another and interact in multiple domains, they often produce behavioral spillovers. Behavioral spillovers can be defined as behavior in one domain that is influenced by behavior in another domain. These behavioral spillovers can produce aggregate behavior that differs from what would occur were each interaction considered in isolation (Bednar & Page 2007). Agents who are cooperative or trusting in one interaction may be more likely to be cooperative or trusting in another, a finding in experiments with human subjects (Bednar et al. 2012).

Cross-domain spillovers can be thought of as multiple general-equilibria effects. An ABM can include general-equilibrium models of politics, economics, disease, and environmental quality in the same instantiation. Why might this matter? Political and economic decisions affect global temperatures, and global temperatures, in turn, affect epidemics. ABMs can combine all these models incorporating feedbacks from one domain to others. Or, to give a slightly different example, networks form for many reasons: economic, political, social, etc. An economics-only or politics-only model might not be capable of capturing dynamics that resembles real-world, multipurpose networks.

We should be clear that there are limits to realism. If a model includes too many domains, it may become difficult to understand and has little value. Yet, this recognition of limits on realism does not imply that models that link domains cannot produce meaningful insights (Lustick & Miodownik 2000, Lustick et al. 2004). One could argue, for example, that a model that combined transportation systems with demographic data or that linked bargaining in coalition governments with electoral concerns would produce insights on turnout. There is thus a tension between limiting a domain so that one can derive tractable models and omitting relevant dependencies and behaviors. Choosing the right point on this spectrum should be a matter of the power of the results in explaining real-world behavior and not a matter of modeling aesthetics.

Fourth, ABMs can include heterogeneity not only in location, beliefs, information, preferences, and ability, but also in learning rules, perspectives, mental models, behavioral repertoires, and cognitive framing. Moreover, this heterogeneity can be matched to actual distributions for parameters of interest. Income distributions, life expectancies, and numbers of friends can be calibrated based on data. Heterogeneity for realism's sake is nice but is not the end in itself. The key questions that need to be asked are, When does heterogeneity cancel out, so that only the mean matters, and when do the actual distributions matter? ABMs can help to answer those questions (Bruch & Mare 2006, Page 2011).

Fifth, as already mentioned, ABMs prove capable of producing messy contingent outcomes and a range of phenomena: randomness, equilibria both static and distributional, patterns and complexity, and the complex outcomes can vary in their amount and type of path dependence (Page 2006). Contrast this with game-theoretic models, which for the most part derive equilibrium (though they need not!).

Equilibrium is a nice benchmark, but what if the system does not equilibrate? Or what if time to equilibrium is so long that the equilibrium is not of much scientific interest? In such instances, equilibrium analysis may be misguided, and ABMs can add substantial value (Laver & Sergenti 2011).

Sixth, analysis of ABMs tends to focus on outcome robustness, i.e., whether the system can keep doing what it needs to do (Jen 2005) as opposed to the existence and efficiency of equilibria. For example, ABMs can be used to assess the robustness to the entrance of new types, the exit of existing types, changes in environmental parameters, and even changes to the rules that agents follow (Laver & Sergenti 2011).

Bednar (2009) argues that robustness should be central to our consideration of political institutions. Recent events in financial markets demonstrate that the costs of the occasional large failure can far outweigh small efficiency gains. When considering allowing certain types of trades, efficiency considerations should be balanced against concerns for system-level robustness, which can in turn depend on the difficulty of the learning environment (Andreoni & Miller 1995).

Finally, ABMs allow for exploring neighborhoods of models (de Marchi 2005, de Marchi & Page 2009). Any model makes a set of assumptions. Those assumptions can be altered to create a set of neighboring models. Consider, as an example, Axelrod's (1997) culture model. Neighboring models can assume different behaviors on the part of agents or different network topologies and can even place the attributes in a network structure. Depending on one's imagination and the dimensionality of the model's assumptions, the set of neighboring models can number from the dozens to the hundreds of thousands.

Of course, when the set of neighboring models is small, it is possible to consider each alternative. Such will be the case when considering different voting rules, where one might compare three possibilities. Larger sets of possible models can be sampled and analyzed with the tools of applied statistics. By exploring the neighboring models, the researcher can gain a deeper understanding of the core model.

AGENT-BASED MODELS: THE BASICS

We now describe the building blocks of ABMs. Our description consists of three parts. We begin with a discussion of model primitives focusing on agent attributes and behaviors. We then turn briefly to how these models are constructed in practice and distinguish between simple intuition engines and more complicated high-fidelity ABMs. We conclude with a brief description of the types of outcomes that ABMs produce, starting with a discussion of equilibrium outcomes and then turning to other types of outcomes less familiar to social scientists.[5]

Agent Attributes and Behaviors

ABMs begin with a set of objects referred to as agents. Agents can include human actors, governments, media, and even constitutions. One agent might be a politician. Another agent might be a voter. The agents are similar to the players in a game-theoretic model. Each agent is characterized by a vector of attributes and behaviors, and these are often partitioned into fixed and mutable categories. For example, a voter might possess a fixed ideology and a mutable preferred party. Attributes are often the same as would be found in a game-theoretic model. These include beliefs, actions, payoffs, and so on. In ABMs, attributes also often include the location of the agent. We can formally define the set of attributes as follows:

$$\{X_1, X_2, \ldots X_M\}, \text{ where each } X_i \text{ is a countable set, i.e., } X_i = \{x_{i1}, x_{i2}, x_{i3}, \ldots\}.$$

To instantiate a model, we assign a value to the agent for each attribute. Attributes can be set valued, which allows an agent to have the empty set as a value if the attribute does not apply to the agent. For example, an attribute could represent gender; voters would be assigned genders and political parties would not.

Agents may change their attributes, and we refer to the attribute values at a moment in time as the agent's state. To keep track of time, we introduce an index t so that we can denote the state of an agent at time t as well as the configuration of the model itself. In what follows, we assume a fixed set of N agents. This framework can be extended to allow for changing numbers of agents.

$$\text{The state of an agent } j \text{ at time } t: \mathbf{a}_j^t = (x_{j1}^t, x_{j2}^t, \ldots x_{jM}^t), \text{ with } x_{ji}^t \in X_i.$$

The collection of all of the agents' states can be referred to as a configuration.

$$\text{The configuration of the model at time } t \text{ is } A^t = \{\mathbf{a}_1^t, \mathbf{a}_2^t, \ldots \mathbf{a}_N^t\}(x_{j1}, x_{j2}, \ldots x_{jM}), \text{ with } x_{ji} \in X_i.$$

Typically, an ABM will have between 2 and 10,000 agents depending on what is being modeled. Models of spatial party competition in the United States may have as few as two agents if only the major political parties are considered and many more agents if voters are instantiated in the model. Models of riots may have thousands of agents. Though ABMs with millions of agents can be built, they demand substantial computational time. Models with thousands and even millions of agents have grown in popularity, given advances in data availability. Notably, Geanakoplos et al. (2012) have constructed a model of the Washington, DC, housing market that includes every household. Their model shows the importance of leverage ratios to boom and bust cycles.

As stark as this setup may seem, it offers enormous flexibility. Attributes can include tags, such as identity or party identification, and geographic or network locations. Many ABMs assume

[5]This section focuses on the concepts but not the particular tools involved. There is a wide range of software packages used by researchers, and choosing the best set of tools is a nontrivial enterprise. A good starting point, however, is Python, with its extensions to call R functions, and the open-source scientific library SciPy .

some locality to interactions. For example, in a spatial collective-action model, agents might be placed on a checkerboard and only interact with those neighbors in adjacent squares. Though the terminology is cumbersome, we can refer to such restrictions as state-mediated contact structures.

ABMs in which agents interact with geographic neighbors or over a network as well as models in which agents interact only with agents of their same or opposite type are all examples of models with state-mediated contact structures. In each case, rather than each agent interacting with all other agents or agents interacting randomly (this is formally known as random mixing), agents interact with others contingent on their states. When implemented in computer code, whether or not two agents interact depends on the vector of attribute values, which could include location, social standing, gender, or ideology.

Thus, from a computational perspective, constructing agents with geographic locations who interact locally is not dissimilar to constructing agents with spatially represented ideologies who interact only with like-minded individuals. In Axelrod's (1997) model of cultural boundary formation, agents change their ideology to match that of geographic neighbors. This model can be seen entirely as a model of geographic segregation (with the caveat that some dimensions are fixed and others are mutable). We might add that taking a geographic perspective makes his findings more intuitive.

We now turn to agent behaviors. These behaviors take the form of rules that tell the agents what to do. Simple behavioral rules may depend only on the current configuration. More sophisticated rules may depend on past configurations as well and may even take into account the potential actions of other agents. This would be the case for strategically minded candidates trying to locate winning platforms. Most often, ABMs rely on standard learning rules borrowed from other disciplines, such as game theory's fictitious play (Brown 1951), mathematical psychology's best-response dynamics and Hebbian learning (Bendor et al. 2003), or cognitive psychology's heuristic shortcuts (Kim et al. 2010).

It is common to refer to these systems of rules and the rules themselves as adaptive because the agents' behavior often adapts to the current state of the configuration. Intuitively, there would seem to be a distinction between fixed rules like grim trigger in the prisoners' dilemma game and adaptive rules in which the agent learns how to play a game. However, even adaptive rules are instantiated with a fixed algorithm that is open to inspection and replication (Miller & Page 2007).

A relevant distinction does exist between individual-based learning rules and population-based learning rules. The latter are referred to as selection processes if the learning involves killing off poorer-performing agents and replacing them with copies of more successful agents. Individual-based learning rules, such as best-response learning, are more likely to locate Nash equilibria (Vriend 2000). The intuition behind this general finding is straightforward. If agents adapt their actions as individuals, no equilibrium that is not a Nash equilibrium would be sustainable: Some agent would find a superior strategy with positive probability. The same need not be true of models in which agents follow a population-based learning rule. In population-based rules, reproduction depends on relative performance. For example, under replicator dynamics, agents copy more successful neighbors. Therefore, the agents can lock into suboptimal behaviors.

The final piece of an ABM is the implementation of the interactions. Once the agent's attributes and behaviors have been defined, the model must fully describe the timing of those interactions. This is more complicated than it might seem. In early models, most of the rules were applied simultaneously. But this becomes problematic once agents can change location. A particular problem is how to prevent multiple agents from moving to the same location (Gaylord & D'Andria 1998). When sequential timing fails or does not make sense given the context, ABMs execute the agents' rules sequentially. This ordering can be random, based on tags or geography, or endogenous, with those agents who most benefit by calling their rules acting first (Page 1997).

Intuition Engines and High-Fidelity Models

ABMs can range from extremely simple models of binary agents strung on one-dimensional arrays to models of millions of agents placed in real geographic space. These two types of models serve very different scientific purposes. The simple models are often used to generate intuition and to elucidate core principles. The high-fidelity models are most often used to craft policy, run counterfactuals, and function as calibration exercises.

We describe, in brief, a simple model and then a high-fidelity model to give a hint of the range of models that can be constructed. We also highlight the distinct scientific contributions of these two models. The first model is a spatial prisoners' dilemma model (Nowak et al. 1994). In this model, the players are placed on a two-dimensional lattice with N rows and N columns. The top of the lattice is connected to the bottom and the right edge is connected to the left edge to create a torus. Each agent plays a one-shot prisoners' dilemma game against its four neighbors to the north, south, east, and west. In the prisoners' dilemma, the optimal action is to defect (D), but your opponent does better if you choose to cooperate (C). The socially optimal outcome is for both players to choose to cooperate.

The attributes of the agents would be their locations and their actions. The behavioral rule could be anything. Nowak et al. (1994) assume the following behavior: Each agent plays one-shot prisoners' dilemma games with its four neighbors. The agent's fitness is the sum of its payoffs across those four games. In the next period, agents copy the action of one of their four neighbors if and only if that neighbor had a higher fitness.

Given the spatial nature of play, even though defecting is always a better action, cooperative players can get higher payoffs. For example, consider the canonical payoffs where both cooperating pays R, both defecting pays $P < R$, and a defector playing a cooperator earns $T > R$, while the cooperator gets $S < P$. A cooperator who plays three cooperators and one defector earns a payoff of $3R + S$. A defector who plays three defectors and one cooperator earns $3P + T$. For many values of $T > R > P > S$, such as $T = 5$, $R = 3$, $P = 1$, $S = 0$, the cooperator will do better and cooperation will spread in the population. A computational analysis of this spatial model shows that for many sets of parameter values it produces regions of cooperators and defectors with chaotic boundaries (see Nowak & May 1992).

This simple model produces an important insight about the emergence of cooperation. Sustaining cooperation in a prisoners' dilemma had been thought to require punishment, or the threat thereof. Punishments require repeated interactions. This model, however, produces clusters of cooperation through spatial interactions. It provides a new mechanism for cooperation—local spatial heterogeneity combined with social learning. This provides a partial answer to why researchers like Ostrom (1990) find that cooperation (i.e., solving common-pool resource problems) exists in some locations but not in others, even when these locations are pursuing identical strategies. For example, Ostrom found that the ability of Swiss officials to levy fines and keep half of these fines for themselves produced effective cooperation; yet, similar tax schemes existed throughout Europe in the same time period and failed dramatically given the obvious incentives for the official.

ABMs can also be very realistic. They can include realistic depictions of physical space and numbers of agents that approximate the real world. These high-fidelity ABMs are primarily used in business and policy applications. They have proven useful in the study of epidemics. Until recently, epidemic modeling was dominated by formal models. For reasons of tractability, the models tended to assume random mixing of people and not to include real geographic space.

ABMs can include houses, schools, workplaces, and transportation systems. They can assign each person a state that includes not only whether or not the person has the disease but also a susceptibility, virulence, and duration. The only behavior that is relevant for the spread of a disease is where people travel and with whom they come into contact. Those behaviors, including key features such as social networks, can be approximated from data. In this way, ABMs can be used to predict how fast a disease will spread or to infer the virulence of a disease that has begun to spread. They can also be used to evaluate intervention scenarios. It should come as no surprise that these high-fidelity ABMs of epidemics have become the status quo for researchers (Epstein et al. 2006).

As mentioned in the introduction, this potential representational richness is both a blessing and a curse. On one hand, if details have been left out of mathematical models solely for reasons of tractability, those details can be included and the result may well be an improved model. On the other hand, agent-based modelers can become obsessed with producing accurate agent-level detail at the expense of producing a model that has any value.

Outcomes in Agent-Based Models

ABMs generate equilibria from the bottom up (Epstein 2007). It is important to draw bright lines between generation, existence, and stability. An equilibrium exists if the assumptions about behavior hold and the configuration remains fixed. Nash equilibrium, as well as its many refinements, assumes that the behavior is optimizing given a payoff function. An equilibrium can exist but not be stable. For example, in a pure coordination game, each player can take one of two actions. If both choose the same action, each receives a positive payoff. If they choose opposite actions, each receives a payoff of zero. The pure coordination game has three equilibria: two pure-action equilibria, where both players choose the same action, and one mixed-strategy equilibrium in which both play each action with equal probability.

The mixed-strategy equilibrium is often informally described as unstable. Technically, stability requires an assumption about behavior. If the players best respond, then it is true that the mixed strategy is unstable. If either player chooses one of the two actions a little more often than the other, then the opposing player would best respond by always choosing the more likely action. However, if the behavior is to always return to pure randomization, then the mixed strategy will be stable and any deviation will be corrected. Thus, stability is a function of the behavioral rule.

In an ABM, the starting points are often random, or in the case of high-fidelity models, tuned to reality to the extent possible. If we keep track of where the model ends up from each initial configuration, assuming equilibria exist, then we can map out basins of attraction. The basin of attraction for an equilibrium is the set of initial configurations that eventually lead to that equilibrium. For the purposes of this section, we will assume that the dynamics are deterministic, so that if the model starts in the same configuration it always goes to the same equilibrium.

Stability properties often do not depend on the choice of behavior so long as agents take improving actions. Agents take actions in the next period that would have fared better than their current action in this period. That is not true of the mixed strategy described above for coordination games. Best response, replicator dynamics, and other learning rules enforce improving actions, and so most equilibria that are stable under best response are stable under replicator dynamics (see Golman & Page 2009, 2010).

The same cannot be said for the basins of attraction. Replicator dynamics and best-response behaviors can have basins of attraction with no overlap (Golman & Page 2009). This means that starting from the same point, they will always lead to different equilibria. This occurs even though the two learning rules have the same set of stable equilibria. It is just that the learning rules drive the population in different directions. This intuitive finding—that agent behavior affects where the system ends up—provides a primary reason for constructing ABMs if one cares about generating empirically testable results.

One final point on equilibrium in ABMs: Many of the equilibria produced are statistical equilibria or what we have called equilibrium distributions. By that we mean that the ABM keeps churning but that on average the behavior is the same. Imagine a simple model in which agents choose to grow either apples or pears. Each period they take their goods to a bazaar to trade. If apple trees and pear trees produce fruit at equal rates but with a stochastic component and if agent preferences weight the two goods equally, a statistical equilibrium will emerge in which agents produce approximately equal numbers of apples and pears. This statistical equilibrium in which the average price paid for a pear equals the average price paid for an apple differs markedly from a static equilibrium in which each agent pays the same price for each good (Epstein & Axtell 1996).

COMPLEXITY AND EMERGENCE

We now turn to the potential for ABMs to produce nonequilibrium outcomes and complex phenomena. These outcomes often include self-organization or emergence. Self-organization refers to the creation of some pattern or form from the bottom up. Here, the best examples occur outside of the social realm. The formation of crystals and the flocking of birds are the classic examples of self-organized behavior. Crystals do not have party bosses telling them how to line up. Their symmetric arrangement emerges from physical forces. The term "emergence" applies to more general structured phenomena that include higher-order functionalities. For example, an efficient allocation of goods emerges in an ideal market even though agents pursue individual interests.

ABMs of exchange often produce complex outcomes. Kirman (1997) and Vriend (2000) get rid of the central market maker and allow individual buyers and sellers to form relationships with one another. With this added realism, customers exhibit diverse behaviors. Those with higher values for the goods show less loyalty than those buyers with low values. Interactions between agents produce a complex web of connections and reputations.

ABMs are often used to study event-sized distributions. When aggregate variables equal the sum of individual variables, then under mild conditions the aggregate variable will be normally distributed. But when individual attributes' values are interdependent or when there exist positive feedbacks, the aggregate variables can take on more extreme distributions. Given that ABMs are well suited for including interdependencies and behavioral feedbacks, they are often used to study environments that produce nonnormally distributed outcomes.

The most common method for analyzing aggregate outcomes is to plot the distribution of event sizes and to compare it to a normal distribution. For example, when one builds an ABM of the stock market, one can plot the daily price fluctuations and find evidence of bubbles and crashes that are not consistent with normally distributed returns (LeBaron 2001). In other cases, such as models of traffic flows, the entities in the model can self-organize into a critical state, where small events can have large impacts. Thus, in those critical states, the event sizes tend to have long-tailed, and even power-law, distributions (Bak 1996).

AGENT-BASED AND MATHEMATICAL MODELING: A COMPARISON

In our appraisal of ABMs, we do not mean to suggest that they should replace deductive mathematical modeling. Nor do we believe it to only have value as a complement. The practice of ABMs differs in spirit from mathematical modeling in four important ways. First, in the construction of an ABM, analytic tractability is neither a primary nor a background concern. Instead, the focus is on identifying the most relevant parts of the domain and on capturing how those features interact. Tractability does not matter to the agent-based modeler because almost anything can be implemented with good computer science skills. Although simplicity and understanding one's own model are important, it is recognized that complexity is often worth the cost, particularly when deductive models fail to predict phenomena of interest.

Second, researchers in the game-theoretic and computational models differ, but not as much as one might think. In creating a formal model, the researcher must choose what to explore and what to prove. Is there an equilibrium? Is there a symmetric equilibrium? Is the equilibrium efficient? What type of refinement should be employed? With an ABM, the researcher inspects the results with visualization tools and applied statistics. If a symmetric equilibrium does not emerge, the researchers tend not to focus on this fact or change the assumptions to produce an equilibrium. Instead, the modeler's attention is naturally drawn to the dynamics that do occur. These differences are important, but the way in which agents are characterized and models are instantiated is not a substantially different intellectual process—and many examples exist where computer science has overlapped with game theory (Conitzer & Sandholm 2006).

Third, ABMs, like deductive models, often begin with preferences and their associated utility functions, but as mentioned, they can also begin with sets of rules that evolve over time. In many cases, given a rule, it is possible to infer a utility function for which the rule represents a best response, in which case the distinction between optimizing and rule following reduces to semantics. However, in many cases, no simple utility function rationalizes the rule. In those cases, rule-based behavior and utility-driven behavior must be considered distinct and it becomes incumbent on the modeler to convince the reader of either the empirical plausibility or the methodological benefit of using that rule.[6]

Fourth, as mentioned earlier, even though ABMs and mathematical models are capable of studying the same phenomena, practitioners tend to use them to address different types of phenomena. Thus, if a scholar believes a system is in equilibrium, she may opt for a game-theoretic model. If someone seeks to understand market crashes, he may rely on computational models. Often ABMs are used when the outcome of a model cannot be pinned down analytically. If an analytic approach to solving a model is not possible or practical, an ABM can give insights into what happens.

The flexibility of ABMs explains why they have more detail and, as a result, are more likely to produce nonequilibrium outcomes. ABMs tend to include geographic space, networks, adaptive learning rules, and heterogeneity so that they produce different types of results. Consider a simple model of coordination in which players choose either action A or action B. The standard game-theoretic model would assume random mixing of agents. The results of such a model without noise would be convergence to an equilibrium, either all players choosing A or all choosing B. An ABM of the same coordination game would most likely situate the agents in space, perhaps on a grid or in a social network. The result of this model would likely be clusters of agents playing A and clusters of agents playing B. Note that the agent-based modeler could have written a model with random mixing and if so would have found global convergence. And the mathematical modeler could have written a model on a network and (perhaps with a lot of work) shown how the size and shape of clusters depend on network characteristics. Our point is that this typically does not happen. It usually does not happen because networks can be difficult to analyze analytically, though advances are being made Jackson (2008). Increasingly, papers that analyze networks combine mathematical and computational approaches.

[6]It is worth noting that even purely deductive approaches face related issues. For example, equilibrium refinements, such as trembling-hand equilibrium, are fundamentally behavior assumptions that may or may not be appropriate given the context to which a model is applied. There is no way to justify the use of these refinements based on deductive math; they are essentially statements about psychology and thus are not different in kind from the rules utilized in computational models.

A natural question to ask is when to use an ABM and when to use a game-theoretic model. Experience suggests that in almost all cases a combination of the two produces the deepest insights. We make that claim for two reasons. First, the analytic approach is best for demonstrating existence of equilibria. The agent-based approach tests whether a particular physics of the system (as captured by the agent's rules) leads to an equilibrium or whether cycles or complexity results. But this is the first step in understanding real-world phenomena, and a natural question to ask is whether or not the result of one particular set of assumptions is generalizable.

Second, what counts as a natural assumption in a mathematical framework may not be natural in a computational setting. Game-theoretic models rely on an uncertainty framework, whereas ABMs assume complex environments (Page 2008). To give one specific example, mathematical models of incomplete information rely on generated signals, i.e., they draw from probability distributions. Computational models often employ interpreted signals, i.e., they endow agents with limited perceptions of the whole.

Consider a standard spatial-voting model. Assuming each voter sees the spatial position of the candidates with a small error produces generated signals. Assuming each voter samples only a subset of the issue domains produces interpreted signals. The former assumption is more typical of mathematical models, and the latter is more typical of computational models. We do not claim that either assumption is better than the other at capturing voter behavior. We merely note that each assumption is made with an eye toward its tractability given a particular methodology. Error distributions lend themselves to mathematical analysis. Random subsets are easily generated computationally. Surprisingly, these two approaches lead to quite distinct correlational properties of signals (Hong & Page 2009). Therefore, what is true in a mathematical model may not be true in a computational analog of that model owing to differences in how signals are formalized.

As should be clear, we advocate using both mathematics and computation. We also argue that as a general rule, initial versions of models should not include an abundance of detail. Therefore, even if the context involves diverse agents interacting on a network and the modeler knows that deriving an equilibrium is not likely, starting with a game-theoretic model often makes the most sense. Why? Game theory imposes rigor and structure by implying a set of questions: Who are the agents? What are their interests? What is the space of possible actions? What is the timing of events? What information do agents have?

Once the game is set up and proves intractable, there are two paths to take. The game can be simplified down to the point it is tractable. Or, the game can be converted into an ABM and one can push for more general results or test whether or not the assumptions involved are robust to changes. We have made this sound a little too simple. In constructing the ABM, there remains the question of agent-level behavior. In game theory, the agents choose optimally. This is true in game-theoretic models of institutions as well (Diermeier & Krehbiel 2003).

In an ABM, any one of a number of learning rules—either individual or population based—might be used. Many first-time users of ABMs use best-response rules. This rule, though convenient and consistent with game-theoretic thinking, has two problems. First, it often produces overshooting, or what might be more illustratively called emergent ping-pong. As an example, imagine a simple one-dimensional spatial model of elections with the median voter at position zero on the real line. Randomly select a first candidate to be placed at negative five, and place the other at positive five. The best response for the first candidate will be to move from five to five minus ε, and the best response for the other candidate will be to move from minus five to minus five plus ε. Subsequent iterations will lead to more back-and-forth. A more gradual behavioral rule would enable the candidates to smoothly converge to the median.

Second, best-response calculations often chew up a lot of clock time. In an analytic model, best responses can be found by taking derivatives. In an ABM, they may require exhaustive search of the set of possible actions, which is at odds with the spirit of many computational models that attempt to instantiate agents with cognitively plausible heuristics (Epstein 2014).

In brief, the richer the environment under study (either because the environment is complex or because it involves embedded games), the more likely it is that an ABM will be necessary. This is not to say that ABMs should replace games. To the contrary, the findings from the richer ABMs should be compared and contrasted with findings from game-theoretic models. Agent-based versions of game-theoretic models should also be thrown into the mix given that they enable an analysis of attainability and stability of equilibria. This is where real learning will occur. Lest the reader think we are piling on the work, ABMs are easy to write, analysis is often visual, and data can be produced in any amount. They are an econometrician's dream.

These labels—bottom-up, generative—contrast the inductive agent-based approach with the deductive game-theoretic approach to modeling in which existence results are derived by the model (Judd 1997, Epstein 2007). Consider a model both analyzed using a game-theoretic model and implemented in an ABM (Bednar & Page 2007). The game-theoretic model can be used to derive conditions under which an equilibrium exists and to identify necessary and sufficient conditions for an equilibrium to be efficient. More sophisticated analytic techniques might be used to show that an equilibrium is unique or to show the unique outcome that satisfies some refinement criteria.

The ABM can be used to investigate whether a class of learning rules generates the derived equilibrium. If, for example, in each of 5,000 runs of the model the agents reach equilibrium or something near it, then we can say with statistical confidence that the assumed model generates outcomes in some neighborhoods of the equilibrium.[7] If we can also show that this finding holds for other classes of learning rules as well, then we can make some claim to the robustness of the result. Notice the distinct contributions of the two approaches. The game theorist proves that an equilibrium exists and derives properties of that equilibrium. The agent-based modeler demonstrates whether that equilibrium can be generated by rule-following agents.

EMPIRICAL TESTING OF AGENT-BASED MODELS

The parameter spaces in many ABMs are quite large compared with conventional models in the social sciences, and as a result many researchers are accustomed to examining the effects of different parameters using the tools of applied statistics. This experience has led to less of a gulf between many computational researchers and statistical tests of their models. As a result, there is an increasing emphasis on testing computational models, both because one cares about the power of models to predict real phenomena and because empirical work allows computational researchers to choose parameter values based on data rather than treating them as nuisance parameters.

In bargaining, Laver et al. (2011) and Cutler et al. (2013) have used a computational model to demonstrate that parties in proportional representation systems rely on minimum integer weights to determine membership in winning coalitions and raw weights to assign cabinet seats once membership is secured. Golder et al. (2012) built a computational model of government formation that produces outcomes that match real-world data, including surplus coalitions, minority governments, and minimum winning coalitions. In studies of legislatures, Fowler (2006a,c) has demonstrated how cosponsorship networks between legislators affect voting behavior. Laver & Sergenti (2011) have provided models that describe party birth and death and evolving policy locations.

[7]This is a simplification—see Laver & Sergenti (2011).

In international relations, Weidmann & Salehyan (2013) built an ABM instantiated with GIS-coded data and showed that sorting is the most plausible reason for the reduction in violence in Iraq. Gartzke & Weisiger (2013) built an ABM to show how the prevalence of different identities in the interstate system creates new cleavages for conflict between nations.

Although there are other examples of applications of ABMs, let us conclude this section with one in-depth treatment that will provide some of the flavor of how computational models are tested. In elections, perhaps the most studied area in computational modeling to date, the original Kollman et al. (1992) model has been shown to have empirical support: There is an incumbent advantage, and it depends on the complexity of the electoral landscape.[8] This work is important because it shifts the frame for elections from equilibrium-seeking behavior, where there is an abundance of chaos results (e.g., the general nonexistence of a median in all directions, the expectation of low or zero turnout, and the expectation that incumbents will always lose reelection), to a model of optimizing, incompletely informed parties who must depend on their prior successes and polling to search for good platforms.

Testing Kollman et al.'s (1992) model directly is difficult, other than the obvious: Incumbents, even if one controls for challenger quality and resources, do not always lose reelection. But is there a more direct test of computational modeling and the emphasis on optimization being appropriate rather than equilibrium-seeking behavior?

Fortunately, the answer is yes. Ensley et al. (2009) and Wichowsky (2012) have provided results for the American case, and both have shown that measures of district complexity (i.e., how difficult an electoral landscape is to search) explain why incumbents win or lose. In the case where landscapes are complex—there are multiple salient issues and a lack of broad agreement in the electorate—incumbents perform well in their reelection campaigns. Incumbents not only have learned (optimized) from prior campaigns but also have (perhaps randomly) arrived at an effective platform, as signified by the fact that they have won at least once (see Alchian 1950 for similar analysis of the firm).

Laver (2005) has provided additional support for Kollman et al.'s (1992) theory with his research on the Irish case. In this work, Laver pinpoints the search algorithms used by the various political parties in that system; in effect, this matches hypothesized algorithms from computer science (e.g., hill-climbing is one algorithm) with human decision making in complex environments.[9] Laver's work echoes that of Kollman et al. (1992), who considered alternative heuristics for locating policy positions. They compared hill-climbing candidates, candidates who chose positions randomly, and candidates who used sophisticated search strategies based on genetic algorithms, finding only modest differences between the three.

DISCUSSION

In sum, ABMs have much to offer for the study of politics. They enable us to construct richer behaviors and environments with greater dimensions including networks, coalition seeking

[8]Visually, complex landscapes are jagged and have many optima. An ant with limited vision looking for the highest point (assuming height = more utility) would have a tough time of it. Simple landscapes are smooth and have few local optima, and one can more easily follow a gradient in them to a good outcome. This idea of fitness landscapes is derived from work in optimization theory and population genetics (see Kauffman 1993).

[9]Laver's analysis is instructive of how testing ABMs differs from testing of game-theoretic models. Agent-based models vary in the assumptions they make about behavior. Therefore, it is natural to test the assumptions of the model or to run a horse race between different sets of assumptions about behavior. This opportunity to test both assumptions and outcomes is both a blessing and a curse.

(Axelrod et al. 1995), and learning. They permit us to enlarge the set of questions and to consider alternative types of outcomes, including complexity. At the same time, ABMs can become unwieldy and less useful if they include too much detail. Hence, some balance must be reached that includes just enough, but not too much, of the real world.

ABMs can help us make better sense of complex systems and how institutions and individuals cope with that complexity and can help us to create institutions and rules that will generate robust outcomes. Complex systems cannot be optimized per se, but, to invoke the phrasing of Axelrod & Cohen (2000), they can be harnessed.

We have reason to be hopeful. Applications of ABMs to the physical, chemical, medical, epidemiological, biological, and ecological sciences provide evidence for skeptical optimism. In each of those fields, ABMs have proven capable of shedding light on the causes and implications of complexity (Epstein et al. 2006). In several domains of inquiry, such as models of immunity, the impact has been substantial.

Those successes notwithstanding, our optimism for the ultimate contribution to the study of politics is constrained. Ironically, our reservations are rooted in the complexity of political systems. In systems where ABMs have made substantial contributions to date, the fundamental actors—particles, chemicals, and flora and fauna—follow rules that with sufficient data can be captured with high fidelity. In addition, those rules are fixed and invariant across the agents. Social actors possess the ability to construct models of their worlds and to derive and take actions based on those models, adding a layer of complexity to social systems. People learn. We adapt. Politicians do the strangest things.

That said, not all social behavior has the same dimensionality of responses. And, to the extent that a social system consists of low-dimensional behavior, ABMs may prove more than up to the task. In contexts that admit rich behavioral repertoires, modeling becomes more problematic.[10]

Perhaps ABMs can be best seen as another tool at our disposal. Hyperbole from advocates aside, at their core, ABMs extend and expand game-theoretic models. Both models have the same parts: agents (or players) with interdependent payoffs. ABMs explore these environments using rule-based computer code, and game theorists use mathematics.

Though they occupy the same sandbox, agent-based modelers and game theorists build quite different castles in the sand. That is partly because they are playing in different parts of the box. Game theorists stick to the edges. Most game-theoretic models assume either two, three, or an infinite number of agents. Most analytic models assume full rationality. Those that do not tend to have all agents following the exact same rule. Game theorists take to the extremes not because they believe all interesting problems lie at the edge, but because the edge is tractable.

Agent-based modelers tend to play in the middle of the box, with intermediate numbers of actors. Agent-based modelers ignore the edges for the same reasons that game theorists stick to them; they tend to converge to equilibrium. More interesting, complex outcomes occur with moderate numbers of agents interacting in richer models.

To sum up, we would say that nearly a quarter century after their introduction, ABMs' contribution to the study of politics has been modest but important. We have highlighted many influential papers from across the discipline. What we have not fully credited is the extent to which ABMs have pushed the discipline toward a greater appreciation of features of our political world—networks, learning, diversity, and interactions—promoting new ways of thinking about

[10]Game theory avoids the difficulties inherent in capturing human behavior by assuming that agents optimize and that this optimizing behavior is common knowledge. Under relatively mild assumptions, game-theoretic models produce equilibria, albeit sometimes in mixed strategies only.

the world—recognizing complexity, seeing the difference between stability and robustness, distinguishing between tipping points, path dependence, and chaos—and enabling us to see the value and limitations of formal modeling from a broader perspective.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

Alchian AA. 1950. Uncertainty, evolution, and economic theory. *J. Polit. Econ.* 58.3:211–21

Anderson P. 1972. More is different. *Science* 177:393–96

Andreoni J, Miller JH. 1995. Auctions with artificial adaptive agents. *Games Econ. Behav.* 10(1):39–64

Ansolabehere S, et al. 2005. Voting weights and formateur advantages in the formation of coalition governments. *Am. J. Polit. Sci.* 49(3):550–63

Axelrod R. 1997. *The Complexity of Cooperation*. Princeton, NJ: Princeton Univ. Press

Axelrod R, Axtell R, Epstein J, Cohen MD. 1996. Aligning simulation models: a case study and results. *Comput. Math. Organ. Theory* 1:123–41

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

Axelrod R, Mitchell W, Thomas RE, Bennett DS, Bruderer E. 1995. Coalition formation in standard-setting alliances. *Manag. Sci.* 41:1493–508

Bak P. 1996. *How Nature Works: The Science of Self-Organized Criticality*. New York: Springer-Verlag

Bankes S. 2002. Agent-based modeling: a revolution? *Proc. Natl. Acad. Sci. USA* 99:7199–200

Baron DP, Ferejohn JA. 1989. Bargaining in legislatures. *Am. Polit. Sci. Rev.* 89:1181–206

Bednar J. 2009. *The Robust Federation: Principles of Design*. New York: Cambridge Univ. Press

Bednar J, Bramson A, Jones-Rooy A, Page SE. 2010. Emergent cultural signatures and persistent diversity: a model of conformity and consistency. *Ration. Soc.* 22(4):407–44

Bednar J, Chen Y, Liu X, Page SE. 2012. Behavioral spillovers and cognitive load in multiple games: an experimental study. *Games Econ. Behav.* 74(1):12–31

Bednar J, Page SE. 2007. Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Ration. Soc.* 19:65–97

Bendor J, Diermeier D, Ting M. 2003. A behavioral model of turnout. *Am. Polit. Sci. Rev.* 97(2):261–80

Billings D, Papp D, Schaeffer J, Szafron D. 1998. Poker as a testbed for machine intelligence research. Work. Pap., Dept. Comput. Sci., Univ. Alberta. **http://poker.cs.ualberta.ca/publications.html**

Brown GW. 1951. Iterative solutions of games by fictitious play. In *Activity Analysis of Production and Allocation*, ed. TC Koopmans, pp. 374–76. New York: Wiley

Bruch E, Mare R. 2006. Neighborhood choice and neighborhood change. *Am. J. Sociol.* 112:667–709

Camerer. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton Univ. Press

Carpenter DP, Esterling KM, Lazer DMJ. 2004. Friends, brokers, and transitivity: Who informs whom in Washington politics? *J. Polit.* 66(1):224–46

Cederman LE. 1997. *Emergent Actors in World Politics*. Princeton, NJ: Princeton Univ. Press

Cederman LE. 2005. Computational models of social forms: advancing generative process theory. *Am. J. Sociol.* 110(4):864–93

Conitzer V, Sandholm T. 2006. Computing the optimal strategy to commit to. *Proc. Seventh ACM Conf. Electron. Commer.* New York: Assoc. Comp. Mach.

Cutler J, de Marchi S, Gallop M, Hollenbach F, Laver M, Orlowski M. 2013. *Cabinet formation and portfolio distribution in European multi-party systems*. Work. Pap., Dept. Polit. Sci., Duke Univ.

de Marchi S. 1999. Adaptive models and electoral instability. *J. Theor. Polit.* 11(3):393–419 de Marchi S. 2005. *Computational and Mathematical Modelling in the Social Sciences*. New York: Cambridge Univ. Press

de Marchi S, Page S. 2009. Agent-based modeling. In *Oxford Handbook of Political Science Methodology*, ed. JM Box-Steffensmeier, HE Brady, D Collier, pp. 71–96. New York: Cambridge Univ. Press

Diermeier D, Krehbiel K. 2003. Institutionalism as methodology. *J. Theor. Polit.* 15(2):123–44

Ensley MJ, Tofias MW, de Marchi S. 2009. District complexity as an advantage in congressional elections. *Am. J. Polit. Sci.* 53(4):990–1005

Epstein JM. 2007. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton Univ. Press

Epstein JM. 2014. *Toward Neurocognitive Foundations for Generative Social Science*. Princeton, NJ: Princeton Univ. Press

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

Epstein JM, Goedecke DM, Yu F, Morris RJ, Wagener DK, Bobashev GV. 2006. *Pandemic influenza and the value of international travel restrictions*. CSED Work. Pap. 46, Brookings Inst., Washington, DC

Fowler J. 2006a. Connecting the Congress: a study of cosponsorship networks. *Polit. Anal.* 14(4):456–87

Fowler J. 2006b. Habitual voting and voter turnout. *J. Polit.* 68(2):335–44

Fowler J. 2006c. Legislative cosponsorship networks in the U.S. House and Senate. *Soc. Netw.* 28(4):454–65

Fowler J, Laver M. 2008. A tournament of party decision rules. *J. Conf. Resolut.* 52(1):68–92

Fowler J, Smirnov O. 2005. Dynamic parties and social turnout: an agent-based model. *Am. J. Sociol.* 110:4

Gartzke E, Weisiger A. 2013. Fading friendships: alliances, affinities and the activation of international identities. *Br. J. Polit. Sci.* 43:25–52

Gaylord R, D'Andria LJ. 1998. *Simulating Society: A Mathematical Toolkit for Modeling Socioeconomic Behavior*. London: Springer

Geanakoplos J, Axtell R, Farmer JD, Howitt P, Conlee B, et al. 2012. Getting at systemic risk via an agent-based model of the housing market. *Am. Econ. Rev.* 102(3):53–58

Gell-Mann M. 1994. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. New York: Freeman

Glaeser E, Sacerdote B, Scheinkman J. 1996. Crime and social interactions. *Q. J. Econ.* 111:507–48

Golder M, Golder S, Siegel D. 2012. Modeling the institutional foundation of parliamentary government formation. *J. Polit.* 74(2):427

Golman R, Page SE. 2009. Basins of attraction and equilibrium selection under different learning rules. *Evol. Econ.* 20:49–72

Golman R, Page SE. 2010. Individual and cultural learning in stag hunt games with multiple actions. *J. Econ. Behav. Organ.* 73(3):359–76

Hoff PD, Ward MD. 2004. Modeling dependencies in international relations networks. *Polit. Anal.* 12:160–75

Holland J, Miller J. 1991. Artificial agents in economic theory. *Am. Econ. Rev. Pap. Proc.* 81:365–70

Hong L, Page S. 2009. Interpreted and generated signals. *J. Econ. Theory* 144:2174–96

Huckfeldt RR, Johnson PE, Sprague JD. 2004. *Political Disagreement: The Survival of Diverse Opinions Within Communication Networks*. Cambridge, UK: Cambridge Univ. Press

Jackson M. 2008. *Social and Economic Networks*. Princeton, NJ: Princeton Univ. Press

Jen E. 2005. *Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies*. Oxford, UK: Oxford Univ. Press

Jervis R. 1997. *System Effects: Complexity in Political and Social Life*. Princeton, NJ: Princeton Univ. Press

Jones-Rooy A, Page SE. 2012. The Complexity of Systems Effects. *Crit. Rev.* 24(3):313–42

Judd K. 1997. Computational economics and economic theory: complements or substitutes? *J. Econ. Dyn. Control* 21(6):907–42

Kauffman SA. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford Univ. Press

Kim S-Y, Taber CS, Lodge M. 2010. A computational model of the citizen as motivated reasoner: modeling the dynamics of the 2000 presidential election. *Polit. Behav.* 32:1–28

Kirman A. 1997. The economy as an interactive system. In *The Economy as a Complex Evolving System II*, ed. WB Arthur, S Durlauf, D Lane, pp. 491–533. Reading, MA: Addison Wesley

Kollman K. 2012. The potential value of computational models in social science research. In *Oxford Handbook in Philosophy in the Social Sciences*, ed. H Kincaid. New York: Oxford Univ. Press

Kollman K, Miller J, Page S. 1992. Adaptive parties in spatial elections. *Am. Polit. Sci. Rev.* 86:929–37

Kollman K, Miller J, Page S. 1997. Political institutions and sorting in a Tiebout model. *Am. Econ. Rev.* 87:977–92

Kollman K, Miller J, Page S, eds. 2003. *Computational Models of Political Economy*. Cambridge, MA: MIT Press

Kollman K, Page S. 2005. Computational methods and models of politics. In *The Handbook of Computational Economics*, ed. L Tesfatsion, K Judd, pp. 1433–63. Amsterdam: Elsevier/North Holland

Lamberson PJ, Page SE. 2012. Tipping points. *Q. J. Polit. Sci.* 7(2):175–208

Laver M. 2005. Policy and the dynamics of political competition. *Am. Polit. Sci. Rev.* 99(2):263–81

Laver M, de Marchi S, Mutlu H. 2011. Negotiation in legislatures over government formation. *Public Choice* 147:285–304

Laver M, Sergenti E. 2011. *Party Competition: An Agent-Based Model*. Princeton, NJ: Princeton Univ. Press

Lazer D. 2001. The co-evolution of individual and network. *J. Math. Sociol.* 25:69–108

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

Leombruni R, Richiardi M. 2005. Why are economists sceptical about agent-based simulations? *Phys. A* 355(1):103–9

Lodge M, Sung-Youn K, Charles T. 2004. *A computational model of political cognition—the dynamics of candidate evaluation in the 2000 presidential election*. Presented at Annu. Meet. Midwest Polit. Sci. Assoc., Chicago

Lustick I, Miodownik D. 2000. Deliberative democracy and public discourse: the agent-based argument repertoire model. *Complexity* 5(4):13–30

Lustick I, Miodownik D, Eidelson RJ. 2004. Secessionism in multicultural states: Does sharing power prevent or encourage it? *Am. Polit. Sci. Rev.* 98(2):209–29

Miller J. 1998. Active nonlinear tests (ANTs) of complex simulations models. *Manag. Sci.* 44:820–30

Miller J, Page S. 2007. *Complex Adaptive Systems: Computational Models of Social Life*. Princeton, NJ: Princeton Univ. Press

Mitchell M. 2009. *Complexity: A Guided Tour*. Oxford, UK: Oxford Univ. Press

Newman MJ. 2003. On the logic and structure of complex networks. *SIAM Rev.* 45:167–256

Nowak M, Bonhoeffer S, May R. 1994. Spatial games and the maintenance of cooperation. *Proc. Natl. Acad. Sci. USA* 91:4877–81

Nowak M, May RM. 1992. Evolutionary games and spatial chaos. *Nature* 359:826–29

Ostrom E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge Univ. Press

Page SE. 1997. On incentives and updating in agent based models. *Comput. Econ.* 10:67–87

Page SE. 2006. Essay: path dependence. *Q. J. Polit. Sci.* 1(1):87–115

Page SE. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Teams, Schools, and Societies*. Princeton, NJ: Princeton Univ. Press

Page SE. 2008. Uncertainty, difficulty, and complexity. *J. Theor. Polit.* 20:115–49

Page SE. 2011. *Diversity and Complexity*. Princeton, NJ: Princeton Univ. Press

Palmer R, Arthur W, Holland J, Lebaron B, Tayler P. 1994. Artificial economic life: a simple model of a stock market. *Phys. D* 75:264–74

Richards D. 2000. Nonlinear dynamics in games: convergence and stability in international environmental agreements. In *Political Complexity: Nonlinear Models of Politics*, ed. D Richards, pp. 173–206. Ann Arbor: Univ. Michigan Press

Schelling T. 1978. *Micromotives and Macrobehavior*. New York: Norton

Smirnov O, Fowler JH. 2007. Policy-motivated parties in dynamic political competition. *J. Theor. Polit.* 19(1):9–31

Tesfatsion L. 1997. How economists can get alife. In *The Economy as a Complex Evolving System II*, ed. W Arthur, S Durlauf, D Lane, pp. 533–65. Reading, MA: Addison Wesley

Vriend N. 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *J. Econ. Dyn. Control* 24:1–19

Weidmann N, Salehyan I. 2013. Violence and ethnic segregation: a computational model applied to Baghdad. *Int. Stud. Q.* 57:52–64

Wichowsky A. 2012. District complexity and the personal vote. *Legis. Stud. Q.* 37(4):437–63

Zinkevich M, Bowling M, Johanson M, Piccione C. 2007. *Regret minimization in games with incomplete information*. Presented at Neural Inf. Process. Syst., Dec. 3–6, British Columbia, Can. **http://books.nips.cc/nips20.html**