Oxford Handbooks Online

Agent-Based Modeling

Scott de Marchi and Scott E. Page

The Oxford Handbook of Political Methodology
Edited by Janet M. Box-Steffensmeier, Henry E. Brady, and David Collier

Print Publication Date: Aug 2008
Online Publication Date: Sep 2009
Subject: Political Science, Political Methodology, Comparative Politics
DOI: 10.1093/oxfordhb/9780199286546.003.0004

[–] Abstract and Keywords

This article provides a discussion on agent-based modeling. Two examples that show the ability of computational methods to extend game-theoretic results are presented. It then discusses modeling agents, modeling agent interactions, and system behaviour. In addition, it describes how agent-based models differ from and complement mathematical models and concludes with some suggestions for how one might best leverage the strengths of agent-based models to advance political science. Most mathematical analyses of game-theoretic models do not look into the stability and attainability of their equilibria and would be made richer by complementing them with agent-based models that explored those properties. The ability of computational models to test the robustness of formal results would be reason alone to add them to tool kits. As a methodology, agent-based modeling should be considered as in its infancy, its enormous potential limited only by the scientific and creative talents of its practitioners.

Keywords: agent-based modeling, game-theoretic models, computational methods, modeling agents, system behaviour

1 Computational Models as Numerical Methods

In mathematics, engineering, and the biological and physical sciences, computational methods have proliferated in recent years due in part to the increase in computing power available to researchers. The ability of computational methods to build upon deductive results has expanded the set of problems researchers can investigate. As a result, in almost every field outside of the social sciences, computational methods now occupy substantial space in the working scientist's bag of tools, and one can see this everywhere from textbooks to the proliferation of clustered computing on campuses. The growth in computational models is also partly caused by demand. Pressing problems ranging from the global (climate change), to the social (disease spread), to the microscopic (protein folding, <http://folding.stanford.edu>) outstrip current mathematical expertise. Owing to their complexity, leading scholars now model each of these problems primarily with computers, eschewing analytics. For more esoteric, academic pursuits such as solving differential equations (Strogatz 2001), producing statistical measures (Gentle 2002), and locating optimal solutions (Judd 1998 and Ballard 1997), using computational methods has also become mainstream.

(p. 72) The question then for social science is not if, but how computational methods can help extend existing results. For statistical work, the answer is easy. Bayesian statistics in our discipline [1] along with preexisting research programs such as Poole and Rosenthal's work on Congress (see Poole 2006 for an overview), illustrate the utility of computational methods. For formal models, however, the answer to the foregoing question remains open to debate within our discipline. There is some belief (particularly among formal theorists themselves; see Tesfatsion 1997; Leombruni and Richiardi 2005) that noncooperative game theory is distinct from other modeling enterprises; investigations of nonequilibrium dynamics or of alternative solution concepts such as the core are not

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

seen as valuable. Baron and Ferejohn (1989), for example, argue that noncooperative bargaining theory takes institutions "seriously" and (unlike prior cooperative models) provides "definite predictions about legislative processes." Diermeier and Krehbiel (2003) extend this argument by asserting that behavioral rules should be fixed (i.e. based on rational choice assumptions), and that "the need for an equilibrium concept [i.e. Nash and its extensions] that is general and flexible enough to model a broad variety of institutions and institutional choices, strongly favors non-cooperative game theory" over other methods.

The case for analytics rests on their *generality* and their capacity to reveal *cause and effect*. Some analytic results hold for classes, e.g. convex or single peaked, of functions. Others hold for parametric families of functions such as Cobb-Douglass utility functions. Computational results hold only for the values tested or, in the best case, statistically for parametric families of functions. Thus, a case can be made that analytic results are more general (e.g. the justly famous Arrow result is both general and explores substantively interesting assumptions). Second, deriving results often reveals cause and effect. It shows how and why assumptions matter. By comparison, computational models can be black boxes.

These critiques have limited merit. Clearly, if a general analytic result can be shown, then by all means it should be. But general results rarely exist and so the standard approach is to simplify the model rather than abandon the methodology. As for the critique that math provides clarity and computational models obscure, reality suggests the opposite. Computational models often provide greater transparency because they can reveal the value of any variable or combination of variables.

At the risk of being provocative, we sense some oppose agent-based and computational methods for other, less scientifically based reasons. Some dislike the methodology because it makes modeling too easy. With a computer, anyone can construct a model and many of these models will be uninformative. Computational models are more flexible, which we see as a good thing, but what this critique misses is that they are also logically consistent owing to the fact that they are programmed. Others question computational approaches because they equate analytic techniques, Nash (p. 73) Equilibrium, and purposeful behavior. In fact the latter two are in no way implied by nor do they require mathematics. Yes, equilibria can be derived with mathematics, but they can also emerge computationally.

To date, there has been insufficient commerce between mathematical and computational modelers, so much so that an outside observer might conclude that computational and formal methods are incompatible. Given that this supposed incompatibility of methods does not exist in mathematics or the physical sciences, it is our position that it is not a necessary feature of the social sciences either, but rather a disciplinary choice that will likely be short-lived. Computational models enable scholars to explore richer environments (see, for example, Fowler and Smirnov 2005); they can move beyond Nash Equilibrium to include coalitional deviations (see Page forthcoming), and they can incorporate more realistic behavioral assumptions (Bendor, Diermeier, and Ting 2003). Computational models, therefore, should be seen as a tool that extends the problem domain. In our view, models are models, regardless of the method employed to investigate them.

We provide two examples that will illustrate this point, that highlight the ability of computational methods to extend game-theoretic results. The first example is as simple as possible, so as to provide some general intuition. The second considers a more sophisticated model and better shows the power of computational models to deepen our understanding.

1.1 The Currency Game: Computational Models Used to Expand and Complement Deductive Results

Consider the following example borrowed from Young (2001), who used it as a model to investigate path dependency. N actors in a society must decide among two possible currencies: gold and silver. Initially, the N actors will be randomly assigned gold or silver with equal probability. Let p_t be the proportion of gold users in the population at time t and $(1 - p_t)$ be the silver users. At each subsequent time period, one actor will be chosen by a uniform draw from the population and will make a new decision according to the following rule:

With probability $(1 - \epsilon)$, the actor chooses the more common currency. Formally, if $p_t > 0.5$ (i.e. gold is the dominant currency) the actor chooses gold or remains a gold user if one already; else. The actor behaves similarly

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

if silver is dominant. If $p_t = 0.5$ exactly, the actor continues with whatever currency he or she was using previously. With probability $\epsilon > 0$, a shock occurs and the actor randomly chooses one of the currency. Thus, with a 50 percent probability, he or she switches to the nondominant currency, i.e. makes a mistake. In this model, a natural dependent variable would be the average number of regime shifts one sees over any given time period *t*. For example, if the game goes on for 100 iterations, how many times would the currency change from one standard to another? This outcome variable is an integer with a (p. 74) range from [0.*t*]. The two main explanatory variable are ϵ (the mutation rate) and *N* (the population size) with presumably more changes with a higher mutation rate and fewer changes with a larger population.

One can go about solving this problem in two different ways, as outlined in de Marchi (2005). The first approach uses game theory. Using this approach, the currency game generates relatively long periods of convergence characterized by sudden shifts to the alternative currency[2] Young proves a theorem that states that if you have a suitably large population, a majority of the agents will choose the same currency. Additionally, this majority is largest when ϵ is small. This provides intuition about the currency game, but many readers might not be certain about what the properties of the currency game are for different values of *N* and ϵ. How large is a "suitably large" population? How small does epsilon have to be? Interpreting this theorem can thus be elusive, even though it is deductively true and even though the game has only two parameters. Young, for example, claims that the currency game is useful in studying path dependent processes, such as competing technologies (see Arthur 1989):

> After an initial shakeout, the process converges quite rapidly to a situation in which most people are carrying the same currency—say, gold. This norm will very likely stay in place for a considerable period of time. Eventually, however, an accumulation of random shocks will "tip" the process into the silver norm. These tipping incidents are infrequent compared to the periods in which one or the other norm is in place (assuming ϵ is small). Moreover, once a tipping incident occurs, the process will tend to adjust quite rapidly to the new norm. This pattern—long periods of stasis punctuated by sudden changes of regime—is known in biology as the *punctuated equilibrium effect*. (Young 2001,11–12; emphases in original)

To complement his mathematical results, Young provides graphs of numerical experiments that illustrate multiple regime shifts. For the experiments he details in the text, he relies upon small population sizes (N =10) and a very high mutation rate (ϵ = .5) to generate his intuition about the properties of the game. He runs 30,000 iterations of the model and considers a very small subset of the parameter space in his initial forays into modeling the game. For $N = 10$ and $\epsilon = .5$, $E(\cdot) = 0.015$—in 1,000 iterations of the game, one would expect to see 15 regime shifts. If one considers a larger population and a smaller mutation rate—which more closely mirrors the nature of the real world problem Young uses the currency game to investigate[3]— one quickly discovers that regime shifts never, ever happen. For $N = 100{,}000$ (still a tiny population) and $\epsilon = .05$ (a reasonably high mutation rate), E (regime shifts) = $6.89E$—1534—that is a lot of 0's.

What Young has done is build a simple computational model, which in this case is a straightforward process of translating the dynamics of the game into a programming (p. 75)

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

Table 4.1. Results for the currency game language

Poisson regression				**Number of obs = 1600**		
				LR chi2(2) = 19492.69		
				Prob > chi2 = 0.0000		
Log likelihood = −2244.3969				**Pseudo R2 = 0.8128**		
Num_reg	**Coef.**	**Std.Err.**	**Z**	**P> |z|**	**[95 percent Conf.**	**Interval]**
N_pop	−.0351948	.0007884	−44.64	0.000	−.03674	−.0336496
Mutat_rate	15.4024	.9942932	15.49	0.000	13.45362	17.35118
_cons	−4.148296	.4970227	−8.35	0.000	−5.122443	−3.17415

like C, python, or Perl—in many respects, this is much easier than constructing a game-theoretic model of comparable complexity. While choices such as the particular computer language are a matter of personal preference, computational models must obey the following strictures:

1. encode the rules of the game as accurately as possible;
2. iterate the computational model through a wide range of parameter values;
3. present a set of statistical results that explain the relationship between the parameters of the model and the dependent variables of interest.

Once this process is accomplished, the results of the computational model produce data in a format that is familiar to social scientists. Each observation is the product of running the model using different parameter choices for N and ϵ, and these data can be used to investigate the role of N and ϵ in predicting mean regime shifts.

As in any statistical work, the problem of choosing a reasonable empirical specification must be faced, though in the case of computational models this is somewhat offset by the ability to generate novel data to test various specifications. One can then use out-of-sample data-sets and make multiple attempts to find a good match. In the case of Young's currency game, a good model is a Poisson regression, which for a large number of observations spans much of the parameter space described by N and ϵ (Table 4.1).

For a Poisson regression, the predicted values are given by e^{xb}. By examining representative values, we can gain a more complete understanding. The output table shows that as N increases, universal adoption of a single currency becomes extremely likely because the likelihood of regime shifts decreases, and regime shifts almost never occur.

Contrasting Young's deductive approach with the computational model, we see that the deductive result gives us a baseline understanding of what the currency game looks like in equilibrium, but it does not tell us much about the dynamics of the system, or even what the system would look like for different parameter values. The computational model extends and deepens our understanding of the deductive result, and generates an empirical model that is both familiar to social scientists (due (p. 76) to the ubiquity of statistics training) and useful in generating intuition about different parameter values. Seen in this way, the deductive and computational approaches are complementary.

1.2 Baron-Ferejohn Bargaining: Computational Models Used to Solve Hard Problems, Search Parameter Spaces, and Investigate Equivalence Classes

Bargaining models occupy a central place in the political science and economics literatures. This prominence is due to two factors. First, these models can be applied in a range of contexts, ranging from cabinet formation in PR

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

systems to coalition behavior within legislatures to alliance structures in the international system. Second, bargaining models produce testable hypotheses which can be explored in both the laboratory setting (Fréchette, Kagel, and Morelli 2005) as well as with macro-level political data (Warwick and Druckman 2006).

Here, we consider the Baron-Ferejohn bargaining model. To call it a single model is not correct. In their original 1989 paper in the *American Political Science Review*, they presented a number of qualitatively distinct games, and those models have, in turn, been extended by other scholars. For example, their infinite-horizon, closed- rule model has received the most attention in the literature. Recently, Snyder, Ting, and Ansolabehere (2005) have advocated for the use of minimum integer weights as a proxy for bargaining power rather than raw votes scores.[4] Subtle changes in assumptions about weights, timing, and parameter create an ensemble of possible models. A crude characterization of these models suggests that they must include:

- discount rate for the value of time
- horizon for the number of periods that bargaining occurs
- minimum integer or raw voting weights
- minimum winning or minimum weight coalitions
- open vs. closed rule for changes in the agreement
- recognition rule (i.e. who proposes a coalition in each time period)
- reversion point if bargaining fails
- stationarity (or not) so that agents do (or do not) incorporate history in their strategies.

The core insight of the formal results derived in the Baron-Ferejohn literature is that the proposer has an advantage (in government coalitions, this is named the *formateur* advantage), but this of itself is not very precise. Ideally, we could make more precise statements of exactly how much of an advantage we would expect to see for a *formateur*.

(p. 77) It is important to note that these bargaining games produce brutal combinatorics when one wishes to extent these results to greater than three players. Calculating minimum integer weights, for example, is NP-hard and necessitates the use of computational methods (Strauss 2003). We thus have computational work as a handmaiden to the formal results, but the larger question remains. Has the formal work in this literature yielded a consensus opinion given the large parameter space?

The short answer is “no.” Diermeier and Morton (2003), for example, use experiments to test whether subjects adhere to behavior that matches the Baron-Ferejohn model of bargaining. In this model, equilibrium play dictates that the proposer will take a disproportionate share of the pie; in contrast, it might be the case that subjects instead take a piece of the pie proportional to their voting weight in the coalition, regardless of whether they were the proposer or not. Their design relies upon three “different” bargaining games with the following vote totals (all adding to 99) for each of three groups: {34,33,32}, {49,33,17}, and {46,44,9}; and uses a proportional rule for recognition with a finite number of proposal periods. Given these assumptions, each game has a different equilibrium. The problem, of course, is that it is not obvious that these three voting contexts are in fact different games. Their minimum integer weight representations are the same—{1,1,1}—and thus the expected continuation values (the values that each player can expect to receive) would also be the same: 1/3 for each player regardless of their differing vote totals.

Other parameter choices such as an infinite horizon, different recognition rules for selecting the proposer, or a different discount rate would also qualitatively change their expectations. We thus find ourselves in an uncomfortable situation, where minor parameter choices produce a range of dissimilar outcomes that make it difficult to describe Baron-Ferejohn bargaining as a unitary group of models.

In large part, much of the opacity in the Baron-Ferejohn inspired literature can be explained by the fact that in many settings, researchers use simple examples and do not worry overmuch about the parameter choices. Individual modelers pick a set of parameters, prove a result, and fill in a single pin-prick of light in a vast, dark void. Just as with the foregoing example of path dependency and currency adoption, the quality of the intuition generated by this approach to modeling can be quite lacking and computational models are one way to fill in more of this void.

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

As Laver, de Marchi, and Mutlu (2007) show, computational models extend the reach of existing deductive results by testing different parameterizations. Questions like:

- does a greater vote share always help a party gain cabinet seats?
- what happens if we assume parties prefer coalitions with minimum weight rather than minimum winning partners?
- are nonhomogeneous and homogeneous games the same?
- what happens if the reversion point is not the null vector?

are not easily answered by existing formal models. One needs to span more of the parameter space to address them, and doing this proof by proof is often inefficient.

(p. 78) Additionally, if one plans on testing a formal model, computational models are the glue that makes this enterprise possible.[5] Unless one is completely comfortable with point values chosen in an ad hoc fashion for the above parameters, it is necessary to span as much of the parameter space as possible, if one wants to connect formal models to empirical work. Simply put, measures have error terms, and one needs to know if one's model is immune to such perturbations.

2 Computational Models as Agent-based Models

Unlike the foregoing examples, where computational models are portrayed as a systematic and efficient numerical method for extending the reach of deductive methods, computational models can also investigate problems that are intractable for deductive approaches. Many of the problems that confront the social sciences quickly become very difficult for purely deductive approaches, and it is exactly this kind of complexity that has ushered in a computational era in mathematics and the physical sciences. Faced with a choice between constructing a simple (albeit solvable) game and making a loose analogy to the real-world puzzle we are interested in, and constructing a more realistic computational model (that we can only investigate with statistics), many researchers may wish to follow the latter course. Every form of cooperation is not described by the iterated prisoner's dilemma, and every bargaining situation is not a Rubinstein game. Moreover, every political or economic agent is not able to perform backwards induction in all circumstances.

The main architecture for dealing with these sorts of complex problems is agent- based modeling. This methodology consists of agents who interact according to a set of rules instantiated by a computational algorithm. Often, agents are drawn from a set of heterogeneous types and are situated in a geographic or social structure that takes the form of a network. Agents interact with other agents, and these interactions and the agent information sets are constrained by the networks. Unlike game- theoretic models where the theorist has an end in sight and builds a model that achieves that end point in equilibrium, agent-based modeling starts with the rules and environments that govern behavior, and examines the dynamics that result from these assumptions. This kind of social science is grown from the bottom up, and more emphasis is typically placed on achieving verisimilitude with the assumptions of the (p. 79) computational model (Palmer et al. 1994; Epstein and Axtell 1996; Arthur et al. 1997; Epstein 2006; Miller and Page 2007).

It is important to understand that in most cases these are dynamic models in which agents behave purposefully, but do not necessarily follow best response strategies nor do they necessarily choose what would be stage game equilibrium strategies. These agent-based models may produce path dependent outcomes, tipping points, self-organized criticality, or chaos, but in all cases the dynamics of how (or if) you arrive at an equilibrium is of substantive interest. Questions such as how likely the model is to achieve an equilibrium, or how sensitive that equilibrium is to small perturbations in the model, may be investigated with this approach.

Agent-based models thus offer vast possibilities, which can be both a blessing and a curse. On the one hand, agent-based models allow researchers to capture any situation or problem that can be represented in a computational algorithm. Because of this, they offer the flexibility of narrative accounts coupled with the logical consistency of analytic approaches (Holland and Miller 1991). Using more realistic assumptions and building dynamic models allows a wider range of phenomena to be investigated. One recent example is Lodge, Kim, and

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

Taber (2004), who use a state-of-the-art-model of cognition and Annenberg data to predict voter responses to campaign events such as party conventions. In this case, the model's logical consistency enables Lodge, Kim, and Taber to use a far more interesting (and plausible) model of decision-making than is the norm, with enough transparency to see why the voters in the model react to new information in the way that they do.

A risk, as mentioned earlier, is that computational models can contain too many moving parts. The potential even exists for "rigging results" in which the modeler adds assumptions until generating a preferred result. The same critique, of course, also applies to game-theoretic models. One can produce almost any outcome (a Nash equilibrium, a particular set of dynamics, whatever) with the "right" assumptions. A difference between the two is that unlike game-theoretic models, computational models are instantiated in a programming language, and while they look like proofs, they most often contain hundreds or thousands of statements and can be difficult to assess for most researchers. Proving the robustness of a computational result is hard work and reliable methods for accomplishing this is an active research area in the field (Axelrod 1997; Miller 1998; de Marchi 2005; Laver and Sergenti 2007).

As with any methodology, proper usage requires great care and sophistication. Used properly, agent-based models offer substantial potential; nevertheless, methodological concerns exits. Below, we discuss three core concerns: *modeling agents, modeling agent interactions*, and *analyzing system behavior*. We describe how agent-based models differ from and complement mathematical models and conclude with some suggestions for how we might best leverage the strengths of agent-based models to advance political science. Throughout, we will use the example of computational models of elections (following Kollman, Miller, and Page 1992; de Marchi 1999; and Laver 2005) to make our points clear.

(p. 80) 2.1 Elections as Optimization Problems: Computational Models Used to Describe Candidate and Voter Behavior

Starting with Kollman, Miller, and Page (1992), elections have been a very active area of research in the computational modeling field. Unlike analytic models of elections, where candidates and voters use best response functions and react strategically, computational models of elections treat the problem as a nonstrategic optimization problem for voters and candidates. The electorate has preferences over platforms, but without complete information it is difficult for candidates to know which platforms are preferred. Moreover, the bulk of work on mass electorates indicates that the policy space is at least two-dimensional, and may have additional complexities such as nonseparable or nonsingle peaked preferences over some issues. The fact that candidates resort to polling, focus groups, and political consultants illustrates the immediacy of this problem in the real world, so how would one model candidate behavior?

The famous chaos results of Arrow, McKelvey, Plott, Schofield, and others indicate that, in general, there is no equilibrium, and most efforts to tame this problem have been controversial. Computational models of elections look at candidate behavior from a different perspective. If boundedly rational candidates, with limited computational ability, resources, and incomplete information, had to look for a platform, how would they do it, and what sorts of electoral outcomes would one see?

The first step in building a model of this kind is to structure mass preferences as an electoral landscape. In general terms, this means that the policy space is represented as n dimensions, and an additional dimension captures the aggregate voter utility of each platform. Visually, this electoral landscape forms a geography where high points represent advantageous platforms, and low points disadvantageous platforms. Candidates traverse this landscape with their limited resources using heuristics taken from the optimization literature, and choose the best platform they "see" given limited resources.

One such optimization technique is hillclimbing. Hillclimbing candidates look in a small neighborhood around their current platform, then change platforms to the best neighbor they find. Iterating this algorithm allows candidates to optimize locally and improve their chances in an election, but it is important to realize that local optima can trap such an algorithm. These algorithms are not strategic in a game-theoretic sense, insofar as there is no term in the optimization function for the other candidate's behavior.

This framework was extended by de Marchi (1999; 2003) to include voters with limited attention and varying levels

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

of sophistication using Zaller's RAS model as inspiration. Laver (2005) added parties with more realistic optimization strategies, and mapped the outcomes in the computational model to data from Irish elections. Broadly, however, this class of models has demonstrated a real advantage to incumbency (which stands in contrast to game-theoretic accounts)—like Alchian's evolutionary theory of firm competition, incumbents by definition have been successful at least one time. Whether by luck or because they have a solid underlying (p. 81) algorithm, it is difficult for challengers to unseat them because of the problems posed by searching complicated electoral landscapes. Ensley, Tofias, and de Marchi (2007) provide empirical support for this claim using data from the US case. Thus, we have agents (in this case, candidates), rules (their optimization/search algorithms), and a system (modeled as electoral landscapes) which produces behaviors such as an incumbency advantage. With this example in mind, we will now examine each of these aspects of agent-based models in more detail.

2.2 Agents as Rule-based Objects

An agent-based model consists of a set of computational objects. An agent is defined by a vector of *state variables* along with a set of *rules*. For convenience, we can refer to both states and rules as *attributes*. It will be important to distinguish those state variables that are seen by the other agents from those that are known only to the agents themselves. We refer to the former as *external* and the latter as *internal*. Using this nomenclature, an agent's *action* and *geographic location* belong to its external state, but an agent's *preferences* belong to its internal state. The distinction between internal and external state variables need not be a bright line. An agent's private information could be internal but subsequently learned by other agents.

Agents also contain *rules*. We distinguish between three types of rules: *external rules, internal rules*, and *rule-changing rules*. External rules describe how an agent changes its external state variables. These rules depend on the environmental variables, the external states of other agents, and the external and internal state of the agent. For example, if we use the above example of *n*-party spatial competition, the external state of an agent, in this case a political party, would be the spatial location of its platform. Its external rule would be its rule for changing its platform as a function of its preferences (its internal state), its current position (its external state), the position of the other party (an external state of another agent), and the preferences of the voters over platforms (external states of still other agents).

Agents also have internal rules. These rules change the agent's internal state variables and often use the same arguments as the external state rule. In our model of political competition, a voter's internal state might include a mapping from policies to outcomes. This mapping would determine the voter's preferences. If we assume that winning parties implement their preferred policies, then voters would learn how policies map into outcomes. This learning would be an example of an internal rule. Thus, over time, the model might produce voters who update their policy mappings (Cohen and Axelrod 1984).

Finally, agents can have rule-changing rules. These meta-rules allow the agents to change their internal and external rules, allowing the system to adapt (Vriend 2000). As an example, consider a model of K communities each of which contains *N* agents who confront a collective action problem. Within each of these communities, agents have external rules that tell them how much to contribute as a function of how much other agents contributed. Selection could occur at the level of agent, where (p. 82) more successful agents reproduce at higher rates. Or selection could occur at the level of community, where more successful communities replace less successful ones. In the latter case, the external rules of the agents in one community are replaced by the external rules of agents in another community. Selection thus changes the distribution of rules that exist in the population.

2.3 Following Rules versus Maximizing Utility

In an agent-based model, the rules governing agent behavior are a focal point of the enterprise (Riolo, Axelrod, and Cohen 2001; Miller and Page 2007). Most agent-based models make no explicit reference to a payoff function. These include Schelling's Tipping Model (1978), the game of life, Axelrod's culture model (1997), Lustick's identity model (Lustick and Miodownik 2000), and some versions of Epstein and Axtell's Sugarscape model (1996). What exactly distinguishes agent-based models from game- theoretic models? On the surface, the difference is not vast. We can, for example, embed agents who use best response functions in an agent-based model. Thus, every equilibrium game-theoretic model can be written as an agent-based model. However, that agent-based model is not very interesting. If we start the model in equilibrium, it stays there. If we do not, best response dynamics often

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

overshoot the equilibrium. Thus, to really compare rational choice models to agent-based models, we should step away from equilibrium models and consider learning models.

So, is there a difference? Let's reconsider political parties competing in a multidimensional issue space and compare two possible behavioral rules that an incumbent might apply in their last chance to modify their platform prior to an election. The first rule polls some sample of voters, computes the median of their ideal points, and moves the candidate's position in the direction of that median. The second rule tests a new platform by polling voters in a neighborhood of the existing platform and uses that information to adapt the platform. Both rules seem plausible empirically, yet only the second is consistent with a payoff function. Why? Because even good rules sometimes make stupid mistakes. Assume that the candidate's payoff function is to maximize its probability of winning an election (or its vote total). Consider the last platform modification prior to the election. The incumbent should choose a policy that it believes based on its information produces an expected increase in the probability of winning. With positive probability, the first rule will cause the candidate to choose a platform that it knows—based on previous polls—will do worse than its current platform. It may know this with certainty or it may expect it to be true. This can occur because the candidate is not taking into account what it has previously learned. Thus, in some cases, the first rule results in the incumbent taking an action inconsistent with a utility function. The second rule does not necessarily optimize or seek a best response to the challenger, but it will increase expected utility in all cases.

Our point is a simple one: *actors who follow purposeful rules may sometimes take actions that are not rational given a utility function approach*. Thus, a true adherent to **(p. 83)** the principles of game theory may therefore see any particular rule-based behavior as ad hoc. We disagree. In many cases, we may know the rules that people follow. And, we might add, that if other players follow rules instead of adhering to utility function- based behavior, it's not clear that following a utility function-based rule makes sense.

Our discussion at this point bears some similarity to the lengthy debate about off- equilibrium play in the centipede game (see Binmore 1996). In this game, imagine your opponent moves first, and does not employ the subgame equilibrium of ending the game—what do you do? Most game theorists would recommend that you follow the equilibrium strategy and end the game, but there is some sense (entirely justified) that playing this strategy would cost you money.

It is worth summarizing how one applies game theory to a problem. First, one needs a representation of the problems that agents confront, and ideally this encoding represents what players know at different points in the game. In game theory, this instantiation is most often an extensive form game, where the innovation of information sets provides a vehicle for understanding how knowledge impacts play. Second, one needs explicit utility functions that represent how players evaluate the outcomes of the game. Last, one needs a solution concept and an algorithm that "solves" a given problem. Game theory typically employs a Nash equilibrium (or a refinement) as the solution concept, and given an extensive form and a utility function, one can apply backwards induction to test for the existence of a Nash equilibrium.

Humans are in many ways limited or less adept than backwards inducing *homo economicus*, and it is also the case that humans typically make choices in difficult environments with limited information. Kahneman and Tversky style experiments point to our obvious flaws in relatively simple tasks; but up the ante a bit by presenting humans with more complicated games and the tables turn quite dramatically. The Ing prize in the boardgame Go, for example, offered just over one million dollars to the first computer player to beat a master level human; when the prize was discontinued in 2000, the best computer player was having difficulty beating eight-year-olds (albeit good eight-year-olds). We "outperform" rational choice players, and there is no reason to suspect that game-theoretic models have much to say about certain (complex) classes of human games and decision contexts. Thus, while computer players modeled after rational choice players have enjoyed significant success in recent years, these successes have been greatest in two-person, zero-sum games.[6]

If the goal is to model people, political parties, organizations, and nations, however, it is fair to ask if the game-theoretic approach is useful in all contexts. Few psychologists believe that utility functions underpin behavioral rules except in a loose sense. Behavior of an organization is probably not best described as optimal relative to an **(p. 84)** underlying payoff function but as the expression of tried and true routines and rules that are thought to produce outcomes consistent with broad goals. We do not mean to imply that game-theoretic assumptions of utility

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

function maximization never apply, just that in many contexts, people, organizations, and nations follow rules and these rules need not be maximizing a specific utility function.

We might add that when agents represent collections of individuals, the existence of a utility function or an objective function may be problematic. Diverse preferences need not aggregate into a collective preference ordering. Overall, the premise that we might learn something by saying "here's what we think nations do, let's set those rules loose and see how they aggregate" seems as plausible as saying "here are their payoffs, let's solve for equilibrium."

To summarize, the largest differences between game theory and agent-based modeling are in how agents are characterized and in the solution concepts employed. Game theorists use deductive mathematics. Agent-based modelers use computer science and statistics. Some modelers use both techniques (Bendor, Diermeier, and Ting 2003). But there is no necessary divide between the two methodological approaches. To the extent these differences are seen as vast, it is often due to culture more than anything else.

For example, both game theorists and agent-based modelers capture selection using the replicator equation which assumes that the survival of a particular type of agent increases with its fitness relative to the average fitness in the population. Game theorists use this to solve for *evolutionarily stable strategies*, an equilibrium refinement. Agent-based modelers use the replicator dynamic to create interesting dynamic patterns. According to the replicator equation the probability of a strategy in time $(t + 1)$, p_{t+1}, is proportional to its probability in time t, p_t, times its relative fitness. Formally,

$$p_{t+1} = p_t(f_t/a_t)$$

where f_t equals the fitness of the strategy and a_t equals the average fitness in the population.

We might elevate the general mathematical proofs in this area of research above the computational based models. Were the models the same, we would, but they are not. The agent-based models tend to have more moving parts and include greater realism, and the results of agent-based models include the dynamics of the system (and not just an equilibrium, if one exists).

2.4 Networks, Spaces, Types, and Movement

Agent-based models also allow people and organizations to be placed in social structures or geographies. Cederman (2005) refers to these as *social forms*. A social form is a configuration of actors and rules for interaction along with the structures (possibly multiple) in which they are embedded. These social forms can include social networks and geography and are most commonly modeled using the tools provided by graph (p. 85) theory. Depending on what is being modeled, agents could potentially interact only with agents who either belong to their social network or who are geographically close, or they might only interact with other agents who are both socially connected and geographically close. Alternatively, a modeler might assume that agents have types, and that interactions are type contingent.

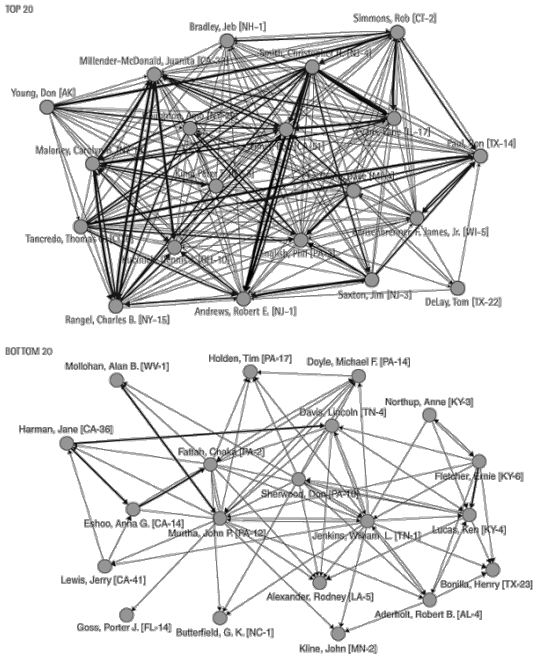
The existence of types, or what some call tags, leads to a surprising observation: *any model with agents situated in networks or in geographic space is equivalent to a model in which agents have types and interactions are type contingent*. This observation follows from the fact that an agent's type can include its location in a network. This logic also works in reverse. Types can be seen as points in geographic space. Consider for example Axelrod's culture model (1997). In that model, adjacent agents become less likely to interact the further apart their types are. If we map the agent types into Euclidean space, then this assumption becomes equivalent to saying that agents who are far apart do not interact. Type space becomes geographic space.

If geographic space can be recast as a set of types, then agents who move in space can be reconsidered as agents who change their type. Geographic location and types are thus factors in an agent's *contact structure*. The ability of agent-based models to include contingent contact structures is one of their great advantages. Consider two examples. Most mathematical models of markets assume that all economic activity takes place on the head of a pin, which lacks verisimilitude for many markets. Even in virtual markets not everyone is at the computer terminal at all times. And, in markets, distance matters. It can provide rents (Kirman 1997), for example. Similarly, mathematical models of disease often make an assumption of random mixing. These models in effect assume that

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

some random sample of people flies to O'Hare airport. These people interact—some get the disease and some don't—then they fly home. The next week the process repeats itself. In contrast, agent-based models of disease transmission can include geographic and social networks as well as constraints on network interactions (such as flight schedules), and make more accurate predictions than previous models (Epstein et al. 2006).

These contact structures or social forms need not be hardwired into the model; rather, they can emerge from the micro-level agent interactions (Lazer 2001). Cederman's (1997) model of the emergence of nation states provides a good example of emergent contact structures, as does Lustick, Miodownik, and Eidelson's (2004) model of identity formation, Bhavnani's (2003) model of ethnic conflict, and Axelrod et al.'s (1995) model of alliances.



Click to view larger

Fig.4.1. Most and least connected legislators in the 108th House, taken from Fowler (2006). These graphs show connections between the 20 most connected and 20 least connected members of the 108th House. Connections between these two groups and the other legislators are not shown.

These networks often have empirical relevance. For example, Fowler (2006) examined the Poole and Rosenthal model and asked whether or not co-sponsorship networks mattered in predicting member roll call votes. The null hypothesis would be that they do *not* matter; it is clearly possible that nominate scores explain all the systematic variance. Fowler created a measure of connectedness based on co- sponsorship (see Newman, Barabasi, and Watts 2006 for an overview). Using computational methods, he was then also able to create a visual representation of legislator behavior. Figure 4.1 displays the twenty most and least connected members of the 108th House. The pictures prove to be worth a thousand words. Fowler's empirical (p. 86) analysis reveals that connectedness mattered a great deal, above and beyond the information contained in the nominate scores. In fact, "changing the connectedness of the sponsor by 2 SDs (e.g., from the 95th to the 50th percentile—from very high to average) would change the final passage outcome in 16 percent of the House votes and 20 percent of the Senate votes."

(p. 87) 2.5 System Outputs: Complexity and Equilibrium

As we have mentioned, agent-based models are iterative. The modeler defines a set of agents, rules, and the framework or environment they exist in and lets it run. These micro-level interactions produce macro-level behavior in Schelling's sense. The macro-level behavior in turn alters the micro-level interactions that in turn produce new macro-level behavior. These sequences of micro to macro transitions are the output of the model, and as such, they may be investigated by statistical work that ties parameter settings to the results of the computational model through time.

These outputs may or may not be equilibria, however. Given the amount of complexity we see in the real world, we might expect social scientists to leap at the chance to use a tool that might explain or make sense of that

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

complexity. But attachment to purely deductive, equilibrium-based modeling runs deep. The presumption of equilibrium in game-theory models by social scientists rests on strong logical foundations: namely the fixed point theorems of Brouwer and Kakutani which can be used to show that Nash equilibria exist in a wide range of circumstances. These theorems state that a nice map (a continuous function in the case of Brouwer and an upper hemi continuous correspondence in the case of Kakutani) from a nice set (bounded and topologically convex) into itself has a fixed point: a point that maps into itself. Thus, if you swirl your coffee with a wooden stick, when the motion stops at least one coffee molecule will end up in the exact same place that it started.

These theorems can be used to prove the existence of equilibria by thinking of agents' actions as functions of agents' actions. We thus have a map from the space of agents' actions into the space of agents' actions. Formally, we have a mapping from X into X. Provided the set and the map are well behaved, the system has a fixed point: a set of actions such that every agent's action is optimal given the actions of the other agents. These fixed points need not be unique. A system can produce multiple equilibria. For this reason, many analyses rely on equilibrium refinements, which are usually not justified with the same sort of rigor as the rest of the apparatus of game theory.

Proving the existence of equilibria is a helpful first step, but existence alone should not be seen as sufficient for understanding a problem. An equilibrium should also be *stable* and *attainable*. A pencil balancing on its tip is an equilibrium, but not a stable one, and it is difficult to achieve in any case. It is easier to balance a pencil on its eraser, and still easier to balance a square pencil on its side (although not a round one). Stability can be local or global. An equilibrium is *locally stable* if small perturbations in the agent's action return to the equilibrium—this will certainly be true of the square pencil balanced on its side and it is even true for very small perturbations when it is balanced on its eraser, but not on its point. An equilibrium is *globally stable* if any perturbation leads back to that equilibrium. For an equilibrium to be globally stable it must be unique, thus a pencil does not have a global equilibrium

In social science, stability must be defined relative to a behavioral rule. Any local or global equilibrium can be made stable by creating a behavioral rule that says "return to the equilibrium"—friction terms are added to many models to achieve just this (p. 88) purpose, but rules of this sort presume that equilibria are stable. Strict adherence to the idea that actors are self-serving would suggest instead that actors should *naively best respond*, that they take the action that maximizes their payoffs holding everyone else's action fixed. But this rule often overshoots the equilibrium. As mentioned, the prevalence of overshooting is so widespread that most agent-based models rely on a replicator dynamic or a muted best response.

If a system has multiple equilibria, then each will have a basin of attraction given a behavioral rule. For a pencil balanced on its eraser, for example, this basin of attraction is very small since it will remain balanced for only small perturbations, but for a square pencil on its side, bigger perturbations will be possible. The size of the basin of attraction gets us to the issue of attainability. Even if an equilibrium exists, and even if it is stable, it may not be easily attainable. A system with N agents each of whom can take any one of A actions has A^N possible states, that is ignoring the agents having types which influence how they interact. If we add in that each agent can have one of T types, then the system has $A^N T^N$ possible states. We should not assume that a system converges to a particular equilibrium even if it is stable. Thus, if we drop 1,000 round pencils sharpened at both ends on a frictionless surface, we should not expect that they will ever come to an equilibrium—certainly not on their tips and probably not on their sides. They will just keep rolling and rolling around.

As a thought experiment, suppose we take a game-theoretic model and construct an agent-based implementation where the agents will adapt rules with the goal of achieving a best response function. The agents may reach equilibrium. If so, the agent-based model tells us something about equilibrium selection. However, the result instead might be a pattern or a long transient. Does this mean that the agent- based model failed? No. An equilibrium need not be a vortex that attracts the agents. The dynamics created by the behavioral rules could create complex dynamics that take a long time to settle down. If the system is churning, some agents must not be doing well or there may be complex patterns of interactions between agents. The crux of the argument is that if the system is not at an equilibrium, then some agent can improve its payoff by changing its action. But these improvements are defined relative to the status quo. If a lot of agents adapt simultaneously, then not all of the adaptations need to be improving.

Thus, the existence of an equilibrium may not alone be sufficiently informative. We may also want a model that

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

produces the equilibrium based on plausible rules detailing the agent behavior in question. Epstein (2006) refers to this as generative social science: If you didn't grow it, you didn't show it. Concepts from the differential equations/dynamic systems literatures are useful, including techniques such as a Lyapunov function. Many agent-based models with fixed rules have an associated Lyaponuv function. For example, under some assumptions, Schelling's segregation model does, as does the majority rule cellular automata model (Page 1997). Thus, we can prove theorems about the convergence to equilibrium of these agent-based models. Yet, even if we *cannot* prove a theorem, it is far better to have an agent-based model (or two or three; see Bednar and Page 2006) that produces the equilibrium, then to accept the equilibrium's attainability on faith.

(p. 89) 3 Concluding Remarks

Most mathematical analyses of game-theoretic models do not explore the stability and attainability of their equilibria and would be made richer by complementing them with agent-based models that explored those properties. Doing so would identify models in which symmetric equilibria have small basins of attraction or are unstable. We might also identify new equilibria with large basins of attraction that can be solved for analytically. For example, Page and Tassier (2004) used a computational model to show that the Groves Ledyard mechanism had literally billions of inefficient boundary equilibria and that under some assumptions these equilibria had enormous basins of attraction. As their paper shows, the process of creating a computational version of a model forces us to confront hidden assumptions related to boundaries, symmetry, timing, and information.

Agent-based models also permit richer empirical testing. In any situation where we expect to test our model against data, it is desirable to have equivalence classes in the parameter space of our model (de Marchi 2005). If, for example, a small perturbation of any of the parameters or assumptions in our model qualitatively changes the results of the model, it is difficult to imagine we would be able to conduct meaningful empirical tests of these results. All social science data is measured with noise, and to a large degree, we must be able to depend on the results of the model subject to such noise. And as the number of parameters or assumptions in a model increases, the chances we will go wrong (i.e. there is no equivalence class in the model's parameter space) also increases—the combinatorics of ever-larger parameter spaces makes statistical examinations of the results more and more difficult. This gloomy phenomenon is appropriately named the curse of dimensionality, and it applies with equal force to all models, whether they are computational, game-theoretic, or statistical.

Mathematical theorems reveal logical implications of assumptions. If A and B are true, then so is C. Given the often unrealistic assumptions employed by social scientists, the theorems may logically follow but one cannot assume they are of substantive interest. All too often, assumptions are chosen to produce an end result; e.g. we observe warfare or strikes, so we write a bargaining model where asymmetric or imperfect information can lead to these behaviors. Unfortunately, the space of deductive models that produce any given result is infinite; something more is needed to convince us that the particular model settled upon has value. Thus, mathematical models uncover fundamental truths about mathematical objects and not about the real world. The question then becomes not what is a general truth and what is a bunch of examples, but what do we do with them? How are these models—mathematical or computational—useful? To answer these questions we have to return to the question of why we built the models in the first place.

The essence of good modeling is leaving out unnecessary details. However, ignoring important features of real-world phenomena—such as timing or geographic space or the fact that most conflicts are not dyadic—hardly counts as general. The (p. 90) result is only general if these limiting assumptions do not matter. And often they do. Thus, we should focus on the quality of the results and their generality, rather than the method that is employed.

Returning now to the question of usefulness, the bright-eyed person in the front row will respond that one reason we build these models is to empirically test them. A mathematical theorem may say something like *if A is true then B is larger than C*, or it may say that the coefficient on D has a positive or negative value, or it may say that the interaction term E*F has a coefficient that differs significantly from zero. Whatever empirical hypotheses the model produces, we can subject these results to empirical falsification.

Contrast this with a finding from an agent-based model. Suppose that our agent- based model always produces a given phenomenon. We do not have a formal proof that this phenomenon occurs, but this does not mean that we

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

cannot subject the phenomenon to empirical testing (Laver 2005). Nor does it mean that the agent- based model doesn't generalize (Page 2006*a*). Most of our mathematical models of electoral politics consist of single-dimensional issue spaces. We might prefer having a multitude of examples in higher-dimensional spaces over an equilibrium result in one dimension. Thus, even though it is true that a theorem gives a result for infinite numbers of parameters or even functional forms, and that agent-based models can at best show something to be true for a finite number of cases, the finite cases can span substantial volume in the space of possible assumptions. In sum, if our criterion is empirical validity, the only reason that we'd place more faith in the mathematical model is the possibility of a counter-example to the agent-based model. Perhaps some future run of the model won't produce the phenomenon that we claimed to be true. If this happens, we might ask ourselves whether this makes the model any less valid if most of the time the model produces a phenomenon that also happens to occur in the real world.

Another criterion that we might use to compare agent-based models and mathematical models is whether they provide insights into social processes. Existence theorems rarely provide a lot of transparency. In contrast, an agent-based model may provide intuition. We can probe and test the model to see why it produces what it does. The effort might be especially worthwhile in those cases where a game- theory model relies upon synoptic rationality and the agent-based model makes a more limited rule-based or boundedly rational assumption. Or in cases where the game theorist assumes symmetry for convenience. If the game-theory model cannot account for the phenomenon, then we can ask (and test) whether it makes sense to presume that in this situation people would be more likely to use a simple rule rather than a complex rational calculation or if we can really presume symmetry. A third criterion we might use is how well these models help us to design institutions. Here again, we see the complementary nature of the two approaches. Mathematical models may give us general results, but we must be aware that those are mathematical and not social truths. A mathematical theorem establishes that the revenue generated by an auction does not depend on the auction rules provided that the highest (p. 91) valued bidder wins the object. Thus, an ascending bid auction, a first price auction, and a second price auction all yield the same revenue for the government. Experiments with human subjects and real-world tests show that the auction rules matter (Andreoni and Miller 1995; Hossain and Morgan 2006). Revenue varies depending on the auction design. Empirical regularities need not align with mathematical truth.

The ability of computational models to test the robustness of formal results would be reason alone to add them to our tool kits. But to restrict their use to be as extenders of what we can prove analytically is to ignore their potential to explore the implications of dynamic, spatial models with type-contingent interactions and diversity. Game-theoretic models must be kept analytically tractable and therefore simple. Mathematics proves most powerful when we assume one or two agents or an infinite numbers of agents. In the former case, math works because the problem is simple. In the latter, we can rely on limiting distributions and vanishing marginal effects. But mathematical analysis becomes bogged down in details when we have handfuls to hundreds of heterogeneous agents interacting in a complicated network. This space between two agents and an infinite number of agents happens to be where most social science takes place (Anderson, Arrow, and Pines 1988; Arthur, Durlauf, and Lane 1997; Miller and Page 2007).

For example, models that allow for coalitions of players to deviate are often mathematically intractable. These models lie in this region between two and infinity. If one looks at the formal literatures on bargaining, interstate conflict, or parliamentary behavior, one would come to the conclusion that all interactions are dyadic in nature. In part, this is because the *core*—outcomes for which no subset of agents can collectively improve by deviating—is often empty for interactions with more than two people. The intuition for why it is empty follows from the combinatorics of groupings. Fifty people can form over a million billion coalitions. The core treats each coalition as possible, which is why it has limited purchase. In a computational model, coalitions can emerge, but only limited numbers of them are likely given the constraints introduced by plausible behavioral rules and social or geographical networks. These models need not always focus on elites such as members of Congress. They can also focus on members of ethnic groups who choose whether or not to riot or commit acts of genocide (Backer and Bhavnani 2000; Kuran 1989).

To summarize, if we look at the world and ask ourselves whether it looks to be in equilibrium or to be complex, we cannot help but think it is complex (Jervis 1997). Yet the discipline prizes simple models, which are not likely to produce interesting, complex behavior (Page 2006*b*). Given that real-world complexity exists and given that we want to understand it, we should have a methodology that allows us to construct models that generate complexity. A methodology consists of both science and art. As a methodology, agent-based modeling should be seen as in its

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

infancy, its enormous potential limited only by the scientific and creative talents of its practitioners.

Bibliography

References

ANDERSON, P., **ARROW**, K., and **PINES**, D. (eds.) 1988. *The Economy as an Evolving Complex System*. Redwood City, Calif.: Addison-Wesley.

ANDREONI, J., and **MILLER**, J. 1995. Auctions with adaptive artificial agents. *Games and Economic Behavior*, 10: 39–64.

ARTHUR, B. 1989. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99: 116–31.

ARTHUR, W. B., **DURLAUF**, S., and **LANE**, D. 1997. *The Economy as an Evolving Complex System II*. Reading, Mass.: Addison-Wesley.

—— **HOLLAND**, J. H., **LE BARON**, B., **PALMER**, R. G., and **TAYLOR**, P. J. 1997. Asset pricing under endogenous expectations in an artificial stock market. In *The Economy as an Evolving Complex System II*, ed. W. B. Arthur, S. Durlauf, and D. Lane. Redwood City, Calif.: Addison-Wesley.

AXELROD, R. 1997. The dissemination of culture: a model with local convergence and global polarization. *Journal of Conflict Resolution*, 41: 203–26.

—— **MITCHELL**, W., **THOMAS**, R. E., **BENNETT**, D. S., and **BRUDERER**, E. 1995. Coalition formation in standard-setting alliances. *Management Science*, 41: 1493–508.

BACKER, D., and **BHAVNANI**, R. 2000. Localized ethnic conflict and genocide: accounting for differences in Rwanda and Burundi. *Journal of Conflict Resolution*, 44: 283–307.

BALLARD, D. 1997. *An Introduction to Natural Computation*. Cambridge, Mass.: MIT Press.

BARON, D., and **FEREJOHN**, J. 1989. Bargaining in legislatures. *American Political Science Review*, 83: 1182–202.

BEDNAR, J., and **PAGE**, S. 2006. Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Rationality and Society*, 19: 65–97.

BENDOR, J., **DIERMEIER**, D., and **TING**, M. 2003. A behavioral model of turnout. *American Political Science Review*, 97: 261–80.

BHAVNANI, R. 2003. Adaptive agents, political institutions, and civic traditions in modern Italy. *Journal of Artificial Societies and Social Simulation*, 6.

BINMORE, K. 1996. A note on backward induction. *Games and Economic Behavior*, 17: 135–7.

CEDERMAN, L. E. 1997. *Emergent Actors in World Politics*. Princeton, NJ: Princeton University Press.

—— 2005. Computational models of social forms: advancing generative process theory. *American Journal of Sociology*, 110: 864–93.

COHEN, M., and **AXELROD**, R. 1984. Coping with complexity: the adaptive value of changing utility. *American Economic Review*, 74: 30–42.

DE MARCHI, S. 1999. Adaptive models and electoral instability. *Journal of Theoretical Politics*, 11: 393–419.

—— 2003. A computational model of voter sophistication, ideology and candidate position-taking. In *Computational Models in Political Economy*, ed. K. Kollman, J. Miller, and S. Page. Cambridge, Mass.: MIT Press.

—— 2005. *Computational and Mathematical Modeling in the Social Sciences*. Cambridge: Cambridge University

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

Press.

DIERMEIER, D., and KREHBIEL, K. 2003. Institutionalism as a methodology. *Journal of Theoretical Politics*, 15: 123–44.

—— and MORTON, R. 2004. Proportionality versus perfectness: experiments in majoritarian bargaining. In *Social Choice and Strategic Behavior: Essays in the Honor of Jeffrey S. Banks*, ed. D. Austen-Smith and J. Duggan. Berlin: Springer.

(p. 93) DURLAUF, S. 2001. A framework for the study of individual behavior and social interactions. *Sociological Methodology*, 31: 47–87

ENSLEY, M., TOFIAS, M., and DE MARCHI, S. 2007. Electoral selection and the congressional incumbency advantage: an assessment of district complexity. Presented at the American Political Science Association meetings.

EPSTEIN, J. 2006. *Generative Social Science*. Princeton, NJ: Princeton University Press.1

—— and AXTELL, R. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, Mass.: MIT Press

—— GOEDECKE, D. M., YU, F., MORRIS, R. J., WAGENER, D. K., and BOBASHEV, G. V. 2006. Pandemic influenza and the value of international travel restrictions. CSED Working Paper No. 46.

FOWLER, J. 2006. Connecting the Congress: a study of cosponsorship networks. *Political Analysis*, 14: 456–87.

—— and SMIRNOV, O. 2005. Dynamic parties and social turnout: an agent-based model. *American Journal of Sociology*, 110: 1070–94.

FRÉCHETTE, G., KAGEL, J. H., and MORELLI, M. 2005. Gamson's law versus non-cooperative bargaining theory. *Games and Economic Behavior*, 51: 365–90.

GENTLE, J. 2002. *Elements of Computational Statistics*. New York: Springer-Verlag.

HOLLAND, J., and MILLER, J. 1991. Artificial agents in economic theory. *American Economic Review Papers and Proceedings*, 81: 365–70.

HOSSAIN, T., and MORGAN, J. 2006. &Plus shipping and handling: revenue (non)equivalence in field experiments on eBay. *Advances in Economic Analysis and Policy*, 6.

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

JUDD, K. 1997. Computational economics and economic theory: complements or substitutes? *Journal of Economic Dynamics and Control*, 21: 907–42.

—— 1998. *Numerical Methods in Economics*. Cambridge, Mass.: MIT Press.

KIRMAN, A. 1997. The economy as an interactive system. In *The Economy as a Complex Evolving System II*, ed. W. Arthur, S. Durlauf, and D. Lane. Reading, Mass.: Addison-Wesley.

KOLLMAN, K., MILLER, J., and PAGE, S. 1992. Adaptive parties in spatial elections. *American Political Science Review*, 86: 929–37.

—— —— —— 1997. Political institutions and sorting in a Tiebout Model. *American Economic Review*, 87: 977–92.

KURAN, T. 1989. Sparks and prairie fires: a theory of unanticipated political revolution. *Public Choice*, 61: 41–74.

LAVER, M. 2005. Policy and the dynamics of political competition. *American Political Science Review*, 99: 263–81.

—— and SERGENTI, E. 2007. Agent-based models of party competition: analysis and/or exploration? Presented at the American Political Science Association meetings.

—— DE MARCHI, S., and MUTLU, H. 2007. Bargaining in N-party legislatures. Presented at the *American Political Science Association* meetings.

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

LAZER, D. 2001. The coevolution of individual and network. *Journal of Mathematical Sociology*, 25: 69–108.

LEOMBRUNI, R., and **RICHIARDI**, M. 2005. Why are economists sceptical about agent-based simulations? *Physica A*, 355: 103–9.

LODGE, M., **KIM**, S., and **TABER**, C. 2004. A computational model of political cognition: dynamics of candidate evaluation. Presented at the Midwest Political Science Association meetings.

(p. 94) **LUSTICK**, I., and **MIODOWNIK**, D. 2000. Deliberative democracy and public discourse: the agent-based argument repertoire model. *Complexity*, 5: 13–30.

—— and **EIDELSON**, R. J. 2004. Secessionism in multicultural states: does sharing power prevent or encourage it. *American Political Science Review*, 98: 209–29.

MARTIN, A., and **QUINN**, K. 2002. Dynamic ideal point estimation via Markov Chain Monte Carlo for the US Supreme Court. *Political Analysis*, 10: 134–53.

MILLER, J. 1998. Active nonlinear tests ANTs of complex simulations models. *Management Science*, 44: 820–30.

—— and **PAGE**, S. 2007. *Complex Adaptive Systems*. Princeton, NJ: Princeton University Press.

NEWMAN, M., **BARABASI**, A.-L., and **WATTS**, D. J. 2006. *The Structure and Dynamics of Networks*. Princeton, NJ: Princeton University Press.

PAGE, S. E. 1997. On incentives and updating in agent-based models. *Computational Economics*, 10: 67–87.

—— 2006*a*. Agent-based models. In *The New Palgrave Encyclopedia of Economics*, ed. S. Durlauf and L. Blume. New York: Macmillan.

—— 2006*b*. Type interactions and the rule of six. *Economic Theory*, 30: 223–41.

—— 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press.

—— forthcoming. Uncertainty, difficulty, and complexity. *Journal of Theoretical Politics*.

—— and **TASSIER**, T. 2004. Equilibrium selection and stability for the Groves Ledyard Mechanism. *Journal of Public Economic Theory*, 6: 311–35.

PALMER, R., **ARTHUR**, W., **HOLLAND**, J., **LEBARON**, B., and **TAYLER**, P. 1994. Artificial economic life: a simple model of a stock market. *Physica D*, 75: 264–74.

POOLE, K. 2006. *Spatial Models of Parliamentary Voting*. New York: Cambridge University Press.

RIOLO, R., **AXELROD**, R., and **COHEN**, M. 2001. Evolution of cooperation without reciprocity. *Nature*, 414: 441–3.

SCHELLING, T. 1978. *Micromotives and Macrobehavior*. Toronto: George J. McLeod.

SNYDER, J. M., **TING**, M., and **ANSOLABEHERE**, S. 2005. Legislative bargaining under weighted voting. *American Economic Review*, 95: 981–1004.

STRAUSS, A. 2003. Applying integer programming techniques to find minimum integer weights of voting games. Department of Electrical Engineering and Computer Science, MIT.

STROGATZ, S. 2001. *Nonlinear Dynamics and Chaos*. Cambridge, Mass.: Perseus.

TESFATSION, L. 1997. How economists can get A–Life. In *The Economy as a Complex Evolving System II*, ed. W. Arthur, S. Durlauf, and D. Lane. Reading, Mass.: Addison-Wesley.

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: 1–19.

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014

WARWICK, P. V., and **DRUCKMAN**, J. N. 2006. The portfolio allocation paradox: an investigation into the nature of a very strong but puzzling relationship. *European Journal of Political Research*, 45: 635–65.

YOUNG, P. 2001. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.

Notes:

(1) For an application to the Supreme Court, see Martin and Quinn (2002).

(2) The formal statement of the theorem is as follows: "Let G be a 2 × 2 symmetric coordination game with a strictly risk dominant equilibrium, and let Q_N, ϵ be adaptive learning in the playing the field model with population size [N], complete sampling, and error [mutation] rate ϵ, $0 < \epsilon < 1$. For every $\epsilon' > \epsilon$, the probability is arbitrarily high that at least $1 - \epsilon'/2$ of the population is playing the risk dominant equilibrium when [N] is sufficiently large" (Young 2001, 76).

(3) For example, the adoption of competing technologies such as a Wintel computer versus an Apple.

(4) Informally, minimum integer weights are the smallest integer weights that generate a given set of coalitions. E.g. the three-person legislatures {1,1,1}, {2,2,1}, and {99,99,1} generate identical winning coalitions; {1,1,1} is the minimum integer weight representation.

(5) The most prominent school of thought on this issue is the *Empirical Implications of Theoretical Models* (EITM) approach to modeling. As put forth in the original 2002 National Science Foundation report, the goal of EITM is to rework the discipline so that the chasm between formal modelers and empirical researchers is bridged, with the hopes that this synthesis will lead to better models that have clearly testable empirical hypotheses.

(6) See, for example, the "Man vs. Machine" 2007 poker tournament (<http://poker.cs.ualberta.ca>). Expert humans narrowly defeated the Nash-equilibrium playing computer algorithms designed by the University of Alberta team, but significant limitations were imposed on play. The game was two-person, fixed limit hold-em; no limit, *n*-person, or more complex games such as Omaha would be vastly more challenging.

Scott de Marchi

Scott De Marchi is Associate Professor of Political Science, Duke University.

Scott E. Page

Scott E. Page is Professor of Political Science, Research Professor in Center for Political Studies, and Associate Director of Center for the Study of Complex Systems, University of Michigan, Ann Arbor.

PRINTED FROM OXFORD HANDBOOKS ONLINE (www.oxfordhandbooks.com). (c) Oxford University Press, 2014. All Rights Reserved. Under the terms of the licence agreement, an individual user may print out a PDF of a single chapter of a title in Oxford Handbooks Online for personal use (for details see Privacy Policy).
Subscriber: Duke University; date: 26 November 2014