Agent Based Models

Scott de Marchi[1]
Scott E Page[2]
Tue May 28 16:49:23 2013

Agent based models provide a methodology to explore systems of interacting, adaptive, diverse, spatially situated actors. Outcomes in agent based models can be equilibrium points, equilibrium distributions, cycles, randomness, or complex patterns; these outcomes are not assumed but instead emerge from the interactions of actors in the model. What occurs in a model depends on the assumptions and behaviors built into that 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, agent based models offer the potential to advance social sciences and to help us better understand our complex world.

Keywords:computational; complexity; game theory; social science.

[1] demarchi@duke.edu. Department of Political Science, Duke University
[2] spage@umich.edu. Center for the Study of Complex Systems, Departments of Political Science and Economics, University of Michigan.

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

In an ABM, agents receive inputs from their environment and take actions in response to those inputs; as such, many ABM's 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, or bottom-up (Epstein 2006, Richards 1997).

The outcomes in agent based models can be temporal: agent based models of two party electoral competition produce time series of vote totals for the parties (Kollman, Miller Page 1992, 1997, de Marchi 1999, Laver 2005, Fowler and Smirnov 1995, Smirnov and Fowler and 2007), they can be spatial, as in the case in Axelord's culture model (1997), and they can be classic equilibria, such as in Schelling's racial segregation model (1974).

Unlike much of the modeling work done in the social sciences, ABM's often concern themselves with 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 and Sergenti, 2011).

In the minds of many, agent based models are synonomous with complexity. This is an unfortunate conflation. Agent based models are a methodology, a tool that scientists from many disciplines use to construct models (Kollman, Miller and Page 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 agent based models 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 and Page 2012). Many scholars believe that the political world is becoming increasingly complex and that to progress, political science will need a deeper understanding of complexity (Jervis 1997, Cederman 1997). That may well be true. And if it is, agent based models, which are well suited to make sense of complex systems, may well grow in popularity, provided that appropriate standards are developed (Miller 1998, de Marchi and Page 2009, Axelrod, et al 1996).

The uptake of agent based modeling within political science has been steady and growing (Kollman and Page 2005). And, yet, information about and knowledge of agent based models among political scientists varies widely and misperceptions are common. As previously mentioned, many people incorrectly conflate agent base models 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.,

2007), which cannot be easily explained, described, or predicted (Mitchell 2009, Page 2011) become a natural are of research to apply agent based models.

There exist 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 agent based models 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).[3] Game theorists do devote substantial attention to the robustness of the equilibirum 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 backwards 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, Taber, and Lodge, 2010). Though some see rationality asssumptions as stronger than heuristic assumptions, this intuition ignores the fact that models that assume rational behavior must, by

[3] There is no natural way 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, lack verisimilitude, and the conclusions are overly responsive to these assumptions.

definition, be solveable. Therefore, rationality can only be assumed when the model's author has the capacity to derive what optimal behavior is.

On 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 non-cooperative game theorists – has most often been fruitful when deductive approaches have been combined with computational approaches (Billings, Papp, Schaeffer, and Szafron, 1998).[4]

Advocates of agent based approaches, among whom we count ourselves, emphasize the resonance between agent based models – 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. Political processes also produce a mix of equilibria, patterns, and seeming randomness (Palfrey, 2005).

ABM's offer the possibility of increasing the reach of existing models and to make 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),

[4] In games in 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 both limit the size of the extensive form by adding constraints (e.g., limiting the number of betting opportunities or the range in bet size) or by computing aproximations to the Nash (e.g., see the work on regret minimaztion in Zinkevich, et. al., 2007). In the social sciences, bargaining theory has also benefitted from a combination of non-cooperative game theory and computational approaches – see Ansolabehere (2005) and Laver, de Marchi, and Mutlu (2010). In this work, 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.

their ability to inclusion of space and networks (Newman 2003), their potential to link multiple domains (Epstein and Axtell 1996), their potential for greater agent heterogeneity (Page 2007, 2011), their production of outcome contingency, their emphasis on outcome robustness (Jen 1995), and their plasticitiy, 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). While these are serious concerns, they result from the fact that ABM's must be explicit about the choices that are made regarding timing, interactions, and adaptation. To be concrete let us take baragaining models as an example. As Laver, de Marchi, and Mutlu (2010) point out, non-cooperative models of bargaining in the Rubinstein and Baron-Ferejohn tradition assume that the reversion point to bargaining is for all players to receive a payoff of 0; for a strict sequence to be observed in offers rather than an open-outcry assumption[5]; for an exogenous, uniform, and independent selection mechanism to exist that chooses the proposer in each time period; and that agents adhere to stationarity. While the result of models in the non-cooperative 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.[6] 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 agent based models. We then provide a basic primer on how to construct

[5] 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.

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

an agent based model. More complete soures exist (Epstein and Axtell 1996, de Marchi 2005, Miller and 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 agent based models, we then compare and contrast them with traditional deductive techniques. We claim that political science would improve if modelers combine the two techniques but that agent based models also have stand alone value. We conclude with a discussion of empirical testing of agent based models 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 agent based models and their relevance to social scientific inquiry. Given the diversity of computational and agent based models, we will 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.

At the most basic level, agent based models 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 aproaches. Agent based models allow researchers to construct this micro to macro mapping and to test the empirical validity of the results produced (Fowler and Laver 2008).

In many cases, models that assume psychologically based behavioral rules may better fit experimental and empirical data than do rational actor models (Camerer 2003). ABMs can encode these behaviors and help us understand how they aggregate. This capacity to encode any behavioral rule leads some to see agent based models as opposed to the notion of rational choice theory and as wedded to the behavioral approach. This is a misrepresentation. Agent based models rely on behavioral rules. Those rules could be best response functions given some utility function and 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 of itself; rather, researchers resort to computational models when purely deductive models produce poor outcomes that are useful in generating results on substantive problems (Leombruni and Richiardi 2005, Tesfatsion 1997).

The literauture on turnout is one prominent example. Bendor, Diermeier, and Ting (2003) built a computational model where they demonstrate how a learning rule with more verisimilitude than purely deductive models accounts for positive turnout; Fowler (2006) 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, ABM's 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 and Ward, 2004). People segregate by race, income, religion and other characteristics, and (evidence suggests) individual behavior depends partly on that of people around them (Glaeser, Sacerdote, and Schienkman 1993). Further, political representation is often geographically based. These three stylized facts combined with with evidence from rich empirical case studies (Huckfeldt, Johnson, and Sprague, 2004) suggest that even if it's not true that all politics are local, what goes on spatially probably does matter sufficiently for it to be included in our models.

Relatedly, agent based models also can also include networks (Newman 2003). These could be policy networks (Carpenter, Esterling and Lazer 2003), interaction networks (Lazer 2005) or friendship networks. Given the growth of network measures and statistical techniques, testing network effects has become much easier (Fowler 2006a, 2006b).

Third, agent based models 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 agent based model, general equilibrium spillover

effects can extend to the behavioral realm and cross domains. ABM's with these features have been used extensively for research into policy resistant problems by the National Institutes of Health (under the Systems Science RFA) 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 and Page 2007). Agents who are cooperative or trusting in one interaction may be more likely to be cooperative or trusting in another, a finding that has been shown in experiments with human subjects (Bednar et. al. 2012).

Cross domain spillovers can be thought of as multiple general equilibriums 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 economics decisions effect global temperatures and global temperatures, in turn, effect epidemics. Agent based models 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 or politics only model might not be capable of producing model that resemble real world, multi-purpose networks.

We should be clear that there are limits to realism. If a model includes too many domains, it may become as 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 and Miodownik 2000, Lustick, Miodownik and Eidelson 2004). One could argue,

for example, that a model that combined transportation systems with demographic data or that linking 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, agent based models can include heterogeneity in not only 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. Heterogeneity for realism's sake is nice but is not the end in itself. The key question that needs to be asked is: when does heterogeneity cancel out, so that only the mean matters, and when do the actual distributions matter? Agent based models can help to answer those questions (Page 2011, Bruch and Mare 2006)

Fifth, as already mentioned, agent based models 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 doesn't equilibriate? Or what if time to equilibrium is so long that the equilibrium isn't of much scientific interest? In such instances, equillibrium analysis may be misguided, and agent based models can add substantial

value (Laver and Sergenti, 2011).

Sixth, analysis of agent based models tends to focus on outcome robustness, i.e. can the system keep doing what it needs to do (Jen 2005) as opposed to the existence and efficiency of equilibria. For example, agent based models can be used to assess the robustness to the entrance of new types, the exit of existing types, changes in environmental parameters, and even to the rules that agents follow (Laver and 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 occassional large failure can far outweight 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 and Miller 1995)

Finally, agent based models allow for exploring neighborhoods of models (de Marchi and Page, 2005). Any model makes a set of assumptions. Those assumptions can be altered individually and in pairs to create a set of neighboring models. Consider Axelrod's culture model (2007). Neighboring models can assume different behaviors on the part of agents, different network topologies, and 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. When the set of neighboring models is small, they can all be examined. Larger sets can be sampled and analysed 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 agent based models. 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 agent based models. We conclude with a brief description of the types of outcomes that agent based models produce starting with a discussion of equilibrium outcomes and then turning to other types of outcomes less familiar to social scientists.[7]

Agent Attributes and Behaviors

Agent based models 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 may 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 agent based models, attributes also often include the *location* of the agent. We can formally define the set of attributes as follows:

[7] This section focuses on the concepts but not the particular tools involved. There is a wide range of software packages utilized by researchers and choosing the "best" set of tools is a non-trivial enterprise. As a starting point, however, python with its extensions to call R functions and the open-source scientific library scipy is a good starting point.

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

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.

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

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

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

Typically, an agent based model will have between two and ten thousand agents depending on what's being modelled. Models of spatial party competition in the US may have

two agents. Models of riots may have thousands. Though agent based models with millions of agents can be built, they demand substantials computational time. Of late, large models have grown in popularity given advances in data.

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

Agent based models 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 assuming that each agent interacts with all other agents or that agents interact with others randomly (this is formally known as *random mixing*), agents interact with others contingent on their states. When implemented in the computer code, whether or not two agents interact depends on the vector of attribute values. Whether those values represent location, social standing, gender, or ideology makes no difference to the computer.

Thus, from a computational perspective, constructing agents with geographic locations who interact locally can be equivalent 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 perspecitve 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, agent based models 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, Diermeier, and Ting 2003), or cognitive psychology's heuristic shortcuts (Kim, Taber, and Lodge, 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 and 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 Nash 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 agent based model 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 and D'Andria 1998). When sequential timing fails or does not make sense given the context, agent based models call 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

Agent based models 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 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, Bonhoeffer, and May 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 do defect (D) but your opponent does better if you choose to cooperate (C). The social 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, Bonhoeffer, and May (1993) assume the following behavior: Each agent plays a 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 spead in the popultion. 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 and 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 does 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.

Agent based models 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 agent based models 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. They did not include real geographic space.

Agent based models can include houses, schools, workplaces, and transportation systems. They can assign each individual 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 to whom they come into contact. Those behaviors, including key features such as social networks, can be approximated

from data. In this way, agent based models can be used to predict how fast a disease will spread or to infer the virulence of a disease that's begun to spread. They can also be used to evaluate intervention scenarios. It should come as no surprise that these high fidelity agent based models of epidemics have become the status quo for researchers (Joshua, Goedecke, Yu, Morris, Wagener, and Bobashev, 2006).

As mentioned in the introduction, this potential representational richness is both a blessing and a curse. 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

Agent based models generate equilibria from the bottom up (Epstein 2006) 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 the 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's true that the mixed strategy is unstable. If either player choose 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 the pure randomization, then the mixed strategy will be stable and any deviation will be corrected. Thus, stablity is a function of the behavioral rule.

In an agent based model, 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'll 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's 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 and 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 zero percent overlap (Golman and 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's just that the learning rules drive the population in different directions. This obvious result is that agent behavior affects where the system ends up, providing a primary reason for constructing agent based models if one cares about generating empirically testable results.

One final point on equilibrium in agent based models. Many of the equilibria produced are *statistical equilibria* or what we have called *equilibrium distributions.* By that we mean that the agent based model 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 apples and pear trees produce fruit at equal rates (on average) and if agent preferences weight the two goods equally, a statistical equilibrium will emerge in which approximately equal numbers of apples and pears are produced. 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 payss the same identical price for each good (Epstein and Axtell 1996).

Complexity and Emergence

We now turn to the potential for agent based models to produce non equilibrium 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 includes higher order functionalities. For example, an efficient allocation of goods emerge in an ideal market even though agents pursue individual interests.

Agent based models of exchange often produce complex outcomes. Kirman and Vriend (2001) 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 produces a complex web of connections and reputations.

Agent based models 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 ABM's are well suited for including interdependencies and behavioral feedbacks, they are often used to study environments that produce non normally distributed outcomes.

The most common method for analyzing aggregate outcmes is to plot the distribution of event sizes and to compare it to a normal distribution. For example, when one builds an agent based model 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 agent based models differ in spirit from mathematical modeling in four important ways. First, in the construction of an agent based model, analytic tractability is neither a primary nor a background concern. Instead, the focus is on identifying the most relevant parts of the domain and in 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. While simplicity and understanding one's own model are important, there is a recognition 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 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 tends 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 difference 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 and Sandholm, 2006).

Third, agent based models, like dedutive 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.[8]

Fourth, as mentioned earlier, even though agent based models and mathematical models are capable of studying the same phenomena, practioners 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 computatonal models. Often ABM 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 agent based model can have give insights into what happens.

The flexibility of ABM's explains why they have more detail and, as a result, why they are more likely to produce nonequilibrium outcomes. ABM tend to include geographic space, networks, adaptive learning rules, and heterogeneity that they produce different types of

[8] It's worth noting that even purely deductive approaches face related issues. E.g., 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 than the rules utilized in computational models.

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 agent based model 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 depends on network characteristics. Our point is that this typically doesn't happen. It doesn't because networks can be difficult to analyze analytically, though advances are being made Jackson (2008). Increasingly, papers that analyze networks combine mathematical and computational approches.

A natural question to ask is when to use an agent based model 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 agents rules) lead to an equilibria or whether cycles or complexity result. But, this is the first step in understanding real world phenomena and a natural question is to ask 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 computatonal setting. Game theoretic models rely on an uncertainty framework, while agent based models assume complex environments (Page 2008). To give one specific example, mathematical models of incomplete information rely on *generated signals*, i.e. draws from probability distributions. Computational models often employ *interpreted signals,* i.e. the 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 more typical of compuational model. We don't 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 themselve to mathematical analyis. Random subsets are easily generated computationally. Surprisingly, these two approaches lead to quite distinct correlational properties of signals Hong and Page (2009). Therefore, what is true in a mathematical model may not be true in a computional 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 would 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 its tractable. Or, the game can be converted into an agent based model and one can push for more general results or test whether or not the assumptions involved are robust to changes. We've made this sound a little too simple. In constructing the agent based model, 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 and Krehbiel 2003).

In an agent based model, any one of a number of learning rules -- either individual or population based -- might be used. Many first time users of agent based models 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. Place on candidate randomly at plus five and the other at minus five. The best response of the first candidate will be to move to plus five minus epsilon and for the other candidate to move to minus five plus epsilon. 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 agent based model, 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.

In brief, the richer the environment under study (either because the environment is complex, or because it involves embedded games), the more likely that an agent based model will be necessary. This is not to say that agent based models should replace games. To the contrary, the findings from the richer agent based models should be compared against and contrasted with findings from game theory models. Agent based versions of game theoretic model should also be thrown into the mix as they enable an analysis of attainability and stability of equilbria. This is where real learning will occur. Lest the reader think we're piling on the work, agent based models are easy to write, analysis is often visual, and data can be produced in any amount. They're 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). Consider a model analyzed using both a game theoretic model and implemented in an agent based model (Bednar and Page 2006). The game theory 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 too show that an equilibrium is unique, or the unique outcome that satisfies some refinement criteria.

The agent based model can be used to investigate whether a class of learning rules generates the derived equilibrium. If, for example, in each of five hundred runs of the model, the agents reach equilibrium or something near to it, then we can say with statistical confidence that the assumed model generates outcomes in some neighborhoods of the equilibrium. 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 theorists *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 ABM's are quite large compared to 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 also because empirical work allows computational researchers to choose parameter values based on data rather than treating them as nuisance parameters.

In bargaining, Laver, de Marchi, and Mutlu (2010) 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, Golder, and Siegel (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, 2006b) has demonstrated how cosponsorship networks between legislators affects voting behavior. Laver and Sergenti (2011) have provided models

that describe party birth and death and evolving policy locations.

In international relations, Weidmann and Salehyan (2013) built an ABM instantiated with GIS-coded data and shown that sorting is the most plausible reason for the reduction in violence in Iraq. Gartzke and Weisiger (2012) build an ABM to show how the prevalence of different identities in the interstate system creates new cleavages for conflict between nations.

While there are other examples of applications of ABM's, 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, Kollmlan Miller, and Page (1992) model has been shown to have empirical support: there is an incumbent advantage and it depends on the complexity of the electoral landscape.[9] This work is important because it shifts the frame for elections from equilibrium seeking behavior, where there are an abundance of chaos results (e.g., the general non-existence 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 KMP's model directly is difficult, other than to state the obvious: incumbents, even one controls for challenger quality and resources, do not always lose reelection. But is there a more direct test that computational modeling and the emphasis on optimization rather than equilibrium-seeking behavior is appropriate?

Fortunately, the answer is "yes". Ensley, Tofias, and de Marchi (2009) and Wichowsky

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

(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. Not only have incumbents learned (optimized) from prior campaigns, the mere fact that they have won at least once signifies they (perhaps randomly) have arrived at an effective platform (see Alchian, 1950, for similar analysis of the firm).

Laver (2005) has provided additional support for KMP's 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., hillclimbing is one algorithm) with human decision-making in complex environments.[10] Laver's work echoes that of Kollman, Miller, and Page (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, agent based models have much to offer for the study of politics. They enable us to construct higher dimensional environments and richer behaviors including networks, coalition seeking (Axelrod, et al 1995), and learning. They permit us to to ask a larger set of questions

[10] Laver's analysis is instructive of how testing ABM differs from testing of game theoretic models. Agent based models vary in the assumptions they make about behavior. Therefore, it's 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.

and to consider alternative types of outcomes, including complexity. At the same time, agent based models 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.

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

We have reason to be hopeful. Applications of agent based models to the physical, chemical, medical, epidemiological, biological, and ecological sciences provide evidence for skeptical optimism. In each of those fields, agent based models 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 notwithstaniding, our optimism for the ultimate contribution to the study of politics is constrained. Ironicaly, our reservations are rooted in the complexity of political systems. In systems where agent based models have made large 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 posses 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, agent based models may

prove more than up to the task. In contexts that admit rich behavioral repertoires, modeling becomes more problematic.[11]

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

Though they occupy the same sand box, agent based modelers and game theorists build quite different castles in the sand. That's partly because they're playing in different parts of the box. Game theorists stick to the edges. Most game theory models assume either two or three agents or an infinite number. Most analytic models assume full rationality. Those that don't tend to have all agents following the exact same rule. Game theorists take to the extremes not because they believe all interesting problems belong lie as 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 outcomess occur with moderate numbrs of agents interacting in richer models.

To sum up, we would say that nearly a quarter century after their introduction, agent based models' contribution to politics has been modest, but important -- we have highlighted many influential papers from across the discipline. What we have not fully credited is how

[11] 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 theory models produce equilibria, albeit sometimes in mixed strategies only.

agent based models 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 the world: of recognizing complexity, seeing the difference between stability and robustness, of distinguishing between tipping points, path dependence, and chaos; and enabling us to see the value and limitiations of formal modeling from a broader perspective.

References

Alchian, Armen A. (1950) "Uncertainty, evolution, and economic theory." The Journal of Political Economy 58.3: 211-221.

Anderson, P. (1972). “More is Different.’’ Science, 177, 393-396.

Andreoni, James, and John H. Miller (1995). “Auctions with Adaptive Artificial Agents” *Games and Economic Behavior* 10:1 pp 39-64.

Ansolabehere, Stephen, et al. "Voting weights and formateur advantages in the formation of coalition governments." *American Journal of Political Science* 49.3 (2005): 550-563.

Axelrod, Robert (1997). *The Complexity of Cooperation* Princeton University Press. Princeton.

Axelrod, R, W. Mitchell, R.E. Thomas, D. S. Bennett, and E. Bruderer (1995). "Coalition Formation in Standard-Setting Alliances," *Management Science*, 41, pp. 1493-1508.

Axelrod, R, R. Axtell, J. Epstein, and M. D. Cohen (1996). "Aligning Simulation Models: A Case Study and Results,", *Computational and Mathematical Organization Theory*, 1, pp. 123-141.

Axelrod, Robert and Michael Cohen (2000). *Harnessing Complexity Organizational Implications of a Scientific Frontier* Free Press.

Bak, Per (1996). *How Nature Works: The Science of Self Organized Criticality* Springer Verlag, New York.

Ballard, Dana. (1997). *An introduction to natural computation*. MIT University Press. Cambridge, MA.

Bankes, S. (2002). Agent-based modeling: a revolution? *Proceedings of the National Academy of Sciences* 99, 7199Ð200.

Baron, David P. and John A. Ferejohn (1989). “Bargaining in Legislatures” *American Political Science Review* 89 pp 1181–1206.

Bednar, Jenna (2009). *The Robust Federation: Principles of Design*, Cambridge University Press.

Bednar, Jenna and Scott E Page (2007). “Can Games Theory Explain Culture?” *Rationality and Society* pp 65-97.

Bednar, Jenna, Yan Chen, and Xiao Liu and Scott E Page (2012). ``Behavioral spillovers and cognitive load in multiple games: An experimental study." *Games and Economic Behavior* 74:1 pp 12-31.

Bendor, Jonathon, Daniel Diermeier and Michael Ting (2003). “A Behavioral Model of Turnout” *American Political Science Review* Vol 97 No 2 pp 261-280.

Billings, Darse, Denis Papp, Jonathan Schaeffer, and Duane Szafron (1998). “Poker as a Testbed for Machine Intelligence Research” Proceedings of AI'98, *Canadian Society for Computational Studies in Intelligence.* pp 387-389.

Bendor, J., D. Diermeier, and M. Ting (2003). “A Behavioral Model of Turnout.”American Political Science Review 97 (2): 261–80.

Brown, G.W. (1951). "Iterative Solutions of Games by Fictitious Play." pp 374-376 in T.C. Koopmans (ed.) *Activity Analysis of Production and Allocation* New York: Wiley.

Bruch, Elizabeth and Robert Mare (2006). “Neighborhood Choice and Neighborhood Change. *American Journal of Sociology* 112: 667-709.

Camerer (2003). *Behavioral Game Theory: Experiments in Strategic Interaction* Princeton University Press.

Cederman, L.E. (1997). *Emergent Actors in World Politics*, Princeton University Press.

Cederman, L.E (2005). “Computational models of social forms: Advancing generative process theory” *American Journal of Sociology* 110 (4): 864-93.

Conitzer, V., & Sandholm, T. (2006). “Computing the optimal strategy to commit to. In Proceedings of the 7th ACM conference on Electronic commerce”. ACM.

Cutler, Josh, Scott de Marchi, Max Gallop, Florian Hollenbach, Michael Laver, and Matthias Orlowski (2013). “Cabinet formation and portfolio distribution in European multi-party systems”, working paper, 2013.

de Marchi, Scott (1999). “Adaptive Models and Electoral Instability” *Journal of Theoretical Politics* 11(3): 393Ð419.

de Marchi, Scott and Scott Page (2009). “Agent Based Models” in Oxford Handbook of Political Science Methodology. Cambridge University Press.

de Marchi, Scott (2005). *Computational and Mathematical Modelling in the Social Sciences*. Cambridge University Press.

Diermeier, Daniel and Keith Krehbiel (2003). “Institutionalism as Methodology” *Journal of Theoretical Politics* 15(2): 123Ð144.

Ensley,M, MW Tofias, S de Marchi (2009). “District complexity as an advantage in congressional elections” *American Journal of Political Science. CITE*

Epstein, Josh (2006). *Generative Social Science: Studies in Agent-Based Computational Modeling* Princeton University Press.

Epstein, Joshua M., D. Michael Goedecke, Feng Yu, Robert J. Morris, Diane K. Wagener, and Georgiy V. Bobashev (2006). Pandemic Influenza and the Value of International Travel Restrictions. CSED Working Paper No. 46 2007.

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

Fowler, James (2006). “Habitual Voting and Voter Turnout” The Journal of Politics.

Fowler, James. (2006a). “Connecting the Congress: A Study of Cosponsorship Networks.” *Political Analysis* 14 (4): 456-487.

Fowler, James. (2006b). “Legislative Cosponsorship Networks in the U.S. House and Senate’’ *Social Networks 28* (4): 454-465.

Fowler, J, and Oleg Smirnov. (2005). "Dynamic Parties and Social Turnout: An Agent-Based Model." American Journal of Sociology 110:4.

Fowler, James and Michael Laver (2008). ``A Tournament of Party Decision Rules.’’ *Journal of Conflict Resolution* 52 (1): 68-92.

Erik Gartzke and Alex Weisiger (2013). Fading Friendships: Alliances, Af□nities and the Activation of International Identities. British Journal of Political Science, 43, pp 25-52

Gaylord, Richard and Louis J. D’Andria (1998). *Simulating Society: A Mathematica Toolkit for Modeling Socioeconomic Behavior* Springer, London.

Gell Mann, Murray, (1994). *The Quark and the Jaguar: adventures in the simple and the complex*. W.H. Freeman. New York.

Golder, Matt, Sona Golder, and David Siegel (2012). “Modeling the Institutional Foundation of Parliamentary Government Formation.” *Journal of Politics*.

Golman, Russell, and Scott E Page (2009). ``Basins of Attraction and Equilibrium Selection Under Different Learning Rules" *Evolutionary Economics* 20: 49-72.

Golman, Russell, and Scott E Page (2010). ``Individual and Cultural Learning in Stag Hunt Games With Multiple Actions'' *Journal of Economic Behavior and Organizations* 73(3): 359-376.

Holland, John and John Miller (1991). “Artificial Agents in Economic Theory” American Economic Review papers and proceedings, 81 pp 365-370.

Hong Lu, and Scott Page (2009). ``Interpreted and Generated Signals" *Journal of Economic Theory* 144: 2174-2196.

Huckfeldt, R. Robert, Paul E. Johnson, John D. Sprague (2004). *Political Disagreement: The Survival of Diverse Opinions Within Communication Networks* Cambridge University Press, Cambridge, England.

Jackson, Matt (2008). *Social and Economic Networks* Princeton University Press, Princeton N.J.

Jen, Erica (2005). *Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies*, Oxford University Press, Oxford, UK.

Jervis, Robert (1997). System Effects: Complexity in Political and Social Life Princeton University Press.

Judd, Ken (1997). "Computational Economics and Economic Theory: Complements or Substitutes?" *Journal of Economic Dynamics and Control*.

Kauffman, Stuart (1993). *The origins of order: Self organization and selection in evolution*. Oxford University Press.

Kim, S, CS Taber, M Lodge (2010). “A computational model of the citizen as motivated reasoner: Modeling the dynamics of the 2000 presidential election - Political Behavior.

Kirman, Alan (1997). “The Economy as an Interactive System” in *The Economy as a Complex Evolving System II* W. Brian Arthur, Steven Durlauf, and David Lane eds. pp 491–533. Addison Wesley, Reading, MA.

Kollman, Ken (2012). “The Potential Value of Computational Models in Social Science Research.” in Harold Kincaid, ed. *Oxford Handbook in Philosophy in the Social Sciences.* New York: Oxford University Press.

Kollman, Ken, John Miller, and Scott Page (1992). “Adaptive Parties in Spatial Elections,” *American Political Science Review*, 86:929–37.

Kollman, Ken, John Miller, and Scott Page (1997). ``Political Institutions and Sorting in a Tiebout Model.'' *American Economic Review.* 87:977-92.

Kollman, Ken, John Miller, and Scott Page, (2003). (eds.) *Computational Models of Political Economy* MIT Press.

Kollman, Ken and Scott Page (2005). “Computational Political Economy” in *The Handbook of Computational Economics*, Leigh Tesfatsion and Kenneth Judd (eds) pp 1433-1463. North Holland

Lamberson, P.J. and Scott E Page (2012). “Tipping Points”, *Quarterly Journal of Political Science* Vol. 7:No 2, pp 175-208

Laver, Michael, · Scott de Marchi and · Hande Mutlu, (2010). "Negotiation in legislatures over government formation," Public Choice.

Laver, Michael and Ernest Sergenti (2011). *Party Competition: An Agent-Based Model*, Princeton.

Laver, Michael (2005). “Policy and the dynamics of political competition.” *American Political Science Review* 99:2 (May) 263-281.

Laver, M. (2007). "Agent-based Models of Party Competition: Analysis and/or Exploration?" Presented at the American Political Science Association meetings.

Lazer, D. (2001). ``The Coevolution of Individual and Network,'' *Journal of Mathematical Sociology, January.*

LeBaron, Blake (2001) ``Stochastic volatility as a simple generator of apparent financial power laws and long memory'', *Quantitative Finance,* 1:621-631

Leombruni, R. and Richiardi, M. (2005). ``Why are economists sceptical about agent-based simulations?'' *Physica A* 355, 103Ð9.

Lodge, Milton., Kim, Sung-youn. and Taber, Charles. (2004). A Computational Model of Political Cognition - Dynamics of Candidate Evaluation Paper presented at the annual meeting of the The Midwest Political Science Association, Palmer House Hilton, Chicago, Illinois.

Lustick, Ian and Dan Miodownik. (2000). ``Deliberative Democracy and Public Discourse: The Agent Based Argument Repertoire Model.'' *Complexity* Vol. 5, no.4 (2000) pp. 13-30.

Lustick, Ian, Dan Miodownik and Roy J. Eidelson, (2004). "Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It?" *American Political Science Review*, Vol. 98, no. 2: 209-229.

Miller, John (1998). ''Active nonlinear tests ANTs of complex simulations models.'' *Management Science* 44, 820-830.

Miller, John and Page, Scott (2007). *Complex Adaptive Systems: Computational Models of Social LIfe* Princeton University Press.

Mitchell, Melanie (2009). *Complexity: A Guided Tour*. Oxford University Press, Oxford, UK.

Newman, M J (2003). “On the Logic and Structure of Complex Networks, *SIAM Review*, 45, 167-256.

Nowak, Martin and Robert M. May (1992). "Evolutionary Games and Spatial Chaos", *Nature*, vol: 359, pp 826-829

Nowak, M., Bonhoeffer, S. and May, R. (1994). “Spatial games and the maintenance of cooperation.” *Proceedings of the National Academy of Sciences* 91, 4877Ð81.

Ostrom, E (1990). Governing the Commons. Cambridge University Press.

Page, Scott E. (1997). ``On Incentives and Updating in Agent Based Models.'' *Computational Economics*. 10 pp 67-87.

Page, Scott E. (2006). “Essay: Path Dependence" *Quarterly Journal of Political Science* Vol 1, No1 pp 87-115.

Page, Scott (2007). *The Difference: How the Power of Diversity Creates Better Groups, Teams, Schools, and Societies* Princeton University Press. Princeton, NJ.

Page, S E. (2008). ``Uncertainty, Difficulty, and Complexity'', *Journal of Theoretical Politics*. 20: pp. 115 – 149.

Page, S E. (2011) *Diversity and Complexity*, Princeton University Press, Princeton, NJ.

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.

Richards, Diana (1997). “Nonlinear Dynamics in Games: Convergence and Stability in International Environmental Agreements” mimeo University of Minnesota.

Schelling, Thomas (1978). *Micromotives and Macrobehavior* W.W. Norton, New York.

Smirnov, O and JH Fowler (2007). “Policy-Motivated parties in dynamic political competition”, Journal of Theoretical Politics 19 (1), 9-31.

Strogatz, Steven (2001). *Nonlinear Dynamics and Chaos*. Perseus Books.

Tesfatsion, Leigh (1997) “How Economists Can Get A-Life” in *The Economy as a Complex Evolving System II* W. Brian Arthur, Steven Durlauf, and David Lane eds. pp 533–565. Addison Wesley, Reading, MA.

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, pp 1Ð19.

Weidmann, Nils and Idean Salehyan (2013). “Violence and Ethnic Segregation: A Computational Model Applied to Baghdad”. International Studies Quarterly.

Wichowsky, Amber (2012). “District Complexity and the Personal Vote”, Legislate Studies Quarterly.

Zinkevich, Martin, et al. (2007). "Regret minimization in games with incomplete information." *Advances in neural information processing systems*.