Review Article

Abstractions, Ensembles, and Virtualizations Simplicity and Complexity in Agent-Based Modeling

Ian S. Lustick and Dan Miodownik

The Unique Analytic Leverage of Agent-Based Modeling

Models are analytically focused metaphors. For example, "the moon is a ghostly galleon" is a metaphor. Thinking of the moon as if it were a ghostly galleon evokes particular moods and images associated with the moon and not others. Analyzing a polity as an array of competing factions arising from social and economic interests models the polity in a particular and limited way. Portraying a political struggle in exact terms as a game of chicken involves application of a formal model to highlight certain simple, highly abstract, strategic elements in a complex multilevel reality. The distinctive advantage of a formal model is that it can be expressed unambiguously. That is, the model can be written as a computer program (as can any proper algebraic formula) and run successfully without the exercise of human interpretation or discretion.

Accordingly, to translate problems of scientific interest into a language a computer can understand is to model them formally. Agent-Based Modeling (ABM) is a computer-assisted methodology that allows researchers to design, analyze, and investigate formal models realized as artificial worlds inhabited by agents that interact with each other following prespecified simple rules.[1] Agent-based models vary considerably, but they all consist of arrays of autonomous, myopic units. Whether these units are modeled as states, individuals, corporations, ethnic groups, villages, or kinship groups is up to the experimenter and his or her theoretical domain of interest. Units in an ABM environment seek to adapt to their environment as they see it based on whatever goals they are programmed to pursue. Although individual actions are wholly determined at the micro level, when large numbers of such agents operate by interacting with one another, the macro state of the array as a whole is not predictable (though patterns in multiple "runs" of the array can be). As is the case in the "real" world, the specific trajectory the array will take is an empirical matter. It results from the initial condition of the array, the rules implemented to govern individual behavior of units, the complexly interdependent effects of adaptive behavior by the units, and whatever exogenous perturbations are streamed toward the array by the experimenter. Standard procedure is to generate large numbers of trajectories of these virtual worlds by perturbing them randomly or introducing randomized adjustments in initial conditions. By carefully calibrating these experiments, crucial experiments can be designed to bring reliable data into contact with isolated, theoretically informed claims.

To the extent that the rules governing agent behavior are derived from clear, corroborated or widely accepted theories, this methodology offers a powerful technique for refining, evaluating, and testing specific theoretical claims. The visual, controlled experimental, and transparent nature of the technique, as well as the ease with which it exports results in tabulated, statistically analyzable form, also supports its use as an idea pump for the development of theory.[2] Traditional formal modeling methods rely on algebraic formulas to translate simple relationships into mathematical expressions. But the limits of algebra prevent such techniques from incorporating many of the things known to be true of most of the worlds social scientists find interesting, including their multidimensionality, the presence of large numbers of interacting and autonomous units, and the predominance of highly irregular but nonrandomized patterns in the distribution of traits or interaction styles. The constraint of algebraic solvability therefore limits the ability of traditional formal modeling methods to capture the richness of interesting and even well-established substantive theories. Reliance only on traditional formal modeling techniques (such as game theory and rational choice) thus often entails ignoring what the modelers actually believe to be true about analytically crucial parts of the world. However, ABM, and the computer-assisted bottom-up simulations it produces, can be designed to capture beliefs about the real world embedded in or expressed by good substantive theories, thereby providing researchers with new opportunities to examine possible and probable outcomes associated with specific theory-based claims.

The revolutionary potential of this technique is associated with the fact that very large numbers of alternatively possible "futures" (or "histories") can be produced by varying initial conditions or a specific parameter setting of interest or by subjecting the theoretically specified model to random perturbations. Because of the automatic operation of the computer, every trajectory produced by a given model is consistent with the assumptions and propositions instantiated in the operating rules of the computer programs. Arising from the specific assumptions of the model, these unique counterfactual outcomes can yield reliable data about the implications of changing assumptions, boundary conditions of claims, sufficient conditions for particular outcomes, and robustness of results. The degree to which that potential is realized is a function of the empirical validity of substantive models and the degree to which these theoretical ideas have been implemented clearly and accurately.

No formal tool can substitute for empirical work based on the deep study of a single case, structured comparisons of a small number of cases, and/or the statistical analysis of data in the large N tradition. But formal techniques, and especially computer-assisted ABM simulation, are powerful complements to those techniques. An important attraction of ABM is that it can be used to explore the justification for many claims that ordinarily would be impossible to evaluate, in particular in regard to rare but interesting events. Indeed, computer-assisted ABM is the most effective technique available for conducting very large numbers of complicated but disciplined thought experiments about futures (or pasts) that could occur or could have occurred according to available theories. Careful design of artificial worlds, including exact control of all theoretically significant parameters, allows systematic manipulation of putatively significant independent variables and precise observation of the results. Using standard techniques of controlled experimentation and producing batches of theoretically consistent but unique outcomes, social scientists can expand the realms within which their variable-rich theories encounter numbers of cases much larger than the number of variables contained in the theories.

However, the method of computer assisted agent-based modeling also raises new challenges, especially in the realm of research design. As is the case with informal verbal models, there are limits to the amount of complexity it may be useful to include in a study. For example, a researcher studying political participation may well believe that participation is affected by culture, aspirations to move into elite positions, beliefs about the efficacy and responsiveness of government, or alternatives available if participation opportunities are foregone. But the research design adopted will require taking only a strictly limited subset of variables into account. If the researcher were required to specify and integrate the best model he or she considered relevant for political culture, elite recruitment, government responsiveness, or exit opportunities, the complexity of the task would paralyze the project.

No form of modeling should hold out as its goal the complete theoretical specification of all relevant variables and constants. On the other hand, for many problems the availability of computer-assisted simulations allows researchers not to settle for radically stylized models with key parameter values stipulated purely for analytic convenience. Computer assisted ABM can be used to integrate multiple theoretical modules, thereby leveraging the considerable knowledge available. This capability represents an enormous opportunity, but also a challenge as unfamiliar to formal modelers wedded to algebraic solvability as it is familiar to most other researchers—the problem of balancing parsimony with verisimilitude. A model that includes too many interacting modules explosively increases both the internal complexity of the model and the size of the space of possible outcomes. This sacrifice of parsimony can interfere with both confidence in theoretical operationalizations and in the interpretation of experimental results.[3]

Typically, the question of how much complexity to include in the design of a game theoretic or rational choice model is solved by the radical simplicity enforced on problem definition, dimensionality, and other features by the need for algebraic tractability. The claim or hope is then offered, implicitly or explicitly, that the solution of the game applies equally to more complex settings.[4] However, once the move is made to computational modeling, the "algebraic tractability" constraint vanishes. The research design options open to ABM modelers are even richer thanks to the availability of cheap computer power and modeling platforms that do not require computer programming skills.

The wider horizons for formal modeling opened up by ABM make the problem of research design both more important and more difficult. ABM researchers are forced to decide whether, for a particular project, the formal model to be built should be simple, complex. This article aims to help researchers cope with this problem. After consideration of the turn toward ABM taken by many political scientists, a typology of three different kinds of ABM research designs used to address crucial problems in political science—problems otherwise resistant to formal modeling strategies—will be illustrated. This typology—abstractions, ensembles, and virtualizations—will then be applied to work done on a family of problems pertaining to the evolution of collective identities and norms and their contribution to collective action. This illustrative treatment will present the range of ways political scientists using ABM have reacted to the greatly increased degrees of freedom this method of investigation permits them. We hope that this typology will encourage wider and more systematic use of ABM techniques and help improve the correspondence between the problem under consideration and the research design adopted.

Growth in Use of Agent-Based Modeling by Political Scientists

The use of agent-based modeling methodology has spread in the last ten years as reflected in the number of papers delivered in major political science conferences, as well as in the growth of peer-reviewed articles published by mainstream political science journals. The number of ABM presentations rose from one at the 1998 annual meeting of the American Political Science Association (APSA), to twelve papers combined at the 1999 annual meetings of the APSA and the Midwest Political Science Association (MPSA), to twenty-five papers at the annual APSA, MPSA, and International Political Science Association conventions in 2007.

In 2001 Lars-Erik Cederman surveyed the emergence and contribution of ABM in political science since the work by Schelling in the 1970s and Axelrod in the 1980s.[5] His review surveyed 108 items, including very influential books, as well as contributions in the field of economics and sociology. [6] However, only nine of the items included in the 2001 study appeared in major political science journals. Our own survey of journals in political science and related disciplines between 2002 and 2006 found thirty-four journal articles using ABM methodology to study an issue falling within the domain of political science. Fifteen of the articles were published in major political science journals. Others appeared in more specialized journals, including those devoted primarily to the exemplification of computer simulation techniques.

Parsing the Space of ABM Research Designs: Abstractions, Ensembles, Virtualizations

In a recent book on theoretical and computational biology, Volker Grimm and Steven F.

Railsback justify their work by the need they have identified for "additional guidelines for coping with the complexity of bottom-up simulation models."[7] Grimm and Railsback offered their judgment against the background of a well-established tradition of computational modeling in the biological and related sciences. Parallel to Grimm and Railsback, this article will identify research design decisions made by political scientists using ABM techniques and assess the relationship of these decisions to aspects of the problems and the ambitions of the modelers. For any given project, how much complexity has the researcher decided to include in the model, how was that level of complexity depicted, and what kinds of justifications were used to support it?

If successful, this trichotomy—abstractions, ensembles, and virtualizations—can guide development and construction of different kinds of models for different purposes and for understanding what criteria to apply when learning from and evaluating this kind of formal work. The categories of our typology are intended to provide flexibility in the design of ABM models, precision in the accounts and justifications for formal modeling research designs, and a broader warrant for the deployment of formal simulation models by guiding the modulation of their complexity with distinctive kinds of problems and modeling justifications. To limit this effort, the analysis will concentrate on the related topics of culture, norms, and political action. The reciprocal relations among these domains comprise a multifaceted field of inquiry and argument stretching across the boundaries of all contemporary social sciences.

Abstractions The intention of an abstract model is to leverage discrete and exact understanding of very simple relationships by observing the impact of slight variation at the micro level on behavior at the macro level. Grimm and Railsback describe such models as "minimal" and "intended to explore a concept without reference to a particular species or place. These models are not intended to make testable predictions, or even to be applied to specific real systems."[8] For example, an abstract computational model constructed to study the social world might simulate a large group of interacting simple agents (representing people as nothing but myopic, minimally adaptive, cognizing units) to investigate the impact, at aggregate levels, of varying to the sensitivity of agents to the information available in their social network, the number of attributes of others individual agents can monitor or the radius of influence agents exert over other agents.

For political scientists a typical question that might be posed at this level of abstract modeling is how alliance structures and patterns of change in those structures might vary in relation to whether interacting units forming the alliances were themselves united or divided. In general, given a large array of standardized and similar interacting units, an abstract modeling research design seeks to understand how varying the algorithms that produce behavior of individual units, or how varying the number and distribution of types of these units, affects the behavior of the entire array. Such work is not designed as a stylization of a specific problem or even of a real kind of problem. Instead, abstractly designed models have as their target the behavior of the elements from which portions of the real world of interest are constructed. The kind of general insights that can be produced by such abstract ABM work, although incapable of explaining outcomes in any particular case or even in any particular kind of case, can inform, illuminate, enrich, and simplify an understanding of the real world by showing that very simple relationships at the micro level can account for complex patterns at the macro level. Because of their generality, the insights derived from this kind of ABM design, much like those generated from abstract game theory, are blunt but readily transportable as patterns of recognized relationships across the borders of many scientific domains and problem sets.

For example, in a series of studies assembled in his 1997 book *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Axelrod discusses work on the emergence, stability, and diffusion of norms as a result of interaction among people with limited rationality.[9] In these studies Axelrod abided by his influential “KISS” maxim, “Keep It Simple Stupid.” In our schema what this means is that Axelrod’s work was organized almost entirely within abstraction-type research designs. For example, in Chapter 7 of the book, “Disseminating Culture,” Axelrod explores the determinants of the spread of cultural traits and the way different degrees of homogeneity can appear in a society despite a great deal of latent heterogeneity.[10] Using a basic social influence model of tendencies to conform among interacting individuals, Axelrod finds that the number of distinctive traits (strategies, behaviors, or identity signals) activated by agents in stable homogeneous regions of such activated traits decreases with increases in the size of the repertoire, or portfolio, of traits that each agent has at its immediate disposal but increases with the overall number of traits present in the society.

As is typical for an abstraction research design, the clarity and simplicity of patterns exhibited at the macro level by the populations of agents are determined, in nonobvious ways, by characteristics measureable only at the micro level. Here, the abstract relationship being investigated is the effect of interacting size of agent repertoire and the size of the overall pool of cultural traits present in the repertoires of any agent, to explain variation in the collective exhibition of traits. An added finding in the study is linked to Axelrod’s manipulation of the size of the population, regardless of the traits that define its complexion. He finds that increasing the size of the population (not the pool of traits within it) beyond a certain point decreases the heterogeneity of that population as measured by the number of different distinctive and stable regions of cultural activation. By studying the relationship between adjustments at the micro level and emergent properties at the macro level across a large population of agents, Axelrod leverages an abstraction research design to suggest underlying reasons for the unpredictability of individual state boundaries and the conditions under which theories predicting cultural homogenization as a result of increased communication are more likely to be correct. His work leads to the hypothesis that the production of a culturally diverse public space is not an automatic result of rendering individual citizens more cosmopolitan and that in fact it is precisely in societies comprised of individuals with a limited array of cultural resources at their disposal that the most diversity (or cultural fragmentation) at the collective level is to be expected.

Others have followed in Axelrod’s footsteps by exploring relationships between the emergence of cooperation norms linked to simple characteristics of agent and simple aspects of the environment within which agents interact. Macy and Skvoretz, use a computational model to explain trust and cooperation between strangers in the absence of any type of formal or informal social control.[11] Without endowing human characteristics to their agents apart from the fact of individuals linked together within distinctive networks, they postulate that behavior "as if" based on trust can emerge more readily in networks involving some interaction with strangers than within one single isolated network.

Majeski, Linden, Linden, and Spitzer explored a similar question and come to somewhat different, but not necessarily inconsistent, conclusions.[12] They ask if and how cooperation evolves in social worlds characterized by the presence of selfish agents engaged in repeated relations without central authority. With their ABM simulation they demonstrate that cooperation can evolve in a harsh environment even if agents have quite limited capacities. Their main argument is that the key condition for transition from uncooperative to cooperative worlds is the emergence of small, relatively stable networks of cooperating agents.

Shibani, Yasuno, and Ishiguro examined the impact of nonlocal information on patterns of convergence.[13] They used an agent-based model focused on variation in the capacity of agents to obtain and process global information. They showed how this variable leads to the operation of a mechanism producing both local convergence and populationwide diversity. Similar effects of globally available information were reported in studies by Latané, Nowak, and Liu.[14]

The abstraction type of research design was also used by Lustick in his work on the basic mechanisms hypothesized by constructivist identity theorists in their general explanations of collective identity formation and change.[15] Lustick noted that from a Lakatosian point of view the constructivist identity theory research program (in political science and other related disciplines) had defeated its primordial or essentialist rival. However, it faced degeneration if it could not escape from decreasingly interesting demonstrations of its hard-core negative heuristic—that permanent, primordialist, or essentialist identities did not exist—or of its core positive heuristic—that patterns of collective identity change require explanation in terms of fluid and deployable identity resources or attributes.[16]

Lustick identified the main obstacle to moving the research program in a progressive direction as the difficulty of acquiring the finely grained data required to investigate theoretically crucial mechanisms and answer questions about the dynamics of collective identity prompted by the program's positive heuristic. For example, how much more difficult or disruptive should constructivists expect it to be for individuals or groups to display an alternative identity than to acquire a new one? What would different degrees of relative difficulty associated with these operations imply for the way populations of constructivist individuals or communities behave? Similarly, what effects might result from variation in the size of repertoires, or the impact of different proportions of a population with different sized repertoires.

Lustick's Agent-Based Identity Repertoire (ABIR) simulation tool was designed to answer these and other related questions. It was a relatively simple platform for the production of computational models known as cellular automata.[17] By populating these models with agents with repertoires of activatable identities, governed by updating rules corresponding to different interpretations of constructivist identity theory, Lustick's experiments went beyond previous work along these lines that studied clustering and segregation without drawing explicitly on constructivist identity theory.[18] Experiments enabled controlled comparisons to explain variation in the consequences of endowing agents with differently sized repertoires, assuming larger or smaller gaps between the ease of rotating an "identity" versus acquiring a new one or distributing different ratios in the population as a whole of agents with different characteristics.

The display in Figure 1 shows a snapshot of a grid of agents drawn from one of these studies. Each square is an agent activated on one particular color (identity) drawn from a specific repertoire of colors available to each agent drawn as a subset of the total spectrum of colors present in the array as a whole. For example, the agent indicated by the arrow on the upper grid in Figure 1, is "activated" on a "red" identity, shared with some but not all of its immediate neighbors. It happens also to have a repertoire (that is, cache) of five additional identities: blue, green, brown, purple, and orange. Whether it maintains its activation on red, rotates another one of its repertoire identities into view instead of red, or acquires a new identity in its repertoire by discarding one not currently displayed, will depend on the values implied by constructivist theory for updating triggers, the influence levels of agents whose behavior it is monitoring, and the signals it receives at every time step about the general popularity or attractiveness of alternative identities (colors).[19] Indeed, an image of the same grid after one-hundred time steps (the bottom grid in Figure 1) suggests that the agent altered both its activated identity and the composition of its identity repertoire.

This type of model falls within the category of abstraction. It was neither inspired by nor tested against particular kinds of political problems in which collective identity dynamics of a certain configuration were thought to drive outcomes. Instead, it was used to explore the impact of a range of simple relationships implied by the theory, primarily the effects of variation in the number of identities available in the landscape (polity) and in the repertoires of agents (individuals) on the pattern of clustering and of identity extinction and concentration. For these experiments the algorithms driving agent behavior were maintained exactly across conditions as were all parameter values reported controlled for all other variables except size of the repertoire of identities available to each individual (agent) in the population. Lustick's main findings were twofold. First, his data revealed a curvilinear and nonobvious relationship between increasing size of identity repertoires and the likelihood of trajectories featuring more concentrated clustering of collective identity formation on a small number of favored identities. Second, he found that small amounts of unevenness in the size of identity repertoires had a powerful and positive impact on the likelihood of identity extinction and of outcomes featuring highly concentrated clustering of a small number of identities.[20]

In a subsequent study, Lustick and Miodownik investigated whether mechanisms within constructivist identity theory could account for hegemonic patterns of institutionalization—the crystallization in dynamic and uncertain worlds of narrative structures so resistant to change that they become effectively permanent.[21] They were able to associate different kinds of initial conditions in the distribution of identity attachment,

Figure 1 Abstraction-Agent-based Identity Repertoire (ABIR) Model

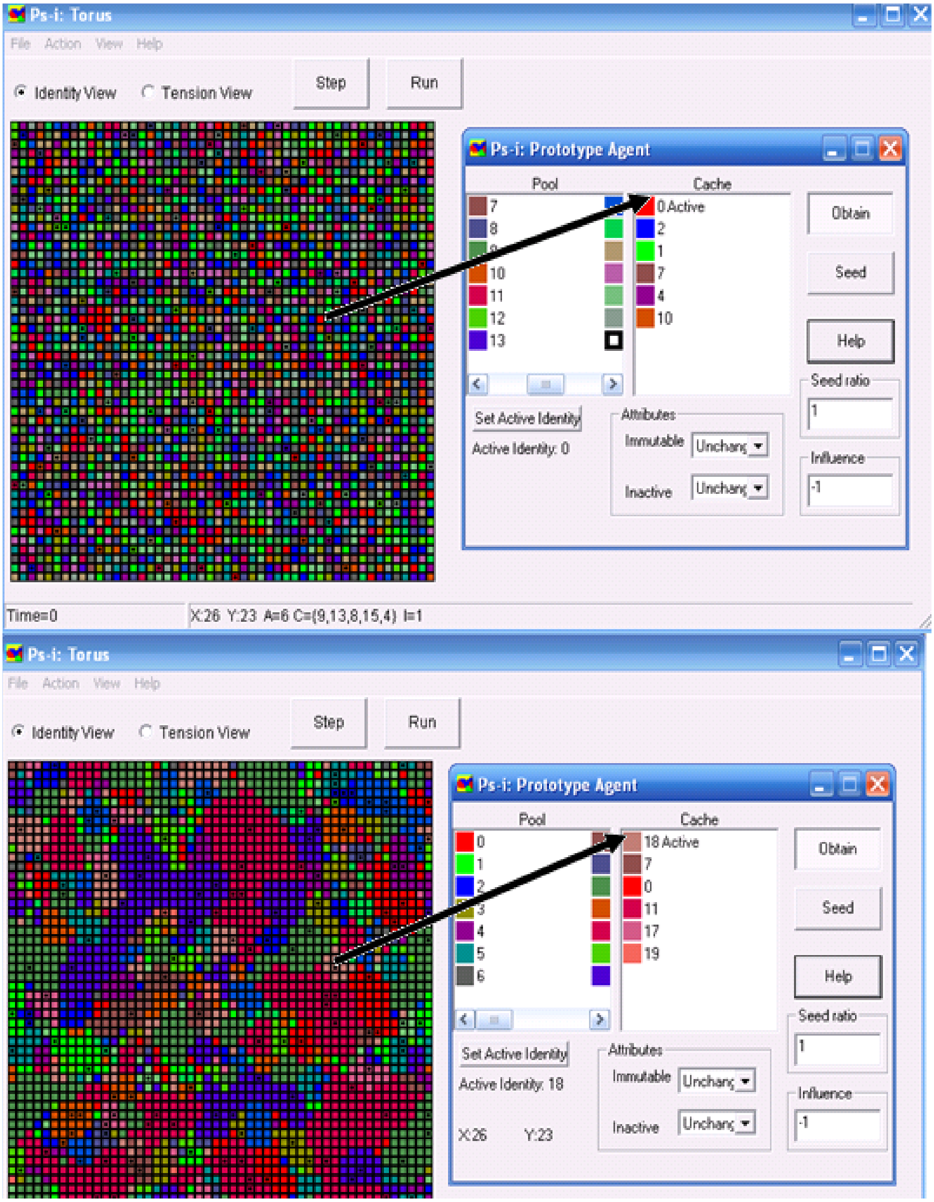

the presence or absence of identity entrepreneurs, and different rates of change in the objective attractiveness of different identities with the emergence of a "threshold effect" leading to identity "crystallization" and its institutionalization. Among their results was the finding that identities enjoying even marginal advantages over others were significantly more likely to secure majority positions in the population and that a threshold effect was present within delimited zones that triggered very large payoffs to marginal increases in initial margins of superiority.

Ensembles Despite the prominence of the KISS principle as advocated by Axelrod and its reinforcement of agent-based modeler attention on abstraction as a preferred research path for ABM simulation, there is no one best way to design ABM research. While some researchers may have a strong taste for the elegance and power achievable by studying the emergent effects of simple abstractions, others may aspire to greater verisimilitude or be attracted to wider, more complex horizons. They may also wish to use ABM to contribute directly to social science debates featuring specific theoretical claims and counterclaims about dynamics within social fields comprised of both complex agents and distinctive institutions.

As understanding of the dynamics of simple mechanisms in general theoretical contexts increases, it is natural for agent-based modelers to think about how more specific but still generic problems can be modeled. By building stylized models of types of empirical domains, competing hypotheses advanced to explain patterns in those domains can be evaluated. For instance, one can design ABM models of culturally divided states to examine their susceptibility to secessionism, model authoritarian polities that may (or may not) be liable to experience national mobilizations or democratic bubbles, simulate the conditions under which international systems may be more or less disposed to stability, or explore authoritarian regimes' vulnerability to latent challengers. We label this type of research design "ensemble."

The intention of an ensemble model is to create a replicable simulation of a specific kind of problem. The building blocks of an ensemble are not only simple algorithmic mechanisms, but operationalizations of theoretical positions with enough empirical corroboration to attract the interest of the researcher. The goal is to identify inconsistent suppositions, explore the logical implications of interactions among theories, and/or locate the range of possible and likely variation in outcomes in the real world domain of interest.

The ensemble category corresponds closely to what Grimm and Railsback mean by "synthetic" models.[22] Such studies are not designed to illuminate a particular case but are focused instead on a class of cases or a kind of problem. The distinction is the same as that between seeking to predict the implications for Tunisia if its long-time autocratic leader, Zine El Abidine Ben Ali, were suddenly to pass from the scene and studying the implications of different types of authority structures within authoritarian systems for the behavior of the system in the aftermath of a succession crisis.

The ensemble design is illustrated by a series of studies conducted by Lustick, Miodownik, Britt A. Cartrite, and Roy J. Eidelson on issues pertaining to the sources and outcome of collective-identity-based mobilization in multicultural countries.[23] Using PS-I, a considerably more powerful and elaborate modeling platform than ABIR, but one based on the same premises of constructivist and evolutionary theory, these researchers produced a family of models ("Beita," "Ethniland," and "Virtualstan") to capture common features of multicultural polities.[24] Among the topics addressed in these studies are the contributions of cultural and economic incentives to demands for regional autonomy, the impact of political decentralization on regional mobilization, and the effects of power sharing as well as other accommodation strategies on ethnopolitical mobilization, secessionism, and state partition.[25]

While abstraction studies produce conclusions applicable to many empirical domains and even many disciplines, ensemble studies can yield insights into specific issues of interest as defined within a discipline or a disciplinary subfield. For example, Miodownik suggested that different aspects of identity mobilization (support for the minority group, identity clustering, and emergence of political boundaries) correlate with the intensity of ethnic differences and the scale of economic incentives. Miodownik and Cartrite's study of the relationship between political decentralization and regional mobilization indicated that methods of power devolution may result in increasing identification with regional minorities but that decentralization has a nonmonotonic relationship with the emergence of political boundaries. Even more specifically, Lustick and his collaborators examined how specific strategies of power distribution affect the likelihood of ethnopolitical mobilization, secessionism, and secession. In this regard they examined the effect of state repression, bureaucratic responsiveness to minority demands, minority representation in state institutions, and the creation of semiautonomous, minority-led institutions. The results of the study indicated that both repression and accommodation decrease the rate of secessionism and secession. However, accommodation brings with it higher levels of ethnopolitical mobilization, while repression, which requires heavy investment of resources, reduces mobilization, but only extreme repression reduces secessionism and secession.

Figure 2 is an example of an ensemble model dubbed "Beita," used to study the consequences of power-sharing to secessionism and secession.[26] The top grid captures in composite some of the characteristics of a mutlicultural and multiregional state that may, under some conditions, experience secessionism and/or secession in at least one of its regions. A disgruntled minority—an identity systematically excluded from the repertoires of the agents in the authority structure—is most prevalent in the southeast quadrant. The magnified image of this quadrant appears on the right hand side of the figure. This article is not the place to justify all the coding decisions. Suffice it to say that in accordance with what can be expected from an ensemble research design the model included dynamic, rule-based operationalizations of real world entities, including regime authority structures, regional minority elites, and agent states expressing public identification with the regional minority. The bottom of Figure 2 shows one of many possible trajectories of agents' interactions under varying parametric conditions. This particular run features consolidation of a space controlled by the regime, an area of strong mobilization and minority identification, and the emergence of boundaries that can be taken as an indication of secessionism.

Figure 2 Ensemble-BEITA-Multicultural and Multiregional State

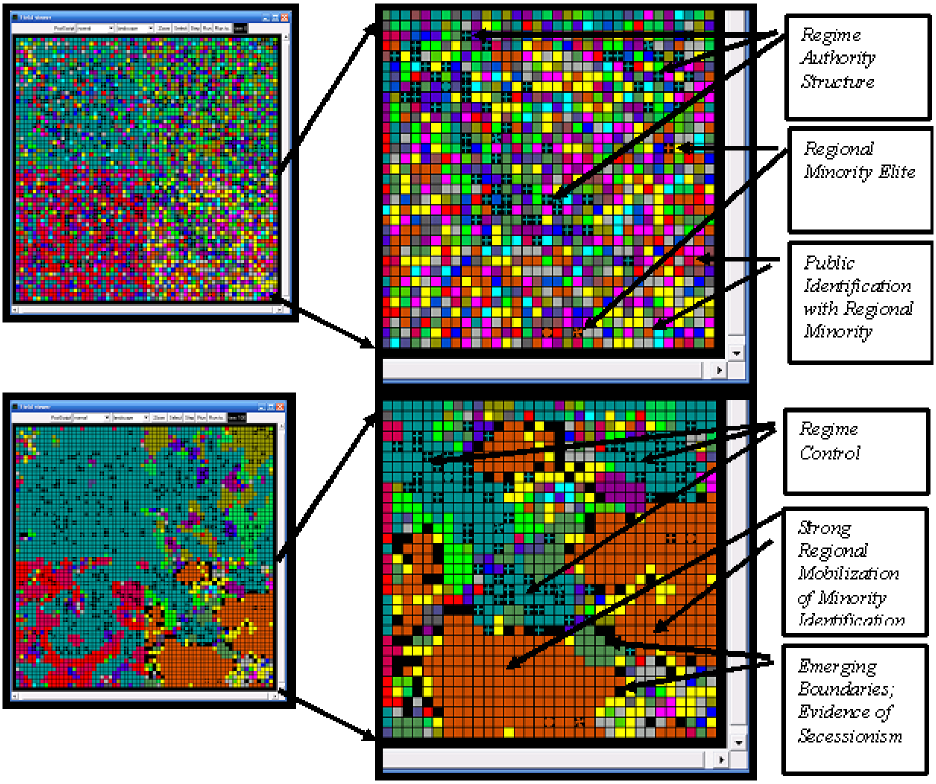

The basic challenge of an ensemble research design is to translate the independent variables of competing theories into formal and dynamic simulations and then to use virtual world outcomes to make inferences about real world patterns and the correctness of competing theories. For instance, Rousseau and van der Veen's study furthers understanding of the impact of shared identities in the international system.[27] Their strategy was to use a model of the diffusion of collective identity across individuals within a society as the basis for building an ensemble model of the operation of similar variables among states to transform the international system from one governed by balance of power to one governed by norms or identities. Their model is more complex than those used by Lustick because agents (states) can see more about other agents than their public face. Rousseau and van der Veen observed a curvilinear association between the complexity of an agent's views and the emergence of a shared identity at the aggregate level, with maximum sharing associated with moderately sized repertoires (a finding similar to Lustick's), a greater likelihood for the emergence of shared identity under more turbulent environmental conditions (a finding similar to Lustick and Miodownik's), and a decrease in the emergence of shared identity in the presence of influential entrepreneurs (a finding somewhat different from Lustick and Miodownik). They also found that the interaction among the three variables in their model (complexity of agents, stability of environmental conditions, and presence of leadership) contributes to patterns of polarization.

Lars-Erik Cederman's book *Emergent Actors in World Politics* used the affinity between ABM and constructivist identity theory to criticize prevailing theories in international relations. Cederman stressed that the radical separation in prevailing international relations theory between processes of nation-state emergence and nation-state behavior drastically reduces the power of this body of work. Cederman's approach is explicitly based on an overarching evolutionary and constructivist framework that disaggregates national states as the fundamental units of analysis. Instead, his models are populated by regions, competing identity groups and formulas, leadership and economic resources, different kinds of military technologies, and problematically established state boundaries.

Cederman built his agent-based models by bundling a number of mechanisms inspired by theories of state and nation development with some key constructivist mechanisms based on conformity oriented adaptation among actors with overlapping but nonidentical repertoires of identities. His ensemble models are designed to study specific kinds of problems in international relations, including the circumstances under which large national states can consolidate; the conduciveness of offensive or defensive military technologies to unipolarity, bipolarity, and multipolarity; and the impact of defense alliances on balance of power outcomes. In one chapter on "Nationalist Coordination," however, Cederman presented an ensemble model—in a comparative politics vein—of the process by which a divided state can become more or less consolidated around a unifying national identity.[28] This model featured interacting mechanisms expressing predominance by a center over a divided periphery, the distribution of resources among peripheral regions, degrees of overlap in identity repertoires across regions within the space of the state, and the contribution of entrepreneurial leadership or its absence. He found that the more asymmetric the relationship between the center and a divided periphery, the more inclusive the emergent identity of the new political unit will be, but that this relationship is strongly affected by the resource balance and intensity of competition across rival regions.

Although more detailed than the other models, this work of Cederman still falls within the ensemble category. Instructively, however, he concludes the chapter on nationalist mobilization by applying the categories and findings to the real case of Yugoslavia. For our purposes it is important to note that he does not attempt what we would term (see below) a "virtualization" of the Yugoslavian case. That is, he does not provide or claim to provide a stylized version of the Yugoslavian polity at a particular point in time in order to simulate its breakup and specifically identify the determinants of that outcome as opposed to others. The analytic step he does take toward virtualization can be seen as equivalent to that made by a game theoretic analyst who aligns a specific game with a specific case in the real world by stylizing the case until it matches the categories of the game, by a researcher with confidence in a regression model who makes an out of sample prediction, or by a scholar using propositions refined with a small-N study to predict new facts about an unstudied case by generating hypotheses about it based on categories used and results obtained from data drawn from the primary cases.

Indeed, it is natural for agent-based modelers using ensemble designs to project the implications of their findings onto real cases, especially in regard to the dynamics of severe ethnopolitical conflicts. For example, Srbljinovic, Penzar, Rodik, and Kardov developed an ensemble model including combinations of mechanisms believed to affect propensities to engage in ethnic mobilization.[29] Agents vary by group membership, level of mobilization, and degree of grievance experienced. They are embedded in social networks and exposed to mobilization appeals. Although the authors represented their model as a virtualization of the former Yugoslavia, the data used to drive the model were abstractly produced. Their work is therefore too stylized to count as a virtualization, and it should be regarded instead as an ensemble model of processes that are, in fact, prominently featured in the Yugoslav case and that inspired development of the model by the authors.

As Srbljinovic, et al., were inspired by the Yugoslav case, so Bhavnani and Backer's simulation model was inspired by episodes of vicious ethnic conflict in Rwanda and Burundi.[30] They sought to explain variation in the scale of violence across episodes of ethnic conflict. The authors' computer simulation incorporated several empirically validated assumptions, exploring the independent and interactive effects of four theoretical factors. The presentation of the simulation suggests that they had in mind a virtualization research design in which their simulation models could be taken to be realizations of Rwanda and Burundi rather than a kind of situation or a structure of circumstances shared by those two countries. However, since no empirical data from Rwanda and Burundi were used to build their model, the research design employed by Bhavnani and Backer falls into the ensemble category. With this distinction the use they make of Rwanda and Burundi to motivate the construction of a formal model of medium complexity can be recognized without requiring them or their readers to treat the model as the substantially more complex virtualizations they would be if Bhavnani and Backer were actually presenting some version of a virtual Rwanda or Burundi.

In his study of the Rwandan genocide, Bhavnani presented a distinct but related ensemble model.[31] His purpose was to study the effects that a specific behavioral norm, defined in ethnic terms, would have on participation in a genocidal campaign. The testing of the norm was accomplished by inserting it into a dynamic array that also incorporated several other theoretically relevant mechanisms affecting agents' behavior, including the impact of distributions of individual heterogeneity, within-group enforcement, types of behavioral adaptation, structures of group networks, and the presence of ethnic entrepreneurs. The simulation results were then used to help explain the mass participation of reluctant Hutu in the mass killing of Tutsi in Rwanda.

Inaccurate representations of their own work by agent-based modelers pushing beyond abstraction model designs suggest that the boundary between ensemble work and virtualizations is not well understood. This problem is easily corrected. Most comparative politics specialists are familiar with the distinction between theoretically informed comparative work and case studies. Success in the former, equivalent to an ensemble research design, implies increased knowledge about kinds of circumstances or types of entities and the leverage into specific cases such knowledge may afford. Successful case studies, on the other hand, equivalent to what we call virtualizations, can clarify the logic, in specific empirical contexts, of what would otherwise be regarded as arbitrary or highly contingent outcomes. Good case studies, or virtualization models, can also improve assessment of the validity of competing interpretations of a case's past and the likelihood of alternative visions of the future.

Virtualizations Generally, ABM researchers have focused most of their attention at the level of abstractions and ensembles. But there are strong reasons for those who learn techniques of computer assisted bottom-up modeling to use them to generate theory or policy relevant findings by studying specific cases. This inclination is especially strong in relation to assessments of contingent outcomes, risk analysis in specific empirical cases, and scenario production liked to alternative policy choices or historical counterfactuals. These tasks standardly entail thought experiments based on privately and partially specified mental models linked to un-self-consciously accepted theories and assumptions.[32] Such exercises, however, are notoriously susceptible to political, psychological, and cognitive biases. By using virtualization research designs, agent-based modelers can systematically probe the space of the possible for insight into the likely attributes of the actual. This is accomplished by producing large batches of scenarios, with each consistent with expert knowledge. Indeed, the objective of a virtualization research design is to surface, specify, accumulate, and implement the assumptions, theoretical relationships, and initial conditions corresponding to the best work experts could do with their mental models, if those experts could identify them and then discipline their use.

To be sure, virtualization models will still be incomplete, that is, they will still be models, not parts of the world itself. Just as the mental model of a problem in an expert's head—for example, a question pertaining to the effect of a tsunami on ethnic conflict in Sumatra—would be an incomplete "operationalization" of the real world and of relevant aspects of it, so an ABM virtualization of the problem capable of posing that question and generating possible trajectories into the future will not be an exact replica of the world. But in contrast to the expert's mental model, a computerized virtualization is exact in its incompleteness, and it is replicable. Moreover, to the extent that the expert's mental model can be specified, and to the extent that those specifications can be translated into algorithms implementable in a virtual world, the model thereby created can serve as the basis for controlled experiments about implications of variation in assumptions, theoretical relationships, parameter values, or initial conditions for the probability of different kinds of futures. Undoubtedly, ABM virtualization research design is demanding. But it is not more demanding than the actual (if usually unacknowledged) tasks required of an expert asked to make a counterfactually based judgment about the effect of a policy option or a specific set of circumstances or a historical institutionalist judging the causal implications of a historical antecedent for an observed state of affairs.

For Grim and Railsback virtualizations are "predictive" models. It should be emphasized, however, that, no matter how complex, such models can never be considered capable of "point predictions." Rather, their predictions are in the nature of cones or distributions of probabilities of different kinds of outcomes under stipulated conditions.[33] For example, an expert panel with access to good data about the city government of New Haven and excellent understanding of how the city works would still not be sure of the exact implications of differing policies for change in the distributions of possible outcomes. However, if a dynamic model of the city of New Haven were to be constructed, based on the factors believed by experts to be important for the shaping of that future, then by subjecting the model to small random perturbations that could arise from beneath the analytic horizon, large numbers of possible futures for the city could be simulated. Patterns in that distribution of counterfactual futures would be the basis for probes of the city's actual future based on the best knowledge and expertise available. Alternatively, if theories were substituted for the panel of experts, experiments could be run with slightly different theoretical assumptions about how cities like New Haven behave, or existence proofs could be sought for competing hypotheses about what processes may have led to current outcomes (an exercise analogous to the prediction of the trajectory of an automobile from analysis of the distribution of debris following an accident).

In a series of papers on contemporary Pakistan that explores possible and more or less likely future trajectories for the Pakistani political system, Lustick created a virtualization of the Pakistani polity. "Virtual Pakistan" (VirPak) is a representation of the dominant factors understood by most Pakistan specialists to be driving events and trends in that country. These include military interests and dispositions, rural landowning elites, ethnic loyalties, kinship and religious affiliations, regional rivalries, economic stratification, urban-rural divisions, political communication patterns, fundamentalist oppositionism, central and regional bureaucratic authority structures of variable efficacy, and corruption. In their fine details, of course, no agreement exists as to exactly how these various factors relate to one another, but every Pakistan expert operates with some set of implicit theories of how they do, even if it is rather clumsily arranged, changeable, and difficult to specify.

VirPak instantiates a dynamic model of the collective imagination of Pakistan specialists. It is a transparent, albeit complex, but internally consistent model—necessarily so, since it runs successfully as a computer program. Indeed, it is probably more detailed, stable, replicable, and comprehensive than any mental model of an entire polity can possibly be. Its construction entailed blending knowledge about Pakistan produced by ethnographers, historians, sociologists, political scientists, and other Pakistan specialists; theories used by social scientists to understand countries like Pakistan; and demographic, social, economic, and political data about the real country of Pakistan. At the theoretical core of VirPak are the same constructivist principles animating the mechanism and ensemble studies described above, plus other well-corroborated theories of bureaucracy, ethnic conflict, and civil military relations.

The screenshots in Figure 3 show VirPak at time step 0 and time step 100 illustrating a typical trajectory of an unfolding future. Unique features of each trajectory are produced by emergent processes of interaction under the impact of small streams of exogenous shocks (events below the analytic horizon). Agents across the array update their activation based on myopic assessments of the attractiveness of different political, cultural, or

Figure 3 Virtualization- Virtual Pakistan

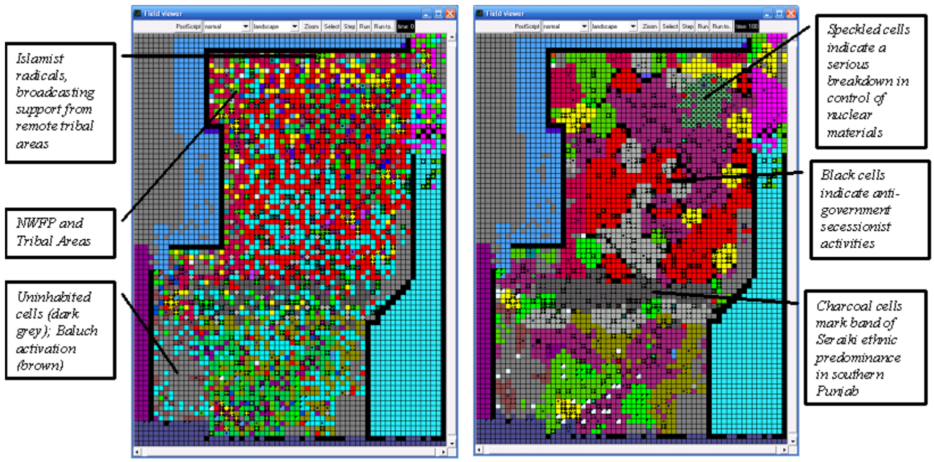

economic affiliations and the effects of networks of elites representing overlapping and partially integrated elite structures. At this point in this trajectory the strength of the Pakistani regime is reflected in the purple clustering in the heavily populated Islamabad, Lahore, Multan area of northeastern Punjab and in the prevalence of manifestations of support for it in Karachi and other portions of the southern Sindh region. Public Punjabi identification (red) predominates in the rural areas of the Punjab with landed oligarchs (grey circle icons) exploiting their influence over local religious elites and provincial legislators (cross icons) to dominate the largely peasant population (note peasant [blue] activation of many "Punjabi" cells at t=0). While the Northwest Frontier Province and Tribal Areas are divided between agents activated on their "Pushtun" identity and those manifesting support for militant Islam (dark green), Sindhi "sons of the soil" movements (support registered in khaki) are dominant in the rural areas of the eastern and northern Sindh, while different versions of Islam Muslim Fundamentalist (dark green and light green) pose a challenge to the government and to Sindhi nationalists in the region that reflects its conflict-laden character. Yellow clusters and icons indicate centers of military influence.

The central question Lustick addressed in his 2002 simulation experimentation with VirPak was whether, over a two to three year time frame, it was more reasonable to view the country as poised on a volcano about to explode or as stabilized by a logic that flies in the face of the daunting array of disabilities and dangers the country faced.[34] His findings (reported in early 2003) suggested the latter to be a more plausible view than the former. Regardless of the level of fundamentalist mobilization and the amount of apathy experimentally introduced into the society, the core capabilities of the regime to control the central urban areas remained intact across most experimental conditions and the largest portions of the distributions of future outcomes arising from hundreds of simulation runs. To be sure, regime presence in the rural areas was often light or even nonexistent, but even under conditions of fundamentalist strength extremist Islamism was never able to mobilize more than 50 percent of VirPak. This pattern speaks to the entrenched divisions in the society—regional, ethnic, religious, class, and other. It also reflects a fundamental alienation between the state and society.

No point predictions were offered as a result of experimentation with VirPak with respect to President Musharraf's or his government's survival. However, simulation results did suggest his government to be typical of the most likely kind of government the Pakistani regime and political system are likely to produce. Based on a close but fundamentally uneasy relationship between the military and the national bureaucratic, Punjabi, and rural oligarchic elites, the government was described as having a hold on power that was very likely to allow it to survive through 2006, likely but not very likely to avoid serious regime challenges, and unlikely to establish its rule on a wider and more secure basis.

Virtualizations such as VirPak, unlike abstractions and ensembles, are still quite rare in political science, although more common in ecology, migration studies, and epidemiology. Virtualization models in these fields are commonly constructed by mapping geographical, demographic, economic, political, or other relevant data onto a grid of potentially interactive cells constructed as a cartograph of a particular country or region. Interactions among agents or cells are then structured by rules that may exploit both their locational and state characteristics.[35] One successful ABM virtualizations was developed by a team of social scientists using the SWARM platform. Anthropologists, archeologists, political scientists, and computer scientists collaborated to build and experiment with a virtualization of the society of prehistoric inhabitants of the Long House Valley, located in the Black Mesa area of northeastern Arizona.[36] The region was inhabited by the Kayenta Anasazi, ancestors of the contemporary Pueblo people, from circa 1800 B.C. to circa A.D. 1300. The authors described their overall research program as including extensive experiments with what they described as "a multi-agent computational model of this society that closely reproduces main features of [the Anasazis'] actual history, including population ebb and flow, changing spatial settlement patterns, and eventual rapid decline." The purpose of the research study was to test theories of the Anasazi's evolutionary trajectory based on archeological data. Experimentation with the model, whose agents are monoagriculturalists, was conducted by "systematically altering demographic, social, and environmental conditions, as well as the rules of interaction [among households]."

Current work underway by Lars-Erik Cederman's research group illustrates another direction that ABM virtualization research designs may take.[37] With the proliferation of data sets for exploring relationships between a host of country-level factors (topographical, demographic, economic, cultural, political) and armed conflict, civil war, and violence, analysts have increasingly worried that the mechanisms implied by the correlations may not, in actual cases, be the causes for the outcome predicted by the regression model. The main objective of GROWLab (Geographic Research on War Laboratory) is to help guard against this danger by building "more accurate computational models of civil wars."[38] These models, still under construction, will use geographic information systems and data about population densities in specific regions by cultural group to produce accurate visual maps as frameworks for dynamic agent-based models.

Although not available in the public domain, it is known that large-scale virtualization projects using ABM techniques have been undertaken by the United States and its allies for intelligence purposes. One danger associated with such projects is that the technical capacities to achieve one-to-one modeling ratios with the "real world," to animate elements within the simulation, or to multiply elements and levels of analysis that may appear relevant to a detailed rendering of a case will drive this activity to ignore the discipline that should be imposed on the enterprise by corroborated social science theories and techniques of internal and external validation. The result would be to recapitulate, in the virtual world, the same "bloomin, buzzin confusion" that normally afflicts analysts faced with the impossible challenge of making "point predictions" about the trajectory of the real world. For both social science and public policy purposes it is therefore of great importance to formulate best practices for virtualization modeling. These will require, of course, that virtualization is recognized as a specific type of research design for ABM whose requirements and limits need to be understood with as much care and analytic rigor as are applied to abstraction and ensemble designs.

Conclusion

One reason for the sustained emphasis on simple models and abstract data may be the seminal influence of Robert Axelrod's work in this area and his maxim for agent-based modelers, "Keep It Simple Stupid." Though a salutary and helpful principle for computational modelers in the infancy of the field, it may now serve to conceal the robustness of this approach and inhibit its development. It is likely that adherence to this principle has dissuaded researchers from manipulating virtual worlds to conform with theoretical expectations. Meanwhile, confined to simple models with small numbers of moving parts by the imperative of algebraic tractability, closed-form formal modelers lack training or experience in techniques of research design. At the same time, scholars employing qualitative methods—historical institutionalists, political ethnographers, and small-N comparativists—seldom imagine that the nuance they value and the external validity they prize may be achieved with the rigor, clarity, and disciplined opportunities for refinement associated with formal modeling techniques.

The increasing availability of ABM means that formalists will have opportunities for which they have not prepared—opportunities for experimenting with more ambitious research designs with substantial verisimilitude, meaning respect for consensual expert knowledge. For that same reason, qualitatively oriented scholars, accustomed to fine-tuning the design of research projects in relation to the scale and granularity of the questions they wish to ask and the data available, will have opportunities for disciplining their work by expressing the rigor of their thinking through computers and heuristically and experimentally powerful simulation techniques. Standing in the way of these opportunities are misconceptions about the epistemological status of computer assisted agent-based models, tendencies to believe that formal modeling means closed form (algebraically solvable) modeling, and beliefs that attempts to introduce nuance and case or theory-specific knowledge into agent-based models would contaminate them with a level of complexity that only human brains, unassisted by computers, can handle.

We have presented the abstraction, ensemble, virtualization typology of ABM research designs to help inform political scientists unfamiliar with this methodology of its opportunities and to suggest guidelines for the development of research designs for those who are developing their skills with the technique. In the process we have tried to combat a number of scholarly prejudices, but there is much more work to be done. In the end, the proof of the concept and the usefulness of this typology will be in the answers social scientists using ABM techniques in different ways give to questions that stump researchers using other methods and in the excitement generated by the questions those answers prompt.

NOTES

An earlier version of this article was presented at the annual meeting of the American Political Science Association, Chicago, August 30–September 2, 2007. The authors wish to thank Britt Cartrite and David Laitin for useful comments and Chanan Cohen for assistance in the research supporting this paper. The authors are also grateful for support received from the Hebrew University of Jerusalem and the University of Pennsylvania.

1. For general introductions to agent-based modeling for the social sciences, see Joshua M. Epstein and Robert Axtel, *Growing Artificial Societies: Social Science from the Bottom Up* (Washington, D.C./Cambridge, Mass.: Brookings Institution Press/ MIT Press, 1996); Robert Axelrod, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Princeton: Princeton University Press, 1997); Lars-Erik Cederman, "Agent Based Modeling in Political Science," *Political Methodologist*, 10 (Fall 2001), 16–22.

2. Axelrod; Ian S. Lustick, "Agent-Based Modeling of Collective Identity: Testing Constructivist Theory," *Journal of Artificial Societies and Social Simulations*, 3 (January 2000), http://jasss.soc.surrey.ac.uk/3/1/1.html; Michael W. Macy and Robert Willer, "From Factors to Actors: Computational Sociology and Agent-Based Modeling," *Annual Review of Sociology*, 28 (August 2002), 143–66.

3. On parameter spaces and limiting complexity of ABM models, see Scott de Marchi, *Computational and Mathematical Modeling in the Social Sciences* (New York: Cambridge University Press, 2005).

4. Most limits to complexity (restriction to two players, omission of spatial relationships, calibrations of knowledge of others available to actors, and variability of preferences) are standardly adopted without explicit mention. For explicit attention, albeit without further remark, to the omission from their iterated prisoners' dilemma mode of multidimensionality and certain contextual features, see, for example, Avner Greif and David Laitin, "A Theory of Endogenous Institutional Change," *American Political Science Review*, 98 (November 2004), 648, 648n.

5. Lars-Erik Cederman, "Agent-Based Modeling in Political Science," *Political Methodologist*, 10 (Fall 2001), 16–22.

6. Thomas C. Schelling, *Micromotives and Macrobehavior* (New York: W. W. Norton, 1978); Epstein and Axtell; Axelrod; Lars-Erik Cederman, *Emergent Actors in World Politics: How States and Nations Develop and Dissolve* (Princeton: Princeton University Press, 1997).

7. Volker Grimm and Steven F. Railsback, *Individual-based Modeling and Ecology* (Princeton: Princeton University Press, 2005), p. 37.

8. Grimm and. Railsback, p. 367.

9. Axelrod.

10. Originally published by Robert Axelrod in "The Dissemination of Culture: A Model with Local Convergence

and Global Polarization," *Journal of Conflict Resolution*, 41 (April 1997), 203–26.

11. Michael W. Macy and John Skvoretz, "The Evolution of Trust and Cooperation between Strangers: A Computational Model," *American Sociological Review*, 63 (October 1998), 638–60.

12. Stephen J Majeski, Greg Linden, Corina Linden, and Aaron Spitzer, "Agent Mobility and the Evolution of Cooperative Communities," *Complexity*, 5 (1999), 16–24.

13. Yasufumi Shibani, Satoko Yasuno, and Itaru Ishiguro, "Effects of Global Information Feedback on Diversity: Extensions to Axelrod's Adaptive Culture Model," *Journal of Conflict Resolution*, 45 (February 2001), 80–96.

14. Bibb Latané, Andrzej Nowak, and James H. Liu, "Measuring Emergent Social Phenomena: Dynamism, Polarization, and Clustering as Order Parameters of Social Systems," *Behavioral Science*, 39 (1994), 1–24; Bibb Latané and Andrzej Nowak, "Self-Organizing Social Systems: Necessary and Sufficient Conditions for the Emergence of Clustering, Consolidation and Continuing Diversity," in G. A Barnett and F. J. Boster, eds., *Progress in Communication Sciences*, 13 (Norwood: Ablex, 1997), 43–74.

15. Ian S. Lustick and Dan Miodownik, "The Institutionalization of Identity: Micro Adaptation, Macro Effects, and Collective Consequences," *Studies in Comparative and International Development*, 37 (Summer 2002), 24–53; Lustick, "Agent-Based Modeling of Collective Identity."

16. The key idea referenced here is that science advances not by staggering from one consensual paradigm for investigation to another in times of radical and chaotic uncertainty, but as a result of a constant process of competition for attention and resources among conceptually inconsistent "research programs." Moving a research program in a "progressive" as opposed to a "degenerative" direction means generating research that produces and reproduces empirical and theoretical questions for further research interesting enough to attract the attention of talented scholars. Imre Lakatos, "Falsification and the Methodology of Scientific Research Programmes," in Imre Lakatos and Alan Musgrave, eds., *Criticism and the Growth of Knowledge* (Cambridge: Cambridge University Press, 1970).

17. Lustick, "Agent-Based Modeling of Collective Identity." A cellular automata is a grid or an array featuring elements within cells. The cells themselves do not move, but depending on the microrules implemented to animate behavior within these cells and the networks of communication available among them, the array of cells can exhibit movement of traits and behaviors from one unit to another and thereby from one part of the grid to another.

18. See, for example, Schelling, *Micromotives and Macrobehavior*; Axelrod, *The Complexity of Cooperation*; Latané, Nowak, and Liu, "Measuring Emergent Social Phenomena."

19. Information about these changing signals is not included in this display but is directly available as the model runs.

20. Sensitivity studies of various sorts were also run to rule out the artifactuality of these results. See Lustick.

21. Lustick and Miodownik, "The Institutionalization of Identity."

22. Grimm and Railsback, *Individual-based Modeling and Ecology*, pp. 367, 369. Other terms used in ecology and biology for this type of model include heuristic and explanatory. Ibid., p. 366.

23. Ian S Lustick, Dan Miodownik, and Roy J. Eidelson, "Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It?," *American Political Science Review*, 94 (May 2004), 209–30; Dan Miodownik, "Cultural Differences and Economic Incentives: An Agent-Based Study of Their Impact on the Emergence of Regional Autonomy Movements," *Journal of Artificial Societies and Social Simulation*, 9 (Fall 2006), http://jasss.soc.surrey.ac.uk/9/4/2.html; Dan Miodownik and Britt Cartrite, "Does Political Decentralization Exacerbate or Ameliorate Autonomy Mobilization? A Test of Contesting Propositions" (manuscript, 2006).

24. "PS-I: A User-Friendly Agent-Based Modeling Platform for Testing Theories of Political Identity and Political Stability," *Journal of Artificial Societies and Social Simulations*, 5 (June 2002), http://jasss.soc.surrey.ac.uk/5/3/7.html. The PS-I platform is freely downloadable, with documentation. To download the platform and for information on ABIR and PS-I studies cited in this article, including larger versions of the figures, go to http://www.polisci.upenn.edu/ps-i/.

25. Miodownik, "Cultural Differences and Economic Incentives," p. 213; Miodownik and Cartrite, "Does Political Decentralization Exacerbate"; Lustick, Miodownik, and Eidelson.

26. Lustick, Miodownik, and Eidelson.

27. David L. Rousseau and Maurits van der Veen, "The Emergence of a Shared Identity: A Simulation," *Journal of Conflict Resolution*, 49 (October 2005), 686–712.

28. Cederman, *Emergent Actors in World Politics*, pp. 184–212.

29. Armano Srbljinovic, Drazen Penzar, Petra Rodik, and Kruno Kardov, "An Agent-Based Model of Ethnic Mobilisation," *Journal of Artificial Societies and Social Simulation*, 6 (January 2003), http://jasss.soc.surrey.ac.uk/6/1/1.html.

30. Ravi Bhavnani and David Backer. "Localized Ethnic Conflict and Genocide: Accounting for Differences in

Rwanda and Burundi," *Journal of Conflict Resolution*, 44 (June 2000), 283–306.

31. Ravi Bhavnani, "Ethnic Norms and Interethnic Violence: Accounting for Mass Participation in the Rwandan Genocide," *Journal of Peace Research*, 43 (November 2006), 651–69.

32. For the best available study of how experts use scenario-building techniques and how unreliable the process is from a scientific point of view, see Philip C. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton: Princeton University Press, 2005).

33. Grim and Railsback, pp. 49, 366. For a discussion of the closely related question of external validity and ABM, see Lustick, Miodownik, and Eidelson, pp. 211–13.

34. See Ian S. Lustick, "Virtual Pakistan and Its Futures: An Exercise in the Deployment of Constructivist Theory," draft prepared for the Workshop on Constructivist Approaches to Ethnic Identity, Philadelphia, August 26–27, 2003, http://www.psych.upenn.edu/sacsec/abir/_private/publications/VirtualPakistan_CAEG.doc. See also Jeffrey Rothfeder, "Terror Games," *Popular Science*, 264 (February 2004), 82–91, 115.

35. See, for example, Chris L. Barrett, Stephen G. Eubank, and James P. Smith, "If Smallpox Strikes Portland," *Scientific American*, 292 (March 2005), 42–49; F. Bousquet, C. LePage, I. Bakam, and A. Takforyan, "Multi-agent Simulations of Hunting Wild Meat in a Village in Eastern Cameroon," *Ecological Modelling*, 138 (2001), 331–46.

36. George J. Gumerman, Alan C. Swedlund, Jeffrey S. Dean, Joshua M. Epstein, "The Evolution of Social Behavior in the Prehistoric American Southwest," http://www.sarweb.org/home/gumerman/articles/evolsocfinal.pdf;
See also R. L Axtell, J. M. Epstein, J. S. Dean, G. J. Gumerman, A. C. Swedlund, J. Harburger, S. Chakravarty, R. Hammond, J. Parker, and M. Parker, "Population Growth and Collapse in a Multiagent Model of the Kayenta Anasazi in Long House Valley," in B. J. L. Berry, L. D. Kiel, and E. Elliott, eds., *Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-based Modelling*, *Proceedings of the National Academy of Sciences*, 99 (2002), 7275–79.

37. Lars-Erik Cederman and Girardin Luc, "Toward Realistic Computational Models of Civil Wars: A Roadmap to Realistic Computational Models of Civil Wars," paper prepared for the Annual Meeting of the American Political Science Association, Chicago, August 30–September 2, 2007; Nils B. Weidmann, "From Sparks to Prairie Fires: Spatial Mechanisms of Group Mobilization," paper prepared for the Annual Meeting of the American Political Science Association in Chicago, August 30–September 2, 2007.

38. Nils B. Weidmann and Luc Girardin, "GROWLab: A Toolbox for Social Simulation," presented at *The First World Congress on Social Simulations*, Kyoto Japan, August 21–25, 2006; see also http://www.icr.ethz.ch/research/growlab.