The Emergence of a Shared Identity: An Agent-Based Computer Simulation of Idea Diffusion

David Rousseau and A. Maurits van der Veen

Journal of Conflict Resolution 2005; 49; 686
DOI: 10.1177/0022002705279336

The Emergence of a Shared Identity

AN AGENT-BASED COMPUTER SIMULATION OF IDEA DIFFUSION

DAVID ROUSSEAU
Department of Political Science
University at Albany (SUNY)
A. MAURITS VAN DER VEEN
Department of International Affairs
University of Georgia

Liberal theorists and, more recently, constructivists have argued that a shared sense of identity will decrease threat perception and increase international cooperation. But what processes lead to the emergence (or collapse) of a shared sense of identity? Drawing on individual-level experiments of identity change and societal-level theories of idea diffusion from the constructivist literature, the authors develop an agent-based computer simulation of identity change. Four findings emerge from the analysis: (1) there is a curvilinear relationship between the complexity of agents' views and the emergence of shared identity, (2) highly unstable environments encourage the emergence of a shared identity, (3) the presence of leaders (i.e., agents with greater influence in the population) in the neighborhood decreases shared identity, and (4) the interaction of complexity, stability, and leadership produces extremely polarized societies.

Keywords: *identity; agent-based modeling; constructivism; threat perception*

Realists have long argued that power considerations should dominate perceptions of threat in international relations. In an anarchic world characterized by self-help, states must constantly be wary of more powerful states. From the realist perspective, threat is an objective measure that can be calculated from the material balance of power (Gulick 1955; Waltz 1979). Liberals such as Kant ([1795] 1971) and constructivists such as Wendt (1999) have challenged this view by arguing that a shared sense of identity can reduce, or even eliminate, perceptions of threat posed by power asymmetries. For example, Wendt's structural social theory of international politics predicts that a world populated with Kantian states would view other states as friends rather than enemies or rivals (Wendt 1999, 298).

AUTHORS' NOTE: A previous version of this article was presented at the annual meeting of the American Political Science Association, August 30 to September 2, 2001, San Francisco. The simulation code, a parameter dictionary, and a technical appendix describing the model are available from the authors (maurits@uga.edu) or at http://www.yale.edu/unsy/jcr/jcrdata.htm.

JOURNAL OF CONFLICT RESOLUTION, Vol. 49 No. 5, October 2005 686-712
DOI: 10.1177/0022002705279336

From an empirical perspective, constructivists have *described* many different collective identities, including organizational cultures (Kier 1997), national cultures (Berger 1998), and global norms (Finnemore 1996). However, research on the *process* of identity construction has lagged far behind the descriptive research (Kowert and Legro 1996, 46).[1] Why do particular identities, such as antimilitarism or racial superiority, become dominant in particular societies? What causal mechanisms help sustain or undermine the maintenance of these dominant identities? What role do individuals play in the construction and destruction of particular social identities?

This article begins to fill this void in the literature by developing a model of the process of identity formation that explicitly links the individual level (i.e., micro level) with the societal level of analysis (i.e., the macro level). We then test the model using an agent-based computer simulation. Four findings emerge from the analysis. First, there is a curvilinear relationship between the complexity of our views of ourselves and others and the emergence of shared identity. Second, we find that highly unstable environments encourage the emergence of a shared identity. Third, the simulation reveals that the presence of leaders (i.e., agents with greater influence in the population) in the neighborhood decreases shared identity. Finally, the interaction of complexity, stability, and leadership produces extremely polarized societies.

The remainder of the article is divided into five sections. In the next section, we review the literature on identity and identity formation in international relations and comparative politics, highlighting core areas of agreement. Next we develop an individual-level model of the diffusion of identities. Then we derive the hypotheses we will test with our computer model. The fourth section presents the simulation and describes the results of our simulation runs. A concluding section presents our conclusions and their implications for the international relations literature.

IDENTITY IN INTERNATIONAL RELATIONS AND COMPARATIVE POLITICS

In the 1990s, international relations took a "constructivist" turn (Checkel 1998; Katzenstein 1996). This turn can be viewed in many ways as a backlash against neorealism and neoliberalism, which came to dominate discourse in the 1980s. The structural and materialist orientations of neorealism and neoliberalism strove to remove identity from the equation, even though identity had played a central (but often implicit) role in many traditional realist and liberal theories. For example, power transition theory suggests that rising challengers that are satisfied with the status quo are unlikely to trigger a hegemonic war. Although Organski (1968, 362) does not refer to

1. We are not arguing that all constructivist research needs to have a process focus. If a researcher is interested in studying the consequences of a collective belief in a society (e.g., anti-militarism), then he or she need not immediately explore the process through which agents produced the belief. In the short term, one may want to simply assume that the collective belief is reproduced by the society. However, if one is interested in the long-term emergence and decline of collective beliefs, a process-oriented model that links agent interaction to the emergence of societal structures is required (see Sil 2000, 380).

shared identity explicitly, he does state that satisfaction can be a function of a common culture, history, and/or language. Similarly, studies of the balance of power often conceded that a shared culture can facilitate balancing by making the behavior of others more interpretable (e.g., Gulick 1955). Finally, identity has always played a central role in classical liberalism: a shared liberal identity reduces incentives for conflict and perceptions of threat (Kant [1795] 1971).

Within the emerging constructivist literature in international relations, identity has come to the fore again. Shared identity, in particular, features in a wide variety of arguments. Hopf (1998) argues that a shared identity may reduce the security dilemma. Kahl (1998-1999) suggests that shared identity may be a third causal mechanism (in addition to norms and structures) behind the democratic peace. Wendt (1999) claims that anarchy does not condemn states to a Hobbesian world of fear and conflict because some shared identities can decrease the expectation of violent conflict. Empirically, the experimental analysis of Rousseau (2002, forthcoming) has shown that a shared identity is associated with a reduction in both threat perception and the salience of relative gains.

Thus, shared identity plays a central role in a wide variety of traditional and recent theoretical arguments in the international relations literature. However, although shared identities are often described in great detail, they tend to be treated as exogenous variables. To help us develop a general model of the *process* of identity formation linking individual and societal levels, we must turn to the comparative politics literature.

The process of identity construction has been most intensively explored within the fields of ethnic conflict and nationalism. In a review of this literature, Chandra (2001, 7) argues that the primordialist view, which claims identity becomes fixed once it is acquired (either socially or biologically), has been largely replaced by the constructivist view over the past several decades. While all constructivists agree that individuals have multiple identities and that causal factors can trigger a shift from one identity to another, they disagree on the relative importance of these causal factors (e.g., industrialization, colonialism, economic opportunism, or entrepreneurship) (Chandra 2001, 8).

Despite continuing debate, Lustick and Miodownik (2002, 25) contend that a consensus has emerged with respect to six central elements of constructivism in comparative politics. First, as already mentioned, identities can change across time and space. For example, politics in present-day Zambia are a function of linguistic identities that emerged over the past seventy years through missionary activity, colonial policy, and economic incentives. Bemba, which was only spoken by about 7 percent of the population in 1930, is now the first or second language of close to 40 percent of the population (Posner 2003, 129). While it was by no means inevitable, shared language has become the dimension used to draw the line between "us" and "them" in Zambian politics.

Second, individuals possess multiple identities and often shift from one identity to another. "Multiple identities are common in social life," and the context of the situation can alter which identity moves into the foreground (Laitin 1998, 23). For example, a resident of Tallinn with Estonian, Russian, and European identities may shift to a

European identity when giving a paper in New York City and to a Russian identity when voting in local elections. While Laitin (1998) points out that the emergence of conflict among the multiple identities can force a choice (e.g., can a real Estonian also be a Russian?), most of the time, individuals operate quite efficiently with a collection of (mostly latent) identities.

Third, political, economic, and social incentive structures can influence the adoption of specific identities. Individuals adopt identities in response to economic incentives such as access to patronage in the form of jobs, markets, and land. In this situation, identity is a choice based on existing interests and rational calculations (e.g., Bates 1974). Along similar lines, Lustick, Miodownik, and Eidelson (2004) argue that political institutional structures can promote the adoption of some identities and discourage the adoption of others. Moreover, some observers have argued that the adoption of new identities may result from mechanisms even more basic than instrumental rationality, such as simple algorithmic reactions or even reflexive/instinctual behaviors (Lustick and Miodownik 2002, 25).

Fourth, the selection or emergence of an identity from a collection of possible identities depends on social interaction. This aspect of identity formation is particularly prominent in the literature on nationalism. Thus, Gellner (1983) argues that industrialization and the subsequent rise of literacy has shifted and deepened the links in social networks. The altering of the social networks undermined existing local identities and facilitated the emergence of national identities. Similarly, Anderson (1983) contends that an imagined community requires both a shared vocabulary and access to common information. A social network that allows human interaction is necessary for the emergence of a vocabulary, the spread of information, and eventually the emergence of a collective identity.

Fifth, the identities prevalent within a society are disproportionately shaped by political entrepreneurs. Intellectuals and politicians can occupy "a site that is privileged in formatting and transmitting discourses" (Verdery 1991, 15). Greater access to resources and the means of communication give such entrepreneurs greater persuasive power in the social discourse. Thus, Brass (1974, 1997) argues that entrepreneurs attempting to mobilize groups for their own political purposes can have a dramatic impact on identity formation. Similarly, in her study of identity formation in Romania, Verdery (1991, 15) shows that entrepreneurs play a central role in shaping values and legitimating social orders.

Finally, and implicit in most discussions of identity formation, constructivists assume that actors and structures are mutually constituted (Lustick and Miodownik 2002; Wendt 1999, 139; Hopf 1998, 172). This implies that the interaction of actors creates a social structure that, in turn, regulates and constitutes the actors. The mutually constitutive assumption makes it impossible to talk about actors without reference to structure (and vice versa). Laitin (1998, 33) applies this logic in his "tipping model": as an identity becomes more dominant in the population, it can reach a tipping point in which it spreads rapidly throughout a population and becomes firmly entrenched (or hegemonic). Whether such a tipping point is reached depends on a wide variety of contextual factors. Thus, Laitin argues that the linguistic tipping point was reached in the Ukraine but not in Kazakhstan.

The model developed in the following sections is consistent with all six of these central claims from the comparative politics literature. However, our model differs from this literature in two ways. First, the ethnic conflict and nationalism literature focuses on the formation of a single identity within a country or region. In contrast, our model focuses on two identities: the construction of the collective Self (e.g., what is America?) and the construction of the collective Other (e.g., what is China?).[2] Although the process is similar in each case, the international relations field is more concerned with the perception of a *shared* identity (e.g., is China becoming more like America?). Second, in many instances, the theoretical literature provides little or no guidance for critical model parameters because the high level of abstraction in the texts does not compel researchers to address each element of the complex process. For example, while many authors claim that actors have multiple identities, they provide little guidance on the number of identities available or the specific thresholds triggering a shift between identities. Our approach has been to develop a model as consistent as possible with the literature and to parameterize any variables (e.g., number of identities) that are either contested in, or ignored by, the literature.

CONSTRUCTING SHARED IDENTITY AT THE INDIVIDUAL LEVEL

In keeping with the constructivist emphasis on mutual constitution, our model privileges neither agents nor structures. Following Giddens (1984) and Onuf (1989), we contend that any complete model of identity formation and identity evolution must incorporate multiple levels of analysis. Specifically, a complete model must address how individuals construct ideas within their minds (i.e., the individual level) and how ideas diffuse across a population though social interaction (i.e., the societal level).[3] The construction of a bridging model ensures that neither the agents nor the structures are analytically prior and/or fixed across time and space. While these ideas have been discussed in general terms in the literature for literally generations (e.g., Weber [1914] 1978; Durkheim [1897] 1951), most of the recent work on identity in comparative politics and international relations has focused on the societal-level component of the model. Thus, we turn to the theoretical foundations of the individual-level component.[4]

Psychologists have long been interested in identity construction. As reflective humans, we constantly evaluate the question "Who am I?" (Baumeister 1998). The answer to this simple question serves as a foundation for organizing relations with others (Brewer and Brown 1998). If you define yourself as a Republican (or homosexual or Catholic or white), the definition influences how you respond to both members of the in-group (e.g., other Republicans) and out-groups (e.g., Democrats or Libertari-

2. Although we do not repeatedly use the term *collective* identity throughout the article, we focus exclusively on individual assessments of the in-group and the out-group.

3. Sil (2000) identifies several other authors attempting to construct a micro-macro link, including Collins (1981), Giddens (1984), Touraine (1985), Bourdieu (1998), and Sewell (1992).

4. For example, Smith (2000, 161) critiques Wendt for his failure to address the individual level in his theory.

ans). Within the context of international relations, individual definitions of the Self (such as American or French or Indian) can influence the definition of the Other (such as Chinese or German or Pakistani) and the sense of shared identity produced by these two definitions.

But how do individuals construct state identities, and how stable are these constructions across time? Consider the process by which individuals construct opinions in response to survey questions. Survey researchers have long been interested in not only the opinions of respondents but also how respondents go about answering a question. As the cognitive revolution in psychology diffused into the field of survey research, survey researchers began to test alternative models of answer construction. Sudman, Bradburn, and Schwarz (1996, 58) present a model with four steps. First, the respondent must interpret the question. While survey researchers have long known that the wording of the question can have a profound impact on interpretation, a number of other factors influence how we comprehend a question, including the subject matter of prior questions, the context of the situation, and the nature of the interviewer-subject relationship (Converse and Presser 1986; Sudman, Bradburn, and Schwarz 1996). Second, the respondent must generate an opinion. In some cases, the individual simply retrieves the opinion from memory. This occurs for simple factual questions (e.g., how many brothers do you have?) or questions about which the individual has thought deeply and holds strong views (e.g., do you support a woman's right to have an abortion?). However, in many, if not most, instances, individuals do not hold readily accessible responses for survey questions.[5] When faced with this situation, individuals compute a response using accessible information. How this is done is still a matter of great controversy. Moreover, the process may well vary from situation to situation (Gigerenzer, Todd, and the ABC Research Group 1999). The traditional explanation proposes that individuals simply balance salient considerations for and against the proposition retrieved from memory. A key element of Zaller's (1992, 49) widely praised model of attitude change is the "response axiom" that claims "individuals answer survey questions by averaging across considerations that are immediately salient or accessible."[6] Third, the respondent must format the response for closed-ended questions. Finally, the respondent must decide whether to edit the response. Individuals may feel reluctant to report to the interviewer that they watch twenty-seven hours of television a week, that they hold racist views, or that they engage in risky sexual behavior.

We believe that a similar process takes place when constructing identities of the Self and Other. Suppose a survey research organization such as Gallup phoned you and asked the following question: "Should Japan become a permanent member of the

5. "Most of the answers that we record in surveys reflect judgments that respondents generate on the spot in the specific context of the specific interview. They are therefore strongly influenced by the information that is accessible at that time and this is in part a function of the preceding questions" (Sudman, Bradburn, and Schwarz 1996, 70).

6. For an interesting critique of this balancing model, see Gigerenzer, Todd, and the ABC Research Group (1999, 22). They suggest that individuals use heuristics or "cognitive shortcuts" to speed the processing of information and reduce cognitive demands. Contrary to much of the research in the area, Gigerenzer, Todd, and the ABC Research Group claim that these heuristics might be as efficient as more complex decision procedures.

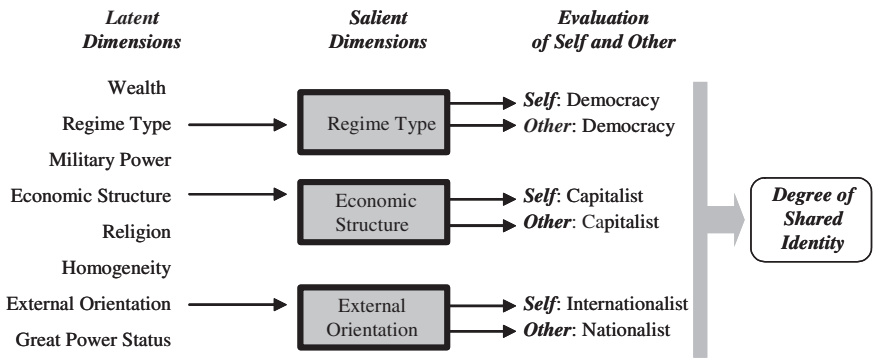

Figure 1: Jane Doe's Construction of an Identity of the "Other"

United Nations Security Council?" Immediately, you must construct some image of Japan. Is it a "good" country? Is it trustworthy and cooperative? If Gallup were able to ask you this question repeatedly across time, we would find that you relied on a handful of dimensions to categorize the other state. That is, while you would probably not construct the same image of Japan on every occasion, there would be a discernible pattern to your responses across time.

In Figure 1, we present the hypothetical case of "Jane Doe." Jane tends to use eight different elements to evaluate Japan, ranging from wealth to great power status. However, these dimensions are *latent* in that they are in memory but not necessarily immediately available. On any given day, only a subset of the latent dimensions will be *salient* or readily accessible. Only dimensions that are salient influence the construction of the opinion (e.g., influence the aggregation process). In Figure 1, we see that on this particular day, three dimensions are salient: regime type, economic structure, and external orientation. Jane Doe evaluates both her own country and the other country using these salient dimensions. She concludes that both countries are democratic and capitalist but that the United States is more internationalist than Japan. Her net assessment is that the two states are pretty similar (but not identical).

The process is iterative in that the individual's assessment of the Self may make certain dimensions more salient for the construction of the Other and vice versa (Hopf 2002). So when Jane Doe thinks of the United States, certain dimensions come to mind, and when she thinks of Japan, other dimensions come to mind. Ultimately, a comparison will be made on each dimension because *both* countries are salient. For example, if Jane Doe believes Japan is different because it is Buddhist, she is implicitly claiming that the United States is non-Buddhist.[7]

7. Although more Japanese profess adherence to the Shinto religion, this "objective" fact is irrelevant to the subjective assessment of Jane Doe. If she believes Japan is a Buddhist country, her behavior will reflect this belief.

Rousseau (forthcoming) explored this model of the construction of the Self and Other within the context of international relations in an experimental setting. Several findings emerged from the individual-level experiments. First, the experiments demonstrated that individuals vary greatly in the number of dimensions used to evaluate other states. While the mean number of dimensions used was 5.4, the number of dimensions employed by respondents varied from two to over a dozen. Second, the results indicated that more knowledgeable individuals employed more dimensions to evaluate the Other. High-knowledge respondents identified an average of 2.6 similar dimensions (e.g., the United States and Japan are both democracies) and 3.1 dissimilar dimensions (e.g., the United States and Japan have different dominant religions). In contrast, low-knowledge respondents identified only 1.7 similar and 2.5 dissimilar dimensions, on average. Third, the analysis indicated that states viewed as similar were seen as less threatening from a military perspective. Finally, priming individuals to view states as similar increased respondents' willingness to cooperate with this Other, even if cooperation resulted in relative losses. Overall, the experimental results strongly supported the micro-level model of identity construction.

HYPOTHESES

By linking the societal level with the individual level, we can pose a number of interesting hypotheses. Our goal is to explore how changes in the *attributes of the individual* and/or *elements of the environment* influence our dependent variable: the degree of shared identity in a population. It is important to emphasize that we are interested in the spread of ideas across a population rather than convergence on "truth." Whether the Other is capitalist or democratic or isolationist is irrelevant because if an individual or group accepts the claim as "truth," it can influence behavior toward the out-group.

Our first hypothesis, which focuses on an attribute of the individual, predicts that the greater the number of dimensions used to compare the Self to the Other, the lower the level of shared identity in the population. While very simple views of the Self and Other (e.g., a democracy or not) can spread rapidly through a population, more complex views (e.g., some combination of regime type, economic structure, religion, ethnicity, military power, wealth, *and* external orientation) are likely to spread slowly and lead to a decline in shared identity. Greater complexity raises the number of hurdles that must be jumped for a sense of shared identity to emerge. Moreover, complexity increases the likelihood that a stable sense of shared identity will collapse as latent identities suddenly become salient.

The second hypothesis focuses on the environment within which individuals are embedded. We predict that fluctuations in the incentives associated with particular identities will increase the level of shared identity in the population. Rigidly controlled status quo societies, which discourage new ideas and alternative conceptualizations, are less likely to alter constructions of the Self and Other. In contrast, societies in which ideas continuously fall in and out of favor are more likely to settle upon a shared identity. While the fluctuations in incentives may slightly undermine a shared identity

once established, they greatly enhance the likelihood of a shared identity emerging in the first place.

Finally, hypotheses 3 and 4 focus on the impact of specific, influential individuals on identity formation within the population. We expect that the presence of leadership will increase the level of shared identity in a population. In general, political and social leaders have disproportionate power in terms of their ability to transmit ideas to others. They have greater access to financial and human resources as well as links to communication channels extending across a population. By increasing the homogeneity of views in the population, leaders increase the probability of shared identity. Our model distinguishes between two types of leaders: leadership with more power in the local neighborhood and leaders with greater reach across the population. Hypothesis 3 predicts that leaders with more power in the local neighborhood will increase shared identity. Similarly, hypothesis 4 predicts that leaders with the ability to reach beyond the immediate neighborhood will increase shared identity.

SIMULATIONS

These four hypotheses are general claims that could be tested using a wide variety of empirical methods. The complexity hypothesis, for example, could be tested using an experimental survey measuring the number of dimensions used to evaluate other countries and a self-report of the sense of shared identity (Rousseau, forthcoming). The leadership claims could be tested using a content analysis of elite beliefs and surveys of mass opinion (Zaller 1992). Finally, fluctuations in the desirability of particular identity dimensions could be examined using case studies stretching over years or decades (Hopf 2002). In fact, "triangulation" through the use of multiple methods of investigation is the best way to establish the robustness of these causal claims (Abdelal et al. 2004).

In this article, we test the four hypotheses using an agent-based computer simulation. Agent-based simulations are "bottom-up" models that probe the microfoundations of macropatterns. Agent-based models assume that actors are autonomous, interdependent, simple rule followers and adaptive or backward looking (Macy and Willer 2002, 146). These models are ideally suited for complex, nonlinear, self-organizing situations involving many actors. Macy and Willer (2002, 148) claim that agent-based models are "most appropriate for studying processes that lack central coordination, including the emergence of organizations that, once established, impose order from the top down." Thus, agent-based models appear to be ideally suited to examine the emergence of collective identity from the interaction of thousands of independent actors. While agent-based models have been used in a variety of identity-related applications in recent years—for example, the emergence of norms (Axelrod 1997), the transformation of identity (Lustick, Miodownik, and Eidelson 2004), the evolution of cultural cooperation (Macy and Skvoretz 1998), and the emergence of nationalism (Cederman 1997)—to our knowledge, no other computational model simultaneously addresses the process of identity formation, multidimensional identities, latent identities, saliences, incentive structures, and leadership.

As with all methods of investigation, computer simulations have strengths and weaknesses.[8] On the positive side of the ledger, five strengths stand out. First, as with formal mathematical models, simulations compel the researcher to be very explicit about assumptions and decision rules. Second, simulations allow us to explore extremely complex systems that often have no analytical solution. Third, simulations resemble controlled experiments in that the researcher can precisely vary a single independent variable (or isolate a particular interaction between two or more variables). Fourth, while other methods of inquiry primarily focus on outcomes (e.g., do democratic dyads engage in war?), simulations allow us to explore the processes underlying the broader causal claim (e.g., how does joint democracy decrease the likelihood of war?). Fifth, simulations provide a nice balance between induction and deduction. While the developer must construct a logically consistent model based on theory and history, the output of the model is explored inductively by assessing the impact of varying assumptions and decision rules.

On the negative side of the ledger, two important weaknesses stand out. First, simulations have been criticized because they often employ arbitrary assumptions and decision rules (Johnson 1999, 1512). In part, this situation stems from the need to explicitly operationalize each assumption and decision rule. However, it is also due to the reluctance of many simulation modelers to empirically test assumptions using alternative methods of inquiry. We attempt to minimize this problem by drawing on experimental studies of identity construction and theoretical models of idea diffusion in the survey research, cognitive psychology, and the constructivist literatures. Second, critics often question the external validity of computer simulations. While one of the strengths of the method is its internal consistency, it is often unclear if the simulation captures enough of the external world to allow us to generalize from the artificial system we have created to the real world we inhabit. However, this shortcoming is not limited to agent-based modeling; all models (e.g., formal and statistical) are simplifications designed to isolate the essential elements of complex systems or processes. Moreover, we probe the external validity of our model using a case study of a contemporary international conflict elsewhere (Rousseau, forthcoming).

THE SIMULATION MODEL

Figure 2 illustrates the basic structure of our shared identity simulation.[9] The agent-based model consists of a population of "agents" that interact on a "landscape." In the simulation runs described below, the population consists of a 50×50 square grid of agents that wraps around the edges both horizontally and vertically (i.e., a torus). Each of the 2,500 individual agents is surrounded by eight immediate neighbors, forming a

8. For a more extensive discussion of strengths and weaknesses of agent-based modeling, see Rousseau (forthcoming), Johnson (1999), Axtell (2000), and Lustick, Miodownik, and Eidelson (2004).

9. Our simulation was originally inspired by the agent-based identity repertoire (ABIR) model developed by Lustick (2000; Lustick and Miodownik 2000, 2002; Lustick, Miodownik, and Eidelson 2004). The original ABIR model has been expanded and improved. The new more general model, which is referred to as PS-I, is available at www.sas.upenn.edu/~Lustick.

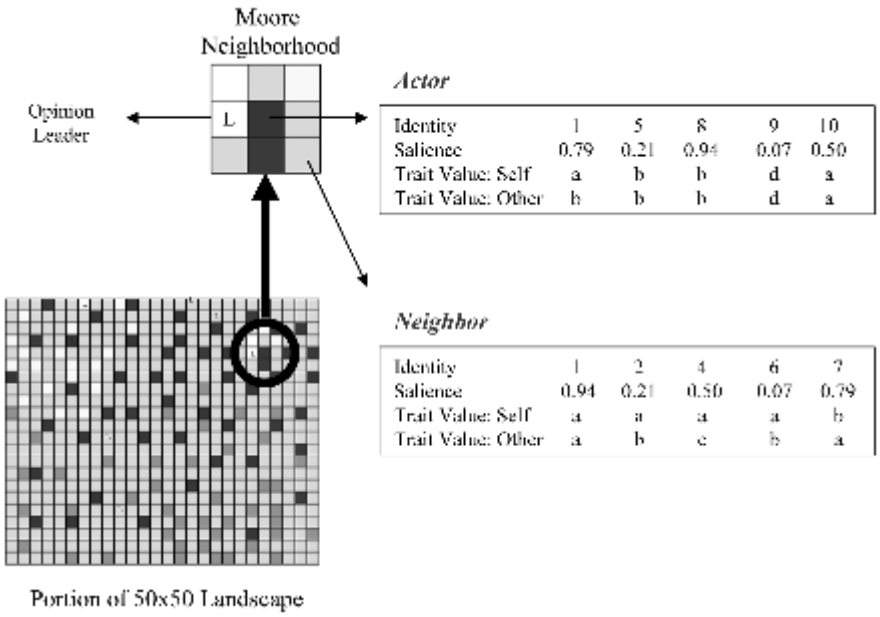

Figure 2: Illustration of the Model

local social network (referred to as the Moore 1 neighborhood in the simulation literature). Each agent in the landscape has an identity repertoire composed of five identity dimensions. As the figure indicates, the salience of each dimension varies from zero to one. Each identity dimension also has four trait values (labeled *a* through *d* in the figure), and these trait values are used by agents to determine similarities and differences between the agent and its neighbors. When calculating similarity, agents are influenced by immediate neighbors, leaders, and global bias. The identities of the agents, which are randomly assigned at the start of the simulation, are updated in each of the 1,000 iterations of a run. All the key features of the model have been parameterized to allow the user to alter variables such as the size of the landscape, the composition of agents, and/or the interaction rules. The model, which was written in a publicly available Java-based simulation package (called Repast, version 1.4.1) developed at the University of Chicago, can be downloaded for use from the *Journal of Conflict Resolution*'s replication Web site. Although the technical appendix on the Web site contains a full description of the simulation, we will briefly discuss seven important features of the simulation here.

The organization of our hypothetical society into a square grid with only neighbor-to-neighbor interactions is, of course, a simplifying assumption. However, it is critical to understand that the grid represents social connections among agents, not necessarily geographical ones. We are capturing the empirical fact that for most people, most of

their friends are also friends with one another. Similarly, most of an individual's closest colleagues in a department also interact with other colleagues in the department. In other words, most social connections are, in a way, "local," even if none of these friends or colleagues "lives" close to one another.

Repertoires. Although each agent has a repertoire size of 5, the universe of possible identity dimensions is set at 20 for the simulations reported here. Thus, each agent has only a small subset of the possible identity dimensions. Moreover, over the course of the simulation, identity dimensions often drop out of the landscape entirely as they become less and less prevalent among agents. The sizes of the repertoire and identity spectrum in the baseline model are based on experimental research with human participants. Rousseau (forthcoming) showed that respondents employed on average 5.4 identity dimensions to compare the United States to other countries. Moreover, he found that the respondents employed just over 20 identity dimensions to evaluate other countries in open-ended questions. The most common dimension—regime type —was only used in approximately 14 percent of the cases. This diversity implies that there is no consensus on the most appropriate dimensions to evaluate the Other. This same research has indicated that knowledge of international affairs is positively correlated with the number of dimensions used to evaluate the other. Therefore, varying repertoire size (or the complexity of images of the Self and Other) could be seen as varying the knowledge of the population.

Trait values. Each identity dimension can take on a number of trait values. For example, a hypothetical "religion" dimension could take on trait values such as Catholic, Muslim, Buddhist, or Other. For the sake of generality, we simply label these trait values *a*, *b*, *c*, *d*, and so on in the model and illustrations. In the baseline simulations, there are four possible trait values for each identity dimension. Each agent has two trait values: one indicating the agent's perception of its own group's trait value on that particular dimension and one indicating the agent's perception of the value for the "other." Thus, an agent with a highly salient "religion" dimension might see the Self as a Catholic country and the Other as a Muslim country. At the start of the simulation, initial trait values are chosen at random for the agents' own identities.[10] Initial beliefs about the trait values of the Other are determined using "priors." Each agent has its own prior belief about the degree to which the Other is similar or different. These priors are initialized to a random value between 0 and 1 at the start and are subsequently updated as the agent updates its repertoire and finds itself sharing more or less of its own repertoire with its perception of the Other's repertoire.

Salience. Laitin and Posner (2001, 15) argue that the salience of different identity dimensions can vary over time. Although some simulations treat salience in binary

10. Although traits are randomized in the results discussed here, the user can alter the distribution of traits in a wide variety of ways (e.g., 80 percent Catholic, 5 percent Muslim, 5 percent Buddhist, and 5 percent Other).

terms, as either 0 or 1 (e.g., Lustick 2000), in our model, saliences range between 0 and 1 and sum to 1 for each actor. Moreover, the saliences are scaled by a logistic function so that only a few dimensions will be of high salience to an agent at any particular time. Identity dimensions whose salience ranks them at the bottom of an agent's repertoire will tend to be eliminated over time.

Bias. During the course of the simulation, each agent receives information from two sources: (1) agents in the Moore neighborhood and (2) global bias. Global bias captures the fact that societies reward some identities (i.e., positive bias) and punish others (i.e., negative bias). For example, during the cold war, Americans could choose "Communism" on the political ideology dimension, but only at great personal cost. There was a bias against the expression of this identity, and this probably inhibited its spread relative to other identities. Global biases could be fixed (e.g., always favor one identity), random (e.g., alter bias after every five time steps), or endogenous (e.g., as the identity becomes prevalent in the population, the positive global bias increases). In our baseline simulation, bias ranges from +3 (i.e., the identity is three units more attractive) to –3 and changes randomly to capture the fact that certain ideas wax and wane in a population across time for reasons independent of the variables in the simulation. In the baseline simulations, the biases have a 1/100 chance of changing in each iteration.[11]

Leaders. There are two types of agents in our model: basic agents and leaders. All leaders differ from basic agents in that leaders update their repertoires and trait values in each iteration before any of the basic agents begin updating. In addition, the simulation allows leaders to differ from basic agents in a number of ways, including (1) greater power, (2) greater range of interactions, (3) degree of factionalization, (4) more ideological (i.e., less likely to change identities), or (5) more pragmatic (i.e., more likely to change identities). In this article, we focus on power and range. In one set of runs, leaders have an influence level twice that of basic agents (i.e., 2 rather than 1). This implies that interacting with a leader is similar to interacting with an additional agent whose repertoire and trait values are identical to those of an agent already in your neighborhood. In the second set of runs, leaders are able to connect with basic agents beyond the Moore 1 neighborhood.

Updating. The updating of repertoires, saliences, and trait values is the heart of the simulation. In each iteration of the simulation, an agent must decide whether to update the contents of its identity repertoire by observing its neighbors and the global bias. For example, if neighbors, on balance, consider the dimension of religion important, the agent is more likely to add religion to its own repertoire or, if the dimension is already present in the repertoire, to increase its salience. It is vital to note that we do not believe that if most of your neighbors are Christian (or Jewish or Muslim), you will

11. Although trait values also have biases in our simulation, we restrict our focus to identity dimension bias in this article.

become a Christian (or a Jew or a Muslim). Rather, if most of your neighbors define the United States as a Christian society, then you are more likely to adopt this value on the religion dimension than if your neighbors did not hold this belief.

In each iteration, the agent composes a list of the overall value of each possible identity dimension in the neighborhood by summing up the saliences associated with each dimension in its own repertoire and in that of each of its neighbors. The agent then adjusts the list by adding the global bias associated with each dimension. If an identity dimension exceeds the agent's "Add Dimension Threshold," the popular identity is inserted into the repertoire (and the least salient identity is dropped). If the identity dimension exceeds the agent's "Reinforce Dimension Threshold," the salience of the popular identity is increased. Finally, extremely popular identities can be both added and immediately increased in salience. Trait values for the Self and Other are adjusted in a similar manner.[12]

Using the identity dimensions, saliences, and trait values, we can calculate the sense of shared identity between the Self (i.e., the in-group) and the Other (i.e., the out-group). The overall sense of shared identity is calculated by summing the saliences for the dimensions in which the self-trait value and the other-trait value are identical and dividing by the sum of the saliences of all dimensions in the repertoire. The resulting value will range from 0 (no identical trait values) to 1 (all trait values identical). The average shared identity can be used to calculate the polarization of society, the variance of identities across time, and the degree to which identities cluster spatially.

The preceding discussion reveals that the simulation captures all six of the central elements of identity construction highlighted in the ethnic conflict and nationalism literatures. First, identities are allowed to vary across time and space (Posner 2003). The interaction of the agent with its Moore neighbors can alter the contents of the repertoire, the salience of dimensions, and the values of traits. Second, agents can have multiple identity dimensions, and the environmental context can shift agents from one identity to another (Laitin 1998). Third, the environment can provide incentives (and disincentives) for the adoption of particular identities (Bates 1974). In our model, the global bias alters the incentives associated with particular identity dimensions. Fourth, identities evolve through social interaction at the individual level (Risse 2000). In the simulation, the social network is the Moore neighborhood that captures high-contact social circles (rather than geographic proximity). Fifth, the model captures the fact that entrepreneurs will have more impact in the process (Brass 1974). Finally, the agents and the structure are mutually constituted (Lustick and Miodownik 2002). Changes in agents can trigger the emergence of an identity structure that then constrains the identity opportunities for the individual agents.

SIMULATION RESULTS

The general hypotheses discussed above can be restated using the language of the parameters in our simulation. In this section, we test these hypotheses individually

12. The technical appendix on the replication Web site describes the updating process in detail.

before examining interactions between the variables. The power of the simulation is most evident in exploring these complex interactions because it is often impossible to deduce the consequences of interaction variables in a complex environment of thousands of individual interactions. All the hypotheses were tested using thirty simulations runs, each of which lasted 1,000 iterations or rounds.

Hypothesis 1: The larger the *repertoire size* of the basic agent, the lower the average shared identity in the population.
Hypothesis 2: The greater the range of the *global bias*, the higher the average shared identity in the population.
Hypothesis 3: The introduction of leaders with greater *power* than basic agents will increase average shared identity.
Hypothesis 4: The introduction of leaders with greater *reach* than basic agents will increase average shared identity.

In the following figures, three graphs are presented for each simulation run similar to those shown in Figure 3. The top row is a landscape display of all agents using color coding to measure the degree of shared identity: darker red indicates lower shared identity, white represents neutrality, and darker blue indicates higher shared identity. The second row of the figure provides a histogram of the distribution of shared identity. The first bar counts the number of agents with a shared identity between 0 and 10 percent, the second between 10 and 20 percent, and so on. Finally, the lower row in the figure measures average shared identity across time (where 1.0 is equal to 1,000 iterations). In most figures, we present time plots of the average of thirty runs to show general trends. For the runs reported in Figure 3, all actors were "basic" agents with a repertoire size of five and four trait values for each identity dimension.

Figure 3 displays the results of the simulation with the baseline settings at iterations 1, 500, and 1,000. At the initialization of the simulation, identity dimensions and trait values are randomly distributed among agents. The random scattering of color in the landscape image in the upper left reveals no distributional pattern. The histogram in the middle row on the left reveals that slightly more agents view the actors as different (i.e., slightly skewed to the left), but this is simply due to chance at initialization. Finally, the left-hand side of the time plot at the bottom reveals that average shared identity is just over 40 percent at the start of the simulation. By iteration 500, we see that many more agents view the Other as very different from the Self, as indicated by the large number of dark red agents in the landscape plot and the left-skewed histogram in the middle of the figure. However, by iteration 1,000, there is a noticeable increase in the perception of shared identity, indicated by the increase in the number of dark blue agents and the more symmetrical histogram. The time plot at the bottom reveals that shared identity dropped to a mere 20 percent by iteration 200 before recovering to about 43 percent by iteration 1,000. Finally, the landscape plots reveal that identities tend to cluster spatially over time as neighborhoods reach agreement on the nature of the Self and Other.

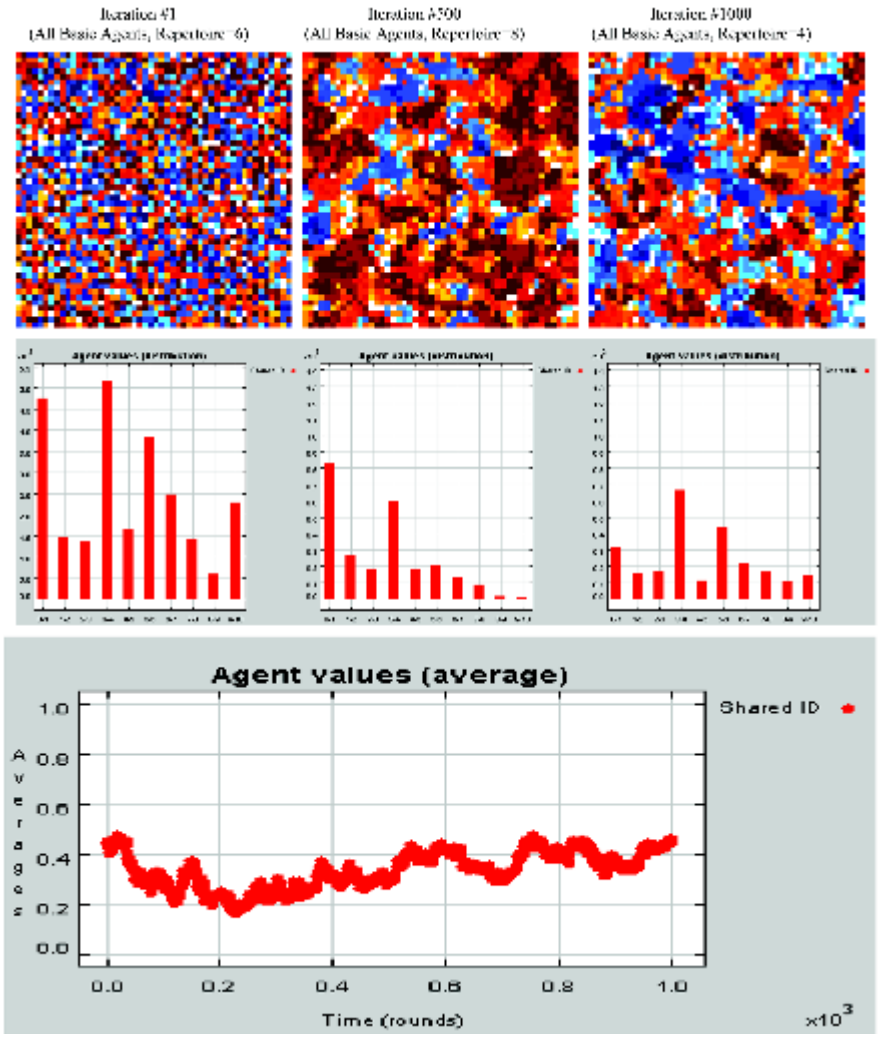

Figure 3: Shared Identity at 1, 500, and 1,000 Iterations

HYPOTHESIS 1: VARYING REPERTOIRE SIZE

Hypothesis 1 predicts that increasing the repertoire size of the agents in the population will reduce the sense of shared identity. We test this by examining the average shared identity in populations with repertoire sizes of 3, 5, and 7. Figure 4 displays "representative" landscape plots and histograms for runs at 1,000 iterations for repertoire sizes 3 (left), 5 (middle), and 7 (right). While a single run cannot tell the observer much about general patterns, we selected a "typical" run after looking at the averages

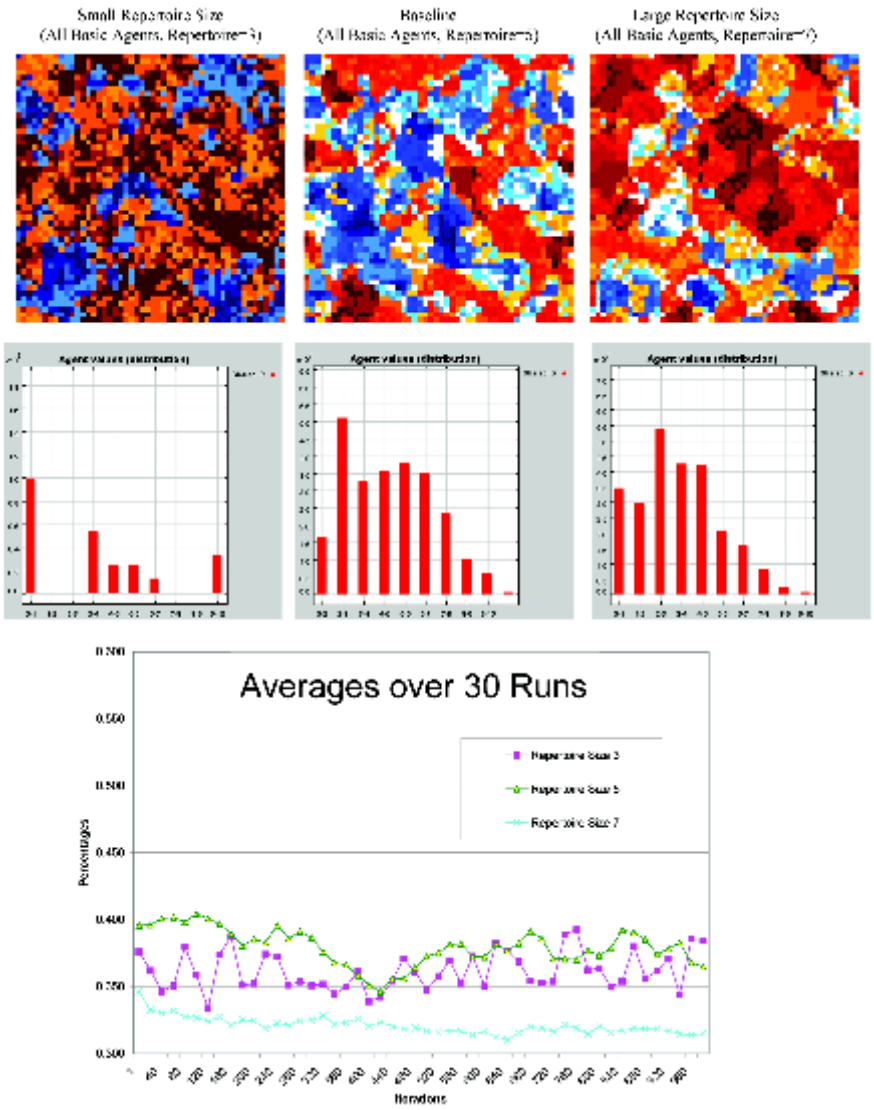

Figure 4: Shared Identity with Various Repertoire Sizes

over the thirty runs. In contrast, the time plot at the bottom of the figure displays average values over thirty runs across 1,000 iterations. The *T* tests and *F* tests reported below are based on the thirty runs of the simulations.

The spatial plots at the top of Figure 4 reveal the impact of altering the repertoire size of agents in the landscape. The high percentage of dark red agents in the low-repertoire case (left) and high-repertoire case (right) hints at a curvilinear relationship—

the highest level of shared identity occurs in the medium-repertoire size of the baseline simulation. While all the landscape plots indicated a great deal of spatial cluster (i.e., clustered reds and blues), none of the landscapes reveals the emergence of hegemony. The left-leaning histograms for the low- and high-repertoire cases indicated a lower shared identity in these particular runs. Finally, the time plot of averages over the thirty runs at the bottom of the figure also indicates a curvilinear pattern because the line for the middle case reaches the highest average shared identity. Moreover, the time plot reveals that there is much more variance within runs (i.e., the lines moving up and down) for the lower repertoire sizes. The runs with a repertoire size of 3 show a great deal of oscillation. When an agent has a repertoire size of 7, adding a new identity dimension is unlikely to profoundly alter a sense of shared identity. In contrast, adding a new dimension with a repertoire size of 3 can shift your sense of shared identity considerably.

The means again reveal an interesting curvilinear pattern: the mean shared identity for repertoire sizes 3, 5, and 7 is 36.4, 38.2, and 33.4, respectively. A *T* test of means reveals that the difference in means between repertoire sizes 3 and 5 is statistically significant at .007, and the difference between sizes 5 and 7 is statistically significant at better than .001. An examination of variances over the thirty runs reinforces the findings from the representative runs. The variance for repertoire sizes 3, 5, and 7 is .007, .006, and .001, respectively. An *F* test of variances reveals that all the variances are statistically different at least at the .05 level. In sum, the simulation reveals that increasing the repertoire size has a curvilinear impact on the mean of shared identity and has a negative impact on the variance of shared identity across time.[13]

Are these results consistent with others from the constructivist literature employing agent-based models? Using the agent-based identity repertoire (ABIR) model, van der Veen, Lustick, and Miodownik (2001) also find a curvilinear relationship as repertoire size increases. They argue that low levels of diversity in an identity promote local clustering but inhibit the spread of a single identity across the population. This clustering is evident in the dark blue and dark red agents in the low-repertoire condition. At medium levels of diversity in the repertoire, single-identity dimensions are likely to be found in the vast majority of actors. Finally, when repertoires are large and very diverse, it is more difficult for a single pattern to emerge in all actors. Thus, they find an inverted U-shaped relationship between repertoire size and identity.

HYPOTHESIS 2: ALTERING IDENTITY DIMENSION BIAS

In our second set of experiments, we examined the impact of environmental volatility on the emergence of shared identity. As described above, agents use both *local signals* in the form of information from the eight immediate neighbors in the Moore neighborhood and *global signals* in the form of identity dimension bias when calculating the salience of identities. The signals are "global" in the sense that they influence the calculations of everyone in the landscape. The global signals capture the fact that

13. Due to the unequal variances, we estimated the *T* tests assuming heteroscedastic variances.

institutional structures and societal norms tend to provide advantages for the propagation of some identities and disadvantages for others. In some societies, the global biases change extremely slowly because new beliefs are suppressed by either the government or powerful societal actors. In other more open and/or less traditional societies, global bias can change more rapidly (e.g., as a result of changing administrations in democratic polities).[14]

The impact of bias is difficult to predict due to countervailing pressures. Wide bias ranges increase the likelihood that similar identity dimensions will be added to repertoires or reinforced within repertoires. Widespread agreement regarding the level of shared identity will collapse as shifting bias triggers the insertion and reinforcement of new identities. In the baseline model, identity dimension bias has a range of –3 to +3 and a volatility of .01 (i.e., a 1 percent chance of a change in any given iteration). In the results presented below, we vary the range of the bias from narrow (–1, +1) to wide (–5, +5). Volatility is kept constant throughout.

Figure 5 indicates that increasing the range of identity dimension bias in the environment slightly increases the convergence on a common identity. The spatial plot for the narrow-identity dimension bias condition on the left reveals a great deal of spatial clustering that does not appear in the high-bias landscape. The narrow range allows locally homogeneous groups to emerge (e.g., the deep blue cluster in the upper left and the deep red cluster in the lower left). In contrast, the wide-bias case shows almost no clustering at all because embryonic clusters shatter as very positive biases flips to become very negative biases (and vice versa). The time plot averaged over thirty runs at the bottom of the figure indicates, as expected, that low bias leads to very little oscillation within the landscape over time. In contrast, the high-bias case oscillates above and below the baseline results.

Shifting to the averages over thirty runs of the simulation, the mean shared identity is 36.2 percent with low bias, 38.2 percent with the baseline range, and 38.6 percent for the high-bias case. *T* tests reveal that only the difference between low- and high-bias conditions is statistically significant at better than the .05 level. An *F* test of variances reveals a statistically significant difference between the low-bias and medium-bias condition at the .001 level. Surprisingly, the variance for the high bias and medium bias is almost identical. In sum, a shift in identity dimension bias from a narrow- to high-bias range increases the mean shared identity and increases the variance of shared identity, as predicted by hypothesis 2. A narrow-bias range creates fewer opportunities for a convergence in identities by injecting new identity dimensions into the mix. The findings imply that societies with some openness are more likely to develop a shared sense of identity than closed societies.

14. In the short run, bias is likely to shift more frequently in an open society. However, closed societies may experience rare but extreme shifts in bias due to coups or regime changes. For example, while bias ranges may be narrow and volatility may be limited in both absolute monarchies and communist dictatorships, the shift from the first to the second will trigger a dramatic shift in biases.

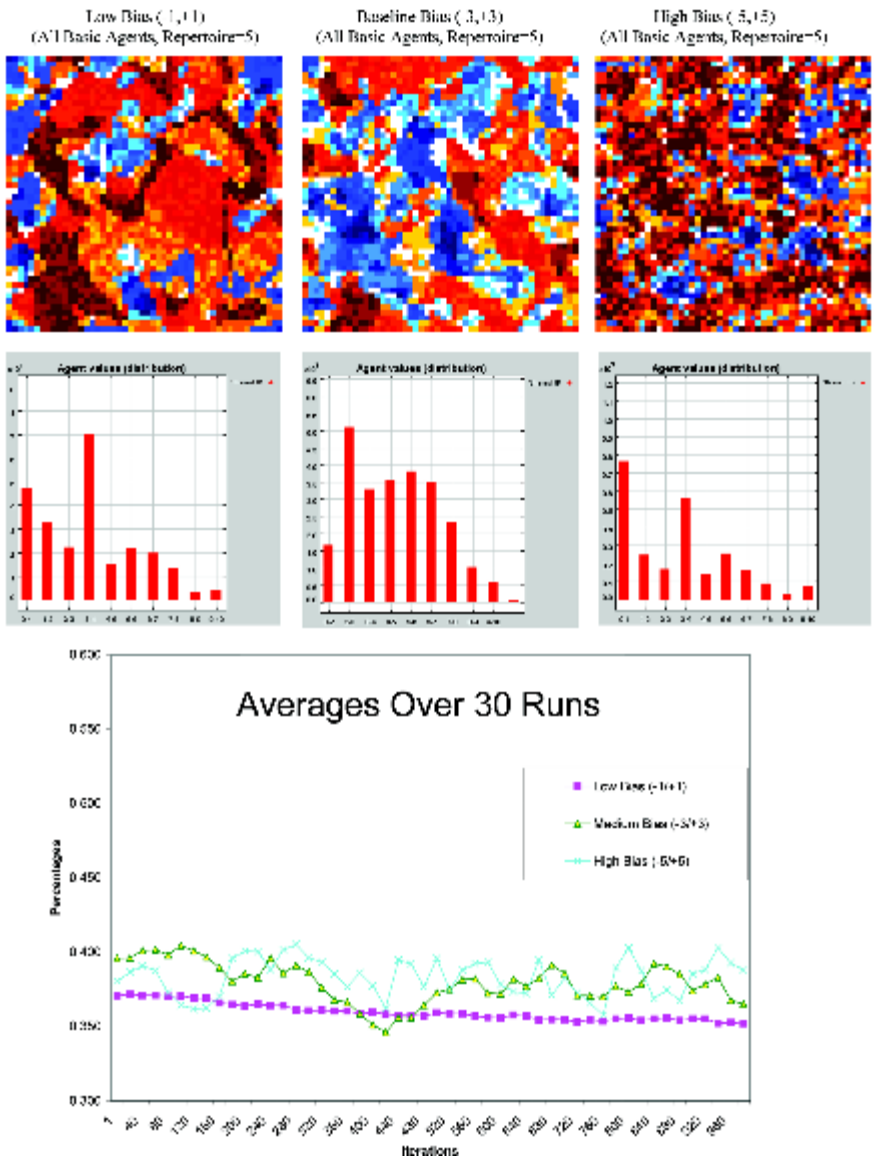

Figure 5: Shared Identity with Various Bias Levels

HYPOTHESES 3 AND 4: PROBING INFLUENCE OF LEADERSHIP

Hypothesis 3 predicts that adding leaders with more power than basic agents to the landscape will cause a rise in average shared identity. In contrast, hypothesis 4 predicts that leaders with the power to reach out beyond the immediate Moore 1 neighborhood

will lead to a rise in shared identity. We explored these questions by varying the composition of actors in the simulation. The Powerful Leadership analysis increases the percentage of leaders in the population from 0 percent (the baseline case) to 5 percent and gives each leader twice the influence of a basic agent. In the Broad-Reach Leadership analysis, we again increase the percentage of leaders to 5 percent. However, rather than increasing the power of the leaders, we increase their range by allowing them to contact up to eight actors beyond the Moore 1 neighborhood.

The results in Figure 6 reveal that power has a more important impact than reach in our simulation. The presence of powerful leaders reduces the average shared identity from the baseline simulation from 38.2 to 35.2 percent. This result is statistically significant at the .004 level. In contrast, broadening the reach of leaders has no discernable impact on the average shared identity. The average shared identity is virtually identical to the baseline (and it fails to achieve statistical significance). However, broadening the reach of leaders does significantly increase the variance of simulation runs compared to the baseline condition. The variance almost doubles from .006 to .011, and the difference is significant at the .04 level. Thus, reaching out beyond the neighborhood can trigger oscillations in the landscape as new identities penetrate neighborhoods faster than in the baseline case. However, it does not appear to affect the mean around which the population oscillates.

EXAMINING INTERACTIONS

The power of simulations over other techniques is most pronounced in cases of nonlinearity and interaction. Although the simulation can be used to probe literally hundreds of interactions, we explore one interaction for demonstration purposes. The previous analysis has demonstrated that repertoire size, global bias, and leadership power all influence the prevalence of shared identity. Although each had an interesting substantive and statistically significant impact, the results did not fundamentally transform the landscape. This raises an interesting question: could a combination of moderately powerful factors trigger a transformation? The results, which appear in Figure 7, employ the following parameters: small repertoire size (3), low bias (–1, +1), and leaders with twice the power of regular agents.

The time plot, averaged over thirty runs at the bottom of the figure, reveals that the combination of factors does not alter the average shared identity of the landscape. Although the three-way interaction produces less oscillation than the baseline trend, the means are very similar (38.2 vs. 37.2). However, the mean only captures one dimension of the data. The landscape plot and the histogram demonstrate that the interaction does have a profound impact on the spatial distribution of shared identity. Specifically, the interaction of repertoire size, bias, and leadership leads to an extremely polarized population in comparison with the baseline simulation. In the landscape plot, the population of agents is sharply divided into densely packed pockets of red and blue. The boundaries are sharp, and an examination of them as the simulation runs indicates that they are quite stable once established. A quick glance back at Figures 3 through 6 indicates that the polarization never emerges in the prior runs. The histogram also illustrates extreme polarization as the low shared-identity bin (0-10

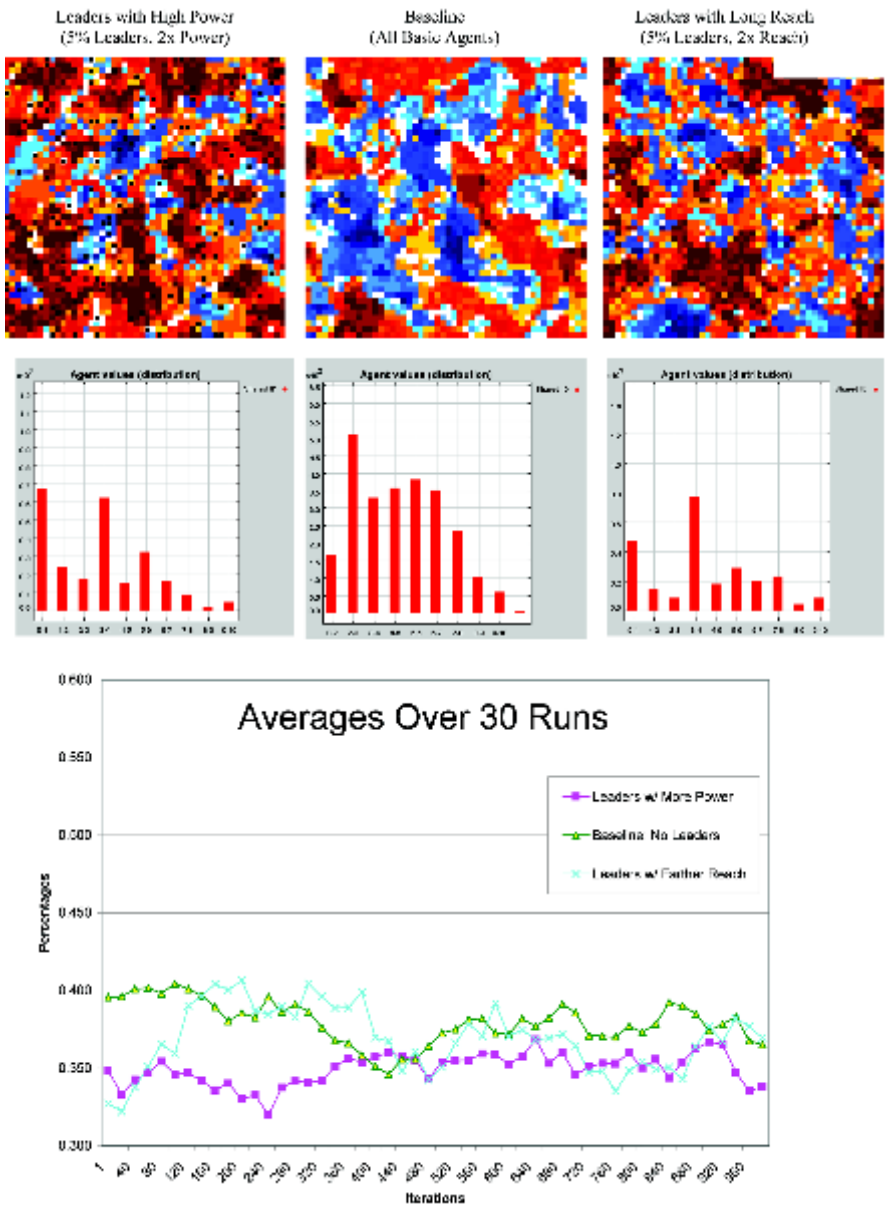

Figure 6: Shared Identity with Various Leader Types

percent or 0-1 in the figure) and the high shared-identity bin (90-100 percent or 9-10 in the figure) capture the lion's share of the agents. The mean level of polarization rises from .02 in the baseline simulation to .15 in the interaction simulation; this difference

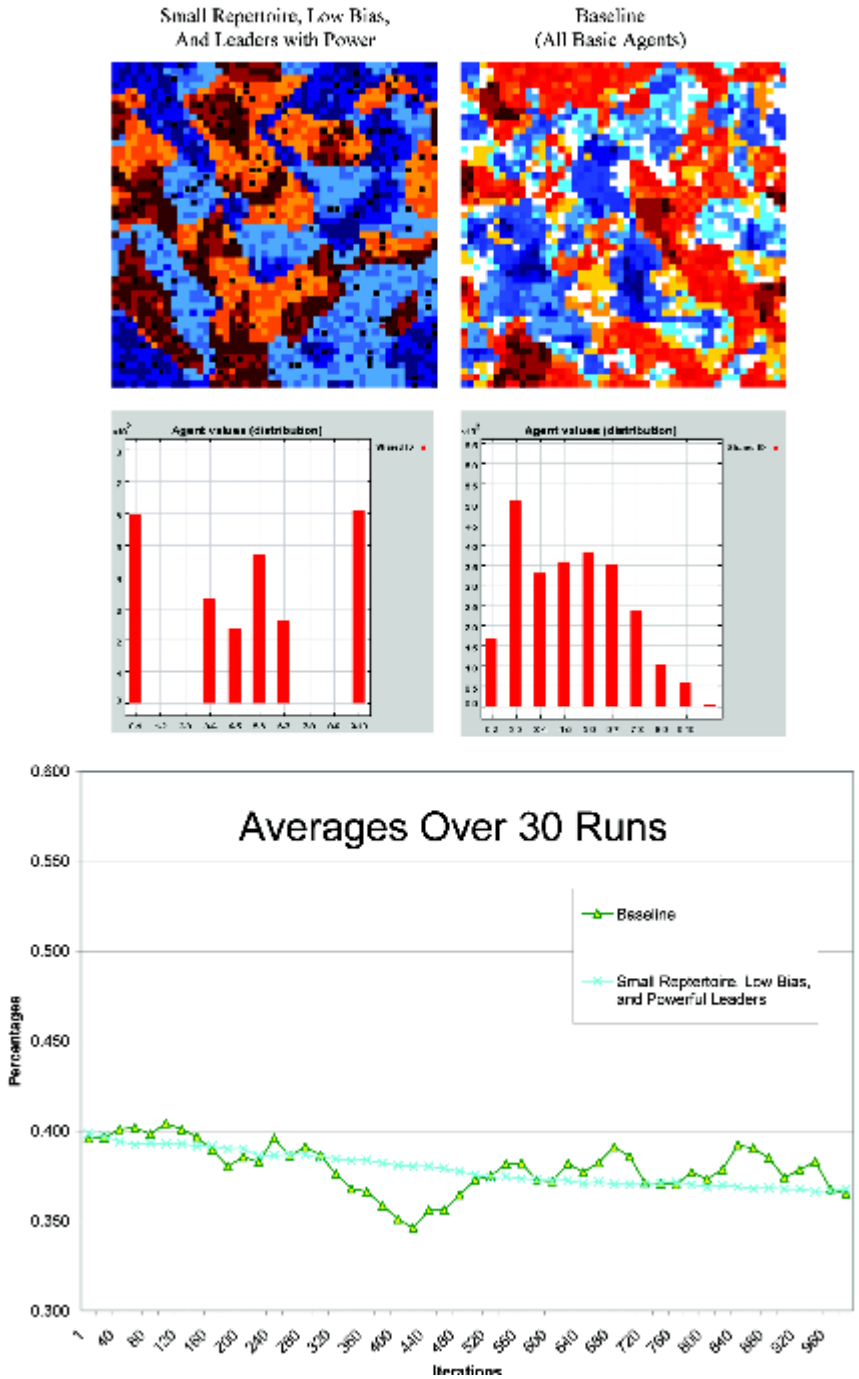

Figure 7: Shared Identity with Interactions

is statistically significant at better than the .001 level. The interaction nicely illustrates one of the great strengths of agent-based models: relatively simple micro-level processes can produce the emergence of very surprising macro-level patterns.[15]

CONCLUSIONS

This article represents an initial attempt to model how conceptions of the collective Self and Other spread throughout a population. In much of the constructivist literature in international relations, scholars begin with the observation that society X possesses belief Y or attribute Z (e.g., the Japanese are anti-militarists or Americans are individualists). We hope to push the investigation back one step by exploring the *process* through which identities such as anti-militarist and individualist are spread throughout a population and maintained across time. By linking our simulation parameters to laboratory experiments with human participants, we hope to further our understanding of the process through which individuals continually construct a sense of Self and Other.

Four broad conclusions emerge from the analysis. First, there is a curvilinear relationship between the complexity of our views of ourselves and others and average shared identity. The highest levels of shared identity emerged with moderately sized identity repertories. Moreover, higher complexity was also associated with greater stability in views across time. Second, shared identity was higher in volatile environments than in stable environments. Third, the presence of powerful political entrepreneurs decreased the likelihood of the emergence of a shared identity. In contrast, entrepreneurs with a broad reach across the population only appeared to affect the rate of oscillation of the population. Finally, the interaction of complexity, stability, and leadership produced extremely polarized societies.

The processes proposed and tested in our model further our understanding of international politics in a number of ways. First, we know that societies vary greatly in their level of knowledge about other states in the system. The model predicts that Americans, who possess limited knowledge of the world compared to their counterparts in industrialized democracies in Europe and Asia, will have more volatile views of other countries. Second, the model can help us understand why specific beliefs and identities can become hegemonic. While this is sometimes relatively straightforward (e.g., Japan's military defeat undermined militarism after World War II), in other cases, the causal links are much less obvious. The simulations point to three areas of interest for developers of historical case studies: the complexity of images, the volatility of the environment, and the attributes of leaders. Third, the results strikingly illustrate the importance of *perception*. It is important to note that there is no "true" value for either the self or the other identity in our model: all the differences and patterns that emerge

15. We measure polarization using the following formula: minimum(A, B) • 2 • (A + B), where A is the percentage of agents in the lowest decile and B is the percentage of agents in the highest decile. Polarization can be both an input into the simulation model via parameter settings and an output of the model via the measurement of extreme views in the population. In the simulations reported here, we restrict our focus to polarization as output.

thus arise not from different degrees of playing up (or obscuring) real differences but rather from the spread of perceptions that are completely independent from any real underlying differences. Fourth, the model provides the micro-foundations for "tipping" models of identity formation (e.g., Laitin 1998). Interestingly, the interaction analysis illustrates how tipping can occur in both local clusters and across the entire landscape. Finally, the simulation demonstrates the feasibility of linking the oft-neglected individual level of analysis to broader structural models. This is particularly important for constructivism because social construction must begin with human interaction.

REFERENCES

Abdelal, Rawi, Yoshiko M. Herrera, Alastair Iain Johnston, and Rose McDermott. 2004. Identity as a variable. Paper presented at the Harvard Identity Project Conference, December, Boston.

Anderson, Benedict R. 1983. *Imagined communities: Reflections on the origin and spread of nationalism*. London: Verso.

Axelrod, Robert. 1997. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton, NJ: Princeton University Press.

Axtell, Robert. 2000. Why agents? On the varied motivations for agent computing in the social sciences. Working Paper No. 17, Brookings Institution, Center on Social and Economic Dynamics, Washington, D.C.

Bates, Robert H. 1974. Ethnic competition and modernization in contemporary Africa. *Comparative Political Studies* 6 (4): 457-84.

Baumeister, Roy F. 1998. The self. In *The handbook of social psychology*, 4th ed., edited by Daniel T. Gilbert, Susan T. Fiske, and Gardner Lindzey, 680-740. Boston: McGraw-Hill.

Berger, Thomas U. 1998. *Cultures of anti-militarism: National security in Germany and Japan*. Baltimore: Johns Hopkins University Press.

Bourdieu, Pierre. 1998. *Practical reason: On the theory of action*. Stanford, CA: Sanford University Press.

Brass, Paul R. 1974. *Language, religion and politics in North India*. Cambridge, UK: Cambridge University Press.

———. 1997. *Theft of an idol*. Princeton, NJ: Princeton University Press.

Brewer, Marilynn B., and Rupert J. Brown. 1998. Intergroup relations. In *The handbook of social psychology*, 4th ed., edited by Daniel T. Gilbert, Susan T. Fiske, and Gardner Lindzey, 554-94. Boston: McGraw-Hill.

Cederman, Lars-Erik. 1997. *Emergent actors in world politics: How states and nations develop and dissolve*. Princeton, NJ: Princeton University Press.

Chandra, Kanchan. 2001. Introduction: Constructivist findings and their non-incorporation. *APSA-CP: Newsletter of the Organized Section in Comparative Politics of the American Political Science Association* 12 (Winter): 7-11.

Checkel, Jeffrey T. 1998. The constructivist turn in international relations theory. *World Politics* 50 (2): 324-49.

Collins, Randall. 1981. Micro-translation as a theory building strategy. In *Advances in social theory and methodology: Toward an integration of micro- and macro-sociology*, edited by K. Knorr-Cetina and A. V. Cicourel. London: Routledge Kegan Paul.

Converse, Jean, and Stanley Presser. 1986. *Survey questions: Handcrafting the standardized questionnaire*. Beverly Hills, CA: Sage.

Durkheim, Emile. [1897] 1951. *Suicide: A study in sociology*. Translated by John A. Spaulding and George Simpson. Edited by George Simpson. Reprint, Glencoe, IL: Free Press.

Finnemore, Martha. 1996. *National interest in international society*. Ithaca, NY: Cornell University Press.

Gellner, Ernest. 1983. *Nations and nationalism*. Ithaca, NY: Cornell University Press.

Giddens, Anthony. 1984. *The constitution of society: Outline of the theory of structuration*. Cambridge, UK: Polity.

Gigerenzer, Gerd, Peter Todd, and the ABC Research Group. 1999. *Simple heuristics that make us smart*. Oxford, UK: Oxford University Press.

Gulick, Edward Vose. 1955. *Europe's classical balance of power.* New York: Norton.

Hopf, Ted. 1998. The promise of constructivism in international relations theory. *International Security* 23:171-200.

———. 2002. *Social construction of international politics: Identities and foreign policies, Moscow 1955 and 1999*. Ithaca, NY: Cornell University Press.

Johnson, Paul E. 1999. Simulation modeling in political science. *American Behavioral Scientists* 42 (10): 1509-30.

Kahl, Colin H. 1998-1999. Constructing a separate peace: Constructivism, collective liberal identity, and democratic peace. *Security Studies* 8 (2): 94-119.

Kant, Immanuel. [1795] 1971. Perpetual peace: A philosophical essay. In *Kant: Political writings*, 2nd ed. Edited by Hans Reiss. Reprint, Cambridge, UK: Cambridge University Press.

Katzenstein, Peter J., ed. 1996. *The culture of national security: Norms and identity in world politics*. New York: Columbia University Press.

Kier, Elizabeth. 1997. *Imagining war: French and British military doctrine between the wars*. Princeton, NJ: Princeton University Press.

Kowert, Paul, and Jeffrey Legro. 1996. Norms, identity, and their limits: A theoretical reprise. In *The culture of national security: Norms and identity in world politics*, edited by Peter J. Katzenstein. New York: Columbia University Press.

Laitin, David D. 1998. *Identity in formation: The Russian-speaking populations in the near abroad*. Ithaca, NY: Cornell University Press.

Laitin, David D., and Daniel Posner. 2001. The implications of constructivism for constructing ethnic fractionalization indices. *APSA-CP: Newsletter of the Organized Section in Comparative Politics of the American Political Science Association* 12 (Winter): 13-17.

Lustick, Ian S. 2000. Agent-based modeling of collective identity: Testing constructivist theory. *Journal of Artificial Societies and Social Simulations* 3 (1). http://jasss.soc.surrey.ac.uk/3/1/1.html.

Lustick, Ian S., and Dan Miodownik. 2000. Deliberative democracy and public discourse: The agent based argument repertoire model. *Complexity* 5 (4): 13-30.

———. 2002. The institutionalization of identity: Micro adaption, macro effects, and collective consequences. *Studies in Comparative International Development* 37 (2): 24-53.

Lustick, Ian S., Dan Miodownik, and Roy J. Eidelson. 2004. Secessionism in multicultural states: Does sharing power prevent or encourage it? *American Political Science Review* 98 (2): 209-30.

Macy, Michael, and John Skvoretz. 1998. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review* 63 (5): 638-61.

Macy, Michael W., and Robert Willer. 2002. From factors to actors: Computational sociology and agent-based modeling. *American Review of Sociology* 28:143-66.

Onuf, Nicholas G. 1989. *The world of our making: Rules and rule in social theory and international relations*. Columbia: University of South Carolina Press.

Organski, A. F. K. 1968. *World politics*. 2nd ed. New York: Knopf.

Posner, Daniel N. 2003. The colonial origins of ethnic cleavages: The case of linguistic divisions in Zambia. *Comparative Politics* 35 (2): 127-46.

Risse, Thomas. 2000. Let's argue! Communicative action in world politics. *International Organization* 54 (1): 1-39.

Rousseau, David L. 2002. Motivations for choice: The salience of relative gains in international relations. *Journal of Conflict Resolution* 46 (3): 394-426.

———. Forthcoming. Identifying threats and threatening identities: The social construction of realism and liberalism. Stanford, CA: Stanford University Press.

Sewell, William. 1992. A theory of structure: Duality, agency and transformation. *American Journal of Sociology* 98:1-29.

Sil, Rudra. 2000. The foundations of eclecticism: The epistemological status of agency, culture, and structure in social theory. *Journal of Theoretical Politics* 12 (3): 353-87.

Smith, Steve. 2000. Wendt's world. *Review of International Studies* 26 (1): 151-63.

Sudman, Seymour, Norman M. Bradburn, and Norbert Schwarz. 1996. *Thinking about answers: The application of cognitive processes to survey methodology.* San Francisco: Jossey-Bass.

Touraine, Alain. 1985. Introduction to the study of social movements. *Social Research* 52 (4): 749-87.

van der Veen, A. Maurits, Ian S. Lustick, and Dan Miodownik. 2001. Studying performance and learning with ABIR: The effects of knowledge, mobilizing agents, and predictability. *Social Science Computer Review* 19 (3): 263-80.

Verdery, Katherine. 1991. *National ideology under socialism: Identity and cultural politics in Ceau[S1]escu's Romania.* Berkeley: University of California Press.

Waltz, Kenneth N. 1979. *Theory of international politics.* Reading, MA: Addison-Wesley.

Weber, Max. [1914] 1978. *Economy and society.* Reprint, Berkeley: University of California Press.

Wendt, Alexander. 1999. *Social theory of international politics.* Cambridge, UK: Cambridge University Press.

Zaller, John R. 1992. *The nature and origins of mass opinion.* Cambridge, UK: Cambridge University Press.