A HYBRID MODEL OF DECISION-MAKING IN CLOSED POLITICAL REGIMES

R. Bhavnani, Department of Political Science, University of Illinois Urbana-Champaign*
D. Backer, Department of Political Science, University of Michigan, Ann Arbor
R. Riolo, Center for the Study of Complex Systems, University of Michigan, Ann Arbor

ABSTRACT

In this paper, we present the first cut at an exploratory agent-based modeling project where the aim is to develop a tool for intelligence analysts to employ in studying decision-making processes in closed political regimes such as Iraq, North Korea and Syria. Our hybrid of the landscape metaphor (Kollman, Miller, & Page 1992, 1998; Axelrod & Bennett 1993) and the rule-based system approach (Holland et al 1986) captures the trade-offs leaders of closed regimes face in attempting to balance power, policy preferences and regime support—components of a utility fitness function—as well as differences in leaders' types that can result in departures from a strict notion of rationality and utility maximization. Based on the results of two simple experiments concerning succession, we observe that peaks in the utility landscape can arise in surprising places away from the ideal positions of both the leader and other elites—as the result of compromises they are willing to make. Moreover, the interplay between the leader's basic tendency to maximize utility and his risk-sensitive heuristic rules leads to a high level of instability in the regime.

* *Corresponding author address*: Ravi Bhavnani, Department of Political Science, University of Illinois, 361 Lincoln Hall, 702 S. Wright Street, Urbana, IL 61801–3696; email: rvib@uiuc.edu

INTRODUCTION

In this paper, we report upon our preliminary efforts to apply agent-based modeling methodologies to the study of decision-making processes in closed political regimes such as Iraq, North Korea and Syria.[1] Such an approach is plausible because these processes exhibit certain classic characteristics of a complex system. Chief among these is the difficulty of mapping outcomes to precipitating factors: leaders with ostensibly similar characteristics and interests act in ways that appear to be highly conditional on subtle variations in circumstances. At the same time, given the heterogeneity of actors and the political settings in which they operate, such a system becomes impractical to model statistically (with multivariate regression), mathematically (with differential equations) or using a game-theoretical approach. Moreover, due to the secretive and insular nature of these regimes, gathering detailed, reliable information on their members, their interactions and institutional structures presents a challenge for conventional analytical techniques. Thus, we opt instead to conduct computer simulations using an exploratory agent-based model (Bankes 1994; Casti 1997), using a hybrid of the landscape metaphor (Kollman Miller Page 1992, 1998; Axelrod Bennett 1993) and the rule-based systems approach (Holland et al. 1986).

The model we develop explores the trade-offs faced by leaders in altering the composition and policy stances of the regime to shore up support from other elites, while at the same time keeping sight of their own preferences and power. The dimensions of the landscape correspond to an issue space, modulated by the distribution of power and interests among the agents. The basic goal of agents is to maximize their utility, as a function of the power they enjoy as well as the proximity of their preferences to those of the coalition to which they belong; leaders also care about support for the regime. Our objective is to explore how different levels of rationality or strategic play on the part of leaders affect the dynamics of regime survival.

Section 2.0 outlines an archetype of closed regimes, focusing on characteristics that are central to the structure and mechanics of decision-making processes. Section 3.0 discusses our efforts to develop a model of these processes, as well as our expectations regarding its evaluation and application. Section 4.0 offers a description of the model. Sections 5.0 and 6.0 present and discuss some preliminary results from our first exploratory runs of the model, focusing on scenarios of leadership succession.

AN ARCHETYPE OF CLOSED REGIMES

Closed regimes share at least four characteristics that have a direct bearing on decision making processes—concentrated power, an authoritarian command structure, high threat perception, and extreme secrecy. Individual regimes will display these characteristics to varying degrees, and may be distinguished by other characteristics, but we consider these four to be fundamental to our modeling exercise.

[1] This paper documents work in progress on a project entitled "Simulating Closed Regimes Using Agent-Based Models," which is funded by the Department of Defense (DOD) and administered under the auspices of the Center for the Study of Complex Systems (CSCS) at the University of Michigan, Ann Arbor. The views expressed in this paper are those of the authors alone, and do not reflect the opinions of the DOD or CSCS.

Power is concentrated in a closed regime, such that the regime is ordinarily comprised of a small number of key decision makers, e.g., a paramount leader or the members of a ruling council or military junta. Whether one, several or a group of individuals occupy positions of authority, they typically monopolize decisions on all matters—political, military, economic, social, cultural, etc. This circumstance magnifies the importance of these key individuals, each of whom is likely to have a significant area of jurisdiction and considerable resources at his or her disposal. Consequently, a suitable model of decision-making in this context could logically contain relatively few agents. Yet, they would almost inevitably be specified in a manner that is more complex than a traditional agent-based model, both in terms of the number and variety of traits as well as the actions they can take.

Related to the concentration of power among a small coterie of elites is the lack of institutional barriers to decision-making. Often, there are no consequential distinctions between executive, legislative and judicial authority: they are effectively collapsed into a single command-and-control institutional structure. This has several implications. Most notably, one or a few powerful agents are in a position to quickly implement their preferred policy choices and to execute other significant decisions. In fact, there may be no formal decision-making processes, or those that nominally exist are trumped by informal norms and arrangements, if not the edicts of key actors. As a result, closed regimes are capable of taking swift action in many realms, not least political and military affairs, even though they may rank low on conventional measures of state capacity.

The inclination to consolidate authority is also inter-related with a heightened sense of insecurity among members of closed regimes. Such anxiety and uncertainty has several relevant by-products. To begin with, cohesion among the governing elites will tend to vary, thus invalidating assumptions of unitary action and undermining the prospect of stable long-term cooperative arrangements. In fact, there is a likelihood of punitive actions within the ranks, in the form of abrupt changes to the governing elites. More generally, closed regimes are prone to aggression, whether reactive or proactive, against challengers to the authority of the leader and other key figures. Consistent with this orientation, these regimes often display a well-ingrained political, religious, or ethno-cultural identity that emphasizes confrontation with and endorses the exclusion of non-adherents and other perceived outsiders. With this mindset, the members are prone to misperceptions, including overstating the extent of threats to authority. Responses, in turn, are highly subjective and may even appear irrational—e.g., retaliatory measures that immediately backfire.

In addition, the leaders of closed regimes typically seek to insulate themselves from potential risks by restricting the flow of information. As we observed earlier, the resulting secrecy and concealment complicates analysis. Unconventional means are often necessary to compile relevant details about leaders, events and institutional arrangements, yet still one's understanding of these aspects will inevitably be incomplete or uncertain, if not speculative. From a modeling perspective, the information constraints also imply that agents—even including those in positions of authority within the regime—are rarely equipped with the comprehensive information required to make fully rational decisions all the time. Instead, with at least some regularity they will have to (1) rely on proxies or guesses; (2) limit the number of options they consider; and/or (3) discount the future. One would expect, therefore, that they are subject to misperceptions and other miscalculations that reflect the noisiness of the information they have at their disposal, leading to mistakes in judgment.

MODEL DEVELOPMENT, USE AND EVALUATION

Model Development

The process by which we approach our task merits further discussion, not least because of the unusual context. From a substantive perspective, we are dealing with an information-poor setting involving a small number of agents who are generally in a position to assert their will and are also prone to aggressive and seemingly irrational acts. Yet, the practical objective is to build a tool that will help intelligence analysts to fill in some of these gaps, to build and challenge their intuitions, and to attempt to make sense of various complex behaviors and processes.

At the outset, we concluded that far too many factors affect decision-making in closed regimes for anyone to generate realistic (and defensible) predictions of the type "a regime will go to war with 100% certainty", much less "it will invade its neighbor on Tuesday." Hence, we do not engage in a data-mining exercise to find a model that best fits a particular case or a set of cases, nor do we endeavor to construct a consolidative model. Instead, we develop an exploratory model to capture the key characteristics and behaviors of members of closed regimes, i.e., how elites respond to changes in the political, economic, and social environment, as well as each other's actions. Our objective is relatively modest: to afford additional insight into dynamics that *may* lead to a particular result or engender a certain phenomenon, with an emphasis on the *process* rather than the outcome itself.

We build the model from the ground up, relying heavily upon input from analysts as to what is important in terms of both foundational characteristics as well as problems and contexts that deserve in-depth study. Here, we wish to underscore the distinction between the direct use of such data in analysis, versus translating it into a more abstract form that shapes one's conceptualization and implementation of a model. Our approach follows the latter route, i.e., harnessing the analysts' distinctive knowledge and expert interpretations in order to map between empirical data and our stylized representation of closed regimes, and vice versa.

We have addressed this task in a number of ways.

First, we gathered pertinent details on each of the standard components of an agent-based model. At the broadest level, we asked a selection of intelligence analysts to identify (1) which factors are salient—and in what respect—to their evaluation of closed regimes; (2) which of these factors could reasonably be combined into general properties; (3) which factors might be considered exogenous, at least under particular circumstances or during certain time frames; and (4) which factors could be omitted from the model without much loss in precision. With respect to the *agents*, we requested information on the following:

- the number of key decision-makers,

- their positions/roles/status,
- generic traits (e.g., beliefs, motivation, background, interpersonal decision style, degree of paranoia, proclivity to use violence) and unique characteristics (e.g., agent has life-threatening illness),
- the constraints they face, particularly vis-à-vis obtaining information from other agents and the non-agent environment (e.g., specific agents may be more clued into military issues),
- the heuristics they use to process information and make choices,
- the actions they can take (commands, messages, etc.), and
- their capacity to affect outcomes (e.g., control over the security apparatus, ability to influence economic policy).

In terms of the *connections among agents*, we focused on these factors:

- which agents interact with which others and how often (e.g., formal vs. informal regime structures, fixed vs. dynamic structure, military junta vs. religious tribunal),
- the nature of decision-making (e.g., degree of consensus, relative weight of electorate vs. selectorate),
- the information environment (e.g., sources and accuracy of information, degree of conflicting/reinforcing information, level of monitoring within and beyond member of inner circle), and
- the level of support along the chain of command (e.g., loyalty, compliance, defiance).

Notable *environmental factors* included stable aspects of the military, economic, political, and social realms, as well as shocks, triggers and other short-run variables.

Second, with an eye towards capturing analysts' understandings of closed regimes—and as best possible, formalizing where such structure does not yet exist—we requested information on the sorts of mental models analysts currently employ in their work. In particular, we asked them to provide us with descriptions of the following:

- *Informal models*—perspectives that employ intuitive, issue-dependent reasoning and therefore lack an explicit, consistent structure.
- *Semi-formal models*—approaches that organize information in terms of at least some specific salient factors; analytical tools such as lists of actor

attributes, classification schemes, and discrete if-then rules or other heuristics attributed to the actors.

- *Formal models*—explicit, well-defined paradigms, comprised of a standardized framework, established typologies and classification schemes, proven relationships among factors, and/or concrete theories, that guide information processing and analytic reasoning.

Most of what the analysts discussed fell into the first two categories, especially the second. As a result, there were not any well-defined models that we could implement directly, but rather a large number of discrete fragments, some better elaborated than others.

Third, we then sought to consolidate this knowledge, bridge the various regional and substantive perspectives of the analysts, and usher things in the direction of developing a generic model of decision-making processes in closed regimes. With this in mind, we asked them to specify and rank model variants and mechanisms, brainstorm about the types of problem areas they deal with most frequently, the regime environments in which they play out, and the mechanisms that are potentially at work.[2] We then looked for any indications of overlap, if not consensus, as to what is critical.

While the basic progression has been from the particular to the general, several caveats are in order. To begin with, our working assumption has always been that this area of intelligence analysis exhibits too much variety—not to mention intricacy and uncertainty—for a single, comprehensive model to be a practical objective. Instead, a more plausible outcome would be a modular set of models, each pertaining to a specific context (e.g., succession, repression, invasion) or class of phenomena (e.g., coalition politics). Certain underlying features and parameters could conceivably be shared by more than one model. These points of commonality, however, might be limited by the necessity of employing a number of different approaches to agent-based modeling, tailored to the relevant context.

At the same time, what we are developing is intended as a tool to aid analysts and should logically be calibrated and attuned to their perspective. An analyst, for example, might have a series of related hypothetical scenarios that she wishes to study—e.g., how the same regime responds to different economic, political, military and cultural triggers. That is, the basic components of the model are essentially the same, but the analytical lens is directed at separate aspects of the environment. Or maybe the analyst wishes to focus on a specific outcome—e.g., regime collapse—and thus needs to vary components such as its composition and structure. One of our ultimate objectives, therefore, is to ensure that the end user has the freedom to vary agent characteristics and decision rules, organizational structures, environmental conditions and dynamics, and perhaps even adaptive algorithm parameters, as well as to repeat actions under the same or different conditions and to test strategies that push the system to its extremes.

---

[2] Since regular meetings with DOD analysts were not feasible, we will also rely heavily on the domain experts hired for this project. Two of the domain experts are Middle East specialists—Ellen Lust-Okar (Political Science, Yale University) and Moshe Maoz (Political Science, Hebrew University of Jerusalem).

The importance of flexibility is not limited to user specification of starting conditions or interventions during a run. Ideally, analysts should have the capacity—in terms of both access to the source code and the knowledge required to make the appropriate changes—to easily modify virtually all the model mechanisms. The utility of this is multifold. The models may be too simple or abstract to be of value for analysts, thus necessitating additional elaboration and sophistication. Moreover, developers cannot possibly include all the variants that different users will want to explore. By the same token, one cannot anticipate what new questions will be raised as users explore the dynamics produced by a particular version of the model. Constructing an intuitive framework and then providing the user with an "open box", should ensure that at least some analysts will be able to devise their own distinctive and individually authentic versions of the model, given only modest programming skills.[3]

In sum, our goal in this project is to build a family of models, each of which attuned to analysts' understanding of individuals and their behaviors in some closed regime setting (or set of related settings). Thus, we are trying to balance the need for (somewhat) general models, which can be applied to a range of related situations, with the need for a model that is not so abstract as to become divorced from analysts' views of the world.

Model Use and Evaluation

One important advantage of the agent-based modeling approach is that the model can be validated at multiple levels—at the level of individual agent behaviors and properties, as well as in terms of emergent structure, relationships and dynamics at the aggregate or system level. Given the complexity of the processes we are trying to model, however, validation of models in this area will be especially difficult.

Bearing this in mind, an analyst would map her data and knowledge about a regime into input parameters that describe the agent traits and decision rules, agent organization, environment and adaptive processes. If, for instance, the regime was comprised of elites from the dominant ethnic group who shared military backgrounds and beliefs about an ethnically pure state, similar interpersonal and decision styles, and a proclivity for violence, then the agent traits would be set to mirror these real-world attributes. One of the challenges of creating models that will be useful to analysts, then, is to make it easy for them to see the relationship between the relatively limited menu of model components and mechanisms and the diverse set of factors that are believed to play important roles in the decision-making processes of closed regimes.

For that reason, our initial approach to evaluating the models we produce is to establish their face validity. In essence, the model should "behave" as experts would expect under a variety of simple conditions. We also intend to assess the validity of the model using standard, model-independent techniques, such as exploring the model's sensitivity to variations in parameter values. For example, we would want to establish that changes to a particular parameter yield the anticipated responses among agents and aggregate behavior. Not only

[3] This need for flexibility is a key reason why we are using an Object-Oriented Programming (OOP) approach and building our model using an existing package of classes designed for constructing adaptable, powerful agent-based models—e.g., RePast.

should an association exist, but the sign, direction and/or magnitude of this relationship should be consistent with analysts' expectations. This is not to suggest that all of the results generated by the model must match analysts' repository of prior knowledge. Rather, there are certain critical "facts"(e.g., A generally leads to B, Y never follows X, Z always increases over time) that can be used to discern whether the model is consistent with reality. Many other aspects of the model may not be subject to definitive confirmation, yet still must be within the bounds of plausibility—in other words, analysts are not certain that something operates in a particular way, but it would make sense if it did. One must also allow for the possibility of counter-intuitive results, despite nearly everything else being consistent with what analysts understand to be the case. Thus, our criteria for evaluating the model include whether it (a) corresponds in certain fundamental respects to the current understandings of intelligence analysts, (b) helps them to fill in some of their information gaps with plausible details and explanations, (c) builds and challenges their intuitions, thereby aiding the process of understanding various complex behaviors and processes.

MODEL DESCRIPTION

In this section, we describe our first attempt at developing a model of decision-making processes in closed regime. To reiterate what we discussed earlier, our efforts are informed and inspired by the needs of intelligence analysts for a simple, (somewhat) general, yet still useful exploratory tool. At this stage, we opted to employ a hybrid landscape/rule-based approach for a variety of reasons.

To start with, the landscape metaphor nicely captures analysts' basic intuitions about closed regimes. In particular, the heterogeneity of key actors and policy options creates a complex set of alternative states. The task of the leader and other influential elites is to navigate this difficult terrain to their advantage. With the landscape approach, we can depict—both conceptually and visually—the tension between the leader's desire for power, his wish to set policy to mirror his own preferences, and his concern for support (so as not to be overthrown). Similarly, we can represent other agents' trade-offs between acquiring power and achieving policy goals.

By also including a set of rule-based heuristics, we are able to attribute simple characteristics to the leader and explore the resulting variation in behavior. While we could capture the diversity in leadership styles by using more complicated utility functions, analysts deemed the approach of introducing heuristics to capture exceptions from "rational" (i.e., utility-maximizing) behavior to be more intuitive because individuals in these positions often operate with first-order norms or principles (e.g., survival).

Agents

There are three sets of agents in the model: the leader, individuals who hold formal positions in the regime hierarchy (Tiers 1 and 2), and other elites (Tier 3). The tiers are generally a short hand for the amount of influence an agent enjoys in the regime, including his/her status and capabilities.[4] At this stage in the model development, the tiers do not entail anything as far as relationships among agents, nor are there any limitations in terms of movement both within and across these different levels. Such considerations are reserved for future extensions.

The agents have several traits. First are preferences [0,10] and salience weights [0,1] on each of a set of M policy issues. (For the experiments reported in this paper, M=2.) For the sake of simplicity at this early stage of model development agents' preferences and saliencies are specified by the user at the start of a run to reflect the "state of the regime" the user wishes to explore.[5] Second is the power, if any, an agent enjoys from a formal position in the regime hierarchy (r_power). Third is an agent's "independent" or extra-regime power (i_power). All power is issue specific, to allow for the prospect that an agent might be more capable of influencing particular policy decisions, e.g., military generals would be strong in their own domain but weaker in others. Unlike preferences, power does not remain constant, but rather can vary as the result of the leader's actions as well as exogenous shocks. Initial values for each power value (r_power and i_power) on each of the M dimensions are drawn from [0,1].

Agents' Utility Functions

The basic goal of every agent is to maximize his/her utility. The leader's utility function is comprised of three components:

- D, the distance between his/her preferences and the policies of the regime (weighted by the salience of the issue)
- P, his total power (both regime-based and independent)
- S, the total support for the regime (the aggregation of the salience- and power-weighted policy distances of all agents, including the leader)[6]

---

4 That being said, Tier 3 agents can enjoy more power than those who hold formal positions in the regime hierarchy.

5 Our current implementation of this model allows the user to drag agents to new preference positions and to use RePast probes to change other agent characteristics during the course of a run, e.g., to explore how particular changes would alter the utility or support landscapes. Future versions of the model may introduce endogenous variation in issue preferences and weights.

6 In detail, total support for a regime policy position R is the sum of the support from all agents, including the leader. Support is calculated independently over each policy dimension, so here we just assume M=1 dimensions. The support s for an agent A with policy preference I, salience W, and total power (i_power + r_power) P is s = diff(I,R) * I * W * P where diff (I,R) is a function of the difference between the agents preferred position I , and the regime policy, R (on the dimension in question). That is, diff(I,R) indicates how much the agent "likes" the policy (independent of the agent's salience and power). For R=I, diff(I,R) is 1.0, providing for maximum support, and for R maximally distant from I, diff(I,R) = 0, providing no support. We use a diff(I,R) function that has two parameters that allow adjustment from a simple linear function to sigmoid-shaped functions with various thresholds and degrees of steepness. In particular, for d = |R - I|, $\text{diff}(d) = (1 + h^s) * (1 - d)^s / ((1-d)^s + h^s)$. For the experiments reported here, h = 2.0 and s = 3.0, which gives a steep drop, with diff(2.0) ~ 0.5.

Each of these components is weighted [0,1] as part of a linear utility function for the leader L:

$$U\_L\ (D,P,S) = \alpha_D * D + \alpha_P * P + \alpha_S * S$$

In all the experiments reported here, the alpha coefficients are all 0.33—e.g., the leader cares equally about policy distance, power and support for the regime. All other agents differ from the leader in that they do not take account of support for the regime. As such, their utility function is the weighted sum of just two components—policy distance and total power. In the experiments we conduct, they likewise weight the components equally.

The Landscapes

Our initial focus is on the trade-offs faced by the leader of a closed regime in making personnel and policy decisions. We help the model user understand these trade-offs by means of two M+1 dimensional landscapes, one for leader utility and the other for regime support, where M is the number of policy/issue dimensions. Again, for simplicity, in this paper we consider models where M=2.[7]

Thus, the x and y dimensions of the landscape correspond to two issues, each with a range of policy options [0,10]. While the specific nature of these two issues is irrelevant, it is logical that they would be at least somewhat unrelated, so that it is possible for an individual agent to have more or less extreme preferences on a given dimension, irrespective of what he feels about the other issue.[8] Where the agents array themselves in this policy space, as well as where the leader locates the regime's position, determines the support that the regime enjoys from elites and, in turn, affects the leader's utility. On the support landscape, therefore the z values at policy x, y indicate the support the regime would have if the regime position was set to x, y (all else unchanged). Similarly, the z values on the leader utility landscape indicate the utility the leader would have when regime policy is set to the corresponding x, y positions.

Note that the agents in the model generally do not know the shape of these landscapes. For this to be possible, they would need to know the preferences, saliencies, power and utility weights for all the other agents in the model. While such knowledge might be viewed as an upper bounds on the "rationality" an agent could exhibit, for the domains we are interested in, the agents have much less knowledge and are not able to do all the requisite information-gathering calculations necessary to search such landscapes.

---

7 The model implementation supports an arbitrary number of dimensions, although for landscape displays, of course, the user can only display two at a time.

8 One dimension, for instance, might reflect an agent's preferences with respect to religious authority.

Agents' Actions

In the experiments described here, we focus entirely on the leader's actions. All non-leader agents are essentially without volition. They are not permitted to refuse an offer from the leader to join the regime hierarchy, much less to challenge a demotion or firing. Nor do they exercise any direct control over the position of the regime, though obviously their preferences do factor into the leader's decision-making. We readily acknowledge that these sorts of actions do occur in practice, and we do intend to endow agents with such capabilities in future iterations of the model. For now, however, the model can be thought of as a simplification, applicable only to those settings where the leader enjoys what amounts to absolute authority, subject only to the presence of other elites who exert some influence via the policy stances they hold, the independent power they possess and the formal positions they occupy within the regime.

Specifically, the leader can take three types of actions. The first is to shift the policy of the regime. The second is to reallocate regime power, whether by removing an agent from their formal post ("fire") and bringing in another agent as a replacement ("hire"), or by demoting an agent and thereby consolidating power in his own hands. The third is to do nothing on a given step. The key question is how does the leader decide which of those actions—he is limited to just one— to take at each step? In the experiments described here, we explore two factors that characterize the leaders decision-making style. One is the leader's "rationality", i.e., how often and how effectively the leader takes actions that increase his utility? The other is the leader's "risk type", i.e., does the leader make decisions based on simple heuristics (if-then rules) that reflect a propensity to either risk-averse or risk-taking decisions, without being guided by explicitly considering the effect those decisions will have on the leader's utility. The risk-averse (RA0) and risk-taking (RT0) heuristics are described below. Clearly there are many agent characteristics and cognitive processes that could influence these decision-making styles. The model described here uses just a few mechanisms and parameters to capture at least some of the important aspects.

In short, at each time step, the leader does the following: (1) check to see if conditions are such that RA0 or RT0 (for risk-averse and risk-taking leaders, respectively) should be executed, and if so, take the action specified by the appropriate rule; and (2) if a heuristic is not activated, then try to make a "rational" decision, i.e., one that increases his utility. In particular, a highly rational leader tries to find an alternative regime policy or distribution of r_power that increases his utility. However, a less rational leader sometimes ends up making random changes in policy or power distribution.[9] Risk-type also affects how the leader looks for higher-utility options, as described below.

RA0. Risk-averse leaders are disposed to perceive other powerful agents in the regime as posing a threat to their authority. As a result, they will demote (or failing that replace) agents whose

[9] The reasons for such random moves can include a lack of information, mistaken calculation or judgment, and errors in policy implementation.

power exceeds a certain threshold.[10] In the experiments described here, if an agent has total power that is equal to or greater than 75% of the leaders total power, the risk-averse leader will try to reduce the agent's r_power to the point where his total power is equal to or less than 50% of the leader's power. As a result of this “demotion”, the leader transfers this r_power to himself.

RT0. Risk-averse leaders, by contrast, are prone to shift regime policy to satisfy their personal preferences, regardless of what such a move will cost them in terms of support for the regime. In particular, RT0 is activated when the regime's position deviates too far from the leader's ideal. In the experiments reported here, the difference threshold is 5.0, i.e., 50% of the range of possible preference values. When RT0 is activated, the leader moves the regime policy so as to reduce the difference between his preferences and the regime policy by 75%.

Highly Rational Leaders. A highly rational leader finds alternative regime policies or distributions of r_power that increases his utility. A fully rational leader (in the *homo economicus* sense) would consider all possible alternative actions, and then pick the one that would maximize utility. As such, varying N could be one way to adjust the "rationality" of the agents. In the experiments reported here, the leaders only consider only one alternative per step. We consider leaders who have bounded rationality—they can only consider N alternatives at a time. If that alternative would result in higher utility, the highly rational leader makes that change.[11] Thus, this kind of search for a higher utility alternative state of the regime can be thought of as a kind of (weak) hill-climbing algorithm.

Each time the leader tries to find a better alternative, the leader has an equal probability of considering either an alternative policy or the effect of reallocating r_power from one agent (i.e., fire that agent) to some other agent who has no r_power (i.e., hire that agent). When looking for an alternative policy to evaluate, the leader picks one candidate by adding Gaussian noise (mean=0, StdDev=policySD) to the current regime policy (on each dimension). Note that the policySD parameter controls the range of policy alternatives considered. Thus in the experiments reported here, a risk-averse leader has a low policySD (0.05), i.e., he only considers alternative policies that are close to the status quo. On the other hand, risk-taking leaders use a higher policySD (2.0), meaning that they may consider policy alternatives that are relatively far from the status quo.

When looking for alternative distributions of r_power, the leader randomly picks one agent with r_power for possible firing, and one agent without r_power for possible hiring. The leader then calculates the utility that would result from that transfer of r_power and makes the change if his utility is increased. The leader's risk-type affects this process only for risk-averse agents. In particular, a risk-averse leader will only accept hiring an agent if the agent’s total power is less than 75% of the leaders total power.

10 All of the thresholds, rates, and other values that are used in the leader's decision processes are parameterized in the model, so they could be subject to exploration in future experiments.

11 We assume the leader can calculate the expected utility of an alternative. Adding errors of various kinds to this calculation would be another way of adjusting the rationality of agents.

Low Rationality Leaders. As mentioned above, there are many ways to adjust leader "rationality," i.e., the ability to find higher utility alternative regime configurations. In the model described here, we took a very simple approach: each time a leader tries to increase his utility by the search process described above, there is a probability that the leader will instead just make a random change, i.e., accept a policy or hire/fire alternative without considering the consequences for his utility. In particular, for the experiments described below, low rationality leaders make random moves 80% of the time; only 20% of the time do they consider utility before accepting a change to the regime policy or r_power distribution.

Exogenous Shocks

We introduce random exogenous shocks to include the effects of windfalls and losses on each agent's independent power. Every time step, each agent (including the leader) has a probability (probExogShockRate) of having the independent power (i_power) on one randomly selected policy dimension increased or decreased. The amount of change is obtained by making a random draw from a zero-mean, normal distribution, with standard deviation of probExogShockSD). For the experiments reported here, probExogShockRate=0.1 and probExogShockSD=0.2. If a shock would cause i_power to go below 0, that i_power is set to 0. Note that because i_power is bounded by zero, this exogenous shock process induces a bounded random walk on power; as a result agent’s power will slowly increase over time. Because the runs we do here are relatively short, this long-run change in total power is not important. However, in future versions we intend to normalize power, to ensure that total power does not change (and thereby distort the balance of factors in the utility calculations.)

DESCRIPTION OF EXPERIMENTS

Our objectives in conducting experiments at an early stage in model development are to explore the basic dynamics of the model, to identify measures and tools that will be useful for analysis (e.g., the landscape displays, the ability to move agents by dragging), to conduct some early "face validation" exercises to determine if the model runs plausibly, and to examine the effects of various parameters. To accomplish these tasks, we conduct two simple experiments to explore the dynamics of leadership succession in a closed regime.

FIGURE 1 Experiment 1: Screenshot at Step 1

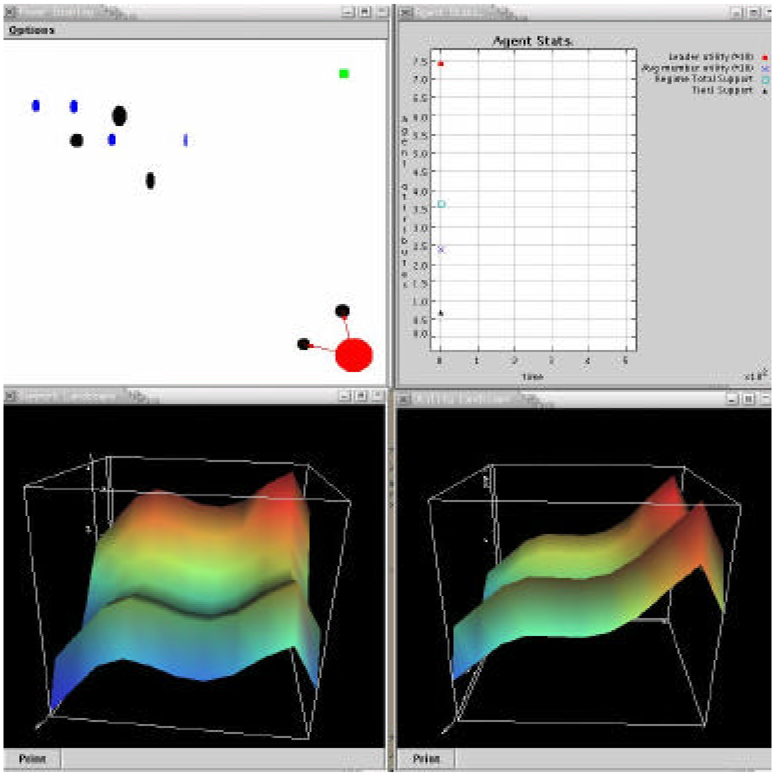

Both experiments are comprised of multiple scenarios. In each scenario, the starting conditions are the same: the incumbent, who is risk averse and highly rational, enjoys strong support from Tier 1 agents (those who hold high-level formal posts in the regime and have low levels of independent power), and weak support from both Tier 2 agents (those who hold low-level posts in the regime and have high levels of independent power) and Tier 3 agents (elites who hold no formal posts in the regime but have high levels of independent power). In the first experiment, we alter the type of the leader, whereas in the second experiment we leave the leader's type fixed, but alter the relative balance of power between the leader and his Tier 1 inner circle, on the one hand, versus all of the other (Tier 2 and 3) agents, on the other.

Figure 1 shows the initial state of the model (i.e., step 1) for all three scenarios of experiment 1. The display labeled "Power" shows all 10 agents, arrayed in the two-dimensional space based on their respective policy preferences. The size and shape of the symbol used to represent an agent conveys the total power (i.e., regime-related plus independent) that he/she enjoys in each policy dimension: the larger the symbol, the higher the power; ovals imply more power on one dimension, corresponding to the longer axis, while circles imply equal power on both dimensions. The leader is colored in red, those with formal positions in the regime are colored black and all others are colored blue. In this experiment, the leader is located in the lower right hand corner; the red lines show "connections"—which are purely illustrative at this stage—to the two Tier 1 agents. The regime's position, indicated by a green square in the upper left corner, is also included for reference purposes.[12]

The display labeled "Support Landscape" depicts the support all agents would afford—the *z* dimension—for a regime position at any one of the potential locations in the policy space—the *x* and *y* dimensions—given their salience weights and preferences, as well as the power they currently enjoy in each policy dimension. It is worth noting that both local and global maxima are possible, reflecting the particular distribution of agents' preferences, the non-symmetric nature of their salience weights, and the non-linear shape of each individual's support function.

The display labeled "Utility Landscape", in turn, depicts the leader's utility if he was to shift the regime's position around in policy space—the *x* and *y* dimensions—given his preferences, his current power and support. What is significant here is that there are two maxima with very similar utilities. Both entail extreme positions on one dimension, i.e., to the far right of the landscape. The leader is effectively indifferent, however, between setting the regime's position on the other dimension to be equivalent to his own (and close to that of the Tier 1 agents), or instead accepting a compromise on that dimension that satisfies the Tier 2 and 3 agents. This circumstance can arise when there is a divergence between what the leader and his close allies, on the one hand, and the other elites, on the other, consider to be the most salient issue(s). This heterogeneity of interests likewise explains why the utility landscape is not symmetrical.

The display labeled "AgentStats" tracks four measures at each time step: the leader's utility, average utility of the agents, total support for the regime, and support from the two Tier 1 agents. For the sake of simplicity, we only report on aggregate support and leader utility in this paper.

One objective of the simple experiments we conduct is to explore the trade-offs a leader faces in attempting to balance two separate considerations:

- maximizing his utility—especially vis-à-vis obtaining more power and bringing regime policy in line with his own policy preferences,
- bolstering support—a component of the leader's utility function

In addition, support also constitutes an important measure of how likely the regime is likely to survive. Of course there are many other measures that are also of interest, e.g., the average utility of the elites represented by agents in the model, which could serve as an alternative measure of potential disenchantment.

[12] We also generate a "Preferences" display, not presented here due to space limitations, which again shows all 10 agents arrayed in the two-dimensional space based on their respective policy preferences. The only difference is that the size and shape of the symbol used to represent an agent conveys the salience weights that he/she accords to each policy dimension: the larger the symbol, the higher the weights; ovals imply a greater weight on one dimension relative to the other, while circles imply equal weights on both dimensions. The regime's position, indicated by a square, is also included for reference purposes.

TABLE 1 Three Succession Scenarios

Experiment	Leader Characteristics	
	Risk	**Rationality**
Scenario 1	Risk Taking	High
Scenario 2	Risk Taking	Low
Scenario 3	Risk Averse	Low

RESULTS

Our first experiment assesses the impact of different types of successors[13] In scenario A, the successor is far more prone to taking risks than his predecessor. That said, both individuals share a disposition to make calculated moves—and therefore display high levels of rationality—as well as the same policy preferences. In scenario B, we instead introduce a successor who is again risk-taking, but far less likely than his predecessor to make rational choices. And in scenario C, the successor shares his predecessor's aversion to risk, but is far more prone to making choices without regard for whether they are rational. These scenarios, which are summarized in Table 1, permit us to explore the effects of the two moving parts in the model: risk and rationality.

Results from the first experiment support our basic intuitions concerning leadership types and levels of rationality. Our results are presented in Table 2 and are based on averages drawn from 30 runs of the model. For each measure (e.g., leader utility, steps 90-99), we report four numbers. To obtain these numbers, we first average leader utility over the specified 10 steps, and so have a mean m_i and standard deviation s_i over those ten steps. In the set of four numbers for leader utility (steps 90-99) the first number (0.7442) is the mean of the m_i's from the 30 runs, the second number (0.0025) is the standard deviation over those m_i's, the third number (0.0071) is the mean of the s_i's from the 30 runs, and the fourth number (0.0051) is the standard deviation over those s_i's. As such, the first number indicates the "average" value of the measure over the specified steps, the second number indicates how much that average varies from run to run, the third number indicates how variable the measure is within the sampled steps of one run (the in run standard deviation), and the fourth indicates how much that variability differs from run to run.

[13] At this stage, whether or not the incumbent dies, is killed, or is forcibly removed from office has no bearing on the nature of succession.

TABLE 2 Experiment 1 Results

Scenarios		Leader Utility *Mean*	*In-Run StdDev*	Regime Support *Mean*	*In-Run StdDev*
A					
steps 90-99	Risk Averse, Rational	0.7442 (0.0025)	0.0071 (0.0051)	4.0932 (0.4357)	0.0611 (0.0444)
steps 100-109	Risk Taking, Rational	0.7695 (0.0061)	0.0094 (0.0060)	3.7623 (0.4999)	0.1920 (0.1385
steps 190-199	Risk Taking, Rational	0.8081 (0.0169)	0.0065 (0.0054)	4.0392 (0.5969)	0.0784 (0.1201)
B					
steps 90-99	Risk Averse, Rational	0.7432 (0.0026)	0.0053 (0.0036)	4.0468 (0.4680)	0.0571 (0.0376)
steps 100-109	Risk Taking, Irrational	0.7639 (0.0142)	0.0156 (0.0063	3.4713 (0.3967)	0.2769 (0.1543)
steps 190-199	Risk Taking, Irrational	0.7627 (0.0206)	0.0171 (0.0065)	3.3656 (0.5129)	0.2736 (0.1605)
C					
steps 90-99	Risk Averse, Rational	0.7435 (0.0032)	0.0084 (0.0059)	3.8048 (0.4211)	0.0791 (0.0344)
steps 100-109	Risk Averse, Irrational	0.7427 (0.0028)	0.0070 (0.0049)	3.7495 (0.3531)	0.0761 (0.0366)
steps 190-199	Risk Averse, Irrational	0.7427 (0.0017)	0.0048 (0.0025)	3.8269 (0.6307)	0.0568 (0.0258)

Thus, in scenario A—where a risk-averse and highly rational leader is succeeded by an equally rational, risk-taking leader—we observe the successor's utility increase slightly immediately after the leadership change (steps 100-109) and more substantially as time unfolds (steps 190-199). By contrast, regime support drops after the change, though it later recovers, albeit not to the level enjoyed by the risk-averse predecessor. To explain these results, we look at the underlying dynamics more closely. The risk-averse leader remains stuck on smaller hill in the upper right portion of utility landscape, given that he only shifts policy minimally. Figure 2 displays the support and utility landscapes for this run at step 99, just prior to succession. At step 100, as depicted in Figure 3, the risk-taking successor assumes power and is able to shift policy (using his RT0 heuristic) sufficiently to increase his own utility; hence the upward tick in the AgentStats display. Support for the regime falls, however, because policy now veers away from the preferences of Tier 2 and Tier 3 agents. Moreover, this shift effectively induces more variation in support for the regime after the change. As the rational risk-taking leader restructures the regime by hiring and firing agents to make up for the loss in support, the contours of the support landscape change dramatically to reflect the increased influence of new regime members with different preferences, salience weights and independent power than the agents they replaced in formal posts, which effectively leads to more variation in support for the regime after the change (See Figure 4).

FIGURE 2 Experiment 1, Scenario A: Pre-Succession Screenshot (Step 99)

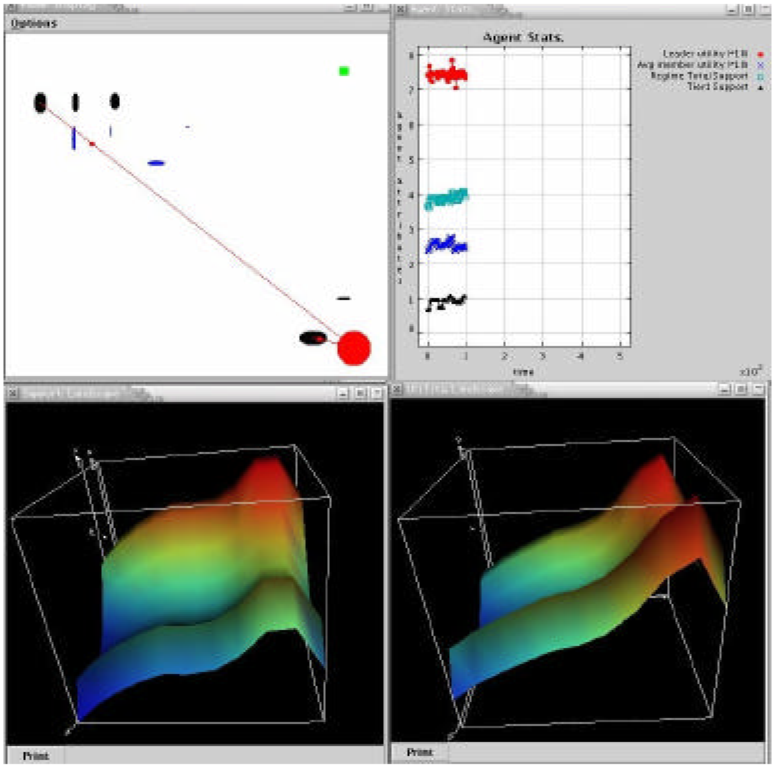

FIGURE 3 Experiment 1, Scenario A: Post-Succession Screenshot (Step 100)

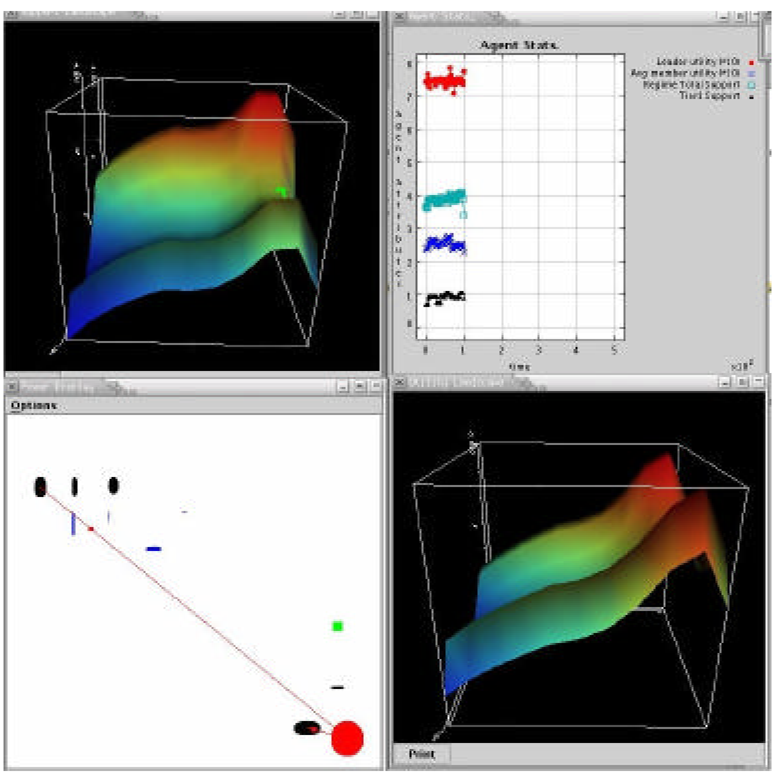

FIGURE 4 Experiment 1, Scenario A: Post-Succession Screenshot (Step 200)

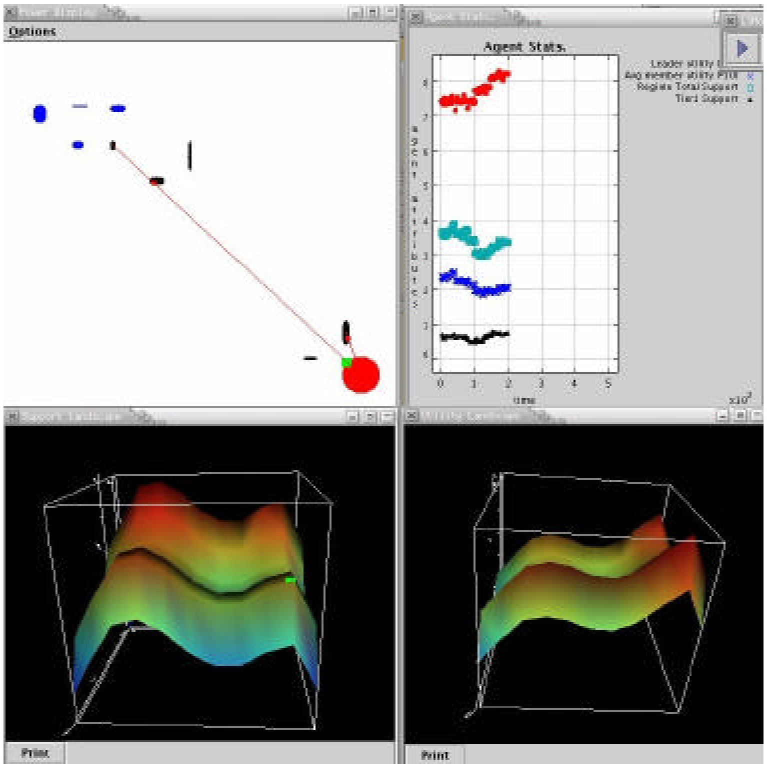

In scenario B, where the successor is now more prone to taking risks and less rational than his predecessor, the successor's utility increases after the change (as does the variance in utility), whereas support for the regime decreases (with a corresponding increase in variance). The successor's propensity to shift policy to mirror his preferences, as well his low rationality, as expressed through random changes in policy and allocation of r_power, accounts for the increased variation in utility and support after the succession.

Finally, in scenario C we hold the risk propensity of the leader and successor constant, and simply examine the effect of a change in rationality. We find that the leader's utility remains stable after the change, as does regime support. The fact that the successor does not consider dramatic changes in the composition and position of the regime offsets his vulnerability to make decisions that are not to his benefit

In our second experiment, we vary the ratio of power in the "leaders corner" (i.e., the cluster in the lower right formed by the leader and the two Tier 1 agents) to power in the upper left corner (Tier 2 and Tier 3 agents). This permits us to further explore the tension between the leader's policy preferences and the location of the support peaks. As in scenario A of experiment 1, the incumbent is risk averse and highly rational, and is succeeded by an equally rational risk-taking leader. The results from this experiment are summarized in Table 3. Note that a ratio of 1.00 indicates that the leader and tier 1 agents have combined total power equal to tier 2 and 3 agents in the upper left corner of the landscape. As this ratio decreases, it reflects the increase in total power of Tier 2 and 3 agents.

TABLE 3 Experiment 2 Results

Scenarios		Regime Support *Mean*	*In-Run StdDev*
Ratio of LR/UL = 1.00			
steps 90-99	Risk Averse, Rational	2.5725 (0.4376)	0.0660 (0.0437)
steps 100-109	Risk Taking, Rational	2.7452 (0.4243)	0.1704 (0.1347)
steps 190-199	Risk Taking, Rational	3.5724 (0.6805)	0.1058 (0.1530)
Ratio of LR/UL = 0.75			
steps 90-99	Risk Averse, Rational	2.7661 (0.4179)	0.0566 (0.0340)
steps 100-109	Risk Taking, Rational	2.8522 (0.3271)	0.1636 (0.1434)
steps 190-199	Risk Taking, Rational	3.5865 (0.6902)	0.1134 (0.1271)
Ratio of LR/UL = 0.50			
steps 90-99	Risk Averse, Rational	3.4217 (0.5392)	0.0665 (0.0294)
steps 100-109	Risk Taking, Rational	3.3440 (0.4810)	0.3154 (0.2603)
steps 190-199	Risk Taking, Rational	4.1976 (0.7027)	0.1994 (0.1918)
Ratio of LR/UL = 0.33			
steps 90-99	Risk Averse, Rational	4.0798 (0.5385)	0.0564 (0.0240)
steps 100-109	Risk Taking, Rational	4.0051 (0.5317)	0.5190 (0.3241)
steps 190-199	Risk Taking, Rational	4.4978 (0.6937)	0.2851 (0.3774)

We find that when the ratio is 1.0, initially support peaks in leader's corner; however, a local maximum exists in the upper right corner where policy starts. At step 100, the jump in policy taken by the risk-taking successor (as a result of RT0) to the lower right corner therefore improves support slightly just after the change, and more at steps 199-200. As the initial ratio of power falls (i.e., Tier 2 and 3 agents have more power), we see a trend toward a smaller increase or even decrease in support just after the successor changes policy, and a smaller increase in support for the regime at the end of the runs. We also see a dramatic increase in the instability of support (the in-run standard deviation) for lower in/out-power ratio cases. [14]

Figure 5 presents a screenshot from one set of runs at an in/out-power ratio of 0.50.[15] Support peaks in the upper left corner where tier 3 agents are located at (and before) step 99. Once the risk-taking successor assumes power (step 100), his propensity (via heuristic RT0) to shift policy to reflect his preferences to the lower right corner comes into conflict with rational (utility-hill climbing) policy shifts. Thus right after the succession, the successor's RT0 is activated, which moves policy to the lower right corner, near the leader's position. Then utility-increasing policy shifts move policy back to one or the other of the peaks in the other three corners of the landscape (see Figure 6). But eventually policy moves far enough away from the successor's preferences to activate RT0 again, moving policy back to the lower right. These

[14] The leader's total power if fixed at 3.5, and tier 1's total power is fixed at 1.5. For a ratio of 1.0, the total power of tier 2 and 3 agents is 5.0; for a ratio of 0.33, the total power of tier 2 and 3 agents is 15. Note that the ratio in our experiment 1 is about 0.75.

[15] Note the in/out-power ratios reflect the initial state of the model. As the leader begins to reallocate power, and as exogenous changes in power occur, these ratios will change.

FIGURE 5 Experiment 2: Pre-Succession Screenshot (Step 99)

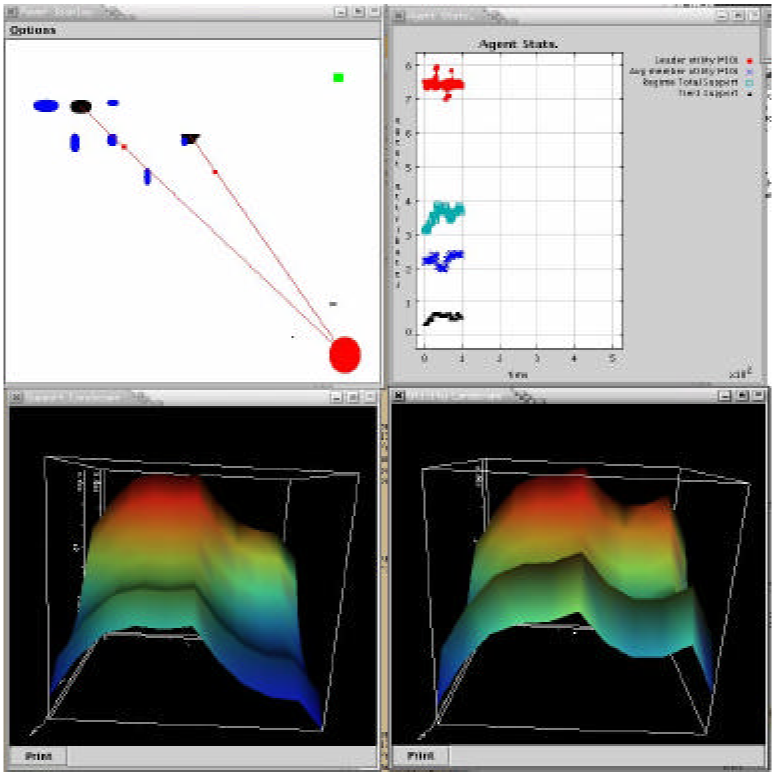

FIGURE 6 Experiment 2: Post-Succession Screenshot (Step 157)

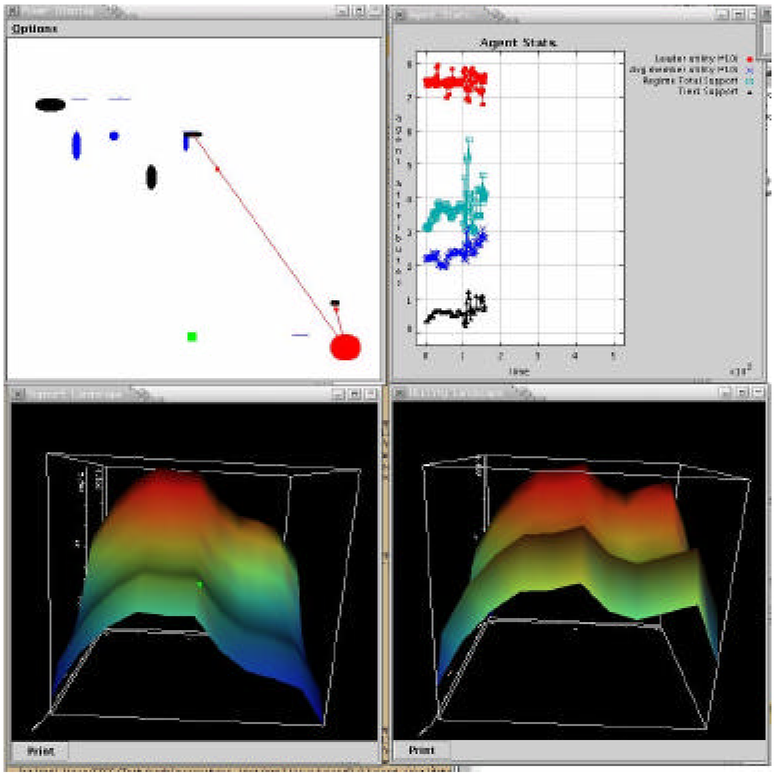

conflicting policy moves continue for the rest of the run, generating high in-run support variance (0.3154 for in/out-power ratio = 0.5 at steps 100-109). However, by the end of the run (steps 190-199), the successor also reallocates power to mitigate the loss in support from shifting policy to reflect his preferences.

To summarize, leader behavior does seem plausible in the initial set of experiments we conduct, given the composition of the regime and the decision rules and behaviors the leaders have in this model. Our observation of a number of runs reveals a variety of histories and outcomes, showing the kind of perpetual novelty that is a hallmark of complex adaptive systems (Holland, 1995; Axelrod and Cohen, 2000). In this model, that novelty is driven by the exogenous shocks and the resulting actions of the leader, i.e., moving the regime policy and reallocating r_power. There clearly is much more that remains to be done to understand the dynamics of this model under different conditions. For instance, the rule for the risk-averse leader (RA0) was seldom used in our runs, and thus had little effect on these results. Perhaps other parameter settings (e.g., a lower threshold for triggering RA0) or other initial allocations of power or exogenous shock dynamics would result in runs in which RA0 was more of a factor.

DISCUSSION

What we have presented is the first cut at an exploratory modeling project, where the aim is to develop a tool for intelligence analysts by applying methodologies for representing non-linear complex adaptive systems to the study of decision-making processes in closed political regimes. We opt for a hybrid of the landscape metaphor and the rule-based system approach, to capture the trade-offs that leaders face in attempting to balance power, policy preferences and regime support—the three components of a utility fitness function—as well as differences in leaders' types that can result in departures from a strict notion of rationality and utility maximization.

Our goals at this early stage of model development are modest: to construct a basic framework and to begin to assess its validity and relevance by means of simple experiments that seek to represent real-world phenomena of interest to intelligence analysts. The experiments reported in this paper focus on the dynamics of succession. In particular, we devise a set of simulations that look at how changes in the leader's type, rationality and level of strategic play via this transition of authority affect the leader's utility and the support for the regime from other elites. Our efforts are clearly not exhaustive; rather, these are merely an initial set of experiments that employ one specific version of the model. Nevertheless, even this foray yielded some interesting dynamics. Of note, we observe that peaks in the utility landscape can arise in surprising places—away from the ideal positions of both the leader and other elites—as the result of compromises they are willing to make. Moreover, the interplay between the leader's basic tendencies to maximize utility and his risk-sensitive heuristic rules also led to much more instability in the regime. The leader would sometimes change the regime’s policy to make it close to his preferences, and other times change policy and reallocate power among other agents to raise his utility by garnering support for the regime.

In the immediate future, we plan to continue the process of exploration and validation of this and related agent-based models of closed regimes. For the current model, we need to test the dynamics using extreme parameter settings, different arrangements of agents preferences, different distributions of power among agents, different weights between issues and among the components of utility, and different levels of rationality, both in the context of succession and in other applied settings. Our efforts to date have been aided in no small measure by the feedback we received from analysts and domain experts who have seen demonstrations of earlier versions of this model. Not only did these meetings stimulate thought and provide us with an opportunity to discuss which factors are important, which are not, and how to include important factors and mechanisms in simple models, but the response also indicated that the model matched the way (some) analysts view (some) regimes, and that our notion of leadership behavior and priorities seemed plausible.

Ultimately, we intend to modify the current model to reflect the observations and ideas we have received (and continue to receive) from analysts and domain experts. Some factors and mechanisms that are not in the current model but which were identified as crucial for understanding decision making in closed regimes include formal and informal organizational structure, as well as making agents' power a function of the relationships they have with other agents. In addition, we expect to represent political groupings within the regime itself (i.e., factions and coalitions) as well as among those who are outside the regime, including opposition elements, to reflect influences of competing forces and alternatives to the regime.

References

Axelrod, R. and Bennet, S. 1993. "A Landscape Theory of Aggregation," *British Journal of Political Science* 23 (April): 211-233.

Axelrod, R. and M. Cohen. 2000. *Harnessing Complexity.* Free Press.

Bankes, Steve. 1994. Exploratory Modeling for Policy Analysis. Operations Research, 41:3: 435-449.

Casti, J. 1997. *Would-Be Worlds: How Simulation is Changing the Frontiers of Science.* NY.

Holland, J., K. Holyoak, R. Nisbett and P. Thagard. 1986. *Induction: Processes of Inference, Learning, and Discovery.* The MIT Press, Cambridge, MA.

Kollman, K., J. Miller, and S. Page. 1992. "Adaptive Parties and Spatial Elections." *American Political Science Review.* 86:929-937.

Kollman, K., J. Miller, and S. Page. 1998. "Political Parties and Electoral Elections." *British Journal of Political Science.* 28:139-158.