

Simulating Closed Regimes with Agent Based Models

This article describes efforts to develop an exploratory agent-based model as a tool for studying decision making in political regimes such as Iraq, North Korea, and Syria. Our hybrid of the landscape metaphor and the rule-based system approach captures the trade-offs leaders face in balancing components of a utility function, plus risk profiles that allow departures from conventional utility maximization. Two simple experiments concerning succession demonstrate the surprising compromises both leaders and elites are willing to make, as well as the instability of these bargains. © 2008 Wiley Periodicals, Inc. Complexity 14: 36–44, 2008

R. BHAVNANI, D. BACKER, AND R. RIOLO

R. Bhavnani is affiliated with the Department of Political Science, Michigan State University, East Lansing, Michigan 48824 (e-mail: bhavnani@msu.edu)

D. Backer is affiliated with the Department of Government, College of William and Mary, Williamsburg, Virginia

R. Riolo is affiliated with the Center for the Study of Complex Systems, University of Michigan, Ann Arbor, Michigan

1. INTRODUCTION

Autocratic governments have long drawn interest from both scholars and government analysts, seeking to understand better how they operate and potentially to anticipate the sorts of decisions they are likely to make in domestic and foreign policy contexts. Over recent years, specific attention has been paid to countries such as Afghanistan, Burma, China, Cuba, Iran, Iraq, Libya, North Korea, Syria, and Zimbabwe. During previous eras, the focus at different points was on the fascist systems in Germany and Italy; the Communist monopolies of the Soviet Bloc; and military juntas and other dictatorships across Latin America, as well as in Africa and Asia. A common feature of these governments is severe restrictions on openness, pluralism and competition in many spheres of society.

Comprehending the dynamics of such closed political regimes presents a challenge for traditional analytical methods. As a general matter, these regimes exhibit outcomes that are difficult to explain or trace back to their causes. In particular, the actions of their leaders often appear to be conditional on subtle variations in circumstances, obscure, unpredictable, or even illogical. When one factors in the secretive, insular nature of these regimes and the intricacies of and differences

This is an invited paper for the Special Issue—Security and Complexity, Guest Editor: Jürgen Scheffran.

among the constituent actors and the settings in which they operate, it can be difficult to achieve adequate explanatory leverage with statistical, mathematical, and game-theoretical techniques, informed by the standard means of collecting quantitative and qualitative primary data.

We opt instead for an analytical strategy that involves considering closed regimes as complex adaptive systems (CAS). In our view, this alternative makes sense because these regimes have the hallmarks of CAS, which include (a) heterogeneous agents, (b) interacting with each other and their environment in nonrandom ways, (c) often using non-linear rules to decide what actions to take, (d) both creating and responding to positive and negative feedback processes, and (e) sometimes adapting and learning to achieve better results over time. Various types of computer simulations permit one to capture these features and to explore their association with the dynamic properties of the system.

For purposes of studying closed regimes, we develop an exploratory agent-based model [1–3], employing a hybrid of the landscape metaphor [3–5] and the rule-based systems approach [6], while also building upon prior work concerning coalition formation [7], heuristic-driven decision-making [8–13], the cognitive dispositions of leaders [14–16], and expected utility [17–19]. Sections 2–4 outline the foundations and components of the model. Sections 5–7, in turn, describe the results of a pair of experiments—designed in part for purposes of testing and validation—where we examine the phenomenon of leadership succession and the implications of competing priorities for regime stability.

2. AN ARCHETYPE OF CLOSED REGIMES

The characteristics of closed regimes vary. Nonetheless, these regimes tend

to share several attributes that are relevant to our modeling efforts.¹

First, decision-making authority is concentrated. A closed 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, with the power to actively manage the regime (i.e., hire, fire, reassign, promote, demote, etc.). Yet, they would almost inevitably be specified in a manner that is more elaborate than a traditional agent-based model, both in terms of the number and variety of traits plus their repertoire of actions.

Second, in closed regimes, the overwhelming influence of a single individual or small coterie of elites is typically reinforced by a lack of institutional barriers to decision-making. In the extreme, there may be no consequential distinctions between executive, legislative and judicial authority, which are

¹*Political regimes are already the subject of an extensive literature. For our purposes, the material that is devoted to describing the distinctions among different types of regimes, as well as understanding the transitions between these types, is largely peripheral. Instead, we are principally concerned with the internal dynamics of a particular class of regimes. The archetype draws generally on discussions of autocratic regimes, plus incorporates insights provided by area studies 'domain experts' on specific empirical cases.*

effectively collapsed into a unified command-and-control institutional structure. This fusion 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—in the sense of requiring consultation if not voting—or those that nominally exist are trumped by informal norms and arrangements, as well as the ability of key actors to issue binding edicts. 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 certain measures of state capacity. In the exploratory model developed here, we opt to focus on the actions of a dominant leader who exercises absolute discretion as far as making decisions about the composition of the regime.

Third, the consolidation of authority is often associated with both insecurity and aggression. 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 belligerence, 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 cultural (i.e., ethnic, religious, etc.) or ideological identity that emphasizes confrontation with and endorses the exclusion of nonadherents 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 the model, we consider potential sources of affinity and division among elites, as well as permit punitive actions directed against perceived threats.

Fourth, the leaders of closed regimes typically seek to insulate themselves from potential risks by restricting the flow of information. As we indicated 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 holding 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.

3. MODEL DEVELOPMENT

We believe decision-making processes in closed regimes are too intricate and insufficiently documented to construct a consolidative model that can generate precise predictions. Instead, our model is exploratory, capturing key characteristics and behaviors of agents in these contexts. The aim is to afford insight into dynamics that may lead to a particular result or engender a certain phenomenon, with an emphasis on mechanisms rather than outcomes.

We built the model from the ground up, incorporating guidance from intelligence analysts, who shaped our stylized representation of a closed regime, rather than supplying data for direct analysis.

In particular, they provided background on types of agents (number, traits, status, influence, actions, heuristics, constraints) and the connections among them (political and social structures, nature of interactions), as well as factors in the political, social and economic environment (constants, variables, triggers, shocks). These inputs were complemented by input from area studies domain experts. Our subsequent efforts were then devoted to developing a modular set of models, each pertaining to a specific substantive context (e.g., succession, repression, invasion) or class of phenomena (e.g., coalition politics).

4. MODEL DESCRIPTION

We employ a hybrid approach for several reasons. The landscape metaphor captures analysts' basic conception of elites seeking to navigate to their advantage among a complex set of alternative states, as well as provides a visual depiction of the tensions among competing objectives. The incorporation of rule-based heuristics reflected analysts' intuitions that individuals in these regimes often operate with first-order norms or principles (e.g., survival), as well as permitted us to explore the impact of leadership styles that depart from traditional utility-maximizing behavior. Here, we provide a conceptual-level description of the model components, their micro-mechanisms and the overall dynamics (for further details, see [20]).

4.1. Agents

The model consists of a leader, individuals who hold formal positions in the regime hierarchy (Tiers 1 and 2), and other elites (Tier 3). The tiers are effectively shorthand for the status of an agent.² All

²At this stage, the tiers do not entail anything in the model as far as relationships among agents, nor is movement across these different levels limited.

agents have several traits. To begin with, they have preferences [0,10] and salience weights [0,1] on each of a set of M policy issues, which are specified at the start of a run to reflect the regime configuration of interest. An agent may also derive power from a formal position (*r_power*). In addition, agents can have power independent of the regime (*i_power*).³ All power is issue specific (i.e., an agent potentially influences particular policy decisions more than others) and—unlike preferences—can vary, as the result of the leader's actions and exogenous shocks.

4.2. Utility Functions

The utility function of the leader *L* is specified as follows:

$$U_L(D, P, S) = \alpha_D \times D + \alpha_P \times P + \alpha_S \times S$$

where:

- *D* is the distance between the leader's preferences and the regime's policies (weighted by issue salience);
- *P* is the leader's total power (both regime-based and independent);
- *S* is the aggregate support for the regime, reflecting the salience- and power-weighted policy distances between the regime's policies and the preferences of all agents, including the leader; and
- α_P , α_D , and α_S are weights [0,1] that reflect the leader's relative interest in power, policies that mirror his personal preferences, and support for the regime.

The other agents differ in that their utility functions ignore regime support.

4.3. The Landscapes

The trade-offs the leader faces in making personnel and policy decisions can be depicted via two $M + 1$ dimen-

³Note that Tier 3 agents, despite their lack of formal status, can actually possess more power than those with positions in the regime.

sional landscapes—one for leader utility, the other for regime support. For ease of visual representation, we consider models where $M = 2$. Thus, the x and y dimensions correspond to two issues, each with a range of policy options $[0,10]$. The exact nature of these issues is not critical, though logically they should not be entirely collinear, so an agent can exhibit contrasting preferences. Where the agents array themselves in this policy space, as well as where the leader locates the regime's policies, determines the support for the regime, which, in turn, affects the leader's utility. On the two landscapes, the z values indicate the leader's utility and the regime's support, respectively, if the regime's policies are set to x, y .

Because of information restrictions, comprehensive knowledge of the traits of all of the other agents in the model is unrealistic. As a result, each agent's rationality is bounded: rather than performing the calculations necessary to define and search the entire landscape, the individual agents are constrained to contemplating a more limited range of options.

4.4. Agents' Actions

Our current analysis focuses on a dominant leader, who can take one of three actions at each time step. First, he can shift the regime's policies. Second, he can reallocate power, either by removing an agent from their position ("fire") and bringing in a replacement ("hire"), or by demoting an agent and consolidating authority. Third, the leader can opt to do nothing.

Various agent characteristics and cognitive processes could influence decision-making. Our model considers two integral factors. One is the leader's general level of rationality, i.e., how often and effectively he takes actions that increase his utility. The other is his inclinations towards risk-averse or risk-taking decisions, irrespective of their impact on his utility.

At each time step, therefore, the leader initially considers whether con-

ditions are such that he should act on one of the following (depending on his type) risk-oriented heuristics:

- RA0. A risk-averse leader is prone to perceive other agents in the regime as posing a threat to his authority and security. Hence, he demotes—or failing that, replaces—agents whose power exceeds a certain threshold.
- RT0. A risk-taking leader is willing to alter regime policies simply to suit his own interests. Once the regime's policies deviate too far from his ideal, he will shift them to be closer to his personal preferences.

If a heuristic rule is not activated, then the leader's decision depends on his level of rationality:

- A highly rational leader identifies alternative regime policies or distributions of r_power that increase his utility. A fully rational leader would consider all possible options and pick the one that would maximize utility. We consider leaders whose rationality is bounded: they can only consider N alternatives at a time (but then opt for the one that affords the greatest gain in utility).⁴
- A low-rationality leader may instead implement a random change, i.e., choose a policy or hire/fire alternative without considering the consequences for his utility.⁵

At the same time, risk type also affects how the leader looks for higher-utility options.

⁴We 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.

⁵Such random moves can be due to a lack of information, miscalculation and errors in implementation.

5. DESCRIPTION OF EXPERIMENTS

Next, we sought to explore the basic dynamics of the model, to conduct validation exercises to determine if it runs plausibly, to test features of use for analysis, and to assess the impact of various parameters. For these purposes, we conduct two simple experiments to examine the phenomenon of leadership succession, the tradeoffs between maximizing personal utility and bolstering the support of other elites,⁶ and the resulting implications for future regime stability.

Both of the experiments share the same starting conditions: a risk-averse, highly rational leader enjoys strong support from Tier 1 agents (high-ranking regime members with limited independent power) and weak support from both Tier 2 agents (low-ranking members with considerable independent power) and Tier 3 agents (no formal status, but significant independent power). Figure 1(a) shows the initial state of the model, summarized by four visual elements:

- The power display (upper left) arrays all 10 agents in the two-dimensional space based on their respective policy preferences. The symbols indicate the total power that agents exercise in each dimension: the larger the symbol, the greater the power; ovals imply more power on the longer-axis dimension, 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. The red lines are connections—purely illustrative at this stage—between the leader and the two Tier 1 agents. The regime's policies are indicated by the green square.

⁶Alternatively, one could consider average utility as a measure of satisfaction/disenchantment.

FIGURE 1

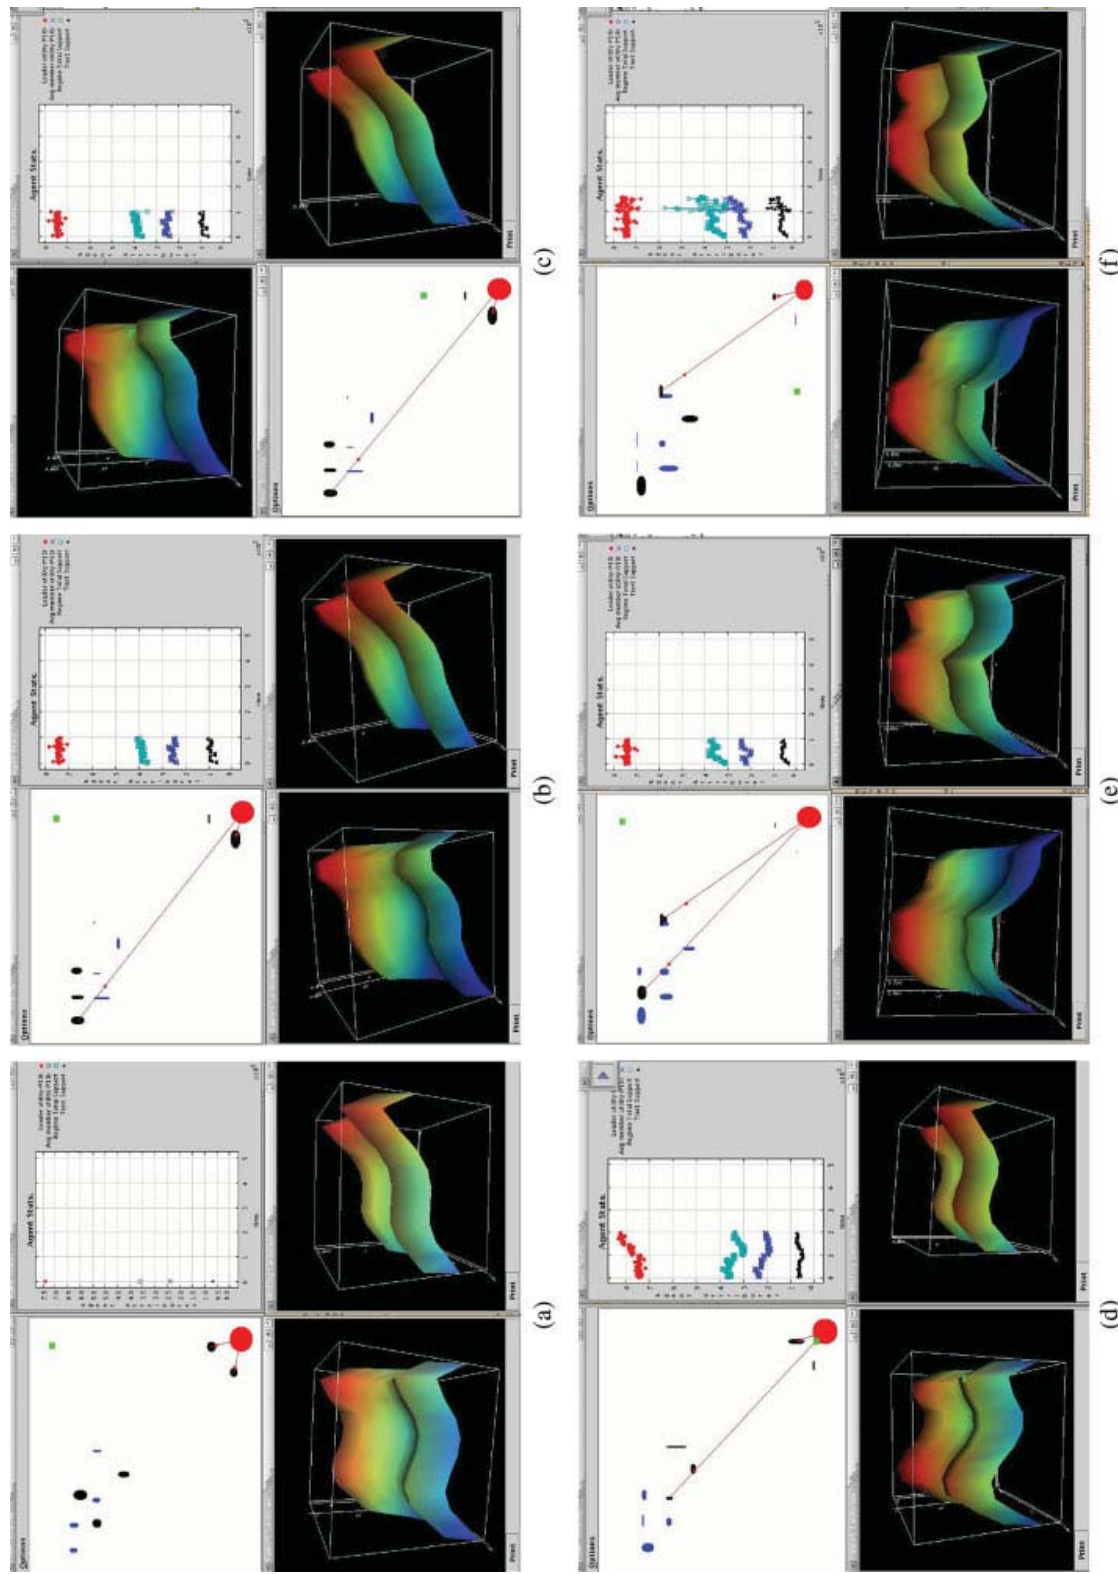


TABLE 1

Three Scenarios of Succession

Experiment	Characteristics of Successor	
	Risk	Rationality
Scenario A	Risk taking	High
Scenario B	Risk taking	Low
Scenario C	Risk averse	Low

- The AgentStats chart (upper right) tracks the leader's utility, the average utility of the agents, aggregate support for the regime, and support from the Tier 1 agents.
- The Support Landscape (lower left) shows how much agents would back potential regime policies. Both local and global maxima are possible, reflecting particular distributions of agents' preferences, the asymmetric nature of their salience weights, and the non-linear shape of their support functions.
- The Utility Landscape (lower right) depicts the leader's utility depending on where he locates the regime in the policy space. Here, two maxima exist with very similar utilities. Each entails an extreme stance on one dimension. The leader is virtually indifferent, however, between setting the regime's policy on the other dimension to be equivalent to his own preferences and close to that of the Tier 1 agents, or accepting a compromise policy that satisfies the other agents. This circumstance reflects disparities in the salience individual agents attach to the issues—a heterogeneity of interests that also explains the asymmetrical utility landscape.
- With some probability exogenous shocks change the i_power of a randomly selected agent
- At a user-specified step (100 in the runs reported here), the leader agent can also be replaced by a successor, with their own heuristics for picking actions
- Every non-leader agent calculates its own utility, and its support for the current regime policy
- The leader agent calculates the aggregate support for regime policy, and its own utility
- The leader uses its heuristics and its utility calculations to choose an action (change policy, fire/hire agents, reallocate r_power).
- The model reports/displays measures (leader utility, regime support, etc.)

These steps are repeated until a given run is terminated by the user.

Experiment 1 considers the impact of different types of successors by manipulating two parameters in the model: the leader's risk and rationality characteristics (Table I).⁷ For *Scenario 1-A*, the successor is far more prone to taking risks than his predecessor, but otherwise the two are identical. For *Scenario 1-B*, the successor is once

⁷In this version of the model, the circumstances that led to succession have no bearing on subsequent events.

again risk-taking, but is also far less likely than his predecessor to make rational choices. For *Scenario 1-C*, the successor shares his predecessor's aversion to risk, yet is far more prone to making irrational decisions.

Experiment 2 assumes the same leadership types that were employed in *Scenario 1-A* (i.e., successor is risk taking, unlike his predecessor; both exhibit high levels of rationality), but varies the ratio of the power held by the 'inner circle' (i.e., the leader and the Tier 1 agents) to the power exercised by the other agents. This manipulation permits us to explore the tension between the leader's policy preferences and what is entailed to maximize regime support.

6. RESULTS

The outcomes of *Experiment 1*—involving 30 runs of the model for each scenario—are consistent with our intuitions concerning the impact of leadership types and levels of rationality. Table 2 reports four measures for both the leaders utility and regime support associated with relevant 10-step intervals in the simulations:

- the mean across the runs of the model (m_i) for a given interval;
- the standard deviation (s_i) of the m_i values across the runs (in parentheses);
- the mean of the s_i values across the runs; and
- the standard deviation of those s_i values (in parentheses).

The first measure is essentially an overall average for the specified steps. The second measure captures the extent to which that average varies from run to run. The third measure indicates how variable the second measure is within the sampled steps of one run (the in-run standard deviation), and the fourth measure how much that variability differs from run to run.

With *Scenario A*, we observe that the risk-taking successor enjoys a slightly higher utility than his risk-

For each model run, after the initial conditions are created, at every time step the following occurs:

TABLE 2

Results of Experiment 1

Scenarios	Leader's Utility		Regime Support		
	Mean	In-Run StdDev	Mean	In-Run StdDev	
1-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)
1-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)
1-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)

averse predecessor in the immediate aftermath of the leadership change (steps 100–109). This utility gap increases over time (steps 190–199). Meanwhile, regime support drops after the change, though it later recovers, albeit not to the level enjoyed by the risk-averse predecessor.

To understand these results, one can look at the underlying dynamics. Figure 1(b) displays the support and utility landscapes for a representative run at step 99, just prior to succession. The incumbent leader remains stuck on smaller hill in the upper right of utility landscape, since he is only capable of modest utility-enhancing shifts in policy as a result of his risk aversion. At step 100 [Figure 1(c)], his successor assumes power and his RT0 heuristic enables him to shift policy sufficiently to increase his utility. Support for the regime falls, however, because regime policy now veers away from the preferences of Tier 2 and Tier 3 agents. This shift then induces subsequent responses that result in further fluctuations. In an effort to make up for the loss in support, the successor restructures the regime by firing

some agents from their formal positions and substituting others. The increased influence of agents whose preferences, salience weights and independent power differ from those they replaced transforms the support landscape [Figure 1(d)].

With *Scenario B*, the leader's utility increases after the succession (as does the variance in utility), whereas support for the regime decreases (with a corresponding increase in variance). These outcomes are attributable to the successor's propensity to shift policy to mirror his own preferences, as well to his low rationality, which manifests in random changes to policy and to the allocation of regime power.

With *Scenario C*, we find that both the leader's utility and regime support remain stable after the succession. The fact that the successor does not consider dramatic changes in the composition and policies of the regime offsets his propensity to make decisions that are not to his benefit.

Turning to *Experiment 2* (Table 3), we find that when the power ratio is 1.0, support initially peaks in the leader's corner; however, a local maximum

exists in the upper right corner, the location of the regime's policies. At step 100, the risk-taking successor alters policy (as a result of RT0) to the lower right corner, which improves support to a marginal extent immediately after the change and to a greater degree by steps 199–200.

With a lower initial power ratio (i.e., Tier 2 and 3 agents possess relatively greater power), we observe a tendency toward a smaller increase or even a decrease in support right 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 example, Figure 1(e) displays step 99 in a run using a power ratio of 0.5 (i.e., Tier 2 and 3 agents have strong sources of independent influence that more than offset their limited formal authority).⁸ Support is highest in the

⁸These ratios reflect the initial state of the model. They change once the leader starts reallocating power and as power shifts exogenously.

TABLE 3

Results of Experiment 2

Scenarios	Regime Support		
	Mean	In-Run StdDev	
Power Ratio = 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)
Power Ratio = 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)

upper left, where the Tier 3 agents are located. Once the new leader assumes power (step 100), his rule-driven propensity to shift policy to reflect his own preferences comes into conflict with his rational utility-maximizing inclinations. Right after the succession, RT0 is activated, which moves policy to the lower right corner. Yet he then seeks to increase his utility by moving policy to one of the peaks in the other corners of the landscape [see Figure 1(f)]. When the policy becomes too distant from the leader's preferences, RT0 is activated again. These alternations persist, generating high in-run variance of support. By the end of the run, however, the leader reallocates power to mitigate the loss in support from shifting policy to reflect his preferences.

7. DISCUSSION

This article presents the first cut at developing a tool for intelligence analysts by applying methodologies for representing non-linear CAS to the study of decision-making in closed political regimes. Our exploratory agent-based model, which incorporates both the landscape metaphor and the rule-based system approach, is designed to capture the trade-offs leaders face in balancing competing priorities, as well as differences among leaders that prompt deviations from

utility maximization. This model, therefore, is a hybrid of several existing frameworks, including aspects of some traditional ways of studying authoritarian politics, but also capitalizes upon a flexible tool that allows us to examine the impact of notable departures from conventional assumptions (e.g., rational behavior, full information).

We begin to assess the validity and relevance of the model design by conducting two simple experiments of succession dynamics that examine the impact of changes in the leader's type and level of rationality and strategic play. Overall, one observes a variety of histories and outcomes across runs of the model—the kind of perpetual novelty that is typical of CAS [21, 22]. In this model, the major sources of novelty are exogenous shocks, in the form of a change of leaders, and the resulting actions of the successor. The behavior exhibited by leaders seems plausible given the actions they can take, the decision rules that they utilize and the composition of the regime.

A number of interesting findings emerged from this limited foray. Most notably, the utility landscape exhibits peaks in surprising places—away from the policy ideals of both the leader and other elites—as the result of compromises they are willing to make, based

in part on differences in priorities. Moreover, the interplay between the incumbent leader's risk-sensitive heuristics and general desire to maximize utility engenders instability: he fluctuates between egocentric actions that favor accumulating power and adjusting the regime's policy to his preferences, and alternative policy changes and reallocations of power designed to bolster support for the regime among other agents.

At the same time, additional work is required to more fully understand the dynamics of this model. For instance, risk-averse leaders rarely invoked their heuristic and demoted or replaced strong rivals. Other parameter settings (e.g., a lower threshold for threat perception) or other initial allocations of power or exogenous shock dynamics might yield runs in which RA0 was more of a factor. The next step, therefore, is to test out the effects of using different sets of agents preferences, distributions of power, weights between issues and among the components of utility, heuristic parameters, and levels of rationality, in the context of both succession and other applied settings of interest to analysts.

Meanwhile, we have also been developing a variant of Axelrod and Bennett's [7] model of alliance formation. In our version, agents with ascriptive traits (e.g., ethnicity, religion, tribe, etc.), predilections (i.e., policy preferences, party/institutional affiliations, inter-personal histories) and influence organize into distinct coalitions so as to minimize their "frustration," which increases when agents with dissimilar traits are in the same coalition, and decreases as trait differences within a coalition are minimized. This extension captures the notion that an agent's influence is a function of both their inherent power and the power they derive from patron-client relationships with other agents, refines our conception of active management by a leader, and permits other agents to be more active players.

REFERENCES

1. Bankes, S. Exploratory modeling for policy analysis. *Oper Res* 1994, 41, 435–449.
2. Casti, J. *Would-be Worlds: How Simulation is Changing the Frontiers of Science*; Wiley: New York, NY, 1996.
3. Miller, J.; Page, S. *Complex Adaptive Systems. An Introduction to Computational Models of Social Life*; Princeton University Press: Princeton, NJ, 2007.
4. Kollman, K.; Miller, J.; Page, S. Adaptive parties and spatial elections. *Am Polit Sci Rev* 1992, 86, 929–937.
5. Kollman, K.; Miller, J.; Page, S. Political parties and electoral landscapes. *Br J Polit Sci* 1998, 28, 139–158.
6. Holland, J.; Holyoak, K.; Nisbett, R.; Thagard, P. *Induction: Processes of Inference, Learning, and Discovery*; The MIT Press: Cambridge, MA, 1989.
7. Axelrod, R.; Bennett, S. A landscape theory of aggregation. *Br J Polit Sci* 1993, 23, 211–233.
8. Christensen, E.; Redd, S. Bureaucrats versus the ballot box in foreign policy decision making: An experimental analysis of the bureaucratic politics model and the poliheuristic theory. *J Conflict Resolut* 2004, 48, 69–90.
9. Dacey, R.; Carlson, L. Traditional decision analysis and the poliheuristic theory of foreign policy decision making. *J Conflict Resolut* 2004, 48, 38–55.
10. DeRouen, K.; Sprecher, C. Initial crisis reaction and poliheuristic theory. *J Conflict Resolut* 2004, 48, 56–68.
11. Goertz, G. Constraints, compromises, and decision making. *J Conflict Resolut* 2004, 48, 14–37.
12. Mintz, A. How leaders make decisions?: A poliheuristic perspective. *J Conflict Resolut* 2004, 48, 3–13.
13. Mintz, A. Foreign policy decision making in familiar and unfamiliar settings: An experimental study of high-ranking military officers. *J Conflict Resolut* 2004, 48, 91–104.
14. Hermann, M. Explaining foreign policy behavior using the personal characteristics of political leaders. *Int Stud Q* 1980, 24, 7–46.
15. Hermann, M. On foreign policy makers, personality attributes, and interviews: A note on reliability problems. *Int Stud Q* 1980, 24, 67–73.
16. Post, J., Ed. *The Psychological Assessment of Political Leaders: With Profiles of Saddam Hussein and Bill Clinton*; University of Michigan Press: Ann Arbor, MI, 2003.
17. de Mesquita, B.B.; Siverson, R. War and the survival of political leaders: A comparative study of regime types and political accountability. *Am Polit Sci Rev* 1995, 89, 841–855.
18. De Mesquita, B.B.; Smith, A.; Siverson, R.M.; Morrow, J.D. *The Logic of Political Survival*. MIT Press: Cambridge, MA, 2003.
19. Miller, R. Regime type, strategic interaction, and the diversionary use of force. *J Conflict Resolut* 1999, 43, 388–402.
20. Bhavnani, R.; Backer, D.; Riolo, R.A. Hybrid Model of Decision-Making in Closed Political Regimes. In: *Proceedings of Workshop on Social Agents: Ecology, Exchange, and Evolution*, Gleacher Center, University of Chicago, October 11–12, 2002.
21. Holland, J. *Hidden Order: How Adaptation Builds Complexity*; Addison Wesley: Reading, MA, 1995.
22. Axelrod, R.; Cohen, M. *Harnessing Complexity*; Free Press: New York, NY, 2000.