A LANDSCAPE MODEL OF RULE-BASED CO-ADAPTATION IN INTERNATIONAL BEHAVIOR

Philip A. Schrodt
Department of Political Science
University of Kansas
Blake Hall
Lawrence, KS 66045
USA

voice: 785-864-9024
fax: 785-864-5700
email: p-schrodt@ukans.edu

February 1994

An earlier version of this paper was presented at the 1993 meeting of the International Studies Association, Acapulco. This research was supported in part by a subcontract to National Science Foundation grant SES90-25130 (Data Development in International Relations). My thanks to Shannon Davis, Deborah J. Gerner and Judith L. Weddle for assistance in developing the Middle Eastern event data series.

ABSTRACT

This paper discusses a model of rule-based adaptive behavior in foreign policy based on Wright's concept of an "adaptive landscape" and Kauffman's recent work on the characteristics of optimization on such a surface. The model accounts for several regularities observed in international behavior. First, behavior governed by a landscape model will show long periods of stability punctuated by periods of rapid change. Second, the existence of stable rule regimes is virtually inevitable; in the presence of a hegemon these regimes are established more quickly but have a utility lower than that likely to occur in the absence of a hegemon. Third, incremental strategies are rational and precedent is useful as a guide to policy. Finally, innovation occurs primarily in the context of crises.

Kauffman and Johnsen suggest that a system whose level of rule interactions optimizes utility on a landscape will exhibit a power-law distribution of changes. This prediction is strongly supported in an examination of WEIS event data for the Middle East (1982-1992) and COPDAB data for the USA→USSR (1948-1978). A tuned system should also show a 1/f power spectrum. This spectrum is found only in dyads exhibiting protracted conflict; departures from the 1/f pattern seem to be related to the degree of cooperation in the dyad.

Introduction

This article will explore the implications of a model of co-adaptation, the process by which two or more interdependent states adapt their foreign policies over time. The model studied is Kauffman's co-evolutionary "NK model" (Kauffman and Levin 1987, Kauffman and Johnsen 1992), a variant on the "landscape model" originally purposed for the study of biological coevolution by Wright (1931). Despite the differences between organizational and biological optimization, the information constraints inherent in foreign policy decision making make this model appropriate and the deductions of the model are consistent with many characteristics of international systems, for example incremental and precedent-based decision making and the existence of stable international regimes. The model is also useful in explaining the otherwise puzzling combination of stability and instability found, for example, in the outbreak of WWI and the collapse of communism in Eastern Europe. Finally, some aggregate characteristics of behavior predicted by the landscape model are supported in a statistical analysis of international event data.

The landscape model is generally consistent with the rational choice model of international behavior (e.g., Ordeshook 1989, Zagare 1990 and Nicholson 1992). Like rational choice, it assumes that international actors have preferences and are goal seeking. Consistent with most contemporary work on decision making in international politics, the model also assumes that decision-makers do not have full information. Unlike rational choice, the goal seeking need not be based on expected utility maximization, though the model does not preclude this possibility.[1]

Adaptation in a landscape model deals with rules rather than behaviors. Foreign policy organizations are generally rule-following: On a day-to-day basis, most activities reflect the implementation of rules established at a prior point in time (in many cases decades earlier) rather than resulting from the novel solution of an optimization problem. Rules are not constant—they are changed as a result of catastrophic failure or incrementally—but they change slowly except in unusual situations such as crises.[2] Rules derive from goal-seeking to maximize utility in a given environment; behavior derives from rules; utility derives from behavior.

[1] Expected utility maximization is arguably the most problematic aspect of the classical rational choice paradigm since a decade and a half of experimental studies (see for example Kahneman, Slovic and Tversky 1982, Hogarth and Reder 1987) have shown that under a wide variety of conditions, individual decision-makers, including experts, make decisions inconsistent with the theory. This does not preclude expected utility as a model: for example expected utility might be a useful first approximation or organizational constraints might cause foreign policy decisions to be made "as if" expected utility were maximized even if individuals are not doing so. However, given the empirical evidence against expected utility, a model circumventing dependence on the assumption is preferable to one with expected utility as a core assumption.

[2] "Rules" are any explicit or implicit factor that causes an organization to respond to a given set of conditions in a consistent fashion over time: these would include explicit rules such as treaties and legislative constraints, bureaucratic constraints such as standard operating procedures and organizational norms, and implicit rules such as organizational codes, shared historical precedents and cultural norms. A variety of studies over the past ten years have shown that relatively complex sets of international behaviors can be simulated using relatively simple sets of

The choice of rules followed in foreign policy is *co-adaptive*: it is dependent on the choices of other actors in the system. Co-adaptation is a familiar assumption, since interdependent choice is at the core of game theory. Axelrod's (1984) study of emergent strategies in the iterated Prisoners' Dilemma provides a good example of co-adaptation in a set of rules—strategies for playing the IPD game—and more generally analytical game theory deals with deducing strategies based on assumptions about the consequences of interdependent choices. Optimization on a landscape involves a very different mechanism than game theory, but the underlying principle of interdependence is the same.

Adaptive Landscapes and the NK Model

The concept of an adaptive landscape was originally proposed by Sewall Wright (1931) as a model for the evolution of specific traits in a species.[1] Figure 1 shows an abstract example of such a surface in two dimensions; an actual adaptive surface would be highly N-dimensional space rather than a 2-dimensional. The dimensions correspond to two different characteristics contributing to the "fitness" of a species—the probability that individuals will pass on their genes.

Figure 1

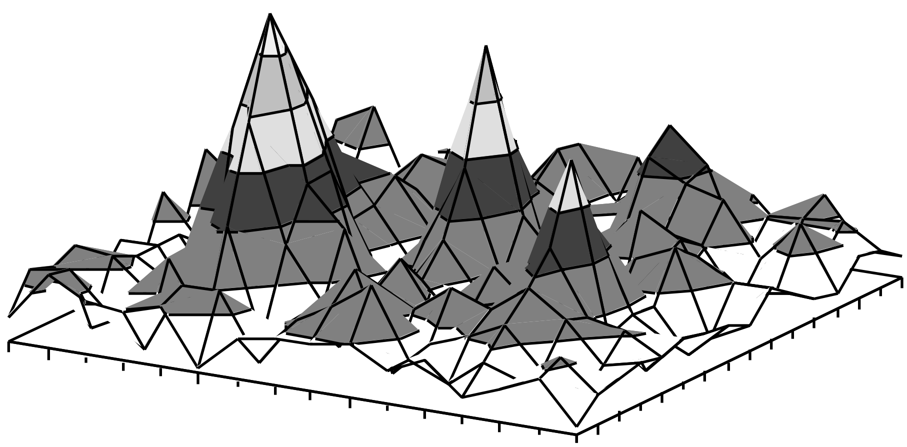

In birds, for example, two characteristics affecting fitness are the average number of eggs produced and the average egg weight. Producing too few eggs or eggs that are too small reduces the likelihood that the chicks will survive. Producing too many eggs or the eggs that are too large weakens the hen and reduces the likelihood that she can provide care for the chicks; this also adversely affects survival. Various other factors such as the availability of food, the amount of care provided by the

rules; see for example Sylvan, Goel and Chandrasekran (1990), Job and Johnson (1991) and Taber (1992), as well as the general use of rule-based models discussed in Sylvan and Chan (1984), Cimbala (1987) and Hudson (1991).

[1] Slatkin (1983) provides a quick introduction to these concepts as well as the concept of coevolution. Axelrod and Bennett (1993) use a landscape model to predict international and economic alliances; the state-space in their model is the set of all possible alliances and the landscape is determined by a measure of the mutual compatibility of the actors in the alliance. The Axelrod and Bennett model works well in predicting the international alliances of WWII and the economic alliances of computer firms attempting to establish a standard for the UNIX operating system.

father, and the presence of predators would favor some combinations of egg quantity and weight over others.

In biological evolution, movement on the surface occurs through random mutations. Under a fairly general set of conditions—R.A. Fisher's (1930) "fundamental theorem of natural selection"—the average fitness of a population will increase under the pressure of natural selection. On the adaptive surface, a species will, over successive generations, engage in "hill-climbing" along the surface due to mutation and selection, since individuals with mutations that improve fitness will produce, on average, more offspring than those who do not. The species stops evolving when it reaches the top of one of the local maxima on the surface, points at which there are no incremental mutations that can improve fitness. This maximum is local rather than global, and the species cannot further improve fitness because of its inability to cross the "adaptive valley" between the peaks.[1]

The application of this model to foreign policy adaptation is straightforward. For genes, substitute rules; for fitness substitute preference; for selective pressure substitute rational goal seeking based on rules.[2] The hill-climbing process occurs much more quickly and efficiently in organizations than in genetic evolution because foresight and planning, rather than random mutation, can be used to explore nearby points on the surface. However, as I will argue in more detail below, organizations are not omniscient about the landscape and are likely to have substantially less information about points that are distant from their current policy—in other words, policies that involve a large number of simultaneous changes—than about points that are close.

For example, the US at the end of WWII could have joined with the USSR in an anti-colonial coalition against the United Kingdom and France, ignoring European interests in order to gain the benefits of developing markets in the Third World. This scenario is not totally far-fetched—the policy emerged temporarily in the Suez Crisis—and arguably it would have been preferable to containment in avoiding the Cold War, the Vietnam War, and leading to earlier economic liberalization in China and southeast Asia. However, the payoff resulting from the required multiple simultaneous policy changes required for this coalition was difficult to evaluate in 1945. In contrast, the payoff from the incremental modifications required for a containment policy—active US

[1] Crossing will occur in two circumstances. First, the landscape itself could change, even temporarily. For example, there is substantial evidence that an asteroid collision with Earth at the end of the Cretaceous period temporarily modified the climate in a fashion that allowed small mammals to gain an advantage over the large reptiles that had dominated the planet for the previous 135-million years. This temporary change disrupted the adaptive surface for a sufficiently long period that mammals could move off the local maximum they are occupied during the Mesozoic to a position where they could ascend a much higher maximum in the Cenozoic.

Second, multiple simultaneous mutations can produce a "long jump" across a valley. The probability of such a set of mutations is very low but it is not zero. Evolution involves tens of billions of random "experiments" in the form of reproduction each year for millions of years, providing ample opportunity for low probability events to occur. When the extremely low probability crossing of a valley is followed by the high probability hill-climbing to a local maximum, this produces the sudden emergence of a new species. This is the pattern predicted by the theory of "punctuated equilibrium" in contrast to the evolutionary gradualism assumed by Darwin.

[2] Genes are simply rules for making proteins, so the biological model is actually a model of the adaptation of rules.

participation in the UN and Western European security, restoration of the pre-1941 anti-Soviet policies—seemed fairly clear. The US adopted the containment policy.

The landscape model largely consistent with Cyert and March's (1963) classic approach to organizational decision-making. Cyert and March assert that organizational behavior is adaptive, not optimizing:

> Without denying the substantial abilities of organizations as problem-solving and decision-making institutions, … a business firm is constrained by the uncertainty of its environment, the problems of maintaining a viable coalition, and the limitations on its capacity as a system for assembling, storing and utilizing information. As a result, the theory outlined in this volume characterizes the firm as an *adaptively rational* system rather than an *omnisciently rational* system. (pg. 99) [italics in original].

Cyert and March accept the existence of individual and organizational preference but note that organizations will find it difficult to optimally implement those preferences. The history of an organization is one of gradual improvement in an uncertain and changing world:

> Standard operating procedures are the memory of an organization. Like any other memory or learned behavior, [they] change over time at varying rates. Some rules seem to change frequently; others appear to have been substantially unchanged for as long as the organization has existed. (pg. 101)

The basic adaptive mechanism in Cyert and March is corrective feedback under satisficing norms. If the organization is achieving approximately what it desires, policy changes are marginal. Changes occur only when results fall short of minimal expectation due to changes in the external environment, stochastic "bad luck"[1] or the fact that system was not working well in the first place. Cyert and March note that policy changes often take place in the discrete realm of rules and heuristics, not in the realm of continuous variables, even for factors such as price and output that could be continuously adjusted.

The process of rule adaptation is dependent on the number and configuration of the local maxima, or "ruggedness", of the landscape. Kauffman (Kauffman and Levin 1987) has developed a general family of landscape models, the "NK model", where ruggedness can be systematically varied as a function of K.[2] In an NK model a species is assumed to have N genes. Each gene can have a value of 0 or 1; the evolution of a species is determined by randomly changing these values. The landscape is a fitness function $L:B^n \rightarrow [0,1]$, where B^n is the space of N-element binary vectors and [0,1] is the real interval from 0 to 1. K is the number of genes, $0 \leq K \leq N-1$, that interact epistatically (that is, can mask each other's effects) in jointly determining fitness. In the foreign policy context, K would reflect the

[1] For example one of Cyert and March's case studies involved a debate initiated by a fatal accident involving antiquated overhead crane controls in a factory. In the foreign policy arena, the explosion of the Mt. Pinatubo volcano in the Philippines substantially altered the U.S. military basing policy in the Pacific.

[2] The NK model is also discussed exhaustive detail in Kauffman (1993); this article was completed prior to the publication of that book and I therefore am using only the two earlier Kauffman articles, which cover the key features of the model.

extent to which changes in one set of rules (for example those involving the use of military force) affects the ability to use other rules (for example, those involving trade or diplomacy).

K determines the "correlation" of the landscape. If K=0, each gene affects fitness independently and has a unique correct value depending on whether 1 or 0 gives the higher fitness. Because the optimal value of each gene is independent of the values of the other genes, the landscape has no local maxima; its unique global maximum can be reached by independently finding correct value for each gene.

At the opposite extreme, K=N-1, the effect of each gene is related to the values of all of the other genes; any change modifies the effects of the entire configuration. In this case, the landscape is essentially random (uncorrelated) and there are a large number of local maxima. Hill-climbing in an uncorrelated landscape will usually find a local maximum that has a value substantially less than that of the global maximum. The lower levels of the surface shown in Figure 1, which were produced by a random number generator, illustrates the ruggedness of an uncorrelated landscape in two dimensions.

Co-Adaptation

The landscape model is easily extended to co-adapting entities by making the payoff a function of some of the rules of another entity, as well as the entities' own rules. For example, the effect on fitness of increasing the length of an animal's legs depends in part on the characteristics of its predators. If the increased leg length allows the animal to outrun a predator, it is more likely to reproduce so fitness increases. However, increased leg length could decrease fitness if it makes the animal more conspicuous without providing a significant speed advantage against predators. Long legs are useful against lions but not against hawks and eagles. Consequently, antelope lope; mice scurry. In the foreign policy arena, a significant increase in military power may force an opponent to back down and cease trying to challenging a state (e.g. Israel vis a vis Jordan after 1967), but in other circumstances military spending will simply lead to an expensive arms race and provide no strategic advantage (Israel vis a vis Egypt and Syria 1967-1973).

Formally, if one is looking at two organizations, O_1 and O_2 , the payoff functions are L_i:$B^{n+c} \rightarrow [0,1]$ where i=1,2 and C is the number of rules that are interdependent. An organization can adapt to the landscape by adjusting the values of the N rules it controls directly, but its payoff is also dependent on the C parameters of the competing organization:

Organization 1	n_{11}	n_{12}	n_{13}	n_{14}	n_{15}	n_{16}	n_{17}	c_{11}	c_{12}	c_{13}	c_{14}
	↓	↓	↓					↑	↑	↑	↑
Organization 2	c_{21}	c_{22}	c_{23}	n_{21}	n_{22}	n_{23}	n_{24}	n_{25}	n_{26}	n_{27}	n_{28}

Co-adaptation involves O_1 modifying n_{1j} in a hill-climbing fashion but the payoff depends on some of the n_{2j} of the other organization. This situation is similar to a two-person game: the row player controls the choice of row strategies but the payoff is dependent on the choices of both row and column.

Interdependence modifies the dynamics of adaptation considerably. Whereas the isolated organization simply finds a local maximum and stays there, a co-adapting organization may be forced to jump around on the landscape due to changes in c_{1j} induced by its opponent's changes in n_{2j}. This more complex process may nonetheless have an equilibrium, provided there are points where both organizations are at local maxima. Following the terminology of Maynard-Smith (Maynard Smith and Price 1973, Maynard Smith 1982), this is called a Nash equilibrium on the co-adaptive landscape, since it is a situation similar to a Nash equilibrium in game theory.[1] Howard's (1971) metagame approach, where a game is played over strategies rather than behaviors, is an example of this approach in the game theoretic paradigm. Metagame analysis, for example, shows that tit-for-tat is a Nash equilibrium among the set of strategies for Prisoners' Dilemma.

Kauffman and Levin (1987) observe that in an uncorrelated landscape with K>>0, there are a large number of local maxima and the likelihood of finding a Nash equilibrium is high, but the magnitude of these local maxima is relatively low. In such a situation, the organizations would usually find a set of mutually compatible rules without searching much of the landscape to find a high payoff.[2] Kauffman and Johnsen also find

> The waiting time to encounter Nash equilibria depends upon N, K, and C. For K>C, Nash equilibria are encountered rapidly, for K<C the waiting time to find Nash equilibria becomes very long. … The main point to note is that as K increases relative to C, the waiting time to hit a Nash equilibrium decreases. … Presumably this reflects the increased number of local optima in NK landscapes as K increases for a fixed N. (pg. 334)

L_1 and L_2 are not necessarily in a zero-sum relationship, so the co-adaptive Nash equilibrium need not be a classical zero-sum saddle point; similarly O_1 and O_2 are simply interested in maximizing their own payoffs, not minimizing the payoff of the opponent. During the Cold War, the defense industries in both the USA and USSR found substantial rewards in advocating increased defense spending irrespective of the threat from the other side, but these expenditures still had international implications.

Experimenting with a series of simulations, Kauffman and Johnsen show that for optimal fitness, it is necessary to "tune" K depending on C.

> When C is high, increasing K has two beneficial effects. First, Nash equilibria are encountered more rapidly… Second, fitness during [the search for equilibrium] is higher. This it is advantageous to any player to increase K in a high C environment. Perhaps equally remarkably, in the biologically reasonable case of random mutations, such a move by one species *also helps the second species*. Each has a higher pre-Nash fitness and finds Nash equilibrium sooner. [It

[1] A co-adaptive Nash equilibrium is local—there are no alternative sets of rules accessible from the current set without crossing an area of lower fitness, given the current rule choice by the opponent. In contrast, a Nash equilibrium in game theory places no restrictions on the proximity of other strategies.

[2] If rules are changed sequentially—a plausible assumption in organizational behavior—then it is possible to have cyclical sets of rules involving a local pattern similar to that found in a mixed-strategy game. If an organization knew this part of the landscape it could be expected to use the equivalent of a mixed strategy.

> therefore] seems clear that there are reasonable selective advantages to a species as a whole to "tune" K to match C (pg. 338; italics in the original)

In the foreign policy context, the value of C is determined by the extent to which the policies of an opponent directly affect one own policies. C cannot be controlled, but the value of K can be modified by adjusting the extent to which some sets of rules cancel out or modify the effects of other rules within the organization. The "turf battles" endemic to bureaucracies—where one part of the bureaucracy attempts to assert partial control over the activities of another—are attempts to raise the value of K; compartmentalization of policy lowers K.

Low K systems improve their performance very slowly, since rules must be changed one by one without synergistic effects. High K systems can be changed more rapidly because the change in one rule can affect a large number of other rules. For example, linking human rights to military aid affected a large number of military aid programs that otherwise would have needed to be debated one at a time. Ironically, however, high K systems also freeze into bureaucratic gridlock at low levels of utility compared to what a low K system can do in the long term. Carter used the human rights issue to achieve a high K; Reagan, unhappy about the perceived effects of the Carter policy in specific cases such as Nicaragua and Iran, lowered the impact of human rights on other policies.

The optimal value of K depends on C. If C=0, then if K=0 the organization will eventually find the configuration of maximum utility since a K=0 landscape has only a single, global, maximum. If C>0, this no longer works, because uncontrollable changes in c_{ij} can continually change the landscape. For a fixed c_{ij} and K=0 there is still a single maximum, but if c_{ij} are changing, this is irrelevant.

When C>0, having a high value of K gives the organization a wide variety of choices for movement from whatever point has been determined by the c_{ij}. If K=0, the landscape is correlated and the change in utility resulting from the change in any single rule is small, but if K is high, the landscape is less correlated and there is a higher probability that a point of improved utility can be reached with a small number of rule changes.

As K increases, however, the average height of these local maxima decrease. Therefore, to maximize utility, the value of K needs to be high enough to make it likely that points of higher utility can be reached with a small number of rule changes but not so high that the organization will get stuck on a local maximum with low utility. This optimal level of K puts the system into a state of "self-organized criticality"; this will be discussed in more detail below and used as the basis of an empirical test of the model.

The mechanism through which K is adjusted would be similar to that used on the rules themselves: incremental changes with the objective of improving utility. K is the product of "meta-rules"—rules that affect the operation of other rules—and changing K results in changing the distribution of local maxima rather than changing locations on the landscape, but with these changes of scale the same principles of maximization apply.[1] For example the management structure of United States Department of Defense underwent a fundamental reorganization with the National Security Act of 1947; major legislative amendments refined this structure in 1953 and 1958, and McNamara made a final set of incremental modifications in the early 1960s. At this point the Secretary of Defense/Joint Chiefs of Staff system remained largely fixed for about two decades (despite Vietnam, curiously), until it faced increasing calls for reform by the early 1980s, particularly after the coordination failures in Grenada and Lebanon. That reform movement culminated in the Goldwater-Nichols Reorganization Act in 1986. The meta-rule changes occur at a much slower rate than the basic rule changes—contrast for example the dramatic swings in Department of Defense personnel and budgeting rules during the 1965-1985 period—but still show a pattern of incremental improvement, stasis, and then dramatic change.

Information and Adaptation

In biological evolution, species move incrementally across the landscape through random mutations. Human organizations, in contrast, can plan and anticipate the consequences of some of their actions, and therefore would be expected to optimize on a landscape more quickly than DNA can. However, in most situations organizations possess relatively little information about the landscape beyond those areas they have experienced earlier and those points in the immediate vicinity of their current situation. Organizations do not have a "god's eye view" of their utility landscape, much less that of their competitors, and consequently they must search out equilibria by hill-climbing rather than by deducing a globally optimal equilibrium. Foresight and planning modify the *scale* of the landscape but not the underlying optimization process.

There are several reasons for this limitation. The first is the high dimensionality of the system. "Expert systems" that simulate organizational expertise typically contain about 1000 to 10,000 rules[2],; rules governing the detailed behavior of an organization might add another order of magnitude to that. The consequences of some rule changes are better known than those of others—for example one usually knows more about the preferences of one's allies than one's opponents (though as the Suez crisis illustrated, even this can be badly flawed).

Second, information is scarce in politics because there is no carrier of political information comparable to the role of prices in economic systems. In particular, political preferences cannot be directly observed, but can only be inferred through behavior. Even behavior is oftentimes of

[1] In other words, the process is nested: one can derive a landscape of meta-rules based on the local maxima of landscape of rules controlled by those meta-rules, and in a hierarchical political structure there may be several such levels of landscapes whose values are determined by other landscapes. This is how the system has evolved in biological systems, where the role of some genes is simply to control the expression of other genes.

[2] See for example discussions in Turban and Watkins (1988) and Forsyth (1989).

questionable value in inferring preferences because incentives exist to deceive. For example in a Chicken situation, a state trying to bluff an opponent into backing down has an interest in making an opponent think the state has greater resolve than it actually has.

Being good intuitive psychologists, politicians discount verbal pronouncements; because talk is cheap, talk will be deceptive. Instead, most exploration of the landscape is done through actual experimentation. But experimentation can be dangerous: having taken a convincing action, one must live with the consequences. There is a great deal of difference between the statement "If you don't back down, I will mobilize my troops" and the confirmed behavior "I have mobilized my troops".

Third, feedback concerning the effects of a policy is typically very slow, with a substantial delay between the time a rule has changed and the time the consequences of that change are known. The onset of World War I (WWI) provides a classical example of this problem. Clearly the rules of the European system in June 1914 were in a configuration with low global utility. If European leaders had known the consequences of that configuration—a war bringing the eventual destruction of the empires of all of its principal participants—it would have been rational to change the rules to move the system to a Pareto-superior configuration. Globally such configurations must have existed.[1]

European leaders, however, made any number of incorrect assumptions about the rules operating for other actors[2]. Based on those assumptions, the system was in a Nash equilibrium. Only after the *behaviors* of other actors revealed the true rules was their actual location on the landscape recognized. At this point, a cascading series of changes occurred that revealed additional information about the landscape, and based on this revealed information the core actors of the system to shifted to war. In contrast, in the Cuban Missile Crisis the revealed rules—and some movements in the rule space, such as the withdrawal of US missiles from Turkey—probably left the system in a much more stable configuration than it had been prior to the crisis.

An organization explores the landscape with an imperfect altimeter and a weak flashlight: One knows roughly one's location, one typically knows where one has been, and one can see the relative height of nearby points, but one does not know the complete surface. Since the implications of complex changes are generally more difficult to predict, this "light" becomes dimmer with distance.

In the absence of the god's eye view of the adaptive surface, organizations can still stumble into Nash equilibria, just as a species can evolve to higher genetic fitness through random mutations. The effect of conscious policy evaluation and foresight is primarily one of changing the *scale* of the landscape. In contrast to genetic evolution, organizations need not experiment with policy combinations that are clearly disadvantageous; the decisions of organizations focus on the difficult problems. If an

[1] The failure of the Hermann and Hermann (1967) attempt to simulate the outbreak of WWI in the INS simulation environment is further evidence of the unusual nature of the European rule system in July 1914.

[2] Van Evera (1986) and Lebow (1981) are among the *many* discussions of this.

organization is confident a change will lead to improved utility, it will be implemented quickly. Organizations will look for improved performance where the light is best (i.e. where they are confident of their assessments)—one observes this behavior frequently in intelligence organizations—and it is likely that the steeper the gradient of the surface, the faster the change. The time scale at which an organization changes is much more rapid than that of biological evolution so organizations probably spend relatively little time away from the local maxima in the landscape.

In short, in the absence of omniscience, conscious planning affects the micro-level characteristics of the landscape model but not its overall appropriateness. In order to engage in hill-climbing, organizations need only to know the values of points adjacent to their current location in the n_{ij} dimensions—in other words, the consequences of changes in their own rules.[1] In order *not* to engage in globally maximizing behavior—in other words, to have adaptive valleys between Nash points—organizations need only to be unaware of the consequences of *complex* changes in rules and to be uncertain about the landscapes of the organizations with which they are co-adapting.

Implications for International Behavior

The behaviors implied by the NK model of foreign policy adaptation provides an explanation for a number of regularities observed in international behavior. In this section I will discuss several of these in an informal fashion.

1. Minor modifications of rules have consequences on all scales; the system is both stable and unstable.

Kauffman and Johnsen argue that a system where K has been adjusted to an optimal value vis a vis C will be in a state of "self-organized criticality", a concept due to Bak (Bak, Tang and Weisenfeld 1988, Bak and Chen 1991):

> …many composite systems naturally evolve to a critical state in which a minor event starts a chain reaction that can affect any number of elements in the system. Although composite systems produce more minor events than catastrophes, chain reactions of all sizes are an integral part of the dynamics. According to the theory, the mechanism that leads to minor events is the same one that leads to major events. Furthermore, composite systems never reach equilibrium but instead evolve from one meta-stable state to the next. (Bak and Chen 1991: 46)

International politics consists of a large number of crises waiting to happen, and on a day-to-day basis, many small crises actually do happen. Some of these are absorbed by the system and produce no further consequences, for example the Iraqi attack on the USS Stark in 1987; others cascade into crises of substantial duration, for example the Iraqi attack on Kuwait in August, 1990; and a few cascade into events of major magnitude, such as the assassination of the Archduke Ferdinand in June 1914. A single system, however, produces events at all scales. The post-Bismarckian European system that resolved

[1] This mechanism requires only an ordinal landscape—the organization needs only to know whether it prefers one rule configuration to another rather than assigning a numerical value to the configuration.

the crises at Fashoda (1898) and Agadir (1911), and which responded with disregard to the First and Second Balkan Wars in 1912-13 was the same system, with the same mechanisms and largely the same individuals and preferences, that self-destructively collapsed in 1914. One rule change—the Austro-Hungarian decision to suppress Serbian nationalism in 1914—seems largely responsible for triggering the catastrophic series of changes (Lebow 1981:26-29).

2. Stable regimes are likely because landscapes are rugged.

A "regime" is a set of rules that are voluntarily adhered to for an extended period of time. In a rational choice perspective, the existence and establishment of regimes is a collective goods problem; collective goods solutions in a quasi-anarchic system should be unstable, and therefore a substantial theoretical literature exists trying to explain their presence (e.g., Oye 1986, Keohane 1986).

In an adaptive landscape, in contrast, regimes are not a puzzle but a near certainty. The complex interactions between foreign policy rules and between actors ensure that in international relations K and C are relatively high. A regime is simply a Nash point on the landscape, and when K+C>>0, the landscape is rugged and contains a large number of local maxima. Myopic self-interest and imperfect information actually contribute to the stability, though not the optimality, of such regimes.

Changes in regimes can come from two sources. First, the landscape itself can change; in the past three centuries, this probably has been primarily due to technological factors. Second, an internal political change in a state may modify its preferences and the external consequences of those changes can shift the system as a whole: the effects of the Nazi ascendancy in Germany would be an example.

The specific nature of a regime—the actual rules involved—will depend on the landscape and the history of the system. But the existence of *some* regime, some fixed set of rules adhered to for a long period of time, is an almost inevitable consequence of the ruggedness of the landscape. A rule-based interdependent system governed by self-interest and goal-seeking will, in general, spontaneously settle into a system of order, not anarchy.

These results can be extended to explain some of the characteristics of hegemonic stability. Because the rules of a hegemon impede upon the behaviors of the other states in the system, a hegemon increases the value of C. Kauffman and Johnsen do a series of simulations on the effects of raising C and find three general results.

First, for a fixed value of K, the system is more likely to converge on a Nash equilibrium as C increases. This is consistent with most hegemonic regime theories. and hegemonic regimes are stable for the reasons predicted by the NK landscape. The hegemon imposes and holds fixed a set of c_{Hj} that are a local maximum for the hegemon, and then the other actors in the system optimize within that constraint. The hegemon has effectively reduced the dimensionality of the landscape, allowing a Nash point to be found quickly. Also consistent with the theory, if the power of the hegemon declines—interpreted either as a shift in the hegemon's landscape or a decrease in C—one would expect to see a great deal of change in the system.

Second, Kauffman and Johnsen find that the average equilibrium value at the Nash point is lower in a high C system. Arguably, this is how many non-hegemonic states interpret the effects of hegemony. The prediction of lower utility is also consistent with the empirical observation that hegemonic behavior is expensive (at least in the sense of opportunity costs), though the mechanism generating that expense is quite different in the NK model than it is in rational choice theories.

Finally, Kauffman and Johnsen find "In general, but not uniformly, as C increases the optimal mutation rate decreases" (pg. 338), arguably also consistent with observed behavior. Hegemonic systems generally suppress change, and the likelihood of a state successfully changing the regime (in other words, finding an acceptable new set of rules) is lower than in a system without a hegemon.

3. In a stable environment, incremental strategies are prudent.

In an uncorrelated landscape (high K) Kauffman shows under fairly general conditions[1],

> Along an adaptive walk, the expected number of *fitter neighbors* decreases by *half* when each improved variant is found. That is, the *numbers of ways "uphill"* is cut in half at each improvement step (Kauffman 1988: 129; italics in original)

This, as well as the limited information about distant points, argues for incrementalism. Because adaptive systems tend to be near local maxima, most major changes in policy are likely to be deleterious to organizational performance. The exception occurs with minor, disconnected rules that might have been ignored earlier in the optimization, or responses to changes by other actors.

If organizations punish failure, the successful individuals in a bureaucracy will tend, over time, to be those characterized as "risk averse". This is not because bureaucracies uniquely attract risk averse individuals; it is rather that, being adaptive, bureaucracies create an environment that rewards risk aversion. In the words of a Chicago machine politician, to get ahead in the organization, "Don't make no waves; don't back no losers." (Rakove 1975)

An organization must retain some ability to innovate in order to cope with changes in the landscape or changes in the policies of other actors. However, the level of innovation must match the level of change in the system. Too much innovation in an essentially static system will simply result in a series of failed initiatives; the Carter foreign policy with respect to the Soviet Union provides an example of this. Too little innovation will cause policies to become outdated; the organization will "lose the initiative" and be continually away from local maxima due to more rapid changes by other

[1] Kauffman and Levin note that a correlated landscape with large step sizes (i.e. changes in multiple rules are allowed) is equivalent to an uncorrelated landscape with small step sizes. The former situation—correlation but large allowable steps—probably characterizes most co-adaptation situations in international politics.

organizations. The second terms of the Eisenhower and Reagan presidencies would be examples of this problem.[1]

4. Policy innovation is most likely to occur as a result of crisis.

While intentional movement by O_1 in the n_{1j} space is limited by the constraints of local maximization, movement induced by the modification of the c_{1j} rules is not, since these are local changes in n_{2j} for O_2. A shift in c_{1j}, rather than n_{1j}, can cause a transition to a very different point on the landscape.

This is a "crisis" in two respects. First, O_1 will rapidly modify its rules n_{1j} in order to hill-climb to a new local maximum. Second, O_1 will initiate experimental activity to determine the values of neighboring points in the landscape. These changes affect c_{2j} and can trigger a series of additional changes by O_2; this continues until the system locates a new Nash point.[2] Thus, consistent with Schelling's "crisis = danger + opportunity," a crisis involves a new location with a lower short term utility, but that location may provide a route to a point of higher utility. Crises also create a high information load and necessitate modification of existing standard operating procedures.

In international behavior, crisis adaptation tends to be sequential because O_1 does not know it is at a new point on the landscape until some behavior by O_2 alerts O_1 to the change in c_{1j}. For example, most evidence indicates that the Soviet Union abandoned the "Brezhnev Doctrine" with respect to intervention in Eastern Europe sometime in mid-1987. This change had substantial implications for Eastern European domestic politics, but the modification was not fully appreciated (or believed) until the summer of 1989, largely as the result of experimentation by individual Poles and East Germans. Once the rule change had been verified, Eastern Europe moved rapidly to a new equilibrium substantially different than the one it had occupied for the previous forty years.

5. Precedent is a plausible decision mechanism.

Because of the lack of information concerning the adaptive surface, precedent is a useful guide to decision making.[3] *Ceteris paribus*, an organization's history provides a relatively accurate map of the value of points on the surface visited in the past. Positive precedents (Marshall Plan; Lebanon 1958) show points that are desirable; negative precedents (e.g. Munich, Vietnam) show points to avoid. It is

[1] A more dramatic example of the problems of lack of innovation comes from the computer industry in the 1980s, where multi-billion-dollar bureaucratic giants such as IBM, DEC and Wang were upset by innovative start-ups such as Apple, Compaq, Sun and small Taiwanese computer companies.

[2] The extent to which a crisis sets off a series of changes depends entirely on which n_{1j} must be modified to find the local maximum. If these are uncoupled from the c_{2j}, the new Nash point is simply the nearest local maximum for O_1 and requires no additional changes in the rules of O_2. In other cases, the required modifications in n_{1j} could be tightly coupled to some c_{2j}, so one would see an avalanche of mutually adaptive behaviors.

[3] Formal models of precedent based reasoning are widely used in the computational modeling literature, for example Alker, Bennett and Mefford (1980) and Hudson (1991). In the policy literature, precedent is embodied in the "lessons of history" and case studies; Neustadt and May (1986) provide a recent discussion of this.

generally safer and less information intensive to use known points to set policy rather than trying to explore new policy combinations, just as it is unwise to explore an area characterized by steep cliffs and an unstable surface with a weak flashlight.

Precedent failures—for example the Bay of Pigs or Lebanon 1982-83—occur when one follows directions that were plausible at one point on the landscape but inappropriate for the point one actually occupies. Dead reckoning only works if one knows the starting point. Precedent and organizational memory may also restrain the ability to *search* the landscape because of the perceived dangers of experimentally evaluating certain points.

Empirical Evidence

The discussion thus far has been anecdotal. This section will explore statistical evidence that international behavior as measured by event data shows patterns consistent with those predicted by the NK landscape model. The tests are based on the assumption that a tuned NK system is in a state of self-organized criticality. Such systems exhibit two distinctive statistical characteristics. First, the marginal distribution of the magnitude of events follows a power law; second, the time-series of events has a distinctive "1/f spectrum". Since event data reflect behavior rather than rules, the test is not direct; one must instead assume, in this exploratory test, that changes in rules are consistently linked to observed behavior.

Data

Two data sets are used. The first is a WEIS-coded event data series for the Middle East covering 1982-1992 based on Reuters news service leads.[1] Each event was converted to a numerical value using Goldstein's (1992) which has -10 for the most conflictual event and +10 for the most cooperative; these values were totaled for each month. Eight directed dyads are examined: Israel→Palestinians, Palestinians→Israel , Israel→Lebanon, Syria→Lebanon, USA→Israel, USA→Syria, USA→Egypt, and Egypt→USA. The first four dyads show primarily conflictual behavior; the latter four show a combination of conflictual and cooperative behavior.

For comparison purposes, the analysis is also applied to the USA→USSR COPDAB data set (1948-1979, monthly aggregations) used in Goldstein and Freeman (1990). This series is completely distinct from the Middle East series—it involves a different dyad, different period of time, different scaling, different coding method, and different source—so it provides a relatively stringent test of the robustness of the results.

[1] This data set is described in Gerner et al. (1994) and Schrodt and Gerner (1994)

Power Law Distribution

In simulation experiments, Kauffman and Johnsen find that when the value of K is set to optimize the level of fitness of the species for a given value of C, the distribution of the size of the "avalanches" of change in the system approximates a power-law distribution

$$f(x) = cx^a$$

where f(x) is the number of occurrences of x. Power law distributions are quite common in natural phenomena; Schroeder (1991: chapter 4) provides a number of examples.

To test the fit of the power law, a histogram for the behavior of each dyad was created using ten histogram bins of equal size, $\frac{max_v - min_v}{11}$, where max_v and min_v are the maximum and minimum values found in the series. Figures 2 and 3 show the natural log of the frequency for each decile; these are aligned on the bin containing the maximum f(x); this contains the values with a magnitude around zero in each case.

If the power law holds, the plot of the log of f(x) with the log of the magnitude should be a line:

$$\ln(f(x)) = \ln(c) + a \ln(x)$$

As Figure 2—which shows the four dyads with conflict and cooperation— indicates, the ln(f(x)) plotted against x, rather than ln(x), is a line.[1] The slopes of these lines are almost identical across the various dyads except in the tails of the distributions; the slope of the line in the cooperative events is shallower than that for the conflictual. The consistency of the slope is quite surprising since the US relations with Israel, Syria and Egypt are quite different, but there is no obvious reason why this would occur as a statistical artifact. The separate slopes in the cooperative and conflictual events would be expected given earlier work on event data showing that conflict and cooperation tend to be two separate dimensions rather than a continuum.

The pattern for the conflictual dyads shown in Figure 3 is less consistent than that for the dyads in Figure 2, but the log of f(x) is still generally a linear function of x. The ISR→PAL and PAL→ISR lines have roughly the same slope, as do the ISR→LEB and SYR→LEB lines, but the slopes of those pairs is quite different. The conflictual tails of the distributions are quite discordant; this is probably a reflection of a few months of unusually intense violent activity in all four series.

Goldstein's scale is derived from expert rankings of the magnitude of various types of events and thus is somewhat arbitrary. If the Goldstein scaling reflects the log, rather than the magnitude, of behaviors that are generated by rule adjustments, then the power law behavior predicted by the model assuming self-organized criticality appears to hold quite strongly. This is a plausible interpretation of

[1] In Figure 2, the UAR→USA data contains one additional bin with a value 0.69 located two categories to the right of the graph.

the Goldstein scale: for example it maps criticism to -2.2, a demand to -4.9, a threat to -7.0, nonmilitary destruction to -8.7 and a military engagement to -10.0. Alternatively, the likelihood that an activity will be reported in the international media—and thus coded in an event data collection—may be function of the log of the number of changes underlying the activity.

Figure 2

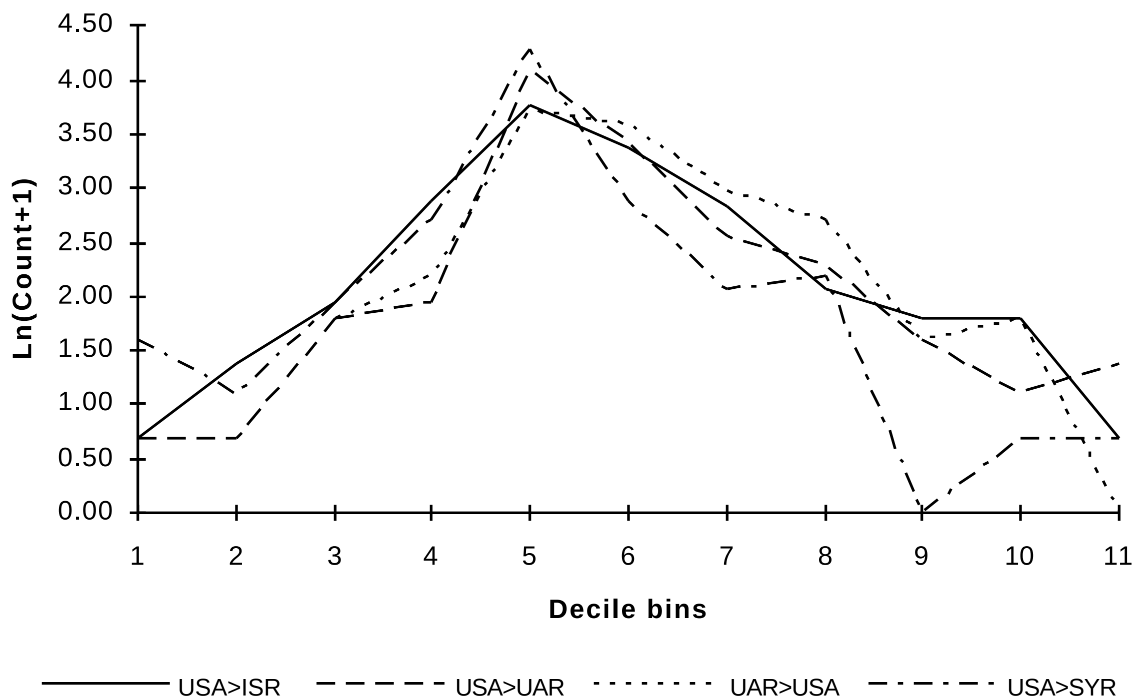

Figure 3

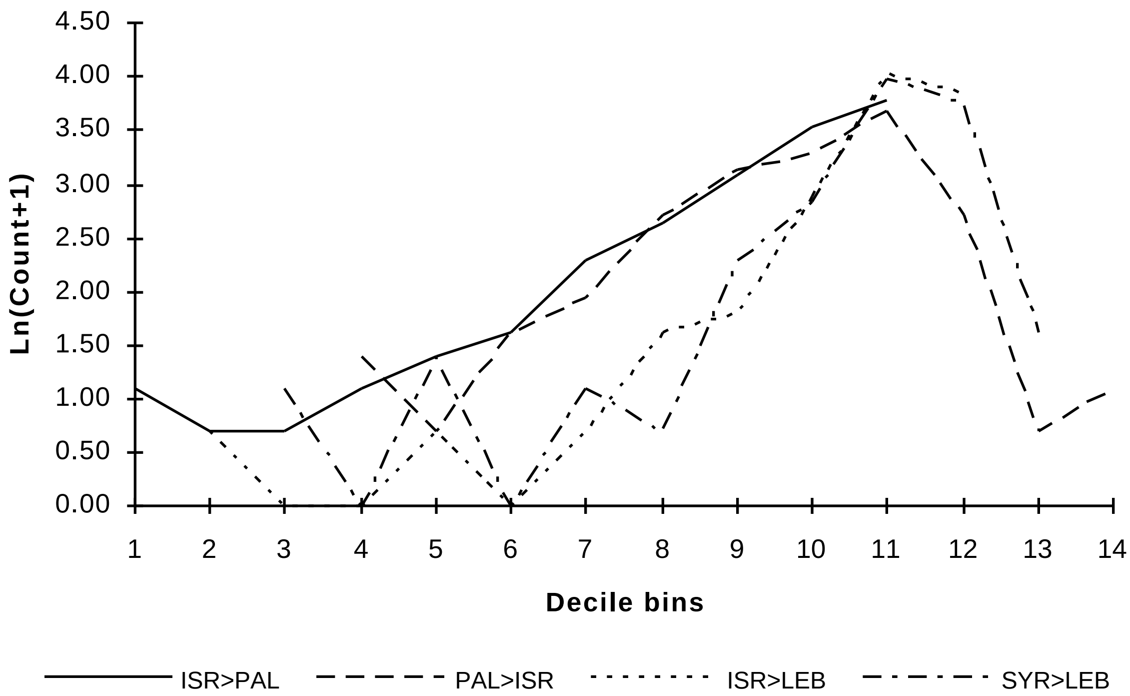

Further evidence of this empirical regularity can be found by comparing the Middle East data with COPDAB data on USA-USSR interactions. As Figure 4 shows, the COPDAB series matches the power law distribution of the non-conflictual Middle East dyads very closely, particularly on the conflict dimension. The solid line in Figure 4 is the average of the logged decile counts for all of the dyads except the purely conflictual ISR/PAL dyads. The dotted line shows the logged value, minus 1.03, of the decile bins of the USA→USSR data reported in COPDAB.[1] The lines on the conflictual behavior correspond almost exactly; as in the comparison among the Middle East dyads, the cooperative behavior of the USA→USSR dyad corresponds less well. This would also suggest that the conflictual events in COPDAB are scaled in a manner comparable to those of the Goldstein scale; the cooperative events may be scaled differently, though still as a log of the true magnitude.

[1] The COPDAB series contains 372 points; 1.03=LN(372/132) and adjusts for the larger magnitude of the counts in the COPDAB series.

Figure 4

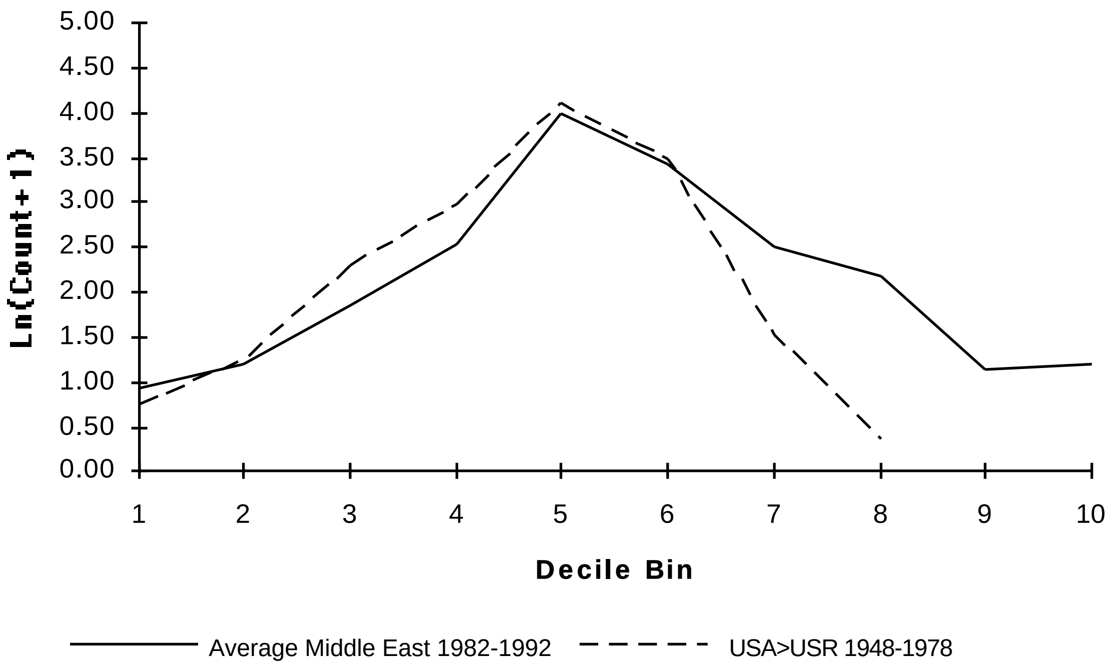

Frequency

Systems exhibiting self-organized criticality have a distinctive characteristic called 1/f noise:

> If one graphed the [behavior of a self-organized critical system] versus time, one would see a very erratic signal that has features of all durations. Such signals are known as flicker noise, or 1/f noise. Scientists have long known that flicker noise suggests that the dynamics of a system are strongly influenced by past events. In contrast white noise, a random signal, implies no correlation between the current dynamics and past events.
>
> Flicker noise is extremely common in nature. ... Indeed, the ubiquitousness of flicker noise is one of the great mysteries of physics. The theory of self-organized criticality suggests a rather general interpretation: flicker noise is a superposition of signals of all sizes and durations — signals produced when a dynamic system in the critical state produces chain reactions of all sizes and durations. (Bak and Chen: 48)

1/f noise has a distinctive power spectrum where the power at each frequency is proportional to the inverse of the frequency. Schrodt (1991) found that the power spectrum of USA→USSR interactions in the COPDAB series had such a spectrum.

For the Middle East dyads, the periodogram—which is proportional to the power spectrum—for each time series was estimated using the Fast Fourier transform (FFT) in SYSTAT 5.1. The FFT works with series whose length is a power of 2 and therefore the 132-point monthly series was truncated slightly to 128 points. The FFT produces estimates of the phase and magnitude of cycles with frequencies ranging from 1/N to 1/2 in steps of 1/N, where N=number of points in the time series. The frequency is the inverse of the cycle length so, for example, in the Middle East series, where the data points are monthly, a frequency of 0.015 (= 2/128) corresponds to a cycle of 64-months, whereas a frequency of 1/2 corresponds to a cycle of 2 months[1]. The periodograms have been smoothed using a four-frequency (0.031) moving average.

Figures 5, 7 and 8 give the periodograms for the eight dyads; these show three general patterns. The three conflictual dyads show the predicted 1/f pattern. The anomalously high power centered on the frequency 0.083 corresponds to a twelve-month cycle, which is unsurprising given the seasonal nature of military conflict and the propensity of Israeli-Palestinian violence to escalate on anniversaries of political events. Most of the other "bumps" in the power curve correspond to harmonics of the twelve-month cycle—these are centered on 6 months (0.16), 4 months (0.25), 3 months (0.33) and 2.4 months (0.42). Even accounting for the harmonics, the PAL→ISR and ISR→LEB periodograms show more power in the high frequencies than the 1/f model would predict.

Figure 6 compares the periodogram of the USA→USSR series with the periodogram of the ISR→PAL dyad. As in Figure 4, the two lines coincide quite closely.[2] The ISR→PAL and USA→USSR dyads are both stable protracted conflicts, albeit with quite different levels of violence, and that stability may generate the 1/f behavior.

Figure 7—the periodogram for USA→UAR, UAR→USA and USA→SYR—is approximately flat, with substantial power in the higher frequencies (note that the vertical scale of Figures 7 and 8 is half that of Figure 5). Figure 7 shows some seasonality—the USA→UAR interactions have a strong peak at 12 months and UAR→USA at 6 months—but unlike Figure 5, the power in the higher frequencies is not concentrated on the harmonics of the annual cycle.

[1] See Gottman (1981) chapters 15-17 for a discussion. The periodogram is not a standard measure: for example "periodogram" reported in SYSTAT's printout is 2N*amplitude^2, though in the discussion the SYSTAT manual simply uses amplitude^2; Gottman (205) also suggests dividing this measure by N to measure "power". I'm using amplitude^2 standardized by the total power in order to compare the various distributions.

[2] Both dyads show a bump centered on 0.15 cycles/month, which corresponds to a 6-month cycle. This is probably just a harmonic from the annual cycle with the annual cycle itself being indistinguishable from the high power in the low frequencies in the 1/f pattern.

Figure 5

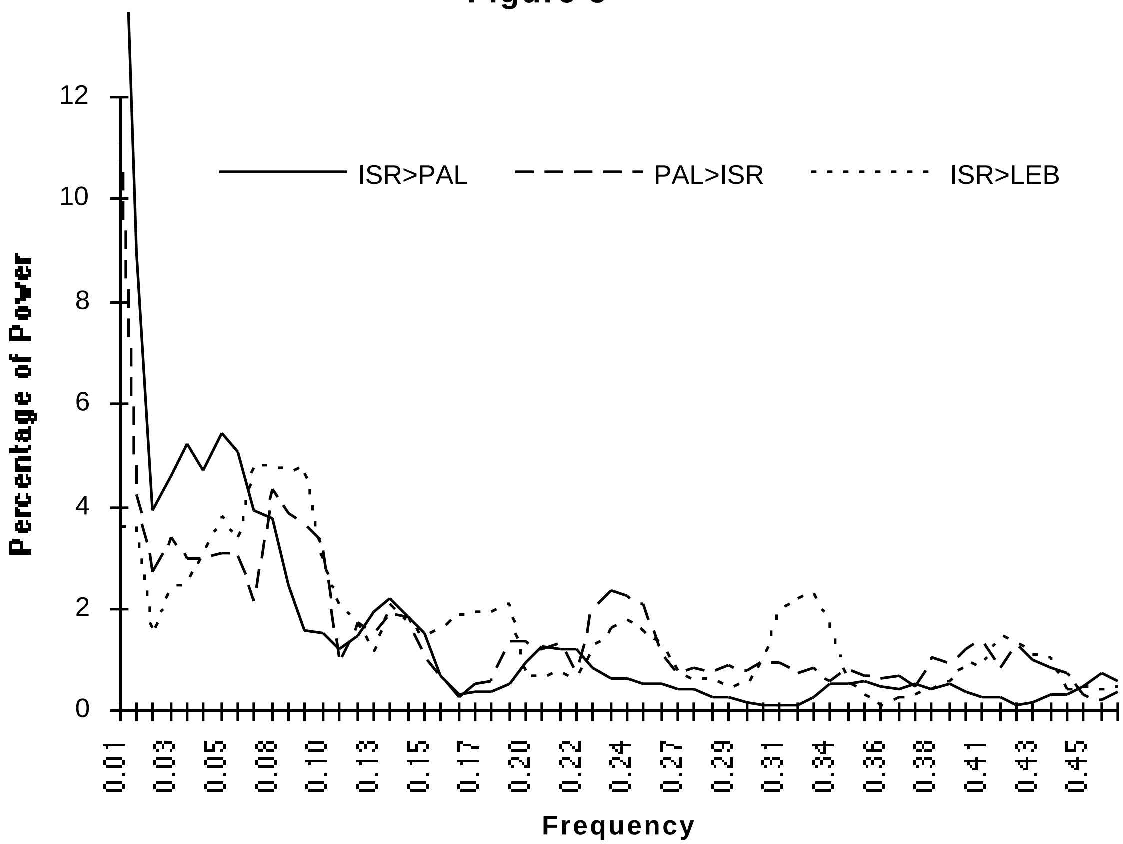

Figure 6

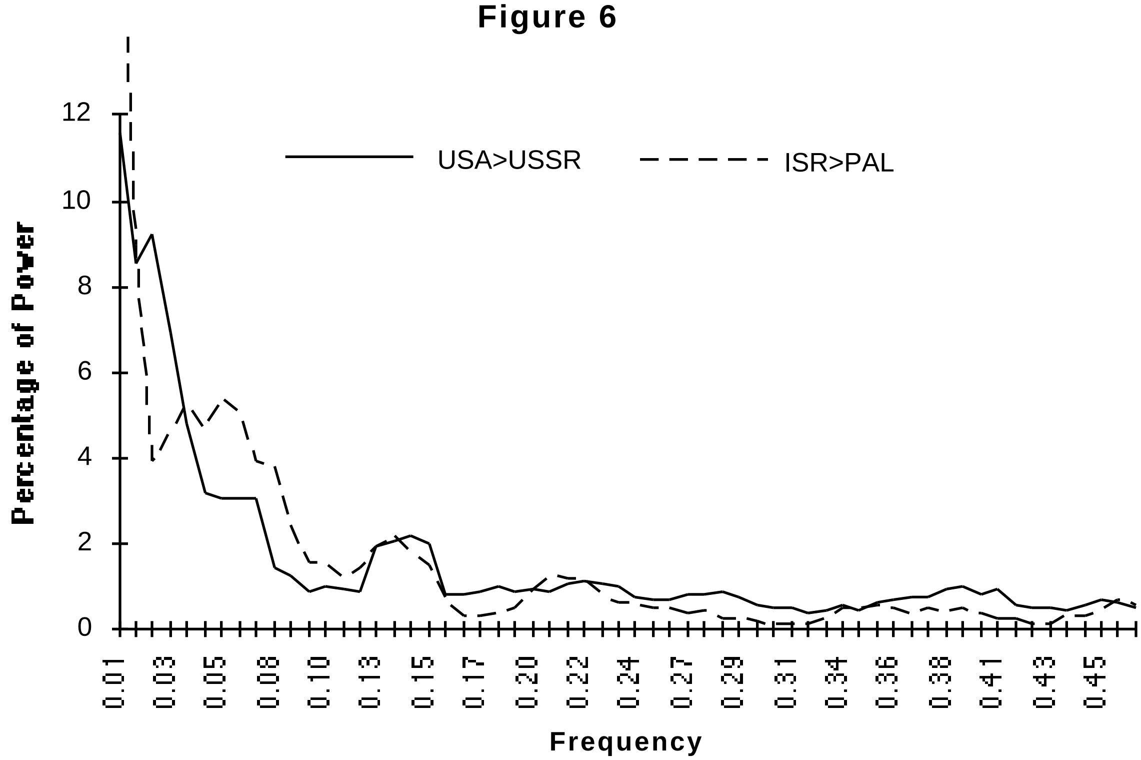

Figure 7

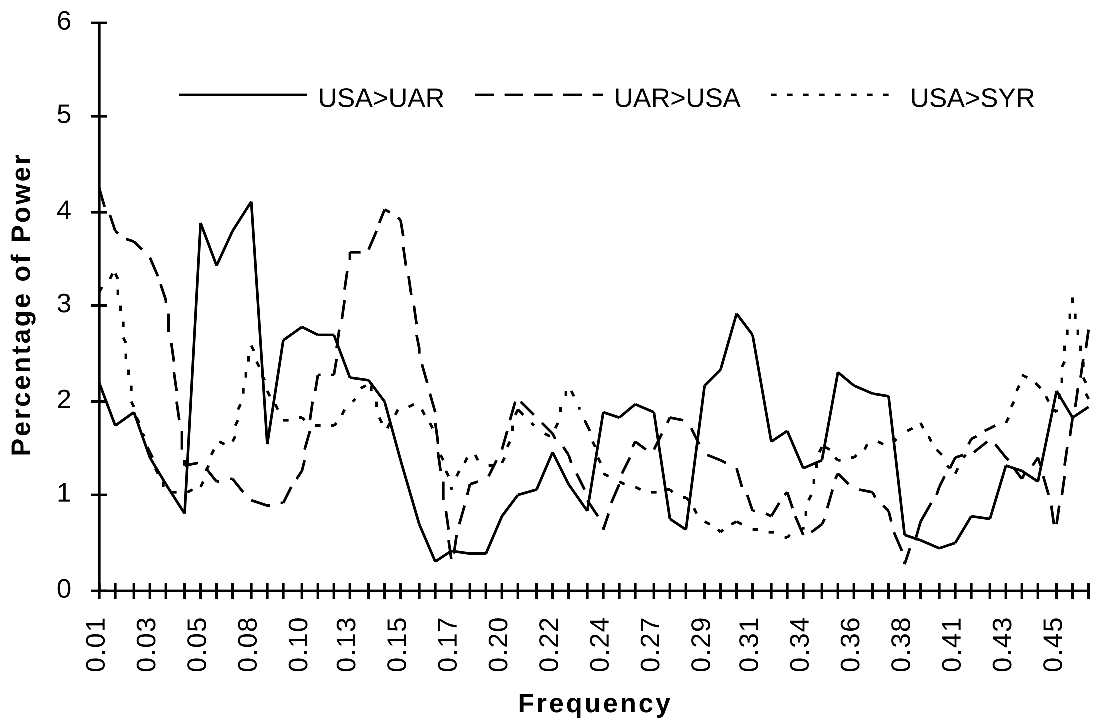

Figure 8

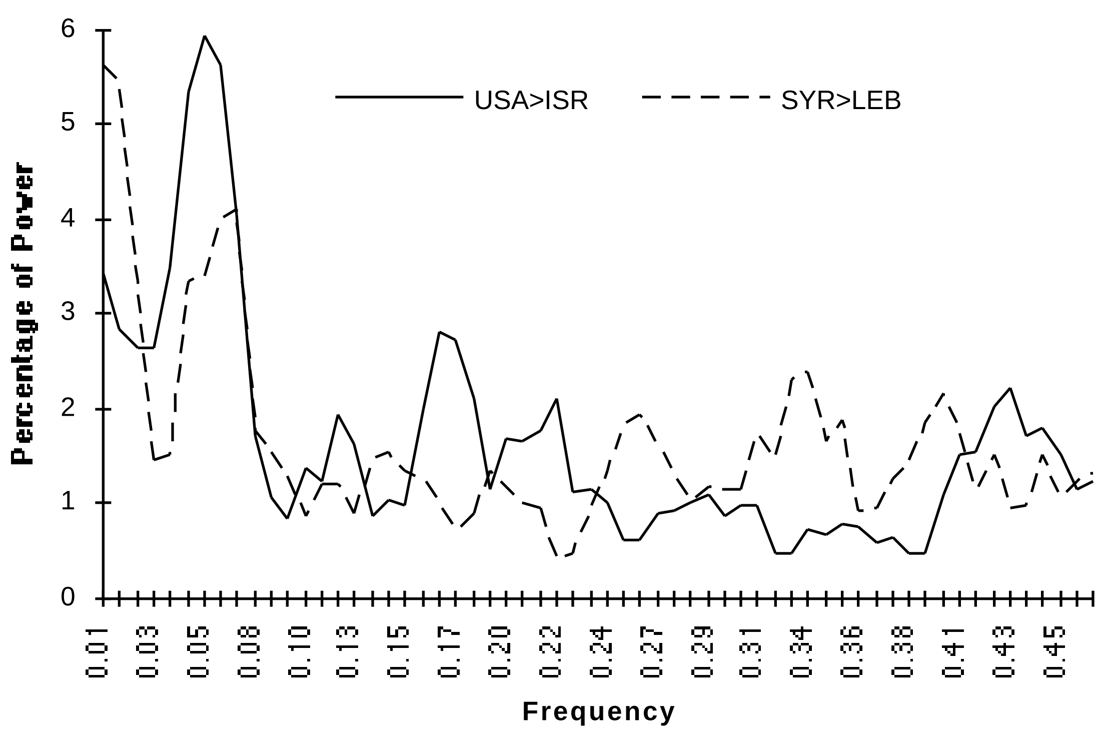

Figure 8—USA → ISR and SYR → LEB —shows a pattern intermediate between the other two. It has much more power in the high frequencies than the 1/f pattern in Figure 5, but has a stronger peak in the low frequencies than Figure 7. Both of these dyads are close alliances, but both have also experienced considerable conflict (verbal in the case of USA → ISR; violent in the case of SYR → LEB). The USR → ISR periodogram also has a very strong seasonal component; this is probably a result of US criticism of Israeli actions against Palestinians, those actions being seasonal.

If the eight periodograms are smoothed further and juxtaposed in Figure 9—which gives the percentage power in each of six equal sized frequency bins—a rough pattern emerges. The dyads are arrayed in a rough order going from the most conflictual (ISR → PAL) to the most cooperative (UAR → USA). On this dimension, there is a general tendency of the periodograms to shift from emphasizing low frequencies to high frequencies, and the distinctive 1/f pattern of the ISR → PAL interactions moves to the relatively flat periodograms of the USA/UAR dyad. This suggests that in event data, conflictual events may follow a long-term dynamic while cooperative events follow a short-term dynamic.

Figure 9

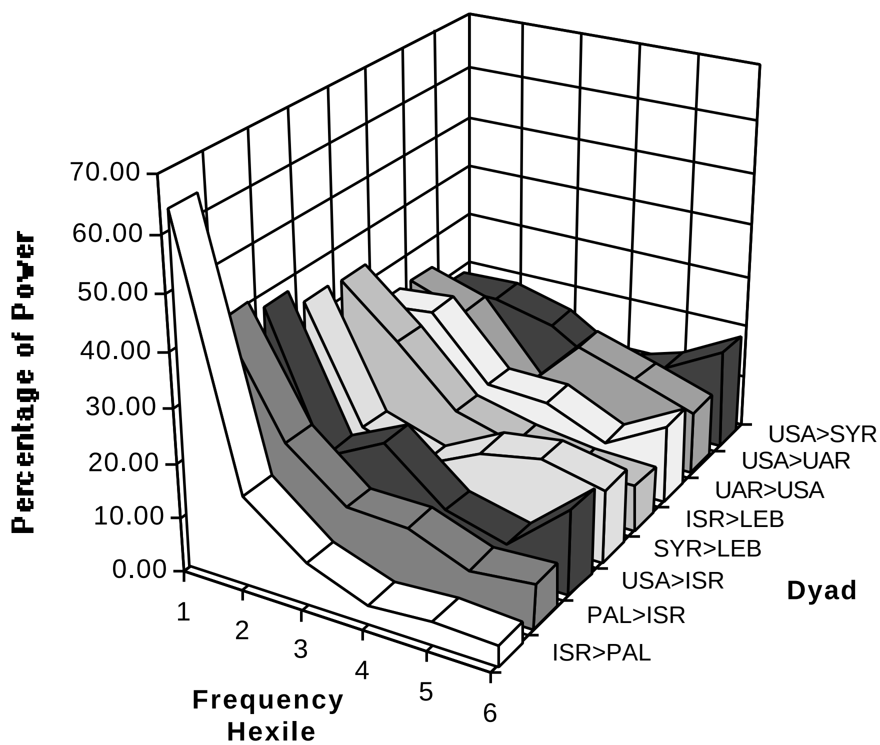

Discussion

The power law analysis shows that the pattern of international behavior measured by event data is generally consistent with that expected from a system in a state of self-organized criticality. Self-organized criticality is not the only explanation for this regularity, since a variety of different processes can produce power law distributions. and the test is only suggestive, not definitive, because the power law distribution under self-organized criticality applies to *rule* changes while event data measure *behaviors*. Still, some general process seems to be generating a remarkably regular statistical pattern in a number of disparate dyads; this pattern is evident in event data produced using two different, though intentionally comparable, event data sets.

The spectral analysis shows a weaker correspondence with the predictions of the theory. The predicted 1/f pattern is clearly evident in the dyads involved in protracted conflict, but the 1/f pattern is not found in situations that combine conflict and cooperation. The periodograms, however, show a fairly smooth transition from 1/f behavior to a relatively flat behavior as one moves from conflictual to cooperative dyads. This transition is not predicted by the theory but might be explained by the differences in reporting of conflictual and cooperative events. Conflictual events are usually reported in the international media; routine cooperative events are not. When the United States delivers foreign aid to Israel, this does not make the news, but if the United States suspends aid to Israel (or threatens this action), it is widely reported.

Conclusion

The complex systems found in political behavior have multiple equilibria and can shift between them. For example, the centuries of nearly continuous internal war in the Muromachi and Momoyama periods in Japan were followed by 265 years of total political stability under the Tokugawa Shoganate. Both patterns occurred under substantially the same constraints of resources, technology, population, religion and culture. An observer in 1500 would have concluded that Japan was doomed to perpetual political chaos and samurai warfare; an observer in 1700 would have concluded that Japan was a system of intrinsic political stability. Neither generalization is correct: the conditions in Japan could accommodate either configuration.

The predictions of the co-adaptive landscape model generally correspond to observed characteristics of international behavior. Most of the assumptions of the model are consistent with theories of rational choice; the remainder are consistent with theories of bounded rationality and make no demands for perfect information, super-rationality or expected utility calculations.

The landscape model scales easily: there is no particular reason that the organization generating c_{ij} needs to be another nation-state; it could as easily be another bureaucracy within the nation-state. Given the hierarchical structure of political control, lower levels of a bureaucracy control only a few rules, but they still control some rules. In a landscape perspective, individuals high in a bureaucracy determine the general location on a surface, then individuals further down optimize within those constraints. Depending on the landscape, low-level optimization may go quite far in reversing the impact of the rules higher up. In Vietnam in the 1970s, optimizing behavior at low levels in the system (for example drug use and "fragging" by enlisted personnel) substantially restricted the options available to decision-makers in Washington and Saigon; the Soviet Union encountered similar problems in Afghanistan. Thus the classical problem of policy implementation it can also be seen as a situation of co-adaptation.

The obvious weakness of the landscape model is its dependence on unobservables. The landscape surface, like the individual preference functions of rational choice theory, is not directly observable, though it is partially revealed through behavior. In addition many latent rules are unobservable or at least difficult to observe. Had Britain supported Germany in WWI, or the USA invaded Cuba during the 1962 crisis, rules would have been revealed. The historical record gradually, if imperfectly, reveals some of those, but never as clearly as the rules revealed by actual historical events.

This analysis has been ambiguous in the distinguishing between rules and behavior. Some actions are themselves changes in rules—for example the war/peace transition—while others are the *consequences* of rules. In fact, most of the behaviors in the international system are the consequences of existing rules rather than changes in the rules themselves, though these routine events generally are not recorded in event data sets. The rule/behavior dichotomy itself is probably too simple, as one can distinguish at least three levels of interactions: signals, actions and rules, with signals being actions taken for their information value alone. While the existing event data sets such as WEIS and COPDAB are not designed with this distinction in mind, it may be possible to construct new event coding systems that would make such distinctions and thereby enable more direct tests of the model.

Finally, some characteristics of the landscape model might provide insights relevant to crisis forecasting. Knowledge of the full landscape itself is inaccessible—the dimensionality is too high and too many rules are latent. However, by observing the interdependent reactions by actors in the international system to events, one might gain some insight as to how changes could proliferate through the system. For example, Schrodt and Mintz (1988) used a conditional probability analysis of relations in the Middle East to identify Kuwait as a point of high sensitivity in the region, an observation that seemed counterintuitive at the time; the VAR technique used by Goldstein and Freeman (1990) to analyze super-power relations emphasizes the study of how changes in the behavior of one actor trigger changes in other actors. A more sophisticated model of the dynamics of co-adaptation, particularly in a state of self-organized criticality, might provide insights to refine these statistical forecasting models.

Bibliography

Alker, Hayward R., James Bennett and Dwain Mefford (1980) "Generalized Precedent Logics for Resolving Security Dilemmas", *International Interactions* 7:165-200.

Axelrod, Robert (1984) *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, Robert and D. Scott Bennett (1993) "A Landscape Theory of Aggregation", *British Journal of Political Science* 23: 211-233.

Azar, Edward E. (1982) *The Codebook of the Conflict and Peace Data Bank (COPDAB).* College Park, MD: Center for International Development, University of Maryland.

Bak, Per and Kan Chen (1991) "Self-Organized Criticality", *Scientific American* (January, 1991): 46-53

Bak, Per, Chao Tang and Kurt Weisenfeld (1988) "Self-organized criticality", *Physical Review A* 38,1:364-374.

Cimbala, Stephen (1987) *Artificial Intelligence and National Security.* Lexington, MA: Lexington Books.

Cyert, Richard M. and James G. March (1963) *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice-Hall.

Fisher, R. A. (1930) *The Genetical Theory of Natural Selection.* Oxford: Claredon Press.

Forsyth, Richard. (1989). *Expert Systems: Principles and Case Studies.* New York: Chapman-Hall.

Gerner, Deborah J., Philip A. Schrodt, Ronald A. Francisco, and Judith L. Weddle. (1994) The Machine Coding of Events from Regional and International Sources, *International Studies Quarterly* 38:91-119.

Goldstein, Joshua S. and John R. Freeman (1990) *Three-Way Street: Strategic Reciprocity in World Politics*. Chicago: University of Chicago Press.

Goldstein, Joshua S. (1992) "A Conflict-Cooperation Scale for WEIS Events Data", *Journal of Conflict Resolution* 36,2:369-385.

Hermann, Charles F. and Margaret G. Hermann (1967) An Attempt to Simulate the Outbreak of World War I. *American Political Science Review* 61:400-414.

Hogarth, Robin M. and Melvin W. Reder, eds. (1987) *Rational Choice: The Contrast between Economics and Psychology*. Chicago: University of Chicago Press.

Howard, Nigel (1971) *Paradoxes of Rationality.* Cambridge: MIT Press.

Hudson, Valerie, ed. (1991) *Artificial Intelligence and International Politics*. Boulder: Westview

Janis, Irving and Leon Mann (1977) *Decision Making: A Psychological Analysis of Conflict, Choice and Commitment.* New York: Free Press.

Jervis, Robert (1976) *Perception and Misperception in International Politics.* Princeton: Princeton University Press.

Job, Brian L. and Douglas Johnson (1991) "UNCLESAM: The Application of a Rule-Based Model to US Foreign Policy" in Valerie Hudson, ed. *Artificial Intelligence and International Politics*. Boulder: Westview.

Kahneman, Daniel, Paul Slovic and Amos Tversky (1982) *Judgement Under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

Kauffman, Stuart A. and Simon Levin (1987) "Towards a General Theory of Adaptive Walks on Rugged Landscapes", *Journal of Theoretical Biology* 128: 11-45.

Kauffman, Stuart A. (1988) "The Evolution of Economic Webs", pp. 125-146 in Philip W. Anderson, Kenneth J. Arrow and David Pines, eds. *The Economy as an Evolving Complex System*. New York: Addison Wesley.

Kauffman, Stuart A. and Sonke Johnsen (1992) "Co-Evolution to the Edge of Chaos: Coupled Fitness Landscapes, Poised States and Co-Evolutionary Avalanches" in Langton et al. *Artificial Life II*. Redwood City, CA: Addison-Wesley: pp. 325-371.

Kauffman, Stuart A. (1993). *The Origins of Order*. Oxford: Oxford University Press.

Keohane, Robert O. ed. (1986) *Neorealism and Its Critics*. New York: Columbia University Press.

Lebow, Richard Ned (1981) *Between Peace and War: The Nature of International Crises*. Baltimore: Johns Hopkins.

Maynard-Smith, J. and G.R. Price (1973) "The Logic of Animal Conflict", *Nature* 246:16-16

Maynard-Smith, John (1982) *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

McClelland, Charles A. (1976) *World Event/Interaction Survey Codebook*. (ICPSR 5211). Ann Arbor: Inter-University Consortium for Political and Social Research.

Most, Benjamin A. and Harvey Starr (1984) "International relations theory, foreign policy substitutability and 'nice' laws", *World Politics* 36: 383-406

Neustadt, Richard E. and Ernest R. May (1986) *Thinking in Time: The Uses of History for Decision Makers*. New York: Free Press.

Nicholson, Michael (1992) *Rationality and the analysis of international conflict*. Cambridge: Cambridge University Press.

Ordeshook, Peter C., ed. (1989) *Models of Strategic Choice*. Ann Arbor: University of Michigan Press.

Oye, Kenneth, ed. (1986) *Cooperation Under Anarchy*. Princeton: Princeton University Press.

Rakove, Milton (1975) *Don't Make No Waves, Don't Back No Losers: An Insider's Analysis of the Daley Machine*. Bloomington: Indiana University Press.

Schrodt, Philip A. (1991) "Feedforward and Adaptive Rationality in Political Behavior." Paper presented at the American Political Science Association, Washington.

Schrodt, Philip A., and Alex Mintz (1988) "A Conditional Probability Analysis of Regional Interactions in the Middle East", *American Journal of Political Science*. 32: 217-230.

Schrodt, Philip A., and Deborah J. Gerner. (1994) "Statistical Patterns in a Dense Event Data Set for the Middle East, 1982-1992", *American Journal of Political Science* 38.

Schroeder, Manfred. (1991) *Fractals, Chaos and Power Laws*. New York: W.H. Freeman.

Simon, Herbert A. (1982) *The Sciences of the Artificial* (2nd ed.) Cambridge: MIT Press

Slatkin, Montgomery (1983) "Genetic Background", in Douglas J. Futuyma and Montgomery Slatkin, eds. *Coevolution*. Sunderland, MA: Sinauer Associates.

Sylvan, Donald A. and Steve Chan (1984) *Foreign Policy Decision Making: Perception, Cognition and Artificial Intelligence*. New York: Praeger.

Sylvan, Donald A., Ashok Goel and B. Chandrasekran (1990) "Analyzing Political Decision Making from an Information Processing Perspective: JESSE." *American Journal of Political Science* 34:74-123.

Taber, Charles S. (1992) "POLI: An Expert System Model of U.S. Foreign Policy Belief Systems", *American Political Science Review* 86:888-904.

Turban, Efraim and Paul R. Watkins. (1988) *Applied Expert Systems*. New York: North-Hall.

Van Evera, Stephen (1986) "Why Cooperation Failed in 1914", in Kenneth Oye, ed. *Cooperation Under Anarchy*. Princeton: Princeton University Press.

Vertzberger, Yaacov Y.I. (1990) *The World in their Minds: Information Processing, Cognition and Perception in Foreign Policy Decision Making.* Stanford: Stanford University Press.

Wright, Sewall (1931) "Evolution in Mendalian populations", *Genetics* 16:97-159.

Zagare, Frank C., ed. (1990) *Modelling International Conflict*. Special issue of *International Interactions* 15,3/4.