UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

PREDICTING WAR, PREDICTING PEACE

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

WILLIAM EDWARD HANSON
Norman, Oklahoma
2009

PREDICTING WAR, PREDICTING PEACE

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF POLITICAL SCIENCE

BY

Dr. Suzette Grillot, Chair

Dr. Thomas James

Dr. Gregory Miller

Dr. Jozef Raadschelders

Dr. Rickey Thomas

This dissertation is dedicated to my mother,

who has been and will always be my hero in life.

Acknowledgments

No work of this magnitude gets completed without a great deal of help, and this dissertation is no exception. First and foremost, I am deeply grateful to Professor Suzette Grillot, who was both my PhD and dissertation chair from start to finish, and provided invaluable guidance and counsel throughout. Special thanks goes to Professor Rick Thomas, who gave me my first formal introduction into the judgment and decision-making literature, introduced me to memory-based models, and got me started with Mathematica. I could never have gotten the project off the ground without his help. My other committee members, Professors Jos Raadschelders, Tom James, and Greg Miller all gave me much needed help and encouragement along the way and I owe them all a debt of gratitude. I would also like to thank Professor Charles Taber, SUNY Stony Brook, for his comments and suggestions on an early version of this work.

Finally, and most important, my lovely and talented wife Jean kept me going through the long push with love, encouragement, and the occasional piece of chocolate. I could not have finished without her.

All of these people, along with my classmates, friends, and colleagues at the University of Oklahoma contributed greatly to the success of this dissertation. Any errors are solely mine.

Table of Contents

List of Tables

List of Figures

Abstract

This dissertation examines how the structure of a given situation and the way in which decision makers mentally process information interact to determine decisions. Decision are affected by the surrounding ecology, information stored in memory, and the hypotheses chosen to access information in the surrounding ecology and memory. This theory draws on and extends existing cognitive simulation models to show that judgments, along with observed heuristics, biases and affect can be predicted probabilistically using models of memory processes. A Mathematica multiple-trace simulation of cognitive processes of memory and hypothesis generation which learns and reacts to various opponents: Memory-Based Ecological Model of Relationship Interactions between States, version 2 (MEMORIS-2) is tested in both game-theoretic settings and in a simplified multi-player system of international interaction.   Results show that the memory-trace simulation results in equal or better outcomes in both static (PD) and shifting-payoff (alternating PD and Stag Hunt) game theoretic situations than previously-developed strategies and heuristics (e.g., Tit-For-Tat, Grim Trigger, etc.). Analysis of the multi-player international interaction game shows that results show specific reactions to increasing threats and break into definite clusters of results which could be used to probabilistically predict state behavior. The simulation is then used to model two cases: Europe during the period 1885-1915 and the Cold War between 1945 and 1980, with results being compared to the historical record. Results show that the MEMORIS-2 simulation has the ability to probabilistically predict emergent properties in the international system, such as the likelihood of conflict, the players in that conflict, and the structure of alliances.

Preface

This paper is a step on a long professional and intellectual journey. When I graduated from the Air Force Academy as a newly-minted lieutenant more than 30 years ago, I did not think very much about the process of decision making – it was just something you did. The decisions of the members of my military chain of command were simply givens. My job was, at the time, to try to be the best B-52 pilot that I could be, and to be prepared to carry out the decisions of the "National Command Authority". In practice, I spent one week out of every three sitting with a plane loaded with nuclear weapons, ready to carry out the decisions of my nation.

The long hours spent in anticipation of an order that (thankfully) never came gave me ample time to consider the logic of those decisions. In particular, the question of just how the President could come to the time and place where he would actually consider unleashing the many thousands of nuclear weapons in the US arsenal, and what thought process would go into that decision preoccupied me. How would he go about deciding an action which would likely result in unimaginable death and destruction, if not the very end of life on earth?

About that time, one of my commanders told me about a program where I could learn to do something called "studies, analysis and gaming". I wasn't too sure exactly what that was, but it sounded interesting. So I called the director of the program and told him I was interested, and was (quite frankly) surprised to be notified a few weeks later that I had been admitted to the "Strategic and Tactical Sciences" program at the Air Force Institute of Technology, where I was to get a MS degree in something called "operations research". There, I learned for the first time the theology of nuclear deterrence, learned to model the effects of nuclear weapons, and how to optimally assign those weapons to targets.

Armed with this knowledge, I soon found myself immersed in the process of planning and developing nuclear deterrent forces, first at Strategic Air Command and later in the Pentagon. I had developed the ability and computer modeling skills to calculate the results of nuclear warfare to three significant decimal places, and my skills were in great demand to support the Pentagon's decision-making process.

Yet there was something about the whole process that I found puzzling. My superiors would often put great credence in small changes in the model results, which even though they could spit out numbers of arbitrary precision, were in reality approximations stacked upon other approximations. At other times, results of the analysis process simply would be ignored. As I returned to the Pentagon on two subsequent assignments, that puzzle grew: analysis was sometimes used and sometimes ignored, consensus products of the staffing process were sometimes used and sometimes ignored, different sets of factors seemed to shape the decision each time – why? This puzzle has seized me for my entire professional career.

At the end of my military career, that puzzle still bothered me, and I decided to pursue my PhD in an effort to answer it. When I started, I did not even know how to frame the question, nor how I might pursue even a small part of the answer. Now, I can at least think of a way to frame the question and to approach a piece of the answer.

Chapter One: Introduction

Decision-making reflects who we are. It is a human enterprise, accomplished by people who share all the glory and folly of humanity. When H.G. Wells introduced the alien invaders in *The War of the Worlds*, his first sentence described them as "intellects vast and cool and unsympathetic" – aptly capturing their non-humanity.

When studying the process of decision, it is important to keep in mind not only who we are as humans, but also who we are not. This dissertation proceeds from that point of view – that the act of making a decision reflects how people actually process the information surrounding that decision. Expecting decision-making to proceed from other modes is to force decision-makers into a semantic cage – and to be surprised when the occupants persist on breaking the bars.

The Research Question

The core research question of this dissertation is simply stated: how does the way our minds process data affect a decision maker's actions? If the interaction of the situation and mental processing is the key to understanding decision-making behavior, can that interaction be calculated and predicted probabilistically? This leads to the answers to critical questions in international relations and foreign policy decision-making theory and practice: When do states make war and engage in conflict? When do states make peace and engage in cooperation? Despite the spilling of much ink over the question, the issue of state interaction and decision-making is one which continues to be debated in both scholarly and practitioner literature. Clearly, the ability to predict, even probabilistically, the outbreak of war and the probability of peace will provide an important tool for policy makers, and an important starting point for future research.

Specifically, I am attempting to answer the following questions:

- Given starting conditions, can the probability of conflict be predicted?
- If there is conflict, who are the likely participants, and on which side?
- What is the structure of alliances?
- What are the probabilities for each variation of the above results?

As I will discuss in more detail below, decisions, whether for war or peace; are shaped by the interaction between the *decision space*, the decision-makers' *situation recognition*, and the *proposed alternatives*. The intersection between these three items is the way in which decision-makers mentally process that information and use their memory to aid that process. Thus, memory lies at the heart of decision making. Accordingly, I have created a simulation model which explicitly models how individuals process information in a given situation. This memory-based simulation will be validated in game theoretic situations, in a simplified multi-player national security game, and then will be tested against two historical cases. These historical cases will show that the process of memory can help predict the likelihood of conflict in a given situation, the probably states involved, and the probable structure of alliances.

The Argument in a Nutshell

The act of decision is shaped in two ways. The first shaping factor is the surrounding ecology. This includes the challenge at hand along with all of the known and unknown factors which affect the decision. One way to think of this is that the decision itself is shaped by the *decision space* in which the decision-maker must operate.

The second way that decisions are shaped is through the way in which information about that decision is processed by the decision-maker. This includes generating hypotheses about the situation, or *situation recognition*, the process of generating possible courses of action, or *proposed alternatives*, and the selection of the *decision* itself.

This process of understanding the decision space, situation recognition, generating proposed alternatives, and the ultimate decision all are controlled by the way in which individuals mentally process information. This requires information to be brought into working memory, matched up to existing semantic concepts in memory, and then checked against long-term memory. This process continues until a satisfactory conclusion is reached, or the time for decision runs out. The exact way in which this iterative process occurs within the brain is not fully understood, but it is well enough understood to mathematically model the general process.

This is different than the process assumed by rational choice theory. Of course, rational choice theorists readily admit this is so, but argue that it is sufficient to be able to assume instrumental rationality – that people act *as if* they were operating according to the tenets of rational choice. In this setting, it is enough to understand or be able to infer the preference structure of the decision-maker, and to be able to fit the decision within a particular game. Decision-makers may have incomplete information, may be affected by misperception, may be affected by biases which cause them to frame or see the issue in different ways, but all of these can be dealt with either mathematically or by inferring that certain biases or misperceptions are operative based on *post hoc* analysis.

The rational choice mode of analysis requires that the game being played, preferences, perceptions and biases be either inferred or assumed. However, such inferences and assumptions can lead to misleading conclusions. This dissertation will demonstrate that specific assumptions about game structure and player preferences are not required. Instead of being fixed, preferences are malleable, and the way in which preferences shift based on the situation can be modeled and, hence, predicted. Intuitively, this can be seen in daily life. There may be a general preference, along the lines of "do as well as I can", but that general preference is operationalized via specific, situationally-based preferences. The constantly changing situation of life causes people to apply these differing sets of preferences based on their recognition of the current situation, and even to change their preferences within a given set. A person's strong preference for maintain a desirable weight can (too easily!) be changed when faced with warm chocolate brownies. The memory of the pleasure associated with previous brownies causes the recognition of the situation to change (I'll just eat one tonight) and the preference structure accordingly shifts in favor of eating.

This implies that situation and preference structure are linked, and that the linkage is defined by the way information is mentally processed. Thus, given a situation and a good idea of how the information is processed, the operative preferences can be determined and probable decisions can be predicted. In the context of international relations, a leader may want to "do well", but the way in which the leader goes about "doing well" changes from time to time and place to place. At one point the preference may be for economic strength, at another, military strength. In a dangerous situation where options are vanishing, a normally peaceful leader may well decide that this is Sun Tzu's proverbial "death ground", and that war is the only answer.[1] Decision-makers face choices all the time, and sometimes choices take the form of choosing between desirable outcomes. More often, a choice takes the form of choosing the "least bad" course of action, and as such, choice will be critically dependent on which set of preferences is dictated by the situation.

The MEMORIS-2 simulation created for this dissertation shows how decision space, situation recognition, alternatives, and choice all are linked by the way people mentally process information. Taking this knowledge, it then is possible to probabilistically predict outcomes given a set of situational information.

Methodology and Procedures

As mentioned briefly above, a three-step process is used to gain traction on this problem. First, the concept of memory-based simulation will be introduced via analysis of simple two-sided (dyadic) interactions. This simulation will then be used, following Axelrod (1984), in a tournament vs. a number of other competing strategies in an iterated Prisoners Dilemma (PD) game. The tournament will then be repeated in a game which iterates dynamically between a PD and a Stag Hunt game. This will demonstrate that the memory-based simulation is indeed a viable approach in both static and dynamic game theoretic situations.

[1] "In death ground, fight." *The Art of War*, Chapter 11: "The Nine Situations"

In the second step, the simulation will be expanded in two ways. First, the simulation will be expanded to accommodate three players. Next, the simulation will be used to model the interaction of these players in a simplified game of international relations. This three-player baseline will be used to assess model behavior. Specifically, it will demonstrate how the simulation results vary, and how that variation is related to the level of the threat. The outputs of the three-player game then will be analyzed using clustering techniques to show that the results fall into three distinct categories, which I characterize as “war”, “armed standoff”, and “peaceful alliance”. Further, the distribution of these clusters changes as the general threat level increases, shifting from mostly peaceful alliance at low threat levels, to armed standoff at moderate threat levels, to war at high threat levels. These results show that the memory-based simulation is a viable way to analyze the impact of varying situations in a small-scale setting.

With those results in hand, a more complex simulation will be developed and used to analyze two different cases, the situation in Europe which led to World War I, and the Cold War, which did not lead to war between the superpowers. In each case, a simulation baseline will be developed to show how each case reacts to increasing levels of threat and to validate the cluster analysis performed previously. Then, the simulations will be run using the actual historical data, which are drawn from the Correlates of War (COW) datasets.

The simulation uses a simplified model of state interactions, which requires reducing the variables in the COW National Military Capabilities (NMC), Militarized Interstate Dispute (MID), Alliances, and Total Trade (TT) datasets to four dimensions:

combat power, economic power, war, and alliance. Factor analysis is used to reduce the MID and TT data into two factors (combat and economic power), and then war and alliance data are recoded to fit the requirements of the simulation.

Using 20 years of data as initialization, the simulations will be run year by year to show the change in probability of war and alliance, along with the patterns of conflict and alliance for each year of the appropriate case. For the World War I case, the simulations start every year from 1905-1913, with each simulation being initialized with 20 previous years of actual data. For the Cold War, the simulations will run yearly from 1968-1987, with 20 years of actual data for initialization. Thus, the 1968 Cold War case is initialized with data from 1948-1967 before running. To capture the early years of the Cold War, an alternative initialization strategy will be used. In both cases, the patterns of conflict and alliance match up well with the historical record.

Results and Analysis

Along with the data above, the simulation model will provide the following data for each country and each year of the simulation:

- Whether the state is in conflict, and with which other states;
- Whether the state is in alliance, and with which other states; and
- Descriptive data about state characteristics and capabilities.

When considering the expected results from the models, it is important to understand how validity of results will be judged. Obviously, there is an existing historical record of conflict and alliance leading to WWI and during the Cold War, which can be compared directly to the record of conflict and alliance year by year.

Thus, the primary measure of merit will be *outcome* validity, meaning how well the predicted outcomes match the historical record. Another way to consider validity is via *process* validity, which asks whether the model faithfully replicated the process of decision, as suggested by Taber & Timpone (1996). As this particular simulation is focused on the individual level of analysis, it does not model (for example) group-level behavior within a particular state's decision-making process. Modeling of these group-level processes will be a potential avenue for future research.

This effort will attempt to answer the research questions by a three-step process: first building a simulation which is based on the way people mentally process information. The second step will be testing the simulation, first in a game-theoretic setting and then in a simplified model of international interaction. The final step will be using the simulation to examine two cases, the run-up to World War I and the Cold War. The results show that the memory-based simulation accurately reflects the historical record, including accurate predictions of which states are most likely to engage in conflict, which states are most likely to ally with each other, and the increase or decrease in the probability of major war.

While few would deny that decision making in real life is an imperfect exercise of imperfect judgment by imperfect people, rational choice remains the dominant paradigm for analyzing decision making. One reason for this is the existence of powerful quantitative tools of analysis and modeling which can be applied within the rational choice paradigm. On the other hand, tools for performing quantitative analysis of the cognitive process of situational recognition and the application of appropriate rules have been lacking. As will be discussed below, the emergence of new tools for modeling the process of situational recognition and rule choice allows the ability to rectify this situation. This dissertation will use ideas from cognitive psychology as the micro-foundation for analysis of decision making as an alternative to using tools which draw their foundation from economic analysis. Clearly, there is plenty of room for improvement.

This allows for a new line of analytical approach, instead of using tools such as expected utility modeling based on the rational choice paradigm, it is possible to model situational recognition and the process of decision from the cognitive angle. The application of this framework returns us to the aim of this dissertation: to show how the relationship between the situation and how that information is processed affects the actual decision, and how well can modeling this interaction can probabilistically predict the decisions.

Roads Not Taken

To be sure, and as pointed out above, memory-based models are not the only way to model this set of interaction. Clearly, expected utility modeling has been used to good effect, as has agent-based simulation (Bueno de Mesquita 2002; Cederman 2002). Given more time and computing power, I could explore a variety of spatial models, which take their cues from Downs' (1957) insights on how politicians adjust their positions to appeal to voters and apply them to issues of international relations and security (Hug 1999; Morgan 1984, 1990 and Morrow 1985, see also Liu 2006). I could also follow Richardson's (1960) equation based approach to arms races and model state interactions via dynamic modeling (Li & Thompson 1978; Muncaster & Zinnes 1983; Zinnes & Muncaster 1984). There also are models which attempt to simulate the process of decision, such as Taber's Poli (Policy Arguer) model (1992) and Guetzkow's Inter-Nation Simulation (1962) along with its follow-ons, such as Hermann & Hermann (1968). In the cognitive world, I could also explore the potentially promising terrain suggested by Decision Field Theory, which models shifts in preference structure based on how the brain processes information (Busemeyer & Townsend 1993; Roe, Busemeyer & Townsend 2001). Direct comparisons between these approaches and memory-based models are beyond the scope of this dissertation, and will have to wait for follow-on efforts.

The Cases: Europe 1880-1915 and The Cold War

The situation in Europe between 1880 and 1915, and the Cold War provides a useful pair of cases to assess the memory-based simulation approach. These cases, while certainly not an exhaustive set in the sense of a comparative case study, provide exemplars that I use to show that the simulation produces usable results. In these situations, there are a relatively small set of states with a wide variety of conflict and alliance interactions, high levels of tension and militarized disputes, and significant arms races. One led to a conflict involving most, but not all, of the participants, while the other did not lead to full-scale war. While it is well beyond the scope of this dissertation to analyze or predict the causes of World War I and the Cold War (see, among many others: Farrar, 1972; Gaddis 2005; Gordon, 1974; Holsti, 1965; Hermann and Hermann, 1968; Kaiser, 1983; Kennedy, 1984; Levy, 1990; Maier, 1988; Holsti, North and Brodie, 1968; Keegan, 2000; Kissinger 1994; Massie, 1991, Trachtenberg, 1990; Tuchman, 1962; Van Evera, 1984; Williamson, 1988), the results of the simulation provides validation for both the model and its underlying logic.

These cases are a useful test for the cognitive modeling approach, as many of the actions taken by actors during the run-up to the war turned on their recognition of the situation and the application of rules for action, such as decisions to pursue arms races, militarized disputes, mobilization and, either war or peace. Further, the decision makers in the situation reverse their actions during key points, an eventuality which has tended to be problematic for rational choice approaches but can be modeled using this cognitive simulation approach.

Two cases, even though they contain thousands of individual data points, are not of themselves decisive. They do, however, provide the opportunity to demonstrate the viability of a cognitively-based modeling approach. This would be a significant benefit on several levels. First, by providing a quantitative alternative to rational choice modeling, it adds a significant new tool for future analysis. Second, it would provide a potential foundation for analysis of political decisions that is not based on the economic ideas which underlie much of existing IR and decision theory. Finally, it provides a potential foundation for the analysis of constructivist modes of interaction.

Outline of the Dissertation

The remainder of the dissertation will proceed as follows: Chapter Two will review the relevant literature in the general areas of International Relations theory, foreign policy-making, and judgment and decision making. Chapter Three will be devoted to developing the memory-based simulation model, and will use a building-block approach where successive iterations of the model will add more and more complexity, until a memory-based, multiple-player model is ready for use. The simulation first will be validated in a game-theoretic setting, and then in a simplified multi-player game. This will create the simulation and the baseline data from which the remainder of the dissertation will flow.

Chapters Four and Five will be devoted to the analysis of the two cases: Chapter Four will cover the run-up to World War I and Chapter Five will cover the Cold War. Chapter Six will provide the summary, conclusions, and directions for future work.

Chapter Two: Review of the Literature

War, Peace, and Decision

The questions of why nations choose war or peace, alliance or isolation, draw on distinct, although converging, streams of literature: international relations, foreign policy and judgment and decision-making theory. Even though these sets of literature often address different questions, the focus for this effort will be the nexus of foreign-policy decision-making in the conduct of international relations, and as such it is useful to situate this effort within these sets of literature.

International Relations

Approaches to this question can be defined on several levels. Following Waltz' images, the question of assessment and decision can be approached on three levels: individual, state, and system (1959). Others have effectively added a fourth level, that of interaction or relationships between various actors or between the actors and the system (see, Bueno de Mesquita and Lalman, 1988, 1992; Vasquez, 1998, among others). As such, in this literature, the reasons that states choose war and alliance can be found in either the structure of the system, the characteristics of a state, or the qualities of the individual.

Systemic, Intermestic, and State-Level Models

Realism and Its Variants

Systemic models of international relations include realism (Carr, 1946; Morgenthau, 1948 [1993]), along with its various offshoots of neorealism, such as defensive realism (Waltz, 1979), and offensive realism (Mearsheimer, 2001).

Considered broadly, realism and its variants see states as unitary, rational actors within an anarchic environment. In general, realist approaches see power, security and threats are the most important issues, and states must either engage in "self-help" or attempt to band with others in various "balance of power" formulations. The key here is the rational (or realistic) pursuit of national interest. The concepts of realism have been subsequently expanded to incorporate a wide variety of approaches, including the incorporation of national intentions and identity (see, among others; Walt, 1987; Schweller, 1994; Glaser, 1994; Oye, 1986; Ruggie, 1983; and Vasquez, 1998 for a summary).

War and Alliance in Realist Theory

In the various realist formulations, the choices for war and alliance are determined by the need to pursue national interests, however defined. At least as far back as Thucydides, the realist view of war and alliance comes down to fear and power. In the Melian Dialogue, the Athenians demand that the Melians join their alliance, arguing that the failure of Melos to ally with Athens would hurt the Athenians by making other states believe that Athens was weak. Out of fear, the Athenians offered a stark choice: alliance or destruction (Thucydides, 1951). Further, because realism assumes an anarchical system, then nothing can stop states from using force as the *ultima ratio* to solve their differences (Waltz 1959). Thus, states might turn to force to settle disputes, out of misunderstandings of relative power, as a preemptive step, or due to private information leading to perceived advantage (Fearon 1995). Application of game theory in realist theory often leads to explain the interaction between nations as a form of a Prisoners Dilemma (see Chapter 3 for a complete explanation) where, while it may be possible to gain by cooperation, the lowest risk choice can lead to conflict.

To some theorists, the power transition is a particularly dangerous time. At this point, the state with the most power is being overtaken by other states, and war may break out as the aspiring hegemon challenges the incumbent (Oganski, 1958; Modelski and Thompson, 1989). As above, this can easily be seen in the traditional sense of the interplay between power and fear.

This combination of fear and power also leads to a security dilemma. By building up military force for its own security, State A also becomes more threatening to its neighbors, who in turn build up their forces, which then decreases the security of State A. Thus, the desire to increase security can, paradoxically, lead to decreased security. The opposite choice, decreasing forces, also leads to decreased security. How then to proceed? One way taken by realist theorists has been to redefine the aims of states within the systemic landscape. Mearsheimer's original concept was based on the idea that the pursuit of absolute power – the more the better – was the most important aim. That idea has been successively refined to the idea of the pursuit of relative power, and particularly a balance of power by Waltz (1979), then to a balance of threat (Walt, 1987) and more recently, a balance of interest (Schweller, 1994).

This leads to the other way that various shades of realist theorists believe that states should manage security threats – through alliances. Alliances are (potentially) a way to circumvent the security dilemma. States can engage in balancing alliances, where weaker states ally against the stronger, or can take a bandwagoning approach, where they join with the strongest state to reap some anticipated reward or to avoid the prospect of being conquered themselves. Finally, as seen in the Melian Dialogue above, strong states can engage in chain-ganging to harness weaker states to the cause. In all of these cases, alliances are necessarily shifting and temporary structures, lasting only as long as all the parties believe it to be in their interest. Further, it is by no means certain who should align with whom. As Snyder puts it,

> … in a multipolar system there is a general incentive to ally with *some* other state or states … that is generated by the structure of the system. Who aligns with whom results from a bargaining process that is theoretically indeterminate. The indeterminacy is reduced, though not eliminated, by the prior interests, conflicts, and affinities between states and their internal political make-up. (1994, 465-6, italics in original)

The Role of Reputation

One way that states try to resolve this indeterminacy is via reputation. Following Schelling, one thing that helps resolve games is the presence of a focal point that narrows the range of choice (1960). In this sense, reputation can provide a focal point. If a state has a reputation for backing up its threats, or for punishing those who cross it, then other states should make their choices accordingly. Thus, reputation becomes a potential manifestation of state power. As such, the logic goes that states should act to protect their reputation, even if it means short-term losses. As Mercer notes, US Presidents have often invoked the specter of the impact of a lost reputation for resolve when taking military action (1996).

However, Mercer takes a different view of reputation, arguing from social psychology that desirable behavior by members of the *in-group* and undesirable behavior by the *out-group* is subject to situational attribution – the friend stood firm and the enemy defected because they had no choice. Opposing actions are thus subject to dispositional attribution – an enemy stands firm because it has resolve, or a friend defects due to lack of resolve. Thus, Mercer concludes that others will judge reputation by our status as a member of the in- or out-group, so there is nothing to be gained by taking actions to protect reputation (1996). Effectively, this is a *least-regret* type of viewpoint. Enemies are expected to resist to the utmost and friends are expected to cut and run, and it is a pleasant surprise when the opposite occurs.

Mercer focused on a single type of reputation – that of resolve. Miller argues that other types of reputation, and in particular the reputation for reliability, are also important in alliance choices. In this analysis, a reputation for reliability or lack thereof can affect the likelihood of a state being chosen for alliance, and can affect the structure of the agreement for alliance. Thus, Miller's view offers a more nuanced view of reputation, which allows for both situational and gradational aspects of reputation (2003).

Adding State-Level Variables: Liberalism

One thing that the literature above shows is the increasing role in the characteristic of the state, whether it is interests or reputation, in the literature. This leads naturally to more liberal IR concepts of the interaction between states.

With the perceived decline of American hegemony and accompanying questions about the utility of power, scholars in the late 1970s began to look at state interaction in a different frame of reference. If realism and its offshoots focus on why states often engage in violent conflict, liberalism focuses on why states cooperate. Rather than focusing on war and conflict, liberalism focuses on the interactive, interdependent, and generally peaceful nature of international relations (Keohane, 1982, 1984; Moravcsik, 1997; Ferguson & Mansbach, 2003).

An important difference between liberalism and realism is that liberalism looks inside realism's "black box" decision model to find the internal processes that often drive a state's decision-making process. Beyond power and security, the economy has an important role to play in liberal theories of IR. Liberal theory adds international and supra-national institutions to the mix as well, trading the parsimony of realism for a more complete look at the interplay of various factors. As Baldwin (1993) points out, the anarchical state of the world is an opportunity for cooperation as well as competition. Liberalism thus leavens the essential anarchy of the world system with an interdependent web of formal and informal relationships, and thus, structure is an important consideration in liberal and neo-liberal theory (see, among others; Baldwin, 1993; Barnett and Duvall, 2005; Barnett and Finnemore, 2004, 2005; Keohane, 1982, 2002; Lipson, 2003; Milner, 1991; Moravcsik, 1997).

War and Alliance in Liberal Theory

Instead of seeking systemic reasons for conflict, Liberal IR Theory looks to the state. This can be explicitly seen in Kant's *Perpetual Peace*, where the causes of war are laid at the feet of despotism and the hope of peace is found in republican forms of government. To liberal theorists, the interdependence caused by free trade, free exchange of ideas, and a worldwide cosmopolitan society of free people is the way in which war would be eventually abolished (Kant, 2001 [1795]).

War then, occurs for two reasons. The first is the natural tendency of those states which have not developed the necessary conditions to join the worldwide society of republican nations. These states will be aggressive and will cause war in their natural desire to increase influence and power. This leads to the second reason for war, which is the way in which ordinarily peaceful states will band together to repel aggression in the name of collective security.

This results in a situation where non-democracies (or non-republics) fight each other, democracies fight non-democracies, but democracies do not fight each other. This empirical fact leads to Democratic Peace theory, as well as various formulations based on the interplay of domestic politics and regime types. In these approaches, the dynamic of domestic politics and decision-making is the primary determinant of state behavior (Lipson, 2003).

The reader will have probably noted that the above approaches are not purely systemic. Instead, they cross the boundary from the systemic level to incorporate parts of state behavior. This "intermestic" approach is explicitly taken in two-level game theory. In this approach, a country's leader essentially sits down to bargain at two tables simultaneously. At the international table, the interlocutors are heads of state, while at the domestic table the interlocutors are those important groups who have the ability to either ratify or veto the international agreement – or, in some cases to remove the executive from office entirely. The leader is faced with the task of balancing those potentially competing agendas. Arguably, this can lead to conflict by limiting the leader's freedom of action so much that an otherwise reachable agreement is not possible, and war ensues instead. As Putnam puts it, this can limit the allowable "win space" where the parties' acceptable positions overlap. At the same time (and familiar to those who have actually engaged in high-level negotiations), a savvy executive can use the threat of non-ratification as a tool to extract additional concessions (Putnam, 1988). Here, the clash of interests is not just between states, but also within the states themselves.

Not surprisingly, the liberal IR explanation of alliances follows the same lines as that of war. If war is a collective problem, then it needs collective action to solve it. Alliances and other agreements are primarily aimed at solving collective problems, and as such require states to provide public goods (military force in this case) to protect the collective interest, as well as mechanisms to ensure there are not too many free-riders of defectors (Olson, 1971). As such, reputation is also important in this circumstance. That said, the earlier discussion of reputation is also applicable to the liberal conception of war and alliance, so I will not repeat it here.

Marxist and Critical Theoretic Approaches

Just as Liberalism sees the international structure being driven by the interaction between liberal republics and other states, Marxist theory views the international structure as a product of historic and material factors – and in particular the control of the means of economic production. To Marx and Lenin, the current world system is sharply divided along economically-based class lines. Capitalism, as the current dominating economic system, effectively controls the world system, and capitalism also carries within itself the seeds of its eventual destruction (Marx, 1859; Lenin, 1939).

The failure of the "inevitable" proletarian revolution to occur has caused some redefinition of the original Marxist-Leninist ideas. Immanuel Wallerstein tried to explain this state of affairs by redefining the world structure from a two-class division between the capitalists and the exploited masses to a three-layer division of core states (the exploiters), peripheral states (the exploited) and a buffer of semi-peripheral states which both exploit and are in turn exploited. This structure, Wallerstein argues, essentially allows exploitation at arm's length, so the contradiction between the exploiters and the exploited is not so sharp as to generate open conflict (1974). This basic analysis has also been shared by *world systems theory* or *dependency theory*, which posits a situation of stratification, where states in the periphery are purposely kept dependent on the core. This allows a continuation of the current world system (see, for example: Buzan, 1994; Frank, 1966; Galtung, 1971).

War and Alliance in Marxist Approaches

As the world structure for Marxists is determined by economic factors, so is the outbreak of war and the conclusion of alliance. To Lenin, imperialism and subsequent imperialist wars of conquest was a necessary stage of capitalism, to be followed by an inevitable war between capitalist states once all imperialist conquests were completed (1939). Wallerstein (1974) and Buzan (1994) view the domination of the center as nearly unshakeable, and that conflict will either be between states in the periphery, or more likely through war and subversion on the part of the core states to maintain the current world system *status quo*.

Similarly, alliances lie along economic lines, and are generated by economic interests (Lenin 1939). Any alliance between the core and peripheral states would naturally be for the purpose of maintaining the current state of affairs, and would naturally be broken if economic interests require it.

The Origin of Interests: Constructivism

It is fair to say that realism, liberalism, and Marxism tend to view the world through monochromatic lenses – realism via conflict, liberalism via cooperation, and Marxists via economics. However, all of these approaches view the fundamental interest of actors within the system as a given. Thus, states seek power, states seek to solve problems, and states are products of economics, respectively.

This fixed set of interests (however defined) generally should result in a stable or slowly evolving system, but changes and upheavals occur, and constructivism is an attempt to explain why things change and how and why state interests and the system structure come to exist. Constructivism is not in itself a coherent IR theory, but it is a way of looking at how interests and the relations between states are formed. In the constructivist idiom, meaning and structure shift constantly as states and other actors create an endlessly changing social reality through their actions and especially through their chosen means of communication (Wendt, 1992, 1999; Lebow, 2001; Onuf, 1989).

Put simply, the interests and identities of states are not exogenously given. Instead, interests and identities are constructed via social interaction. It is much like two strangers meeting for the first time; they have no relationship or reality between them until they begin to interact. Along with the relationship, states also build up, or constitute, identities as they interact. Thus, constructivism is concerned with the process of forming identity, interest, and structure (Hasenclever, et al., 1997; Hopf, 1998; Onuf, 1989; Reus-Smit, 2001; Wendt, 1992, 1999; among others).

While there are several approaches to constructivism, they all take their cue from Hedley Bull's (1977) observation that states form a society. Social relationships are built up through social interaction, and the structure of these interactions forms the structure of the international system. To constructivists, this socially-created structure is more important than the material aspects of the world structure in determining state behavior (Wendt, 1992, Reus-Smit, 2001). Once states have socially determined their identities, interests, and the associated structure, they may pursue these socially-constructed interests in much the same fashion as dictated by rational choice, except that the value structure is socially determined, rather than a given. Alternatively, actors may play socially-determined roles within the system (Wendt, 1999; Barnett & Duvall, 2005; Steger, 2008).

Since both state identities and the structure of their interaction are constantly evolving, the constructivist system can be thought of as the interplay between complex, interdependent dynamic systems. The branch of mathematics known as complexity theory tells us that such a system will experience *punctuated equilibria;* that is, periods of relative stability punctuated by large shifts in alignment. This is also sometimes referred to as *emergent behavior* (Axelrod, 1997, 4). This is often mirrored by the general situation in the world system. As Steger puts it, the world's "social space" can be constructed and reconstructed, and "such change can occur with lightning speed and tremendous ferocity." (2008, 7) However, neither complexity theory nor constructivism can tell us *when* or *why* a shift will occur or in what direction things will change.

War and Alliance in Constructivism

Constructivism has little to say directly about the outbreak of war. Instead, the character of conflict depends on the particular set of interests and structure that has been developed via the interaction of the various actors. Following the logic in the previous section, if the actors develop power or security as an interest, then wars and alliances will occur in the same way predicted by realism, and so on (Wendt, 1992).

Alliances and other agreements follow the same pattern, except that the socially-determined set of interests can create system-wide learning through such devices as *epistemic communities*, which in turn creates new norms and sets of interests. These can push states toward alliances and other control regimes that would not be possible without this sense of collective norms (Hasenclever, et al., 2001)

Critique of Systemic, Intermestic, and State-Level Approaches

If asked to answer the question: "which one of the above approaches does a good job in predicting the possibility of war or the probability of alliance?" the answer would have to be, "none of the above". This is because, at the root, all of these theories are normative. They make prescriptions for how states *should* act rather than describe how states *do* act. Waltz makes this point with refreshing candor in his spirited defense of the viability of structural realism after the Cold War – but also makes it clear that states may choose to fight, or not, and may choose to balance, bandwagon, or not (2000).

Similarly, Liberal theory indicates that democratic states should not fight each other, but may fight other states in the name of collective action. There is currently strong empirical support for the democratic peace theory, but this begs the question of whether this relatively short-lived phenomenon is nothing more than a blip in the data. While democracies have not fought democracies, they have caused regime change (sometimes back to authoritarian governments) in other democracies, and who is to say that this *sub rosa* conflict will never break out into the open? Further, like realism, liberalism does not predict when states will fight or ally, only that states should or should not do so under certain circumstances.

Along with the fundamental issues of the Marxist critique of capitalism (such as the Marxist view of how value is determined), which I will not expand upon, the main weakness of this type of theory is that the predicted conflict along class lines simply has not happened. Theorists such as Wallerstein have had to result to Ptolemaic devices such as the invention of the "semi-periphery" in an attempt to get around this failure. Even with this epicyclical addition, this set of theories basically predicts a *status quo*, which does no good when trying to understand sudden and emergent behaviors such as war.

As mentioned earlier, constructivism is not a coherent theory of how states act, but instead focuses on how interests and structure are derived in the international system. The bottom line with this set of theories is that they give prescriptions of how states should act, and try to explain how states determine their interest, but none of them predict (or to be fair, even purport to predict) *when* and *how* states will act.

Individual-Level Models

If state and system-level theories are not helpful, there may be more traction on the problem of prediction in in the individual level of analysis. Psychological approaches realize that an individual decision-maker is often the critical link in determining what a state will or will not do. Decisions are not simple results of some arbitrary black box process. Instead, a decision results from a very human and fallible process. The psychological approach assesses all the factors which bear on the decision-making process, including cognitive and personality issues, types of information, and the historical, social and organizational context to attempt to isolate those factors which may or may not prove to be decisive in a particular decision (See, among others; Dyson and Preston, 2006; Holsti, 1965, 1972; Horowitz, McDermott and Stam, 2005; Jervis, 1976, et seq.; Kanwisher, 1989; Keller, 2005; Lasswell, 1930, 1935, 1950; Vertzberger, 1986, 1990).

States still exist and operate in an anarchic environment, and leaders still attempt to make "correct" decisions. What is meant by "correct" is sometimes questionable, but is usually defined as being somehow "logical" and in accordance with "objective reality" (however defined). This often means that authors use this approach to critique leaders for making decisions which are not "logical" or "correct" – meaning that the author thinks the question should have been decided differently (for example, see Vertzberger's (1990) critique of Israeli decision-making) This approach does have the advantage of pointing out that it is often difficult for a decision-maker to discern what facts are actually true, meaningful and relevant, and even more difficult sometimes for correct data to work its way through all the various filters between the sender and the receiver (Vertzberger 1990, among others).

Psychological approaches can be roughly divided into four categories, individual characteristics, perceptual, affective and social. The individual characteristics approach is perhaps the oldest of these, going back at least as far as the "great man" theory of history. Here, a decision-maker's performance or actions are a result of various characteristics, whether they be personality type or personality disorder (Lasswell, 1930 et seq.), decision style (Keller, 2005), cognitive complexity (Shapiro and Bonham, 1973; Dyson and Preston, 2006) or physical characteristics such as age as a surrogate for testosterone level (Horowitz, McDermott and Stam, 2005).

Another group of authors also began to consider the idea that the way in which decision makers perceived and evaluated their environment was crucial. Ole Holsti (1965, et seq.) looked at the role of stress and time pressure and their effects on decision making. Robert Jervis, in his landmark study *Perception and Misperception in International Politics*, proposed that the need for cognitive consistency caused decision makers to either filter or misperceive contrary data (1976). This book led to several similar analyses, including Yaacov Vertzberger, who posited decision makers as "practical-intuitive historians" who were often led astray because they became wedded to their intuitive ideas about a given situation, and were unable to properly integrate conflicting data (1985, see also Khong 1993). Vertzberger later expanded this drive for cognitive consistency into a more nuanced concept of decision-making schema (1990).

Perceptual approaches draw strength from an exhaustive coverage of the variables affecting both communication and processing of information. We have a better appreciation for why decisions are sometimes wrong or confused, because we can see that leaders are not creatures of vast and cool intellect, serenely making decisions based on precise data. Instead, leaders are immersed in an environment of "complex uncertainty" where they are fallible, groping, often blinded by competing streams of confusing data, and all too human (Verztberger, 1990).

The social psychology literature studies the effects of social group dynamics on decision making. As Tetlock (1985) argues, social factors play a pivotal role in decision making, and numerous studies have tried to quantify that effect. Probably the most well known effort in this realm was Irving Janis' article and subsequent book "Groupthink", where he looked at decision making in a variety of situations, most notably the ill-fated "Bay of Pigs" invasion. Here, Janis showed that a variety of small group interactions could create an overwhelming "need for consensus" where contrary information and viewpoints are withheld from the process. This type of dynamic creates the conditions for disastrous failure (1972). While the idea of groupthink has entered the popular lexicon and has often been thrown around as a critique, efforts to expand the concept have met with mixed success (see, for example 't Hart et al 1997).

Indeed, the picture of whether groups help or hurt decision-making effectiveness is decidedly mixed. Plous (1993) reviews a great deal of literature on group decision making and finds support for both sides of the coin. Groups can often outperform the average individual, but the best member of the group often outperforms the group. At the same time, heuristics and biases which affect individuals are often mirrored in groups, and group discussion often simply reinforces predetermined attitudes. Plous concludes that "collaboration is no guarantee of a successful outcome" (1993, 214; see also March 1994).

Decision-Making

This psychological literature has led us to the literature on decision-making. While much of the literature on rational choice decision making has been noted above, it is useful to consider the mechanisms of rational choice before turning to other decision-making approaches.

Rational Choice Models

Rational choice approaches attempt to provide more rigor to the analysis of decisions by assuming that an individual decision-maker will be goal-oriented toward maximizing his or her utility function (however defined) within a given structure of incentives or "game". Everyone in the system has a distinct set of ordered preferences from which they work, and they analyze the situation and choose the desired alternative (Bueno de Mesquita 1984 et seq.; Kahler, 1998; Schelling, 1960; Snyder, 1960). Individuals are also aware of other individuals, and they will take their observed or probable actions into account using the thought processes like those described in game theory. Since these processes proceed from a defined set of logical principles, they are amenable to computer-based modeling and analysis (Bueno de Mesquita 1984 et seq.; Kahler, 1998, see also Zinnes and Gillespie, 1976 for a variety of statistically-based approaches). These ideas have been used to considerable effect by scholars such as Bruce Bueno de Mesquita, who has demonstrated significant predictive power using a combination of game theory and expected utility modeling. This approach requires the analyst to understand or make assumptions about the game structure, actor preferences, the salience of issues, and the amount of power that actors control (see, among others, Jervis, 1978; Bueno de Mesquita, 1984; 1985; 2000; 2002; Bueno de Mesquita, et al., 2003).

This mathematical rigor is one of the strengths of rational choice theory, because it requires causal relationships and judgment mechanisms to be explicitly addressed. There is also a common-sense appeal to the idea that everyone will attempt to choose what is best for them. It also recognizes the key role of the individual in the process. This approach is thought to be most appropriate when the stakes are high and the choices are distinct, such as nuclear deterrence (Kahler, 1998; Kahn, 1960; Schelling, 1960; but see Snyder, 1978; Vasquez, 1998 and Ferguson & Mansbach, 2003 for a critique).

Critique of Individual and Rational Choice Models

Rational choice approaches give the illusion of careful consideration of alternatives and the ability to take the best action. Since rational choice theory requires rank ordering of alternatives, this implies that such alternatives are transitive; that is no matter how the ranking is created, it will be the same ranking. Recent (and not so recent) work in cognitive psychology shows that preference rankings are intransitive – the decision changes depending on how the question is framed. As one example, Tversky and Kahneman point out that preferences can show multiple reversals depending on the decision maker's understanding of the relative prospects for gain and loss – even when given different formulations of the same choice (Kahneman & Tversky, 1979; Tversky & Kahneman, 2000b, 2000c). It has been subsequently claimed that preferences are transitive within a given positive or negative frame (Bueno de Mesquita and McDermott, 2004). Other research, however, has indicated that preference orders can indeed be intransitive based on the order of presentation of alternatives, particularly where there are three or more choices (Arrow, 1951; Tversky & Kahneman, 2000c). Indeed, as Verba (1961) points out, decision makers are rarely completely aware of their own value systems, to say nothing of neatly ordered preferences.

Further, decision makers are almost never able to determine a complete set of alternatives. Work by Weber, et al. (1993) among others has shown that individuals tend to generate between 2-4 hypotheses about a given situation, even when many other hypotheses are possible. One implication of this is that decision makers generally greatly over-estimate the probability that a given hypothesis is correct. For example, say there are five possible hypotheses about a given situation, with true probabilities of being correct as 0.4, 0.3, 0.1, 0.1, and 0.1. If the decision maker only generates two hypotheses, then any estimate made of correct probabilities (say .7/.3) will wildly over-estimate the probability that the leading hypothesis is correct (in this case, .7 vs .4) (Dougherty & Hunter, 2003)

Additionally, decision makers will tend to generate hypotheses based on what are considered the most probable outcomes, rather than trying to cover the total range of possibilities. This has significant implications in how decision-makers search for information, as this information search is guided by the hypothesis being considered. Rather than looking for information that is diagnostic – meaning information that will disprove a hypothesis, the information search tends to be confirmatory – meaning that the decision-maker naturally searches for information to confirm the existing "leading contender" (Weber, et al., 1993, Dougherty & Hunter, 2003, Gettys, et al., 1987)

These factors (and others) when taken together show that individuals cannot be expected to judge or rank order preferences in the way required for rational choice theory. This can be scaled from the individual to the organizational or state level. Even though states are not individuals, they are made up of collections of individuals, and the ultimate decisions are also made by individuals. Indeed, as noted above, groups can generate limitations in the ability to process information that mirror well-known individual biases (see Allison 1971, Janis 1972, Jervis 1989, Plous 1993 and March 1994). Further, the assumption that individual choice can be aggregated to explain societal choice is one shared by many theoretical paradigms, such as Adam Smith's "invisible hand", public choice theory, and various flavors of liberalism. While it will be useful in the future to empirically test this assumption, it will be used for this effort.

Finally, and often overlooked, is the fact that the specific set of incentives which define the "game" must be either inferred or assumed. For example, an analyst may assume that the game is a "Prisoners Dilemma" rather than a "Stag Hunt", and will then draw conclusions (e.g., Jervis, 1978) which can lead to policy recommendations based on distinctions which may or may not be operative in the actual policy context (e.g., "offensive" vs. "defensive" weapons). The choice of which game to use is crucial, as one will get quite different results if the game is assumed to be a "Battle of the Sexes" vs. "Stag Hunt" vs. "Prisoners Dilemma" vs. "Chicken" or one of the many other games in the literature (see Hasenclever, et. al, 1997 for some examples).

In a way, the terminology "rational choice" is in itself a value-laden judgment that a person's maximum utility is equal to rationality which is then equal to correct decisions. As Tversky and Kahneman (2000b) point out, normative and descriptive analysis of choice behavior are two different enterprises.

Heuristics and Biases

With the original publication of Tversky and Kahneman's "Judgment Under Uncertainty" in 1974, political science scholars began to use the new ideas of heuristics and biases in an attempt to better understand decision making. Despite the inability of the heuristics and biases paradigm to predict *what* decisions would be made, the ever-expanding menu of various heuristics and biases offered an ever-increasing feast of explanations for why decision makers did not behave in accordance with the tenets of rationality.

Voting behavior has been seen as a rich field for the application of heuristics and biases. Given the limits of information and choice, the voting booth is a prototype arena of bounded rationality. Gant and Davis made first use of this paradigm by linking party affiliation to the representativeness heuristic (1984). This has been followed by numerous applications of existing heuristics and the creation of new voting heuristics. This includes Mondak and McCurley's analysis of what the "coattail" heuristic, where voters were more likely to vote for congressional candidates that were in the same political party as the Presidential winner (1994). More recently, Lau and Redlawsk have proposed five "political heuristics" to explain voter behavior. These included political party, ideology, endorsements, poll strength, and appearance (2001). Interestingly, all of these ideas are simply repackaging "conventional wisdom" about why people vote for a given candidate into heuristics and biases terminology. My other observation here is that the authors have been unable to predict *which* heuristic might be used – instead it takes the form of existence proofs: i.e., voters choose by party, but if they don't do that, they choose by ideology, if they don't do that, they choose by endorsement, and so on.

Moving to national security decision making, there is a long history of trying to predict individual leader behavior through the use of heuristics and biases, such as Nancy Kanwisher's explicit links between various heuristics and their associated biases to several current policy ideas. For example, the "hot hand" heuristic was seen as the basis for the so-called "domino theory", and that the availability heuristic was shown to account for judgments of a given scenario's viability (Kanswisher, 1989).

Another way of analyzing elite behavior is via "operational code" analysis. Nathan Leites (1951) published a seminal study for RAND which attempted to predict the behavior of the Soviet Politburo. Alexander George (1969, 1979) followed up and systematized Leites' original work into a series of philosophic and instrumental questions, the answers to which determined the operational code of a particular leader or group. At this point, it is important to note that neither Leites nor George saw this operational code as a direct link between beliefs and action. Instead, existing operational codes serve to bound the decision space for the particular actor by providing general guidance on key goals or perceived relationships that affect the decision at hand (Walker, 1990). This basic paradigm was expanded by Holsti (1976, 1977) to provide a set of six master operational beliefs. Since then, many authors have attempted to mine that vein for clues about leadership behavior, using more and more powerful software tools. A large number of world leaders and other elite groups have been subject to this type of analysis, in a variety of contexts, such as Jerrold Post's

(2004) compendium of terrorist and “rogue state” leader profiles (see also Walker: 1990 and 2000 for two useful snapshots).

This effort to normatively critique decision making via the heuristics and biases paradigm continues. As mentioned above, Vertzberger brought together an entire menu of psychological factors to explain why decision makers deviate from the rational choice norm. The author combines a variety of factors, including heuristics and biases, scripts and schema, perceptual filtering, social factors, and organizational factors into a grand theory of decision maker behavior. As above, the problem with this approach is the inability to show or predict which factors are operative, only that deviation can be explained by some combination of these factors (Vertzberger, 1990).

As an aside, one of the weaknesses of Vertzberger’s conclusions, which is shared by other theorists in this vein, is that the normative yardstick is unrealistic. Vertzberger states that decisions are correct when they are congruent with the objective facts of the situation (1990). In practice, that seems to mean either a decision that has a good result, or the decision that the author *thinks* should have been made. As our fragmented political system demonstrates, the norms in politics or political science are often quite debatable.

Prospect Theory Approaches

Another alternative to rational choice has been prospect theory (Kahneman and Tversky, 1979, Levy 1992 et seq). As mentioned previously, this theory shows that decision-makers view the utility of gains or losses quite differently, depending on how the decision is framed. This theory has been used to explore a variety of issues. (See McDermott, 2004a and 2004b, for a summary). This theory, while intellectually attractive, has proved difficult to operationalize. Further, as McDermott and others have pointed out, prospect theoretical analysis has often been unable to rule out competing explanations. While the quest has been elusive, one effect, that of framing, has been shown to be fairly robust in its explanatory power. Essentially a refinement of Tversky & Kahneman's original "anchoring and adjustment" bias, the role of framing in shaping subsequent work is highly significant. (See, for example, Berejikian, 2002; Boettcher, 2004; Kanner, 2004; Taliaferro, 2004)

Affective Approaches

Lately, the analysis of political decision making has turned to the "affect" paradigm. The majority of this work has been in the area of voting behavior, where affect and heuristics are often combined in an effort to explain voter preferences, although some authors have studied behavior in other areas of individual decision making (Bueno de Mesquita and McDermott, 2004).

Lodge and Tabor (2000) established a three-part paradigm to explain voter behavior, which combined affect and heuristics. First, the receipt of information is affectively charged, which triggers a number of heuristics and automatic associations about the information. This "connectionist" type of analysis assumes a "semantic network" of connections between concepts, so the triggering of one node also triggers the connected nodes. This method is often modeled using neural networks. Thus, there is essentially a confirmation bias about information on political candidates based on the affective valence of the information. The authors followed this hypothesis up with a laboratory study designed to simulate the complex information environment of the political campaign, reaching the conclusion that affect is both primary and automatic in effecting voting behavior (Lodge & Taber, 2000, 2005).

Redlawsk (2002) expanded on this original concept by adding the online processing paradigm from Pennington and Hastie (1997). Here, voters are hypothesized to keep a constant running assessment of information and affect about a given candidate and issue, and that affect changes the way in which information is processed. Voters are motivated reasoners, but in the affective rather than Bayesian sense, therefore positively valenced information is more readily absorbed into the semantic network, while voters have difficulty with the opposite (Redlawsk, 2002).

Again returning to national security issues, McDermott (2004c) applied the finding of Lerner, et al. (2003) to political decision making, again drawing normative conclusions about the efficacy of decisions taken under conditions of either anger or fear. Further, the author posits, following Lerner, that specific biases can be tied to specific emotions.

Critique of Heuristics, Biases, and Affect

When reading the literature on the affect paradigm, one is struck by the similarity to the heuristics and biases paradigm. They both appear to be robust effects, and can be readily produced in laboratory settings. Further, it is self evident that affect has some impact on behavior. However, the causal chains are often ill-defined and even interchangeable, which implies that even though the effect exists, the overall causal paradigm as currently used can be incorrect. For example, Lerner, et al. (2004), posits two different causal linkages in their study of the endowment effect. The authors propose an "expulsion" cause that is triggered by disgust (get rid of the item), and a “change” cause that is triggered by sadness (change the situation). However, we can readily make a logical case if we switch the causes: it could well be that disgust triggers the desire to change circumstances (in particular, avoid the disgusting situation) and that sadness triggers the desire to get rid of whatever is causing you to be sad – often experienced by country music singers and their ex-wives.

Further, some of the effects are tautological, such as the “how do I feel?” affect, or the “likeability” heuristic. If we say that we like someone because that person is likeable, we have not explained much. Alice Isen likened the impact of affect to the framing heuristic (1993), and that may be a valid way to think about it – affect as just another manifestation of heuristics and biases. These approaches explain real effects, but these must be seen in the context of the surrounding environment. This literature generally mines its ore from a single vein, that there exists some sort of psychosocial function, or series of functions, which mediate information processing and cause deviations from the rational choice ideal. Unfortunately for the theory, each bias or affect seems to rely on a different psychosocial mediating function, and thus it is difficult to explain which affect or heuristic is responsible for a given bias, even post hoc.

In a related vein, the idea of semantic networks has also been controversial in the literature. Gary Marcus notes that such networks raise a number of practical issues. For example, if the network would have the nodes: Brad – bought – green – apples, and then Brad buys green pears the following week, there is a potential issue with crosstalk or ambiguous meaning when comparing the first set of nodes with: Brad – bought – green – pears. Further, there is the issue of multiple instances, what if Brad buys fruit on multiple days? There would need to be additional nodes for today, yesterday, last week, and so forth. Finally, there is the issue of creating connections between previously unrelated sets of nodes. If the owner of the grocery store where Brad shops runs for Mayor, then does a new set of nodes connecting Greg Grocer – candidate – running for – Mayor and the place where Brad shops need to be created? Marcus instead (simply stated) proposes the need for sets of recursive rules which manipulate information rather than a pre-determined semantic network (2003).

Further, it has proven difficult to operationalize psychological approaches in such a way that clearly links heuristics and biases, operational codes, prospect theory, or affect to the actions taken or the policy selected – the actual impact of these factors (where they can be isolated) has tended to be quite modest. Ferguson and Mansbach (2003) sum up the conventional wisdom about psychological approaches – in their view, psychological explanations often force us to abandon the scientific approach in favor of judgment and "stubbornly non-empirical" theories (135).

Just because wisdom is conventional, however, does not mean it is correct. At the aggregated, large-n level, it is entirely possible to assume that psychological factors can be treated as an omitted variable. In that case psychological factors simply become part of the error term, and standard methodological procedures for dealing with the error can be applied. If variation due to psychological factors is normally distributed, then that variation simply adds to the "noise", but does not affect the predictions of the overall model. If variance due to psychological factors is systematic, or as Ariely (2008) puts it, we are "predictably irrational", some way to account for that omitted variable needs to be found.

Information Processing Approaches

In another way to gain traction on the problem, recent cognitive psychology literature has examined the way that the human brain processes information. This literature uses mathematical models of how information is stored and recalled in memory to show that the "biases" demonstrated in the previous literature are in fact caused by the way in which we process information, and particularly by how we store and recall information (see Hintzman, 1988; Dougherty, et al., 1999; Dougherty and Hunter, 2003). The clear implication of this literature is that the process of memory itself creates the many observed biases, and not some hypothetical internal library of heuristics, or pre-determined semantic networks.

While a significant portion of the literature on foreign policy decision making deals with group processes, the focus of this effort will be on the individual process of decision making. As Deutch (1966) notes, the process of decision making requires the existence of actual decision points. This notion has been reiterated by Bueno de Mesquita when he stated that "…states are only metaphorically decision-makers. States do not, in actuality, choose policies or have goals; leaders do." (2002, 8)

James March usefully divides individual decision-making approaches into two categories: those involving rational choice, and those involving rule following – or the logic of calculation vs. the logic of appropriateness. He notes that these two approaches make quite different demands on the decision maker. The rational choice approach requires the consideration of goals, how actions will impact the attainment of those goals, along with actions that will likely be taken by other actors – a demanding task. The logic of appropriateness requires recognition of the situation and some mechanism for selecting which rule to apply – a cognitive task that can be efficiently performed using simple heuristics (March 1994).

Conceptually, it can be seen that the second approach might be more attractive. The rational choice requires a lot of complex calculation on the part of the decision maker – Clausewitz quotes Napoleon approvingly that “the decisions faced by the commander-in-chief resemble mathematical problems worthy of the gifts of a *Newton* or an *Euler*” (1989 [1874], 112 – italics in the original). The incidence of geniuses of that rank among the world’s leaders is most likely just as small as it is among the general population, and it appears unlikely that the average decision maker is capable of performing those calculations well – if he or she can perform them at all.

Further, the rational choice paradigm draws its underlying ideas from economics, and thus incorporates the underlying constitutive logic of the marketplace. While rational choice theorists, as shown above, have argued by analogy that conditions faced by decision makers in the political or foreign policy realm can be mapped to conditions in the marketplace, such argument can be suspect. Even Gilpin, in his seminal explication of the idea of political economy made the clear distinction between the logic of the market and the logic of the state (1987). This distinction has been often raised by constructivist theorists, who draw the distinction between rational or utilitarian theories that draw their micro-foundations from economics and those theories which draw their micro-foundations from social interaction between actors (see Adler 2002 for a summary).

Results: A Decidedly Mixed Bag

As seen above, neither realist nor liberal theories predict emergent properties or change within the international system. It was the failure of these theoretical approaches to predict or even account for the upheaval of the end of the Cold War that showed the value of approaches such as constructivism. Constructivism, however, offers only a menu of general approaches to the idea of change, and lacks the tools to precisely define and test the linkages between interaction and subsequent systemic change.

Aside from the above observations, the record of any theoretical paradigm or approach mentioned above in actually predicting the outset of war is decidedly mixed. Bennett and Stam recently conducted a large-n comparison of a wide variety of realist, liberal, state-level and rational choice formulations, and found that most approaches did not produce strong correlations between their theoretical predictions and the outset of war, and those with significant correlations "have no direct linkage to the causal processes generating the correlations" (2004, 207). For example, while the "war equilibrium" state of Bueno de Mesquita and Lalman's (1992) expected utility measures roughly doubled the risk of the onset of war compared to the "status quo" equilibrium (and was thus statistically significant given the very large n of the study), the actual probability of the onset of war was .00006 for a "status quo" equilibrium and .00013 for a "war" equilibrium (Bennet & Stam 2004, 132). This result, albeit a statistically significant correlation, is hardly diagnostic and certainly does not rise to the level of practical decision-making significance.

Further, war itself is an emergent event, where the existing order can be upset and new follow-on conditions are created. All of the approaches mentioned above have difficulty with predicting the *existence* of emergent behavior, so say nothing of predicting the *timing* or the *outcome* of that emergent behavior.

Needed: A New Approach

All of the IR and decision-theory literature has to take into account that, for all the complexity of the specific situation, that in general, the actual decision is shaped by both the situation and the process of decision. Further, decisions for war or for alliance occur rarely – and sometimes never for some states.

With rare events, it is tempting to simply bet on the *status quo*. After all, that is what happens most of the time. If considering playing the lottery, a simple check of the odds shows that the *status quo* and most likely outcome is that the player will not win – and so the best bet is simply to keep the money and not play. In predicting the choice between war and peace, simply predicting "peace" is correct over 95% of the time, so high accuracy solutions may not be the best.

Instead, the challenge is far tougher, to predict when unlikely events will happen and which actors will be involved, without the requirement to make untenable assumptions about either the structure of the game or the interests of the players.

This returns me to the main line of analytical approach, that of the interaction between situation and how decision-makers process information, which can be modeled from the cognitive angle. The application of this framework will be shown to provide a way of not only understanding this interaction, but using that interaction to predict the actions of decision-makers, given a specific set of inputs.

Chapter 3: An Indirect Approach – Memory and the Rules of Decision

Introduction

One of the most common descriptions of what it's like to be a pilot says that flying is "hours of boredom punctuated with moments of stark terror". This is also true when considering the history of war and alliance. Both of these are fairly rare, and states have existed without experiencing either. A simple look at the Alliance or Militarized Interstate Dispute datasets reveals that peace and non-alliance are far more likely than conflict and alliance.

The challenges this raises for other approaches to international relations and decision-making in particular have been discussed in previous chapters, so I will not dwell on this point, other than to reiterate that understanding the relationship between the situation and the process of decision – how these two interact to produce a particular result – is critical to understanding the phenomenon of why these unlikely events pop up from time to time.

In the previous chapter, I briefly looked at the idea of punctuated equilibrium. This refers to the property of the interaction of complex adaptive systems where long periods of relative calm are punctuated by periods of upheaval, to be followed by a new, often changed, equilibrium. One of the properties of this type of situation is that a certain point exists where the system becomes unstable. At this point, the behaviors that produced equilibrium suddenly start to produce disequilibrium.

To take a simple example, consider the behavior of a car in a tight turn. Normally, when the turn is not tight enough, the driver can simply turn the steering wheel a little bit tighter, and the car will respond. Turn the wheel too much, and the car enters a skid, where turning the wheel more will simply cause the skid to get worse and the driver will lose control. The driver must take corrective actions to consciously dampen the out-of-control situation by turning the wheel in the opposite direction.

The point at which this occurs is sometimes referred to as the departure point, as the car is in departure from a controlled path. From dynamics, the car has transitioned from a position of stability to a position of instability. Of course, the situation in international relations is far more complex, and the process of changing from peace to war is not immediate. Nevertheless, it is fairly rapid, as can be seen from consistent observations in the literature on crises and the outbreak of war of things like "events were gathering steam" and "we could feel things slipping away".

Similarly, the well-known simulation of rabbits and hawks exhibits the same characteristics. In the simulation, rabbits eat grass, hawks eat rabbits, and both reproduce. Populations grow when food is plentiful and vice versa. Normally, this system settles into long-term states of equilibrium, but a slight change in the food supply or reproductive rates will cause the system to rapidly become unstable with alternating population booms and busts until the system finds a new equilibrium. Note that the exact same behaviors and choices (eating and reproduction) on the part of actors within the system lead to both stability and instability, depending on the situation. This situation is equally true with more complex adaptive systems, such as the one used in this dissertation.

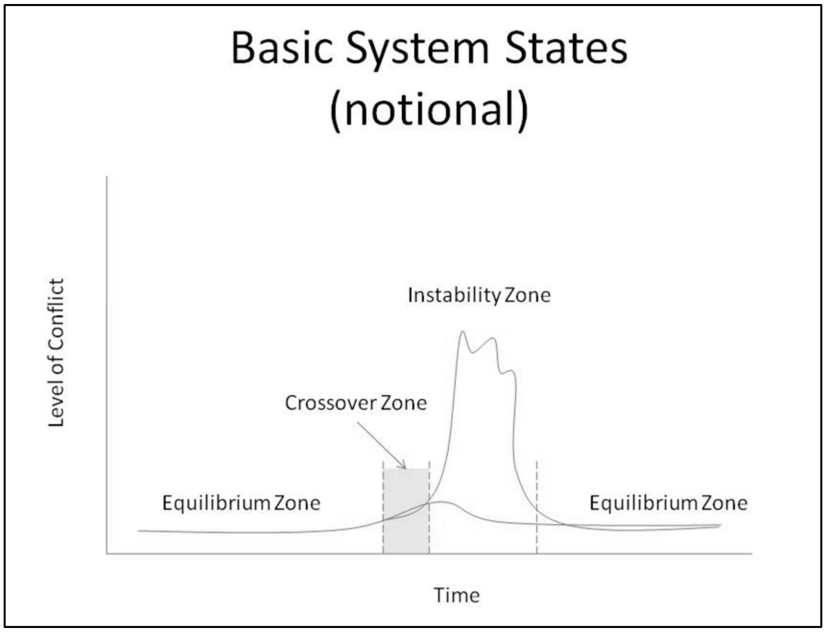

Figure 1: Notional System States

Looking at the outbreak of war as a problem of punctuated equilibrium, or of departure from the position of stability, I posit three basic states, as shown in Figure 1, above: the *equilibrium zone*, the *instability zone*, and a *crossover zone*, where the system is starting to become, but not fully yet unstable. Once events reach that point, continuing actions that seem reasonable and prudent can instead lead to disaster.

It is one thing to posit such a relationship, it is another to calculate and predict it. That decisions are based on the situation is obvious – but how? The literature contains all sorts of possible reasons, mostly descriptive or normative. Some, such as game theory and expected utility modeling, require assumptions about the game structure or judgments about actor preferences, power, and salience of issues.

Results of decisions create new reality. This interaction is also critical to the path of future decisions. To attack this problem, I have created a model of state interaction based on the cognitive processes by which humans remember and recall information. State decision-making is, in the end, human decision-making – someone has to make a decision. Assessment and decision are both human enterprises, and attempts to model those enterprises should proceed from how information is actually processed rather than how one might normatively wish to have it processed. To provide a tool to model this concept, a Memory-based Ecological Model of Relationship Interactions between States, version 2 (MEMORIS-2) has been developed by the author.[1] The MEMORIS-2 model is an extension of memory trace models such as MINERVA-2, MINERVA-DM, and HyGene (Hintzman, 1988; Dougherty, et al., 1999; Dougherty and Hunter, 2003, Thomas, Dougherty, Springer and Harbison, 2008). MEMORIS-2 is implemented using Mathematica version 7.0, and extends these models to work in a dynamic multiple-actor environment necessary to support research in state decision making in a complex interactive system. As a proof of concept, a prototype MEMORIS model was used to simulate state perceptions in Europe between 1885 and 1915 as a brief case study to validate the basic structures of the model (Hanson 2006).

The MEMORIS-2 model rests on two foundations. First, it proceeds from Simon's observation that, far from being the beneficiaries of some Laplacean daemon[2], both individuals and organizations in the real world are required to make

---

[1] I must acknowledge the invaluable guidance and help of Professor Rick Thomas, who provided guidance on memory-based models, algorithms, and computer code, along with jump-starting me on programming in Mathematica.

[2] A Laplacean daemon is a mythical creature with the infinite time, resources and processing power needed to make a purely rational choice decision.

decisions in the realm of bounded rationality. This means that neither organizations nor individuals have the unlimited time, processing power, or resources required to make purely rational calculations. Similarly, and sometimes overlooked, the solution space is also bounded: only certain decisions or behaviors are feasible (Simon, 1945 [2000], see also Verba, 1961).

The second foundation is a broad extension of another of Simon's ideas: organizations exist to process information with the aim of making a decision (1945 [2000]). At the same time, this process of decision is not some mechanistic or disembodied process. This implies the need to understand how organizations, and by extension, individuals, process information. Much of the discussion of this issue has been conditioned on the premises of rational choice theory, which implies an economic mode of decision where individuals survey the situation, rank order the available choices, and then choose the choice which offers the "best" utility (however defined). As noted above, this approach has its limitations, not the least of which is that individuals cannot be expected to operate in consonance with rational choice theory. It has, however, been the normative yardstick for most analysis of information processing and subsequent decision making. As Jervis points out, however, an approach to decision-making theory should have some resemblance to how people actually make up their minds on serious issues and should be based not solely on perceptions, but on how information is actually processed (1976).

Conceptual Basis of the Model

Building from this foundation, the challenge is then to understand just how decision-makers process information within a bounded framework. Based on the cognitive psychology literature noted above, which empirically demonstrates how biases can be replicated using conditional memory trace models, it is unnecessary to posit a multitude of psychosocial mediating functions. Instead, using this increased understanding of the process of memory, and thus understanding how judgment of a given situation can be determined by the current ecology, the sample of information that is stored in memory, and the cue, or hypothesis which is used to probe memory. Importantly, these processes can be explicitly modeled, and judgments can be predicted probabilistically based on human information processing mechanisms without recourse to either the untenable assumptions of rational choice theory, or the need to discern the operative heuristic or affect. Further, the simulation is based on a recursive processing model of mental function rather than the connectionist models mentioned in the previous chapter.

Further, this approach does not require any *a priori* assumptions or judgments about the game structure. Instead, the actors draw conclusions about the game structure based on the information presented and adjust strategy accordingly. As will be seen, the simulation automatically adjusts both strategy and player preferences to the structure of the game and to the actions of other players.

Sampling the Environment

Since decision makers operate in a bounded decision environment, and have limited ability to perceive that environment, they essentially have to work from the information that is available. The availability of information is conditioned by the "natural" sample space of the surrounding ecology – that is, decision makers sample and categorize information in the environment according to fairly simple classifications based on gross features, such as: friend or foe, powerful or weak, and so forth (see Fiedler, 2000 and Gavanski & Hui, 1992). Since decision makers do not have perfect knowledge of their surrounding environment, information is generally tentative and is continually updated in a probabilistic fashion, as are tentative sets of hypotheses about the environment – the more data points which become available, the higher the confidence in the judgment (Gigernenzer, et al., 1991). It should be pointed out that the samples can either be representative of the ecology or biased, and as such the sampling process can produce both adaptive and maladaptive results.

This set of samples essentially forms a set of records, or traces, in the decision maker's memory. These traces are the record of experiences, which may or may not be biased or misperceived. Part of this record is the judgments associated with that set of experiences. As pointed out, these traces are imperfect, and are sometimes contradictory in specific detail, but they build up over time to create a general understanding of the environment.

Since hypotheses are generated based on their probability, this implies that the set of traces in memory conditions the set of possible hypotheses. Further, the search for information within the environment is conditioned by the hypothesis under consideration, which is determined by reference to semantic concepts.

Development of Semantic Concepts

This search is guided by an existing set of semantic concepts. These semantic concepts are simplistic representations of the situation. For example: "falling behind" another state in a given realm, such as military power, or that the state is an "enemy" or an "ally". These semantic concepts can be combined to give nuanced judgments of a given situation, such as the obvious difference between “falling behind an enemy” and “falling behind an ally”.

Further, unlike traces of memory which can be repeated, semantic concepts are singular. That is, there is only one exemplar concept of “falling behind”, one concept of “enemy” and so forth. Thus, the judgment of whether a different semantic concept is active is dependent on the strength of the response from a probe of previous memories, but the *existence* of the concept does not rely on memory. This is important because it allows the decision maker to be able to access semantic concepts for which there are no previous experiences in memory. These concepts could be learned by observing other states, through history, through reasoning, or other means. So, even though a particular decision maker has never directly experienced “getting ahead”, that semantic concept is available for judgment. Allow me to reiterate the point that “semantic concepts” are not the same thing as the semantic networks discussed in Chapter Two.

The case of the Prisoners’ Dilemma (PD) from game theory can be a useful example to illustrate the idea of semantic concepts and how they can help shape decision making. In this situation, two known criminals are apprehended near the scene of a crime. Following convention, the first prisoner is male (Bob) and the other is female (Alice). The prosecutor immediately separates the prisoners and offers each the following deal: you can either remain silent (cooperate with the other prisoner) or implicate the other person (defect). If both Bob and Alice cooperate with each other and remain silent, then they will each receive a light sentence (6 months) for loitering. If Bob implicates Alice while she remains silent, he will be set free and she will receive a heavy sentence (10 years), and vice versa. If both prisoners implicate each other, they will receive a moderate sentence (3 years). Being set free is then the temptation to defect (T), the light sentence is the reward for cooperation (R), the moderate sentence is the penalty for mutual defection (P) and the heavy sentence is the sucker's payoff (S). As a notation convention, Bob's moves will be put in upper case, while Alice's move are in lower case, so if Bob defects and Alice cooperates, the move will be noted as (D, c) and so forth. Taking as the general principle that both Alice and Bob prefer to spend as much time out of prison as possible, this relation can be written as T > R > P > S for Bob and t > r > p> s for Alice, where ">" means that the first result is preferred to the second. Put into tabular form with numerical utilities assigned to each choice, where:

$$U_T = 5, \quad U_R = 3, \quad U_P = 1 \text{ and } U_S = 0 \qquad (1)$$

The canonical view of the PD game is as follows:

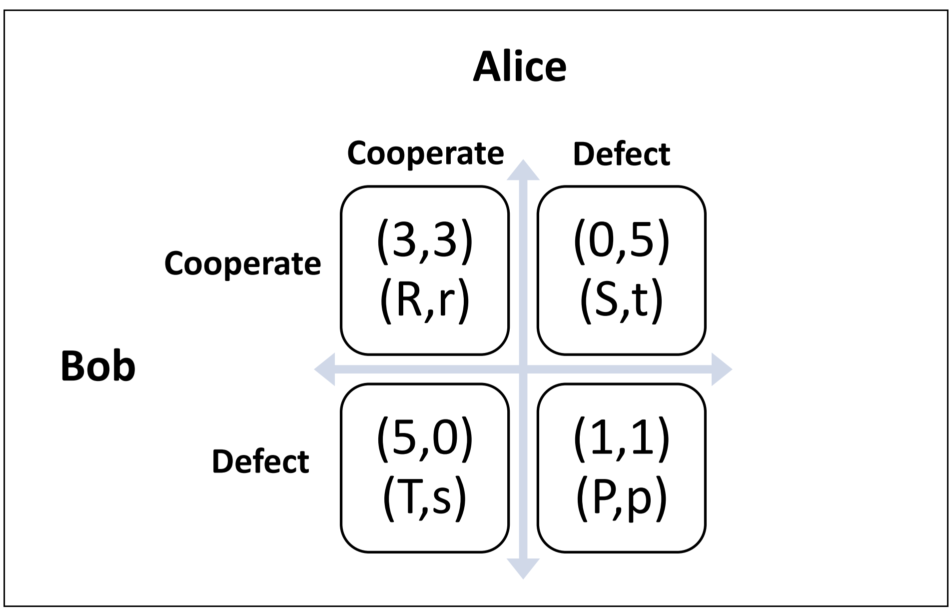

Figure 2: Prisoners' Dilemma Game

One of the points raised by this particular game formulation is that, in a 1-time game, both sides are better off defecting, even though that results in the next-to-last favorable result for both sides. Looking at the game from Bob's viewpoint, no matter what Alice does, he is better off to defect (if the play is (D, c), Bob's payoff is 5 vs 3 if both sides cooperate (C, c), and if Alice defects and Bob defects (D, d), Bob's payoff is 1 vs 0 if Alice defects and Bob cooperates (C, d)). The same is true for all of Alice's moves, so both sides have a dominant strategy of defection. This situation, where both sides defect, and neither side has an incentive to change strategy (or to put it differently, each side is indifferent to the other's strategy) is called a Nash equilibrium (see Myerson, 1991 [1997] for a rigorous derivation).

The PD game also has a Pareto optimal strategy, which maximizes the total utility at the point where both sides cooperate. At this point, neither side can change the strategy without decreasing the total utility of the game. This (C, c) strategy is also a potential equilibrium strategy in repeated games, given certain conditions of either information, utility, time horizon, or discounting of future moves (Morrow 1994). (See, however, Kreps, Milgrom, Roberts and Wilson, 1982, along with Binmore, 1994 for a discussion of multiple equilibria in iterated PD games)

The PD game can also be plotted in terms of its *decision space*, which here is graphically depicted as the results of various combinations of strategies. This way, the possible set of payoffs given various pure and mixed strategies can be seen (Figure 3, below). Thus, the set of payoffs in repeated PD games will fall within the figure $\{T,s\},\{R,r\},\{S,t\},\{P,p\}$. Further, the set of payoffs for pure strategies will lie along the edges of the figure, so that if Alice always cooperates (always C), the set of payoffs for both players will lie along the line {T,s},{R,r} and if Alice always defects (always D), the set of payoffs will lie along the line {P,p},{S,t}, and so forth. If both players pursue mixed strategies, the set of payoffs lies in the interior of the figure. In accordance with the so-called "Folk Theorem", an iterated PD game will also have a set of cooperative equilibria that lie within the square defined by (P, p) and (R, r) (Binmore 1994).

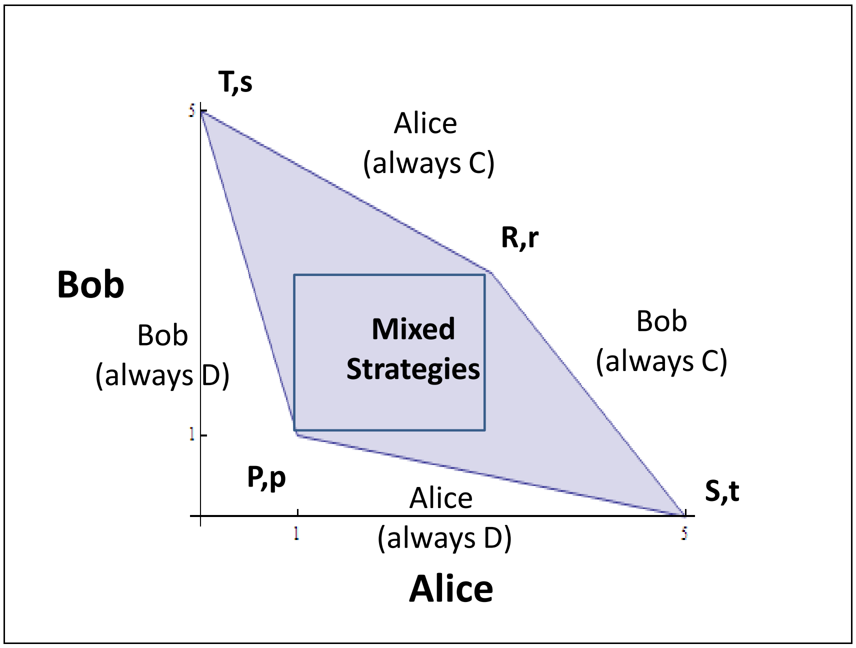

Figure 3: Decision Space Representation of the PD Game

With this diagram in mind, consider the case where both Bob and Alice, being the well-known criminals that they are, are used to being in prison from time to time, and each one thinks that any prison sentence of a year or less is acceptable, but a sentence of more than a year is not. Since in the canonical PD formulation $U_R = 3$ and $U_P = 1$, let the utility of a 1 year sentence be:

$$U_{1\ yr} = 2 \qquad (2)$$

This new utility allows the creation of 2 semantic concepts *SC*: "satisfied" and "unsatisfied", which are denoted by SC_{sa} and SC_{un}. Given a result A:

$$U_A \geq 2 \rightarrow SC_{sa} \quad \text{and,} \quad U_A < 2 \rightarrow SC_{un} \qquad (3)$$

Adding this to the decision space diagram gives the following:

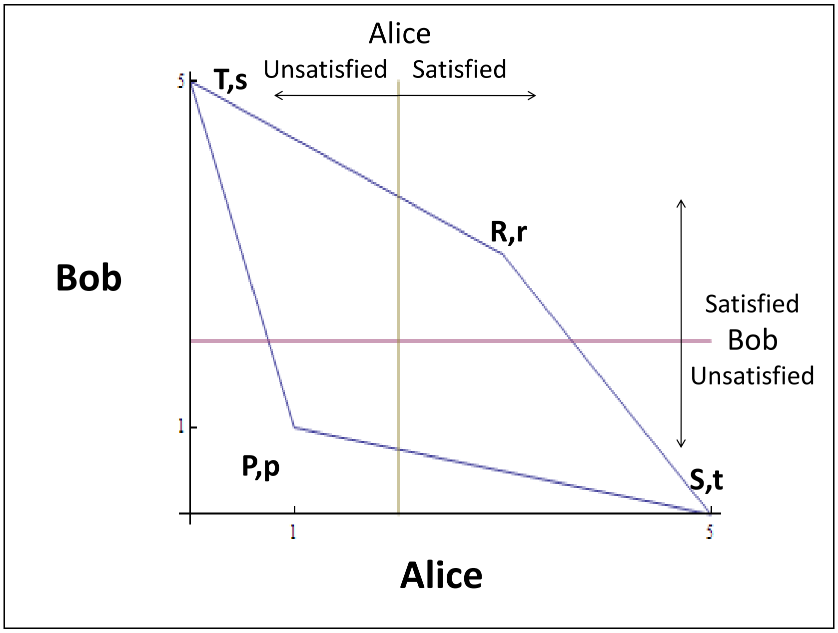

Figure 4: Semantic Concepts Applied to PD Game

Thus a result to the left of the line where Alice's utility is 2 (prison sentence greater than one year) will be unsatisfactory, and likewise for Bob. This greatly constrains the set of likely strategies, as both players are apt to change their course of action if the previous course leads them into a situation where they are unsatisfied. In fact, consideration of the likely options will drive the players into the "cooperate, cooperate" strategy, which as discussed above, is Pareto optimal for both sides. Further, it will create a natural "tit for tat" result if either side defects. If, for example, the sides originally choose to cooperate, and then Bob defects while Alice continues to cooperate, that result will cause Alice to conclude that the situation is unsatisfactory (Point T,s → SC_{un} for Alice and SC_{sa} for Bob). This will cause Alice to change her move (defect) while Bob continues to defect. The result of (defect, defect) is P,p → (SC_{un} for both Alice and Bob). This will cause both sides to change back to R,r → (SC_{sa} for both Alice and Bob). See Figure 5, below.

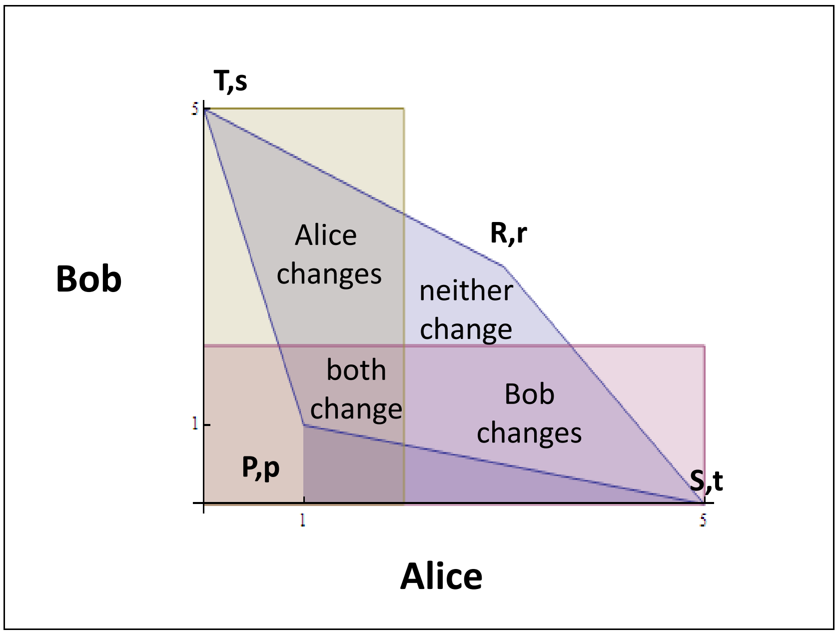

Figure 5: Changes in Strategy Based on Semantic Concepts

This, of course, does not rule out the development of new semantic concepts through experience. This might occur in two modes. The first is a situation where no existing semantic concept provides a close enough match to the given situation. In that case, a decision maker may attempt to develop a new semantic concept which better describes this new situation, and then use this new concept to probe memory in an attempt to confirm this new concept.

For example, say that Alice knows for sure that a third person (Charlie) has committed an unsolved crime and has the ability to prove it. She offers to turn state's evidence against Charlie in return for her freedom no matter what Bob chooses to do. In this case, she now has no incentive to take any particular action because she will always go free. Further, she will have no way of knowing what Bob did since her result is always the same (freedom). Indifferent to the courses of action, she is just as likely to do one as the other.

This puts Bob in a radically different situation – he cannot predict or infer Alice's action, so it is as if he were faced with a random event rather than another player. He can play either strategy, and it will have no effect on what Alice does even in the long run. His previous experience tells him to play "tit for tat" in a repeated game when Alice defects in the expectation that the game will return to a satisfactory state, but playing tit for tat has no impact if Alice is simply playing randomly. This may well lead Bob to add a new semantic concept: that of Random Play (SC_{rand}) and to adjust his play accordingly. Additionally, having come up with SC_{rand}, Bob might also reason that a complementary semantic concept exists: that of Rational Play (SC_{rat}) even if he was not experiencing that particular concept in action.

So, if an existing set of semantic concepts, such as SC_{un} and SC_{sa}, does not closely enough describe the situation faced in a game, a new concept (SC_{rand}) might be needed. This, in turn, could lead to the development of a complementary concept (SC_{rat}) purely through reasoning about the situation.

This, naturally, begs the question of what constitutes "close enough". This would most likely depend on the situation – in some situations "close enough" could cover a very wide stretch of ground, for example whether one is “getting ahead” or “falling behind” an ally is probably a moot point and has a wide margin for error. The same situation for an enemy state, on the other hand, probably has a much smaller margin for error. The acid test here is in the results. Did the judgment lead to a successful or an unsuccessful choice of policy? If the result is good, then future recollections of the event could either conclude that either the judgment was correct (despite doubts at the time), or that the judgment was incorrect (but it really didn’t matter anyway). Either one of these cases would reinforce existing semantic concepts and argue against the need for developing a new semantic concept. This, naturally, may well have negative long-term implications as a wide margin for error or suppression of doubt might lead to eventual disaster.

Given the above discussion, it appears more likely that a policy disaster would lead to new semantic concepts than a simple “it doesn’t quite fit” judgment. If the results of policy are bad, then future recollections could conclude that either the judgment was correct (but the policy was bad, or the breaks went bad) or the judgment was incorrect and should have been something else. The following decision tree diagram shows this in a more compact fashion.

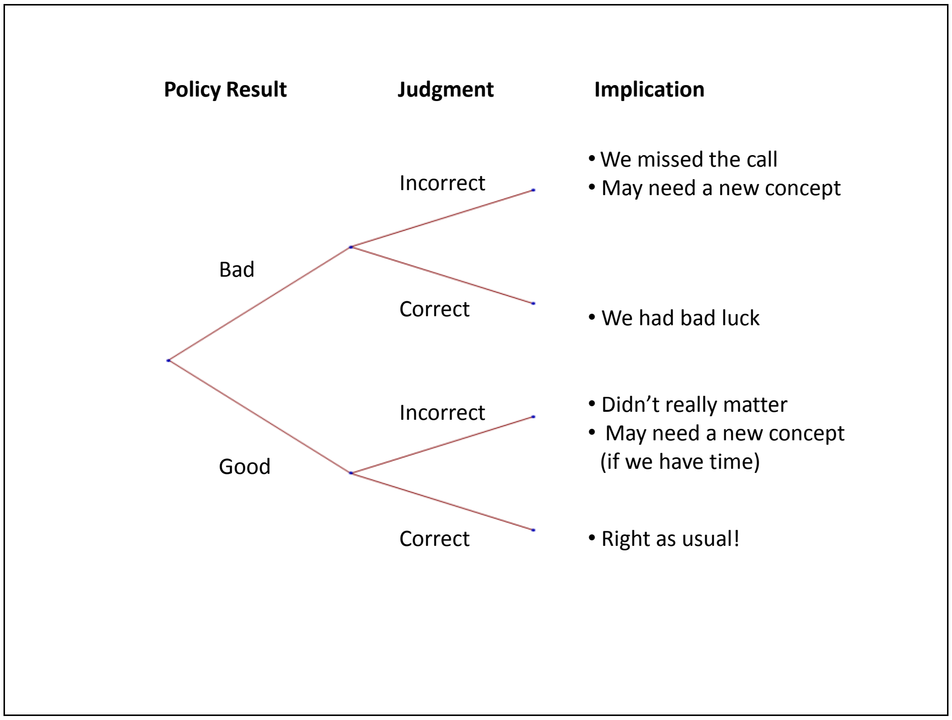

Figure 6: Does the Situation Require a New Semantic Concept?

The introduction of new factors into a game will most likely require new semantic concepts to be developed by the players based on both their experience and reasoning. To further illustrate this point, consider the situation illustrated by the "Stag Hunt" (SH) from game theory. This game, which is based on a situation described by Rousseau, gives the two players a choice. They can either choose to hunt stag, or they can hunt rabbits. To successfully hunt stag, both players must hunt stag (C, c), and both will reap a high reward (R, r) for doing so. Either player can hunt rabbits successfully (D, d), and will reap a small reward (P, p). If a player hunts stag while the other hunts rabbits (C, d), the player attempting to hunt the stag alone will go home hungry while the player hunting rabbits will be successful and reap a small reward (S, t). For a stag hunt, the preference order of the game is:

R > P >= T > S for Bob and r > p >= t > s for Alice, (4)

where >= means "either prefers or is indifferent to"

Put into strategic form, one version of the game looks like this:

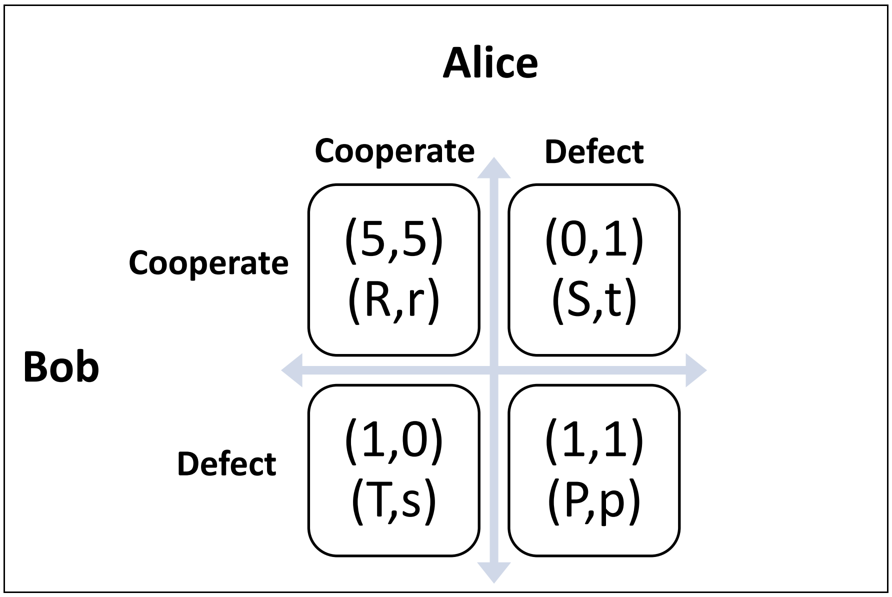

Figure 7: Stag Hunt Game

With this set of preferences, it can immediately be seen that there is no single dominant strategy, as was the situation with the PD game. Here, if Alice chooses to cooperate, Bob's best strategy is to also cooperate (C, c). On the other hand, if Alice chooses to defect, Bob's best strategy is to defect as well (D, d). Thus, there is no pure Nash equilibrium strategy, although the strategy (C, c) is Pareto optimal, as is the same strategy in the PD game. With this situation, for a one-time game, players should play a mixed strategy. Like the PD game, (C, c) is an equilibrium point for a repeated game, and "always defect" is a valid strategy to minimize risk, since it keeps the player from getting the sucker's payoff and going home hungry.

Repeated Play

As mentioned briefly above, the "best" strategy in any game may change if the game is repeated. In this case, either Alice or Bob must somehow account for the fact that actions in one round of play may affect decisions in subsequent rounds. Thus, factors such as reputation, past actions, and the degree to which the players value (or discount) future payoffs may affect the strategies that players will choose to play. As briefly discussed previously, the degree to which reputation affects game play is still controversial in the literature. The following section will discuss the issue of how memories of game play are created.

The Role of Memory

Continuing with the discussion of multiple plays of this combined (or any other) game, the set of individual moves plus their result forms a series of memories, or traces within memory, as shown in the Venn diagram below (Figure 8). The most recent trace, which is the current state of affairs, is compared against the existing set of semantic concepts, and the most likely concept(s) are then brought into working memory as active hypotheses, where they are compared to the traces in memory to either confirm or disconfirm the likely hypothesis. The hypothesis with the greatest support within memory is then chosen as the judgment of the situation, and acted upon accordingly.

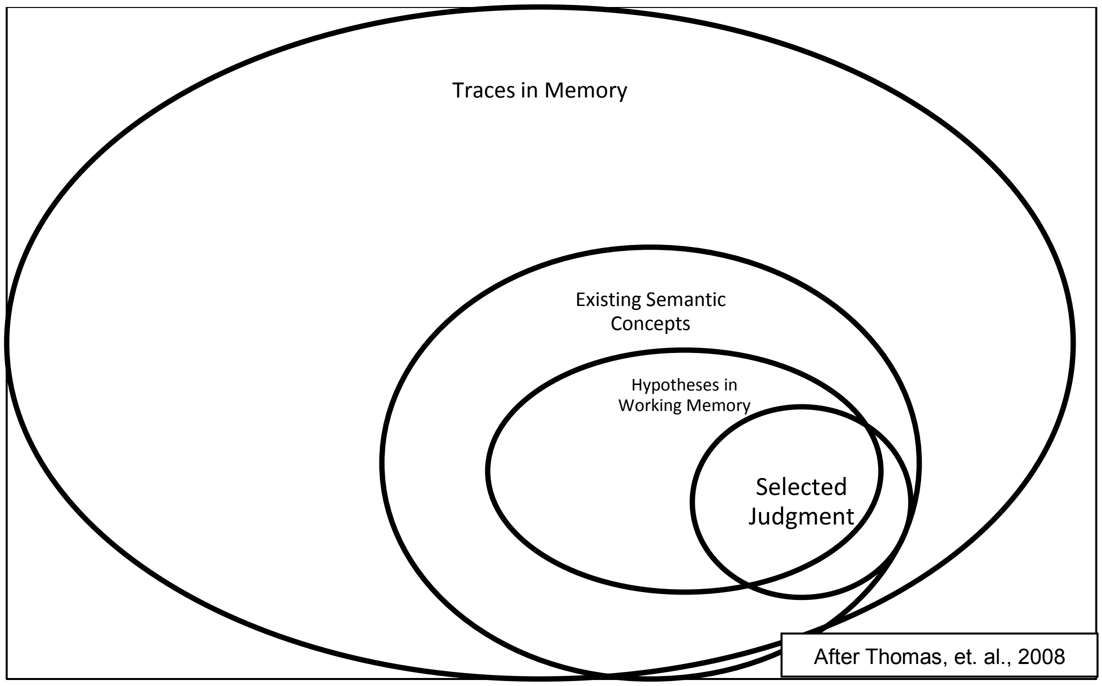

Figure 8: Memory, Semantic Concepts, Hypotheses and Judgment

Note that, as pointed out above, the set of existing semantic concepts is largely drawn from experience, but some may have been drawn from reasoning about the situation or through other means. The hypotheses about the current event are drawn from the store of existing semantic concepts, and the selected judgment will either be drawn from the set of hypotheses immediately selected into working memory, or may be drawn from a new hypothesis if none of the existing set of hypotheses can be confirmed.

Essentially, the representation of the situation that is written as traces into memory is “probed” for confirmation of the hypothesis. Additionally, since the set of traces is imperfect, different hypotheses will lead to different recall of information. This can explain, for example, why different individuals (or groups, for that matter) can come up with completely different interpretations of the same set of events. As will be seen in the application of the MEMORIS-2 model, decision makers may well have very different interpretations of the same data, with one side of an interaction seeing benign intentions and the other seeing hostility. The ecological sample and traces in memory can be very similar, but different probes will elicit different sets of results, and the follow-on search for information will be conditioned to confirm the initial hypothesis.

Summary: Information for Decision

Decision makers thus will receive information in two modes. The first is based on the events and interaction of the surrounding environment. This mode of information forms the sample and the basic traces of memory. The second mode of information is that of a guided search for confirmatory information, which is conditioned by the goal or hypothesis currently under consideration. Given the limitations of individual working memory capacity, the first step of the situation assessment process is to determine the focus of attention. In reality, this focus may often be thrust on the decision maker by conflict or crisis. Once the decision maker's attention is focused on a particular concept, the decision maker then uses existing views of the semantic concept under review to guide the search for information to confirm or disprove existing judgments of that concept. Once these judgments are updated, they become part of the "memory" of the event, so the subsequent memory trace contains both "facts" of the given situation and "values" or judgments associated with that event. These value judgments will then form the starting point for judgments about future events.

Operationalizing the Theory

A Two-Sided Prisoners Dilemma Game

With that groundwork in mind, I created a simulation to extend the idea of using semantic concepts in a two-sided game. In keeping with a building-block approach, this model is used to explore a two-sided PD game, and will be extended to the case of an alternating PD and SH game below. First, the simulation will simply add semantic concepts to the game play, and later the simulation will be extended using a memory-based component.

In this case, the logic of the process for each side is essentially a mirror image of the other. This time, the payoff function is that of the canonical PD game, as presented previously in Figure 2. The process flow of the model is as follows:

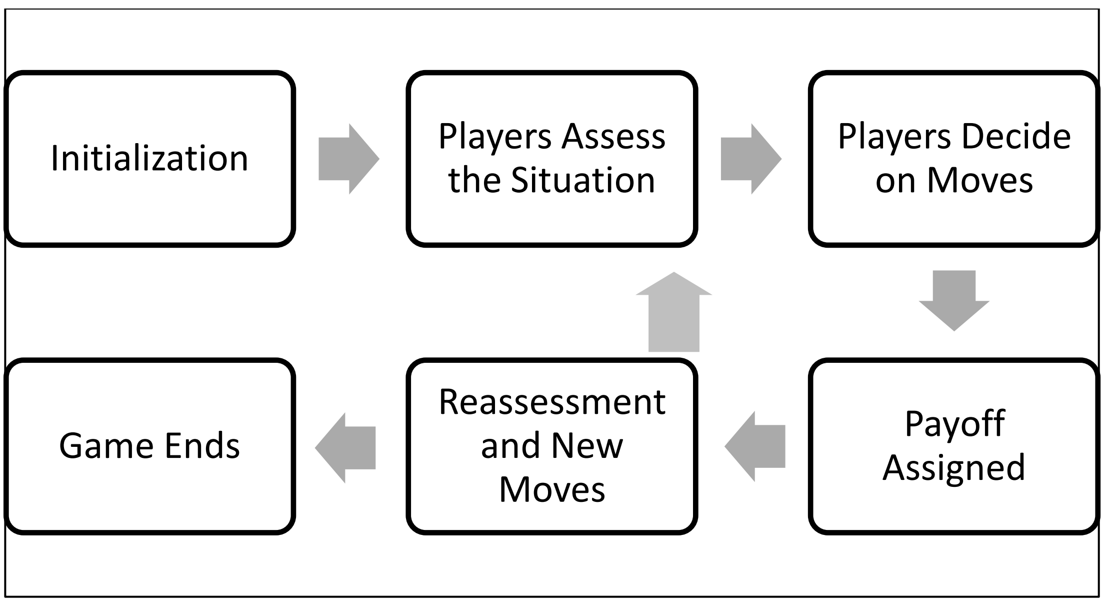

Figure 9: Basic Game Flow

First, the model is initialized by creating a set of 32 random moves. This allows for several repetitions of each possible move combination of the PD game.

This initialization process allows the sides to have a set of memory traces for each move combination, and the random walk process gives the sides a somewhat different starting position for each repeated set of runs.

After the game is initialized, the players assess the situation at the start of the game. For the first set of simulation runs, only two semantic concepts, those previously identified as SC_{un} and SC_{sa}, will be used. As discussed earlier in the PD game description, a dividing line of 2.0 will be used, and so a payoff greater than 2.0 will be judged as satisfactory and payoff less than 2.0 will be unsatisfactory.

The players use a simple logic for this simulation, where if the situation is satisfactory (SC_{sa}), then the players will keep the same move, but if the situation is otherwise (SC_{un}), the players will reverse their move. After the players simultaneously choose their moves, payoff is assigned. Players will then repeat the process until the game ends at some point. Players do not know the length of the game, so they cannot choose to defect at the end of the game to gain advantage. While the number of simulation moves is finite, to the players it is essentially a game with an infinite horizon.

As expected from the simple logic and the previous discussion about the PD game, this leads to a case where the players move immediately to cooperation and do not defect. For 1000 trials, this leads to the following result:

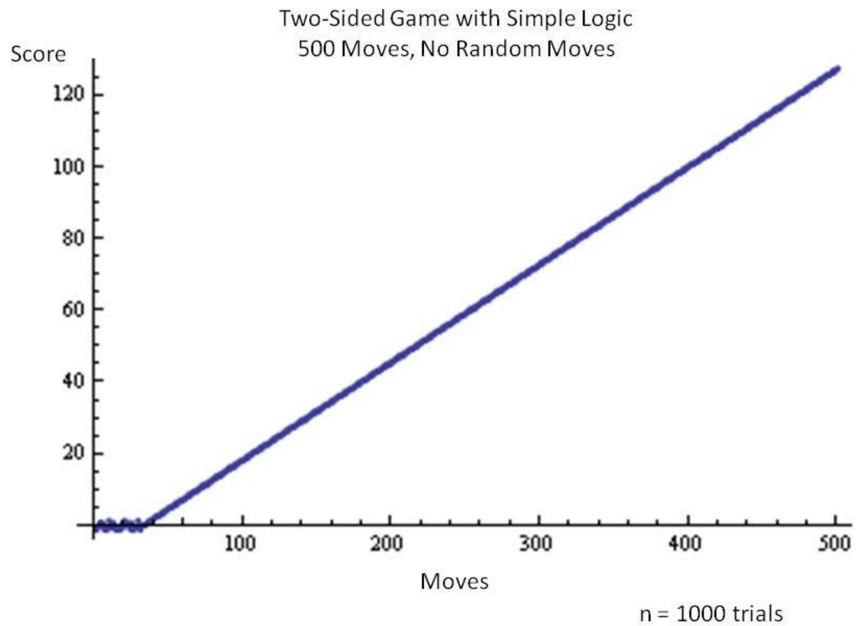

Figure 10: Results of the Two Sided Game with Simple Logic, No Random Moves

Adding a random move parameter, which effectively allows players to defect even though they ordinarily would not, adds some dispersion to the results as expected, but overall the players still prefer to cooperate. As can be seen by the right-hand chart in Figure 11, as the percentage of random moves goes much past 25%, the game will essentially become a random walk between the players.

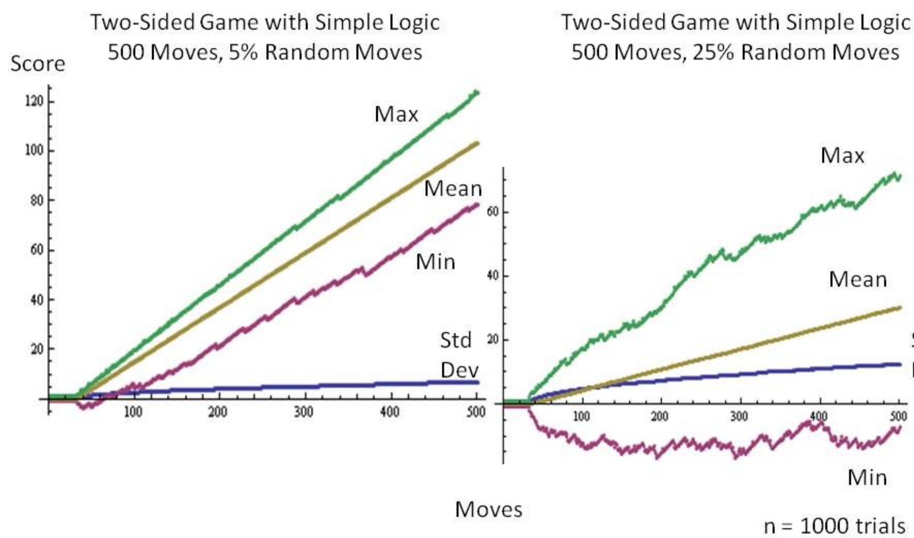

Figure 11: Simple Games with Random Moves

Adding uncertainty to the payoff function provides some interesting results. Until the level of uncertainty approaches the level at which payoffs for various moves begin to overlap, the payoff and play remains fairly robust. Once it becomes possible for the payoff of one move to overlap another, possibly changing an "acceptable" result to an "unacceptable" result, then the play begins to break down and the players soon are moving almost randomly. Combining payoff uncertainty with random moves creates much the same effect, although the combination of both factors causes the game to break down increasingly quickly.

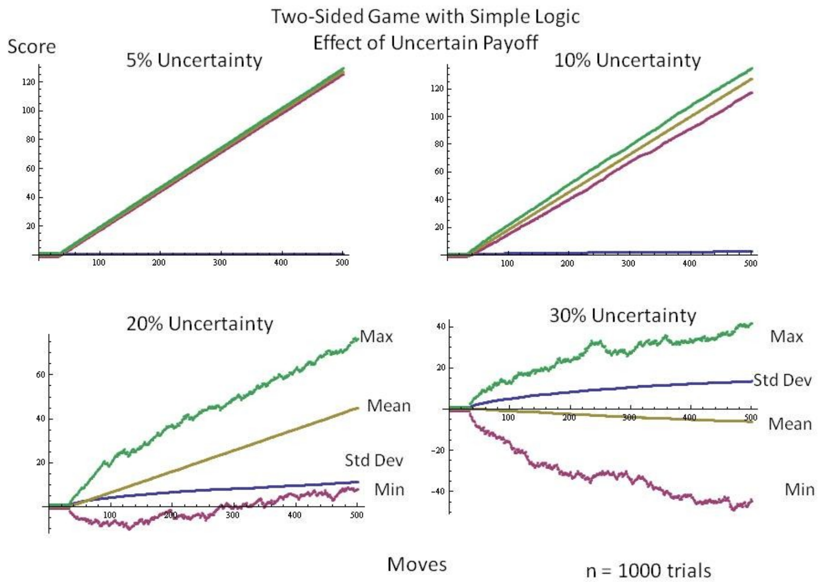

Figure 12: Effect of Increasingly Uncertain Payoffs

Changing the L (learning) parameter also creates considerable uncertainty within the simulation and causes it to quickly become unstable as well. Incorrectly encoding the results of moves causes the uncertainty to effectively overlap even at relatively high rates of correct encoding.

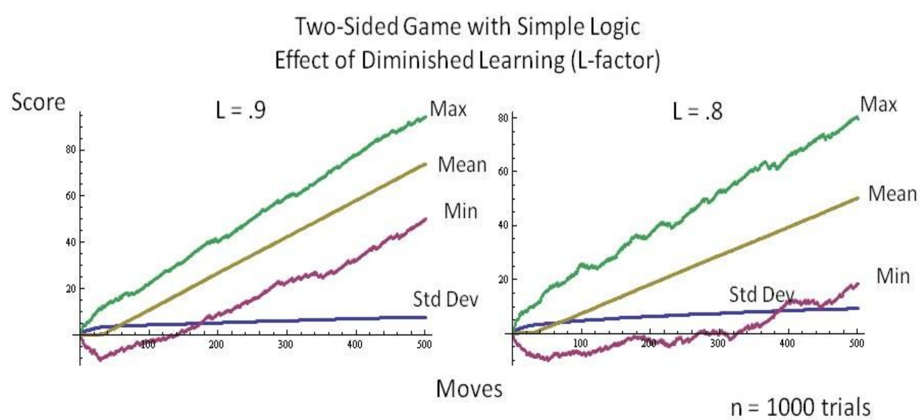

Figure 13: Effect of Incorrect Learning

With the limitations of this set of semantic concepts in mind, adding new characteristics to the mix might serve to further improve model performance. To that end, two new semantic concepts are added, based on the absolute position, either positive or negative. If the absolute position is below the expected value of the game times the number of moves so far, the semantic concept is negative (SC_{neg}) and vice versa if the absolute position is above the expected value thus far (SC_{pos}).

Based on these new semantic concepts, an "optimistic" set of decision rules was used. To the previous logic of keep the last move if the result is satisfactory and change if it was unsatisfactory, add the following:

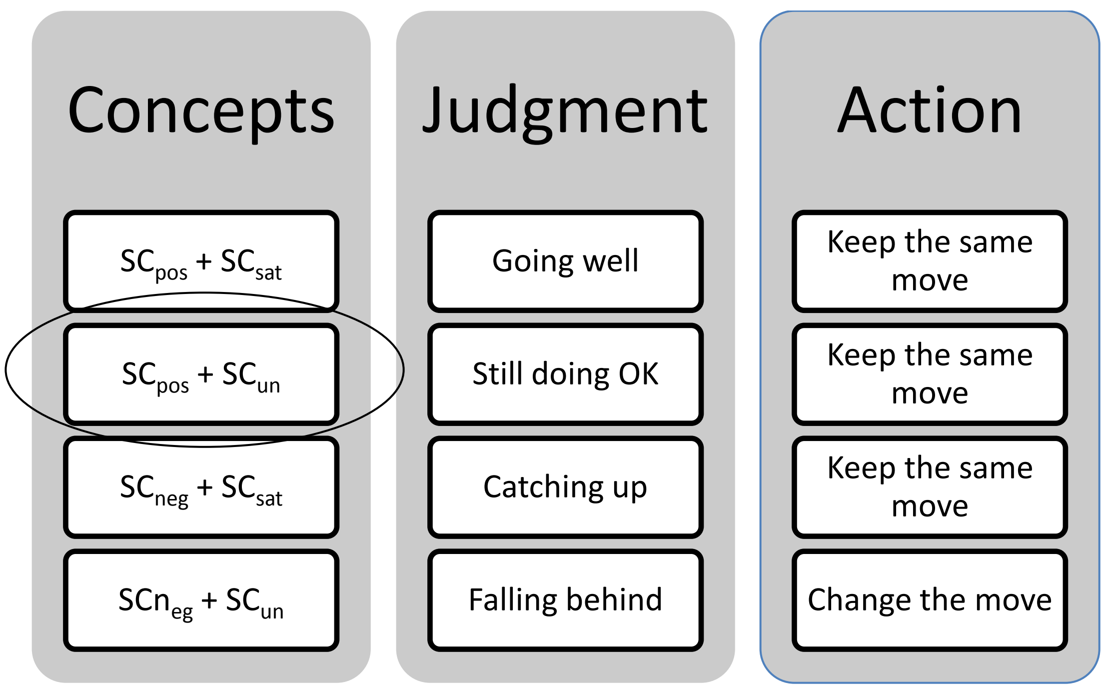

Figure 14: Semantic Concepts, Judgment and Action

As mentioned previously, the above logic adds an "optimistic" component to the decision process. Previously, any satisfactory result for a single move resulted in the choice to keep the move unchanged, but now the move will also remain unchanged as long as the overall picture is still positive. The result of this logic, however, is not to improve the chances of cooperation. Instead, it acts to effectively blunt the ability of a player to punish the other player's defection under certain initial conditions.

Generally, the game will play out in the same fashion as in the simple logic above, but in cases where the random setup of the first 32 moves causes one player to be more than two steps into positive territory while the other is more than two steps into negative territory (where a step is defined by the combination of either two positive or two negative move results), then the situation exists where the logic will cause the player in negative territory to oscillate back and forth between the move choices while the other player continues to play the same move. If the player in positive territory is playing "defect" that will cause the player in negative territory to continue to fall farther and farther behind. Normally, this version of the simulation keeps the two players quite close in terms of their performance, with margins at the end of 500 moves being less than one unit, as did the previous, simpler version of the game. However, under the right initial conditions, the game settles into a definite "win/lose" situation, which is set up as shown below.

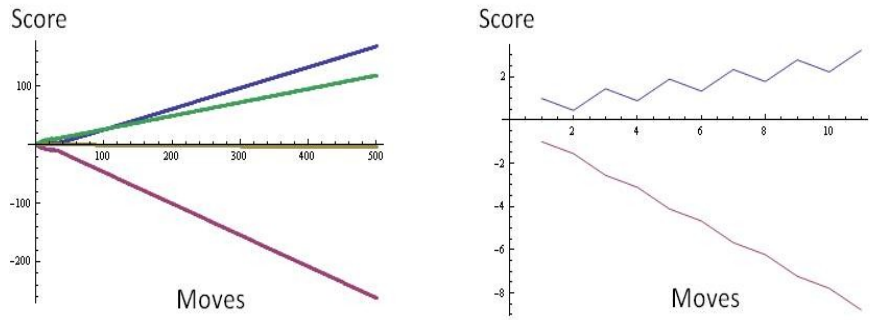

Winners and Losers in the 4-concept Game

Initial Setup for a Losing Game

n = 1000 trials

Figure 15: Results of the 4-Concept PD Simulation

The impact of random moves, payoff uncertainty, and a diminished learning factor follow the same pattern as discussed above for the case of the simple, 2-concept simulation, except that they increase the probability that the game will turn into negative territory.

In an effort to improve the ability of the losing player to recover, I added another semantic concept to cover the situation where one player was falling behind the other. This is based on the relative difference between the cumulative payoffs at each move. If the relative difference between the cumulative payoffs is increasing, the probability of the losing side acting randomly in an effort to “catch up” increases. As expected, this causes the game to quickly become unstable and turn into negative territory as the cycle of mutual random moves increases over time, as shown in Figure 16 below.

One additional factor that plays into all the simulations discussed thus far is the impact of the time horizon. Even though players do not know the time horizon, the dynamics of the game will eventually force the game into negative territory if there is any uncertainty and the time horizon is long enough. The more uncertainty, the shorter the time horizon before the game reaches negative territory. Even with adding a fair amount of uncertainty to the mix, however, the game does tend to stay in positive territory over the short run (less than 200 moves), as shown in the figure below.

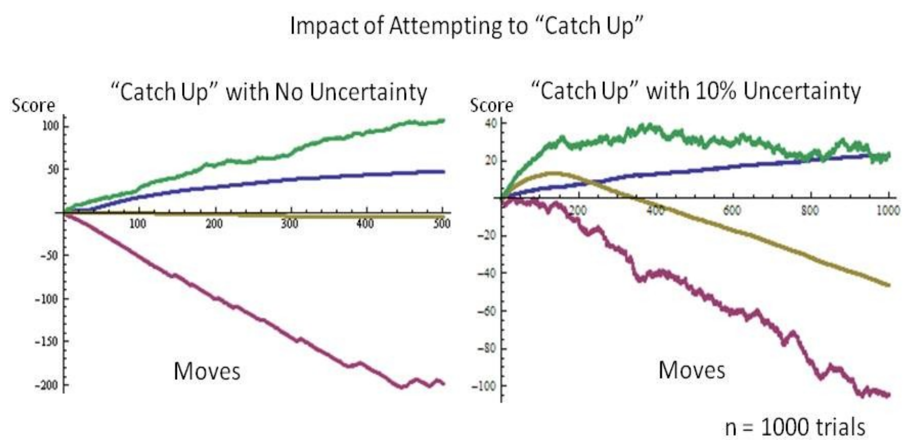

Figure 16: Impact of Trying to Catch Up

Summary of PD Games with Semantic Concepts Only

Clearly, little ground has been broken by simply adding semantic concepts to the mix, other than to show how a {C,c} equilibrium can be reached. Further, these sets of models also show how fragile such cooperation can be – breaking down quickly in the face of uncertain payoffs or faulty information. Also, it has been shown that the use of semantic concepts alone can lead to a semantic trap where the players repeat strategies in a vain attempt to catch up. As discussed earlier, some mechanism needs to be added to capture the impact of learning from the past and the ability to reason about the future. That mechanism, for people, is memory.

A Simple Memory Vector Model

To set up a demonstration of how this works, the following procedure will be used. This procedure will try to confirm the presence or absence of a given semantic concept. Following the idea of Thomas, et al, (2008), I generated a set of 1000 data vectors of eight elements, each generated from an idealized vector corresponding to a given semantic concept is generated. The vectors are generated by randomly adding or subtracting a small factor from each element, so the data is noisy. As an example, the idealized vector corresponding to the semantic concept SC_{ip} (increasing payoff) would be:

$$SC_{ip} = \{0, .02, .04, .06, .08, .10, .12, .14\} \qquad (5)$$

and an example generated vector would be:

{-0.0196, 0.0075, 0.040, 0.080, 0.091, 0.0976, 0.125, 0.122}

Plotting the two vectors together shows that the same general shape is retained, but the data is a noisy representation of the ideal vector.

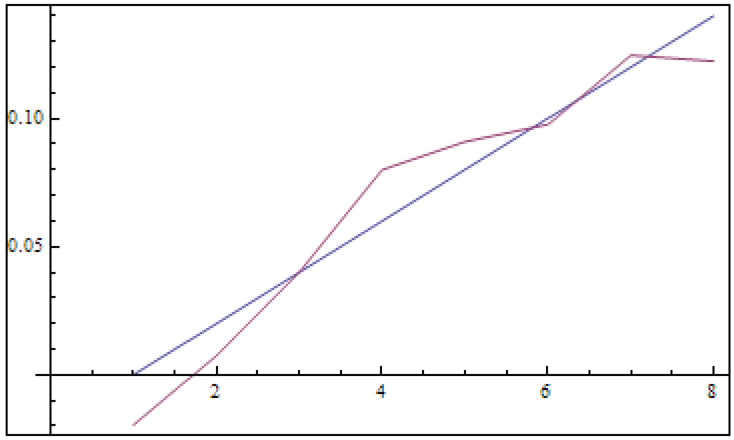

Figure 17: Ideal vs. Generated Vector

Additionally, the vectors were encoded using an "L" or learning factor ranging from .1 (only 10% of vectors encoded correctly vs. 90% encoded as random noise – only 10% learned) to 1 (100% of vectors encoded correctly). The probability of a particular hypothesis vector being encoded was varied between high base rate (7, 2, 2, 2, 2 distribution between the particular hypothesis vector and the alternatives, ≈ 48%) and a low base rate (2, 7, 7, 7, 7 distribution, ≈ 8%). This means that in the first case, vectors are generated at the rate of seven vectors corresponding to the target semantic concept, and two vectors representing each of four competing concepts. In the low base rate case, the proportions are reversed. These sets of vectors were then encoded as memory traces.

Information is then gathered using various probes of the 1000 generated memory traces. The probe is generated using the idealized semantic concept vector, and each concept vector is used in turn to probe memory. This probe essentially allows the model to calculate the similarity between each probe and each trace in memory based on the sum of the squared differences between each vector element. This similarity, S_i, between Probe P and a particular (i) trace T_i is given by:

$$S_i = \frac{\sum_{j=1}^{N} \left(P_j - T_{ij}\right)^2}{N_i} \tag{6}$$

where P_j is the j^{th} position of Probe P, T_{ij} is the j^{th} position of Trace T_i, and N is the number of vector elements in the i^{th} trace.

A perfect match would create a similarity value of 1, while a total mismatch would result in a similarity value of 0. Following earlier models such as MINERVA-DM and HyGene, the similarity value is then cubed to create an Activation value A_i, which effectively decreases the weighting of low similarity values. These activation values are then summed across all M traces to create an Echo Intensity, I.

$$I = \sum_{i=1}^{M} A_i \tag{7}$$

Thus, a situation where the probe matched all 1000 traces in memory would create an echo intensity of 1000. After each of the hypotheses is tested against the traces in memory, the hypothesis with the greatest echo intensity for each trace is chosen as the "best" hypothesis. With suitable tuning for the correct minimum activation value to judge similarity (see Thomas, et. al., 2008 for a much more detailed description of a complex model applied to clinical judgment), the model generates impressive rates of hypothesis recognition, particularly at high levels of learning, as shown if Figures 18 and 19, below.

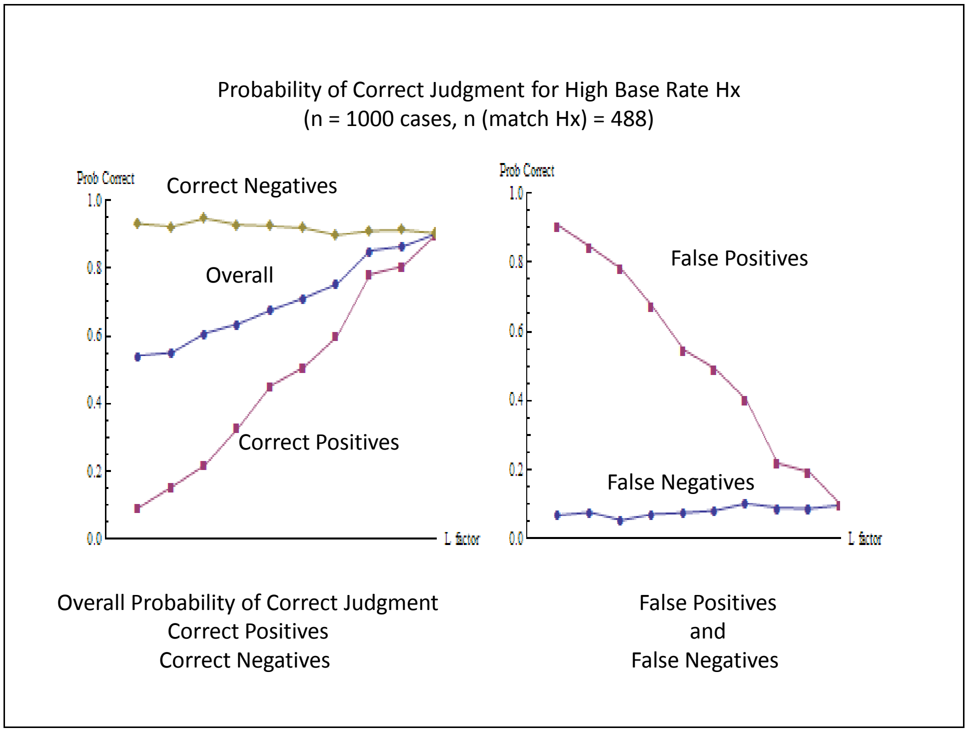

Figure 18: Model Performance for High Base Rate Case

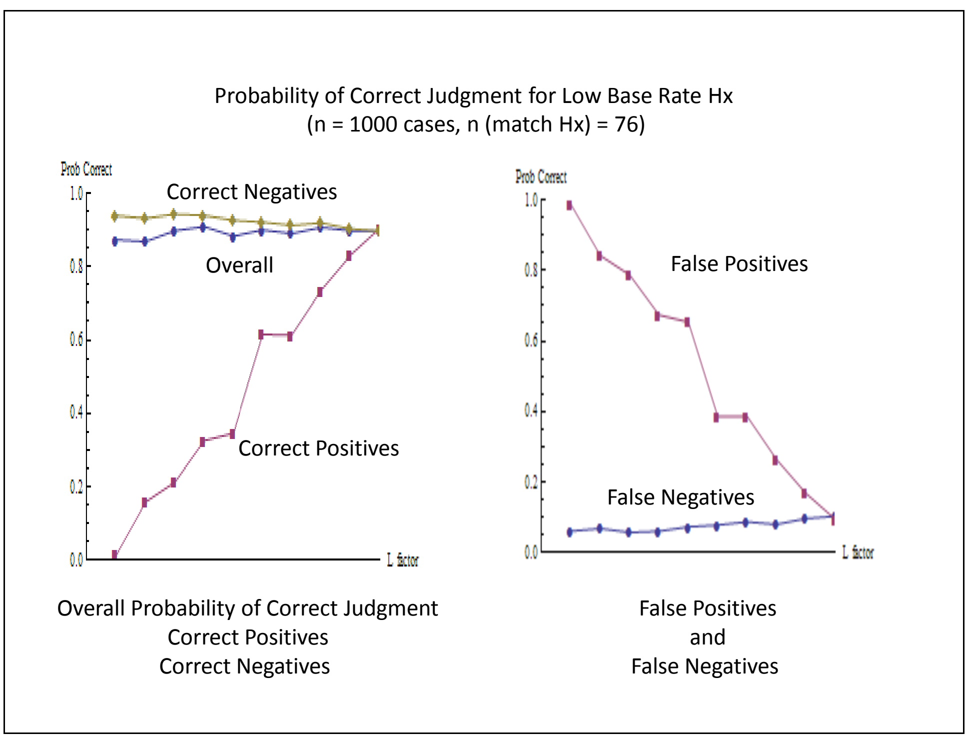

Figure 19: Model Performance for the Low Base Rate Case

In both cases, there are impressive gains in correctly identifying the traces which represent the desired hypothesis as the percentage of correctly learned memory traces increases. In the high base rate case, the overall performance of correct answers ranges from just over 50% in the low-learning state to near unity in the perfect learning state. This matches intuition, since in the high base rate case nearly 50% of the traces are generated from the chosen hypothesis, so simply guessing the hypothesis without even looking at the data would give a 50% chance of being right. As learning increases, the number of correctly-identified traces increases rapidly, although at slight increase in number of positive cases falsely rejected (false negatives).

In the low base rate case, the overall rate of successful prediction remains high throughout. Again, this matches intuition, since the incidence of the chosen hypothesis is less than 10%, so there would be a 90%+ chance of being right by simply rejecting the hypothesis. As learning increases, the false positive rate greatly decreases, again at the cost of a slight increase in the number of false negatives.

Adding Memory to the PD Game

Thus far, all of the simulations have focused on the choice simply given the current situation without regard to the potential impacts on future moves. For example, by playing PD (or other) games repeatedly, it might be more advantageous to consider the possible longer-term repercussions of a given move, rather than simply concentrating on the anticipated rewards of the next move. Thus, considering the longer-term implications may have a positive effect on the overall stability of the game, and serve to keep the game in positive territory for both players rather than degenerating into mutual defection as seen above.

To this end, a new simulation was developed that considers the longer-term consequences of potential moves. A probe of the "memory" of previous moves was added in an attempt to understand the longer-term consequences of a given move. The basic logic flow of this memory probe is as follows:

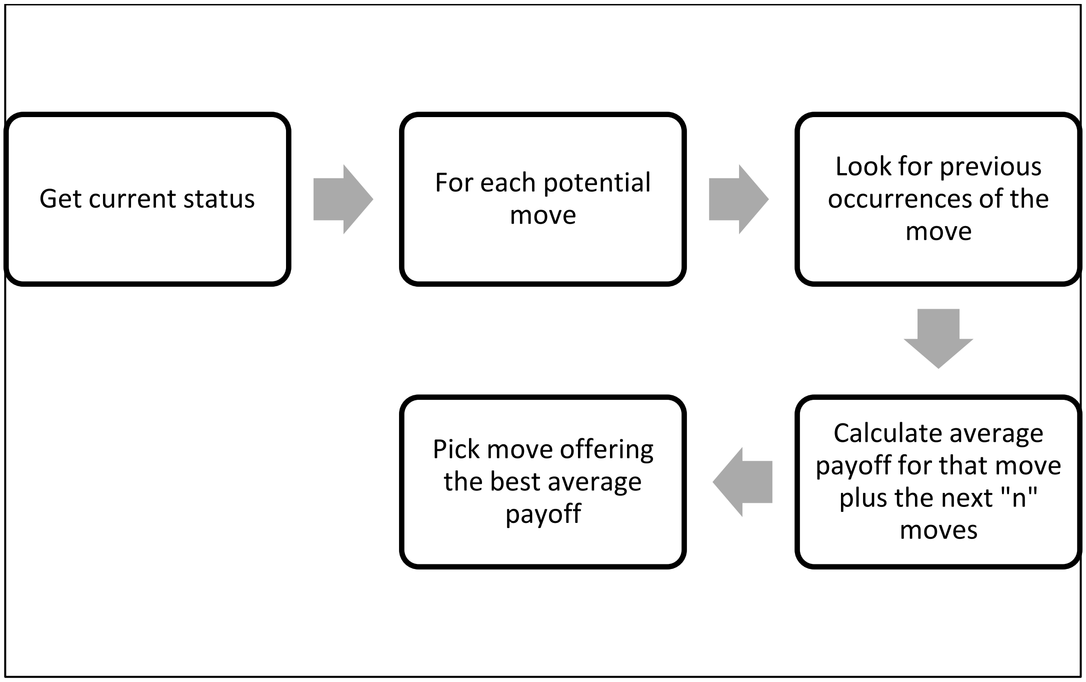

Figure 20: Logic Flow for Probing Memory

The mathematical mechanics for probing memory are as previously discussed. One of the key questions which this logic raises is: "how many moves to look forward?" Intuition tells us that the more information that the process has to work with, the better. However, in this case this turns out not to be so. While developing this simulation, the relationship between the number of moves assessed (that is, how many moves down the line need to be considered when choosing the best move) and the accuracy of game play (was the best move chosen) was evaluated. Interestingly, looking ahead more than 6-7 moves, actually decreases the performance of the simulation. The percentage of incorrect moves increases rapidly from near zero to about 30% and then levels off at around 20 moves of look-ahead.

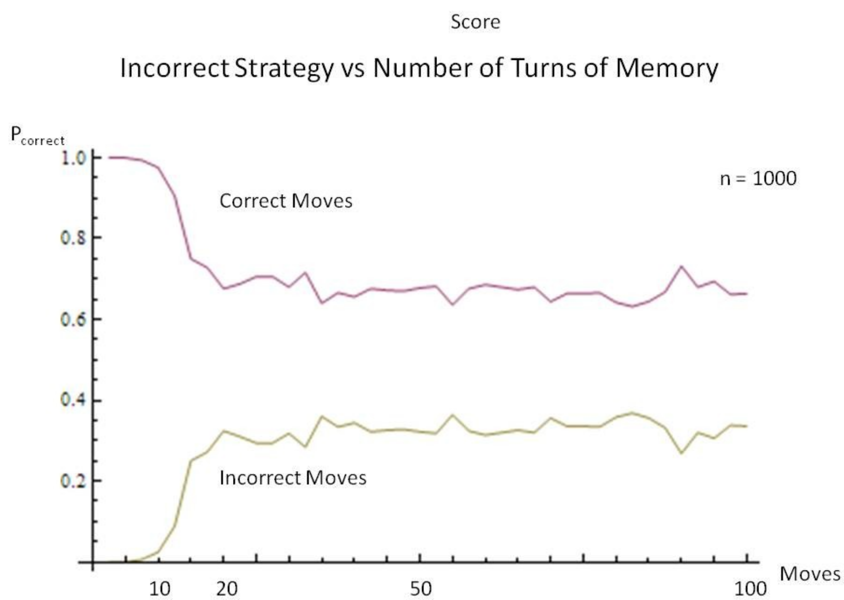

Figure 21: Impact of Looking Ahead Too Far

The reason this occurs is that the repeated PD game is not just a series of 1-time games strung together, but is also a series of cycles between moves (Alpern, 1993). Since there are only 4 move combinations, even a random chance series of moves will return to the same starting state fairly quickly (8 moves). With a purposeful combination of cooperate and defect, the cycle tends to only be 2-3 moves. Thus, by looking farther than 6-7 moves, the logic starts to pick up the results of more than one or two cycles, and the additional information gained beyond that point is mere repetition. Plus, longer-term results gets combined with the results of other moves, so instead of adding to the set of valid information, looking too far ahead actually decreases the percentage of good information as it is submerged in repetitious noise. This result also has a felicitous coincidence with the long-held idea of the size of working memory (see Simon and Newell, 1972). At any rate, assessing potential results beyond the period which we can reasonably expect to predict is probably counterproductive. This point is why many computer firms have a very short planning horizon since, given Moore's law, computer capacity can be assumed to increase 10x within 5 years, so it is foolish and limiting to look more than 3 years in the future (personal conversation with Nathan Myrvold, former CTO, Microsoft Corp.)

Application: A Prisoners' Dilemma Tournament

Following Axelrod (1984, 1997), I set up a PD tournament to test the performance of the memory-based model against a variety of alternative strategies. The tournament was run as a round-robin competition with each strategy facing all of the opposing strategies in turn. Each individual competition consisted of 1000 simulation runs of 1000 moves each. The value of each move combination was in accordance with the canonical PD game listed in Figure 2, above, with payoff values being mathematically transformed by simply subtracting 2 from each result. This causes payoffs for mutual defection and the sucker's payoff to have negative utility (-1 and -2, respectively). This also causes the division between "satisfactory" and "non-satisfactory" semantic concepts (SC_{sat} and SC_{un}) to be made clear. As discussed above, the payoffs were made slightly "noisy" with the addition of a small (2%) random component, so the payoff (for example) of the strategy {D,d} would be in the range of .98 – 1.02.

The following strategies were used for the tournament:

- Memory: A strategy using the memory-based model discussed above.

- Tit-For-Tat: The winning strategy from Axelrod 1984. This strategy will begin by cooperating, but will then mirror the opponent's previous move.

- Always Defect: This strategy plays a pure defection strategy. Defect is the optimal single-game strategy.

- Always Cooperate: This strategy plays a pure cooperation strategy.

- Tat-For-Tit: An inverse strategy to Tit-For-Tat. It starts out by defecting, and will continue to defect until the opponent defects as well, and then it will cooperate until the opponent cooperates, whereupon it will defect again. (Binmore, 1994)

- Grim Trigger: This strategy will cooperate until the first time the opponent defects. According to Morrow (1994), this strategy can be optimal for repeated PD games.

- 70-30: This strategy defects 30% of the time and cooperates 70% of the time.

- RLS: This strategy attempts to predict the opponent's next move using a Recursive Least Squares projection based on the previous sets of moves, and then takes the optimum move given that projection.

- Average Move: This strategy uses a heuristic based on the average opponent move from the previous set of moves, and selects an optimum move based on that heuristic.

- Last Move: This strategy uses a heuristic where the next opponent move is assumed to be the same as the last move, and the selects a move accordingly.

- Random: This strategy picks moves at random.

Each set of contests was scored using three different criteria:

- Won/Loss: Number of wins (better score) vs. number of losses.

- Total Score: Total of scores for all iterations of the contest

- Victory Margin: Total margin of victory (or magnitude of defeat) for all iterations of the contest.

The following three tables give the results of the tournament according to each criterion.

	Memory	Random	AlwaysD	70-30	RLS	AvgPayoff	LastPayoff	TFT	AlwaysC	Tat4Tit	Grim	Win %	Rank
Memory		1000	529	1000	1000	1000	1000	633	1000	1000	581	87.43	**1**
Random	0		0	0	530	0	492	554	1000	0	0	25.76	**10**
AlwaysD	471	1000		1000	1000	1000	1000	589	1000	1000	600	86.60	**2**
70-30	0	1000	0		1000	995	1000	552	1000	0	0	55.47	**5**
RLS	0	470	0	0		116	505	584	745	0	234	26.54	**9**
LastPayoff	0	508	0	0	495	118		552	764	0	268	27.05	**8**
TFT	367	446	411	448	416	447	448		493	399	507	43.82	**7**
AlwaysC	0	0	0	0	255	144	236	507		0	491	16.33	**11**
Tat4Tit	0	1000	0	1000	1000	741	1000	601	1000		0	63.42	**4**
Grim	419	1000	400	1000	766	858	732	493	509	1000		71.77	**3**
Lose %	12.57	74.24	13.40	44.53	73.46	54.19	72.95	56.18	83.67	36.58	28.23		
Rank	**1**	**10**	**2**	**5**	**9**	**6**	**8**	**7**	**11**	**4**	**3**		

Table 1: PD Tournament Wins vs. Losses

Note 1: Read this chart from left to right for wins, and top to bottom for losses. For example (reading from the left column of labels), Memory won vs. AlwaysD 529 times, and won against TFT 633 times. For the losses (reading downward from the top row of labels), Memory lost against AlwaysD 471 times, and lost against TFT 367 times.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

	Memory	Random	AlwaysD	70-30	RLS	AvgPayoff	LastPayoff	TFT	AlwaysC	Tat4Tit	Grim	Total Score	Rank
Memory		228717	-382098	25586	229294	229355	229308	-380483	841369	838914	-380699	1479263	1
Random	-535027		-535309	-179701	-353	-249320	-1138	92	535367	-531710	-532163	-2029261	10
AlwaysD	-382470	230221		24746	229295	228974	229117	-380347	841110	838794	-380399	1479040	2
70-30	-484029	77492	-483810		76390	87	75960	-109529	637258	-478284	-481491	-1169956	7
RLS	-535101	-227	-534994	-178334		-337793	228978	76418	538783	-534045	-150307	-1426623	9
AvgPayoff	-535240	107774	-535049	-151257	228444		228618	-5712	638955	-178524	-353440	-555432	4
LastPayoff	-535194	779	-534925	-178212	228871	-332692		71240	546065	-534096	-154774	-1422937	8
TFT	-382652	-963	-382760	-111063	75044	-6760	70194		229427	-380461	229544	-660451	5
AlwaysC	-688070	-229747	-688207	-382737	-234741	-385079	-245729	229382		-687859	229441	-3083344	11
Tat4Tit	-687362	226799	-687292	23696	229326	-8407	229389	-378273	840807		-685470	-896787	6
Grim	-382415	228444	-382649	23667	228361	227247	228232	229200	229325	837148		1466559	3
Total Opponent Score	-5147559	869289	-5147093	-1083610	1289931	-634389	1272929	-648012	5878464	-810122	-2659757		
Rank	1	8	2	4	10	7	9	6	11	5	3		

Table 2: PD Tournament Total Points

Note 1: Read this chart from left to right for points total for, and top to bottom for total points against. For example (reading from the left column of labels), Memory scored -382,089 points vs. AlwaysD, and scored -380,483 points against TFT. For the total points against (reading downward from the top row of labels), AlwaysD scored -382,470 points against Memory, and TFT scored -382,652 against memory.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

	Memory	Random	Always D	70-30	RLS	AvgPayoff	LastPayoff	TFT	AlwaysC	Tat4Tit	Grim	Avg Margin	Rank
Memory		763.74	0.37	509.22	764.39	764.60	764.50	2.17	1528.94	1526.28	1.72	662.59	2
Random	-763.74		-765.53	-257.19	-0.13	-357.09	-1.92	1.06	765.11	-758.51	-760.61	-289.85	10
AlwaysD	-0.37	765.53		508.56	764.29	764.02	764.04	2.41	1529.32	1526.09	2.25	662.61	1
70-30	-509.22	257.19	-508.56		254.72	151.34	254.17	1.53	1020.00	-501.98	-505.16	-8.59	6
RLS	-764.39	0.13	-764.29	-254.72		-566.24	0.11	1.37	773.52	-763.37	-378.67	-271.66	9
AvgPayoff	-764.60	357.09	-764.02	-151.34	566.24		561.31	1.05	1024.03	-170.12	-580.69	7.90	4
LasPayoff	-764.50	1.92	-764.04	-254.17	-0.11	-561.31		1.05	791.79	-763.49	-383.01	-269.59	8
TFT	-2.17	-1.06	-2.41	-1.53	-1.37	-1.05	-1.05		0.05	-2.19	0.34	-1.24	5
AlwaysC	-1528.94	-765.11	-1529.32	-1020.00	-773.52	-1024.03	-791.79	-0.05		-1528.67	0.12	-896.13	11
Tat4Tit	-1526.28	758.51	-1526.09	501.98	763.37	170.12	763.49	2.19	1528.67		-1522.62	-8.67	7
Grim	-1.72	760.61	-2.25	505.16	378.67	580.69	383.01	-0.34	-0.12	1522.62		412.63	3
Avg Margin	-662.59	289.85	-662.61	8.59	271.65	-7.90	269.59	1.24	896.13	8.67	-412.63		
Rank	2	10	1	6	9	4	8	5	11	7	3		

Table 3: PD Tournament Margin of Victory

Note 1: Read this chart from left to right for average margin of victory, and top to bottom for average margin of defeat. For example (reading from the left column of labels), Memory had of 0.37 average margin vs. AlwaysD, and scored 2.17 average margin against TFT. For the margin of defeat (reading downward from the top row of labels), AlwaysD had an average margin of -0.37 (negative indicates defeat) against Memory, and TFT scored -2.17 average margin against memory.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

The following chart summarizes the ranking in the PD tournament.

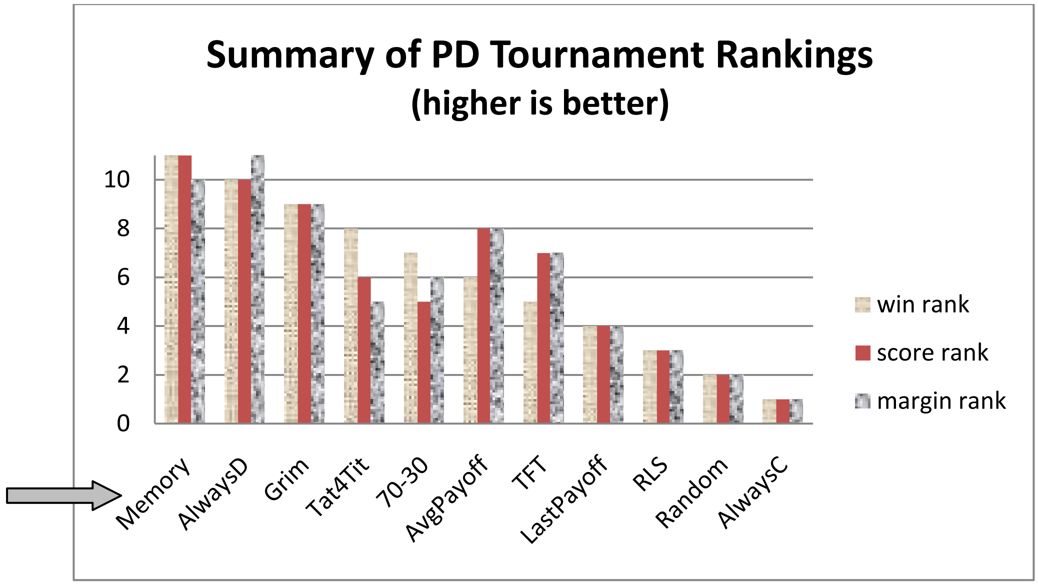

Figure 22: Summary of PD Tournament Rankings

The overall results show that the Memory-based approach had the top rank in both wins and overall score, and was second to the "Always Defect" strategy in overall margin. That said, statistical tests show that Memory and Always Defect ended up in a statistical dead heat, with the difference in the means over 1000 runs was only 0.06, for a one-sided T-statistic of 0.31, and a p-value of 0.38. Comparison to the third place finisher, "Grim Trigger" shows a definite win, with the difference in means over 1000 runs of 1.88, for a one-sided T-statistic of 8.7, and a significance <0.0001. This shows that the memory-based model performed at least as good as, if not better than, the game-theoretic optimal approaches for both single-shot and iterated PD games.

The relatively poor performance of Tit-for-Tat (TFT) in this tournament bears some discussion, as it cuts across the conventional wisdom that TFT is the best strategy. Binmore (1994) points out that TFT is only a viable overall strategy as long as there is a sufficient number of “nice” strategies in the overall mix of opponents, with “nice” being strategies that are generally cooperative. In this particular tournament, five strategies (Memory, AlwaysD, Grim, 70-30, and Tat-for-Tit) tended to be “mean”, which in this context means that they tended to defect more often, the three heuristic strategies tended to play whatever their opponent played, and the random strategy was neutral, so only one “nice” strategy (Always Cooperate) was in the tournament. This explains the middle-of-the-road performance of TFT. Additionally, Axelrod (1984) himself noted that TFT lost every single encounter, and that he only judged the “goodness” of the outcome by the total score.

Alternating Games

Having previously set up both the Stag Hunt (SH) and PD games, imagine a situation where, rather than playing a single game repeatedly, a game shifts back and forth between the two sets of payoffs. This situation could well be more likely in real life than that of the single game played repeatedly. Imagine, for example, the situation where the set of incentives switched back and forth between opportunities for cooperation and temptations for aggression, where the difference is determined exogenously to the players by something like advances in technology (for example, the introduction of things like the rifled musket, barbed wire, the machine gun, or nuclear weapons). On the surface, there is quite a bit of similarity between the games – both have the same Pareto optimal and repeated game equilibrium of (C, c), and both share the same minimum risk strategy of (D, d). However, plotting these two games in decision space reveals they create quite different constraints on the players.

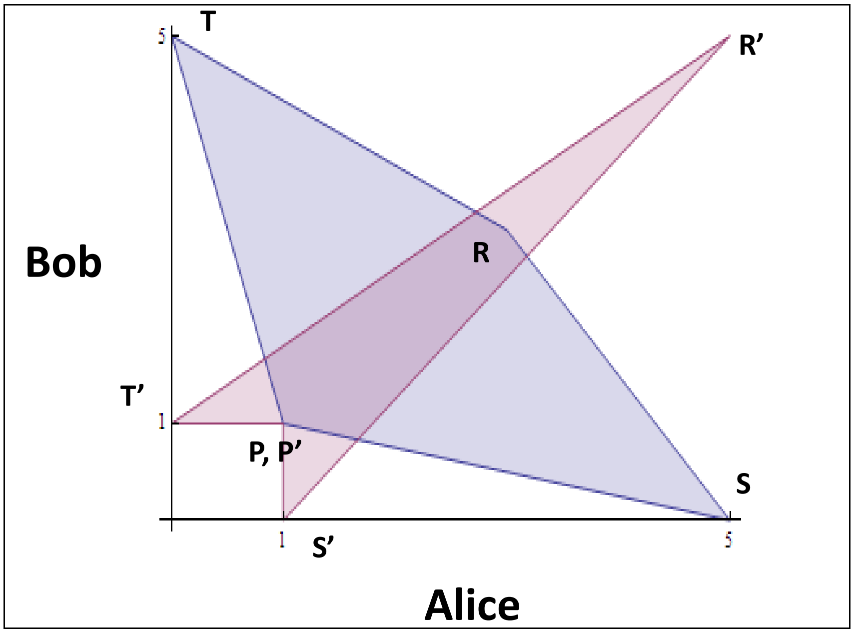

Figure 23: PD and SH Games in Decision Space

For ease of reading, the points in the decision space are denoted according to Bob’s payoff, so the decision space for Bob’s PD game is the figure P-T-R-S and for the SH game the figure is P'-T'-R'-S'. Bob’s payoff in the PD game if he plays a pure “cooperate” strategy is along the line R-S and along the line R'-S' for the SH game. His payoff for an “always defect” strategy is along the line T-P or T'-P'. Similarly, Alice’s payoffs for the two strategies are T-R or T'-R' for cooperate and S-P or S'-P' for defect.

Looking at this figure, it can immediately be seen that the payoff for a given decision is critically dependant on which game is being played. As the game shifts from PD to SH, the decision space rapidly collapses to P'-T'-R'-S', and then expands rapidly once the game goes back to PD. The players will, therefore, have greatly different judgments on the best strategy depending on their judgment of the probability that one game is being played rather than another. If the games oscillate back and forth in a sinusoidal fashion, for instance, the payoff structure might look something like this:

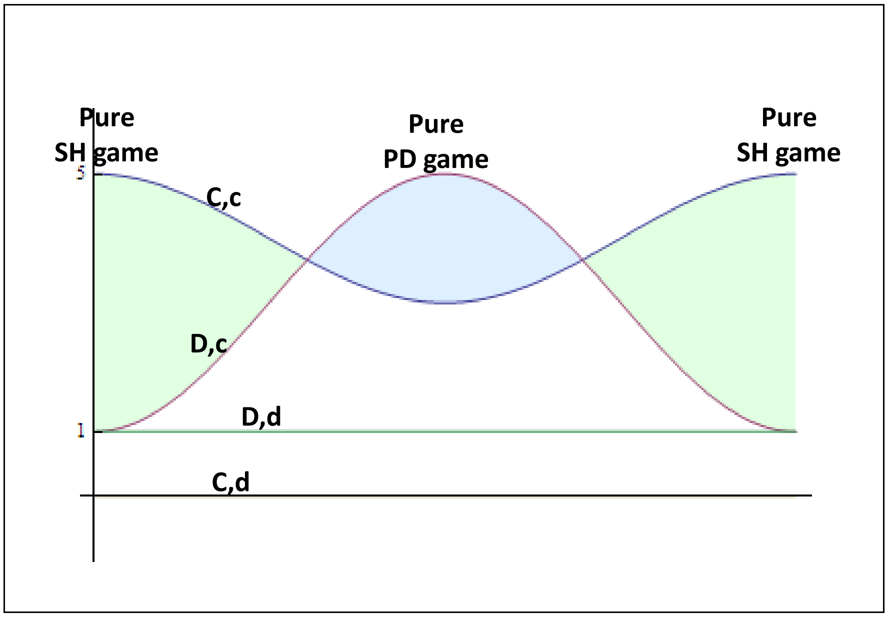

Figure 24: Payoff Structure of an Alternating SH and PD Game

This combined game can be thought of as a single 3-move game where the first move is the chance that the game is a SH game vs. that of a PD game (see, for another example, the "fashionable prisoners dilemma" in Binmore, 1994), and the payoff can be written in terms of the probability of an SH game. To derive this, first look at the game in extensive form.

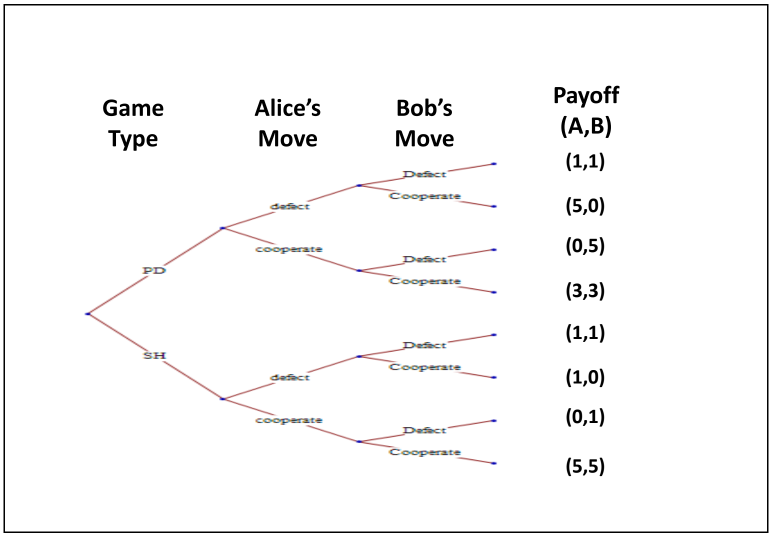

Figure 25: Combined PD and SH Game in Extensive Form

It can be seen that there is some opportunity to simplify the problem by engaging in some judicious pruning of the tree, since the move combination (d, D) gives the same payoff no matter which game is chosen. Further, the payoff for the player on the wrong side of the (c, D) or (D, c) move is always zero, so the game type can be deleted and payoffs assessed in terms of probability p, where

$$p = \text{Probability}_{\text{Stag Hunt}} \quad (8)$$

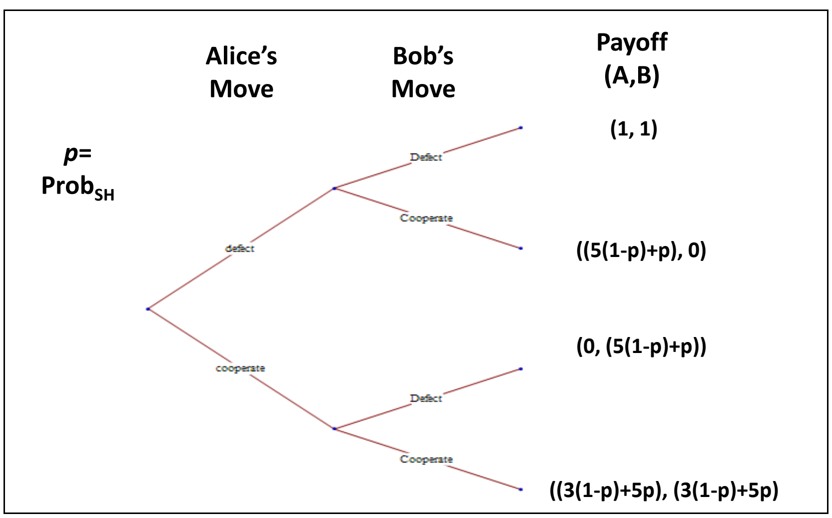

Figure 26: Payoffs for a Combined Game

Simplifying:

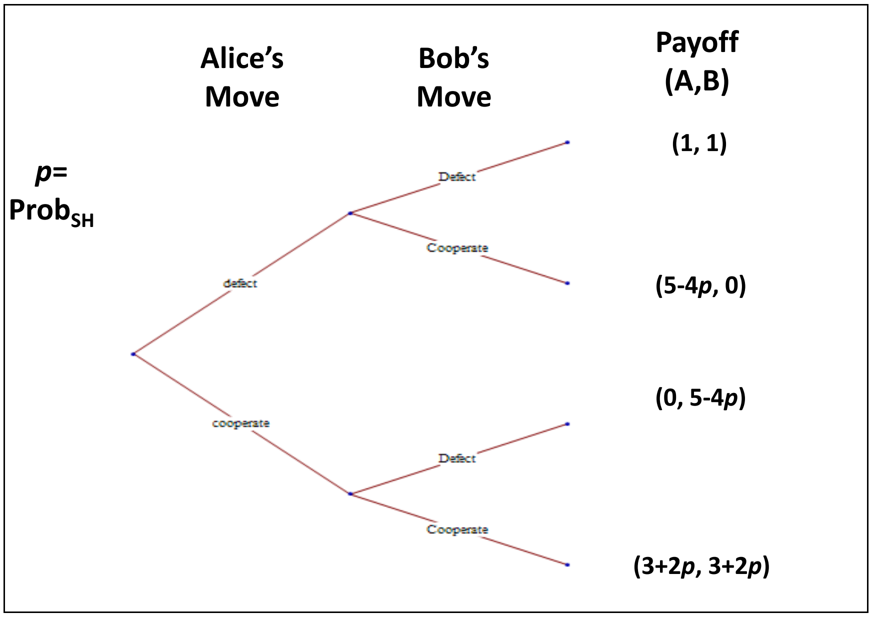

Figure 27: Simplified Payoffs for the Combined Game

Going back to Figure 26, it is intuitive that the payoff for (d, D) will always be preferred to the payoff for (c, D) as 1 > 0. At some point, the preference order between (C, c) and (D, c) will switch. From Figure 24, this will be at the point where the payoff from (C, c) and (D, c) are equal, which can be written as:

$$5 - 4p = 3 + 2p \;\rightarrow\; 2 = 6p \;\rightarrow\; p = 1/3 \qquad (9)$$

Thus, based on the judgment of p, a player will play the game as:

Stag Hunt where $p > 1/3$

Prisoners Dilemma where $p < 1/3$

Indifferent at $p = 1/3$

This leads to the conclusion that one of the important judgments to be made in the combined SH-PD game is the probability judgment that the game is in fact an SH game. This is in accordance to my earlier observation that the judgment of the type of game is critical to the analysis. The difference here is that the judgment is being made by the players within the game rather than being exogenously imposed. This process of reaching a probability judgment can be made possible through the use of 4 semantic concepts, which are shown on the following figure:

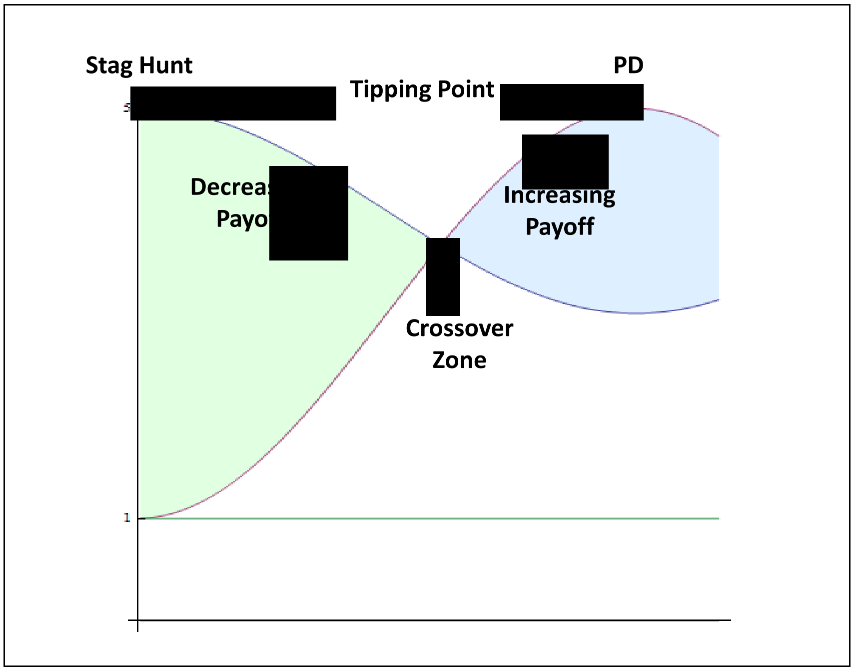

Figure 28: Semantic Concepts to Describe a Combined Game

For example, in a situation where both players were cooperating (c, C) and the payoff was decreasing, the probability that the game was in fact a Stag Hunt will also be decreasing, and players will tend to increase the probability of defection in hopes of being able to spot the inflection point and gain an advantage. Similarly, if one player defects while the other cooperates (d, C) and the payoff is increasing, that shows that the probability of a PD situation is increasing and that of a SH is decreasing. Once the *crossover zone* is past, the players will play the game as a PD. Next, the point where the payoff for (d, C) stops increasing will indicate a tipping point where the probability that the game is a PD starts to decrease and the probability of a SH starts to increase.

Thus, there are 4 semantic concepts upon which to base the judgment of whether the game is a Stag Hunt or a Prisoners Dilemma:

SC_{dp} = Decreasing payoff

SC_{ip} = Increasing Payoff

SC_{xover} = Crossover Zone

SC_{tp} = Tipping Point

Put into table form, this yields:

Moves	Semantic Concept	Judgment	Action
(c, C)	SC_{dp}	P_{SH} is decreasing P_{PD} is increasing	Consider defecting
(c, C)	SC_{ip}	PSH is increasing PPD is decreasing	Continue to cooperate
(c, C)	SC_{xover}	PSH is decreasing PPD is increasing	Change to defect
(c, C)	SC_{tp}	PSH is increasing PPD is decreasing	Continue to cooperate
(d, C)	SC_{dp}	PSH is increasing PPD is decreasing	Consider cooperating
(d, C)	SC_{ip}	PSH is decreasing PPD is increasing	Continue defecting
(d, C)	SC_{xover}	PSH is increasing PPD is decreasing	Change to cooperate
(d, C)	SC_{tp}	PSH is increasing PPD is decreasing	Continue defecting

Table 4: Semantic Concepts in a Combined Game

Application: An Alternating Game Tournament

Having described the environment of an alternating game situation, the performance of the memory-based model can be tested using the same tournament procedures as before. This time, the payoff structure slowly alternates between a PD and SH payoff, as shown in Figure 27 above, and is slightly "noisy" (payoffs vary +/-

2% from the computed value), as were the payoffs in the pure PD tournament. Again, the tournament was a round-robin of 1000 model runs of 1000 turns each for each of the 10 strategy pairs.

In this game, the heuristic approaches (RLS, Average Payoff, and Last Payoff) attempt to discern which semantic concept is operative and then will move accordingly. As before, the memory-based approach has no fixed strategy, and will adjust its strategy based on the memory traces of the opponent's previous moves and the resulting anticipated payoff.

As before, the results were recorded by win/loss, total score, and margin of victory. The results are summarized in the three tables below.

Reversal Tournament: Wins vs. Losses													
	Memory	Random	AlwaysD	70-30	RLS	AvgPayoff	LastPayoff	TFT	AlwaysC	Tat4Tit	Grim	Win %	Rank
Memory		1000	1000	992	1000	996	1000	626	1000	1000	1000	96.14	1
Random	0		0	82	0	0	0	544	1000	0	0	16.26	10
AlwaysD	0	1000		1000	0	1000	0	597	1000	1000	585	61.82	4
70-30	8	918	0		0	0	0	553	1000	3	0	24.82	9
RLS	0	1000	1000	1000		126	500	527	759	1000	734	66.46	3
AvgPayoff	4	1000	0	1000	874		870	566	853	252	136	55.55	6
LastPayoff	0	1000	1000	1000	500	130		514	746	1000	757	66.47	2
TFT	374	456	403	447	473	434	486		514	384	517	44.88	7
AlwaysC	0	0	0	0	241	147	254	486		0	484	16.12	11
Tat4Tit	0	1000	0	997	0	748	0	616	1000		0	43.61	8
Grim	0	1000	415	1000	266	864	243	483	516	1000		57.87	5
Lose %	3.86	83.74	38.18	75.18	33.54	44.45	33.53	55.12	83.88	56.39	42.13		
Rank	1	10	4	9	3	6	2	7	11	8	5		

Table 5: Reversal Tournament Wins vs. Losses

Note 1: Read this chart from left to right for wins, and top to bottom for losses. For example (reading from the left column of labels), Memory won vs. 70-30 992 times, and won against TFT 626 times. For the losses (reading downward from the top row of labels), Memory lost against 70-30 8 times, and lost against TFT 374 times.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

Reversal Tournament: Total Points Scored													
	Memory	Random	AlwaysD	70-30	RLS	AvgPayoff	LastPayoff	TFT	AlwaysC	Tat4Tit	Grim	Total Score	Rank
Memory		259067	-174520	128376	331129	-17657	347343	-116703	571610	-80370	-172784	1075492	6
Random	4345		-273134	-52522	162729	148702	167020	55598	385715	-270175	-268849	59430	8
AlwaysD	-420047	-171211		-294188	-49251	-146712	-50286	-540564	200372	198659	-540638	-1813867	11
70-30	-118198	-19519	-361984		92334	99327	95285	-133066	324862	-356414	-360619	-737993	9
RLS	124977	300332	-15916	193219		422688	571568	454978	594320	-14650	280824	2912340	2
AvgPayoff	-143505	332640	-181033	209149	545137		544730	323417	595418	-16917	-4896	2204140	4
LastPayoff	134856	302945	-16616	196922	571267	422485		459000	592870	-15480	289673	2937920	1
TFT	-119589	54261	-542741	-134285	454135	321988	458703		571861	-540745	572038	1095627	5
AlwaysC	155019	284285	-3302	188118	347708	281546	351662	571640		-2880	571574	2745370	3
Tat4Tit	-405679	-172549	-2965	-296954	-49171	-178048	-50504	-538045	200076		-377	-1494215	10
Grim	-420569	-171997	-542730	-295610	262441	17112	271806	571502	572021	196225		460201	7
Total Opp Score	-1208390	998254	-2114942	-157775	2668458	1371431	2707326	1107758	4609124	-902747	365946		
Rank	2	6	1	4	9	8	10	7	11	3	5		

Table 6: Reversal Tournament Total Points

Note 1: Read this chart from left to right for points total for, and top to bottom for total points against. For example (reading from the left column of labels), Memory scored -174,520 points vs. AlwaysD, and scored -116,703 points against TFT. For the total points against (reading downward from the top row of labels), AlwaysD scored –174,520 points against Memory, and TFT scored -116,703 against Memory.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

Reversal Tournament: Margin of Victory													
	Memory	**Random**	**AlwaysD**	**70-30**	**RLS**	**AvgPayoff**	**LastPayoff**	**TFT**	**AlwaysC**	**Tat4Tit**	**Grim**	**Avg Margin**	**Rank**
Memory		254.72	245.53	246.57	206.15	125.85	212.49	2.89	416.59	325.31	247.79	228.39	1
Random	-254.72		-101.92	-33.00	-137.60	-183.94	-135.93	1.34	101.43	-97.63	-96.85	-93.88	10
AlwaysD	-245.53	101.92		67.80	-33.33	34.32	-33.67	2.18	203.67	201.62	2.09	30.11	3
70-30	-246.57	33.00	-67.80		-100.89	-109.82	-101.64	1.22	136.74	-59.46	-65.01	-58.02	8
RLS	-206.15	137.60	33.33	100.89		-122.45	0.30	0.84	246.61	34.52	18.38	24.39	4
AvgPayoff	-125.85	183.94	-34.32	109.82	122.45		122.25	1.43	313.87	161.13	-22.01	83.27	2
LastPayoff	-212.49	135.93	33.67	101.64	-0.30	-122.25		0.30	241.21	35.02	17.87	23.06	5
TFT	-2.89	-1.34	-2.18	-1.22	-0.84	-1.43	-0.30		0.22	-2.70	0.54	-1.21	7
AlwaysC	-416.59	-101.43	-203.67	-136.74	-246.61	-313.87	-241.21	-0.22		-202.96	-0.45	-186.38	11
Tat4Tit	-325.31	97.63	-201.62	59.46	-34.52	-161.13	-35.02	2.70	202.96		-196.60	-59.15	9
Grim	-247.79	96.85	-2.09	65.01	-18.38	22.01	-17.87	-0.54	0.45	196.60		9.43	6
Avg Margin	-228.39	93.88	-30.11	58.02	-24.39	-83.27	-23.06	1.21	186.38	59.15	-9.43		
Rank	1	10	3	8	4	2	5	7	11	9	6		

Table 7: Reversal Tournament Victory Margin

Note 1: Read this chart from left to right for average margin of victory, and top to bottom for average margin of defeat. For example (reading from the left column of labels), Memory had of 245.53 average margin vs. AlwaysD, and scored 2.89 average margin against TFT. For the margin of defeat (reading downward from the top row of labels), AlwaysD had an average margin of -245.53 (negative indicates defeat) against Memory, and TFT scored -2.89 average margin against Memory.

Note 2: This chart based on a round-robin tournament, 1000 iterations of 1000 moves for each pair of strategies.

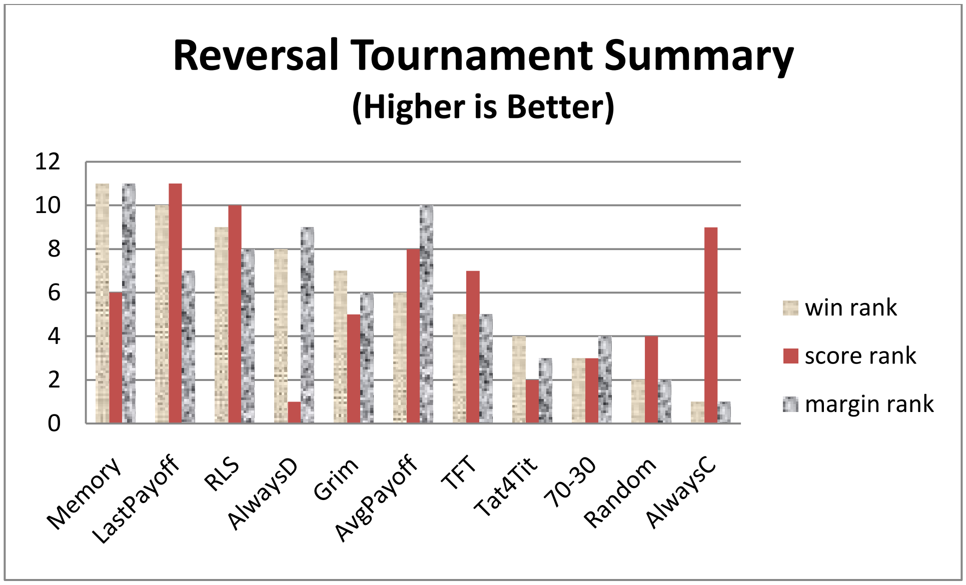

Figure 29: Summary of Reversal Tournament

The above figure shows the overall results of the Reversal Tournament. As can be seen, the memory model had the top scores as far as win percentage and win margin, but was in the middle of the pack in the total score ranking. Further, the performance of the heuristics which attempted to predict the next move by various heuristics was much better than in the static PD game, as would be expected. The old standbys "Always Defect" and "Grim Trigger" also ranked in the top half. This would be expected since strategies of defection are equilibrium strategies in both the PD and SH games. The 70-30 strategy, which was crafted to take partial advantage of the earlier analysis that one should play the game as a PD game 2/3 of the time in an alternating game, did more poorly than expected. Analysis showed that while this heuristic did play as expected, the difference between the 70-30 and 50-50 chances of Cooperate vs. Defect in the 70-30 strategy and Random strategy, respectively, was not enough to make more than a minor difference in the overall outcome.

The middling performance of the Memory model in terms of total score bears some more examination. Looking at Tables 5 and 7, above, the performance of this model in the Won/Loss and Margin measures of merit was well superior to the other models. The Memory-based model scored a convincing 96% win percentage, easily outscoring its nearest rival "Last Payoff", which only had a 66% win percentage. Similarly, the Memory-based model also scored convincing wins across the board in terms of victory margin, with an average margin of over 228 points, vs. its closest rival "Average Payoff", which had an average margin of about 83 points. Further, the Memory-based model had positive win margins against every other strategy, and the best any other strategy could do was to have positive margins in 8 out of 10 contests. However, the Memory-based model ran in the middle of the pack in terms of overall score. Analysis of the moves taken by the model revealed that the model was slower to react to changes in the payoff function than the heuristic models. This caused the memory-based model to play more conservatively (i.e., Defect more often) than the heuristics, which tended to switch quickly to a "Cooperate" strategy in an effort to catch up to the opponent. This often led to large score gains against similarly cooperative strategies, but was punished in head-to-head competition against the Memory-based model and by the strategies which played "Defect" most often.

Combining the Tournament Results

When combining the results of the two tournaments, it is clear that the Memory-based approach is viable across the board – both in the static PD game and in an alternative PD/SH environment. Further, looking at the combined picture shows that the Memory-based model dominates across all categories, as shown in the figure below.

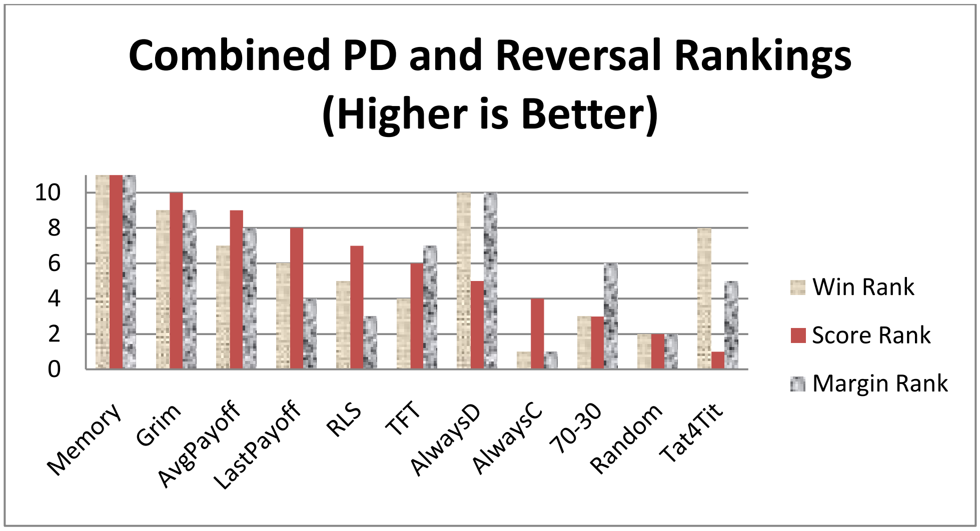

Figure 30: Combined PD and Reversal Rankings

As can be seen, the Memory-based approach dominates the rankings, ahead of both the traditional game-theoretic strategies of "Grim Trigger" and "Always Defect", and ahead of all the heuristic prediction strategies. Arguably, "Always Defect" could be moved up in the rankings based on which way the chart values are sorted. The figure below shows the same chart, except that this time the overall performance in each area has been normalized (1 = best, 0 = worst), which gives a bit more of a nuanced view of the results.

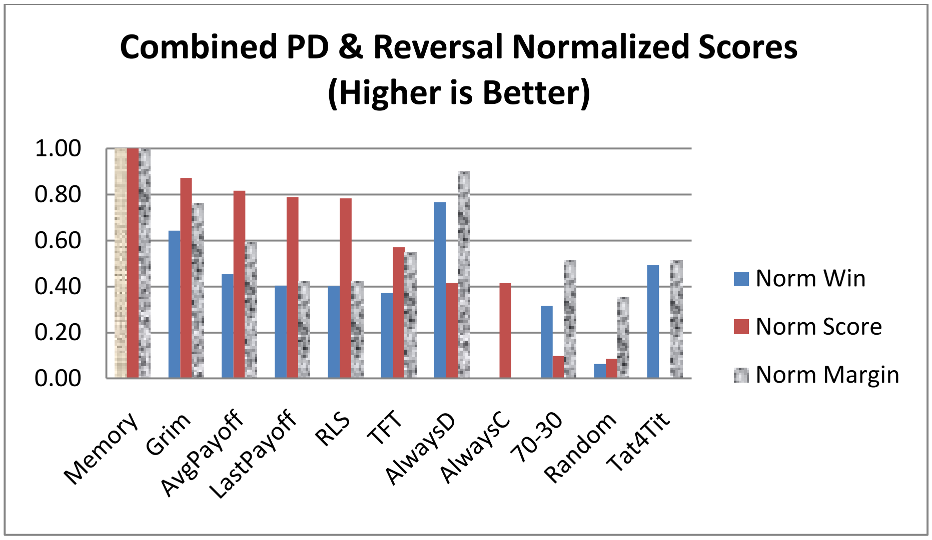

Figure 31: Normalized Combined Scores

This figure shows the degree to which the Memory-based approach dominated the competition in both tournaments. The results above demonstrate fairly convincingly that a memory-based model is a viable and effective approach to iterated game-theoretic situations, both for static and for dynamic games. It is at least as good, if not better than the standard game theoretic or heuristic approaches, and does so without the necessity of making *a priori* judgments about the game structure.

Application of the Model: A Three-Sided National Security Game

Taking the memory-based approach a step farther, and in an effort to refine model concepts and debug any logic errors prior to continuing to analysis of the two test cases, a three-sided game was developed to represent the essence of national security choices – war or alliance, support the economy or build the military. These play out in a two-level game where the choice between war and alliance is aimed outward at the other players in the game, while the choice between building the economy and building the military is primarily aimed inward, but both choices affect the other level. In an effort to keep the game focused on the "big" choices and to limit complexity, these simple choices are stand-ins for the wide variety of actions within a given rubric, so the choice for war (or conflict) covers all actions from threats, to moving troops to the border, to full-scale war. Similarly, the choice for alliance covers the ground between diplomatic consultations, military staff talks, *ententes*, and full-scale military alliance. The "guns vs. butter" choices on the domestic level cover similar ranges. Players can only make one of the four choices, and each choice has different payoffs on the international and domestic level. Players were also required to maintain positive scores in both military and economic realms – if either one dropped below zero, indicating total collapse of the economy or total destruction of the military, then the player was eliminated from the game and any remaining resources are absorbed by the "winning" player.

There is a natural tendency to dismiss such a structure as too simplistic, as there are many shades of meaning within each choice. Further, choosing only one action is clearly an artificiality, as states do not switch their efforts (and budgets) from 100% focus on one action to 100% focus on another action, but instead devote different amounts to support each portion of their military and domestic portfolios. All of these statements are true, but this is an attempt to capture the essence of state decisions while not getting bogged down in the details. There is always an emphasis on one facet of policy or another, as decision-makers' limited attention moves from one crisis or the next. Further, focusing resources on one thing means that other things are not done. Money spent on the military cannot be used (directly) to build the economy. Threatening another state generally means that you cannot offer alliance in the same breath. Finally, the choice to embark on a general course of action implies the possibility of other actions or outcomes within that space. Moving troops to the border risks the outbreak of full-scale war. A tentative offer of talks may be embraced by the other party and military alliance may result. Investing in the economy and attendant prosperity may well create new businesses and constituencies devoted to continual economic expansion.

Further, each player chooses moves in relation to the other two players, so in reality each player has 16 unique move combinations, which allows for considerable latitude. This latitude to conduct nuanced strategy is even more pronounced in the 8-player game discussed below.

One distinct advantage to this approach is that it is parsimonious – it focuses solely on "big picture" decisions. Further, the underlying process of situation recognition, hypothesis generation and decision are all the same process. Thus this approach is parsimonious not only in its view of national security choices, but also in its view of the process underlying those choices.

Game Move and Payoff Structure

The game payoff structure was determined by analysis of Correlates of War (COW) data version 3.02 during the period 1880-1980 for major powers, specifically the National Military Capability (NMC) variables of military personnel, military budget, iron and steel production, energy production, total population and urban population, and the Total Trade (TTrade) dataset. Alliance and conflict values were drawn from the Militarized International Disputes (version 3.02) and Alliance datasets (version 3.03) (Singer, 1987; Ghosen and Bennett, 2003; Small and Singer, 1969;

Stinnett, et al., 2002; Marshall and Jaggers, 2003). The values for move payoffs was based on analysis of the change in NMC and TTrade variables, so the net growth in the simulation stays within the range of the growth or decline of the underlying COW variables. After initial analysis of trends in the variables, including their absolute and relative changes in periods of both build-up and drawdown, the values were adjusted to create a game structure where there was no dominant game-theoretic strategy for either the domestic or international sides of the game, nor was there a dominant strategy for the combined game. General description of the moves and payoffs is as follows:

Build military: This encompasses all actions aimed at increasing military capability short of threatening another state. The payoff in military power for the international game is significant. Any positive economic impacts of the military build-up are more than offset by the opportunity cost to the overall economy, so it causes the economy and the domestic game to suffer a loss.

Build Economy: This choice means that the side is concentrating on building the non-military part of the economy, which generally creates positive gains in the domestic game, unless the state comes under threat. While building the economy can help create latent military power, emphasis on the civilian economy will have a net negative effect in the international game.

Offer alliance: This encompasses all actions in the diplomatic realm, such as consultations, staff talks, commercial negotiations, *ententes*, and alliances. The results are generally positive in both the international and domestic games, except when the other side rebuffs the effort through threats or ignores the effort via concentrating on the economy, and in that case the impact is somewhat negative due to wasted effort and loss of face. If both sides choose to offer alliance, then an alliance is concluded which remains in effect until one side chooses to break the alliance via threats. Both sides will then receive an increase in combat power (international payoff) proportional to their relative strength.

Threaten: This includes all threatening military action, including war. Threats have a generally negative effect in both the international domestic games, as combat power is wasted through wear and tear, and the economy is stressed by such actions. If both sides choose to threaten, then war breaks out, with unpredictable length and consequences, as shown below.

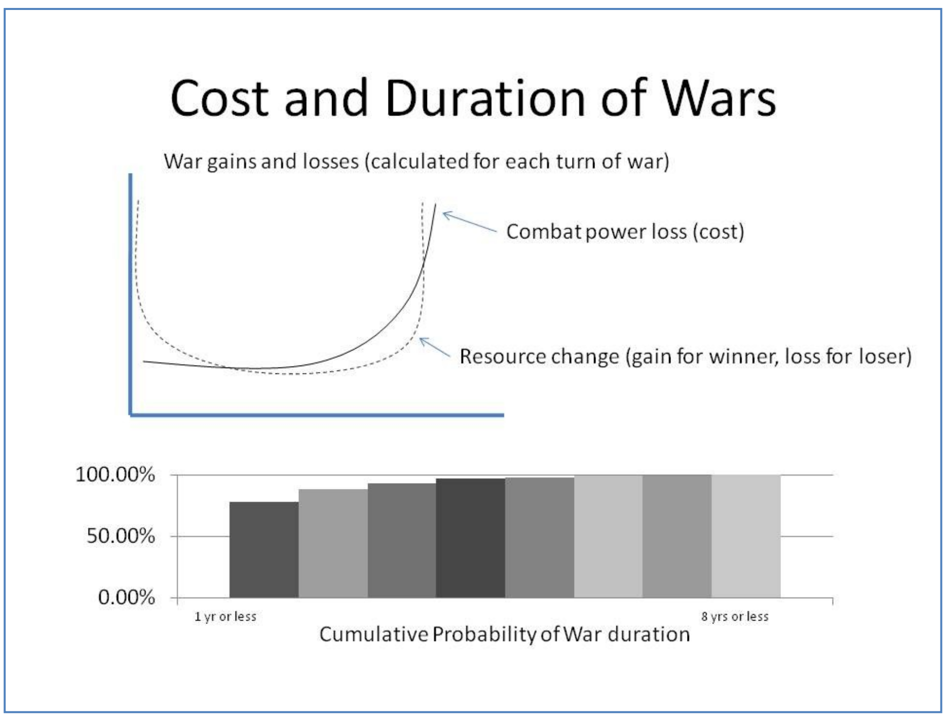

Figure 32: Cost and Duration of Wars

The war duration is determined by a random exponential draw, so most wars will be of one turn duration, but may go as high as seven turns. War results in loss of combat power by both sides, but the winner of the conflict may or may not receive domestic gain through capture of resources or territory, both costs and gains are determined by a random exponential distribution, as shown above.

The summary of the game's payoff structure is shown below.

Game Payoff Structure

External Game (payoff = combat/coercive power)

Player A, Player B	Build Military	Offer Alliance	Threaten	Build Economy
Build Military	2,2	2,1	0,-.25	2,-.25
Offer Alliance	1,2	var,var (pos)	0,-.25	1,-.25
Threaten	-.25,0	-.25,0	var,var (neg)	-.05,-1
Build Economy	-.25,2	-.25,1	-1,-.05	-.25,-.25

Internal Game (payoff = economic/soft power)

Player A, Player B	Build Military	Offer Alliance	Threaten	Build Economy
Build Military	-.75,-.75	-.75,1	0,2	-.75,2
Offer Alliance	1,-.75	3,3	-.25,1	1,2
Threaten	.5,0	1,-.25	var,var(+/-)	.5,-1
Build Economy	2,-.75	2,1	-1,.5	2,2

Figure 33: Game Payoff Structure

Game Logic Flow

During the course of the game, for each turn, each player simultaneously chooses a course of action toward each of the other two players (total of six directed dyads), after which payoffs are assessed. After a short initialization period where moves are drawn randomly to provide a baseline for assessing payoffs, the sides begin to choose their moves based on the existing situation and the memory of previous moves. First, each player assesses the recent history – what sequence of moves by each of the other players led to the current situation? With that sequence in mind, the player then probes the existing set of semantic concepts one at a time to try to find a match. Under crisis or pressure situations, the amount of time that players have to probe semantic memory is limited, and partial matches may be the only results. If there is no complete or partial match, a new semantic concept is created to describe the current situation. If there is a matching semantic concept, then the long-term store of memory is probed to see if this situation has occurred, and if so, to test hypotheses about the next likely move by the opponent, and which move will produce the best outcome given the situation and the opponent's likely move. Once the hypothesis about the predicted opponent move is confirmed in memory, the best possible move for that situation will be chosen.

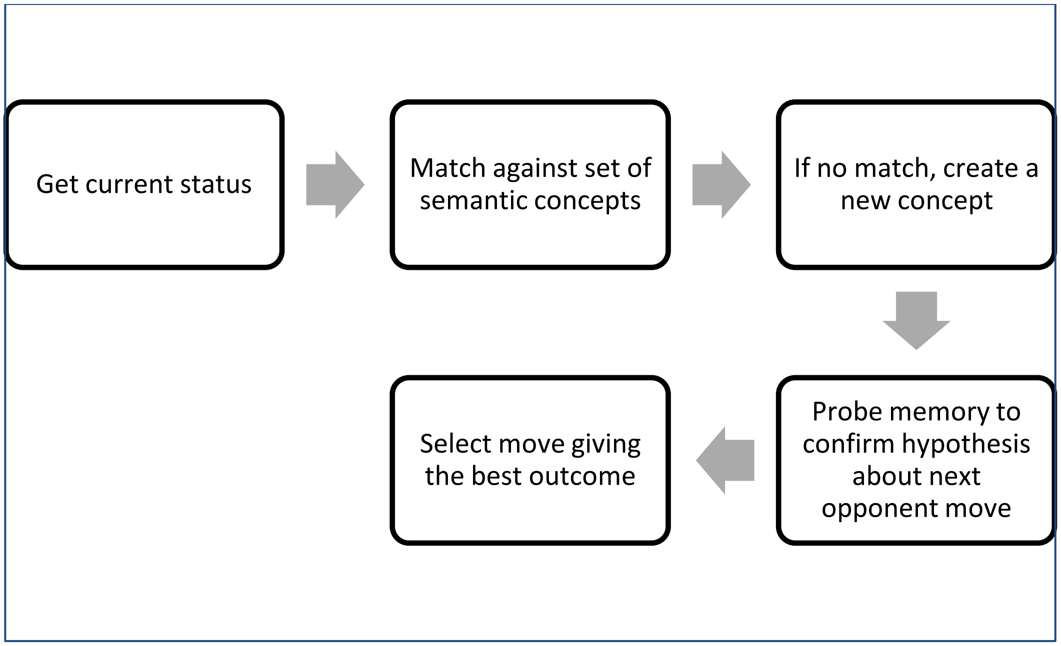

Figure 34: Logic Flow for the 3-player Game

Initially, the order of semantic concepts is random. Once the game begins, and after each move is selected, the accompanying semantic concept is moved to the top of the semantic concept stack, where it will then be the first concept considered the next move. Thus, the decision system, particularly under crisis modes, will tend to give more weight to recent events rather than to events long in the past, and events which occur repeatedly will tend to remain near the top of the stack, which replicates the well-known recency and rehearsal biases in the decision-making literature (Tversky and Kahneman, 2001a [1974]). One can think of the set of semantic concepts as a box of wrenches, and the matching of concept to situation as picking the right-sized wrench for the job. After each wrench is used, it is put back into the pile of wrenches in the toolbox, so the same wrench is on the top when reaching for the next one. The result of this is that when tasks are repeated, the most commonly used wrenches end up near the top, and are thus quickly found. In the figure below, the initial set of semantic concepts in the first move is checked for a match. In this case, SC4 is the best match, and so that semantic concept is moved to the top of the stack for the next move. The end result is that concepts which are used often are kept at or near the top of the stack, making subsequent searches more efficient.

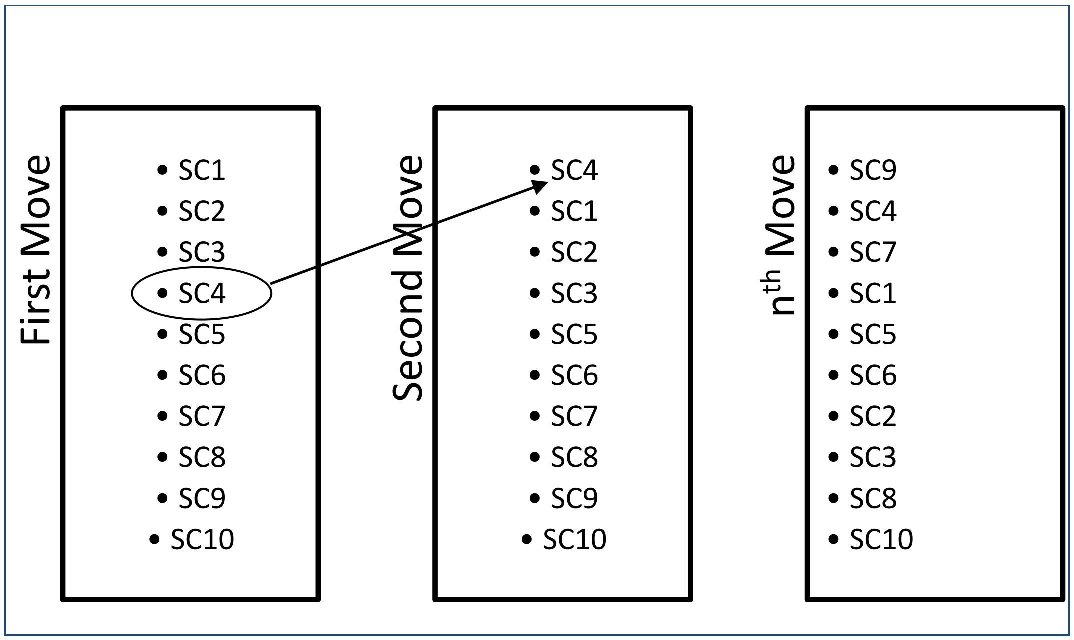

Figure 35: Operation of the Semantic Concept Stack

Further, and as discussed earlier in the chapter, the record of events forms traces in memory, which can be either correctly or incorrectly encoded, and which will become degraded over time (both adjustable within the model via parameters). Thus more recent events will tend to carry more weight, although the entire record of events is available to the decision-maker. Additionally, the model has provisions to degrade or even deny player knowledge of opponent strategy and payoffs, so the game can be played in states of either full or incomplete information, allowing the impact of misperception to be explicitly modeled.   This also allows situations as shown in the figure below to be modeled. In this case, the decision-maker in State A cannot directly perceive State B's intentions, but may or may not have an existing judgment or set of semantic concepts (or inferences) about those intentions. If so, the set of semantic concepts shape the way in which the State A decision-maker views the environment, which is formed by State B's actions, along with the actions of all other actors in the system. The decision-maker for State A then updates judgments about State B based on the interaction between the environment, the sample of the environment, and the existing set of previous judgments about State B. The parameters in the model allow for exploration of the role of existing inferences or concepts, the accuracy with which the decision-maker samples the environment, and the difference between State B's intentions (chosen strategy) and the strategy perceived by State A in the existing environment.

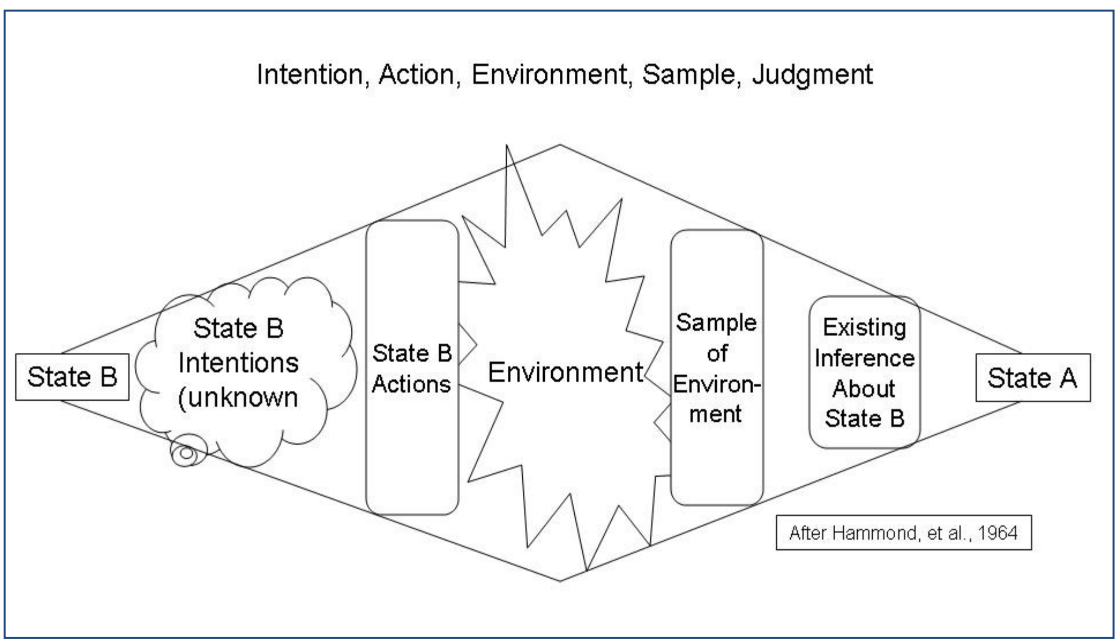

Figure 36: Factors Affecting State Judgments

Additional Model Features

Further, the model also incorporates features to explicitly study the role of various tradeoffs. For example, one important controversy in the IR literature is whether states are more concerned with absolute than relative gains. In a relative gain situation, states are predicted to be more conflict-prone as they jockey for relative position. In the model, the emphasis on relative or absolute gains is modeled probabilistically, with a random draw for each decision point to determine whether the model will make judgments based on relative or absolute gains. This can also be thought of as modeling the difference between a decision-maker satisfying internal requirements (absolute gains) or maintaining position in the international arena (absolute gains), per Bueno de Mesquita (2002). Initial tuning runs of the model reveal exactly that behavior as the emphasis is shifted from absolute to relative gains. Similarly, the alliance behavior can be adjusted across a range between less powerful states seeking alliances in an effort to balance against the more powerful state, and less powerful states may instead emphasize bandwagoning behavior and attempt to ally with the most powerful state (and vice versa for balancing). Initial tuning runs of the model reveal that this feature also works as expected, so that these types of behaviors can also be studied in future efforts.

Adjustments to the Model

MEMORIS-2 uses four additional parameters during memory trace generation and decay: state situation uncertainty, system situation uncertainty, a learning parameter, and an exponential decay parameter. The state situation uncertainty parameter models uncertainty in the assessment of the characteristics of a given state, and the situation uncertainty parameter models uncertainty in the assessment of relations between two states. Both parameters were set fairly high, at .95 and .98 respectively, indicating a high, but not perfect accuracy in assessing the status of the situation. The learning parameter models the probability that the situation will be correctly recognized and controls whether or not the situation is correctly encoded into memory. The exponential decay parameter controlled the rate at which memory degrades, and this was set so that memory traces remained fairly accurate for 15-25 years, then become increasing degraded, with memory becoming totally degraded after 100 years. While all of these values are fairly reasonable, only a limited amount of sensitivity analysis was performed to assess the impact of changing these values, and the model was not "tuned" using these values in an attempt to obtain better results. It should be pointed out that the nature of the dataset, with most values concentrated near the zero end of the 0-1 scale, results in fairly sparse memory trace vectors and consequently fairly low data content within the traces, which should lead to the model being relatively insensitive to changes in the uncertainty parameter values. As discussed earlier, the rate at which memory decays may well have important effects which will bear further investigation in the future.

The 3-state model was tuned by adjusting the various parameters to cause the incidence of conflict within the model system to roughly mirror the incidence of war in the data. Using EUGene 3.03, and using the population of European directed dyads between 1816 and 1970 (cf. Bueno de Mesquita 1992), approximately 5% of the directed dyad-years also had Militarized Interstate Disputes, which in this case, counts each year of continuing conflict as a separate dyad-year. Results of 1000 model runs after adjustments show that the incidence of conflict roughly matches the MID data.

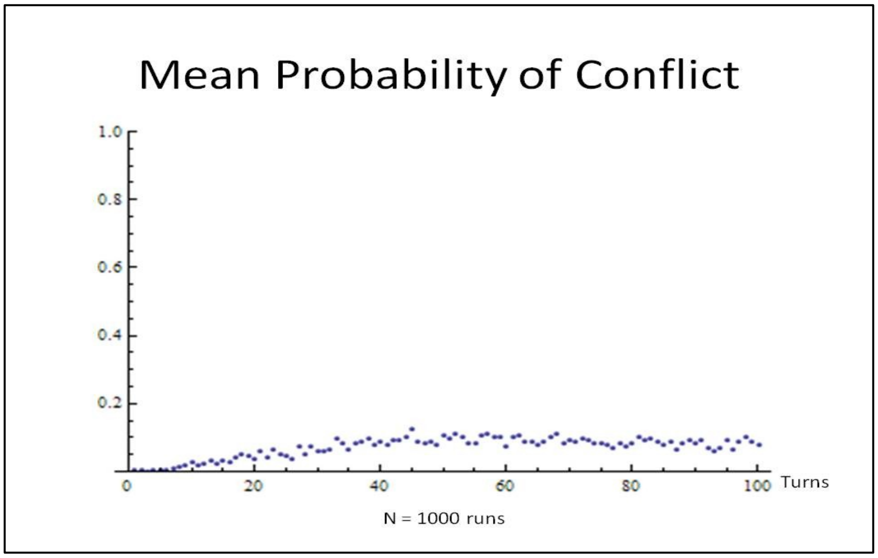

Figure 37: 3-Sided Model Probability of Conflict

Results of the 3-Player Model

With initial model tuning complete, a test series of runs was used to see if the model would yield recognizable patterns or groupings of results. Since the actual simulations will be using Correlates of War data, which is largely year by year, each move in the simulation will be equal to one year of real time. The actual cases will be relatively short periods of history (50 years or less), so the total model run will be 70 years: 20 years of initialization and 50 years of actual simulation of moves.

First Sets of Runs: Random Moves

For these runs, the 20 years of initialization was based on randomly chosen moves. All three players started with individual sets of random moves for 20 years, and then the simulation was allowed to run for an additional 50 years. Each set of moves were divided between those which were generally hostile (1=build up military, 3=threaten) and those which were generally friendly (2=offer alliance, 4=build up economy), and the percentage of each move was varied between 90% hostile/10% friendly and 10% hostile/90% friendly as shown. Each set of runs was replicated 1000 times, for a total of 5000 simulation runs.

	Build Military	Offer Alliance	Threaten	Build Economy
Case 1	45%	5%	45%	5%
Case 2	35%	15%	35%	15%
Case 3	25%	25%	25%	25%
Case 4	15%	35%	15%	35%
Case 5	5%	45%	5%	45%

Table 8: Move Distribution for Simulation

After completing 5000 simulation runs, the resulting overall data was collected:

- Total score (final resource + final combat power)
- Resource score (final resource)
- Combat score (final combat power)
- Ally score (total alliances/turn)
- War score (total war/turn)

The values are summarized in the table below. Notice that the Resource, Ally and War variables are all positively skewed, showing that these results tend to have concentrated values near the low end of the spectrum. The Resource variable also is fairly "peaked", with more values concentrated toward the middle, when compared to the normal distribution, and the other variables tend to be more spread out. This type of relationship will repeat itself throughout the analysis.

Variable	n	Mean	Std Dev	Minimum	Maximum	Skewness	Kurtosis	K-S Stat	p-value
TOTAL	5000	956.74	122.72	526.87	1283.04	-0.36	-0.47	0.05	<.01
CP	5000	550.37	105.72	262.72	868.22	-0.06	-0.69	0.04	<.01
RES	5000	310.23	126.26	149.64	761.38	1.34	0.50	0.22	<.01
ALLY	5000	0.97	0.64	0.00	3.64	0.64	-0.42	0.10	<.01
WAR	5000	0.26	0.20	0.00	1.12	0.82	0.14	0.13	<.01

Table 9: Summary of 3-Player Results

With the first set of data in hand, the results were analyzed to see if specific sets or patterns developed. If the model replicates the dynamics seen in the real world, the results should show specific outcomes matching (at a minimum) conflict, peace and some intermediate state of tension. Similarly, the numbers of wars and alliances within the system should match with the general set of outcomes.

Cluster Analysis: Determining the Appropriate Number of Clusters

The results data was analyzed for patterns of association, or clusters, using a two-stage cluster analysis, as suggested by Hair (Hair, et al., 1998). The SAS k-means clustering procedure (FASTCLUS) was used, as that allowed particular cases to move in and out of a cluster as new points are considered. This is more robust for this type of data than the hierarchical clustering procedure (CLUSTER), which locks a case to a particular cluster. This can often create situations where the particular clustering result is dependent on the order of the data, rather than the actual values of the data. K-means clustering is not without its potential pitfalls, as the clusters are extremely dependent on the number of original clusters chosen by the analyst. To ameliorate this issue, the analysis was performed using a range of starting assumptions in an attempt to discover the most robust approach. Hierarchical clustering was used to determine the number of clusters, followed by k-means clustering to perform the actual clustering. After the data was entered into SAS, the data was normalized using the SAS STANDARD procedure and then a hierarchical cluster procedure was performed using the CLUSTER procedure. Results of the cluster analysis are shown in the table below.

The Cluster Procedure				
Ward's Minimum Variance Cluster Analysis				
Eigenvalues of the Covariance Matrix				
	Eigenvalue	**Difference**	**Proportion**	**Cumulative**
1	2.9922	0.6905	0.3740	0.3740
2	2.3017	0.9804	0.2877	0.6617
3	**1.3213**	0.6028	0.1652	**0.8269**
4	**0.7185**	0.4430	0.0898	**0.9167**
5	0.2755	0.0636	0.0344	0.9512
6	0.2119	0.0557	0.0265	0.9777
7	0.1562	0.1336	0.0195	0.9972
8	0.0226		0.0028	1.0000

Table 10: Hierarchical Cluster Analysis

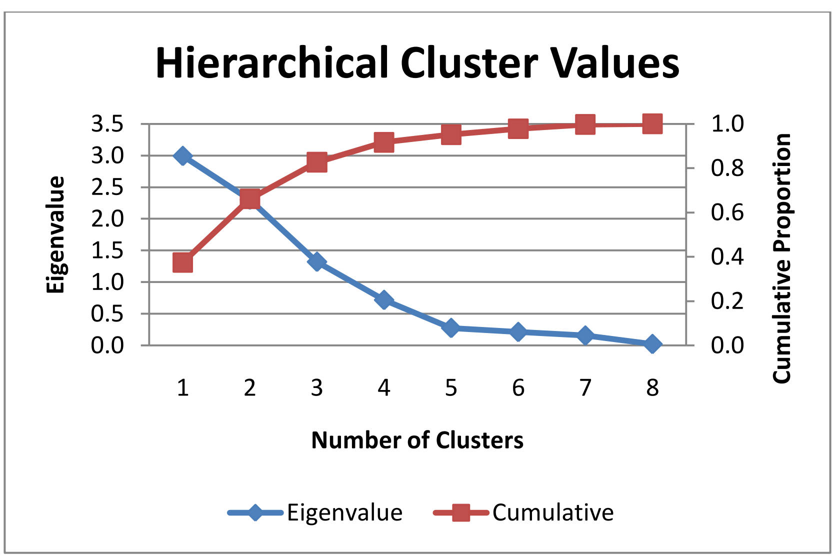

Figure 38: Hierarchical Cluster Analysis

There are no hard and fast rules for choosing the correct number of clusters, but the normal rules of thumb are based on the eigenvalues, the eigenvalue plot, and the cumulative proportion plot, and then analyst judgment of the actual clusters. Based on the results above, I expected that the final result would be either three or four clusters. Following Hair's procedure, I then produced a dendogram, or tree plot of the cluster values. This shows in a different format how the clusters tend to break apart, and can give additional guidance on the best number of clusters.

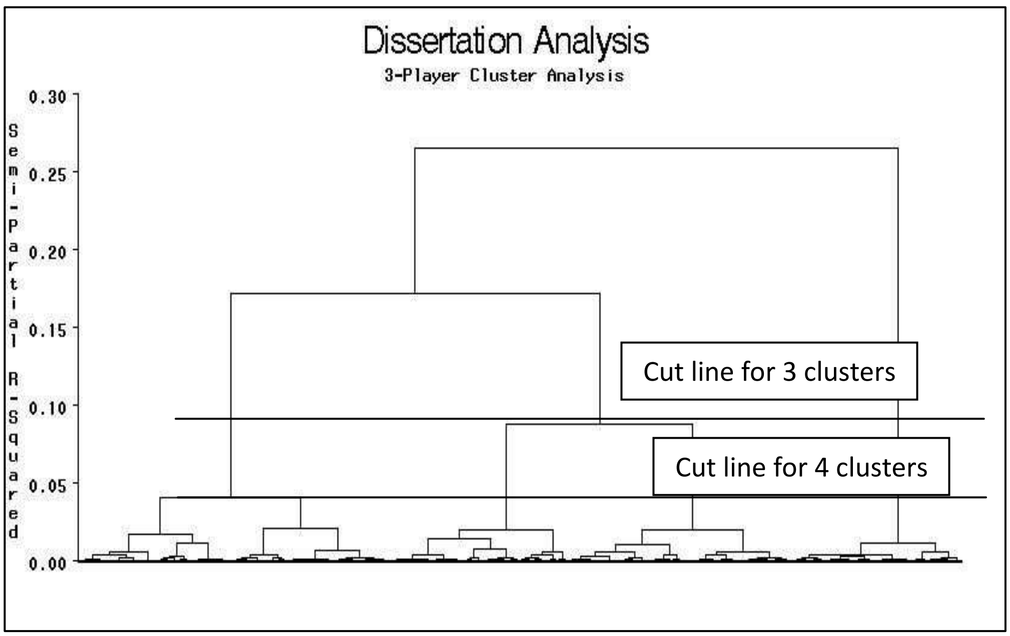

Figure 39: Dendogram of Clusters

The dendogram shows the same basic situation as the eigenvalues. The height of each "limb" of the tree is the amount of pseudo-R^2 or variance that each cluster provides. Once the heights of the limbs get short, that means that the additional clusters are not providing much additional explanation of the variance in the data. The division into two clusters accounts for about 27% of the variance, and the subsequent division into three clusters adds another 18% of variance. When the center cluster divides into two additional clusters, that division picks up about 8% of variance. The "cuts" in the tree are at the point where additional clusters give 10% and 5% of additional variance – which can easily be seen as the points of diminishing returns. These two cuts also show that the number of clusters should be either three or four.

The final criterion is analyst judgment – how well do the clusters actually work in describing the data. For that, I turn to an additional clustering algorithm, k-means clustering. This is a very quick way to generate clusters, but its reliability is critically dependent on picking the proper number of clusters. This is why the two-stage procedure was used – the hierarchical stage generates the proper number of clusters, and then the k-means stage creates the actual clusters for analysis.

Analysis with Three Clusters

To test the three cluster solution, I used the SAS k-means clustering procedure (FASTCLUS), specifying the entire set of variables as data and three clusters. As shown in the figure below, the procedure produced three distinct clusters which I characterize as follows:

- Conflict (1749 cases): High levels of conflict, low levels of alliance and resource, mean levels of combat power.
- Standoff (2182 cases): Low levels of conflict, mean levels of alliance, low levels of resources, and high levels of combat power.
- Alliance (1069 cases): Low levels of conflict, high levels of alliances, high levels of resources, and low levels of combat power.

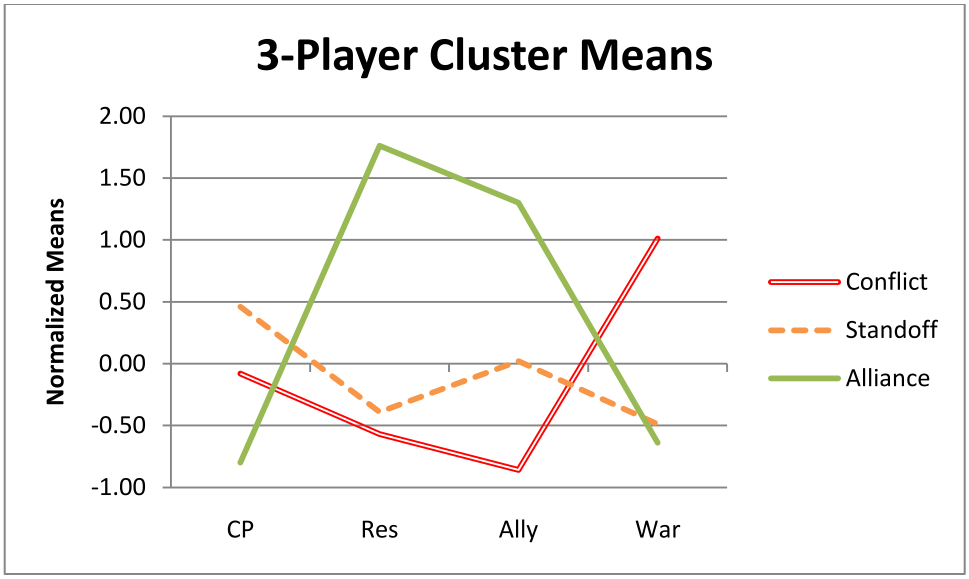

Figure 40: Normalized Means for 3 Clusters (n=5000)

This breakout of clusters appears to have considerable face validity. In both the "Conflict" and "Standoff" cases, the emphasis is on developing combat power rather than resources. In the "Conflict" cluster, players spend a lot of effort developing combat power, and then that power is used up during conflict, so the resultant combat power state tends to stay low. In both the "Standoff" and "Alliance" cases, the number of alliances is at the mean value or much higher, which means that there is a fairly active and long-lasting alliance structure in both cases, where in the "Conflict" case, alliances are rarely made and quickly broken.

The three clusters also break out nicely from a statistical viewpoint. An Euclidean measure of distance between clusters, Mahalanobis distance (D2), was used with an F-test to determine whether the clusters are actually statistically different. As can be seen in the table below, the difference in each of the variables is highly significant across clusters ($p<.0001$), and the Scheffé test shows that the inter-cluster variation is also highly significant (p<.0001) for all cluster combinations (1-2, 1-3 and 2-3).

	n	CP	Res	Ally	War	Scheffé
Conflict	1749	-0.08	-0.57	-0.86	1.01	<.0001
Standoff	2182	0.46	-0.39	0.02	-0.49	<.0001
Alliance	1069	-0.80	1.76	1.30	-0.64	<.0001
R^2		0.65	0.84	0.59	0.55	
F-test		4700	12761	3654	3096	
Significance		<.0001	<.0001	<.0001	<.0001	

Table 11: 3-Player Cluster Statistics

Analysis of Four Clusters

The three cluster results appear solid and provide good explanatory power, but given the results from the hierarchical cluster analysis, four clusters may also be a contender. Recall from the dendogram in Figure 38 that the fourth cluster is a subdivision of the central cluster, so I expected that the four cluster case would keep the “Conflict” and “Alliance” clusters and then subdivide the central “Standoff” cluster.

As before, I used the SAS k-means algorithm, and the results are as follows:

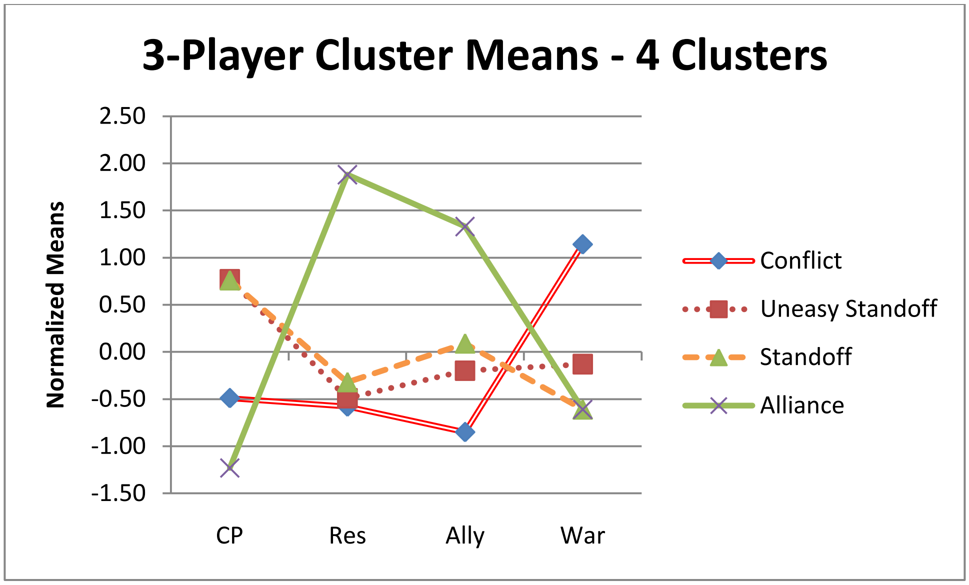

Figure 41: Means for 4 Clusters (n=5000)

The four clusters shown are described below:

- Conflict (1491 cases): High levels of conflict, low levels of alliance and resource, low levels of combat power.
- Uneasy Standoff (928 cases): Mean level of conflict, low levels of alliance and resource, and high levels of combat power.
- Standoff (1605 cases): Low levels of conflict, mean levels of alliance, low levels of resources, and high levels of combat power.
- Alliance (976 cases): Low levels of conflict, high levels of alliances, high levels of resources, and low levels of combat power.

Comparing these results to the 3-cluster analysis in Figure 39 shows that, as expected, the central "Standoff" cluster generally divided into two subdivisions, largely based on the level of conflict. Notice also that the overall level of combat power has increased slightly in the two central clusters, and decreased in the "Conflict" cluster, which shows that about 300 cases with moderate levels of conflict (and consequently lower levels of combat power) that would have originally been in the “Conflict” cluster have now shifted to the two “Standoff” clusters. As expected, the “Alliance” cluster was virtually unchanged, with a difference of less than 100 clusters.

These clusters are fairly descriptive, although the two central clusters share virtually the same levels of combat power, and the “Standoff” and “Alliance” clusters share the same levels of conflict, which diminishes their descriptive power in terms of distinguishing between outcomes. Looking at the statistical results in the table below, each cluster does show highly significant differences in each of the variables ($p<.0001$), and the additional cluster does increase R^2 values by a small amount (1-3%), which is probably simply due to the extra variable in the mix. However, the Scheffé test shows that the clusters themselves are not necessarily distinct, with the overlapping values discussed above making the inter-cluster distances become non-significant between clusters 2 and 3 (Uneasy Standoff – Standoff), and 3 and 4 (Standoff-Alliance).

	n	CP	Res	Ally	War	Scheffé
Conflict	1491	-0.49	-0.58	-0.85	1.14	<.0001
Uneasy Standoff	928	0.77	-0.49	-0.20	-0.13	(2-3)
Standoff	1605	0.76	-0.32	0.09	-0.61	(2-3),(3-4)
Alliance	976	-1.23	1.88	1.33	-0.61	(4-3)
R^2		0.66	0.87	0.56	0.58	
F-test		3292	10776	2255	2347	
Significance		<.0001	<.0001	<.0001	<.0001	

Table 12: Statistical Results for 4 Clusters

Conclusion of Cluster Analysis

Based on the above results, the case for three clusters looks to be solid, and the remainder of the analysis will proceed with three clusters as the best combination of descriptive power and statistical significance. As stated above, if the simulation is working as expected, it should show a changing distribution of results in accordance with the changes in inputs. Recall that the distribution of threat vs. peaceful moves was changed from 90%/10% in the first set of runs through 10%/90% in the final set. Based on that, results should show a change in the distribution of cluster results, with more threatening environments eliciting more conflict, and less threatening environments resulting in less conflict. The actual results are as shown:

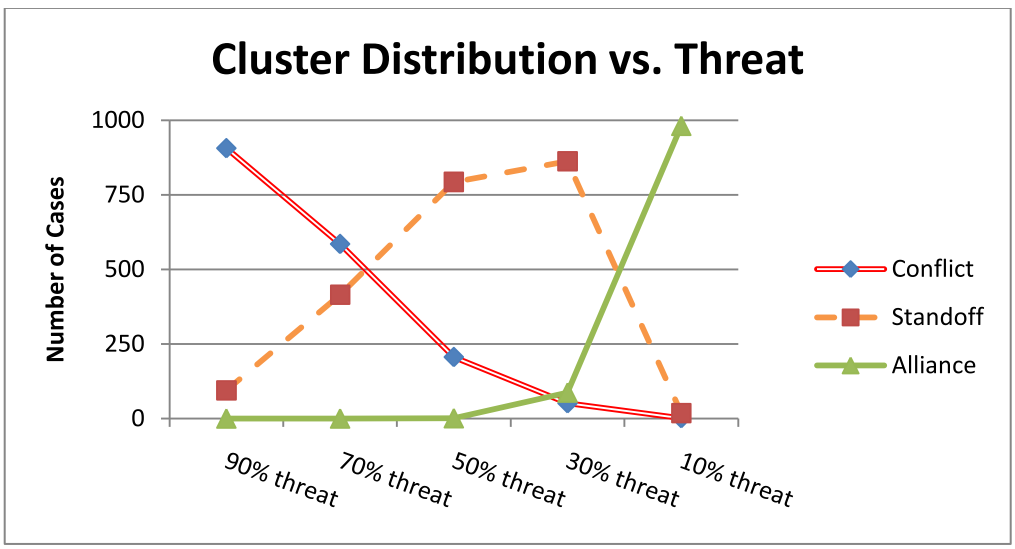

Figure 42: Cluster Distribution vs. Threat (n=5000)

	90% threat	70% threat	50% threat	30% threat	10% threat	Total	Test	Value	Signif.
Conflict	906	585	206	51	1	1749	chi-sq	6549	<.0001
Standoff	94	415	793	862	18	2182	contingency coeff	0.7531	
Alliance	0	0	1	87	981	1069	Cramer's v	0.8093	

Table 13: Statistical Results of Clusters vs. Threat

The above results clearly show that the simulation behaves as expected, and that the results are highly statistically significant. With high or even moderate levels of threat, the cluster results are entirely divided between the Conflict and Standoff clusters, indicating patterns of high armaments, high conflict, or both. However, once the conflict level decreases to lower levels, there is a rapid shift in the distribution to an almost entirely peaceful situation at very low levels of conflict.

These initial results, while by no means definitive, are certainly suggestive that the three-sided model is replicating some of the dynamics predicted in the IR literature and demonstrated in the COW data, particularly in the areas of conflict and alliance behavior. Further, the model creates clusters of results which have face validity, and react to changes in the general threat level as expected. With these initial results in hand, it is time to turn to expanding the three-sided game into a large enough game to model the chosen test cases.

Developing the multi-state simulation model: MEMORIS-2

The MEMORIS-2 model is a straight-forward extrapolation of the three-sided game discussed above, with the exceptions of increased number of players (8), rescaling the payoff structure, and the requirement of a mechanism to initialize the simulation using real-world data instead of through random moves.

It was necessary to rescale the payoff structure to allow for the increased number of players. If the payoff values had remained constant, than the five additional move choices would have (potentially) increased the possible total payoff by a factor of 2.5, which would have created unrealistic growth rates. The overall payoff structure was simply arithmetically rescaled to keep the maximum possible growth or decline rates to +/- 10%.

The increase in number of players is a (relatively) simple programming task, requiring only that the Mathematica simulation logic be expanded to allow for more players, and increasing the size of the several arrays and data structures within the program code.

Data for Initialization

Initialization of the model through existing COW data is somewhat more complex. For each iteration (one year of time), the model will need to input the dyadic data vectors for all states in the system. These vectors are then converted to sets of features, with values transformed from COW data for the given year and rescaled to match the model's payoff structure.

These sets of features form an extended trace vector, composed of (0,1) values. This trace vector represents the episodic memory of a state's interaction with another state for that year. This trace is then read into memory, where is becomes part of the state's total experience with all other states in the system. As stated earlier, the trace can also be encoded incorrectly into memory, which would model the action of either misperception or deception within the system. Once the trace is stored in memory, traces from previous years are degraded using an exponential decay function to simulate gradual degradation of information over time.

The above process is relatively simple, and can be automated. The more difficult aspect is to attribute player strategies to match up with the changes in scores over time. Of course, if MID or alliance data shows that conflict or alliance was initiated in a given year, then the choice is simple, but otherwise the choice of imputed strategy is subjective. The historical record helps with breaking ties, but as with all coding endeavors, the exact coding decision is necessarily tentative and subject to revision.

Once the coding of the initialization data is complete, that data is read into the MEMORIS-2 memory, and the simulation begins from that point. Each player will go through the identical process of situation recognition, comparison to semantic concepts, hypothesis generation and testing, and selection of chosen strategies, as in the previously-discussed three-player model. Once this inference engine completes its operation, MEMORIS-2 then records key variables, completes any required updates of existing memory traces, and then steps to the next year. As before, a Monte Carlo process will be used to determine the outcome of probabilistic events such as war, and to determine whether the player is seeking relative vs. absolute gains, and whether the player is engaging in balancing or bandwagoning alliance behavior.

Having laid the groundwork and successfully developed the model, I now turn to analysis of the two test cases.

Chapter 4: Europe 1885-1914: The Great War That Was

Introduction

As stated in Chapter One, the study of decision offers two competing lines of analytical approach, one of rational choice which can be modeled via the tools of expected utility, and one of situational recognition and use of rules which can be modeled from the cognitive angle. This returns us to how the interaction between the situation and the way in which decision-makers process information affects subsequent actions. The situation in Europe between 1885 and 1914 provides an opportunity to explore this question. In this situation, there are a relatively small set of states with a wide variety of conflict and alliance interactions, leading up to a conflict involving most, but not all, of the participants.

The Situation

World War I was a tragedy of enormous proportions. Despite the optimism of the end of the 19th century (c.f., Angell, 1910), the world soon found itself embroiled in a conflict that was beyond the imagination of any decision-maker, military or civilian. Such was the calamity that the outbreak of World War I is one of the most analyzed periods in history. The number of causes proposed for the outbreak of war is as diverse as the authors writing about the subject. These causes include: the actions of Germany, the existence of Germany itself, interlocking alliance structures, secret alliances, offense-defense imbalances, the cult of the offensive, mobilization races, crisis instability, military plans, misperception, the situation in the Balkans, the naval armaments race, militarism, greed, and the structure of European power. While it is well beyond the scope of this effort to analyze or predict the causes of World War I (see, among many others: Farrar, 1972; Holsti, 1965; Hermann and Hermann, 1968; Holsti, North and Brodie, 1968; Keegan, 2000; Massie, 1991, Tuchman, 1962, Joll & Martell 2007, Strachan, 2004), I will not add a new cause to that already long list, other than to show *probabilistically* how the combination of the external situation and the processing of that information created a greatly increased probability of major systemic war in Europe between 1885 and 1914. I will also show that, even though the incentives for a given action remained constant throughout the period, that the pattern and pace of events created a situation where hostile actions were more attractive than peaceful actions.

Data for the Simulation

The MEMORIS-2 model is able to readily handle eight state players at a time. In theory, the model itself is not arbitrarily limited to a certain number of players. Each increase in the number of players, however, causes an exponential increase in run-times, roughly doubling for every three extra players. With eight players, a normal set of 1000 simulation runs takes roughly eighteen hours on a quad-core desktop computer which I hand-built specifically for running Mathematica.

The eight states chosen were: The United Kingdom, France, Germany, Austria-Hungary, Italy, Russia, Turkey, and Serbia. While this certainly leaves out some potentially key players, such as Romania, Greece, and Belgium, the major decision-makers for the critical events leading to the war are all present. Further, the model shows roughly the same dynamic in the three-player game in Chapter Three as it does in the eight-player European scenario, so it is unlikely that the inclusion of other states would significantly change the results. Further optimization of the model, faster computers, and of course, more time will allow for expanded analysis, but that is another project.

As stated in the previous chapter, this set of simulation runs is drawn from the existing Correlates of War datasets. Specifically, data is drawn from the National Military Capabilities, Militarized Interstate Disputes, Formal Alliances, and Total Trade datasets (Singer, 1987; Ghosen and Bennett, 2003; Small and Singer, 1969; Stinnett, et al., 2002; Marshall and Jaggers, 2003). The specific variables drawn from each dataset are summarized in the Table below:

Dataset	Version	Variable(s)
National Military Capabilities	3.02	Total population (tpop), military personnel (milper), military expenditures (milex), iron and steel production(irst), energy production (energy), urban population (upop)
Militarized Interstate Disputes Dyadic Data	3.02	HostLevA, HostLevB
Formal Alliances	3.03	SSType
Total Trade	**3.0**	ttrade

Table 14: Data variables and sources

It should be noted that the variables above are not normally distributed. Continuous variables form either unimodal or bimodal distributions with most values clustered to one end or the other of the value range. The data from the Militarized Interstate Disputes and Alliance databases are discrete on a 0-4 or 0-5 scale, again with most values clustered at the low end of the scale.

Factor Analysis of NMC Data

As seen in the previous chapter, the existing MEMORIS-2 model uses four variables, Resource, Combat Power, Alliance and War for its calculations. This required a conversion from the above data into a form which was readily usable by the simulation model. My intuition was that the NMC data along with Total Trade measures two basic factors, military strength and economic strength. This led to an initial try at grouping the variables as follows:

Economic strength (Resource) = f(energy, upop, ttrade) (10)

Military strength (Combat Power) =s f(milper, milex, irst, tpop) (11)

To test this intuition, I conducted a factor analysis on a subset of the COW data using SAS FACTOR procedure. I used an original exploratory set of five factors, which resulted in the following eigenvalue matrix and factor patterns for the principal components.

	Eigenvalues of the Correlation Matrix: Total = 7			
	Eigenvalue	**Difference**	**Proportion**	**Cumulative**
1	**4.559**	2.651	**0.651**	**0.651**
2	**1.909**	1.556	**0.273**	**0.924**
3	0.353	0.260	0.050	0.975
4	0.093	0.039	0.013	0.988
5	0.054	0.036	0.008	0.996
6	0.018	0.004	0.003	0.998
7	0.014	---	0.002	1.000

Table 15: Factor Analysis Eigenvalues

Factor Pattern					
	Factor1	**Factor2**	**Factor3**	**Factor4**	**Factor5**
MILPER	**0.5209**	**0.8330**	-0.0274	0.1671	0.0465
MILEX	**0.8193**	0.3442	-0.4502	-0.0284	-0.0687
ENERGY	**0.9118**	-0.3938	0.0293	-0.0375	-0.0432
IRST	**0.8921**	-0.3403	0.2202	0.1674	-0.0986
UPOP	**0.9776**	-0.1066	0.0496	-0.1460	-0.0122
TPOP	0.4284	**0.8382**	0.3120	-0.1167	-0.0050
TTRADE	**0.9223**	-0.3340	-0.0240	0.0149	0.1881

Table 16: Initial Factor Pattern

Variance Explained by Each Factor				
Factor1	Factor2	Factor3	Factor4	Factor5
4.5595	**1.9087**	0.3531	0.0933	0.0540

Table 17: Variance Explained by Factors

The above factor results show that the vast majority of the variance in the data (92.4%) is explained by the first two factors (Table 15, right column). Additionally, the factor pattern, which shows the amount of variance in each variable explained by the corresponding factor, shows that each variable has over 50% variation explained by the first two factors, and then only weakly for the last three (Table 16). This shows that the initial impression of two distinct factors looks pretty good.

The cluster analysis was re-run with just two factors, giving the following results:

	Factor Pattern	
	Factor1	**Factor2**
MILPER	**0.52087**	**0.8330**
MILEX	**0.81934**	0.34418
ENERGY	**0.91178**	-0.39384
IRST	**0.89208**	-0.3403
UPOP	**0.97755**	-0.1066
TPOP	0.42838	**0.83819**
TTRADE	**0.92228**	-0.33402

Table 18: Factor Pattern with Two Factors

MILPER	MILEX	ENERGY	IRST	UPOP	TPOP	TTRADE
0.9652	0.7898	0.9864	0.9116	0.9670	0.8861	0.9622

Table 19: Variance Explained by Factors

The above results confirm that two factors explain the vast majority of the total variance, and that for individual variables, the two factors explain about 80% or better of the variance in each individual variable. However, from the factor pattern, the factor loadings do not follow the initial intuition. This is not unusual for principal components factor analysis, as it tries to account for as much variance with the first factor as possible, so all of the variables except tpop load at 50% or higher on this first factor. The general solution at this point is to use Varimax factor rotation, which orthogonally rotates the factors to produce a situation where each factor will have either high or low correlation with each variable. Thus, Varimax rotation tries to divide the variables between the factors rather than loading the maximum variance on the first factor, as seen earlier in the principal components analysis. With this in mind, Varimax factor rotation was used to produce the final factor set, with the results as follows:

Rotated Factor Pattern		
	Factor1	Factor2
MILPER	0.1083	**0.9765**
MILEX	**0.5891**	**0.6654**
ENERGY	**0.9924**	0.0404
IRST	**0.9514**	0.0801
UPOP	**0.9271**	0.3277
TPOP	0.0227	**0.9410**
TTRADE	**0.9759**	0.0988

Table 20: Rotated Factor Pattern for NMC Data

MILPER	MILEX	ENERGY	IRST	UPOP	TPOP	TTRADE
0.9652	0.7898	0.9864	0.9116	0.9670	0.8861	0.9622

Table 21: Variance Explained by Rotated Factors

As expected, the rotated factor pattern breaks fairly cleanly into two pieces, with milper and tpop in the second factor (combat power) and irst, energy, upop and ttrade into the first factor (resources). Note that the milex variable shows at higher than .5 in both factors (Table 20), but more strongly in the second factor, which follows the intuition that milex should also be in the combat power factor.

Resources = *f*(irst, energy, upop, ttrade) (12)

Combat Power = *f*(milex, milper, tpop) (13)

This is slightly different than the original intuition, which originally had iron and steel production in the combat power column, but clearly such industry is also a big part of the overall economic strength of a state. The showing of military expenditures in both factors is also worthy of note, but I can certainly make a logical case that while military expenditure is primarily related to combat power, such expenditures are often used to buy military hardware, which clearly has an impact on the overall economy. This case is, of course, made quite often by any Congressman who has a defense contractor in his or her district!

Conversion of Variables

With this analysis in hand, converting the COW NMC and Trade data to MEMORIS-2 input data was fairly straightforward. Since the variables had several orders of magnitude differences in scale, all variables were rescaled to a 0-1 scoring range based on the total range of each variable for the period 1885-1905. The individual scores were added together (irst, energy, upop and ttrade for resources and milper, milex and tpop for combat power) and each result multiplied by 500 as a scaling factor. Based on the 1885-1905 data, this left Serbia with a zero score for both factors. Since going below zero causes a state to be eliminated, 25 points were added to all results to keep everyone in the black. This gave the following initial set of scores for 1885 (see Table). Scores were calculated for each year 1885-1914.

	UK	FR	GMY	AUH	ITA	SER	RUS	TUR
Resource	1089	382	467	154	258	25	210	40
Cmbt Pwr	674	424	482	265	282	25	646	254

Table 22: Initial Values for WWI Case

The second part of the conversion process was attributing MEMORIS-2 moves to the actual data. Recall from the previous chapter that each player only has four move choices:

- Move 1: build-up military
- Move 2: offer alliance (results in formal alliance if both sides choose this move)
- Move 3: threaten (results in war if both sides choose this move)
- Move 4: build economy

The process of move attribution required a combination of the resource, combat power, MID, and alliance data. The overall process was relatively straightforward.

Figure 43: Alliance and War Coding

The data was then formatted to create a status vector for each dyadic combination of states for each year, for a total of over 1,500 directed dyadic pairs of input data spanning 30 years of time. This vector will consist of the moves for each state in the dyadic pair, followed by resource change, combat power change, cumulative resources and cumulative combat power for each side, memory-degraded versions of the previous four factors, a 5-digit binary coding variable, flags for alliances, war, and the year. Move data, coding variables and status flags are all binary (0,1) data. For example, a status vector would look like:

[200,300,1,0,0,0,0,1,0,0,3,-2, .5,.9,1080,670,236,260 ,… 0,0,1890]

where the first element is the country code for the first country in the pair (in this case the UK), the second element is the country code for the other country in the pair (Austria-Hungary), player one's move (1,0,0,0 = move 1), player two's move, etc., finishing with alliance=0 (no alliance), war=0 (no war) and the year. This data was then read into a single combined file, and was then used to initialize the simulation.

Drawing from the above data, each year is composed of 56 directed dyad interactions, so each state's interaction with every other state in the system requires a separate calculation and choice. For the first 20 years of the simulation, each state's move choice and result is based on the actual Correlates of War data for that period of time. This allows the historical record to be essentially "read" into the memory of the simulation. Once the 20-year initialization is complete, then the model begins to calculate move choices based on each state's judgment of the given situation, the assessment of the likely moves of the other party, and the anticipated results. As described in the previous chapter, each state's strategy toward every other state is independent, so a state could threaten one state, offer alliance to another, build up forces against a third, and so forth. States do have knowledge of existing wars and alliances, and that knowledge is part of the decision process.

Creating the simulation baseline

Prior to starting the analysis with actual historical data, I accomplished several sets of baseline runs get a feel for how the simulation reacted to various move combinations. The initial baseline runs all started with the initial 1885 values shown above, and the first 20 moves were generated randomly from that starting point. After this initialization, the simulation was allowed to calculate its own moves for an additional 50 moves, making the total length of the simulation 70 moves, or 70 years of time. This was accomplished 1000 times for each set of runs.

The random generation of the 20-move initialization was varied according to the procedure used in the 3-player game in the previous chapter. Moves were divided between those which were generally hostile (1=build up military, 3=threaten) and those which were generally friendly (2=offer alliance, 4=build up economy), and the percentage of each move was varied between 90% hostile/10% friendly and 10% hostile/90% friendly as shown.

	Build Military	Offer Alliance	Threaten	Build Economy
Case 1	45%	5%	45%	5%
Case 2	35%	15%	35%	15%
Case 3	25%	25%	25%	25%
Case 4	15%	35%	15%	35%
Case 5	5%	45%	5%	45%

Table 23: Distribution of Initial Moves

Results of the Simulation Baseline

After completing 1000 simulation runs for each case, the resulting overall data were collected:

- Total score (final resource + final combat power)
- Resource score (final resource)
- Combat score (final combat power)
- Ally score (total alliances/turn)
- War score (total war/turn)

A summary of the statistical properties of the main results is below. As can be seen from the table, the output variables are all roughly normally distributed, positively skewed, and show varying degrees of negative kurtosis. Thus, each output variable is distributed somewhat "flatter" than the normal distribution, and the means are skewed by clusters of observations at the low end of the range.

Variable	N	Mean	Std Dev	Minimum	Maximum	Skewness	Kurtosis	K-S Stat	p-value
TOTAL	5000	4339.92	3179.33	574.38	17493.05	0.68	-0.56	0.17	<.01
CP	5000	522.94	243.54	78.87	1371.45	0.87	-0.05	0.12	<.01
RES	5000	1786.06	1542.27	25.07	6309.90	0.43	-1.04	0.19	<.01
ALLY	5000	11.43	7.54	0.72	30.40	0.13	-1.36	0.17	<.01
WAR	5000	2.08	1.97	0.00	9.08	0.72	-0.84	0.21	<.01

Table 24: Statistical Properties of Results (WWI Baseline)

Cluster Analysis of the Results

These results were then analyzed for distinct clusters, following the basic procedure outlined in the previous chapter. The data were all normalized, and then analyzed both hierarchically and using k-means cluster techniques (SAS procedures CLUSTER and FASTCLUS). As seen in the 3-player case, the results broke into either three or four distinct clusters. Again, the difference between three and four clusters is that the "middle" cluster essentially broke into two separate sub-clusters, which did not offer any new interpretive value. The two sub-clusters were essentially variations of the same trend, so three clusters are used in the remainder of the analysis for the sake of simplicity.

Again, the three clusters show virtually the same dynamic that was shown in the previous chapter. One cluster showed a high level of war, along with high combat power, and low resource and alliance levels, indicating high and continuous conflict, which I characterize as a "Conflict" cluster. The second cluster showed lower levels of war and combat power, with moderate numbers of alliances and high resources, which I characterize as a "Standoff" cluster. The final cluster shows low levels of war and combat power, but very high resource and alliance levels, which I characterize as a "Peaceful Alliance" cluster. A summary of the results is shown below.

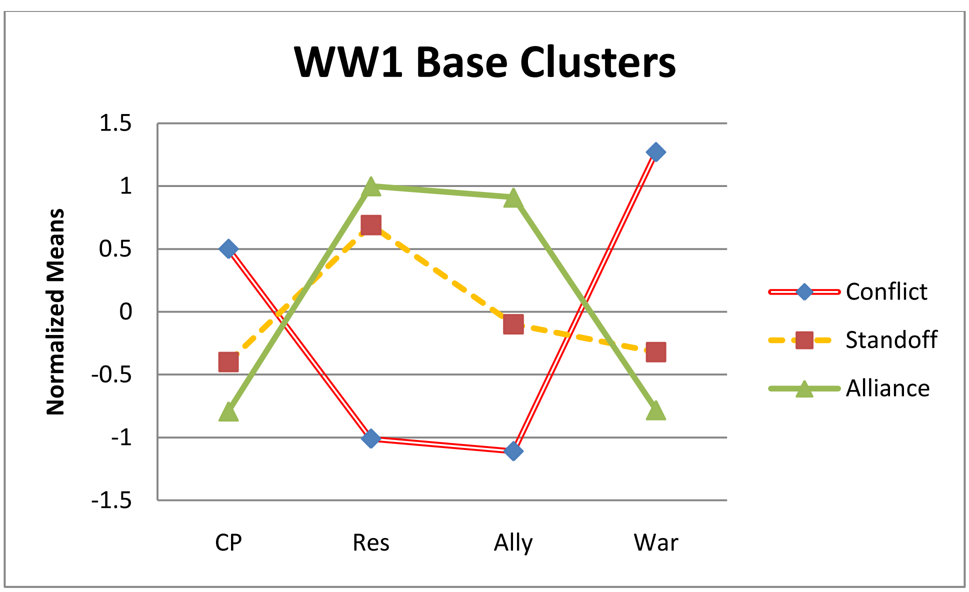

Figure 44: WW1 Base Cluster Values

	CP	Res	Ally	War	Scheffé
Conflict	0.50	-1.01	-1.11	1.27	**<.0001**
Standoff	-0.40	0.69	-0.10	-0.32	**<.0001**
Alliance	-0.79	1.00	0.91	-0.78	**<.0001**
R-squared	*0.47*	*0.80*	*0.75*	*0.81*	
F-test	*3274*	*10255*	*7641*	*10580*	
Significance	*<.0001*	*<.0001*	*<.0001*	*<.0001*	

Table 25: WW1 Base Cluster Statistics

As can be seen, the above table, the clusters are quite distinct, both at the variable level and at the between-clusters level, as shown by the highly significant F and Scheffé tests. With this division into clusters, I can now show how the distribution of results changes as the general level of threat increases.

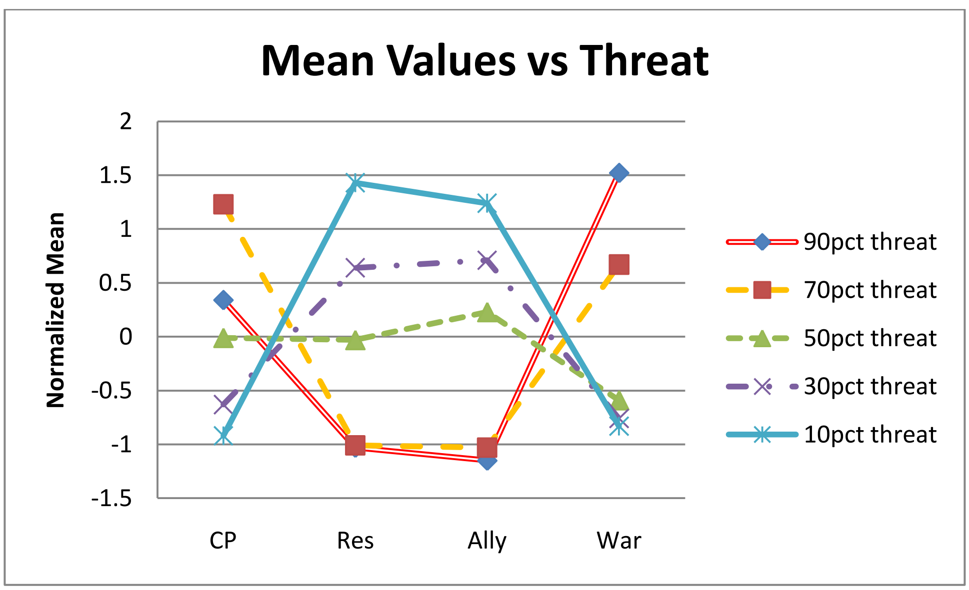

Figure 45: Cluster Values vs. Increasing Threat (WW1 case)

The above chart shows the overall means for each set of runs. Recall that the total dataset was normalized, so the above chart shows the change in the normalized mean for each set of threat-level runs. Note that the means show a nearly-perfect inversion as the threat level changes from low to high. The only exception is for the Combat Power variable, which drops between the 90% and 70% threat cases, then falling as the threat level decreases. The reason for this result is that when players in the model engage in high levels of conflict, it tends to offset the normal buildup in combat power. This is because, as in real life, using your military tends to chew it up. More moderate threat levels (such as the 70% case) cause buildups, but not as much actual conflict.

With the significant change in the mean values shown above, it is not surprising to find a significant change in the number of cases in each cluster as the threat changes. Following the shift in the distribution of threat from Table 25 above gave the following distribution of cases by cluster.

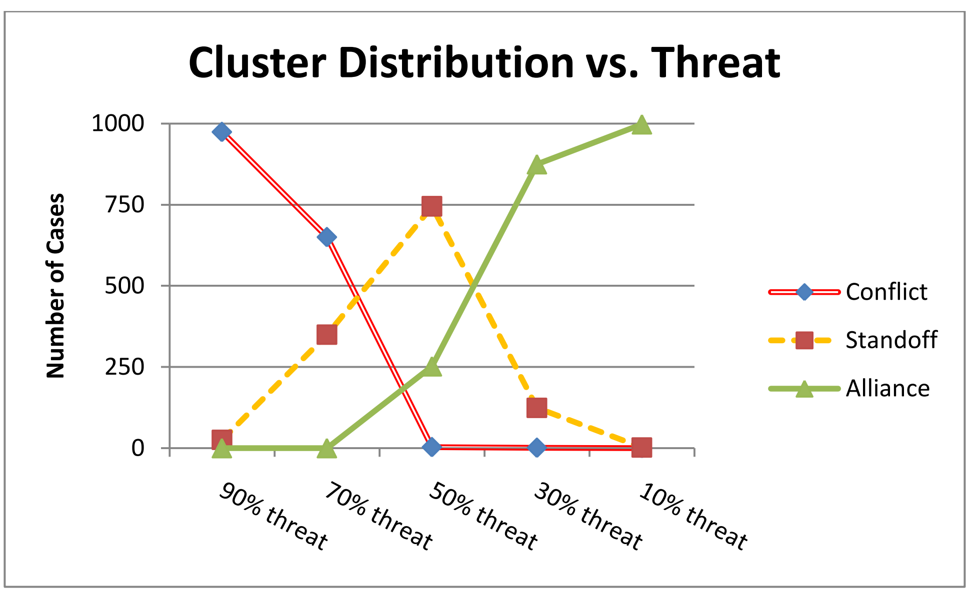

Figure 46: Cluster Distribution vs. Threat (WW1 Base Case)

	90% threat	70% threat	50% threat	30% threat	10% threat	n	Test	Value	Signif
Conflict	974	650	3	1	0	1628	Chi-sq	6287	<.0001
Standoff	26	350	745	124	2	1247	Cont. coeff.	0.7463	
Alliance	0	0	252	875	998	2125	Cramer's v	0.7929	
						5000			

Table 26: Cluster Distribution (WW1 Base Case)

As was the case in Chapter Three, there is also a strong and statistically significant relationship between the amount of threat and the distribution of cluster results. Note that as the threat level drops below 50%, the number of cases with high conflict drops to near zero, while the number of cases with a more peaceful "alliance" pattern increases. In the middle area, the "standoff" pattern dominates. These results give an expected baseline which will now be used to assess the runs which used the actual historical data.

Cluster Example Cases

To give an example of each type of cluster, I have plotted the overall scores (Resources + Combat Power) along with the incidence of conflict (shown as vertical bars). Since these cluster examples are based on random initialization, there is no attempt to tie these particular results with any historical events.

The "Alliance Cluster" is characterized by a high degree of growth and very low conflict (in this case, only two episodes of war). States essentially shift to a maximum peacetime economic growth (about 10% per year).

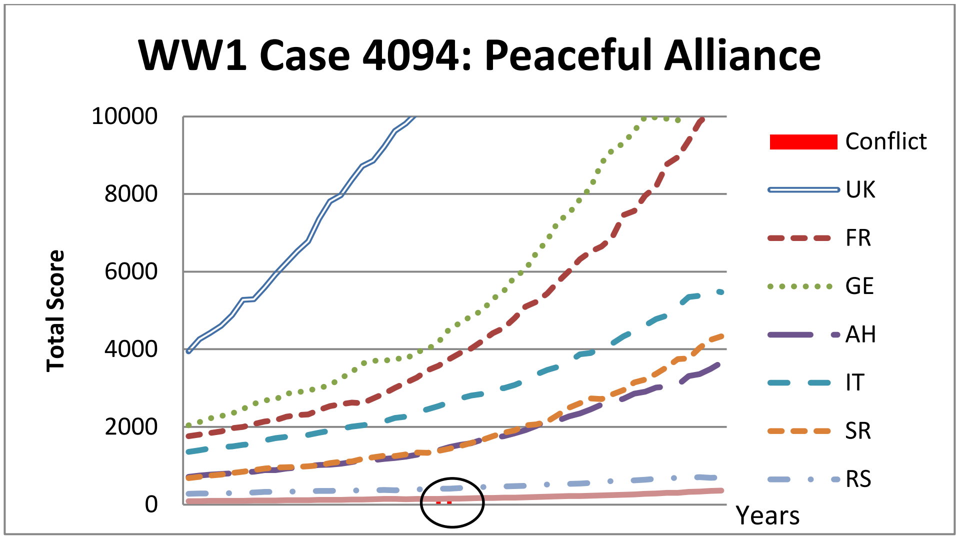

Figure 47: WW1 Alliance Example

The "Standoff" cluster is characterized by low economic growth along with intermittent conflict. Note that the height of the conflict bars indicates the number of states in conflict (2, 4, etc.) and the wider conflict bars indicate multi-year conflicts. States are required to balance their economic and military needs, and therefore the overall growth is much less than in the Alliance case.

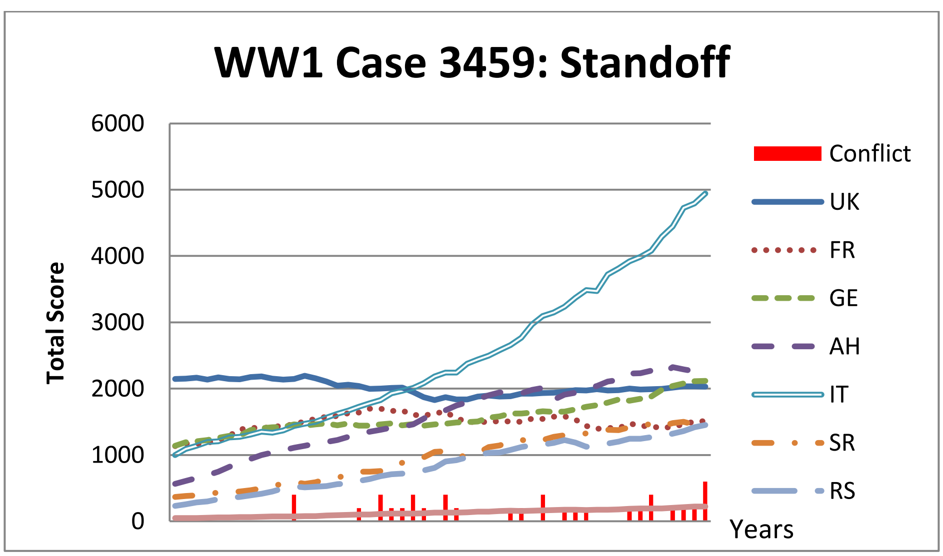

Figure 48: WW1 Standoff Example

The “Conflict” case is marked by high intensity conflict generally involving four or more states, and these conflicts are more likely to be multi-year. As might be expected, this occurs along with generally low growth. Here, most of the effort is put toward military power, and that power ends up being chewed up in constant, large-scale conflict.

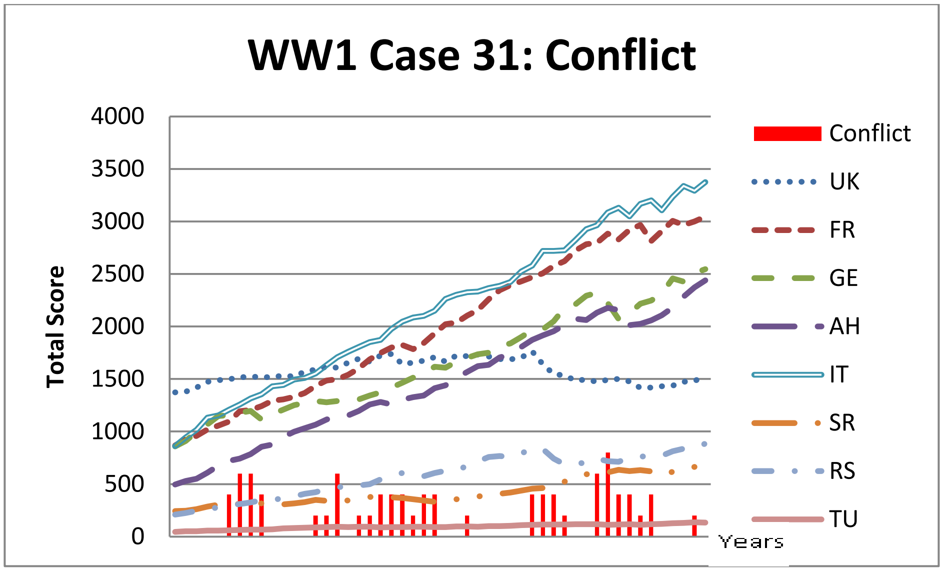

Figure 49: WW1 Conflict Example

Analysis Using Historical Data

The model runs using actual data used the same basic approach as the baseline runs – the simulation was initialized with 20 years of data, followed by 50 years of actual simulation. During initial testing of the model, I discovered that, as would be expected, the range of results resulting from using a specific set of data exhibits much smaller variation than those resulting from random data. This allowed good results with a smaller number of runs, so instead of 1000 runs per case, I used 200 runs. As will be seen, these sets of results are still highly significant, so the 5-fold decrease in run times is a more than acceptable tradeoff.

The first set of simulation runs starts with 1885 and uses actual data from 1885-1904 to initialize the simulation. Subsequently, each set of runs moved one year into the future, with the next set of runs using the 1886-1905 data for initialization, and so forth to the final set of runs using 1893-1912 to initialize. As previously, data was collected year by year for each set of simulation runs. The following table summarizes the statistical properties of the runs.

Variable	N	Mean	Std Dev	Minimum	Maximum	Skewness	Kurtosis	K-S Stat	p-value
TOTAL	1800	7476.11	7970.25	1039.11	10270.76	1.44	1.14	0.261	<.01
CP	1800	1104.38	509.70	169.70	2668.31	0.55	-0.46	0.079	<.01
RES	1800	1623.97	1756.17	125.80	9876.40	1.54	1.86	0.253	<.01
ALLY	1800	7.07	1.06	2.76	10.26	0.08	0.28	0.103	<.01
WAR	1800	0.83	0.98	0.00	6.08	1.29	1.83	0.255	<.01

Table 27: World War I Result Statistics

Cluster Results: World War I Historical Data

The cluster analysis of the nine historical data runs was performed in the same manner as previously discussed. The results are shown in the figure and table below. As expected, the clusters followed the same basic pattern as was seen in both the WW1 baseline case and in the 3-player baseline. Three clusters were found during the analysis, which corresponded fairly well to the three clusters in the baseline of Conflict, Uneasy Standoff, and Peaceful Alliance. The middle “Standoff” cluster is slightly different than the baseline case (Figure 44), but very close to the same values as seen in the 3-player game (ref Figure 40, Chapter 3).

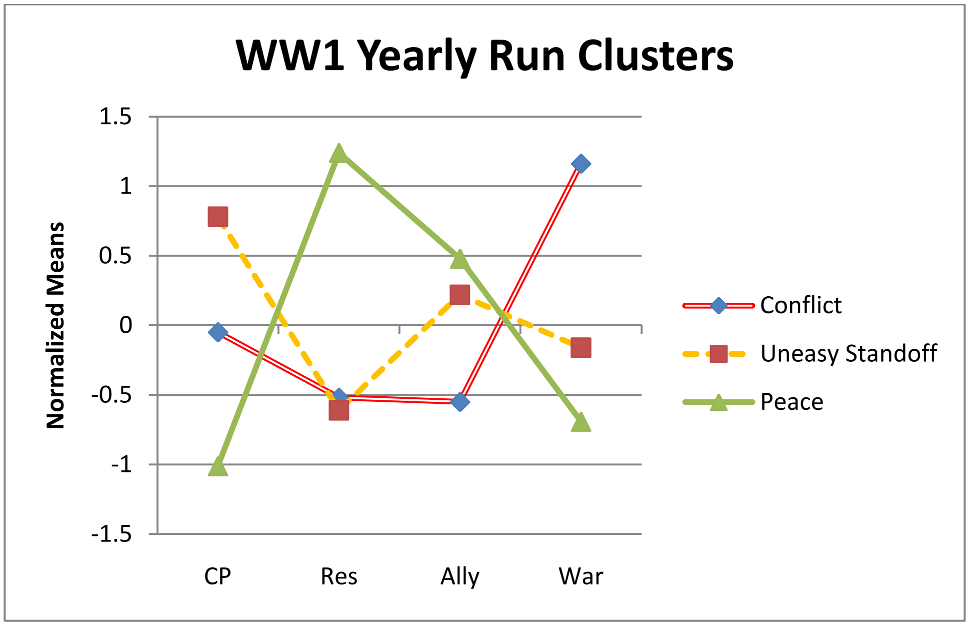

Figure 50 : Cluster Results, WW1 Yearly Analysis

	n	CP	Res	Ally	War	Scheffé
Conflict	448	-0.05	-0.52	-0.55	1.16	<.0001
Uneasy Standoff	779	0.78	-0.61	0.22	-0.16	<.0001*
Peace	573	-1.01	1.24	0.48	-0.69	<.0001*
R2		0.59	0.72	0.18	0.5	
F-test		1288	2351	208	889	
Significance		<.0001	<.0001	<.0001	<.0001	* n.s. for Ally var

Table 28: Cluster Statistics, WW1 Yearly Analysis

Notice in the table above that, while the difference between each of the variables is highly significant for all three clusters, that for the Ally variable, the Scheffé test was non-significant between the "Standoff" and the "Peace" clusters. Looking at the numbers by themselves, one would expect that the Resource variable would also be non-significant, since their means are closer together than the Ally variable. However, from Table 27, recall that the distribution for the Resource variable had high kurtosis, so the distribution of results was much more tightly clustered for this variable than we saw for the Ally variable, which explains why one value is significant and the other is not. Further, this also shows why the Resource variable mean for the "Standoff" cluster was different in this case than in the baseline case. In the baseline case, the Resource variable distribution had negative kurtosis, meaning that the distribution was more spread out than in the yearly case. This is, naturally, due to the fact that these simulation runs were initialized using a single set of historical data, which caused all result variables to be more tightly clustered together.

Nevertheless, the results do follow the initial baseline results quite closely, and the plot of the clusters by year shows the same type of relationship seen in the baseline case, as shown below.

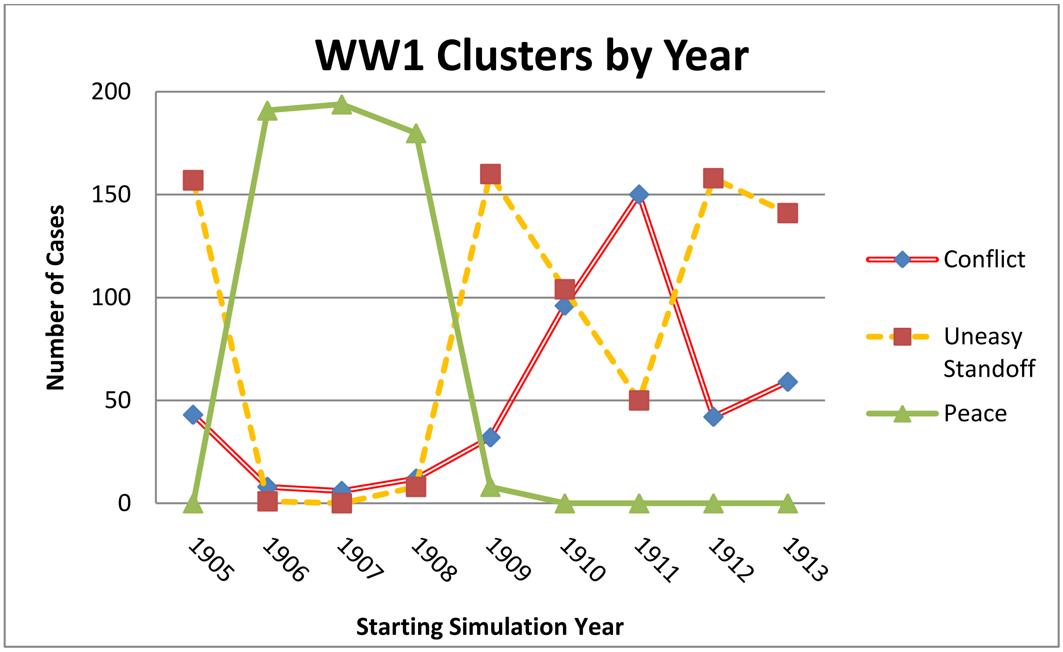

Figure 51: WW1 Clusters by Year

	1905	1906	1907	1908	1909	1910	1911	1912	1913	Total		
Conflict	43	8	6	12	32	96	150	42	59	448	chi-sq	1932
Standoff	157	1	0	8	160	104	50	158	141	779	cont coeff	0.7195
Peace	0	191	194	180	8	0	0	0	0	573	Cramer's v	0.7327
										1800		

Table 29: WW1 Cluster Statistics by Year

Here, we can see that the relationship between the number of clusters in each case is highly related to the year, and that the relationship tracks quite well with the events leading up to WW1. The starting clusters in 1905 correspond to the 1st Morocco Crisis, then there is a short period of relative calm corresponding to the post-crisis period of optimism. By 1908-1909, the threat and probability of war increase greatly through the Agadir Crisis, the Bosnian Crisis, and on into the Balkan wars. By this time, the threat of major conflict remains elevated – and when the threat of major war is significantly high year after year, it only takes one more spark to ignite the tinder, as happened in the summer of 1914.

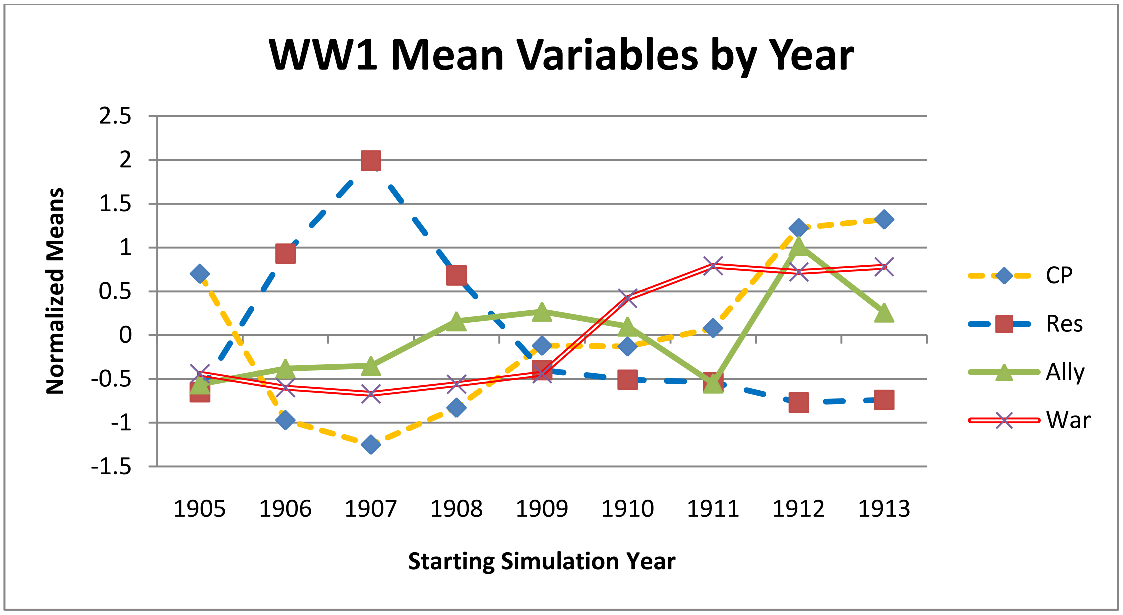

Figure 52: WW1 Variable Means by Year

As shown in the figure above, tracking the means of each of the four variables from year to year shows the same dynamic as the cluster data. In 1905, resource levels were almost one standard deviation above the mean, while alliances, combat power and war were all well below the mean, indicating a peacetime military and economic condition. By 1912, this situation was entirely reversed, indicating an unambiguous shift to a war footing. Starting in 1908-1909, the values of both the combat power and war variables rise significantly, and the incidence of war for 1911-1913 is almost one standard deviation higher than the mean, effectively doubling the probability of the incidence of major war from the pre-1908 period. While the war actually started in 1914, the simulation results indicate that it could have just as easily started in 1912 or 1913.

The same picture is also shown when looking at the percentage of simulation years where there is actual war. This is calculated by dividing the number of simulation dyad-years where war occurred into the total number of dyad years in all of the simulation runs. Each simulation run is 50 years, times 200 simulation runs, times 8 dyads, for a total of 80,000 possible dyad-years. So a probability of .3 means that 24,000 dyad-years out of 80,000 resulted in war – a very high incidence of possible conflict, particularly when repeated over several years.

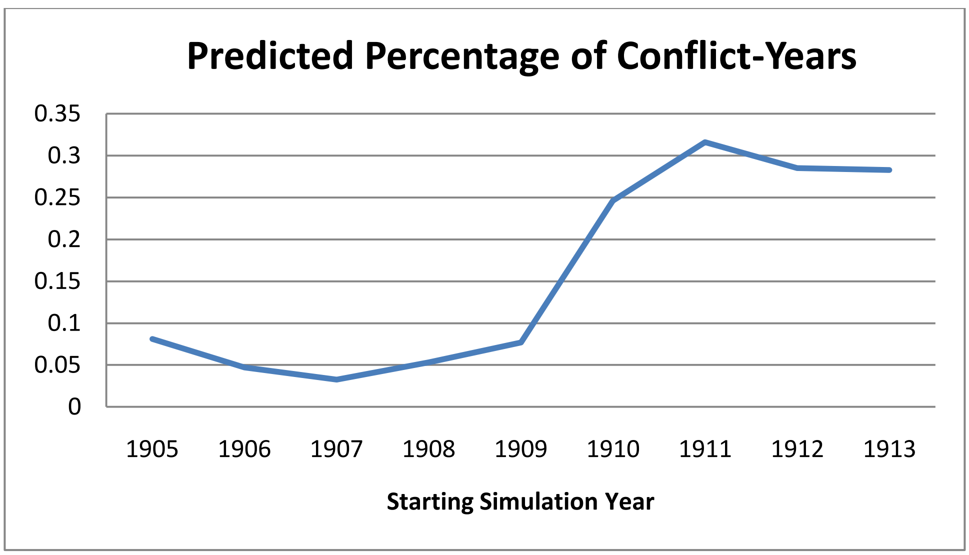

Figure 53: Percentage of Conflict-Years (WW1 Case)

Of course, showing that conflict may break out is one thing and showing *where* that conflict will break out is another. The model results show that the likely players in any conflict are the same ones seen in the actual data. The figure below shows which countries, again by percentage of dyad-years, were predicted to engage in war.

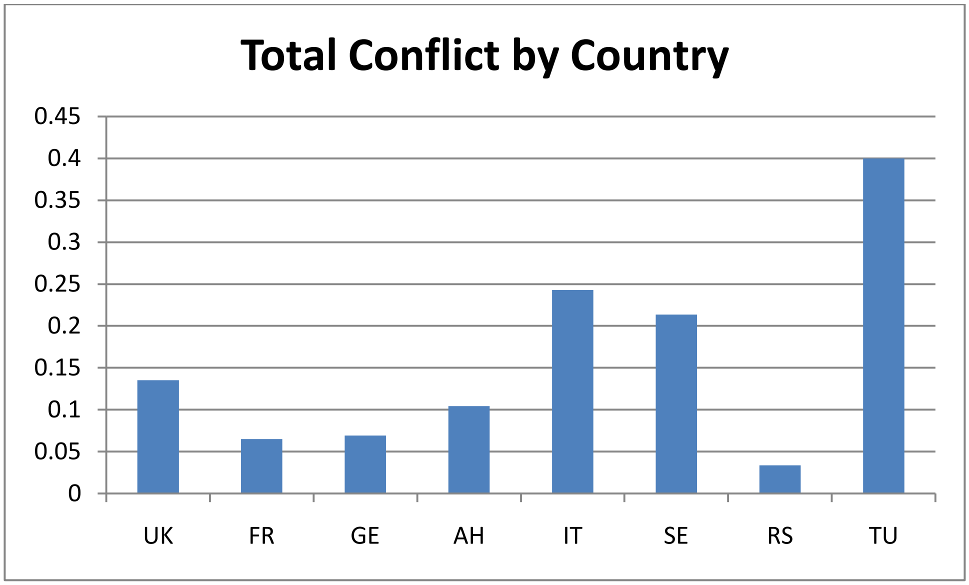

Figure 54: Total Conflict Probability by Country (WW1 Case)

These results are interesting, as they show the locus of conflict to be in Southeastern Europe, rather than from Germany. These results reflect the primary focus on actual conflict in the years leading up to the war, which was indeed primarily in the Balkans and the Italian-Turkish war. There are, of course, many scholars who place the center point of the war's initiation exactly in that spot (see, among others, Williamson, 1988; Joll & Martell, 2007).

The following two figures show the probabilities of each state going to war against each other state, again measured in percentage of possible dyad-years.

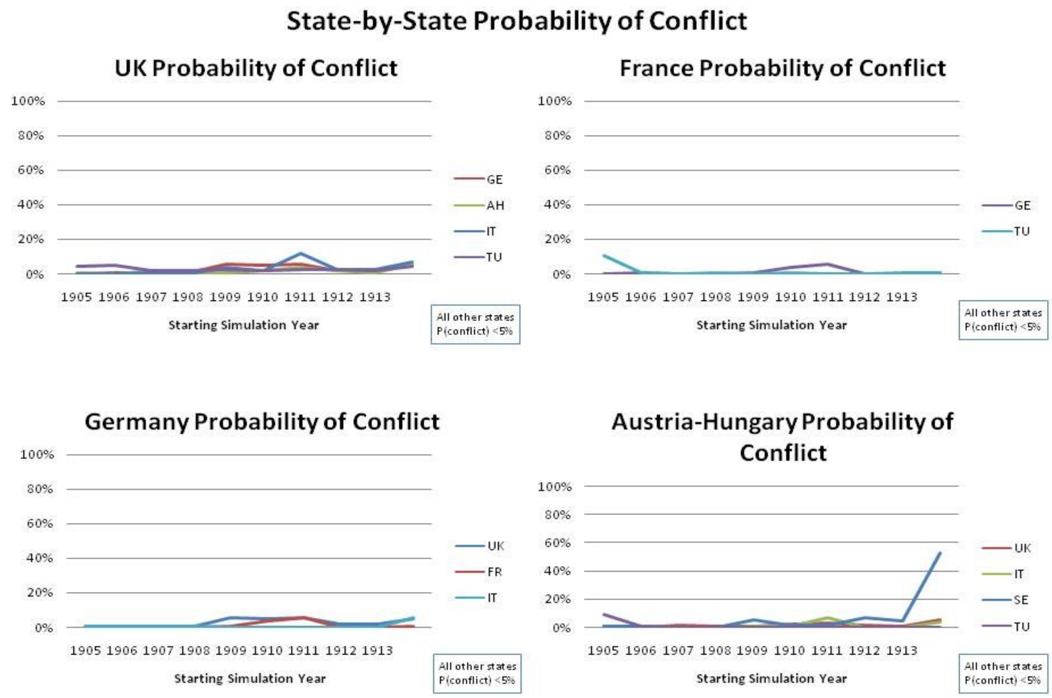

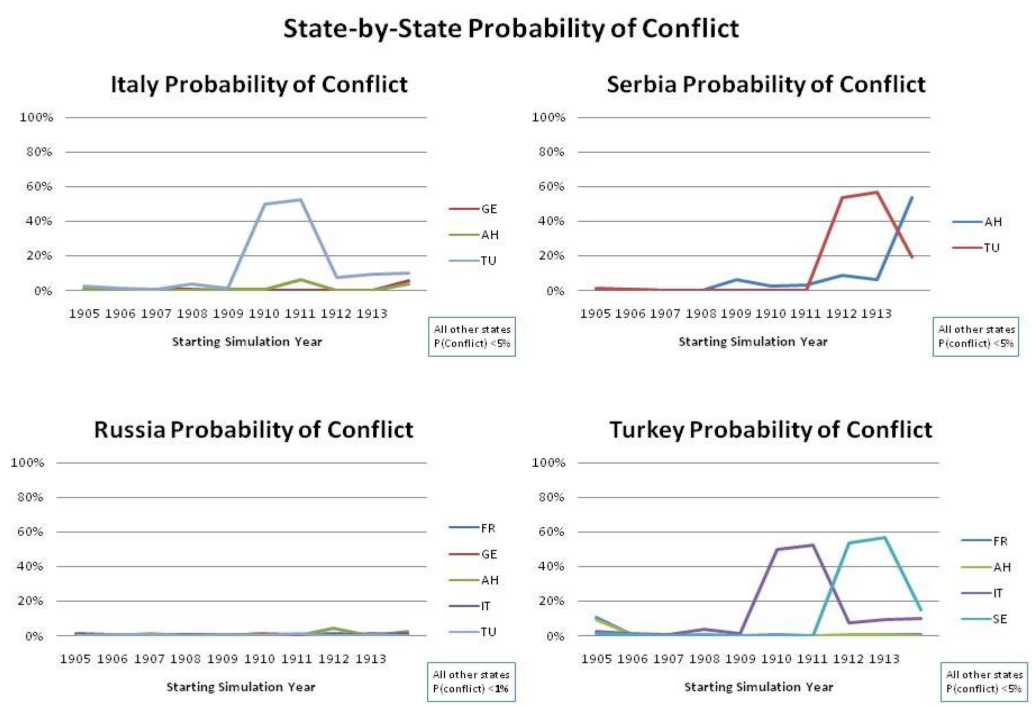

Figure 55: State-State Probability of Conflict (WW1 Case)

The state-by-state results show that, according to the simulation, the locus of conflict is clearly in Southeastern Europe, with the vast majority of the conflict coming between Italy, Serbia, Turkey and Austria-Hungary. This begs the question of why war between Germany and France or the UK rarely happened in the simulation. One of the factors that causes the simulation to drive towards conflict is existing conflict. In the historical data, there is little direct conflict between these states, other than military plans and posturing. Years in which states only engaged in military planning activity were coded as peaceful, while only active threats in the MID data were coded as threats in the model. Thus, the model does not pick up on these latent conflicts, even when conflicts start with states which are allied. Here, as often occurs in real life, the state will not choose war simply because an allied state goes to war. Instead, states make their own judgments about the efficacy of war.

Like the previous conflict charts, the following alliance charts show the percentage of dyad-years where formal military alliance was predicted by the simulation. Even though, as stated in Chapter 3, the *offer* of alliance might incorporate the possibility of actions such as military staff talks or even an *entente,* the data in the following figure shows *only* the probability of formal military alliance.

The patterns of alliance in the simulation are very similar to those found in the historical record. The state-by-state alliance results clearly show the same types of dynamics that occurred in the actual case. For example, the UK largely stays clear of alliance until 1912. Italy backs out of the Triple Alliance and throws in with France. Germany and Austria-Hungary remain staunch allies, and Russia and France increase their probability of alliance throughout.

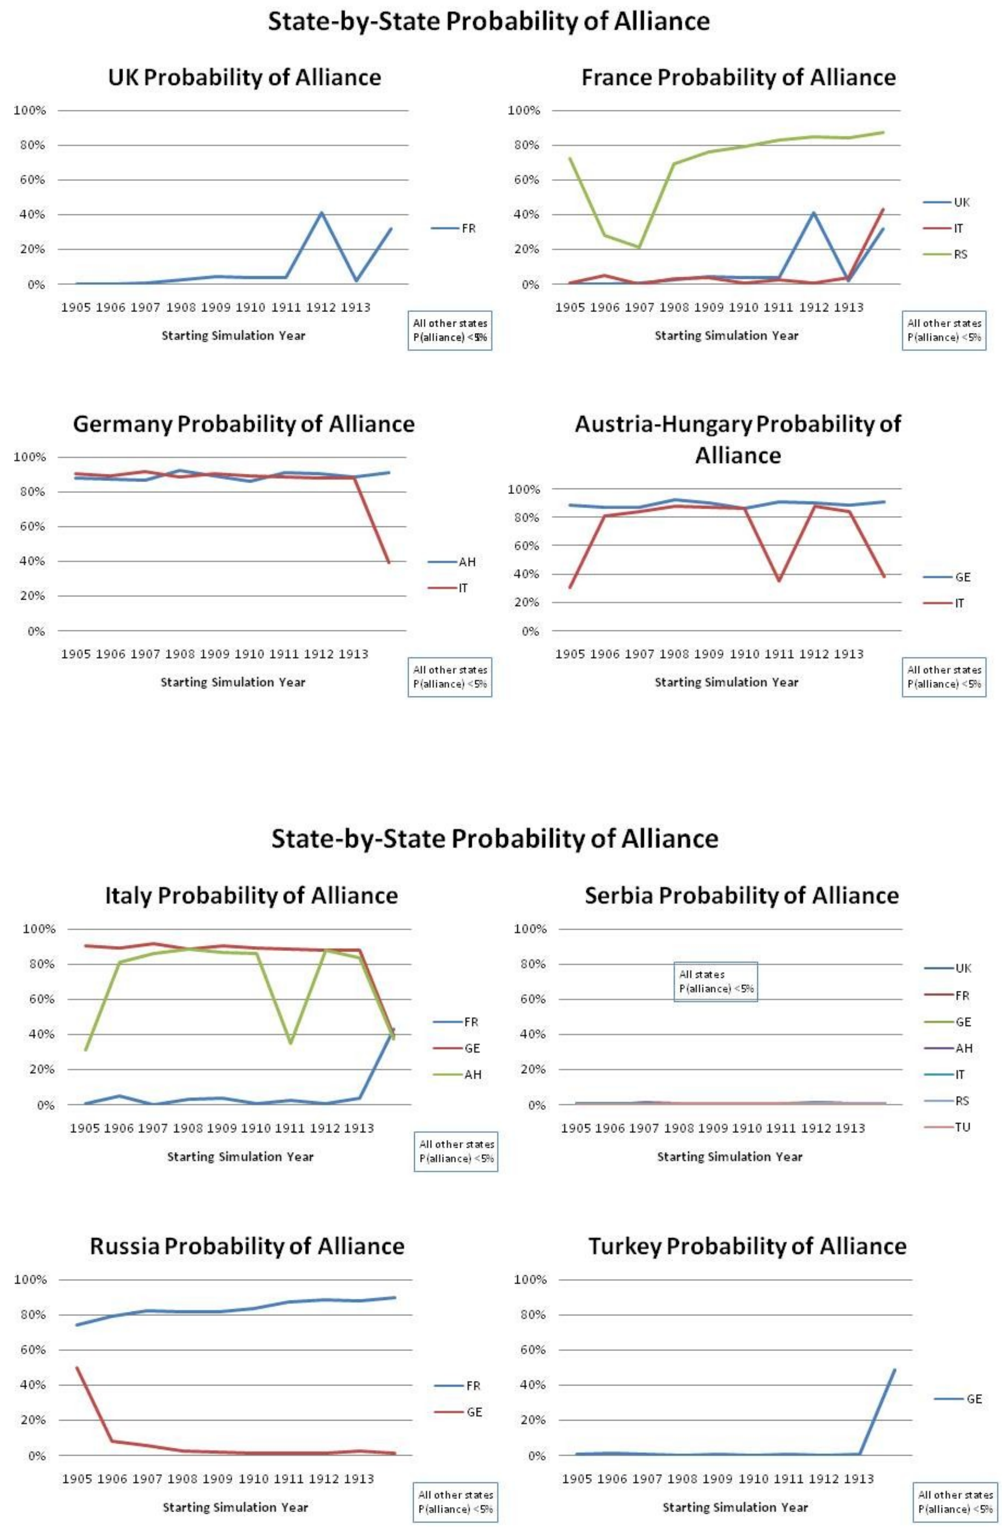

Figure 56: Probabilities of Alliance (WW1 Case)

Conclusions and Observations: World War 1 Case

The MEMORIS-2 simulation performed quite well in predicting the general tenor of the situation along with the patterns of conflict and alliance at the outbreak of World War I. The ebb and flow of tension, levels of armaments, and threat of war generally matches up with the historical record.

Additionally, comparing the yearly-results to the baseline level of threat shows a sudden and dramatic shift in the threat level around 1910. The baseline threat vs. clusters figure below has been annotated with the corresponding cluster results for the years prior to and after 1910. This shows that the corresponding perceived level of threat within the system underwent a dramatic shift around that time, with no "middle ground" between the perceptions of low levels of threat and high levels of threat.

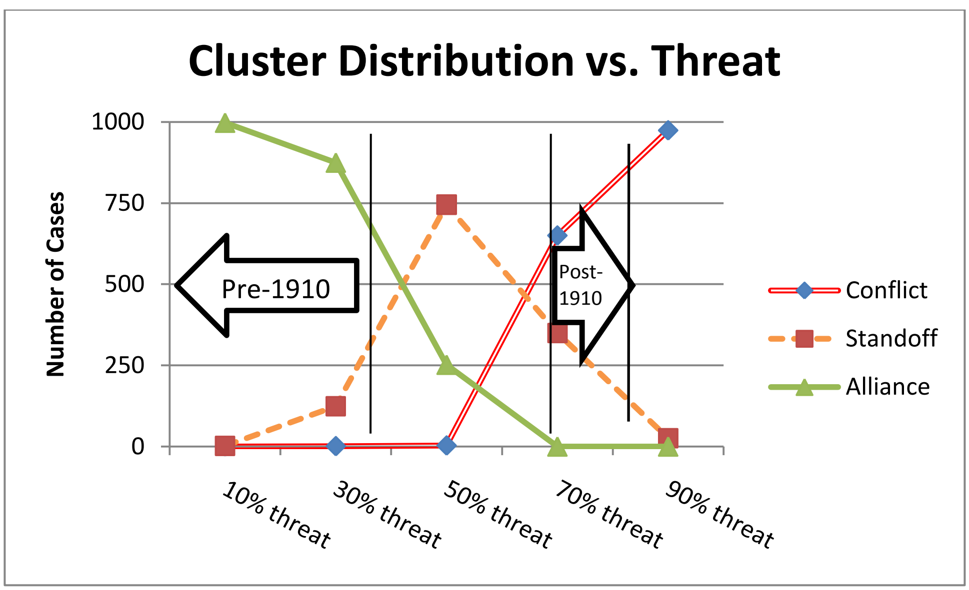

Figure 57: Dramatic Shift in pre-WW1 Threat Level

While the simulation generally did well, the patterns of conflict did not quite match the historical experience. The model accurately picked the players for the outbreak of the war, but did not capture the spread of war from Eastern to Western Europe. This occurred for two reasons. The first was discussed above, that the alliance structure in the model was not as tightly coupled as was the historical case. The UK, France, and Germany, in particular, could choose to "opt out" of the conflict in Eastern Europe, and mostly took that option rather than opting for war. The second reason has to do with the coarseness of the alliance and conflict data. The simulation data takes an all-or-nothing approach. While wars can have differing costs and uncertain length, a state is either at war or not. There is no middle ground of imminent threat between peace and war in the way in which the data were coded. The alliance data shares the same limitation of being all or nothing.

This model, naturally, shares the characteristics of all models: it is a greatly simplified version of reality, and as such any exact prediction is generally exactly wrong. Further, this brief example can in no way be considered a definitive look at WWI – if for no other reason that I made no attempt to perform sensitivity analysis on the results. The model does show fairly convincingly the general outlines of the situation, along with the probabilities of war and formal alliance among the players in the simulation. That said, I think that it is fair to say that the model works as expected in the case of World War I. With that result, I now turn to the Cold War Case.

Chapter Five: The Great War that Wasn't: The Cold War

Introduction

The period of the Cold War was remarkable for what did not happen. Despite sharing many of the facets of the previous case, including great power rivalry, arms races, militarized disputes in far-flung corners of the world, and a lengthy period of armed standoff, the Cold War never did break into the apocalyptic conflict that so many feared. Even though the sides came terrifyingly close to war on several occasions, the match was never lit.

One of the most commonly-cited reasons for the failure to come to blows was the presence of nuclear weapons. When I was at US Strategic Command, the standard briefing given to visitors included a slide showing war deaths as a percentage of population dropping significantly since the end of World War II. To make sure the visitors got the point, the decrease was labeled "Impact of Nuclear Deterrence". I used that slide quite a number of times in my own briefings to various audiences around the country.

Of course, no one knows for sure whether or not that contention is true. Nuclear war did not occur, but why? Waiting for the proverbial "black swan" to appear probably is not a good answer in this particular case. The following analysis will help gain some insight into the probability of war during the Cold War, and this will give some further data points to use to answer the question.

The remainder of this chapter will follow the general procedure outlined in Chapter Four.

Data for the Simulation

Yearly data was drawn from the Correlates of War databases for the period 1947-1967, and was processed into MEMORIS-2 input format following the procedure for the World War I data. The eight states were chosen based on their place in the National Military Capability rankings for that period. The eight states were: US, Soviet Union, China, UK, France, West Germany, Japan, and India. As in the World War I case, this leaves out some important players, but it does capture the major independent players in the Cold War alliance system, and those players have significant interactions with the other players. For example, the next state on the list would have been Brazil, which really didn't have much of a direct role in the Cold War at all.

Creating the Simulation Baseline

As in the previous two chapters, a simulation baseline was created to generate initial clusters of data and to validate that the baseline data matched expectations. The baseline runs started in 1947, with initial data for each state player as determined above. Again, the simulation was initialized using 20 years of random data, which varied in level of threat according to the table below.

	Build Military	Offer Alliance	Threaten	Build Economy
Case 1	45%	5%	45%	5%
Case 2	35%	15%	35%	15%
Case 3	25%	25%	25%	25%
Case 4	15%	35%	15%	35%
Case 5	5%	45%	5%	45%

Table 30: Distribution of Initial Moves

After the 20 years of initialization, the simulation was allowed to run for 50 years. Each set of simulation runs was replicated 1000 times.

Once the simulation results were completed, data was collected as described in the previous two chapters. The statistical summary of the data is shown in the table below.

Variable	N	Mean	Std Dev	Minimum	Maximum	Skewness	Kurtosis	K-S Stat	p-value
TOTAL	5000	1432.71	286.57	547.70	2444.22	-0.169	-0.137	0.036	<.01
CP	5000	972.07	178.79	322.61	1485.29	-0.302	-0.126	0.031	<.01
RES	5000	87.99	79.39	14.25	464.15	1.446	1.041	0.207	<.01
ALLY	5000	9.58	6.80	1.16	26.52	0.847	-0.632	0.158	<.01
WAR	5000	2.55	1.63	0.08	7.20	0.642	-0.790	0.124	<.01

Table 31: Cold War Baseline Statistics

As seen in previous results, the distributions for Resources, Ally and War are all positively skewed, which shows there are a number of values clustered at the low end of the range. Additionally, the Resource value has high kurtosis, so that values of this variable are more concentrated toward the mean compared to the remaining variables, which are somewhat more spread out than the normal distribution.

Cluster Analysis of the Cold War Baseline

The results of the simulation runs were normalized and then analyzed for clusters using the procedure in the previous chapter. The data broke into distinct clusters as seen in the previous chapters. As before, the clusters show similar dynamics to those already seen in this dissertation. One cluster (Conflict) shows high levels of conflict, along with low levels of alliance, resources and combat power. Another cluster (Peaceful Alliance) shows low levels of war, high levels of alliance and resources, and mean levels of combat power. The third cluster (Standoff) takes the middle ground. Note that the levels of combat power for the War cluster are quite low, which again shows the impact of high levels of war on combat power.

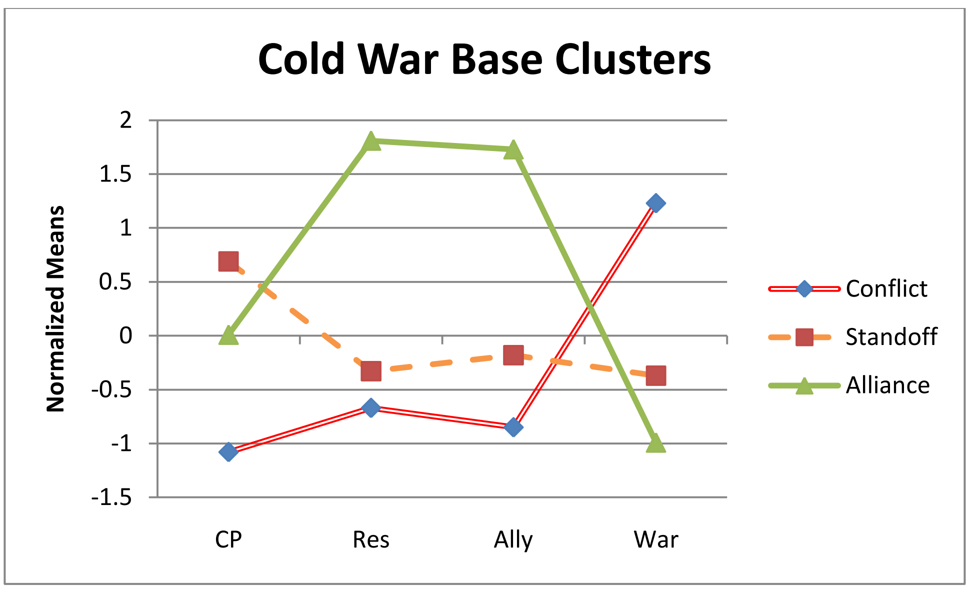

Figure 58: Cold War Baseline Clusters

	CP	Res	Ally	War	Scheffé
Conflict	-1.08	-0.67	-0.85	1.23	<.0001
Standoff	0.69	-0.33	-0.18	-0.37	<.0001
Alliance	0.01	1.81	1.73	-0.99	<.0001
R^2	0.5939	0.8597	0.8472	0.7364	
F-test	3654	1513	13855	6981	
Significance	<.0001	<.0001	<.0001	<.0001	

Figure 59: Cold War Baseline Cluster Statistics

As before, the distribution of clusters is quite distinct, with all variable differences highly significant across clusters, and the Scheffé test shows that the differences between clusters are also highly significant. As in the World War I case, the distribution of clusters also changes as expected with increasing levels of threat, as shown in the figure and table below.

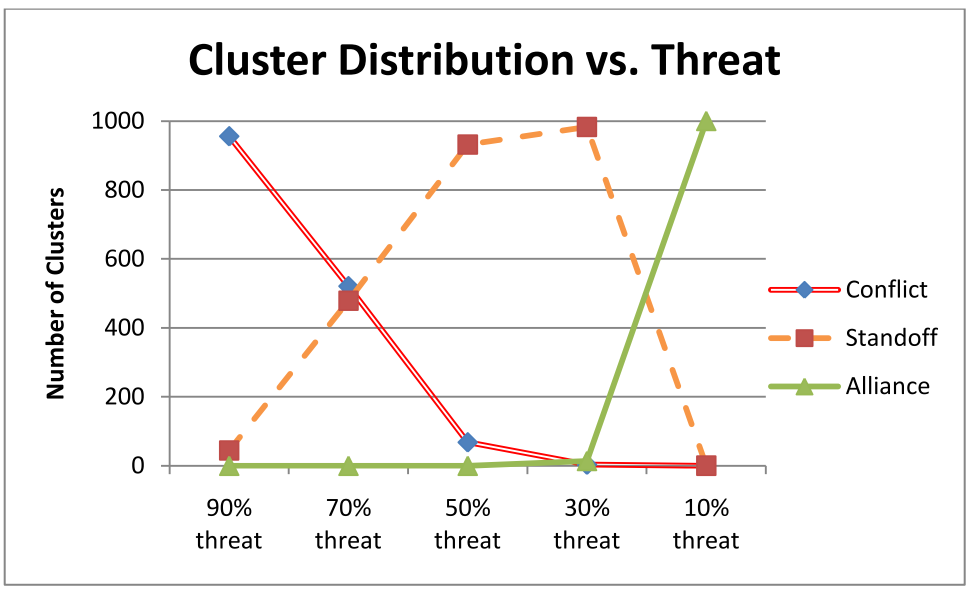

Figure 60: Cluster Distribution vs. Threat (CW Base Case)

	90% threat	70% threat	50% threat	30% threat	10% threat	Total	Test	Value	Signif.
Conflict	956	521	68	3	0	1548	chi-sq	8013	<.0001
Standoff	44	479	932	983	0	2438	contingency coeff	0.7847	
Alliance	0	0	0	14	1000	1014	Cramer's v	0.8952	
						5000			

Table 32: Cluster Distribution Statistics

The distribution of clusters definitely shows the same general trend as in the previous chapters, with a strongly significant shift from War to Standoff to Peaceful Alliance as the threat level decreases. Similarly, looking at the figure below shows how the mean simulation results change with the threat. As seen in previous chapters, there is a near-perfect inversion of the values as the threat level changes from low to high.

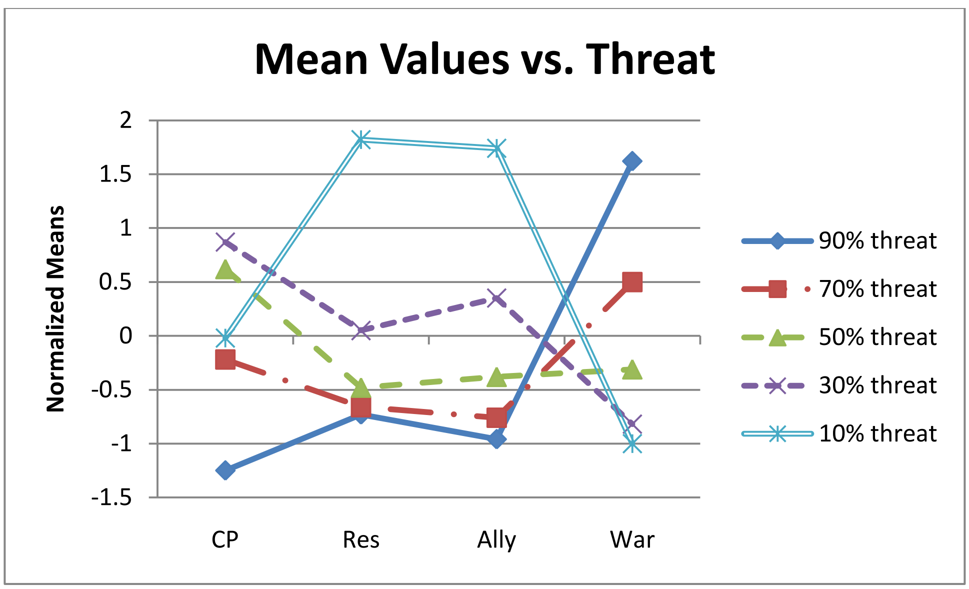

Figure 61: Mean Values vs. Threat (Cold War Base Case)

Cluster Example Cases

The examples for each cluster are quite similar to the examples shown in the previous chapter. The Alliance case example shows the impact of having more than just 1-2 conflicts. Even though conflict is very scattered, it is enough to keep the players from shifting entirely to peacetime rates of growth, as shown below.

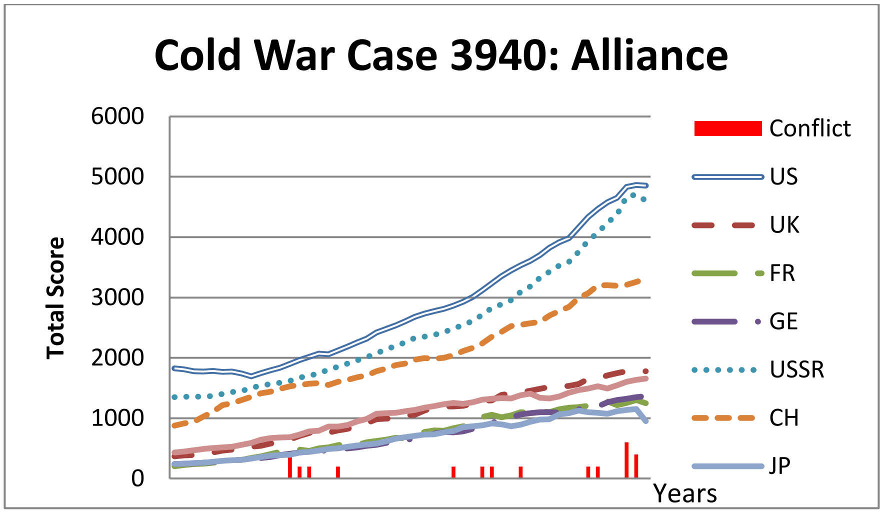

Figure 62: Cold War Alliance Example

The “Standoff” example is virtually the same as seen in the World War I case: intermittent but generally low-level conflict, along with a shift away from economic growth to a balance between combat and economic power. Note that this particular case includes one major conflict, with five players involved.

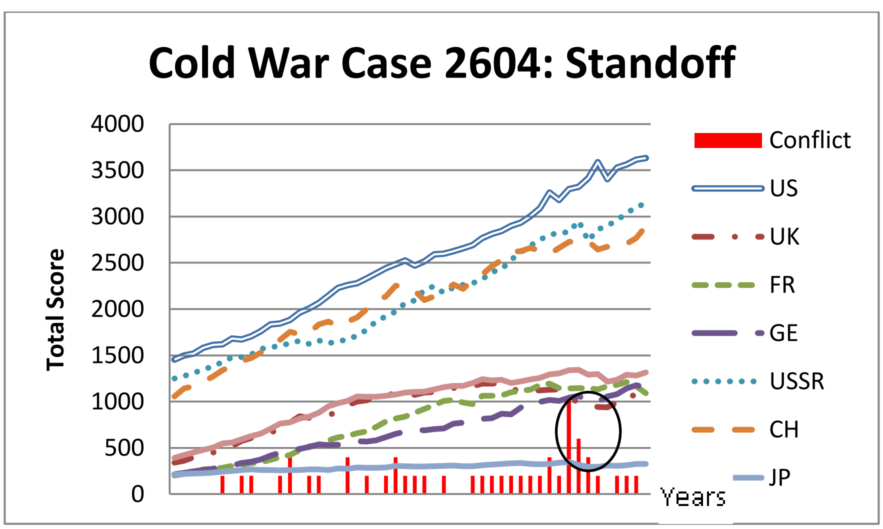

Figure 63: Cold War Standoff Example

The final example case shows the impact of constant conflict with an almost total shift of emphasis toward combat power. The growth of the main players in the conflict is virtually flat, and this gives states who stay out of the conflict (such as India in this example) to surpass the other states.

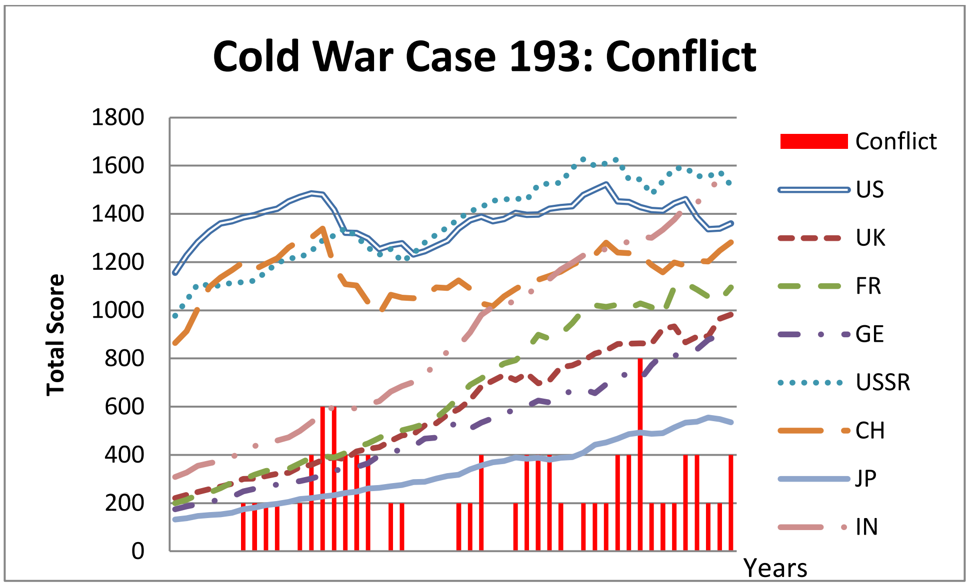

Figure 64: Cold War "Conflict" Example

Analysis Using Historical Data

With the baseline complete, the simulation was run using historical data to initialize each run. As before, 20 years of historical data started each simulation run, and then the simulation continued for another 50 years. At the end of each set of runs, data was collected and analyzed using the procedures from Chapters 3 and 4. The only difference with this set of cases is that a larger number of historical data runs were accomplished (20 vs. 9). In both cases, the number of possible runs was bounded by historical events at either end. The World War case initialization runs had to fit into the years between the unification of Germany and the start of the war.

Similarly, the Cold War runs had to be between the end of World War II and the end of the Cold War. Since 20 years of data were used for each simulation run, that also bounded the possible set of data for the runs, and thus the Cold War simulation runs begin in 1968 (1947 plus 20 years initialization). After the discussion of the initial set of Cold War runs, an alternative method of initialization will be used to capture the initial part of the Cold War, but since the procedure is different, it will be discussed separately.

Statistics for the resulting data are in the table below. The combat power (CP) and Ally variables have negative kurtosis, meaning they are more spread out than the normal distribution, which will affect the performance of the clusters. Note also the high positive skew and kurtosis for the War variable, which indicates that most values are concentrated at the low end of the range (few wars).

Variable	N	Mean	Std Dev	Minimum	Maximum	Skewness	Kurtosis	K-S Stat	p-value
TOTAL	4000	9088.26	6061.22	1256.03	67843.90	1.375	2.013	0.107	<.01
CP	4000	869.40	393.82	137.92	2064.96	0.105	-0.824	0.059	<.01
RES	4000	1904.84	1322.87	89.11	7186.19	0.783	0.074	0.089	<.01
ALLY	4000	11.64	2.24	6.28	16.42	-0.504	-1.117	0.204	<.01
WAR	4000	0.88	0.83	0.08	4.98	1.616	2.399	0.231	<.01

Table 33: Statistics for Yearly Runs (Cold War Case)

Cluster results: Cold War Historical Data

Cluster analysis of the output data was performed using the same procedures in previous chapters. As before, the data broke into three distinct clusters, generally corresponding to the previously described "Conflict", "Standoff" and "Alliance" clusters. The major difference is that the Alliance cluster has higher normalized levels of conflict than seen previously. As pointed out above, the level of conflict was tightly clustered at the low end of the spectrum, so even though it was higher than the Standoff cluster, the Alliance cluster still has a very low level of conflict.

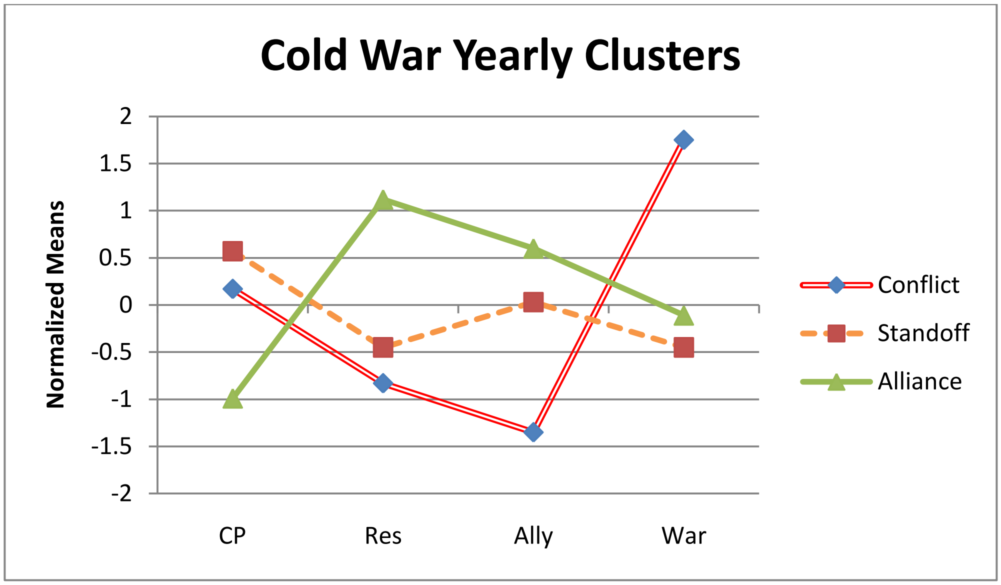

Figure 65: Yearly Clusters (Cold War Case)

	CP	Res	Ally	War	Scheffé
Conflict	0.17	-0.83	-1.35	1.75	<.0001*
Standoff	0.57	-0.45	0.03	-0.45	<.0001
Alliance	-0.99	1.12	0.6	-0.11	<.0001**
R^2	0.4923	0.6242	0.3965	0.5839	*n.s. for CP
F-test	1938	3320	1313	2805	**n.s. for Ally
Significance	**<.0001**	**<.0001**	**<.0001**	**<.0001**	

Table 34: Cluster Statistics (Cold War Years Case)

As expected from the statistical values, even though the results for each variable are highly significant across the clusters, the Scheffé test shows that the difference between the central "Standoff" cluster was not significant in terms of the combat power (CP) variable between the "Conflict" cluster. Similarly, the difference between the “Standoff” and “Alliance” clusters was not significant on the Ally variable. All other variable differences were significant at the .0001 level between clusters, so while not perfect, these clusters are sufficiently well-defined and distinct for analysis.

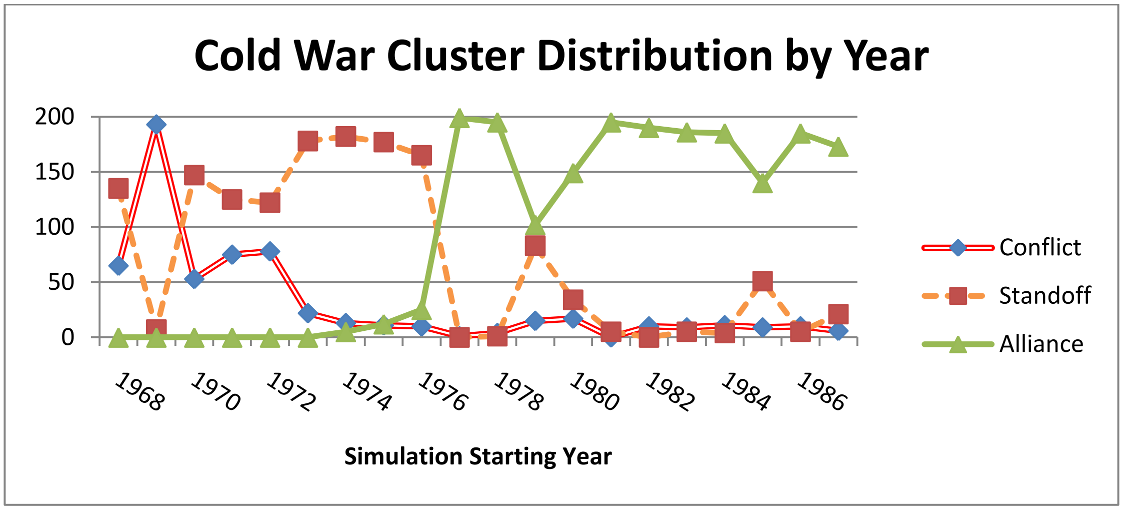

Figure 66: Cold War Clusters by Year

	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	Total
Conflict	65	193	53	75	78	22	13	11	10	1	4	15	17	0	10	9	11	9	10	6	612
Standoff	135	7	147	125	122	178	182	177	165	0	1	83	34	5	0	5	4	51	5	21	1447
Alliance	0	0	0	0	0	0	5	12	25	199	195	102	149	195	190	186	185	140	185	173	1941
																					4000

Test	Value	Signif.
chi-sq	4337	<.0001
contingency coeff	0.7212	
Cramer's v	0.7363	

Table 35: Cluster Distribution by Year (Cold War Case)

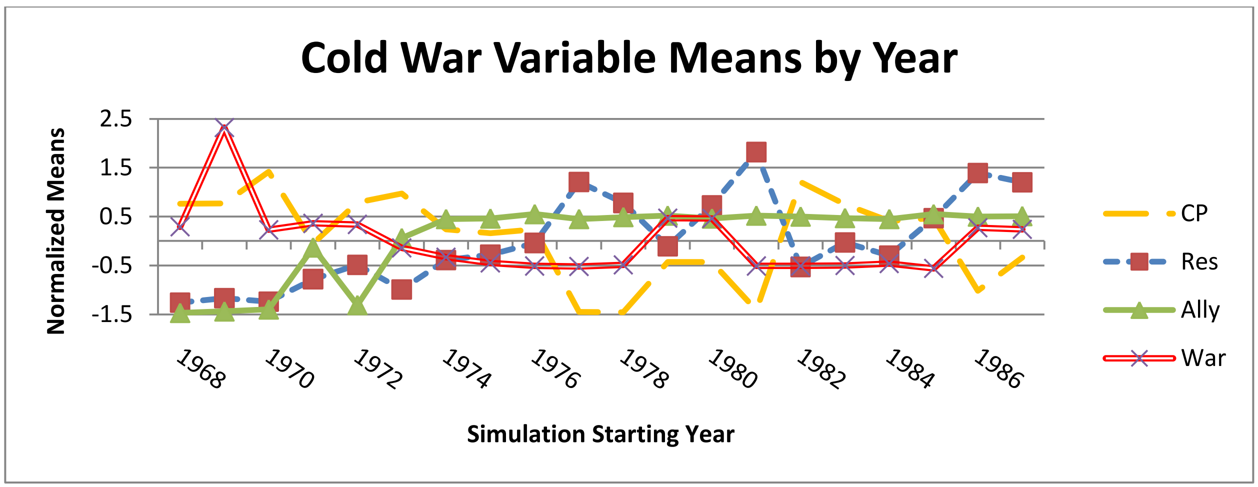

Figure 67: Variable Means by Year (Cold War)

The distribution of clusters in the Cold War case is virtually reversed from that of the World War I case. Here, the first years of the simulation start out with a high degree of threat in the late 60s, but switches to a medium threat level during the 70s, then drops to a low threat level during the 80s. The small peaks in conflict during late 70s and 1986-1987 correspond to Soviet invasion of Afghanistan and subsequent Chinese abrogation of the Sino-Soviet Treaty of Peace and Friendship and China-India border conflicts.

The same general lessening of tension can be seen when looking at the variables themselves in Figure 67 above. By the early 1970s, the general level of conflict is well below the mean, with only intermittent peaks above the mean, but still well below the situation prior to World War I. Further, the alliance structure remains remarkably constant once it is formed. This shows the overall stability of the Cold War alliance system among the major players. Notice that the overall level of resources and combat power also changes at significant historical points, corresponding to the Reagan-era buildup and the subsequent "Peace Dividend".

These same trends are also reflected in the percentages of years of conflict. As in the previous chapter, this was calculated by dividing the number of dyad-years with conflict into the number of total dyad-years. The early part of the chart shows a relatively high (20%) percentage of conflict and then drops to a very low level, with the notable exceptions discussed previously.

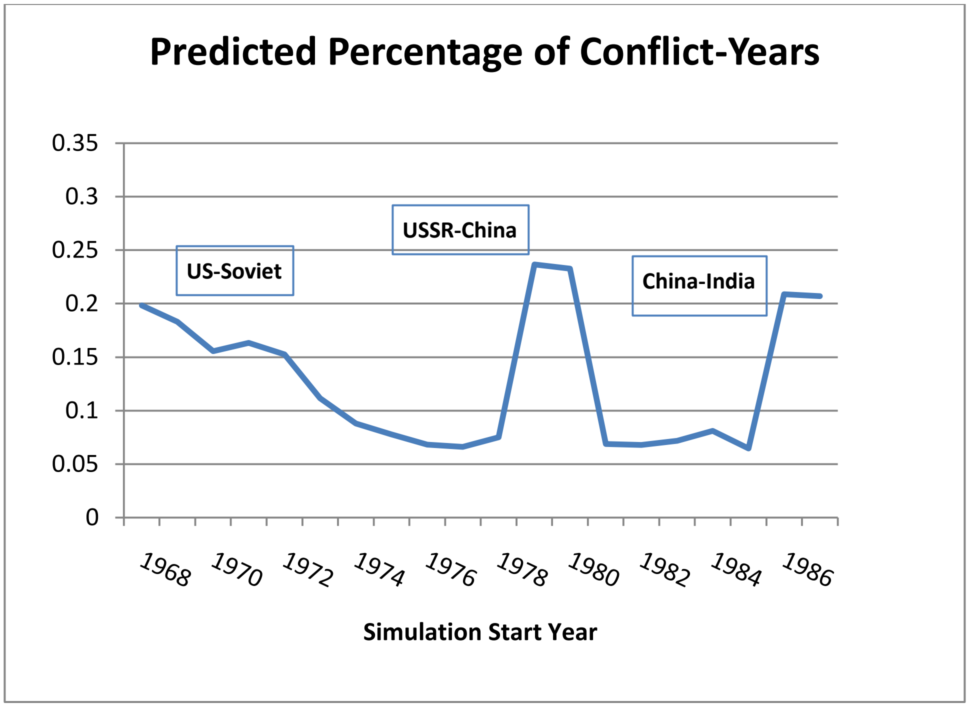

Figure 68: Percentages of Conflict (Cold War)

This is a striking result, which gives the counter-intuitive impression (at least from the American point of view) that while the world tended to focus strictly on the superpower standoff, the real threat of war may well have come under the American radar.

This might also be evidence of the deterrent effect of nuclear weapons. Recall that the model payoff structure did not change between this case and the World War I case, so there is no special way in which nuclear weapons affects the calculations for or against a particular action in the simulation. Thus, the conflict results could be taken as what *would* have been more likely to happen in the absence of nuclear weapons.

The conflict picture by country also shows a somewhat different picture than expected. Here, China and USSR have the highest overall probabilities of war, followed by the India and the US. This reflects the proximity of China to two constant foes in border skirmishes: the USSR and India.

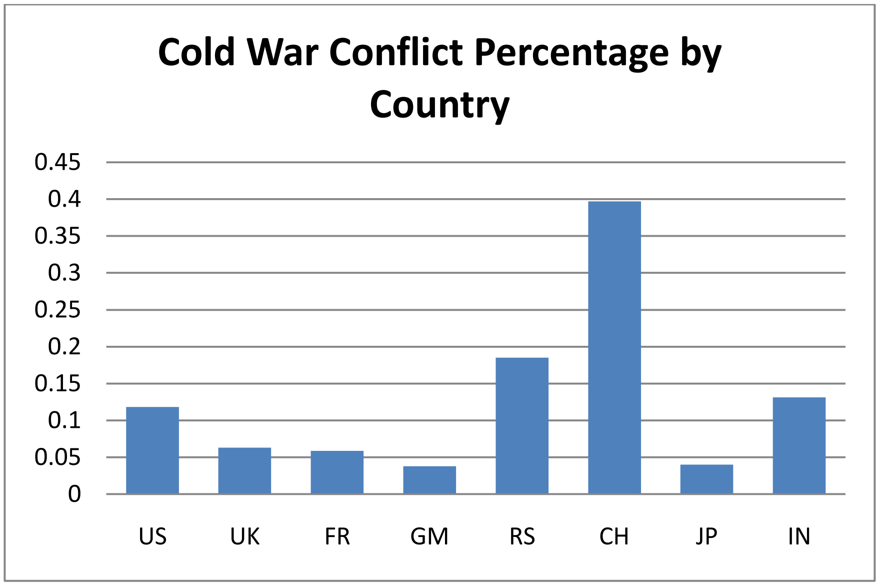

Figure 69: Conflict Percentage by Country

The state-by-state conflict analysis shows essentially the same dynamics as discussed above, but gives more detail on which states were predicted to be at war at a given time.

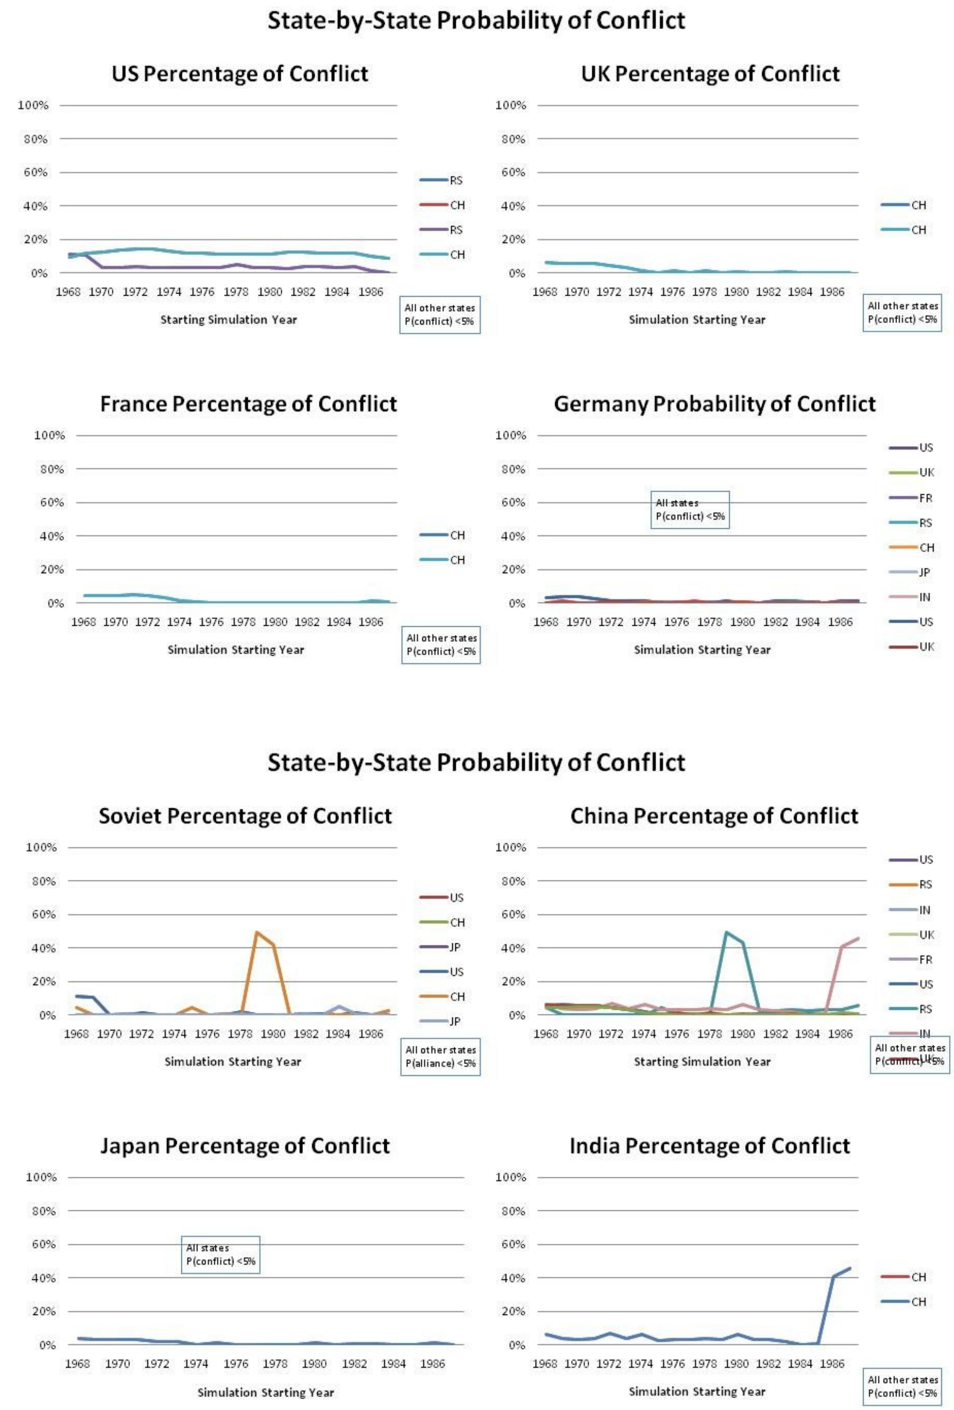

Figure 70: State-by-State Probability of Conflict (Cold War)

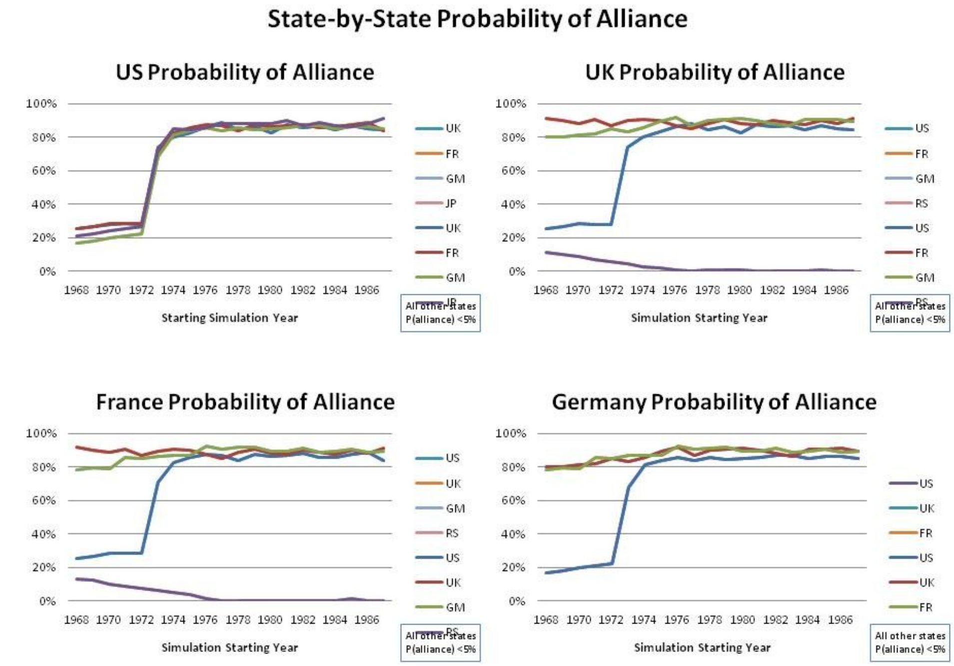

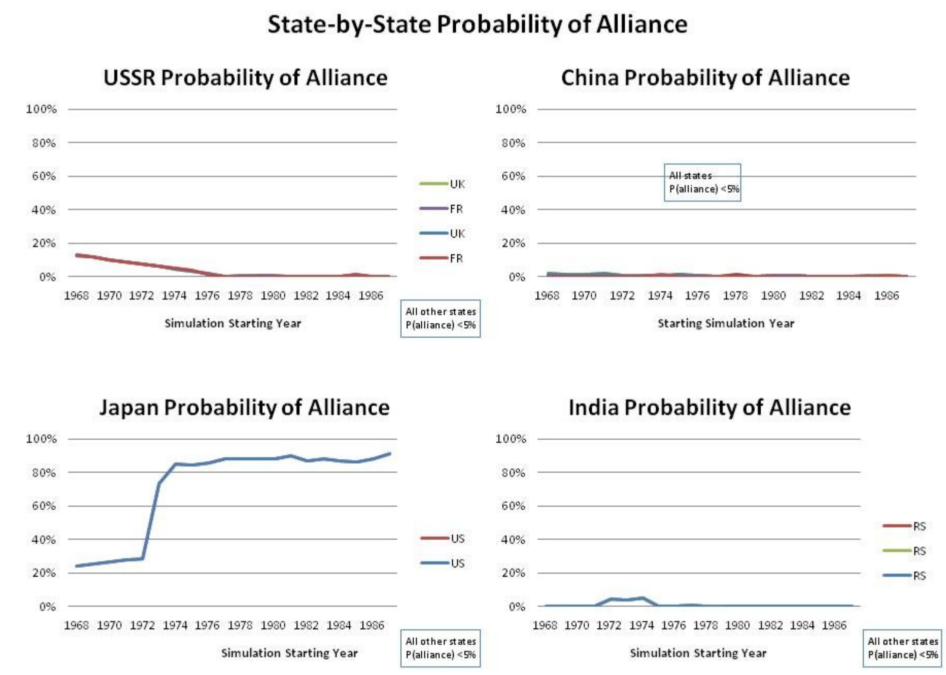

Figure 71: State-by-State Probability of Alliance (Cold War)

As seen in the figure above, the patterns of alliance remain quite constant throughout the simulation period. The anomalous results of decreased probability of western alliance before 1972 are due to a coding issue which caused the historical data not to reflect the already existing alliance data. Notice, however that the simulation effectively regenerated the alliance between NATO partners after just a few years. Also, note that the projected results show France's occasional attempts at dalliance with the Soviet Union, along with India's attempt at *rapprochement* with the Soviet Union in the 1970s.

Capturing the First Years of the Cold War: Short-Initialization Runs

While the simulation results capture the latter part of the Cold War very well, they do not capture the early part due to the need to initialize the model with 20 years of data. One way out of this would be to just start the initialization in 1927 rather than 1947, but the world situation changed so radically during that time that the initialization data would be fairly meaningless.

Instead, I attempted a series of simulation runs with only five years of initialization data. This allows the model to assess the period from the early 1950s on. One problem with this approach is that it produces results which change significantly from year to year. This is because the short initialization period never allows the model to get any long-term data on the behavior of other actors in the system, and so the conditions at the end year of the initialization play an outsize role in where the simulation goes from there. Because of this, the data is presented using 3-year moving averages to smooth out the response and allow for easier interpretation of the results.

Additionally, since this data is not of the same quality as previously discussed, I will only present a few key results instead of the entire set of data given above.

The short-initialization runs created the same basic groups of clusters as seen before: Conflict, Standoff, and Alliance. As before, these clusters were distinct and highly significant across all variables using the F-test (p<.0001 for all variables). Also, as seen previously, the Scheffé test revealed that the clusters were generally distinct from each other, although the center cluster was not distinct from the other two clusters in one of the four variables. Again, since the majority of the variables were statistically significant between clusters (p<.0001), the clusters are judged to be valid for the purposes of the analysis.

The cluster distribution by year was strongly significant (Chi-Sq = 1214, p<.0001) and is shown in the figure below.

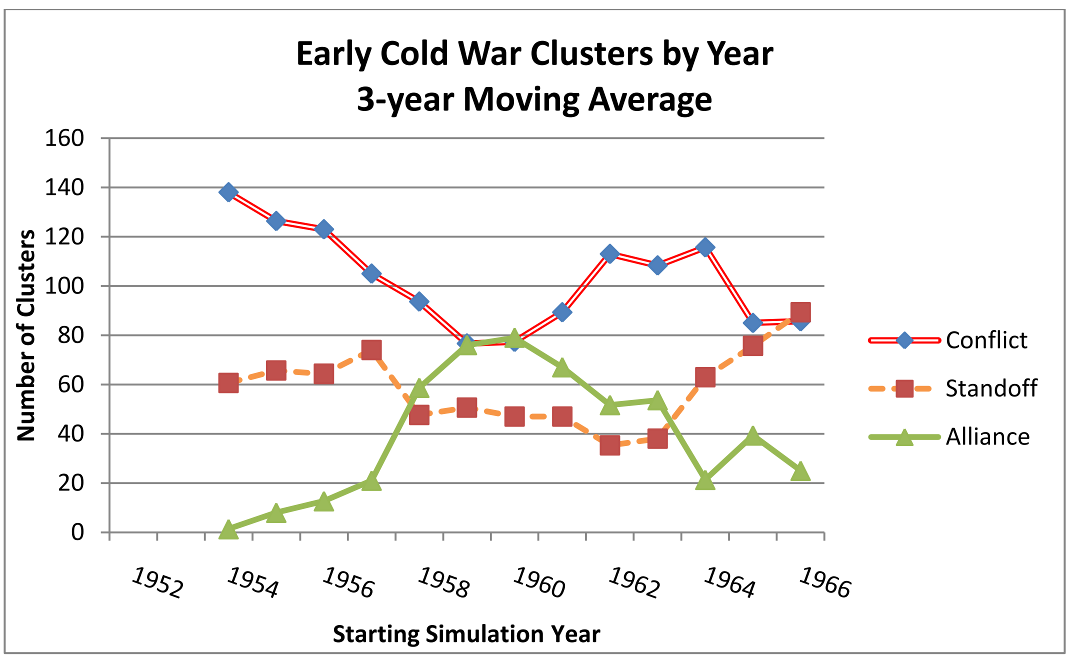

Figure 72: Early Cold War Cluster Distribution

As can be seen, the situation in the early 1950s was very threatening, but the threat levels drops significantly by the late 1950s, only to spike upward with the U2, Berlin Wall, and Cuban Missile Crisis in the early 60s. Plotting the moving average of the mean values for each variable of this time period shows the same basic results.

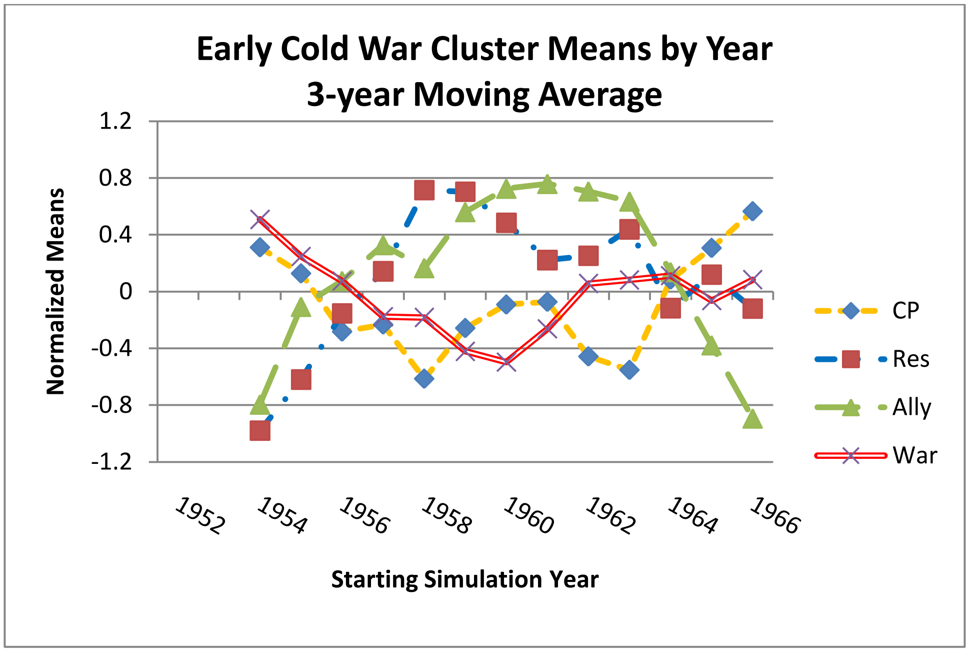

Figure 73: Early Cold War Variable Means by Year

Two things are striking about the above figure. The first is that, despite the crises of the early 1960s, the incidence of war stays pretty close to the mean over the period, which would indicate that even though the threat was high, the actual incidence of war stayed fairly constant. The other notable item is the sharp drop in the number of alliances along with an increase in Combat Power in the mid-60s. The alliance numbers are particularly interesting, given the relative stability of the alliance system in the rest of the Cold War simulation runs. Both of these correspond to the

US military buildup during the Vietnam War, with a significant ramp-up in both US and Soviet military power, and to the strain that the conflict put on American allies.

Similarly, the plot of the probability of conflict shows the pattern of ebb and flow in the level of the general threat.

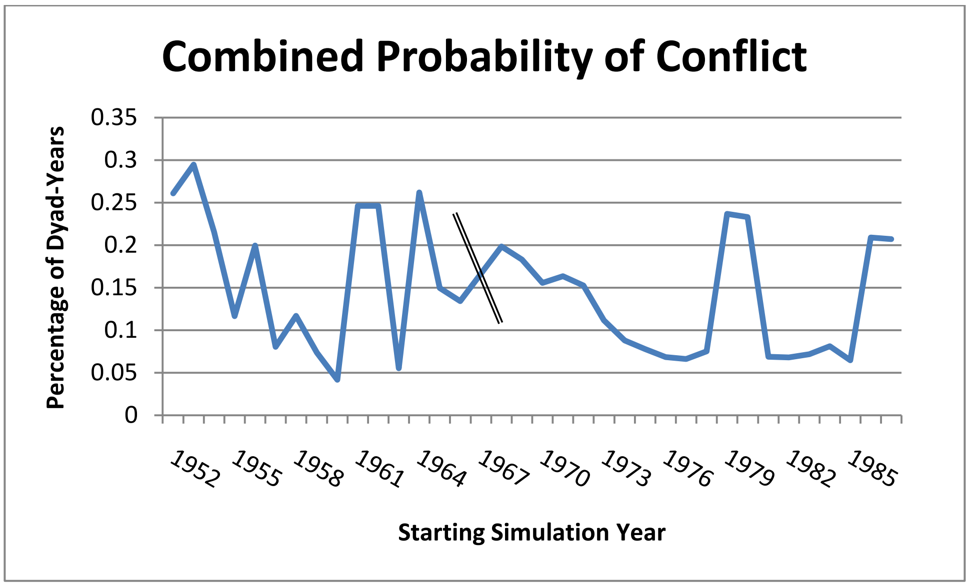

Figure 74: Cold War Probability of Conflict

The above chart combines the probability of conflict for both the main analysis and the short-initialization analysis runs. The double line indicates the break between the two sets of runs. Even though the early part of the chart is more “noisy”, key spikes in the probability of war can be seen after the general decrease in tension after the end of the Korean War. The early 1960s saw a significant probability of war, between the U2 crisis, the Berlin Wall, the Cuban Missile Crisis, and the formal entry of the US into the Vietnam War in 1964.

Conclusion and Observations: The Cold War Case

As seen in the previous chapter, the MEMORIS-2 simulation captures the ebb and flow of conflict and alliance throughout the Cold War period fairly accurately. The increases and decreases in the probability of war match up well to the historical record. Events such as the U2 shoot down, the Berlin Wall crisis, the Cuban Missile Crisis, Vietnam, the Sino-Soviet border wars, and the Sino-Indian wars are all reflected in the simulation predictions.

One thing that is clear is that the dynamics were definitely different between World War I and the Cold War. In the World War I analysis, there was a sudden and definite shift from a state of peace to a high probability of war. During the Cold War, the general trend is a gradually decreasing probability of war, punctuated by intermittent crises. As shown in the figure below, by looking at the distribution of the cluster predictions, it can be seen that the Cold War had three distinct periods. The period of the early Cold War was marked by high-to-moderate, but gradually decreasing threat levels. The period of the 1970s was one of moderate threat, and the latter years of the Cold War were generally low threat, with the exception of the Sino-Soviet and Sino-Indian clashes.

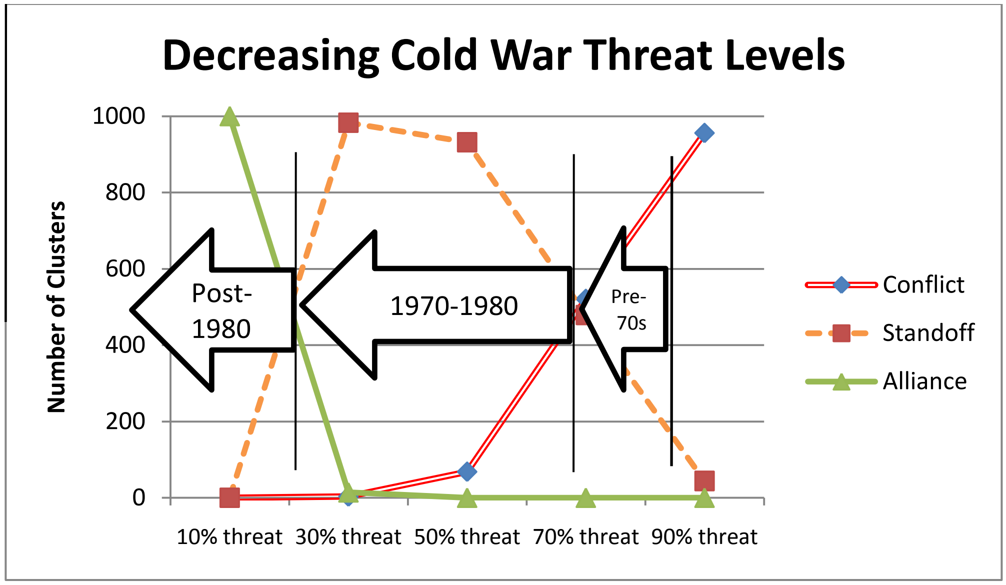

Figure 75: Decreasing Cold War Threat Levels

The simulation did well in pointing out periods of conflict, but as mentioned above, the large spikes for the Sino-Soviet and Sino-Indian conflicts were somewhat surprising. These were surprising not in that they occurred, but that the probability of conflict was as dramatically high as the model predicted. The danger of the Sino-Soviet border conflict spiraling into general war is probably not overstated very much, particularly given that both powers had nuclear weapons and large and capable conventional forces as well. The large amount of conflict associated with the Sino-Indian border wars is more problematic, particularly since both sides kept the conflict limited to a relatively small area. The spike in numbers of dyad-years of conflict is almost entirely due to general war between China and India, and this occurred in virtually every simulation run. There is a constant, low probability of war between the US and China throughout, and a slight increase in the probability of Sino-Soviet war at that time, but the vast majority of conflict during this period was limited to China and India. This is an area that will bear more investigation in future analyses.

As seen in the previous chapter, the simulation did very well in predicting the enduring patterns of alliance throughout the Cold War, along with the strains on US alliances during the Vietnam War. While these results are very encouraging, this set of simulation runs really only included NATO and the US-Japanese alliances as long-term fixtures. The depth and longevity of these alliances is possibly unmatched in history, so future efforts comparing the performance of other alliances, such as the Warsaw Pact and other regionally-based alliances would be potentially quite useful.

The role of nuclear weapons is another area for further study. I speculated earlier that the fact that, in reality the Sino-Soviet and Sino-Indian wars did not spiral into general conflict might well have been due to the deterrent power of nuclear weapons. In the first case, both sides were nuclear weapons states, and in the second, the Chinese did have nuclear weapons, and the Indian government was certainly capable of producing them ever since their so-called “Peaceful Nuclear Explosion” in 1974. This is also an area for further study, and I will have more to say about this during the concluding chapter of the dissertation.

Chapter Six: Observations and Conclusions

Revisiting the Research Question

I started this dissertation in an effort to better understand the process of decision-making, and specifically, to understand how the way our minds process data affects a decision maker's actions. I began the analysis with the intuition that there was an interaction between the decision space and the way in which information was processed by the decision-maker, and further, that this interaction could be simulated. Most important, that once the process could be simulated, the outcome could be probabilistically predicted.

Specifically, once the simulation was built, I wanted to get the answers to the following questions:

- Given starting conditions, can the probability of conflict be predicted?
- If there is conflict, who are the likely participants, and on which side?
- What is the structure of alliances?
- What are the probabilities for each variation of the above results?

The process of building and testing the MEMORIS-2 simulation shows that the way in which people mentally process information is important, if not critical, in determining the resulting decision. This conclusion was validated during the two game theory tournaments in Chapter Three, along with the example cases of World War I and the Cold War in Chapters Four and Five.

Performance of the Simulation

When building models, it is important to remember that they are not reality. Models provide an explicit way to explore key variables and their relationships within a given situation. The MEMORIS-2 simulation is no different. It takes a very limited set of input data and then simulates the interaction of states within that context. The simulation does not provide an "answer" or point prediction of reality. Instead, MEMORIS-2 shows how emergent behavior – which in this case is represented by choices for war or alliance, can emerge from the interaction between available information and the way that information is processed. The tool of simulation allows these interactions to be repeated many times, and thus we can gain insight into both likely results and counterfactuals – particularly by seeing that the same situation can lead to multiple results. Some of these results are more likely than others, and this can lead to better understanding of the relative risks inherent in a given situation or choice of actions. Further, simulation allows the impact of specific factors to be explicitly tested rather than simply assumed.

Nevertheless, any model must have some grounding in reality, and should be able to readily demonstrate some face validity otherwise it will not be accepted. This was accomplished in several ways. In the game theory setting, the memory-based simulation was at least equal to, if not superior to other strategies and heuristics. This, I believe, is reflective of the ability of the human brain to recognize patterns and adapt to new situations. This ability, which was embedded in the simulation, allowed the memory-based model to outperform other models, particularly when the situation was dynamic.

Further, the memory-based simulation approach did not require me to make any particular assumptions about either the game structure or player preferences in order to determine player strategies. The simulation was able to determine the game structure from within, and to adjust play to gain the best results. Preference structure within the game changed as the situation changed, so although there was a general preference to "do well", the manifestations of that general preference changed based on the specific situation.

Additionally, because the action of the memory trace model replicates the heuristics and biases in the decision-making literature, the simulation does not require any assumptions about the presence or absence of such heuristics and biases. In effect, the MEMORIS-2 simulation probabilistically incorporates heuristics and biases. To take a commonplace example, a casino does not need to know which strategy or "system" that a player is using to make bets – all the owners need to know is that, in the aggregate, the house will win a certain, calculable percentage of the money wagered. Similarly, MEMORIS-2 predicts, in the aggregate, the probability of war vs. peace, alliance vs. non-alliance, and who will be involved in either conflict or alliance.

In the two cases, the MEMORIS-2 simulation identified when conflict was likely, identified the initial players in the conflict, and identified the formal alliance structures. As pointed out in Chapters Four, the simulation did not always identify the secondary players in the conflict, so even though the original flash point of the war was correctly identified as Eastern Europe, the other major powers usually sat out the conflict.

Further, the simulation was unequivocal in differentiating between peace and war. The end of the World War I simulation predicted conflict or the continuation of armed standoff with near certainty, while the end of the Cold War simulation predicted a 90% chance of peace (and would have been higher absent the Sino-Indian wars). This compares quite favorably to other results predicted in the literature. For example, as previously noted in Chapter Two, expected utility modeling predicted a doubling in likelihood ratio for states liable to go to war, although the actual percentage was well less than 1%. In contrast, the MEMORIS-2 prediction for war after 1910 is 30% per year, which translates into a six-fold increase in likelihood ratio compared to the pre-1910 state. Similarly, the predicted probability of a peaceful end to the Cold War showed a likelihood ratio of 5 to 1.

Implications for IR Theory

The success of the MEMORIS-2 simulation has several implications for IR theory. As mentioned in Chapter 2, both Realist and Liberal IR theory have well-known problems in accounting for major systemic changes, particularly since both assume that states operate from a set of exogenously-given interests within a defined international system. The results of this dissertation directly challenge these underlying assumptions, and instead demonstrate that conflict and cooperation, along with change and continuity, can result from the same cognitive processes and the same given situation.

Further, the underlying logic of MEMORIS-2 does not require normative assumptions or beliefs about how states *should* act, as is the case with both Realism and Liberalism. Instead, it shows in a replicable way how states *probably will* act in a given situation. The analysis shows that both structure and interest are malleable, even within the context of a game where payoffs for given actions are constant.

Even staying within the realist or liberal paradigm, the MEMORIS-2 simulation offers the opportunity to explicitly test the operation of several mechanisms that have been contested within those paradigms. For example, the question of what causes states to ally with each other can be explored via sensitivity analysis of the various parameters in the model, such as pursuit of relative vs. absolute gains, and the tendency to bandwagon or balance. By varying these factors, the probability that states will or will not ally with others in various situations can be calculated, and this may help clarify the debate on alliance formation within the realist community. Similarly, sensitivity analysis on these factors could provide increased clarity in other contested areas of both realist and liberal theory.

This simulation approach could also be used to explore the working of reputation. In the MEMORIS-2 simulation, reputation is not a separate variable, but the action of reputation is an emergent property which is developed through interaction. Further, the working of reputation appears to be non-monolithic – state reputation in the combat power part of the simulation is different than that in the resource part of the simulation. While the two examples used in this dissertation were not explored in enough depth to make any specific judgments, further analysis could gain traction on whether states with given reputations are more effective than states with opposite reputations.

I think that the main impact of this work will be in the area of Constructivism. One of the main critiques that constructivists have leveled at other theories is that they assume that rational choice, which is based on economic modes of analysis, is the way that best describes how states should act. As has been pointed out above, the assumptions underlying the belief that humans act as *homo economicus* are highly suspect, and even instrumental rationality cannot salvage them.

The issue for Constructivism is that it has previously lacked tools to formally and explicitly explore *how* the interaction between actors shapes actor identities, interests, and the overall structure. By showing how this interaction occurs cognitively, the MEMORIS-2 simulation is a starting point for further development of tools that are not based in untenable assumptions of rational-economic analysis, but instead are based on cognitive processes which are shared by all humans.

Further, this tool can be used to explore the previously undiscovered terrain of how structure and interest are determined, and what set of interests and structure is likely to emerge from a given situation.

Implications for Policy Makers

The most important application of the MEMORIS-2 simulation is that it answers an important question for policy makers. When a situation occurs, there are plenty of voices within the policy process who are more than willing to tell the decision maker what action *should* be taken. While a decision maker may chafe at the unwanted advice, with the (partially warranted) view that it is one thing to propose action and another to bear the consequences of ordering that action, the multitude of viewpoints can serve a positive end by adding to the possibility that important factors will not end up being overlooked.

Nevertheless, the weakness of policy advice based on what *should* be done is that the assessment of what *will* result is inextricably intertwined with the judgment of what should be done. Causes and effects are linked, and so a policy prescription based on a given analysis will undoubtedly predict an outcome based on that same analysis.

What this mode of predicting results misses is that each action and interaction creates a new reality, and the resulting situation may turn on a different axis of policy. If this is not true, then why is the "law of unintended consequences" the only practical maxim that everyone believes in?

The need for decision makers is clear – a way to help them understand the question that no one can answer conclusively: what *will* happen if I choose this action? This requires the ability to assess how the situation will develop dynamically, and to predict how the given set of actions will change the picture for other players in the system. MEMORIS-2 will allow such a decision aid to be developed. By its ability to repeatedly simulate the complex interaction between the situation, information, and how that information is processed, MEMORIS-2 demonstrates that the ability to probabilistically predict the likely outcomes of a given set of actions is within reach. This is not only useful as a set of probable outcomes, but also will help quantify the relative risk of a proposed action.

Beyond the potential to create a decision aid for policy makers, I want to return to the first figure of the dissertation, which is reproduced below. Initially, I posited a dynamic model of emergent behavior for entry into war, as shown in the figure below.

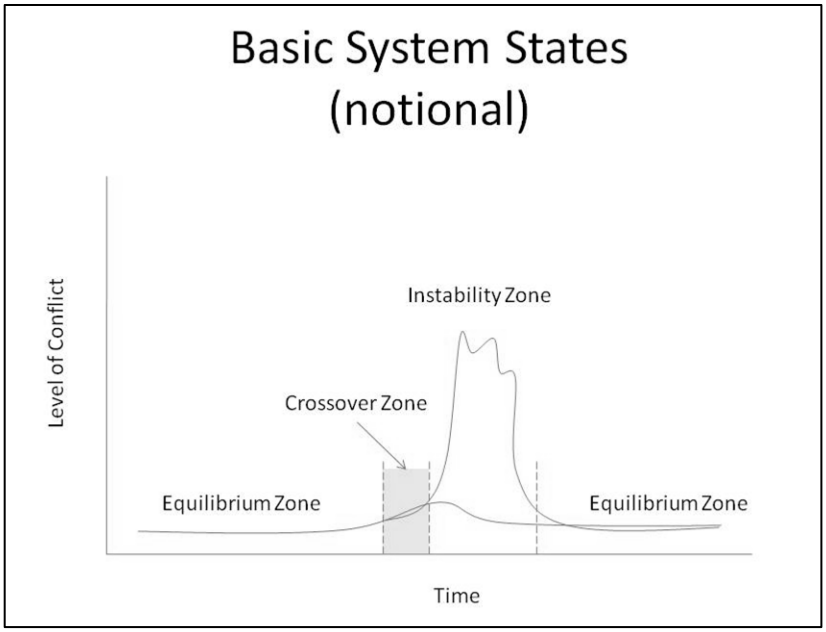

Figure 76: Notional Systems States

My intuition was that as a state or states approaches conflict or other instability, the system enters a "*crossover zone*" when actions which formerly may have been stabilizing now become destabilizing, and the situation can escalate with little or no warning. Here, actions which one side views to be prudent instead are viewed by the other side as threatening. This insight is not new, but the fact that the MEMORIS-2 simulation can show when a group of states enters that zone is new. Further, the insight that players can shift preference structures based on the situation is also new.

This situation can be seen with reference to the World War I example. In the pre-1910 state, the simulation showed that the European system was in a general peaceful state, perhaps tense but with low threat. Within a very short time, the threat escalated to very high levels, and conflict appeared nearly certain – the only question was what would set it off. This sudden shift in the level of threat is shown in the figure below from Chapter Four.

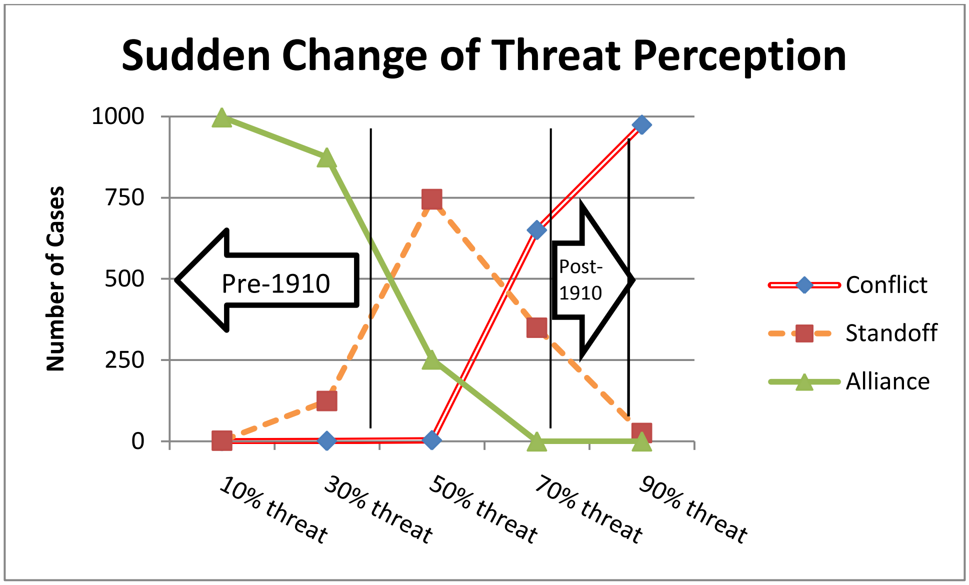

Figure 77: WW1 Sudden Shift in Threat Level

This state of affairs is due to the way in which information is processed. Without re-plowing the ground on the discussion in Chapter Three, the impact of the process of information search and retrieval is generally confirmatory, which pushes the system toward incrementalism and maintaining the status quo as the normal pattern, as judgments will tend to shift incrementally. While moving incrementally, this confirmatory information search also implies a positive feedback loop, causing judgments to diverge from reality as conflicting information is initially overlooked. This positive feedback can occur unchecked for some time. This dynamic can, as seen in the World War I case, create the situation where states end up going to war, not because the decision makers want to, but because they see no other alternative.

This is also in reflected in the shifts in player preferences as the crossover zone is reached. Before the crossover zone is reached, player preference is usually toward non-threatening behavior. When the crossover zone is reached, the perception of a new situation also causes a shift in preferences to more "defensive" actions, such as military buildups and pre-emptive threats.

Change in the system occurs once the set of contradictory information becomes too large, or once the set of confirmatory information is effectively forgotten. Either of these will cause the hypothesis used to probe memory to change. The judgment about a given situation will then rapidly shift as the "crossover zone" is reached and the decision maker sees previous information in a new light. This rapid shift in judgment will likely be accompanied by a rapid shift in policy, which will cascade through the system until a new equilibrium is reached, leading to another period of relative stability.

Similarly, the situation on the back half of the instability curve (Figure 76) shows a move from relative instability and high levels of threat to lower levels of threat as the system reaches a new equilibrium. This is arguably what happens in the Cold War, as predicted by the simulation. This can be seen in the figure below, from Chapter Five.

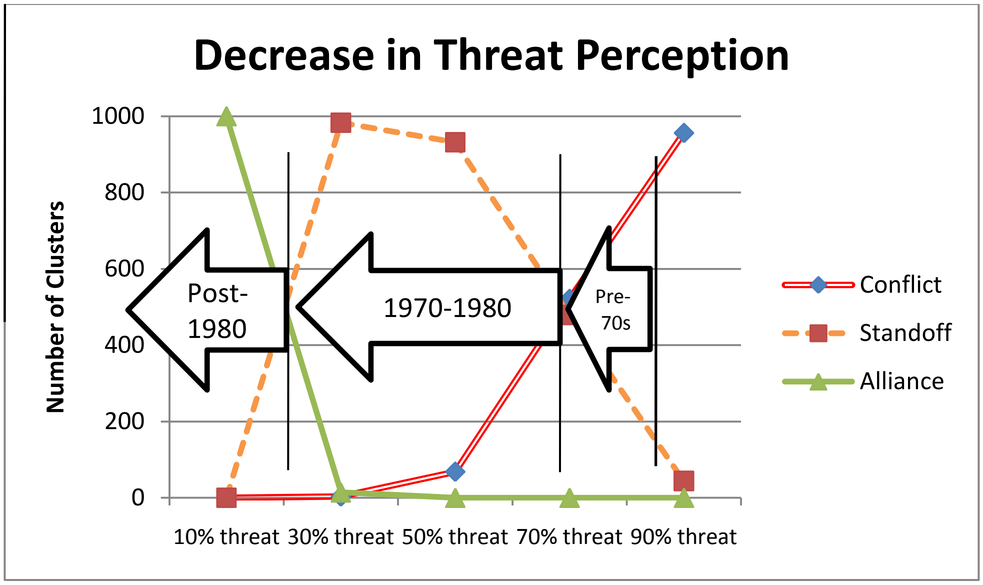

Figure 78: Cold War Threat Distribution

The above figure shows the pre-1970 state of high threat, gradually transitioning through the moderate threats of the 1970s, and then to a relatively stable situation after that.

There are two lessons which can be potentially drawn from this. The first is that actions taken within the crossover zone are very likely to be misunderstood. The MEMORIS-2 simulation has the ability to predict when the system is entering such a zone, and so could be used as a policy-support tool to show when the situation is in danger of getting out of hand. As shown in World War I, actions that seemed prudent to one side were deemed as immediate threats to the other. While this dynamic has been predicted in earlier work on the security dilemma, the difference here is that such a dilemma can emerge spontaneously based on the shifting perception of the situation, rather than from any overt material factors. Again, this occurs because the interaction of the players causes a shift in the structure of interest and therefore a change in the structure of the system. This change only occurs when the perception of the decision space and the menu of possible choices begin to feed back on each other – once this occurs the decision space shrinks as choices are seen in a different light, which further shrinks the perceived decision space, and so on. As seen, this can happen with bewildering rapidity, and once the crossover zone is reached, the decision-makers find themselves girding for war not because they want to, but because they see no other choice.

The decision-maker, then, needs to be very careful when approaching the crossover zone – and further, needs to take actions to actively dampen the crisis rather than "just in case" preparations that may only serve to push the situation over the edge. Mistakes then can happen very quickly, and there may be no pulling back once that happens.

The second lesson is to avoid repeatedly "pushing the envelope". There is a tendency here to think that, now that the crossover zone can be predicted, that an aggressive decision-maker would try to push the situation right up to the edge of the zone in the hope of making the other party blink. In flying, the winner of the fight is often the one who can fly the airplane right to the proverbial "edge of the envelope", where the rumble of buffet in the stick lets you know that you are getting maximum performance out of the airplane, but any further back pressure will put you out of control. The danger here is that you can lose the fight by simply being too aggressive. The danger is even higher in the decision-making arena, particularly when it comes to war. The MEMORIS-2 simulation models cognitive processes, not aerodynamics, so despite the fact that the predictions turn out to be well-behaved numerically, the margin of error is still significant, and so it is better to err on the side of caution than to push the envelope. This runs contrary to the idea of "brinkmanship" from deterrence theory. Instead of showing resolve, the repeated practice of brinkmanship can lead to disaster – the point of no return is probably closer than we know. That said, the use of MEMORIS-2 or a follow-on simulation can help show when that crossover zone is being approached, or even if it has already been entered without knowing it.

Implications for Nuclear Deterrence

One of the most important, if not *the* most important unanswered questions of the Cold War is that of the efficacy of nuclear deterrence. As stated previously, many scholars and policy makers, myself included, have asserted that nuclear weapons kept the peace. I would not have spent the years I did on nuclear alert, ready to deliver those weapons, if I had not believed it to be true. There are, thankfully, no data points to the contrary.

The comparison of the World War I and Cold War results does, however, show some tantalizing evidence to the contrary, as seen in the figure below.

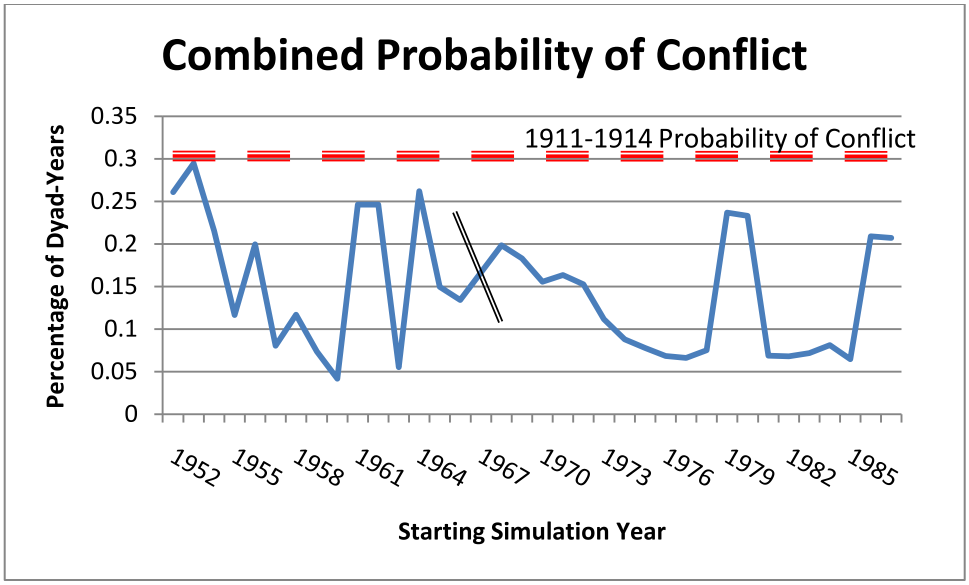

Figure 79: Combined Probability of Conflict

Here, I have overlaid the probability of conflict chart for the Cold War case with the probability of conflict during 1911-1914. As can be seen, the Cold War probability of conflict only reaches that point during the Korean War, but does come within 5% during the early 1960s. This could be interpreted as showing that the probability of war was consistently below, and usually well below the pre-WW1 level for the entire period of the Cold War. Thus, the probability of conflict, although terrifying in spots, never reached the point at which war would have become virtually certain as it did prior to WW1. Taking this reasoning to its conclusion would indicate that since the probability of war was lower, then the role of nuclear weapons in keeping the peace may well have been overstated.

Further, the predicted probability of conflict between the US and USSR was quite low – less than 5% for most of the Cold War period (Figure 70, Chapter 5). This indicates that direct, major conflict between the two superpowers was unlikely during that time, which would also provide ammunition to discount the deterrent value of nuclear weapons.

There is one potential problem with this line of reasoning, and that is that nuclear weapons did exist, and the current simulation does not allow me to unpack their potential role in limiting the extent and violence of the actual crises in the historical record. For example, absent US nuclear weapons, the Soviets may well have decided to simply take Berlin in a *coup de main* rather than settling for blockade in 1948. Whether this would have triggered general war absent nuclear weapons in another issue, as neither side may have had the stomach for large-scale war so soon after the devastation of WWII.

This argument can be stood on its head though. Absent the nuclear standoff, would some of the major crises of the Cold War even have *been* crises? It is hard to imagine that shooting down a reconnaissance plane would be grounds for all-out war. Nor is it easy to imagine how Khrushchev would have been able to covertly deploy significant combat power to Cuba in 1962 absent nuclear weapons.

Nevertheless, the MEMORIS-2 simulation does not capture the existential threat posed by nuclear weapons, and so the tantalizing possibility that nuclear deterrence was not needed to keep the sides from all-out war is just that.

Avenues for Future Research

Although this dissertation answers an important question and advances the understanding of decision-making behavior, it also raises several other questions and avenues for further study.

First, there is the issue of nuclear deterrence. It would be fairly straightforward to try to gauge the impact of nuclear weapons by changing the game payoff structure for nuclear-armed states, and then comparing the results of that simulation with the current Cold War results. If the revised simulation showed lower levels of conflict along with decreased conflict between nuclear-armed powers, it would indicate that the existential threat of nuclear weapons did have a damping effect on the escalation of conflict. Contrariwise, it is entirely possible that higher levels of conflict may result, and that would indicate that, rather than be a stabilizing factor, nuclear weapons may be destabilizing.

Another avenue for research is in the area of alliances. While the MEMORIS-2 simulation was remarkably accurate in predicting enduring patterns of alliance, it did not entirely capture the effect of alliances on bringing other states into World War I. This could mean several things. One is that the entry, in particular, of Germany into the war was not a foregone conclusion. That seems unlikely, given the historical record, but bears further analysis. Another is that the alliance ties between states are not properly implemented, and the simulation logic needs further adjusting to account for such situations. A third and more interesting idea is that there is a difference in capability or focus between states. This would mean, for example, that states may be more or less likely to be concerned with relative gains, and that would in turn impact their propensity for conflict. In another example, states may also have greater or lesser capability to properly assess the situation. Both of these cases are controlled by parameters within the situation. During the dissertation runs, these parameters were kept constant after their initial tuning, but future efforts to perform sensitivity analysis using these parameters would yield more insight into their role in alliance behavior.

As mentioned previously, the simulation could also be used to compare various regionally-based alliances to see if the patterns shown in the two dissertation cases hold for other alliances. It is more likely that other alliances would have differing results, and thus the simulation could perform some interesting comparative analysis.

Similarly, the performance of the MEMORIS-2 simulation could be validated against other regions of interest. There are many historical and possibly future regional wars that could be studied. Also, the model could be scaled up to take on more players, which only takes either faster hardware or more efficient code to make possible. In this vein, the model could also be scaled down to sub-state levels to look at the interaction of actors in revolutionary or ethnic conflicts.

Finally, one of the issues in International Relations theory is the lack of tools such as game-theoretic or expected utility that can operate within the context of Constructivism. Here, the MEMORIS-2 simulation, when extended, offers the opportunity to create a quantitative modeling foundation to explicitly simulate the relationship between social discourse and structure. This could give an avenue to answer the nagging questions of how certain value structures become dominant within the international system.

Final Remarks

There are several important differences between the basis of the MEMORIS-2 simulation and other approaches. The first is that the same process (probes of memory) and the same data (traces in memory and semantic concepts) is used throughout. There is, therefore, no need to determine which bias or heuristic is operational, or to determine the decision-maker's affective state. The judgment about the situation at hand is effectively framed by the probes of memory. Further, there is no need to make *a priori* judgments about the game or player utility functions, as both likelihood and utility judgments are calculated based on existing traces in memory and the way in which those traces are probed. Instead of a rational choice or expected value simulation, it is a way to simulate complex social judgment and interaction.

Further, this type of simulation offers the ability to aid policy makers by allowing probabilistic judgments of what will happen if a given policy choice is taken. Such judgments will not suffer the weakness of judgments which are based on a given policy prescription and thus assume that the set of factors which operated before the action will operate equally after the action. An assessment of what *will* occur, even a probabilistic one, along with the relative risk of things going wrong, will be an invaluable tool to policy makers.

As such, MEMORIS-2 and the follow-ons outlined above offers a solid platform and departure point for future efforts. This dissertation has shown that the tool works in a variety of settings. One tool is certainly not enough – otherwise our tool boxes would be much smaller – but this appears to be a very flexible and extensible tool that can serve a variety of purposes. I look forward to extending that tool in the near future.

Bibliography

‘t Hart, Paul, Stern, Eric & Sundelius, Bengt, (eds.) 1997. *Beyond Groupthink: Political Group Dynamics and Foreign Policy-Making.* Ann Arbor, MI. University of Michigan Press.

Abelson, R. (1963). “Computer Simulation of ‘Hot’ Cognition,” in S. Tomkins and D. Messick (eds.), *Computer Simulation of Personalit,.* pp: 277-298. New York, NY. Wiley.

Adler, Emmanuel (2002). “Constructivism and International Relations,” in Walter Carlsnaes, Thomas Risse & Beth Simmons (eds.) *Handbook of International Relations,* pp: 95-118. London, UK. Sage Publications.

Allison, Graham (1971). *Essence of Decision: Explaining the Cuban Missile Crisis.* New York, NY. Harper Collins.

Alpern, Steve (1993). “Stationary Equilibria for Deterministic Graphical Games,” in Ken Binmore, Alan Kirman & Piero Tani (eds.) *Frontiers of Game Theory,* pp: 95-102. New York, NY. Cambridge University Press.

Angel, Norman (1911). *The Great Illusion: A Study of the Relation of Military Power in Nations to their Economic and Social Advantage.* New York, NY. G.P. Putnam’s Sons.

Ariely, Dan (2008). *Predictably Irrational: The Hidden Forces That Shape Our Decisions.* New York, NY. Harper Collins.

Arrow, Kenneth (1951). *Social Choice and Individual Values.* New Haven, CT. Yale University Press.

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

--------- (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Cooperation*. Princeton, NJ: Princeton University Press.

Baldwin, David (1993). “Neoliberalism, Neorealism, and World Politics.” *Neorealism and Neoliberalism: The Contemporary Debate*, pp: 3-25. New York, NY. Columbia University Press.

Barnett, Michael & Duvall, Raymond (2005) “Power in Global Governance,” in Barnett, Michael and Duvall, Raymond (eds.) *Power in Global Governance*, pp: 1-32. Cambridge, UK. Cambridge University Press. 1-32.

Barnett, Michael & Finnemore, Martha (2004). *Rules of the World.* Ithaca, NY. Cornell University Press

--------- (2005). "The Power of Liberal International Organizations" in Barnett, Michael and Duvall, Raymond (eds.) *Power in Global Governance*, pp: 161-184. . Cambridge, UK. Cambridge University Press.

Bates, Robert (2005). *Beyond the Miracle of the Market, New Edition.* Cambridge, UK. Cambridge University Press.

Bennett, D. Scott, and Allan Stam (2000). "*EUGene*: A Conceptual Manual." *International Interactions* 26:179-204. Website: http://eugenesoftware.org.

--------- (2004). *The Behavioral Origins of War.* Ann Arbor, MI. University of Michigan Press.

Berejikian, Jeffrey (2002). "Model Building With Prospect Theory: A Cognitive Approach to International Relations," *Political Pyschology,* Vol. 23, No. 4: 759-786.

Binmore, Ken (1994). *Game Theory and the Social Contract II: Just Playing.* Cambridge, MA. The MIT Press.

Boettcher, William A. III (2004). "Prospects for Prospect Theory: An Empirical Evaluation of International Relations Applications of Framing and Loss Aversion," *Political Psychology,* Vol. 25, No. 3: 331-362.

Bremer, S.A. & Mihalka, M. (1977). "Machiavelli in Machina: Or Politics among Hexagons," in Deutch, K., et al. (eds) *Problems of World Modeling: Political and Social Implications,* pp. 303-338. Cambridge, MA. Ballinger Publishing.

Bueno de Mesquita, Bruce (1984) "Forecasting Policy Decisions: An Expected Utility Approach to Post-Khomeini Iran," *PS*, Vol. 17, No. 2 (Spring 1984): 226-236.

--------- (1985) "The War Trap Revisited," *The American Political Science Review,* Vol. 79, No. 1 (March 1985): 156-177.

--------- (2000) "Popes, Kings, and Endogenous Institutions: The Concordat of Worms and the Origins of Sovereignty," *International Studies Quarterly,* Vol. 2, No. 2 (Continuity and Change in the Westphalian Order, Summer 2000): 93-118.

--------- (2002) *Predicting Politics.* Columbus, OH. The Ohio State University Press.

Bueno de Mesquita, Bruce (2004). "The Methodological Study of Politics," in Ian Shapiro, et al. (eds) *Problems and Methods in the Study of Politics.* Cambridge, UK. Cambridge University Press. pp. 227-250.

Bueno de Mesquita, Bruce & Lalman, David (1988). "Empirical Support for Systemic and Dyadic Explanations of International Conflict," *World Politics,* Vol. 41 (October): 1-20.

--------- (1992). *War and Reason: Domestic and International Imperatives.* New Haven CT. Yale University Press.

Bueno de Mesquita, Bruce; Smith, Alastair; Siverson, Randolph; and Morrow, James (2003) *The Logic of Political Survival.* Cambridge, MA. MIT University Press.

Bueno de Mesquita, Bruce and McDermott, Rose (2004) "Crossing No Man's Land: Cooperation from the Trenches," *Political Psychology*, Vol. 25, No. 2 (April 2004): 271-287.

Bull, Hedley (1997). *The Anarchical Society: A Study of Order in World Politics.* New York, NY. Columbia University Press.

Busemeyer, Jerome & Townsend, James (1993). "Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment," *Psychological Review,* Vol. 100, No. 3: 432-459.

Buzan, Barry (1994). "The Interdependence of Security and Economic Issues in the 'New World Order'," in Richard Stubbs & Geoffrey Underhill (eds.) *Political Economy and the Changing Global Order*, pp: 89-102. New York, NY. Saint Martin's Press.

Carr, Edward (1946). *The Twenty Years' Crisis 1919-1939: An Introduction to the Study of International Relations.* New York, NY. St Martins Press.

Cederman, Lars-Erik (1994). "Emergent Polarity: Analyzing State-Formation and Power Politics," *International Studies Quarterly,* Vol. 38, No. 4 (Dec 1994): 501-533.

--------- (2002). "Endogenizing Geopolitical Boundaries with Agent-Based Modeling," *Proceedings of the National Academy of Sciences of the United States of America,* Vol. 99, No. 10, Supplement 3 (May 14, 2002): 7296-7303.

--------- (2002a). "Modeling the Size of Wars: From Billiard Balls to Sandpiles," *The American Political Science Review,* Vol. 97 (2003): 135-150.

Clausewitz, Carl von (1989 [1874]). *On War.* Michael Howard & Peter Paret, ed. and trans. Princeton, NJ. Princeton University Press.

Cusack, T.R. & Stoll, R. (1990). *Exploring Realpolitik: Probing International Relations Theory with Computer Simulation.* Boulder, CO. Lynne Rienner.

Deutch, Karl (1966). *The Nerves of Government: Models of Political Communication and Control.* New York, NY. The Free Press.

Dougherty, Michael; Gettys, C.F.; & Ogden, F.E. (1999). "Minerva-DM: A Memory Processes Model for Judgments of Likelihood," *Psychological Review,* Vol. 106: 180-209.

Dougherty, Michael & Hunter, Jennifer (2003) "Probability Judgment and Subadditivity: The role of working memory capacity and constraining retrieval," *Memory and Cognition,* Vol. 31, No. 6: 968-932.

Downs, Anthony (1957). *An Economic Theory of Democracy.* New York, NY. Harper and Row.

Dyson, Stephen & Preston, Thomas (2006). "Individual Characteristics of Political Leaders and the Use of Analogy in Foreign Policy Decision Making," *Political Psychology,* Vol. 27, No. 2 (2006): 265-288.

Etzioni, A. (1967). "Mixed Scanning: A Third Approach to Decision-Making," *Public Administration Review,* Vol. 27, No. 5 (Dec 1967): 385-392.

Farrar, L.L. Jr. (1972). "The Limits of Choice: July 1914 Reconsidered," *The Journal of Conflict Resolution,* Vol. 16, No. 1 (Mar 1972): 1-23.

Fearon, James (1995). "Rationalist Explanations for War," *International Organization,* Vol. 49, No. 3 (Summer 1995): 379-414.

Ferguson, Yale & Mansbach, Richard (2003). *The Elusive Quest Continues: Theory and Global Politics.* Upper Saddle River, NJ. Prentice Hall.

Fiedler, Klaus (2000) "Beware of Samples! A Cognitive-Ecological Sampling Approach to Judgment Biases," *Psychological Review*, Vol. 107, No. 4: 659-676.

Frank, Andre Gunder (1966). "The Development of Underdevelopment," *Monthly Review*, Vol. 18, No. 4 (1966): 17-31.

Gaddis, John (2005). *The Cold War: A New History.* New York, NY. Penguin Press.

Galan, Jose & Izquierdo, Luis (2005). "Appearances Can Be Deceiving: Lessons Learned Re-Implementing Axelrod's 'Evolutionary Approach to Norms'," *Journal of Artificial Societies and Social Simulation,* Vol. 8, No. 3 (Jun 2005).

Gallhofer, Irmtraud & Saris, Willem (1997). *Collective Choice Processes: A Qualitative and Quantitative Analysis of Foreign Policy Decision-Making.* Westport, CT. Praeger.

Galtung, Johan (1971). "A Structural Theory of Imperialism," *Journal of Peace Research*, Vol. 8, No. 2 (1971): 85-91.

Gant, M.M. and Davis, D.F. (1984). "Mental Economy and Voter Rationality: The Informed Citizen Problem in Voting Research," *The Journal of Politics,* Vol. 46, No. 1 (Feb 1984): 132-153.

Gavanski, Igor & Hui, Chun (1992) "Natural Sample Spaces and Uncertain Belief," *Journal of Personality and Social Psychology,* Vol. 63, No. 4: 766-780.

George, Alexander (1969). "The 'Operational Code': A Neglected Approach to the Study of Political Leaders and Decision-Making," *International Studies Quarterly*, Vol. 13, No. 2 (June 1969): 190-222.

George, Alexander (1979). "The Causal Nexus Between Cognitive Beliefs and Decision-Making Behavior: The 'Operational Code' Belief System," in Lawrence Falkowski (ed.), *Psychological Models in International Politics*, pp: 56-124. Boulder, CO. Westview Press.

Gettys, C. F., Pliske, R. M., Manning, C., & Casey, J. T. (1987). An evaluation of human act generation performance. *Organizational Behavior and Human Decision Processes,* Vol. 39: 23 – 51.

Ghosn, Faten, and Scott Bennett. 2003. Codebook for the Dyadic Militarized Interstate Incident Data, Version 3.0. Online: http://cow2.la.psu.edu.

Gigerenzer, Gerd; Hoffrage, Ulrich & Kleinbölting, Heinz (1991) "Probabalistic Mental Models: A Brunswikian Theory of Confidence," *Psychological Review,* Vol. 98, No. 4: 506-528.

Gilpin, Robert (1987). *The Political Economy of International Relations.* Princeton, NJ. Princeton University Press.

Glaser, Charles (1994/95). "Realists as Optimists: Cooperation as Self-Help," *International Security,* Vol. 19, No. 3 (Winter 1994-95): 50-90.

Gordon, Michael (1974). "Domestic Conflict and the Origins of the First World War: The British and German Cases," *The Journal of Modern History.* Vol. 46, No. 2 (June 1974): 191-226.

Guetzkow, Harold (1962). "Inter-Nation Simulation: An Example of a Self-Organizing System," in M.C. Yovits, et al., eds. *Self-Organizing Systems.* Washington, DC. Spartan.

Hair, Joseph; Anerson, Rolph; Tatham, Ronald & Black, William (1998). *Multivariate Data Analysis, 5th ed.* Delhi, India. Pearson Education.

Hammond, KR, Hursch, CJ, & Todd, FJ (1964). "Analyzing the components of clinical inference," *Psychological Review*, Vol. 71, No. 6: (1964): 438-456.

Hanson, William (2006). "Can We Remember Rationally? A Memory-Based Model of State Interactions," paper presented at the American Political Science Association, 2 Sep 2006.

Hasenclever, Andreas; Mayer, Peter & Rittberger, Volker (1997) *Theories of International Regimes.* Cambridge, UK. Cambridge University Press.

Hermann, Charles & Hermann, Margaret (1969). "An Attempt to Model the Outbreak of World War I," in James Rosenau (ed.) *International Politics and Foreign Policy.* New York, NY. Free Press. pp. 622-639.

Hindmoor, Andrew (2006). *Rational Choice.* New York, NY. Palgrave.

Hintzman, D. L. (1988). "Judgments of Frequency and Recognition Memory in a Multipletrace Memory Model," *Psychological Review*, Vol. 96 (1988): 528 – 551.

Holsti, Ole. (1965). "The 1914 Case," *The American Political Science Review,* Vol. 59, No. 2 (June 1965): 365-378.

--------- (1972). *Crisis, Escalation, War.* Montreal. McGill-Queens University Press.

--------- (1976). "Foreign Policy Decision-Makers Viewed Cognitively," in Robert Axelrod (ed.) *Structure of Decision*, pp: 18-54. Princeton, NJ. Princeton University Press.

--------- (1977). "Content Analysis," in Donald Freemen (ed.) *Political Science: History, Scope, and Methods*, pp: 169-186.

Holsti, Ole; North, Robert & Brodie, Richard. (1968). "Perceptions and Action in the 1914 Crisis," in J.D. Singer (ed.) *Quantitative International Politics.* New York, NY. Free Press, pp. 123-158.

Hopf, Ted. "The Promise of Constructivism in International Relations Theory," *International Security*, Vol. 23, No. 1 (Summer, 1998), pp. 171 – 200.

Horowitz, Michael; McDermott, Rose & Stam, Allan (2005). "Leader Age, Regime Type, and Violent International Relations," *Journal of Conflict Resolution,* Vol. 49, No. 5 (Oct 2005): 661-685.

Hug, Simon (1999). "Nonunitary Actors in Spatial Models: How Far is Far in Foreign Policy?" *The Journal of Conflict Resolution,* Vol. 43, No. 4 (Aug 1999): 479-500.

Isen, A. M. (1993). "Positive affect and decision making". In M. Lewis and J. M. Haviland (Eds.), *Handbook of Emotions* (pp. 261-277). New York: Guilford Press.

Janis, L.L. (1972). *Groupthink: Psychological Studies of Policy Decisions and Fiascos.* Baltimore, MD. Houghton Mifflin.

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

--------- (1978). "Cooperation Under the Security Dilemma," *World Politics*, Vol 30, No. 2 (January 1978): 167-214.

--------- (1989). "Rational Deterrence: Theory and Evidence," *World Politics,* Vol 41, No. 2 (1989): 183-207.

Johnstone, Ian (2005) "The Power of Interpretive Communities," in *Power in Global Governance*, ed. by Barnett, Michael and Duvall, Raymond. Cambridge, UK. Cambridge University Press. 185-204.

Joll, James & Martel, Gordon (2007). *The Origins of the First World War, 3rd ed.* New York, NY. Pearson-Longman.

Kahler, M. "Rationality in International Relations." *International Organizations.* Vol. 54, No. 4 (Autumn 1998), pp. 919-941.

Kahn, Herman (1960). *On Thermonuclear War.* Princeton, NJ. Princeton University Press.

Kahneman, D. and Tversky, A. (1979). "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica,* Vol. 47, No. 2 (March 1979): 263-289.

Kaiser, David (1983). "Germany and the Origins of the First World War," *Journal of Modern History.* Vol. 55, No. 3 (Sept 1983): 442-474.

Kanner, Michael (2004). "Framing and the Role of the Second Actor: An Application of Prospect Theory to Bargaining," *Political Psychology,* Vol. 25, No. 3: 213-239.

Kant, Immanuel (2001 [1795]). "To Perpetual Peace: A Philosophical Sketch," in Karen Mingst and Jack Snyder (eds.) *Essential Readings in World Politics*, pp: 14-17. New York, NY. W.W. Norton.

Kanwisher, Nancy. (1989). "Cognitive Heuristics and American Security Policy," *The Journal of Conflict Resolution,* Vol. 33, No. 4 (Dec 1989): 652-675.

Keegan, John (1998). *The First World War.* Toronto, Canada. Random House of Canada.

Keller, Jonathan (2005). "Leadership Style, Regime Type, and Foreign Policy Crisis Behaviour: A Contingent Monadic Peace?" *International Studies Quarterly,* Vol. 49: 205-231.

Kennedy, Paul (1984). "The First World War and the International Power System," *International Security,* Vol. 8, No. 1 (Summer 1984): 7-40.

Keohane, Robert (1982). "The Demand for International Regimes," *International Organizations,* Vol. 36, No. 2, International Regimes (Spring 1982): 325-355.

--------- (1984). *After Hegemony: Cooperation and Discord in the World Political Economy.* Princeton, NJ. Princeton University Press.

--------- (2002). *Power and Governance in a Partly Globalized World.* London, UK. Routledge.

Khong, Yuen Foong (1993). *Analogies at War.* Princeton, NJ. Princeton University Press.

Kissinger, Henry (1994). *Diplomacy.* New York, NY. Simon and Schuster.

Krasner, Stephen D. (1995) "Compromising Westphalia," *International Security*, Vol. 20, No. 3 (Winter, 1995-1996), 115-151.

Kreps, David; Milgrom, Paul; Roberts, John & Wilson, Robert (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, Vol. 27 (1982): 245-252.

Lasswell, H.D. (1930). *Psychopathology and Power.* New York, NY. Viking.

--------- (1935). *World Politics and Political Insecurity.* New York, NY. McGraw Hill.

--------- (1950). *Politics: Who Gets What, When, How*. New York, NY. Peter Smith.

Lau, R.R. and Redlawsk, D.P. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making," *American Journal of Political Science,* Vol. 45, No. 4 (Oct 2001): 951-957.

Lebow, Richard Ned. "Thucydides the Constructivist." *The American Political Science Review.* 95, 3, (Sept 2001), pp. 547-560.

Leites, Nathan (1951). *The Operational Code of the Politburo*. New York, NY. McGraw-Hill.

Lenin, V.I. (2001 [1939]). "Imperialism, The Highest Stage of Capitalism: A Popular Outline," in Karen Mingst and Jack Snyder (eds.) *Essential Readings in World Politics*, pp: 18-24. New York, NY. W.W. Norton.

Lerner, J. S., Gonzalez, R. M., Small, D. A., & Fischhoff, B. (2003). "Effects of fear and anger on perceived risks of terrorism: A national field experiment," *Psychological Science*, 14, 2: 144-150.

Lerner, J. S., Small, D. A., & Loewenstein (2004). "Heart strings and purse strings: Carryover effects of emotions on economic decisions," *Psychological Science*, 15, 5: 337-341.

Levy, Jack (1990). "Preferences, Constraints, and Choices in July 1914," *Interdisciplinary History,* Vol. 15, No. 3 (Winter 1990): 151-186.

Li, Richard P.Y. & Thompson, William R. (1978). "The Stochastic Process of Alliance Formation Behavior," *The American Political Science Review,* Vol. 72, No. 4 (Dec 1978): 1288-1303.

Lipson, Charles (2003). *Reliable Partners: How Democracies Have Made a Separate Peace.* Princeton, NJ. Princeton University Press.

Liu, Xinsheng (2006). *Modeling Bilateral International Relations: The Case of U.S.-China Interactions.* New York, NY. Palgrave.

Lodge, M. and Taber, C. S. (2000). Three steps toward a theory of motivated political reasoning. In A. Lupia, M. McCubbins, & S. Popkin (eds.), Elements of reason: Cognition, Choice, and the Bounds of Rationality, pp. 183–213. Cambridge: Cambridge University Press.

--------- (2005). "The Automaticity of Affect for Political Leaders, Groups, and Issues: An Experimental Test of the Hot Cognition Hypothesis," *Political Psychology,* Vol. 26, No. 3 (2005): 455-482.

Maier, Charles (1988). "Wargames: 1914-1919," *Journal of Interdisciplinary History,* Vol. 18, No. 4 (Spring 1988): 819-849.

March, James (1994). *A Primer on Decision Making: How Decisions Happen.* New York, NY. The Free Press.

Marcus, Gary (2003). *The Algebraic Mind.* Cambridge, MA. The MIT Press.

Marshall, Marty and Jaggers, Keith (2003). "The Polity IV Project: Political Regime Characteristics and Transitions, 1800-2003." Accessed online at http://www.cidcm.umd.edu/inscr/polity/.

Martin, Lisa (2001 [1992]) "Interest, Power, Multilateralism," in *International Institutions: An International Organization Reader,* ed. by Martin, Lisa and Simmons, Beth. Cambridge, MA. The MIT Press.

Marx, Karl (1977 [1859]). *A Contribution to the Critique of Political Economy.* Moscow, USSR. Progress Publishers.

Massie, Robert. (1991). *Dreadnought: Britain, Germany, and the Coming of the Great War.* New York, NY. Ballantine Books.

McDermott, Rose (2004a). “Editor’s Introduction,” *Political Psychology,* Vol. 25, No. 2: 147-162.

--------- (2004b). “Prospect Theory in Political Science: Gains and Losses From the First Decade,” *Political Psychology,* Vol. 25, No. 2: 289-312.

--------- (2004c). “The Feeling of Rationality: The Meaning of Neuroscientific Advances for Political Science,” *Perspectives on Politics,* Vol. 2, No. 4 (Dec 2004): 691-706.

Mearsheimer, John (1994/95). “The False Promise of International Institutions,” *International Security,* Vol. 19, No. 3 (Winter 1994-95): 5-49.

--------- (2001). *The Tragedy of Great Power Politics.* New York, NY. W.W. Norton.

Mercer, Jonathan (1996). *Reputation and International Politics.* Ithaca, NY. Cornell University Press.

Miller, Gregory D. (2003). “Hypotheses on Reputation: Alliance Choices and the Shadow of the Past,” *Security Studies*, Vol. 12, No. 3 (Spring 2003): 40-78.

Milner, Helen (1991). “The Assumption of Anarchy in International Relations: A Critique,” *Review of International Studies,* Vol. 17, No. 1 (Jan 1991): 67-85.

Modelski, George & Thompson, William (1989). “Long Cycles and Global Wars,” in Manus Midlarsky, (ed.) *Handbook of War Studies*, pp: 23-54. New York, NY. Unwin Hyman.

Mondak, J.J. and McCurley, C. (1994). “Cognitive Efficiency and the Congressional Vote: The Pyschology of Coattail Voting,” *Political Research Quarterly,* Vol. 47, No. 1 (Mar 1994): 151-175.

Moravcsik, A. “Taking Preferences Seriously: A Liberal Theory of International Politics.” *International Organization.* 51, 4, Autumn 1997, pp. 513-553.

Morgan, T. Clifton (1984). “A Spatial Model of Crisis Bargaining,” *International Studies Quarterly,* Vol. 28 (Dec 1984): 407-426.

-------------- (1990). “Issue Linkages in Crisis Bargaining,” *American Journal of Political Science,* Vol. 34, No. 2 (1990): 311-333.

Morgenthau, Hans J. and Thompson, Kenneth W. (rev). (1993) *Politics Among Nations: The Struggle for Power and Peace (brief ed.)* Boston, MA. McGraw Hill.

Morrow, James (1985). “A Continuous-Outcome Expected Utility Theory of War,” *Journal of Conflict Resolution,* Vol. 29 (1985): 473-502.

--------- (1994). *Game Theory For Political Scientists.* Princeton, NJ. Princeton University Press.

Muncaster, Robert & Zinnes, Dina (1983). "A model of inter-nation hostility dynamics and war," *Conflict Management and Peace Science*, Vol. 6, No. 2 (

Myerson, Roger (1991). *Game Theory: Analysis of Conflict.* Cambridge, MA. Harvard University Press.

North, M.J., Collier, N.T. & Vos, J.R. (2006). "Experiences Creating Three Implementations of the Repast Agent Modeling Toolkit," *ACM Transactions on Modeling and Computer Simulation,* Vol. 16, No. 1 (Jan 2006): 1-25.

Olson, Mancur (1971). *The Logic of Collective Action.* Cambridge, MA. Harvard University Press.

Onuf, Nicholas (1989). *The World of Our Making: Rules and Rule in Social Theory and International Relations.* Columbia, SC. University of South Carolina Press.

Organiski, A.F.K. (1958). *World Politics.* New York, NY. A.A. Knopf.

Oye, Kenneth (1986). "Explaining Cooperation Under Anarchy," in Kenneth Oye (ed.) *Cooperation Under Anarchy.* Princeton, NJ. Princeton University Press.

Pennington, N. & Hastie, R. (1997). "Explanation-based Decision Making: Effects of Memory Structure on Judgment," in Busemeyr, Hastie & Medin (eds.) *Decision Making, The Psychology of Learning and Motivation: Advances in Research and Theory*, pp. 33-81 San Diego, CA. Academic Press.

Plous, Scott (1993). *The Psychology of Judgment and Decision Making.* Philadelphia, PA. Temple University Press.

Post, Jerrold (2004). *Leaders and Their Followers in a Dangerous World: The Psychology of Political Behavior.* Ithaca, NY. Cornell University Press.

Press, Daryl (2005). *Calculating Credibility: How Leaders Assess Military Threats.* Ithaca, NY. Cornell University Press.

Putnam, Robert (1988). "Diplomacy and Domestic Politics: The Logic of Two-Level Games," *International Organizations,* Vol. 42, No. 3 (Summer 1988): 427-460.

Redlawsk, D.P. (2002). "Hot Cognition or Cool Consideration? Testing the Effects of Motivated Reasoning on Political Decision Making," *The Journal of Politics,* Vol. 64, No. 4 (Nov 2002): 1021-1044.

Redmondino, Marco (2006). "Agent Based Simulation for Management and Economics," Unpublished paper. Accessed online at http://www.cs.unibo.it/~cianca/wwwpages/labspo/Remondino.pdf.

Reus-Smit, Christian (2001). "Constructivism," in Burchill, Scott (ed.) *Theories of International Relations, 2nd ed.* New York, NY. Palgrave. pp. 209-230.

Richardson, Lewis F. (1960). *Arms and Insecurity.* Chicago, IL. Quadrangle Publishing.

Roe, Robert; Busemeyer, Jerome & Townsend, James (2001). "Multialternative Decision Field Theory: A Dynamic Connectionist Model of Decision Making," *Psychological Review,* Vol. 108, No. 2 (2001): 370-392.

Ruggie, John, (1983). "Continuity and Transformation in the World Polity: Towards a Neorealist Synthesis", *World Politics*, Vol. XXXV, January 1983, pp. 261-285.

Samuelson, Douglas (2005). "Agents of Change," *OR/MS Today,* Vol. 32, No. 2 (Feb 2005)

Samuelson, Douglas & Macal, Charles (2006). "Agent-Based Simulation Comes of Age," *OR/MS Today*, Vol. 33, No. 4 (August 2006): 34-38.

Schelling, Thomas. (1960). *The Strategy of Conflict.* New York, NY. Oxford University Press

Schweller, Randall. (1994). "Bandwagoning for Profit: Bringing the Revisionist State Back In," *International Security,* Vol. 19, No. 1: 72-107.

Shapiro, Michael & Bonham, G. Matthew (1973). "Cognitive Processes and Foreign Policy Decision-Making," *International Studies Quarterly,* Vol. 17, No. 2 (Jun 1973): 147-174.

Simon, Herbert A. (1997 [1945]). *Administrative Behavior*, 4th ed. New York, NY. Free Press.

Simon, Herbert and Newell, Alan (1972). *Human Problem Solving.* Englewood Cliffs, NJ. Prentice Hall.

Singer, J. David. (1987). "Reconstructing the Correlates of War Dataset on Material Capabilities of States, 1816-1985" *International Interactions*, Vol. 14: 115-32.

Small, Melvin, and J. David Singer. 1969. "Formal Alliances, 1815-1965: An Extension of the Basic Data." *Journal of Peace Research,* Vol. 6:257-282.

Snyder, Glenn (1960). "Deterrence and Power," *The Journal of Conflict Resolution,* Vol. 4, No. 2 (Jun, 1960): 163-178.

--------- (1994). "The Security Dilemma in Alliance Politics," *World Politics*, Vol. 36, No. 4 (July 1984): 461-495.

Snyder, Jack (1978). "Rationality on the Brink: The Role of Cognitive Processes in Failures of Deterrence," *World Politics,* Vol. 30 (Apr 1978): 345-365.

Steger, Manfred (2008). *The Rise of the Global Imaginary: Political Ideologies from the French Revolution to the Global War on Terror.* New York, NY. Oxford University Press.

Stinnett, Douglas M., Jaroslav Tir, Philip Schafer, Paul F. Diehl, and Charles Gochman. 2002. "The Correlates of War Project Direct Contiguity Data, Version 3." *Conlict Mangagement and Peace Science* Vol. 19. No. 2:58-66.

Stoessinger, John (1978). *Why Nations Go To War.* New York, NY. St. Martins Press

Strachan, Hew (2004). *The First World War.* New York, NY. Viking Press.

Sun Tzu (1983). *The Art of War.* Samuel Griffith, trans. Oxford, UK. Oxford University Press.

Taber, Charles S. (1992). "POLI: An Expert System Model of U.S. Foreign Policy Belief Systems," *The American Political Science Review,* Vol. 86, No. 4 (Dec 1992): 888-904.

Taber, Charles & Timpone, Richard (1996). "Beyond Simplicity: Focused Realism and Computational Modeling in International Relations," *Mershon International Studies Review,* Vol. 40, No. 1 (Apr 1996): 41-79.

Taliaferro, Jeffrey (2004). "Power Politics and the Balance of Risk: Hypotheses on Great Power Intervention in the Periphery," *Political Psychology,* Vol 25, No. 3: 177-211.

Tetlock, Philip (1985). "Accountability: A Social Check on the Fundamental Attribution Error," *Social Psychology Quarterly*, Vol. 48, No. 3 (Sep 1985): 227-236.

Thomas, Rick; Dougherty, Michael; Springer, Amber & Harbison, J.I. (2008) "Diagnostic Hypothesis Generation and Human Judgment". *Psychological Review,* Vol. 115, No. 1: 155-185.

Thucydides (2001 [1951]). "The Melian Dialogue," in Karen Mingst and Jack Snyder (eds.) *Essential Readings in World Politics*, pp: 12-13. New York, NY. W.W. Norton.

Trachtenberg, Marc (1990). "The Meaning of Mobilization in 1914," *International Security,* Vol. 15, No. 3 (Winter 1990-91): 120-150.

Tuchman, Barbara (1962). *The Guns of August.* New York, NY. Random House.

Tversky, Amos & Kahneman, Daniel (2000a [1974]). "Judgment under uncertainty: Heuristics and Biases," in T. Connolly, H. R. Arkes, and K. R. Hammond (eds.). *Judgment and Decision Making: An Interdisciplinary Reader* (Second Edition). 35-52.

--------- (2000b). "Rational Choice and the Framing of Decisions," in Kahneman & Tversky (eds.) *Choices, Values, and Frames.* Cambridge, NY. Cambridge University Press. pp. 209-223.

--------- (2000c). "Advances in Prospect Theory: Cumulative Representation of Uncertainty," in Kahneman & Tversky (eds.) *Choices, Values, and Frames.* Cambridge, NY. Cambridge University Press. pp. 44-66.

Van Evera, Stephen (1984). "The Cult of the Offensive and the Origins of the First World War," *International Security,* Vol. 9 (Summer 1984).

Vasquez, John (1998). *The Power of Power Politics.* Cambridge, UK. Cambridge University Press.

Verba, Sidney (1961). "Assumptions of Rationality and Non-rationality in Models of the International System," in Knor, Klaus & Verba, Sidney (eds.) *The International System: Theoretical Essays.* Princeton, NJ. Princeton University Press. pp. 93-117.

Vertzberger, Y.I. (1986). "Foreign Policy Decisionmakers As Practical-Intuitive Historians: Applied History and Its Shortcomings," *International Studies Quarterly,* Vol. 30, No. 2 (Jun 1986): 223-247.

--------- (1990). *The World in Their Minds: Information Processing, Cognition, and Perception in Foreign Policy Decisionmaking.* Stanford, CA. Stanford University Press.

Wallerstein, Immanuel (1974). "The Rise and Future Demise of the World Capitalist System: Concepts for Comparative Analysis," in Karen Mingst and Jack Snyder (eds.) *Essential Readings in World Politics*, pp: 37-145. New York, NY. W.W. Norton.

Walker, Stephen (1990). "The Evolution of Operational Code Analysis," *Political Psychology,* Vol. 11, No. 2 (June 1990): 403-418.

Walt, Stephen (1987). *The Origin of Alliances.* Ithaca, NY. Cornell University Press

Waltz, Kenneth N. (1959). *Man, the State, and War: A Theoretical Analysis.* New York, NY. Columbia University Press.

--------- (1979). *Theory of International Politics*. Reading, MA. Addison-Wesley Publishing.

--------- (2000). “Structural Realism after the Cold War,” *International Security,* Vol. 25, No. 1 (Summer 2000): 5-41.

Weber, Elke; Böckenholt, Ulf; Hilton, Denis & Wallace, Brian (1993) “Determinants of Diagnostic Hypothesis Generation: Effects of Information, Base Rates, and Experience,” *Journal of Experimental Psychology,* Vol. 19, No. 3: 1151-1164.

Wendt, Alexander (1987). “The Agent-Structure Problem in International Relations Theory,” *International Organization,* Vol. 41, No. 3 (Summer 1987): 335-370.

--------- (1992). “Anarchy is what States Make of it: The Social Construction of Power Politics.” *International Organization*, Vol. 46, No. 2. (Spring, 1992), pp. 391-425.

--------- (1999). *Social Theory of International Politics.* Cambridge, UK. Cambridge University Press.

Williamson, Samuel (1988). “The Origins of World War I,” *Journal of Interdisciplinary History,* Vol. 18, No. 4 (Spring 1988): 795-818.

Zinnes, Dina and Gillespie, John, eds. (1976). *Mathematical Models in International Relations*. New York, NY. Praeger Publications.

Zinnes, Dina & Muncaster, Robert (1984). “The Dynamics of Hostile Activity and the Prediction of War,” *The Journal of Conflict Resolution*, Vol. 28, No. 2 (Jun 1984): 187-229.