# Agents in Conflict: Comparative Agent-Based Modeling of International Crises and Conflicts

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

David P. Masad Bachelor of Science The University of Chicago, 2009

Director: Dr. Robert Axtell, Professor Department of Computational and Data Sciences

> Spring 2016 George Mason University Fairfax, VA

Copyright © 2016 by David P. Masad All Rights Reserved

## Dedication

To Dani; to my family; and in memory of my father.

## Acknowledgments

This dissertation was made possible by the support, advice and guidance of my committee members, which began long before the dissertation itself. It also would not have been possible without the assistance and friendship of Karen Underwood, the central-most node in the CSS network.

I am grateful to George Mason University itself, which funded this dissertation via a Presidential Scholarship, a summer research grant, and a dissertation completion grant.

I am indebted to friends and colleagues who provided ideas, feedback, and information throughout the development of this dissertation. In particular, several practitioners provided valuable insight from their own experience with the expected-utility modeling methodology which helped point me in what I hope was the right direction. Errors, of course, are mine alone.

Finally, a share of the blame for the existence of this dissertation must go to Sarah Wise, who helped get me into this whole computational social science thing in the first place.

## Table of Contents

				Page
Lis	t of T	ables		. vii
Lis	t of F	igures		. viii
Ab	brevia	ations		. x
Ab	stract			. xi
1	Intr	oductio	n	. 1
	1.1	Backg	round	. 1
	1.2	Curren	nt Methods in Conflict Modeling	. 4
	1.3	Resear	ch Contribution	. 11
	1.4	Data S	Sources	. 16
	1.5	Organ	ization of the Dissertation	. 19
2	Lear	rning ai	nd War	. 22
	2.1	Backg	round	. 22
	2.2	Model	Description	. 28
		2.2.1	World Sub-Model	. 29
		2.2.2	Interaction Sub-Model	. 30
		2.2.3	Agent Sub-Model	. 31
	2.3	Model	Variants	. 31
		2.3.1	Small Crisis Game	. 32
		2.3.2	International Interaction Game	. 34
	2.4	Decisio	onmaking Models	. 41
		2.4.1	Rational Model	. 43
		2.4.2	Weighted-Attraction Agents	. 43
		2.4.3	Pure Reinforcement Learning	. 46
		2.4.4	Case-Based Learner	. 47
	2.5	Experi	iments and Results	. 50
		2.5.1	Small Crisis Model	. 50
		2.5.2	International Interaction Game	. 63
	2.6	Discus	sion	. 73
<b>3</b>	BDI	M's EU	M as ABM	. 78

	3.1	Introd	uction $\ldots \ldots .$	
	3.2	Model	Description	
		3.2.1	Overview	
		3.2.2	Details	
	3.3	Model	Validation and Verification	
		3.3.1	Reproduction of Prior Results	
		3.3.2	Deep Dive	
	3.4	Model	Discussion	
		3.4.1	Expected Utility 103	
		3.4.2	Conflict and Coercion	
		3.4.3	Risk Propensity	
		3.4.4	Offer Selection	
		3.4.5	Stochasticity and Uncertainty	
	3.5	Sensiti	ivity Analysis	
	3.6	Inform	nation Extraction $\ldots \ldots \ldots$	
	3.7	Compa	aring Model to Event Data	
		3.7.1	Data Sources	
		3.7.2	Model Variants	
		3.7.3	Experiments and Results 124	
		3.7.4	Discussion	
	3.8	Conclu	usions $\ldots$ $\ldots$ $\ldots$ $\ldots$ $128$	
4	Mod	deling the	he Cold War	
	4.1	Introd	uction $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $\ldots$ $133$	
	4.2	Model	ing Methodology $\ldots \ldots 135$	
		4.2.1	Input Data	
		4.2.2	Model Instantiations	
		4.2.3	Output Data 143	
	4.3	Result	s143	
		4.3.1	Overview	
		4.3.2	Predicted Conflicts	
		4.3.3	Predicted Alliances	
	4.4	Discus	sion $\ldots \ldots 157$	
5	Con	clusion	s	
Bib	Bibliography			

## List of Tables

Table		Page
2.1	Small Crisis Model Utilities	33
2.2	Small Crisis Agent Properties	34
2.3	Small Crisis (Static) – Logistic Regressions	57
2.4	Small Crisis (Dynamic) – Logistic Regressions	61
2.5	IIG Model Equilibrium – Logistic Regressions	70
2.6	Reinforcement Learning Model – Logistic Regressions	71
2.7	Case-Based Learning Model - Logistic Regressions	72
3.1	Agent Variables	84
3.2	Example Model – Agent Data	99
3.3	$Example \ Model - Expected \ Utilities \ \ \ldots $	100
3.4	Accepted Offers	100
3.5	Capability Contribution – Belgium vs Netherlands $\ldots \ldots \ldots \ldots \ldots$	101
3.6	Experiment Inputs	113
3.7	Model Parameters	120
3.8	Agent Data	122
3.9	ICEWS Conflict Events – Linear Regressions	125
3.10	ICEWS Conflict Events – Logistic Regressions	125
4.1	Original Inputs	136
4.2	Alliance Coding	138
4.3	Updated Inputs	140
4.4	Experiment Summary	142
4.5	Militarized Interstate Disputes – Linear Regressions	151
4.6	Militarized Interstate Disputes – Logistic Regressions	151
4.7	Top Predicted Conflict Dyads	152
4.8	Alliance Type – Multinomial Choice Logistic Regression	155
4.9	Alliance Network – Exponential Random Graph Models	156

## List of Figures

Figure		Page
1.1	Modeling Methodology	12
2.1	Model Architecture	29
2.2	Small Crisis Game Tree	33
2.3	International Interaction Game Tree	35
2.4	EUGene Data by Year	41
2.5	Small Crisis (Static) – Traces	53
2.6	Small Crisis (Static) – Quality Distribution	54
2.7	Small Crisis (Static) – CB and RL Differences	54
2.8	Small Crisis (Static) – Quality by Learning Rate	55
2.9	Small Crisis (Static) – Logistic Regression Coefficient Plots $\ldots \ldots \ldots$	56
2.10	Small Crisis (Dynamic) – Traces	58
2.11	Small Crisis (Dynamic) – Quality Distribution	59
2.12	Small Crisis (Dynamic) – CB and RL Differences	59
2.13	Small Crisis (Dynamic) – Quality by Learning Rate	60
2.14	Small Crisis (Dynamic) – Logistic Regression Coefficient Plots	62
2.15	Fraction of Interactions with Equilibrium Outcome	66
2.16	Fraction of Interactions with Correctly-Predicted Outcome	67
2.17	Logistic Regressions Coefficient Plots	69
3.1	Model Architecture	. 87
3.2	Incoming Offer Categories	92
3.3	Original and Replication Model Comparison	95
3.4	Octant Comparison	96
3.5	Chinese Democratization Model Medians	98
3.6	Histograms of Median Positions	115
3.7	Experiment 1 and 2 Comparison	115
3.8	Model-Generated vs ICEWS Conflict Events	126
4.1	Original Inputs	137

4.2	Tau-b Similarities of Actors to United States and Soviet Union	138
4.3	Updated Inputs	139
4.4	Original Model Cumulative Distribution	144
4.5	New Model Cumulative Distribution	145
4.6	Median Position Traces	146
4.7	Example Model Traces – Median Positions	148
4.8	Example Model Traces – Agent Positions	149

## Abbreviations

ABM	Agent-based model
BDM	Bueno De Mesquita
CB	Case-Based
CINC	Composite Index of National Capability
COW	Correlates of War
ERGM	Exponential Random Graph Model
EUM	Expected Utility Model
EWA	Experience-Weighted Attraction
ICEWS	Integrated Crisis Early Warning System
IIG	International Interaction Game
MID	Militarized Interstate Dispute
NATO	North Atlantic Treaty Organization
NMC	National Material Capabilities
ODD	Overview, Design concepts and Details
PITF	Political Instability Task Force
PSTK	Power Structure Toolkit
RL	Reinforcement Learning
US/USA	United States of America
USSR	Union of Soviet Socialist Republics (Soviet
	Union)

### Abstract

#### AGENTS IN CONFLICT: COMPARATIVE AGENT-BASED MODELING OF INTER-NATIONAL CRISES AND CONFLICTS

David P. Masad, PhD

George Mason University, 2016

Dissertation Director: Dr. Robert Axtell

Inter-state conflicts are a key area of study in international relations, and have been approached with a variety of techniques, from case studies of individual conflicts, to formal analysis of abstract models and statistical investigations of all such conflicts. In particular, there are a variety of theories as to how states make decisions in the face of conflicts – such as when to threaten force, when to follow through, and when to capitulate to an opponent's demand. Some scholars have argued that states may be viewed as rational decisionmakers, while others emphasize the role of psychological biases affecting individual leaders. Decisionmaking is challenging to study in part because of its complexity: the decisionmakers may not just be individuals but organizations, following internal procedures and reflecting institutional memory. Furthermore, the decisions are often believed to be strategic, reflecting the decisionmakers' anticipation of multiple other actors' potential responses to each possible decision.

In this dissertation, I demonstrate that agent-based models (ABMs) provide a powerful tool to address this complexity, and advance their use as a bridge between different methodologies. Agents in ABMs, representing countries, can be endowed with a variety of internal decisionmaking models which can operationalize a variety of theories drawn from case studies, psychological experiments or game-theoretic analysis. The specific decision model agents utilize may be changed without altering the sub-models governing how the agents interact with one another. This allows us to simulate the same overall interactions utilizing different decisionmaking theories and observe how the outcomes differ. Furthermore, if these interactions correspond to real-world events, we may directly see how much explanatory or predictive power the outputs of the model variants provide. If one variant's outputs correspond closer to the empirical data, it provides evidence supporting that variant's underlying theory.

I implement two agent-based models, extending well-established prior models of international conflict: the International Interaction Game (Bueno de Mesquita and Lalman, 1992) and the Expected Utility Model (Bueno de Mesquita, 2002). For each, I start with their original agent decisionmaking models, and develop several variants grounded in relevant theories. I then instantiate the models with historic, empirically-derived data and run them forward to generate sets of simulated outcomes, which I compare to empirical data on the relevant time periods. I find that non-rational models of decisionmaking in the International Interaction Game provide similar explanatory power to the purely rational model, and yield rich satisficing behavior absent in the original model. I also find that the Expected Utility Model variant implementing a Schelling (1966)-inspired model of coercion yields richer dynamics and greater explanatory power than the original model.

In addition to providing evidence in support of particular theories and hypotheses, this work demonstrates the power of the comparative modeling methodology in studying international conflict. Future work will involve adding more statistical controls to the model output analysis, comparative analysis between the outputs of the two overall models, and extension of the decisionmaking models for each. The same methodology may also be expanded to other formal and computational models of international relations, and social science more broadly.

### Chapter 1: Introduction

### 1.1 Background

Political conflicts are complex. This is the expected opening in a dissertation such as this, but it is nevertheless true. They involve cooperation and competition between multiple actors who differ from one another in their capabilities, their objectives, and even their decisionmaking processes. Furthermore, these actors themselves are often complex – they are generally political entities (primarily states, though increasingly non-state groups as well) which are ultimately composed not only of individuals but of procedures, institutional memory, and other emergent characteristics. To paraphrase an old joke, it's complexity all the way down.

Much of the study of political science and international relations seeks, naturally enough, to reduce this complexity, and does so with a wide variety of methodologies and tools. General theories of how actors make decisions and interact, derived from and applied to case studies of particular events, are as old as the discipline itself. More recently, such studies have encompassed not just state-level power, security, and institutions, but the incentives and behaviors of sub- and non-state actors, and even the psychological biases of individual decisionmakers. Quantitative and statistical methods have become the mainstays of the discipline, offering powerful tools to identify regularities and test theories in datasets of actor properties, relationships, and events. Game theory grew in part out of Cold War-era security analysis and policymaking. It has provided a language that can formally describe and analyze strategic interactions, and yielded powerful concepts such as the notion of a Nash equilibrium, and stylized, insight-generating models such as the Prisoners' Dilemma.

These lines of research are often somewhat disconnected from one another. Many grand

theories are extremely general, and can only be explored via detailed case studies, or statistically tested one narrow implication at a time. Similarly, psychological theories require notional lab experiments or detailed archival sources on the beliefs and actions of key individuals in a particular crisis. Game-theoretic models are generally relatively abstract, and rarely intended to be applied to data to generate specific testable predictions. Statistical models have gained prominence in part because they can be used to operationalize and test hypotheses developed from various theories. These tests require appropriate and accurate empirical data, which is often difficult to collect, and many statistical tools do not account for the interdependence between cases in the dataset. Statistical models are generally capable of accounting for some top-level outcome variable, but rarely attempt to capture the interactions and decisions by which that outcome was reached; they provide outcome validity, but not necessarily process validity.

Comparative analysis is an important tool, frequently applied across different methodologies. Tables reporting the results of multiple regressions side by side are a common sight in quantitative studies: by comparing the fit and coefficients of several statistical models, we can evaluate which one has more explanatory power, and whether hypothesized relationships in fact occur in the data. Case studies may also be used to compare multiple theories, examining a single case in detail and identifying which events and processes can be explained by (or contradict) each, or reviewing sets of cases together and locating meaningful commonalities. Many theories and models have limited explanatory power on their own, making it hard to assess them without a basis for comparison. This is particularly true for more general theories that abstract away from the actor- and situation-specific details that often drive a great deal of interactions. By comparing two or more theories side by side, we can assess which one does a better job at explaining the cases or data in question.

Agent-based models (ABMs) are a powerful tool for studying complex systems in general (Miller and Page, 2007), and appear particularly suitable to apply to international relations (Cederman, 1997). By definition, ABMs consist of independent actors (agents) interacting with one another and with their environment. The agents are programmed with particular behaviors, which may include, for example, extremely simple *if-then* rules, operationalizations of particular qualitative theories, or complicated optimization methods. Since models are often implemented primarily at the agent level, they can be used to identify emergent, aggregate phenomena arising from the simulated lower-level interactions. Though ABMs have never entered the mainstream of international relations research, there have been several attempts to build and study such models. Many of these attempts (e.g. Axelrod, 1997; Cederman, 1997; Min, 2002) have involved relatively abstract, notional systems. Such models attempt to capture a few key aspects of the behavior of countries, and use these simulated states to grow an international system 'from the bottom up,' exploring how well the hypothesized behaviors lead to the emergence of properties observed in reality. Another set of models explicitly attempt to model real-world actors, and are intended for forecasting and policy support (e.g. Abdollahian et al., 2006; Taylor et al., 2008). These models, in contrast typically do not attempt to test specific theories, though it is hard to determine for certain as they are almost universally proprietary.

This dissertation will attempt to demonstrate that agent-based modeling may be used in a different way, and can serve as a bridge between several of the other methodologies and paradigms (Axelrod, 2006). Such models need not be confined to notional worlds, but can be applied to real-world interactions as well. Empirical application, in fact, offers a way to directly compare the theories of interaction and decisionmaking that the models are based on. In order to demonstrate this, I reimplement two prior formal models of international conflict as agent-based models, arguing that their structure is already amenable to such conversion. Elements of the models are modified, endowing the agents with alternative internal models of decisionmaking; I then instantiate different model variants with the same input data, in order to directly compare their behaviors and outputs. By comparing these outputs' relationships to empirical, historic data, I directly determine which model variant best explains and predicts the observed data.

### 1.2 Current Methods in Conflict Modeling

#### Case Studies

Histories have described the decisionmaking process of key actors in international interactions since long before history was an established academic discipline, and international relations scholarship has long drawn on these descriptions to generate general theories. Modern scholarship has extended this technique into the case study, with the researcher methodically analyzing primary and secondary descriptions of an incident or set of incidents. Such analysis may involve attempts to tease apart commonalities in order to build a new theory; however, more often it involves applying a particular prior theory to the cases in question, examining where the theory can or cannot provide an adequate explanation for the observed events (George and Bennett, 2005). In "The Essence of Decision," Allison and Zelikow (1999) provide a key demonstration of this methodology, laying out three theories of political decisionmaking and analyzing the Cuban Missile Crisis through the lens of each in turn, highlighting decisions that can adequately be explained by one or more of the theories, and (more importantly), decisions which appear to contradict one theory but can be explained by another. Similarly, Jervis (1976) lays out a series of psychological models expected to drive the behavior of individuals and small groups of decisionmakers, and presents historic evidence that such theories indeed explain events observed across multiple past crises. Kaufmann (1994) focuses on one particular case, and uses it to test a similar psychological model against the alternative theory that the decisionmakers were acting rationally, demonstrating that the proposed biases explain divergences between actors that have no purely rational explanation.

There can be little doubt that case studies provide a vital tool for understanding how decisions are actually made and identifying the strengths and weaknesses of existing theories. However, case studies are limited by the need for detailed historical or archival information. Allison and Zelikow (1999) rely on transcripts of discussions among the American leadership during the Cuban Missile Crisis and detailed records of American and Soviet military deployments, while Kaufmann (1994) relies on letters between the key leaders to identify when, and in response to what, each changed his opinion. This requirement makes it difficult to falsify theoretical predictions: for example, an earlier edition of "Essence of Decision" (Allison, 1971) made theory-driven predictions as to the structure and process of the Soviet decisionmaking, which were only proven to be incorrect decades later, when the fall of the Soviet Union made the relevant historic documents available. Furthermore, case studies require 'cases' – for our purposes, crises or conflicts which come to the attention of the analyst. However, a theory may correctly predict many instances where no conflict occurs, and if indeed no conflict occurs then no case for examination will be generated. Achen and Snidal (1989) further argue that case studies alone are insufficient to truly test theories, as they lend themselves too readily to highlighting case-specific exceptions and cannot adequately measure the power of a theory across multiple cases. Such rigorous comparison, they argue, can only be done using statistical analysis.

#### Statistical Models

Statistical and econometric modeling have become the most common methodology in international relations research (Colgan, 2016). In general, researchers will develop a qualitative theory or take one from prior literature, and operationalize it as an econometric model attempting to explain a dependent variable (e.g. the outbreak of war) in terms of several independent variables (such as the system of government and military strength of the belligerents). The researchers then fit the model to an appropriate dataset. If the model fits the data well (e.g. high  $R^2$ , low mean squared error), and the coefficients are statistically significant and have the hypothesized signs and magnitudes, this is taken as evidence of the theory's validity. Alternatively, the coefficients that are determined to be significant can be used to develop new theories which account for the observed relationship. A model can be further tested by applying it outside of the sample used to fit it, either by withholding a subset of the data for testing, or (better yet) applying it to future cases outside the dataset altogether.

An important reason that statistical analysis has become a core research tool is that "it basically works," (Schrodt, 2004). There are many, many examples of statistical models applied to international relations in general, and to conflict in particular. For example, Bremer (1992) tests numerous factors hypothesized to make wars between dyads of states more likely, showing that many factors are less important than previously thought; while Hagan (1994) estimates how states' domestic politics affect their propensity to engage in wars. Wayman and Diehl (1994) review multiple statistical tests of the qualitative predictions of Realist theories, arguing that they are insufficient to explain much of the international system. Such models can be applied specifically to forecasting. The work of Goldstone et al. (2005) with the Political Instability Task Force (PITF) can generate accurate predictions as to which countries will experience internal instability over a two-year period, and is used for government planning and policymaking. Similarly, Ward et al. (2013b) uses statistical methods to generate regularly-updated forecasts for the outbreaks of civil wars. More recently, such models have been applied not just to the outbreak of conflicts but to the counts of specific events. For example, Yonamine (2013) applies statistical models originating in finance to predict the intensity of conflict in Afghanistan, while Brandt, Freeman, and Schrodt (2011) propose a general time-series model for the number of individual events within any given ongoing conflict, with an initial focus on Israeli-Palestinian interactions.

While the statistical approach may indeed provide what Schrodt (2004) refers to as 'outcome validity', by and large it is not intended to provide 'process validity'. Goldstone et al. (2005), for example, can predict where instability will occur, but not which particular actors will initiate or be involved in it. Furthermore, one of the variables with the greatest predictive power for internal political instability is infant mortality (Goldstone et al., 2005). Yet high infant mortality is not directly the cause of most instability, nor is it exogenous to the system as a whole (Marshall, 2008). Similarly, though many statistical analyses (e.g. Gartzke, 2007; Maoz and Abdolali, 1989; Thompson and Tucker, 1997) have confirmed that democracies are less likely to go to war with one another (so-called democratic peace theory), they do not offer a consensus as to why this is the case. These difficulties point to another challenge facing the use of statistical techniques. Political interactions are generally believed to be strategic: actors on all sides can anticipate the other actors' likely responses to possible actions they may take, and make their own decisions accordingly. These strategic considerations can lead to censoring in the data we observe in the world (Smith, 1999), and to non-linearities in the relationship between input variables and expected outcome (Signorino, 1999). To provide a simplified example: if countries are highly likely enter wars when the balance of forces is above a certain threshold, and highly unlikely to otherwise, the population of observed wars would only include cases where the actors' force ratio falls into the range between these thresholds. However, a logistic regression predicting the probability of war from both sides' forces would not detect those sharp break-points.

#### Game Theory

In order to overcome the challenge of analyzing data generated from strategic interactions, Signorino (1999) and Smith (1999) independently proposed statistical models explicitly informed by game theory. Game theory is the mathematical study of interactions (games) between two or more 'players' whose payoffs depend not only on their own choices but those of the other players as well. Much of modern game theory was developed explicitly for the study of international conflict (Myerson, 2009), particularly by researchers at the RAND Corporation (Gates and Humes, 1997). Game-theoretic or formal models are widely used in political science, where they are used to represent conflict (e.g. Powell, 2006), negotiation (e.g. Brams, 2003), elections (e.g. Coughlin and Nitzan, 1981) and more. Gintis and Helbing (2015) have even recently proposed to apply game theory as a foundation to an attempt to derive a mathematical model for all of human society.

Statistical models explicitly incorporating game-theoretic elements are still relatively rare. Many formal models are not intended to represent specific historically-situated events or interactions; they are notional and abstract, meant to capture common characteristics encountered in multiple cases (Snidal, 1985). By solving the models analytically and identifying and parameterizing their equilibria, researchers can identify general theories of strategic behavior and interaction. There have also been several attempts to apply gametheoretic models to empirical data in addition to the statistical models referenced above. Bueno de Mesquita and Lalman (1992) apply an extensive-form model of crisis escalation to historic dyads, an effort extended substantially by Bennett and Stam (2000a). Bueno de Mesquita (1994, 1997, 2002) presents a game theory based model of negotiation and coercion which is explicitly intended as a forecasting tool. These models will be the focus of this dissertation, and so will be described at length later. Additional examples include Tsebelis (1988) proposing a multilateral formal model to explain the outcome of French elections, and Metternich et al. (2013) using equilibria in a network-based model to explain the intensity of opposition group activity in Thailand.

A fundamental assumption of game theory is that the players are rational, which here has a formal definition: that they take actions in order to maximize their utility, aware of and taking into account the other players' rationality. This leads to equilibria, the strategies leading to the best result players can achieve while assuming that all other players are playing their equilibrium strategy as well. The assumption of rationality allows the modeler to analyze and solve the game, as it implies that the actors are conducting the same analysis themselves, and are acting accordingly.

The rationality assumption is, not surprisingly, the source of much of the criticism of game theory. Behavioral economics has shown that human players' deviations from perfect rationality are not merely a white-noise error term, but systemically driven by psychological (Tversky and Kahneman, 1981) and cultural (Henrich et al., 2005) factors. Recently, behavioral game theory has begun to incorporate such insights into game-theoretic models (Camerer, 2003) in order to model the behavior of players who are not formally rational. Much like individual humans, organizations may not be rational either, and their decision-making is characterized by their own heuristics (March and Simon, 1993) and systemic biases (Shapira, 2002). While defenders of the rationality assumption such as Friedman (1953) argue that actors need not be procedurally rational (Simon, 1976) for the equilibrium outcomes to serve as useful predictors, both statistical analyses (Bennett and Stam, 2000a) and detailed case studies (Allison and Zelikow, 1999; Kaufmann, 1994) suggest that real decisionmaking deviates systemically from the formal predictions. In laying out an agenda for the future of game theory research, Fudenberg (2010) suggests a relaxation of this assumption is necessary to allow models to capture the full range of observed behaviors.

#### **Computational Methods**

Computational methods are not a well-defined family of methodologies like the ones discussed in the previous sections. I use the phrase very broadly, to mean any method which in practice relies on a computer to implement, and does not fall into the categories above. In practice, there are roughly three broad subcategories of computational methods: dynamical systems models, machine learning, and agent-based models.

Dynamical systems models represent components of interest as stocks and flows connected to one another and often exhibiting feedback loops (Gilbert and Troitzsch, 2005). While many such models can ultimately be represented as systems of differential equations, in practice their complexity often requires them to be approximated numerically. Early dynamical systems models relevant to international relations include the Lanchester (1916) models of ancient and modern warfare, and the Richardson arms race model (Rapoport, 1957), both of which emerged from the First World War. The introduction of computers, and the growth of interest in systems theory, led to computational systems dynamics models of international relations, such as the Globus Model (Bremer, 1987). The Richardson model has been widely tested, and found to be insufficient to explain the dynamics of arms acquisition even between pairs of well-established rivals (Dunne, Nikolaidou, and Smith, 1999). This approach in general appears to have fallen out of favor in recent years.

Machine learning is, essentially, an extension of statistical methods to leverage expanded computational resources, and this is largely how it has been used in political science. Goldstone et al. (2005) experimented with using methods such as neural networks or Random Forests (Breiman, 2001) before returning to the logistic regression. Schrodt (1997) used a hidden Markov model to predict the dynamics of conflict in the Middle East, and presents several other approaches (decision trees, neural networks and sequence matching) as prediction methods as well (Schrodt, 2004). While these approaches allow for more sophisticated relationships between predictors and outcomes, they suffer from many of the same problems as the more traditional statistical approaches, particularly the 'process' versus 'outcome' validity distinction. Furthermore, these approaches often trade accuracy for interpretability, resulting in models with strong predictive power but which cannot be directly mapped to a qualitative theory.

Finally, and most relevantly for this dissertation, are agent-based models (ABMs). In these models, each actor is represented by a discrete agent, and the agents interact with one another within some sort of environment. In general, ABMs are operationalizations of a particular qualitative theory from which the agent behaviors are derived. The agent behavior may consist extremely simple rules such as tit-for-tat (Axelrod, 1980; Hudson, Schrodt, and Whitmer, 2004); more complicated heuristics (Cederman, 1997); artificial intelligence-derived planning algorithms (Taylor et al., 2008); or adaptive behavior (Hoffmann, 2005; Rouleau, 2011). The level of abstraction of these ABMs also varies widely. Many are 'artificial societies' (Epstein and Axtell, 1996), testing how simple assumptions and behaviors lead to the emergence of key features of the international system. For example, Cederman (1997) models an international system as consisting of political actors on a square grid, using simple theory-derived heuristics to decide whether to go to war with their neighbors, and demonstrates how this leads to state and empire-like entities, and later (Cederman, 2003) under what conditions the power-law distribution of conflict scales (Richardson, 1948) arises. Other models are intended explicitly for forecasting and policy analysis. For example, Taylor et al. (2008) describe a generic (and non-public) model of political conflict, called the Power Structure Toolkit (PSTK), meant to allow subject-matter experts to generate and test scenarios with specific agents endowed with heterogeneous capabilities and goals. The objective of this model is as a decision support tool, rather than a test of any particular theories or hypotheses.

Abstract ABMs serve similar purposes to formal models, presenting relatively sparse theory implementations in order to yield general insights. By relaxing the requirements for perfect rationality and analytic tractability, they can describe richer interactions with a greater number of actors, and explore the models' properties and the emergent phenomena they may give rise to. Such ABMs must qualitatively ground their assumptions, and require additional analysis due to the number of moving parts to ensure that their results are robust, and not merely artifacts of the model. Something as simple as the agent activation order may have profound and unanticipated impacts on the model results (Comer, 2014). Furthermore, there is not always a clear mapping between the outputs of an abstract ABM and empirical data.

Empirical ABMs appear to be even less common in international relations. Such models often require significant amounts of data, or assumptions when necessary data is missing. Similarly, the more complicated the agent behavior, the more coding is required, adding to the assumptions, the parameters to be estimated or explored, or potential points of failure. These requirements help suggest why many empirical ABMs of international relations appear to be primarily developed in the commercial or government sectors, including Taylor et al. (2008), Abdollahian et al. (2006) and Chaturvedi et al. (2000).

### **1.3** Research Contribution

In this dissertation, I develop the use of agent-based modeling as a methodological bridge between several of the paradigms and tools I have described above. The methodology I propose is illustrated in Figure 1.1: I begin by implementing established, well-validated models as ABMs, and introduce changes in particular sub-models in order to generate several different model variants. Empirical data for a particular starting point (such as alliances and material capability data, described in Section 1.4, below) is used to generate sets of agents, which instantiate all the model variants. Multiple simulation runs then generate several output datasets associated with each model variant. By statistically comparing these outputs to empirical data on the corresponding time period, I can assess which model variant has the most explanatory power, which in turn offers evidence for the hypotheses the variant is based upon.



Figure 1.1: Modeling Methodology

The two models I reimplement and adapt are the International Interaction Game (IIG) (Bennett and Stam, 2000a; Bueno de Mesquita and Lalman, 1992), and the Expected Utility Model (EUM) (Bueno de Mesquita, 1994, 1997, 2002). Though I did not intend to focus on the works of Bruce Bueno de Mesquita in particular, these two models are among the few game-theoretic models of international relations that have been instantiated with data intended to capture real-world scenarios and have had their results compared to empirical data for those scenarios. Both are also amenable to implementation as ABMs: they consist of unitary actors, interacting with each other in discrete steps within a broader environment of other actors. Both are ultimately models of the actors' decisionmaking, and keep many

other factors (such as the actors' coercive capabilities) exogenous. Yet both also make it possible to separate the actors' actions from the process by which they decide on them. This allows me to replace the original decisionmaking models with alternative ones based on different theories, repeat the same interactions, and see where the results correspond or diverge.

I implement both models using an object-oriented approach, keeping the sub-models as separate as possible. Each model has an overall World class which contains the agents, schedules their activation, and collects data on their interactions, where each World object is a particular model instantiation, used for one model run. The World class may also have inner classes for different sub-models. The agents are implemented such that they are comprised of two classes: the outer class stores the agents' public, common-knowledge properties, and sends and receives the interactions to the world and the other agents; the inner class controls the actual decisionmaking, storing internal, private properties and implementing methods for selecting among possible choices. This architecture allows for a clear separation between the interaction model and the decisionmaking model, and facilitates the substitution of different decision models between model instantiations.

For each model, I explore sub-model and decisionmaking variants anchored in different theories. As a baseline, I attempt to replicate the original models, and implement the rational decisionmaking assumption they incorporate. I use the International Interaction Game to focus on agent decisionmaking, and implement two new agent models. One is driven by the idea that states, the agents in the model, make decisions using standard operating procedures (Allison and Zelikow, 1999; Levy, 1986) which are updated based on simple feedback from past interactions (Steinbruner, 1976). The second model adds organizational memory, which the agents use to store the lessons learned from each previous interaction (March and Simon, 1993); when facing a new interaction, the agents draw on the lessons of the most analogous past case (Khong, 1992; Schuman and Rieger, 1992). Within the Expected Utility Model, I focus on the model's theory of coercion: embedded in the original model is an assumption that actors can use their power to directly impose their preferences on others, rather than simply inflict costs if their demands are not met. I implement and test a variant which brings the model in line with the understanding of coercion described by Schelling (1966). Other variants I implement bring the agents' risk acceptance sub-model in better line with Prospect Theory (Kahneman and Tversky, 1984; McDermott, 2001), test the effect of the sequence of agent actions, and allow the agents to attempt to balance between potential threats.

Both the IIG and EUM have been applied to explain historic international conflicts, by attempting to use the models to predict the conflicts using earlier data. In Bueno de Mesquita and Lalman (1992) the IIG is applied to European country dyads from the years 1816-1970, which Bennett and Stam (2000a) extended to all global dyads between 1816 and 1992. Similarly, Bueno de Mesquita (1998) instantiates the EUM with data on world and regional powers in 1948, and attempts to predict the course of the Cold War between the United States and Soviet Union. In both cases, the models' output is compared to data on the corresponding actors and period, with the similarity between them provided as empirical evidence that the models are effectively capturing relevant elements of the real world. Importantly, Bueno de Mesquita and Stokman (1994) directly compare the outputs of early EUM implementation (Bueno de Mesquita, 1994) with those of another model, the Stokman and Van Oosten (1994) Logrolling Model, which is instantiated with the same input data as the EUM but implements very different theoretical assumptions – namely, that outcomes are reached via multi-issue compromises rather than coercion. I apply a similar methodology, and extend it in two ways. For each model variant, I run multiple instantiations in order to generate a distribution of outcomes which I can describe and explore. I also investigate the application of the EUM to machine-coded high-resolution political event data: to the best of my knowledge, the first such application. By comparing the output of each model to the empirical data, I can not only assess their fit individually but compare them to one another. If one model variant generates outputs which fit the empirical data better than another, this is evidence that the theories the variant implements have more explanatory and predictive power than the ones the other is based on. Similarly, if two variants have similar predictive power, it suggests that neither is necessarily better than the other.

Simply attempting to reproduce the results of previous models can be a valuable test, and provide a new perspective on the previous results (and indeed, in Chapter 3 I discuss the implications of failing to reproduce prior results). Additionally, this research provides several novel contributions. By anchoring agent-based models to accepted and well-studied models in international relations, I demonstrate that ABMs are not necessarily a separate and alien methodology in the discipline, but can be utilized as an extension to established research tools. In fact, as I will argue in Chapter 3, the EUM appears to be best understood as an ABM itself rather than a formal game-theoretic model. Furthermore, ABMs can implement theories of behavior which heretofore have been difficult to examine outside of resource-intensive case studies, allowing them to be tested statistically and comparatively on large-scale data. To summarize several conclusions, I find that non-rational theories have similar explanatory power to the theory of rational behavior, and that the Schelling (1966)based coercion sub-model yields better predictions in the EUM than the original version.

This research also advances the discipline of agent-based modeling itself in several ways. It demonstrates the value of designing models to be modular, in order to facilitate experiments incorporating different sub-models but controlling for other factors. It further demonstrates how such models can be anchored to real data, and how such data can be used to comparatively test the explanatory power of model variants, and the theories they operationalize. As I argue in Chapter 3, this is particularly important when models are not simply operationalizations of previous supported theories, but are themselves a new theory or hypothesis about the system being modeled. In this way, it extends Epstein (2008), utilizing ABMs for prediction not as a way of knowing the future but as a test of the underlying theory.

Chapter 2 represents, to the best of my knowledge, the first application of reinforcement learning to a multi-agent model of agents dyadically playing an extensive-form game -a narrow combination, to be sure, but one with relevance that likely goes beyond international

relations. Indeed, though my substantive focus here is on international relations, the overall methodology of extending formal models into agent-based ones, implementing different submodel variants, and testing them comparatively against empirical data, is likely applicable across other social science fields as well. This also contributes to the predictive game theory research agenda presented by Fudenberg (2010).

### **1.4 Data Sources**

An important feature of the work in this dissertation is that it utilized historic, real-world data, both to provide the model inputs and to test the outputs. Several data sources will recur across chapters, and are thus worth examining in detail.

The majority of the datasets I use here originate from the Correlates of War (COW) project. This project is an informal collaboration of scholars who since 1963 have been facilitating the collection and dissemination of datasets useful for understanding wars and their causes (Maoz et al., 2015). These datasets have become a primary tool in the quantitative study of conflicts, and of international relations more broadly. While there are several more detailed and specific datasets of interstate relationships (e.g. Hendrix and Salehyan, 2013; Raleigh et al., 2010), there are no others with comparable breadth of scope. Applying the COW datasets to the computational models I present here is particularly fitting. In the first paper summarizing the project's early years, Singer (1972) described the project's long-term ambitions:

"[T]he natural next step would be [the conversion of analytical sub-models] to computerized, dynamic form. Then, by representing the most promising and powerful of these limited models as sub-routines of an ultimately integrated large-scale computer model, we can move to parameter estimation and sensitivity testing in a particularly effective fashion."

The first dataset used to derive inputs across all three chapters is the Formal Alliances

dataset (Gibler, 2009; Small and Singer, 1969). This dataset collects formal, written defenseoriented agreements between pairs of countries, including the dates that each agreement started and ended. Information on these treaties is collected from a variety of historical sources (Gibler and Sarkees, 2004) and coded manually into several categories, with a spotchecking procedure used to ensure consistent coding. The most recent version of the data codes four categories of alliances, from defensive pacts requiring one party to intervene militarily if the other is attacked, to ententes which may merely pledge consultation when one party faces a crisis.

In the models I present here, alliance data is generally not used directly – that is, states are not necessarily assumed to come to the aid of allies they have defensive pacts with, for example. Instead, the models follow the approach introduced by Bueno de Mesquita (1975) and treat alliances as indicators of states' unobserved interests. The more similar the sets of alliances of two states (their alliance portfolio), the closer their interests are considered to align. This similarity is estimated via Kendall's rank correlation coefficient (Kendall's tau-b) (Abdi, 2007).

The next dataset is the National Material Capabilities (NMC) (Singer, 1972, 1988), an estimation of states' 'hard power' (Wilson, 2008), their military and coercive capabilities, for each year. The NMC data consists of six indicators: population size, urban population, energy consumption, iron and steel production, military expenditure, and number of military personnel. These indicators are then aggregated into a single Composite Index of National Capability (CINC) by averaging each state's share of each indicator for all states in the dataset for that year. This normalizes the data across the different units of measurement, and gives equal weighting to each indicator. It also means that CINC provides a measure of relative strength, which is only valid within a particular year. This data is collected from a variety of historic sources, as detailed in the most recent codebook (Grieg et al., 2010); when data for a particular year is unavailable, it is estimated or interpolated. The indicators are intended to estimate the CINC rather than be used on their own; as such, the data collection methodology is directed towards broad data sources which allow direct comparison between states rather than narrower, more precise specific estimates which cannot necessarily be directly aggregated or compared.

The two datasets above are used primarily as model inputs. The main dataset I will compare the model outputs to is the Militarized Interstate Disputes dataset (Palmer et al., 2015). Militarized Interstate Disputes (MIDs) are cases where one state threatens or uses military force against another<sup>1</sup>. The dataset is human-coded from media and historic sources, and records the states involved in each dispute, noting which are the initiators. For each case, the dataset includes several variables coding the intensity of the conflict on each side, in particular the highest category of hostility reached by each side – from threat of force through declaration of war. I use this dataset as the ground truth for empirical crises and conflicts, and will measure the various models' power to predict them.

Additionally, in Chapter 3 I experiment with using more detailed event data, consisting of machine-coded media records of individual actions taken by governments and other politically-relevant actors. Each event consists of a source, a target, an action, a timestamp (generally a day) and possibly a location: more colloquially, who did what to whom, when, and where. An important advantage of event data is that it captures the microlevel dynamics of conflicts, as well as the normal patterns of activity when conflicts are not occurring. There are several different event data sets available, including the Global Data on Event, Location and Tone (GDELT) (Leetaru and Schrodt, 2013) and the Phoenix data system (Schrodt, Beieler, and Idris, 2014), both of which update on a daily basis; and the Integrated Crisis Early Warning System (ICEWS) (Boschee et al., 2015) dataset, a US government-funded system which updates daily but is only made available to the public on a year's delay. Among the datasets with global, daily coverage, Phoenix remains relatively untested, and ICEWS has been assessed as more reliable than GDELT (Ward et al., 2013b). For this reason, I choose to utilize ICEWS in order to explore whether Expected Utility Model variants can predict the volume of conflict events between different international

<sup>&</sup>lt;sup>1</sup>Strictly speaking, MIDs are such cases which fall short of war; however, the MID dataset itself includes wars as well.

actors.

#### 1.5 Organization of the Dissertation

The body of this dissertation is in the following three chapters. Chapter 2 is intended to stand on its own; Chapters 3 and 4 and more closely related, with Chapter 4 utilizing the Expected Utility Model variants described in more detail in Chapter 3, and serving to provide additional validation.

In Chapter 2, I focus on modeling the decisionmaking of states in bilateral crises, using the International Interaction Game. I restate the model as described in Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a), and describe how I reimplement it as an agent-based model. Using this implementation, I introduce a notional model I will use to test the agent decision sub-models, using simplified agent properties and payoff calculations and a reduced game tree. Two new decisionmaking models are presented, which can be applied to these games: one is a pure reinforcement learning model, meant to capture a stylized version of states' learning and application of standard operating procedures which guide their decisionmaking in crises; the other adds a case-based component to the reinforcement learning model, representing how decisionmakers utilize lessons learned from previous similar interactions when facing new crises. Agents are endowed with both decisionmaking models, first in the notional simplified model and then in the IIG using data on historic dyads and disputes. I compare the outcomes of the games utilizing those two decisionmaking models to the rational equilibria, and measure the degree to which the agents collectively learn to reach equilibria outcomes. Finally, the IIG model results are compared to the historic observed event associated with each interaction, and estimate whether these models provide a better explanation of them than the equilibria.

Next, in Chapter 3 I describe the Expected Utility Model, and my implementation of it as an ABM. I focus on identifying the sub-models the EUM is composed of, at both the level of agent decisionmaking (where agents choose what actions to take towards one another), and at the level of the model 'world' (how the agents' actions are sequenced, and how conflicts are resolved). The chapter describes the assumptions embedded in these sub-models, many of which are not explicitly articulated in the original descriptions of the models, and builds on those by proposing and implementing alternatives. I generate notional instantiations in order to conduct a sensitivity analysis and quantify the uncertainty captured by the original model, which has not previously been well-studied. Finally, six model variants are instantiated, each comprised of a different combination of sub-model variants, with input data drawn from the Correlates of War datasets for key world powers in 2004; each model variant is run multiple times, in order to test whether the conflicts generated by the models are useful predictors for ICEWS conflict events for the years 2005-06 following the input data. This allows for a comparison off the predictive power of the various model variants, which in turn provides evidence for which ones are the best models of the actual underlying dynamics between the real actors.

Chapter 4 extends the analysis of the EUM as a tool for modeling international relations by attempting to replicate, and expanding on, the model's previous application to simulate the course of the Cold War (Bueno de Mesquita, 1998). The original paper proposes that, contra Gaddis (1992), the end of the Cold War did not violate established theories of international relations, but can be understood as emerging from them – and that a model, such as the EUM, can capture the complex interactions which lead to its emergence. The procedures described in the original paper is repeated in order to generate model input data from the Correlates of War datasets and to compare the new results to the originals. I then run two variants of the EUM: my reimplementation of the model described in Bueno de Mesquita (2002), and the model variant identified as having the most predictive power in the ICEWS analysis in Chapter 3. Multiple instantiations of each are run, using both the original and updated inputs, and collect the conflicts generated by each run. Following the methodology in Bueno de Mesquita (1998), I examine the resulting distribution of outcomes, and sample agent position traces which lead to each, comparing them to historic intuition about the dynamics of the Cold War. Additionally, I test whether the the conflicts generated by each model are a useful predictor of the number of MIDs recorded for each dyad. Finally, I analyze the distances between agents at the end of each model run, and test whether they are useful predictors of the alliance network observed in 1998, corresponding to the model's end-point.

Finally, Chapter 5 concludes the dissertation. I summarize the results of the previous chapters, and discuss their broader implications. With these findings in hand, I can more concretely describe how this work forwards the broader comparative analysis and agentbased modeling research areas, and propose extensions and experiments which can further build on the work presented here.

## Chapter 2: Learning and War: Applying Learning Models to the International Interaction Game

#### 2.1 Background

An important topic in the study of international relations is how the actors in the international system (primarily countries) make decisions when interacting with one another. One common view, most often associated with the Realist schools of thought (Korab-Karpowicz, 2013) is that states<sup>1</sup> are rational: they act in a way that attempts to maximize their utility, often defined as security from the threat of attack by other states (Keohane, 1986). Many critiques of Realist theories, particularly from the Liberal perspective (e.g. Keohane and Nye, 1987), focus on disputing the constraints and goals faced by states and on whether states are even the correct actors to concentrate on, but not that the relevant actors themselves are rational. Nevertheless, there has also been substantial scholarship arguing against the idea of states as rational, goal-directed actors, focusing both on the differing objectives and relationships of the internal actors involved in the decisionmaking process (Allison and Zelikow, 1999; Singer, 1961; Wendt, 1987), and on the psychological biases of the individuals who are, ultimately, the ones making the decisions (Jervis, 1976; Kaufmann, 1994; Khong, 1992; North, 1967).

Crises, and the precipes of conflict between countries, are particularly important phenomena. The decision to threaten or use force is one of the most consequential external decisions a state may make, and the substantial costs and consequences of wars at the national and human levels make them compellingly important to study, explain, and possibly even predict or avert. In "The Essence of Decision," Allison and Zelikow (1999) closely

<sup>&</sup>lt;sup>1</sup>I will use states and countries interchangeably to refer to sovereign, territorial political entities.

examined one particularly important example: the 1962 Cuban Missile Crisis, in which the world's two superpowers, the United States and the Soviet Union, faced a series of decisions to initiate or escalate a military confrontation with one another that had a real possibility of leading to a nuclear exchange. The book reviews three models of state decisionmaking: rational choice, where the decisionmakers attempt to evaluate various possible courses of action and choose those which have the highest expected utility, while taking the adversary's responses into account; organizational behavior, wherein the states' decisions are driven by standing institutional procedures; and governmental politics, in which decisions are driven by competition between internal factions and actors with different interests and degrees of influence. The book demonstrates that all three models can provide convincing explanations of the behavior of the United States and Soviet Union over the course of the crisis, and in particular that the latter two models can explain actions that do not make sense as strategic, security-maximizing decisions. Similarly, Kaufmann (1994) conducted a detailed case study of another international crisis, the 1905-06 Moroccan Crisis (wherein Germany attempted to undermine French influence in Morocco and brought Europe to the brink of war), and used archival sources to argue that the key German decisionmakers updated their views and positions during the crisis in response to events and new information in a way that is inconsistent with pure rationality, but consistent with the psychological biases expected to affect them.

Both of these examples demonstrate the power of comparative analysis of different decisionmaking models, and especially the application of multiple models to the same case. However, both share a reliance on qualitative analysis of detailed primary sources detailing not only the ultimate behavior of the state-level actors but the internal deliberations of their decisionmaking individuals and organizations. This makes it difficult to test these models across different cases, or to use them to make broad hypotheses as to the conditions under which we can expect conflicts or predict how they will be resolved. In this, they are nearly the complete opposite of game theory, perhaps the other methodological pole in the study of international relations. Game theory is the mathematical study of models of interactions (games) between multiple actors (players), where the payoff each player receives depends on the combination of their own actions and those of the other players. Game theory generally takes utilitymaximizing rationality as axiomatic; this allows games to be analyzed in order to identify and characterize their equilibrium – the outcome of a game when all players are perfectly rational, know that the other players are perfectly rational as well, and figure out the best possible course of action given that the other players are doing the same thing. Recently, however, behavioral game theory (Camerer, 2003) has sought to use experiments and psychological insights to study how humans actually approach such games, and how their behavior differs from theoretical, rational expectations.

Game-theoretic models tend to be highly abstract, and may provide useful analogies and theory-generating insights rather than specific explanations or predictions applicable to a particular case (Snidal, 1985). An important exception is included in "War and Reason" by Bueno de Mesquita and Lalman (1992), which presents the International Interaction Game: a game-theoretic, extensive-form model of inter-state crises. The game consists of two agents sequentially choosing whether to escalate confrontation or attempt to deescalate, with outcomes ranging from the maintenance of the status quo (when neither actor escalates at all) to war (when both actors escalate as far as possible). While much of the book is devoted to conventional formal analysis of the model and its equilibria under different parameters, the authors also test the theory empirically. They propose a methodology to estimate the game's payoffs for any given pair of states, using data on the states' alliances to measure the conflict stakes, and states' material capabilities to measure the strength each side and its allies can bring to bear. Using this methodology with historic data from the Correlates of War project, they estimate the game payoffs for 707 historic European dyads (pairs of countries in a particular year) from the 1816-1970 time period. They find the equilibrium outcome for each, and show that these outcomes are useful predictors of the corresponding event observed between those countries for that year. Bennett and Stam (2000b) developed a software tool expanding the estimation to all historic dyads across the
entire international system between 1816 and 1992, largely confirming that the estimated equilibria can be used to predict the observed event for the corresponding dyad – though the strongest predictions required controlling for distance, great power status, and other relevant factors.

Agent-based models have seen relatively limited use in the study of international relations, and have generally been applied similarly to game-theoretic models. Many (e.g. Axelrod, 1997; Cederman, 1997; Min, 2002) involve simulations of notional, simplified international systems, with agents acting in accordance with some particular theories. These models provide stylized insights into what patterns emerge (or do not) from these theories, but are not meant to directly describe specific real-world countries or conflicts. One important example is the Axelrod (1980) Prisoners' Dilemma Tournament, which explicitly sought to compare the performance of different behavior rules playing an iterated game. While this research utilized a simple theoretical model, it was done with an eye towards the study of, among other issues, international conflict and cooperation. Similarly, though less robustly, Cederman (1997) involves competition between agents with two sets of behavioral rules, those content with the status quo and those seeking to expand at the expense of their neighbors. These experiments demonstrate a powerful feature of agent-based models, which has heretofore been underutilized in political models: the ability to have agents using different behavior models interact and compete with one another within a common context.

While it is not presented as such, the Bueno de Mesquita and Lalman (1992) empirical application of the International Interaction Game bears some similarity to an ABM. Indeed, the empirical application of the model is only feasible computationally. It consists of distinct, unitary actors with well-specified decision rules, carrying out discrete interactions with one another. These interactions, in turn, generate a series of cooperative and conflictual behaviors, which can be compared to empirical data. Nevertheless, as previously implemented by Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a,b), the agents (such as they are) only feature a single decisionmaking model, rationally playing the equilibrium moves. The way I will extend the International Interaction Game (IIG) should now be clear: I propose to reimplement the model explicitly as an agent-based model, and to do so in a way that allows me to endow agents with different decisionmaking models and rules, drawn from different theories of state behavior. Utilizing the IIG game tree and payoff estimates keeps the model anchored to a well studied, empirically validated foundation. I then use this framework to simulate the same set of interactions using different decisionmaking models, obtaining sets of outcomes for each variant. These outcomes can be compared to one another to see where and how they differ, and to empirical data, to see which provides the best fit.

In this framework, formal rationality will be the initial decisionmaking model I will implement. As discussed above, this is the default model of decisionmaking used in gametheoretic models. It serves as an operationalization of the qualitative rational choice model as described by Allison and Zelikow (1999) and many others. The analysis in Bennett and Stam (2000a) compares the explanatory power of the overall IIG to dyad-level variables (e.g. geographic distance, government types) but not to other models of decisionmaking. Here, in contrast, using rationality as a baseline will allow us to directly test whether the additional decisionmaking models provide stronger explanatory power than the rational decisionmaking model itself. Furthermore, reimplementing the rational choice model will serve as a form of verification, ensuring that my model framework is producing the same results as the Bennett and Stam (2000a,b) implementation.

I implement two alternative models of decisionmaking within this framework. The first variant is a stylized representation of the organizational procedure model described by Allison and Zelikow (1999) and Levy (1986): agents develop standard operating procedures, which they apply regardless of the specifics of each interaction they find themselves in. These procedures are updated based on the outcome of the new interaction, such that successful courses of action will be more likely to be executed in the future, and unsuccessful ones less likely.

The second variant implements a stylized version of the theory described by March and Simon (1993) and Khong (1992) by which organizations, and states in particular, make decisions by comparing their present circumstances to previous similar cases they encountered. In this variant, the agents maintain a library of past interactions, which they draw on to make decisions about each new interaction. If a course of action was successful in a similar previous interaction, the agent will be more likely to use it again; otherwise, they will be more likely to attempt an alternative.

These decisionmaking models are both implemented using reinforcement learning. This methodology has several advantages: it originates from psychology, but has had success as a machine learning technique (Sutton and Barto, 1998). Reinforcement learning agents have been shown to be capable of learning to play extensive-form games (Akramizadeh, Menhaj, and Afshar, 2009; Laslier and Walliser, 2005), as well as to learn to cooperate and compete in multi-agent environments (Claus and Boutilier, 1998; Littman, 1994). They have been used to model the learning of individuals and groups of decisionmakers (Kocher and Sutter, 2005); however, they are not always successful at learning to play the equilibrium of abstract, normal-form games, and in fact can yield unexpected and even chaotic behavior (Galla and Farmer, 2013). Furthermore, Khong (1992), Allison and Zelikow (1999), and Cederman (2001) have observed that states as a whole (rather than just individual decisionmakers) exhibit a degree of learning from past experiences. In short, reinforcement learning provides a plausible, general model of decisionmaking which is not procedurally rational (Simon, 1976), but introduces minimal additional assumptions as to the form the behavior takes. Additionally, reinforcement learning makes minimal assumptions as to the structure of the game the agents are playing, meaning that a single behavior model may be implemented once and used to drive agents across different games and other interaction models.

The question I set out to explore here is whether reinforcement learning agents interacting with one another over a series of simulated international crises can converge to behavior resembling the equilibria for those interactions. If this is the case, it will suggest a degree of convergence between the rational choice, psychological, and organizational models of decisionmaking. If it is not, however, it will diverge, generating different results which may then be tested against empirical data. In the course of addressing these questions, I will also lay out a more general architecture and methodology for extending game-theoretic models of international crises into agent-based models, which can be used to implement and test additional models of decisionmaking. As Allison and Zelikow (1999) write, "because simplifications are necessary, competing simplifications are essential." In this chapter, I will present a methodology for extending the competing simplifications from the realm of case studies to computational modeling.

This chapter is organized as follows. In Section 2.2, I present the general architecture for an agent-based model of dyadic, extensive-form interactions with variable agent decisionmaking. I describe two simulated worlds: a simplified, notional international system (Section 2.3.1), and a more complicated data-grounded historic simulation based on Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a) (Section 2.3.2). Next, I present several behavior variants, and ground them in political and multi-agent theory (Section 2.4). In Section 2.5.1, I apply these behavior variants to the notional model and verify their behavior, and then apply them to the full historic data in Section 2.5.2. I compare their behavior to the purely rational decisionmaking, and to the observed historic outcomes of each simulated interaction.

# 2.2 Model Description

The overall model I present here consists of three sub-models intended to capture distinct levels of phenomena: the 'World' consisting of a set of agents with certain properties, which may be fixed or change over time; the Interaction sub-model, consisting of an extensive-form game tree for pairs of agents to play against each other; and the Agents themselves. The Agent class, in turn, is composed of two parent classes: the World Agent, which defines the properties and methods specific to the World, and a Decisionmaking Model, which the agents utilize to play the extensive-form game at each interaction. All these sub-models are implemented initially as abstract classes, which include the methods common across all model variants; these abstract classes are then sub-classed and implemented for each model variant. Figure 2.1 shows a UML class diagram tying together these different classes.



Figure 2.1: Model Architecture

## 2.2.1 World Sub-Model

The World sub-model acts as the overall controller for the entire model. It stores and manages the agents who will interact with each other over the course of each instantiation, generating a particular notional history. A World is instantiated with a set of agents, as well as a subclass of the generic Interaction sub-model class, described below. The World implementation can also accept additional data used to generate agents, schedule their interactions, and store and export the outcomes of those interactions.

The World's key method is the 'step', which advances the model by one simulated year, chosen since that is the level at which the IIG data is available. Each year consists of some set of interactions, depending on the model variant. For example, these may be some number of random pairings of agents, all possible agent dyads, or pairings based on crises occurring in the given year as recorded in the input data. The agents do not directly choose with whom or when to interact; the interaction scheduling is exogenous to the agents themselves. Since the player order matters in extensive-form games, the interaction dyads are directed – that is, an interaction between A and B is not the same as one between B

and A.

Finally, the World also collects and stores data on each agent, and on the result of each interaction – that is, which terminal node the interaction ended up on, as explained below. A World implementation may also perform some evaluation on this interaction outcome. In particular, the World sub-models I develop compare the outcome of each interaction to that interaction's equilibrium outcome.

# 2.2.2 Interaction Sub-Model

This is the core of the overall model: a single interaction between agents playing through a standard, two-player, extensive-form game (Shoham and Leyton-Brown, 2009). An interaction model is implemented with a game tree, and methods for computing the payoff to each agent at each terminal node. The game tree here is assumed to be a direct, acyclic graph with perfect information (that is, there are no information sets of more than one node, and agents know exactly the node where the game is currently). The player agents take turns in fixed order, and there are also no moves by nature.

Each agent has a step method, which takes the interaction state as an input and emits a move. The sub-model makes sure that the move is valid for the current node, and advances the game's current node to the appropriate result of the given move. If this is a terminal node, the interaction is over, and payoffs are distributed to all the agents. Otherwise, the next player agent is called.

I use two game trees here, which are shown in Figures 2.2 and 2.3. The simpler one is drawn from Signorino (1999), and will be used for the notional model; the more complicated one is from Bueno de Mesquita and Lalman (1992) and will be applied to the historic data. The utility calculations associated with both will be provided in more detail in the appropriate sections below.

## 2.2.3 Agent Sub-Model

Finally, there are the agents themselves. Each agent is meant to represent one actor in the international system, as they interact with one another as described above. Each agent class has two parent classes, as shown in Figure 2.1: one is the World Agent, for the particular model variant which provides its properties; the second is the Decisionmaking Model, which contains the methods and properties the agent uses to make decisions at each step of the game. All agents in a given World must have the same World Agent parent, which is compatible with the Interaction sub-class. However, different agents within a world may have different decisionmaking parents.

In general, the agent properties deriving from the model variant parent class are used to compute the utilities the agent expects from each terminal node in the game tree. The decisionmaking parent, in turn, provides the 'step' method, with which the agent chooses its next move through the game tree during an interaction. This structure allows decisionmaking rules to be modeled once, verified, and applied across different specific interaction and world models with minimal additional development.

# 2.3 Model Variants

I present two model variants below. Each one consists of an implementation of the three submodels described above: a World, Interaction model, and agent properties. The first variant, the Small Crisis Game, is a simple notional model used for testing the agent behaviors under controlled conditions. The second, the International Interaction Game, is a full model of the international system extending Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000b), at a level of one agent per state, attempting to capture historic interaction dyads.

## 2.3.1 Small Crisis Game

This model variant is a purposeful simplification: in the number of agents, their properties, and the game tree of their interactions. It is adapted from Signorino (1999), which in turn adapts the game tree shown in Figure 2.2 from the game trees in Bueno de Mesquita and Lalman (1992), as well as Kim and Bueno de Mesquita (1995) and Fearon (1995). At every node of this game tree, agents (Players 1 and 2, or P1 and P2) may escalate the conflict (fight, F) or refrain from escalating (not fight, !F). If neither agent chooses to escalate, the result is the Status Quo; if one agent escalates, the other agent faces a choice between not escalating and yielding to the escalator's demand (Capitulating), or fighting, leading to War. These outcomes are the tree's terminal nodes.

The payoffs associated with these terminal nodes are dependent on agent properties. Each agent *i* in this model is defined by its Assets  $(a_i)$ , indicating what they have that other agents might want (whether territory, material, or more abstract positions or concessions); their Capability  $(c_i)$ , indicating the strength they have to coerce other agents and defend themselves from coercion by others; and a Bloc  $(b_i)$ , where two agents are geopolitically aligned (and have more incentive to cooperate) if their Bloc properties match. These properties are used to compute the utilities associated with each terminal node, as shown in Table 2.1. These utilities are also adapted from Signorino (1999): when agents capitulate, they yield the assets being contested, with the other agent gaining them. If the agents enter into conflict, they will win with probability based on the ratio of their capability to the other agent's; if they win, they secure the other agent's assets – and if they lose, they lose not only the value of their own assets, but of their expended capability as well. The Status Quo leads to no change in utility if the agents belong to different blocs; however, if the agents belong to the same bloc, cooperation (i.e. not entering into conflict) has a positive payoff<sup>2</sup>.

 $<sup>^{2}</sup>$ In Signorino, 1999, there is a Democracy property instead of a Bloc property, with only dyads of democracies receiving the cooperation payoff; the form I present here is more general, and does not introduce the assumption of democratic peace. However, initial experiments using the original specifications produced substantively similar results to the ones I present here.



Figure 2.2: Small Crisis Game Tree (adapted from Signorino, 1999)

Table 2.1: Small Crisis Model Utilities

Node	$U_1$	$U_2$
Status Quo (SQ)	$20\delta(b_1, b_2)$	$20\delta(b_1, b_2)$
Capitulate 1 $(C1)$	$-a_1$	$a_1$
Capitulate 2 $(C2)$	$a_2$	$-a_2$
War (W1, W2)	$p_1a_2 + (1-p_1)(-a_1 - m_1)$	$p_2a_1 + (1 - p_2)(-a_2 - m_2)$
$p = \operatorname{Prob}(\operatorname{win})$	$\frac{c_1}{c_1+c_2}$	$\frac{c_2}{c_1+c_2}$

For every instantiation of the CrisisWorld, N agents are generated at random, with random properties drawn uniformly as shown in Table 2.2. These agents will persist across the entire run; that is, agents are not added to or removed from the model instance. In each instantiation, all agents will share the same decisionmaking model, as specified below.

Each year of the model (that is, a step of the world), all the agents will be paired up, in random order. Since the number of agents being simulated is substantially smaller than the

Property	Distribution	Change
$a_i$	$Unif\{1, 100\}$	$\mathcal{N}(0,5)$
$c_i$	$Unif\{1, 100\}$	$\mathcal{N}(0,5)$
$b_i$	Bernoulli(0.5)	1 - Bernoulli(0.05)

Table 2.2: Small Crisis Agent Properties

total number of agents in the international system, this can be thought of as representing a particular region (analogous to the focus on European dyads in Bueno de Mesquita and Lalman (1992)), or a subset of politically-relevant actors within a larger system.

I create two variants of the Small Crisis model. In the first, the Static variant, the agent properties are fixed across each run. In the second, the Dynamic variant, the properties of each agent change stochastically at the beginning of every simulated year, with the changes distributed as shown in Table 2.2. This means that under the first variant, the utilities (and hence the equilibrium) of an interaction between each dyad of agents are constant, whereas in the second variant they may drift from step to step.

Note that the agent attribute values do not change due to the results of their interactions. While this is a simplifying assumption, it also corresponds to the dynamics observed in the National Material Capabilities dataset (Grieg et al., 2010; Singer, 1988), wherein the majority of conflicts do not substantially impact states' estimated material capabilities or alliance relationships – which are used in the empirical International Interaction Game to operationalize the capabilities of states and the stakes of conflict, respectively, as described below.

### 2.3.2 International Interaction Game

The International Interaction Game, as described in Bueno de Mesquita and Lalman (1992) is as follows. It involves two actors, one who initiates the interaction and one who responds. The game proceeds following a tree, as shown in Figure 2.3. As in the simplified model, at each node the actors (Players 1 and 2, P1 and P2) have a choice between two moves:

escalating the confrontation (fight, F) or not escalating (!F). There are more nodes in this tree, leading to a wider range of outcomes. If neither actor escalates, the result is still the Status Quo, while if both consistently escalate the result is War. If one agent yields to the other's escalation early on, the outcome is Acquiescence, while if they yield after a greater degree of escalation, the outcome is Capitulation, with some additional cost incurred. If both actors initially escalate and then deescalate, the result is a negotiated compromise.



Figure 2.3: International Interaction Game Tree (from Bueno de Mesquita and Lalman, 1992)

While each interaction is dyadic, the utilities of each outcome for both actors are determined by the overall international system the interaction occurs within, and specifically the network of alliances. Alliances are used as a proxy for states' interests. The COW alliance data (Small and Singer, 1969) records the type of each formal alliance each state has with any other state in a given year; this is that state's alliance portfolio. As Bueno de Mesquita (1975) argues, two states with similar interests can be expected to have similar alliance portfolios – that is, many of the same types of alliances with the same other states. Bueno de Mesquita and Lalman (1992) uses a standard statistical similarity measure known as Kendall's tau-b to compute this similarity, though Signorino and Ritter (1999) argues that this measure is inadequate, and proposes an alternative. To maintain consistency with Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a), I will use Kendall's tau-b, denoted below as  $K_{i,j}$  – the similarity between actors *i* and *j*'s alliance portfolios. However, other similarity measures, such as those proposed by Signorino and Ritter (1999) may be substituted as well.

The estimated similarity between the two states is used to determine the two base utility rates for the corresponding actors, under the assumption that the more dissimilar two states' interests are, the more each has to gain or lose from a conflict.  $U^A(\Delta_A)$  is Actor A's utility if its own objective is achieved, and  $U^A(\Delta_B)$  is its utility if the other actor B achieves its objective instead.

$$U^{A}(\Delta_{A}) = 2 - 4 \left(\frac{2 - (1 - K_{B,A})}{4}\right)^{r_{A}}$$
(2.1)

$$U^{A}(\Delta_{B}) = 2 - 4\left(\frac{2 - (K_{B,A} - 1)}{4}\right)^{r_{A}}$$
(2.2)

Where  $r_A$  represents each actor's risk tolerance, also estimated from the alliance network as described in Bueno de Mesquita (1985) and Bennett and Stam (2000b)<sup>3</sup>.

If the two actors have identical alliance portfolios, and hence approximately identical interests, they will both be indifferent as to who will 'win' a confrontation – meaning that a confrontation will certainly not occur.

In addition to these utilities, the international alliance network determines each actor's probability of prevailing in negotiations or war, based on the support they can muster from

<sup>&</sup>lt;sup>3</sup>Risk tolerance is computed by comparing an agent's total probability of victory in all possible conflicts (as detailed below in Equation 2.4) to the maximum and minimum possible total victory probabilities, across all possible alliance portfolios. This is by far the most computationally intensive step; the Bennett and Stam (2000b) EUGene software, which implements the IIG, uses pre-computed results from a single six-month optimization done using genetic algorithms. This methodology is also almost certainly yielding an inaccurate estimate. Recent work such as Maoz (2010) and Cranmer, Desmarais, and Menninga (2012) shows that alliances are best understood as mutually dependent network edges, meaning that they ought not be permuted independently as was done here. A better methodology would estimate risk tolerance from structurally plausible networks only, which would likely have the added advantage of greatly reducing the search space. However, a full implementation of this methodology is outside the scope of this dissertation, and must be left for future work. Additional future work along the same lines may also conduct a sensitivity analysis on risk tolerance in order to estimate the degree to which different values lead to substantively different equilibria or behavior.

the rest of the actors in the system. Specifically, it determines their *perception* of their probability of prevailing. For each actor in the international system, the model computes their own interest in the bilateral conflict under consideration, based on the similarity of that actor's alliance portfolio to the two parties to the conflict. They then allocate some of their resources, estimated by the COW National Material Capabilities index (Correlates of War Project, 2010) to one side or the other. Agent A's estimate of Agent C's preference for A over B is:

$$Pref_{A}^{C}(A,B) = \frac{K_{C,A} - K_{C,B}}{2} \exp(R_{A}(K_{C,A} - K_{C,B}))$$
(2.3)

 $(R_A \text{ is a rescaling of } r_A).$ 

Then, each actor in the conflict has a probability of prevailing equal to its share of the total resources allocated. Agent's A's perception of its probability of prevailing against Agent B is:

$$P^{A} = \frac{\sum_{c \in \Psi} \Lambda^{c} Pref_{c}^{A}(A, B)}{\sum \Lambda^{c} Pref_{c}^{A}(A, B)}$$
(2.4)

Where  $\Lambda^c$  is c's capabilities (operationalized as NMC) for every actor C, and  $\Psi$  is the subset of actors who prefer A to B (that is,  $Pref_A^c(A, B) > 0$ ). Note that this formulation is similar to the model of bilateral conflicts found in the multilateral expected utility model (Bueno de Mesquita, 1997, 2002; Scholz, Calbert, and Smith, 2011) described in further detail in the following chapter.

Note the  $R^A$  term in Equation 2.3 – this is the risk tolerance factor of one of the actors in the conflict, not the third-party actor who may be contributing resources. This is implicit mirror-imaging, a perception documented intelligence in both the cognitive (Meltzoff and Decety, 2003) and political senses (Heuer, 2001) where the actor assumes that everyone shares its own attitudes, in this case towards risk. What Equation 2.4 is capturing, then, is not an objective probability of victory in a possible conflict, but each actor's own perception of this probability. Thus, the utilities appearing below are agent A's expected utilities from each outcome.

Using these probabilities, we can finally write the utilities for all the outcomes of the IIG, which are:

$$U^{A}(StatusQuo) = 2 - 4(\frac{2}{4})^{r^{A}}$$
(2.5)

$$U^{A}(Acqiesce_{B}) = U^{A}(\Delta_{A})$$
(2.6)

$$U^{A}(Acqiesce_{A}) = U^{A}(\Delta_{B})$$
(2.7)

$$U^{A}(Negotiate) = P^{A}U^{A}(\Delta_{A}) + (1 - P^{A})U^{A}(\Delta_{B})$$
(2.8)

$$U^{A}(Capitulate_{B}) = U^{A}(\Delta_{A}) - P^{A}$$
(2.9)

$$U^{A}(Capitulate_{A}) = U^{A}(\Delta_{B}) - (1 - P^{A})$$
(2.10)

$$U^{A}(War) = P^{A}(U^{A}(\Delta_{A}) - P^{A} - (1 - P^{A})) + (1 - P^{A})(U^{A}(\Delta_{B} - P^{A} - (1 - P^{A})))$$
(2.11)

Using the methodology described above, Bueno de Mesquita and Lalman (1992) test the IIG model by computing the utilities and finding the equilibria for 700 dyads between European countries. Bennett and Stam (2000a) extended this analysis by developing the EUGene (Expected Utility Generation) software tool<sup>4</sup>, which is capable of computing utilities for all possible dyads in the international system based on the then-current versions of the Correlates of War datasets (Jones, Bremer, and Singer, 1996; Sarkees and Schafer, 2000). EUGene also modifies the form of the probability of victory in Equation 2.4 to

<sup>&</sup>lt;sup>4</sup>Code and binaries available at http://www.eugenesoftware.org/

account for states' declining ability to project force over greater distances. Finally, as mentioned in Footnote 3, EUGene includes pre-computed risk tolerances, which are prohibitively expensive computationally to reestimate<sup>5</sup>.

While the Correlates of War project has since updated its datasets from the ones used by Bennett and Stam (2000b), yielding slight changes particularly in the alliance networks used to estimate utilities and risk coefficients, I continue to use the EUGene output for this analysis. This serves two purposes: it allows me to directly compare my results to previous ones relying on the same data; and it allows me to concentrate on the comparative behavioral modeling as opposed to the utility calculations. However, future work utilizing the methodology described here for more accurate historic comparison and forecasting should probably recompute the utilities from more recent and complete data.

The EUGene software generates a dataset of 1,610,478 rows, with each row a country dyad by year. I filter out those rows for which, for various reasons, the EUGene model could not compute utilities for those pairings, leaving 1,027,692 rows. Each row of the dataset represents two actors playing the international interaction game at a particular point in time, with particular interests and in a particular context. Each row records the conflict event, if any, that occurred between the two actors in that year; the event categories correspond to the terminal nodes in the IIG tree (see Figure 2.3) and are categories based on the level of military force each actor used against the other, as detailed in Bennett and Stam (2000b). Since the levels of force are sequential, each one corresponds approximately to an escalation move in the IIG. This assumes that each directed dyad of states will have at most one crisis per year. In the majority of cases, neither side will have used any military force against the other, corresponding to a Status Quo outcome.

The World sub-model creates one agent for each unique state which appears in the

<sup>&</sup>lt;sup>5</sup>The task of estimating a single actor's risk tolerance for a single year requires finding the best and worst alliance portfolios it may hold; this is conceptually very similar to the optimization performed in fitting exponential random graph models, which has been shown to be #P-hard (Bannister, Devanny, and Eppstein, 2014). This optimization must then be performed for each actor in each year, for a total of approximately 15,000 times.

dataset, as identified by the country code variable. These agents persist across the entire model instantiation. This is a simplifying assumption, which assumes continuity for countries which may have experienced radically different governments over their lifetimes; however, it is in line with the approach of Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a), and indeed of the COW data itself which does not distinguish between different regimes presiding over the same state.

In agent-based modeling terms, we can use this dataset as the agent scheduler, iteratively advancing through the table and running the Interaction sub-model for the dyad of agents in each row. By default, the data is sorted by year, and then first and second agent country codes. While the year ordering is obviously temporal, the code ordering is not. Since the decisionmaking models presented here are path-dependent (agents learn from the outcome of one interaction, influencing their decisions in the next interaction) it is important to reduce the effect of potential artifacts of particular interaction sequences. In order to do so, the rows of the data table are shuffled within each year for each instantiation. The model then loops over the shuffled rows. For each row, it retrieves the agents specified in that row, and has them play the IIG, with the payoffs on the terminal nodes specified in the row. Note that this implicitly introduces the assumption of one interaction (or potential conflict) per directed dyad per year; however, again, this assumption is in line with the approach taken by Bennett and Stam (2000a), and indeed with the way the COW data itself is organized (Palmer et al., 2015). Finally, and most problematically, since states' attributes are taken directly from the data, this methodology assumes that the particular outcome of each interaction does not directly change them. For example, if an interaction results in a War outcome when in reality the states remained at Status Quo for that dyad-year, the war will not change the agents' capabilities or alliances. This is a simplifying assumption, to be sure, but allows us to focus on modeling the decisionmaking itself, rather than developing and incorporating additional models of the material and multilateral consequences of wars.

Figure 2.4 shows the number of new states present in the dataset each year, and the number of interactions between them. Since the number of potential interactions goes as the

square of the number of states, it is not surprising that the number of interactions per year increase substantially in the latter part of the dataset. Model results below will be reported by annual averages for ease of historic intuition and to account for in-year shuffling; but it is important to keep in mind that later years will include more interactions than earlier years.



Figure 2.4: EUGene Data by Year

This methodology allows us to run an agent-based model against historic data. Each agent interaction within each model instantiation corresponding to a particular observed or potential interaction, and the outcome of the model interaction can be compared to the event observed for that interaction. Importantly, as agents only are influenced by their previous interactions, each model outcome is an out-of-sample prediction.

# 2.4 Decisionmaking Models

Agents in the two models described above have different properties and interact with each other via different games. Over the course of the interactions themselves, however, the agents face the same structure of an extensive-form game, where they must choose one of several moves at each node, with a final payoff depending on the moves both they and their counterpart have made. In most applications of game theory, this model is implicitly perfect rationality; in practice, the equilibrium is found not by agents but by the external modeler (with the so-called gods'-eye view of the situation). The architecture of the agent-based framework I have developed and described here requires an agent-level implementation of any decision rule, including rationality. Any such implementation operationalizes a theory or set of theories of actor decisionmaking.

The framework I describe here makes a few high-level assumptions as to the agent decisionmaking models. The agents have access to the current node of the game, and know all the possible actions they may take at that node. The rational agents described below have access to the full game tree, including the various node payoffs. I also assume that all interactions within a given model implementation share the same game tree structure. Finally, agents do not have direct access to the decisionmaking model of the other agent.

In this section, I describe the three decisionmaking models I implement. The first is classic rationality with perfect information, which serves as a baseline. The second and third are simple reinforcement learning, and a case-based learning extension; both use an experience-weighted attraction implementation so that the actual decisions are stochastic, and vary based on how the attraction weights associated with each move are chosen and updated. In the Reinforcement Learning (RL) decision model, the agents maintain a single set of attraction weights, which they use for all interactions. In the Case-Based Learning (CB) decision model, agents maintain a library of weights learned from past cases, and for each interaction select the weights associated with the most similar past case.

In the next section (Section 2.5) I apply these decision models to the Small Crisis Model and International Interaction Game, described above. My initial hypotheses are that, as suggested by previous work on reinforcement learning (e.g. Akramizadeh, Menhaj, and Afshar, 2009; Galla and Farmer, 2013; Laslier and Walliser, 2005), both the RL and CB models will converge to producing outcomes equal to the rational model's (the equilibrium outcomes) more often than we would expect purely at random; I also expect the CB model to outperform the RL model and result in more common equilibrium outcomes, since it adds a degree of coarse-grained strategic reasoning, with agents adjusting their behavior based on their counterpart in each interaction – implicitly anticipating that agent's own decisions.

#### 2.4.1 Rational Model

Before implementing additional model variants, I start by reproducing the original rational decisionmaking model. Rational agents have full access to the game tree, and both agents' utilities at each terminal node. They identify the subgame-perfect equilibrium by backwards induction (Russell and Norvig, 2009), finding the correct moves for both players to make at each node and propagating the corresponding utilities up the tree. This leads to a best-response move at all nodes of the tree, not just those that are on the main equilibrium line; thus, even when playing against an opponent who is not playing the equilibrium moves, this agent will be able to choose the game-theoretic best response. In the remainder of this chapter, the equilibrium of an interaction is defined as the result of the interaction when played between two such rational agents.

## 2.4.2 Weighted-Attraction Agents

The game-theoretic model of rationality described above is not time dependent: the particular history of the agent does not matter; the agents, already perfectly rational, do not learn. Furthermore, the agents' decisions are wholly deterministic. Finally, the agents have full awareness of where they are in the game tree, and are capable of planning ahead to the end of the tree, see the plan through, and costlessly replan if the other agent deviates from their expectations.

The two decisionmaking models that follow all begin from different assumptions. In describing the organizational model of decisionmaking, Allison and Zelikow (1999) note that, contrary to the rational model's assumption of perfect lookahead, "[o]rganizations exhibit great reluctance to base actions on estimates of an uncertain future," and instead prefer to utilize existing procedures. As Levy (1986) describes, military organizations in particular tend to prepare plans and procedures which may be inflexibly put into practice in crisis situations. Organizations in general, as Lundberg (1995) reviews, tend to learn in the form of adopting new routines based on past experiences. Finally, March and Simon (1993) note that organizations do not necessarily seek the optimal course of action, and often satisfice – find a course of action which meets some threshold of acceptability. I distill these characteristics into the following assumptions:

- The agents play the game myopically, choosing each move without explicit look-ahead in the game tree.
- The agents' choice of move at each node is based on their past experience.
- At the end of each interaction, the agents only have knowledge of the actual outcome, and cannot directly compare it to the 'road not taken.'

These assumptions are implemented via experience-weighted attraction (EWA), a particular form of reinforcement learning (Galla and Farmer, 2013; Sutton and Barto, 1998). At each node of the game tree, agents associate a certain weight with each potential action they can take. They choose among the actions stochastically, in proportion to the weight they place on each. If the weights are equal, all actions may be chosen with equal probability; if one action has overwhelmingly more weight than the others, it will be chosen with probability approaching 1. These weights are learned based on the agent's past experiences, and in particular the utilities they gained at the end of previous interactions. The specific ways the weights are learned are explained in the subsections below.

This model form is taken proximately from Galla and Farmer (2013), who in turn adopt it from several other sources (Camerer, Ho, and Chong, 2002; Ho, Camerer, and Chong, 2007; Kocher and Sutter, 2005). Most importantly for our purposes, the experimental results reported by Kocher and Sutter (2005) indicate that EWA is an effective model for strategic learning not just by individuals but by groups tasked with making a decision jointly.

Formally, the model is as follows. Within a particular interaction, an agent *i* has an attraction function  $Q_i : s_{x,k} \mapsto w$  which maps from  $s_{x,k}$ , a move available at node *x*, to

a numeric weight w. In practice, this is implemented as an associative array, with nodemove tuples as keys and weights as values. At node x, the agent chooses move  $s_{x,k}$  with probability:

$$P_i(s_{x,k}) = \frac{exp(\beta Q_i(s_{x,k}))}{\sum_{s_{x,m}} exp(\beta Q_i(s_{x,m}))}$$
(2.12)

Where  $\beta$  is a fixed parameter which further weights the choice weights. When  $\beta = 0$ , all moves are equally likely regardless of weight. The larger  $\beta$  is, the more likely a move with a smaller weight advantage is to be chosen. In general, I set  $\beta = 1$  for all agents; however, as this is an agent-level parameter, there is nothing preventing agents from having heterogeneous  $\beta$  values.

In Python-style pseudocode, the agent class's EWA method is implemented as follows:

Listing 2.1: Experience-Weighted Attraction

```
def step(model):
    current_state = model.current_node
    choices = {}
    for move in model.tree[current_state].children:
        choices[move] = self.Q[key]
    choice = self.weighted_choice(choices, self.beta)
    self.past_moves.append((current_state, choice))
    return choice
```

This decision framework can be thought of in several ways. It can simply represent the relevant decisionmakers' view of the relative merit of each possible course of action (Kocher and Sutter, 2005). It can also represent the relative influence or credibility of different factions or sub-actors advocating for different courses of action, or standing procedures developed by the relevant organizations, such as reserve call-ups or troop movements (Levy, 1986). How to interpret it will rest partially on the particular learning rule, which will determine how the weights are learned and updated between interactions. The two learning rules I use are described next.

## 2.4.3 Pure Reinforcement Learning

This model variant is the closest to traditional EWA, as described in Galla and Farmer (2013) and elsewhere. Under this decisionmaking model, agents choose actions based on their total past experiences, without taking into account any of the properties of the specific agent they are interacting with, or the broader environment. At the end of every interaction, agents update their attraction to each move they made based on the payoff they received. Moves which led to higher utilities will be more likely to be played again, while ones that led to lower utilities (and in particular, to negative utilities) will be less likely to be played again.

In terms of crisis behavior, this variant can be thought of as a stylized version of the organizational procedure model described by Allison and Zelikow (1999). The agents are developing 'standard operating procedures' they will follow without evaluating their appropriateness for the particular crisis they are facing at the moment. When the crisis is resolved, they will update their procedures based on the crisis's outcomes. To be sure, this is an exaggerated view of the organization procedure model, as nobody claims that states act solely on the basis of preset, inflexible procedures. Yet if so, it may also be no more exaggerated than the game-theoretic operationalization of the rational choice model.

Formally, each agent has a single experience table  $Q_i$ , encoding the utility it has gained from making decisions at each node in the past. This experience table drives the decisionmaking at each node of the game tree, until the interaction ends. At that point, the table is updated at the end of the interaction with the corresponding utility.

At the end of an interaction, let the set of moves agent *i* has decided to play be  $S_i = \{s_{x_0,k}...s_{x_l,k}\}$  for each node in the game tree. Let  $\Pi_i$  be the payoff associated with the terminal node reached by  $S_i$  and  $S_j$ . The updating rule is then:

$$Q_i^{t+1}(s_{x,k} \in S_i) = \alpha \Pi_i + (1 - \alpha) Q_i^t(s_{x,k})$$
(2.13)

Parameter  $\alpha$  is the learning rate, where  $\alpha \in [0, 1]$ . When  $\alpha = 1$ , the agent has effectively

no memory, and associates a move only with its most recent payoff. Conversely, if  $\alpha = 0$  the agent never updates its attraction weights. In practice, obviously, this value will be somewhere in between these extremes. The  $\alpha$  value is set by the modeler; I will explore its effects in Section 2.5.1.

This model differs from some other, similar models, in that it attempts to capture realistic learning from experience rather than serve as an optimization algorithm. Thus, it does not incorporate fictitious play (Shoham, Powers, and Grenager, 2003), and agents have no ability to access other potential payoffs from decisions not taken. Importantly, due to this, the weights on moves not taken do not update, meaning that entire branches of the game tree may be unaffected by a particular interaction where play does not pass through that branch. Similarly, the weight on any move not taken is not updated; thus, whether a particular update will substantially change the probability of a move being chosen in the future depends not just on the move's own updated weight, but the weight present on the moves not chosen.

In Python-style pseudocode, the agent's experience update method is provided below. Note that this uses the *past\_moves* list updated in Code Block 2.1.

Listing 2.2: Reinforcement Learning

```
def add_utility(utility):
    for move in self.past_moves:
        self.Q[move] *= (1 - self.learning_rate)
        self.Q[move] += self.learning_rate * utility
```

## 2.4.4 Case-Based Learner

In the previous model variant, agents' past experiences were undifferentiated: the agents maintain and update a single experience table across all their interactions. This means that in a new interaction, the agent will utilize the same decision weights whether facing, for example, a strong ally or a weak rival. However, individuals and organizations are capable of maintaining and recalling and utilizing discrete memories of specific circumstances they have previously encountered. Gilboa and Schmeidler (1995) argue that case-based reasoning provides a cognitively plausible model of how individuals reason under uncertainty; similarly, March and Simon (1993) describe one of the modes of organizational decisionmaking as choosing actions based on recognizing a given situation's similarity to previously-encountered ones stored in organizational or individual memory, and retrieving rules associated with those situations. In the context of international crises and conflicts, the relevant situations are most often prior crises and conflicts. Historic analogies help shape decisionmakers' view of the stakes of a crisis, as well as the probabilities and risks associated with various courses of action (Khong, 1992). These analogies influence the public debate surrounding potential conflicts (Noon, 2004; Schuman and Corning, 2006; Schuman and Rieger, 1992), as well as the strategies and tactics which may be brought to bear (e.g. Nagl 2002 and too many military 'Lessons Learned' to cite). Furthermore, lessons learned from some conflicts may also lead to the reevaluation of previous conflicts (e.g. Hagerman 1992 or Förster and Nagler 2002, among many examples), sometimes exaggerated to lead to thought experiments such as a reevaluation of the American Revolutionary War in light of the lessons of the United States's Iraq and Afghanistan wars (Exum, 2008).

In order to model such reasoning by historic analogy, I extend the reinforcement learning model into a case-based decisionmaking model (de Mantaras, 2001; Jiang and Sheng, 2009; Ram and Santamara, 1997). Instead of each agent having a single experience table, they now possess a library of such tables representing previous interactions (the cases), indexed by key salient features of each case. At the beginning of each interaction, the agents query their case library for the experience table associated with the most similar case, and use it to drive their decisionmaking for this interaction – in other words, when facing a new case, agents approach it based on the lessons learned from the most similar previous case they have faced. At the end of the interaction, this previous table is updated, and a copy of the updated table is added to the case library, indexed at the coordinates of the current case. Thus, the lessons learned from the current case are propagated back to the previous case as well. The more frequently a case is the most-similar one, the more information (encoded as move weights) the agent has about it. This captures the fact that actors' understanding of

the past is not static, and is updated based on their further experiences. While the agents in this model variant are still not 'strategic' in the game-theoretic sense, they are retrieving different past experiences based on the properties (and hence, expected behavior) of the agent with whom they are currently interacting.

More formally this decisionmaking model is as follows. Let  $c_i(j)$  encode a tuple of coordinates describing agent *i*'s interaction with agent *j*. Each agent *i* has a case library of coordinates and associated attraction functions  $C_i = \{(c_{i,0}, Q_{i,0}), (c_{i,1}, Q_{i,1})...\}$ . When presented with a new case  $c_{new}$ , the agent chooses  $c_* = argmin_{c \in C_i} dist(c_{new}, c)$  and the corresponding  $Q_*$ .  $Q_*$  becomes the active attraction function for the given interaction, as described above.

At the end of the interaction,  $Q_*^{t+1}$  is updated as described above. The overall case library is then updated such that  $c_* \mapsto Q_*^{t+1}$  and  $c_{new} \mapsto Q_*^{t+1}$ .

This decision model is implemented as follows. At the beginning of each interaction, the agent selects the active case, and corresponding experience table:

Listing 2.3: Case-Based Learning – Set Case

```
def set_new_case(coords):
```

```
self.current_coords = coords
self.current_case = self._get_nearest_neighbor(coords)
self.Q = self.case_library[self.current_case]
```

At the end of the interaction, the case library is then updated as follows:

```
Listing 2.4: Case-Based Learning - Update Weights
def add_utility(utility):
    for move in self.past_moves:
        self.Q[move] *= (1 - self.learning_rate)
        self.Q[move] += self.learning_rate * utility
        self.past_moves = []
        new_Q = self.Q.copy()
        self.case_library[self.current_coords] = new_Q
```

The coordinates used for the cases are as follows:

Small Crisis Model:

$$c_i(j) \equiv (\Delta m, \Delta a, 100\delta(b_i, b_j)) \tag{2.14}$$

where  $\Delta m = m_i - m_j$  and likewise for  $\Delta a$ , and  $\delta(b_i, b_j) = 1$  if  $b_i = b_j$ , and 0 otherwise.

**International Interaction Game:** 

$$c_i(j) \equiv (Stakes_i, Stakes_j, P^i) \tag{2.15}$$

Where  $Stakes_i$  is defined in Bueno de Mesquita and Lalman (1992) as  $U^i(\Delta_i) - U^i(\Delta_j)$ (see Equations 2.1 and 2.2), and  $P^i$  is defined in Equation 2.4.

# 2.5 Experiments and Results

We now have all the necessary components in hand: two models of agent interaction, and two new models of agent decisionmaking in addition to the baseline rational model. Now we can combine them and observe the results. I first test the Reinforcement Learning (RL) and Case-Based (CB) agent decisionmaking models with the Small Crisis Model and compare their performance to one another and to the game-theoretic equilibrium. I next apply the RL and CB models to the full International Interaction Game and the EUGene data, analyze the models' behavior against both the equilibrium outcomes and the observed events, and test the model outputs' power as predictive variables.

### 2.5.1 Small Crisis Model

Before we apply the decisionmaking models to historic data, it is worthwhile to test them in a simpler, notional environment. This allowed me to test and verify the decisionmaking models themselves; as the decision trees are relatively small, and the payoff calculations straightforward, it allowed me to manually verify that the decisionmaking code matched the model descriptions, and was behaving as expected. Furthermore, this model provides a baseline against which to judge the decisionmaking models' performance. If agents in the notional system, when using decision models other than game-theoretic rationality, behave very differently from the equilibrium expectation, this leads us to hypothesize that a similar divergence should exist in reality. If, however, the agents' behavior converges to something similar to the equilibrium, it would make the real-world decision models harder to distinguish, but potentially provide support for the use of the equilibrium model – as an approximation of state behavior rather than a direct description of it.

I carry out the experiments as follows. I instantiate Small Crisis models of 10 agents, as described in Section 2.3.1: first using the Static variants, then Dynamic one. For both model variants, I vary the learning rate  $\alpha$  from 0.1 to 1, in increments of 0.1. At each parameter value, I generate 100 model instantiations. Each instantiation is done using a fixed random seed, and generates two model runs: one with agents utilizing the Reinforcement Learning (RL) decisionmaking model, and one with the Case-Based Learning (CB) decisionmaking model. Using the same random seed ensures that the random elements are fixed between the model runs – in particular, that the agent properties are the same<sup>6</sup>. The models are instantiated with agents with the same properties; these agents will interact in the same order; even the random numbers drawn when agents choose between moves will be the same.

For each model run, I record the outcome of each interaction, as well as the equilibrium outcome for that interaction. In particular, I focus on the fraction of interactions within every simulated year where the interaction outcome matches the equilibrium outcome. I also define run quality as the mean of those fractions over the latter 50 steps of each run; this allows for a burn-in period as the agents gain initial information, discarding the early interactions where the agents' moves are approximately random. This metric allows for comparison between the behavior of the learning models and the rational model, and specifically measures the correspondence between them.

I begin with several hypotheses: first, as mentioned in Section 2.4, that across both the Static and Dynamic variants, both the RL and CB agents will converge to near-equilibrium

<sup>&</sup>lt;sup>6</sup>In some cases, individual interactions between agents result in different-length paths through the game tree, which will lead to a divergence in the sequence of random numbers.

behavior – that is, they will learn to play similarly to agents utilizing the purely rational decisionmaking model. I further expect that the CB agents will exhibit higher run quality (as defined above), since it allows agents to incorporate additional information and learn different weights – that is, different likely moves – for different interactions, implicitly incorporating a greater degree of strategic anticipation of the other agent's moves. Finally, in the Static variant, the mix of equilibria for each agent, and the equilibrium of each agent dyad, are fixed for the duration of each run, giving the agents more time to converge on the corresponding moves; in contrast, in the Dynamic variant these equilibria may change, meaning that the move weights learned from prior interactions may no longer be as likely to reach each interaction's equilibrium. Thus, I expect both the CB and RL decisionmaking models to perform worse in the Dynamic variant than in the Static one.

### Static Variant

Initially, let us examine the Static variant runs. Figure 2.5 shows the traces of the equilibrium frequency at each step across all the model runs, for both agent types. For the pure RL agents, we see a variety of behaviors emerging: in some cases, the models rapidly converge to near-equilibrium behavior, while in others appear to converge to a behaviors further away from equilibrium than we would expect solely due to chance. Figure 2.6 shows the distributions of the run qualities (mean fraction of interactions where the outcome matches the equilibrium) across the different runs. We immediately see that both distributions are skew normal. The CB model produces better (more consistently closer to equilibrium) behavior: its mean is greater, while its standard deviation and skew are both smaller. This is in line with my earlier hypothesis in Section 2.4, and implies that the CB model provides agents with more strategic behavior.

Since each model instantiation is run twice with the same seed, once with RL agents and once with CB agents, we can directly compare the equilibrium frequency for each seed and find the difference between the CB and RL mean equilibrium frequencies for each. The distribution of the differences is shown in Figure 2.7. This distribution is also approximately



Figure 2.5: Small Crisis (Static) – Traces

skew normal, with a mean of 0.18 (*p*-value indistinguishable from 0). This allows us to conclude that the case-based learning decisionmaking model produces behavior which is robustly better than the pure reinforcement learners when controlling for other factors – exactly as we would expect.

Another important question is whether different learning parameters yield substantially different behaviors. Figure 2.8 shows box-and-whiskers plots comparing the distribution of run qualities for each learning rate parameter. Overall, the differences are small, though statistically significant (with one- and two-way ANOVA tests yielding *p*-values indistinguishable from zero). Linear models of run quality to learning rate place a negative coefficient



Figure 2.6: Small Crisis (Static) – Quality Distribution



Figure 2.7: Small Crisis (Static) – CB and RL Differences

on the learning rate, indicating that lower learning rates tend to lead to higher expected qualities<sup>7</sup>. In other words, agents learn best when adjusting their behavior relatively slowly

<sup>&</sup>lt;sup>7</sup>Note, however, that there is a discontinuity at  $\alpha = 0$  where the initial weights would never change, meaning that all decisions will be made with equal probability.



Figure 2.8: Small Crisis (Static) – Quality by Learning Rate

and placing most of the weight on the total of their past experiences rather than each new one – again, an intuitive but satisfying result. Visual examination of the results in Figure 2.8 suggests that performance does not have a simple linear relationship with learning rate, and that both decisionmaking models have an optimal learning rate, which are different for each. The median quality for the RL model where the learning rate is 0.2 is substantially higher than the maximum of the inter-quartile range (the box) for the rest of the values, indicating that the maximum performance is attained somewhere near  $\alpha \approx 0.2$ . There is no similar obvious optimal learning rate for the CB model; a learning rate of 0.3 produces the highest median quality, but the distributions produced by values between 0.3 and 0.5 are essentially indistinguishable from one another, though higher than the distributions produced at the extreme values.

Even the best-performing model runs do not result in equilibrium play on all interactions; in general, the majority of model instantiations see equilibrium outcomes on the order of half the interactions. However, as we will discuss in more detail below, the majority of historic real-world interactions do not result in the equilibrium either. Suppose that the decisionmaking process of the model agents were an unknown black box, as real states' decisionmaking processes are. We could propose the game-theoretic equilibrium as a model of the agent decisionmaking, and test whether the equilibrium outcome is a useful predictor of the interaction outcomes. Following the methodology of Bennett and Stam (2000a), I use logistic regressions to test whether the equilibrium is indeed a good predictor of observed outcomes.



Figure 2.9: Small Crisis (Static) – Logistic Regression Coefficient Plots

The results of these regressions are reported in Table 2.3, and shown visually in Figure 2.9. In line with the results of Bennett and Stam (2000a), across both decisionmaking models, the equilibrium is a strong predictor of a corresponding observed outcome, while other equilibria have smaller, often negative, coefficients. This effect is consistently more

		. ,			
	StatusQuo	Capitulate1	Capitulate2	War1	War2
Capitulate1	-0.04***	$2.73^{***}$	$0.76^{***}$	-1.49***	$1.19^{***}$
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Capitulate2	-0.42***	$1.79^{***}$	$1.43^{***}$	-0.93***	-0.99***
	(0.00)	(0.01)	(0.00)	(0.01)	(0.02)
War1	-0.08***	-0.52***	-0.41***	$0.13^{***}$	$1.30^{***}$
	(0.00)	(0.02)	(0.00)	(0.00)	(0.01)
Const.	$1.05^{***}$	-5.63***	-2.55***	-1.55***	-4.79***
	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)

Table 2.3: Small Crisis (Static) – Logistic Regressions

(a) RL Model

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

(b) CB Model

	StatusQuo	Capitulate1	Capitulate2	War1	War2
Capitulate1	-0.44***	$4.28^{***}$	$0.70^{***}$	-1.36***	$1.88^{***}$
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Capitulate2	-1.08***	$0.70^{***}$	$3.51^{***}$	-1.98***	-3.09***
	(0.00)	(0.02)	(0.00)	(0.01)	(0.07)
War1	-1.74***	$0.46^{***}$	-0.36***	$1.67^{***}$	$0.96^{***}$
	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)
Const.	$0.48^{***}$	-5.35***	-3.23***	-0.70***	-4.89***
	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

powerful for the CB model, with larger equilibrium coefficients and smaller coefficients on other variables for each regression. In other words, when the agents are using case-based learning, the equilibria is a stronger predictor of their actual behavior than when the agents are using simple reinforcement learning. This is in line with the hypothesis that the CB agents will behave more rationally than the RL agents. Furthermore, outcomes closer to the equilibrium in the game tree are stronger predictors than further outcomes, suggesting that the agents are playing sequences of moves that are close to the equilibria sequences, deviating only in the step.



Figure 2.10: Small Crisis (Dynamic) – Traces



Figure 2.11: Small Crisis (Dynamic) – Quality Distribution



Figure 2.12: Small Crisis (Dynamic) – CB and RL Differences

# **Dynamic Variant**

Next, I repeat the above experiment and analyses with the dynamic model variant, where agent properties change from step to step. Figures 2.10 and 2.11 show the traces and

distribution of outcomes, respectively. We can immediately see that overall the agents are less likely to play the equilibrium outcome compared to the Static variant, across both decisionmaking models. This is exactly in line with our intuitive expectation – since each dyad of agents do not have a fixed set of moves to reach equilibrium, and indeed the mix of equilibrium moves across the entire population is changing, the lessons of previous interactions will be less applicable to future ones, making moves based on the learned weights less likely to result in equilibrium. Examination of the outcomes across different learning rates, as shown in Figure 2.13 suggests that they do not yield significant differences. This, in turn, suggests that the speed at which agents adjust their behaviors does not substantially affect the behavior the model converges to.



Figure 2.13: Small Crisis (Dynamic) – Quality by Learning Rate
Table 2.4: Small Crisis (Dynamic) – Logistic Regressions

	StatusQuo	Capitulate1	Capitulate2	War1	War2
Capitulate1	-0.13***	2.13***	0.47***	-0.51***	1.30***
	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)
Capitulate2	-0.58***	$1.39^{***}$	$1.47^{***}$	$0.02^{***}$	-0.35***
	(0.00)	(0.02)	(0.01)	(0.01)	(0.03)
War1	-0.24***	$-2.17^{***}$	-0.21***	$0.42^{***}$	-2.46***
	(0.00)	(0.08)	(0.01)	(0.01)	(0.07)
War2	$0.51^{***}$	$0.72^{***}$	-0.65***	-0.95***	$1.32^{***}$
	(0.01)	(0.03)	(0.01)	(0.01)	(0.02)
Const.	$2.49^{***}$	-6.33***	-4.18***	-2.82***	-5.64***
	(0.00)	(0.02)	(0.01)	(0.00)	(0.01)

(a) RL	Model
--------	-------

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

	StatusQuo	Capitulate1	Capitulate2	War1	War2
Capitulate1	-0.17***	$1.26^{***}$	-0.57***	-0.29***	0.87***
	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)
Capitulate2	-0.37***	$0.15^{***}$	$1.22^{***}$	-0.33***	-0.17***
	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)
War1	-0.88***	-1.20***	$0.92^{***}$	$0.62^{***}$	-0.69***
	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)
War2	-0.49***	$1.05^{***}$	-1.21***	$0.17^{***}$	$1.16^{***}$
	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)
Const.	$0.05^{***}$	-3.28***	-2.32***	-0.77***	-3.09***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

(b) CB Model

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

A more surprising result is that the CB decisionmaking model does not outperform the RL model. Figure 2.12 shows the distribution of differences between the model qualities; it has a mean of -0.06, indicating that on average the CB agents perform very slightly worse than the RL agents. This is confirmed by the logistic regressions, reported in Table 2.4 and shown visually in Figure 2.14. As with the regressions done on the Static variant results, the equilibria are still positive predictors of a corresponding outcome. Unlike with



Figure 2.14: Small Crisis (Dynamic) – Logistic Regression Coefficient Plots

the Static variant, however, the equilibrium is not a stronger predictor of the outcome of the CB models than of the RL models. This contradicts the hypothesis that the CB decisionmaking model is broadly superior to the the RL decisionmaking model in terms of approaching rational play.

In the previous model variant, the CB agents were implicitly learning best moves against each particular agent. In this variant, since agent properties are changing, there is no single best response to be learned. Additionally, with ever-changing agent properties, there are more potential coordinates populating the case history space more densely. The equilibrium moves exhibit discontinuities across this space; however, the updating rule, whereby a new interaction updates both its own coordinate point and those of the nearest previous case, mean that moves appropriate for one equilibrium but not the other may propagate across these discontinuities – essentially, agents consider two cases more similar than they actually are.

#### 2.5.2 International Interaction Game

The previous, simple model allowed me to verify that the decision models were implemented properly, with the code corresponding to the model description. More substantively, the results reported above suggest that multiple agents are capable of learning to play extensiveform games against each other in a manner approaching rational behavior, in line with my initial hypothesis. Furthermore, the difference in performance of the RL and CB models between the static and dynamic variant complicates expectations for the real world, which is itself dynamic. With these results in mind, I move from the notional, simplified international system to the full International Interaction Game (IIG) with the EUGene data, as described in Section 2.3.2.

I carry out these experiments as follows. For each of the RL and CB decisionmaking models, I instantiate 105 model runs<sup>8</sup>, with all agents utilizing the same decisionmaking model with fixed learning rates: 0.1 for the RL model, and 0.2 for the CB model. Each instantiation includes one agent for each unique country in the EUGene dataset, and is run as described in Section 2.3.2. I record the outcome for each interaction from each instantiation, and compare them to the equilibrium outcomes and observed events. There are thirteen possible outcomes of the extensive-form game tree, meaning that if agents were to play at random there would be a 7.7% chance of observing an outcome equal to the equilibrium. Figure 2.15 shows the fraction of interactions by year ending in the IIG equilibrium outcome. There are several things immediately obvious: as with the simplified model, we never see convergence to perfect equilibrium behavior. Unlike with the simplified model, however, there is not an obvious rapid initial burn-in period followed by relative stability. Instead, there is a broad range of traces showing different rates of improvement, with occasional sharp dips and upward spikes.

It appears that the rate of improvement across the majority of CB runs is initially greater than for the RL runs; in particular, all CB runs rise above the fraction of equilibrium outcomes expected by chance earlier than the RL runs. However, there is not a clear

<sup>&</sup>lt;sup>8</sup>The non-round number is due to model runs being executed in parallel on 15 cores on 16-core servers.

stable performance rate reached for either decisionmaking model, with multiple substantial shifts in the rate of equilibrium play. In fact, across all the different runs, these shifts occur at similar points, in some cases in different directions. Note in particular the very obvious effects of the two world wars. Those conflicts themselves are outside the scope of the IIG and its variants: they were heavily multilateral, took place across multiple years, and radically changed the states present in the international system and alliance network between them. It appears to be the case that the shocks of the world wars substantially changes the distribution of equilibrium moves for the different agents. In some runs, this leads to a significant drop in equilibrium performance, while in others the agents either successfully rapidly adapt or have previously learned distributions of behaviors which prove to be improvements in the new environment. In contrast, following World War II, there is no similar improvement, with the majority of traces demonstrating rapid deterioration in equilibrium play.

Across both the pure reinforcement learning and case-based learning decisionmaking model variants, there appear to be several overall attractors, with multiple model instantiations demonstrating similar behavior. These attractors are visible in Figure 2.15 as the darker regions where multiple line traces come together. Both variants have one attractor with a fraction of equilibrium play at approximately 0.25, which emerges early on. Across both variants, this attractor shows a sharp falloff approximately during the 1930s, then rebounds somewhat, and stabilizes from 1950 until the end of the model timeframe. Both variants also exhibit a second attractor which emerges later, at a substantially higher equilibrium fraction. In the RL variant, the second attractor emerges most during the mid-1930s; in the CB variant, it emerges earlier, beginning from the mid-1860s onward, though the attractors' traces converge to a tighter band after 1950 here as well. Note, however, that even the lower attractors show substantially better mean performance than if the agents were choosing moves at random.

These results indicate that the agents are learning to play the equilibrium strategies with one another over time, albeit very imperfectly. The next question, then, is whether

this behavior resembles the observed behavior of real actors. Figure 2.16a shows the annual fraction of dyads by year where the observed event, as recorded by the COW Militarized Interstate Dispute data (Palmer et al., 2015), matches the IIG equilibrium outcome. As we can see, the equilibria outcomes are rarely observed directly – the overall mean across all dyads is 12%, or only slightly better than random. It is important to keep in mind that the vast majority of these dyads exhibit Status Quo outcomes and are not politically relevant, involving distant states with few if any directly overlapping interests or ability or desire to start a conflict with each other. Similarly, the majority of model runs yield very small fractions of correctly-predicted outcomes: only a mean of 2% for the RL decisionmaking model (worse than random), and 7% (approximately the same as by chance) for the Case-Based decisionmaking model, across all model runs for each. Figure 2.16 shows the fractions of annual dyads where the outcome matches the observed event for all model runs. While the majority of model instantiations generate extremely low rates of correct outcomes, there are some where the rate of correctly-predicted outcomes is substantially higher by the end of the run. Only 18 of the RL instantiation runs (17%) generate the observed outcomes at a higher rate than random by the final year of the data, and 3 do so in over half of all interactions. Better, 26 of Case-Based runs (25%) generate better-than-random predictions, while 15 (14%) generate the observed outcome over half the time. The traces in these cases indicate that the rate of observed outcomes is high across the majority of the run history, and tends to increase over time. This suggests that the higher accuracy is a robust phenomenon within these instantiations, rather than a statistical fluke within a single year. While the fraction of runs generating these predictions are a minority, they are not insubstantial. These results indicate that the observed outcomes, in the aggregate, can be derived from one of the sets of behaviors the agents converge to across several sets of decisionmaking models and instantiations. This, in turn, suggests that the observed outcomes are satisficing: they yield utilities which encourage all the agents to continue, with increasing frequency, to play the moves which will lead to those outcomes.



(b) Case-Based Learning Model

Figure 2.15: Fraction of Interactions with Equilibrium Outcome



Figure 2.16: Fraction of Interactions with Correctly-Predicted Outcome

With these results in hand, we can proceed to test whether the model outcomes are a statistically useful predictor of the observed events. As demonstrated in Figure 2.16a, the IIG equilibrium is not a direct prediction of observed events; instead, Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000a) use it as a statistical predictor, arguing that estimating a particular equilibrium makes it more likely (though far from certain) that we will observe a corresponding event in the historic data. I apply a similar test to the model outputs. Since the learning models take some number of interactions for their initial learning, it is sensible to drop these early interactions from our analysis as a burn-in period. I could not identify a formal method to determine how long this period ought to be. Based on the traces in Figure 2.15, the main attractors across the decision models begin to appear at approximately 1850; for this reason, I choose this year as the cutoff for the burn-in period.

I start by replicating the findings of the prior studies. Table 2.5 shows the results of logistic regressions with the observed events of each post-1850 dyad as the dependent variables and the IIG equilibria as the independent variables <sup>9</sup>. This is a simplified model, without many of the controls (e.g. for polity type, or length of peace) introduced by Bennett and Stam (2000a). These regression results are comparable to those reported in Bennett and Stam (2000a), providing moderate support for the hypothesis of the IIG equilibrium as an outcome predictor prior to adding in the various controls. More importantly, these results provide a baseline against which we may compare the results of similar regressions using the decisionmaking model interaction results.

Tables 2.6 and 2.7 show the results of the same logistic regressions on the outcomes of the post-1850 interactions in the RL and CB models, respectively. The coefficients for both models as well as the equilibrium model are shown visually in Figure 2.17, below. We can see that the learning model coefficients are similar to those of the equilibrium models, showing moderate support for the hypothesis that the model outcomes are a predictor of the observed

<sup>&</sup>lt;sup>9</sup>Strictly speaking, the correct regression for this case is a multinomial logit, as in Bennett and Stam (2000a), since the outcomes of each interaction are mutually exclusive. However, the datasets produced by the 105 instantiations of each model are too large for a multinomial logit to be computationally tractable. In order to keep the model tractable, I treat each outcome as an independent variable. This method provides a worse fit, but allows the coefficients to be interpreted the same way.

outcome. The level of support varies between outcome types. For Acquiesce events, the coefficients on the corresponding model outcomes are significant and positive, while those on the opposite outcome (the acquiescence of the other party) are negative. However, this is not the case for Capitulate events, where capitulation outcomes in both directions have a positive coefficient. Importantly, War and Negotiation outcomes are both positive predictors of observing corresponding events. However, War outcomes are also predictors of observing Negotiation events – which suggests that the model is correctly capturing cases where a conflict is likely, but that some of those conflicts are in fact resolved via negotiation rather than war. Similarly, note that the Capitulation outcomes are also positive predictors of War events, and vice versa. These outcomes are adjacent to one another on the game tree, which provides support for the structure of the game tree itself. It suggests that the model agents are playing down to the correct subgame of the game tree, with one agent making a single move different from what was observed.



Figure 2.17: Logistic Regressions Coefficient Plots

	Capitulate A	Capitulate B	Acquiesce A	Acquiesce B	Negotiation	War	StatusQuo
Acquiesce A	-0.40***	-0.54***	$0.54^{***}$	-0.45**	-0.92***	-0.34***	0.34***
	(0.13)	(0.12)	(0.17)	(0.22)	(0.24)	(0.07)	(0.05)
Acquiesce B	0.08	-0.29***	-0.73***	$0.34^{**}$	-0.41**	-0.54***	$0.31^{***}$
	(0.11)	(0.10)	(0.23)	(0.17)	(0.19)	(0.07)	(0.05)
Negotiation	0.01	-0.38***	-0.08	0.06	-0.56***	-0.20***	$0.20^{***}$
	(0.08)	(0.07)	(0.14)	(0.14)	(0.14)	(0.05)	(0.03)
War A	-0.48***	-0.52***	-0.19	-0.52***	-0.86***	-0.57***	$0.53^{***}$
	(0.11)	(0.09)	(0.17)	(0.18)	(0.19)	(0.06)	(0.04)
War B	0.81	0.44	-24.50	$1.79^{*}$	$2.36^{***}$	0.72	-0.98***
	(1.00)	(1.00)	(512568.77)	(1.01)	(0.72)	(0.58)	(0.36)
Const.	-6.52***	-6.16***	$-7.51^{***}$	$-7.51^{***}$	-7.38***	-5.33***	$4.59^{***}$
	(0.08)	(0.06)	(0.13)	(0.13)	(0.12)	(0.04)	(0.03)

Table 2.5: IIG Model Equilibrium – Logistic Regressions

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

	Capitulate A	Capitulate B	Acquiesce A	Acquiesce B	Negotiation	War	StatusQuo
Acquiesce A	-0.17***	-0.15***	$0.46^{***}$	-0.43***	-0.14**	0.02	$0.04^{***}$
	(0.03)	(0.03)	(0.04)	(0.05)	(0.06)	(0.02)	(0.01)
Acquiesce B	-0.05**	$0.50^{***}$	-0.42***	$0.47^{***}$	0.02	$0.40^{***}$	-0.29***
	(0.02)	(0.02)	(0.04)	(0.04)	(0.05)	(0.01)	(0.01)
Capitulate A	0.02	$0.41^{***}$	$0.49^{***}$	-0.31***	-0.14	$0.71^{***}$	-0.44***
	(0.05)	(0.05)	(0.07)	(0.10)	(0.12)	(0.03)	(0.02)
Capitulate B	0.01	0.13**	-0.73***	$0.21^{***}$	-0.49***	$0.36^{***}$	-0.16***
	(0.05)	(0.06)	(0.13)	(0.08)	(0.14)	(0.03)	(0.02)
Negotiation	$0.17^{***}$	$0.45^{***}$	$0.26^{***}$	0.20***	$0.32^{***}$	0.44***	-0.36***
	(0.02)	(0.02)	(0.03)	(0.03)	(0.04)	(0.01)	(0.01)
War	0.31***	$0.62^{***}$	0.41***	$0.56^{***}$	$0.44^{***}$	$0.58^{***}$	-0.52***
	(0.02)	(0.02)	(0.03)	(0.03)	(0.04)	(0.01)	(0.01)
Const.	-6.80***	-6.99***	-7.85***	-7.83***	-8.22***	-6.06***	$5.21^{***}$
	(0.02)	(0.02)	(0.03)	(0.03)	(0.04)	(0.01)	(0.01)

Table 2.6: Reinforcement Learning Model – Logistic Regressions

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

	Capitulate A	Capitulate B	Acquiesce A	Acquiesce B	Negotiation	War	StatusQuo
Acquiesce A	-0.02	-0.11***	$0.56^{***}$	-0.26***	$0.15^{***}$	-0.02	-0.01
	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)	(0.01)	(0.01)
Acquiesce B	$0.07^{***}$	$0.45^{***}$	-0.29***	$0.54^{***}$	-0.02	$0.27^{***}$	-0.26***
	(0.02)	(0.02)	(0.04)	(0.03)	(0.04)	(0.01)	(0.01)
Capitulate A	-0.17**	$0.31^{***}$	$0.53^{***}$	-0.27**	-0.30**	$0.47^{***}$	-0.29***
	(0.07)	(0.06)	(0.09)	(0.13)	(0.15)	(0.03)	(0.02)
Capitulate B	$0.14^{**}$	$0.11^{*}$	-0.54***	$0.74^{***}$	-0.33**	$0.23^{***}$	-0.18***
	(0.06)	(0.06)	(0.15)	(0.08)	(0.15)	(0.04)	(0.03)
Negotiation	0.32***	$0.38^{***}$	$0.44^{***}$	0.32***	0.33***	0.34***	-0.35***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)
War	0.44***	$0.55^{***}$	0.39***	$0.75^{***}$	$0.44^{***}$	0.41***	-0.47***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.01)	(0.01)
Const.	-6.91***	-6.88***	-7.97***	-7.91***	-8.22***	-5.92***	$5.17^{***}$
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)	(0.00)

Table 2.7: Case-Based Learning Model - Logistic Regressions

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

## 2.6 Discussion

In this chapter, I have presented an attempt to use agent-based modeling as a bridge between several schools of thought and lines of research in the study of states' decisionmaking in the face of international conflicts. The ABM perspective allows us to build on a previous, wellstudied model of international relations, and to view the rational-choice, equilibrium-playing behavior as only one of many possible models of decisionmaking which may be implemented and tested within a computational framework. I have presented two such decisionmaking models, which are both simplified but psychologically and organizationally plausible. I then used these models to drive the decisionmaking of agents within two simulated international systems: a simplified notional one, and one based on historic data and a well-validated interaction model.

The notional system allows us to focus on the properties of the decisionmaking models themselves by controlling the complexity of the system. Across both the static and dynamic world variants, we see that both the reinforcement learning (RL) and case-based learning (CB) decisionmaking models converge to sets of behaviors that play the rational equilibrium strategies far more than randomly, though far less than perfectly. Close examination of interactions and model runs, as well as preliminary work not presented here, indicates that the agents play close to equilibrium when the equilibrium utilities are substantially greater than those from the non-equilibrium outcomes; when these utilities are closer, convergence is reduced. This suggests an important observation: these decisionmaking rules are producing satisficing outcomes. They collectively find distributions of strategies that yield positive utilities, and thus have little incentive to explore and arrive at the equilibrium strategy with only slightly higher payoffs.

Examination of the model traces shown in Figures 2.5 and 2.10 suggests that the convergence to near-equilibrium play happens either very quickly or not at all. This provides validity to the model's 'cold start' assumption of having all agents start each run with no information or initial weights, and for discarding only a relatively small fraction of early interactions as a burn-in period. This, in turn, helps address the objection that the real decisionmakers in the early periods covered by the model already had previous experience to draw on, unlike the model agents. Furthermore, comparison of the equilibrium play by learning rate shown in Figures 2.8 and 2.13 suggests there is not a qualitative difference between the performance of the agents at the different rates. This reduces the need to repeat the IIG experiments with different learning rates, which is limited by the time and computational resources required for each run.

Finally, and somewhat surprisingly, the the CB model does not consistently outperform the context-free RL model. In the static world variant, the CB model does generally perform better than the RL model, as it allows agents to learn a particular set of weights to use against each other agent. However, in the dynamic world variant, the changing agent properties generate larger case libraries for each agent, which in turn leads to a higher probability of querying a case with a different equilibrium than the current one. While this may appear at first to be an artifact of the model specifications, it may in fact also be capturing a feature present in real-world reasoning: that the most-similar historic analogy to a particular case may not in fact be the correct analogy, in that the lessons learned from one will not lead to the desired or expected outcome in the other.

The International Interaction Game variant offered a greater number of agents, with greater interaction complexity and less structured changes in agent properties and potential payoffs. Despite this, both decisionmaking models again rapidly yield a higher than random share of interactions with an equilibrium outcome. Beginning with no initial weights or other information, the agents appear to rapidly learn weights (essentially mixed strategies) which approach the equilibrium. Whereas the Small Crisis model saw an approximately normal distribution of traces and performance across the quality space, the IIG seems to exhibit several attractor regions which recur across runs and between decisionmaking models. This suggests that the structure of the international system (at least as captured by the data and IIG calculations) drives the agents to learn particular distributions of strategies – and that these strategy sets are satisficing. In particular, the fact that there are two attractors at different quality levels suggests that there are two sets of satisficing strategies that the

agents converge to, with one much closer to the equilibria than the other.

Both the RL and CB models are underpredictive of overall observed events – as, indeed, is the baseline rational model. A large part of the reason for this is that the vast majority of interactions are politically irrelevant dyads between distant countries; even if the IIG utility calculations indicate that one actor has utility to gain from threatening force against the other, this utility gain is largely abstract, with no actual concrete disputes which may lead to war. For this reason, the majority of RL an CB runs, like the rational IIG model, overpredict non status quo outcomes. However, an advantage that the agent-based model has over the game-theoretic original is that its stochasticity produces a range of traces – and, as shown in Figure 2.16, some of those are discovered to be far more predictive. This suggests that the set of strategies that are more likely to produce interaction outcomes actually matching observed events are satisficing as well, despite being far from the equilibrium strategy set. The CB decision model appears to have an advantage over the RL model in terms of both its overall rate of correctly generating observed events, and in terms of the number of iterations consistently generating correct outcomes.

Applying the logistic regression methodology used by Bennett and Stam (2000a) to test the explanatory and predictive power of the different models indicates that the RL and CB models are in line with the original rational model, and do not have radically stronger or weaker predictive power in the aggregate. There are no event categories for which the sign or significance on the corresponding coefficients are different, or where there is a substantial difference in the magnitude of those coefficients. There is no event which is obviously better predicted by the equilibrium than the model results, and the Negotiation equilibrium has a negative coefficient on the corresponding event, in contrast to the positive coefficients both learning models produce.

With these results in hand, we may directly address the underlying question of crisis decisionmaking when addressing potential international conflicts. Multiple agents using both the context-free reinforcement learning decisionmaking model and the case-based learning model do not converge to perfectly rational equilibrium behavior across multiple models of the international system. Nevertheless, the behaviors and outcomes emerging from the agents' interactions often approach the equilibria. If the real-world actors are making decisions in ways that resemble the RL and CB models, these results suggest they will converge to rational-like behavior. This supports the assertion that the assumption of rationality is useful as a simplified model, even if it is not descriptive of the procedure by which decisions are made – a specific case of the broader argument put forward by Friedman (1953). However, the rational model does not appear to be a better predictor of observed events than the learning models. Thus, there does not appear to be strong support for a hypothesis that the rational equilibrium provides a better predictor than the alternative model outcomes; or more generally, for the hypothesis that agents exhibiting perfect rationality provide a better model for states than agents utilizing even simple forms of learning. This means we do not need to rely on the assumption of rationality – we may implement more plausible models of decisionmaking and achieve similar predictive results

Furthermore, the learning models have an advantage over the purely rational model in that they can converge to sequences of outcomes that are non-equilibrium but nevertheless satisficing, under the given utility assumptions. Importantly, we see that the models may converge to outcome sequences which are a better match for historic observations. This provides evidence that, under the IIG's utility calculations, the observed interactions between states may themselves be satisficing, despite often being non-equilibrium outcomes. Finally, this in turn suggests that the learning models provide a plausible explanation for the difference between the equilibria and observed outcomes. If the IIG provides a valid estimate for actors' preferences between outcomes and approximation of the process by which crises escalate, models of decisionmaking driven by past experiences are capable of producing (and hence, explaining) observed history in a way that a strategic rationality model is not.

The work presented in this chapter suggests several opportunities for extension and expansion. There are additional decisionmaking models which could be implemented with relative ease, and tested using the same methods. I conducted initial tests of a learning model in which all agents shared a single case library, allowing them to learn from each other's experience; while this variant appeared to produce greater convergence to equilibrium, its computational costs prevented the full experiments needed to include it here. Other decision models may move away from the learning paradigm entirely in favor of simple rules (e.g. tit-for-tat), expert-derived heuristics, or other methods. I also conducted initial tests with relaxing the 'cold start' assumption. Alternatives include training agents with the correct outcomes (either equilibrium or historic) over an early subset of cases, as well as having only a subset of agents utilizing a learning model while the others are rational. Both training to equilibrium and adding rational-playing agents appeared to generate a higher rate of equilibrium play – though, as we have covered, this is unlikely to improve the models' fit to data.

Future work may also have the model code itself performing the computations for the utility estimations which are currently provided by EUGene. Methods for this are present in the code already, but were not used here – both due to the computational costs of these operations (particularly risk tolerance estimation, as detailed in Footnote 3) and in order to maintain consistency with previous work. The model's object-oriented architecture allows for alteration, and experimentation with, all the elements of the various sub-models, providing a wide range of additional hypotheses to be tested. These may involve something as simple as replacing tau similarity with a different measure, changing the payoff estimations, utilizing a different game tree, or even endogenizing the change in agent properties based on the outcome of each interaction. This means that the code presented here is not simply a tool for testing a few simple models of state decisionmaking, but the basis for a framework for much broader data-anchored agent-based models of international crises, and international relations more broadly.

# Chapter 3: BDM's EUM as ABM: Agentizing the Expected Utility Model

### 3.1 Introduction

In a series of papers and books across several decades, Bruce Bueno de Mesquita (commonly referred to as BDM) has presented and applied several versions of a forecasting model to predicting everything from internal Iranian power struggles (1984) and European regulation negotiations (1994) to the course of the entire Cold War (1998). The model is intended to simulate the iterative negotiation process between multiple actors, identifying where they compromise and where they attempt to coerce one another. It has attracted widespread attention (particularly by the standards of political science models): it was the subject of a special journal issue (Kugler and Feng, 1997), and was featured in a popular-audience TED talk (Bueno de Mesquita, 2009) and book (Bueno de Mesquita, 2010). It can be used not only for prediction but also decision support (Özer-imer, Kugler, and Root, 2013), and has reportedly been successfully applied by the Central Intelligence Agency (Feder, 1992), as well as numerous commercial clients of Decision Insights, Incorporated (a consulting firm BDM co-founded). The model was even the subject of a lawsuit by that company against a rival (US 4th Circuit, 2011) accused of stealing the model.

Such a universal political forecasting model is worth examining in detail. BDM and his collaborators have described the model extensively (Bueno de Mesquita, 1984, 1994, 1997, 2002, 2011), but have never published the full specifications, and in particular never made the underlying code available. Several attempts have been made to replicate it from published material (Kimbrough, 2014; McKibben-Sanders, 2014; Scholz, Calbert, and Smith, 2011), but none of these have successfully fully reproduced its reported results. Indeed,

as Scholz, Calbert, and Smith (2011) document, some details of the model appear to have changed over time<sup>1</sup>, and testimony in the US 4th Circuit (2011) lawsuit states that the commercial version of the model contains trade secrets not reproducible from the public literature. There have also been several attempts to develop similar models: these include an internal CIA variant (Feder, 1992); the Senturion toolkit (Abdollahian et al., 2006) (the competitors' model which triggered the lawsuit mentioned above); at least one more proprietary model that I am aware of; and the Wise, Lester, and Efird (2015a) Toolkit for Behavioral Analysis. Of these, only the last is open source; the rest are proprietary commercial products.

BDM has characterized the model as grounded in micro-economics (Bueno de Mesquita, 1994) and game theory (Bueno de Mesquita, 2002, 2009, 2010), and often refers to it as the Expected Utility Model (EUM) (Bueno de Mesquita, 1984) – borrowing a key term from game theory and economics. Indeed, the EUM bears many of the hallmarks of a formal model. Scholz, Calbert, and Smith (2011) describes the model's underlying interactions as an extensive-form game tree. Several of the model's internal specifications closely resemble those used in the International Interaction Game (IIG) (Bueno de Mesquita and Lalman, 1992), the game-theoretic model of international conflict described in the previous chapter, and Bueno de Mesquita (2011) proposes an update to the model which directly incorporates the IIG's extensive-form tree as well. Yet despite the game-theoretic components, the model bears some clear differences from conventional game theory. It is intended to model and predict particular real-world interactions rather than to identify general patterns and theories (Snidal, 1985). None of the papers discussing the model have analyzed it in closed form to identify and parameterize its equilibria, and it is unclear whether such equilibria exist, or are meaningful. Indeed, the problem of identifying formal equilibria of games with more than two players becomes rapidly intractable (Papadimitriou and Roughgarden,

<sup>&</sup>lt;sup>1</sup>These changes suggest that there is not a single 'original' model to reproduce. Nevertheless, as BDM does not clearly demarcate different generations or variants, I will use phrase 'original model' in this chapter to refer to my overall understanding of the model. My own implementation primarily followed the description provided in Bueno de Mesquita (2002).

2005). Thus, unlike in the majority of game theory models, the players are not assumed to be perfectly rational; they are explicitly boundedly-rational, using a variety of heuristics and looking at most one step ahead (Bueno de Mesquita, 2002). Finally, the model is implemented computationally, with the overall model, and each actor within it, executing set calculations and methods sequentially, and updating accordingly. While the model is largely described in terms of the specific calculations the agents make, I argue that these calculations are best viewed as particular sub-models rather than the model itself. In fact, another model, the Stokman and Van Oosten (1994) logrolling model, which takes the same inputs as the EUM and is intended for comparison with it, is explicitly described as an object-oriented computational model. Similarly, Özer-imer, Kugler, and Root (2013) also describe the model as agent-based. Understanding the model as an ABM emphasizes that the numeric calculations are not necessarily the most important components of it; the particular agent behavior rules, and the way the model sequences them, are equally important. Furthermore, treating the model as an ABM suggests additional methodological tools we can use to analyze it, such as sensitivity analyses and parameter sweeps (Gilbert and Troitzsch, 2005).

Unlike the International Interaction Game presented in the previous chapter, the EUM cannot be viewed as a direct translation of a particular qualitative theory into model form. To be sure, several components of the model are grounded in specific theories. The definition, and significance, of the median voter position are drawn from the Black (1948) economic analysis of collective decisionmaking; the risk propensity calculation, meanwhile, is adapted from Prospect Theory (Kahneman and Tversky, 1984) as detailed further in Section 3.2.2. Beyond that, however, the model follows general assumptions from Realist theories – the political actors are unitary entities, rationally maximizing their utility and cooperating only inasmuch as such cooperation is in their immediate interest – but its details (such as the expected utility calculations and the offer-based interactions) cannot be traced to particular qualitative descriptions of state interactions, or political interactions more broadly. Put another way: the model is itself a theory of political interaction, competition

and decisionmaking.

This chapter attempts to address several inter-connected questions and issues. The EUM has been utilized in multiple papers and developed a receptive audience, making it interesting to study and replicate, in order to test its accuracy independently, understand its driving assumptions, and potentially enable us to use it in future research. In the course of this replication, I argue that the model is best understood as an agent-based model, making it an important example of an ABM with predictive power and both academic and policy relevance. Such models are rare in general, and all the more so in the realm of international relations. Finally, I argue that the model can serve as the basis for comparative modeling of political conflicts. By using fixed inputs and modifying the assumptions and theories embedded in the model code, we can directly compare the effects implied by these theories, and test their explanatory and predictive power. If a model variant utilizing one theory produces more accurate outputs than one using a different theory, this is evidence that the theory is a better representation the system being modeled. Such comparison is particularly important given the model's numerous sub-models and weak grounding in prior theories, as it allows us to test whether the assumptions the sub-models make contribute to or detract from the model's overall explanatory and predictive power.

I have attempted to reimplement the Expected Utility model as an object-oriented, agent-based computer simulation. In this chapter, I restate the model following the ODD (Overview, Design concepts, Details) outline (Grimm et al., 2006), highlighting the entities, variables and sub-models from which the model is assembled (Section 3.2). I describe my attempt to validate the model against previous results (Section 3.3.1) and conduct a deep dive into a single model step to describe and verify its behavior (Section 3.3.2). I then discuss the assumptions embedded in the model structure and implementation (Section 3.4), and propose several sub-model variations which can be easily swapped into the model implementation. In Section 3.5, I describe a sensitivity analysis I conduct using artificial input data sets. In Section 3.6 I describe the output data the model generates, which can be compared to empirical data or used for prediction. Finally, In Section 3.7, I apply the

model in a novel way, implementing several of my proposed model variants and comparing their outputs to micro-level political event data.

## **3.2** Model Description

The following section is written approximately following the ODD protocol (Grimm et al., 2006). I first provide an overview of the model (Section 3.2.1), including its purpose, and its agents and variables; I describe the sub-models, introduce a notation system to concisely reference different sub-model variants, and describe the computational structure used to implement it. Section 3.2.2 then describes the model in more detail: how it is instantiated, and the specifics of the several sub-models which make it up. Unlike in an orthodox ODD description, I leave the model design concepts for Section 3.4, as they will be the points of study for the different model variants.

#### 3.2.1 Overview

The primary objective of the model is to predict the outcome of some 'issue' which multiple different actors conflict on, or are negotiating over. It does this by simulating the iterative process by which these actors attempt to threaten or coerce one another into adopting their own preferred position, update their declared positions in response to these threats, and come into conflict with one another. The issue under consideration may be as broad as the outcome of the Cold War, as in Bueno de Mesquita (1998), or as narrow as the chosen size of the Italian deficit (Feder, 1992). Nevertheless, each instantiation of the model is intended to focus on one particular issue, however defined.

Each model instantiation consists of several actors, all endowed with capabilities they can use to coerce each other or defend themselves from coercion, holding a position on the issue space, and possessing a salience value representing the importance they attach to the issue. The issue space is represented as a continuous one-dimensional line segment; between the maximum and minimum positions, any position represents a possible outcome, which agents can seek. For the sake of convenience, I normalize the issue space for all instantiations to the [0, 1] range.

Each agent is defined by several values, as summarized in Table 3.1. I particular, they have three key values: position, salience and capability.

- **Position** refers to the issue outcome or stance this actor publically supports at the moment, as projected onto the one-dimensional issue space. Agents may change their position as part of each step.
- Salience represents the importance the actor attaches to this particular issue, compared to the other actors. In practice, this refers to the fraction of coercive resources (capability) the actor is willing to devote to this particular issue. Salience is in the (0, 1] range. It will also inform the other agent's belief of how likely the actor is to yield to threats. Unlike position, salience is not updated endogenously, though in some model variants it may change stochastically to simulate exogenous shocks.
- **Capability** represents the actor's coercive capacity, as relevant to the issue at hand. In some cases, it may be military force; however, in cases such as commercial negotiations, financial or regulatory influence may be more important; and in some cases, social capital or persuasive power may be the most relevant source of influence. Capability is continuous, and linearly additive, and the ratio of capabilities on either side of a bilateral conflict determines the probability of either side winning or losing. Since capability values only comes into play as ratios, they can be represented on any positive scale.

In addition, agents have several variables governing their decisionmaking and utility calculations. Key among these is risk propensity, which is used in utility calculations; risk propensity is a derivative property, determined each step by the position of the agent relative to all other agents. Two additional variables model the agents' beliefs about the future if they do not act.  $Q_i$  is agent *i*'s believed probability that if they do not threaten an agent, that agent will nevertheless change its position.  $T_i$  is the agent's belief in the probability that, if that other agents does change position, it will be towards the current median

Table 3.1: Agent Variables

Variable	Description
$x_i$	Current position of agent $i$
$s_i$	Salience of agent $i$
$c_i$	Capability of agent $i$
$r_i$	risk propensity of agent $i$
$Q_i$	Agent $i$ 's belief that any other agent will move absent an offer from $i$ .
$T_i$	Belief that other agent's movement will be towards median.

position. These last two variables are generally constant and identical for all agents in a model instantiation. Values of 0.5 indicates maximum uncertainty, while 1 and 0 indicate certainty that agents will or will not move, and that that movement will be towards the median, respectively.

The model treats all conflicts as bilateral; however, other agents not directly involved in a particular bilateral conflict may contribute some of their own capability to one side or another, as specified in Equation 3.8. Each side of the conflict then has some total capability behind it, and each side's probability of victory is its share of the total capability contributed by all actors to the conflict (see Equation 3.9).

The model proceeds in discrete rounds, or steps. During each step, all actors may make 'offers' (which I will also refer to as threats) to one another; the actors then evaluate all their incoming offers, choose at most one to accept, and change their positions accordingly. Finally, under certain conditions, conflicts may occur, as actors follow through with their threats.

Each of these steps is effectively a sub-model. Later in this chapter, I will propose several variant sub-models, which may be combined in multiple permutations. In order to keep track of these permutations, I introduce a notation system inspired by Epstein and Axtell (1996). Each sub-model is designated by a letter, with the variant noted as a subscript. A model is then denoted by the set of sub-models it utilizes. My implementation of the original sub-models is denoted with the subscript 0.

Each step of a model run, the sub-models are executed in the following sequence:

- All agents compute their current risk propensity, based on their current position and the position of every other actor. (R, for risk update)
- 2. All agents evaluate their expected utility from threatening all other agents, and send an offer to any other agent where that expected utility is positive. (**O**, for offers)
- 3. Agents evaluate all their incoming offers, and choose at most one; this may involve choosing to change their position, or choosing to enter a conflict with another agent.
  (M, for move; and A, for attack)
- 4. Conflicts, if any, are resolved here. (W, for war)
- 5. All agents update their stated positions.

Note that the final step does not have a sub-model notation associated with it. That is because the agents' new position has already been determined by their offer selection and conflicts (sub-models  $\mathbf{M}$  and  $\mathbf{W}$ ). Whether the agents' position changes immediately or (as shown above) at a later point is governed by the model's schedule sub-model,  $\mathbf{T}$  for time. Note also that sub-models  $\mathbf{R}$ ,  $\mathbf{O}$ ,  $\mathbf{M}$  and  $\mathbf{A}$  are internal to the agents, while  $\mathbf{T}$  and  $\mathbf{W}$  are at the level of the overall model. There are also several sub-model variants for updating the agent saliences ( $\mathbf{S}$ ) with the baseline rule  $\mathbf{S}_0$  being that saliences are exogenous, unchanged from their input values.

There is no theoretical reason that a model instantiation may not have a mix of agents utilizing different decision sub-models, but in this chapter I will always set all agents to use the same sub-models. Using this notation, the original model comprises all the baseline sub-models and is denoted  $\{\mathbf{T}_0, \mathbf{S}_0, \mathbf{W}_0, \mathbf{R}_0, \mathbf{O}_0, \mathbf{M}_0, \mathbf{A}_0\}$ 

#### Architecture

The model is implemented via a modular, object-oriented framework, as shown in the UML diagram in Figure 3.1. There are three classes: NegotiationModel, which controls the overall model instantiation and contains the set of agents; NegotiationAgent is an individual

agent, including the agents' properties and methods for their interactions with other agents; these methods, in turn, serve as pass-throughs to its DecisionModel. Each agent has an instantiation of a DecisionModel, which implements the rules and heuristics by which the agents decide how to behave, and stores any agent-specific decision parameters and variables (such as the risk propensity).

A particular model instantiation is a single NegotiationClass object, containing the set of agents to be modeled. The model's scheduler iterates through the NegotiationActor methods, calling the same method for each agent before moving onto the next method. These methods correspond to the sub-models described in the previous section.

In the *initialize* method, agents orient themselves to the current state of the model. In practice, this involves the risk propensity update calculation. For *send\_threats*, agents iterate over all other agents; the DecisionModel's *decide\_threat* method evaluates each, and returns a boolean determining whether to send a threat to that agent. In *resolve\_threats* agents' DecisionModel runs *evaluate\_threats*, deciding how to react to each threat. It is here that offer selection occurs. If agents choose to change their position, this change does not take effect until *finalize* is called. In *resolve\_attacks*, agents iterate over threats they have issued, evaluating each with the *execute\_attack* method, which determines whether to follow through with the threat or not.

Conflicts are the one element of the model which occurs externally to the agents themselves, as they involve the interaction between multiple agents. This is why the model class itself has a *resolve\_attacks* method, which handles conflict resolution. This method is called after all the agents have decided whether to follow through with threats and turn them into conflicts. For each conflict, it iterates over all the model agents, calling their *allocate\_support* method and giving them an opportunity to allocate some fraction of their capability to one side or another of the conflict. The model then resolves the conflict based on the share of support each side received, and calls the winner and loser's *win\_conflict* and *lose\_conflict* methods, respectively.

The NegotiationActor class is not intended to be subclassed; different agent types, and



Figure 3.1: Model Architecture

different model variants, are implemented by changing the DecisionModel. Negotiation-Model subclasses primarily alter two things: the sequencing of the stages, and the conflict resolution sub-model.

#### 3.2.2 Details

#### **Initialization and Inputs**

Each model instantiation is meant to simulate a particular issue under contention. Thus, we begin by identifying the issue of interest, and the actors we need to represent. In some cases, identifying the actors is simple: for example, all the states who are party to a particular formal negotiation process, or (in a retrospective simulation) all the states who participated in a particular conflict. In other cases, qualitative expertise (the modeler's own, or elicited from subject-matter experts) may be needed to determine who the relevant actors are.

For every actor, we must assign position, salience, and capability values. Like the list of actors, the actor properties may be derived from different sources. We may elicit subject-matter experts for their estimates, or utilize various quantitative datasets more objective or replicable measures. Capabilities are perhaps the easiest to measure objectively; the National Material Capabilities (NMC) dataset (Singer, 1972, 1988) provides several calibrated measures of national power, aggregated into the Composite Index of National Capability (CINC). Position is harder to measure. One method is computing the similarity of states' alliance portfolios (Bueno de Mesquita, 1975; Bueno de Mesquita and Lalman, 1992; Signorino and Ritter, 1999): the more similar the alliances that two states hold are, the closer their interests are presumed to align. There has not been, to the best of my knowledge, an attempt to formally estimate salience from data. It can either be set based on expert estimates, or treated as an unobserved random variable.

#### Risk Propensity Update $(R_0)$

Conceptually, actors wish to balance between their preferred issue outcome and their desire to be a part of a winning coalition. At one extreme, an actor may adhere to its own position regardless of the positions of the rest of the actors, and no matter how much power is arrayed against it. At the other extreme, an actor may completely abandon its preferred position and seek to position itself as close as possible to the current consensus or winning position. In practice, most actors are likely to fall somewhere between these two poles.

In order to estimate each actor's risk propensity, the model compares the actor's current security to the security of all other actors in the model. More formally:

$$R_{i} = \frac{2\sum_{j} E^{i}(U_{ji}) - \max_{j} \sum_{k} E^{i}(U_{kj}) - \min_{j} \sum_{k} E^{i}(U_{kj})}{\max_{j} \sum_{k} E^{i}(U_{kj}) - \min_{j} \sum_{k} E^{i}(U_{kj})}$$
(3.1)

In this equation,  $E^{i}(U_{ji})$  refers to agent *i*'s perception of *j*'s expected utility of threatening *i*; the full expected utility calculation is specified in Equation 3.11. This result is then normalized as follows:

$$r_i = \frac{1 - R_i/3}{1 + R_i/3} \tag{3.2}$$

In the model as described in Bueno de Mesquita (2002) and Scholz, Calbert, and Smith (2011), the risk propensity is recalculated at every step, with an initial assumption of  $r_i = 1$  for all actors. This means that risk propensity is not a fixed property of the actor, but emerges from the overall model state. This can be thought of as the risk propensity required to maintain the actor's current position.

#### Expected Utility Calculation $(O_0)$

The expected utility (EU) calculation is important enough to have given the model the name it is most often known as – the Expected Utility Model. In each step of the model, the first decision facing each agent is whether to send an offer to each other agent, which they will do only if their expected utility is positive.

In order to compute the expected utility of threatening a particular target j, actor i computes the utilities of the following potential outcomes:

- 1. The actor threatens the target, and:
  - (a) Enters into conflict and wins:

$$U_s = 2 - 4(0.5 - 0.5|x_i - x_j|)^{r_i}$$
(3.3)

(b) Enters into conflict and loses:

$$U_f = 2 - 4(0.5 + 0.5|x_i - x_j|)^{r_i}$$
(3.4)

- (c) The target capitulates without a conflict ( $U_s$  again, as computed in Equation 3.3)
- 2. The actor does not threaten the target, and:
  - (a) The status quo continues:

$$U_{sq} = 2 - 4(0.5)^{r_i} \tag{3.5}$$

(b) The target moves towards the current median position  $\mu$ :

$$U_b = 2 - 4(0.5 - 0.25(|x_i - \mu| + |x_j - \mu|))^{r_i}$$
(3.6)

(c) The target moves away from the current median position  $\mu$ :

$$U_w = 2 - 4(0.5 + 0.25(|x_i - \mu| + |x_j - \mu|))^{r_i}$$
(3.7)

Next, the actor assigns probabilities to each outcome. The probability of winning and losing the conflict are this actor's perception of the fraction of support allocated to its position as compared to the target's. For a conflict between agents j and k, actor i's support (or 'votes') is computed as follows:

$$v_{jk}^{i} = c_{i}s_{i}(|x_{i} - x_{k}| - |x_{i} - x_{j}|)$$
(3.8)

When  $v_{jk}^i$  is positive, *i* supports *j*; negative, and *i* supports *k*. In light of this, *i*'s perceived probability of a victory by *j* in a conflict between *j* and *k* is:

$$P_{jk}^{i} = \frac{\sum\limits_{h|v_{jk}^{h}>0} v_{jk}^{h}}{\sum\limits_{h}|v_{jk}^{h}|}$$
(3.9)

This probability calculation is also used to determine the median position. The model extends the Black (1948) median voter calculation to account for the fact that each actor has a different 'voting' power, and may cast a different number of votes (allocate a different level of support) in different bilateral contests. The median position  $x^*$  is thus:

$$x_i^* \equiv P_{*k}^i > 0.5 \ \forall \ k \in N \tag{3.10}$$

Or, plainly, the position which is likely to defeat all other positions in bilateral contests.

The target's salience is used as the probability that it will yield to the threat, rather than enter into a conflict. Finally, the actor's perception of the probability of the target not moving at all is denoted  $Q_i$ ; and of the target moving towards the median position if it does move is  $T_i$ . These appear to generally be held fixed for all agents in a given model instantiation, and as such become model parameters.

Putting all the terms together gives us the full expected utility calculation. For convenience, the equation below omits all the full perception designation terms. The utilities are i's perception of j's utility of winning, losing, etc. P is i's perception of j's probability of winning the conflict.

$$E^{i}(j,k) = E^{i}(U_{jk}) = s_{k}(PU_{s} + (1-P)U_{f}) + (1-s_{k})U_{s}$$

$$-Q_{i}U_{sq} + (1-Q_{i})(T_{i}U_{b} + (1-T_{i})U_{w})$$
(3.11)

This EU value comes into play in several ways: first, at each step, agents issue an offer when the expected utility is positive – that is, the agents expect a better outcome from issuing a threat than not issuing it. Second, it determines how each target of a offer responds, as the next section describes.

#### Responding to Offers $(M_0, A_0)$

Once all the actors have sent their offers, they decide on at most one offer to accept. The offer selection mechanism is perhaps the least specified of all the model components in any of the prior literature. I utilize the description provided by Scholz, Calbert, and Smith (2011), who argue that agents choose offers based on two criteria: the expected utility associated with each offer, and the distance it requires the target actor to move.

In this approach, agents initially evaluate the distance between their current position, and the position associated with the offer. Agents will choose offers where this distance is the smallest, preferring to accept offers requiring them to change position the least.



Figure 3.2: Incoming Offer Categories

If there is a single closest offer, it is selected. However, if there are multiple offers with the same distance associated with them, the agent selects the most enforceable one – that is, the one with the highest perceived expected utility on the part of the sender.

How agents actually change their position based on their accepted offer is determined by the combination of the source and the target's perceived expected utility of threatening one another. The recipient of a threat evaluates how much it believes it stands to lose or gain from a conflict, and how much it believes the sender of the threat stands to lose or gain, and decides accordingly. The combinations of perceived utilities may fall into one of several octants, as shown in Figure 3.2.

**Capitulate:** This octant is the region of the EUxEU space where the target's EU is negative, and has an absolute value greater than the positive EU of the source. More colloquially, the target stands to lose more than the attacker stands to gain. This indicates a decisive advantage for the source in case of a conflict; in the face of some advantage, the target yields completely, adopting the source's position as its own. **Compromise:** In this octant, the target's EU is negative, but smaller in magnitude than the source's. In this case, the sender cannot enforce complete surrender, but can nevertheless coerce the target into changing its position towards that of the sender. The size of the step is determined by the ratio between the magnitudes of both actors' expected utilities, or more formally:

$$x'_{i} = x_{i} + (x_{j} - x_{i}) \left| \frac{E^{i}(U_{ji})}{E^{i}(U_{ij})} \right|$$
(3.12)

Thus, as the source's EU increases, the target agrees to increasingly better compromises.

**Conflict:** The conflict quadrant consists of the region of the space where both actors' expected utility from a conflict is positive; in other words, conflict can occur only when both actors expect to gain from it. When actors choose to enter into a conflict with one another, the Conflict sub-model is called, as described below.

Bueno de Mesquita (2002) notes that "Each player is free to renege on a proposed deal so long as it can enforce another agreement or so long as someone else can enforce an agreement on it." In practice, the latter part of the sentence refers to the fact that actors choose only one incoming offer to accept; I interpret the first half to mean that actors need not change position if another actor accepted their offer and is changing position towards them. This resolves a situation I observed while developing and experimenting with the model, where actor A capitulated to actor B, while actor B simultaneously capitulated to actor C – in effect, B coerces A into holding a position A itself is simultaneously abandoning.

#### Conflicts $(W_0)$

Conflicts occur when both actors see positive expected utility from a conflict with the other, leading them to choose each other's offer, and when neither agent reneges due to another agent accepting their offer. For each conflict occurring in a given step, all agents contribute some fraction of their capabilities to one side or the other, as described in Equation 3.8, yielding the probability given in Equation 3.9. The model stochastically chooses the winner with that probability, and the loser updates their position to that of the winner. The assumptions regarding conflicts are discussed in much more detail in Section 3.4.

## 3.3 Model Validation and Verification

The description above appears, to the best of my ability to ascertain, to be a faithful representation of the BDM Expected Utility model, as presented in the prior literature. However, as documented by Scholz, Calbert, and Smith (2011), not only has the model never been fully specified, but several components of it have been modified over time. In this section, I attempt to validate my implementation against two sets of inputs and outputs from the prior literature – and I should note up-front that in neither of these cases do I successfully fully reproduce the original reported results. In light of this, I conduct a deep dive into one of the model runs and provide a detailed walk-through of a single model step. This serves to verify that the model is working as expected, as well as to build the reader's intuition and understanding of the model's behavior and dynamics.

#### 3.3.1 Reproduction of Prior Results

A natural way to validate that the model described here successfully replicates the original is by reproducing prior results from known inputs. There are several models where the inputs are fully specified (or very nearly so) to allow for such replication attempts. For example, Scholz, Calbert, and Smith (2011) use the data from Bueno de Mesquita (1994), an application of the model to European vehicle standard negotiations, and demonstrate the reproduction of several model outputs. I attempt to reproduce this replication, as well as the China democratization example examined in Bueno de Mesquita (2002, chapter 6). Since the latter model in particular involves a stochastic element, we cannot reproduce the exact same results without knowing outcomes of the various random draws; instead, I examine the distribution of results across multiple model runs.



Figure 3.3: Original and Replication Model Comparison

The case examined by Scholz, Calbert, and Smith (2011) involves ten members of the European Community (the precursor to the European Union), negotiating over when to apply new emission regulations to mid-size automobiles (the full inputs are provided in Table 3.2). This case provides a good starting point, since there is minimal stochasticity in the model implementation – the inputs, including the saliences, are all fixed. The Scholz, Calbert, and Smith (2011) replication was validated in two ways: by comparing the median position, and the quadrants pairs of agents fall in, to those produced by the original model. The Scholz, Calbert, and Smith (2011) model perfectly reproduced the median position changes reported in Bueno de Mesquita (1994). Figure 3.3, compares those median positions to those generated by my own implementation. The comparison highlights the fact that the original model produces cyclic behavior which mine does not: certainly a richer dynamic, but not one with any external reason for us to believe it to be correct. (The observed historic outcome, 8.833 years, corresponds to approximately 0.8 on the normalized issue space, which neither model converges to).

Stronger evidence still that the Scholz, Calbert, and Smith (2011) model replicates the BDM results comes from the octant comparison, comparing the expected utilities of between pairs of agents at the first step of the model. As the paper notes, there is a vanishingly-small probability of the same octant configuration resulting by chance. My own model does not



(a) Bueno de Mesquita (1994) (b) Scholz, Calbert, and Smith Expected Utilities (2011) Expected Utilities (c) Model reimplementation

Figure 3.4: Octant Comparison, with Belgium as Focal Agent

replicate these octant configurations. One example is shown in Figure 3.4, with Belgium as the focal actor – the position of each other agent is Belgium's expected utility from challenging it, and its own expected utility from challenging Belgium.

A close comparison of the models suggests that the Scholz octant configuration is driven by far smaller probabilities of victory that the model produces, normalized by all possible votes rather than just the ones on the conflict at hand. This probability form appears to contradict the form of probability calculations provided in Bueno de Mesquita (1997, 2002), and (more importantly) a basic rule of probability: if there are two possible outcomes (in this case, actor i wins or actor j wins), their probabilities must sum to 1. There are, then, two alternatives: one is that the Scholz model is successfully replicating an erroneous BDM model implementation; the other is that there is some other mismatch between the models which is producing matching octants. There are two, somewhat contradictory, preliminary conclusions we can take from here: that the overall model is sensitive to changes in specific sub-models, and that it is nevertheless robust enough to produce similar results even when these sub-models differ, and produce substantively different interactions.

Furthermore, while the Scholz model successfully replicates specific features of this one
particular model instantiation, this success does not translate to other cases. Scholz, Calbert, and Smith (2011) note that their implementation fails to replicate another case, presented in Bueno de Mesquita (1997). Experiments on my part confirm that it also fails to replicate the following cases, which I will describe next.

The next case I use is the Chinese democratization example, presented in chapters 3 and 6 of Bueno de Mesquita (2002). In this case, the model is being instantiated is attempting to predict the dynamics of internal political reform in the People's Republic of China. The agents here are a mix of individuals, organizations, loose factions (e.g. "Students and Intellectuals"), and foreign countries. Bueno de Mesquita (2002) reports model median positions by step for several model instantiation. However, as part of this particular case, agents' saliences are randomized each step, with 20% probability each; the reported results do not specify which saliences were changed and how. This makes it effectively impossible to attempt to replicate the exact same model run. To overcome this, I generate many more model runs: 50 of the Scholz model, and 100 of my own. I start each run with the initial conditions provided in Bueno de Mesquita (2002), and randomize agent salience values as described. This yields a distribution of models, and in particular a distribution of medians. Figure 3.5 shows the distributions of these medians, as well as a dotted line indicating the reported median<sup>2</sup>. In this case, the models produce similar, though not identical, distributions to one another – and neither comes close to the reported values. This indicates that the Scholz model is not a perfect replication of the original model as implemented. Furthermore, it suggests that we may need to be cautious about using any single set of reported model results as 'ground truth.'

In light of these results, and of the reported changes in the variations of the original model, I believe that validating the model reimplementation directly against reported model outputs is not a tenable option. Nevertheless, I believe that the model I have implemented correctly captures the core sub-models of the original, and provides a sufficiently solid basis

 $<sup>^{2}</sup>$ This median value was taken from Bueno de Mesquita (2002, chapter 3); chapter 6 provides only a chart of medians from several runs, but not precise values. However, the chart shows values close to the one used here.



(a) Scholz, Calbert, and Smith (2011) Median Positions

(b) Reimplemented Model Median Positions

Figure 3.5: Chinese Democratization Model Medians, at the End of Step 2

to proceed with this analysis.

## 3.3.2 Deep Dive

In order to understand how the model behaves, it is valuable to take a close look at the sequence of actions in a single iteration. This will also serve as a method of verification, ensuring that the model is not producing unexpected or incorrect behaviors. In order to do so, I will continue to use the case presented above, of the negotiation over European Community vehicle standards. I choose this example in particular since it has already been well-studied: first by Bueno de Mesquita (1994) directly, then by Scholz, Calbert, and Smith (2011) as the key case for testing their replication of the original BDM model; and finally by McKibben-Sanders (2014) testing the replication of the replication. Furthermore, the number of agents is small enough to make the interactions amenable to direct human examination, and the position space has a straightforward, numeric interpretation: the years before the regulations are applied. The agent initial positions, capability and saliences were

Name	Capability	Position	Salience	$r_{i,0}$
Netherlands	0.08	0.0	0.8	1.12
Belgium	0.08	0.5	0.4	1.00
Luxembourg	0.03	0.0	0.2	0.50
Germany	0.16	0.0	0.8	1.12
France	0.16	1.0	0.6	1.02
Italy	0.16	1.0	0.6	1.02
UK	0.16	1.0	0.9	2.00
Ireland	0.05	0.5	0.1	0.65
Denmark	0.05	0.0	1.0	1.45
Greece	0.08	0.5	0.7	1.56

Table 3.2: Example Model – Agent Data

all determined by consultation with subject-matter experts.

Initially, the model computes the risk propensity coefficients shown in Table 3.2. Note that indeed one agent, the UK, has a value of r = 2.0 and another one, Luxembourg, has r = 0.5; as agents are computing their risk propensity in comparison with one another, at least one agent will always have each of these values, in any step of any given model instantiation.

Next, the agents estimate their expected utility of threatening all other agents, sending offers where those expected utilities are greater than zero. These EUs are given in Table 3.3. The important thing to note here is that the majority of directed dyads have a positive EU value – which means that all those dyads will have threats issued across them.

Table 3.4 shows which offer each agent accepts, and what category each one falls into. Immediately, we notice that five agents all choose to enter into conflict with Belgium. Note that Belgium's position is exactly in the middle of the issue space, and it has lower salience than Greece, the other agent holding the same position, making it an appealing target. For a conflict to occur, Belgium must choose one of them – and indeed chooses the Netherlands, which leads to a conflict. Meanwhile, Denmark chooses to compromise with Greece and update its position towards it, while Greece chooses to compromise with the UK. However, recall that agents renege on offers if another agent has accepted their own offer in a more

Source \Target	Bel.	Den.	Fra.	Ger.	Gre.	Ire.	Ita.	Lux.	Net.	UK
Belgium	—	0.27	0.40	0.41	0.00	0.00	0.40	0.85	0.41	0.11
Denmark	0.39	—	0.08	0.00	-0.00	0.79	0.08	0.00	0.00	-0.61
France	0.60	0.28	—	0.62	0.29	0.90	0.00	1.64	0.62	0.00
Germany	0.47	0.00	0.46	_	0.08	0.86	0.46	0.00	0.00	-0.23
Greece	0.00	0.13	0.27	0.28	_	0.00	0.27	0.74	0.28	-0.04
Ireland	0.00	0.30	0.42	0.43	0.00	—	0.42	0.80	0.43	0.17
Italy	0.60	0.28	0.00	0.62	0.29	0.90	—	1.64	0.62	0.00
Luxembourg	0.46	0.00	1.45	0.00	0.18	0.74	1.45	_	0.00	0.76
Netherlands	0.47	0.00	0.46	0.00	0.08	0.86	0.46	0.00	_	-0.23
UK	0.35	-0.70	0.00	-0.36	0.04	0.65	0.00	0.66	-0.36	_

Table 3.3: Example Model – Expected Utilities, step = 0

Table 3.4: Accepted Offers

Agent	Accepted Offer	Category
Netherlands	Belgium	conflict
Belgium	Netherlands	conflict
Luxembourg	Belgium	conflict
Germany	Belgium	conflict
France	Belgium	conflict
Italy	Belgium	conflict
Ireland	Netherlands	conflict
Denmark	Greece	compromise
Greece	UK	compromise

preferable way. In this case, since Denmark has chosen to compromise with Greece, Greece in turn backs out from its decision to compromise with the UK, and will not update its position after all.

The conflict between Belgium and the Netherlands happens before Denmark updates its position. Each agent contributes capability to one side or the other, as governed by Equation 3.8; these contributions are shown in Table 3.5. Note that even the two actual actors entering into the conflict do not allocate their full capabilities: they are scaled once by their salience, and again by the distance between their positions. Putting the total values into Equation 3.9 yields a probability of victory of 0.63 for Belgium, and thus 0.37 for the Netherlands. The model uses these probabilities to randomly select the winner; if Belgium wins, the Netherlands updates its position to 0.5, and if the Netherlands wins Belgium updates its own position to 0. At this point, Denmark will update its own position to compromise with Greece. Then, the next round begins.

Table 3.5: Capability Contribution – Belgium vs Netherlands

Agent	Belgium	Netherlands
Netherlands	0	0.032
Belgium	0.016	0
Luxembourg	0	0.003
Germany	0	0.064
France	0.048	0
Italy	0.048	0
UK	0.072	0
Ireland	0.0025	0
Denmark	0	0.025
Greece	0.028	0
Total	0.2145	0.124

## 3.4 Model Discussion

Examining the description of the model above, we can see clearly that it is in fact an agent-based model: the actors are the agents, interacting with each other in discrete steps

according to a set structure. Each agent is endowed with several properties, some public and some private, making it natural to model them using object-oriented design. Indeed, the focus on the utility calculations in much of the prior literature elides the deeply computational nature of this model. It is anchored not only in game-theoretic logic, but on embedded assumptions which have a substantial effect on the model's behavior, and which, when articulated in code, are easy to experiment with.

The model, like many agent-based models, can be separated into two layers: the agents' external behavior with regard to one another, and their internal decisionmaking process.

Externally, agents send threats to some other agent. A threat has no information specifically attached (though of course the target has access to the public information on the source of the threat) – it is simply boolean; either a source has issued a threat to a target, or it has not. Next, all agents choose how to change their public positions in response to the incoming threats – and these choices are public, and simultaneously become common knowledge. When two agents have chosen conflict with one another, conflicts occur and are resolved. Finally, agents update their stated positions based on the offers and conflict outcomes, before the sequence starts over. These steps happen in sequence with one another, but that all agents act simultaneously within each step.

Note that at no point in this sequence do utilities, expected or otherwise, come into play. The utilities are computed, and applied, internally by agents to drive their decisionmaking, choosing who to send threats to and how to respond to them. We can easily imagine other beliefs and heuristics guiding agent decisions besides the ones that have been chosen here. Even the sequence of events itself is not necessarily fixed.

In this section, I will discuss several of the model assumptions in greater depth. I also propose sub-model variants which implement alternative assumptions. The object-oriented nature of the model means that a single sub-model to be replaced without modifying any other aspect of it. This, in turn, allows us to instantiate different model variants with the same initial conditions and directly compare their dynamics and outputs – and, in turn, use these to understand how the assumptions drive the model as a whole.

## 3.4.1 Expected Utility

While the expected utility calculation (Equation 3.11) is important enough to give the model its name, I note that this calculation is largely heuristic in nature. The probability of winning or losing P (as computed in Equation 3.9) is based on the best possible approximation. The inclusion of the salience term as a probability, however, is much more obviously a heuristic. The source of a potential threat uses salience as the probability of the target yielding to that threat. However, salience does not play a similar role in the threat response side of the model; in fact, it is possible for an agent to have low salience and high probability of winning a conflict, and hence high EU, due to support from other agents. Simply put: agents are modeled as making an estimation that is plausible on its face, but that nevertheless does not reflect the reality of the model itself.

Q and T are more clearly broad heuristics: they are not as specifically tied to the behavior of a particular agent, but are approximate, coarse-grained estimates of the behavior of the world overall. While Bueno de Mesquita (2002) suggests that agents update their belief about the world in response to the outcome of each round, no such learning rule is provided. Q and T provide candidates for such updating: for example, the agents could use the fraction of agents that moved last round as Q, or update their belief in a Bayesian fashion based on a prior probability distribution for Q. Alternatively, we may note that any given model potentially represents one issue among many which the actors are addressing at once or in rapid sequence, and suggest that Q and T represent heuristics developed across all issues.

### 3.4.2 Conflict and Coercion

The Conflict outcome deserves special attention here. Conflict, and especially war, plays an outsize role in the study of international relations, and indeed in world history. Being able to predict the outbreak of conflicts would be a particularly powerful result from this model; in fact, the risk propensity and probability of victory sub-models were originally developed precisely for such predictions (Bueno de Mesquita, 1985). For these reasons, it is important

to understand precisely when the model predicts conflicts will occur, and how to interpret such predictions.

There are several ways of handling a conflict and its outcome. The method put forward in Scholz, Calbert, and Smith (2011) involves comparing both agents' expected utility from the conflict; the winner is the one with the greater EU, and the loser changes position to match the winner's. I believe this approach is incorrect, for several reasons. EU incorporates a specific term for probability of winning, but also the terms for movement absent a threat – which, if a conflict is actively breaking out, are no longer relevant. Furthermore, if a conflict is already occurring, it is unclear why the salience should continue to have a probability-like effect. The calculation for P (Equation 3.9) clearly implies that the probability of winning and losing is tied primarily to the total capability brought to bear, not the expected utilities associated with them. Furthermore, the P term itself strongly implies that conflicts are stochastic, while the EU comparison method makes them determinisitc.

In fact, while several of the original model descriptions (in particular Bueno de Mesquita, 1997, 2002) do not specify how conflicts are resolved or what their outcome is, Bueno de Mesquita (2011) explicitly clarifies a conflict sub-model consistent with the one described in Section 3.2.2 above. In this sub-model, conflicts are resolved stochastically, with probability based on the ratio of the capabilities contributed by all agents to each side. Conflicts have only two possible outcomes, with one side winning and the other side losing, and the loser adopting the winner's current position.

This sub-model itself contains two sub-components: how conflicts are resolved, and what the consequences of this resolution are. The conflict sub-model described above assumes that conflicts result in one agent coercing the other to accept their own preferred position. However, changing position is, ultimately, a decision that is internal to an actor. Allowing one agent to directly change another agent's position breaches the division between agents' external actions and internal decisionmaking. More importantly, this is not how coercion has generally been understood. Coercion, as described by Schelling (1966), involves one actor presenting another, explicitly or implicitly, with a choice of complying with some demand or facing some punishment. Schelling describes this punishment primarily as 'pain' or 'hurt' – in other words, disutility. In this view, the punishment need not degrade the target's capabilities, or eliminate it as a threat; it must simply impose a cost that is greater than the cost of complying with the given demand. If we examine militarized conflicts in particular, only a minority end with one side completely changing its orientation – for example, the Second World War led to Japan's conversion by force from being an enemy of the United States to its ally (Schaller, 1997), and Tanzania's invasion of Uganda in 1978-79 successfully replaced the regime of Idi Amin with one friendly to Tanzania's interests (Acheson-Brown, 2001). However, in the majority of cases, even those involving what Schelling describes as 'brute force' where one side wrested material gains from the other (such as in the Falklands or Gulf Wars), the other side's overall geopolitical orientation remained largely unchanged; certainly, they did not become allies of the power which defeated them.

The expected utility calculation, and its role in the model, is compatible with this view. There is an explicit negative utility to losing a conflict, which the agents attempt to avoid. The greater this threat – the lower the overall expected utility of a conflict – the more likely the agent is to capitulate entirely. Conflict cannot reduce agents' capabilities, or remove them from the model, suggesting that the disutility of losing bears more resemblance to Schelling's concept of 'punishment' than 'brute force.' Though the utilities associated with winning and losing a conflict are directly tied to the difference between the agent positions positions, this does not necessarily imply that the agents must change position upon losing. If the cost of losing a conflict is solely having one's position changed, what motive do actors in the 'Capitulate' octant have to indeed capitulate, rather than embrace a chance, however slight, of changing the source's position through conflict? If a conflict does not change agent positions, it follows that these utilities are exogenous to the utility directly gained by an agent to have support at its current position. And indeed, in general, political actors do gain utility from winning (or succeeding in) a conflict: in the form of increased internal public support, prestige, credibility, and potentially an improved power ratio compared to the loser, while failing at a conflict carries the inverse costs (Bratton, 2005). Furthermore, both the means and the stakes of conflict will not be the same for close allies and rivals. Agents may attempt to coerce an ally, but are unlikely to use harsh measures which will risk harming the ally substantially and thus undermining their own position; similarly, the gains from inflicting such lesser costs on an ally will be less than those of inflicting greater losses on an enemy.

We must note, of course, that the interpretation of conflict varies based on the particular scenario and issue under consideration. In modeling a negotiation over regulatory issues, conflict is highly unlikely to be military. However, in modeling geopolitical contests, conflicts may range from diplomatic snubs to economic sanctions to all-out war. Even military conflicts range in severity. A border clash, or targeted airstrikes, may inflict some costs on the target, but are unlikely to have direct system-wide effects. However, armed conflict between nuclear superpowers, as is a potential outcome in the Bueno de Mesquita (1998) Cold War model (discussed in detail in the following chapter), runs the risk of consequences far outside the scope of this model.

Finally, the model elides the difference between when agents decide to engage in conflict and when the conflict, in fact, occurs. An implication of the offer-choice mechanic is that conflict occurs only when both parties choose to engage in it. This seems to remove the possibility of unilateral attacks. Furthermore, both parties to a conflict estimate the capability contribution of all other actors based on the parties' own risk propensity, and the agents' current positions. It seems reasonable that, once a conflict is initiated, other agents actually contribute capability based on their own decision rule – and in particular, based on their own risk propensity. Here the precise scheduling of the conflicts compared to the rest of the model come into play: if conflicts occur *after* agents change their positions in response to threats, their contribution of resources, and even the side to which they will contribute, may differ from their contributions before moving.

Based on this discussion, I propose two new sub-model variants. As these variants appear better grounded in prior theory than the baseline, I hypothesize they will produce better correspondence with empirical data.

Conflict without position changing  $(W_1)$ : Under this conflict sub-model variant, agents' positions are not changed when they lose a conflict.

Positions update before conflicts  $(T_1)$ : Under this schedule variant, conflicts occur only after agents have moved in response to offers. This means that by the time a conflict occurs, the balance of power has shifted from what it was when the original threat was made.

One-way attacks  $(A_1)$ : This attack decision model variant removes the requirement for both agents to choose the other's offer for a conflict to be initiated. In this variation, an agent can attack any other agent it had threatened previously in this model step. However, it adds an additional decision-point, before attacks actually occur. At this decision point, the agent evaluates a modified form of the expected utility, which removes the uncertainty factors as to the target's response to the threat.

$$EU'_{i}(i,j) = PU_{s} + (1-P)U_{f} - U_{s}q$$
(3.13)

Depending on when the agents new positions take effect, P,  $U_s$  and  $U_f$  may all be different from the values calculated at the offer-sending phase. If either the source or the target have changed position, the utilities of victory and defeat will be different as well; furthermore, if other agents have changed position, the direction and magnitude of support they are expected to provide will also have changed, changing the value of P as well.

## 3.4.3 Risk Propensity

The risk propensity model attempts to estimate the risk an agent is willing to endure in pursuit of a preferred outcome; however, this is not directly risk propensity in the traditional economic sense (i.e. a relationship between a probability and the potential payoff an agent would require to wager on it), but in the political sense: how exposed the agent is willing to be to the possibility of being attacked, or simply left out of a winning coalition, in order to advocate for its preferred position or outcome. Bueno de Mesquita (2002), citing Lamborn

(1991), describes this as a tradeoff between political satisfaction and policy satisfaction. When an agent is very secure, it is assumed to have chosen that secure position due to a low risk propensity; when it is insecure, this is taken as a revealed higher risk propensity.

Risk propensity is a key place where agents mirror-image: assuming that the other agents will make decisions the same way as they themselves do (Heuer, 2001). When agents estimate other agents' expected utility from an offer (Equation 3.11) they do not have access to the other agent's risk propensity, and so use their own instead, as detailed in Equations 3.3–3.7. Mirror-imaging is a well-documented phenomena in humans (Meltzoff and Decety, 2003), and is also one of the cognitive traps intelligence analysts are frequently warned against (Heuer, 2001) – this suggests that it is not an unreasonable assumption for capturing real behavior. Nevertheless, risk propensity in this model is not an inherent property of the agent but emerges from the current configuration of positions. Thus, it is not at all clear whether we ought to expect organizations (or the individuals who compose them) to mirror-image on a dynamic, second-order property.

The form of the risk propensity calculation itself is explicitly grounded in Prospect Theory (Kahneman and Tversky, 1984; McDermott, 2001) in that it compares the agent's current security to two framing values: the best and worst possible security values for the agent. This helps explain why this value is recalculated each step: the best and worst security values may well have changed. This is expressed in a more generalized from of Equation 3.1, using a non-specific security metric:

$$R_{i} = \frac{2security_{i} - \max_{j}(security_{j}) - \min_{j}(security_{j})}{\max_{j}(security_{j}) - \min_{j}(security_{j})}$$
(3.14)

The original model measures security level as the sum of incoming expected utilities against an agent, or more formally:

$$security_i = \sum_{j \in N} EU_i(j, i)$$
 (3.15)

Using this measure, a higher number (worse security) indicates more, and more powerful, potential attacks at the agent's current position.

I propose an alternative security level metric: the summed probabilities of successful attacks. This has the advantage of not relying on other's potential gain and loss; furthermore, unlike EU, probabilities are strictly positive, and thus negative values cannot reduce this value. This metric was also adopted by Wise, Lester, and Efird (2015a), as well as by another similar model variant reported to me by private communication. Formally, this security metric is:

$$security_i = \sum_{j \in N} P_i(j, i) \tag{3.16}$$

Next, consider the range of alternatives. In the original, this range is the security levels of all the other agents, or:

$$(\min_{j \in N}(security_j), \max_{j \in N}(security_j))$$
(3.17)

This method assumes that agents' risk propensities are relative to those of other agents. In particular, it has the consequence that one agent will always have the maximum possible risk propensity, while another will always have the minimum value. The alternative I introduce compares each agent's positions not to those of other agents, but to other possible positions the agent itself could hold, while holding all other agent positions constant. This follows the methodology used in Bueno de Mesquita (1985), as well as Bueno de Mesquita and Lalman (1992) and Bennett and Stam (2000b), and appears true to the description in Bueno de Mesquita (2002) of examining "the most political welfare... and the least the actor could have realized." Thus, the minimum and maximum security positions are defined as:

$$(\min_{x \in [0,1]} (security_{i|x_i=x}), \max_{x \in [0,1]} (security_{i|x_i=x}))$$
(3.18)

This follows from an assumption that if an agent is holding a position which reduces their security they are more risk-accepting, while a position which increases their security indicates less risk propensity. Note, however, that it still cannot distinguish between an actor whose security-maximizing position is due to risk aversion from one which simply has a preference for that position separate from its degree of risk propensity.

Putting these two variants together, I designate the risk-acceptance updating sub-model which utilizes Equations 3.16 and 3.18 as  $\mathbf{R}_1$ .

## 3.4.4 Offer Selection

Much like the expected utility calculation itself, the mechanism for offer selection and position updating is also a heuristic. It assumes that agents associate a particular potential position with each incoming offer, and will only select their new position from among those options. Furthermore, the model assumes that agents whose offers are not selected, or whose offers are reneged on, will simply accept this – even when they have the coercive capability to force the target to change position. Yet this reneging rule has the effect of substantially reducing the frequency of agents actually changing their position.

I experiment with two alternative heuristics which address these issues. The first is simply stated: agents choose offers according to the original rules, but do not renege even if another agent has accepted their own offer. This is sub-model  $M_1$ .

One possible interpretation of the reneging rule is that agents are able to renege with the support of their new allies – the agent or agents who accepted their offer. In order to capture this dynamic explicitly, I implement an additional offer selection heuristic: threat balancing, which I label  $M_2$ . In this variant, agents do not choose among their incoming offers directly. Instead, each agent attempts to balance between those threats, so as to maximize the support it will receive in any conflict. For each incoming offer, the agent evaluates its own EU, and estimates that of the sender. The agent selects the offers where the estimated EU is positive, and its own EU is negative – that is, credible threats from agents it wants to avoid entering into a conflict with. The agent then updates its position by the weighted mean of these offers, where the weights are the absolute value of the agent's (negative) EU from each. Simply put: the agents do not choose a single offer; instead, they respond to all credible threats, in proportion to how threatening they are.

### 3.4.5 Stochasticity and Uncertainty

The original model is largely deterministic; the only source of uncertainty is in the outcome of conflicts, which (as Section 3.3.2 suggests) are relatively rare. This determinism has been one of the sources of criticism of the model (Brandt, Freeman, and Schrodt, 2011), supported by the fact that the stochasticity of conflicts is rarely explicitly spelled out, and by the model outputs frequently being reported as point forecasts rather than distributions.

The inputs to the model are based on imperfect, uncertain estimates. Elicitation of estimates from subject-matter efforts is often prone to biases (Morgan, 2014; Tetlock, 2005); even absent such biases, careful methodology is required to capture the experts' uncertainty (Gill and Freeman, 2013). This uncertainty is particularly important since it seems highly unlikely that even experts can meaningfully estimate the difference between, for example, an actor having a salience of 0.69 and 0.71. Nor does replacing expert opinion with quantitative sources remove the uncertainty. Positional measures based on alliance networks are only approximate estimates themselves (Signorino and Ritter, 1999), while the Composite National Material Capabilities dataset (Grieg et al., 2010) (used to provide capabilities in several model instantiations) is a linear combination of several different measures, and as such only provides a very coarse-grained estimate of national power. Even such datasets rely primarily on human coders and archival sources, and regularly update their methodology (Gibler and Sarkees, 2004), again suggesting that they cannot be taken as certain estimates. Yet given the complexity which characterizes this model (and indeed, agent-based models more generally), it is possible that even small changes to model parameters will lead to radically different outcomes, even if the system is otherwise deterministic. In fact, this appears to be the case with the Scholz, Calbert, and Smith (2011) model (Masad, 2014). Many of the model's results are reported as point estimates, suggesting certain, deterministic outcomes. This certainty reduces the model's usefulness, however. The uncertainty of the world is a core assumption in forecasting, both political and otherwise. The purpose of a forecasting model is not to eliminate this uncertainty, but to characterize it, place bounds on it, and make predictions in spite of it. A way of reincorporating this uncertainty into the model is by sensitivity analysis: running multiple instantiations of a model while slightly varying different parameters, and recording the results. Such an analysis is conducted and reported in detail below.

Furthermore, even if the model and agent parameters are perfectly accurate, they will not necessarily remain static for the entire timeframe the model is attempting to forecast. The agent's capabilities, saliences, and potentially even positions may all be subject to exogenous changes and shocks due to processes and events outside the scope of the issue being modeled. Such factors may range from regular, long-term differences in growth rates of national GDP or military spending, to unexpected events such as major terrorist attacks or natural disasters which alter actors' priorities and views of the world.

One way to address this is to exogenously and stochastically change agent properties. Salience is the least concrete of the properties, the hardest to estimate, and more likely than capability to undergo exogenous or random changes; thus, it is not surprising that it is the one most likely to be perturbed. Salience updating is the only area where multiple variants are provided within some of the original BDM papers. Bueno de Mesquita (1998) has all agents' saliences determined randomly at the beginning of each step of the model, while the Chinese democratization model in Bueno de Mesquita (2002, chapter 6) randomizes agents' saliences with a probability of 30%. I designate the sub-model variant where agent saliences are randomized every step with probability p as  $\mathbf{S}_{1,\mathbf{p}}$ . Similarly, I suggest a model variant  $\mathbf{S}_2$  where agent saliences are randomized once at the beginning of a run, and held constant thereafter.

Variable	Value or Distribution
$\overline{x_i}$	$\sim Unif(0,1)$
$s_i$	$\sim Unif(0,1)$
$c_i$	$\sim Unif(0,1)$
$Q_i$	0.5
$T_i$	0.5
Number of agents	10
Number of steps per run	10

Table 3.6: Experiment Inputs

## 3.5 Sensitivity Analysis

As described above, the model is stochastic only in the outcome of conflicts, and is otherwise deterministic. In fact, in much of the literature, model outputs are reported as only point predictions. This makes it unclear whether these point predictions reflect one particular model run, the median or mode of multiple model runs, or whether the model is functionally deterministic and converges to the same results across multiple runs. Furthermore, these results do not account for uncertainties in the input data – which, as I discuss above, may be inaccurate or approximate. Finally, observing the model address each case and instantiation separately makes it difficult to identify whether the model has any artifacts or behaviors which tend to emerge independent of the specifics of each case.

In order to address these questions, I conduct a sensitivity analysis using artificial data, as follows. I first generate 100 sets of 10 notional agents each, with the agent properties drawn randomly as shown in Table 3.6. Each of these will be an artificial study case, which will serve as the basis for a set of model instantiation.

I conduct two experiments with this data. With each artificial study case, I create 100 model instantiations for each experiment, and run each for 10 steps before recording the median position at the final step. In Experiment 1, all instantiations of each case are run with the same initial properties; the only source of stochasticity is the outcome of conflicts. In Experiment 2, each instantiation of each case is perturbed slightly. For each parameter  $t_i \in (x_i, c_i, s_i)$  for all agents, I perturb the parameter by a small random value:

$$t'_i = t_i + \epsilon$$
, where  $\epsilon \sim \mathcal{N}(0, 0.05)$  and  $t_i$  is clamed such that  $t'_i \in [0, 1]$  (3.19)

Each experiment then yields 100 sets of 100 outcomes each. Initially, we can simply examine the distributions of these outcomes. Figure 3.6 shows histograms for the distributions of outcomes (that is, final median positions) for the first nine cases. Since these are instantiated from the same artificial case, we can directly compare the corresponding histograms in the two sub-figures (e.g. the histograms in the top-left corners both come from the same data, likewise the top-center histograms, etc.) Several things are immediately obvious. Several of the cases in Experiment 1 result in a single outcome across the instantiations. Despite the model's stochasticity, these runs all converge to the exact same median outcome. In fact, across all 100 artificial cases, Experiment 1 produces convergent results in 29 of them. This suggests that there are two qualitatively different categories of cases: those where the model identifies no uncertainty, and those where it does. In contrast, Experiment 2 produces no convergent results, and shows wider distributions of outcomes even when the corresponding cases in Experiment 1 are not convergent.

Next, we can attempt to characterize the resulting distributions. Visual examination suggests that many of the resulting distributions are approximately normal. Applying the d'Agostino (1971) normality test to each dataset shows that only 25 of the cases generate normal distributions of outcomes in Experiment 1, while 43 do so in Experiment 2.

Figure 3.7 compares the mean and standard deviation of outcomes from each case across both experiments. The key thing to observe is that the means of both are nearly identical. This means that though the perturbations of Experiment 2 increase the range of possible outcomes, they remain centered on the same outcome value; at no point does the perturbation appear to drive the model to a distribution qualitatively different from the non-perturbed outcomes. Additionally, the perturbation generally – though not always – increase the uncertainty of the model outcome, as measured by the outcomes' standard



Figure 3.6: Histograms of Median Positions from Nine Artificial Cases



Figure 3.7: Experiment 1 and 2 Comparison

deviation. Furthermore, cases with similar degrees of uncertainty without perturbation (including no uncertainty, in the case of convergent cases) can yield different degrees of uncertainty once perturbed.

## **3.6** Information Extraction

At its most basic level, this model is meant to predict the outcome of the issue under negotiation. The outcome is generally defined as the median voter positions (Black, 1948): the position currently held by one of the actors which would defeat all other positions in bilateral contests. However, this may not be the most useful piece of information the model provides, particularly since many issues do not end up being resolved by a 'vote,' or indeed resolved at all. In this section I will describe several additional outputs which can be extracted from the model, analyzed in their own right, and compared to empirical data.

One key piece of information the model generates is the set of the positions held by each actor at each iteration, and the trajectories of these positions. In cases where the position values represent concrete outcomes, the model provides a prediction (either a point prediction or a distribution) of the position each specific actor will hold. By examining how each actor's position changes, we can estimate the firmness of their position within the broader system being modeled.

Furthermore, actor positions represent not just particular desired outcomes, but alignment with regard to other actors in a particular model run. If the model predicts two or more actors' positions moving together, we can interpret this as a prediction of cooperation, and potentially a formal or informal alliance, between these actors. Certainly, it is a prediction of convergent interests. This interpretation is shown in Bueno de Mesquita (1998), where certain model runs have agents' positions converging in ways that qualitatively resemble real-world alliances. This methodology can be formalized, by using the simulated distances between agent positions as a predictor for alliances or other relationships between the corresponding states. These distances can be used as an independent variable in several models, with alliances as the dependent variables: most simply a logistic regression, or network methods such as exponential random graph models (ERGMs) (Robins et al., 2007) or latent space models (Hoff, Raftery, and Handcock, 2002).

The model does not just generate positional predictions; each step generates discrete events from each interaction between a dyad of actors. While there is not have a clear mapping between model steps and real-world time, model-generated events can be compared to historic events within a window of interest; alternatively, the sequence of events can be compared, without taking into account the time-difference between them. Across multiple model instantiations, we can examine the frequency at which different events occur, either across an entire run or within a specified number of steps.

The most obvious events, of course, are conflicts. Every model run may generate sets of conflict events when one or both of a pair of agents (depending on the model variant) choose to engage in conflict with each other. Conflict events are the most well-documented type of event in international relations, recorded in multiple specialized datasets (e.g. Hendrix and Salehyan, 2013; Jones, Bremer, and Singer, 1996; Leng and Singer, 1988) and are over-represented in media-coded event datasets (Schrodt, Simpson, and Gerner, 2001). This is particularly true when the conflict is overt between major political actors: military action is more likely to be recorded, while verbal conflict inside of a closed negotiation session less so.

Each conflict also creates a set of secondary events, with regard to the actors contributing capability to one side or another. The model does not specify what form this contribution takes; in reality, corresponding events may be positive actions towards the actor being supported (e.g. transfers of weapons or other materiel, favorable economic considerations), or negative actions towards the other side of the conflict (e.g. sanctions, withdrawal of support or access to resources). There may even be cooperative events observed between actors who are both supporting the same third party in its own conflict.

Offers are also events, though their correspondence with real events is also difficult to interpret. The model generates a much larger number of offers than end up being chosen, or even having a direct effect on the decisionmaking of the receivers. However, inasmuch as a conflict can only happen in the presence of threats, these threats may be treated as low-intensity conflict events as well.

# 3.7 Comparing Model to Event Data

The section above described how the model generates dyadic events, which can be compared to observed data. In this section, I describe implementing such an analysis in practice. I will walk through the instantiation of a set of geopolitical model runs using multiple different model variants, and comparing the conflicts they generate with observed micro-level event data in order to test and compare how well each explains the observed events.

I set up several model variants, each using some different combinations of sub-models, based on different assumptions. I instantiate each with the same country-level agents with data from 2004, a year chosen arbitrarily for being well-covered by event data sets and lacking any major geopolitical shocks in the following years. I run each instantiation of the model for 24 steps, corresponding to approximately two years of 'real' time. This mapping is somewhat arbitrary, but is based on the conventions which have developed for aggregating event data (Yonamine, 2011). In general, states do not change position or decide to begin or end confrontations with others on a day-by-day basis; however, neither do they do these things only once per year. Running these models multiple times produces a set of simulated events. I then run several tests to measure whether these simulated events are useful as predictors of real events. Comparing the predictive power of the model variants to one another can allow us to assess which variant is more predictive – and thus, be provide evidence in support of their underlying assumptions.

### 3.7.1 Data Sources

#### **Event Data**

Real-world event data will serve as the dependent variable we are attempting to predict. The event data source I use for this is the Integrated Crisis Early Warning System, or ICEWS.

This is a dataset of international political events, and specifically actions, extracted and machine-coded from media sources (Boschee et al., 2015). Each event consists of a Source (who carried out the action), a Target (who the action was directed at), and a code for the action itself. Actors (sources and targets) are coded to country, sector (e.g. government, civilians, rebels) and when possible to unique entities. Actions are assigned an Intensity score, where positive and negative values are associated with positive and negative actions, respectively. ICEWS has been assessed as being the most complete and reliable of several similar event datasets currently available (Ward et al., 2013a); however, on principle this analysis could be conducted with other event datasets as well.

In order to extract events which can be directly compared to model output, I subset the data and select events meeting the following criteria:

- Occurring in the years 2005-2006, so that they are strictly following the input data.
- Where both the Source and Target actors are coded with the Government sector. This excludes events involving civilians, rebels, and other actors not explicitly captured in the model.
- Where the Source and Target countries are different, excluding internal events.

With these criteria, I aggregate the events into counts of event type by directed dyad. For the purposes of this analysis, I am not looking at the ordering of these events with respect to one another, but merely that they occur.

#### Input Data

Following the methodology in Bueno de Mesquita (1998), I set agent properties based on the Correlates of War (COW) dataset: specifically, the alliance (Gibler and Sarkees, 2004) and national material capabilities (NMC) (Grieg et al., 2010; Singer, 1988) datasets.

Using the material capabilities dataset, I select the 50 states with the highest Composite Index of National Capability scores. These are the model agents, and their CINC values are used as their capabilities. Since capabilities are always used in the model to compute ratios, there is no need to normalize these scores any further. Salience values will be set randomly, as described below.

Next, I use the COW alliance data to estimate positions. Taking advantage of the socalled unipolarity of the international system (Wohlforth, 1999), I will set each actor's initial position as the similarity of its alliance portfolio to that of the United States. Similarity of alliance portfolio has been used an approximation of the alignment of two countries' interests (Bueno de Mesquita, 1975), and alignment with the system's most powerful actor serves as a rough estimate of geopolitical position.

Using the alliances in effect in the selected instantiation year of 2004, I assemble an undirected network of state-to-state relationships, with the relationship coding as the weight<sup>3</sup>. Using the resulting network's adjacency matrix, I compute the tau-b similarity (Bueno de Mesquita, 1975) between country rows, and specifically between each country's row and that of the United States.

Using the NMC and estimated positional data, I instantiate a model with 50 agents, one per state. There are 8 states which do not appear in the alliance network; at the beginning of each instantiation, I arbitrarily assign each a position of  $0.5 + \epsilon$ , where  $\epsilon \sim Normal(0, 0.1)$ . The random perturbation is in order to prevent them from becoming a single alliance bloc.

Variable	Value
$s_i$	$\sim Unif(0,1)$
$Q_i$	0.5
$T_i$	0.5
Number of steps per run	24
Number of instantiations	100

Table 3.7: Model Parameters

Unlike position and capability, there is no dataset giving us a-priori information on

 $<sup>^{3}</sup>$ Note that this network includes all the actors present in the COW data, not just the top 50 actors who will be included in the model: as we are using the network to estimate similarity of positions, even countries which are not powerful enough to be worth including as agents can serve as indicators of countries' interests, and thus worthwhile to include in the network

the agents' salience, particularly when the issue is relatively broad, as it is here. Thus, I apply salience updating rule  $S_2$ : for each model instantiation, I set all the agents' saliences randomly, such that  $s_i \sim Unif(0, 1)$ . As it is unlikely that shocks changing agents' salience will occur on a monthly basis, I hold these saliences constant throughout all 24 steps; this simplifying assumption removes one source of noise. This also corresponds to the assumption in Bueno de Mesquita (1998) (which is also used in the following chapter), whereby agents' saliences are re-randomized on timesteps corresponding to approximately two years.

Name	Position	Capability	Name	Position	Capability
Algeria	0.007	0.005	Mexico	0.734	0.012
Argentina	0.732	0.005	Morocco	0.007	0.004
Australia	0.192	0.007	Myanmar	_	0.007
Bangladesh	_	0.008	Netherlands	0.548	0.006
Belgium	0.548	0.005	Nigeria	0.059	0.007
Brazil	0.734	0.024	North Korea	0.229	0.013
Canada	0.907	0.011	Pakistan	0.092	0.013
Chile	0.732	0.003	Philippines	0.192	0.006
China	0.115	0.183	Poland	0.45	0.007
Colombia	0.734	0.006	Romania	0.271	0.004
Democratic Republic of the Congo	0.072	0.004	Russia	0.141	0.045
$\operatorname{Egypt}$	0.192	0.009	Saudi Arabia	0.007	0.009
Ethiopia	0.17	0.004	South Africa	0.192	0.007
France	0.485	0.02	South Korea	0.17	0.024
Germany	0.55	0.027	Spain	0.548	0.011
Greece	0.514	0.004	Sweden	_	0.005
India	0.324	0.07	Syria	0.008	0.004
Indonesia	_	0.013	Taiwan	_	0.008
Iran	0.168	0.012	Thailand	_	0.008
Iraq	0.007	0.006	Turkey	0.475	0.014
Israel	0.192	0.004	Ukraine	0.088	0.014
Italy	0.548	0.019	United Kingdom	0.548	0.021
Japan	0.192	0.047	United States of America	1	0.143
Kazakhstan	0.056	0.003	Venezuela	0.734	0.005
Malaysia	_	0.004	Vietnam	_	0.008

Table 3.8: Agent Data

### 3.7.2 Model Variants

There are multiple possible permutations of the various sub-model variants presented in Sections 3.2. In this section, I will test a subset of them, attempting to cover a range of possible differences. In light of the coercion question I discuss in 3.4.2, I vary the conflict sub-models  $\mathbf{W}_0$  and  $\mathbf{W}_1$  across several variants. Recall that  $\mathbf{W}_0$  has the loser in each conflict adopt the winner's position, while in  $\mathbf{W}_1$  conflicts lead neither actor to change position.

- **Model 1a** : This is the attempted replication of the original BDM model, as described in detail in Section 3.2. It comprises the sub-models  $\{T_0, S_2, W_0, R_0, O_0, M_0, A_0\}$
- **Model 1b** : This model variant is identical to the one above except for the conflict submodel. This variant will test the hypothesis that removing explicit coercion would make the model better reflect reality, and hence improve its explanatory and predictive power.  $\{T_0, S_2, W_1, R_0, O_0, M_0, A_0\}$
- **Model 2a** : This model implements several of the model variants together, while adhering to the overall structure of the original. Importantly, it implements the modified conflict schedule, with conflicts occurring only after agents have changed position in response to offers.  $\{T_1, S_2, W_0, R_1, O_0, M_1, A_0\}$
- $\label{eq:Model 2b} Model 2b : Similarly, this model implements the previously-discussed variants, changing the conflict model. \{T_1, S_2, W_1, R_1, O_0, M_1, A_0\}$
- **Model 3a** : This model variant is the most different from the baseline. In addition to the new schedule and risk sub-models, it implements the balancing and attack follow-through sub-models. Unlike the other model variants here which require both agents in a conflict to choose to enter into it, this variant allows for unilateral attacks.  $\{T_1, S_2, W_0, R_1, O_0, M_2, A_1\}$
- $\label{eq:Model 3b} \textbf{Model 3b} : Finally, once again, this variant is identical to the previous one, but uses the alternative conflict sub-model. \\ \{\textbf{T_1}, \textbf{S_2}, \textbf{W_1}, \textbf{R_1}, \textbf{O_0}, \textbf{M_2}, \textbf{A_1}\}$

### 3.7.3 Experiments and Results

For each model variant, I instantiate 100 model runs with the data in Table 3.8. I then run each for 24 steps. For each model run, I extract the Conflict events that the model generates. I aggregate these events across all runs for each dyad, and compute the mean conflict events for each dyad. These mean conflicts are the independent variables.

The dependent variables are the conflict-coded events for each dyad drawn from the ICEWS dataset for the 2005-06 period, as described above in Section 3.7.1. Figure 3.8 shows the scatter-plots of total model-generated conflicts versus ICEWS conflict events for each dyad. It is immediately obvious that a large portion of these two result sets are orthogonal: many dyads with ICEWS conflicts do not have model conflicts, and vice versa. I next test whether, for each model variant, the simulated events it generates have any predictive power for the volume of conflict events predicted for each dyad. I conduct two sets of regressions, on two variables: a linear model on the number of observed conflict events (Table 3.9), and a logistic regression on whether any conflicts were observed across each dyad (Table 3.10). In both cases, the mean model-generated conflicts are be the independent variable.

These are relatively simple statistical models; in particular, we have no reason to expect there to be a linear relationship between simulated and observed conflicts. However, these regressions will allow us to identify whether a relationship exists at all between the simulated and observed events, and what direction that relationship is – in general, we expect that if a model is generating plausible conflicts, there will be a positive relationship between simulated and observed conflict events. Furthermore, since the dependent data is the same across all model variants, we can directly compare the coefficients and fit to determine which model is closer to reality.

Examining Table 3.9 immediately shows us that there are substantial differences between the predictive power of the different models. The output of Model 1a, the original model, seems to bear no measurable relationship to the observed events. Models 1b, 3a and 3b have negative mean conflict coefficients, indicating that the more simulated conflicts we observe, the fewer real conflicts we can expect to see – the opposite of the relationship we were hoping

	1a	1b	2a	2b	3a	3b
Const.	2.82***	$3.91^{***}$	2.27***	1.40***	20.37***	21.80***
	(0.35)	(0.44)	(0.35)	(0.33)	(2.10)	(2.10)
Mean Model Conflicts	6.82	-7.90***	$33.58^{***}$	94.31***	-6.20***	-6.64***
	(8.61)	(1.99)	(4.57)	(4.38)	(0.73)	(0.73)
$R^2$	0.00	0.01	0.02	0.16	0.03	0.03

Table 3.9: ICEWS Conflict Events – Linear Regressions

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01

Table 3.10: ICEWS Conflict Events – Logistic Regressions

	1a	1b	2a	2b	3a	3b
Const.	-1.19***	-1.03***	-1.30***	-1.32***	1.22***	$1.47^{***}$
	(0.05)	(0.06)	(0.05)	(0.05)	(0.30)	(0.31)
Mean Model Conflicts	$2.27^{**}$	-1.15***	$5.52^{***}$	8.01***	-0.86***	-0.94***
	(1.03)	(0.30)	(0.62)	(0.81)	(0.10)	(0.11)
Pseudo $R^2$	0.00	0.01	0.03	0.05	0.03	0.04

Standard errors in parentheses. \* p < .1, \*\* p < .05, \*\*\* p < .01



Figure 3.8: Model-Generated vs ICEWS Conflict Events

for. Only Models 2a and 2b have coefficients which are significant and positive, showing that the more conflicts these models predict, the more we can expect to see. Between these two, Model 2b has a substantially higher  $R^2$  value, indicating a better fit.

The same pattern holds when examining Table 3.10, which coarse-grains the ICEWS data further, asking whether more simulated conflicts make it more likely that we will observe any conflicts at all between the corresponding dyad. The coefficients show the same pattern of signs, indicating that only Models 2a and 2b generate simulated conflicts which are useful for predicting real conflicts. Between these, again, Model 2b shows a better fit, though the advantage is less stark here.

### 3.7.4 Discussion

With these results in hand, we can draw several conclusions. At the crudest level, we see that across all the model variants, the number of conflict events generated is much smaller than the number of observed conflict events. Many dyads with recorded conflict events are not predicted by the models, while the models in turn predict conflicts across many dyads where none are observed. The overall relationships between the model outputs and observed data are weak.

Despite this, we must recall that there are countless issues that at any given time are driving the micro-level interactions between countries in the world. Across all the model variants presented here, we have attempted to capture only those issues reflected in states' alliances, and further compressed them onto a single one-dimensional space. Given the simplicity of the model itself, and the global scope that these particular instantiations are addressing, and it appears remarkable that there is any statistical relationship between the model output and reality at all. And indeed, several of the model variants – particularly 2a and 2b – appear to be valid, if weak, predictors of real-world conflicts, at least as captured by the ICEWS data.

Furthermore, this methodology allows direct comparison between model variants, each comprising combinations of different sub-models. Examination of the scatter-plots of Figure 3.8 shows us that different variants generate different patterns of events. We can immediately see, in particular, that Models 3a and 3b generate a higher volume of conflict events,

and that the quantities for each dyad are approximately normally-distributed, conditional on there being any conflicts at all for those dyads. We can also immediately identify that variants 2a and 2b outperform the other variants as predictors, and that between the two of them variant 2b has superior predictive power. This, in turn, is in line with my hypothesis that the conflict sub-model  $W_1$ , where conflicts do not result in agents changing their position, is will capture the underlying system better than the baseline variant  $W_0$ .

# 3.8 Conclusions

In this chapter, I have argued that Bruce Bueno de Mesquita's Expected Utility Model is best understood not as a game-theoretic model but as an agent-based one. The model and its agents are composed of several sub-models, with the agents interacting with one another following certain behavior rules. These rules, in turn, are based on heuristics and embedded assumptions, many of which are not addressed in prior papers describing the model and its applications. I have described my reimplementation of the model as an object-oriented agent-based model. While doing so, I was able to identify its sub-models and heuristics, and discuss how they fit together – and how they may be altered and experimented with.

My reimplementation of the model was unable to reproduce several results generated by the original model and reported in the prior literature, in particular the results of Bueno de Mesquita (1994) and Bueno de Mesquita (2002, chapter 6). It is possible that despite my best efforts, my own model code misimplements one of the formulas from the prior literature, or has another unidentified bug which is the cause of the difference between its outputs and the prior results. It is also possible that I have implemented one of the submodels which is not well described in the prior literature (particularly the offer acceptance sub-model) differently from the original model. Without access to the original application or source code, it is impossible to fully isolate the source of the differences. Nevertheless, I do have access to the source code of the Scholz, Calbert, and Smith (2011) model implementation, as well as to a replication of that implementation (McKibben-Sanders, 2014), both of which successfully reproduce the Bueno de Mesquita (1994) results. As I discuss in Section 3.3.1, I identify a key way in which this implementation contradicts both the stated form and underlying assumptions of the original model, by having the probabilities of victory in a bilateral conflict not sum to 1. It may be the case that by coincidence this model implementation featured multiple errors which lined up to produce the right results for the wrong reasons; however, as Scholz, Calbert, and Smith (2011) themselves note, the odds of generating correct results by chance are extremely low. It seems more likely that they have inadvertently identified an error in a particular implementation of the original model – one which, presumably, was corrected in later versions. Importantly, the Scholz, Calbert, and Smith (2011) implementation was unable to replicate results reported in other publications. This tells us that their implementation cannot be considered a perfect replication either – and that published original model results cannot necessarily be treated as 'ground truth' either.

In Section 3.5, I described and implemented a sensitivity analysis methodology using randomly generated cases. This analysis indicates that the original model has little meaningful stochasticity, and that in many cases converges to a single median outcome. In fact, it suggests that there are two distinct classes of input sets, those which have deterministic outcomes and those that do not. This distinction should be investigated further, in particular in applications of the model to the analysis of real-world cases. Even relatively slight perturbation of the model inputs introduced an additional source of stochasticity, eliminating the determinism. It also does not appear to be the case that convergent models have less uncertainty once the inputs are perturbed, suggesting that the convergence is an artifact of the specific input parameters rather than evidence that the system itself has an inherent attractor. However, the means of the perturbed runs tend to remain very close to the unperturbed mean; the perturbation doesn't push the system into a meaningfully different region of the outcome range. This finding, in turn, lends justification to practice of original model results being reported as point predictions rather than distributions.

The model itself is driven largely by heuristics, with the model and agent behavior not always grounded in specific theories or sources. This on its own is not necessarily a major flaw, and indeed is common for other agent-based models of international relations (e.g. Axelrod, 1997; Taylor et al., 2008). By reimplementing the model as an ABM, I was able to highlight these heuristics and assumptions to a greater extent than had been previously reported. More importantly, this implementation makes it straightforward to test the model with different sub-models and heuristics reflecting different assumptions. I have proposed several such variant sub-models, as well as a notation system to quickly describe different model variants comprising different sub-models.

By enabling multiple model variants to be run with the same input data, I also test the predictive power of those variants side by side. Political event data provides a source of relevant real-world interactions, and to the best of my knowledge, no EUM-type model has been applied to event data before. However, as I argue in Section 3.6, the dynamics and outputs of the model correspond naturally to categories present in event data. Indeed, I demonstrate in Section 3.7 that even using a coarse instantiation without a well-defined issue of contention, several variants of the model generate sets of simulated events bearing weak but statistically significant predictive power for corresponding real-world events.

More important than the specific findings of these models, however, is a demonstration of how such models can be applied comparatively. By using fixed input data and the same target output data, we can run multiple instantiations of different model variants, generating sets of model traces – and especially sets of simulated conflict events – for each. This allows us to find the different patterns which emerge from different models, such as the difference in conflict distributions between model variants 1, 2 and 3 visible in Figure 3.8. By regressing the model outputs against observed data, we can test which among several models has more explanatory or predictive power. Comparing the performance of model variants which differ only by a single sub-model (such as 3a and 3b above, which implement conflict rules  $W_0$ and  $W_1$ , respectively) can highlight the effects of those sub-model variants, and allow us to test the hypotheses they describe.

The model and methodology I have described here mark the beginning, not the end, of a program of study. I have explored only a subset of possible combinations of the sub-models

I have presented here, which in turn are only a subset of plausible sub-model variants. Applying the sensitivity analysis to additional model variants will allow us to characterize inherent dynamics or artifacts that the variants exhibit, which in turn can help isolate meaningful predictions when instantiated with real data. Testing these models against different cases and sets of political actors will help identify whether one mode variant is strictly superior, or whether different variants are appropriate for different categories of cases and issues. It may also be the case that ensembles of model variants may provide greater predictive power than any single variant alone.

There are still assumptions embedded in the current architecture, independent of any sub-model. For example, conflicts are currently strictly bilateral, and agents contributing capability to one side or another face no effects if the principals win or lose, unlike in the Wise, Lester, and Efird (2015a) model framework. Similarly, the model as presented here only models issues as one-dimensional ranges; however, Wise, Lester, and Efird (2015b) have proposed an extension allowing their model to handle multiple issues, including ones represented non-numerically. The agent activation regime plays an important role in the behavior of agent-based models (Comer, 2014), and the validity of the simultaneous activation assumption may depend on the specifics of the scenario being simulated. The simultaneous activation the model currently features may be appropriate for modeling negotiators in close proximity (perhaps in the proverbial smoke-filled room) or when the actors are states acting over months or years, but may be less so for intermediate scales. Varying and sequencing agent activation may allow the model to incorporate time more explicitly, improve modeling of interactions between agents with different paces of decisionmaking, and endogenize incomplete information.

There are likely other methods of utilizing models within the framework described here. I have conducted preliminary experiments with applying the model deductively: given observed sequences of events, to identify initial inputs which are most likely to give rise to similar sequences. This may allow us to estimate the actors' underlying preferences and capabilities in the absence of reliable quantitative data or even subject-matter expertise. My initial experiments suggest that such an application is computationally plausible, but that multiple, substantially different inputs may yield similar event sequences.

Finally, the findings I have described here provide reasons for both appreciation and skepticism of the original BDM model. Building on the published descriptions of the model have allowed me to build a generalized, flexible agent-based model of political interaction with rich outputs and surprising predictive power despite its relative simplicity. The sensitivity analyses I have conducted lend some credence to the practice of reporting the model's point results, though also highlight the degree to which this elides uncertainty that is indeed present in the model. The difficulty in reproducing the model from published descriptions should give us pause. I have identified numerous assumptions and heuristics in the model description, very few of which are explicitly stated, much less supported either theoretically or empirically. This makes it all the more important to implement not just a single model but multiple variants, in order to rigorously test and evaluate them. If we do so, however, it promises to offer a valuable tool for operationalizing and testing competing heuristics, intuitions and theories of competing political actors in the international system and beyond.
# Chapter 4: Modeling the Cold War: An Application of the Expected Utility Model

# 4.1 Introduction

The Cold War between the United States and the Soviet Union (also referred to interchangeably as the Union of Soviet Socialist Republics, or USSR) was the defining feature of the international order in the second half of the 20th century, helping drive everything from local insurgencies to space exploration. The conflict was rooted both in ideological differences and a desire for security. Indeed, both sides possessed not only powerful conventional militaries with global reach, but arsenals of nuclear weapons capable of guaranteeing socalled Mutually Assured Destruction. Each side posed an existential risk not only to the other's system of government, but to the lives of a large proportion of their citizens. Both sides also had stable blocs of aligned countries; in Europe this divide was geographic, with the 'Iron Curtain' separating the west and east. The apparent parity of the two sides, and the stability of their alliances, led to an expectation that the cold war between them would persist indefinitely, terminated perhaps only in apocalyptic conflict.

Yet it was not the world but the Cold War, and the Soviet Union itself, which ended – "Not with a bang but with a whimper," as Dobson and Marsh (2007), Prados (2011) and others quote Eliot (1934). The end of the Cold War took policymakers and scholars both by surprise. In the aftermath, many political scientists (e.g. Gaddis, 1992; Lebow and Risse-Kappen, 1995) suggested that this was evidence of shortcomings in the state of international relations theory and methodology. Bueno de Mesquita (1998) offered an alternative explanation: that the end of the Cold War *was* predictable by contemporary international relations theory, but that it was an emergent phenomenon: an unexpected higher-level outcome generated from the combination of multiple lower-level interactions which themselves are well-understood. Indeed, if the international system is complex, it should not be surprising for it to generate emergent phenomena.

Bueno de Mesquita (1998) tests this claim by simulating the Cold War using the Expected Utility Model. Computational modeling is uniquely suited to generating macro-level emergent phenomena from micro-level interactions; thus, if the model is well-grounded in theory and gives rise to the end of the Cold War, this is evidence that the theory itself was not mistaken, but rather that it was not properly applied, or its implications not fully understood. Furthermore, if the model can outperform experts and predict a major historic outcome from the relatively minimal data it is instantiated with, this would be strong evidence in favor of the model's utility and predictive power. Indeed, this paper argued that such a model, and the EUM in particular, can predict the end of the Cold War, with a victory for the US and its allies.

In the previous chapter, I presented my own reproduction and extension of the Expected Utility model that Bueno de Mesquita (1998) uses. In this chapter, I attempt to apply my model variants to the same Cold War case study. This serves several purposes. It attempts to replicate the previous study, while also providing an additional test of my reproduction: if my direct reimplementation model produces similar outcomes to the original, this is evidence for an accurate replication. Additionally, as with the original model, if my models are capable of successfully predicting the end of the Cold War from historic data, this is evidence of their usefulness. As discussed in the previous chapter, my own models extend the original by recording not just outcomes but key events that lead up to them. Comparing these events – and in particular, conflicts – to historic data provides an additional external test of the models' validity. Finally, by comparing the predictive power of different model variants, I can attempt to identify the one with the most explanatory power.

The rest of this chapter is organized as follows. I describe the input data (Section 4.2.1), how I use it to instantiate runs of two model variants (Section 4.2.2), and data to which I will compare to model outputs (Section 4.2.3). In Section 4.3, I present the results of four experiments, each one a combination of an input dataset with a model variant. I describe the results of each experiment qualitatively, and test their output's power as a predictor of observed conflicts and post-Cold War alliances. Finally, In Section 4.4, I discuss the results more broadly, highlighting and attempting to explain the models' strengths and weaknesses.

# 4.2 Modeling Methodology

The procedure for this analysis follows a similar process to that described in the previous chapter. Initially empirical data is used to generate a set of actors, with position and capability values for each. These actors then serve as inputs to the model variants, each of which is instantiated and run multiple times, generating collections of outputs. Finally, I analyze these outcomes and compare them to real-world data.

#### 4.2.1 Input Data

To instantiate a model, we require several things: a list of agents, which in this case are countries active in the international system during the Cold War; and for each agent, an initial position on the one-dimensional position space and a capability value. There is no need here to attempt to estimate salience, since this will be one of the independent values which will be randomized for all agents, as described below.

Bueno de Mesquita (1998) generated the model inputs from data collected by the Correlates of War (COW) project. The list of actors, and their capabilities, are drawn from the National Material Capabilities dataset (Singer, 1972, 1988), with the model using the 36 most-powerful actors in the starting year of 1948; their capabilities are normalized so that all capabilities sum to 100. Their positions are derived from the COW alliance dataset (Small and Singer, 1969), measuring each agent's closeness to the security preferences of the United States or Soviet Union. This closeness is defined as the tau-b measure of similarity in alliance portfolios (Bueno de Mesquita, 1975), normalized such that a value of +100 is full match with the United States, and -100 is a full match with the Soviet Union. In order to make this data compatible with my implementation of the model, I simply rescale the

Country	Capability	Position	Country	Capability	Position
Argentina	0.972	0.948	Italy	2.426	0.507
Australia	0.889	0.507	Mexico	0.774	0.948
Belgium	1.182	0.514	Norway	0.23	0.507
Brazil	0.993	0.948	Netherlands	0.836	0.514
Bulgaria	0.345	0.000	Pakistan	1.485	0.507
Canada	1.61	0.781	Philippines	0.408	0.507
China	11.941	0.507	Poland	3.273	0.045
Czechoslovakia	1.401	0.045	Romania	0.606	0.045
Denmark	0.24	0.507	Saudi Arabia	0.125	0.514
Egypt	0.408	0.516	South Africa	0.68	0.507
England	7.863	0.518	USSR	18.256	0.000
France	3.597	0.514	Spain	1.683	0.507
Greece	0.418	0.507	Sweden	0.648	0.507
Hungary	0.45	0.045	Syria	0.104	0.514
India	2.468	0.507	Thailand	0.414	0.507
Iran	0.491	0.512	Turkey	1.347	0.512
Iraq	0.157	0.519	USA	29.956	1.000
Israel	0.125	0.507	Yugoslavia	0.891	0.000

Table 4.1: Original Inputs (Bueno de Mesquita, 1998) with Positions Rescaled

starting positions so that they fall on the [0, 1] range, with 0 being the Soviet Union's starting position and 1 being that of the United States. The original model inputs are reported in Bueno de Mesquita (1998) and reproduced in Table 4.1, with the rescaled positions. I term this the Original Input data. Figure 4.1 visually demonstrates the starting distribution of power: the United States and its core allies (those whose positions  $\approx 1$ ) have a slight advantage in total capability over the Soviet Union and its core allies (position  $\approx 0$ ); however, the largest capability bloc is in the center, not yet committed to either side.

I then attempt to reproduce the input data from the current, up-to-date COW datasets (Gibler, 2009; Grieg et al., 2010); however, this did not result in the same inputs as provided in Bueno de Mesquita (1998). The initial difference is the capability scores – not just the Composite Index of National Capability (CINC) scores themselves of the various countries for 1948, but their rankings as well. This is likely due to updated data and methodologies, but already indicates a difficulty in replicating the previous study. Since in this case I have



Figure 4.1: Original Inputs – Initial Distribution of Capabilities

access to all the countries for which there is 1948 data, I will include all of them in the model, rather than utilize an arbitrary cutoff.

Next, we must assign each agent to a position. To do so, I construct an undirected network from the most recent COW alliance data for 1948, where the edge between each pair of countries is coded based on the alliance type, as summarized in Table 4.2. We can then calculate similarity scores between the the actors' rows or columns in the adjacency matrix, which represent the actors' alliance portfolio. To maintain correspondence with the previous data, I use the tau-b similarity score between portfolios.

It is here that we run into a question unaddressed in Bueno de Mesquita (1998). Comparing each country's alliance portfolio to those of the United States and the Soviet Union produces not a single value per country, but two: one for each superpower. As shown in Figure 4.2, the relationship between the similarities is not completely linear, meaning we need to choose a way to reduce the two-dimensional security position space to one dimension. Since the two dimensions are highly correlated (correlation coefficient -0.81), I use principal component analysis (Tipping and Bishop, 1999) to project the two-dimensional

Alliance Code	Type
1	Entente: "one or both states in the dyad had an understanding that consultations with the other state in the dyad would take place if a crisis occurred."
2	Non-Aggression: "one or both states in the dyad had a non- aggression pact with the other state in the dyad."
3	Neutrality: "one or both states in the dyad had a neutrality pact with the other state in the dyad."
4	Defense: "one or both states in the dyad had a defense pact with the other state in the dyad."

Table 4.2: Alliance Coding (descriptions quoted from Gibler et al., 2013)



Figure 4.2: Tau-b Similarities of Actors to United States and Soviet Union

data onto a one-dimensional space, which I then rescale so that the values are between 0 and 1.

I term the new model inputs, appropriately enough, the Updated Input data; Table 4.3

reports the new positions and capabilities<sup>1</sup>. This input set includes all the countries for which data is available in 1948. Like the Original Input data, it excludes Japan and West and East Germany, which had not yet become sovereign at that point. Figure 4.3 visually shows the initial distribution of power in this scenario: while in this case the United States and its core allies have an even larger advantage over the Soviet Union and its core allies than in the Original Input. The key difference, however, is that instead of a single bloc between the two powers' positions, there are agents at multiple intermediate positions; the single largest power bloc at approximately 0.6, dominated by China and India, is slightly closer to the United States, but there are more agents whose positions are closer to that of the Soviet Union.



Figure 4.3: Updated Inputs - Initial Distribution of Capabilities

<sup>&</sup>lt;sup>1</sup>To make this data directly comparable with the Original Inputs in Table 4.1, I have rescaled the CINC scores so that they sum to 100. Since Capability scores are used solely as ratios, this rescaling has no effect on the model.

Country	Capability	Position	Country	Capability	Position	Country	Capability	Position
Afghanistan	0.23	0.58	Haiti	0.04	0.98	Paraguay	0.04	0.98
Albania	0.09	0.36	Honduras	0.02	0.98	Peru	0.17	0.98
Argentina	0.82	0.98	Hungary	0.41	0.30	Philippines	0.33	0.56
Australia	0.83	0.57	Iceland	0.00	0.56	Poland	2.98	0.30
Belgium	1.09	0.42	India	5.23	0.56	Portugal	0.26	0.46
Bolivia	0.06	0.98	Iran	0.43	0.51	Romania	0.55	0.34
Brazil	1.22	0.98	Iraq	0.13	0.49	Russia	16.36	0.00
Bulgaria	0.31	0.42	Ireland	0.10	0.56	Saudi Arabia	0.11	0.59
Canada	1.26	0.56	Israel	0.14	0.56	South Africa	0.52	0.56
Chile	0.20	0.98	Italy	1.92	0.56	Spain	1.48	0.57
China	11.48	0.57	Jordan	0.02	0.52	Sri Lanka	0.09	0.56
Colombia	0.26	0.98	Lebanon	0.03	0.59	Sweden	0.60	0.56
Costa Rica	0.01	0.98	Liberia	0.01	0.56	Switzerland	0.20	0.56
Cuba	0.17	0.98	Luxembourg	0.29	0.42	Syria	0.09	0.59
Czechoslovakia	1.26	0.34	Mexico	0.66	0.98	Thailand	0.34	0.56
Denmark	0.21	0.56	Mongolia	0.04	0.57	Turkey	1.17	0.51
Dominican Republic	0.06	0.98	Myanmar	0.23	0.56	United Kingdom	7.52	0.54
Ecuador	0.09	0.98	Nepal	0.11	0.56	United States of America	29.39	1.00
Egypt	0.54	0.52	Netherlands	0.83	0.42	Uruguay	0.08	0.98
El Salvador	0.03	0.98	New Zealand	0.09	0.57	Venezuela	0.13	0.98
Ethiopia	0.23	0.56	Nicaragua	0.01	0.98	Yemen Arab Republic	0.06	0.59
Finland	0.19	0.57	North Korea	0.36	0.56	Yugoslavia	0.80	0.32
France	3.25	0.51	Norway	0.16	0.56			
Greece	0.36	0.56	Pakistan	1.18	0.56			
Guatemala	0.07	0.98	Panama	0.01	0.98			

Table 4.3: Updated Inputs

## 4.2.2 Model Instantiations

I use the input datasets above to instantiate two model variants, drawn from the previous chapter's Section 3.7.3. The first corresponds to Model 1a, and is the attempt to reproduce the EUM as originally described across multiple papers (Bueno de Mesquita, 1994, 1997, 2002); I will refer to this as the Baseline model. The second corresponds to Model 2b, chosen because it was the most predictive of the ICEWS events in the previous chapter; I will refer to this as the Updated model.

These model variants differ from the versions in the previous chapter in their salience updating sub-model. As Bueno de Mesquita (1998) notes, there are countless unpredictable elements affecting how states will perceive, and act on, their international security interests in the context of the Cold War, from domestic political changes to natural disasters. Rather than extend the model to explicitly capture more and more of these factors, we implicitly incorporate them by randomizing agent salience values. At the beginning of each step of the model, each agent's salience is randomly drawn from a uniform distribution over the [0, 1] range – this is salience sub-model variant  $\mathbf{S}_{1,1}$ , in the notation presented in the previous chapter. Since capability is always multiplied by salience, this effectively adds a noise factor to the agents' strength as well. While agents with greater capabilities are likely to remain stronger than ones with lower capabilities, there exist configurations of salience values under which the generally weaker agent temporarily has an advantage over a stronger agent.

In the notation introduced in the previous chapter, these models can be denoted as:

## Baseline: $\{T_0, S_{1,1}, W_0, R_0, O_0, M_0, A_0\}$

Updated:  $\{T_1, S_{1,1}, W_1, R_1, O_0, M_1, A_0\}$ 

Taken together, the two model variants and two input datasets yield four experiments in all, as shown in Table 4.4.

For all experiments, I instantiate all agents with parameters Q = 0.5 and T = 0.5indicating maximum uncertainty about the world. Each model is run for 25 steps. Bueno de Mesquita (1998) hypothesizes that this is roughly equivalent to 50 years of history,

 Table 4.4: Experiment Summary

	Baseline Model	Updated Model
Original Inputs	Experiment 1	Experiment 2
Updated Inputs	Experiment 3	Experiment 4

meaning that agents experience salience shocks, engage in conflicts, and change position approximately once every two years. With 1948 as the starting year, the models will generate data corresponding approximately to 1998 as the end year.

To briefly restate the model: during each step, each agent evaluates its expected utility from threatening conflict against each other agents in an attempt to coerce the other agents to adopt its current position; when that expected utility is positive, the agent issues the threat. The agents then evaluate all their incoming threats, choosing the one requiring the least movement, and moving based on their own expected utility (interpreted here as fear of a conflict), either compromising and moving slightly towards the threat sender's position or capitulating and adopting that position outright. In the Baseline variant, but not the Updated one, if agent A chooses B's threat, while B chooses C's threat, B backs out of its decision and does not move.

If two agents threaten each other and do not choose to compromise with or capitulate to another agent, they enter into a conflict. All agents then contribute capability to the side they prefer, in proportion with the strength of their preference. The winner is chosen randomly, with probabilities based on each agent's share of total capability contributed to the conflict. In the Baseline variant, the loser of a conflict adopts the winner's position; in the Updated variant, neither agent changes position, regardless of the conflict outcome.

Due to the large number of independent random variables driving the model, I run 1,000 instantiations of each experiment. This number appears to be sufficiently large to sample from the many different possible configurations of saliences and conflict outcomes<sup>2</sup>. Furthermore, since we are concerned with commonly-emerging outcomes and behaviors,

<sup>&</sup>lt;sup>2</sup>excursions using 10,000 instantiations did not produce meaningfully different outcome distributions.

particular phenomena occurring less frequently than in 0.1% of model runs are unlikely to be of practical interest.

## 4.2.3 Output Data

Each model run generates several outputs. At the beginning of each step of the run, it records the position currently held by each agent, as well as the current median voter position, and capability-weighted mean position. During each step, the model also logs all conflicts which occur between agents. I collect and aggregate these outputs across all 1,000 runs of each experiment.

I then compare these outputs to two Correlates of War datasets: militarized interstate disputes (MIDs) for the 1948-1998 timeframe, and alliance data for 1998. The MID dataset provides an expert-coded listing of all cases where one state threatened or used force against another (Jones, Bremer, and Singer, 1996). While there are many forms of conflict, militarized disputes, which can escalate into wars, are particularly important to the international system, and thus of particular interest for forecasting. I will attempt to see whether the number of conflict events between pairs of actors generated by the model runs is a useful predictor of the number of conflicts between states in the relevant time period.

Similarly, the model runs produce traces of agent positions. Agents with closer positions are more likely to support one another more strongly in conflict with other agents, i.e. they behave the way we would expect allied countries to behave. Thus, I will test whether the distance between agents is a predictor of alliances between the corresponding countries.

## 4.3 Results

## 4.3.1 Overview

The initial question posed by Bueno de Mesquita (1998) was whether the end of the Cold War – that is, a substantial shift in the median international security position to the advantage of the United States – could be generated from the model. Figure 4.4 shows the cumulative distribution of model end positions reported by that paper. Note the sharp increase in the cumulative distribution curve very close to 0 (equivalent to 0.5, in the variant used here), indicating that one of the most common types of outcomes is one where the Cold War persists at the end of the model run, with a median position between the starting position of both superpowers and no clear advantage to either side. However, the distribution is also clearly non-symmetric, with more outcomes favoring the United States (closer to +100) than the USSR. The conclusion of the original paper was thus that the end of the Cold War, with an American victory, was not as unforeseeable as others had argued, and in fact could be emerged from the model.



Figure 4.4: Model Outcome Cumulative Distribution, Copied From Bueno de Mesquita (1998)

Figure 4.5 shows the same cumulative distributions across the four experiments<sup>3</sup>. Immediately, we note that the overall shapes are similar, though not identical. The bulk of outcomes are very close to the center, with substantially more US-victory outcomes than

 $<sup>^{3}</sup>$ Bueno de Mesquita (1998) does not provide all the results of each model run, meaning that we cannot directly superimpose the results shown in Figures 4.4 and 4.5.

USSR-victory ones. In particular, Experiment 1, which attempts to replicate the original model, has almost no USSR-victory outcomes, and more outcomes where the median position does not indicate a victory to either side.



Figure 4.5: Model Outcomes Cumulative Distributions

Let us take a closer look at the behaviors these models are generating. Figure 4.6 shows the superimposed traces of the median position for all the model runs of all experiments. The opacity of each trace line is low, so that the darker a line appears, the more traces exhibited that particular transition. While it is impossible to discern individual runs here, these visualizations serve to highlight recurring patterns. Two patterns are obvious across all four experiments. The first is that, as the cumulative distributions in Figure 4.5 suggest, the most common positions for the model at any given step are at or close to the starting median – by and large, the median does not change very often. The next key pattern is that, across all the experiments, the median position exhibits a sawtooth pattern, diverging from the 0.5 position towards one extreme or the other and immediately returning. In other words, one of the sides (though more frequently the United States) gains a brief advantage, and then the status quo quickly reasserts itself. This sawtooth pattern is particularly distinct and uniform in the Baseline model used in Experiments 1 and 3. Experiment 2, and to a lesser extent Experiment 4, show a wider variety of traces. In particular, close



Figure 4.6: Median Position Traces

examination of the Experiment 2 traces indicates more cases where changes in median are not one-round spikes, but transitions to a new stable position. Finally, Experiment 4 shows a decline in movement as time advances, suggesting that the agents are sorting themselves into a more stable configuration around the same median.

The sawtooth pattern raises a question: do the model runs ending in a US (or Soviet, for that matter) victory represent a true long-term shift in the status quo, or are they simply cases where a temporary spike towards one side or the other coincides with the final step of the model? In fact, a close examination of the model runs suggests that the latter is exactly what is happening in Experiments 1 and 3. Figure 4.7 shows example traces from randomly-selected runs which end with a median position of 0.7 or greater, indicating a US victory. In the traces from Experiments 1 and 3, the median position exhibits several shortlived excursions away from the starting position, quickly reverting after each. There is no reason to suppose that the spike preceding the end of the run represents a true change from this pattern. The trace from Experiment 2, in contrast, shows that the median position had gone up prior to the final step of the run, and remained relatively stable at its new position. The Experiment 4 trace also shows the run achieving a new stable position, with smaller spikes away from it, including one coinciding with the terminal step. These patterns recur across the runs in each experiment that are not visualized here, including those not ending in a US victory.

A thorough examination of model runs suggests that the spikes – including the terminal ones – are driven less by changes in agent positions than by the changes of agent salience values. Importantly, these changes are not accompanied by a similar change in the weighted mean position – which is to say, not driven by similar changes in the positions of the agents themselves. Recall that as detailed in the previous chapter, in Section 3.2.2, the median is computed by checking which agent position has the highest probability of defeating all other agent positions in bilateral conflicts; thus, they are affected by the stochastic salience changes. In contrast, the mean position is weighted by (fixed) capabilities, and thus is not directly affected by salience variations.

This phenomenon is illustrated in Figure 4.8, showing the individual agent positions over time for the same runs as in Figure 4.7. In this figure, each line corresponds to an agent: the y axis indicates each agent's position at each step, shown on the x axis; line widths correspond to the agents' capabilities. The top-most line is always the position of the United States agent and the bottom-most line is the USSR agent. The traces for all the experiments show that the radical, temporary spikes in the median position do not correspond to similar sharp shifts in any agent position. In Experiment 1, the agents exhibit a slight movement towards the center of the position space. Note, however, that the movement by the USSR agent and its allies (at the bottom of the graph) are more substantial than those of the US and its aligned agents at the top. In Experiment 3, however, very little movement occurs at all. In Experiments 2 and 4, in contrast, the agents exhibit sharp, more substantial changes in their positions, which correspond to the changes in the median, highlighting that the median change is not being driven solely here by stochastic changes in salience.

In fact, in the Experiment 2 and 4 traces, we can see a particularly interesting story emerge. Early on, several agents whose positions are closer to the median are drawn closer to the position of the United States; this is followed by several Soviet-aligned agents being drawn 'upward' towards the United States's preferred position across multiple steps. The Soviet Union itself lags behind these agents for a number of steps. However, eventually it too is driven to adjust its position upwards, in the Experiment 4 trace even eventually joining the US-aligned cluster. In effect, this model generates a recognizable notional history of the Cold War leading to a US-dominated unipolar world.



Figure 4.7: Example Model Traces – Median Positions

## 4.3.2 Predicted Conflicts

The models do not only generate overall positional traces, but simulated events, and conflicts in particular. These conflicts are an important part of the simulated history, not only because they can drive substantial changes in agent positions but because they provide useful anchor-points for comparing model runs to real history. We have robust historic data on wars and lesser militarized interstate disputes, and a well-developed understanding of international partnerships and rivalries, which allow us to assess the plausibility of the conflicts the model produces. Furthermore, conflicts are an important area for prediction and forecasting, and it is valuable to assess whether these models can be used for such prediction.



(d) Experiment 4

Figure 4.8: Example Model Traces – Agent Positions

Each model run in each experiment collects the conflict events generated across all run steps; for each, the code computes the mean number of conflicts occurring between each dyad of agents across all models. These dyads are undirected, since in the model variants applied here both parties must mutually decide to start a conflict for it to occur. These simulated conflict counts are then merged with the count of militarized interstate disputes for each dyad, as described above.

The correlations between the predicted and observed number of conflicts are small. Most conflicts generated by the model do not reflect real conflicts, while most real conflicts are not mirrored in the simulated data. Nevertheless, we may still ask whether the model provides any statistically useful information. In order to do so, I run two regressions on each experiment output, with predicted conflicts as the independent variable: a linear regression, with observed MIDs as the dependent variable, shown in Table 4.5; and a logistic regression, with a dependent dummy variable for the presence of any MIDs on the dyad, in Table 4.6. The model generates relatively few conflicts – like real-world wars, these are rare events (King and Zeng, 2001), meaning that the mean conflicts values are small, leading to largemagnitude coefficients.

First, let us examine the linear regressions. Note that we have no particular reason to believe that the relationship between the number of model-generated and observed conflicts is linear. However, if more simulated conflicts do in fact predict more observed conflicts, we would expect the coefficient to be positive. In fact, all the coefficients are positive, and all but the one for Experiment 1 are statistically significant, suggesting that the models are generating some useful information as to the number of observed conflicts.

Further evidence of this is visible in Table 4.6. These regressions test whether more model conflicts between dyads of actors predict *any* conflicts occurring in the relevant time period. In this case the coefficient values are positive and significant across all the models, indicating that when the model predicts more conflicts between a pair of countries, we are indeed more likely to have observed at least one conflict between them in the historic data.

While  $R^2$  is an imperfect measure of the goodness-of-fit of a statistical model, it is an

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Const.	$1.043^{***}$	$1.014^{***}$	$0.392^{***}$	$0.362^{***}$
	(0.183)	(0.173)	(0.049)	(0.044)
Model Conflicts	1.473	8.327***	$17.409^{***}$	$1015.746^{***}$
	(1.235)	(1.950)	(5.090)	(40.768)
Adjusted $R^2$	0.001	0.027	0.004	0.195

Table 4.5: Militarized Interstate Disputes – Linear Regressions

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

Table 4.6: Militarized Interstate Disputes – Logistic Regressions

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Const.	-1.300***	-1.422***	$-2.201^{***}$	$-2.197^{***}$
	(0.101)	(0.111)	(0.067)	(0.066)
Model Conflicts	$1.952^{***}$	$27.658^{***}$	$19.739^{***}$	$654.501^{***}$
	(0.622)	(7.190)	(4.556)	(148.758)
McFadden's $\mathbb{R}^2$	0.015	0.030	0.009	0.016
	-			

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

acceptable way of comparing models of the same type and the same dependent variable (Wooldridge, 2008). The  $R^2$  values across all models are low, indicating that they are explaining only a small fraction of the variance observed in the MID data. Nevertheless, we can see that the  $R^2$  value on the Experiment 4 linear regression is an order of magnitude better than for the other experiments, and that this is not the case for the logistic regression. This suggests that the model is more accurately predicting the number of conflicts, conditional on conflicts occurring at all. Across both types of regressions, furthermore, we can see that Experiment 3 has the least explanatory power.

It is also useful to look at the specific conflicts being predicted. Table 4.7 shows the top ten most frequently predicted conflict dyads for each experiment. One finding that stands out here is that the Baseline and Updated models appear to generate two different classes of

Experiment 1		Experiment 2			
England	Iraq	Bulgari	a	USA	
Egypt	England	USSR		USA	
Belgium	France	USA		Yugoslavia	
Belgium	Iran	China		USA	
Argentina	USA	Englan	d	USSR	
Belgium	Egypt	Austral	ia	USA	
Argentina	Brazil	China		England	
Egypt	Iraq	Poland		USA	
Belgium	Netherlands	Bulgari	a	England	
Argentina	Canada	Czecho	slovakia	USA	
Experiment 3		Experim	Experiment 4		
Egypt	Jordan	USA	USSR		
France	Turkey	USA	United I	Kingdom	
Belgium	Netherlands	USSR	Argentir	ıa	
France	Iran	USA	India		
Belgium	Bulgaria	USA	Yugoslav	via	
Afghanistan	Saudi Arabia	USA	China		
Belgium	Luxembourg	USA	Czechos	lovakia	
Czechoslovakia Romania		China	France		
Poland	Hungary	USA	Portugal	1	
Turkey	Iran	USA	Poland		

Table 4.7: Top Predicted Conflict Dyads

conflict. The Baseline model generates no direct conflicts between the two rival superpowers in either Experiment 1 or 3, and very few even directly involving either. Instead, the conflicts are largely between lower-tier agents, with positions closer to the median. In historic terms, these are proxy wars: the belligerents are not the superpowers themselves, but each draws some support from a different superpower and its close allies. These contributions are often small, as their preference for one side over the other may be relatively weak. The stakes of the conflicts, in terms of the overall orientation of the simulated international system, are correspondingly weak. In contrast, the Updated model sees much more direct involvement by the two main superpowers directly, including direct conflicts between them. The conflicts that are not between the two superpowers directly appear to often take the form of one superpower (most often the USA) attempting to coerce a much weaker country. Several of the predicted conflict dyads stand out as being particularly plausible, or implausible. The United States and Soviet Union, with high conflict counts in Experiments 2 and 4, are the most frequent dyad observed in the MID dataset for the 1948-98 period. However, these disputes never escalated to a full military conflict. Such a conflict, if it had occurred, would likely have had catastrophic consequences far outside the scope of this model. Interestingly, despite not incorporating geography, the Baseline model generates several regional conflict dyads which are in fact observed in reality, including Egypt and Iraq (Podeh, 1995) and Brazil-Argentina (Mariano, 2013) – though, again, these real-world rivalries were of relatively low military intensity.

Some of the implausible conflicts may be artifacts of the model input data. For example, the conflicts between the United States and the United Kingdom (UK) are attributable to the UK's initial position, which in both the Original and Updated data is close to 0.5. However, this is almost certainly an underestimate of its alignment with the United States, particularly in light of the signing of the North Atlantic Treaty (United States et al., 1948) in the following year. Similarly, far from being in conflict with one another, Belgium, France, and the Netherlands were all parties to the Treaty of Brussels of 1948, which included a mutual defense guarantee (Belgium et al., 1948). This highlights the degree to which the year chosen to use to instantiate the model can influence the results, as well as the information missing from the tau-b similarity measurement. Furthermore, as mentioned previously, several countries which would be pivotal to the Cold War, such as Vietnam, and West and East Germany, were not yet sovereign in 1948, and are excluded from the input data – making it impossible to generate any conflicts involving them.

## 4.3.3 Predicted Alliances

In addition to conflicts, the other detailed output the model produces is the position of each agent at each step. The space in which these positions exist is fairly abstract, particularly in this example. While in some cases positions can be assigned concrete meaning on a particular issue (e.g. number of years before a policy is adopted, as in Bueno de Mesquita and Stokman (1994)), in this case, the positions simply indicate alignment with the security interests of the United States or Soviet Union, as anchored in the starting year. The specific issues underlying these interests, however, obviously change substantially over the course of the time period we are concerned with, and assessing the position of each country across such issues is a substantial and ill-defined task in and of itself. Furthermore, both superpowers can and occasionally do change their own position as well, suggesting that they are compromising some interests in order to maximize their security with regard to the current state of the world.

One way around needing to interpret the position space is to examine not the specific positions themselves, but the distance between the positions held by two agents. Within the model, when agents' positions are closer, they are more likely to assist each other more strongly in conflicts. Furthermore, this fact is public – all other agents know it, and explicitly take it into account when estimating the probabilities of winning or losing potential conflicts. This behavior bears a clear resemblance to alliances, an enduring institution in international relations (both in practice and theory). Indeed, the Realist school of thought often holds that alliances and treaties are not meaningful in and of themselves, but only inasmuch as they reflect an underlying alignment of states' interests (Walt, 1987). With this in mind, I hypothesize that if a model is capturing aspects of real state behavior, the closer two agents are, the more likely the relevant countries are to have some sort of alliance or other formal defense or security agreement.

In order to test this hypothesis, I use the Correlates of War (Gibler, 2009; Gibler et al., 2013) alliance data. Recall that this dataset served to generate the model inputs as well; The variable I attempt to predict is alliances in effect in 1998 (corresponding approximately to end-year of the number of steps the model has run for), reducing the risk of circular prediction. Furthermore, for experiments 1 and 2, I am using the Original Input data, while the 1998 alliances are drawn from the most recent COW data; as noted in Section 4.2.1, there appears to be substantial enough difference between the older and updated datasets to further reduce the risk of circular prediction.

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Alliance=1				
Const	-6.828***	-8.060**	-7.387***	-8.484***
	(1.917)	(2.429)	(1.601)	(1.459)
Mean Distance	2.540	9.158	2.586	$50.793^{**}$
	(4.019)	(7.667)	(3.559)	(20.673)
Alliance=2				
Const	-4.902***	-5.003***	-5.660***	-5.211***
	(0.863)	(1.327)	(0.897)	(1.270)
Mean Distance	0.044	0.624	-1.554	-37.990
	(2.472)	(6.523)	(2.921)	(56.118)
Alliance=3				
Const	-4.043***	-4.332***	-4.073***	-3.915
	(0.544)	(0.763)	(0.417)	(0.607)
Mean Distance	0.893	2.843	-2.087	-29.748
	(1.393)	(3.386)	(1.439)	(25.521)
Alliance=4				
Const	-0.974***	-0.056	-0.080	-1.2184***
	(0.150)	(0.329)	(0.086)	(0.122)
Mean Distance	-3.286***	-9.905***	-7.792***	-5.716
	(0.660)	(2.231)	(0.529)	(4.214)

Table 4.8: Alliance Type – Multinomial Choice Logistic Regression

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

I compute the distance between the positions of each pair of agents at the end of each run, and then find the average of these distances across each experiment. I then merge the resulting data from each experiment with the dyadic alliance data, padded to include 0 values for each pair of states for which no agreement is recorded. With these merged datasets, I run two models: a multinomial logistic regression (McFadden, 1984), and an exponential random graph model (Robins et al., 2007).

The logistic regression treats each dyad as independent, and attempts to predict which category the relationship between the two states falls into. Across all four experiment results, shown in Table 4.8 the mean distance between agents is not a significant predictor of ententes, non-aggression or neutrality agreements (alliance categories 1 - 3, as detailed

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
2-star	$-0.188^{***}$	$-0.171^{***}$	$-0.023^{***}$	$-0.031^{***}$
	(0.017)	(0.017)	(0.003)	(0.003)
Mean Distance	$-2.113^{***}$	$-5.613^{***}$	$-6.308^{***}$	$-25.981^{***}$
	(0.607)	(1.245)	(0.615)	(4.589)

Table 4.9: Alliance Network – Exponential Random Graph Models

Standard errors in parentheses.

\* p < .1, \*\* p < .05, \*\*\* p < .01

in Table 4.2). However, in Experiments 1-3, for defense pacts (category 4, the strongest type of relationship), the model mean distance is statistically significant, with a negative sign. This means that the smaller the distance between the agents (which is to say, the closer they are), the more likely an alliance becomes. This is not the case with Experiment 4. Though the coefficient is still negative, it is not significant.

A key assumption of a logistic regression is that the observations are independent. However, we know that this is not the case with regard to interstate relationships in general, and alliances in particular. Rather, the alliances form a network, and should be studied as such (Maoz, 2010). In order to test whether the model distances are a useful predictor of the alliance network, I use exponential random graph models (ERGMs), a methodology developed specifically to account for such networked interdependence. The independent variable here is the simplified alliance network, where the nodes are states, connected by an edge wherever a defense pact (*Alliance* = 4) relationship exists. The key independent variable is the generated Mean Distance values, as with the logistic regression. Following Cranmer, Desmarais, and Menninga (2012), I include a 2-star coefficient (the number of pairs of edges, or trios of connected nodes) in order to capture the structural effect whereby states with more allies tend to attract additional allies (which Cranmer, Desmarais, and Menninga (2012) term the 'popularity effect').

The results of this model are shown in Table 4.9. The key finding here is that the Mean

Distance between actors is still a significant, and negative, predictor of an alliance between them, even when taking the network structure into account. Note that unlike with the logistic regression, here Experiment 4 Mean Distances are significant as well.

# 4.4 Discussion

This chapter set out to accomplish several goals. The first was to attempt to reproduce the results of Bueno de Mesquita (1998), in order to test my reimplementation of the Expected Utility Model's power to explain and predict the long-term course of complex international competition. It also extended the analysis presented in the original paper by collecting additional model outputs and quantitatively testing their correspondence with empirical data. Additionally, this chapter continued to comparatively test the EUM variants presented in Chapter 2, by implementing the Baseline model against the variant with the highest explanatory power on the ICEWS dataset; if the theories embedded in the model variant are in fact more explanatory of the world than the Baseline's, we would expect that variant to have more explanatory power here as well. If it does not, that would indicate either that the results of the previous chapter are due to chance, or that the dynamics driving international relations during the Cold War are different than those driving them in the 2004-06 timeframe that we previously examined.

I was unable to reproduce the model inputs from contemporary data. This highlights two issues. One is the quality of the input data: the alliance network which I used is an updated version of the one used to generate the Original Input data, and as such is intended to be more accurate (Gibler et al., 2013). In order to use the methodology presented here for forecasting, we must explicitly be aware that input data sources may be incomplete or contain errors, and if possible account for that possibility explicitly (e.g. by performing sensitivity analyses on the input data). The other issue is the need for explicit clarification of the pre-processing steps needed to get from the raw input data (in this case, the alliance network) to the actual model input (in particular, the agent positions). As shown in Section 4.2.1, the method described for generating agent positions in fact produces a two-dimensional position space; the paper does not specify the heuristic or projection used to reduce this space to one dimension.

Despite being unable to reproduce the Original Inputs from raw data, I was able to instantiate my model runs using the input data provided in the original paper. The overall distribution of outcomes produced by my reproduction of the original model is qualitatively similar to the output described in the original paper, though not identical. The original paper appears to generate a wider range of outcomes, and in particular more USSR-leaning ones than the reproduction model. Furthermore, the individual run traces do not appear to be a good match to the examples presented in the original paper. While the stochastic nature of the model makes it essentially impossible to reproduce specific model traces without access to the original source code and random seed, the example traces reported in the original appear to show more movement of individual agent positions than characterize my reproduction, while the median position does not have the saw-tooth pattern which characterizes all the experiments, but in particular Experiments 1 and 3. Experiment 3, using the Updated Input data, generates similar behavior by the median and individual agent positions to Experiment 1. This suggests that the lack of major, stable shifts in the median position is a feature of this particular model variant, rather than the particular input configuration of agents, positions, and capabilities. Recall that as demonstrated in the deep dive in the previous chapter (Section 3.3.2), cascades of agents backing out of offers in response to other agents accepting their own offers leads to an overall reduction in agent movement.

As I discussed above and in the previous chapter, a close examination of my reproduction model suggests that the stability of the median across its runs is driven by the rule allowing agents to renege on an accepted offer if another agent has accepted their own offer. This rule is part of the offer selection sub-model, which is one of the least-specified sub-models not only in Bueno de Mesquita (1998) itself but in the other descriptions of the model as well. There are two possibilities: one is that, despite my best efforts, I have not implemented this model behavior to match the description. The second is that the descriptions of this behavior in the prior material do not accurately reflect the behavior of original the underlying source code. As the original code is not available to inspect, it is impossible to verify directly which possibility is correct.

Nevertheless, there is at least anecdotal evidence that the flaw is not in my replication but in the original description. Private communications with scholars familiar with the original paper have suggested that the experiments it describes were implemented using the commercial version of the model software, developed by Bueno de Mesquita's consulting firm, Decision Insights, Incorporated (DII). This software was the subject of litigation, during which Gary Slack, a DII employee who is thanked for "invaluable programming assistance" in Bueno de Mesquita (1998), testified that the model contains proprietary methods not disclosed in the public literature (US 4th Circuit, 2011). This suggests, in turn, that the model description in the original paper, and the prior literature more broadly, is incorrect or incomplete. In this case, the difference between the described and observed model appears to be attributable to a specific sub-model – how agents select offers and change position. I believe that sub-model is indeed a key place where the public and proprietary versions of the model diverge. In fact, the traces in Experiment 2, where agents cannot back out of selected offers, appear to be qualitatively more similar to the traces reported in the original paper than Experiment 1.

While replication is certainly important, a more interesting question is how well the models reproduce the observed history of the Cold War and its end. Historically, while the end of the Cold War was sudden and unexpected, it was not a short-lived phenomenon, at least on the time-scale we are modeling here. In fact, it was followed by a decade or more of 'unipolarity' during which many countries formerly oriented towards the USSR, and often allied with it, rapidly shifted their international position towards the United States and its allies (Wohlforth, 1999). The most notable demonstration of the long-term transition in the international geopolitical orientation is was the countries who were formerly members of the Soviet-led Warsaw Pact alliance joining NATO, the US-led alliance which had been the Warsaw Pact's primary adversary (Waltz, 2000). While there is only one run of history for us to examine, it is worth noting that Cold War-era fiction envisioning Soviet ascendancy (e.g. Milius and Reynolds, 1984) described a similar process happening in reverse, with NATO countries being rapidly pulled into the USSR's orbit. Examination of the example model traces shown in Figure 4.8, and other traces not shown here, indicate that the replication model of Experiments 1 and 3 does not tend to produce such sharp, stable transitions. In contrast, Experiments 2 and 4 do feature relatively sharp changes, with Soviet-aligned agents moving towards the United States followed by the Soviet Union itself. Qualitatively, at least, these models appear to be generating behaviors which mirror the historic record. Note, however, that the majority of model runs across all experiments do not follow this behavior, and indeed do not generate a 'victory' to one superpower or the other.

None of the experiments presented here appear to be be powerful predictors of observed conflicts. While the coefficients are statistically significant, they are nevertheless only weak predictors of observed militarized disputes in the relevant historic period. There are several possible explanations here. One is that model conflicts do not play the same role as do MIDs in the international system, at least at the scale and with the issues being examined here. The model's single dimension means it is most likely to generate conflicts driven by the overall 'issue' – in this case, the competition between the two superpowers. However, many conflicts arise due to issues particular to a state, dyad or region (Senese, 1996). If these issues are not embedded in the initial position-generating alliance network, the model will obviously be unable to account for them, and thus will not generate conflicts driven by them. Furthermore, the model lacks a spatial component. Many conflicts are between neighbors, and may be driven at least in part by disputes arising due to that relationship – with territory being the most obvious example, though not the only one. The ability of other states to assist one side or another in a conflict is also influenced by geography: in terms of sending military forces or assistance, but also in terms of 'softer' power. For example, nearby states trade more than far-away ones (Ramanarayanan, 2011), meaning that overall, a state's ability to affect another by increasing or decreasing trade and commerce is also tied to the distance between them. Finally, by only examining militarized disputes, we may be missing other forms that conflicts take, and that states use to coerce one another, such as via diplomatic or economic means.

Similarly, the models are significant but not powerful predictors of alliances. They appear most capable of predicting defensive alliances, the strongest category: the closer the final positions of two agents are, the more likely they are to have a defensive pact between them, committing one to come to the other's aid if attacked. Of the alliance types coded in the COW data, defensive pacts most closely mirror the behavior coded in the model itself, with agents more likely to contribute more resources to a closer agent when that agent is involved in a conflict. This correspondence suggest that the model is indeed capturing one dimension of states' decisionmaking. Importantly, the dyad-wise position distances are a stronger predictor in the Experiments 2 and 4, where the agents tend to change positions more than in Experiments 1 and 3. This indicates that the dynamics of the models are generating useful information not directly present in the initial data.

We have already touched on comparing the experiments to one another in the various sections above. A common thread is that the Updated model outperforms the Baseline model: The agent behaviors it generates are more qualitatively plausible, and the conflicts it generates are more predictive of real conflicts across both input sets. Note that the one exception is the prediction of alliances in Experiment 4. Though this experiment generates the best predictions of conflicts, it also produces the worst predictions of alliances. It is somewhat surprising that these two metrics do not go together, since we would hope that the most accurate model of the system would generate strong predictions for both. One possible explanation is suggested by the fact that Experiment 4 generates the fewest outcomes with one side or another gaining a decisive advantage. Nearly all of its outcomes involve the Cold War continuing indefinitely, and (as Figure 4.6 shows) becoming more stable. In contrast, the target alliance network is from 1998, reflecting a decisively post-Cold War world. Thus, the experiment may be predicting conflicts likely to occur in a Cold War context, which accounts for the bulk of the time (and the conflicts) being studied; however, since it does not predict an end to the Cold War, it fails to predict the alliance network which follows. This analysis provides more evidence for the value of comparatively applying multiple models. The better fit of the Updated model is consistent with the results of Chapter 3, giving us confidence that, across different input datasets capturing different historic periods at different scales, the assumptions of the Updated EUM variant – particularly agents not changing position due to conflicts, and not reneging on accepted offers – provide better models of actual state behaviors. Furthermore, this chapter demonstrates another form of comparative modeling: the use of different input datasets which all attempt to capture the same underlying state of the world. In this case, these are the two input sets, derived from different versions of the Correlates of War capabilities and alliance data. Using these, we can see the qualitatively different behavior in particular between Experiments 2 and 4, where the same Updated model is instantiated with the two different input datasets. There does not appear to be a simple, reduced explanation for this difference. Rather, the particular configurations of the two inputs yield substantially different distributions of outcomes, highlighting the uncertainty present in the model itself.

Finally, let us pull back to examine the original question which motivated Bueno de Mesquita (1998): whether the end of the Cold War could be explained by contemporary theories of international relations, as operationalized by the Expected Utility Model. The answer here appears to be a qualified 'yes'. Despite the differences in the behavior and dynamics of the different model variants and input data, all experiments generate an overall prediction of an advantage to the United States, despite the initial balance of power being nearly completely even. This suggests that the initial hypothesis is correct, and that the end of the Cold War was potentially predictable from much earlier data. In fact, this analysis provides more evidence for the robustness of the US advantage, as it recurs across several model variants, indicating it is not simply an artifact of one particular model implementation. Nevertheless, the end of the Cold War is far from inevitable, even within these models; if these models are capturing some of the true underlying dynamics, the results suggest that the end of the Cold War was surprising not solely because it was emergent, but because it was in fact an unlikely event. Furthermore, these results show the caution we must exercise in using such model for prediction. While the input data may have been knowable in the starting year, we necessarily used knowledge from future years to assess the model outputs, and determine that the conflicts were non-predictive while the positions were. More research is required to assess whether these are general features of the model, or outcomes specific to this particular case or system scale.

# Chapter 5: Conclusions

Agent-based models have not been in the mainstream of international relations scholarship. Despite this, I believe that ABMs are uniquely suited to capture and model the complex interactions characterizing the international system, which I argue are often driven as much by uncertainty and satisficing heuristics as by formal rationality. In this dissertation, I set out to accomplish two overall goals: to adapt prior, well-established formal models of international political conflict into agent-based models, and to demonstrate how such models can be used comparatively to study the underlying theories they are based on. In doing so, I have incorporated elements from across the methodologies used in the study of international conflicts, including game theory, psychological and organizational decisionmaking theories, econometric models and machine-coded event data.

I did not initially set out to address the work of Bruce Bueno de Mesquita in particular; however, he is the key contributor to the two models this dissertation has concentrated on. The first is the International Interaction Game, presented in Bueno de Mesquita and Lalman (1992): a formal, extensive-form model of how pairs of states choose whether to escalate their use of force against one another in a potential crisis situation. I chose this model since it is one of the few game-theoretic models that has been operationalized to make specific predictions about observed interactions, and as such is frequently cited as the best-tested application of the purely rational decisionmaking model of state decisionmaking (e.g. Allison and Zelikow, 1999; Green and Shapiro, 1996). The second model is Bueno de Mesquita's Expected Utility model, a model of multilateral negotiations and conflict described across numerous papers (Bueno de Mesquita, 1984, 1994, 1997, 2002, 2011) and popular presentations (Bueno de Mesquita, 2009, 2010). This model is of interest not just because of the impressive accuracy attributed to it in the prediction of real-world cases (Bueno de Mesquita, 2009; Feder, 1992), but also because it is another relatively rare example of a formal, strategic model applied to forecasting and historic explanation.

Both of these models were, in practice, heavily computational already. The IIG model was tested most thoroughly via the EUGene software (Bennett and Stam, 2000a,b), which not only performed the data manipulation required to combine several disparate datasets into a single million-record dataset, but applied a genetic algorithm to estimate each state's risk tolerance in each year by comparing its observed alliance portfolio to the most- and least-secure sets of alliances it could possibly hold. The Expected Utility model (EUM) was also implemented computationally – in fact, the challenges of addressing it was that many of its elements are not well-articulated in the published descriptions.

What drew me to these models, however, was that both consist of unitary actors directly interacting with one another, making discrete decisions in sequence until a conclusion is reached. In this way, they both already shared a core structure with agent-based models<sup>1</sup>. Furthermore, the actors are not abstract 'players' but represent specific real-world entities, with their choices, and the overall model outcomes, corresponding to events which may be observed in historic data. This offers the opportunity to not only reimplement these models as ABMs, but to directly compare their outputs against both the outputs of the original models and the empirical outcomes of the cases and interactions being modeled.

Both models can be separated into several sub-models, governing when and how the agents will interact with one another, what choices they face when they do, and how they decide between those choices. It is the latter component I was particularly interested in. Both models originally implement particular theories of state decisionmaking: perfect rationality in the case of the IIG, and a heuristic-based rationality in the case of the EUM. Implementing different decisionmaking models allows me to directly compare the explanatory power of their underlying theories, holding the rest of the model structure fixed.

This comparative approach draws on a broad trend of comparative modeling in the study of international relations. Direct comparison of statistical model results is, of course,

<sup>&</sup>lt;sup>1</sup>In fact, as I discuss in more detail in Chapter 3, the Expected Utility model may be best considered as an ABM already, despite being generally described as a game-theoretic model.

exceedingly common. To name a few other examples I have drawn on directly, Allison and Zelikow (1999) apply several qualitative theories of state decisionmaking to the case of the Cuban Missile Crisis, while Kaufmann (1994) compares the predictions of the rational model to psychological theories of individual decisionmaking in an analysis of German policy in the 1905-06 Morocco Crisis. The Axelrod (1980) Prisoner's Dilemma tournament involved the comparison of multiple decisionmaking models controlling agents' play in repeated prisoner's dilemma games, and was motivated at least in part by the study of international conflict and cooperation. Most relevant, Bueno de Mesquita and Stokman (1994) present two models, an early version of the EUM (Bueno de Mesquita, 1994), and the Stokman and Van Oosten (1994) cooperation-focused Logrolling Model (which is described explicitly as a computational, object-oriented model), both of which implement different theories of collective decisionmaking. They instantiate these models with the same inputs, and compare their explanatory and predictive power with regard to European Community policymaking.

Chapter 2 described the International Interaction Game in detail: both the bilateral, extensive-form game itself and the methodology that Bueno de Mesquita and Lalman (1992) lay out, and Bennett and Stam (2000b) apply, using the Correlates of War alliance network data to estimate the payoffs associated with each terminal node of the tree for each dyad of countries. It also described a simplified variant of the game, adapted from Signorino (1999), which was used for testing and experimentation purposes. It presented an agentbased architecture built around these games, where the games themselves are a sub-model governing the interaction of pairs of agents, from a set of agents that persist across an entire model instantiation; the model instantiation object itself sequences these interactions – at random in the simplified model, and following the EUGene-generated historic dyads for the IIG variant.

Most importantly, I describe three models of agent decisionmaking that control how agents choose an action at each node of the tree. The first is simply full, formal rationality: agents have full access to the payoffs of each of the game tree's terminal nodes, and use backwards induction to find the best move at each node. When two such agents interact with each other, they will arrive at the equilibrium outcome. Next, I describe two learning models where the agents store and update weights on each possible move, which they use to make decisions stochastically. Both are built on a stochastic decision model, experience-weighted attraction, which has been shown to be both an effective machine learning technique and a model of how individuals and groups actually make decisions. The first learning model is the Reinforcement Learning (RL) model, where agents maintain a single set of move weights, which are updated based on the payoff each move leads to. This offers a stylized model of the Organizational Behavior view of decisionmaking (Allison and Zelikow, 1999), with the agents developing standard operating procedures, which they apply without considering whether they are well-suited to the specific interaction. The next learning model is a Case-Based Learning (CB) model, where the agents maintain a library of past interactions, and associated weights learned from each. This model captures the way March and Simon (1993) describe organizations as making decisions by retrieving rules associated with relevant experiences, and specifically the tendency of decisionmakers, and states more broadly, to address new crises by analogy to previous ones (Khong, 1992; Schuman and Rieger, 1992).

I first apply all three decisionmaking models to the simplified interaction model, populating it with sets of heterogeneous agents, whose properties were set stochastically. Despite starting with no initial weights, the agents are capable of learning to play in ways that result in equilibrium outcomes more often than we would expect strictly by chance. Over multiple instantiations, neither the RL nor the CB model show a robust advantage over the other in collective learning of the equilibrium outcomes. When the agent properties are static, the CB model robustly (though not universally) outperforms the RL model, since it allows agents to learn the best-response strategies for each other agent. However, when the agent properties (and hence the equilibria) vary from step to step, the correct moves vary as well, making case-based learning less effective, and the RL model becomes more likely to lead to equilibrium interactions.

Next, I apply these models to the International Interaction Game itself with the historic

dyads generated by the EUGene software, with one agent per state in the international system. Again, despite starting with no information, and despite the added complexity of the game tree, the agents collectively and robustly learn to play such that the equilibrium outcome is reached more often than randomly. Unlike with the simplified model, however, the frequency of equilibrium outcomes between instantiations is not distributed evenly: there are several apparent attractors, equilibrium frequencies which are converged to more often than by chance; furthermore, changes to the international system itself lead the equilibrium frequency to change sharply at similar points across the majority of instantiations. These change-points correspond to major shocks to the international system which are exogenous to this model, particularly the ends of the world wars. Finally, the outcomes of the RL and CB models have comparable power to the equilibrium outcomes in predicting the actual observed events between each country pair in each year. Some runs of both, in fact, converge to substantially better fits than the equilibria.

Case studies such as Allison and Zelikow (1999) and Kaufmann (1994) show that nonrational accounts of decisionmaking can explain the results of particular historic interactions as well or better than a purely rational account of decisionmaking; however, such individual case studies are difficult to conduct at scale, and Achen and Snidal (1989) and others have argued that they are insufficiently standardized and constrained to use for testing general theories. These results suggest that agent-based models can serve as a method of testing non-rational decisionmaking theories across a large number of cases, and anchoring such tests to well-established formal and statistical methodology. In fact, these results appear to confirm the conclusions of Allison and Zelikow (1999), Kaufmann (1994) and others, and indicate that non-rational models can offer similar explanatory power to the rational model, when applied to the same data. The learning models presented here produce the satisficing decisions that are a well-understood feature of organizational behavior (March and Simon, 1993) but are not present in the purely rational model. Furthermore, these models offer an endogenous explanation for apparent errors that is richer than simply adding an exogenous noise factor to agents' utilities (Signorino, 1999). Finally, the results I present here may help
offer one reason for the disagreements between advocates of rational and non-rational views of decisionmaking. The models suggest that lessons and procedures learned from experience will, over time, tend to drive to decisions and behaviors that lead to better outcomes, which will with some frequency be the equilibrium outcome. This means that the equilibrium can in fact be a useful predictor of the outcome we will observe, even if the actors are not reaching it in a procedurally-rational way. Furthermore, when the actors deviate from the equilibrium, these deviations are not solely random mistakes, and may be driven by the application of lessons learned from prior interactions.

The Expected Utility Model (EUM), which I explore in Chapter 3 and apply in Chapter 4, is intended to capture a different level of detail than the International Interaction Game. Where the form and payoffs of the IIG are specifically designed (as the name implies) with country-to-country interactions in mind, the EUM is intended to serve as a general-purpose model of formal or implicit negotiation. While the model is often presented as an direct application of game theory, its agents are not formally rational in the traditional sense. and the model itself is not studied analytically in parameterized form, but applied and solved computationally on a case-by-case basis. In this way, it already appears to be an agent-based model. Furthermore, for the most part, the EUM's sub-models are not direct operationalizations of specific qualitative theories. The model's overall assumptions are drawn from the Realist paradigm, but the model is itself a particular theory of state interactions. The main argument put forward for its correctness is its accuracy – its ability to make correct predictions as to the outcomes of particular negotiations or political contests. One of the main pieces of evidence for its accuracy is the Feder (1992) evaluation of the model's application by the CIA, claiming up to 90% accuracy across multiple cases, though providing minimal details on the methodology used to initialize the model runs or determine what constitutes a correct prediction. Several other papers (e.g. Bueno de Mesquita, 1984, 1997) and general-audience accounts (Bueno de Mesquita, 2010) provide single-issue forecasts or predictions of past events using data further in the past. More recently, Bueno de Mesquita (2011) conducted a multi-case evaluation of a modified variant of the model, using the Thomson et al. (2006) dataset of European Union issues and decisions, showing that it has better predictive power than both a simple non-strategic median voter model and the original EUM.

For my purposes, the most significant of the model's reported successes is its apparent prediction of the course and conclusion of the Cold War between the United States and the Soviet Union (Bueno de Mesquita, 1998), which I discuss in detail in Chapter 4. The importance of this claim is twofold: most obviously, it would show that the EUM can capture not just discrete issues but states' overall alignment in the international system, and how that alignment changes over a timeframe of multiple decades. It would also, in addition, offer evidence that the unexpected end of the Cold War was an emergent phenomenon arising from the complex interplay of interactions which are themselves well-understood by existing theories, and not an argument against those theories as Gaddis (1992) and others have claimed. If that is the case, it would strengthen the case that agent-based models are an effective tool for the study of international relations.

In Chapter 3, I present my reimplementation of the EUM, breaking it down into its different sub-models and highlighting its agent-based character using the Grimm et al. (2006) Overview, Design concepts, and Details (ODD) structure. I attempt to identify and discuss the underlying assumptions and hypotheses embedded in these sub-models, and propose several alternative sub-model implementations. In particular, I highlight one core assumption of the original EUM: that when two actors engage in a direct conflict, the loser will adopt the winner's position. This assumption represents a substantially different view of coercion than described by Schelling (1966) and others.

In both Chapters 3 and 4, I attempt to reproduce results of the EUM for three cases where both the inputs and at least representative outputs are provided: in Chapter 3 the automobile standards from Bueno de Mesquita (1994) and the Chinese democratization model from Bueno de Mesquita (2002), and in Chapter 4 the Cold War model (Bueno de Mesquita, 1998). In none of these cases do I perfectly replicate the originally-published results. This is cause for some concern, as it suggests that I have failed to properly implement

the model as described, or that the original model itself was not implemented as described (with this 'or' not being exclusive). To paraphrase Washington (1796), in reviewing my code I am unconscious of intentional error, I am nevertheless too sensible of my defects not to think it probable that I may have included many errors. The difficulty of reproducing agent-based models from their descriptions is well known (Axtell et al., 1996; Edmonds and Hales, 2003; Wilensky and Rand, 2007), and all the more so when the original papers do not themselves describe the model in computational terms. Furthermore, I verify that the Scholz, Calbert, and Smith (2011) reimplementation reproduces the results of the automobile standards model output, but not the Chinese democratization one – and that the reproduction of the former relies on a probability calculation which appears to contradict not only the form of the equation provided in multiple papers, but its basic underlying interpretation. This, along with the evidence that the Cold War model utilized a proprietary implementation that includes features not described in the open literature (US 4th Circuit, 2011), means that we cannot necessarily take the original results as ground truth either. However, if the theory of interaction and decisionmaking articulated by the EUM provides an accurate model of reality then it should not be oversensitive to specific elements of its implementation, including those that differ between my implementation and the original. Indeed, it appears to be the case that there is no single 'original' model, and that the EUM has undergone multiple changes over time, yet is still described as the same basic model. In two out of the three cases I attempt to replicate, my implementation's results are similar but not identical to those presented in the original papers.

The main argument for the validity (and utility) of the original model is its predictive ability, and this must be the main test of my reimplementation as well. In cases such as the timeline for automobile standards, there is a single, quantifiable outcome which can be compared to the model's median position. In the case of broader international competition, however, there is no such outcome; instead, we must utilize the model's other outputs, and compare them to observed data. I apply several model variants to two test cases, at different time periods and scales. One is the Cold War case discussed above; I extend the analysis in Bueno de Mesquita (1998) to quantitatively test the models' ability to predict both militarized conflicts and mutual-defense alliances. The other is a new attempt to predict the volume of conflict events collected in the Integrated Crisis Early Warning System (ICEWS) dataset. To the best of my knowledge, this is the first application of an EUM-type model to such micro-level event data. In both cases, I instantiate the models using data from the Correlates of War project, providing input data which can be independently recreated without relying on expert elicitation.

I apply several model variants to each case, then instantiate and run each multiple times in order to produce ranges of outputs. I hypothesize that if a model is capturing some of the underlying decisions and interactions, the more frequently an output is produced, the more likely a corresponding event or relationship is to be observed in the data. In both cases, the model outputs have only weak predictive power, though they are better than random. This relationship indicates that the models are in fact capturing at least some of the drivers of state behavior in the international system, particularly since the model is projecting the numerous separate issues driving the system onto a single dimension.

Since the historic data is fixed, I can directly compare the goodness-of-fit measures of the outputs of the different model variants to one another. Across both test cases, the baseline model variant, that attempts to be the most direct reproduction of the original model, has substantially less predictive power than the variant which implements several alternative sub-models. This variant showed the most predictive power on the ICEWS data, on the strength of which I applied it to the Cold War case, where it again generated a better fit than the baseline variant. Importantly, in the Cold War case, though the baseline variant produces a distribution of median-position outcomes which bears resemblance to the distribution shown in Bueno de Mesquita (1998), the actual distribution is driven by a model artifact, one-step shifts in the median position, rather than sustained transitions. The updated model, in contrast, generates plausible dynamics as well, with punctuated, stable transitions by groups of agents in the same direction.

By breaking the EUM into its constituent sub-models, I was able to explicitly test the

assumptions each sub-model is based on. In particular, I was able to test the hypothesis that conflicts will lead the losing agent to change its position and adopt that of the winner. This stood out as questionable while I was developing my model reimplementation, not only because it represents a different view of coercion than is generally found in the literature (e.g. Bratton, 2005; Schelling, 1966), but because it violated my conceptualization of the model architecture, and how it relates to reality: specifically, the strict separation between agents' external interactions and their internal decisionmaking, with agents only able to observe the latter. To test this hypothesis, I implemented a model variant where conflict does not lead the losing agent to change position. Across both the ICEWS and Cold War cases, model variants incorporating that new sub-model showed increased predictive power; this is particularly visible with the ICEWS experiments in Chapter 3, wherein this sub-model increased the power of the overall model across several variants with different combinations of other sub-models.

Taken together, these results give reason for both caution and optimism with regard to applying the EUM to international relations. Despite being initially developed for negotiations over specific issues, the model is capable of generating outputs with a degree of non-random correspondence to observed events. This appears to be true across both the high-resolution ICEWS data and the lower-resolution MID data. Furthermore, the Cold War case provides information useful for predicting the presence of alliances as well as conflicts, evidence that the model is capturing multiple dimensions of the actual dynamics of the international system. However, the predictive power of these output data is nevertheless weak, as indeed we might expect from a relatively simple model that abstracts away most of the issues driving cooperation and conflict.

These results strongly suggest that none of the model variants are likely to achieve the 90% claimed issue-specific accuracy (Bueno de Mesquita, 2002; Feder, 1992) in predicting the broad course of the international system. At present, these models are less accurate than the purely statistical models used to predict event data (e.g. Brandt, Freeman, and

Schrodt, 2011; Ward et al., 2013b). Like those models, the EUM variants present an out-ofsample prediction, as they are instantiated with data strictly preceding the time period to which the output is compared. Unlike those models, however, the input data is taken from a strictly different dataset. Furthermore, while the statistical methodologies are generally applied to one specific conflict at a time, I have applied the EUM model variants to the entire core of the international system, leading to an inevitable loss of resolution, and hence of accuracy. A more significant difference is that the EUM models attempt to explicitly capture the decisionmaking and actions of individual actors, providing not just an outcome prediction but a comprehensible account of how the outcome was reached.

Indeed, this is true of the International Interaction Game models as well; no variant of either model presently is capable of fitting data better than the best statistical models. Yet, though prediction is an important test of the models and their theories, it is not their only objective in and of itself. Both classes of models I have presented here serve as tools for explicitly articulating, and testing, theories and hypotheses as to how states (and organizational actors more broadly) make decisions and interact with one another, either in crisis or merely competitive situations. Instantiating and running the models using notional, randomly-generated data can serve a similar purpose to analytical, closed-form examination of a simpler model, allows us to identify categories of model outputs, examine its general behaviors and properties across multiple specific instantiations, and estimate the uncertainty it captures or excludes. Applying empirical input data and running multiple instantiations of the models generates distributions of events and outcomes. If one model variant produces outputs which correspond better to the observed events, this would provide evidence that the assumptions underlying that variant are a better description of the real processes; in the case of the IIG models, where no variant produces obviously better predictions, this is evidence that the variants themselves remain approximately equally plausible. However, Bennett and Stam (2000a) adds numerous additional control variables to the logistic regressions which are not present here, such as regime type and years of ongoing peace. Adding such controls to the regressions of model outputs against historic events may help separate their explanatory power and highlight whether one provides more useful information than the others. More broadly, incorporating the model outputs within more advanced statistical methodologies may allow us to identify which model is providing useful information not captured by other data sources and methods, as well as improving the forecasting power of the overall statistical models. Finally, it may be the case that different model variants are capturing different aspects of the underlying real-world decisionmaking and dynamics, and that the most predictive power is gained by an ensemble of multiple model variants applied to the same input data.

In this dissertation, I have implemented only a small number of the overall model variants one can imagine, and have by no means exhausted the universe of plausible models of decisionmaking. More variants provide a natural way of extending this research. One natural extension is a unification of the two model frameworks I have described here; Bueno de Mesquita (2011) has already proposed the use of the IIG game tree in an EUM variant in order to better model scenarios (such as intra-European Union negotiations) where conflict is unlikely. More interesting, from my perspective, would be an integration of the types of decisionmaking models applied in both. The learning models used in the IIG, though well-grounded, do not lend themselves to straightforward interpretation, and it is not clear how to align the weighted stochastic choice to detailed qualitative and historical sources on the specific decisionmaking processes observed in particular conflicts. The type of heuristics used in the EUM variants would allow the testing of the rule-based decisionmaking that has been explored by expert systems (Taber, 1992) and simpler (Cederman, 1997) or narrower (Hudson, Schrodt, and Whitmer, 2004) agent-based models. The model architecture would allow different agents to utilize different sets of rules, and even for such rules to be developed via evolutionary computation (Axelrod, 2006). Meanwhile, the application of more rigorous and externally-grounded decisionmaking models to the EUM variants may help move it beyond a collection of effective heuristics. While the model's equilibrium may not be analytically tractable, learning agents may be able to approach it, or identify when an equilibrium does not exist (Galla and Farmer, 2013). Recursive n-order rationality has proven to be an effective way of modeling strategic behavior, and is particularly applicable to modeling organizations (Latek, Axtell, and Kaminski, 2009); implementing it within the EUMs would provide a formal way of introducing bounded but non-heuristic rationality. There are additional artificial intelligence methods which can give the agents the ability to look multiple steps ahead and anticipate agents' responses to different courses of action. For example, Monte Carlo Tree Search (Chaslot et al., 2008) and Markov decision process models more broadly, have proved to be powerful decisionmaking models in multi-agent simulations – including in commercial computer games simulating international conflict and diplomacy (Champandard, 2014)<sup>2</sup>.

I chose to study the IIG and EUM in part because both are focused on modeling actors' strategic decisionmaking, to the exclusion of many other factors driving the international system. In particular, both models require that agents' capabilities be set exogenously: at instantiation in the EUM, and based on annual historic National Material Capabilities data in the IIG. In both models, the outcome of individual interactions and conflicts does not affect the agents' capabilities going forward. The plausibility of this assumption will vary from case to case, and conflict to conflict. While the use of non-military and 'soft' power is unlikely to explicitly diminish it (Nye, 1990), the same is not necessarily true for military force, where factors as mundane as ammunition stocks can change the force an actor can plausibly bring to bear. While most of the EUM variants I describe limit agents to at most two conflicts per step (one which they initiate, one initiated against them), the IIG incorporates no such limit, and allows states (and implicitly their allies, who contribute capability as well) to apply their full power to any number of conflicts simultaneously. Yet when a state mobilizes its power (particularly military, though arguably other forms as well) to credibly threaten or deter another state, this is highly likely to reduce the power it can credibly bring to bear simultaneously elsewhere. Even the United States, at the height of its 'unipolar' power in the post-Cold War world, required an explicit doctrine and careful

 $<sup>^{2}</sup>$ In fact, the model architecture I utilize for the EUM allows agents to make internal copies of the model state, substituting their own decisionmaking model for those of the other agents, in order to facilitate look-ahead models such as Monte Carlo Tree Search or recursive rationality. However, implementing these models was outside the scope of this dissertation.

planning in order to credibly claim to be capable of fighting two wars simultaneously (Gates et al., 2006). Many other models of international systems (e.g. Axelrod, 1997; Cederman, 1997; Min, 2002; Taylor et al., 2008) endogenize agents' power, and in particular have it diminish when it is used in conflict. While the EUM endogenizes agents' position changes (indeed, this is the main purpose of the model), the IIG does not. While incorporating additional dynamics would add more moving parts to the models which would need to be justified, verified and validated, they would also allow the models to produce richer notional histories with each run, generating more outputs for validation or prediction. More importantly, they would add additional constraints and tradeoffs for the decisionmaking models to address, potentially leading to richer models and capturing the effects of longterm planning which may be occurring in reality but is not incorporated into the models at the moment.

This dissertation has also examined variants of only two models, out of many existing ones, and more which may be proposed and developed. The same comparative modeling methodology should be straightforward to apply to additional models as well. For example, Tsebelis (1988) and Metternich et al. (2013) describe and analyzes formal models of multilateral coalition formation within single countries (France and Thailand, respectively) and show that the model equilibria have predictive power for real events. It should be possible to agentize these models as well, endow the agents with different decisionmaking models, and analyze the resulting outputs as I have done here. Furthermore, it is not necessary to begin with a formal model. Agent-based models such as Cederman (1997) are natural subjects for this methodology as well, providing opportunities to implement both alternative heuristics and learning-driven decisionmaking models.

While there are many opportunities for future work, I have accomplished the main goal I set out to achieve. In this thesis, I have restated two well-established models used in the study of international relations, and reimplemented them as agent-based models. In doing so, I highlighted the implicit model of state decisionmaking the models had previously incorporated, and proposed and implemented alternatives. By directly comparing the models

to one another, to the originals, and to empirical data, I was able to test them, identify where they converge and differ, and test their underlying theories. This research demonstrates that ABMs can serve as a bridge between methodologies in international relations, building on formal models to incorporate behavioral insights derived from case studies and comparatively testing them using empirical data. While I have applied this methodology to international relations in particular, it can be applied across other areas of social science as well. ABMs may serve as a way of building upon existing methodological tools, and bridging between them: operationalizing qualitative theories and intuitions, articulating underlying assumptions and proposing alternatives, and then comparatively testing them against empirical data in order to assess their external validity. In fact, it may turn out that in many fields, as I argue is the case in international relations, models which look very much like ABMs are already widely accepted, and are waiting to be agentized.

## Bibliography

- Abdi, Hervé (2007). "The Kendall Rank Correlation Coefficient". In: Encyclopedia of Measurement and Statistics. Ed. by Neil J. Salkind. Sage, pp. 508–510.
- Abdollahian, Mark et al. (2006). Senturion: A Predictive Political Simulation Model. Center for Technology and National Security Policy, National Defense University.
- Achen, Christopher H. and Duncan Snidal (1989). "Rational Deterrence Theory and Comparative Case Studies". In: World Politics 41.02, pp. 143–169.
- Acheson-Brown, Daniel G. (2001). "The Tanzanian Invasion of Uganda: A Just War?" In: International Third World Studies Journal and Review 12, pp. 1–11.
- Akramizadeh, Ali, M.B. Menhaj, and Ahmad Afshar (2009). "Multiagent Reinforcement Learning in Extensive Form Games with Complete Information". In: Adaptive Dynamic Programming and Reinforcement Learning, 2009. ADPRL'09. IEEE Symposium on. IEEE, pp. 205–211.
- Allison, Graham (1971). Essence of Decision: Explaining the Cuban Missile Crisis. Little, Brown.
- Allison, Graham and Philip Zelikow (1999). Essence of Decision: Explaining the Cuban Missile Crisis. 2nd edition. Pearson.
- Axelrod, Robert (1980). "Effective Choice in the Prisoner's Dilemma". In: Journal of Conflict Resolution 24.1, pp. 3–25.
- (1997). The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton University Press.
- (2006). The Evolution of Cooperation: Revised Edition. Revised edition. Basic Books.

- Axtell, Robert et al. (1996). "Aligning Simulation Models: A Case Study and Results". In: Computational & Mathematical Organization Theory 1.2, pp. 123–141.
- Bannister, Michael J., William E. Devanny, and David Eppstein (2014). *ERGMs are Hard*. arXiv preprint arXiv:1412.1787.
- Belgium et al. (1948). Treaty of Economic, Social and Cultural Collaboration and Collective Self-Defence.
- Bennett, D. Scott and Allan C. Stam (2000a). "A Universal Test of an Expected Utility Theory of War". In: International Studies Quarterly 44.3, pp. 451–480.
- (2000b). "EUGene: A Conceptual Manual". In: International Interactions 26.2, pp. 179–204.
- Black, Duncan (1948). "On the Rationale of Group Decision-Making". In: The Journal of Political Economy, pp. 23–34.
- Boschee, Elizabeth et al. (2015). ICEWS Coded Event Data. Harvard Dataverse.
- Brams, Steven J. (2003). Negotiation Games: Applying Game Theory to Bargaining and Arbitration. Vol. 2. Psychology Press.
- Brandt, Patrick T., John R. Freeman, and Philip A. Schrodt (2011). "Real Time, Time Series Forecasting of Inter-and Intra-State Political Conflict". In: *Conflict Management* and Peace Science 28.1, pp. 41–64.
- Bratton, Patrick C. (2005). "When Is Coercion Successful? And Why Can't We Agree on It?" In: Naval War College Review 58.3, pp. 99–120.
- Breiman, Leo (2001). "Random Forests". In: Machine Learning 45.1, pp. 5–32.
- Bremer, Stuart A., ed. (1987). The Globus Model: Computer Simulation of Worldwide Political and Economic Developments. Campus Verlag.
- (1992). "Dangerous Dyads Conditions Affecting the Likelihood of Interstate War, 1816-1965". In: Journal of Conflict Resolution 36.2, pp. 309–341.
- Bueno de Mesquita, Bruce (1975). "Measuring Systemic Polarity". In: The Journal of Conflict Resolution 19.2, pp. 187–216.

- Bueno de Mesquita, Bruce (1984). "Forecasting Policy Decisions: An Expected Utility Approach to Post-Khomeini Iran". In: *PS: Political Science & Politics* 17.02, pp. 226–236.
- (1985). "The War Trap Revisited: A Revised Expected Utility Model". In: The American Political Science Review 79.1, pp. 156–177.
- (1994). "Political Forecasting: An Expected Utility Method". In: European Community Decision Making: Models, Applications and Comparisons. Ed. by Frans N. Stokman and Bruce Bueno de Mesquita. Yale University Press, pp. 71–104.
- (1997). "A Decision Making Model: Its Structure and Form". In: International Interactions 23.3-4, pp. 235–266.
- (1998). "The End of the Cold War Predicting an Emergent Property". In: Journal of Conflict Resolution 42.2, pp. 131–155.
- (2002). Predicting Politics. 1st edition. The Ohio State University Press.
- (2009). A Prediction for the Future of Iran. https://www.ted.com/talks/bruce\_ bueno\_de\_mesquita\_predicts\_iran\_s\_future.
- (2010). The Predictioneer's Game: Using the Logic of Brazen Self-Interest to See and Shape the Future. Random House.
- (2011). "A New Model for Predicting Policy Choices Preliminary Tests". In: Conflict Management and Peace Science 28.1, pp. 65–87.
- Bueno de Mesquita, Bruce and David Lalman (1992). War and Reason: Domestic and International Imperatives. Cambridge Univ Press.
- Bueno de Mesquita, Bruce and Frans N. Stokman (1994). European Community Decision Making: Models, Applications, and Comparisons. Yale University Press.
- Camerer, Colin F. (2003). Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2002). "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games". In: Journal of Economic Theory 104.1, pp. 137–188.

- Cederman, Lars-Erik (1997). Emergent Actors in World Politics. Princeton University Press.
- (2001). "Back to Kant: Reinterpreting the Democratic Peace as a Macrohistorical Learning Process". In: American Political Science Review null.01, pp. 15–31.
- (2003). "Modeling the Size of Wars: From Billiard Balls to Sandpiles". In: American Political Science Review 97.01, pp. 135–150.
- Champandard, Alex J. (2014). Monte-Carlo Tree Search in Total War: Rome II's Campaign AI. AIGameDev.com.
- Chaslot, Guillaume et al. (2008). "Monte-Carlo Tree Search: A New Framework for Game AI." In: Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference.
- Chaturvedi, Alok R. et al. (2000). "Agent-Based Simulation Approach to Information Warfare in the SEAS Environment". In: System Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on. IEEE, 10-pp.
- Claus, Caroline and Craig Boutilier (1998). "The Dynamics of Reinforcement Learning in Cooperative Multiagent Systems". In: AAAI/IAAI, pp. 746–752.
- Colgan, Jeff D. (2016). "Where is International Relations Going? Evidence from Graduate Training". In: International Studies Quarterly.
- Comer, Kenneth W. (2014). "Who Goes First? An Examination of the Impact of Activation on Outcome Behavior in Agent-based Models". PhD thesis. George Mason University.

Correlates of War Project (2010). National Material Capabilities Dataset, v4.0.

- Coughlin, Peter and Shmuel Nitzan (1981). "Electoral Outcomes with Probabilistic Voting and Nash Social Welfare Maxima". In: *Journal of Public Economics* 15.1, pp. 113–121.
- Cranmer, Skyler J., Bruce A. Desmarais, and Elizabeth J. Menninga (2012). "Complex Dependencies in the Alliance Network". In: *Conflict Management and Peace Science* 29.3, pp. 279–313.
- d'Agostino, Ralph B. (1971). "An Omnibus Test of Normality for Moderate and Large Size Samples". In: *Biometrika* 58.2, pp. 341–348.

- de Mantaras, Ramon Lopez (2001). "Case-Based Reasoning". In: Machine Learning and Its Applications. Ed. by Georgios Paliouras, Vangelis Karkaletsis, and Constantine D. Spyropoulos. 2049. Springer, pp. 127–145.
- Dobson, Alan P. and Steve Marsh (2007). US Foreign Policy Since 1945. Routledge.
- Dunne, Paul, Eftychia Nikolaidou, and Ron Smith (1999). "Arms Race Models and Econometric Applications". In: A Conference on The Arms Trade, Security and Conflict. Middlesex University Business School, London.
- Edmonds, Bruce and David Hales (2003). "Replication, Replication and Replication: Some Hard Lessons from Model Alignment". In: Journal of Artificial Societies and Social Simulation 6.4, online.
- Eliot, Thomas S. (1934). "The Hollow Men". In: Poems: 1909-1925. Faber And Faber Limited.
- Epstein, Joshua M. (2008). "Why Model?" In: Journal of Artificial Societies and Social Simulation 11.4, online.
- Epstein, Joshua M. and Robert L. Axtell (1996). Growing Artificial Societies: Social Science from the Bottom Up. Brookings Institution Press.

Exum, Andrew (2008). "1776 and All That". In: The Guardian.

- Fearon, James D. (1995). "Rationalist Explanations for War". In: International organization 49, pp. 379–379.
- Feder, Stanley (1992). "Factions and Policon: New Ways to Analyze Politics". In: Inside CIA's Private World: Declassified Articles from the Agency's Internal Journal, pp. 274– 292.
- Förster, Stig and Jorg Nagler (2002). On the Road to Total War: the American Civil War and the German Wars of Unification, 1861-1871. Cambridge University Press.
- Friedman, Milton (1953). "The Methodology of Positive Economics". In: Essays in Positive Economics. University of Chicago Press.

- Fudenberg, Drew (2010). "Predictive Game Theory". In: American Economic Association, Ten Years and Beyond: Economists Answer NSF's Call for Long-Term Research Agendas.
- Gaddis, John Lewis (1992). "International Relations Theory and the End of the Cold War".In: International Security 17.3, p. 5.
- Galla, Tobias and J. Doyne Farmer (2013). "Complex Dynamics in Learning Complicated Games". In: Proceedings of the National Academy of Sciences 110.4, pp. 1232–1236.
- Gartzke, Erik (2007). "The Capitalist Peace". In: American Journal of Political Science 51.1, pp. 166–191.
- Gates, Robert et al. (2006). *Quadrennial Defense Review Report*. United States Department of Defense.
- Gates, Scott and Brian D. Humes (1997). Games, Information, and Politics: Applying Game Theoretic Models to Political Science. University of Michigan Press.
- George, Alexander L. and Andrew Bennett (2005). Case Studies and Theory Development in the Social Sciences. MIT Press.
- Gibler, Douglas M. (2009). International Military Alliances, 1648-2008. Correlates of war series. CQ Press.
- Gibler, Douglas M. and Meredith Reid Sarkees (2004). "Measuring Alliances: the Correlates of War Formal Interstate Alliance Dataset, 1816-2000". In: Journal of Peace Research 41.2, pp. 211–222.
- Gibler, Douglas M et al. (2013). Release Notes for Version 4: Correlates of War Formal Interstate Alliance Dataset, 1816-2012. Correlates of War Project.
- Gilbert, Nigel and Klaus Troitzsch (2005). Simulation for the Social Scientist. 2nd edition. Open University Press.
- Gilboa, Itzhak and David Schmeidler (1995). "Case-Based Decision Theory". In: The Quarterly Journal of Economics 110.3, pp. 605–639.
- Gill, Jeff and John R. Freeman (2013). "Dynamic Elicited Priors for Updating Covert Networks". In: Network Science 1.01, pp. 68–94.

- Gintis, Herbert and Dirk Helbing (2015). "Homo Socialis: An Analytical Core for Sociological Theory". In: Review of Behavioral Economics 2.1-2, pp. 1–59.
- Goldstone, Jack A. et al. (2005). A Global Forecasting Model of Political Instability. Political Instability Task Force.
- Green, Donald and Ian Shapiro (1996). Pathologies of Rational Choice Theory: A Critique of Applications in Political Science. Yale University Press.
- Grieg, J Michael et al. (2010). National Material Capabilities Data Documentation, Version 4.0.
- Grimm, Volker et al. (2006). "A Standard Protocol for Describing Individual-Based and Agent-Based Models". In: *Ecological modelling* 198.1, pp. 115–126.
- Hagan, Joe D. (1994). "Domestic Political Systems and War Proneness". In: Mershon International Studies Review 38.2, pp. 183–207.
- Hagerman, Edward (1992). The American Civil War and the Origins of Modern Warfare: Ideas, Organization, and Field Command. Vol. 715. Indiana University Press.
- Hendrix, Cullen S. and Idean Salehyan (2013). Social Conflict in Africa Database (SCAD).
- Henrich, Joseph et al. (2005). ""Economic Man" in Cross-Cultural Perspective: Behavioral Experiments in 15 Small-Scale Societies". In: *Behavioral and Brain Sciences* 28.06, pp. 795–815.
- Heuer, Richards J (2001). Psychology of Intelligence Analysis. Center for the Study of Intelligence, Central Intelligence Agency.
- Ho, Teck H., Colin F. Camerer, and Juin-Kuan Chong (2007). "Self-Tuning Experience Weighted Attraction Learning in Games". In: *Journal of Economic Theory* 133.1, pp. 177– 198.
- Hoff, Peter D, Adrian E Raftery, and Mark S Handcock (2002). "Latent Space Approaches to Social Network Analysis". In: Journal of the American Statistical Association 97.460, pp. 1090–1098.
- Hoffmann, Matthew J. (2005). "Self-Organized Criticality and Norm Avalanches." In: Proceedings of the Symposium on Normative Multi-Agent Systems, pp. 117–125.

- Hudson, Valerie M., Philip A. Schrodt, and Ray D. Whitmer (2004). "A New Kind of Social Science: The Path Beyond Current (IR) Methodologies May Lie Beneath Them". In: Annual International Studies Association Conference, Montreal, Quebec, Canada.
- Jervis, Robert (1976). Perception and Misperception in International Politics. Princeton University Press.
- Jiang, Chengzhi and Zhaohan Sheng (2009). "Case-Based Reinforcement Learning for Dynamic Inventory Control in a Multi-Agent Supply-Chain System". In: *Expert Systems* with Applications 36.3, Part 2, pp. 6520–6526.
- Jones, Daniel M., Stuart A. Bremer, and J. David Singer (1996). "Militarized Interstate Disputes, 18161992: Rationale, Coding Rules, and Empirical Patterns". In: Conflict Management and Peace Science 15.2, pp. 163–213.
- Kahneman, Daniel and Amos Tversky (1984). "Choices, Values, and Frames". In: American Psychologist 39.4, p. 341.
- Kaufmann, Chaim D. (1994). "Out of the Lab and Into the Archives: A Method for Testing Psychological Explanations of Political Decision Making". In: International Studies Quarterly 38.4, pp. 557–586.
- Keohane, Robert O., ed. (1986). Neorealism and Its Critics. Columbia University Press.
- Keohane, Robert O. and Joseph S. Nye (1987). "Power and Interdependence Revisited". In: International Organization 41.04, pp. 725–753.
- Khong, Yuen Foong (1992). Analogies at War: Korea, Munich, Dien Bien Phu, and the Vietnam Decisions of 1965. Princeton University Press.
- Kim, Woosang and Bruce Bueno de Mesquita (1995). "How Perceptions Influence the Risk of War". In: International Studies Quarterly 39.1, pp. 51–65.
- Kimbrough, Steven O. (2014). Making Predictions: The Group Decision Forecasting (GDF) Problem. Social Science Research Network.
- King, Gary and Langche Zeng (2001). "Explaining Rare Events in International Relations".In: International Organization 55.03, pp. 693–715.

- Kocher, Martin G. and Matthias Sutter (2005). "The Decision Maker Matters: Individual Versus Group Behaviour in Experimental Beauty-Contest Games". In: *The Economic Journal* 115.500, pp. 200–223.
- Korab-Karpowicz, W. Julian (2013). "Political Realism in International Relations". In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Summer 2013.
- Kugler, Jacek and Yi Feng, eds. (1997). International Interactions 23.3-4: The Expected Utility Approach to Policy Decision Making: Assessments, Forecasts and Strategies.
- Lamborn, Alan C. (1991). The Price of Power: Risk and Foreign Policy in Britain, France, and Germany. Unwin Hyman.
- Lanchester, Frederick William (1916). Aircraft in Warfare: The Dawn of the Fourth Arm. Constable.
- Laslier, Jean-Franois and Bernard Walliser (2005). "A Reinforcement Learning Process in Extensive Form Games". In: International Journal of Game Theory 33.2, pp. 219–227.
- Latek, Maciej, Robert Axtell, and Bogumil Kaminski (2009). "Bounded Rationality Via Recursion". In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '09. International Foundation for Autonomous Agents and Multiagent Systems, pp. 457–464.
- Lebow, Richard Ned and Thomas Risse-Kappen (1995). International Relations Theory and the End of the Cold War. Columbia University Press.
- Leetaru, Kalev and Philip A. Schrodt (2013). "GDELT: Global Data on Events, Location, and Tone, 19792012". In: International Studies Association Annual Convention. Vol. 2, p. 4.
- Leng, Russell J. and J. David Singer (1988). "Militarized Interstate Crises: The BCOW Typology and its Applications". In: International Studies Quarterly, pp. 155–173.
- Levy, Jack S. (1986). "Organizational Routines and the Causes of War". In: International Studies Quarterly 30.2, pp. 193–222.
- Littman, Michael L. (1994). "Markov Games as a Framework for Multi-Agent Reinforcement Learning." In: *ICML*. Vol. 94, pp. 157–163.

- Lundberg, Craig C. (1995). "Learning in and by Organizations: Three Conceptual Issues".In: The International Journal of Organizational Analysis 3.1, pp. 10–23.
- Maoz, Zeev (2010). Networks of Nations: The Evolution, Structure, and Impact of International Networks, 18162001. Vol. 32. Cambridge University Press.
- Maoz, Zeev and Nasrin Abdolali (1989). "Regime Types and International Conflict, 1816-1976". In: Journal of Conflict Resolution 33.1, pp. 3–35.
- Maoz, Zeev et al. (2015). Corelates of War History. http://www.correlatesofwar.org/ history.
- March, James G. and Herbert A. Simon (1993). Organizations. Wiley.
- Mariano, Karina Lilia Pasquariello (2013). Two to Tango: An Analysis of Brazilian-Argentine Relations. SciELO Brasil.
- Marshall, Monty G. (2008). Fragility, Instability, and the Failure of States. Working Paper. Council on Foreign Relations, Center for Preventive Action.
- Masad, David (2014). BDM Model Sensitivity Analysis. http://davidmasad.com/blog/ bdm-model-sensitivity-analysis/.
- McDermott, Rose (2001). Risk-Taking in International Politics: Prospect Theory in American Foreign Policy. University of Michigan Press.
- McFadden, Daniel (1984). *Econometric Analysis of Qualitative Response Models*. Handbook of Econometrics. Elsevier, pp. 1395–1457.
- McKibben-Sanders, Jeremy (2014). BDM Scholz Expected Utility Model. https://github. com/jmckib/bdm-scholz-expected-utility-model.
- Meltzoff, Andrew N. and Jean Decety (2003). "What Imitation Tells Us About Social Cognition: A Rapprochement Between Developmental Psychology and Cognitive Neuroscience". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 358.1431, pp. 491–500.
- Metternich, Nils W. et al. (2013). "Antigovernment Networks in Civil Conflicts: How Network Structures Affect Conflictual Behavior". In: American Journal of Political Science 57.4, pp. 892–911.

- Milius, John and Kevin Reynolds (1984). *Red Dawn*. Directed by John Milius. United Artists.
- Miller, John H. and Scott E. Page (2007). Complex Adaptive Systems: An Introduction to Computational Models of Social Life. Princeton University Press.
- Min, Byoung Won (2002). "Trade and War in Cellular Automata Worlds: A Computer Simulation of Interstate Interactions". Ph.D. The Ohio State University.
- Morgan, M. Granger (2014). "Use (and Abuse) of Expert Elicitation in Support of Decision Making for Public Policy". In: Proceedings of the National Academy of Sciences 111.20, pp. 7176–7184.
- Myerson, Roger B. (2009). "Learning from Schelling's Strategy of Conflict". In: Journal of Economic Literature, pp. 1109–1125.
- Nagl, John A. (2002). Counterinsurgency Lessons from Malaya and Vietnam: Learning to Eat Soup with a Knife. First edition. Praeger.
- Noon, David Hoogland (2004). "Operation Enduring Analogy: World War II, the War on Terror, and the Uses of Historical Memory". In: *Rhetoric & Public Affairs* 7.3, pp. 339– 364.
- North, Robert C. (1967). "Perception and Action in 1914 Crisis". In: Journal of International Affairs 21.1, pp. 103–122.

Nye, Joseph S. (1990). "Soft Power". In: Foreign Policy 80, pp. 153–171.

- Ozer-imer, Itir, Jacek Kugler, and Hilton L. Root (2013). "Turkey and the West: Bargaining for Realignment". In: Journal of Contemporary European Studies 21.2, pp. 218–251.
- Palmer, Glenn et al. (2015). "The MID4 dataset, 2002-2010: Procedures, Coding Rules and Description". In: Conflict Management and Peace Science 32.2, pp. 222–242.
- Papadimitriou, Christos H. and Tim Roughgarden (2005). "Computing Equilibria in Multi-Player Games". In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete algorithms. Society for Industrial and Applied Mathematics, pp. 82–91.
- Podeh, Elie (1995). The Quest for Hegemony in the Arab World: The Struggle Over the Baghdad Pact. Brill.

- Powell, Robert (2006). "War as a Commitment Problem". In: International Organization 60.01.
- Prados, John (2011). How the Cold War Ended: Debating and Doing History. Potomac Books.
- Raleigh, Clionadh et al. (2010). "Introducing ACLED: An Armed Conflict Location and Event Dataset Special Data Feature". In: Journal of Peace Research 47.5, pp. 651–660.
- Ram, A. and J. C. Santamara (1997). "Continuous Case-Based Reasoning". In: Artificial Intelligence 90.12, pp. 25–77.
- Ramanarayanan, Ananth (2011). "Distance and the Impact of 'Gravity' Help Explain Patterns of International Trade". In: *Economic Letter* 6.7.
- Rapoport, Anatol (1957). "Lewis F. Richardson's Mathematical Theory of War". In: Conflict Resolution 1.3, pp. 249–299.
- Richardson, Lewis F. (1948). "Variation of the Frequency of Fatal Quarrels with Magnitude". In: Journal of the American Statistical Association 43.244, pp. 523–546.
- Robins, Garry et al. (2007). "An Introduction to Exponential Random Graph (p\*) Models for Social Networks". In: Social Networks 29.2, pp. 173–191.
- Rouleau, Mark D. (2011). "A Computational Theory of Endogenous Norm Change: The NormSim Agent-Based Model in MASON". PhD thesis. George Mason University.
- Russell, Stuart and Peter Norvig (2009). Artificial Intelligence: A Modern Approach. 3 edition. Prentice Hall.
- Sarkees, Meredith Reid and Phil Schafer (2000). "The Correlates of War Data on War: An Update to 1997". In: Conflict Management and Peace Science 18.1, pp. 123–144.
- Schaller, Michael (1997). Altered States: The United States and Japan Since the Occupation. Oxford University Press.
- Schelling, Thomas (1966). Arms and Influence. Yale University Press.
- Scholz, Jason B., Gregory J. Calbert, and Glen A. Smith (2011). "Unravelling Bueno de Mesquita's Group Decision Model". In: Journal of Theoretical Politics 23.4, pp. 510– 531.

- Schrodt, Philip A. (1997). "Early Warning of Conflict in Southern Lebanon Using Hidden Markov Models". In: American Political Science Association.
- (2004). "Patterns, Rules and Learning: Computational Models of International Behavior".
- Schrodt, Philip A., John Beieler, and Mohammed Idris (2014). "Three's a Charm?: Open Event Data Coding with EL: DIABLO, PETRARCH, and the Open Event Data Alliance." In: Annual International Studies Association Conference, Toronto, Quebec, Canada.
- Schrodt, Philip A., Erin M. Simpson, and Deborah J. Gerner (2001). "Monitoring Conflict Using Automated Coding of Newswire Reports: A Comparison of Five Geographical Regions". In: *Identifying Wars: Systematic Conflict Research and it's Utility in Conflict Resolution and Prevention*, pp. 8–9.
- Schuman, Howard and Amy D. Corning (2006). "Comparing Iraq to Vietnam Recognition, Recall, and the Nature of Cohort Effects". In: Public Opinion Quarterly 70.1, pp. 78–87.
- Schuman, Howard and Cheryl Rieger (1992). "Historical Analogies, Generational Effects, and Attitudes Toward War". In: American Sociological Review 57.3, pp. 315–326.
- Senese, Paul D. (1996). "Geographical Proximity and Issue Salience: Their Effects on the Escalation of Militarized Interstate Conflict". In: Conflict Management and Peace Science 15.2, pp. 133–161.
- Shapira, Zur (2002). Organizational Decision Making. Cambridge University Press.
- Shoham, Yoav and Kevin Leyton-Brown (2009). Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Cambridge University Press.
- Shoham, Yoav, Rob Powers, and Trond Grenager (2003). Multi-Agent Reinforcement Learning: A Critical Survey. Web manuscript.
- Signorino, Curtis S. (1999). "Strategic Interaction and the Statistical Analysis of International Conflict". In: The American Political Science Review 93.2, p. 279.

- Signorino, Curtis S. and Jeffrey M. Ritter (1999). "Tau-b or Not Tau-b: Measuring the Similarity of Foreign Policy Positions". In: International Studies Quarterly 43.1, pp. 115– 144.
- Simon, Herbert A. (1976). "From Substantive to Procedural Rationality". In: 25 Years of Economic Theory. Ed. by T. J. Kastelein et al. Springer, pp. 65–86.
- Singer, David J. (1961). "The Level-of-Analysis Problem in International Relations". In: World Politics 14.01, pp. 77–92.
- Singer, J. David (1972). "The "Correlates of War" Project: Interim Report and Rationale". In: World Politics 24.2, pp. 243–270.
- (1988). "Reconstructing the Correlates of War Dataset on Material Capabilities of States, 1816-1985". In: International Interactions 14.2, pp. 115–132.
- Small, Melvin and J. David Singer (1969). "Formal Alliances, 1816–1965: An Extension of the Basic Data". In: Journal of Peace Research 6.3, pp. 257–282.
- Smith, Alastair (1999). "Testing Theories of Strategic Choice: The Example of Crisis Escalation". In: American Journal of Political Science 43.4, p. 1254.
- Snidal, Duncan (1985). "The Game Theory of International Politics". In: World Politics 38.01, pp. 25–57.
- Steinbruner, John (1976). "Beyond Rational Deterrence: The Struggle for New Conceptions". In: World Politics 28.02, pp. 223–245.
- Stokman, Frans N. and Reinier Van Oosten (1994). "The Exchange of Voting Positions: An Object-Oriented Model of Policy Networks". In: European community decision making. Models, applications, and comparisons, pp. 105–127.
- Sutton, Richard S. and Andrew G. Barto (1998). Reinforcement Learning: An Introduction. Bradford.
- Taber, Charles S. (1992). "POLI: An Expert System Model of US Foreign Policy Belief Systems". In: American Political Science Review 86.04, pp. 888–904.
- Taylor, Glenn et al. (2008). PSTK: A Toolkit for Modeling Dynamic Power Structures. Soar Technology, Inc.

- Tetlock, Philip (2005). Expert Political Judgment: How Good Is It? How Can We Know? Princeton University Press.
- Thompson, William R. and Richard Tucker (1997). "A Tale of Two Democratic Peace Critiques". In: Journal of Conflict Resolution 41.3, pp. 428–454.
- Thomson, Robert et al. (2006). The European Union Decides. Cambridge University Press.
- Tipping, Michael E. and Chris M. Bishop (1999). "Mixtures of Probabilistic Principal Component Analyzers". In: Neural computation 11.2, pp. 443–482.
- Tsebelis, George (1988). "Nested Games: The Cohesion of French Electoral Coalitions". In: British Journal of Political Science 18.02, pp. 145–170.
- Tversky, A. and D. Kahneman (1981). "The Framing of Decisions and the Psychology of Choice". In: Science 211.4481, pp. 453–458.
- United States et al. (1948). North Atlantic Treaty.
- US 4th Circuit (2011). Decision Insights, Inc. v. Sentia Group, Inc., No 09-2300.
- Walt, Stephen M. (1987). The Origins of Alliances. Cornell Univ Press.
- Waltz, Kenneth N. (2000). "NATO Expansion: A Realist's View". In: Contemporary Security Policy 21.2, pp. 23–38.
- Ward, Michael D. et al. (2013a). "Comparing GDELT and ICEWS Event Data". In: Analysis 21, pp. 267–297.
- Ward, Michael D et al. (2013b). "Learning From The Past and Stepping Into The Future: Toward a New Generation of Conflict Prediction". In: International Studies Review 15.4, pp. 473–490.
- Washington, George (1796). Farewell Address.
- Wayman, Frank Whelon and Paul Francis Diehl (1994). Reconstructing Realpolitik. University of Michigan Press.
- Wendt, Alexander E. (1987). "The Agent-Structure Problem in International Relations Theory". In: International organization 41.03, pp. 335–370.
- Wilensky, Uri and William Rand (2007). "Making Models Match: Replicating an Agent-Based Model". In: Journal of Artificial Societies and Social Simulation 10.4, online.

- Wilson, Ernest J. (2008). "Hard Power, Soft Power, Smart Power". In: The ANNALS of the American Academy of Political and Social Science 616.1, pp. 110–124.
- Wise, Ben, Leo Lester, and Brian Efird (2015a). An Introduction to the KAPSARC Toolkit for Behavioral Analysis (KTAB) Using One-Dimensional Spatial Models. King Abdullah Petroleum Studies and Research Center.
- (2015b). *Multidimensional Bargaining Using KTAB*. King Abdullah Petroleum Studies and Research Center.
- Wohlforth, William C. (1999). "The Stability of a Unipolar World". In: International Security 24.1, pp. 5–41.
- Wooldridge, Jeffrey M. (2008). Introductory Econometrics: A Modern Approach. Cengage Learning.
- Yonamine, James E. (2011). Working with Event Data: A Guide to Aggregation Choices.Working Paper, Penn State University.
- (2013). "Predicting Future Levels of Violence in Afghanistan Districts Using GDELT".
  Working Paper, Penn State University.

## Curriculum Vitae

David Paymar Masad attended the Tel Aviv University Science and Technology High School, and graduated from the University of Chicago in 2009 with a BA in economics and a minor in Near Eastern Studies. He worked as a research assistant at the Institute for Defense Analyses before enrolling in the Computational Social Science doctoral program. David was the recipient of the 2012 GMU Presidential Scholarship, providing three years of tuition and assistantship funding.

As a graduate research assistant, David worked with faculty on building models and analyzing data, covering topics from the Syrian conflict to DC-area mortgages. His own research focused on applying network analysis and agent-based modeling to international relations, with a focus on conflict and event data. In 2013, he attended the Santa Fe Institute's Complex Systems Summer School, researching the co-evolution of network attackers and defenders, and the networks emergent from Wikipedia edits. As a summer intern at Caerus Analytics and the MITRE Corporation, David led internal research efforts to develop agent-based models for policy support.

In addition to the Presidential Scholarship, David received university funding to attend numerous conferences where he presented his research. He received NSF grants to attend the 2014 SBP Conference, and the 2014 Political Networks Conference, where his presentation of *Dyadic Competition in the International Arms Trade Network* was awarded a Best Methodology Poster award. Over the course of his Ph.D. research, David received a 2015 summer research grant, and a 2016 dissertation completion grant.

David is a candidate for a PhD in Computational Social Science, in the Computational and Data Sciences department at George Mason University.