

Information Market Failure: Lapses in Public Consent and Their Consequences for Democratic Crisis Diplomacy


by

Neill A. Mohammad

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Political Science)
in The University of Michigan
2012

Doctoral Committee:

Professor James D. Morrow, Chair
Professor Ted Brader
Professor Allan Stam
Professor Michael W. Traugott



© Neill Mohammad 2012
All Rights Reserved

To C, R, and J.

ACKNOWLEDGEMENTS

I spent a long time in graduate school, which has left with a long list of people to acknowledge and appreciate in this space. In roughly chronological order, I am indebted first to Bill Zimmerman, my first advisor at Michigan, who was impressively patient as I managed to flit between post-Soviet comparative studies, international security theory and complex systems analysis during my first few semesters in Ann Arbor. Jim Morrow, who wound up chairing my committee, also demonstrated an impressive amount of faith in the idea that I would be able to produce a worthwhile project, even after I had spent over two years in management consulting without having really considered the idea that I might return to finish my degree. Ted Brader, Mike Traugott and Allan Stam have been similarly gracious with their time and energy. They deserve special credit for buying into my recruitment effort that began in the summer of 2009, when I was finally able to start on a prospectus in earnest as I was winding down my consulting work. Mike was indispensable in helping me work through the pilot tests of the experimental chapter, in particular when the first such test fell flat on its face. I also relied on the good spirits of Ken Kollman and Bill Clark, who supervised my effort to write my preliminary examination paper in research methodology. Most of that paper was written on Waikiki Beach in Honolulu, five time zones and seemingly a million miles (and 30 years of economic and social development) from Ann Arbor. That project would become the foundation of the modeling and theory chapter presented here.

I am also grateful for the support of the Department of Political Science staff, and

in particular Michelle Spornhauer, Joann Nemeth, Kathy Halleck, Kristin Romelhardt, and Vicky Horvath. I received valuable research funding from the Rackham Graduate School and the Gerald R. Ford Fellowship fund, as well as gainful employment from the Rackham Information and Technology Services from 2010-11. The Rackham and Ford awards funded the two experimental waves presented in Chapter 4 as well as travel to conferences and other incidentals. I was also privileged to work with Edie Goldenberg and Margaret Howard in Washington, DC for the “Michigan in Washington” program from 2011-12. Elsewhere, Alexander Nestorov gathered most of the data used in Chapter 5, and went above and beyond the call of duty in searching out a number of sources across the Michigan campus.

The seed notion of this project—that something is amiss with the way that international politics scholars think about elections and their relationship to democratic security behavior more widely—was originally developed with Scott Helfstein and Dominick’ Wright in a paper presented at the 2005 Annual Meeting of the American Political Science Association, and again in an internal seminar at Michigan the following spring. Fighting tooth-and-nail with Scott and Dom during that process over what, precisely, the project was *actually about* was an invaluable experience that shaped quite a bit of how I would later come to think about both my dissertation and the field itself. I’m grateful for that experience even if it—and this actually happened and is not an exaggeration—put me in an MRI tube at the University of Michigan medical school with unexplained chronic migraine headaches.

The headaches mysteriously disappeared after APSA.

I am also thankful for a number of other friends and colleagues who were kind enough to share their feedback and expertise on various seminar papers, conference presentations and the general ennui that research can sometimes produce. This unordered and by-no-means-exhaustive list includes Alex Von Hagen-Jamar, Shanna Kirschner, Michelle Allendoerfer, Paul Poast, Gary Uzonyi, Charles Doriean, David

Smith, Matt Wells, Cassie Grafstrom, Cali Ellis, Spencer Piston, Nathan Kalmoe and many more participants in the Nameless Internal Seminar and Inter-disciplinary Workshop in American Politics series. Andrew Kerner was helpful in giving me early feedback on the financial event study chapter, including the sage and prescient advice that relying on it as my only empirical chapter would be a huge mistake.

David Smith, Greg Lewis, Charles Doriean and Steve Meserve have all been great friends throughout this process and important sources of support, as has Andrea Benjamin (who has a promising career in life-coaching if she ever leaves the academy). This dissertation would not have been completed without an enormous amount of support and understanding from Joanne Hsu. Finally, Steve Meserve and Carie Steele were exceptionally gracious during some trying times near the end of the writing process.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	viii
LIST OF TABLES	x
ABSTRACT	xii
CHAPTER	
1. Introduction	1
1.1 Sketch of Research Design(s)	8
1.1.1 Multiple Concepts & Multiple Methods	11
1.2 Outline of Remaining Sections	13
2. Related Literature, Descriptive Theory, and Empirical Chal- lenges	14
2.1 Information Market Intervention in Motion	15
2.2 “MindWar” in a Political Science Framework	16
2.3 The Relationships between Regime, Crisis Bargaining, and War	18
2.4 Information and Voting	22
2.5 Voting on Foreign Policy	25
2.6 “Information Market Intervention”: a Theoretical Sketch	28
2.7 Empirical Obstacles: The “Rally ’Round the Flag”	32
3. Formalizing Information Market Intervention	41
3.1 A Synthesis of Public Opinion and Bargaining, and a Theory	41
3.2 The Model	45
3.3 Equilibrium Analysis under Complete Information	50

3.4	Equilibrium Analysis Under Incomplete Information	53
3.5	Interpretations	66
4.	Radio Free USA: Does Information Market Intervention Work?	69
4.1	Prior Experimental Work	71
4.1.1	Treatment Design and Hypotheses	74
4.2	Experimental Procedure	78
4.3	Results (Part 1)	84
4.3.1	Means by Treatment Conditions	84
4.3.2	Econometric Estimations	86
4.4	Results (Part 2)	97
4.5	Concluding Thoughts	105
5.	The Wages of Sin: When Does Information Market Intervention Occur?	109
5.1	Literature Review: Financial Economics, Political Information, and the Event Study	110
5.2	Understanding Financial Markets and their Link to Information Market Intervention	113
5.2.1	A Primer on Finance, Uncertainty and Information	113
5.2.2	Uncertainty and War	115
5.2.3	Financial Markets and Conflict Onset	120
5.3	Data and Method: an Event Study Approach to American Crisis Diplomacy	123
5.3.1	Estimation	124
5.3.2	Data	128
5.3.3	Selection of Conflict Events	132
5.4	Analysis and Results	139
5.4.1	All Firms, and All Events	139
5.4.2	“Core” Defense Firms	141
5.4.3	Regression Analysis	145
5.5	Concluding Thoughts	147
6.	Concluding Thoughts and Observations	150
6.1	Observations About the Study of Foreign Policy	156
	BIBLIOGRAPHY	159

LIST OF FIGURES

Figure

2.1	The extensive form of the immediate deterrence game. Strategies are labeled in <i>italics</i> along each edge (<i>Challenge</i> , <i>Resist</i> , <i>StandFirm</i>), with outcomes (Status Quo, Defender Concedes, Challenger Backs Down, War) at the end of each branch as appropriate.	30
2.2	The timing of information market intervention.	32
2.3	Rally events graphed by percentage-point change. The shaded area indicates the average joint margin of error.	39
3.1	The Crisis Censorship game under complete information. Incomplete information adds a second branch to the game with an initial move by Nature that determines the Challenger’s cost term, such that the Defender only has two information sets: one after observing a challenge with prior intervention (<i>Ch, C</i>) and one after observing a challenge without such intervention (<i>Ch, ¬C</i>).	49
3.2	Boundary equilibrium conditions for the restricted Crisis Censorship game for varying parameters	53
4.1	Confidence in various government and media institutions, from the General Social Survey. In the cumulative data file spanning 1972 to 2008, confidence in institutions was measured with just three categories. “A great deal” of confidence is coded 1, while “only some” confidence is coded 2 and “hardly any” confidence is labeled 3. The trend lines indicate the mean confidence rating across all respondents in each year the question was asked.	73
4.2	“Information market intervention” revealed.	75
4.3	The opening lines of the NSA-authored leak, as reported in the <i>Wall Street Journal</i> , August 25th, 1986.	76
4.4	Stimulus for the “Military Expert” condition.	80
4.5	Time spent viewing stimuli, by treatment group, June 2011.	81
4.6	Self-reported interest in news from Iraq, June 2011	85
4.7	Marginal effects of treatment assignment, conditional on partisanship, for “Iran a threat?”	90
4.8	Marginal effects of treatment assignment, conditional on partisanship, for “Support escalation?”	91

4.9	Marginal effects of treatment assignment, conditional on partisanship, for “Handle relations?”	92
4.10	Marginal effects of treatment assignment, conditional on partisanship, for “Trust the president?”	93
4.11	Marginal effects of treatment assignment, conditional on partisanship, for “Made more secure?”	94
4.12	Mean response to whether or not the Iraq War was “the right thing to do,” separated by treatment groups, August 2011.	101
4.13	Mean response to whether or not the Iraq War was “worth it,” separated by treatment groups, August 2011.	104
5.1	Daily Mentions of ‘Iraq,’ 1988-1991	119
5.2	Daily Mentions of ‘Iraq,’ 2000-2003	119
5.3	Daily returns for a selection of prominent defense firms, related to the Gulf and Iraq Wars	121
5.4	An event-history model. The baseline estimation period runs for τ days, ending k_1 days before the event. Equity behavior is measured in a window about the event; the longest possible event window would be $k_1 + k_2$ days. The event of interest falls at t_0	123

LIST OF TABLES

Table

2.1	Model 3, Chapman and Reiter (2004) – replications with modification. . .	36
2.2	Distribution of rally events in Chapman and Reiter (2004), after re-coding.	38
2.3	Marginal effects of UNSC Authorization in the ordered probit estimation.	40
4.1	Means for control variables, separated by treatment condition, for June 2011 (Iran) data.	83
4.2	Randomization check for June 2011 data. The critical value from the F-distribution characterizes a null hypothesis that the mean of each variable is identical across all 4 treatment groups.	84
4.3	Means and standard deviations of five dependent variables, by treatment condition.	85
4.4	Three dummy variables to distinguish between four experimental conditions (the control condition serves as a baseline).	86
4.5	Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions (*: $p < .05$; **: $p < .01$; ***: $p < .001$) . .	88
4.6	Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions and interactive effects between treatment conditions and partisan identity (*: $p < .05$; **: $p < .01$; ***: $p < .001$)	89
4.7	Pairwise correlation scores for party identification and evaluations of the Iraq War, for June 2011 panel	96
4.8	Model restriction tests for June 2011 panel. Both sets of restrictions tested against the full regressions shown above. “Is Iran a threat?” response modeled with logistic estimator, rather than linear.	96
4.9	Means for control variables, separated by treatment condition, for June 2011 (Iran) data.	100
4.10	Randomization check for August 2011 data. The critical value from the F-distribution characterizes a null hypothesis that the mean of each variable is identical across all 4 treatment groups.	101
4.11	Estimations of each dependent variable from the Venezuelan experiment, with dummy variables distinguishing between treatment conditions (*: $p < .05$; **: $p < .01$; ***: $p < .001$)	102

4.12	Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions and interactive effects between treatment conditions and partisan identity (*: $p < .05$; **: $p < .01$; ***: $p < .001$)	103
4.13	Model restriction tests for August 2011 panel. Both sets of restrictions tested against the full regressions shown above. “Is Venezuela a threat?” response modeled with logistic estimator, rather than linear.	103
5.1	Contrasting values for crisis / dispute starting dates.	116
5.2	Sales totals for selected firms. Total sales and % revenue from arms taken from SIPRI; U.S. Department of Defense sales from author’s collection from FPDS. Additional sales information from non-SIPRI-ranked firms taken from annual investor reports. All revenue reported in millions of 2000 US\$.	130
5.3	Representative product offerings from included American defense contractors, reproduced from SIPRI yearbooks.	132
5.4	Intra-war events, from 1962 onward.	135
5.5	Intra-war events and events falling on non-trading days, from 1962 onward.	137
5.6	All included crises from the International Crisis Behavior dataset, from 1962 onward.	138
5.7	Cumulative abnormal returns and associated rank statistics for American international crises since 1962.	141
5.8	Cumulative abnormal returns and associated rank statistics for American international crises since 1975.	142
5.9	Cumulative abnormal returns and associated rank statistics for American international crises since 1983.	142
5.10	Cumulative abnormal returns and associated rank statistics for American international crises since 1986.	143
5.11	Cumulative abnormal returns and associated rank statistics for American international crises since 1962, core defense firms only.	144
5.12	A time-series cross-sectional estimation of daily abnormal returns in each crises’ event window. All American crises since July 1962 are included, with the daily S&P 500 index serving as the only portfolio control for the purposes of constructing the abnormal return series.	146

ABSTRACT

Information Market Failure: Lapses in Public Consent and Their Consequences for Democratic Crisis Diplomacy

by

Neill Mohammad

Chair: James D. Morrow

Current understandings of the effects of domestic regimes on international security behavior rely on voting behavior in at least two ways. First, the need for democratic executives to provide public goods to large groups of supporters in order to retain office produces clear *monadic* behavioral patterns. These include a preference for military quality over quantity, and in most cases incentives against rash or speculative foreign adventures. Second, because democratic leaders are accountable to the popular will—and, most importantly, recognize that their counterparts in other democratic states face the same incentives they do—we can also observe certain *dyadic* behaviors between democratic states. The most well-known of these, of course, is the “democratic peace.”

However, these arguments implicitly rely on voting publics that can punish and reward their leaders appropriately. In this dissertation, I argue that the assumptions that world politics scholars make about voting behavior are at odds with what we know about voters’ conduct from the study of American politics. The clear partisan divides that are necessary to generate meaningful public opposition to government

policy are rare at the outset of a conflict or dispute. When we combine this observation with the theoretical insight that public support for a government's security policy is most useful at the outset of a crisis, then we are left with a puzzle: democratic executives have the most latitude in shaping public sentiment toward a potential conflict at the exact moment in time when that latitude is the most strategically useful.

I refer to this influence as “information market intervention” and develop this argument in three distinct applications. First, I amend a common crisis-bargaining model to account for the possibility of endogenous public support. Second, I present results from a survey experiment that demonstrate how the sources of voters' evaluations of prospective foreign policy are distinctly different from evaluations of prior events. Finally, I examine the behavior of common equity in prominent American defense firms to isolate those historical episodes in which elite manipulation of public opinion was more likely to have occurred.

CHAPTER 1

Introduction

“We must be the great arsenal of democracy. For us this is an emergency as serious as war itself. We must apply ourselves to our task with the same resolution, the same sense of urgency, the same spirit of patriotism and sacrifice as we would show were we at war.”

Franklin D. Roosevelt, December 29, 1940.

The “arsenal of democracy” referred to both the destroyers-for-bases agreement with Great Britain, concluded on September 2nd of that year, and the forthcoming Lend-Lease Program, which would be completed on March 11th, 1941, as measures that would keep the United States out of an European war by equipping the British (and, eventually, the Soviets) to win the fight on their own. We now know, however, that the United States was already engaged in what would become the Battle of the Atlantic by 1939, when President Roosevelt ordered Navy and Coast Guard patrols to broadcast the positions of any suspicious craft in “plain English” and to follow them as long as possible while maintaining such a broadcast (Bailey and Ryan 1979). The U.S. Navy, in other words, was acting as an auxiliary of the Royal Navy a bit more than two years before officially entering the war. This was notionally meant to police an American “zone of exclusion” in its sphere of influence from the various European belligerents. If the British, sailing American-built destroyers, sank or captured Axis ships near American waters after hearing an unencrypted, English-language re-

port broadcast by the American Navy, then—in Roosevelt’s formulation—that would be a regrettable but nonetheless unavoidable consequence of America defending its neutrality against foreign aggression.

Roosevelt’s reluctance to ask Congress for a formal declaration of war until after the attack at Pearl Harbor is typically described as a product of the shackles of isolationism, both in the Capitol and across the broader nation. Nonetheless, not even the Neutrality Acts of 1936, 1937 and 1939 were enough to prevent the U.S. from not only arming the British, first under the “cash-and-carry” principle and then later, when Britain no longer had the merchant shipping capacity necessary to “carry” much of anything herself, by transporting goods and war materiel themselves. Ultimately, America provided her own armed escorts for supply convoys across the Atlantic. In April 1941, the American destroyer *Niblack* opened fire at what her crew thought was a German U-boat sonar signature. In September, the *Greer* fired on another U-boat after taking torpedo fire from an unidentified craft. It is unsurprising, then, that by Autumn 1941 Americans overwhelmingly believed that they were already at war “for all intents and purposes” (*Gallup and Fortune Polls* 1941). The Republican Party’s isolationist platform had found its limits long before the “day that would live in infamy.”

The growing involvement of the Navy in the convoy business was in fact broadly supported by the public, which had also come to view Hitler as a serious threat to U.S. security by 1941. Certainly no small part of this change in opinion was a reaction to events as they unfolded, as Germany conquered Poland, Norway, the Low Countries and France before invading the Soviet Union. However, Roosevelt was not content to let ordinary Americans reach just any conclusion of their own choosing regarding which side, if any, the United States should favor. In his Navy Day address in October 1941, Roosevelt claimed to have a copy of German plans for partitioning South America—in all likelihood, a British forgery—as well as the

liturgy of an “International Nazi Church” that would replace Christianity after an Axis victory (Bailey and Ryan 1979, p. XX). He also reassured a pool reporter during a weekly press conference after the start of Operation Barbarossa that the Soviet Union was actually quite liberal when it came to matters of religious expression, and thus by implication a worthy would-be ally should the U.S. enter the war against Germany.¹ By hook or by crook, American opinion was solidly behind a prospective war effort even before December 7, 1941 and remained so until the final victory in 1945 (Berinsky 2007).



Modern research into matters of war and peace has found that liberal and illiberal states behave very differently from one another on the international stage. Both liberal and illiberal governments fight wars. However, liberal states win a greater percentage of their wars than do their illiberal counterparts, even when accounting for the self-evident point that liberal states tend to be wealthier and to enjoy a concomitant advantage in military quality. Liberal states seem to engage in fewer speculative challenges against other states. And, most famously, liberal states do not fight one another. This latter observation produced a vital research agenda in world politics that is now thirty years old, and the vulgar characterization that democracy makes peace has become a rule of thumb among the American foreign policy establishment. While considerable disagreement remains about the mechanism by which liberal states have realized Kant’s prescription of a “separate peace” among themselves (Gartzke 2007; Dafoe 2011), there is little sound research to suggest that the empirical observation of the liberal peace is up to mere chance or a temporary and exogenous coincidence of preferences.

¹To be fair, this required somewhat less prevarication than one might imagine. The 1936 Soviet Constitution was among the world’s most progressive, as it provided for universal suffrage, direct elections and improved rights for women in addition to religious freedoms and a host of other advances. That Stalin had just murdered more than a million people in the Great Purge was, at least syntactically, a small omission.

The resilience of this observation about the patterns of war and peace among liberal and illiberal states produced, in turn, a derby to find a rigorous theory which could explain it. The most promising theories all depend on the clearest difference between liberal and illiberal states: elections. Elections allow voters to choose leaders *prospectively*—which candidate is likely to prove best able to handle an uncertain and dangerous road ahead?—as well as *retrospectively*—how did the incumbent perform in their prior term? Although there is some disagreement over the extent to which either mode of thought dominates in most voters’ minds (e.g. Fiorina 1981), both the prospective and retrospective voting models suggest that democracies ought to be led by more prudent leaders who, as a consequence, will be better-equipped to avoid the *ex post* inefficiencies (and horrors) of war when dealing with like minds.

World politics scholars have developed progressively more nuanced presentations of the above logic in recent years. The prospect of conducting foreign policy in the public eye—what Aldrich, Sullivan and Borgida (1989) prosaically described as “waltzing in front of an audience”—has led researchers to conclude that regime-dependent international behavior might be explained by leaders’ anticipation of political costs. Fearon (1994*a*), perhaps most famously, has argued that the conduct of diplomacy in the public eye generates “audience costs” that will penalize elected leaders who engage in foreign adventurism only to have to abandon a bluff. In this analysis, retrospective voters offer a specter of future electoral punishment that then ought to keep their leaders acting in good faith with respect to legitimate national interests. Other formulations (e.g. Bueno de Mesquita et al. 2003; Schultz 1998) have identified certain institutional features of liberal regimes that imply pecuniary incentives to worry about electoral backlash over poor policy performance, which supplements Fearon’s original, and somewhat nebulous, conception of audience costs, which imagined that ordinary voters would be upset at tarnishing the “national honor” (Fearon 1994*a*, p.581).

But, crucially, *voters do not evaluate policies, or leaders, in a vacuum*. Voters may be “pretty prudent” on average (Jentleson 1992), but the process by which individuals—voters—translate whatever beliefs they have about the political world around them into opinions that they can express at election time is complex and subject to all sorts of outside influences. This process depends not just on voters’ own attitudes, but also upon the balance of messages and appeals they receive from friends, family, and, inevitably in the modern period, the mass media. At best, citizens rely on some set of off-the-cuff thoughts, or “considerations” (Zaller 1992) in formulating responses to the issues of the day, whether in a public poll or on a ballot, and a wealth amount of research shows that political elites enjoy a great deal of influence over what those considerations are.



In this dissertation, I attempt to establish a research agenda at the crossroads of two distinct literatures: the study of domestic regimes and war, which relies on voting rules and electoral punishment to gain theoretical traction; and the study of voting and political communication, which suggests that the sort of electoral punishment assumed by international relations scholars is in fact difficult to obtain. The challenge for retrospective voting in this context is even more severe when we consider a few stylized facts about foreign policy² as an issue area.

First, foreign policy is a supremely technical field. There are relatively few confirmatory signals that voters can use to judge policy performance or their own welfare in light of such, at least compared to more common and usually more relevant issues like unemployment and other broad economic conditions. Even straightforward notions

²Here, and likely not for the last time, I use “foreign policy” to refer to “defense policy” or “security behavior.” This usage mirrors what classical international relations scholars and military historians called “high politics,” or what Wolfers (1960) described as the “pole of power.” While foreign policy can describe any of a state’s interactions with other governments, whether those issues are characterized by conflict or cooperation, this dissertation is solely concerned with security issues, or with that type of interaction in which the use of coercive violence is a plausible solution.

of victory and defeat are open to considerable interpretation, which has produced, in the American experience in particular, striking disconnects between victory on the field and victory in political discourse (Johnson and Tierney 2006).

Second, an incumbent government enjoys a distinct information asymmetry relative to voters or the private media in matters of war, and in the short-term this advantage is nearly absolute. Norms against public criticism often prevail, such that the government's preferred frame of an event will dominate against any challengers (Bennett, Lawrence and Livingston 2006). Given that narrative frames exert a powerful effect on political beliefs, our assumption that media and voters can exert their watchdog role may be on shaky footing.

Third and finally, foreign policy is conducted largely in the realm of rational expectations. In other words, a successful crisis negotiation is one that is not, in the language of the bargaining literature, resisted by the defending state. The most effective threats are those that change someone's behavior without having to follow through on them. By contrast, the most politically salient international disputes³ are those that have "gone wrong" in some sense, in that a threat was made and dismissed, eventually producing a conflict whose outcome might be less attractive to both the victor and the defeated than a settlement may have been.⁴

These points in mind, I argue that not only *can* democratic executives influence public opinion about foreign policy under some circumstances, but we should *expect leaders to attempt to exert this influence* given the bargaining framework that characterizes international crises. As a state actor's ability to claim with credibility that they will follow through on a threat grows, so too does the probability that their

³Which, problematically, are the ones on which our understanding of mass political opinion's reaction to foreign policy are based.

⁴Here I focus on just one set of rationalist explanations for war, in which incomplete information and actors' incentives to misrepresent their true reservation prices prevents them from finding peaceful outcomes that are preferable to war to all parties. The inability to commit to the terms of such an agreement is another powerful, and perhaps more important, explanation (Powell 2006) but for the purposes of this project I focus on initial positions in crisis bargaining and associated short term motives.

threat will be successful. This stems from one of the conventional explanations for war, which is that fighting is produced by the parties' inability to communicate their true level of "resolve," or willingness to fight, to one another. A voting public that accepts the need to fight is *prima facie* evidence of a resolute (democratic) actor, everything else being equal. I refer to any sort of public diplomacy meant to stoke the public's appetite for war as "information market intervention."⁵ For the purposes of this project, information market intervention refers to public diplomacy within democracies. Illiberal or autocratic regimes do of course, under many circumstances, face their own incentives to build their own cases in the courts of public opinion,⁶ but given that nearly all prior work on popular opinion and foreign policy has been developed in relation to democratic institutions, I restrict my argument at present to those same circumstances.

Second, I point out that this sort of intervention must necessarily occur *before* war occurs. To the extent that existing research has seriously considered the role of misinformation among voters in shaping attitudes toward war, it has done so with respect to wars and not crises. A number of studies, for example, have considered whether accurate estimations of casualty figures during the Iraq War make any difference for voters' estimations of President George W. Bush or of the war effort more generally (Gelpi, Feaver and Reifler 2006; Gartner 2008; Berinsky 2009). However, if we take seriously the notion that war is the product of a failed bargaining process (Schelling 1960),⁷ then it must be the case that any attempts to stake out a better

⁵The term "information market" implies the free provision of thoughts and ideas among a polity such that the best rise to the top; the term was most famously used by Milton in the *Areopagitica* (1644), who described a "marketplace of ideas" as essential for liberal government. The term "intervention" reflects that the equilibrium opinion expressed by that market is changed by government appeals.

⁶The classic example in this regard is the Soviet development of the "combat task" at the core of the international Communist movement, which was meant to legitimize totalitarian government and a long-term focus on military production at the expense of consumer production. More recent examples note that a variety of autocratic contexts exist wherein cadre elites can impose costs on a single recognized leader for policy failure (Weeks 2008, e.g).

⁷There is considerable disagreement in the field on the extent to which war and peace ought to be thought of as exclusive states. Most canonical work presents war as the ultimate "outside

position in crisis bargaining must be made before that crisis unfolds. Misinformation may endure after war begins—and the persistence of such misinformation and the effect of such on later voting decisions is clearly pertinent—but is logically distinct from elite-driven misinformation on the road to war. When information market intervention is successful in allowing a democratic challenger to command a favorable deal at the bargaining table, we may in fact not even observe that a conflict happened at all by the coding standards of the available political data.

Although my theoretical argument applies to any liberal regime, both of the empirical applications in Chapters 4 and 5 are taken from the American post-War record. This was done both to be more conversant with established literature concerning public opinion and warfighting, most of which has been motivated by American behavior during this time period, as well as data availability. Chapter 5, for example, examines stock market data for defense firms, and the United States is a large arms producer; additionally, stock prices are only available in most cases from the mid-1950's onward, and in some cases from 1973 onward.

1.1 Sketch of Research Design(s)

Any empirical effort, then, must satisfy both of these insights: that democracies are in fact able to persuade the public under at least some circumstances to support war efforts that they would *not* have supported otherwise, and that this persuasion occurs *prior* to the onset of a recognized crisis. A direct test of my theory would

option”: unlike bargaining as a process of economic exchange (Rubinstein 1982), either party can resort to the use of force to seize an outcome that is not in the offing at the negotiation table. Here, consensual outcomes are the “inside option,” and indeed in a formal setting the accepted agreement is presumably an interior solution given either party’s objective function. War, by contrast, can lead to an outcome that at least one party would not have agreed to within the bargaining protocol. Some more recent work in the international conflict literature has addressed the empirical point that there are many smaller, limited conflicts that do not satisfy classical definitions of a war by representing conflict itself as an inside option in which something short of a full-fledged war is used to tease out an adversary’s level of resolve in a fashion similar to traditional bargaining (Wagner 2000; Slantchev 2003). In this formulation, conflicts are drawn from some distribution where major, overt wars between states are quite rare but limited conflicts are quite common.

show that intervention, when used, produces a higher ex ante probability of winning the dispute at the crisis stage, short of fighting.

Finding empirical evidence to this point, however, would require several components: a measure of the amount, or even the presence or absence of intervention prior to a particular dispute, and a measure of an intervention's success or failure in changing the public's attitude toward a potential conflict relative to some unobservable baseline. In order to capture the domestic constraints an executive faces in deciding whether or not to intervene in the information market prior to an international dispute, we would also have to measure the strength of media institutions: their ability to uncover the presence of an intervention. The amount of political backlash inflicted on the executive for a revealed intervention might also be at stake. Even measuring the degree of public support for an executive's foreign policy performance in the abstract is difficult to do with any precision. As part of the literature review in Chapter 2, I examine the "rally 'round the flag" research agenda and discuss several under-appreciated problems with measurement and sampling error.

On a more fundamental level, operationalizing what "intervention" refers to, in practical terms, is difficult. Depending on the circumstances of a particular dispute, an effort to produce additional public support in anticipation of a later war might be an overt campaign to demonize the enemy and persuade voters of the need for immediate action, as was the case in early 2003 before the American invasion of Iraq. Intervention might also be obfuscatory, however, hiding an unpopular dispute from public view while the crisis negotiation unfolds. These behaviors would appear to be quite different in the historical record even if they serve the same *functional* purpose of ensuring that short-term public opinion is more favorable than it would be otherwise, leaving aside for the moment the associated intractability of estimating a counterfactual level of support against which to compare them.

I intend to overcome at least some of these obstacles by employing a multiple

methods approach. First, I present a formal model that expands and codifies the theoretical logic I proposed above. Comparative statics in this model suggest that the option of intervention, whether or not it is actually used, increases the range of realized preferences for which war results. Rival states are driven to respond in force to a wider range of challenges from a democratic state, precisely to protect themselves against weak, high-cost, “irresolute” opponents taking advantage of them.

From there, I present two novel empirical tests. The first is a survey experiment, in which I use manipulations to a print news article to test what sort of appeals are the most effective in building public support for conflict. In a neat bit of irony, the experimental treatments are closely based on a real-life propaganda effort developed by the American National Security Agency in the mid-1980’s. The objective of this study is to examine whether individual perceptions of the case for war react in response to elite appeals. By relying on a real historical case of an elite appeal meant to build public support for conflict, I address the common external-validity criticism leveled at many experimental studies of political beliefs and behavior. This particular case also allows me to test two prominent factors associated with changes in political beliefs in response to persuasion: the role of expertise and credibility, and the role of emotional appeals and anxiety. While I find limited evidence of either persuasive effect, I find strong indications that the widely-accepted role of partisanship in predicting foreign policy opinion only obtains when voters are asked to react to general, and retrospective, policy performance. Questions about support for *potential* conflicts do not yield any relationship between party identification and stated opinion at all. But since the logic of crisis diplomacy is inherently prospective, voters’ support for future conflicts is in fact what’s most important for executives.

The second empirical chapter explores which American foreign policy crises in the superpower era were most likely to have been associated with public campaigns meant to build additional public support. Because of a variety of measurement problems en-

demic to the use of public opinion polling data on foreign policy as it has traditionally been used by international relations and foreign policy scholars, I examine American financial markets' response to the onset of international disputes. I argue that when a government has engaged in information market intervention, that their domestic market for defense equities will have observed that campaign and inferred—correctly, given the results of my formal model—that there is an increased likelihood of conflict from having done so. In such cases, investors should expect greater future dividends to those firms' stock due to this risk of war.

Under these conditions, the trigger or onset of an official dispute should not provide any new information, because investors will have already observed a public affairs campaign that, per the formal model I advance in Chapter 3, has produced a more war-prone political environment. By contrast, in the absence of intervention—whether because of a more resilient media environment that is more likely to resist intervention, or because the balance of power is tilted enough in one direction or another that the additional resolve produced by a potential intervention isn't needed during the crisis phase—then the onset of a dispute will come as a surprise. That surprise will produce an *abnormal return* associated with a dispute's official onset, and the magnitude of this surprise can be estimated in an event study framework.

1.1.1 Multiple Concepts & Multiple Methods

The relationship between the various chapters described above merits one additional remark. Popular-accountability models of crisis behavior involve at least two strategic relationships. First is the relationship between voters and their government. Individual voters observe the state of the world and the success or failure of their government's policies during times of crises, and update their beliefs about their government accordingly. When the government performs well, then any individual voter ought to become more supportive of the incumbent government from that moment

forward, *ceteris paribus*. Knowing this, the government must act in ways that are likely to earn approval from their electorate.

The second relationship, of course, is the confrontation between any two government actors who are embroiled in a crisis. Here, the strategic logic is well-understood and frequently explored. Any state, in confronting a democracy abroad, must understand that the democratic executive's behavior will be tempered by his or her own expectations about their voters' likely reaction, and that will shape the first state's strategy. The democratic state, knowing that its rival *knows* that it is a democracy, and knowing that its own strategic motives are known, will in turn incorporate all of *those* understandings in its own strategy, and on and on in keeping with the recursive nature of equilibrium behavior.

As we will see, the choice of empirical strategies here was motivated in large part by the shortcomings that characterize existing studies of democratic crisis behavior and public approval. I describe and present those shortcomings in Chapter 2. However, these approaches also map themselves neatly onto different units of inquiry, or what was originally called the "levels of analysis" problem in world politics (Singer 1961). The strategic logic between states at the highest diplomatic level is addressed through a formal model. From there, we will turn our attention to the individual voter, and examine how voter perceptions of a potential conflict change in response to appeals meant to convince them that war is necessary. Finally, the study of equity markets in Chapter 5 is ultimately a study of aggregate public opinion by an unusual means. Seen in this way, the dissertation examines the behavior of both individuals, the large group of individuals that comprises the American public, and then the influence that the mass public has on foreign relations through the potential electoral cost it can impose upon the sitting government.

1.2 Outline of Remaining Sections

In Chapter 2, I review a few different literatures that are relevant to understanding the interplay between public opinion, elite messaging, and foreign policy. I then offer a descriptive theory of information market intervention as well as a replication exercise demonstrating the shortcomings of observational data for addressing this theory. In Chapter 3, I formalize my argument with a game-theoretic presentation of bargaining with endogenous public opinion. The extensive form of the game I describe is a refinement to what Slantchev and Tarar (2011) call the “standard” immediate crisis bargaining game (Morrow 1985; Fearon 1994*a*), in which a dissatisfied democratic state has the opportunity to play a second lottery to pursue information market intervention that will produce additional resolve relative to its adversary. This generates several testable hypotheses, albeit ones which are subject to the limitations I described above. I conclude Chapter 3 by discussing some of the challenges of measurement and operationalization that make a direct empirical estimation of the game’s equilibrium conditions so difficult. In Chapters 4 and 5, I turn to two alternate, but complementary, empirical strategies that are able to test second-order predictions of my theory. The event-study framework in Chapter 5 uses stock market data to determine when information market intervention is likely to have preceded security disputes from the American record. Chapter 4 uses a survey experiment to determine whether intervention, as I describe it here, can actually prove effective in changing individual-level opinion about the desirability of a potential conflict and an incumbent president’s foreign policy success and competence. Chapter 6 concludes and, in keeping with tradition, speculates on possible extensions or subsequent installments of this research.

CHAPTER 2

Related Literature, Descriptive Theory, and Empirical Challenges

“When I came back from Viet Nam, I had just had the greatest brainwashing that anybody can get when you go over to Viet Nam. Not only by the generals but also by the diplomatic corps over there, and they do a very thorough job, and, since returning from Viet Nam, I’ve gone into the history of Viet Nam, all the way back into World War II and before that. And, as a result, I have changed my mind.”

Gov. George Romney, September 1967

In Chapter 1, I began by describing the ongoing campaign by Franklin Roosevelt to persuade the American public of the necessity and wisdom of opposing German military expansion in Europe. The considerable disconnect between Roosevelt’s rhetoric—the “International Nazi Church”—and the legal status of the United States *vis-a-vis* the war in Europe—neutral, per three separate acts of Congress—paired with the Navy’s *de facto* entry into the conflict in 1939 suggests an overt attempt to change public opinion in the expectation of a future conflict. In this chapter, I will demonstrate that this sort of campaign is not only possible, but indeed inevitable given the incentives that democratic governments face with respect to crisis diplomacy. These incentives are pieced together from separate research agendas in conflict theory, the empirical foreign policy record, mass public opinion, and the behavioral

underpinnings of voting behavior. I begin by briefly sketching another, and even more dramatic, example of what I call information market intervention, this time taken from a very recent conflict event involving the United States. I then move into separate literature reviews that span both international relations and American politics, before identifying two distinct shortcomings of extant research on foreign policy and public consent: one theoretical, and one empirical. This discussion is meant as a springboard to the formal model presented in Chapter 3.

2.1 Information Market Intervention in Motion

General William L. Nash, formerly of the United States Army, found the military briefings he and a group of other retired officers received during a 2003 Iraq tour so implausible that “he joked to another group member that they were on ‘the George Romney memorial trip to Iraq’ (Barstow 2008).” Romney famously sabotaged his own presidential aspirations in 1967 by volunteering to an interviewer during a televised appearance in Detroit that his previous support for the Vietnam War came as the result of brainwashing during a tour of Saigon years earlier.

Nash was more correct than he may have realized. As a lengthy *New York Times* investigation later revealed, tours like the one Nash described were part of an overt campaign to marshal American public opinion behind the military occupations in Afghanistan and Iraq. Retired officers like Nash were at the center of this strategy. Since many of them were already paid contributors to various news programs, they were the ideal means by which the Department of Defense could attempt to influence the narrative of media coverage of their military efforts, and by extension the public’s support for the continued use of force.¹ The bureaucrats responsible for administering

¹Political scientists may have reached very different conclusions about the likely success or failure of this campaign given that the public audience for the sort of news programs these officers frequented is a small and idiosyncratic slice of the wider electorate. I visit the professional literature regarding this question later in the chapter.

the program supplied these officers with specific talking points to raise during their media appearances and secured a marketing firm to track how many appearances each analyst made and the degree to which their language matched the Pentagon's preferred line.

From the perspective of the Defense Department, the analyst program was a “force multiplier,” in which an investment of modest means might produce an extremely valuable end: keeping troops in the field in both Iraq and Afghanistan. The subtext behind the program was military elites' collective disillusionment with the erosion of public support for the Vietnam War.² From the perspective of the analysts themselves, the program was extremely financially lucrative. Most officers were already paid small honoraria for their appearances, while some were also active in the Washington policy establishment that churns out column inches for the *New York Times*, *Wall Street Journal*, and *Foreign Affairs*. The access that the analyst program required, however, was worth more than its weight in gold. Analysts were frequently invited for conference calls and meetings with Secretary Donald Rumsfeld, which made them valuable consultants for the defense contracting industry. A number of former officers developed thriving businesses in advising firms on the pursuit of Pentagon contracts, which could only happen with their newfound status as part of Rumsfeld's circle.

2.2 “MindWar” in a Political Science Framework

Much of what Barstow describes as the “Pentagon's Hidden Hand” amounts to an unwitting folk replication of established findings in political communication and

²Subtext for some, and text for others. Army General Paul Vallely, one of the most enthusiastic participants in the program judging from the emails that were received by the *Times* as part of its Freedom of Information Act request, had spent much of his career after 1975 developing a strategic concept of “MindWar.” MindWar was meant to direct psychological operations and foreign and domestic audiences simultaneously in order to “strengthen our national will to victory” (Barstow 2008).

media effects. “MindWar,” or what Torie Clarke, the Assistant Secretary of Defense for Public Affairs, called “information dominance” (Barstow 2008) amounts to an exercise in narrative framing. Clarke, Valley, and others all intuitively understood that public consent to the war effort might be affected if only they could ensure that their preferred frame came to be the one that dominated coverage.

Viewed on its own, the notion that political actors would work to secure public support for their political aims seems uncontroversial. Rumsfeld’s “long war” was not the first and will certainly not be the last episode in which elites attempt to secure public support, or at least acquiescence, to some public policy. However, Clarke’s “information dominance” program began, by all indications, *before* either the Iraq or Afghanistan campaigns did. In fact, Clarke began recruiting likely candidates for the program shortly after George W. Bush took office, and well before the September 11th attacks (Barstow 2008). This suggests that organized, purposeful attempts to manage the public debate are not limited to those conflicts which have already begun, such as in Iraq, Afghanistan, or in earlier years Vietnam.

In this chapter, I explore two different bodies of academic literature: international relations research on the relationship between regimes and war, and American politics research on public opinion and voting. These reviews will expand on some of the broad outlines I offered in Chapter 1, which identified two separate sources of tension between these two literatures. I also work through an example of the empirical dilemma regarding large- n measures of public support for foreign policy performance through the lens of the “rally-’round-the-flag” tradition. I show that these earlier findings ought to be viewed with some degree of suspicion because of endemic issues of measurement and sampling error for any use of public opinion data to indicate proximate public support for a foreign policy venture. Finally, I provide a verbal formulation of a theory of information market intervention. This latter discussion continues in Chapter 3, in which this argument is formalized and turned into a series

of comparative-statics assertions about the role that endogenous public opinion can be expected to play in security behavior.

2.3 The Relationships between Regime, Crisis Bargaining, and War

Dating from Wilson and the creation of the League of Nations, through the Truman and Eisenhower Doctrines and beyond, the United States has maintained that democracy abroad contributes to security at home. This position found a particularly transparent endorsement during the presidency of Bill Clinton, who asserted in his 1994 State of the Union address that “ultimately, the best strategy to ensure our security and to build a durable peace is to support the advance of democracy elsewhere. Democracies don’t attack each other” (Clinton 1994). Then-Undersecretary of State, Strobe Talbott, later wrote in *Foreign Affairs* that “the larger and more close-knit the community of nations that choose democratic forms of government, the safer and more prosperous Americans will be, since democracies are demonstrably more likely to maintain their international commitments, less likely to engage in terrorism or wreak environmental damage, and less likely to make war on one another” (Talbott 1996).

This belief, inspired by research in world politics, offers at least a partial explanation for American interventions in Haiti and Kosovo, among other engagements. The relationship between democracy and peace also proved to be among the final justifications referenced by George W. Bush for the 2003 invasion of Iraq. This connection between scholarly research and its implicit policy prescription was made plain by Bush immediately before the 2004 presidential election. Bush argued in a press conference that “the reason why I’m so strong on democracy is democracies don’t go to war with each other. And the reason why is the people of most societies don’t like

war, and they understand what war means” (Bush 2004).

Of course, the balance of evidence suggests that Bush overstated the case in claiming that people in democratic societies simply do not care for conflict and violence. A number of studies have found that, as a group, democracies find themselves embroiled in crises and wars just as often as any other type of regime.³ And, even if established democracies were in fact pacifist, the transition to liberal government might itself lead to additional violence under certain conditions (Gleditsch and Ward 2000). In any event, current evidence suggests that the democratic peace is a conditional peace: pairs of democracies seem to be able to avoid conflict between themselves, but this logic does not seem to extend to other regimes’ interactions.

The earliest findings on the relationship between domestic regimes and war were driven by inductive observation rather than theory (Babst 1972; Rummel 1975-1981; Small and Singer 1976). Modern theoretical work on the democratic peace can be traced to Doyle (1986), who asserted that the Immanuel Kant’s separate republican peace (Kant 1970*a,b*) was the logical wellspring of then-President Reagan’s goal of spreading democracy abroad. Although there are multiple explanations for the democratic peace in the current literature, the most powerful link domestic political institutions with the realization of elite preferences. Fearon (1994*b*) has argued that democratic leaders’ desire to stay in office means that they must avoid backing down on public commitments. These leaders’ decisions to escalate by means of direct military action or the issuing of public appeals—tempered by their incentive to defend their reliability and reputation—sends a more accurate signal of resolve than similar acts by autocratic leaders, because of the potential fallout they face for failure or dishonesty.⁴ By this logic, democratic states should have greater bargaining power in

³Alexander Hamilton noted in *Federalist Number 6* that “Few nations, nevertheless, have been more frequently engaged in war” than parliamentary Great Britain, and that “the wars in which that kingdom has engaged have, in numerous instances, proceeded from the people.”

⁴This broad formulation covers three separate empirical considerations, all of which have been the subject of additional scrutiny. Are democratic voters actually able to impose audience costs on their leaders (Tomz 2007)? Are democracies truly able to generate larger audience costs than other

escalating crises and be more likely to settle them on favorable terms, even if they are at a disadvantage in relatively minor disputes in which they cannot credibly claim a large investment.

Bueno de Mesquita and Lalman (1992) and Bueno de Mesquita et al. (2003) also examine the domestic political incentives of leaders who, at the mercy of the reputation game described by Fearon, assign differing values to the procurement of public versus private goods based on selection institutions. Bueno de Mesquita et al. observe that democratic leaders must earn the support of a greater number of people in order to maintain office. Autocratic leaders need only satisfy a small minority; generally a coalition of military and business elites as well as party functionaries. Dictators can earn support by bribing their supporters with spoils or private goods, but this approach cannot work for a democratic leader because of greater scale involved. These regime-specific incentives generate potential costs for leaders that depend on their respective regimes, and push them to act in much the same manner as Fearon (1994*a*) even though the sources of those costs are completely different. When confronted with war, democratic leaders will try harder, relatively speaking, to win because of the risk that war poses for public goods allocation in addition to their reputation for making credible commitments. Implicit to this understanding of democratic crisis behavior is the idea that threats to any of a set of core principles (e.g., core public goods) (Bueno de Mesquita et al. 2003) will provoke leadership into devoting relatively more resources (and in effect fight harder) to achieve victory. In the selectorate model of electoral competition, victory also produces a surplus that can be allocated to potential supporters in the next election. The public goods dynamic enters play here as well; not only do democratic leaders have a strong incentive to not suffer any loss in their public goods allotment, but victory that produces additional public goods will aid in recruiting supporters as well.

regimes (Weeks 2008)? Are democracies in fact better at making coercive threats than other regimes (Downes and Sechser 2012)?

Schultz (2001) also leverages Fearon's audience cost metaphor but focuses on the role of political competition. Rather than address the effects that democracy might have on leaders' preferences, he notes that in open, competitive electoral systems, a leader has not only to contend with a would-be rival abroad but also their own domestic political competition. Schultz's primary contention is that a self-interested opposition party functions as an additional signal to the government's position since the opposition presumably knows more about the government's preferences than does the foreign rival. As such, when the government does make a challenge, the existence of the opposition party allows the rival to better estimate the government's willingness to follow through on that threat independent of the government's existing credibility from its domestic audience. Schultz illustrates how certain internal processes, which are unique to liberal democracies, impact the strategic behavior of these states and their international rivals on the world stage. He also makes a subtle distinction about why conflicts involving at least one democracy will lead to different patterns in bargaining—rather than relying on the blanket incentive for success and prudence that are the product of democratic audience costs, Schultz essentially argues that opposition parties force the government to hew closer to its true reservation point in the bargain, which makes speculative challenges somewhat more moderate and considerably more rare.

These approaches all make a few key assumptions about voting. First, and most importantly, they assume that voters are able to evaluate foreign policy performance—wins versus losses—with some degree of precision. Their voting decisions are clearly endogenous in all models, as they must reflect the realized outcomes of international bargains; however, the weighting for policy performance in their decision, or their understanding of the degree to which their leadership has succeeded or failed, cannot also be at issue. A related assumption is that if democratic leaders possess divergent preferences from the electorate, they do not pursue them contrary to the national

interest. Bueno de Mesquita et al. (2003) allude to such a possibility by noting that the leaders of liberal regimes, with large required winning coalitions to retain power and large selectorates from which to assemble that coalition, may overemphasize the public goods nature of an espoused policy while downplaying its potentially more prominent private-goods character. This is an initial step toward acknowledging the tension that may exist between a democratic leader and their public, but it does not go so far as to explicitly evaluate the likely behavior given these differences.

These explanations, if accepted to be true, also advance a theory of democratic decision-making and crisis behavior. Elections that are free, in which all eligible voters are able to pass judgment on the government's performance, and that are fair, in which multiple parties contend for political power, result in deliberative logic that is qualitatively different for different types of polities even if their results are (as they must be) similar in the international arena. The earlier reference to the generally-accepted finding that democracies are, in expectation, as war-prone as non-democracies⁵ is just one example. Voters must do "good enough" at the ballot box for the microfoundations of these theoretical arguments to hold.

2.4 Information and Voting

Research in American public opinion suggests that this is unlikely. Indeed, the history of research in public opinion and voting behavior reflects a long-term cynicism about voter competence that has only begun to be rolled back very recently. In their pioneering *The American Voter*, Campbell et al. (1960) suggested that the American public was rarely informed about the issues of the day and could not make rational connections between issues, their own positions, and those of the two major political parties. This consensus was eventually challenged, as some of the "response

⁵Originally this was identified by Small and Singer (1976); explaining this regularity was a key motivation for Maoz and Russett (1993) and Bueno de Mesquita et al. (2003).

instability” that researchers found in the earliest waves of the American National Elections Study was attributed to problems with the survey instrument itself, rather than dramatic and inchoate swings in voter attitudes. Page and Shapiro (1992) argued, for example, that while American voters are collectively under-informed, their deliberations are sequentially rational given available information. The distinction is subtle but vital: whether voter incompetence is driven by shortcomings in their information environment, or by inflexible cognitive limitations, has very different implications for democratic governance.

Zaller (1992) complements and extends this line of thought by noting two ways in which the informational environment might effect measures of public opinion. The first is through survey questions themselves, in which different formulations might provoke a given respondent to think about any political issue in one of a multitude of different lights, or what Zaller called “considerations.” By privileging one consideration over another, different polls might reach different conclusions about the public’s attitude toward the same core issue. Second, the breadth of public discourse might, in its own way, privilege certain considerations over others and in so doing lead to one sort of position or another becoming dominant across the breadth of the public sphere. Either dynamic reflects the same basic model of political behavior, which Zaller called the “Receive-Accept-Sample,” or RAS, model. In his formulation, some number of citizens—voters—are inherently predisposed to engage with any particular issue domain. Of those voters, some fraction are disposed to listen to, or accept, any particular message from political elites and entrepreneurs; respondents will tend to process those signals that conform to what they understand to be their own core political beliefs, and they will tend to reject those that reflect opposition beliefs. And, finally, exposure to messages and cognitive apprehension of same will, over time, make it more likely that an individual will rely on one consideration or another in answering a survey question—and, presumably, the same logic obtains during election season.

The “rational public” framework and Zaller’s RAS model both reflect an electorate that does the best it can to come to terms with the issues of the day but one which is at the mercy of whatever messages they happen to encounter. To the extent that voters hear and process political messages, the subjective rationality at the heart of the mechanisms identified by Page and Shapiro, Zaller and others amounts to an argument about *framing*: information about politics is not important only through its presence, but also in the way that it is packaged, labeled, and presented. Framing is “one of the most important concepts in the study of public opinion” (e.g. Druckman 2001), and differences in how an issue or topic is framed can lead to dramatic shifts in how the mass public evaluates an issue. There is a wealth of experimental evidence from behavioral politics on this latter note; a series of studies have found, for example, that respondents’ expressions of permissiveness toward a hypothetical Ku Klux Klan rally depend on the extent to which a media description of the event frames the rally in terms of norms of free speech, or in terms of incipient danger to public safety (e.g. Druckman 2001; Chong 1993). Even an issue with some of the most immediate and visceral attachments possible can find either support or opposition through distinctions in framing.

Seen in this way, the expressed positions of even a rational public are inherently elite-driven. To the extent that any political appeals reach voters, the balance of frames in the broader informational environment surrounding voters ought to predict the balance of their positions. The content of frames themselves is driven by a predominantly elite-centered contest between government officials, opinion leaders, journalists, and other opinion leaders. Voters are not blank slates, open to just any suggestion that an elite happens to offer, but clearly all of us are dependent on outside expertise in the process of forming our opinions at least some of the time (Druckman and Parkin 2008).⁶

⁶Before adopting or amending a position in light of some received message, an individual must both hear the message, and judge it to be credible and informed. The first two stages of Zaller’s

2.5 Voting on Foreign Policy

Other researchers have explored the overlap between voting and foreign policy, particularly in the American context. In many ways, America's wars have helped define the study of public opinion. Some of the earliest methodological advances in survey analysis were made by Hadley Cantril's Office of Public Opinion Research at the behest of the Roosevelt administration during World War II. Roosevelt relied on OPOR polls to determine not only the conduct of American grand strategy during the war—the American war planners frequently deferred to Roosevelt's reluctance to build troop commitments in the North African and Italian theaters beyond what public opinion could be expected to bear, to the exasperation of their British allies⁷—but also the character of the post-war international order.

In the immediate post-war period, political scientists such as Gabriel Almond and Sidney Verba concluded that the American public was fickle, feckless and incompetent based largely on Americans' unwillingness to remain as preoccupied with foreign policy as they themselves were.⁸ These findings, of course, mirrored the early consensus among students of domestic political attitudes regarding American voters' failings (Campbell et al. 1960; Converse 1964). Among foreign policy scholars, however, voter incompetence took a much more sinister edge. Not only were voters undermining their democracy at home, but in these experts' view, their lack of engagement with the

model reflect this process of “receiving” and “accepting” some unit of political communication. I return to the importance of speaker expertise and credibility in Chapter 4.

⁷At times, the specter of American public opinion was a useful “out” within the Combined Chiefs of Staff Committee. Whenever Churchill and Alan Brooke demanded that the United States commit forces to defend India, Roosevelt's personal envoy, Harry Hopkins, blamed the American public for such a plan being completely out of the question (Roberts 2009, p. 156).

⁸Almond told the National War College in 1956 that “For persons responsible for the making of security policy these *mood* impacts have a highly irrational effect. Often public opinion is apathetic when it should be concerned, and panicky when it should be calm.” Almond had a deeply sympathetic audience in this regard; foreign policy specialists, then and now, often take a dim view of the role of public and Congressional consent in shaping foreign policy. George F. Kennan's time as deputy head of the American mission in Moscow left him with a certain curious admiration of the Soviet Union's lack of popular political constraints. In a series of 1950 lectures, he famously likened American voters to “...one of those prehistoric monsters with a body as long as this room and a brain the size of a pin” (Holsti 1996).

international stage endangered the rest of the free world as well.

Subsequent scholarship proceeded along three distinct tracks. The first was to establish that the American public did indeed pay attention to world affairs and that response instability over longer periods of time often reflected measurement error rather than genuine changes in attitudes. Many of these results were demonstrated by the same works cited in the previous section (Page and Shapiro 1992; Zaller 1992); in fact, reappraisals of polling data from the Korean and Vietnam Wars were instrumental in identifying survey instrument problems in the first place. Some modest assessments of public attention to foreign affairs simply relied on the point that large numbers of voters proved willing to cite international politics as a serious concern when given the option. Aldrich, Sullivan and Borgida (1989), for instance, found that voters correctly identified Ronald Reagan's foreign policy as more conservative than either Jimmy Carter or Walter Mondale's in the 1980 and 1984 ANES waves, and that roughly one third of respondents identified relations with the Soviet Union as the most serious issue facing America in those elections. In a survey of Western European voters, Isernia, Juhasz and Rattinger (2002) find that mass publics demonstrate some predictability in their response to major international events, rather than the wild swings suggested by Campbell et al. (1960) or Almond (1960). Knopf (1998) reaches some of the same conclusions for American voters, noting that public opinion seems to change in the face of international events, but does so in a relatively cautious manner that he labels "rational." Jentleson (1992) has argued that the public is in fact quite nuanced in how it reacts to the use of force, supporting efforts to deter revisionist states but not to replace those governments wholesale.

The second track might be described as the ability of American presidents to lead public opinion simply by engaging it. Shapiro and Jacobs (2000), for example, argue that the end of the Cold War has paradoxically given the executive more latitude in opinion-shaping, insofar as the Soviet threat allowed the public to place some

demand for action on the president, rather than the other way around. Berinsky (2004) shows that the choice by social leaders to converse with international issues can itself promote or restrain public opinion on the same. Early in the Vietnam War, dovish voters reported “don’t know” or “no opinion” regarding the conflict at high rates. As elites began to question the war openly, particularly on the left wing of the Democratic Party, these respondents began to communicate their own opposition more and more openly. This dynamic was most noticeable within the African-American community.

From the other end of the lens, Baum (2004) found that presidents act strategically in their own choices to engage their mass audience—to “go public” with a crisis—because they are mindful of the possibility of blowback should the crisis not be resolved on favorable terms, recalling Fearon’s audience cost model. When elite divisions exist, then the balance of media coverage and by proxy public opinion both follow those divides, which supports the long-held “indexing” model of media and politics (e.g. Entman 2004). However, those divides tend to appear only well into a crisis, or even after the beginning of a war (Baum and Groeling 2008*b*), by which point they will be of little significance to the government’s original decision to pursue a crisis. Significant norms against criticism of war policy work against public dissent, and these norms are strongest earliest in a crisis (Baum and Groeling 2008*a*).

A third track concerns the much stickier business of *affective*, rather than objective, information about international politics. Even liberal regimes enjoy significant short-term advantages in framing the terms of public understanding of an international crisis, above and beyond their decision to engage or avoid public interest in the first place. In some cases, as with the early Vietnam War data examined by Berinsky (2004), that advantage might be interpreted as an artifact of the Cold War, which was such an imposing frame that other interpretations could scarcely be covered by news media, let alone register with individual voters. However, evidence

suggests that framing dominance did not disappear after 1991, and that this helps explain the Bush administration’s early successes in shaping the narrative of the Iraq War (Entman 2004, Ch. 4).⁹ Bennett, Lawrence and Livingston (2006) extended this logic even further, arguing that the government’s dominance in framing limited public criticism of the Abu Gharaib prisoner scandal.

2.6 “Information Market Intervention”: a Theoretical Sketch

The major contention that these reviews suggest surrounds the question of voter competence. Citizens in democracies face difficult, non-trivial obstacles in fulfilling the watchdog role that international relations theory asks of them. The foreign policy issue domain is often more technically demanding than other areas, and its consequences are often further removed from voters’ day-to-day lives than are other potential topics. Short of a major war involving the United States, foreign policy issues enjoy much less coverage and visibility in the mass media, and by extension in the electorate’s mind, than other issues (Shapiro and Jacobs 2000). Furthermore, the most prominent means by which voters can adopt positions without the benefit of first-hand expertise, endorsements and cues from party leaders, are not always available with respect to security crises.

To see why, consider the available evidence for public backlash against foreign misadventures that follow in the wake of elite divisions. Berinsky (2004) found that the emergence of mass-level opposition to Vietnam trailed elite divisions, while Groeling and Baum (2008) and Baum and Groeling (2008*b*) found much the same for the Iraq War, albeit in these latter cases with particular attention to elite splits *within* the presidential party. The other commonly-cited mechanism by which voters might learn about policy failure and punish their leaders—wartime casualties—has also been

⁹It is worth pointing out here that the framing success described by Entman overlays the Pentagon’s analyst program, which I described at the beginning of the chapter.

developed through the record of American public opinion in three major wars: Iraq and Vietnam, as well as the Korean War¹⁰ (Mueller 1973*a*; Gelpi, Feaver and Reifler 2006).

However, most crises do not produce wars. Per the crisis bargaining literature described above, crises begin when some actor in the international system attempts to coerce another actor into changing the status quo. These demands are commonly referred to as “revisionist challenges.” Some subset of those challenges become overt wars. Of course, defining a bright line between crisis and war is difficult. The standard operational definition of a war is an exchange between the armed forces of recognized governments that results in at least 1,000 battle deaths (Jones, Bremer and Singer 1996). Crises are defined in somewhat more nebulous terms.¹¹

Between 1946 and 2003 (the last year for which the Correlates of War project has been updated), the United States participated in 218 militarized disputes but just 7 wars. For the entire international system, the comparable counts were 1,527 disputes and 38, respectively (Ghosn, Palmer and Bremer 2004). Even though most disputes do not produce wars, the bulk of our evidence regarding the electoral accountability mechanism for democratic crisis behavior is taken from large-scale conflicts. In essence, much if not most of the evidence derived by Americanist and foreign policy scholars has tested a mechanism that ought to obtain in all international interactions by examining evidence from a tiny sliver—and an unrepresentative sliver, at that—of all such interactions.

The timing of our measures of public opinion is also problematic. In what the bargaining literature has come to call the “standard” immediate deterrence model

¹⁰Though newer evidence strongly suggests that even casualty figures are simply lightning rods for existing partisan divides; see Berinsky (2009).

¹¹The usual standard comes from COW’s Militarized Interstate Dispute project, in which disputes must include an “explicit threat, display, or use of force” between two members of what COW recognizes as the international system. This standard over-counts some displays of force, like those between coast guards and illegal fishing operations, while under-counting other types of implicit threats; on this latter note, see Fordham and Sarver (2001) and Downes and Sechser (2012).

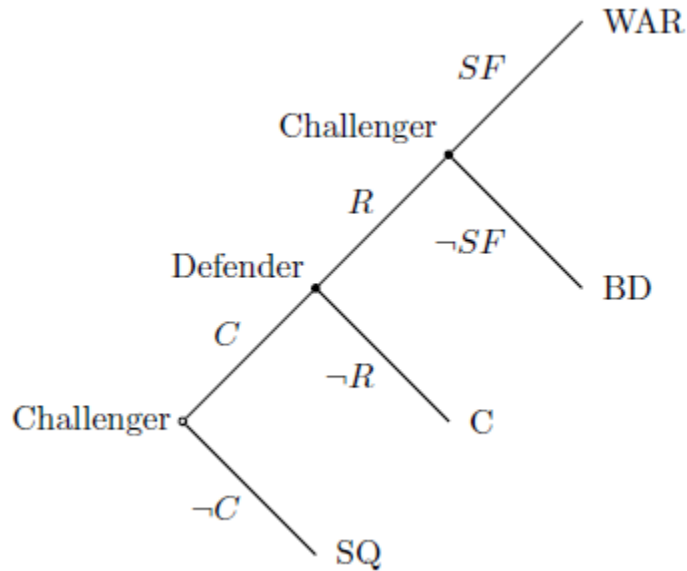


Figure 2.1: The extensive form of the immediate deterrence game. Strategies are labeled in *italics* along each edge (*Challenge*, *Resist*, *StandFirm*), with outcomes (Status Quo, Defender Concedes, Challenger Backs Down, War) at the end of each branch as appropriate.

(Morrow 1985; Fearon 1994*a*), successful deterrence for either the challenging or defending state amounts to making a credible display of their willingness to fight, or “resolve.” If the actor *receiving* the challenge is certain that the revisionist state is truly willing and able to make good on their threat, then they will be more likely to concede the issue at hand without a fight, *ceteris paribus*. The same logic obtains from the perspective of the revisionist state; if the defender state is certain to resist any challenge, then the challenger state is more likely to cave in after their initial probe, or even to refrain from making any challenge at all.

The more willing someone is to fight, the more likely they are to get their way without *having* to fight, whether in a high school lunchroom or in international politics. Irrespective of a state actor’s domestic regime, the willingness to fight determines the revisionist state’s decision at the last opportunity they have to make a decision in the simple game above. After they have attempted to extract some concession from a rival, only to have their challenge resisted, are they in fact willing to resort to violence

to seize their objective by force? If they are, and the rival state knows this as well, then the ultimate decision switches hands: it is up to the defending state to decide whether they prefer to fight than concede, in keeping with the logic of backwards induction (Schelling 1960, p. 124).

In any regime, resolve hinges in part on a value judgment of the executive as well as the governed—are the benefits from fighting worth the cost? In democracies, of course, the relevant base of support are the voters. If voters think that the issue under dispute in a crisis is in fact worth fighting over, then their leaders ought to have a much stronger hand at the bargaining table, because they can adopt a more recalcitrant tone with the blessing of their constituents.¹² However, we can see from the discussion above that the question of whether a public policy is “worth it” is conditional on the messages received by the mass public from government figures and opinion leaders, and that nowhere is this more true than on issues of foreign policy and national security.

And, of course, the short term is exactly the time when this calculation matters most. If a supportive electorate gives additional bargaining power to a democratic regime facing a rival abroad, then the crisis negotiation might be resolved (presumably, on more favorable terms to that democratic state) before the regime’s informational and framing advantage expires. Marshaling public opinion is of course important during ongoing wars as well, particularly as war exhaustion begins to set in (Reiter and Stam 2002, Chap. 7), but successful threats avoid war altogether.

The following figure presents a simplified model of this interaction. The choice to engage in information market intervention is distinct from the decision to make any revisionist challenge at all, but when used it must occur before any potential war in order to have any beneficial effect on the outcome of crisis negotiations. Its

¹²This is reminiscent of the “two-level” dynamic proposed by early models of international cooperation, in which negotiators might leave ratification to their own domestic legislature so as to prevent themselves from being able to offer a better deal to their partners (e.g. Putnam 1988).

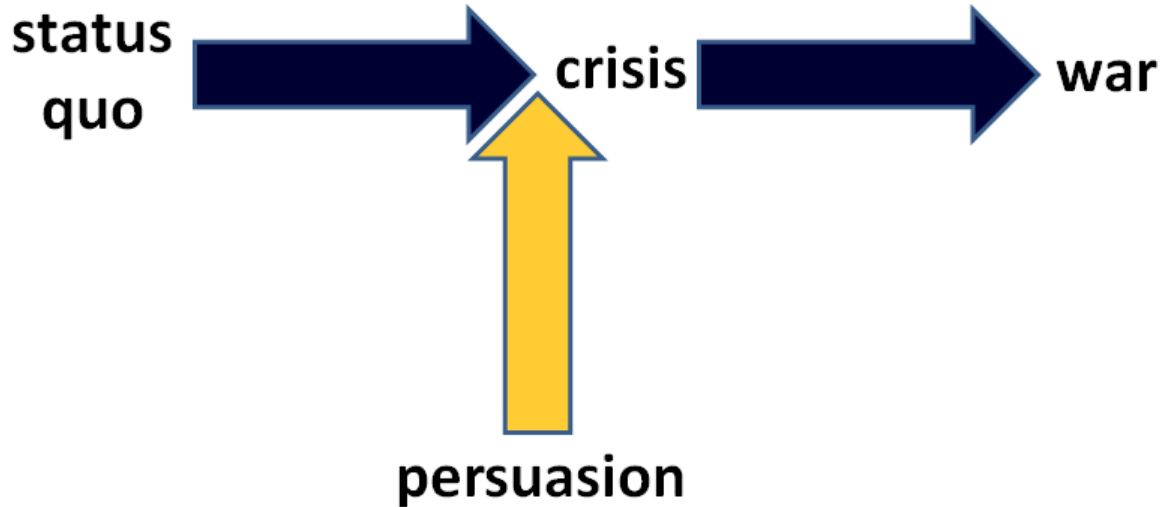


Figure 2.2: The timing of information market intervention.

substantive character may also vary from case to case, reflecting either the decision to act covertly or to accentuate the danger posed by a particular adversary.¹³

The following three chapters examine three separate implications of this relationship between leaders, voters and the public will. Chapter 3 explores the implications of this endogenous relationship between democratic executives and their constituency for crisis bargaining between states. Chapter 4 explores the effect of early crisis media coverage on support for a potential war with a survey experiment. Chapter 5 attempts to determine which crises from the American record are more likely to have been associated with information market intervention in the past with financial data. As a precursor to the two empirical chapters, I present an empirical critique of current research in public opinion and international crisis behavior in the section below.

2.7 Empirical Obstacles: The “Rally ’Round the Flag”

The theory of information market intervention claims that democratic executives can use their short-term framing dominance to build additional public support for

¹³As I argue in Chapter 3, these are functionally equivalent because either would reduce the “cost of fighting” parameter in the state actor’s utility function.

a conflict. A long list of studies have examined the American public's attitudes toward the use of force abroad, beginning with Mueller (1973*a*), who found that public approval of the president surges when troops are deployed, only to slowly decay afterwards with mounting casualty reports and growing interest in other issues of the day. Mueller described this relationship as the mass public "rallying 'round the (American) flag," which has come to be a shorthand for this entire research agenda. Rally events refer to a change in mass public opinion around an international crisis.

At first blush, then, it might be the case that observed variation in rallies could be a promising way of operationalizing my theory of information market intervention. Intervention, after all, is meant to increase political support because it strengthens one hand during a dispute, and the literature has interpreted the rally effect as a burst of public support for their government when faced with an international crisis. We might reasonably expect that intervention might produce larger rallies, or at least produce larger (or more durable) responses in the relevant audience when successful.

However, the empirical estimation of rally effects has been problematic. Most analyses match some sources of public opinion data against the crises identified by one of the major international dispute data. Even when focusing on a single country, such as the United States, polling data is not readily available that is actually tailored to any particular dispute. Some disputes are so minor that they never achieve the sort of salience necessary for inclusion in a poll; and even when they do, they are only included once the crisis has begun to unfold, which makes a proper measurement of any change in the public pre- and post-crisis impossible. Other, more nuanced questions are only asked in low-frequency panels such as the American National Election Studies. Polls of course have become much more frequent in the past few years, which allows for clever aggregations between multiple polls' time series so long as the questions asked are sufficiently similar (Jacobson 2007); however, this approach would still significantly limit the number of included crises. For these reasons, most studies

rely on presidential job approval as a workhorse dependent variable. Job approval is the only measure taken with sufficient frequency and consistency in framing such that it can be interpreted as a plausible measure of foreign policy sentiment.

There are a number of serious drawbacks to this measure, however, despite its ongoing popularity (Baum 2002; Chapman and Reiter 2004; Colaresi 2007; Oneal and Bryan 1995; Baker and Oneal 2001). The first deals with the set of conflict events to be included. The most-cited source of observations of international crises, the Militarized Interstate Dispute data maintained by the Correlates of War project, include any use of force or under certain circumstances in the implied threat of same. However, this coding standard includes actions taken on the initiative of local authorities, rather than the national executive, as well as police actions under certain circumstances (Fordham and Sarver 2001, p. 459). In the case of the United States, the cited set of MID events includes a number of disputes with Canada and illegal fishing incidents with a litany of minor Caribbean and South American states, all of which made only fleeting appearances in contemporary news media and none of which could ever have generated a rally event consistent with a theory of public engagement and reaction.¹⁴

Chapman and Reiter's 2004 paper, for example, theorizes that the American public ought to be more supportive of uses of force that are accompanied by UN Security Council authorization. They argue that the Security Council ought to act as a second, honest agent to the electorate's principal given that voters cannot be expected to judge the worth of a military operation on their own. If the UN gives their consent, however, voters ought to be assured that their government is acting in the public's legitimate interest, and express their support (measured, again, by positive responses to presidential job approval polls) by greater margins.

In order to determine the extent to which case selection is a problem for rally studies, I replicated Chapman and Reiter's analysis. I assembled narrative descriptions

¹⁴I return to the thorny issue of case selection in Chapter 5.

for these American MIDs, with details from both Correlates of War’s own descriptions that were published with version 3.0 of the MID data¹⁵ and an updated set of descriptions from Fordham and Sarver (2001). Of 231 militarized disputes since 1946, 35 described fishing rights disputes or similar police actions that seem exceptionally unlikely to have entered the public consciousness, much less run the risk of starting a more serious war. Six events were bilateral disputes between the United States and Canada, though these were all fishing disputes in their own right.¹⁶

Although excluding these events from Chapman and Reiter’s analysis makes only a modest impact on their results, the effect of their main explanatory variable, UNSC authorization, is no longer discernible from the null hypothesis of no effect. If we take this exploratory analysis one step further and treat the September 11th - Afghan War crisis as a proper outlier, given that the authors used ordinary least squares regression in their analysis, then the effect of this variable diminishes further. The coefficients and their standard errors for these three models—Chapman and Reiter’s original analysis, and the two restrictions I suggest here—appear in columns 1 through 3 of Table 2.1, respectively.

The reported coefficient for the UNSC variable in the third column of Table ?? is nearly “significant” ($p < .069$). However, this is actually worse news than it appears. The second, more pernicious problem with observational studies of the rally effect is that the dependent variable is measured with error whose own structure is not typically modeled. Consider that the public’s support (or opposition) to the president’s foreign policy is operationalized as the change in presidential job approval, as reported by Gallup (or similar polls), before and after the crisis’ coded start date. The point estimate of the change in public approval, then, is estimated with sampling error related to the size of the sample in each individual poll. We can express the

¹⁵http://www.correlatesofwar.org/COW2Data/MIDs/MID_v3.0.narratives.pdf.

¹⁶One exception is MID #2952, which appears to have involved a Vietnam-era draft evader who was arrested by US Customs agents on the Canadian side of the border, and then later released after Canadian diplomatic protests (Fordham and Sarver 2001).

	M3	M3, no fishing disputes	M3, no fishing, & 9/11 dummy
	b/(se)	b/(se)	(b/se)
Prior Popularity	-0.111* (0.04)	-0.091* (0.04)	-0.097* (0.04)
Bipartisan Support	1.619* (0.80)	1.692* (0.80)	1.491* (0.75)
Admin. Statement	1.688** (0.59)	1.732** (0.63)	2.070*** (0.60)
Time to Next Election	0.013 (0.03)	0.004 (0.03)	-0.001 (0.03)
NYT Coverage	-0.424 (0.44)	-0.419 (0.48)	-0.599 (0.47)
UM Consumer Confidence	0.028 (0.02)	0.015 (0.01)	0.017 (0.01)
USA Revisionist	1.440* (0.67)	2.083** (0.71)	1.726** (0.64)
War Dummy, inc. Kosovo	11.634 (5.96)	11.416 (5.80)	5.858 (3.50)
Major Power Opp.	-1.085 (0.60)	-0.879 (0.66)	-0.899 (0.66)
# of US Allies	-0.135 (0.10)	-0.112 (0.10)	-0.054 (0.10)
Wartime MID	-1.211 (0.93)	-1.228 (1.10)	-1.201 (1.07)
UNSC Authorization	4.809 (2.60)	4.257 (2.59)	3.164 (1.88)
Other UN Action	-0.707 (0.94)	-0.659 (0.94)	-0.311 (0.85)
Regional Org. Action	-0.745 (1.99)	-0.681 (2.00)	-0.448 (1.99)
Crisis Severity	0.740 (0.55)	0.681 (0.62)	0.605 (0.61)
9/11			21.280*** (3.68)
constant	2.125 (1.45)	2.452 (1.55)	2.552 (1.51)
Obs	198	166	166
R2	0.345	0.410	0.487
dfres	182	150	149

* p<0.05, ** p<0.01, *** p<0.001

Table 2.1: Model 3, Chapman and Reiter (2004) – replications with modification.

margin of error for an estimate of the change in two independent polls as follows:¹⁷

$$\text{Margin of Error}_{p_2-p_1} = 1.96 * \sqrt{\left(\frac{p_1 * q_1}{n_1} + \frac{p_2 * q_2}{n_2}\right)} \quad (2.1)$$

Where p_i is the proportion of respondents answering a yes/no question—in this case, “Do you approve of the president’s job performance?”—in the affirmative, q_i is the complement of that percentage, and n_i is the total sample size for each poll.

Since the presidential approval data (and related sample sizes for each observation) are readily available,¹⁸ we can calculate the above formula for each observation and determine whether or not the estimated coefficient for Chapman and Reiter’s model falls outside this range. The average margin of error for presidential approval, across all the militarized disputes included by Chapman and Reiter, is +/- 3.58%, which is larger than the UNSC authorization variable after accounting for fishing disputes and the unique character of September 11th. In other words, for any pair of otherwise identical militarized disputes, the marginal contribution of UN Security Council authorization would not be distinguishable from sampling noise. The set of all rally events for the United States for which polling sample sizes are available are charted in Figure 2.3, with the average margin of error across the period included as well.

The distribution of the rally events themselves suggests another way to examine this dynamic. With only a few exceptions, the rally events sort themselves neatly into three distinct categories: positive rallies, between 5 to 10 positive percentage points, negative rallies, between 3 to 5 percentage points, and then those rally events which fall within the joint margin of error as defined above. Given the distribution of these data points, we might do more justice to the underlying sampling error in

¹⁷See Charles Franklin, <http://abcnews.go.com/images/PollingUnit/MOEFranklin.pdf>.

¹⁸http://webapps.ropercenter.uconn.edu/CFIDE/roper/presidential/webroot/presidential_rating.cfm

these data, while preserving the rough order of the rally events themselves, with an ordered choice estimation.

I recoded the rally events into three categories and estimated the probability that any given rally episode would fall within, above or below the average margin of error with an ordered probit model. Discrete ordered-outcome models are commonly interpreted as models of individual choice among alternatives, where the latent utility function of some respondent cannot be observed directly (Greene and Hensher 2010, p. 106). In this sense an ordered outcome estimation is throwing out information, because we are introducing threshold censoring where none existed before. However, these discrete categories nonetheless capture the point that for the bulk of rally events, we cannot distinguish small effects from no change in public opinion at all. The distribution of rally events after recoding is reported in Table 2.2.

Rally type	Frequency	%
Significant negative rally	42	21
Insignificant rally	130	66
Significant positive rally	26	13
Total	198	100

Table 2.2: Distribution of rally events in Chapman and Reiter (2004), after recoding.

Table 2.3 reports the marginal effect of UN Security Council authorization on the probability of observing a positive rally of sufficient magnitude that it could be discerned from sampling error. All other covariates in Model 3 from Chapman and Reiter (2004) were held at their sample medians. Given that the estimator was an ordered probit, the marginal effects are reported as percentage point changes. Unlike in the OLS estimations replicated above, UNSC Authorization fails to achieve statistical significance in any of the three sets of cases tested.

Finally, a closer inspection of the individual cases themselves provide another cause for concern. the single largest negative rally associated with an American MID is a 16-point drop in presidential approval rating connected to a dispute between the

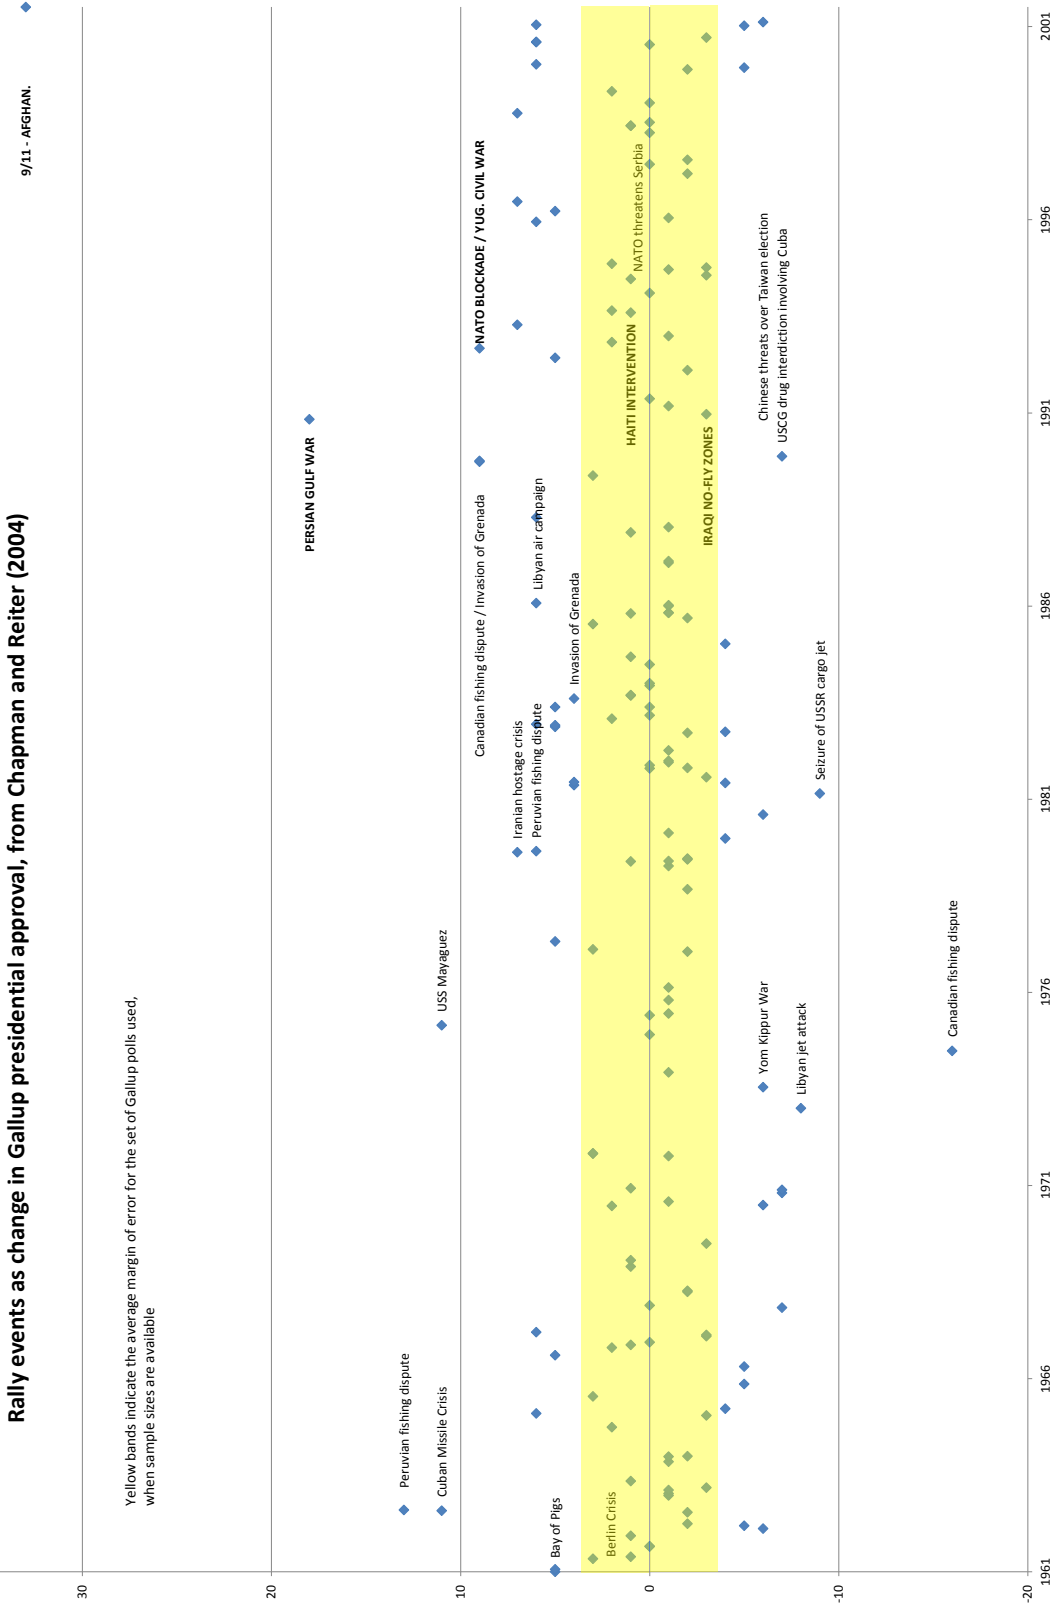


Figure 2-3: Rally events graphed by percentage-point change. The shaded area indicates the average joint margin of error.

	$\Delta Pr(\text{Pos. Rally})$, all Events		$\Delta Pr(\text{Pos. Rally})$ exc. fishing		$\Delta Pr(\text{Pos. Rally})$ exc. fishing & 9/11	
Marginal effect,	Pr	z-score	Pr	z-score	Pr	z-score
UNSC Auth.	6.0%	0.39	8.0%	0.55	8.6%	0.69

Table 2.3: Marginal effects of UNSC Authorization in the ordered probit estimation.

US and Canada on September 14th, 1974. This event was originally coded by Baker and Oneal (2001), and has since been included in a number of other studies including Chapman and Reiter (2004) and Colaresi (2007). This dispute was derived from a fairly vague wire report that five Canadian ships were detained by the United States Coast Guard for illegal fishing and then released the next day for \$5,000 in bond each.¹⁹

It is possible that the American public was particularly outraged by this development, and angry enough that they punished Ford with the largest single backlash in the post-war period for any event in the MID dataset. It is also possible that the public was reacting to something else that happened that same week: on September 8th, 1974, Ford issued a general pardon for Richard Nixon.

¹⁹Unfortunately, narrative documentation on both the MID and Correlates of War data are thin, and mostly non-existent for MID 2.0 and earlier. The only supporting documentation on the Canadian fishing dispute is the note “New York Times” in the file at http://www.correlatesofwar.org/COWData/MIDs/MID2.1_SpecificSources.csv.

CHAPTER 3

Formalizing Information Market Intervention

In this chapter, I build on the informal theoretical discussion offered at the end of Chapter 2 and develop a game-theoretic treatment of my argument. The game is, at its core, a signaling game between two states. The revisionist state, or “Challenger” is explicitly democratic with electoral accountability included in their utility function. The other state, which is satisfied with the status quo (“Defender”) is the prospective target of some challenge. The Defender’s own domestic regime is unspecified in the interest of tractability. The extensive form of the game is similar to the standard immediate-deterrence game that has been widely used to describe the period of bargaining immediately prior to a potential war. Although the single-play form of this game implies that it is not a model of bargaining *per se*, with a potential back-and-forth of offers between the two belligerents, many if not most of this latter class of models behave as single-offer games along the equilibrium path anyway (Powell 1999).

3.1 A Synthesis of Public Opinion and Bargaining, and a Theory

The discussion in Chapter 2 is important for two reasons. First, while voters might be more reasonable than the early foundational work in public opinion suggested, they rarely receive the full story of a foreign policy question from the media or their elected

leaders. This is where the bulk of behavioral political opinion research is located, and the available evidence, both empirical and experimental, seems to establish that censored, limited, primed or framed informational environments can generate swings in expressed opinion and judgment. Second, despite their vulnerability to information constraints, voters do make an effort to respond to the messages they receive in consistent ways. Although they might not always have the resources available to come to the “correct” decision as we might define it *ex post*, they will nonetheless respond.¹

These two factors, in conjunction with institutional, democratic peace-driven explanations of democracies’ foreign policy, suggest that failures or aberrations in democratic policy making are not simply attributable to voters themselves. It is the democratic executive who, confronted with the constraints of open politics, actually faces an incentive to restrict information available to citizens. Autocrats who rule at whim have comparatively little reason to change people’s hearts and minds, so long as they have sufficient military might to maintain their hold on power. With no other way to maintain power *other* than changing hearts and minds, democrats must, paradoxically, fear free speech the most. To the extent that winning bargaining surpluses against international rival aids in the provision of public goods, furthermore, democrats face an additional incentive to influence public opinion.

The interest, willingness and ability of a government to intervene in the information market, as well as the consequences of its intervention, require further specification. Under what conditions might a government choose to intervene in the information market and limit or manipulate information? What impact does this have on international outcomes? How does the government’s ability to shape the informational environment affect electoral support?

Previous modeling attempts have ignored the impact of endogenous information

¹See, however, the discussion in Chapter 4, in which subjects appeared to prioritize partisan cues in opinion formation over information about a specific event.

provision on democratic foreign policy. In this model's limited context, a democratic, challenging state may choose to intervene in the domestic information market, and this decision is observed by the foreign defender, whose regime is not specified. This intervention might take the form of burying the existence of current military action, such as the United States' Atlantic campaign prior to World War II, or by waging a proactive campaign to inflate the danger from a potential adversary, such as in the months and years leading to the US invasion of Iraq in 2003. In either event, the decision to intervene in the information market is a costly lottery, in much the same way that war itself is. If voters discover an intervention, they may choose to penalize the government, since government intervention is unpalatable. However, discovery of intervention is contingent upon the quality and freedom of the information market itself within a given country. With a well-established and protected media establishment, intervention is more likely to be discovered because the number of interested media players should increase. Trivially, a larger number of media should stand a better chance on stumbling across the truth than a smaller one, *ceteris paribus*.

Given the costs specified above, it is important to specify why a government might find information market intervention beneficial. There are two distinct ways of describing the value of intervention to a government. The simpler reason is that by intervening, the government might avoid being held accountable for failure. Given that even apologists for voters' interest in foreign policy, who claim that efficient voting behavior implies only paying attention during crises², admit that foreign policy enjoys far less salience than other issues, leaders might rationally expect that a publicized failure might carry an inflated weight with the public as compared to a number of less-well-known successes. In other words, leaders might fear post hoc retrospective voting that has a particularly short time horizon.³

²Shapiro and Jacobs (2000)

³In practice, of course, even "failed" cases of intervention nonetheless delayed the electoral consequences. The Pentagon Papers, for example, brought the abuses and secrecy of the Kennedy and Johnson administrations to light but did not implicate the incumbent at the time of publication,

The second reason, which is more complicated but also closer to existing literature on crisis escalation, is that intervention might increase the range of policy options available to the government. If a government can successfully manipulate the information market, they may have a greater ability to escalate conflict without risking a large loss of support. This additional resolve, generated by the government as an instrument for diplomacy, should yield some greater ability to demand concessions from a rival given that their expected costs of war⁴ would now be smaller. The ability to escalate a crisis more quickly with a supportive (albeit under-informed) public should enhance credibility and the chances of a favorable outcome by encouraging a defender to back down (Fearon 1994*a*).

This adds a wrinkle to the institutional theory of Bueno de Mesquita et al. (2003), whose model assumes that the most important preference of leaders is to stay in office. As they show, that objective implies that regime-specific differences in foreign policy and crisis behavior are largely a product of strategic selection, in which leaders are careful to only risk their prestige in safe bets such that their forward-looking credibility to prospective winning coalitions remains intact. Here, information market intervention might actually *create* new conflicts in which democratic executives can win spoils that better their odds of retaining office. While my model appears to give leaders a much greater degree of agency over their decision-making insofar as they can affect the public support for their policies in advance of those policies coming to fruition, their motivation in doing so is consistent with selectorate theory.

The ability to influence information markets and its associated cost-benefit relationship yield a number of hypotheses about both monadic and dyadic behavior. First, robust information markets that are relatively resistant to manipulation will encourage peaceful outcomes. Since strong information markets increase the probability

Richard Nixon.

⁴Operationalizing the costs of war as some combination of material costs (money, men, equipment) and psychological costs (national morale, popularity of the government) is a common crutch in the modeling literature; see Powell (1999).

of getting caught, which imposes a cost on the executive, such governments are less likely to use intervention and thus less likely to engage in unnecessary provocation. Second, democracies may have a relatively strong incentive to enter the information marketplace. An autocratic government that ignores the popular will can embark on whatever policy it deems fit.⁵ Democratic governments who care about maintaining support may have an incentive to manipulate information to gain support for certain policies, as described above. Strong institutions supporting free speech thus introduce a second strategic tension: they deter leaders from intervening, but their high expected cost also makes them more credible signals when they are used.

3.2 The Model

The modeling effort here simplifies many of the above possibilities in the interest of tractability. Rather than a two-sided interaction where both governments have the opportunity to intervene, in this model only the challenger has that option. Implicitly, this means the model assumes a democratic challenger, since in a dictatorship there is not a relevant public to bother misleading in most circumstances.⁶ Although the defender is left without the opportunity to censor, it should not be interpreted as an autocratic state.

This model adapts a crisis escalation model to include the possibility of information market intervention. At the outset of the crisis, the challenging government can choose to maintain the status quo ($-Ch$), challenge the status quo and influence the information market available to its voters about the challenge (Ch, C), or challenge

⁵As I mentioned in Chapter 2, however, scholarship on this point appears to be changing. Comparative politics research has identified that some “competitive autocracies” do exhibit a few of the trademarks of open political competition, whether within the ruling coalition or in carefully managed local elections, even while politics and government remain dominated by one party or faction.

⁶This, of course, only refers to mass publics, who may be activated in non-liberal regimes but not as a matter of course. Elites, whose relative importance is outsized in such a regime, are presumed to arrive at their conclusions on government performance via other means; on the life-cycles of nondemocratic leaders, see Goemans (2000).

the status quo without such intervention ($Ch, \neg C$).⁷ If there is no challenge, the game ends and each state receives payoffs equal to their existing share of the good under dispute. An important, albeit common, simplification is to assume that the good under dispute is only held by one of the two states at the outset, and that any resolution short of conflict will retain that all-or-nothing distribution.

$$U_{Chall}(\text{Status Quo}) = 0 \tag{3.2.1}$$

$$U_{Def}(\text{Status Quo}) = 1 \tag{3.2.2}$$

The instrumental value of intervention is that when successful, it further reduces the challenger’s material war costs. Since resolve, or the willingness to go to war, is a function of the costs of doing so, successful intervention can be seen as an added bonus to resolve. After observing the challenging state’s decision to challenge as well as whether or not to censor, the defending state decides whether or not to mobilize in response ($M; \neg M$). If the defender doesn’t mobilize, then the game ends and the challenger wins the entirety of the good being disputed.

$$U_{Chall}(\text{Defender Acquiescence}) = 1 \tag{3.2.3}$$

$$U_{Def}(\text{Defender Acquiescence}) = 0 \tag{3.2.4}$$

⁷This notation is carried forward from an early version of this paper, in which I described elite influence over public opinion as “censorship” of the information market, rather than the less-charged “intervention.”

If the defender does mobilize, then the challenger has a decision on whether to stand firm (SF) or back down ($\neg SF$). If the challenger backs down, then both sides are left with their status quo payoffs, but the challenger pays an additional reputation cost, r , which is bounded $[0,1]$.

$$U_{Chall}(\text{Chall. Acquiescence}) = (0 - r) = -r \quad (3.2.5)$$

$$U_{Def}(\text{Chall. Acquiescence}) = 1 \quad (3.2.6)$$

If the challenger continues the dispute, then both sides receive payoffs from the war lottery lottery, such that

$$U_{Chall}(\text{War}) = (w - c_{ch}) \quad (3.2.7)$$

$$U_{Def}(\text{War}) = (1 - w - c_{def}) \quad (3.2.8)$$

The challenging state has some chance of winning the conflict, w ; their utility for winning, then, is w times the value of the good under dispute (1), minus the cost of that conflict. The defending state has the complementary probability of winning, $(1 - w)$, the same value for the good and a separate cost of war. Both costs are bounded $[0, 1]$. Each state faces a different cost term, and these costs are common knowledge. They are also specified exogenously, rather than arising endogenously

from one another's strategy or typing. Note also that that the cost terms are scalar, and as such both actor's beliefs about the other's type will be expressed as simple weighted sums rather than the cumulative area under a smooth distribution function.

Finally, in those cases where the challenger decides to intervene, they face an additional lottery in which their domestic polity may discover censorship. When intervention is successful, the challenging government pays a lower cost for conflict. X parameterizes the strength of the challenger's media and speech institutions; higher values of x indicate a higher probability that an unresolved challenger will reduce their war costs to that of the resolved type. In other words, x is the probability that intervention works. $(1 - x)$ is the strength or independence of media, and as such is the complementary probability that intervention fails. When $x = 0$, the challenging government has no ability to reduce their cost of conflict, and the "independence" measure $(1 - x)$ is at its maximum value, 1. Successful intervention then reduces the challenging government's cost of war from c_{ch} to c^*_{ch} , where $c_{ch} > c^*_{ch}$. The value of $(c_{ch} - c^*_{ch})$ reflects the importance that additional political support, gained through intervention in the information market for voters, has for political calculations abroad.

An additional parameter, l , reflects the political penalty to being caught in a lie. While there is only some chance, $(1 - x)$, that intervention is uncovered during the conflict (restricting the challenger to its high cost military option, c_{ch}), intervention will always be revealed after the conflict has passed. By contrast, there is only $(1 - x)$ probability that attempted intervention is revealed after the fact in the event of a non-violent outcome. The intuition here is that the government's position will be more closely examined after the outbreak of open conflict than it would be otherwise. One simplifying assumption is added where the challenging government avoids a penalty for intervention when it wins an acquiescence from the defending government.

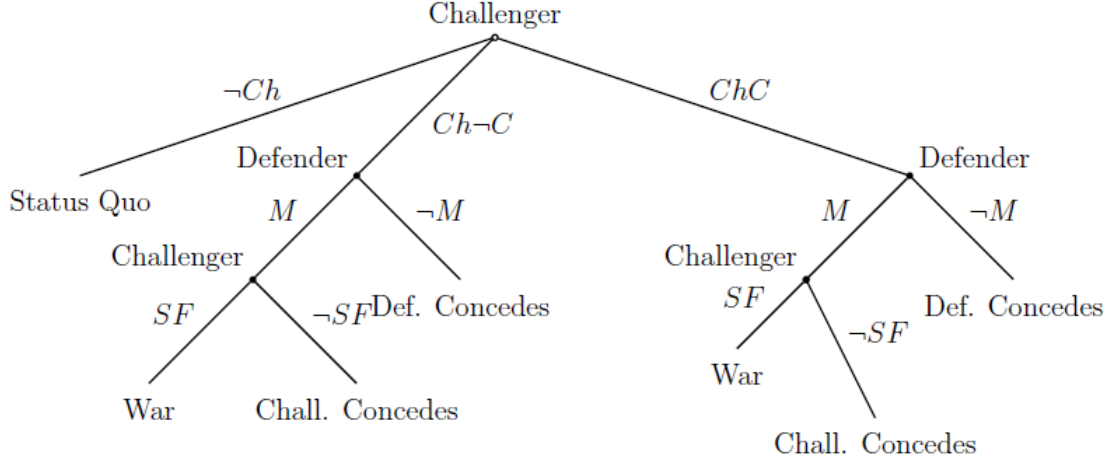


Figure 3.1: The Crisis Censorship game under complete information. Incomplete information adds a second branch to the game with an initial move by Nature that determines the Challenger’s cost term, such that the Defender only has two information sets: one after observing a challenge with prior intervention (Ch, C) and one after observing a challenge without such intervention ($Ch, -C$).

$$U_{Chall}(\text{War, Intervene}) = w - (x(c_{ch}) + (1 - x)(c^*_{ch})) - l \quad (3.2.9)$$

$$U_{Chall}(\text{Chall. Acq, Intervene}) = -r - (1 - x)l \quad (3.2.10)$$

$$U_{Chall}(\text{Def. Acq, Intervene}) = U_{Chall}(\text{Def. Acq, -Interv.}) = 1 \quad (3.2.11)$$

Although earlier versions of this model included a third player, the challenger’s median voter, who had the opportunity to choose whether to support the government after the crisis is resolved, it has been omitted here. Including a voter is intuitive, but did not add any value to the model’s conclusions. In generalized payoffs, specifying the value of an unknown challenger amounts to some exogenous scaling parameter, m , that ultimately falls out of all relevant equilibrium conditions. As such, the backlash from a “voter” who has been lied to is captured in l .

3.3 Equilibrium Analysis under Complete Information

For the complete information case, we will explore a limited version of the model which further simplifies the interaction between challenger and the defender. In this version of the model, the challenger can either pay a high cost of war (c_{ch}) or a low cost of war (c_{ch}^*), depending on whether or not they elect to engage in censorship. Implicitly, this model assumes an irresolute challenger whose cost of war depends on their decision to intervene. Their international rival, the defending state, can choose whether or not to respond to the challenge in force. The challenging government has the final decision on whether or not the crisis ends in war. In the event that they have previously intervene, the payoff to war for the challenging state is

$$U_{Chall}(\text{War, Censorship}) = w - (x(c_{ch}^*) + (1 - x)(c_{ch})) - l \quad (3.3.1)$$

while concessions yield

$$U_{Chall}(\text{Chall. Acq.} | \text{Censorship}) = -r - (1 - x)l \quad (3.3.2)$$

The challenging government will then stand firm and fight a war after censorship when

$$w - (x(c_{ch}^*) + (1 - x)(c_{ch})) - l > -r - (1 - x)l$$

$$w - xc_{ch}^* - c_{ch} + xc_{ch} - l > -r - (1 - x)l$$

$$w + x(c_{ch} - c_{ch}^*) - c_{ch} - l > -r - l + xl$$

$$(w - c_{ch}) + x(c_{ch} - c_{ch}^*) - xl > -r \quad (3.3.3)$$

Put simply, when the expected utility of uncensored war plus the average benefit of intervention—the difference between the high and low cost of war—and the penalty for intervention is better than the alternative outcome of reputation loss, then the challenging state will elect to stand firm and fight a conflict at the final node in the game. When the challenging government has not intervened, they face a relatively simpler trade-off. Here, war yields

$$U_{Chall}(\text{War}) = w - c_{ch} \quad (3.3.4)$$

while concessions produce

$$U_{Chall}(\text{Chall. Acq.}) = (0 - r) = -r \quad (3.3.5)$$

When $w - c_{ch} > -r$, then war will result after the challenging state has played ($Ch \rightarrow C$). Additionally, in those cases where $w - c_{ch} > -r$, the challenger will also fight provided that $-r < -xl$. To see why, we return to the earlier inequality describing the challenger's final decision:

$$(w - c_{ch}) + x(c_{ch} - c_{ch}^*) - xl > -r$$

Should the first term, $(w - c_{ch})$, which expresses the uncensored utility of conflict,

be greater than $-r$, then the only other important consideration is the relationship between $-xl$ and $-r$, since $x(c_{ch} - c_{ch}^*)$ must always be non-negative.

The value of $-r$, then, is the primary factor in determining the challenger's equilibrium behavior. When it is less than the challenger's value for war in the unintervened case as well as the expected loss due to intervention, then the challenger will stand firm in both subgames. When it is greater than both of those terms, it will always accept the status quo. The intermediate case, when $-r$ is greater than the challenger's utility from conflict but less than their disutility from intervention, requires additional specification.

When this condition holds and the defending state has a low value for war, such that

$$(1 - w - c_{def}) < 0 \tag{3.3.6}$$

Then the challenging state will intervene and challenge, and the defender will acquiesce. The crucial test, however, occurs when $(1 - w - c) > 0$. In that case, the challenging government is limited to two strategy profiles: (ChC, SF) or $(\neg Chall, \neg SF)$. They cannot challenge without censorship and stand firm because the reputation cost $-r$ is clearly less than 0, the payoff to the status quo that was available at the beginning of the game. Challenging, intervening, and fighting the war is possible depending on the comparison of their payoff to war to the reserve payoff, as detailed above. When that condition holds—for large values of x , the government's ability to "get away with it," and small values of l , voter retaliation for discovering censorship—the government will challenge and war will result.

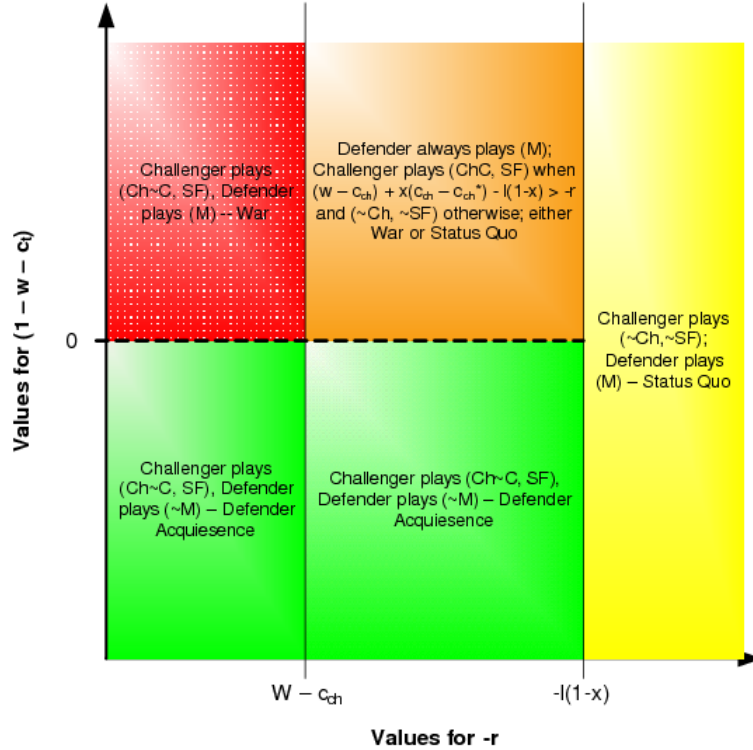


Figure 3.2: Boundary equilibrium conditions for the restricted Crisis Censorship game for varying parameters

Figure 3.2 summarizes these conditions. Note that I assume here that $(w - c_{ch})$ is less than $-xl$ for the sake of visualization.

3.4 Equilibrium Analysis Under Incomplete Information

The incomplete information game builds upon the structure established in the complete information case above. Where the complete information game simplified the distinction between different possible costs for war dependent on whether the challenger elected to make a challenge after intervention, we now establish two different types of challenger, where type is decided by Nature prior to any decision by either player. An irresolute challenger pays c_{ch} for conflict without intervention, and c_{ch}^* after intervention, where $c_{ch} > c_{ch}^*$. A resolute challenger will pay c_{ch}^* irrespective of intervention. The irresolute challenger's decision whether or not to intervene in

the information market will depend on x and l , the two parameters that capture the penalty for the challenging government lying to its citizens. Since the game only features two types of challenger, I use the typical transformation (Harsanyi 1967) in representing the game as a game of imperfect information wherein Nature determines the type of challenger in accordance with some given probability.

The relative probability of each type of challenger is γ and $(1 - \gamma)$, where γ is distributed uniformly over the range $[0,1]$. In any candidate equilibria, the defender's updated beliefs must satisfy sequential rationality given γ , the challenger's observed actions and known relative values of their utilities for the potential outcomes.

In the complete information case, we saw that for any values of w and c_{def} where $1 - w - c_{def} > 0$, the defender will always mobilize at their decision step. Similarly, when $w - c_{ch} < -r$, then the challenging state will never make a challenge since they will never be willing to stand firm at their last decision. The challenger is guaranteed a payoff of 0 if they play $(-Ch)$, irrespective of their type. If they challenge without censorship, $(Ch-C)$, then the defender's consistent belief about the expected utility of that action, in the event of conflict, to the challenger is

$$\gamma(w - c_{ch}^*) + (1 - \gamma)(w - c_{ch}) \tag{3.4.1}$$

The defender must believe that the challenger will receive the payoff for a resolute type with probability γ - their starting prior about the distribution between resolute and irresolute types - and the payoff for an irresolute type with the complementary probability $(1 - \gamma)$. The defender must also believe that that expression is greater than both the payoff to capitulation at the final node, $-r$, as well as the status quo payoff 0 for the challenger to have played that strategy. By construction, the status quo payoff is larger.

Therefore,

$$\gamma(w - c_{ch}^*) + (1 - \gamma)(w - c_{ch}) > 0 \quad (3.4.2)$$

This inequality is true when

$$\begin{aligned} \gamma w - \gamma(c_{ch}) + w - c_{ch}^* - \gamma w + \gamma c_{ch} &> 0 \\ w - c_{ch} + \gamma(c_{ch} - c_{ch}^*) &> 0 \\ \gamma(c_{ch} - c_{ch}^*) &> c_{ch} - w \end{aligned}$$

$$\gamma > \frac{(c_{ch} - w)}{(c_{ch} - c_{ch}^*)} \quad (3.4.3)$$

The inequality in (3.4.2) gives us a condition on the defender's updated beliefs given $(Ch \rightarrow C)$ and relative to the status quo payoff. This is indicative of the set of inequalities that must be satisfied in order for the defender's beliefs to be sequentially rational. The defender should further believe, however, that the expected utility to the challenger for this action must also be greater than the expected utility from challenging under intervention, (ChC) . The defender's expectation of that utility to the challenger is

$$\begin{aligned} \gamma(w - c_{ch}^* - l) + (1 - \gamma)[w - (x(c_{ch}^*) + (1 - x)(c_{ch})) - l] \\ \gamma(w - c_{ch}^* - l) + (1 - \gamma)[w - (xc_{ch}^* + c_{ch} - xc_{ch}) - l] \\ \gamma(w - c_{ch}^* - l) + (1 - \gamma)[w - (-x(c_{ch} - c_{ch}^*) + c_{ch}) - l] \end{aligned}$$

$$w + \gamma(c_{ch} - c_{ch}^*) - \gamma x(c_{ch} - c_{ch}^*) + x(c_{ch} - c_{ch}^*) - c_{ch} - l \quad (3.4.4)$$

We can find a lower bound for γ such that expression (3.4.1) is greater than expression (3.4.4), which will yield a belief threshold that must obtain for the defender's beliefs after observing $(Ch \neg C)$ to be sequentially rational.

$$\begin{aligned} w - c_{ch} + \gamma(c_{ch} - c_{ch}^*) &> w + \gamma(c_{ch} - c_{ch}^*) - \gamma x(c_{ch} - c_{ch}^*) + x(c_{ch} - c_{ch}^*) - c_{ch} - l \\ l &> x(c_{ch} - c_{ch}^*) - \gamma x(c_{ch} - c_{ch}^*) \\ l &> x(c_{ch} - c_{ch}^*)(1 - \gamma) \\ \frac{l}{x(c_{ch} - c_{ch}^*)} &> 1 - \gamma \end{aligned}$$

$$\gamma > 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} \quad (3.4.5)$$

Both (3.4.3) and (3.4.5) give lower bounds on γ that must be satisfied for sequential rationality if the challenging state plays $(Ch \neg C)$. We can determine which of the bounds is larger with respect to the other open parameters of the model:

$$\begin{aligned} 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} &> \frac{c_{ch} - w}{c_{ch} - c_{ch}^*} \\ (c_{ch} - c_{ch}^*) - \frac{l}{x} &> c_{ch} - w \\ w - c_{ch}^* &> \frac{l}{x} \end{aligned}$$

$$x > \frac{l}{(w - c_{ch}^*)} \quad (3.4.6)$$

For sufficiently large values of x , the probability of successful censorship, expression (3.4.5) is larger than (3.4.3) and the relevant boundary on the defender's beliefs is derived from the comparison between the censored and uncensored challenge. This critical threshold is the penalty for being caught in the act of censorship divided by the expected utility of war to the resolute type. When x falls below this threshold, then the relevant comparison is between the uncensored challenge and the status quo. For sufficiently small values of x —for sufficiently small probabilities of successful censorship by the challenger—the possibility of censorship is no longer relevant in shaping the defender's updated beliefs.

In either event, the defender will use their beliefs in weighting the possible outcomes to their own military mobilization. When inequality (3.4.6) is true, the defender's beliefs are derived from expression (3.4.5). In order to mobilize, the defender's expected utility from doing so must be greater than 0, the payoff to acquiescence:

$$\begin{aligned} & [1 - (\frac{l}{(x(c_{ch} - c_{ch}^*)})}](1 - w - c_{def}) + [1 - (1 - (\frac{l}{(x(c_{ch} - c_{ch}^*)})))](1) > 0 \\ & (1 - \frac{l}{x(c_{ch} - c_{ch}^*)})(1 - w - c_{def}) + (\frac{l}{x(c_{ch} - c_{ch}^*)}) > 0 \\ & 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} - w + \frac{wl}{x(c_{ch} - c_{ch}^*)} - c_{def} + \frac{c_{def}l}{x(c_{ch} - c_{ch}^*)} + \frac{l}{x(c_{ch} - c_{ch}^*)} > 0 \end{aligned}$$

$$(1 - w - c_{def}) + \frac{l(c_{def} + w)}{x(c_{ch} - c_{ch}^*)} > 0 \quad (3.4.7)$$

The second left hand side term, $\frac{l(c_{def}+w)}{x(c_{ch}-c_{ch}^*)}$, modifies the defender's default utility from war. The numerator will always be positive, and the left hand side of the equation will decrease with larger values of $x(c_{ch}-c_{ch}^*)$, given that $(1-w-c_{def}) < 0$. A larger value of $(c_{ch}-c_{ch}^*)$, which indicates a wider gulf between the two types of challenger, will make it less likely that the defender will mobilize, as will a larger probability of "getting away" with intervention, x . Recall that (3.4.7) determines defender mobilization when x is large, such that the defender's belief set is reflective of the difference in utility to the challenger from intervening or not intervening. With a large enough value of x to satisfy (3.4.6), a challenger that fails to intervene is declining a relatively sure option to "trade up" from c_{ch} to c_{ch}^* . This serves as an indicator to the defender that the challenger is almost certainly resolute, meaning that they should not mobilize.

Equilibria for the incomplete information game will be given in the form (Challenger's strategy profile if resolute, Challenger's strategy profile if irresolute; Defender's strategy profile; Defender's belief that the Challenger is resolute given $(Ch-C)$; Defender's belief that the challenger is resolute given (ChC)). The defender's updated belief after witnessing $(Ch-C)$ is labeled p ; their belief after (ChC) is q . When the inequalities in (3.4.6) and (3.4.7) are true, then the equilibria are

$$\{(Ch-C, SF), (-C, \neg SF); (M); p > 1 - \frac{l}{x(c_{ch}-c_{ch}^*)}; q\} \quad (3.4.8)$$

The defender's off-path belief after observing the non-equilibrium action (ChC) can be defined as any belief that will support its equilibrium strategy of mobilization.

$$\begin{aligned}
EU_{def}(M|ChC) &= q(1 - w - c_{def}) + (1 - q)(1) > 0 \\
q - qw - qc_{def} + 1 - q &> 0 \\
q(w + c_{def}) &< 1
\end{aligned}$$

$$q < \frac{1}{w + c_{def}} \quad (3.4.9)$$

As such, the full equilibrium specification is

$$\{(Ch \neg C, SF), (\neg Ch, \neg SF); (M); p > 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; q < \frac{1}{(w + c_{def})}\} \quad (3.4.10)$$

With the defender's belief set such that it will mobilize in the face of the uninter-vened challenge, the irresolute challenger is forced to stand down lest they trade a sure outcome of 0 for their war payoff $w - c_{ch}$. When (3.4.7) is violated, then the defender will not mobilize. The equilibrium becomes pooling, with both types of challenger challenging. This equilibrium is specified as $\{(Ch \neg C, SF), (Ch, \neg SF); (\neg M); p > 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; q > \frac{1}{w + c_{def}}\}$, with the defender's off-path belief specified as above.

When the domestic information market in the challenging state is stronger, such that x is relatively small and the inequality in (3.4.6) is violated, the lower bound on the defender's belief is then derived from (3.4.3). In that case, the defender's expected utility to mobilization in the face of $(Ch C)$ becomes

$$\frac{c_{ch} - w}{c_{ch} - c_{ch}^*}(1 - w - c_{def}) + (1 - \frac{c_{ch} - w}{c_{ch} - c_{ch}^*})(1) \quad (3.4.11)$$

This expected value must be greater than the reservation payoff of 0 in order to support mobilization in equilibrium. We can derive a new expression to explore when that is the case:

$$\begin{aligned} & \left(\frac{c_{ch} - w}{c_{ch} - c_{ch}^*}\right)(1 - w - c_{def}) + \left(1 - \frac{c_{ch} - w}{c_{ch} - c_{ch}^*}\right)(1) > 0 \\ & \frac{(c_{ch} - w) - w(c_{ch} - w) - c_{def}(c_{ch} - w) - (c_{ch} - w) + (c_{ch} - c_{ch}^*)}{c_{ch} - c_{ch}^*} > 0 \\ & \frac{(-w - c_{def})(c_{ch} - w) + (c_{ch} - c_{ch}^*)}{c_{ch} - c_{ch}^*} > 0 \end{aligned}$$

$$(-w - c_{def})(c_{ch} - w) + (c_{ch} - c_{ch}^*) > 0 \tag{3.4.12}$$

Since $c_{ch} > c_{ch}^*$ by definition, and $(-w - c_{def})$ will always be a negative value given that w is bounded $[0, 1]$, (3.4.12) would hold anywhere in the parameter space should $c_{ch} < w$. However, we assumed earlier that $w - c_{ch} < 0$ such that the irresolute challenger does in fact prefer to capitulate than fight, *ceteris paribus*. That in mind, (3.4.12) can only be true for sufficiently large differences between the irresolute and resolute types. Recall that for relatively weak information markets, larger values of $(c_{ch} - c_{ch}^*)$ led to capitulation from the defender, because the decision to forego censorship when it was attractive indicated resolve from the challenger. Here, with a relatively strong information market, a large gap between the resolute and irresolute types actually provokes a response from the defender; indeed, as the costs of war to the resolute challenger fall, the defender is more likely to fight. The defender has to worry about the possibility of making an incorrect concession to an irresolute foe who can attempt to piggy-back on the strength of the resolved type now that it is costly for them to do so with a smaller value for x . When (3.4.12) is satisfied, then the defender will mobilize and the irresolute defender will once again be deterred from a

challenge, leaving a separating equilibrium of

$$\{(Ch \neg C, SF), (\neg Ch, \neg SF); (M); p > \left(\frac{c_{ch} - w}{c_{ch} - c_{ch}^*}\right); q < \frac{1}{w + c_{def}}\} \quad (3.4.13)$$

When (3.4.12) is not satisfied, then we return to a pooling equilibrium,

$$\{(Ch \neg C, SF), (Ch, \neg SF); (\neg M); pp < \left(\frac{c_{ch} - w}{c_{ch} - c_{ch}^*}\right); q > \frac{1}{w + c_{def}}\} \quad (3.4.14)$$

I pause now to summarize the various boundary conditions that lead to equilibria with non-censored challenges. When $x > \frac{l}{w - c_{ch}^*}$, and the information market of the challenger is relatively weak, the equilibrium can be described as follows:

$$\begin{aligned} & (Ch \neg C, SF), (\neg C, \neg SF); (M); p > 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; q < \frac{1}{w + c_{def}} \\ & \text{if } 1 - w - c_{def} + \frac{l(c_{def} + w)}{x(c_{ch} - c_{ch}^*)} > 0 \\ & (Ch \neg C, SF), (Ch \neg C, \neg SF); (\neg M); p > 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; q > \frac{1}{w + c_{def}} \\ & \text{if } 1 - w - c_{def} + \frac{l(c_{def} + w)}{x(c_{ch} - c_{ch}^*)} < 0 \end{aligned}$$

When $x < \frac{l}{w - c_{ch}^*}$, meaning that media and institutions of the challenging state are stronger, the equilibrium is then

$$\begin{aligned}
& (Ch \neg C, SF)\}, (-C, \neg SF); (M); p > \frac{c_{ch} - w}{c_{ch} - c_{ch}^*}; q < \frac{1}{w + c_{def}} \\
& \text{if } (-w - c_{def})(c_{ch} - w) + (c_{ch} - c_{ch}^*) > 0 \\
& (Ch \neg C, SF)\}, (Ch \neg C, \neg SF); (\neg M); p < \frac{c_{ch} - w}{c_{ch} - c_{ch}^*}; q > \frac{1}{w + c_{def}} \\
& \text{if } (-w - c_{def})(c_{ch} - w) + (c_{ch} - c_{ch}^*) < 0
\end{aligned}$$

There is a second set of candidate equilibria where the challenging government elects to censor. In the complete information game, information market intervention was exclusively a tool of the weak, irresolute government. With a weak enough information market, the irresolute challenger could credibly commit to fighting at the end of the game, possibly deterring the defender from their own mobilization. In the incomplete information case, the defender once again updates its beliefs after observing the challenging government's decision to challenge and censor. Recall that the defender's belief about the expected utility of that action was given in expression (3.4.4): $w + \gamma(c_{ch} - c_{ch}^*) - \gamma x(c_{ch} - c_{ch}^*) + x(c_{ch} - c_{ch}^*) - c_{ch} - l$.

The defender's updated value for γ must reflect an expectation that the challenger's expected utility as defined in (3.4.4) is greater than both the payoff to $(Ch \neg C)$ and the reserve status quo payoff, 0. We begin by determining when (3.4.4) is greater than 0:

$$\begin{aligned}
w + \gamma(c_{ch} - c_{ch}^*) - \gamma x(c_{ch} - c_{ch}^*) + x(c_{ch} - c_{ch}^*) - c_{ch} - l &> 0 \\
w + \gamma(1 - x)(c_{ch} - c_{ch}^*) + x(c_{ch} - c_{ch}^*) - c_{ch} - l &> 0 \\
\gamma(1 - x)(c_{ch} - c_{ch}^*) &> c_{ch} + l - w - x(c_{ch} - c_{ch}^*)
\end{aligned}$$

$$\gamma > \frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} \quad (3.4.15)$$

Expression (3.4.15) is the lower bound that the defender's beliefs must satisfy in this case, and is derived from the comparison in the challenger's expected utilities for censored conflict and the status quo. The second boundary condition, from the comparison between censored and uncensored war, is simply the reverse of expression (3.4.5), which found the critical value for γ such that un-intervened conflict was preferred to intervention for the challenger:

$$\gamma < 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} \quad (3.4.16)$$

Taken together, (3.4.15) and (3.4.16) define a specific range that the defender's updated value for γ must satisfy in order to maintain sequential rationality after observing (*ChC*). We can further determine under what conditions (3.4.15) and (3.4.16) are properly ordered and able to support such a belief:

$$\begin{aligned} \frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} &< 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} \\ c_{ch} + l - w - x(c_{ch} - c_{ch}^*) &< (c_{ch} - c_{ch}^*)(1 - x) - \frac{l(1 - x)}{x} \\ -x(c_{ch} - c_{ch}^*) - (1 - x)(c_{ch} - c_{ch}^*) + c_{ch} + l - w &< \frac{l(1 - x)}{x} \\ -xc_{ch} + xc_{ch}^* - c_{ch} + c_{ch}^* + xc_{ch} - xc_{ch}^* + c_{ch} + l - w &< \frac{l(1 - x)}{x} \\ c_{ch}^* + l - w &< \frac{l(1 - x)}{x} \end{aligned}$$

$$x < \frac{l}{c_{ch} + l - w} \quad (3.4.17)$$

When (3.4.17) is satisfied, we can examine the expected utility to the defender of mobilization at both the lower and upper bound on its updated beliefs. At the lower bound, (3.4.15), that utility must be greater than 0 in order to support mobilization:

$$\begin{aligned} EU_{def}(M) &= \\ &= \left(\frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} \right) (1 - w - c_{def}) + \left(1 - \frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} \right) \\ &= \frac{(1 - w - c_{def})(c_{ch} + l - w - x(c_{ch} - c_{ch}^*))}{(c_{ch} - c_{ch}^*)(1 - x)} \\ &+ \frac{+c_{ch} - c_{ch}^* - xc_{ch} + xc_{ch}^* - c_{ch} - l + w + xc_{ch} - xc_{ch}^*}{(c_{ch} - c_{ch}^*)(1 - x)} \\ &= \frac{(1 - w - c_{def})(c_{ch} + l - w - x(c_{ch} - c_{ch}^*)) + w - l - c_{ch}^*}{(c_{ch} - c_{ch}^*)(1 - x)} > 0 \end{aligned}$$

$$(1 - w - c_{def})(c_{ch} + l - w - x(c_{ch} - c_{ch}^*)) + w - l > c_{ch}^* \quad (3.4.18)$$

Knowing that censorship is a tool of the weak, low resolute war costs to the challenger actually induce the defender to mobilize. Similarly, at the upper bound on the defender's beliefs, there is another cap on the value of c_{ch}^* :

$$\begin{aligned}
EU_{def}(M) &= \left(1 - \frac{l}{x(c_{ch} - c_{ch}^*)}\right)(1 - w - c_{def}) + \left(1 - 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}\right)(1) \\
&= \frac{(x(c_{ch} - c_{ch}^*) - l)}{x(c_{ch} - c_{ch}^*)}(1 - w - c_{def}) + \left(-\frac{l}{x(c_{ch} - c_{ch}^*)}\right) > 0 \\
&\quad \frac{x(c_{ch} - c_{ch}^*)(1 - w - c_{def}) - l(1 - w - c_{def}) - l}{x(c_{ch} - c_{ch}^*)} > 0 \\
&\quad (1 - w - c_{def})(x(c_{ch} - c_{ch}^*) - l) > l \\
&\quad x(c_{ch} - c_{ch}^*) - l > \frac{l}{(1 - w - c_{def})} \\
&\quad (c_{ch} - c_{ch}^*) > \frac{l(1 - w - c_{def}) + l}{x(c_{ch} - c_{ch}^*)}
\end{aligned}$$

$$c_{ch}^* < c_{ch} - \frac{l(1 - w - c_{def}) + l}{x(1 - w - c_{def})} \quad (3.4.19)$$

When (3.4.17), (3.4.18) and (3.4.19) are true, then the defender's beliefs fall within the expected range for γ , and they will mobilize in response to any challenge. The resolute challenger's strategy is unchanged, since their utility from conflict is still strictly larger than that from the status quo. The irresolute challenger will be deterred. This equilibrium is specified as

$$\left\{ (Ch \neg C, SF), (\neg Ch, \neg SF); (M); p < 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; \frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} < q < 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} \right\} \quad (3.4.20)$$

By contrast, when (3.4.18) and (3.4.19) are such that the defender will not mobilize, the equilibrium is

$$\left\{ (Ch-C, SF), (ChC, \neg SF); (M); p < 1 - \frac{l}{x(c_{ch} - c_{ch}^*)}; \right. \\ \left. \frac{c_{ch} + l - w - x(c_{ch} - c_{ch}^*)}{(c_{ch} - c_{ch}^*)(1 - x)} < q < 1 - \frac{l}{x(c_{ch} - c_{ch}^*)} \right\} \quad (3.4.21)$$

Most importantly, there are conditions under which an irresolute challenger can find itself in a conflict with the defender. We have just demonstrated the qualifications on c_{ch}^* such that the defender will mobilize in the face of a censored challenge. Taking the restriction on c_{ch}^* from the lower bound on the defender's updated belief after this event and substituting into the irresolute challenger's utility for conflict yields:

$$w - [x((1 - w - c_{def})(c_{ch} + l - w - x(c_{ch} - c_{ch}^*)) + w + l) + (1 - x)(c_{ch})] - l \quad (3.4.22)$$

When (3.4.22) is greater than zero, then the irresolute challenger will censor and challenge, even in the face of the defender's own mobilization.

3.5 Interpretations

I have highlighted several candidate equilibria, conditioned on the open parameters of the model, where the option to intervene in the information market prior to a crisis results in counterintuitive crisis behavior for the defending state. Critically, the option to censor means that the defender is more likely to resort to war when the costs of conflict associated with the irresolute type are low. Conventional deterrence theory suggests an actor with low costs should be more likely to avoid conflict because its implicit threat to resort to conflict is more credible. Here, however, the possibility of a irresolute (high cost) type improving its lot actually produces more cases in which the

defender will mobilize. Low cost types appear to suffer in the shadow of information market intervention insofar as they are much less able to communicate their resolve to their adversary, which would otherwise provide a concession that both sides would prefer.

The strength or weakness of the information market was less relevant than I expected. In prior versions of the model, weak information markets led to conventional deterrence behavior (the defender would mobilize against the challenger when the resolute type had high costs of war, and back down against low-cost challengers) while strong information markets would lead to the more orthodox strategy. Under incomplete information, both strong and weak information markets led to a mobilization decision on the part of the defender that was increasing in c_{ch}^* , as in expressions (3.4.7) and (3.4.12). For the equilibria possibilities that included a censorship decision on the part of the irresolute type, larger values of c_{ch}^* meant that the defender's mobilization inequality was less likely to be satisfied, as in expressions (3.4.18) and (3.4.19).

The most important condition on theoretical grounds is expression (3.4.22), which describes the possibility of a defender mobilizing in the face of an irresolute challenger who nonetheless, because of the censorship option, intends to stand firm at their final decision node. Here, the second lottery of building additional support for conflict in the hope of securing a bargaining surplus without a fight actually backfires, entrapping irresolute, high-cost types in a gamble that does not manage to cow the adversary.

Given the importance of public consent in the literature on regime type and war, the effect of electoral backlash is particularly interesting. The partial derivative of expression (3.4.22) is $x(2 - w - c_{def}) - 1$, which of course means that the effect of the backlash term, l , is itself dependent on particular realization of the other parameters. By inspection, however, the partial derivative with respect to l is itself decreasing in the defending state's cost term, c_{def} . This implies that the problem facing democratic leaders is even more stark: as c_{def} increases, and the defending state's bargaining

position deteriorates, the effect of *more* prospective electoral punishment is to make it more likely that (3.4.22) is greater than 0 and thus avoid the trap of “failed” intervention. However, electoral punishment is often visited on successive leaders rather than the actual incumbent at the time of the decision (Croco 2011), whether this would actually discourage intervention in the pre-crisis phase for such types is still unclear.

Aside from the question of the extent to which different equilibrium configurations are realized for different types of conflicts from the historical record, from here I turn to two very simple questions: is the sort of public diplomacy that would affect resolve—reduce a challenger’s cost term from c_{ch} to c_{ch}^* in fact possible? And, in drawing on the American record, is there any variance at all in the set of realized equilibrium strategies? These questions motivate Chapters 4 and 5, respectively.

CHAPTER 4

Radio Free USA: Does Information Market Intervention Work?

“Joker, I’ve told you, we run two basic stories here. Grunts who give half their pay to buy gooks toothbrushes and deodorants—Winning of Hearts and Minds—okay? And combat action that results in a kill—Winning the War. Now you must have seen blood trails... drag marks?”

Lt. Lockhart, *Full Metal Jacket*

In the previous two chapters, I outlined intuitive, and then formal, examinations of the overlap between the study of individual voter behavior and democratic crisis diplomacy. The formal analysis implies that the ability of democratic executives to pre-condition their mass public’s perceptions of the need for war changes the strategic calculus facing a foreign adversary. Even when information market intervention is not used, the mere option itself incentivizes other states to respond more aggressively to crisis challenges than they otherwise would.

However, this model relies on a number of open parameters, and as I discussed in the introduction, many of them are difficult or impossible to measure accurately or reliably. And, even if such were not the case, any long-term observational research design would be forced to rely on poor measures of public sentiment toward foreign policy performance. This would bring us back to the series of dilemmas that I

presented at the end of Chapter 2 regarding the use of presidential job approval data.

In this chapter, I use a novel survey experiment¹ to answer a simple preliminary question: how do individual perceptions of a *hypothetical* conflict, measured during a pre-crisis phase, change in response to different political messages meant to describe the case for war? If information market intervention is possible in practice, then we must be able to find some evidence of parallel behavior in a controlled test. This experimental framework also addresses the persistent issue that most studies of the public’s understanding of the use of force do so in the context of wars that have already begun, as with the body of research on American voters’ sensitivity to casualties. Voters’ reactions to crises that have yet to occur are much more consistent with the forward-looking nature of crisis bargaining as it is understood and modeled in international relations.

I test whether two broadly-accepted elements of political persuasion—expertise and emotional appeals—can produce changes in support for a potential conflict. The experimental instrument is adapted a known real- life case of information market intervention involving the United States and Muammar Gadhafi’s Libya late in the Cold War, while the persuasive modes themselves are taken from the 2003 “Hidden Hand” scandal that I described in Chapter 2. While I do not find direct evidence for the ability of former military officers to persuade voters of the cause for war, I do identify a clear distinction between the sources of voter attitudes for prospective and retrospective evaluations of foreign policy performance. Retrospective evaluations, which are commonly used in observational research designs because of the need to rely on mass polling data for empirical identification,² are well-predicted by voters’ own partisan affiliation. Democratic respondents answered questions that referred specifically to President Obama’s handling of foreign relations much more positively

¹All research for this chapter was conducted with approval by the University of Michigan Institutional Review Board (Project # HUM00044024).

²Most famously, Gallup’s presidential job approval question, which is at the heart of the rally ‘round the flag research agenda that I presented at the end of Chapter 2

than did Republican respondents, and this is of course entirely consistent with our understanding of political attitudes in other contexts.

However, two questions that referred specifically to the use of force against a rival state, without mention of the presidential administration, were not associated with respondents' partisan affiliation. For these two questions, the respondents' self-reported evaluation of the Iraq War were much more important.³ I present results from a series of model restriction tests that show that partisanship and sentiments toward the Iraq War were in fact separate attitudinal dimensions among my respondents.

4.1 Prior Experimental Work

In the earlier literature review, I summarized a number of studies that examined the electoral connection underlying American foreign policy. These works established rough correspondences between international events and certain polling responses in the American mass public, such as the ability to discern between presidential candidates' positions (Aldrich, Sullivan and Borgida 1989) or identify the most dangerous rival facing the United States. And, of course, a number of scholars have attempted to extend the "rally 'round the flag" phenomenon to a host of international interactions, with mixed success; some have found an even more responsive public by examining public reactions to rally events on a partisan basis (Baum 2002).

However valuable these studies, they tell us relatively little about the public's appetite for war itself. As the earlier discussion of crisis bargaining literature in Chapter 2 revealed, the extent of support for the use of coercive force is crucial for understanding how crises are resolved. The response of voters to foreign policy platforms during peacetime, and their sensitivity to the costs of war once it begins, are both important areas of inquiry but neither help us understand the role of public

³I discuss the questionnaire design in greater detail below; the two Iraq-specific questions asked whether the "Iraq War was worth the cost" and whether "The Iraq War was a good idea," both of which were taken from the Iraq War Casualty Experiment developed by Berinsky (2009).

consent during a crisis interaction.

A small but growing number of studies have attempted to address the public's appetite for war. Some, like Herrmann, Tetlock and Visser (1999), present respondents with any of a series of real-world events and then ask whether the use of force was justified. Herrmann, Tetlock and Visser explore various combinations of real-world American allies and adversaries and manipulate certain key elements of state power, such as nuclear arsenals, in an attempt to identify those contexts in which the American public will in fact support the use of force. Their findings, in brief, are largely consistent with Jentleson's "pretty prudent public:" their respondents are more likely to support the use of force to defend a friendly ally, like Israel, than an erstwhile stranger, like Cambodia.

However, the bulk of Herrmann et al's findings are drawn from highly stylized treatments in which respondents are simply told that, say, one country has invaded its neighbor and that American interests are indeed at stake. Then they are asked whether they would support an intervention. This is a common approach as it avoids inadvertently priming respondents with other considerations beyond what the experiment itself is meant to evaluate (e.g. Tomz 2007). The cost of this approach, of course, is external validity. Voters are rarely told in such direct, terse, and authoritative terms what the two or three most salient points are about an issue and then prompted for their opinion.

In order to create my own experimental manipulations, I returned to the Pentagon's military analyst program described in Chapter 2. Ultimately, information market intervention is an argument about well-timed and strategically chosen attempts at political persuasion. As described by Lupia (2002), persuasion requires two separate conditions. The first is a reputation for expertise, or a belief on the part of the persuaded that the actor persuading them actually has knowledge that they ought to listen to. The second is a reputation for honesty or reliability. In order

for persuasion to be successful, the persuaded party has to believe that the actor persuading them is revealing what they truthfully know about the state of the world.

As it happens, the military is widely respected across the American political spectrum, and in this regard the Defense Department’s analyst program was either cannily assembled or particularly fortunate. While the most liberal respondents in the General Social Survey time series are slightly less trusting in the military, and the most conservative respondents slightly more trusting, these differences are fairly slight. Persuasive speech attributed to decorated military officers is likely to be effective.

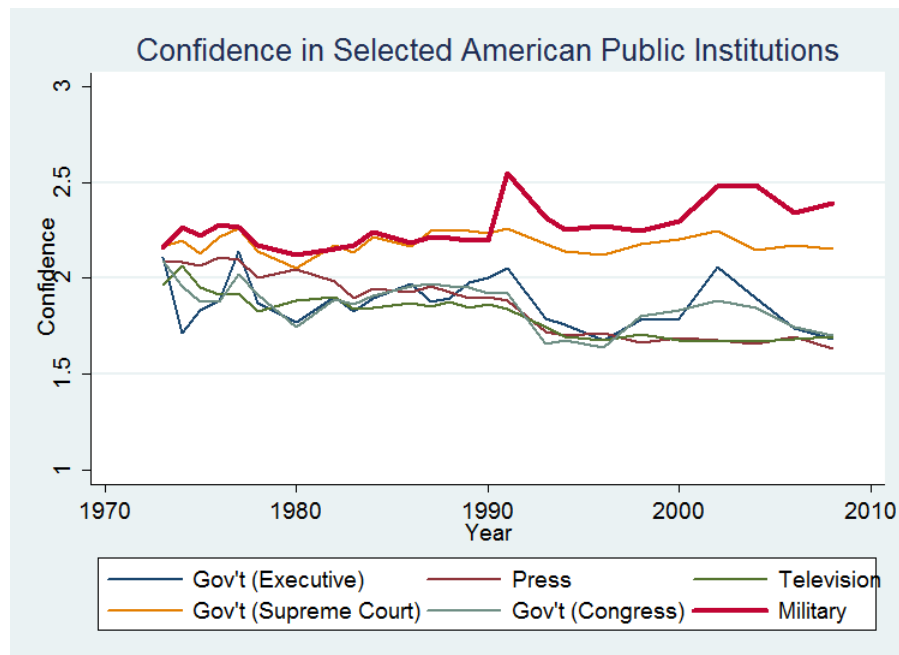


Figure 4.1: Confidence in various government and media institutions, from the General Social Survey. In the cumulative data file spanning 1972 to 2008, confidence in institutions was measured with just three categories. “A great deal” of confidence is coded 1, while “only some” confidence is coded 2 and “hardly any” confidence is labeled 3. The trend lines indicate the mean confidence rating across all respondents in each year the question was asked.

A second factor widely associated with political persuasion is emotional anxiety. In their path-breaking *Affective Intelligence and Political Judgment*, Marcus, Neuman and MacKuen (2000) argue that different types of cognition and reasoning are

engaged by voters' different emotional reactions to the world around them. In particular, anxiety is associated with the tendency to seek new information about an issue; in other words, anxiety and fear are associated with a greater propensity to be persuaded. By contrast, appeals that produce a sense of enthusiasm in the recipient are associated with a wholly different kind of reasoning, in which people are satisfied with their current beliefs and understanding of the world, making them less likely to be persuaded. Anecdotally, of course, media coverage of security policy is extremely fear- and anxiety-centric, making this a natural factor to consider in the search for manipulated public opinion.

4.1.1 Treatment Design and Hypotheses

In order to test the influence of these factors, I rely on a separate historical case of information market intervention. I derive four separate test conditions through alterations of an *Wall Street Journal* news story describing a confrontation between the United States and Libya in August 1986. I picked this particular story because it was originally planted in the American media by the National Security Council (Woodward 1986).

It also bears a few hallmarks of an attempt to grow support for a new confrontation with the Gadhafi government: it relies on a variety of anonymous quotes detailing the danger posed by Libya, and paints Gadhafi as an exceptionally dangerous, irrational adversary. Most importantly for my argument about the importance of pre-crisis “information market intervention,” it did not occur in conjunction with a crisis as defined by either the Militarized Interstate Dispute or International Crisis Behavior data. Indeed, while the United States had bombed Libya in April of that year, killing members of Gadhafi's immediate family, the two states are not defined as being in crisis in August, and the content of the article strongly suggests that it was meant to engender support for a second round of military strikes. An anonymous source in

Gadhafi Target of Secret U.S. Deception

Elaborate Campaign Included Disinformation That Appeared as Fact in America

By Bob Woodward
Washington Post Staff Writer

In August the Reagan administration launched a secret and unusual campaign of deception designed to convince Libyan leader Moammar Gadhafi that he was about to be attacked again by U.S. bombers and perhaps be ousted in a coup, according to informed sources and documents.

The secret plan, adopted at a White House meeting on Aug. 14, was outlined in a three-page memo that John M. Poindexter, the president's national security affairs adviser, sent to President Reagan.

"One of the key elements" of the new strategy, the Poindexter memo said, "is that it combines real and illusionary events—through a disinformation program—with the basic goal of making Gadhafi *think* [word underlined in the original] that there is a high degree of internal opposition

to him within Libya, that his key trusted aides are disloyal, that the U.S. is about to move against him militarily."


It was an elaborate plan: "a series of closely coordinated events involving covert, diplomatic, military and public actions," according to Poindexter's memo. Military officers expressed some reservations about the plan, and intelligence specialists were deeply divided about its potential efficacy. The plan was the latest phase of the administration's policy, first adopted last year, to try to topple Gadhafi, a known instigator of terrorist acts targeted by the administration as a threat that has to be removed.

Beginning with an Aug. 25 report in The Wall Street Journal, the American news media—including The Washington Post—reported as fact much of the false information generated by the new plan. Published articles described renewed Libyan backing for terrorism and a looming, new

U.S.-Libya confrontation. Few officials had actually concluded Gadhafi was "quiescent" on "confrontation" was the on administration plan, according to administration planning papers.

During September, however, agencies assembled evidence begun planning a significant attacks, and some senior of that this is in part a response to the latest campaign against Gadhafi. The administration's latest campaign against Gadhafi is still considered reliable but still in a direct hand in the Sept. 5 American World Airways Flight airport in Pakistan and pre-

See LIBYA, A1



JOHN M. POINDEXTER
... "real and illusionary events"

Figure 4.2: "Information market intervention" revealed.

Collision Course
**New Signs That Libya
Is Plotting Terrorism
Bring Quick Response**

**U.S. Readies Air-Raid Plan,
Three-Pronged Program;
Naval Maneuvers Begin**

Looking for a Smoking Gun

By JOHN WALCOTT
And GERALD F. SEIB
Staff Reporters of THE WALL STREET JOURNAL
The U.S. and Libya are on a collision
course again.

Figure 4.3: The opening lines of the NSA-authored leak, as reported in the *Wall Street Journal*, August 25th, 1986.

the story goes so far as to suggest that the “Libyans might need another lesson” as well as a reminder of the April campaign.

Once the deception was revealed, the *Journal* was forced to run a correction to its original report:

The reporting which produced our Aug. 25 story came from multiple sources in multiple agencies of the U.S. government as well as several foreign governments, including material gathered by Seib. The reporting turned up much indisputable information, including the plan to send an emissary to Europe to seek new sanctions, and discussed the difficulties in pinning the exact blame for terrorist events. The Journal subsequently revealed in a story on Sept. 2 that the U.S. planned to ‘promote reports in the Middle East of growing opposition to Gadhafi.

If our government also mounted a complex disinformation campaign, involving multiple sources here and abroad aimed at the U.S. press, we knew nothing about it. If, indeed, our government conducted such a domestic disinformation campaign, we were among its many victims.

I test two general propositions in this experiment. The first is whether persuasive appeals in general, of the kind reminiscent of the “Pentagon’s Hidden Hand” scandal, are effective in generating support for conflict. The second is whether persuasive appeals that are built around military expertise and appeals to anxiety are uniquely

effective in doing so. Including a pure control condition leads to a four-cell design, as follows:

1. Control: No story is presented to the subject, and they are immediately directed to the foreign policy questionnaire.
2. Information Only: A short, factual story is presented to the subject, which does not identify a credible expert or include anxiety-provoking language.
3. Civilian Expert: The story presented in the “Information Only” condition is buttressed by the testimony of an academic foreign policy scholar, who provides the rationale for a conflict as given in the planted *Wall Street Journal* story.
4. Military Expert: The story is presented as in the “Civilian Expert” condition, but the foreign policy expert is now identified as a retired Army officer. In addition, reference is made to the same anxiety-producing language as in the original planted story.

This design implies a few easily-tested propositions:

- H1: Persuasive appeals regarding the danger of a potential adversary will produce greater support for conflict. (The “civilian expert” and ”military expert” cells will both report greater support for military escalation and engagement.)
- H2: The presence of an emotional appeal (highlighting threat / danger) will produce more support for escalation than in the cases where it is absent, all things equal. (The ”military expert” cell will also report greater support for military escalation than the “civilian expert” cell.)

Note here that the treatment assignments do not scale in a linear fashion. In particular, the two persuasive appeals are separated on two distinct dimensions. Respondents in the fourth cell were given expert analysis by a former military officer, rather than a civilian, and with the presence of additional combustible language taken from the original *Journal* article. This reflects a series of pilot tests that I conducted in late 2010 with an undergraduate sample drawn from the Department of Communications at the University of Michigan. In those studies, I fielded a true 4-cell factorial design in which the expert’s identity and the presence or absence of incendiary language were both manipulated. Because those studies found no meaningful effect in the subjects’ eventual support (or opposition) to a war, I collapsed the two treatment

dimensions here in order to produce a stimulus that, *ex ante*, was much stronger than in these previous tests.

I amended the *Journal* story to describe tensions between the United States and Iran over the latter’s nuclear weapons program. This was initially meant to avoid the possibility of respondents knowing that Muammar Gadhafi was still in charge of Libya, and by implication that the American military threat against him described in the stimuli had failed. That proved to be a better decision than I initially realized after the Libyan revolt began in February 2011.

4.2 Experimental Procedure

The above design was administered to 1,262 adults between June 16th and July 2nd, 2011. These respondents were recruited through Amazon Mechanical Turk, a popular tool used for “crowdsourcing” large data-intensive tasks that are not well-suited to computational solutions, such as semantic content analysis. Mechanical Turk⁴ has come into vogue as a low-cost, quick-turnaround method to field survey experiments. Studies of this sort can be fielded for as little as 50 cents per respondent, compared to \$5 - \$10 when fielded through popular survey companies, or even more for in-person convenience samples.⁵ The treatments and questionnaire themselves were hosted by Qualtrics, who provide an easy-to-learn graphical interface for designing web surveys.

Respondents, or “Workers” in the MTurk parlance, are presented with a list of tasks that are currently recruiting participants and for which they are qualified when they log in to the MTurk web client. For the purposes of this study, Workers were required to be at least 18 years of age and located in the United States, to provide

⁴Henceforth, “MTurk.”

⁵For more on using Mechanical Turk for web-based survey experiments, see Adam Berinsky, Gregory Huber and Gabriel Lenz’s excellent introduction: http://huber.research.yale.edu/materials/26_paper.pdf

at least a rough facsimile of the American voting age population.⁶ MTurk tasks are given a short text description and a series of keywords to aid Workers in finding jobs for which they are suited. Both of the experiments described below were advertised as a “short survey about politics,” and participants were not informed until their debriefing that the study included random assignment to a treatment condition as well as deception in those treatments’ content.

After advancing through a short disclaimer, the subjects were told that:

“The following news article describes an ongoing international dispute between the United States and a foreign country. After reading the article, we will ask about your reaction and whether or not you support the particular approach that this leader has adopted.”

and then presented with a short vignette fitting each of the treatments described above. For example, the “military expert + anxiety-provoking language” stimulus read as follows:

Although I did not include an explicit compliance check in the questionnaire, the web survey software I used did allow me to record how much time subjects spent with the news story in their browser window. Figure 4.5 reports a “box-and-whisker” plot of the median time spent on the news story and the other usual percentile ranks, separated by treatment group.

After reading a story, subjects were presented with five questions, each of which was meant to probe for different conceptualizations of the degree of support they would offer for a potential conflict with Iran. The first two questions framed the issue on its own merits; the final three questions were designed to look more like the sort of proxy variables on which empirical studies of public opinion on foreign policy typically rely. Questions 2 through 4 were presented with the same sort of branching logic that

⁶Of course, both of these are easy to fake, though IP addresses can be logged and checked against internet registrars to determine physical location. Berinsky, Huber and Lenz find that while MTurk samples are reliably younger and more female than the general population, that these differences are smaller than one might otherwise expect.

Iran Suspected Over Nuclear Program

WASHINGTON (UPI) --- An unnamed source said Thursday that Iran's nuclear program is a cover for clandestine military activity. The comments were likely to increase pressure on Iran amid the international dispute over its nuclear activities, which Tehran insists are purely civilian but U.S. leaders fear are aimed at building nuclear weapons.

U.S. Army Maj. General Anthony Archer (ret.), formerly a deputy commander of the Pacific Command and a frequent contributor to United Press International on matters of foreign policy, speculated that any effort would combine covert activities and tightening the economic screws on Iran with plans for new military action.

Back in Business

General Archer told United Press International that the Army and military intelligence agencies have become convinced that the volatile Iranian leader is moving back into the nuclear business. "In the event that new movements of nuclear material are traced to Iran, the Pentagon will need to have a contingency plan in place," Archer said.

General Archer described a concentrated effort by the United States to exploit what is seen as Ahmadinejad's domestic vulnerability, melding

unspecified covert activities and efforts to encourage U.S. allies to tighten an economic embargo against Iran to undermine his rule.

Contingency Plans

Archer's position has been supported by a group of senior Air Force officers, who told UPI that "We know the Iranians may need another lesson." They went on to say that the unpredictable Ahmadinejad "seems to have gone off his rocker again."

The report coincided with the beginning of U.S.-Iraqi joint military exercises in the Persian Gulf. However, U.S. officials said that the exercises would not be conducted near the Iranian border at al-Faw.

The U.N. Security Council is to consider Iran's nuclear program next month. France, Britain and Germany have led European negotiations that have failed to persuade Iran to suspend parts of its nuclear program.

Iran confirmed on Monday that it had resumed small-scale uranium enrichment last week, and Iran's President Mahmoud Ahmadinejad visited the Natanz uranium enrichment plant Wednesday.

Figure 4.4: Stimulus for the "Military Expert" condition.




Figure 4.5: Time spent viewing stimuli, by treatment group, June 2011.

is common to academic polling, such as the American National Elections Study. For example, Question 3 would initially be posed to a respondent as simply whether they “approve” or “disapprove” of the president’s policy toward Iran, and then depending on whether they voiced approval, disapproval or ambivalence would be asked to refine their position as “strong,” only “somewhat,” or whether they leaned in one direction or another in the case of an ambivalent initial response.

1. In your opinion, does Iran represent a threat to U.S. national security? (0=No; 1=Yes)
2. In light of recent events, do you support or oppose the use of military force against Iran in order to prevent it from collecting nuclear material? (1=Strongly Oppose ... 4=Strongly Support)
3. Do you approve or disapprove of the president’s handling of relations with Iran? (1=Strongly disapprove...4=Neither approve nor disapprove...7=Strongly approve)
4. Do you trust or distrust the president’s conduct of foreign policy in general? (1=Distrust a great deal...4=Neither trust nor distrust...7=Trust a great deal)

5. Would you say that the current presidential administration has made the United States much more secure from its foreign enemies, less secure, or hasn't it made much difference either way? (1=Made a great deal less secure...4=Hasn't made much difference...7=Made a great deal more secure)

After answering the above, respondents were asked an additional series of control questions, including demographic characteristics, partisan identification, and their self-reported interest in news and politics. They were also asked two specific questions about their thoughts on the Iraq War, a handful of objective political knowledge questions,⁷ and how often they view Fox News, if applicable. This latter question tracks with a growing number of studies (e.g. Baum and Groeling (2008a)) that suggest Fox viewership carries its own explanatory power for conservative attitudes, above and beyond identifying as a Republican partisan. Summary statistics for both the demographic and political attitude controls follows in table 4.1. While some of the control measures are similar to what Berinsky, Huber and Lenz found in their benchmark study of Mechanical Turk, such as years of education and political interest, the sample as a whole differs in some substantial ways from those findings. The sample here is substantially less female, with roughly 50% of the subject pool women rather than the 60% found by Berinsky and colleagues, and more racially diverse as well. Blacks comprise nearly 8% of the sample, which is only half of their representation in the real American population but twice the contribution found by Berinsky et al. This sample is also more diverse with respect to partisanship, with the mean political identification in each group falling almost halfway between Democrats and Republicans.

The success or failure of random assignment along these dimensions can be estimated through an F -test, which reports the likelihood that the means of each variable are the same across each of the four treatment groups. The pertinent values for this

⁷The Iraq War questions, as mentioned before, were taken from Berinsky (2009). The political knowledge variable is a five point index, and the individual questions were split between objective knowledge about American government and then-current events in the Middle East.

Treatment	n	Party ID	Political Know.	Political Interest	Iraq Interest	Iraq Right?	Iraq Worth it?	
Baseline	326	mean	3.80	3.01	2.53	2.08	1.88	
		sd	1.28	0.86	0.76	1.22	1.17	
Minimal info.	317	mean	3.74	2.95	2.43	1.96	1.80	
		sd	1.34	0.93	0.81	1.18	1.10	
Mil. Exp.	309	mean	3.79	3.01	2.45	2.14	1.84	
		sd	1.31	0.85	0.77	1.19	1.13	
Civ. Exp.	310	mean	3.85	3.06	2.36	2.14	1.88	
		sd	1.25	0.87	0.78	1.25	1.15	
Treatment		Fox News Viewer	Fox News Freq.	Gender	Black	Hispanic	Age	Education
Baseline		mean	4.43	0.51	0.07	0.06	31.03	14.71
		sd	0.50	0.50	0.26	0.24	10.55	2.37
Minimal info.		mean	4.22	0.45	0.06	0.07	30.29	14.79
		sd	0.49	0.50	0.24	0.25	10.28	2.21
Mil. Exp.		mean	4.53	0.52	0.09	0.07	31.11	14.53
		sd	0.49	0.50	0.29	0.26	10.87	2.09
Civ. Exp.		mean	4.42	0.49	0.08	0.06	30.87	14.70
		sd	0.49	0.50	0.28	0.25	11.13	2.13

Table 4.1: Means for control variables, separated by treatment condition, for June 2011 (Iran) data.

test appear in table 4.2. By traditional standards, we can maintain the null hypothesis that there is no difference in the true means for any of these variables across the four treatment groups with one exception: self-reported interest in ongoing developments in Iraq. The differences between groups in this regard are quite modest, however. Respondents assigned to the “no information” condition reported being somewhat more interested in events in Iraq, while respondents who were assigned to the “civilian foreign policy expert” condition reported themselves to be somewhat less interested than their peers.

	F statistic	P value
Party identification	1.57	0.20
Knowledge scale (index)	0.41	0.75
Political Interest	0.85	0.47
Interest in Iraq	2.52	0.06
Iraq the Right Thing?	1.44	0.23
Iraq Worth It?	0.35	0.79
Fox News viewer	0.31	0.82
Fox News frequency	0.93	0.43
Gender	1.04	0.38
Black	0.98	0.40
Hispanic	0.16	0.92
Age	0.37	0.78
Education	0.76	0.52

Table 4.2: Randomization check for June 2011 data. The critical value from the F-distribution characterizes a null hypothesis that the mean of each variable is identical across all 4 treatment groups.

4.3 Results (Part 1)

4.3.1 Means by Treatment Conditions

With no other apparent issues with randomization, the most straightforward analysis of the results involves examining the sample means for each treatment. Table 4.3 shows the mean by treatment condition for each of the five dependent variables de-




Figure 4.6: Self-reported interest in news from Iraq, June 2011

scribed above.

Treatment		Iran a threat?	Escalate	Handle relations	Trust	Secure
Control	mean	0.63	2.38	3.97	4.35	4.45
	sd	0.48	1.03	1.74	2.12	1.74
Info. Only	mean	0.77	2.32	4.00	4.38	4.40
	sd	0.42	1.03	1.62	2.07	1.53
Military Expert	mean	0.74	2.48	4.09	4.48	4.39
	sd	0.44	1.02	1.58	1.89	1.54
Civilian Expert	mean	0.73	2.42	4.26	4.51	4.35
	sd	0.45	1.05	1.52	2.00	1.51

Table 4.3: Means and standard deviations of five dependent variables, by treatment condition.

Only a cursory examination is needed to determine that the subjects' responses do not seem to be adhering to the hypotheses I laid out before. While the control condition is associated with relatively pacific responses to the question of whether Iran presents a threat to national security, respondents assigned to the control condition are indistinguishable from their counterparts across the other cells of the design for the remaining dependent variables. Likewise, respondents who were presented with

persuasive testimony from a military expert did not express any greater support for a new conflict than did those who were instead presented with a civilian expert.

4.3.2 Econometric Estimations

That in mind, I proceed with a series of simple regressions to determine whether the effect of each treatment, or individuals’ information environments, could be separated from the effect of other political traits that might be expected to influence political attitudes. I coded dummy variables to distinguish between treatment conditions in the manner shown in Table 4.4 and estimate the following equation using ordinary least squares regression:⁸:

Group	Info.	Intervention	Mil. Exp.	n
No Story	0	0	0	326
Minimal Desc.	1	0	0	317
Civ. Expert	1	1	0	309
Mil. Expert	1	1	1	310

Table 4.4: Three dummy variables to distinguish between four experimental conditions (the control condition serves as a baseline).

$$\text{Outcome} = \beta_0 + \beta_1 \times \text{PartyID} + \beta_2 \times \text{Info} + \beta_3 \times \text{Intervention} + \beta_4 \times \text{Mil. Expert}$$

$$+ \sum_{i=5}^{13} \beta_i \times X_i$$

The vector of control variables X_i includes self-reported general political interest, post-hoc judgments about the Iraq War, self-reported interest in ongoing media coverage of Iraq, Fox News viewership, the additive political knowledge scale, gender, race (white vs. all non-whites) and a categorical education variable. The coefficients for

⁸This means that the dichotomous “Is Iran a threat?” variable is estimated under the assumption of linear, as opposed to non-linear, probability. The linear probability model coefficients are easier to interpret, and the same set of factors are deemed statistically significant whether or not a linear or logistic model is fitted.

this set of estimations are presented below, for each of the five dependent variables.

The results in Table 4.5 demonstrate that even after controlling for an array of trait variables, the treatment conditions seem to matter relatively little for any of these measures of support for a conflict or of trust and satisfaction with the presidential administration's foreign policy. Being exposed to any of the three news stories causes respondents to be quite a bit more concerned about a potential threat from Iran, but receiving either of the experts' persuasive appeals has no additional effect. To the extent that these persuasive appeals were effective, they worked in opposite directions; taken together the persuasive appeals made respondents more trusting of the administration's policy, but those receiving an appeal from a military expert were actually less trusting, supportive, or fearful of an Iranian threat than were those who received an appeal from a civilian.

One other possibility is that the effect of these treatments is not just moderated by trait variables, but is in fact conditional on them. To explore this possibility, I estimate the same models above with the addition of interactive terms between party identification and the three dummy variables representing treatment conditions:

There are some encouraging signs here—the coefficients on the “Military Expert” dummies are somewhat larger in magnitude, albeit in the wrong direction—but of course in a model with interaction terms the effect of any one factor has to be interpreted as the sum of its main and partial effects. Figures 4.7 through 4.11 show graphs with 95% confidence intervals for the effect of each treatment dummy on respondents' likelihood of answering each measure of support for the administration's foreign policy, across the self-reported left-right partisan self-identification space. In each figure, subfigure A presents the marginal effect of receiving any news story at all, while B and C show the effect of receiving a persuasive appeal and then finally a persuasive appeal from a military expert, respectively.

Many of the marginal effects graphs merely confirm that the treatments themselves

Variables	Iran Threat		Escalate		Handle Relations		Trust		Secure	
	(ols)	b	(ols)	b	(ols)	b	(ols)	b	(ols)	b
Information Dummy	0.15***		0.00		0.00		-0.06		-0.14	
Intervention Dummy	-0.05		0.02		0.31*		0.23		0.07	
Military Expert Dummy	0.00		0.04		-0.26*		-0.16		-0.09	
Party ID (unit)	-0.06		0.07		-2.09***		-3.21***		-2.06***	
Gen. Political Interest	-0.01		-0.06		-0.04		-0.11		-0.08	
Iraq War Right?	0.07***		0.28***		0.09		0.03		-0.02	
Iraq War Worth It?	0.05**		0.22***		0.13*		0.06		0.04	
Knowledge Index	0.04***		-0.04*		-0.04		0.02		-0.03	
Iraq Interest	0.01		0.02		-0.01		-0.02		0.03	
Fox News Viewer	0.09**		0.10		0.02		-0.18		-0.11	
Male	-0.03		0.01		0.10		0.30**		0.10	
White	0.02		-0.18**		-0.33**		-0.49***		-0.23*	
Age (Years)	0.00		-0.00		0.00		-0.00		-0.01**	
Education (Cat.)	0.01		-0.01		-0.02		0.10*		0.01	
Constant	0.14		1.82***		4.93***		5.90***		6.03***	
r2	0.12		0.35		0.16		0.28		0.20	
N	1207		1207		1207		1203		1205	

Table 4.5: Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions (*: $p < .05$; **: $p < .01$; ***: $p < .001$)

Variables	Iran Threat (ols) b	Escalate (ols) b	Handle Relations (ols) b	Trust (ols) b	Secure (ols) b
Information Dummy	0.16** (0.06)	-0.09 (0.11)	-0.05 (0.20)	0.12 (0.23)	-0.20 (0.19)
Intervention Dummy	-0.06 (0.06)	-0.07 (0.11)	0.27 (0.20)	0.24 (0.23)	0.06 (0.19)
Military Expert Dummy	0.02 (0.06)	0.22* (0.11)	-0.36 (0.20)	-0.56* (0.23)	-0.17 (0.19)
Party ID (unit)	-0.05 (0.08)	-0.09 (0.16)	-2.30*** (0.28)	-3.13*** (0.33)	-2.24*** (0.27)
Info X Party ID	-0.03 (0.11)	0.23 (0.21)	0.12 (0.38)	-0.44 (0.44)	0.15 (0.37)
Intervention X Party ID	0.04 (0.11)	0.21 (0.22)	0.11 (0.38)	0.02 (0.44)	0.03 (0.37)
Mil. Expert X Party ID	-0.05 (0.11)	-0.45* (0.22)	0.27 (0.39)	0.99* (0.45)	0.22 (0.37)
Gen. Political Interest	-0.01 (0.02)	-0.05 (0.03)	-0.04 (0.06)	-0.13 (0.06)	-0.08 (0.05)
Iraq War Right?	0.07*** (0.02)	0.28*** (0.03)	0.08 (0.06)	0.02 (0.06)	-0.02 (0.05)
Iraq War Worth It?	0.05** (0.02)	0.22*** (0.03)	0.14* (0.06)	0.07 (0.07)	0.05 (0.06)
Knowledge Index	0.04*** (0.01)	-0.05* (0.02)	-0.04 (0.04)	0.02 (0.04)	-0.03 (0.04)
Iraq Interest	0.01 (0.02)	0.02 (0.03)	-0.01 (0.06)	-0.02 (0.06)	0.03 (0.05)
Fox News Viewer	0.09** (0.03)	0.11* (0.05)	0.02 (0.10)	-0.18 (0.11)	-0.11 (0.09)
Male	-0.03 (0.03)	0.00 (0.05)	0.10 (0.09)	0.32** (0.10)	0.10 (0.08)
White	0.02 (0.03)	-0.19** (0.07)	-0.34** (0.12)	-0.48*** (0.14)	-0.23* (0.11)
Age (Years)	0.00 (0.00)	-0.00 (0.00)	0.00 (0.00)	-0.00 (0.01)	-0.01** (0.00)
Education (Cat.)	0.01 (0.01)	-0.01 (0.02)	-0.01 (0.03)	0.10* (0.04)	0.01 (0.03)
Constant	0.14 (0.09)	1.89*** (0.17)	5.02*** (0.30)	5.88*** (0.34)	6.11*** (0.28)
r2	0.12	0.35	0.16	0.29	0.20
N	1207	1207	1207	1203	1205

Table 4.6: Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions and interactive effects between treatment conditions and partisan identity (*: $p < .05$; **: $p < .01$; ***: $p < .001$)




Figure 4.7: Marginal effects of treatment assignment, conditional on partisanship, for “Iran a threat?”




Figure 4.8: Marginal effects of treatment assignment, conditional on partisanship, for “Support escalation?”




Figure 4.9: Marginal effects of treatment assignment, conditional on partisanship, for “Handle relations?”




Figure 4.10: Marginal effects of treatment assignment, conditional on partisanship, for “Trust the president?”




Figure 4.11: Marginal effects of treatment assignment, conditional on partisanship, for “Made more secure?”

appeared not to make much difference to the respondents' perceptions of the need for conflict with Iran or their evaluation of the president's general foreign policy performance, and all suffer from a relative dearth of self-identified strong Republicans that would strengthen any estimated relationships between positions and partisanship.

Figure 4.7a restates the finding from the regression table above that respondents in any informative treatment condition were moved to see Iran as more of a threat than those in the baseline control group, which is consistent with a simple application of priming theory. A few other intuitive possibilities exist, such as the potential shown in figures 4.9a, 4.11a, and 4.10a that the most Republican-leaning subjects were somewhat more willing to respond positively to the (Democratic) administration on the three retrospective measures after being exposed to a former military officer's testimony, but again, there were too few of those respondents in the general sample to be confident about such a conclusion.

However, the earlier regression results showed that the subjects' evaluations of the Iraq War did seem to have a relationship with potential support for a conflict with Iran. Respondents who viewed the Iraq War more favorably were more likely to view Iran as a threat, and more likely to support the potential use of force against same. These views of Iraq did not, by contrast, seem to matter for perceptions of the current president's general foreign policy performance.

This finding is consistent with an earlier generation of research on the rational or prudent public (Jentleson 1992, e.g.), which claimed that American voters were in fact able to pass judgment on specific conflict areas and disputes and express opinions that reflected their evaluation of the international stage. However, more recent work has been clear in finding that American's positions in this issue are merely another vehicle by which their extant partisanship is demonstrated (Berinsky 2007; Baum and Potter 2008). And, of course, to the extent that Americans' views of the Iraq War are informed by their partisanship, it might be foolhardy to view these as separate

dimensions in the first place. In other words, the relationship between partisanship, Iraq, and prospective versus retrospective evaluations of policy performance might be driven by multicollinearity.

Since the simple correlation score between the two Iraq War measures and party identification was indeterminate, I performed a series of likelihood ratio tests to determine whether they were, in fact, separate dimensions. Likelihood ratio tests compare goodness-of-fit criteria between full and reduced models to produce a χ^2 test statistic. This statistic characterizes whether the omitted variables in fact reduce the original model’s explanatory power.

Pairwise Correlations	”Iraq Right?”	”Iraq Worth It?”
Party Identification	0.41	0.41

Table 4.7: Pairwise correlation scores for party identification and evaluations of the Iraq War, for June 2011 panel

Panel:	June 2011 (Iran)					
DV	Removing Party ID			Removing 2 Iraq Measures		
	Δ LL	LR χ^2	$p > \chi^2$	Δ LL	LR χ^2	$p > \chi^2$
Iran a threat?	-0.96	1.91	0.17	-49.42	98.84	0.00
Support escalation?	-0.34	0.67	0.41	-167.88	335.76	0.00
Handle relations	-82.84	165.67	0.00	-11.49	22.97	0.00
Trust	-139.53	279.06	0.00	-1.46	2.91	0.23
Secure	-87.81	175.61	0.00	-0.32	0.63	0.73

Table 4.8: Model restriction tests for June 2011 panel. Both sets of restrictions tested against the full regressions shown above. “Is Iran a threat?” response modeled with logistic estimator, rather than linear.

Estimations of the “threat” and “escalate” questions suffer from the absence of the Iraq War attitude controls, but not from the absence of self-reported party identification. By contrast, the three retrospective measures exhibit complimentary behavior. It does not appear that evaluations of Iraq were simply proxies for partisanship in this sample.

4.4 Results (Part 2)

If this distinction holds, then it would be an interesting finding. Recent scholarship in both the experimental (Berinsky 2007) and observational (Baum and Groeling 2008*a*) modes suggests that foreign policy attitudes are determined by party loyalties. Voters are thought to observe the positions of their favorite political partisans, and then adopt the same foreign policy positions as those elites in order to confirm their party bona fides. The particulars of a given policy, or clear attribution of credit or blame for policy success or failure, matter less than do voters' personal attachments. In such a case, the extent to which the public can place constraints on their leader's foreign policy decision-making is less a question of policy success or failure than it is about the durability of the leader's coalition. Even in the case of a failed policy, party loyalists might believe that blame truly lies with the opposition, and refuse to inflict any punishment on the incumbent.

Early in crisis periods, though, partisan cues are less accessible (Bennett 1990) and reporting is less likely to be independent. Media may not even take an interest at all in an unfolding crisis before it reaches a certain salience threshold, which would mean that the narrative surrounding a dispute is both more open to interpretation for two separate reasons: because the facts of the case are not yet in the public domain, and because cues do not yet exist to offer a framework for interpretation.⁹ Such periods might be the rare time in which foreign policy attitudes are not anchored by party politics.

The results of the "Iranian Nuclear Program" experiment described above are broadly consistent with this latter view. The "threat" and "escalate" questions were posed without an informative partisan frame, while the remainder were asked in relation to the sitting president. When cues are not available, both because the

⁹Baum and Groeling (2010) coined a particularly artful phrase to describe this dynamic: "the elasticity of reality."

conflict described is hypothetical and because no one was primed to think of the president in association with foreign policy, subjects instead turned to an alternate source of understanding in order to determine their response to the questionnaire. In order to test whether this divide between prospective and retrospective foreign policy attitudes holds up across potential rival states, I fielded the same experimental design described above a second time, between August 27th and August 29th, 2011. I recruited 1,226 subjects through Amazon Mechanical Turk in the same manner as before. The shorter window required to recruit that many subjects was made possible by offering more lucrative compensation; each eligible subject was paid \$1.25, rather than \$0.50.

This second experiment was different in one crucial respect. Instead of presenting respondents with various treatments of the threat of a nuclear Iran, the experimental treatments featured Venezuela instead. I chose Venezuela because it is a plausible rival state for the the US that has not featured nearly as prominently as Iran in the popular media in recent years, which implies that respondents will have fewer easily-accessible considerations on the matter from sources outside the experiment itself. At the same time, it is not completely implausible that Venezuela would be suspected of a nuclear weapons program. Russia agreed to build a nuclear reactor in Venezuela in the fall of 2010. While the deal was later scrapped, it nonetheless establishes a measure of external validity. The stimuli presented in this version were otherwise identical to those in the Iran edition, albeit with proper names changed as necessary to be appropriate.

I also introduced another slight manipulation. In order to test whether the failure of partisan identity to predict prospective judgments on the need for military escalation against Iran was simply a product of *not* framing the question as a referendum on the American president rather than the force-versus-approval divide identified earlier, I randomly assigned all respondents, irrespective of the core treatment they received,

to one of two versions of the questionnaire. Half of respondents saw the same wording of the “threat” and “escalate” questions as before, while the other half saw new wordings that included explicit references to the current presidential administration.

To wit:

1. “Threat”: “In your opinion, does Venezuela represent a threat to U.S. national security?” **versus** “In your opinion, does Venezuela represent a threat to the president’s national security policy?”
2. “Escalate”: “In light of recent events, do you support or oppose the use of military force against Venezuela in order to prevent it from collecting nuclear material?” **versus** “In light of recent events, do you support or oppose the president’s use of military force against Venezuela in order to prevent it from collecting nuclear material?”

The demographic traits of this second sample follow in table 4.1, with the same randomization check as before presented in table 4.10:

Here, two potential randomization failures have occurred, and unfortunately they are the same variables that had been objects of interest in the Iran study. And, even more pointedly, the group that demonstrated the greatest approval of the Iraq War happened to be those assigned to the “military expert” treatment, which I had originally hypothesized would elicit more hawkish positions toward a new use of force.

The difference in means between respondents exposed to the “military expert” treatment and all others is not dramatic, even if it is larger than the difference associated with interest in news about the war in the randomization check in the June study. The statistical significance reported by the χ^2 test is largely driven by placing more than 300 respondents in each cell. In any event, lacking any means by which I could identify whether these distinctions were in fact driven by the experimental treatment itself or simply exceptionally poor luck, I proceed with the same sequence of analyses as above. However, I do not present the confidence interval graphs for the interactive model as they they were inconclusive in the same way as the Iran study.

This sample’s support for a prospective conflict with Venezuela strongly resembles

Treatment	n	Party ID	Political Know.	Political Interest	Iraq Interest	Iraq Right?	Iraq Worth it?	
Baseline	317	mean	3.28	2.97	2.42	2.07	1.79	
		sd	1.36	0.81	0.75	1.26	1.13	
Minimal info.	306	mean	3.52	2.96	2.47	2.14	1.89	
		sd	1.93	0.87	0.79	1.26	1.14	
Mil. Exp.	305	mean	3.57	3.04	2.43	2.28	2.01	
		sd	1.95	0.85	0.78	1.26	1.23	
Civ. Exp.	298	mean	3.50	3.01	2.48	2.02	1.81	
		sd	1.88	0.86	0.76	1.19	1.11	
Treatment		Fox News Viewer	Fox News Freq.	Gender	Black	Hispanic	Age	Education
Baseline		mean	0.48	0.44	0.09	0.09	32.11	14.78
		sd	0.50	0.50	0.28	0.28	11.67	2.25
Minimal info.		mean	0.49	0.41	0.08	0.07	32.23	14.68
		sd	0.50	0.49	0.26	0.25	10.65	2.15
Mil. Exp.		mean	0.54	0.41	0.10	0.08	31.80	14.86
		sd	0.50	0.49	0.30	0.27	11.25	2.25
Civ. Exp.		mean	0.46	0.41	0.07	0.06	32.66	14.82
		sd	0.50	0.49	0.26	0.23	11.75	2.30

Table 4.9: Means for control variables, separated by treatment condition, for June 2011 (Iran) data.

	F statistic	P value
Party identification	1.39	0.25
Knowledge scale (index)	0.36	0.79
Political Interest	0.59	0.62
Interest in Iraq	0.43	0.73
Iraq the Right Thing?	2.36	0.07
Iraq Worth It?	2.22	0.09
Fox News viewer	1.29	0.28
Fox News frequency	1.40	0.24
Gender	0.36	0.78
Black	0.50	0.67
Hispanic	0.88	0.45
Age	0.30	0.83
Education	0.29	0.83

Table 4.10: Randomization check for August 2011 data. The critical value from the F-distribution characterizes a null hypothesis that the mean of each variable is identical across all 4 treatment groups.

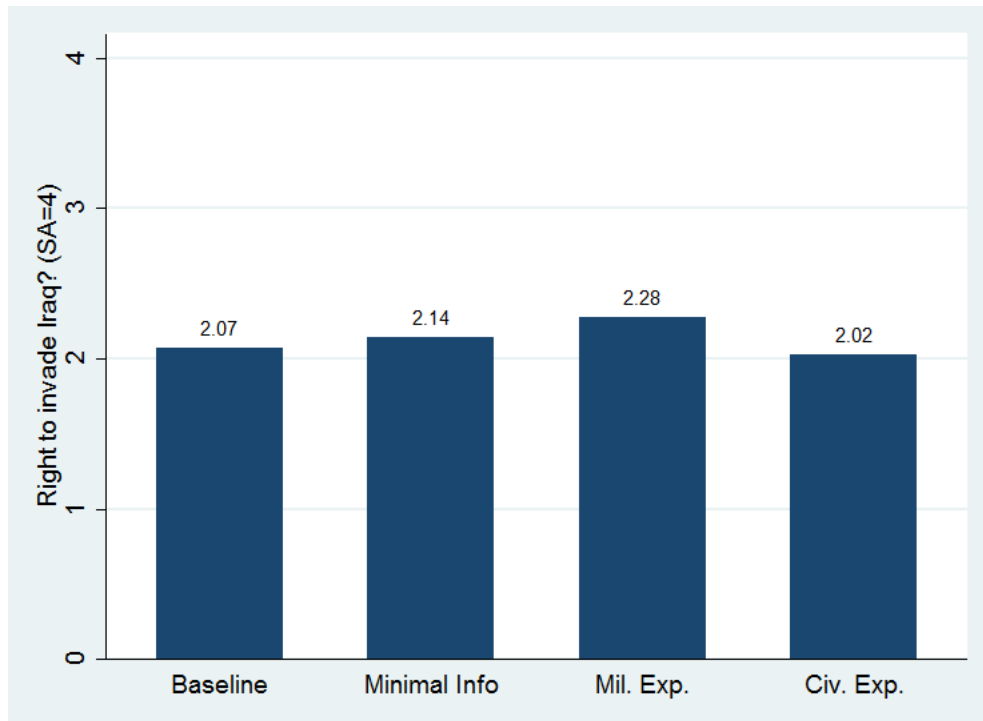


Figure 4.12: Mean response to whether or not the Iraq War was “the right thing to do,” separated by treatment groups, August 2011.

Variables	Ven. Threat (ols)		Escalate (ols)		Handle Relations (ols)		Trust (ols)		Secure (ols)	
	b		b		b		b		b	
Information Dummy	0.23***		-0.13		-0.09		-0.17		-0.29*	
Intervention Dummy	0.07		0.16*		0.14		0.18		0.34**	
Military Expert Dummy	0.01		0.07		0.06		0.03		-0.14	
Party ID (unit)	0.11*		0.02		-0.94***		-2.71***		-2.13***	
Pres. Format	0.10***		0.22***		-0.03		0.15		0.06	
Gen. Political Interest	0.02		-0.04		-0.06		-0.09		-0.04	
Iraq War Right?	0.04*		0.27***		0.25***		0.09		0.07	
Iraq War Worth It?	0.08***		0.12***		-0.02		0.10		0.06	
Knowledge Index	-0.01		-0.10***		-0.01		0.03		0.00	
Iraq Interest	-0.04*		-0.03		-0.03		-0.00		0.06	
Fox News Viewer	0.06*		0.19***		-0.05		-0.32**		-0.38***	
Male	-0.03		0.05		-0.01		0.11		0.07	
White	-0.01		0.02		-0.04		-0.18		-0.18	
Age (Years)	0.00		-0.00		0.01		0.01		-0.00	
Education (Cat)	-0.00		-0.03		-0.01		0.05		0.03	
Constant	-0.01		1.97***		4.22***		5.00***		5.03***	
r2	0.20		0.35		0.05		0.20		0.22	
N	1212		1212		1210		1209		1210	

Table 4.11: Estimations of each dependent variable from the Venezuelan experiment, with dummy variables distinguishing between treatment conditions (*: $p < .05$; **: $p < .01$; ***: $p < .001$)

Variables	Ven. Threat (ols) b	Escalate (ols) b	Handle Relations (ols) b	Trust (ols) b	Secure (ols) b
Information Dummy	0.29*** (0.06)	-0.19 (0.10)	-0.05 (0.18)	-0.33 (0.23)	-0.20 (0.18)
Intervention Dummy	0.04 (0.06)	0.04 (0.11)	0.08 (0.19)	0.03 (0.25)	0.22 (0.19)
Military Expert Dummy	0.00 (0.06)	0.07 (0.11)	-0.05 (0.19)	0.07 (0.25)	-0.17 (0.19)
Party ID (unit)	0.17* (0.08)	-0.25 (0.15)	-1.00*** (0.26)	-3.18*** (0.34)	-2.15*** (0.26)
Info X Party ID	-0.14 (0.11)	0.17 (0.20)	-0.09 (0.36)	0.42 (0.45)	-0.20 (0.36)
Intervention X Party ID	0.08 (0.12)	0.29 (0.21)	0.13 (0.37)	0.36 (0.47)	0.30 (0.37)
Mil. Expert X Party ID	0.01 (0.12)	-0.01 (0.21)	0.27 (0.37)	-0.09 (0.47)	0.05 (0.37)
Pres. Format	0.10*** (0.03)	0.22*** (0.05)	-0.03 (0.08)	0.14 (0.10)	0.05 (0.08)
Gen. Political Interest	0.02 (0.02)	-0.05 (0.03)	-0.06 (0.05)	-0.10 (0.07)	-0.04 (0.05)
Iraq War Right?	0.04* (0.02)	0.27*** (0.03)	0.25*** (0.06)	0.08 (0.07)	0.07 (0.06)
Iraq War Worth It?	0.08*** (0.02)	0.13*** (0.03)	-0.02 (0.06)	0.10 (0.07)	0.06 (0.06)
Knowledge Index	-0.01 (0.01)	-0.10*** (0.02)	-0.01 (0.04)	0.03 (0.05)	0.00 (0.04)
Iraq Interest	-0.04* (0.02)	-0.03 (0.03)	-0.03 (0.05)	-0.01 (0.07)	0.05 (0.05)
Fox News Viewer	0.06* (0.03)	0.20*** (0.05)	-0.05 (0.09)	-0.31** (0.12)	-0.38*** (0.09)
Male	-0.03 (0.03)	0.05 (0.05)	-0.01 (0.08)	0.11 (0.11)	0.07 (0.08)
White	-0.01 (0.03)	0.03 (0.06)	-0.04 (0.11)	-0.18 (0.14)	-0.18 (0.11)
Age (Years)	0.00 (0.00)	-0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	-0.00 (0.00)
Education (Cat.)	-0.00 (0.01)	-0.03 (0.02)	-0.00 (0.03)	0.05 (0.04)	0.03 (0.03)
Constant	-0.03 (0.09)	2.08*** (0.16)	4.25*** (0.28)	5.20*** (0.35)	5.05*** (0.28)
r2	0.20	0.35	0.06	0.20	0.22
N	1212	1212	1210	1209	1210

Table 4.12: Estimations of each dependent variable, with dummy variables distinguishing between treatment conditions and interactive effects between treatment conditions and partisan identity (*: $p < .05$; **: $p < .01$; ***: $p < .001$)

Panel: August 2011 (Venezuela)							
DV	Removing Party ID			Removing 2 Iraq Measures			
	Δ LL	LR χ^2	$p > \chi^2$	Delta LL	LR χ^2	$p > \chi^2$	
Iran a threat?	-2.49	4.97	0.02	-38.44	76.80	0.00	
Support escalation?	-0.02	0.03	0.85	-127.85	255.69	0.00	
Handle relations	-19.40	38.80	0.00	-18.88	37.76	0.00	
Trust	-93.64	187.28	0.00	-5.90	11.80	0.00	
Secure	-94.10	188.20	0.00	-4.61	9.22	0.01	

Table 4.13: Model restriction tests for August 2011 panel. Both sets of restrictions tested against the full regressions shown above. “Is Venezuela a threat?” response modeled with logistic estimator, rather than linear.

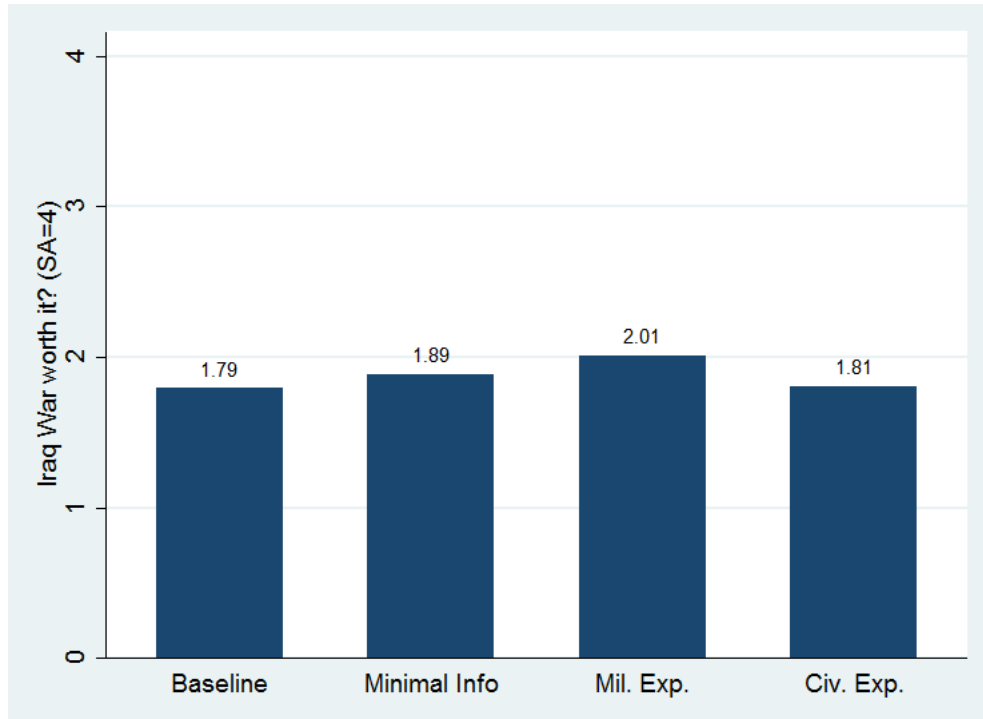


Figure 4.13: Mean response to whether or not the Iraq War was “worth it,” separated by treatment groups, August 2011.

the earlier sample’s support for a prospective conflict with Iran. Here, partisanship does predict a given subject’s perception of a threat from Venezuela, and subjects who felt more strongly that the Iraq War was justified were more inclined to approve of the president’s handling of relations with foreign states in light of the Venezuelan threat. If we set those two exceptions aside, however, the original dynamic remains: partisanship is strongly tied to general measures of foreign policy performance, while Iraq War sentiment is strongly tied to a feeling of international threat and support for the use of force against that threat.

The results from the same series of model restriction tests are quite different, however. Dropping the Iraq War measures produces a significant change in goodness-of-fit for all five outcome variables, including the three more general measures of support which they do not predict according to the estimations presented in Table 4.12. I cannot currently explain this result, though to the extent that subjects who are sym-

pathetic to the Iraq War were sorted at an exaggerated rate into the “military expert” condition, some additional explanatory power may be represented by the Iraq War variables and the large change in model fit reflects the absence of that power when those variables are dropped.

Unsurprisingly, the wording of the first two questions does seem to matter. What is surprising is that we can estimate a main effect to presidential framing that is not conditional on a subject’s own identification with either major American party. Whether or not a respondent is a Republican or Democrat, we can say with quite a degree of confidence that they are more likely to consider Venezuela a threat, and more likely to support the use of force against the country, when these questions are explicitly framed as references to the current president. This effect is clear and significant even though I have not modeled any interaction between this frame and a subject’s own partisanship. Furthermore, this effect is limited to the two items for which question wording was manipulated. Framing the first two questions about threat and support for force did not appear to cause a priming artifact that was connected to responses to the following three questions.

4.5 Concluding Thoughts

I began this chapter with the expectation that certain types of political speech, through the mass media, would shape and influence individual voters’ support for war. That expectation was borne out of several different sources. First, a casual survey of public opinion surrounding the Iraq War suggests that voters did in fact internalize several of the narratives proposed by the Bush administration in support of a war effort. More importantly, however, even those claims which were later to be shown to be untrue, such as the threat from Saddam Hussein’s weapons of mass destruction program, remained a strongly-held belief among some sections of the voting public even well after the initial phase of the war had concluded. To the extent that this

dynamic might characterize other conflicts, it suggested that world politics scholars' expectations about voter competence might be unrealistic. And, in light of what we know about voting behavior in other contexts from the discussion in Chapter 2, the notion that political messaging about the need for war ought to work seemed not just reasonable, but inevitable.

Furthermore, other historical cases clearly show that government actors believe this sort of persuasion to be possible as well. Voters did internalize beliefs about the Iraq War that were ultimately false, but American decision-makers were more than complicit in this case. The "Pentagon's Hidden Hand" scandal, with which I began this chapters, clearly demonstrates the logic of information market intervention. Now that former administration officials have begun to write their memoirs about that period, they have confirmed that these attempts to influence public opinion about war were not an accident. Press secretary Scott McClellan, for instance, told ABC News in 2008 that "in the fall of 2002, Bush and his White House were engaging in a carefully orchestrated campaign to shape and manipulate sources of public approval to our advantage ... And through it all, the media would serve as complicit enablers."¹⁰ In the course of researching potential stimuli, the 1986 NSA propaganda effort against Colonel Gadhafi came to light and also seemed to fit the overall trend.

And yet the evidence from two large online survey experiments did not support these hypotheses. The simplest explanation, of course, is that Secretary Rumsfeld, Secretary McClellan and their colleagues in the Pentagon's analyst program were wrong. Voters might not be as amenable to persuasion-by-insinuation as political agents expect. This conclusion would be a victory for other scholars in world politics who rely on voter accountability to underpin theories of domestic regimes and international security behavior. If individual voters are unmoved by appeals even by expert figures, then this would inoculate them against the sort of public diplomacy

¹⁰Accessed at<http://abcnews.go.com/blogs/politics/2008/05/mclellan-media/>

that the United States has practiced in recent years, even if that reliance on their own prior beliefs over new information implies a certain perverse conclusion for voter competence writ large.

However, if all of that were true it would be extremely surprising. Given the wealth of research that has shown that framing itself matters for opinion formation, and that source credibility and perceived expertise do seem to be important for convincing voters to adopt new positions, or at least to lend additional support for a particular policy, the notion that none of this is true for foreign policy and war beggars belief. An alternate explanation, which I am not yet able to evaluate, is that the experimental design itself was inadequate in some way to measure the effect I expected. Voters in their day-to-day lives are exposed to all manner of cues and information. Politically-active or aware voters are exposed to even more, and the metaphorical—not to mention literal—volume of these appeals are often quite loud indeed. The political message embedded in the stimuli I used here was rather quiet by comparison. Even after collapsing the source-credibility and inflammatory language conditions into a single treatment, subjects were still essentially asked to change their position about the use of force after reading just a few additional lines of text in a brief news story. The “Hidden Hand” campaign was both richer in content, as most of the paid analysts’ work was put into televised appearances, and more direct, in that what newspaper articles they were involved in tended to be direct op-ed columns. I intend to run additional installments of this research with richer, more direct, and “louder” manipulations.

All of that in mind, in the process of looking for the influence of framing and persuasion on the public’s support for war, I found what might be a more important contribution. The short-term advantage that governments have in framing foreign policy crises has been the subject of considerable speculation in the literature (Baum and Groeling 2010; Potter and Baum 2009), but seemed to be dismissed by findings

that foreign policy attitudes are so neatly consistent with individual's own partisan identities and the position of each (American) party's respective position on these issues. By contrast, when we introduce measures of support for conflict that are much more in line with the international relations literature's description of crisis bargaining and war, the influence of partisan identity melts away.

A strategic understanding of war and diplomacy clearly demonstrates that public support, or at least acquiescence, is most important in the opening stages of a crisis. In these stages, partisan cues are often not available, and so we would expect voters, to the extent that they are aware of these events at all, relying on other sources of understanding in order to make sense of them. In this study, partisan cues were not available either, other than the involvement of the presidential administration. As we would expect, questions that specifically referenced the presidency tended to activate partisan loyalties, while those that did not reference the president and simply asked for a position on the potential war were not associated with partisanship.

This implies at least two conclusions. One is that persuasion is still very likely to be an option for leaders early in crises because of the absence of partisan cues, which tends to mitigate against the impact of new information.¹¹ The second is that the emerging consensus among public opinion scholars that partisanship matters for foreign policy is simply an artifact of the observational data that their research has used. Not only do the public opinion data suffer from more sampling error than is commonly understood, as I demonstrated in Chapter 2, but they likely load the dice in favor of finding partisan structures beneath the partisan debate as well.

¹¹From opposition figures, anyway. One could argue that persuasion is still likely so long as it comes from partisan confederates, but then again, the power of partisan ties suggests that voters were already likely to support anything suggested by their own party in the first place, which would make persuasion somewhat superfluous.

CHAPTER 5

The Wages of Sin: When Does Information Market Intervention Occur?

“In 2009, [the Missile Systems division] booked \$645 million for AMRAAM systems for international customers and the U.S. Air Force ... and \$402 million for Phalanx Weapon Systems. MS also booked \$384 million on Standard Missile-2 for international customers and the U.S. Navy, \$318 million for Standard Missile-3 for the MDA and \$294 million for Tactical Tomahawk cruise missiles for the U.S. Navy.”

2010 Form 10-K Annual Report, Raytheon Company.

In Chapter 4, I presented evidence that suggests subjects do in fact rely on different considerations in response to questions about foreign policy and the use of force. When questions refer specifically to the presidential office, then respondents rely on their partisanship in reporting their support for foreign policy performance. When questions refer to the use of force itself—and a hypothetical use, at that—then respondents cannot rely on their partisanship and instead look to other attitudes in shaping their support or opposition to prospective government action.

However this approach, like any other survey experiment, can only explain the observed historical record by implication. In this chapter, I investigate empirical evidence for information market intervention. Given the large and endemic measurement error imposed by the traditional measure of the public’s reaction to foreign

policy events, presidential job approval, I instead present a *financial event study*, wherein I use common stock prices for American defense firms as an indirect measure of public engagement before a crisis. I include a short literature review of the event study method below before outlining the data, methods and results of my own application.

5.1 Literature Review: Financial Economics, Political Information, and the Event Study

Event studies are common in business and financial economics, where researchers are often interested in the effect of firm-level developments on investor confidence in those firms' future prospects. Quarterly and annual disclosures mandated by the Securities and Exchange Commission, for example, might include positive news, such as new products or market access, or negative news, like worse-than-anticipated declines in marketshare. In either case, these developments ought to effect the market price for equity in that firm (Fama et al. 1969; Ball and Brown 1968; MacKinlay 1997). In conceptual terms, event studies resemble a natural experiment in which a baseline correlation between a firm and the wider market is established and then compared to that same correlation during a particularly "newsworthy" time. If broader market conditions do a particularly poor job of explaining a firm's market price in the immediate aftermath of a new piece of sensitive information reaching the market, then presumably that information was in fact relevant to a firm's future performance.

However, the functional form of an event study is easily generalized to other types of information as well, and as such it has begun to be adopted by economists outside of business schools and, indeed, by political scientists as well. Dube, Kaplan and Naidu (2011), for example, used an event study framework to better understand the use of covert action by the United States during the early period of the Cold War. A

substantial fraction of the Central Intelligence Agency's workforce was recruited from the financial sector, including its highest-level officials. Allen Dulles, the Director of Central Intelligence who oversaw the American coup against Guatemala's Jacobo Arbenz Guzman, had previously been a member of the board of United Fruit. Likewise, Thomas Dudley Cabot had served as both the Department of State's Director of International Security Affairs as well as United Fruit's CEO (Dube, Kaplan and Naidu 2011, p. 1379). Given the prevalent connections between the men in charge of the American security apparatus and firms which were exceptionally invested in the outcome of covert actions abroad, Dube, Kaplan and Naidu suspected that the secret planning and authorization of the coups in Guatemala, Cuba, Iran and elsewhere would have left their fingerprints in the broader market. And, in fact, such was the case. Using newly-declassified timelines of CIA activity between 1953 and 1961, Dube, Kaplan and Naidu found that *authorizations* of covert action produced a boost of 10% or more in the common equity price of the most-affected firms. The *public revelations* of the *result* of these coups were much smaller by comparison. The authors infer that well-connected investors were able to act on their credible, private information about CIA planning to purchase shares of affected firms in advance with the knowledge that their future prospects were about to improve dramatically thanks to American intervention abroad.

Other work has found a series of relationships between political events and financial markets in a variety of other dramatic circumstances. Guidolin and La Ferrara (2007) connected a timeline of the Angolan Civil War and the relative prosperity of mining companies with significant investments in Angolan diamond production. The publically traded stock of companies with a large degree of involvement in the Angolan market responded much more markedly to politically salient events, such as the death of UNITA's Jonas Savimbi, than did similar mining firms without such exposure to Angola. (DellaVigna and La Ferrara 2010), working in a similar vein,

examined countries subject to international arms embargoes and found that increases in conflict intensity—which would, perforce, increase demand for arms—positively inflated the trading prices of arms producers headquartered in corrupt states with poor reputations for transparency with respect to the weapons trade. These firms would be most likely to be able to circumvent international embargoes, and if such were the case, their earnings would grow when they supplied those belligerents who more scrupulous firms were prevented from serving.

Event studies can also be used to estimate changes in broader market conditions, rather than as a means to investigate the particular fortunes of a small number of firms. Applications in political science include Bernhard and Leblang (2006), who examined market responses to a series of different democratic political developments, from coalition building in parliamentary democracies, where the presence or absence of strong parties implies a lesser or greater degree of uncertainty for later economic prospects among traders, to the disputed presidential election in the United States in 2008, where the ebb and flow of observers' certainty about the winning party are reflected in the volume of individual after-hours trades. Bernhard and Leblang found, for example, that strong parties, by reducing the uncertainty associated with negotiating a new ruling coalition, were associated with a relative premium to stock prices across the board. More risky negotiations, which are associated with the absence of strong parties and the emergency of a multitude of smaller parties in their place, encouraged investors to reallocate their portfolios away from stocks in the pursuit of less risky alternatives.

5.2 Understanding Financial Markets and their Link to Information Market Intervention

5.2.1 A Primer on Finance, Uncertainty and Information

The theoretical framework for the use of an event study to understand political behavior rests on the efficient markets hypothesis, or EMH.¹ In equilibrium, the price of any security ought to reflect the discounted present value of the future cash flows associated with holding that security. In the case of common stocks, the inherent value of holding a stock is the income received in the form of periodic dividends issued by the company in which an investor now holds an ownership stake. As a firm's business prospects improve, the market price of its stock ought to grow because the likelihood that the firm will issue dividends to shareholders, and larger dividends at that, improves as well.² The investor's attitude toward portfolio risk will also help determine the price he or she is willing to pay at market to hold any given security, but we would expect equity prices to track with their likely future income performance.

This expected relationship between performance, return, price and risk underpins the various forms of market efficiency that financial economists have defined. If investors are rational, or if there are merely enough rational investors who can profit through arbitrage to take advantage of poor decisions by irrational investors, then the only thing that can affect a security's price is relevant news about that security's future performance. Per Shleifer:

...when news about the value of a security hits the market, its price should react and incorporate this news both quickly and correctly. The 'quickly' part means that those who receive the news late—for instance by reading it in the newspapers or in company reports—should not be able to profit from this information. The 'correctly' part means that the price adjust-

¹Except where otherwise cited, the discussion in this section draws from Shleifer (2000).

²Like many other models, this is clearly an abstraction. Apple (NASDAQ: AAPL) did not make any dividend payments between October 1995 and July 2012, over which time its share value grew by over 400%.

ment in response to the news should be accurate on average: the prices should neither underreact or overreact to particular news announcements.

Fama et al. (1969) first defined efficient financial markets as those in which current prices “always reflect [all] available information.” This might be the case for a number of reasons, from strong assumptions about rational investors (if investors know all information at all times about the securities which they are trading, then the clearing price at market ought to reflect such), to the expectation that irrational investor strategies will “cancel one another out” so long as the irrational agents’ own strategies are not correlated with one another. Other researchers subsequently defined progressively weaker forms of the EMH: weak efficiency, in which no trader can profit from knowledge of past prices and returns, and semi-strong efficiency, in which no trader can profit from *any* public information. Semi-strong efficiency implies that traders might do better than the expected, risk-adjusted return on a security through the use of private information—insider trading—which among other findings is consistent with Dube, Kaplan and Naidu (2011), as described above.

These various approaches to market efficiency are important to understand because they force us to more carefully consider the effect of information on security pricing. Not even the most restrictive definition of efficiency excludes the possibility of prices reacting to events and information, but instead suggests that any new information about the future performance of a stock or bond will be quickly incorporated into the clearing price at market, rather than effecting a slow change over time. News becomes “stale” in very short order given the profit-seeking incentives facing traders large and small.

For that reason, if political developments in the public eye are to influence any market outcomes, then those developments have to be unexpected. Bechtel and Schneider (2010) rely on this insight in their examination of the market for European defense stocks and its response to summits on European Union military integration.

They found that agreements on expanding cross-border military cooperation and, most pointedly, increasing military spending, did on average produce positive shocks to the market capitalization of European arms manufacturers. This much makes sense: as agreements to increase European military capabilities occur, the expected future revenue of the firms that will sell those capabilities to their respective governments must also increase, and this growth in expected future revenue also requires that the net present value of holding stock in those firms has increased as well.

However, summits that merely discussed the possibility of amending the European Security and Defense Policy were not meaningful to investors. Only those summits that produced real and binding agreements—and, in particular, those agreements that were novel and less anticipated in advance—generated new investor demand for European defense firms' equities. These summits were able to produce attendant responses in the market precisely because their outcome could not be known with certainty in advance. If they had, then the market's response (whether a reduction in value if European defense integration was walked back, or growth due to an increase in integration and spending) would have occurred *when the outcome became known*, and not with the summit itself. Analogous reasoning holds for the various political events investigated by Bernhard and Leblang.

5.2.2 Uncertainty and War

Another political arena in which outcomes are not often known with certainty in advance is crisis bargaining. Consider two recent conflicts involving the United States, one of which was a surprise and one of which was not: the Gulf War (1991) and the Iraq War (2003). It merits repeating here that international relations theory portrays a clear chronological sorting between crisis and conflict, even if we accept Wagner's (2000) contention that limited wars might be intertwined within the crisis stage. Two major conflict data sets report the start date of either crisis as follows:

	Actor Level			System Level		
	ICB	COW	ICB	ICB	COW	COW
Conflict	Crisis Starts	War Starts	Crisis Starts	War Starts	Crisis Starts	War Starts
Gulf War	10/30/1990	na	7/24/1990	1/16/1991	8/2/1990	5/28/1990
Iraq War	1/13/2003	na	na	na	9/12/2002	na
						8/2/1990
						na

Table 5.1: Contrasting values for crisis / dispute starting dates.

Clear disagreement exists between these coding systems, despite both conflicts being salient events. The International Crisis Behavior project defines a crisis in three parts: a “threat to basic values” that requires a time-sensitive response and carries with it a “heightened probability” of military conflict. In the Gulf War, the first such threat occurred on May 28th, 1990, when Saddam Hussein accused the Kuwaitis of “waging economic war against Iraq by overproducing oil and causing a collapse in world oil prices.”³ However, the United States is not coded as a crisis actor until October 30th, when President Bush approved a military campaign against Iraq to take place the following January. A third alternative is to code the onset of a threat to the prevailing international order; the ICB reports that this occurred on August 2nd, with the Iraqi surprise invasion and occupation of Kuwait.

By contrast, the Correlates of War project’s Militarized Interstate Dispute data supposes that the United States first entered the crisis stage *vis-a-vis* the Iraqi-Kuwaiti conflict on July 24th, when Iraq committed six regular divisions to the frontier with Saudi Arabia, in preparation for the invasion of Kuwait proper. Militarized Interstate Disputes must, per Jones, Bremer and Singer (1996), entail a “collision of interests” in which “the threat, display or use of military force short of war by one member state is explicitly directed toward the government, official representatives, official forces, property, or territory of another state.” This standard is at once more inclusive than that of the ICB coding—the implicit threat to the United States by Iraq connoted by a troop mobilization directed at their client states in the Persian Gulf counts—and more selective, in that the original Iraqi verbal provocation of Kuwait is no longer enough to merit coding a new dispute.

While the MID data has not yet been extended beyond 2001, a similar dynamic exists within the ICB data. Iraq is faced with a crisis as early as September 2002, when President George W. Bush appeared in front of the United Nations General

³Adapted from the International Crisis Behavior profile of this event.

Assembly and announced the conditions under which Iraq ought to immediately end its alleged unconventional weapons programs. For the United States, a crisis does not obtain until January 2003, when Bush privately decides that the UN-guided weapons inspection regime has not been effective.⁴

In either case, the theory of information market intervention suggests that the formal start of hostilities ought, under some circumstances, to be preceded by a period of public engagement meant to build domestic support for conflict in the expectation that this will improve a state actor's chances of winning concessions during the crisis phase. As it happens, recent software advances have made it easier than ever before to measure the behavior of mass media. The following figures graph the frequency of the word "Iraq" in all sections of the *New York Times*, for a period of three years before both the beginning of major combat operations in both the Gulf and Iraq Wars:

The resulting pictures are, of course, quite different, particularly in how each set of observations demands a dramatically different scale. Even a few short weeks before the Gulf War begins with the Iraqi invasion of Kuwait, a given day's publication of the *Times* likely contained no more than five mentions of Iraq. On the other hand, as early as the fall of 2001 no daily issue of the *Times* contained *fewer* than five mentions of that country. Six months before the American invasion, an average issue contained 20 mentions of Iraq or more.

With the benefit of hindsight and a number of investigative reports, we can easily explain the discrepancy: the Bush administration had decided early in their first term that regime change in Iraq was a prudent course of action, and began laying the groundwork necessary for the war effort soon after the 9/11 attacks. That is

⁴In this case, ICB rests their coding on a book by Bob Woodward, who reported the conversation between Bush and Secretary of State Colin Powell. This is consistent with ICB's stated coding rule, which is to capture the first moment at which a state's critical decisionmakers perceive a serious threat to their values or position that requires an immediate, or at least prompt, response. Of course, the sources from which ICB determines such times vary considerably, and one wonders at the incentive for Bush and Powell to report such a conversation in the way that it truly happened.




Figure 5.1: Daily Mentions of 'Iraq,' 1988-1991

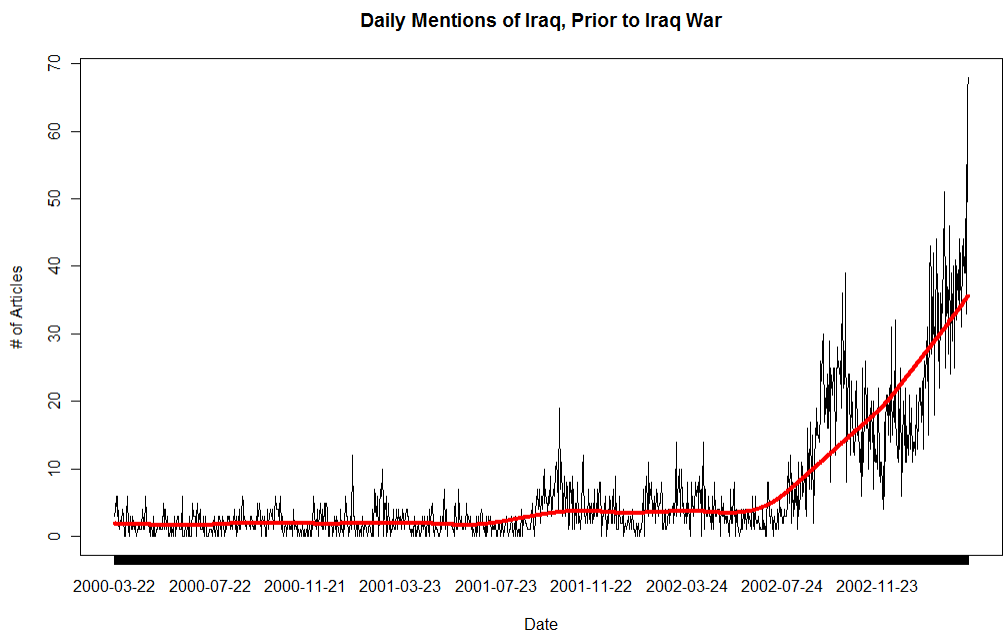


Figure 5.2: Daily Mentions of 'Iraq,' 2000-2003

entirely consistent with the theoretical logic I developed above, and this pattern is consistent with media coverage as measured by mentions of the word ‘Iraq’ in the *Times*’ coverage.

At the same time, this sort of evidence amounts to a very limited comparative case study, with all of the potential pitfalls that case studies imply for the study of conflict behavior (Achen and Snidal 1989; Geddes 1990). A more accurate content analysis of this kind would have to account for the content of these stories, and not just their relative frequency, which would present a considerable burden. In order to determine whether similar patterns of public engagement obtain over the longer course of the American foreign policy record, some manner of quantitative analysis is necessary.

5.2.3 Financial Markets and Conflict Onset

When war breaks out, the future financial prospects of arms manufacturers ought to improve for a similar reason to that identified by Bechtel and Schneider (2010): war presents an opportunity for growth in future revenue streams. Where Bechtel and Schneider argued that unanticipated improvements in European defense integration, and thereby unanticipated growth in European defense budgets, present growth opportunities to the major European defense contractors that supply those governments, the outbreak of war should fulfill a similar role for those firms that are prominent suppliers to the belligerents in any particular conflict.

However, per the efficient markets hypothesis, this can only be true in a very particular sort of case: when the conflict is in fact unanticipated. The omnipresence of Iraq in the mass media as portrayed by the previous graph tracks with the reporting of Woodward (2004) and others, who claim that the search for *casus belli* against Iraq among the circles who would go on to populate Bush’s foreign policy apparatus predated the 2000 election, let alone the September 11th attacks or the disclosure of

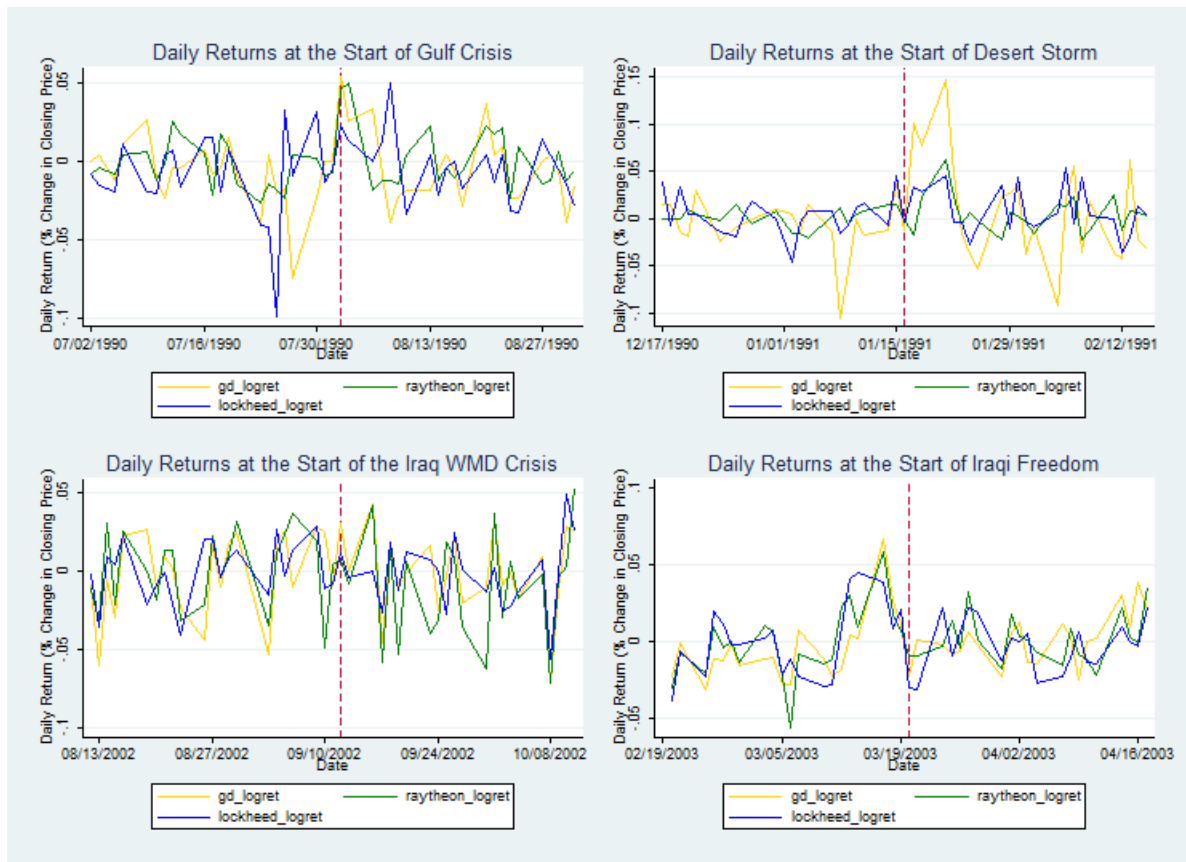


Figure 5.3: Daily returns for a selection of prominent defense firms, related to the Gulf and Iraq Wars

any number of alleged charges with respect to WMD programs. The prevalence of this sort of discourse before the invasion of Iraq suggests that the outbreak of war itself ought not to have had any effect on investors' beliefs about future revenue for prominent American defense contractors. And, indeed, a casual glance at the nominal prices for a selection of major American contractors would appear to bear this out.

The above figure charts the daily change in closing price⁵ for three major American defense contractors. The top row examines pricing behavior associated with two key dates related to the Gulf War; August 2nd, 1990, when Iraq invades Kuwait and triggers a systemic crisis by the standards of the International Crisis Behavior coding

⁵Here, and throughout the rest of the chapter, the day-on-day change in pricing is expressed as a log-difference, $R_{it} = \ln(\text{price}_{it}) - \ln(\text{price}_{i,(t-1)})$. This ensures that the series of daily returns is stationary.

rule, and January 16th, 1991, when the United States and its allies begin the air war over Iraq. Both events, at first blush, appear to be informative to the market. The market price for Raytheon, Lockheed, and General Dynamics all move in a positive direction more or less in unison, gaining between 3 and 4 points in value, when the Iraqi invasion occurs. The American invasion is extremely beneficial for the future prospects of General Dynamics, with more modest gains produced for the remaining two firms.

The second row examines the systemic crisis date associated with the Iraq War, September 12, 2002, and the beginning of this second American invasion, March 20, 2003. A somewhat different story emerges in these cases. The systemic crisis—President Bush reporting Iraq to the UN General Assembly—seems to not make any difference to the market. The invasion itself appears to be newsworthy, despite the months of discussion of war plans beforehand, though these gains are largest on the Friday *before* the surprise attack gets underway on March 20th, which was a Wednesday.

Owing to the difficulty in choosing a consistent measure of pre-crisis engagement that obtains over multiple crises, in the following sections I present two more preliminary tests. First, I determine the next unexplained variance in daily prices for a selection of American defense firms for a variety of crises. The unexplained variance for an equity price net of a series of other controls in a synthetic portfolio is an *abnormal return*. I then go on to follow Bechtel and Schneider (2010) in predicting the magnitude of the abnormal return associated with a crisis with a selection of covariates taken from the model presented in Chapter 3.

5.3 Data and Method: an Event Study Approach to American Crisis Diplomacy

Event studies have two primary components (MacKinlay 1997). First, we define an event of interest and a window around that event in which we expect stock prices to change. Second, we calculate a baseline covariance between an individual firm's equity and some larger market of reference. Finally, we predict the daily returns of each share during the event window using the estimation from the baseline period. In other words, if a particular event was not especially relevant news for a given equity from the market's point of view, than the out-of-sample prediction from the baseline period ought to be a reasonable predictor of equity behavior during the event window.

The difference between the predicted daily returns and the actual daily returns during the estimation are *abnormal returns*. Additionally, we can sum the abnormal returns over the duration of a multiple-day event window to generate an event's *cumulative abnormal return*. If the cumulative abnormal return over the event window can be distinguished from 0, then we can plausibly say that the event itself had some effect on the




Figure 5.4: An event-history model. The baseline estimation period runs for τ days, ending k_1 days before the event. Equity behavior is measured in a window about the event; the longest possible event window would be $k_1 + k_2$ days. The event of interest falls at t_0 .

5.3.1 Estimation

There are a number of different ways to estimate the difference in asset pricing behavior between the estimation and event windows. In the most naive possible functional form, the constant mean return model, the price of a single asset of interest, i , is modeled as a function of some underlying asset-specific, time-invariant return and a stochastic component:

$$R_{it} = \mu_i + \varepsilon_{it} \quad (5.3.1)$$

While the constant mean return model is surprisingly effective under certain conditions (Campbell and MacKinlay 1997, p.154), it will suffer when analyzing the effect of events over a long time horizon because of the increasingly likely case that assets—in this application, firms—are likely to experience fundamental changes in value that would make the constant mean assumption untenable. On a related note, abnormal returns will demonstrate a greater sample variance under this specification, which will lead to a reduction in power.

Progressively more restrictive specifications reduce the portion of an asset's return that is driven by variance in other related investments. The “market model” described by Campbell and MacKinlay (1997) adds an additional term to the constant mean return model representing the return to a broad market index, such as the Standard and Poor 500. With some fraction of observed variance in daily price behavior for an asset explained by prevailing conditions, the researcher is more likely to distinguish event effects when they do in fact exist. MacKinlay (1997) defines the market pricing model,

$$R_{it} = \alpha_i + \beta_i R_{mt} + \varepsilon_{it} \quad (5.3.2)$$

with R_{mt} the return of the broad market index on a particular day, defined in the same fashion as the individual equity term. From the perspective of financial economics, this is an application of the Capital Asset Pricing Model, or CAPM, which assumes that any arbitrage opportunities are exploited instantly and that all substitute investment opportunities are accounted for by the market variable against which the individual equities are compared. If arbitrage opportunities are always realized, then the only way in which the return series for an individual asset can differ from that of the wider market is in its inherent risk, which is captured in the previous equation by β . Any substantial deviations from the difference in risk profiles for a given asset and the market during the event window ought to then be driven by the information that the event introduced to investors.

The most straightforward weakness of this approach, that the scale for β will change depending on the market index used (Bernhard and Leblang 2006), is not a concern here because I am not, strictly speaking, interested in the coefficient estimated as much as I am the magnitude of the cumulative abnormal return itself. The CAPM can itself be supplanted by models derived from Arbitrage Pricing Theory (APT), which supposes that the return to an asset is related to a number of other investment vehicles, all of which are described by their own risk profiles. In practical terms, the APT replaces the single market return term in the previous occasion with a series of covariates, supplementing the daily return of a basket of stocks with the daily return to broad indices for commodities, exchange rates, and the like.

In either case, these various approaches to estimating the abnormal return to a security that can be explained by the shock of an exogenous event differ only in the

number of control variables included, and indeed even more convoluted applications of APT models still require the researcher to determine the appropriate financial assets on the right hand side of their equation on the basis of something resembling a hunch (Bernhard and Leblang 2006, p.63). For that reason, I follow Bechtel and Schneider (2010) in implementing a modest version of the APT that includes controls for the American stock market, oil and gold prices, and the exchange rate between dollars and euro⁶.

The daily abnormal return—the return to an asset during the event window that is not explained by the relationship between the firm and these other factors as estimated in the baseline period—is given by

$$AR_{i\tau} = R_{i\tau} - [\hat{\alpha}_i + \sum_{j=1}^J \hat{\beta}_j R_{j\tau}] \quad (5.3.3)$$

where j denotes each of the control portfolio elements in term. In other words, the abnormal return associated with any single day of the event window is merely the residual between an asset's actual return on that day and the out-of-sample prediction offered by an equation like (5.3.2). Both the APT and CAPM-derived models can be estimated with ordinary least squares, which remains consistent and efficient even with daily data (Campbell and MacKinlay 1997).

Testing for statistical significance is a different matter entirely. Parametric methods exist for event studies but as Guidolin and Ferrara (2007) point out, standard distributional assumptions about variance can be undermined given the inter-asset correlation in returns if multiple firms' equity are considered in response to an event. Furthermore, useful estimation and event windows are not likely to be long enough to satisfy large-sample conditions for the convergence of traditional tests. Corrado

⁶For dates before January 1st, 1999, the dollar/deutschmark exchange rate is used instead.

and Zivney (1992) find that t-tests to compare the means in the estimation and event windows over-reject the null hypothesis of no effect from the event, which was corroborated by Campbell and Wesley (1993). As such, I follow Guidolin and La Ferrara’s decision to rely on a non-parametric test first identified by Corrado and Zivney (1992) for use with event studies using daily stock data and short estimation windows. The “rank test” constructs a ranking for each firm of all observations of an abnormal return, across both the estimation and event windows. When an event has in fact produced some effect on a firm, then the average rank in the event window will be significantly different than the expected value of any period’s rank *ex ante*, which is simply the average rank $\frac{\tau+k+1}{2}$. With multiple event days, the test statistic is:

$$\frac{\sum_{i=1}^n [\frac{1}{k+1} \sum_{j=t_0-k}^{t_0+k} (\kappa_{ij} - \frac{\tau+k+1}{2})]}{\sqrt{\frac{\tau+k+1}{2} \sum_{j=t_0-\tau}^{t_0+k} [\sum_{i=1}^n (\kappa_{ij} - \frac{\tau+k+1}{2})]^2}} \quad (5.3.4)$$

where k is the number of additional event window days after the day of the event, τ is the number of the days in the estimation window, n is the number of firms in the sample, and κ is the rank of each abnormal return observation across both windows. The rank statistic has the additional virtue of producing easy-to-interpret results. It produces a z-score of the difference in means of the estimation and event windows; in the results below, I report both the test statistic value, and one- and two-tailed p-values.

In the second series of analyses below, I return to Bechtel and Schneider (2010) and present a generalized least squares time series estimation using the firm-event-day as the observational unit. GLS allows us to consider the issue of conditional variance while retaining the option of conventional regression-based hypothesis testing.

5.3.2 Data

5.3.2.1 Relevant Firms

I argued at the outset of this chapter that it was sensible to examine investors' reaction to the onset of conflict events in order to determine the likely presence or absence of information market intervention because of a straightforward model of financial self-interest: the prospects of conflict suggest that the future earnings of firms that sell weapons have improved, and to the extent that a government has spent time persuading its voters of the need for war in advance, those additional prospective earnings ought to have already influenced investor decisions before the crisis begins.

Beyond describing the estimation process itself, then, careful selection of both crisis events and relevant firms is required. Owing to the size of the American defense industry, there are any number of firms to choose from whose equity pricing could conceivably move in response to American conflict events. I used the Federal Procurement Data System⁷ to determine the total amount spent by the U.S. Department of Defense from its contractors at five year intervals in FY 1990, 1995 and 2000.⁸ FPDS provides accounts payable summary data for all prime contracts awarded by the federal government, which can be separated by functional area. I then created a correspondence table between the vendors reported by the FPDS and their ultimate corporate parent, to account for slight spelling and punctuation changes in the FPDS records. This yielded an accurate picture of vendors representing 60 -70% of all Defense Department spend in each of the target years.

I compared these figures to firms' own arms sales figures as listed in the "Top 100 Arms Producers" report compiled by the SIPRI yearbook series⁹ for the same

⁷<http://www.fpds.gov>.

⁸Although the FPDS website can produce contractor reports as far back as FY 1978, I found that for the years before 1990, more than 95% of spending was attributed to "Unknown," or was listed as classified. The government's fiscal year begins October 1st.

⁹SIPRI, the Stockholm International Peace Research Institute, began compiling this list in 1990,

three base years. SIPRI also reports what percentage of firms' total sales came from military contracts. As the ratio of a given contractor's total business with the United States increases relative to its total revenue, then its sensitivity to American foreign policy can said to have grown and therefore it is a likely candidate to exhibit abnormal returns around unanticipated crisis events.

Table 5.2 lists the firms in my sample. Aside from concentrating on large, publicly-traded firms, there were a few other sampling issues to consider. First, I attempted to choose firms that would cover as many categories of military goods as possible, as defined by SIPRI. Second, I attempted to include a few companies that have exceptionally large civilian operations in addition to their military contracting business, such as Boeing and General Electric, to provide an additional benchmark for detecting the effect of the onset of a dispute; I expect that these firms will experience smaller shocks, if any shocks at all, because any elevated probability of future arms sales will represent a relatively smaller fraction of future revenue. Third, only firms for which closing prices were available for the entire period containing the conflict sample were considered. Depending on how certain mergers between defense firms were handled, this means that some qualify but not others.

The stock pricing series themselves are dependent on data from the Center for Research in Security Prices (CRSP), administered by the Booth School of Business at the University of Chicago. This database is available electronically through the Wharton Research Data Service, which in turn is provided at the University of Michigan by the Ross School of Business. CRSP also reports historical, daily spot prices for oil, gold, and international exchange rates, although these are not available over the entire time period. Oil prices are available as daily series beginning in 1983, while the \$/DM exchange rate is only available from 1986 onwards. Daily pricing for gold

which covered the global arms industry for 1988. The extent to which SIPRI reports between 1988 and 1992 reflect the true distribution of arms production across the globe is unclear because of the difficulty in measuring Soviet economic production, but in any event American firms dominate the list for all available years.

Constant Prices in 2000 US \$MM										
Firms	1990					1995				
	SIPRI Rank	Total Arms Sales	Total DoD Sales	Arms Sales %	SIPRI Rank	Total Arms Sales	Total DoD Sales	Arms Sales %	SIPRI Rank	Arms Sales %
Boeing	9	6,681	2,256	18%	8	4,746	2,483	22%		22%
Concorp	58	1,035	5,778	45%	68	554	1,661	28%		28%
General Dynamics	2	10,873	4,995	82%	15	3,311	2,462	96%		96%
General Electric	6	8,450	1,736	11%	26	1,921	339	7%		7%
Honeywell	100	472	3,425	6%	72	520	6,858	7%		7%
Lockheed	4	9,825	2,037	75%	1	15,594	2,317	60%		60%
Martin Marietta	11	6,026	7,658	75%	2	10,871	9,057	67%		67%
McDonnell Douglas	1	11,816	7,774	80%	6	6,441	1,974	84%		84%
Northrop	4	9,956	366	27%	75	497	449	14%		14%
Gruzman	61	917	3,521	57%	10	4,475	2,109	34%		34%
Olin	7	7,205	681	24%	NR	1,808	946	16%		16%
Raytheon	NR	2,489	1,571	20%	NR	418	418	16%		16%
Sturm Ruger	29	2,620	1,968	20%	NR	418	418	16%		16%
Textron	27	5,371		19%	11	4,125	1,494	16%		16%
Unisys	13									
United Technologies										

Constant Prices in 2000 US \$MM										
Firms	2000					2000				
	SIPRI Rank	Total Arms Sales	Total DoD Sales	Arms Sales %	SIPRI Rank	Total Arms Sales	Total DoD Sales	Arms Sales %	SIPRI Rank	Arms Sales %
Boeing	2	16,900	3,579	33%						
Concorp	56	480	1,976	46%						
General Dynamics	6	65,200	1,492	63%						
General Electric	20	1,600	870	1%						
Honeywell	1	18,610	14,673	73%						
Lockheed	1	18,610	8,348	87%						
Martin Marietta	1	18,610	2,405	21%						
McDonnell Douglas	5	6,660	6,687	60%						
Northrop	5	6,660	1,96	82%						
Gruzman	11	2,880	156	9%						
Olin	11	2,880	180	9%						
Raytheon	4	10,100								
Sturm Ruger	4	10,100								
Textron	27	1,200								
Unisys	27	1,200								
United Technologies	11	2,880								

Table 5.2: Sales totals for selected firms. Total sales and % revenue from arms taken from SIPRI; U.S. Department of Defense sales from author's collection from FPDS. Additional sales information from non-SIPRI-ranked firms taken from annual investor reports. All revenue reported in millions of 2000 US\$.

contracts are available from 1975. In keeping with standard practice among financial analysts, I use the West Texas Intermediate grade to measure the price of oil.

The reduction in defense spending after the end of the Cold War led to a number of mergers, particularly involving firms with significant aerospace interests. These most prominent mergers include:

1. *Lockheed-Martin*: Merged March 15, 1995. CRSP provides a single series for Lockheed and Lockheed-Martin that begins 12/9/1939 and runs until the present day, while the historical series for Martin Marietta ends on the date of the merger. Since Martin Marietta is in or around the top ten in arms sales volume until its acquisition, I include their pricing series as a separate firm for inclusion with those crises that begin before the merger date.
2. *Northrop-Grumman*: Merged May 19, 1994. In a similar fashion to Lockheed-Martin, the Northrop series covers both Northrop as an independent concern as well as the merged entity. Since Grumman is a distinctly smaller firm, I do not include them as a separate subject as I did with Marietta Martin.
3. *McDonnell-Douglas*: Merged May 1, 1967. Since the merger occurs so early in the time series, I rely on the default series from CRSP, which lists McDonnell as the originating firm.

One other point is to consider more carefully how the likelihood of conflict would improve earnings estimates for any particular defense vendor. A quick examination shows that, by total sales, producers of sophisticated electronics, satellites and other such equipment are among the largest firms by total sales. However, this is somewhat misleading—they all achieve these sales by selling small numbers of very expensive items that are not likely to be damaged in fighting or need to be replaced. It is unlikely that any single conflict would be, *a priori* sufficient to build investor anticipation of

new future revenue for these firms. Unfortunately, the companies that sell defense-specific “consumables”—in other words, firearms and ammunition, which would need to be replaced—tend to be privately-held, so they do not lend themselves to this sort of analysis. Sturm-Ruger and Olin are two exceptions, which is why they are included here. Colt, which makes the M16 and M4 rifles, is not publicly-held.

Firms	Products
Boeing	Aircraft, electronics, missiles
Gencorp	Electronics, engines, small arms
General Dynamics	Aircraft, military vehicles, electronics, missiles, ships
General Electric	Aircraft engines
Honeywell	Electronics, missiles
Lockheed	Aircraft, electronics, missiles
Martin Marietta	Missiles
McDonnell Douglas	Aircraft, electronics, missiles
Northrop Grumman	Aircraft, electronics, missiles, small arms
Olin	Small arms
Raytheon	Electronics, missiles
Sturm Ruger	Small arms
Textron	Aircraft, electronics, aircraft engines, military vehicles
Unisys	Electronics
United Technologies	Aircraft, electronics, missiles

Table 5.3: Representative product offerings from included American defense contractors, reproduced from SIPRI yearbooks.

5.3.3 Selection of Conflict Events

5.3.3.1 Relevant Events

Selecting a reasonable set of conflict events against which to test the market’s reaction is also difficult. Earlier in this chapter I described the subtle but important distinctions between major datasets in defining the start and end of even the largest international conflicts. These different coding systems also present implications for

understanding the distinctive phases of crisis and war.

In terms of the formal literature on conflict initiation, wars are expected to be preceded by some period of crisis. This distinction is stylized, of course. In some cases, a formal demand is made by one belligerent of another, as with Soviet demands for Finnish territory before the 1939 Winter War. In other cases, the war itself comes as a surprise but nonetheless reflects some prior exchange of give-and-take, either in secret or by implication from other policy choices, as with the prelude to the Pacific War between the United States and Japan in the late 1930's. In either event, what we consider as an international crisis and dispute ought to be a superset of what we describe as war. Crises are important because they stand some non-trivial chance of becoming shooting wars.

And yet this is not always reflected in the available empirical data. Disputes over fishing rights and maritime boundaries, for example, are included in the Correlates of Wars' Militarized Interstate Dispute data. Fishing disputes fit the MID coding rule because COW considers "a single military action involving an explicit threat, display, or use of force by one system member state towards another system member state" to comprise a dispute (Ghosn, Palmer and Bremer 2004); fishing disputes often qualify because they often if not exclusively involve at least one state's coast guard or navy. While some fishing disputes do ultimately produce more serious conflicts, or at least overlap them—the long-standing territorial dispute between Peru, Chile, and Ecuador comes to mind—they are the exception rather than the rule.

For that reason, I focus here on events drawn from version 10 of the International Crisis Behavior data (Brecher 1977). ICB's sampling rules are more strict than MID, which is reflected in the number of events each dataset include; for example, MID lists 1524 total disputes across all actors between 1946 and 2001, while ICB only includes 346 crises over the same period of time.¹⁰ Of 36 ICB events involving the use of

¹⁰Other interesting points of distinction between MID and ICB abound. Per MID, the initiator of the Cuban Missile Crisis was the United States, because it "engaged in a show of planes against

regular American troops after 1946, the United States is the “crisis trigger” of record in just 5 of them. Restricting the discussion to only those events which occurred after January 3rd, 1986 produces a list with 14 crises, of which the United States was the aggressor in just two, the Iraq War and the second territorial stand-off between the United States and Libya in March 1986.¹¹

In specific terms, the following estimations include only those crises where the United States played at least an indirect military role, rating at least a 3 on the 4-point scale that ICB uses to code American involvement. Furthermore, only those events where the United States is a named crisis actor are included. Because of its preponderance of power, the US is often cited as a participant in crises but not as an outright actor since its own values have not been adequately threatened. A number of client-state entanglements, such as crises between Ethiopia, Eritrea and Somalia, exhibit this behavior.

Two particular types of events are obstacles for the overall estimation strategy. First, ICB includes some events that fit their coding definition of a new crisis even if they occur in the context of some ongoing war. These events are extraneous from the perspective of this examination, because they cannot tell us anything about the informational environment that characterized the original international challenge. For example, the United States experiences a new crisis onset in December 1972, when the North Vietnamese government rejects a number of concessions meant to end the Vietnam War.¹² I exclude all intrawar crises from consideration.

Cuba and the Soviet Union” in January 1962 (Ghosn, Palmer and Bremer 2004) and because the Soviet missile deployment is not an actionable item in the MID coding scheme. By contrast, ICB reports that the Soviet Union has the instigator of that crisis. ICB goes on to name Grenada as the revisionist party in the American invasion of that state in 1983. In this latter case, ICB coders determined that the coup by Marxist hard-liners against Maurice Bishop in early October constituted a “perceived threat to basic values” of the United States.

¹¹Which, of course, was the prelude to the American / Libyan stand-off described in the propaganda piece published by the National Security Agency and used as the experimental treatment in Chapter 4.

¹²In response, the United States escalates its bombing campaign once again, which is recorded as the beginning of this crisis from the Vietnamese perspective.

<i>ICB #</i>	<i>Crisis</i>	<i>Date</i>	<i>Day of Week</i>	<i>Initiator</i>	<i>Interval</i>	<i>Mid-War?</i>	<i>textitPrior Event</i>	<i>Prior Initiator</i>
213	PLEIKU	07feb1965	Sunday	Non-state Actor (ICB)	187	7	CONGO II	Non-state Actor (ICB)
225	TE/TOFFENSIVE	30jan1968	Tuesday	Non-state Actor (ICB)	9	7	PUEBLO	North Korea
230	VIETNAMSPRING OFF.	22feb1969	Saturday	Vietnam	389	7	TETOFFENSIVE	Non-state Actor (ICB)
237	INVASION OFCAMBODIA	13mar1970	Friday	Cambodia	332	7	EC-121 SPY PLANE	North Korea
246	VIETNAMSPORTS MINING	30mar1972	Thursday	Vietnam	561	7	CIENFUEGOS SUB.BASE	Russia
249	CHRISTMAS BOMBING	23oct1972	Monday	United States of America	207	7	VIETNAMSPORTS MINING	Vietnam

Table 5.4: Intra-war events, from 1962 onward.

Other crises begin on weekends, or on other trading holidays. These events pose a challenge insofar as much of the reaction to a crisis in the financial markets is likely to take place on the day of the event itself.¹³ In those cases, I re-code the crisis trigger date reported by ICB to fall on the next possible trading day, which reflects the next opportunity that investors would have to react to relevant foreign policy news. In addition, some crises occur within 80 trading days of one another. In these cases, the estimation window for one event would include another event. The list of these cases follows; a few of the crises in this table are intra-war events as well.

The included list of events, then, is presented in table 5.6. The final list makes some slight modifications to the above set of rules. First, the Cienfuegos Submarine Base affair is included despite nominally overlapping with another event. In this case, the prior event was the Black September crackdown on Palestinian militants in Jordan, and the two crises happen to begin on consecutive days. Since the United States is a named actor with respect to the Cienfuegos affair but not with respect to Black September, it seems reasonable to include the former rather than the latter.

Second, a number of crises include the United States as a named actor even though the systemic ICB data does not reference any significant military involvement by the US. I include these cases as well in the subsequent analyses, reasoning that they all reflect significant American activity and therefore may potentially be associated with significant reactions in the equity markets. These cases include: the Panama Flag protests (1964), the Six Day War (1967), and a trio of nuclear program standoffs (Iran in 2006, and North Korea in 2002 and 2006).

¹³The September 11th attacks occurred on a Thursday, and quickly closed the market. I code the date of this crisis as September 17th, which was the day that the New York Stock Exchange opened.

<i>ICB #</i>	<i>Crisis</i>	<i>Date</i>	<i>Day of Week</i>	<i>Initiator</i>	<i>Interval</i>	<i>Mid-War?</i>	<i>textitPrior Event</i>	<i>Prior Initiator</i>
211	CONGO II	04aug1964	Tuesday	Non-state Actor (ICB)	5	1	GULF OF TONKIN	Republic of Vietnam
213	PLEIKU	07feb1965	Sunday	Non-state Actor (ICB)	187	7	CONGO II	Non-state Actor (ICB)
215	DOMINICAN INTERVENTN.	24apr1965	Saturday	Non-state Actor (ICB)	76	1	PLEIKU	Non-state Actor (ICB)
225	TETOFFENSIVE	30jan1968	Tuesday	Non-state Actor (ICB)	9	7	PUEBLO	North Korea
230	VIETNAMSPRING OFF.	22feb1969	Saturday	Vietnam	389	7	TETOFFENSIVE	Non-state Actor (ICB)
233	EC-121 SPY PLANE	15apr1969	Tuesday	North Korea	52	1	VIETNAMSPRING OFF.	Vietnam
237	INVASION OF CAMBODIA	13mar1970	Friday	Cambodia	332	7	EC-121 SPY PLANE	North Korea
239	CIENFUEGOS SUB.BASE	16sep1970	Wednesday	Russia	1	1	BLACK SEPTEMBER	Jordan
246	VIETNAMSPORTS MINING	30mar1972	Thursday	Vietnam	561	7	CIENFUEGOS SUB.BASE	Russia
249	CHRISTMAS BOMBING	23oct1972	Monday	United States of America	207	7	VIETNAMSPORTS MINING	Vietnam
260	WARIN ANGOLA	12jul1975	Saturday	Non-state Actor (ICB)	61	2	MAYAGUEZ	Cambodia
441	N KOREANUCLEARII	04oct2002	Friday	United States of America	22	1	IRAQ REGIMECHANGE	United States of America

Table 5.5: Intra-war events and events falling on non-trading days, from 1962 onward.

<i>ICB #</i>	<i>Crisis</i>	<i>Date</i>	<i>Do W</i>	<i>Interval</i>	<i>Trigger</i>	<i>Ext. Actor</i>	<i>Violence</i>
343	INVASION OF GRENADA	19oct1983	Wednesday	1445	Grenada	55	3
206	PANAMA FLAG	09jan1964	Thursday	450	United States of America	95	2
391	INVASION OF PANAMA	15dec1989	Friday	359	Panama	95	3
430	KOSOVO	20feb1999	Saturday	112	Yugoslavia	345	3
255	OCT.-YOM KIPPUR WAR	05oct1973	Friday	347	Multiple Actors (ICB)	365	1
210	GULF OF TONKIN	30jul1964	Thursday	203	Republic of Vietnam	365	3
292	SHABA II	11may1978	Thursday	632	Non-state Actor (ICB)	365	1
354	NICARAGUA MIG-21S	06nov1984	Tuesday	384	Russia	365	1
222	SIX DAY WAR	17may1967	Wednesday	753	Egypt	365	1
196	CUBAN MISSILES	16oct1962	Tuesday	1	Russia	365	1
239	CIENFUEGOS SUB.BASE	16sep1970	Wednesday	1	Russia	365	1
363	GULF OF SYRTE II	24mar1986	Monday	503	United States of America	620	3
386	LIBYAN JETS	21dec1988	Wednesday	1003	Libya	620	2
309	US HOSTAGES IN IRAN	04nov1979	Sunday	542	Iran	630	1
448	IRAN NUCLEAR II	10jan2006	Tuesday	1194	Iran	630	1
393	GULF WAR	02aug1990	Thursday	230	Iraq	645	4
429	UNSCOM II	31oct1998	Saturday	85	Iraq	645	3
412	IRAQ DEPLOY./KUWAIT	07oct1994	Friday	1527	Iraq	645	1
440	IRAQ REGIME CHANGE	12sep2002	Thursday	366	United States of America	645	4
422	UNSCOM I	13nov1997	Thursday	439	Iraq	645	1
419	DESERT STRIKE	31aug1996	Saturday	694	Iraq	645	2
427	US EMBASSY BOMBINGS	07aug1998	Friday	267	Non-state Actor (ICB)	700	2
434	AFGHANISTAN/US	11sep2001	Tuesday	934	Non-state Actor (ICB)	700	3
450	N KOREAN NUCLEAR III	05may2006	Friday	115	North Korea	731	1
274	POPLAR TREE	17aug1976	Tuesday	402	North Korea	731	1
224	PUEBLO	21jan1968	Sunday	249	North Korea	731	1
259	MAYAGUEZ	12may1975	Monday	584	Cambodia	811	3

Table 5.6: All included crises from the International Crisis Behavior dataset, from 1962 onward.

5.4 Analysis and Results

5.4.1 All Firms, and All Events

The dependent variable is the cumulative abnormal return over the event window, which represents the average effect of the event (or, more precisely, the effect of the information provided by the event, which is observed by investors) on the equity of any particular firm. I estimated the cumulative abnormal returns for each of the American defense contractors listed in table 5.2. The estimation period, which is a model of “normal” firm performance outside of any newsworthy period, is forty trading days, or roughly two calendar months, ending a month before each dispute observation. Although a longer estimation period would provide a better fit between the estimation and event windows and possibly a more robust estimation, longer periods also mean that more events will fall within some other event’s estimation window. Since the crisis events are public knowledge, I expect that any such abnormal returns would be realized quickly. I specify a two day event window, which includes the date on which the crisis began and the following day. The cumulative abnormal return for each crisis period is simply defined as the sum across all firms that are “alive” for each event¹⁴, and calculated as in equation 5.3.4.

The theory of information market intervention suggests two testable hypotheses: that intervention (from the perspective of the United States, anyway) ought to only be possible for those cases in which the United States is the revisionist party, and should be more tempting when facing progressively stronger rival states. With respect to the market dynamics explored here,

H1: The cumulative abnormal return characterizing the American defense industry will be smaller when the United States is the revisionist state in an international challenge, as opposed to the target state, *ceteris paribus*.

¹⁴The impact of more recent events is then measured through the performance of slightly fewer firms, as Martin Marietta and McDonell Douglas leave the sample.

H2: The cumulative abnormal return characterizing the American defense industry will be larger as the other major party in a dispute grows, *ceteris paribus*.

Since the non-parametric Corrado-Zivney rank test does not provide a mechanism for classic hypothesis testing, I merely produce the cumulative abnormal return associated with each of the crisis events in table 5.6 in turn, before adopting a regression-based approach in the following section. Since each of the portfolio controls described above is only available for certain periods of time, I re-run the underlying calculation described in equation 5.3.3 and calculate the associated rank statistics for each of four separate periods:

1. Controlling for the S&P 500 daily closing value, from July 1962 onward.
2. Controlling for the S&P 500 and the daily trading price of gold bullion¹⁵, from January 1975 onward.
3. Controlling for the S&P 500, gold price, and the daily closing spot price of crude oil, from January 1983 onward.
4. Controlling for the S&P 500, gold price, the daily closing spot price of crude oil, and the prevailing spot exchange rate between American dollars and German deutschmarks¹⁶ from January 1986 onward.

Since abnormal returns represent the otherwise unexplained change in price to all of a company's outstanding common equity, even changes of one percentage point represent enormous changes in market capitalization. Given that the figures in table 5.7 represent the joint market capitalization of upwards of 14 large firms, then, the dollar value of the change in investor beliefs in the table above is enormous. That said, only one crisis, the 1990 Gulf War, produced an abnormal return that is distinguishable from a null effect at a traditional threshold of statistical significance. The same is

¹⁵Provided by the Bloomberg financial data service.

¹⁶And after 1999, the dollar to euro rate.

<i>ICB #</i>	<i>Event Name</i>	<i>Date</i>	<i>Revisionist</i>	<i>CAR</i>	<i>Rank Stat</i>	<i>p value</i>
196	Cuban Missile Crisis	10/16/1962	Russia	-0.01	-0.48	0.63
206	Panama Flag	1/9/1964	United States of America	-0.01	-0.72	0.47
210	Gulf of Tonkin	7/30/1964	Republic of Vietnam	-0.01	-0.48	0.63
222	Six Day War	5/17/1967	Egypt	0.01	0.90	0.37
239	Cienfuego Sub Base	9/16/1970	Russia	-0.00	0.39	0.70
255	Yom Kippur War	10/5/1973	Multiple Actors (ICB)	-0.01	-1.00	0.32
259	USS Mayaguez	5/12/1975	Cambodia	0.02	0.10	0.92
274	Poplar Tree (Paul Bunyan)	8/17/1976	North Korea	-0.01	-0.77	0.44
292	Shaba II	5/11/1978	Non-state Actor (ICB)	0.01	0.70	0.48
343	Invasion of Grenada	10/19/1983	Grenada	-0.00	0.09	0.93
354	Nicaragua MiG-21s	11/6/1984	Russia	-0.01	-1.27	0.20
363	Gulf of Syrte II	3/24/1986	United States of America	0.01	1.37	0.17
386	Libyan Jets	12/21/1988	Libya	-0.01	-0.65	0.51
391	Invasion of Panama	12/15/1989	Panama	-0.00	-0.38	0.71
393	Gulf War	8/2/1990	Iraq	0.03	1.95	0.05
412	Iraq Deploy. to Kuwait	10/7/1994	Iraq	-0.00	-0.11	0.91
419	Desert Strike	9/3/1996	Iraq	0.00	0.05	0.96
422	UNSCOM I	11/13/1997	Iraq	-0.01	-0.36	0.72
427	US Embassy Bombings	8/7/1998	Non-state Actor (ICB)	0.01	1.32	0.19
429	UNSCOM II	11/2/1998	Iraq	0.01	0.33	0.74
430	Kosovo	2/22/1999	Yugoslavia	-0.01	-0.29	0.77
434	Afghanistan / 9-11	9/17/2001	Non-state Actor (ICB)	0.04	0.92	0.36
440	Iraq War	9/12/2002	United States of America	0.00	0.18	0.86
448	Iran Nuclear II	1/10/2006	Iran	0.00	0.29	0.77
450	N. Korea Nuclear III	5/5/2006	North Korea	-0.01	-0.49	0.63

Table 5.7: Cumulative abnormal returns and associated rank statistics for American international crises since 1962.

largely true when confining the sample of events and introducing additional portfolio controls, as tables 5.8 through 5.10 show.

5.4.2 “Core” Defense Firms

The reason why the results in the previous section are so poor becomes clear when examining the patterns of abnormal returns between events. Even the most shocking international crisis of our times, the September 11th attacks, failed to generate abnormal returns to the defense firms that could be distinguished from background noise. Many of the firms in the sample described in table 5.2 have a large degree of exposure to the civilian market as well as arms production. The civilian business prospects of these firms, such as Boeing, suffered as a result of the 9/11 crisis. Decomposing the joint cumulative abnormal returns for each event and generating a return series for each firm independently reveals that Boeing lost some 20% of its share value over the first two trading days after September 11th, as did United Technologies. As such, the

<i>ICB #</i>	<i>Event Name</i>	<i>Date</i>	<i>Revisionist</i>	<i>CAR</i>	<i>Rank Stat</i>	<i>p value</i>
259	USS Mayaguez	5/12/1975	Cambodia	0.02	0.08	0.94
274	Poplar Tree (Paul Bunyan)	8/17/1976	North Korea	-0.01	-0.81	0.42
292	Shaba II	5/11/1978	Non-state Actor (ICB)	0.01	0.58	0.56
343	Invasion of Grenada	10/19/1983	Grenada	-0.00	-0.02	0.99
354	Nicaragua MiG-21s	11/6/1984	Russia	-0.01	-1.57	0.12
363	Gulf of Syrte II	3/24/1986	United States of America	0.01	0.94	0.35
386	Libyan Jets	12/21/1988	Libya	-0.01	-0.71	0.48
391	Invasion of Panama	12/15/1989	Panama	-0.00	-0.72	0.47
393	Gulf War	8/2/1990	Iraq	0.03	1.59	0.11
412	Iraq Deploy. to Kuwait	10/7/1994	Iraq	-0.00	-0.14	0.89
419	Desert Strike	9/3/1996	Iraq	0.00	0.05	0.96
422	UNSCOM I	11/13/1997	Iraq	-0.01	-0.19	0.85
427	US Embassy Bombings	8/7/1998	Non-state Actor (ICB)	0.01	1.31	0.19
429	UNSCOM II	11/2/1998	Iraq	0.01	0.29	0.77
430	Kosovo	2/22/1999	Yugoslavia	-0.00	-0.21	0.83
434	Afghanistan / 9-11	9/17/2001	Non-state Actor (ICB)	0.04	0.73	0.47
440	Iraq War	9/12/2002	United States of America	0.00	0.17	0.86
448	Iran Nuclear II	1/10/2006	Iran	0.00	0.08	0.94
450	N. Korea Nuclear III	5/5/2006	North Korea	-0.01	-0.47	0.64

Table 5.8: Cumulative abnormal returns and associated rank statistics for American international crises since 1975.

<i>ICB #</i>	<i>Event Name</i>	<i>Date</i>	<i>Revisionist</i>	<i>CAR</i>	<i>Rank Stat</i>	<i>p value</i>
343	Invasion of Grenada	10/19/1983	Grenada	-0.00	-0.01	0.99
354	Nicaragua MiG-21s	11/6/1984	Russia	-0.01	-1.52	0.13
363	Gulf of Syrte II	3/24/1986	United States of America	0.01	0.76	0.44
386	Libyan Jets	12/21/1988	Libya	-0.01	-0.72	0.47
391	Invasion of Panama	12/15/1989	Panama	-0.01	-1.17	0.24
393	Gulf War	Iraq	Iraq	0.03	1.52	0.13
412	Iraq Deploy. to Kuwait	10/7/1994	Iraq	-0.00	-0.22	0.82
419	Desert Strike	9/3/1996	Iraq	0.00	0.32	0.75
422	UNSCOM I	11/13/1997	Iraq	-0.00	-0.03	0.97
427	US Embassy Bombings	8/7/1998	Non-state Actor (ICB)	0.01	1.31	0.19
429	UNSCOM II	11/2/1998	Iraq	0.01	0.15	0.88
430	Kosovo	2/22/1999	Yugoslavia	-0.00	-0.20	0.84
434	Afghanistan / 9-11	9/17/2001	Non-state Actor (ICB)	0.04	0.87	0.39
440	Iraq War	9/12/2002	United States of America	0.00	0.29	0.77
448	Iran Nuclear II	1/10/2006	Iran	0.00	0.12	0.90
450	N. Korea Nuclear III	5/5/2006	North Korea	-0.01	-0.42	0.67

Table 5.9: Cumulative abnormal returns and associated rank statistics for American international crises since 1983.

<i>ICB #</i>	<i>Event Name</i>	<i>Date</i>	<i>Revisionist</i>	<i>CAR</i>	<i>Rank Stat</i>	<i>p value</i>
363	Gulf of Syrte II	3/24/1986	United States of America	-0.00	0.04	0.97
386	Libyan Jets	12/21/1988	Libya	-0.01	-0.83	0.41
391	Invasion of Panama	12/15/1989	Panama	-0.01	-1.06	0.29
393	Gulf War	8/2/1990	Iraq	0.03	1.50	0.13
412	Iraq Deploy. to Kuwait	10/7/1994	Iraq	-0.00	-0.27	0.79
419	Desert Strike	9/3/1996	Iraq	0.01	0.33	0.74
422	UNSCOM I	11/13/1997	Iraq	-0.00	-0.14	0.89
427	US Embassy Bombings	8/7/1998	Non-state Actor (ICB)	0.01	1.17	0.24
429	UNSCOM II	11/2/1998	Iraq	0.01	0.17	0.87
430	Kosovo	2/22/1999	Yugoslavia	-0.00	-0.21	0.83
434	Afghanistan / 9-11	9/17/2001	Non-state Actor (ICB)	0.04	0.87	0.38
440	Iraq War	9/12/2002	United States of America	0.01	0.46	0.65
448	Iran Nuclear II	1/10/2006	Iran	0.00	0.14	0.89
450	N. Korea Nuclear III	5/5/2006	North Korea	-0.01	-0.43	0.67

Table 5.10: Cumulative abnormal returns and associated rank statistics for American international crises since 1986.

total cumulative abnormal return for the event, across all sampled firms, is attenuated because the effect of this serious crisis ran in opposite directions for the civilian and military economies.

That in mind, I repeated the above estimations while restricting the included firms to only those which derived at least 50% of their revenue from arms production across the 10 year period covered by the SIPRI yearbook series. Admittedly, this is not a perfect solution. SIPRI only began their “Top 100” arms production report in 1988, meaning that the extent to which any of these firms concentrated on military production for the previous twenty years amounts to an educated guess. This approach likely undervalues the military production of firms that were prominent defense contractors in the 1960s and 70s but are no longer, including General Motors and General Electric. This sampling refinement cut the number of relevant firms to just six: General Dynamics, Lockheed / Lockheed-Martin, Martin Marietta, McDonnell Douglas, Northrop-Grumman, and Raytheon.

As the results were again largely the same no matter which portfolio controls were used, I report only the series beginning in 1962 here so as to include the full set of possible events.

As with the expanded sample of firms, the key distinction between the 1991 Gulf

ICB #	Event Name	Date	Revisionist	CAR	Rank Stat	p value
196	Cuban Missile Crisis	10/16/1962	Russia	-0.01	-0.57	0.57
206	Panama Flag	1/9/1964	United States of America	-0.00	-0.47	0.64
210	Gulf of Tonkin	7/30/1964	Republic of Vietnam	-0.00	0.06	0.95
222	Six Day War	5/17/1967	Egypt	-0.00	0.00	1.00
239	Cienfuego Sub Base	9/16/1970	Russia	-0.02	-0.64	0.52
255	Yom Kippur War	10/5/1973	Multiple Actors (ICB)	-0.02	-0.88	0.38
259	USS Mayaguez	5/12/1975	Cambodia	0.01	-0.56	0.57
274	Poplar Tree (Paul Bunyan)	8/17/1976	North Korea	-0.02	-1.24	0.21
292	Shaba II	5/11/1978	Non-state Actor (ICB)	0.01	1.23	0.22
343	Invasion of Grenada	10/19/1983	Grenada	-0.00	-0.08	0.94
354	Nicaragua MiG-21s	11/6/1984	Russia	-0.01	-0.73	0.47
363	Gulf of Syrte II	3/24/1986	United States of America	0.03	1.45	0.15
386	Libyan Jets	12/21/1988	Libya	-0.02	-1.45	0.15
391	Invasion of Panama	12/15/1989	Panama	-0.00	-0.33	0.74
393	Gulf War	8/2/1990	Iraq	0.08	3.16	0.00
412	Iraq Deploy. to Kuwait	10/7/1994	Iraq	0.00	0.18	0.85
419	Desert Strike	9/3/1996	Iraq	-0.01	-0.43	0.67
422	UNSCOM I	11/13/1997	Iraq	-0.02	-1.10	0.27
427	US Embassy Bombings	8/7/1998	Non-state Actor (ICB)	0.01	1.08	0.28
429	UNSCOM II	11/2/1998	Iraq	-0.01	-0.61	0.54
430	Kosovo	2/22/1999	Yugoslavia	0.01	0.23	0.82
434	Afghanistan / 9-11	9/17/2001	Non-state Actor (ICB)	0.18	2.05	0.04
440	Iraq War	9/12/2002	United States of America	0.03	0.68	0.49
448	Iran Nuclear II	1/10/2006	Iran	0.00	0.02	0.98
450	N. Korea Nuclear III	5/5/2006	North Korea	-0.00	-0.06	0.95

Table 5.11: Cumulative abnormal returns and associated rank statistics for American international crises since 1962, core defense firms only.

War and 2003 Iraq War is preserved. In addition, the September 11th attacks are rightly treated as a surprising development, consistent with conventional wisdom, and two additional events flirt with traditional levels of statistical significance: the second Gulf of Syrte stand-off with Libya, and the subsequent crisis surrounding Libya's support for international terrorism in 1988. Oddly enough, however, the event window including the Lockerbie bombing saw a *decline* in market pricing for core American defense firms, indicating investor belief that the probability of war had in fact declined.

The series of crises for which there was no discernible change in equity pricing leads lends itself to two very distinct interpretations. First, some conflicts from the American record were preceded by years of public debate. The Iraq War certainly fits in this category, as does the 1989 invasion of Panama. In a similar fashion, other crises emerged as conclusions to much longer developments, even if they count as discrete events per ICB's coding standards. The Kosovo conflict, for example, came

on the heels of a number of related conflicts related to the dissolution of Yugoslavia, and as such its impact may have been muted. And, of course, some crises were not public knowledge at the time they occurred, depending on ICB's estimation of the proper crisis start point. The Cuban Missile Crisis, for example, was not immediately public, and as such defining the event window via ICB's crisis trigger date might load the dice against finding expected market behavior.

A separate interpretation is that some crises simply would not significantly affected the future earnings potential for even the most prominent defense contractors, whether or not that crisis produced a war. The 1983 invasion of Grenada, despite all of its embarrassing gaffes, was never in danger of turning into a protracted, expensive war for the United States. As such, rational investors may well have deemed that the Grenadine crisis was largely irrelevant to the prospects of General Dynamics, Raytheon, and the like. The series of distinct crises with Iraq during the 1990s might play a similar role, given that the practical ability of Iraq to defend itself from one-off American strikes after 1991 was so limited.

5.4.3 Regression Analysis

Although the non-parametric Corrado Zivney test appears to be the correct choice of estimation strategy given the short estimation and event windows that I have defined here, as well as a limited sample of firms, this does not allow us to directly test either hypothesis derived in section 5.4.1. An alternative approach is laid out by Bechtel and Schneider (2010), who predicted abnormal returns to European defense firms with a number of covariates describing European defense coordination summits. In their framework, a generalized time series estimation serves to account for the heteroskedasticity implied by the abnormal returns approach—when prices respond to new events, presumably that price change is accompanied by a greater number of transactions on those days, and thus possibly changing the variance of the underlying

Variables	Coef./se	Z	$p > Z $
USA Revisionist	-0.0029 (0.0058)	-0.51	0.61
Rival's capabilities share	-0.0207* (0.0084)	-2.45	0.01
Severity of violence	0.0078*** (0.0016)	4.88	0.00
Constant	-0.0117* (0.0046)	-2.55	0.01
N		266	

Table 5.12: A time-series cross-sectional estimation of daily abnormal returns in each crises' event window. All American crises since July 1962 are included, with the daily S&P 500 index serving as the only portfolio control for the purposes of constructing the abnormal return series.

data process on those days as well—while retaining the ability to directly test the contribution of any of their variables to the observance of an abnormal return.

I follow their example by modeling the daily abnormal return for each firm-event pairing in the event window of each crisis using only the “core firms” as defined above, where the abnormal return is still defined as in equation 5.3.3. The firm-event pairs serve as the panel identifiers in a time-series, cross-sectional analysis, meaning that the effective time series is only the two days of the event window, regardless of when they occur. I use ICB's measure of the “severity of violence” associated with a crisis to account for the likelihood that bloodier crises are likely to be seen as better investment opportunities for defense firms, regardless of the character of any associated information market intervention at the time.

Although this is clearly a very brief analysis, the results are broadly supportive of the hypotheses above. As the capabilities share of any rival state grows, the size of the associated cumulative abnormal return falls, which is consistent (but not uniquely so) with the logic of pre-crisis engagement. Events in which the United States is the revisionist party ought to be associated with indeterminate abnormal returns because of information market intervention, and while the reported coefficient shows as much, that could just as well be an artifact of the point that the United States is rarely

coded as the aggressor party in the ICB data.

5.5 Concluding Thoughts

What can we take away from this analysis? The results are somewhat mixed, as is to be expected when adapting a methodological tool to address a different research agenda. On one hand, the market's reaction after the Gulf War and 9/11 crises is what we would expect, with large price premiums associated with the defense industries were stoked by unanticipated news that caused observers to adjust their expectations for the possibility of an armed conflict involving the United States. On the other hand, the market's reaction for many of the other disputes included in my sample was quite muted.

I consider a few possible interpretations below:

Conflation with other financial events: My decision to use financial data as a measure of public engagement or disengagement in advance of a crisis was dictated by the deep-seated problems with the use of public opinion data, which I presented in Chapter 2. The type of analysis I present here might share an analogous flaw, however, in that it might be affected by contemporaneous *financial* events that overlap with these particular disputes. These could be quarterly financial reports that were either more or less optimistic than expected, accounting errors and financial restatements, or problems with a firm's own particular products—in other words, the types of questions that event studies were originally meant to examine. An additional coding effort to chart the baseline financial performance of the firms I included here might help distinguish between changes in investors' expectations that are driven by international disputes versus those driven by financial disclosures.

ICB crisis coding: I would also argue that this study also reveals some considerable issues with how we treat crisis and dispute data itself. I chose the International Crisis Behavior as a source of cases because its coding rules exclude the most spurious

cases, and is closer to capturing crisis negotiations as they are typically modeled in the theoretical literature. The virtues of avoiding, say, fishing-rights disputes between the United States and Canada came at a price, however. A dispute enters the ICB *systemic* sample on the first date that any party to the dispute perceives a threat to its security, or “core values.” Other parties to the dispute may not have perceived that same threat at the same time, however, and indeed the *actor*-level data reports a variety of onset dates for each party to a dispute, as each country’s own leadership perceives an external threat at different moments.

In this study, I relied on the system-level data and considered only those events that saw at least a limited American military response or redeployment. The implication is that the crisis trigger dates provided for a dispute, and in particular for those disputes that were triggered by an actor other than the United States, might not be the same date on which American leadership truly viewed themselves as embroiled in a crisis. In these cases, the market reaction might have truly spiked on some day immediately before or after the date reported by ICB. There is no clear *ex ante* method to predict that difference, other than to use the financial data to identify the most likely date in *association* with our existing sources of information about the American foreign policy record—but then, in this case, the same information would essentially be used to identify variables on both the left- and right-hand sides of the eventual estimation.

Note that relying on the actor-level ICB data, rather than the system-level version, would not resolve the issue either. Consider that, as I reported in Table 5.1, the actor-level ICB data indicates that the United States entered the Gulf War crisis on October 30th, 1990, and entered the war itself on January 16th, 1991. Yet the market became much more optimistic about the future earnings for American defense firms on August 2nd, 1990, the date of the Iraqi invasion of Kuwait. In this case, the market moved first, well before ICB’s account is resolved that the United States has entered

a crisis. That such wide differences should exist for as conceptually simplistic as the time at which a state enters crisis is worrying for all observational research into crisis diplomacy and war, including conventional studies that rely on public approval data.

CHAPTER 6

Concluding Thoughts and Observations

“[Mr. President]... you’ll need to scare the hell out of the American people.”

Sen. Arthur Vandenberg (Possibly apocryphal).

The notion that the need for public consent pacifies democratic states in their international conduct is seductive. While American presidents have long endorsed the spread of freedom and democracy abroad, this support was often merely cover for the support of friendly tyrants. As I pointed out in Chapter 2, however, in recent years this support for democratization has drawn more directly from research into the democratic peace: witness Bill Clinton’s claim that “democracies don’t attack one another.” The overwhelming majority of empirical research in the field has indeed found that democracies do not fight one another, and furthermore that democracies fight wars on a more selective basis and are more effective soldiers in the field. This latter set of findings on democracies’ warfighting ability comprise what Reiter and Stam (2002) call “democracy’s fourth virtue.”

The logical path by which representative government produces the necessary incentives to guide democratic leaders in this way, however, remains unclear. International relations scholars expect that free elections produce strong incentives for leaders. Because leaders hope to maintain office, they must presumably be mindful

of electoral consequences in the course of conducting foreign policy. As such, they ought to be more prudent in their foreign policy, and avoid the sort of rash challenges that often characterize illiberal regimes. They also ought to invest in military quality, and not quantity, so as to avoid casualties when wars do occur and thereby retain as much political support during those times as they possibly can. Of course, public consent can only shape leaders' decisions to the extent that the withdrawal of consent imposes real political costs. In order for elections to produce the incentives that we expect in international politics, voters must make the "right choices" at the ballot box.

One principal motivation for this dissertation project was the seeming incredibility of this claim. How can voters pose a credible political threat to their leaders to make these incentives tenable? Research into American voting behavior demonstrates that, no matter how well-intentioned the mass public may be, the conditions under which ordinary voters can make a reasoned choice about their candidate that fully and accurately evaluates leaders' prior decisions are difficult to satisfy. Voters know less about politics than do politicians. Worse yet, the primary methods by which voters can impose limits on political behavior in other policy areas are often not available to evaluate foreign policy matters. Few non-political signals, such as those connected to the state of the economy, are available at the onset of an international crisis; neither is party competition, as incumbent leaders enjoy wide discretion in limiting public knowledge of these events early in a crisis (Baum and Groeling 2010; Bennett, Lawrence and Livingston 2006). In other words, the same factors that can lead voters astray in the domestic arena—framing, priming, and gatekeeping—are no less likely in the domain of foreign policy.

And so, in practice, sometimes the structure imposed on foreign policy by the need for public consent breaks down. At times, the mass public might find out about poor or counter-productive decisions after the fact, but there is no guarantee that

such delayed electoral consequences will affect later governments' own behavior. The release of the "Pentagon Papers" in 1971 revealed American misdeeds in Indochina over the course of twenty years, but if anything the scandal only served to drive later decisions about the use of force further underground. In the United Kingdom, revelations that Tony Blair's "September Dossier" on Saddam Hussein was essentially a fraud resulted in a firestorm of controversy, a searing inquest by Parliament, and the eventual suicide of a Ministry of Defence official. Even then, much if not most of the punishment transpired under the regime of Blair's successor, Gordon Brown. By the time that voters had the chance to evaluate their government's behavior, the war had already happened. The prospect of electoral punishment *ex ante* was not enough to deter the behavior of Blair's government, and the full-bore realization of that punishment *ex post* was visited upon an entirely different leader.

With these thoughts in mind, in this dissertation I attempted to answer a straightforward question—can leaders talk their citizens into supporting wars they otherwise would not?—through a multiple methods approach. Chapters 3, 4, and 5 addressed the overarching question by examining markedly different social arenas. These different approaches addressed the distinct implications that the breakdown of public consent in democratic foreign policy-making can have and three different levels of analysis.

First, what I call "information market intervention" could affect diplomatic interactions between states. The extent to which it does is best explored by a formal model given that inter-state bargaining is a strategic game, which empirical studies are not well-suited to explore. While the model in Chapter 3, as presented, is quite plain, I found that under certain conditions information market intervention matters quite a lot. The option to build (or attempt to build) additional public support by misleading the public perverts the typical relationship between resolve and deterrent ability in some cases. There also appear to be configurations of state interest wherein

irresolute, “weak” democracies might engage in information market intervention even though it is not going to produce the desired concession short of fighting from their rival. In these cases, intervention manages to both avoid public consent before conflict but also generate additional costs to the executive after the fact.

At times, formal models make assumptions about behavior that are not grounded in empirical findings. For that reason, in Chapter 4, I moved from a discussion of international politics between state actors to a more modest question about individuals. In order for the key assumption underpinning my general theory and model to hold—that leaders are able to influence the public’s belief about a prospective war—it must be the case that individuals’ opinions do in fact change in response to their informational environment.

In practice, this required me to evaluate an entirely different research design, conducted at the personal level. No observational study would have allowed me to measure changes in public opinion in response to clear and distinguishable changes in the prevailing informational environment. In the course of looking for suitable measures of support for war, I was fortunate enough to find two seeming examples of information market intervention in action, first with respect to Libya in the 1980s and then with respect to Iraq in the 2000s. Finding that the Pentagon had already tried the sort of public diplomacy campaign that my view of the security and public opinion literatures suggested should happen was a sort of *prima facie* evidence that I was correct. These episodes thus served as motivation for the overall project, in addition to helping refine my choice of measures and stimuli for the experiment.

However, the experimental manipulations I designed seemed to matter relatively little for subjects’ evaluations of the case for war. The only stimulus that mattered over the course of two related panels was the least interesting one, wherein subjects who received any information at all about a threat from abroad were more willing to perceive that threat and to support the potential use of force in dealing with it. My

primary expectation, that appeals credited to figures from the military would be more effective than those associated with purely civilian expertise, was not supported; as I noted above, this expectation was borne out of both the relevant academic literature and the real-life information market intervention ahead of the Iraq War. On balance, I suspect the culprit in this regard is almost certainly the subtlety of the manipulations, which were delivered as short passages of text. Other experimental papers that have examined responses to public and war are much more blunt, with extremely stylized questionnaires that dispense with concerns about external validity in favor of more cleanly identifying the causal mechanism in question. It may be the case that mimicking that approach in testing my hypotheses would generate stronger results.

That said, I did find that different measures of public support seem to draw on different inherent attitudes. In particular, measures of general support for an incumbent president's foreign policy record draw primarily on subjects' own partisan loyalties. This is consistent with the newest research into public opinion and foreign policy, which holds that foreign policy is in fact quite similar to other issue domains. However, if we adopt a strategic view of international relations, then it becomes immediately apparent that some prevailing, ambiguous level of support is not what is most important for informing a leader's behavior during crisis; support for the use of force is. Encouragingly—from an empirical if not ethical point of view—subjects do not appear to rely on their partisanship in evaluating the use of force. Instead, they rely on other recent experiences with the costs of war, and in Chapter 4 I went on to find very suggestive evidence that opinions about the Iraq War are in fact a separate underlying attitude than is partisanship. This implies that an identification of elite influence on the public's support for war is possible given a new experimental protocol.

Chapters 3 and 4 took place at two polar extremes: the international level and

the individual level. The final component of the dissertation explored the behavior of the mass public in aggregate, which we might think of as addressing a level of analysis somewhere between the two. I showed in Chapter 2 how the most common approach to the study of public consent and foreign policy—investigations into the “rally ’round the flag” effect—does not appear to measure the public’s evaluation of its leaders with any reliability at all. Many events that are included with the Militarized Interstate Dispute data do not appear to be interactions that run any risk of triggering a wider conflict. Furthermore, even if such were not the case, many rally events and most rally predictors cannot be distinguished from sampling noise, in no small part because the measure that the rally literature uses is an evaluation of the president’s popularity, and not his or her conduct in any particular issue area.

Like public opinion polls, the equity pricing data in Chapter 5 is a measure of the public’s aggregate sentiment toward political issues. By selecting a particular class of firms to examine, we can in some sense pick and choose the underlying policy domain toward which we would like to evaluate public opinion. In this case, financial self-interest serves to ensure that the sort of public opinion we measure is in fact motivated by genuine concerns with the political item at hand, as opposed to other political events which might be of no interest at all.

However, because trends in equity pricing are one step conceptually removed from the series of individual events that I explored, the results I found in Chapter 5 can only only support my theoretical expectations by implication. *If* information market intervention has occurred, then interested parties (whether individuals, firms, or institutional investors) face a keen financial incentive to purchase equity in arms manufacturers so that they might share in the spoils of a potential war. When information market intervention has *not* occurred, then the onset of a crisis ought to come as a surprise. If we apply an event study framework and observe a price shock about a conflict, then we are led to believe that financial actors *did not know* about

these new potential earnings ahead of time, and therefore intervention must not have occurred. The precision with which we can estimate an effect comes at the cost of a certain amount of uncertainty about how the necessary political conditions would have unfolded at the time. In cases where the financial market failed to move in reaction to an international dispute, we cannot tell from the pricing trends themselves whether that was because intervention occurred ahead of time, or because the prospective conflict itself would not have been sufficiently costly to create a financial upside for investors. I should note here that this issue might be specific to cases involving the United States, all of whose disputes in the modern period have been with much weaker adversaries.

6.1 Observations About the Study of Foreign Policy

This last point suggests one of a few insights into the broader research agenda in foreign policy and public opinion that the dissertation project has provided me. The first is the extent to which this research question has been dominated by scholars examining the United States in isolation. The theoretical logic tying together elections and war, of course, was developed in the democratic peace literature which speaks to the entire set of democratic regimes over time. The study of public opinion and war, by contrast, has been developed more often than not by scholars trained in Americanist research, and not international relations. In practice, however, public opinion in the United States should matter relatively less for the conduct of its foreign policy because most prospective conflicts are so inexpensive in material terms. Since the publication of *War, Presidents and Public Opinion* (Mueller 1973b), researchers have been primarily concerned about the role of public opinion in large and costly conflicts such as those in Korea, Vietnam and Iraq, even though a cursory understanding of theoretical research into security and war shows that those conflicts are the exception, and not the norm. When threats are credible, then wars do not occur; in most

cases, the sheer military strength of the United States is enough to make its threats credible.

The second point I would like to raise is an epistemological concern about the relationship between crisis and war. Our theoretical framework for understanding why wars occur is very clear. Some exogenous change leads to a state being unsatisfied with the prevailing international status quo. That state, under certain conditions, will present a challenge to that status quo, leading to a circumstance that we call “crisis.” Depending on the conduct of negotiations at this stage, any given crisis might lead to a war.¹

In practical terms, however, identifying the distinction between crisis and war is difficult. In the course of replicating studies of the “rally ’round the flag” effect for Chapter 2, I was forced to conclude that this distinction is much more difficult than scholars typically understand. The most common practice in the literature is to treat the Militarized Interstate Dispute data as a set of all crises, and then identify those that produce a war from the Correlates of War’s interstate conflict data. However, upon closer examination, it is clear that a number of the “crises” identified by the MID data only qualify on a tenuous basis. The series of fishing disputes recorded between the United States and its neighbors, however genuine they might be, simply cannot be thought of as crises that may or may not produce a war.

The International Crisis Behavior data represents in part an attempt to deal with that issue, and as such I relied on it for the analysis in Chapter 5. However, this data suffers from its own particular shortcoming in that its coding rules are often baffling. The United States, as I mentioned previously, is only listed as joining the Gulf War crisis in October 1990, two months after the Iraqi invasion of Kuwait. Elsewhere, Grenada is coded as the revisionist party in its own invasion by American

¹This is not to ignore the contributions regarding limited war (e.g. Wagner 2000; ?). These applications are helpful, but do not reject the three-rung ladder of peace, crisis and war so much as they smooth out the transition from crisis to war over a series of intervening periods.

marines. Neither, I would argue with historians' view, nor even the viewpoint of the voting public at the time. The precise date that a crisis begins is of course especially important for an event study, but it also matters quite a bit for scholars relying on public opinion polls. Disagreement about dates by only a day or two could lead researchers to use different polling numbers in constructing a measure of public support for war, and as I showed in Chapter 2 there are more than enough problems with this measure in conceptual terms without introducing more error with respect to which particular polling numbers happen to be used.

BIBLIOGRAPHY

BIBLIOGRAPHY

- Achen, C.H. and D. Snidal. 1989. "Rational deterrence theory and comparative case studies." *World Politics* 41(2):143–169.
- Aldrich, John H., J.L. Sullivan and E. Borgida. 1989. "Foreign affairs and issue voting: Do presidential candidates "waltz before a blind audience?"." *The American Political Science Review* 83(1):123–141.
- Almond, Gabriel. 1960. *The American People and Foreign Policy*. New York: Praeger.
- Babst, Dean V. 1972. "A Force For Peace." *Industrial Research* 14:55–58.
- Bailey, T.A. and P.B. Ryan. 1979. *Hitler vs. Roosevelt: the Undeclared Naval War*. Free Press.
- Baker, William D. and John R. Oneal. 2001. "Patriotism or Opinion Leadership?: The Nature and Origins of the "Rally 'Round the Flag" Effect." *The Journal of Conflict Resolution* 45(5):661–687.
- Ball, Ray and Philip Brown. 1968. "An Empirical Evaluation of Accounting Income Numbers." *Journal of Accounting Research* 6(2):159–178.
- Barstow, David. 2008. "Behind TV Analysts, Pentagon's Hidden Hand." *New York Times* p. 1.
- Baum, Matthew. 2002. "The Constituent Foundations of the Rally-Round-the-Flag Phenomenon." *International Studies Quarterly* 46(2):263–298.
- Baum, Matthew. 2004. "Going Private: Public Opinion, Presidential Rhetoric and the Domestic Politics of Audience Costs in U.S. Foreign Policy Crises." *Journal of Conflict Resolution* 48:603–631.
- Baum, Matthew A. and Philip B.K. Potter. 2008. "The Relationships Between Mass Media, Public Opinion, and Foreign Policy: Toward a Theoretical Synthesis." *Annual Review of Political Science* 11(1):39–65.
- Baum, Matthew A. and Tim Groeling. 2008a. "Crossing the Water's Edge: Elite Rhetoric, Media Coverage, and the Rally-Round-the-Flag Phenomenon." *The Journal of Politics* 70(04):1065–1085.

- Baum, Matthew A. and Tim Groeling. 2008*b*. "Shot by the Messenger: Partisan Cues and Public Opinion Regarding National Security and War." *Political Behavior* 31(2):157–186.
- Baum, Matthew A and Tim Groeling. 2010. "Reality Asserts Itself: Public Opinion on Iraq and the Elasticity of Reality." *International Organization* 64(3):443–479.
- Bechtel, Michael M and Gerald Schneider. 2010. "Eliciting Substance from 'Hot Air': Financial Market Responses to EU Summit Decisions on European Defense." *International Organization* 64(Spring):199–223.
- Bennett, W. Lance. 1990. "Toward a Theory of Press-State Relations in the United States." *The Journal of Communication* 40(2):103–127.
- Bennett, W. Lance, Regina G. Lawrence and Steven Livingston. 2006. "None Dare Call It Torture: Indexing and the Limits of Press Independence in the Abu Ghraib Scandal." *Journal of Communication* 56(3):467–485.
- Berinsky, Adam J. 2004. *Silent Voices: Public Opinion and Political Participation in America*. Princeton: Princeton University Press.
- Berinsky, Adam J. 2007. "Assuming the Costs of War: Events, Elites, and American Public Support for Military Conflict." *The Journal of Politics* 69(4):975–997.
- Berinsky, Adam J. 2009. *In Time Of War: Understanding American Public Opinion From World War II To Iraq*. University of Chicago Press.
- Bernhard, William and David Leblang. 2006. *Democratic Processes and Financial Markets*. New York: Cambridge Univ Press.
- Brecher, M. 1977. "Toward a Theory of International Crisis Behavior: A Preliminary Report." *International Studies Quarterly* 21(1):39–74.
- Bueno de Mesquita, Bruce, Alistair Smith, Randolph Siverson and James Morrow. 2003. *The Logic of Political Survival*. Cambridge: MIT Press.
- Bueno de Mesquita, Bruce and David Lalman. 1992. *Reason and War*. New Haven: Yale University Press.
- Bush, George Walker. 2004. "White House Press Conference." Accessed at <http://www.whitehouse.gov/news/releases/2004/11/20041112-5.html>.
- Campbell, Angus, Philip E. Converse, Warren E. Miller and Donald E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Campbell, Cynthia J. and Charles E. Wesley. 1993. "Measuring security price performance using daily NASDAQ returns." *Journal of Financial Economics* 33(1):73–92.
- Campbell, J.Y., Lo A.W. and A.C. MacKinlay. 1997. *The econometrics of financial markets*. Princeton University Press.

- Chapman, Terrence L. and Dan Reiter. 2004. "The United Nations Security Council and the 'Rally 'Round the Flag Effect'." *The Journal of Conflict Resolution* 48(6):886–909.
- Chong, D. 1993. "How people think, reason, and feel about rights and liberties." *American Journal of Political Science* 37(3):867–899.
- Clinton, William Jefferson. 1994. "State of the Union Address." Accessed at <http://www.washingtonpost.com/wp-srv/politics/special/states/docs/sou94.htm>.
- Colaresi, Michael. 2007. "The Benefit of the Doubt: Testing an Informational Theory of the Rally Effect." *International Organization* 61(01):99–143.
- Converse, Phil. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, ed. D.E. Apter. New York: Free Press.
- Corrado, Charles J. and Terry L. Zivney. 1992. "The Specification and Power of the Sign Test in Event Study Hypothesis Tests Using Daily Stock Returns." *Journal of Financial and Quantitative Analysis* 27(03):465–478.
- Croco, Sarah E. 2011. "The Decider's Dilemma: Leader Culpability, War Outcomes, and Domestic Punishment." *American Political Science Review* 105(03):457–477.
- Dafoe, Allan. 2011. "Statistical Critiques of the Democratic Peace: Caveat Emptor." *American Journal of Political Science* 55(2):247–262.
- DellaVigna, Stefano and Eliana La Ferrara. 2010. "Detecting illegal arms trade." *American Economic Journal: Economic Policy* 2(November):26–57.
- Downes, A and T Sechser. 2012. "The Illusion of Democratic Capability." *International Organization* 66(Summer):457–89.
- Doyle, Michael. 1986. "Liberalism and World Politics." *American Political Science Review* 80:1151–1169.
- Druckman, James N. 2001. "On the Limits of Framing Effects: Who Can Frame?" *The Journal of Politics* 63(04):1041–1066.
- Druckman, James N. and Michael Parkin. 2008. "The Impact of Media Bias: How Editorial Slant Affects Voters." *The Journal of Politics* 67(04):1030–1049.
- Dube, Arindrajit, Ethan Kaplan and Suresh Naidu. 2011. "Coups, Corporations, and Classified Information." *The Quarterly Journal of Economics* 126(3):1375–1409.
- Entman, Robert M. 2004. *Projections of power: Framing news, public opinion, and US foreign policy*. University of Chicago Press.
- Fama, Eugene F., Lawrence Fisher, Michael C. Jensen and Richard Roll. 1969. "The Adjustment of Stock Prices to New Information." *International Economic Review* 10(1):1–21.

- Fearon, James. 1994a. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* 88:577–92.
- Fearon, James. 1994b. "Signaling versus the Balance of Power and Interests: An Empirical Test of a Crisis Bargaining Model." *Journal of Conflict Resolution* 38:236–269.
- Fiorina, M.P. 1981. *Retrospective voting in American national elections*. Yale University Press.
- Fordham, Benjamin O. and Christopher C. Sarver. 2001. "Militarized Interstate Disputes and United States Uses of Force." *International Studies Quarterly* 45(3):455–466.
- Gallup and Fortune Polls. 1941. *The Public Opinion Quarterly* 5(3):470–497.
URL: <http://www.jstor.org/stable/2745138>
- Gartner, Scott Sigmund. 2008. "The Multiple Effects of Casualties on Public Support for War: An Experimental Approach." *American Political Science Review* 102(01):95–106.
- Gartzke, E. 2007. "The capitalist peace." *American Journal of Political Science* 51(1):166–191.
- Geddes, Barbara. 1990. "How the cases you choose affect the answers you get: Selection bias in comparative politics." *Political analysis* 2(1):131.
- Gelpi, Christopher, Peter D. Feaver and Jason Reifler. 2006. "Success matters: Casualty sensitivity and the war in Iraq." *International Security* 30(3):7–46.
- Ghosn, Faten, Glenn Palmer and Stuart A. Bremer. 2004. "The MID3 Data Set, 1993-2001: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science* 21(2):133–154.
- Gleditsch, Kristian S. and Michael D. Ward. 2000. "War and Peace in Space and Time: The Role of Democratization." *International Studies Quarterly* 44(1):1–29.
- Goemans, H. E. 2000. "Fighting for Survival: The Fate of Leaders and the Duration of War." *The Journal of Conflict Resolution* 44(5):555–579.
- Greene, W.H. and D.A. Hensher. 2010. *Modeling ordered choices: A primer*. Cambridge Univ Press.
- Groeling, Tim and Matthew A. Baum. 2008. "Crossing the Water's Edge: Elite Rhetoric, Media Coverage, and the Rally-Round-the-Flag Phenomenon." *The Journal of Politics* 70(04):1065.
- Guidolin, Massimo and Eliana La Ferrara. 2007. "Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms?" *American Economic Review* 97(5):1978–1993.

- Guidolin, Massimo and Eliana La Ferrara. 2007. "Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms?" *The American Economic Review* 97(5):1978–1993.
- Harsanyi, John C. 1967. "Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model." *Management Science* 14:159–182.
- Herrmann, Richard K., Philip E. Tetlock and Penny S. Visser. 1999. "Mass public decisions to go to war: A cognitive-interactionist framework." *American Political Science Review* 93(3):553–573.
- Holsti, O.R. 1996. *Public opinion and American foreign policy*. Cambridge University Press.
- Isernia, Pierangelo, Zoltan Juhasz and Hans Rattinger. 2002. "Foreign Policy and the Rational Public in Comparative Perspective." *Journal of Conflict Resolution* 46:201–224.
- Jacobson, Gary. 2007. "A divider, not a uniter: George W. Bush and the American people."
- Jentleson, B.W. 1992. "The pretty prudent public: Post post-Vietnam American opinion on the use of military force." *International Studies Quarterly* 36(1):49–73.
- Johnson, Dominic D.P. and Dominic Tierney. 2006. *Failing to Win: Perceptions of Victory and Defeat in International Politics*. Harvard University Press.
- Jones, D., S. Bremer and J. D. Singer. 1996. "Militarized Interstate Disputes, 1816–1992: Rationale, Coding Rules, and Empirical Patterns." *Conflict Management and Peace Science* 15(2):163–213.
- Kant, Immanuel. 1970a. Idea for a Universal History with a Cosmopolitan Purpose [1795]. In *Kant's Political Writings*, ed. Hans Reiss. Cambridge: Cambridge University Press.
- Kant, Immanuel. 1970b. Perpetual Peace: a Philosophical Sketch [1784]. In *Kant's Political Writings*, ed. Hans Reiss. Cambridge: Cambridge University Press.
- Knopf, Jeffrey. 1998. "How Rational is the 'Rational Public?'" *Journal of Conflict Resolution* 42:544–71.
- Lupia, Arthur. 2002. Who can persuade whom? Implications from the nexus of psychology and rational choice theory. In *Thinking About Political Psychology*, ed. James H. Kuklinski. Cambridge: Cambridge Studies in Public Opinion and Political Psychology pp. 51–88.
- MacKinlay, A. Craig. 1997. "Event Studies in Economics and Finance." *Journal of Economic Literature* 35(1):13–39.

- Maoz, Zeev and Bruce Russett. 1993. "Normative and structural causes of democratic peace, 1946-1986." *American Political Science Review* 87(3):624-638.
- Marcus, G.E., W.R. Neuman and M. MacKuen. 2000. *Affective Intelligence and Political Judgment*. University of Chicago Press.
- Morrow, James D. 1985. "A continuous-outcome expected utility theory of war." *Journal of Conflict Resolution* 29(3):473-502.
- Mueller, John E. 1973a. *War, Presidents, and Public Opinion*. New York: John Wiley & Sons.
- Mueller, John E. 1973b. *War, Presidents and Public Opinion*. Number 408 New York: John Wiley and Sons.
- Oneal, John R. and Anna Lillian Bryan. 1995. "The Rally 'Round the Flag Effect in U. S. Foreign Policy Crises, 1950-1985." *Political Behavior* 17(4):379-401.
- Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public*. Chicago: University of Chicago Press.
- Potter, Philip K. and Matthew Baum. 2009. "The Democratic Peace, Domestic Audience Costs, and Political Communication." *Political Communication* pp. 1-35.
- Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton: Princeton University Press.
- Powell, Robert. 2006. "War as a commitment problem." *International Organization* 60(01):169-203.
- Putnam, R.D. 1988. "Diplomacy and domestic politics: the logic of two-level games." *International Organization* 42(3):427-460.
- Reiter, Dan and Allan C. Stam. 2002. *Democracies at War*. Princeton University Press.
- Roberts, Andrew. 2009. *Masters and Commanders: How Four Titans Won the War in the West, 1941-1945*. Harper.
- Rubinstein, A. 1982. "Perfect equilibrium in a bargaining model." *Econometrica: Journal of the Econometric Society* 50(1):97-110.
- Rummel, R.J. 1975-1981. *Understanding Conflict and War*. Vol. 1-5 Beverly Hills: Sage.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Harvard University Press.
- Schultz, Kenneth. 1998. "Domestic opposition and signaling in international crises." *The American Political Science Review* 92(4):829-844.

- Schultz, Kenneth. 2001. *Democracy and Coercive Diplomacy*. New York: Cambridge University Press.
- Shapiro, Robert Y. and Lawrence R. Jacobs. 2000. Who Leads and Who Follows? US Presidents, Public Opinion and Foreign Policy. In *Decisionmaking in a Glass House*, ed. Brigitte L. Nacos, Robert Y. Shapiro and Pierangelo Isernia. New York: Rowman and Littlefield.
- Shleifer, A. 2000. *Inefficient markets: An introduction to behavioral finance*. Oxford University Press.
- Singer, J David. 1961. "The level-of-analysis problem in international relations." *World Politics* 14(1):77–92.
- Slantchev, Branislav L. 2003. "The power to hurt: Costly conflict with completely informed states." *American Political Science Review* 97(1):123–133.
- Slantchev, Branislav L. and Ahmer Tarar. 2011. "Mutual Optimism as a Rationalist Cause of War." *American Journal of Political Science* 55(1).
URL: <http://escholarship.org/uc/item/4p39g3px.pdf>
- Small, Melvin and David Singer. 1976. "The War Proneness of Democratic Regimes." *Jerusalem Journal of International Relations* 1:50–69.
- Talbott, Strobe. 1996. "Democracy and the National Interest." *Foreign Affairs* 75(6):47–64.
- Tomz, M. 2007. "Domestic Audience Costs in International Relations: An Experimental Approach." *International Organization* 61(4):821–840.
- Wagner, R. Harrison. 2000. "Bargaining and War." *American Journal of Political Science* 44(3):469–484.
- Weeks, JL. 2008. "Autocratic audience costs: regime type and signaling resolve." *International Organization* 62(1):35–64.
- Woodward, Bob. 1986. "Gadhafi Target of Secret U.S. Deception Plan—Elaborate Campaign Included Disinformation That Appeared as Fact in American Media." *The Washington Post* p. A1.
- Zaller, John. 1992. *The Natures and Origin of Mass Opinion*. New York: Cambridge University Press.