

Bureaucracies, Advisers, and Interstate Crises: Affiliation Affects Information Provision, Not Policy Prescriptions*

Robert Schub[†]

This Draft: September 8, 2018

Abstract

Whether international crises end peacefully depends upon the information that leaders possess. To understand how leaders acquire this information, I develop and test an informational theory of bureaucracies. Time-constrained leaders delegate information collection and provision tasks to senior advisers heading the state's national security bureaucracies. A division of labor between bureaucracies breeds comparative specialization among affiliated advisers. Some specialize in adversaries' political attributes (e.g., resolve) which are relatively hard to observe while others emphasize military attributes (e.g., force quantities) which are easier to observe. The type of information an adviser provides follows from her bureaucracy's specialization. The argument challenges canonical accounts in which parochial interests cause advisers to endorse clashing policy prescriptions. Rather, where you sit affects the informational content you provide and uncertainty you express, not the policy you promote. To systematically evaluate the theory, I construct a corpus of over 1,100 memos and transcripts from US Cold War crises. Using automated text analysis and qualitative coding to measure advisory input, I find strong support for the argument. The paper reveals benefits to expanding analytical attention beyond leaders to incorporate advisers. Additionally, the findings suggest that marginalizing bureaucracies curtails information flows which fosters poor decision making.

*For fruitful discussions on this and related projects, I thank Muhammet Bas, Julie Faller, Jeff Frieden, Jeff Friedman, Connor Huff, Iain Johnston, Tyler Jost, Josh Kertzer, Andy Kydd, David Lindsey, John Marshall, Noah Nathan, Steve Rosen, Anne Sartori, Dustin Tingley, Alex Weisiger, Ariel White, Keren Yarhi-Milo, George Yin, and audience members at APSA 2017.

[†]University of Nebraska–Lincoln, robert.schub@unl.edu

1 Introduction

Information is often the root cause of international conflict. Crisis diplomacy can devolve into violence when opposing leaders have access to private information that they cannot credibly reveal to one another (Morrow 1989; Fearon 1995; Powell 1999; Fey and Ramsay 2011; Slantchev and Tarar 2011) or disagree about the implications of available information (Blainey 1988; Johnson 2004). Given the importance of information for explaining war and peace, it is vital to understand how leaders actually acquire it. Time constraints mean they cannot do this on their own and rarely have the requisite information at their fingertips. Instead, they delegate information collection tasks to their senior advisers who convey this information during advisory processes. Advisers thus play a critical, though often overlooked, role in international politics (George 1980; Janis 1982; Saunders 2017). This paper brings advisers to the fore, stretching the locus of analysis beyond leaders which have assumed a privileged position in a wave of recent scholarship on their backgrounds (Colgan 2013; Horowitz and Stam 2014; Dafoe and Caughey 2016), institutional incentives (Chiozza and Goemans 2011; Croco 2011; Weeks 2014), and dispositions and beliefs (Hermann 1980; Saunders 2011; Yarhi-Milo 2014; Kertzer 2016; Whitlark 2017).

The senior advisers that leaders rely upon typically head the core bureaucracies of the national security apparatus. The bureaucratic role an adviser occupies affects the input she provides to the leader, but in a vastly different way than previously theorized. In my account, bureaucratic role affects the type of information advisers provide and uncertainty they express, not the policies that they promote. Leaders demand information; senior advisers can supply it. Akin to many principal-agent setups, agents (advisers) develop informational advantages over the principal (leader).¹ In addition to being better informed than the leader, advisers are *differentially* informed compared to one another. Informational differences emerge from a division of labor across bureaucracies. Functional specialization between bureaucracies causes affiliated advisers to become differentially informed. Military advisers in, say, a ministry of defense, become relative experts on an adversary's military attributes or characteristics. Foreign policy advisers in, say, a ministry of foreign affairs, become relative experts on an adversary's political attributes or characteristics. Advisers, consequently, emphasize different informational content when participating in deliberative processes. To be sure, differentiation is partial, not absolute and the walls of bureaucratic silos are permeable—military agencies employ political experts and vice-versa. Some bureaucracies, such as intelligence agencies, fall along a continuum between these ideal types. Granting partial substantive overlap, bureaucratic affiliation heavily shapes the information advisers acquire and content they emphasize to the leader.

Bureaucratic position affects the uncertainty advisers express as well as the content they em-

¹See Bendor, Glazer and Hammond (2001) and Gailmard and Patty (2012) for reviews of these vast literatures.

phasize. These two implications are directly related. Some domains are relatively information rich while others tend to be information poor. I posit that information is more readily available on an adversary's military traits, making them subject to less uncertainty. Analysts can often observe force quantities, postures, and locations. Political characteristics—such as an adversary's resolve or domestic political landscape that influences whether battlefield accomplishments can be translated into political objectives—are harder to assess and prone to greater uncertainty.

This paper challenges the canonical account of how bureaucratic affiliation guides advisers' input, as originally stated in Allison (1969). Allison's Bureaucratic Politics Model contends that bureaucracies' parochial interests, for instance on budgetary matters, cause advisers' preferences to diverge. Advisers engage in a game of "pulling and hauling" to sway policy outcomes to favor their bureaucracy. Famously echoing Miles' Law, "where you stand depends on where you sit" (Allison 1969, p. 711). Foreign policy advisers in a ministry of foreign affairs promote diplomatic solutions while military advisers (in or out of uniform) at a ministry of defense advocate militarily aggressive policies. I join a chorus of critics and challenge this account (Krasner 1972; Art 1973; Snyder and Diesing 1977; Welch 1992; Bendor and Hammond 1992) with a particular emphasis on its inappropriateness for the policy advisory process during crises. Its inappropriateness stems from two factors. First, adviser preferences tend to converge during crises. Second, advisory structures are hierarchical, not flat. Leaders sit atop this hierarchy equipped with tools to mitigate preference divergences among advisers.

I test the theoretical implications by examining advisory processes during US Cold War crises. I collect and analyze a corpus of documents from security deliberations at the highest levels of decision making, largely drawn from the Foreign Relations of the United States series (FRUS) which is the official historical record of the State Department. Employing qualitative methods and a variety of automated text analysis tools, I generate measures capturing each bureaucracy's advisory input for each crisis. Coupling speakers' original and private words with capabilities to analyze texts at scale offers a new means to measure advisers' emphases and beliefs. Specifically, I measure a bureaucracy's content (political vs. military), uncertainty, and relative hawkishness.

Three findings emerge from the analyses. First, foreign policy advisers, as embodied by those at the State Department, were twice as likely to discuss an adversary's political attributes as compared to colleagues from other bureaucracies. Second, these advisers express more uncertainty than those with different affiliations. Third, an adviser's bureaucratic role does not predict the policy positions he advocates. Diplomats are no more likely than their counterparts from Defense to endorse the less militarily aggressive policy response. The findings strongly suggest bureaucratic role matters, but it matters in a different way than the conventional wisdom suggests. Consistent with my theory, bureaucratic affiliation affects the type of information advisers provide and the certainty accompanying that advice, but not the hawkishness of the policies they champion.

The theory and results elucidate (1) when advisory processes are likely to produce suboptimal decisions and (2) the mechanics of information transmission. First, though a full explication of how advisory processes can contribute to war is beyond this paper's scope, it does show the danger of leaders marginalizing key bureaucracies in the national security apparatus. Because bureaucracies bring distinct expertise, curtailing any one bureaucracy's input will cause leaders to miss important pieces of information. For instance, preceding the Iraq War George W. Bush limited the State Department's input and neglected its expertise. Marginalization distorted the administration's assessments of postwar stability.² The administration, as a result, marched to war with an incomplete and inaccurate conception of what political end-states could be achieved through force and at what cost. Second, an informational account of bureaucracies provides a real-world grounding for how information guides leaders' strategic choices. Much past work suggests that a leader's information follows from the strategic environment—such as, whether information is complete and what signals opponents send (for instance, [Lake and Powell \(1999\)](#); [Powell \(2017\)](#)). These accounts assume leaders acquire and use whatever credible information is available, subject to cost constraints ([Arena and Wolford 2012](#)). My account elucidates the mechanics of how leaders might actually go about acquiring this information. Others have illuminated information sources that leaders privilege ([Hall 2015](#); [Yarhi-Milo 2014](#)), advisory ordering principles ([George 1980](#); [Bendor and Hammond 1992](#)), leader and adviser personal attributes ([Saunders 2017](#)), and civil-military dimensions of information provision ([Huntington 1957](#); [Feaver 2003](#); [Brooks 2008](#)). I believe an emphasis on bureaucratic role, especially among civilian advisers, is a novel one.

The rest of the paper proceeds as follows. It first explicates Allison's Bureaucratic Politics Model, the implications that follow, and reasons to doubt those implications enjoy empirical support. Next, it develops the informational theory of bureaucracies, explaining why bureaucratic role shapes advisory content and uncertainty. The following section presents the research design, describing case and document selection criteria as well as the measurement strategies, which include supervised learning, dictionary approaches, and qualitative coding. I then present results, consider alternative explanations and generalizability, and provide intuitions for when advisory processes might increase the risk of conflict. A final section concludes, tracing the implications of leaders who do not value information or expertise.

2 Bureaucratic Preferences: Sitting and Standing

What does existing scholarship tell us about how bureaucracies affect advisory processes during interstate crises? Many studies speak to crisis advisory processes generally, analyzing how

²For instance as contained in the *Future of Iraq Project*. On flaws in postwar planning more generally, see [Ricks \(2006\)](#); [Bensahel \(2007\)](#); [Lake \(2010\)](#); [Rapport \(2012\)](#); [Saunders \(2017\)](#).

decision-making bodies are structured (Hermann and Preston 1994; Hermann 2001) or endorsing optimal arrangements such as multiple advocacy (George 1980; George and Stern 2002), overlapping responsibilities (Rudalevige 2005; Dickinson 2005), or having leaders and advisers with extensive experience (Saunders 2017). Others document pathologies that might emerge such as groupthink (Janis 1982; Herek, Janis and Huth 1987). While speaking to advisory processes, actual bureaucratic affiliations play a limited role in these accounts. Accordingly they do not generate clear implications for how input from a secretary of defense and a secretary of state might differ.

The most prominent theory that does stipulate how bureaucratic affiliation affects advisory input assumes preference heterogeneity among advisers. As first developed in Allison (1969), and reframed and expanded upon in subsequent work (Allison 1971; Allison and Halperin 1972; Allison and Zelikow 1999), advisers' parochial interests dictate or, more modestly, affect the policies they endorse.³ Competing preferences drive advisers and even leaders to engage in a political battle to shift policy in directions that advantage their bureaucracies.

The Bureaucratic Politics Model produces testable implications for the behavior of advisers during interstate crises. Intuitively, advisers should endorse policies that advance their bureaucracies' interests. This rolls back the inquiry to one of discerning bureaucracies' interests. Interests on budgetary matters are straightforward: advisers endorse policies that will bring more money to their bureaucracy. Interests on crisis policy choices are somewhat less immediate (though those policy choices might have budgetary implications). A fair reading might conclude that diplomats or foreign policy advisers, such as from a ministry of foreign affairs, should prefer relatively peaceful policies. In contrast, military advisers, such as from a ministry of defense, should prefer relatively militaristic policies. An even more generous reading might restrict the latter category to exclude advisers in uniform, focusing only on civilian defense advisers. Several studies contend that military officials are not systematically more hawkish than civilian counterparts (Huntington 1957; Betts 1991), especially for initial (as opposed to escalatory) force employment decisions and for missions with humanitarian objectives (Feaver and Gelpi 2004). Focusing on gaps between civilian advisers affiliated with different bureaucracies skirts the complications introduced by these documented civilian-military differences and stacks the deck in favor of Allison's account.

Critics might view this reading of the Bureaucratic Politics Model's implications as overly sweeping, noting that parochial interests are contextually dependent. Marsh and Jones (2017) provides a nice example. The study documents competing interests weighing on the State Department before the Afghanistan surge and Libya strikes to shed light on Secretary Clinton's hawkish stances. This nuance, while inevitably valuable for any given case, undercuts the theoretical power

³For Allison and Halperin (1972) (p. 44), "government leaders have competitive, not homogenous interests; priorities and perceptions are shaped by positions." Allison and Zelikow (1999, p. 307) tempers the prediction, stipulating that where you stand "is substantially affected by" where you sit.

and utility of the Bureaucratic Politics Model. At best it circumscribes our ability to make testable predictions. At worst it renders the theory unfalsifiable, forcing scholars to infer bureaucratic interests from post hoc analyses of their policy prescriptions.

Allison's parochial theory of bureaucratic politics yields the following testable implication:

Parochial Theory Expectation 1: When advising leaders during crises, civilian foreign policy advisers (i.e., those in foreign affairs ministries) endorse less militarily aggressive policies than civilian advisers with different substantive assignments.

Critics have raised many reasons to doubt the merits of the Bureaucratic Politics Model (Krasner 1972; Art 1973; Snyder and Diesing 1977; Welch 1992; Bendor and Hammond 1992).⁴ Most crucially for this paper, several factors mitigate parochial preference divergences, especially during the advisory processes for international crises.⁵ First, when contemplating the use of military force, preferences are especially likely to converge on some conception of promoting general welfare of the state. Allison (1969) (p. 711) himself hints at this relationship between stakes (wars vs. budgets) and parochialism (low vs. high). Second, leaders who can select their advisers choose allies with aligned preferences which can supersede bureaucratic affinities (Krehbiel 1991; Bendor and Hammond 1992). Krasner (1972) makes the point succinctly, "The most important "action-channel" in the government is the President's ear. The President has a major role in determining who whispers in it." Senior advisers—e.g., those at the cabinet level—are thus likely "attuned" to the leader's interests (Betts 1991, p. 40).⁶

Third, even if preferences do diverge due to parochialism, leaders have tools to mitigate the problem. Sanctioning, for instance, can elicit truthful advisory input. If advisers care whether the leader thinks them competent or retains their services, then incentives exist for them to offer the best advice given the information they have, rather than pursue narrow bureaucratic interests

⁴For Welch (1992) and Bendor and Hammond (1992), it is unclear whether the Bureaucratic Politics Model is actually a model of bureaucracies or of the individuals in senior positions. Allison (1969, p. 709) assigns causal importance to the fact that "metabolisms differ." This sounds like an individual rather than bureaucracy-level explanation, barring intriguing selection or cafeteria effects.

⁵My expectations apply within the scope of advisory input during interstate crises. They do not speak to long term strategy choices, such as investing in nuclear deterrence at the expense of conventional forces. Nor does it apply to the implementation of a leader's orders during crises, such as whether a diplomat conveys offers in accordance with the leader's demands. What the theory does apply to—the information leaders obtain when making choices than can escalate to war—is important.

⁶The relevance and utility of adviser selection diminishes under some conditions. Leaders may be constrained in their ability to select advisers—for instance, in a parliamentary coalition government (Preston and 't Hart 1999). Also, a competence versus loyalty tradeoff can undermine preference convergence. Adviser preferences may diverge if such a tradeoff exists and leaders prioritize competence over ideological congruence. Lewis (2011) provides an overview of this tradeoff in US appointments.

(Meirowitz 2006). Leaders can dismiss advisers whose performance they deem inadequate. President Kennedy's dismissal of two senior CIA officials, Allen Dulles and Richard Bissell, after the Bay of Pigs is illustrative.⁷ Fourth, beyond deductive logic, prior empirical analysis (Snyder and Diesing 1977; Rhodes 1994) and several anecdotes in which the posited preference divergences run in the reverse direction encourage further skepticism of the canonical account. Secretary of State Clinton's hawkish stance on intervening in Libya as compared to Secretary of Defense Gates' dovish one is a recent counter-example.

In sum, when stakes are high, shared objectives such as enhancing national security supersede particular interests, such as maximizing budgets. Accordingly, bureaucratic position is an unreliable predictor of an adviser's policy prescriptions. *I anticipate that the parochial theory receives little empirical support.*

3 An Informational Theory of Bureaucracies and Crises

I now develop an affirmative informational theory of bureaucracies. It specifies the types of information leaders need, how easy each type is to collect, and who provides it to leaders.

The Information Leaders Need

Richard Neustadt (1990, p. 128-129) wrote of the US context, "A president is helped by what he gets into his mind. His first essential need is information." This is especially true during international crises when leaders must assess what each side can expect to achieve in a military conflict and at what cost. With war as an outside option, optimal strategy decisions—such as how generous or stingy to be in crisis diplomacy—hinge upon expected payoffs to fighting. The probability of military victory and the costs of fighting are key determinants of these payoffs and thus affect which peaceful settlements are preferable to conflict (Blainey 1988; Fearon 1995; Powell 1999). Further unpacking these elements offers a helpful framework for understanding advisory processes. Both the probability of victory and war costs depend on military *and* political attributes. A distinction along these broad lines—military versus political—elucidates the information leaders must collect to guide their policy choices during crises. Crucially, I contend the political is harder to assess than the military and thus subject to more uncertainty.

Military attributes affect the probability of victory and the costs incurred during fighting. These attributes include the quantity and quality of an adversary's military assets (Friedberg 1987), its

⁷Extreme sanctioning, such as personalist dictators assassinating advisers, alters the dynamics. Advisers will likely converge on a common prescription aligned with their perception of the leader's preferences. Former regime members in Saddam's Iraq, for example, indicate that no one dared challenge Saddam's narrative before the Iraq War (Lake 2010).

defense spending (Lebovic 1995), its personnel and their training, its doctrine and force posture (Biddle 2004; Narang 2014), and a net assessment of how these characteristics interact with a state's own military capabilities and doctrine (Posen 1984; Cohen 1988; Rosen 1991). Leaders need information on each of these elements to develop an accurate picture of the crisis landscape.

The observable nature of many military characteristics aids the information gathering and assessment process. While they are imperfectly observable, imperfect signals can yield valuable information. Military demonstrations (e.g., North Korean tests of new missile technologies) and publicly available documentation (e.g., Soviet pronouncements of military personnel reductions in the mid-1950s) facilitate assessments, whether quantitative or qualitative. Aerial surveillance technologies provide information for simple bean-counting or monitoring opponent positioning and posture (e.g., US detection of Soviet missiles in Cuba). Assessing opponent doctrine poses a tougher estimative challenge, but states can glean insights from recent conflicts and pronouncements. Hence, the tradition of non-combatants monitoring military engagements (e.g., US observers of the Crimean War evaluating the effect of technological developments on combat dynamics) and carefully studying adversary military doctrine. A net assessment of opposing forces, which moves beyond one-sided analysis, is inherently more difficult. States rely on qualitative insights, historical studies, and simulations. I do not claim that information on military characteristics is always abundant. States vary in transparency levels and may even have incentives to mask information (Meirowitz and Sartori 2008; Slantchev 2010). While acknowledging variation, I contend that military information is often available and obtainable. Greater information allows advisers and leaders to form higher certainty assessments over these attributes.

An adversary's political attributes similarly affect both the probability of victory and war costs. I consider three forms of these political attributes and highlight difficulties in assessing each.

(1) *Resolve*, or a state's steadfastness to a policy despite "temptations to back down" (Kertzer 2016), is the prototypical political attribute that affects expected conflict payoffs. Resolve encompasses an array of concepts, such as a state's willingness to use force at all or the costs it will endure during the course of using force. Resolve can also relate to the value the adversary places on the issue in dispute. A high valuation—akin to high resolve—diminishes the utility loss of suffering wartime costs. Unfortunately for those tasked with making assessments, resolve is typically unobservable *ex ante*. Leaders and advisers employ various techniques to assess opponent resolve before conflict, often with limited success. First, decision makers may look to an adversary's verbal claims. Unfortunately, these claims yield valuable information only under circumscribed conditions given the adversary's incentives to claim a high valuation of the issue (Fearon 1995).⁸

⁸Cheap talk can, under some conditions, credibly convey information (Sartori 2002; Kurizaki 2007; Trager 2010). Alternatively, adversaries may send costly signals to convey resolve, such as sinking costs by building a military base in the region or tying their hands via public pronouncements that invoke the leader's reputation (Fearon 1994, 1997;

Second, leaders may be able to rely on dispositional traits of the adversary's leader (Kertzer 2016). For instance, if we could infer personal attributes of the adversary's leader—e.g., risk tolerance or patience—we could be more precise in the assessment. However, this difficult task is subject to high uncertainty, particularly in the midst of a crisis. Third, an opponent's prior behavior could shed light on future behavior. However, it is unclear whether leaders can (and should) infer an adversary's resolve today based on its past actions (Press 2005; Weisiger and Yarhi-Milo 2015). Contextual differences across space and time may limit the applicability of prior interactions for gauging today's likely responses. Assessing resolve is difficult even after conflict begins because it is unclear what the opponent is willing to endure. Seven years after US combat troops arrived, Chairman of the Joint Chiefs Moorer was still surprised that the "North Vietnamese seem to be able to take unlimited losses."⁹ Resolve is rarely amenable to high-certainty assessments due to a general dearth of information.

(2) *Unity*. Great powers repeatedly attempt to influence which regimes hold power in states throughout the world. Domestic unity is central to assessing the prospects of military efforts toward this end. Widespread domestic opposition to an incumbent increases the prospects for his successful displacement. So too do hostile elite factions. For instance, high-level military defections against Arbenz's regime were critical for victory in CIA's 1954 operation in Guatemala. Tribal dynamics in Afghanistan have proven critical in US efforts to support a functioning central government in Kabul. As with resolve, domestic unity is frequently difficult to observe. Adversaries, particularly autocracies, often mask or inflate their domestic support. Limited information means that domestic unity is a source of high uncertainty.

(3) *Translation*. An opponent's broader political landscape affects post-conflict planning which has ramifications for translating military accomplishments into political gains. Force is an instrument to achieve desired political end-states. Military success must be translated into the objectives, such as territory or preferred policy outcomes, for which the conflict was fought. Theories of war often abstract away this complication and assume combat winners can impose their preferred political settlement. Most conflicts fall short of the ideal type of absolute war with absolute victory. Spoils rarely automatically accrue to the combat victor. Leader beliefs about an adversary's political landscape inform estimates about translating military gains into political gains. For instance, estimates of the US payoff to the Iraq War rested on beliefs about what postwar policies it could install in Baghdad. This in turn depended on assessments of Iraqi military attributes, crucial for *displacing* Saddam, but also Iraqi political attributes such as the post-conflict political landscape,

Yarhi-Milo, Kertzer and Renshon forthcoming). However, states often struggle to provide relevant costly signals that eliminate the problematic source of uncertainty during a crisis.

⁹Minutes of a Washington Special Actions Group Meeting, May 1, 1972, *FRUS*, Volume VIII, Vietnam, 1969-1976, Document 107.

crucial for *replacing* Saddam. Translation difficulties between military and political outcomes are hard to assess beforehand.¹⁰

In sum, leaders must assess military and political attributes of the adversary to develop a complete picture of the expected value of post-conflict settlements. International politics is often a realm of informational scarcity. Such scarcity is especially pronounced for political attributes. Resolve, domestic unity, and translation difficulties between military and political gains are typically harder to observe than military posture and military assets. Accordingly, uncertainty is typically higher for an adversary's political attributes than for its military attributes.

The Information Advisers Provide

Leaders need information on both military and political attributes of the adversary to assess what can be achieved through force and at what cost. Due to time and attention constraints they typically lack the requisite information. Burdens across policy domains mean leaders can rarely be self-reliant during crises. Instead, they turn to their senior advisers. Thus, a leader (principal) delegates information collection and provision tasks to advisers (agents). Compared to the leader, these agents become better, though not necessarily completely, informed about adversary attributes. That is, advisers know more than the leader but are still frequently uncertain about adversary traits because uncertainty permeates much of international politics. Information asymmetries between a principal and agent are fixtures of politics as scholars of Congress and many other domains have long recognized. Yet, these informational dynamics play a limited role in the systematic study of international security with the civil-military literature (Feaver 2003) and Saunders (2017) being some notable exceptions.

In addition to being better informed, advisers are *differentially* informed. These advisers lead or reside within the core bureaucracies of the national security apparatus. A prototypical setup entails a division of labor between bureaucracies. Functional specialization between bureaucracies causes advisers to become differentially informed about adversary attributes. As an abstraction, we can say military advisers from a ministry of defense are relative experts on an adversary's military attributes (e.g., the quality of its capabilities) while foreign policy advisers from a ministry of foreign affairs are relative experts on an adversary's political attributes (e.g., its resolve or domestic political landscape). In their efforts to enhance national security, military advisers study opponents' doctrines, arms development programs, troop levels, force locations, and fighting power. Soviet Defense Minister Malinovsky reportedly stressed to Khrushchev that Cuban forces could

¹⁰Political objectives vary in their translation uncertainties. Installing a new form of government in Baghdad is subject to greater uncertainty than occupying newly acquired territory in the Sinai. See Edelstein (2008) and Sullivan (2007) among others for more on securing political objectives through force.

withstand a US assault for only three or four days during the Cuban Missile Crisis. To provide this crucial input, military advisers accrue information and expertise on an adversary's military characteristics. Foreign policy advisers acquire a distinct information set to perform their assigned tasks. Embassies abroad provide insights into jockeying among domestic factions. Ambassador Sullivan's cable, "Thinking the Unthinkable," casting (belated) doubt on the Shah's long term viability is a famous example.¹¹ Due to their substantive remit, foreign policy advisers develop expertise in an adversary's political characteristics.

An adviser's informational expertise and emphasis follows from the substantive mission of her bureaucracy or institution. Granted, there is expertise overlap between bureaucracies. Foreign policy advisers analyze military information in addition to more purely political information. And the reverse holds as well. Some agencies sit between these ideal-type extremes. The theory's implications have less bite when advisory structures involve duplicative and overlapping responsibilities, which some states may legally require and some leaders may prefer. Franklin Roosevelt's competitive system is a possible example in the US. However, in expectation bureaucracies serve differentiated tasks and develop differentiated expertise as a result.

The process of providing information to the leader augments these expertise divergences. What advisers emphasize when speaking to a leader likely follows from the adviser's assigned task. Limited "face time" with the executive incentivizes advisers to speak to the issues in which they have specialized, potentially unique, knowledge. Beyond an adviser's incentives to stick to his area of expertise, leaders may reprimand those who do not. President Obama reportedly criticized Secretary of State Kerry for bringing him military plans for Syria (Goldberg 2016), which is outside the purview of Kerry's role. Straying beyond core competencies invites blowback and, even worse, can provide leaders with inaccurate or distorted information.

In total, delegation generates informational asymmetries between leaders and their advisers. Comparative specialization among the advisers generates differentiated areas of expertise and emphasis. A testable implication follows:

Informational Theory Expectation 1: When advising leaders, foreign policy advisers discuss political (as opposed to military) attributes of the adversary more often than advisers with different substantive assignments.

Advisers tasked with analyzing and conducting foreign policy are more likely to discuss political, as opposed to military, attributes of the adversary when given an audience with the leader. A relative dearth of information makes political attributes subject to greater uncertainty than mil-

¹¹Telegram from American Embassy Tehran to Secretary of State, 11/29/78, NLC-16-57-3-21-6, Jimmy Carter Library.

itary attributes. It follows that foreign policy advisers typically discuss topics prone to higher uncertainty.

Informational Theory Expectation 2: When advising leaders, foreign policy advisers express greater uncertainty than advisers with different substantive assignments.

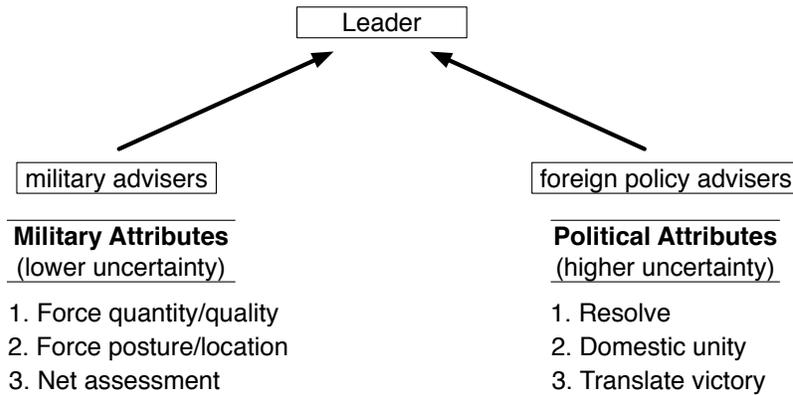


Figure 1: Delegation of information collection and provision tasks to advisers with differentiated domains of expertise. Information availability, and thus uncertainty, vary across domains.

Figure 1 summarizes the theorized adviser delegation and specialization as well as the variation in uncertainty levels for assessing political versus military attributes. To fix ideas, consider this implication in the US context. State Department officials are who I term foreign policy advisers. The first expectation claims they are more prone to discussing political dimensions of an adversary and the second posits that this substantive emphasis, coupled with the inherent difficulty of forming high-certainty assessments on the matter, makes State Department officials express a low degree of certainty. Bureaucracies thus diverge in the information they provide, not the policies they champion.

4 Cases, Evidence, and Measures of Adviser Input

This section describes the research design for testing the empirical implications. Generating measures for empirical analyses proceeds in three steps. The first identifies a set of cases. The second specifies documentary evidence on the advisory processes during these cases. The third uses these documents to generate measures for three dimensions of advisory input: content, certainty, and policy prescription.

Cases

The analysis draws on 61 US crises during the Cold War from the Eisenhower through the Carter administrations. Bounding the analysis to the US during the Cold War has several advantages. To start, it holds potential confounding variables relatively constant. First, keeping the state fixed mitigates concerns that variation in institutional structures drives the results. US institutional arrangements were relatively static as the analysis begins after the 1947 National Security Act created the Defense Department, CIA, and NSC, and ends before the 1986 Goldwater-Nichols Act altered the military chain of command. Second, the background demographics of the individual advisers were remarkably stable. Gender, education, and race, all of which could affect advisory input, were near constants. Third, bipolarity prevailed throughout this period, albeit with stretches of relative detente and animosity.

Practical considerations—such as the availability of evidence—also recommend a US Cold War focus. The analysis requires access to information about advisory processes at the highest levels of decision making. Classification considerations mandate a lag between the events and their study. Pertinent declassified documents exist through the Carter administration. Additionally, the US maintains and provides a uniquely rich documentary record. More broadly, the global ramifications of US behavior make its advisory characteristics of general interest. Nevertheless, the context raises questions about generalizability which I address later.

I identify *opportunities to use force* which represent instances in which the president and senior advisers seriously considered and discussed using US force to strike abroad. The case qualification criteria is closest in spirit to that in [Betts \(1991\)](#), which similarly studies adviser behavior during times when both peace and conflict were realistic options. Relatively strict qualification criteria help preclude smaller incidents that do not precipitate leader-level deliberations and thus are ill-suited for evaluating the hypotheses.

Two sources contribute the majority of cases. First, I include crises from the International Crisis Behavior (ICB) data set if the US was an actor ([Brecher and Wilkenfeld 1998](#)). ICB defines a crisis as a “situation deriving from change in a state’s internal or external environment which gives rise to decision makers’ perceptions of *threat to basic values, finite time for response, and the likelihood of involvement in military hostilities* [emphasis in original]” ([Brecher and Wilkenfeld 1982](#), p. 383). Second, I include major military mobilizations that [Blechman and Kaplan \(1978\)](#) (p. 50) code as level one and two uses of force. The authors define these as:

- *Level 1*: Use of strategic nuclear unit plus at least one “major” force component;
- *Level 2*: Two or three “major” force components used, but not strategic nuclear units,

where a “major” unit is (1) two or more aircraft carrier task groups, (2) more than one army battalion, or (3) one or more combat wings. Directing many military assets toward a possible

conflict location indicates serious discussions about using force. Importantly, asset mobilization does not imply the actual use of force. Some mobilizations ended peacefully while others did not. Together, these two sources contribute 46 cases. I supplement the sample with 15 additional observations with the key qualification criteria being the presence of leader-level deliberations about using force. These include some lower-level military mobilizations, militarized interstate disputes (not all of which included the use of force) (Ghosn, Palmer and Bremer 2004), and debates concerning preventive strikes against the Soviet Union and China's nascent nuclear programs.

The Online Appendix lists the opportunities to use force (which I use interchangeably with 'crises') in the sample and provides information about each, such as its qualification criteria and the precipitating context. Some cases concern potential new uses of force—such as the 1958 turbulence in Lebanon, 1969 North Korean EC-121 spy plane incident, and 1975 Cambodian *Mayaguez* seizure. Others capture pivotal moments to potentially expand the scope of ongoing uses of force—such as 1953 deliberations surrounding the Korean War armistice and 1971 invasion of Laos during the Vietnam War. Sample censorship, in the sense of excluding cases where using force never received leader-level scrutiny, could introduce bias favoring the hypotheses if the relationship of interest is even more pronounced under conditions favoring sample inclusion. However, this inferential threat is likely limited. Advisers' *relative* areas of expertise and emphasis remain static regardless of presidential scrutiny.

Evidence: Constructing a Corpus

Measuring advisory input requires access to details about crisis advisory processes. While no existing data set systematically provides this information across many cases, rich qualitative case studies do an admirable job reconstructing advisory processes.¹² Inspired by these examples, I use a new approach to directly capture the private words of senior advisers at scale. Specifically, I collect documentary evidence from deliberations at the presidential level about using force. Elites are the theoretical locus of analysis and ought to be for the empirical analysis as well. Declassified documents published through the Foreign Relations of the United States (FRUS) series, the official documentary record of the State Department, constitute the core sources for the analysis. FRUS volumes supplement State Department files with materials from presidential libraries, the Department of Defense, National Security Council, Central Intelligence Agency, and other agencies and individuals involved in crafting US foreign policy. Documents include National Security Council meeting minutes and transcripts, intra-elite memos (such as from Kissinger to Nixon), and minutes from elite conversations in unstructured (such as ad hoc White House meetings) and structured

¹²Among many others, some fine examples include George and Smoke (1974); Van Evera (1984); Betts (1991); Marsh (2014); Saunders (2017).

settings (such as presidential briefings from the Joint Chiefs of Staff).¹³ FRUS provides access to *private* statements that likely reflect sincere beliefs. For instance, when opting to intervene in the Laotian Civil War in 1964, Johnson told advisers “we should go ahead with the mission but that he had doubts about the action.”¹⁴ Presidents rarely voice such reservations publicly.

I select and prune documents to include only those portions directly pertinent to the theory. Retained portions concern political or military attributes relevant for assessing potential conflict outcomes. I restrict the corpus to memos sent to or by the president and transcripts from meetings which the president attended or was briefed on. Given his primacy in the decision-making process, I focus on only those advisory opinions to which he was privy.¹⁵ All transcripts are split into speaker-specific texts; thus, a Laird-Kissinger dialogue produces a Laird text and a Kissinger text.

In total, the corpus includes over 287,000 total words in 1,138 speaker texts generated from 382 distinct documents with 176 unique speakers. 85% of speaker texts come from meeting transcripts with the remainder from memos. As shown in the Online Appendix, the extent of adviser input within the corpus accords with the historical record. For instance, text volume in the corpus mirrors the well-known prominence of key advisers, such as John Foster Dulles, Kissinger, and Brzezinski.

Two potential drawbacks to using FRUS—classification and selective document inclusion—merit attention. Redacted segments within available texts pose little inferential risk because they typically concern targeting location details or covert sources rather than broad strategic considerations. Another concern is that entire documents remain classified or that FRUS historians excluded certain types of documents. Assessing either possibility is difficult at scale. A highly publicized dispute over declassification of materials for the FRUS volume on the 1954 Guatemala operation alleviates some concerns.¹⁶ It is reassuring that this one case is well known and researcher complaints generally have not reached a similar volume since. More generally, to bias results in favor of the hypotheses, excluded documents must consistently involve State Department officials speaking on military (vs. political) matters and/or expressing low uncertainty.

¹³Because FRUS volumes for the Carter-era are incomplete, I gathered documents from the Carter Presidential Library and online materials from the National Security Archive. Documents akin to those in FRUS were widely available with partial exceptions for the fortnight preceding the April 1980 Iran Hostage rescue attempt and a general deficit of JCS materials. I also supplement FRUS with transcripts from the Executive Committee of the National Security Council (ExComm) from the Cuban Missile Crisis, which provide a unique resource worth exploiting.

¹⁴Memorandum of Conference with President Johnson, June 8, 1964, *FRUS*, Volume XXVIII, Laos, 1964-1968, Document 83.

¹⁵The 1976 North Korean Tree Trimming Incident meets the sampling criteria but is excluded because FRUS does not contain documents involving the president.

¹⁶See [McAllister et al. \(2015\)](#) for a description of the controversy and the politics of FRUS.

Measuring Advisory Input

Raw inputs—memos and transcripts—must be transformed into analytically useful measures. Text analysis tools, which provide a replicable means of distilling advisory input from a mass of private statements, enable this transformation for measuring content and certainty. However, these tools are not well suited for measuring advisory policy prescriptions. Accordingly, I qualitatively code the relative hawkishness of State and Defense Department officials for each case. The following sections describe the measurement strategies for the three measures.

I generate measures at the bureaucracy-crisis level by aggregating the input of advisers affiliated with five bureaucracies: State, Defense, the Joint Chiefs of Staff (JCS), Central Intelligence Agency (CIA), and National Security Council (NSC) or White House staff.¹⁷ State Department officials are the foreign policy advisers who specialize in an adversary’s political attributes. Defense Department officials and members of the Joint Chiefs of Staff, who I distinguish between for reasons that will become clear, are specialists in the adversary’s military attributes. These officials represent the ‘ideal’ types. CIA and the National Security Council staff (plus other White House advisers) lie along a continuum stretching between them and consequently produce ambiguous empirical expectations.¹⁸

Adviser Content: Political or Military

Testing the informational theory of bureaucracies requires a measure of whether advisory content concerned military or political considerations. Supervised learning techniques, which classify texts into pre-specified categories, offer an accurate and scalable method for this measurement task (Hastie, Tibshirani and Friedman 2001). Supervised methods follow a straightforward process. Human coders classify a subset of texts as either political or military. An algorithm uses this *training set* to “learn” which words of texts are associated with each category. The algorithm then uses these features to sort the remaining texts between the categories.

I use Naive Bayes, one of the simplest supervised learning methods (Maron and Kuhns 1960). The Online Appendix describes the method in greater detail—following the presentations in Grimmer and Stewart (2013) and Nielsen (N.d.). In brief, we wish to estimate the probability that text T belongs to the military class M , given the words in T . By Bayes’ Rule, this desired quantity $P(M | T)$ is equal to $\frac{P(T|M)P(M)}{P(T)}$. Let $P(T | M)$ be the independent product over all $P(w_i | M)$ where w_i represents each word in text T . The independence assumption is clearly

¹⁷The final category includes vice presidents. Treasury officials do feature in some deliberations, but assumed a far more limited role than those from the five listed bureaucracies during the cases under scrutiny.

¹⁸On NSC responsibilities and their expansion over time, see Preston (2001); Destler (1980); George and Stern (2002).

wrong; words used in a text are highly correlated. Nonetheless, Naive Bayes provides accurate classifications despite the assumption violation. I use the observed word frequencies in the training set to estimate $P(w_i | M)$. With substitution, we can rewrite the original equation as $P(M | T) = \frac{P(M)}{P(T)} \prod_i P(w_i | M)$. Because texts are either military or not (in which case they are political), we can similarly write $P(M' | T) = \frac{P(M')}{P(T)} \prod_i P(w_i | M')$. Following Nielsen (N.d.), these two quantities generate a log likelihood ratio which serves as the raw text score:

$$\text{Raw Text Score} = \sum_i \log \frac{P(w_i | M')^{w_i}}{P(w_i | M)^{w_i}}.$$

Positive scores indicate political texts (M') and negative scores indicate military texts (M). For the analyses, I generate a continuous *Content Text Score* which divides the raw score by the number of words in text T . Variation in text lengths drives scoring discrepancies, with longer texts producing more extreme scores. Standardizing for length adjusts for this correlation, though results are similar regardless of the measure used. I also generate a binary variable, *Political*, which codes texts as 1 if the *Content Text Score* is positive, and 0 otherwise.

Training Set. Speaker texts in the full corpus range from a single word spoken during a meeting to memos exceeding 4,000 words. Variation poses a coding challenge as short texts render it difficult to discern the expressed content while long texts exceed coders' realistic attention spans. To provide a degree of uniformity, I split longer texts into portions containing between 30 and 250 words. This length proved sufficient to judge content yet remain manageable. The training set consists of a random subset of memos and transcripts. In total, it has 475 texts containing 58,945 words, which is approximately 20% of the corpus' total word count. To guard against coding bias, I remove identifying information on the speaker's name and bureaucratic affiliation and then classify each text as either political or military. Proper nouns, such as "Castro," remain in the texts for readability. Military and political texts constitute 61% and 39% of the training set, respectively.

Classification rules followed from the theory and were refined through iteratively applying these rules to the actual texts. Military texts address military capabilities, the quality of forces, force locations, and how the interaction of opposing forces might unfold. In practice, texts frequently referred to opponent, allied, and the United States' own military forces and options. Political texts address the opponent's commitment to the issues at stake (resolve), the domestic political landscape in the relevant state, the popularity of that state's leader, or challenges involved with translating the application of force to desired political end-states. Some documents refer to US resolve, as well as the reputational effects for the US of abstaining from versus entering in to conflict. Example texts in the Online Appendix typify content from each category.

Validation. Two validity checks provide strong evidence that the Naive Bayes classifier captures the substantive content of interest. The first check is descriptive and suggestive, the second

more systematic and definitive. The first examines the words that most distinguish between the categories. This offers a sanity check that military documents do indeed pertain to military factors and the same is true for political documents. Figure 2 plots the 25 terms with the highest *relative* frequency for each category—that is, those with the greatest differentials between the categories.

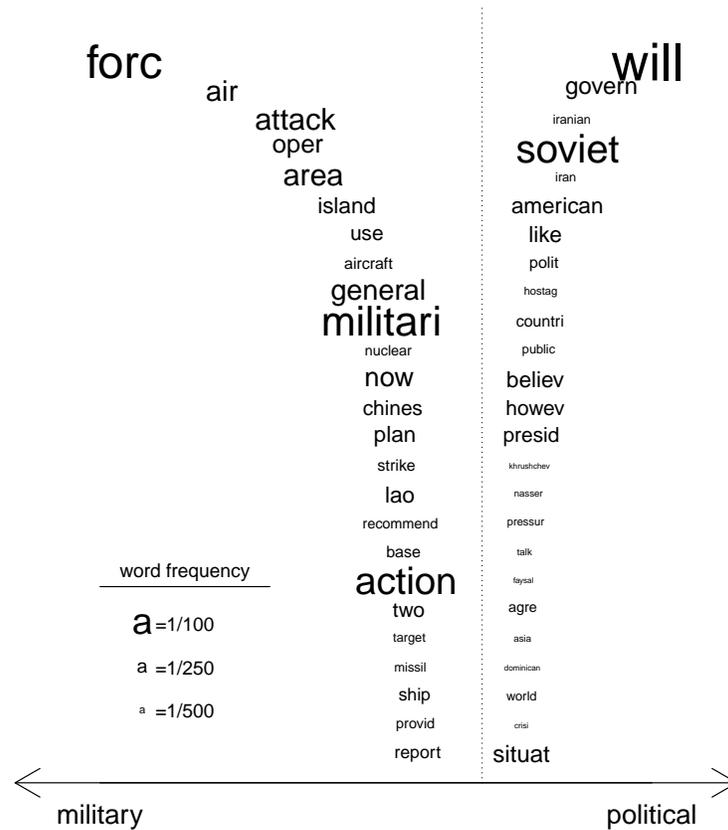


Figure 2: Terms with greatest frequency differences between military and political documents. Words to left are relatively more frequent in military documents; words to right are relatively more frequent in political documents. Size of text indicates overall term frequency in all training documents.

Figure 2 provides surface plausibility for the measurement strategy. The left side of the plot shows terms, or word stems, that most distinguish military texts. The stems indicate a high degree of measurement validity: “forc,” “attack,” “general,” “nuclear,” “target,” and “strike.” Word stems to the right are those most indicative of political texts. Given that the political category spans a wider range of substantive topics, words tend to be less discriminating as compared to the military category. Instead, state leaders—“Khrushchev,” “Nasser,” “Faysal”—help distinguish these texts, as do references to “govern” and “talk.” The most telling term for political texts, “will,” has two

connotations. The first relates to projections of how actors will respond and the second to their willpower or resolve. Each connotation concerns political, as opposed to military, content.

Cross-validation, the second validity check, offers systematic evidence supporting the measure’s appropriateness. I perform 10-fold cross-validation which randomly partitions the training set into 10 equally sized groups. The exercise sequentially drops each group, trains the model using the remaining nine groups, and then applies the classifier to the withheld group. This produces out-of-sample predictions for each text in the training set and avoids overfitting the classifier.

The classifier performs well on several metrics. It accurately categorizes 88% of the texts. Other diagnostics for evaluating classification performance are similarly supportive. F -scores extend beyond accuracy, incorporating other metrics such as precision and recall. For classifying military texts, precision and recall respectively refer to the fraction of documents classified as military that were hand coded as military and the fraction of the hand coded military texts that were classified as military texts. The F -score is 0.90. Finally, the cross-validation exercise generates the previously discussed *Raw Text Scores* for each document. Figure 3 plots these scores on the x-axis and hand coded classifications on the y-axis. Beyond high accuracy (solid points as a fraction of all points), the figure shows that inaccurate classifications were rarely far off the mark, clustering around the dashed vertical line at 0. The solid line plots fitted values from a bivariate logistic regression. Its steepness further attests to the classifier’s ability to discriminate between military and political texts.

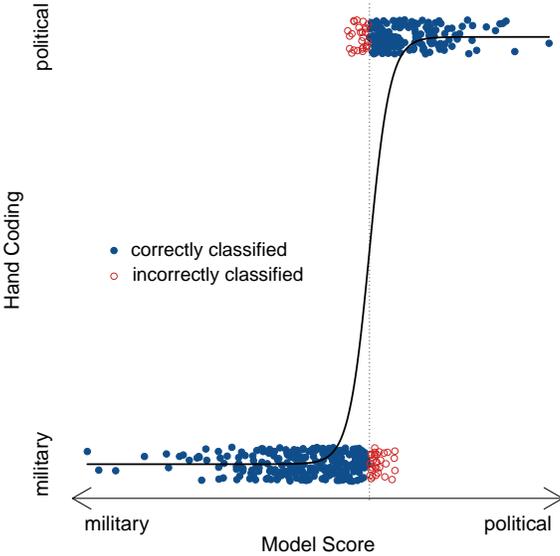


Figure 3: Cross-validation results. Dashed vertical line indicates *Raw Text Score*= 0.

Application. After validating the approach on the training set, I apply it to the full corpus

at the bureaucracy level of analysis. I aggregate all text from a bureaucracy’s officials during a single crisis—for instance, from all Defense Department officials during the *Mayaguez* seizure—and apply the classifier to this aggregated block of text. Each bureaucracy-crisis observation gets a continuous *Content Text Score* and dichotomous *Political* content indicator.

Adviser Uncertainty

The analysis requires a measure of assessment certainty in addition to the measure of assessment content. Dictionary methods offer a simple means for summarizing text attributes. I use the “If” dictionary from the Harvard General Inquirer, which linguists constructed to measure uncertain tones in speech and writing. It includes words such as “approximately,” “doubt,” “possibility,” and “unpredictable.”¹⁹ Two reasons argue in favor a dictionary method over a supervised learning method for the measurement task. First, there is a pre-existing dictionary of appropriate words that connote uncertainty across contextual domains. Second, the classification concept is far subtler compared to the military/political distinction. Words delineating between military and political content are frequent and stark, as seen in Figure 2. Words distinguishing certain from uncertain assessments are less frequent and obvious, making it harder for a learning method to discern distinguishing features. The 96 words in the uncertainty dictionary constitute just over 5% of words in the corpus. In comparison, the 96 words that most distinguish military texts constitute nearly 22% of the words in the corpus training set. The relative scarcity of uncertainty words means that other elements of the texts likely swamp the uncertainty features, thus limiting classification accuracy.

Following prior work summarized in Grimmer and Stewart (2013), I generate an uncertainty measure by assigning a score of one to words in the dictionary, sum across a text, and divide by the total number of words in the text (see McManus (forthcoming) for a related approach). The Online Appendix includes an example text that highlights words in the uncertainty dictionary. Overall, speaker texts with at least 30 words have a mean uncertainty score of 5.3% and standard deviation of 3.0%. To produce measures at the bureaucracy-crisis level, I aggregate texts from all speakers within a specified bureaucracy during a given case and apply the dictionary to these aggregated texts.

Dictionary methods are difficult to validate because they lack an analogue to k -fold cross-validation (Grimmer and Stewart 2013). I adopt two approaches, discussed further in the Online Appendix. The first hand codes a subset of all documents (roughly 5%) from the corpus into three levels of uncertainty. Human coding strongly correlates (53%) with the dictionary method uncertainty scores. The second approach leverages texts where there are clear expectations for the *relative* degree of uncertainty. For instance, during the EC-121 incident Defense Secretary Laird

¹⁹Dictionary is available at <http://www.wjh.harvard.edu/~inquirer/>. I add “risk” to the dictionary because decision makers often voice uncertainty by invoking risks.

wrote to Nixon suggesting that the JCS was overly certain in its estimate of what outcome US forces could impose on North Korea. He wrote, “If U.S. losses occur in the strike (and I believe there is more chance they may than the JCS papers indicate). . .”²⁰ Laird’s memo should, and does, register a higher uncertainty score than the JCS communications (6.2% vs. 2.6%).

Adviser Policy Prescriptions: Hawks and Doves

I take a qualitative approach, measuring *relative* hawkishness on a three-point scale, to evaluate whether bureaucratic affiliations drive advisers’ policy prescriptions. I restrict the focus to *civilian* advisers from the State and Defense Departments. This offers a fair test of the theoretical contention. State and Defense represent the extreme bureaucracies in terms of parochial benefits to diplomacy versus force. CIA and NSC have mixed interests that produce ambiguous implications. Additionally, a strictly civilian approach sidesteps the aforementioned complicating issues from civil-military relations, which scholars have already studied (Betts 1991; Feaver and Gelpi 2004).

I qualitatively code relative aggressiveness for each opportunity to use force based on original documents such as those from FRUS, the secondary literature, and memoirs of participants when necessary. What constitutes the hawkish position varies by crisis. Some debates concern using conventional versus nuclear forces. Others concern the size of arms shipments to send to an ally. The concept is not amenable to automated text measurement because subsets of words that represent the aggressive policy in one case may represent the dovish posture in another. For each case, I evaluate the positions of the Secretaries of State and Defense as well as senior officials in each bureaucracy. Following Betts, the measure takes one of three values. Either Defense is more hawkish, there is no substantive difference, or State is more hawkish. The Online Appendix describes all sources and rationales so readers can evaluate the coding choices.

5 Results: Adviser Content, Uncertainty, and Prescriptions

The quantitative findings presented in this section strongly support an informational theory of bureaucracies, rather than a parochial theory. The four main variables in the analyses capture (1) bureaucracy via a binary indicator for *State Department*; (2) advisory content measured with the continuous *Content Text Score* and binary *Political* score; (3) advisory *Uncertainty* based on the dictionary approach; and (4) advisory *Aggressiveness* using the three-level coding comparing the policy preferences of Defense and State officials.

Ordinary least squares (OLS) regression is the primary method of analysis with standard errors

²⁰Memorandum from Secretary of Defense Laird to President Nixon, April 18, 1969, *FRUS*, Volume XIX, Part 1, Korea, 1969-1972, Document 17.

clustered on the case. I present specifications without and with case fixed effects, which account for unobserved invariant components of each case.²¹ Models with fixed effects show how bureaucratic affiliation affects advisory input when holding the case fixed (say, to the Cuban Missile Crisis). Each bureaucracy thus assesses the same adversary at the same time. Differences between State and other bureaucracies are not attributable to underlying differences in the cases they participate in. All OLS models are bivariate because any case-level control variables are definitionally collinear with the fixed effects. To ensure that text-based measures reflect substance rather than noise, I limit the sample to the 167 bureaucracy-crisis observations with at least 100 words, though results are robust to using alternative cutoffs.²²

State Department Officials Provide Political Information

The results presentation begins with Informational Theory Expectation 1, which contends that State Department officials are more likely to emphasize political attributes compared to counterparts from other bureaucracies. Non-overlapping functional responsibilities between the bureaucracies encourages comparative specialization. Tasked with conducting foreign policy, State Department officials specialize in an adversary's political characteristics.

I examine descriptive patterns before turning to regression results. Overall, advisory input skews toward military rather than political content with fewer than 30% of observations coded as political on the binary measure. This is unsurprising given that cases only qualify for inclusion if the president and senior advisers discuss the potential use of force. However, there is important heterogeneity among bureaucracies. State Department officials are twice as likely (46% vs. 23%) to discuss political content compared to officials from other bureaucracies.

Figure 4, further decomposes the data, plotting summary statistics for each of the five bureaucracies. The left and right panels respectively show the mean *Content Text Score* and *Political* indicator. Beyond confirming State's tendency to emphasize political content, the figures confirm another expectation: advisers in the Defense Department and JCS tend to emphasize military content. On the binary measure, members of the JCS strictly limit themselves to military content. CIA and NSC/White House officials occupy a middle ground.

Descriptive plots show that functional roles affect the type of information advisers provide to the president. Regressions generally confirm this conclusion. Table 1 reports results from

²¹To calculate clustered standard errors with fixed effects, I follow the advice in Cameron and Miller (2015). Each cluster in the data set has a small number of observations (maximum of five). For small clusters, Cameron and Miller (2015, p. 331) recommend the 'within' fixed effects estimator, rather than the dummy variable fixed effects estimator.

²²There would be $5 \cdot 61 = 305$ observations if the corpus contained sufficient text for each bureaucracy during each crisis. Due to either limitations of the corpus or marginalization of key advisers, many potential observations do not exist.

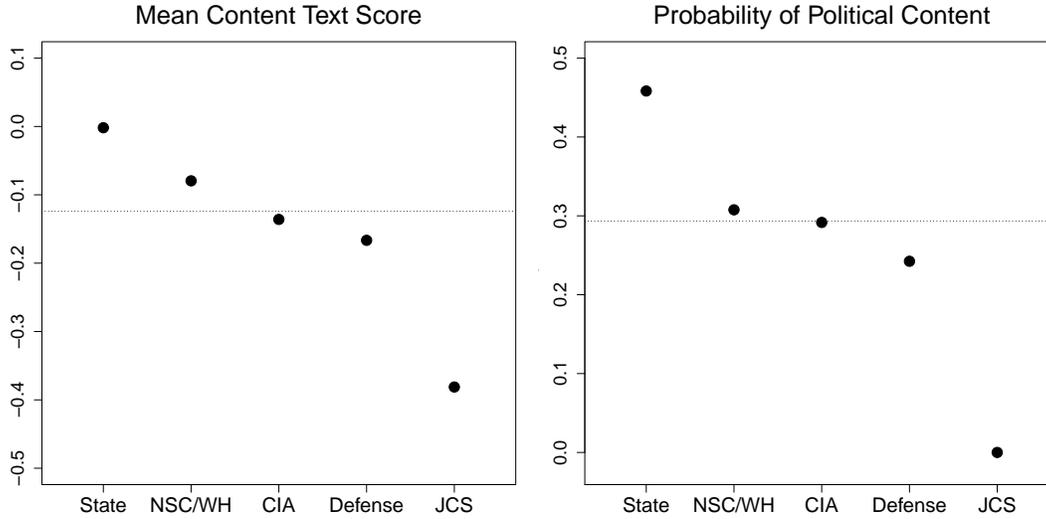


Figure 4: Text content: All Bureaucracies. NSC/WH indicates staff from the National Security Council or White House. *Left*: Mean *Content Text Score* with negative values indicating military content and positive values indicating political content. *Right*: Percentage of observations coded as political.

specifications with content as the outcome variable and bureaucratic affiliation as the explanatory variable. Models vary the outcome variable coding (continuous in 1-3, binary in 4-6), the inclusion of case fixed effects, and sampling criteria. Logistic regression for the binary outcome coding provides similar results. Models 3 and 6 prune the sample to only the ‘ideal’ type bureaucracies—State, Defense, and JCS. If the theory is correct that specialization dictates advisory content, then divergences should be especially pronounced within this subset. Again, models do not include control variables because variables measured at the case level (e.g., geographic proximity) are collinear with the fixed effects.

The *State Department* explanatory variable has a positive coefficient, indicating more political content, across all specifications. This effect is statistically significant in all three models using the continuous content score. Whereas the average bureaucracy input concerns military content (*Content Text Score* = -0.12), State Department officials stress, on average, either neutral or political content (coefficients ≥ 0.11). This effect becomes more pronounced when restricting the comparison to Defense and JCS officials, as reported in Model 3. Consider the crisis sparked by North Korea shooting down an EC-121 reconnaissance plane. On average, bureaucracies had a *Content Text Score* of -0.15 , which is just below the median for the full sample. The model predicts that State’s score would be -0.04 , or substantially more balanced between military and political content. State’s actual input was even more skewed toward political matters with a score of 0.24 . These patterns in the data accord with the historical record. Defense officials focused on the advantages of

Table 1: Bureaucratic Role and Advisory Content

	<i>Content Text Score</i>			<i>Political</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
State Department	0.17*** (0.03)	0.11*** (0.03)	0.20*** (0.03)	0.23*** (0.07)	0.09 (0.07)	0.22*** (0.08)
Constant	-0.17*** (0.03)	-0.16*** (0.01)	-0.23*** (0.01)	0.23*** (0.05)	0.27*** (0.02)	0.18*** (0.04)
Observations	167	167	104	167	167	104
Case Fixed Effects	N	Y	Y	N	Y	Y
Only ‘Ideal’ Types	N	N	Y	N	N	Y
Outcome Mean	-0.12	-0.12	-0.14	0.29	0.29	0.29

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS with bureaucracy observations as unit of analysis. Standard errors clustered on the case. Fixed effects not shown for models that include them. Models 3 and 6 limit the sample to State, Defense, and JCS observations.

possible airstrikes compared to mining North Korean harbors or establishing a blockade. In contrast, State officials considered Soviet reactions and likely North Korean responses to diplomatic overtures.

Using the binary outcome coding produces similar, though less robust results. In Model 4, moving from any other bureaucracy to the State Department doubles the probability of political content (23% points \pm 14% at the 95% confidence interval). When including fixed effects, the substantive effect diminishes and is no longer statistically significant. However, when restricting the sample to the ideal types, I again find a strong and statistically significant effect. Further analysis confirms that while State Department provides the most political content of all bureaucracies, the divergence from CIA and NSC input is relatively small, at least when using the binary outcome coding and fixed effects. The result further suggests that the latter two bureaucracies lie in the middle of the continuum of substantive assignments and can encroach upon or, more optimistically, contribute to the State Department’s substantive remit. Robustness checks in the Online Appendix address small sample size concerns using the diagnostic tests outlined in [Gelman and Carlin \(2014\)](#).

State Department Officials Express More Uncertainty

Informational Theory Expectation 2 anticipates that State Department officials express greater uncertainty than their counterparts during advisory processes. Officials from State are more likely to discuss an adversary’s political attributes and these attributes tend to be more uncertain due to

limited information.

Consistent with the expectation, *Uncertainty* is 10% higher in State Department observations compared to other agencies (5.6% vs. 5.1%). Decomposing the sample by bureaucracy and content yields additional insights, shown in Figure 5. First, State is the most uncertain of all five bureaucracies when averaging across the full sample (left panel of figure). JCS observations offer the starkest contrast with an average score below 4.5%. Second, content domain affects bureaucracies' uncertainty (center and right panels). State's uncertainty is especially pronounced in the realm of political content, approaching 6%. In contrast, State's uncertainty over military attributes is comparable to several agencies. Troublingly, those outside of State express a relatively high degree of certainty when they wade into political matters, at least compared to the specialists inside State. Political assessments from advisers speaking beyond their core competency may exhibit unwarranted certainty. Third, competency widens the political-military uncertainty gap. When the ideal type bureaucracies—State, Defense, and JCS—discuss their areas of expertise, they provide dramatically different degrees of uncertainty as depicted in the shaded regions.

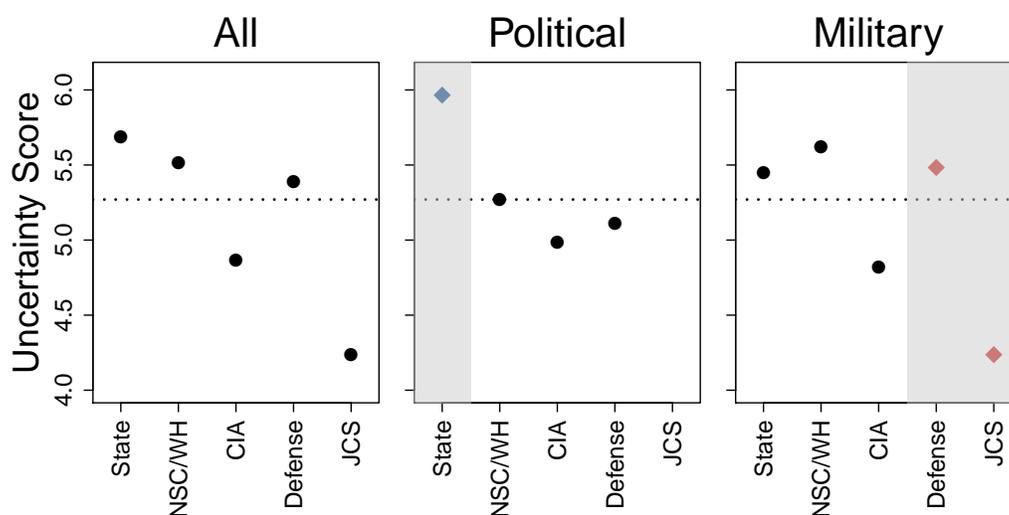


Figure 5: Mean uncertainty score by bureaucracy. Dashed horizontal line shows overall mean uncertainty. Diamonds in shaded regions indicate ideal type bureaucracies speaking in their area of expertise. Left, right, and center panels respectively plot all observations, observations discussing political content, and observations discussing military content.

Models 1 and 2 in Table 2 regress *Uncertainty* on the *State Department* indicator without and with case fixed effects. Fixed effects guard against the possibility that State officials are only participating in information-poor cases. Both specifications produce similar results: State officials express approximately 10% ($\pm 10\%$) more uncertainty relative to other advisers. Deliberations before the 1965 intervention in the Dominican Republic help illuminate overall patterns. The average

bureaucracy *Uncertainty* score was 6.0%. Regression results predict a State Department score of 6.6%. In reality, State officials expressed even more uncertainty, producing a score of 7.3%, or a 20% relative increase over other bureaucracies. Documents from the crisis corroborate this pattern. Defense officials expressed limited doubt that two or fewer US divisions could stabilize the island. In contrast, State Department officials, attuned to political factions on the ground, conveyed a lack of certainty. Undersecretary of State Mann told Johnson that the “loyalties of the troops outside the capital are still uncertain.”²³

Table 2: Bureaucratic Role and Certainty

	<i>Uncertainty</i>			
	(1)	(2)	(3)	(4)
State Department	0.58*	0.61*	0.96**	1.33**
	(0.29)	(0.31)	(0.39)	(0.60)
Constant	5.10***	5.10***	4.83***	4.81***
	(0.16)	(0.09)	(0.18)	(0.19)
Observations	167	167	104	70
Case Fixed Effects	N	Y	Y	Y
Only ‘Ideal’ Types	N	N	Y	Y
Only ‘Expert’ Types	N	N	N	Y
Outcome Mean	5.27	5.27	5.27	5.23

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS with bureaucracy observations as unit of analysis. Standard errors clustered on the case. Fixed effects not shown for models that include them. Models 3 and 4 limit the sample to ‘ideal’ types—State, Defense, and JCS observations. Model 4 further restricts the sample to ‘expert’ observations where an ‘ideal’ type discusses its area of expertise.

State’s relative uncertainty becomes even more pronounced when subsetting the sample to ideal-type bureaucracies. Model 3 presents results where Defense and JCS serve as the comparison set. Substantively, advisory input from State conveys 19% ($\pm 15\%$) more uncertainty than input from military advisers. Limiting the sample to ‘expert’ observations—i.e., State discusses political and Defense/JCS discuss military attributes—produces even larger substantive effects (27% $\pm 25\%$). See the Online Appendix for robustness checks assessing the implications of the relatively small sample size for effect magnitudes (Gelman and Carlin 2014).

²³Telephone Conversation Between the Under Secretary of State for Economic Affairs (Mann) and President Johnson, April 26, 1965, *FRUS*, Volume XXXII, Dominican Republic; Cuba; Haiti; Guyana, Document 22.

An alternative explanation for the results might suggest that individuals who select into the State Department systematically differ from others. Evidence suggests that they are more liberal (Milner and Tingley 2015) and more likely to be generalized trusters (Rathbun 2012). This could drive the results if either characteristic makes individuals more prone to express uncertainty. Given the impossibility of randomly assigning senior advisers to different bureaucracies, I cannot definitively rule out this alternative. That said, it is telling that State Department officials are notably uncertain when assessing political attributes but not notably uncertain when assessing military attributes. It appears that recognition of informational scarcity on political traits, as opposed to something specific about pre-existing individual dispositions, drives the elevated uncertainty.

Bureaucratic role affects advisory uncertainty as the theory anticipates. When writing or speaking to the president, State Department officials express greater uncertainty than their peers. Limiting the comparison to officials at Defense and the JCS magnifies this effect, as does restricting the analysis to observations within a bureaucracy's core competence.

Searching for Dovish Diplomats

The analysis now evaluates Parochial Theory Expectation 1. Are State Department officials less militarily aggressive than their civilian counterparts at Defense as the canonical parochial account suggests? For Allison and Zelikow (1999, p. 307), “[k]nowledge of the organizational seat at the table yields significant clues about a likely stand.” In crises, however, I expect parochial interests dissipate. Presidents wield tools—selection, sanction, and removal of advisers—that further encourage preference alignment.

Consistent with my expectation and contrary to Allison, ‘organizational seat’ offers little insight into the policy positions that officials adopt. Table 3 shows the overall pattern. State and Defense advocate substantively indistinguishable positions in nearly half of the cases. In those cases with a difference, Defense and State are almost equally likely to assume the hawkish position. Civilians in the Office of the Secretary of Defense call for more aggressive actions in 30% of cases. Those at State are the more aggressive actor in 25% of cases. A variety of statistical tests confirm what the table makes immediate: State Department officials are not especially dovish. Whether treating the outcome codings as linear, ordinal, or categorical, all tests fail to reject the null hypothesis of no difference between the bureaucracies.

To an extent, aggressiveness depended upon who headed the bureaucracy. Cyrus Vance never adopted the most aggressive posture while in the Carter administration. His consistency is atypical though. Other secretaries varied depending on the context. Rogers, for instance, opposed bombing Cambodia in 1969 while Secretary of Defense Laird advocated, or at least approved, it.²⁴ Rogers

²⁴Editorial Note, *FRUS*, Volume VI, Vietnam, January 1969-July 1970, Document 41.

Table 3: Counts of Relative Aggressiveness by Bureaucracy

	<i>More Aggressive Bureaucracy</i>		
	Defense	Similar	State
Overall Count	18	28	15
<i>One sample tests: p-values</i>			
T-test	0.61		
Wilcoxon	0.61		
Chi-squared	0.70		

Notes: No test precludes the null hypothesis of no difference in hawkishness between the bureaucracies. T-test codes outcomes as 0, 1, and 2. “Similar” serves as median value for Wilcoxon test. The Chi-squared test assigns probabilities across the three categories of 25, 50, and 25.

and Laird flipped positions three years later. Laird strongly opposed using B-52s to bomb industrial complexes in Hanoi and Haiphong in late 1972²⁵ whereas Rogers consented to the *Linebacker II* bombing, albeit hesitantly. Little daylight existed between the State and Defense positions in the modal case. Sometimes this followed naturally from close State-Defense collaboration, as occurred during the Berlin Wall crisis. Other times, they advocated different but not discernibly more aggressive policies. After much wavering, John Foster Dulles abandoned diplomatic efforts at the UN and suggested US forces intervene in Lebanon in 1958 (George and Smoke 1974; Little 1996).²⁶ Donald Quarles, the Deputy Secretary of Defense, similarly preferred that the UN play a role because he wanted the moral sanction that comes with UN authorization. He ultimately acquiesced to the intervention despite no UN authorization.²⁷ Reasonable people could disagree, but I view their respective positions as similarly aggressive.

6 Discussion, Generalizability, and Implications

The results indicate that a foreign policy adviser—State Department official—can simultaneously emphasize political attributes, express uncertainty, and, on occasion, endorse hawkish poli-

²⁵Memorandum From Secretary of Defense Laird to President Nixon, December 12, 1972, *FRUS*, Volume IX, Vietnam, October 1972-January 1973, Document 166.

²⁶Though see Hoopes (1974) on the sincerity of Dulles’ diplomacy.

²⁷Memorandum for the Record of a Meeting, July 14, 1958, *FRUS*, Volume XI, Lebanon and Jordan, Document 123 and Memorandum of a Conference With the President, July 14, 1958, *FRUS*, Volume XI, Lebanon and Jordan, Document 124.

cies. It is not contradictory that they do so because each component is distinct. For instance, knowledge of political attributes encourages hawkish positions when the available information indicates the opponent is unresolved or easily displaced from power. Placing more uncertainty around the estimate may temper, but not reverse, the policy posture one adopts. Secretary Rusk was attuned to political dimensions while simultaneously advocating hawkish positions. Rusk was not a Vietnam dove. From 1961 onwards he contributed to escalatory decisions. Nonetheless, he interrogated the strength, viability, and resolve of leaders in Saigon. For both Diem and Khanh, Rusk questioned whether they were worth supporting and how to use military leverage to extract governance concessions from them (Pelz 1981; Freedman 1996). Months before the Tonkin Gulf incidents, Rusk said that “we do not go to war for a government which is more interested in quarreling than in fighting the Viet Cong.”²⁸ The viability of a supported regime calls for the State Department’s political judgment. Rusk answered the call yet still saw fit to increase US involvement in Indochina. Cognizance of political considerations does not imply an aversion to using force.

Generalizability

How generalizable are the findings? A systematic evidence-based answer is beyond the scope of this paper. We can though specify theoretical conditions needed for the implications to hold in other states and time periods. Five elements merit attention. First, a leader’s ability to select and sanction advisers facilitates preference convergence. If leaders do not enjoy this authority, as in parliamentary coalition governments, then bureaucratic interests could inform policy prescriptions. Second, leaders must incentivize advisers to perform their responsibilities. The possibilities of poor monitoring (Saunders 2017) and excessive sanctioning, such as killing dissenting advisers, will curtail the free flow of information during advisory processes. If a foreign minister believes a leader does not want to hear about political obstacles during a crisis, then that foreign minister might keep that information to herself.

Third, leaders must opt to be informed rather than ignorant. If leaders do not demand information than advisers have limited means to supply it. Differences in advisory input becomes moot.

Fourth, multiple bureaucracies must enjoy access to the leader. If a single institution dominates the information provision process, then affiliated advisers may not stick to their areas of core competence. Regimes with civil-military balances that vastly privilege the military might exhibit this tendency. Fifth, there must be a division of labor among bureaucracies. The divide in advisory content (political vs. military) likely shrinks as substantive overlap increases among bureaucracies.

²⁸Summary Record of the National Security Council Executive Committee Meeting, May 24, 1964, *FRUS*, Volume I, Vietnam, 1964-1968, Document 172.

Leaders face a tradeoff. Assigning duplicative responsibilities to advisers might prove inefficient, but it also provides them with multiple sources for each type of information. The dramatic expansion of the US NSC over the past four decades generates this kind of substantive overlap. NSC staff numbers have exploded from the mid-20s under Carter to an estimated size of over 500 under Obama (Ries 2016, p. 32-35). This expansion offers presidents an alternative information source that often wades into the State Department's traditional territory. Consequently, State's relative emphasis may not stand out as much in recent administrations as it did during the Cold War period I examined. I suspect divides remain compared to Defense and the JCS, but State's expertise may prove less differentiated than it once was.

Implications for War and Peace

What do the findings mean for war and peace? A theory that embeds advisory outcomes into a strategic account of conflict onset is the most compelling way to link this paper's results to war and peace. Space constraints preclude this possibility. Instead, I offer two tentative conjectures on the implications of an informational theory of bureaucracies on crisis outcomes. First, adviser input matters and should not be ignored. Preference convergence facilitates truthful and effective communication between advisers and leaders. Advisers do not offer strategic messages that leaders rationally discount. If they did, advisory processes would amount to interesting sideshows orthogonal to state behavior during crises.

Second, a leader marginalizing a bureaucracy is deeply problematic. Due to specialization, each bureaucracy offers a potentially unique piece of information and perspective to leaders. Curtailing any bureaucracy's input leaves leaders with distorted, or at least incomplete, views of the strategic environments they face. In some cases, these distortions contribute to conflict. As noted, Bush sidelined the State Department before the Iraq War, which arguably contributed to his poor assessment of the prospects for postwar stability. This assessment increased the apparent value of war, reducing the relative appeal of the status quo. President Kennedy similarly marginalized his State Department advisers before authorizing the Bay of Pigs invasion. Relying on the plan's architects from CIA, he believed widespread latent opposition to Castro would emerge in sympathy with the beachhead landing. This distorted belief inflated the operation's expected value, again rendering the status quo less attractive on a relative basis. In other cases, marginalization may promote peace—for instance, when neglected information would have made war appear more attractive. But even if contributing to peace, an informational deficit of this sort impairs a state's crisis outcome. With that additional piece of information, a leader would know she could push for more favorable bargains. Absent that information, leaders may settle for less than was available. Specialization among bureaucracies mandates that leaders solicit input from each bureaucracy. A

failure to do so impairs a leader's ability to optimize payoffs for her state.

7 Conclusion

Leaders need information during crises to optimize strategic choices that have the potential to lead to war. Leaders typically require input from others to guide these choices. Recognizing the group-based nature of such processes, I expand the analytical lens to encompass those at a leader's side. This paper analyzes how leaders acquire the necessary information and the role that bureaucratic affiliations play in shaping the information that advisers provide. Bureaucracies matter during crises, but not in the way Allison contends. Advisers subordinate parochial to national interests. As a result, their bureaucratic positions offer little insight into the policies they endorse. Prospects for a bigger budget do not consistently drive advisers to favor war or peace. Instead, bureaucratic position affects the type of information they provide. The type of information, in turn, affects the extent of uncertainty they express. Differentiation in their substantive tasks drives advisers to specialize in different substantive areas. In broad strokes, those in a foreign ministry emphasize an opponent's political traits when advising leaders. Because information is scarce for these traits, they convey ample uncertainty in their advice. The inverse conditions hold for those leading a defense ministry. Going straight to the private words of senior advisers, I construct an original corpus of documents from Cold War crises. Employing qualitative methods and tools to use text-as-data, I measure each bureaucracy's advisory content, uncertainty, and hawkishness. As theorized, the canonical preference-based account enjoys little evidentiary support. Instead, empirical patterns support an informational account of bureaucracies.

Future work assessing whether these findings hold in other time periods or places would be valuable. Additionally, efforts to explicitly link advisory content and processes to crisis outcomes would connect the dots and assure the analytical benefits to expanding scholarly attention beyond leaders to include advisers.

The findings carry implications for the literatures on informational accounts of war and the role of bureaucracies. First, the paper grounds the process by which leaders collect and make use of the information that is available in the broader strategic environment. Doing so is an essential step in the causal chain that links informational environments (e.g., complete or asymmetric) to leader beliefs to strategic choices and ultimately to war. Second, the findings suggest "where you stand depends on where you sit" is a poor heuristic for analyzing crisis advisory processes. A bureaucracy to preference mapping of that sort may be appropriate in other contexts, such as long term strategic planning. However, it lacks a systematic empirical basis during crises and is a dubious point of departure for future analyses.

For policymakers, the findings speak to the danger of curtailing a bureaucracy's involvement

during advisory processes. Due to their comparative specializations, marginalizing any one institution means forfeiting unique expertise. Marginalization breeds informational blindspots (i.e., ignorance). Informational blindspots lead to faulty decision making, potentially with catastrophic consequences given the setting. More sweepingly, and obviously, leaders ought to seek to be informed rather than ignorant. If they do not, then advisory processes are unlikely to resemble those I studied. Instead, they will be abbreviated and ill-structured, if existent at all. Emerging ignorant after defective advisory processes is a recipe for a leader to make poor decisions that reduce his country's welfare.

References

- Allison, Graham and Philip Zelikow. 1999. *Essence of Decision: Explaining the Cuban Missile Crisis*. 2nd edition ed. Pearson.
- Allison, Graham T. 1969. "Conceptual Models and the Cuban Missile Crisis." *American Political Science Review* 63(3):689–718.
- Allison, Graham T. 1971. *Essence of Decision: Explaining the Cuban Missile Crisis*. Boston: Little, Brown.
- Allison, Graham T. and Morton H. Halperin. 1972. "Bureaucratic Politics: A Paradigm and Some Policy Implications." *World Politics* 24:40–79.
- Arena, Phil and Scott Wolford. 2012. "Arms, Intelligence, and War." *International Studies Quarterly* 56(2):351–365.
- Art, Robert J. 1973. "Bureaucratic Politics and American Foreign Policy: A Critique." *Policy Sciences* 4(4):467–490.
- Bendor, Jonathan, Ami Glazer and Thomas H. Hammond. 2001. "Theories of Delegation." *Annual Review of Political Science* 4:235–269.
- Bendor, Jonathan and Thomas H. Hammond. 1992. "Rethinking Allison's Models." *American Political Science Review* 86(2):301–322.
- Bensahel, Nora. 2007. Mission Not Accomplished. In *War in Iraq: Planning and Execution*, ed. Thomas G. Mahnken and Thomas A. Keaney. New York, NY: Routledge.
- Betts, Richard K. 1991. *Soldiers, Statesmen, and Cold War Crises*. New York: Columbia University Press.
- Biddle, Stephen. 2004. *Military Power: Explaining Victory and Defeat in Modern Battle*. Princeton, NJ: Princeton University Press.
- Blainey, Geoffrey. 1988. *The Causes of War*. London, Macmillan.

- Blechman, Barry and Stephen Kaplan. 1978. *Force without War: US Armed Forces as a Political Instrument*. Washington, DC: Brookings Institution.
- Brecher, Michael and Jonathan Wilkenfeld. 1982. "Crises in World Politics." *World Politics* 34(3):380–417.
- Brecher, Michael and Jonathan Wilkenfeld. 1998. *International Crisis Behavior Project, 1918-1994*. Inter-University Consortium for Political and Social Research.
- Brooks, Risa. 2008. *Shaping Strategy: The Civil-Military Politics of Strategic Assessment*. Princeton, NJ: Princeton University Press.
- Cameron, A. and Douglas Miller. 2015. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources* 50(2):317–372.
- Chiozza, Giacomo and H.E. Goemans. 2011. *Leaders and International Conflict*. Cambridge: Cambridge University Press.
- Cohen, Eliot A. 1988. "Toward Better Net Assessment: Rethinking the European Conventional Balance." *International Security* 13(1):50–89.
- Colgan, Jeff D. 2013. "Domestic Revolutionary Leaders and International Conflict." *World Politics* 65(4):656–690.
- Croco, Sarah. 2011. "The Decider's Dilemma: Leader Culpability, Domestic Politics, and War Termination." *American Political Science Review* 105(3):457–477.
- Dafoe, Allan and Devin Caughey. 2016. "Honor and War: Souther US Presidents and the Effects of Concern for Reputation." *World Politics* 68(2).
- Destler, I.M. 1980. "National Security Management: What Presidents have Wrought." *Political Science Quarterly* 95(4):573–588.
- Dickinson, Matthew J. 2005. "Neustadt, New Institutionalism, and Presidential Decision Making: A Theory and Test." *Presidential Studies Quarterly* 35(2):259–288.
- Edelstein, David M. 2008. *Occupational Hazards: Success and Failure in Military Occupation*. Cornell University Press.
- Fearon, James D. 1994. "Political Audiences and the Escalation of International Disputes." *The American Political Science Review* 88(3).
- Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3):379–414.
- Fearon, James D. 1997. "Bargaining Over Objects that Influence Future Bargaining Power." Annual Meeting of the American Political Science Association.
- Feaver, Peter and Christopher Gelpi. 2004. *Choosing Your Battles: American Civil-Military Relations and the Use of Force*. Princeton, NJ: Princeton University Press.

- Feaver, Peter D. 2003. *Armed Servants: Agency, Oversight, and Civil-Military Relations*. Cambridge, MA: Harvard University Press.
- Fey, Mark and Kristopher W. Ramsay. 2011. "Uncertainty and Incentives in Crisis Bargaining: Game-Free Analysis of International Conflict." *American Journal of Political Science* 55(1):146–169.
- Freedman, Lawrence. 1996. "Vietnam and the Disillusioned Strategist." *International Affairs* 72(1):133–151.
- Friedberg, Aaron L. 1987. "The Assessment of Military Power: A Review Essay." *International Security* 12(3).
- Gailmard, Sean and John Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15:353–377.
- Gelman, Andrew and John Carlin. 2014. "Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors." *Perspective on Psychological Science* 9(6):641–651.
- George, Alexander. 1980. *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice*. Boulder, CO: Westview Press.
- George, Alexander L. and Erik K. Stern. 2002. "Harnessing Conflict in Foreign Policy Making: From Devil's to Multiple Advocacy." *Presidential Studies Quarterly* 32(3):484–505.
- George, Alexander L. and Richard Smoke. 1974. *Deterrence in American Foreign Policy: Theory and Practice*. New York: Columbia University Press.
- Ghosn, Faten, Glenn Palmer and Stuart Bremer. 2004. "The MID3 Data Set, 1993-2001: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science* 21(2):133–54.
- Goldberg, Jeffrey. 2016. "The Obama Doctrine." *The Atlantic* April.
- Grimmer, Justin and Brandon Stewart. 2013. "The Promise and Pitfalls of Automatic Content Analysis Methods for Political Text." *Political Analysis* 3(267-297).
- Hall, Todd H. 2015. *Emotional Diplomacy: Official Emotion on the International Stage*. Cornell University Press.
- Hastie, Trevor, Robert Tibshirani and Jerome Friedman. 2001. *Elements of Statistical Learning*. New York, NY: Springer.
- Herek, Gregory M., Irving L. Janis and Paul Huth. 1987. "Decision Making During International Crises: Is Quality of Process Related to Outcomes?" *Journal of Conflict Resolution* 31(2):203–226.
- Hermann, Margaret G. 1980. "Explaining Foreign Policy Behavior Using the Personal Characteristics of Political Leaders." *International Studies Quarterly* 24(1):7–46.

- Hermann, Margaret G. 2001. "Leaders, Groups, and Coalitions: Understanding the People and Processes in Foreign Policymaking." *International Studies Review* 3(2):47–81.
- Hermann, Margaret G. and Thomas Preston. 1994. "Presidents, Advisers, and Foreign Policy: The Effect of Leadership Style on Executive Arrangements." *Political Psychology* 15(1):75–96.
- Hoopes, Townsend. 1974. *Devil and John Foster Dulles*. Deutsch.
- Horowitz, Michael C. and Allan C. Stam. 2014. "How Prior Military Experience Influences the Future Militarized Behavior of Leaders." *International Organization* 68(3):527–559.
- Huntington, Samuel P. 1957. *The Soldier and the State: The Theory and Politics of Civil-Military Relations*. Belknap Press.
- Janis, Irving L. 1982. *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*. Boston: Houghton Mifflin.
- Johnson, Dominic. 2004. *Overconfidence and War: The Havoc and Glory of Positive Illusions*. Cambridge, MA: Harvard University Press.
- Kertzer, Joshua. 2016. *Resolve in International Politics*. Princeton, NJ: Princeton University Press.
- Krasner, Stephen. 1972. "Are Bureaucracies Important? (Or Allison Wonderland)." *Foreign Policy* 7:159–179.
- Krehbiel, Keith. 1991. *Information and Legislative Organization*. Ann Arbor, MI: University of Michigan Press.
- Kurizaki, Shuhei. 2007. "Efficient Secrecy: Public versus Private Threats in Crisis Diplomacy." *American Political Science Review* 101(3):543–558.
- Lake, David A. 2010. "Two Cheers for Bargaining Theory: Assessing Rationalist Explanations for the Iraq War." *International Security* 35(3):7–52.
- Lake, David A. and Robert Powell. 1999. *International Relations: A Strategic-Choice Approach*. In *Strategic Choice and International Relations*, ed. David A. Lake and Robert Powell. Princeton, NJ: Princeton University Press.
- Lebovic, James H. 1995. "How Organizations Learn: US Government Estimates of Foreign Military Spending." *American Journal of Political Science* 39(4):835–863.
- Lewis, David E. 2011. "Presidential Appointments and Personnel." *Annual Review of Political Science* 14:47–66.
- Little, Douglas. 1996. "His Finest Hour? Eisenhower, Lebanon, and the 1958 Middle East Crisis." *Diplomatic History* 20(1):27–54.
- Maron, M. E. and J. L. Kuhns. 1960. "On Relevance, Probabilistic Indexing, and Information Retrieval." *Journal of the Association for Computing Machinery* 7(3):216–244.

- Marsh, Kevin. 2014. "Obama's Surge: A Bureaucratic Politics Analysis of the Decision to Order a Troop Surge in the Afghanistan War." *Foreign Policy Analysis* 10(3):265–288.
- Marsh, Kevin P. and Christopher M. Jones. 2017. "Breaking Miles' Law: The Curious Case of Hillary Clinton the Hawk." *Foreign Policy Analysis* 13(3):541–560.
- McAllister, William B., Joshua Botts, Peter Cozzens and Aaron W. Marrs. 2015. *Toward "Thorough, Accurate, and Reliable": A History of the Foreign Relations of the United States Series*. Washington, DC: Office of the Historian, U.S. Department of State.
- McManus, Roseanne. forthcoming. "Fighting Words: The Effectiveness of Statements of Resolve in International Conflict." *Journal of Peace Research* .
- Meirowitz, Adam. 2006. "Designing Institutions to Aggregate Preferences and Information." *Quarterly Journal of Political Science* 1:373–392.
- Meirowitz, Adam and Anne E. Sartori. 2008. "Strategic Uncertainty as a Cause of War." *Quarterly Journal of Political Science* 3(4):327–352.
- Milner, Helen and Dustin Tingley. 2015. *Domestic Politics and American Foreign Policy: Sailing the Water's Edge*. Princeton, NJ: Princeton University Press.
- Morrow, James D. 1989. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33(4):941–972.
- Narang, Vipin. 2014. *Nuclear Strategy in the Modern Era: Regional Powers and International Conflict*. Princeton, NJ: Princeton University Press.
- Neustadt, Richard E. 1990. *Presidential Power and the Modern Presidents: The Politics of Leaders from Roosevelt to Reagan*. New York: Free Press.
- Nielsen, Rich. N.d. "Jihadi Radicalization of Muslim Clerics." Unpublished typescript.
- Pelz, Stephen. 1981. "John F. Kennedy's 1961 Vietnam War Decisions." *Journal of Strategic Studies* 4(4):356–385.
- Posen, Barry R. 1984. "Measuring the European Conventional Balance: Coping with Complexity in Threat Assessment." *International Security* 12(4):186–202.
- Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton, N.J.: Princeton University Press.
- Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71(S1):S267–S277.
- Press, Daryl G. 2005. *Calculating Credibility*. Ithaca, NY: Cornell University Press.
- Preston, Andrew. 2001. "The Little State Department: McGeorge Bundy and the National Security Council Staff, 1961-195." *Presidential Studies Quarterly* 31(4):635–659.

- Preston, Thomas and Pau 't Hart. 1999. "Understanding and Evaluating Bureaucratic Politics: The Nexus Between Political Leaders and Advisory Systems." *Political Psychology* 20(1):49–98.
- Rapport, Aaron. 2012. "The Long and Short of It: Cognitive Constraints on Leaders' Assessments of "Postwar" Iraq." *International Security* 37(3):133–171.
- Rathbun, Brian C. 2012. *Trust in International Cooperation International Security Institutions, Domestic Politics and American Multilateralism*. Cambridge: Cambridge University Press.
- Rhodes, Edward. 1994. "Do Bureaucratic Politics Matter?: Some Disconfirming Findings from the Case of the US Navy." *World Politics* 47(1):1–41.
- Ricks, Thomas E. 2006. *Fiasco: The American Military Adventure in Iraq*. New York, NY: Penguin Press.
- Ries, Charles P. 2016. Improving Decisionmaking in a Turbulent World. Technical report RAND Corporation.
- Rosen, Stephen P. 1991. Net Assessment as an Analytical Concept. In *On Not Confusing Ourselves: Essays in Honor of Albert and Roberta Wohlstetter*, ed. Andrew W. Marshal et al. Boulder, CO: Westview Press.
- Rudalevige, Andrew. 2005. "The Structure of Leadership: Presidents, Hierarchies, and Information Flow." *Presidential Studies Quarterly* 35(2):333–360.
- Sartori, Anne E. 2002. *Deterrence by Diplomacy*. Princeton, NJ: Princeton University Press.
- Saunders, Elizabeth N. 2011. *Leaders at War: How Presidents Shape Military Interventions*. Cornell University Press.
- Saunders, Elizabeth N. 2017. "No Substitute for Experience: Presidents, Advisers, and Information in Group Decision-Making." *International Organization* 71(S1):S219–S247.
- Slantchev, Branislav L. 2010. "Feigning Weakness." *International Organization* 64(3):357–388.
- Slantchev, Branislav L. and Ahmer Tarar. 2011. "Mutual Optimism as a Rationalist Explanation for War." *American Journal of Political Science* 55(1):135–148.
- Snyder, Glenn H. and Paul Diesing. 1977. *Conflict Among Nations: Bargaining, Decision Making and System Structure in International Crises*. Princeton, NJ: Princeton University Press.
- Sullivan, Patricia L. 2007. "War Aims and War Outcomes: Why Powerful States Lose Limited Wars." *Journal of Conflict Resolution* 51(3):496–524.
- Trager, Robert F. 2010. "Diplomatic Calculus in Anarchy." *American Political Science Review* 104(2):347–368.
- Van Evera, Stephen. 1984. "The Cult of the Offensive and the Origins of the First World War." *International Security* 9(1):58–107.

- Weeks, Jessica L. P. 2014. *Dictators at War and Peace*. Ithaca, NY: Cornell University Press.
- Weisiger, Alex and Keren Yarhi-Milo. 2015. "Revisiting Reputation: How Past Actions Matter in International Politics." *International Organization* 69(2):473–495.
- Welch, David A. 1992. "The Organizational Process and Bureaucratic Politics Paradigms: Retrospect and Prospect." *International Security* 17(2):112–146.
- Whitlark, Rachel E. 2017. "Nuclear Beliefs: A Leader-Focused Theory of Counter-Proliferation." *Security Studies* 26(4):545–574.
- Yarhi-Milo, Keren. 2014. *Knowing the Adversary: Leaders, Intelligence Organization, and Assessments of Intentions in International Relations*. Princeton, NJ: Princeton University Press.
- Yarhi-Milo, Keren, Joshua D. Kertzer and Jonathan Renshon. forthcoming. "Tying Hands, Sinking Costs, and Leader Attributes." *Journal of Conflict Resolution* .