Advisers and Aggregation in Foreign Policy Decision-Making*

Tyler Jost† Joshua D. Kertzer‡ Eric Min§ Robert Schub¶
Last modified: July 19, 2023

Word count: 13917

Abstract: Do advisers affect foreign policy and, if so, how? Recent scholarship on elite decision-making prioritizes leaders and the institutions that surround them, rather than the dispositions of advisers themselves. We argue that despite the hierarchical nature of foreign policy decision-making, advisers' predispositions towards the use of force shape state behavior through the counsel advisers provide in deliberations. We test our argument by introducing three original datasets, including an original dataset of 2,685 foreign policy deliberations between US presidents and their advisers from 1947 to 1988. Applying a novel machine learning approach to estimate the hawkishness of 1,134 Cold War-era foreign policy decision-makers, we show that adviser-level hawkishness affects both the counsel advisers provide in deliberations, and the decisions the leader makes: conflictual policy choices grow more likely as hawks increasingly dominate the debate, even when accounting for leader dispositions. These results enrich our understanding of international conflict by demonstrating how advisers' dispositions, which aggregate via deliberation, systematically shape the choices leaders make.

*This research was funded by the Defense Advanced Research Projects Agency (Grant #W911NF1920162). We're also grateful to the Institute of Quantitative Social Science (IQSS) and the Weatherhead Center for International Affairs (WCFIA) at Harvard University for their logistical support, and to Diana Agdashian, Saadi Ali, Perry Arrasmith, Christine Bienes, Florian Bochert, Lexa Brenner, Alexis Carraway, Youmna Chamieh, Erica Colston, Joshua Conde, Vanessa Cordova, Jacinta Crestanello, Caroline Curran, George Dalianis, Charlie Davis, Darius Diamond, Iain De Boer, A.J. Dilts, Jelena Dragicevic, Corbin Duncan, Will Engelmayer, Karla Flores, Kendrick Foster, Angela Fu, John Gehman, John Gleb, Hope Ha, Joolean Hall, Joseph Hasbrouck, Sarah Jacob, Joshua Kelley, Joel Kesselbrenner, Samantha Ku, Jacob Kuckelman, Max Kuhelj Bugaric, Zaki Lakhani, Johannes Lang, Steven LeGere, Isabel Levin, Rachel Levine, Daiana Lilo, Angel Lira, Clare Lonergan, Parker Mas, Holly McCarthy, Alexandra McClintock, Sean McGuire, Deja McKnight, Scarlett Neely, Yihao Niu, Denise Nzeyimana, Christine Ow, Clay Oxford, Nic Pantelick, June Park, Marc Prophet, Kasey Quicksill, Janna Ramadan, Bianca Rodriguez, Kay Rollins, Noah Schimenti, Isabel Slavinsky, Bruno Snow, Matthew Stafford, Randle Steinbeck, Claire Sukumar, Ollie Sughrue, Hannah Suh, Maya Taha, Sal Vidal, Matthew Von, Eileen Walsh, and Taylor Whitsell for research assistance. Authors listed in alphabetical order.

†Assistant Professor, Department of Political Science, Brown University. Email: tyler_jost@brown.edu. Web: http://www.tylerjost.com

‡John Zwaanstra Professor of International Studies and of Government, Department of Government, Harvard University. 1737 Cambridge St, Cambridge MA 02138. Email: jkertzer@gov.harvard.edu. Web: https://jkertzer.sites.fas.harvard.edu

§Assistant Professor, Department of Political Science, University of California, Los Angeles. Email: eric.min@ucla.edu. Web: http://www.ericmin.com

¶Assistant Professor, Department of Political Science, Rutgers University. Email: robert.schub@rutgers.edu. Web: http://www.robertschub.com

The past decade has seen a resurgence of interest in the study of leaders in International Relations (Saunders, 2011; Chiozza and Goemans, 2011; Croco, 2011; Jervis, 2013; Rathbun, 2014; Weeks, 2014; Kertzer, 2016; Renshon, 2017; Whitlark, 2017; Yarhi-Milo, 2018; Holmes, 2018; McManus, 2021). Who assumes the role of president, prime minister, or dictator dramatically shapes the foreign policies a state chooses, in large part because leaders systematically vary from one another: in their experiences before entering office (Horowitz, Stam and Ellis, 2015; Fuhrmann, 2020); in their personalities, leadership styles, and operational codes (Kaarbo, 1998; Hermann and Preston, 1994); and in their traits or dispositions like hawkishness (Yarhi-Milo, Kertzer and Renshon, 2018; Carter and Smith, 2020). Leader characteristics, this body of scholarship has argued, are critically important to understanding when states are prone to interstate conflict.

One of the critiques of the study of leaders in IR, like the study of political psychology more generally, concerns the problem of aggregation: although leaders sometimes act alone, many of the most important decisions in foreign policy are made in small groups (Kertzer et al., 2022). History is replete with images of advisers at their leaders' side at critical moments: Otto von Bismarck during the Franco-Prussian War, Vo Nguyen Giap during Dien Bien Phu, John Foster Dulles during the Berlin Crisis, Moshe Dayan during the Six Days War, P.N. Haksar during the Bangladesh War, and so on.

Do advisers systematically shape foreign policy behavior and, if so, how? Whether in reference to the "best and brightest" in Vietnam, or the "Vulcans" who advocated for the invasion of Iraq, many popular accounts of foreign policy decision-making suggest that critical foreign policy choices often hinge on whether a leader's inner circle is filled with hawks or doves. Yet much of the academic literature in IR has relatively little to say about the role of adviser characteristics in shaping foreign policy. The new wave of first-image IR scholarship predominantly focuses on leaders rather than those that counsel them. This is partly for substantive reasons. Foreign policy decision-making is hierarchical, which might lead us to suspect that leader dispositions dominate those of advisers, that leaders tend to appoint advisers with similar worldviews, and that leaders disregard advice incongruent with their prior beliefs. The asymmetry in focus also stems from methodological considerations: thanks to large-scale data-collection efforts (Chiozza and Goemans, 2011; Horowitz, Stam and Ellis, 2015), we have excellent data on leader-level characteristics, but relatively little data on adviser-level ones.

Moreover, the scholarship that does exist tends to study advisers situationally rather than dispositionally, focusing on the quality of the group's advisory *process* and how leaders and advisers interact, rather than on advisers' predispositions themselves. The bureaucratic politics literature, for example, argues that the recommendations that advisers supply flow from their institutional affiliation; where advisers stand depends on where they sit, rather than stable and enduring predispositions (Allison and Zelikow, 1999). Similarly, a rich literature studies advisory systems from the perspective of institutional design, showing that some decision-making processes ('t Hart, Stern and Sundelius, 1997; Schafer and Crichlow, 2010; Mintz and Wayne, 2016), decision-making units (Hermann and Preston, 1994; LeVeck and Narang, 2017), and institutions (Reiter and Stam, 2002; Weeks, 2014; Jost, 2023) yield more accurate assessments, less biased information provision, and more effective policy outcomes. Where discussions of dispositions arise in this literature, it tends to be how *leaders'* dispositions affect the advisory structures that leaders establish (e.g. George, 1980), rather than how advisers' dispositions affect how advisers behave.[1] As Preston (2001, 267) notes, we have much still to learn about whether and how adviser dispositions shape foreign policy behavior, if at all.

In this article, we develop a theory of foreign policy decision-making that emphasizes adviser dispositions. The uncertainty, complexity, and ill-defined nature of foreign policy decision-making means that leaders turn to advisers for counsel. The kind of counsel advisers offer—the information they share, the analysis they provide, and the policies they recommend—is shaped by advisers' core dispositions about the desirability and efficacy of the use of force. As a result, despite the decision-making authority that leaders retain, foreign policy choices bear the fingerprints of advisers as much as those of leaders. Against the claims of the bureaucratic politics literature, those fingerprints are not reducible to the adviser's institutional role: *who* occupies the most important positions of government affects where those advisers "stand" and the corresponding counsel they supply. In short, our theoretical contribution is to offer a theory of foreign policy decision-making that shifts the focus to *advisers* rather than advisory systems (e.g., George, 1980), to *bureaucrats* rather than bureaucracy (e.g., Allison and Zelikow, 1999).

To test our theory, we leverage big data and machine learning techniques to offer systematic

[1] When studies do take the important next step of incorporating a personal characteristic of advisers, the relevance of the characteristic still depends on how it relates to the leader's own characteristics. Adviser experience, for instance, is argued to affect adviser behavior conditional on the *leader's* experience (Saunders, 2017).

and unusually high-resolution empirical tests on how advisers shape foreign policy. First, we located, collected, digitized, and processed the transcripts of 2,685 US foreign policy decision-making meetings from 1947 to 1988. We compiled the records through a combination of in-person collection at seven libraries and archives, as well as from online repositories. These include all available meeting records of the US National Security Council (NSC), as well as 1,894 informal meetings in which presidents discussed foreign policy issues with their advisers. We segmented each meeting transcript by speaker, meaning that our data identify not only which advisers provided counsel, but what substantive topics they emphasized. We also reviewed each transcript to manually code over 2,500 decisions that leaders made during these meetings, which ranged from diplomatic cooperation to interstate conflict. These include some of the most consequential foreign policy choices in modern American history, such as the decision to blockade Soviet ships during the Cuban Missile Crisis and the decision to engage Mikhail Gorbachev during the Reykjavik Summit. We believe this to be the most comprehensive resource available by which scholars can study the microfoundations of foreign policy decision-making.

Second, we apply a novel machine learning-based approach that estimates (at a distance) leader and adviser dispositions, such as hawkishness, based off an original dataset of the biographical characteristics of every individual who participated in the meetings in our sample. Our biographical data describe the backgrounds and professional experiences of 1,134 individuals ranging from secretaries of state to Pentagon bureaucrats—an adviser-level counterpart to the leader-level datasets that have revolutionized the study of leaders in IR. This innovation allows us to study advisers in US foreign policy on a far larger scale than has traditionally been possible, and to study quantitatively what has traditionally been the preserve of qualitative approaches. Collectively, this approach enables us to study not only whether dispositions shape the types of counsel advisers provide, but also whether the aggregated dispositions of advisory groups shape the choices that leaders ultimately make.

Analyzing these data yields two major findings demonstrating the central importance of advisers in foreign policy decision-making. First, we find that adviser dispositions shape the counsel that leaders receive when making consequential foreign policy choices. American presidents consistently solicit information from advisers during deliberations—and advisers use these opportunities to offer information, perspectives, and policy recommendations congruent with their dispositions. Second,

we show that hawkish advisory groups are associated with more conflictual foreign policies, even after considering several potential pathways for selection effects. Contrary to accounts assuming that decision-making outcomes simply reflect leader dispositions, we find that adviser-level hawkishness has large and systematic effects on foreign policy decision-making—and that appointment to and participation in foreign policy groups does not merely mirror the hawkishness of the leader. The theory and findings collectively illustrate the formidable influence that foreign policy advisers can wield by supplying the counsel that leaders demand.

1 Leaders, Advisers, and Aggregation in Groups

The division between hawks and doves is central to our understanding of why states choose conflict over cooperation (Herrmann, Tetlock and Visser, 1999; Schultz, 2005). Hawks and doves differ in their beliefs about the nature of international politics, adversary motivations, and the efficacy and appropriateness of using force, shaping the foreign policies that individuals support (Rathbun, 2007; Yarhi-Milo, Kertzer and Renshon, 2018; Lin-Greenberg, 2019). Consequently, knowing an individual's *general* hawkishness often predicts their propensity to endorse *specific* conflictual policies.

Despite the centrality of hawkishness to our understanding of policy preferences, there is debate about whether and how individual dispositions like hawkishness aggregate to shape state behavior (Hafner-Burton et al., 2017; Kertzer et al., 2022). Most foreign policy decisions occur in group settings in which leaders and advisers interact. American presidents made key decisions during episodes spanning from the Berlin Crisis to the Persian Gulf War in consultation with advisers ranging from John Foster Dulles to Colin Powell. Prior to the Soviet Union's invasion of Afghanistan, Leonid Brezhnev conferred with an advisory troika consisting of Defense Minister Dmitry Ustinov, KGB director Yuri Andropov, and Foreign Minister Andrei Gromyko. Even Richard Nixon and Mao Zedong—known for domineering decision-making styles—routinely relied on advisers, such as Henry Kissinger and Zhou Enlai, respectively.

Like the leaders they serve, advisers presumably possess stable dispositions like hawkishness. The central question, however, is whether and how these leader- and adviser-level traits aggregate to shape foreign policy outcomes. The existing literature suggests two answers. The first proposes

that the dispositions of group members (leaders or advisers) have no bearing on the foreign policy decisions reached—an assumption shared by several disparate theoretical traditions in IR. In the second, individual traits do matter, but because foreign policy is hierarchical, foreign policy decisions simply follow from leaders' dispositions. We discuss each in turn before presenting our dispositional model of advising.

Figure 1: Three Models of Aggregation in Groups

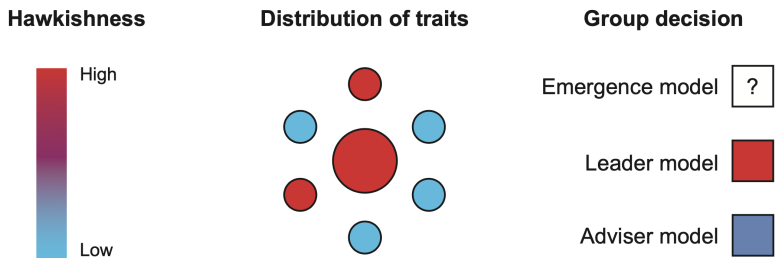

Note: Large circles represent leaders, small circles represent advisers.

1.1 The Emergence Model

Several theoretical traditions in IR assume that the dispositions of group members—leaders or advisers—should not aggregate in systematic or predictable ways. These theories offer markedly distinct justifications for this nonetheless common assumption. First, realist scholarship emphasizes that the structure of the international system creates powerful incentives for states to behave as "unitary actors." In this view, the domestic features of the state, including the types of individuals constituting decision-making groups, exert little effect on state behavior compared to structural variables, such as polarity, the balance of power, and alliances (e.g., Mearsheimer, 2001). Second, some game theoretic scholarship remains skeptical about the challenges of aggregation, suggesting that the aggregation of traits in groups is sufficiently complex so as to warrant a "methodological bet" that they are not worth studying (Powell, 2017). Third, a body of holist or constructivist scholarship argues that group-level outcomes cannot be reduced to attributes of the members comprising the group (Wendt, 2004). Just as international politics is a complex social system (Jervis, 1998), so too is domestic policymaking. While these intellectual traditions differ dramatically, they share an assumption that when studying state behavior, the individual level of analysis is the wrong place to look: neither leader nor adviser characteristics should neatly map onto a group's foreign

policy decisions, as the top model in Figure 1 indicates.

1.2 The Leader Model

A second view of aggregation, evident in much recent work, focuses on the traits of leaders (the middle model in Figure 1). Simply put, some leaders are more hawkish than others, and states led by hawkish leaders are more likely to engage in conflictual behavior (Jervis, 2013, 165, Carter and Smith, 2020). For example, Yarhi-Milo (2018, 82) finds that some American presidents during the Cold War, such as Richard Nixon and Ronald Reagan, tended to exhibit more hawkishness than others, such as John F. Kennedy and Jimmy Carter. Other studies similarly note the division between hawkish and dovish leaders in other countries, such as China and India (e.g., Kennedy, 2011).

The leader model posits that group decisions reflect leader traits for two reasons. First, foreign policy decision-making groups tend to be hierarchical: leaders enjoy more substantive and procedural authority than other group members. Leaders do not just make the final decisions in foreign policy; they also set the rules for decision-making. In this view, leaders might be able to impose their worldview on policy by strategically manipulating adviser appointments or participation in decision-making. As Krasner (1972, 166) argues, "The President chooses most of the important players and sets the rules... These individuals must share his values." Saunders (2011, 26) similarly notes that leaders can shape the decision-making group by "hiring advisers or government officials who share [similar] beliefs." Byman and Pollack (2001, 143) suggest that adviser preferences are often "determined by the leader," rather than the adviser's bureaucratic position or worldview. Leaders might also structure the decision-making process to afford privileged access to advisers with congruent dispositions. Leaders might deliberately manipulate which advisers participate in meetings, steer discussions in directions that suit their preferences, or include a "domesticated dissenter" within meetings to fabricate the appearance of debate when, in fact, their decision has already been made. For example, Lyndon Johnson excluded Vice President Hubert Humphrey from policy deliberations on Vietnam in 1965 after Humphrey expressed opposition to escalation (Burke and Greenstein, 1989). If leaders strictly surround themselves with like-minded advisers, or manipulate the advisory process to ensure the information they receive predominately reflects their worldview, the real causal power comes from the traits of the leader, not those of the advisory

group.

Second, the leader model argues that leader beliefs tend to supersede the input that advisers provide. In this view, leaders enter office with fixed preferences and firm beliefs about optimal strategies to achieve them (e.g., George, 1969; Byman and Pollack, 2001; Saunders, 2011). When making decisions under uncertainty, leaders may privilege their own "mental Rolodex" regarding the nature of international politics (Horowitz, Stam and Ellis, 2015, 10), placing more emphasis on "vivid, personalized, and emotionally involving" information from first-hand experiences, rather than the abstract intelligence provided by their bureaucratic advisers (Yarhi-Milo, 2014, 16). Cognitive barriers, such as the desire for consistency (Jervis, 1976) and motivated reasoning (Kertzer, Rathbun and Rathbun, 2020), may also cause hawkish leaders to prioritize input from hawkish advisers and discount it from dovish confidants—and vice versa. As such, advisers are "influential" only when the input they provide is congruent with what the leader already believes. But if the only adviser input that matters is that which confirms the leader's prior beliefs, the effects of adviser dispositions are once again epiphenomenal to those of leaders themselves.

Perhaps because of the presumed importance of leaders, far more of the existing empirical literature focuses on leaders rather than their advisers (Chiozza and Goemans, 2011; Horowitz, Stam and Ellis, 2015; Kertzer, 2016; Carter and Smith, 2020). Qualitative approaches to studying elite decision-making often consider advisers chiefly to illustrate the leader's importance by providing a counterfactual of what a different individual might have done in the same situation (e.g. Jervis, 2013, 149; Yarhi-Milo, 2018, 42). This asymmetry in focus is also likely a function of methodological considerations: in quantitative IR, we have excellent datasets of leader-level background characteristics, but as of yet, nothing comparable for adviser-level characteristics.

1.3 The Adviser Model

In contrast to the emergence and leader-centric views that much of the existing literature espouses, we develop a model of aggregation emphasizing advisers' dispositions. Our adviser model is based on three simple intuitions. First, the challenges of foreign policy decision-making causes leaders to turn to advisers for counsel. Second, advisers have stable predispositions about foreign policy. Third, these predispositions affect the nature of the counsel that advisers provide during deliberations—and thus the decisions that leaders are likely to make. We discuss each in turn.

Our model begins with an assumption as well known to scholars of foreign policy analysis as it is to leaders themselves: foreign policy decision-making is hard (Snyder, Bruck and Sapin, 1962; Jervis, 1976). International politics is characterized by ill-defined problems in which the nature of the situation, potential solutions, and even the optimal outcome are fundamentally unclear (Sylvan and Voss, 1998). Before making foreign policy decisions, leaders must determine the type of situation they face, what information, interests, and norms are pertinent, and adjudicate between conflicting accounts, all while facing time constraints, information constraints, irreducible uncertainty, and complex value tradeoffs (George, 1980).

It is here that advisers are useful. Advisers in foreign policy do many things: they offer emotional support and companionship to leaders coping with the stress of decisionmaking (George, 1980, 80, Goldhamer, 1978, 8), they bestow the leader with public legitimacy and political cover (George, 1980, 81; Kenwick and Maxey, 2022), and they coerce leaders using the threat of public protest (Saunders, 2018). Yet advisers do not just provide comfort, cover, and coercion: they also provide counsel, which we can understand as consisting of three distinct functions.[2]

First, advisers engage in *information provision*, giving leaders information they need about the state of the world (Schub, 2022). Historically, advisers served as the king's "eyes and ears and hands and feet" (Aristotle, 1920, Book III, 1287b). Today, they monitor intelligence, diplomatic cables, and news reports. Sometimes, even questions as mundane as "What happened?" are not straightforward, even for highly experienced leaders (Katagiri and Min, 2019). During the Gulf of Tonkin crisis in August 1964, Lyndon Johnson and his advisers struggled for hours to determine whether North Vietnam had conducted a second attack on the USS Maddox. The sheer volume of information contemporary leaders command—by the mid-1960s, US ambassadors were sending 400,000 words a day by telegraph (Goldhamer, 1978, 58)—means that advisers do not merely provide information to the leaders they serve, but also screen it (Burke and Greenstein, 1989, 6), choosing what pieces of information to relay and what to hold back.

Second, advisers engage in *problem representation*, helping leaders develop a "definition of the situation" they face (Sylvan and Voss, 1998; Snyder, Bruck and Sapin, 1962). This function, which George (1980, 240) called the "diagnostic" function of advising, and Maoz (1990) refers to as

[2]The three counseling functions we identify—problem representation, information provision, and policy recommendations—map nicely onto Destler's (1972, 10) claim that advisers provide "analysis, information, and advice."

"framing," is less about gathering information than it is about interpreting it. Is the conflict in Korea in 1950 a civil war or an act of communist aggression? Is Ho Chi Minh a local nationalist or a Soviet puppet? Should the uptick in violence in Iraq in 2007 be understood as terrorism or insurgency? This is why analogical reasoning is so powerful in foreign policy, since it offers decision-makers schemas they can use to define what values or interests are at stake in a given crisis (Khong, 1992). Much of what advisers do in foreign policy deliberations consists of offering leaders these schemas or perspectives, as reflected in documents like memos written by McGeorge Bundy and George Ball in 1964 and 1965 with titles like "Vietnam: what is our interest there and our object?" and "How valid are the assumptions underlying our Vietnam policies?" In this sense, advisers do not just provide leaders with information, they also provide them with theories.

Third, advisers provide *policy recommendations*, helping leaders select the optimal strategy given the situation they face (Jost, 2023). Burke and Greenstein (1989) call this "reality testing," and George (1980, 243–244) calls it "option assessment," as decision-makers assess the expected consequences of different courses of action. In the Cuban Missile Crisis, for example, advisers deliberated about whether the United States should respond diplomatically, with air strikes, or with a blockade. Just as strategic scripts follow from the images that precede them (Herrmann and Fischerkeller, 1995), the policy recommendations that advisers advocate for are intimately connected to the problem representations to which they subscribe. Jervis (1976), for instance, noted that most of the debates during the Cold War depended on whether observers viewed US-Soviet relations through the prism of the spiral model or the deterrence model; which model you embraced determined what policies you favored.

Our next assumption posits that advisers have stable and well-defined predispositions that shape the way they view foreign policy. Some, like Curtis LeMay and Donald Rumsfeld, are relatively hawkish, while others, like Cyrus Vance and George Ball, are relatively dovish. The notion that leaders systematically differ from one another in their predispositions—whether operationalized as ideological belief systems, personalities, worldviews, leadership styles, or operational codes—is well established in the foreign policy literature (e.g. George, 1969; Preston, 2001; Whitlark, 2017; Yarhi-Milo, 2018). Much like the leaders they serve (e.g. Horowitz, Stam and Ellis, 2015), adviser traits are forged by early experiences which continue to shape how they behave when in office decades later. Experience planning Allied bombing campaigns against Japan during World War II shaped

Robert McNamara's assessments of the feasibility of the bombing campaign against Hanoi during the Vietnam War. Experiences touring American aircraft carriers in the early 1980s colored the views of Chinese admiral Liu Huaqing during Politburo debates in the 1990s concerning Chinese naval modernization.

Finally, we argue that adviser predispositions shape the counsel advisers provide during deliberations, and thus, the decisions that leaders make. Research in political psychology leads us to expect that an adviser's dispositions affects all three of the counseling functions—information provision, problem definition, and policy recommendations. Scholarship on motivated reasoning and confirmation bias argues that our predispositions affect not only the information decision-makers seek out, but also how persuasive they find that information to be (Taber and Lodge, 2006; Kertzer, Rathbun and Rathbun, 2020). Advisers in the George W. Bush administration, for example, were convinced that Iraq had weapons of mass destruction, so they tasked intelligence officers to look for any indication that Iraqi WMDs existed. Predispositions also affect how decision-makers define the situation they face, responding to ill-defined problems by anchoring on their core dispositions. Hawks and doves facing the same strategic setting tend to perceive the situation in fundamentally different ways (Brutger and Kertzer, 2018), suggesting one reason why hawks after the Cold War continued to perceive the same level of international threat even after the Berlin Wall fell (Murray, 2002). Finally, as the discussion of the spiral and deterrence model above shows, predispositions affect the policies we prefer. This assumption lies at the heart of hierarchical models of foreign policy preferences, which envision our more general orientations towards foreign policy (e.g., hawkishness) shaping our preferences towards the use of force in particular circumstances (Hurwitz and Peffley, 1987; Herrmann, Tetlock and Visser, 1999; Rathbun, 2007).

In sum, our adviser model predicts that foreign policy decisions will at least partially reflect the dispositions of the advisers who participate in the deliberations (the bottom model in Figure 1). The model suggests a number of observable implications. First, leaders should seek advisers' counsel. Leaders should meet with advisers, and in these meetings, leaders should ask questions, solicit advice, and request clarifications, rather than meetings merely being pro forma opportunities for leaders to keep their subordinates informed. Second, the counsel that advisers supply in these meetings should systematically differ based on their dispositions: hawks and doves should emphasize different pieces of information, or interpret the information in different ways, engaging in arguments

and counter-arguments as they compete with one another over the direction they want leaders to take. Third, the dispositions of advisers participating in deliberations should affect the foreign policy decisions reached. The more hawks dominate the discussion, the more conflictual decisions the group should make.

It is worth noting what is distinctive about our approach. First, it explicitly studies advisers dispositionally, rather than situationally. By drawing our attention to how dispositions affects the information, problem representation and policy recommendations that advisers provide, our dispositional focus not only connects the study of advisers to the study of political psychology more generally, but also leads to substantively different predictions. Unlike the bureaucratic politics literature, whose situational focus assumes that where advisers stand is based on where they sit (Allison and Zelikow, 1999), we argue that advisers' recommendations flow from predispositions that are not reducible to their institutional role. Second, recent scholarship has tended to focus on *coercive* pathways to advisory influence, showing that advisers matter because of the costs they can impose on leaders outside of meetings, such as the threat of leaks or public criticism (e.g. Garrison, 1999; Saunders, 2018, 22–23). In contrast, our model rests on a *counseling* pathway to advisory influence: advisers shape decision-making because they fulfill a leader's psychological and informational needs during deliberations. Our model thus complements this recent wave of research by identifying an additional pathway though which advisers matter despite ultimately being subservient to the leaders they serve.

Finally, it is important to emphasize that the three models of aggregation are not necessarily mutually exclusive. The leader model might explain certain decisions, while the adviser model explains others. Sometimes leaders may know both what they want to do and how they want to do it, such that no amount of information or deliberation will sway them — whereupon we would expect leaders to manipulate the decision-making process to obtain their desired outcome. The value of theorizing an adviser model stems from the fact that there are many circumstances in which leaders are uncertain about which direction to take. Ultimately, the performance of each model is an empirical question we seek to test below.

2 Data

To test our adviser model, we systematically collected and analyzed a large set of archival records documenting high-level foreign policy meetings in the United States from 1947 to 1988. We first used these archival records to identify the participants in foreign policy decision-making. We then measured participant hawkishness (our explanatory variable) from a distance using a novel methodological approach. All data were collected specifically for this study. Figure 2 visually summarizes our main datasets and the steps taken to convert them into the measures we ultimately use to test the three models of aggregation. We describe each step below and provide additional details in Appendices §1-4.

2.1 Identifying Group Participants and Decisions using Archival Records

2.1.1 Archival Record Collection, Processing, and Segmentation

There are a wide number of contexts we could use to study trait aggregation in foreign policy, but we focus here on the United States. As a global hegemon with the largest military budget in the world, the US represents a substantively important case. Crucially, the US maintains an unusually well-kept set of historical records from 1947-1988 of both formal and informal meetings, which we assembled from two types of sources. First, a team of research assistants photographed over a half million pages of records in six presidential libraries, the US National Archives II in College Park, and several other print and digitized resources pertinent to foreign policy decision-making across eight presidential administrations from Harry Truman to Ronald Reagan.[3] Second, using an automated text scraping protocol, we collected all records included in the *Foreign Relations of the United States* (*FRUS*) archival database. Using a combination of automated filtering and manual review by research assistants, we extracted all FRUS meeting records that met two scope criteria: presidential participation and no participation by foreign leaders (i.e. no diplomatic exchanges).[4] Note that this excludes informal meetings between advisers in which the president did not participate.

As Appendix §1.3 summarizes, the collection process yielded 2,685 meeting records. 791 (29%

[3] See Appendix §1.
[4] See Appendix §1.2.

Figure 2: Dataset Construction

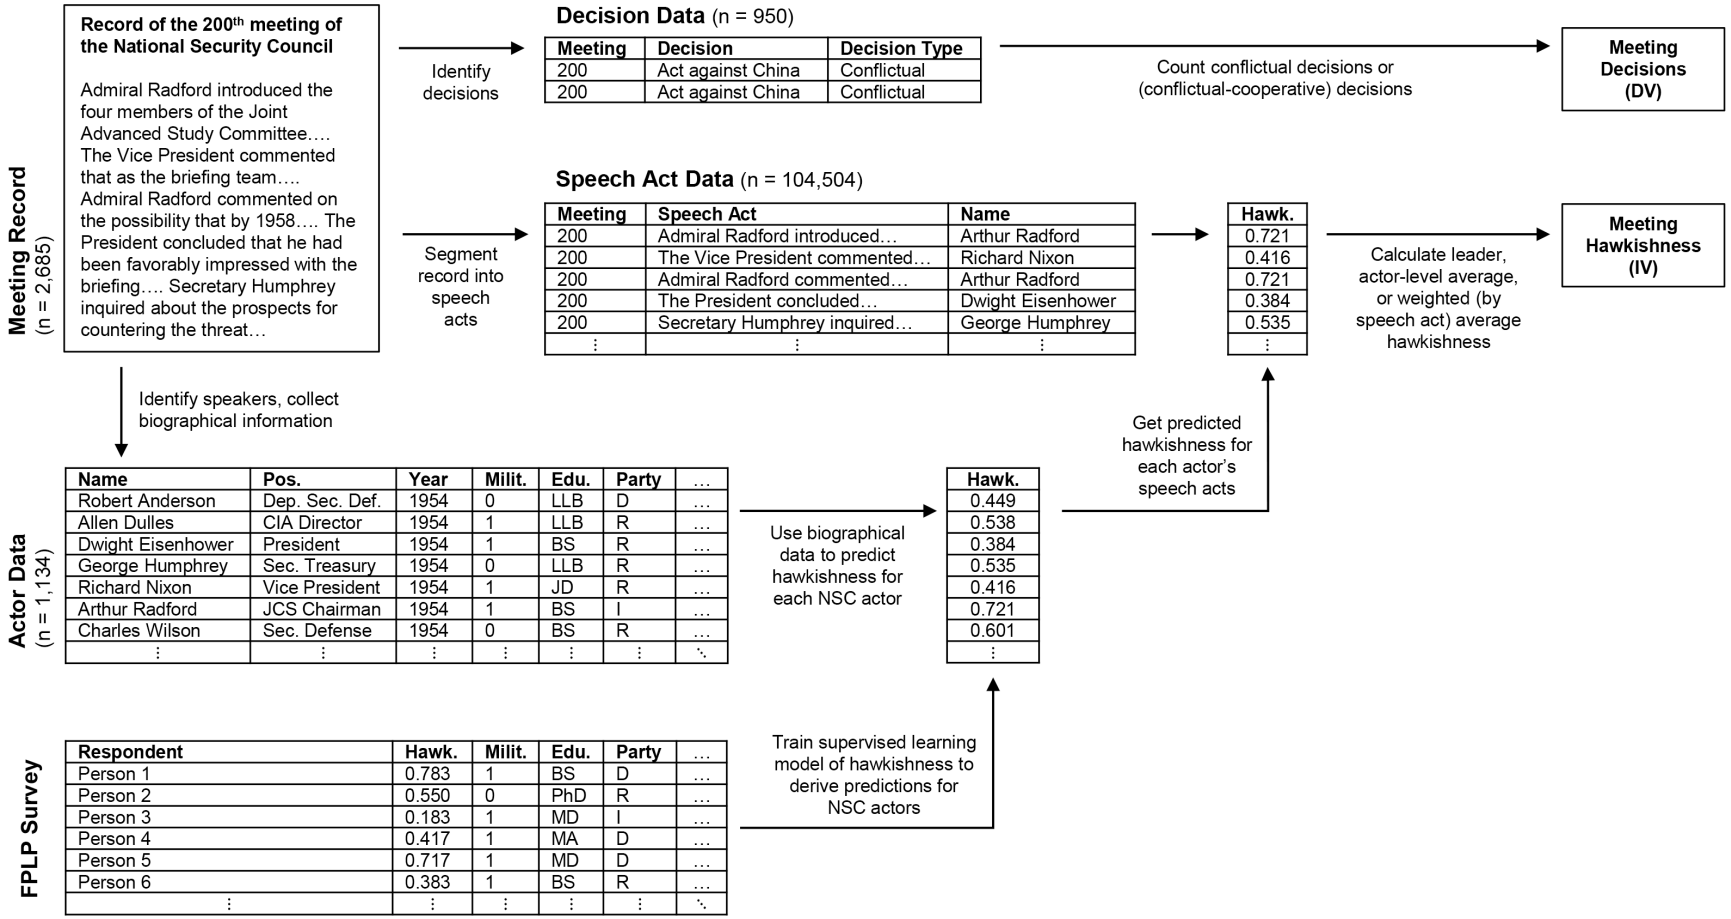

of the sample) of the meetings were formal meetings of the NSC. An additional 1,894 (71% of the sample) were informal meetings. The inclusion of informal meetings in our sample is particularly important because not all presidents have used the NSC in the same way (Jost, 2023). Some informal meetings featured only the president and a single adviser, while others included dozens of bureaucratic officials. We believe these data constitute the most extensive and complete collection to date of foreign policy meeting records in any country—although we do not claim that our sample encompasses all meetings the president attended.[5] The largest single set of records and decisions in our data come from the Eisenhower administration—reflecting the extent to which Eisenhower utilized a highly formalistic advising system in which the NSC played an outsized role. On the one hand, given his formalistic advising structure, Eisenhower's prominence in this data might bias our results in favor of finding evidence that adviser dispositions matter. On the other hand, existing theories might imply that adviser dispositions might be less likely to matter during this period, as Eisenhower's extensive foreign policy experience prior to his presidency may have improved his ability to monitor adviser information provision better than foreign policy novices like Carter and Reagan (Preston, 2001, 21; Saunders, 2017). We discuss the potential ways that classified records (particularly for the Carter and Reagan administrations) and the *FRUS* editorial process might affect our meeting sample in Appendix §1.2.

Research assistants used optical character recognition software to convert the photographs of meeting records into digital text, manually correcting text recognition errors. We then split each meeting record into what we call "speech acts"—the uninterrupted words spoken by a single individual during a meeting.[6] Our records featured 2,685 meetings containing 104,504 speech acts made by 1,134 unique participants.

2.2 Explanatory Variable: Measuring Hawkishness with Biographical Data

To test our dispositional model of advising, we need a measure of hawkishness for each of the 1,134 people identified in the meeting records. A major methodological challenge to the study of elite decision-making is that researchers lack detailed information on the numerous individuals,

[5] Our sample includes approximately 195 meeting records that were omitted from *FRUS*, as well as about 400 records for which *FRUS* provides only partial excerpts. We estimate that our sample includes 81% of all NSC meetings convened during this period.

[6] See Appendix §2.

most of them advisers, in decision-making groups. At present, there are no comprehensive datasets on adviser characteristics comparable to leader biographical data (Chiozza and Goemans, 2011; Horowitz, Stam and Ellis, 2015) to enable researchers to study advisers in a nomothetic way. Moreover, even when researchers are able to identify *which* advisers participate in decision-making, they lack a stable measure to estimate adviser traits and dispositions, such as hawkishness, at a distance. Researchers are often only able to identify an individual as a hawk or dove by observing the position that an adviser takes on a particular issue (e.g., Feaver, 2003).

We introduce a two-pronged methodological innovation to address this challenge. It pairs systematic coding of the biographic characteristics of presidents and advisers with past surveys administered to real foreign policy elites during the Cold War. This allows us to estimate the hawkishness of elite decision-makers at a distance, without inferring hawkishness from the behavioral outcomes we are using it to explain.

2.2.1 Coding Biographical Characteristics of US Decision-Makers

Ideally, we could administer surveys to presidents and advisers who participated in US foreign policy meetings during the Cold War. Since this is impossible, we adopt a biographical approach, building on work using policymakers' background characteristics as a proxy for their unobservable traits (Rathbun, 2014; Fuhrmann, 2020; Horowitz, Stam and Ellis, 2015; Kertzer, 2016; Carter and Smith, 2020).

To start, we identify all presidents and advisers who spoke at least once during the meetings in our sample. Each segmented speech act discussed above is attributed to one unique speaker. We collected information on the backgrounds and careers of these speakers, ranging from cabinet secretaries, to senior bureaucratic officials (e.g., assistant and deputy secretaries), to mid-level bureaucrats in the State Department, Pentagon, Central Intelligence Agency, and other government agencies.[7] Our codings focused on two aspects of the individual's background. First, we recorded their position and the dates on which the position was held. Second, we gathered an array of demographic variables that might affect hawkishness, including gender, birth year, education level, and political party, as well as diplomatic, intelligence, or military experience. Following Horowitz, Stam and Ellis (2015), we also coded combat experience by identifying deployment to a combat

[7] See Appendix §4.1 for coding procedures.

theater during a war involving the US.[8] Table A5 in Appendix §4.1 provides the coding for Henry Kissinger.

2.2.2 Imputing Decision-Maker Hawkishness with Elite Surveys from the Cold War

While we hope our biographical dataset of American foreign policy decision-makers will be useful to researchers in and of itself, our methodological innovation is to use machine learning approaches to measure adviser hawkishness at a distance. Like Carter and Smith (2020), we incorporate machine learning approaches on biographical codings. Our novel contribution is to anchor our measures using information from Cold War-era surveys of American foreign policy elites administered through the Foreign Policy Leadership Project (FPLP) (Holsti and Rosenau, 1984). Crucially, FPLP surveys included a battery of items measuring respondents' levels of militant internationalism, a standard measure for hawkishness in the public opinion literature (e.g., Herrmann, Tetlock and Visser, 1999; Holsti, 2004), as well as demographic questions that mirror those coded in our biographical dataset.

We use this overlap to estimate the hawkishness of meeting participants in our sample based upon the hawkishness of survey respondents in the FPLP with similar biographical characteristics. Our measurement strategy proceeds in three steps. First, we create a measure of hawkishness averaging across a 15-item battery of FPLP questions that tap into respondents' views towards containing communism using force, prioritizing offensive military action over diplomacy or defensive measures, believing that the American effort in Vietnam was too limited, and so on.[9] Second, we identify the individual-level characteristics that exist in both the FPLP and our biographical data. These include *gender*, *birth decade*, *level of education*, *military experience*, *combat experience*, *diplomatic experience*, *current military officer*, *current foreign service officer*, and *political party*.[10] Third, we harness a series of supervised learning models to adduce the relationship between a respondent's biographical characteristics and their hawkishness in the FPLP. We adopt a boosted linear regression model—a form of ensemble learning where many simple linear models are

[8] We are able to match a speaker name and position for 102,720 speech acts out of a total of 104,504 in our dataset – over 98%. Several hundred of these missing values, however, are explicitly listed as "Unidentified" in the record and are therefore impossible to match. Our effective match rate is thus over 99%.

[9] See Appendix §4.2 for details.

[10] As a robustness check in Appendix §5.5, we remove individuals' current affiliation from our statistical models; our results remain the same.

sequentially trained and reweighted until a final model is established—as our primary method of estimating participant hawkishness. The models are fed the FPLP data, which provides explicit information the model can process to understand the relationship between biographical characteristics and individuals' level of hawkishness according to their survey responses. In order to tune the hyperparameters of the boosted model, a five-fold cross-validation process is used to select the model that produces the best out-of-sample predictions. This optimal model is then used to predict hawkishness on new data, which in our case is the full set of presidents and advisers who participated during meetings in our sample. We use a bootstrapping process through which we randomly re-sample the FPLP survey data with replacement 1,000 times. This generates 1,000 predicted hawkishness measures for each individual, and we use the average as an actor's measure of hawkishness.[11]

In these models, one important predictor of hawkishness is an individual's party affiliation. However, we know that the Republican and Democratic parties changed their stances in foreign policy issues during the early Cold War (Fordham, 2007). Democrats went from being more hawkish to dovish, while Republicans did the opposite, leading to partisan positions that are more broadly familiar to us today. If we ignored this shift, our measures would underestimate the hawkishness of Truman-era Democrats and overestimate the hawkishness of Eisenhower-era Republicans. To address this issue, we make temporal adjustments to our hawkishness measure that compensate for the shift in party platforms, using longitudinal measures of partisan hawkishness assembled by Jeong (2018) to make time-conditional adjustments to the estimated coefficients for hawkishness of senior meeting participants.[12] This adjustment produces hawkishness measures that align more closely with historical assessments.

Figure 3 illustrates the predicted hawkishness measures for six senior positions in US foreign policymaking—the president, secretary of state, secretary of defense, the Central Intelligence Agency (CIA) director, Chairman of the Joint Chiefs of Staff (the president's senior military adviser), and the national security advisor—sorted in chronological order. Chairmen of the Joint

[11] See Appendix §4.3 for more information on model selection and performance. Appendix §4.4 compares the performance of the boosted linear model with a simple linear regression. In Appendices §5.3, §5.8.1, and §8.4, we show that our main results hold using a more computationally intensive process where we propagate the uncertainty from our predicted measures of hawkishness through our statistical analyses. Appendix §5.4 demonstrates that our findings are unchanged by using hawkishness measures produced with a simpler linear model.

[12] See Appendix §4.6.

Figure 3: Predicted Hawkishness Measures for Senior US Decision-Makers

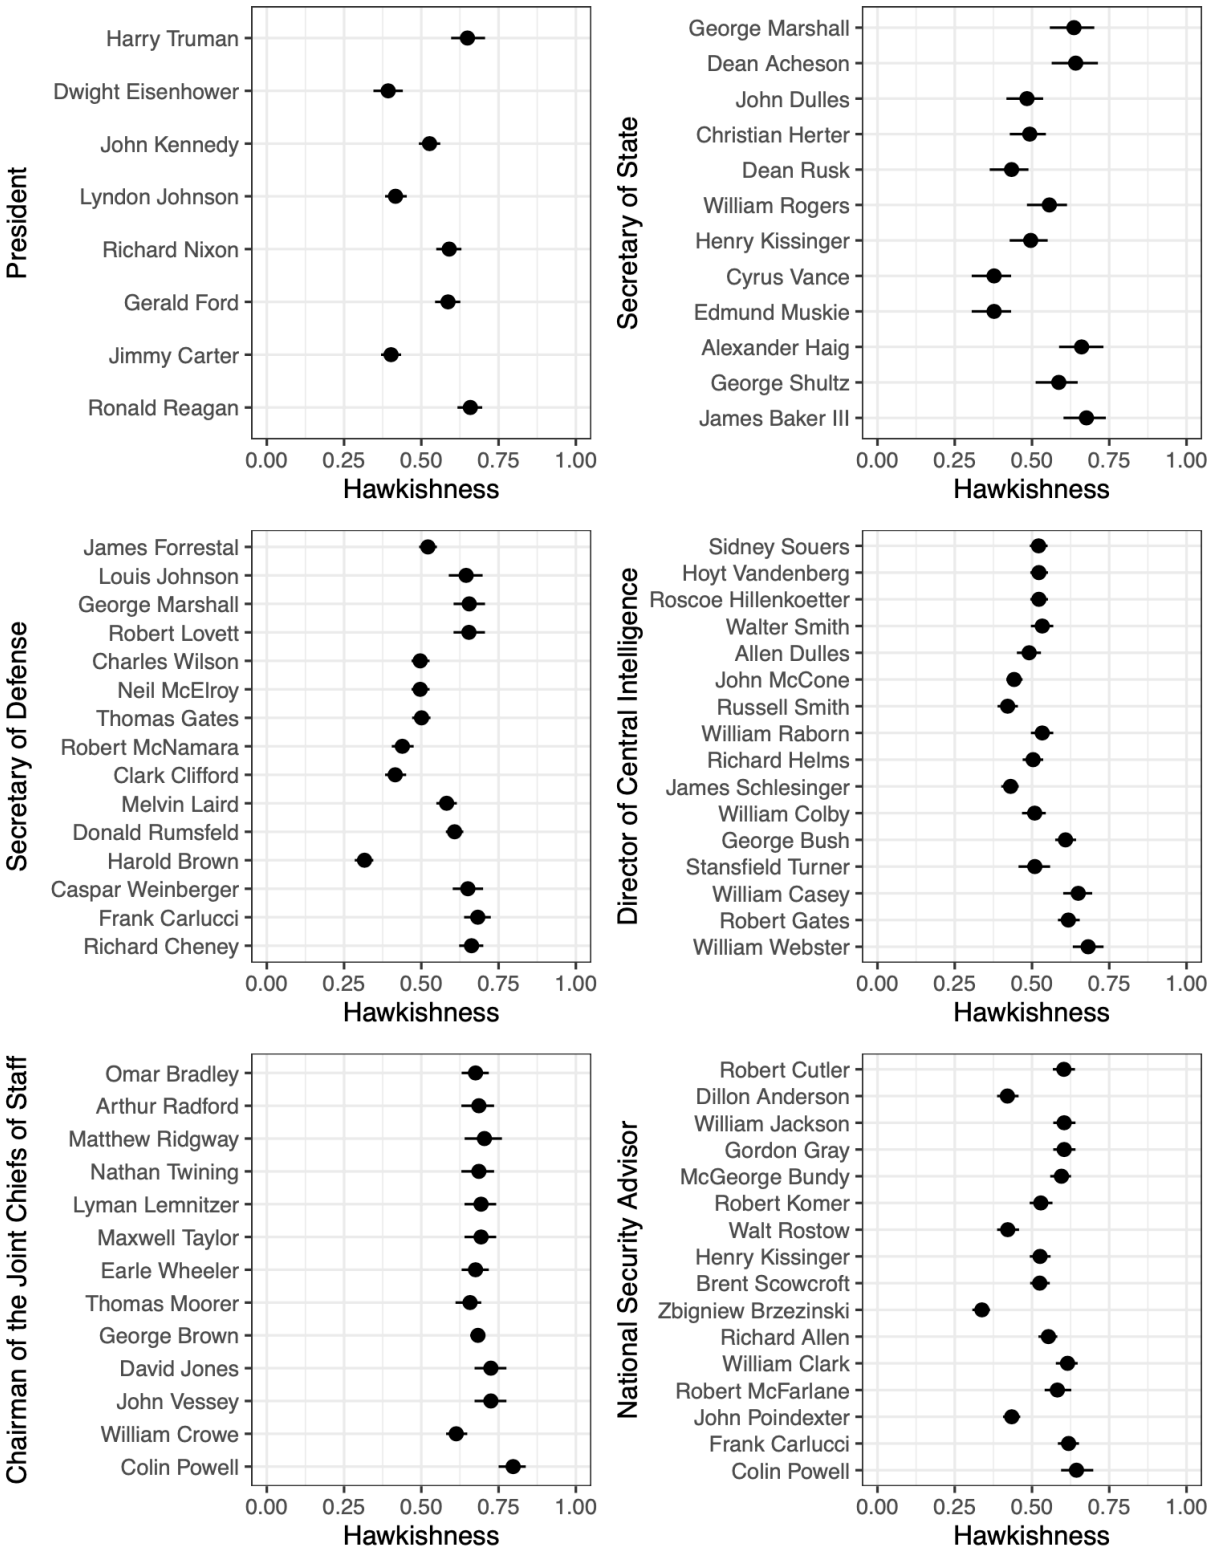

Note: Black dots report mean hawkishness value across $B = 1000$ bootstraps, with 95% confidence intervals

Figure 4: Average Speaker Hawkishness in US Foreign Policy Meetings

Chiefs are generally more hawkish than secretaries of state, but crucially, some secretaries of state (like Alexander Haig) are more hawkish than others (like Cyrus Vance). Figure 4 displays group hawkishness at the meeting level over time by calculating the average hawkishness of all meeting participants.

3 Results I: Testing the Microfoundations of the Adviser Model

We first take advantage of our rich deliberation data to validate the adviser model's microfoundations. The analysis in this section answers two questions. First, do leaders seek counsel from advisers in foreign policy deliberations? And second, does the nature of the counsel that advisers provide depend on their foreign policy dispositions?

3.1 Leaders Seek Counsel during Deliberations

The first assumption of our adviser model is that leaders seek counsel from their advisers. The 2,685 meeting records in our data demonstrate that leaders routinely met with advisers. Our adviser model assumes, however, that these meetings should be genuinely deliberative: leaders should routinely seek counsel from their advisers, rather than merely informing them of decisions that have already been made. Deliberation also requires dissent: advisers should be willing to

Figure 5: Leader Search for Counsel during Meetings

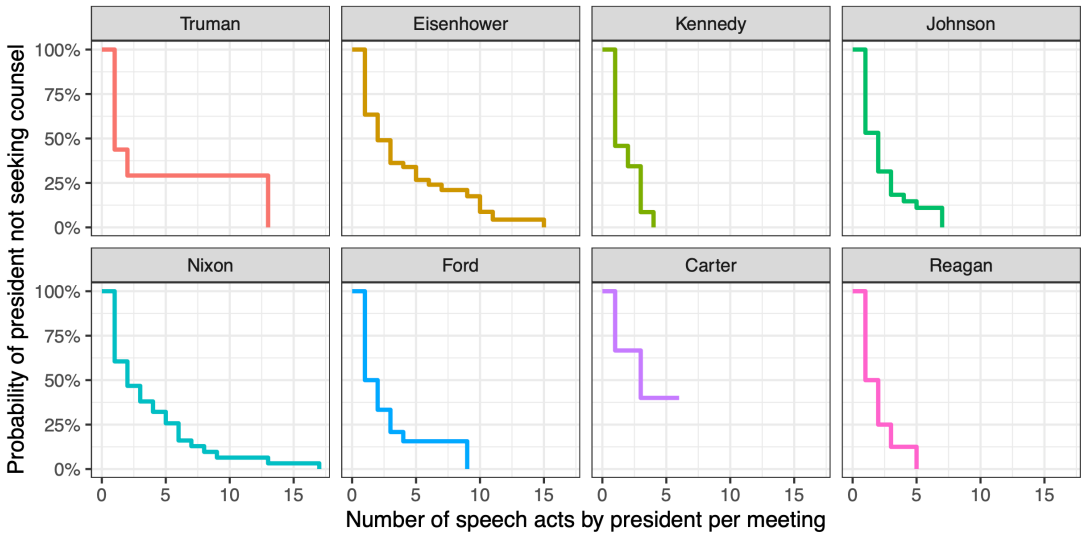

Note: Each panel depicts a survival curve reflecting likelihood of a president seeking counsel as a function of speech acts by the president. The sharper the decline, the *faster* a president sought counsel. See Appendix §8.1 for coding details.

express opinions even when others disagree.

To evaluate whether these assumptions hold, we examine the timing and frequency of such deliberation in meetings. Since automated methods are unlikely to capture the subtleties of deliberation we wish to identify, we generated a stratified random sample of 270 formal and informal meetings across presidential administrations (approximately 10% of the full sample), for which we coded our concepts of interest manually. Drawing upon the study of deliberation elsewhere in political science, we developed a coding scheme (detailed in Appendix §8.1) to identify speech acts that exhibited *seeking counsel*. A research assistant reviewed and coded each of the 12,006 speech acts in the random sample. Seeking counsel was defined as a speaker requesting another participant to introduce new information, ideas, or recommendations into the discussion. Simply stating one's own position does not qualify as seeking counsel. Rather, speakers must have proactively asked *others* for their perspective. A meeting in which participants reiterated their own position over and over again in slightly different ways would be coded as having no search for counsel.

We find that US foreign policy deliberations exhibited a high level of seeking counsel, particularly by leaders. About one in three presidential speech acts—and over one in ten adviser speech acts—queried for more information from advisers. Figure 5 formalizes this intuition through a simple Cox model, in which administration is regressed on the duration of time before the leader seeks

counsel. The model highlights that even the least inquisitive presidents still quickly sought counsel in their deliberations. While some of these queries might be performative, it is clear that presidents expend considerable time and effort soliciting input from advisers. This finding is difficult to reconcile with the leader model's emphasis on fixed and unchanging leader beliefs.

One question, however, is whether these queries simply led to the identical viewpoints being expressed ad nauseam. If leaders manipulated deliberations to ensure that only pre-approved viewpoints would be voiced, as the leader model contends, we would expect little disagreement among participants. To explore this contention, we replicated our coding process for speech acts exhibiting *dissent*, defined as a textual indication that the speaker disagreed with another meeting participant.[13] The data again suggest that leaders and advisers use deliberations to offer conflicting views: 14% of adviser speech acts and 10% of president speech acts in our random sample featured a dissenting opinion, and 64% of meetings featured some form of debate between participants. Appendix §8.6 uses these data to show that the probability of meeting attendees speaking is not a function of their dispositional distance from the president. The prevalence of debate suggests that, on average, leaders do not consistently manipulate deliberations to ensure conformity with their own worldview. We also find that expressions of dissent more frequently come from advisers who *differ* substantially in hawkishness relative to the president, which helps to allay concerns that self-censorship dominated the meetings. Collectively, the findings suggest that leaders seek input from advisers, and that participation allows advisers the opportunity to deliberate.

3.2 Advisers Provide Counsel Congruent with their Dispositions

The second assumption of our adviser model is that the counsel that advisers supply during deliberations depends on their predispositions. As noted in our theory section, we can think of counsel as consisting of information provision, problem representation (or analysis), and policy recommendations. To get at information provision and analysis, we use the rich textual data contained within the collected records, examining the content of speech acts by hawks and doves during the meetings.[14] Drawing on existing research on hawkishness, we identified ex ante five concep-

[13] In some cases, advisers directly identified their dissent (e.g., "disagreed" or "objected"), while in other cases dissent could only be discerned in the context of the meeting. This further motivated our choice to manually code the concept on a subsample of meetings rather than using automated methods on our whole dataset.

[14] We group information provision and problem representation together given difficulties in empirically disentangling them using automated approaches.

tual categories summarizing the types of considerations that hawks and doves should stress during deliberations. First, hawks should emphasize that using violence is an effective and appropriate strategy in international politics (George, 1969; Herrmann, Tetlock and Visser, 1999; Weeks, 2014). Second, hawks should be more likely than doves to emphasize the ubiquity of threats or other "competitive elements" between states (Russett, 1990, p. 516). Third, both hawks and doves may be attuned to the military balance of capabilities, albeit for different reasons: hawks might emphasize the importance of primacy in material strength and power (Herrmann, Tetlock and Visser, 1999), while doves might instead emphasize that the balance of power limits the potential payoffs to violence. Fourth, doves should tend to ascribe greater promise to diplomacy and the need for international cooperation (Russett, 1990). Finally, doves should tend to emphasize the importance of viewing international disputes from the adversary's perspective, recognizing that an adversary may face international or domestic constraints that impede a negotiated settlement (Brutger and Kertzer, 2018). The first three columns of Table 1 summarize the concepts that existing literature suggests hawks and doves should emphasize.

To test whether hawkish and dovish advisers exhibit different speech patterns on these topics, we employ a straightforward dictionary approach. We specify a set of nine to fourteen keywords, some of which are listed in Table 1, that capture each of our conceptual topics.[15] Using this list, we calculate the proportion of words associated with an individual adviser in a specific meeting that overlaps with these keywords (if at least 50 words in total). A total of 11,609 adviser-meeting observations, representing 100,089 speech acts (96% of our data), are processed in this manner.

Several examples suggest that these proportion measures effectively identify text related to our concepts of interest. For example, the highest scoring text for diplomacy comes from a meeting in June 1976, where Chairman of the Council of Economic Advisers Alan Greenspan reported on a recent Puerto Rico economic summit: "It was an extraordinary meeting, especially in the context of other meetings I have attended. There was a real intellectual grappling with major philosophical issues... We may have developed a new form of international institution. We have broken down the formality and protocol of summit meetings so that true dialogue can take place." The text most indicative of military violence comes an October 1958 meeting, in which Chief of Staff of the Air Force Thomas White noted that "our problem was that we must assume that the Soviets will strike

[15] Appendix §8.2 features full lists of terms for each topic, and complete statistical results.

Table 1: Sample of Hawkish and Dovish Terms

Topic	Who Emphasizes?	Source	Distinctive Terms
Military violence	hawks	George (1969); Herrmann, Tetlock and Visser (1999)	attack, bomb, escalate, invade, strike
International threats	hawks	Russett (1990)	hostile, threat, war, aggression, enemy
Military balance	hawks and doves	Herrmann, Tetlock and Visser (1999)	capability, missile, nuclear, silo, tank
Diplomacy	doves	Russett (1990)	diplomatic, negotiate, summit, talk, treaty
Adversary interests	doves	Brutger and Kertzer (2018)	china, communist, soviet, ussr, vietnam

first. If they do we cannot stop them by our Distant Early Warning lines. We must therefore find the number of bombers which it is logical for us to maintain in order to strike back after the initial Soviet attack."

We find that hawks and doves discuss systematically different considerations, consistent with their underlying predispositions. Figure 6 plots the effect of moving from the least to most hawkish speaker within a single administration on expected topic proportions in a speech act. The plotted effects are based on OLS specifications that leverage the hawkishness scores described above, while also including administration fixed effects and controls for whether the meeting was formal or whether the meeting record was a transcript. As shown, hawks and doves differ in the topics they emphasize. Consistent with theoretical expectations, hawks are more likely to address issues related to violence and threat, while doves place greater stress on diplomatic possibilities and adversary interests. Hawks' discussions of military balance appear to overwhelm those of doves. Beyond statistical significance, these differences are substantively meaningful, despite the seemingly small magnitude of the estimated coefficients. An average adviser-meeting observation has a *diplomacy* proportion of 0.0045. Moving from the least to most hawkish individual reduces the expected proportion by 0.0034, or about an 76% decrease.[16]

To test whether predispositions also shape advisers' recommendations, we drew another random

[16]In Appendix §8.5, we replicate this analysis using a keyword-assisted topic model (keyATM).

Figure 6: Effect of Adviser Hawkishness on Topic Proportions during Meetings

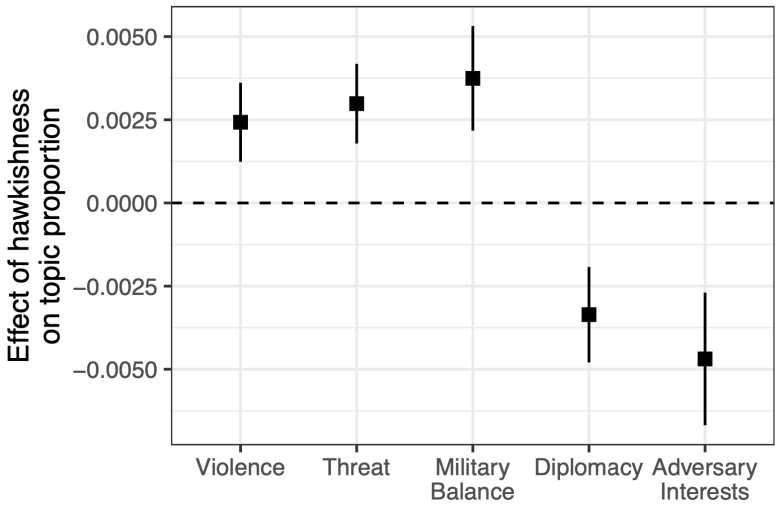

Note: Bands represent 95% confidence intervals

sample of 425 meetings (approximately 16% of all meetings in our dataset) and manually coded the 20,451 speech acts in them for whether each speaker was calling for a highly hawkish policy: the use of coercive force against, or conflict with, an adversary. The analysis confirms the straightforward intuition that hawks tend to be more likely to make conflictual policy recommendations than doves do. Individuals in the 80th percentile of hawkishness are 72% more likely to recommend conflictual policies than individuals in the 20th percentile of hawkishness are. Advisers may have reasons to go "against type" to improve their persuasive power—as hawks advocating against conflict may have more appeal than doves (Saunders, 2018; Mattes and Weeks, 2019)—but on average they do not.

In sum, analyzing tens of thousands of speech acts in foreign policy meetings offers evidence consistent with the microfoundations of the adviser model. Deliberations give advisers the opportunity to offer counsel, and the considerations and recommendations advisers raise depend on their disposition. Hawkish advisers raise considerations emphasizing military violence, while dovish advisers make arguments emphasizing the utility of diplomacy and adversary perspectives. These dispositional differences also hold in terms of the policy recommendations for which hawks and doves advocate. Next, we turn to the question of foreign policy decisions themselves.

4 Results II: Testing the Three Models of Aggregation

The above analysis shows how adviser dispositions affect the counsel they provide, but does not test how these dispositions aggregate to affect decision-making. We now turn to our entire meeting record data to address this question. In order to fully test our adviser model—and compare it to the emergence and leader models—we introduce an additional dataset, which studies the decisions made in each of our deliberation meetings, which we use for our main analysis.

4.1 Outcome Variable: Conflictual Decisions

Given our interest in how hawkishness as a leader- and adviser-level disposition aggregates in foreign policy decision-making, our central outcome of interest concerns policy choices aimed toward US adversaries in each meeting. To construct the outcome variable, a team of research assistants identified and classified all substantive decisions made within these meetings—thereby avoiding the truncation bias implicated by studies of decision-making in IR that focus only on major uses of force (Mitchell and Moore, 2002). To qualify as substantive, a decision must have presidential approval and plausibly observable ramifications for US foreign policy. Examples of substantive decisions include authorizing an increase in military spending, accelerating arms testing, moving military personnel or assets, altering strategic priorities, pushing for diplomatic engagement, and crafting language for public statements. Decisions that would not qualify are those that merely note the policy preferences that meeting participants adopted, call for additional study of a topic, or establish a committee to set policy in the future.

For each substantive decision, a team of coders collected contextual background information, classified the decision into one of several categories, and specified a target of the decision.[17] The pertinent classification categories were *conflictual acts* which could be verbal or material and span from making threats to using force, and *cooperative acts* which could similarly be verbal or material and span from conveying agreement to providing aid.[18] Targets of the actions were the state or political organization, such as a rebel group, most directly affected by the decision.[19] Since the

[17] When the text of the decision itself proved insufficient, as it often did, coders used the full meeting record and contemporaneous policy papers to clarify the decision's context, nature, and target.

[18] Our categories are similar in spirit to the event type categorizations in Goldstein (1992). See Appendix §3.

[19] Given the Cold War context in which these meetings took place, the "Soviet Bloc" typically served as the target for military spending adjustments, or those strategy changes without an explicitly identified target.

Table 2: Decisions Against Adversaries by Administration

Administration	Decisions	Conflictual	Cooperative
Truman	102	88	14
Eisenhower	417	329	88
Kennedy	124	100	24
Johnson	951	65	26
Nixon	75	43	32
Ford	56	20	36
Carter	22	12	10
Reagan	63	45	18
Total	950	702	248

effects of hawkishness are linked to the treatment of adversaries in particular, rather than allies or neutral entities, the analysis focuses on *adversary* targets, with an entity's classification in this category potentially varying across time depending on the state of bilateral relations.

As Table 2 shows, our sample yielded 950 decisions towards adversaries made across formal and informal meetings, 702 of which are conflictual decisions and 248 of which are cooperative.[20] We use these data to produce two measures. The first is a meeting's raw number of conflictual decisions toward adversaries. Because this is a count variable, the corresponding analyses use Poisson regressions.[21] The second measure accounts for both conflictual and cooperative decisions by subtracting the latter from the former. Positive values indicate a meeting that produces more conflictual decisions than cooperative ones. We use ordinary least squares regressions to analyze this variable. Distributions of these two outcomes are reported in Appendix §3.3.

4.1.1 Control Variables

One challenge with studying the effects of adviser hawkishness on foreign policy decision-making is that adviser participation is not randomly assigned. To address these potential selection effects, we employ a two-pronged empirical strategy, beginning with a set of control variables meant to address potential confounding factors for our main analysis, and then proceeding to a more thorough set of robustness tests in Section 4.1.3 below.

One threat to inference is that individuals appointed to the advisory team might reflect the

[20] This does not include other cooperative decisions the United States made towards allies or neutral countries, which intuitively constitute the bulk of US cooperation during the Cold War.

[21] Our results also hold using negative binomial models reported in Appendix §5.7.

leader's preferences regarding the use of force. To address this issue, we include *administration fixed effects* across several specifications to account for unobserved invariant components of each administration, particularly those that may have prompted the leader to choose a hawkish advisory team – or use advisers and advisory institutions in systematically different ways (Hermann and Preston, 1994; Jost, 2023). Models with fixed effects study the effect of group composition while holding the leader constant, which controls for these predilections. We further probe the question of adviser appointment below.

Another threat to inference concerns which advisers are invited to participate in which meetings. Although which advisers attend formal meeting is partly routinized, imagine a topic on which the leader is inclined to authorize conflictual decisions. Based on that inclination, the leader could skew meeting invitations toward hawkish advisers. Leader inclination would confound the relationship of interest because it is a common cause of group hawkishness (the explanatory variable for the adviser hypothesis) and policy decisions (the outcome variable). To address this selection process in our models, we manually coded the *agenda topic* for each of the 104,504 speech acts in our sample.[22] Using these classifications as a control variable helps minimize potential bias in the meeting's invitees.

A third potential concern is that advisers may be predisposed to participate when they anticipate positive reactions to their worldview. For instance, hawkish advisers may speak more when the international environment (e.g. recently being attacked) makes the state predisposed to pursue conflictual policies. While this concern is allayed somewhat by the self-censorship and dissent results presented above, we also include a set of system- and leader-level control variables that may have motivated hawkish advisers to speak more. Following existing research, these include: a variable measuring the *lagged number of militarized interstate disputes* challenging the US in the last five years, as well as the *national capabilities* to measure military strength and economic health.[23]

Additional control variables track the number of meeting participants affiliated with the *Defense Department*, and *intelligence agencies*, *military*, and *State Department*, as bureaucratic interests

[22] See Appendix §9 for more information on agenda items and their impact on group composition.

[23] These controls follow those in Horowitz, Stam and Ellis (2015), dropping domestic characteristics such as regime type and polity score that do not vary for the United States in our time period. We omit characteristics such as a leader's age, time in office, and prior military experience because they were attributes used to develop predicted levels of hawkishness for each president.

may skew hawkish or dovish (Allison and Zelikow, 1999) – as well as the (logged) total years of *military*, *diplomatic*, and *intelligence* experience of meeting participants. Another control captures the *number of attendees* in each meeting (LeVeck and Narang, 2017) and a binary variable indicating whether a meeting was a *formal* meeting of the NSC, as opposed to an informal session outside it.

4.1.2 Results

The emergence model predicts that dispositions of group members should have no systematic effect on policy decisions, whether because structural characteristics of the international system dominate or because group decisions cannot be reduced to individual-level traits. To test this model, we begin with the simplest aggregation procedure: the mean level of hawkishness of all speakers in the meeting.[24] If consequential policy choices are not reducible to the traits of individuals involved in making those choices, as the emergence model suggests, then we ought to observe no relationship between the group's average hawkishness and its policy decisions.

Inconsistent with emergence models, we find that a rise in group hawkishness increases conflictual policy choices toward adversaries. The relationship holds across different specifications, as Models 1 and 2 of Table 3 show. Regardless of outcome variable specification or the presence of control variables, the group's average hawkishness consistently has a positive effect on conflictual policy decisions.[25] The left-hand panel of Figure 7 presents the results graphically. Shifting the group's mean hawkishness from its minimum to maximum values while holding other variables constant more than quintiples the predicted number of conflictual decisions (based on Model 1).

The results demonstrate a clear, consistent, and substantively meaningful relationship between a group's composition and its decisions. Moreover, including a measure of aggregated group traits (mean hawkishness) improves the statistical model fit compared to a sparser specification without this measure. A likelihood ratio test that compares Model 1 of Table 3 with a null model that omits the measure of mean group hawkishness strongly indicates that accounting for hawkishness yields a significant improvement ($p = 0.0004$). A similar exercise using Model 2 produces an analogous result ($p = 0.038$). The findings thus do not support the emergence model.

The leader model predicts that leader hawkishness should systematically affect policy decisions.

[24] Meeting participant here refers to advisers who attend the meeting and speak at least once. We adopt this definition for practical reasons: the full list of attendees is not available for all meetings.

[25] Full results, including bivariate specifications, are provided in Appendix §5.1.

Table 3: Effect of Participant Hawkishness on Foreign Policy Decisions

	Emergence Model		*Leader Model*		*Adviser Model*		*Advisers + Admin. FEs*	
	Conf.	Conf. − Coop.	Conf.	Conf. − Coop.	Conf.	Conf. − Coop.	Conf.	Conf. − Coop.
	Poisson	*OLS*	*Poisson*	*OLS*	*Poisson*	*OLS*	*Poisson*	*OLS*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean Hawkishness	5.098***	0.913**						
	(1.449)	(0.443)						
Advisers' Hawkishness (Acts)					2.995***	0.705**	3.136***	0.687**
					(0.951)	(0.287)	(1.065)	(0.316)
President's Hawkishness			0.240	−0.018	−0.672	−0.127		
			(0.523)	(0.178)	(0.599)	(0.223)		
No. of Attendees	0.011	0.022***	0.033*	0.025***	0.040**	0.027***	0.018	0.024***
	(0.021)	(0.008)	(0.020)	(0.007)	(0.020)	(0.008)	(0.021)	(0.008)
Defense	0.197***	0.072***	0.118***	0.059***	0.139***	0.054**	0.185***	0.071***
	(0.048)	(0.022)	(0.044)	(0.019)	(0.044)	(0.022)	(0.048)	(0.022)
Intelligence	0.293***	0.106***	0.253***	0.106***	0.282***	0.109***	0.289***	0.109***
	(0.097)	(0.039)	(0.091)	(0.034)	(0.093)	(0.039)	(0.097)	(0.039)
Military	−0.010	0.058*	0.055	0.067***	0.030	0.063**	0.043	0.067**
	(0.064)	(0.030)	(0.055)	(0.023)	(0.056)	(0.027)	(0.060)	(0.028)
State	0.043	0.0001	0.001	0.003	0.039	0.006	0.028	0.0004
	(0.054)	(0.021)	(0.051)	(0.019)	(0.052)	(0.021)	(0.054)	(0.022)
Diplomatic Experience	0.096**	0.005	0.079*	0.007	0.064	−0.010	0.097**	0.004
	(0.048)	(0.016)	(0.045)	(0.014)	(0.046)	(0.016)	(0.048)	(0.016)
Intelligence Experience	−0.138**	−0.060***	−0.115**	−0.067***	−0.135***	−0.056***	−0.131**	−0.059***
	(0.054)	(0.020)	(0.047)	(0.017)	(0.049)	(0.019)	(0.054)	(0.020)
Military Experience	0.163*	−0.054**	0.235***	−0.009	0.142*	−0.053**	0.110	−0.067***
	(0.084)	(0.025)	(0.077)	(0.019)	(0.083)	(0.024)	(0.088)	(0.026)
5-Year MID Challenges	−0.386*	−0.100	0.095	−0.039	0.041	0.013	−0.371	−0.107
	(0.231)	(0.078)	(0.147)	(0.040)	(0.149)	(0.045)	(0.232)	(0.079)
US CINC	5.508**	2.898***	2.998***	0.923***	2.452**	1.255***	5.849**	2.960***
	(2.348)	(0.916)	(1.064)	(0.321)	(1.088)	(0.367)	(2.356)	(0.929)
Formal	0.656***	0.115**	0.799***	0.181***	0.793***	0.130**	0.692***	0.122**
	(0.136)	(0.051)	(0.130)	(0.045)	(0.132)	(0.051)	(0.137)	(0.052)
Constant	−7.098***	−1.231***	−4.458***	−0.240	−5.058***	−0.609***	−5.963***	−1.069***
	(1.199)	(0.407)	(0.645)	(0.188)	(0.684)	(0.219)	(1.075)	(0.369)
Administration FEs	✓	✓					✓	✓
Agenda Items	✓	✓	✓	✓	✓	✓	✓	✓
Observations	2,685	2,685	2,685	2,685	2,650	2,650	2,650	2,650

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Figure 7: Summary of Three Models of Trait Aggregation

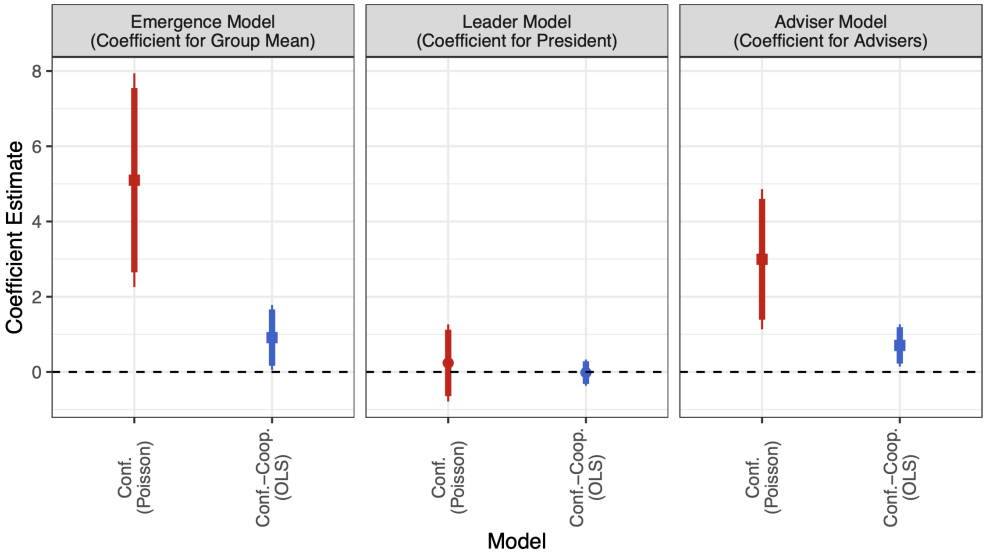

Note: Coefficient plots corresponding to main findings, corresponding to Models 1-6 in Table 3. Thicker and smaller bands represent 90% and 95% confidence intervals respectively. Square points indicate 95% statistical significance.

Whether due to group hierarchy or a leader's reluctance to hear other views, we should observe a positive relationship between leader hawkishness and conflictual policy choices.[26] Results presented for Models 3 and 4 in Table 3 and the middle panel of Figure 7 suggest that presidential hawkishness does not have a systematic relationship with conflictual decisions toward adversaries.[27] Appendix §7 details a range of potential explanations for why leader-level hawkishness may not be positively associated with conflictual decisions, ranging from strategic interaction to the uniquely institutionalized context of the US. Moreover, replications and extensions of Horowitz, Stam and Ellis (2015) and Carter and Smith (2020) in Appendix §7.2 similarly find little evidence that presidential hawkishness predicts American conflict behavior, consistent with Yarhi-Milo's (2018) recent work. We also obtain similar results when measuring presidents' hawkishness using measures derived from codings provided by experts of American foreign policy (Appendix §7.1), as well as using a different operationalization of the dependent variable (Table 5 below). As with most studies of presidential decision-making, a sample of eight leaders limits the conclusiveness of these results, as the idiosyncratic nature of a single president could be exerting an outsized effect on the results. Within these confines, however, our analysis offers little support for the leader model.

[26] Because each president's hawkishness in our data does not appreciably change over time, we omit administration fixed effects.

[27] See Appendix §7.4 for robustness checks.

The adviser model predicts that the hawkishness of advisers affects decision outcomes in deliberations. If advisers exert influence through the counsel they provide, then meetings in which hawkish advisers speak frequently ought to produce more conflictual decisions. For these models, we calculate a weighted average of adviser hawkishness, where each adviser's weight is a function of the proportion of speech acts they contributed to the discussion, reflecting our emphasis on communication as a vehicle for influence.[28]

As shown by Models 5 and 6 in Table 3 and the right-hand panel of Figure 7, we find strong evidence that adviser hawkishness affects decision outcomes. Meetings in which more hawkish advisers speak more tend to adopt more conflictual policies toward an adversary. This pattern appears across specifications. Holding other variables fixed, shifting the group's weighted hawkishness from the minimum to maximum values more than triples the expected number of conflictual decisions (based on Model 5's specification). Models 7 and 8 drop the leader measure so that we can include fixed effects to guard against the possibility that results reflect differences between different presidents' management style and preference for formal (i.e. NSC) or informal meetings.

To further contextualize the substantive effects, Table 4 presents the predicted number of conflictual decisions toward an adversary for all fully-specified models. These calculations shift each relevant measure of hawkishness from its minimum to maximum value while holding other variables fixed at their means, while also presenting the substantive effects of other contextual variables to provide a relative benchmark. The table shows the dramatic effect of both the mean and weighted mean measures of group hawkishness, which respectively cast doubt on the emergence model and provides evidence consistent with the adviser model. In the leader versus adviser models, it is worth noting that even though the president and advisers have cross-cutting effects on decision-making, the more conflictual nature of hawkish advisers appears to outweigh the effects of the president, lending further support to importance of advisers. Collectively, the results suggest that the leader model is incomplete, and that we must consider the dispositions of advisers in the room.

[28] Consequently, meetings with identical lists of participants can produce divergent hawkishness levels. To estimate the independent effect of advisers, the weighted average score excludes the president, and some specifications control for the president's hawkishness (and thus omit administration fixed effects) to assure that advisers exert influence rather than merely serving as proxies for the president.

Table 4: Predicted Number of Conflictual Decisions Towards Adversaries

Model	Model	Variable	Min.	Max.
Emergence	Poisson (Model 1)	Mean Hawkishness	0.093	0.610
	OLS (Model 2)	Mean Hawkishness	0.021	0.357
Leader	Poisson (Model 3)	President's Hawkishness	0.264	0.282
	OLS (Model 4)	President's Hawkishness	0.336	0.331
Adviser	Poisson (Model 5)	Advisers' Hawkishness (Acts)	0.137	0.425
		5-Year MID Challenges	0.236	0.256
		US CINC	0.203	0.323
	OLS (Model 6)	Advisers' Hawkishness (Acts)	0.043	0.309
		5-Year MID Challenges	0.167	0.191
		US CINC	0.079	0.317

4.1.3 Selection and Robustness

Three additional questions regarding selection effects merit consideration. First, above we noted that administration fixed effects help identify the effect of adviser dispositions *within* each administration by holding constant unobserved variables, such as leader-level differences in advisory arrangements. Yet one potential question is whether this methodological choice masks effects that leaders exert through the appointment process. Leaders make appointments for a variety of reasons, including appointee qualifications, personal connections, and public approval. If leaders only appointed advisers who shared their foreign policy worldview (i.e. hawkish presidents only appointed hawks), then the advisory environment—and our results—would simply represent an extension of the leader's disposition. We find no evidence to support this conclusion. Mixed-effect models that include administration random effects find that the intraclass correlation ranges between 0.039 and 0.179: there is approximately 5.6 to 25 times greater variation in hawkishness *within* individual administrations than there is between them. We would expect a far lower figure if hawkish leaders simply hired or invited hawkish advisers into meetings.

One reason why advisers are not simply dispositional mimeographs of the leaders they serve is that adviser appointment can be affected by multiple considerations, such as a candidate's education, experience, qualifications, personal connections, and fit for the position (Jost and Kertzer, 2023)—not just their hawkishness. Moreover, many leaders prefer viewpoint diversity in their advisory group either to improve the quality of foreign policy debates by considering multiple per-

spectives (Mintz and Wayne, 2016) or to bolster their domestic credibility (Saunders, 2018). For every example of hawkish leaders like Ronald Reagan selecting dispositionally similar advisers, like Caspar Weinberger, there are also examples of leaders like Barack Obama selecting dissimilar advisers, like Hillary Clinton. We find further evidence in support of this contention in Appendix §6, where we build a dataset of "runner ups" considered for high-level national security positions since the Truman administration, showing that individuals ultimately appointed were not systematically closer in hawkishness to the president than those who were considered but not selected.

Second, even if leaders do not merely select mimeographs as their advisers, leaders may still have the ability to decide when these meetings take place and which advisers are invited (Krasner, 1972), in ways that controlling for meeting agendas might fail to capture. To address this concern, we replicate our results using an alternative model specification that ignores our meeting data altogether, instead examining the effect of adviser-level hawkishness (limited to NSC principals) on the United States' propensity for being involved in militarized interstate disputes in a given month. Our findings, shown in Table 5 below, remain the same, despite a different unit of analysis (the time-unit, rather than the meeting-level), and a more restrictive dependent variable (militarized interstate disputes, rather than all foreign policy decisions). This strongly suggests that, even if leaders attempt to manipulate the advisory group or fabricate deliberation in ways that accord with their worldview, advisers are still able to sway foreign policy decisions in aggregate. Moreover, Appendix §5.9 replicates our meeting-level analysis for formal gatherings, only using characteristics of NSC principals, who are obligated to have a presence at every meeting. The results remain consistent. Adviser predispositions appear to be significantly associated with interstate conflict, even controlling for a host of international, domestic, and leader-level variables.

Additionally, one might wonder whether our conceptualization of adviser hawkishness was specific to the Cold War—perhaps making our findings an artifact of the highly competitive US-Soviet relationship. Two factors discredit such an interpretation. First, the militant internationalism measure we use to impute decision-maker hawkishness has been widely used since 1991 (e.g., Rathbun, 2007; Brutger and Kertzer, 2018). In fact, Murray (2002) shows that hawkish beliefs among American decision-makers were surprisingly consistent before and after the Cold War. Second, in Appendix §5.6, we run a robustness check where we drop decisions involving the Soviet Union from our analysis and find that results remain broadly consistent.

Table 5: Effect of NSC Principals' Hawkishness on MIDs, Using Monthly Data

	Dependent variable:	
	Onset of MIDs involving US	
	(1)	(2)
Advisers' hawkishness	6.816**	11.198**
	(2.734)	(4.629)
President's hawkishness	−3.510**	−3.795
	(1.593)	(2.832)
War ongoing		−0.377**
		(0.153)
Deaths per capita in last war (logged)		0.163
		(0.111)
Months since last war (logged)		0.294
		(0.198)
Victory in last war		−1.070
		(0.760)
MID challenges to US in last 5 years		−0.035
		(0.049)
Average MID outcome in last 5 years		−55.273
		(33.947)
Economic recession		−0.194
		(0.216)
Unified government		0.404
		(0.258)
US material capabilities		−5.518
		(6.850)
President's tenure (logged months)		−0.016
		(0.090)
Constant	−2.989***	−3.126
	(0.836)	(2.942)
Observations	501	501

Note: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. Advisers' hawkishness reflects average hawkishness score of senior advisers in a given month. See Appendix §5.8 for details.

Finally, given space constraints, our analysis here focuses on establishing that adviser dispositions affect the counsel they provide leaders in deliberations, and the decision the leader makes—rather than the follow-up question of when leaders are more likely to heed advisers' counsel, a question we explore in other research. Nonetheless, important recent work by Saunders (2017, 2018) suggests some possible scope conditions to our findings, such that leaders may be less likely to be swayed by their advisers when leaders are more experienced and when dovish leaders are paired with dovish advisers. In Appendix §5.11, we use our decision-making dataset to test both propositions. We find that, at least when it comes to adviser hawkishness, neither leader-adviser gaps in experience nor leader-adviser gaps in predisposition significantly moderate the effects of adviser traits. Across the Cold War, hawkish advisers in the United States were no less able to push decision-making in a hawkish direction under experienced leaders than inexperienced ones; similarly, hawkish and dovish advisers appear to influence hawkish and dovish leaders alike. Second, one of the virtues of our main analysis is that it avoids aggregation bias by considering the universe of substantive decisions being made at the meetings, but this raises questions about whether our results are the artifact of lower stakes decisions rather than the high stakes decisions made in crises. To ensure this is not the case, Appendix §5.10 shows that the effects of adviser dispositions remain significant both in and out of international crises featuring decision-making on high-stakes issues.

5 Conclusion

Foreign policy decisions are made in groups, but whether for theoretical or methodological reasons, we know much more about the role of leader-level dispositions in shaping foreign policy outcomes than adviser-level ones. In this article, we develop an argument linking adviser dispositions to consequential foreign policy choices about peace and conflict. We test our proposition by introducing a new methodological approach that estimates, at a distance, the hawkishness of over 1,000 American presidents and advisers who participated in over 2,500 of the most important foreign policy meetings from 1947 to 1988 for which archival records exist. Our theoretical and empirical innovations allow us to move beyond conceptualizing advisers as fungible extensions of leaders—and systematically study the ways in which the advisers surrounding leaders matter, particularly for questions of interstate conflict.

The theory and findings suggest that leaders by themselves may be insufficient to explain many foreign policy decisions. While we emphasize that the leader and adviser models are complementary, we show that leaders consistently turn to advisers for counsel during consequential foreign policy meetings, that adviser dispositions shape the type of counsel that leaders receive, and that shifting from a maximally dovish to a maximally hawkish advisory group triples the expected number of conflictual decisions in a meeting. These dynamics illuminate an intuitive and compelling reason that advisers wield such influence: even experienced leaders confront numerous policy challenges on which they are relatively uninformed and hold few preconceived notions. And even when a leader knows *what* they want, there is often room for hearing disparate perspectives concerning the numerous possible strategies for *how* to get it. Advisers supply the information, analysis, and recommendations that leaders demand.

More broadly, our findings cast doubt on a longstanding tradition in IR arguing that the "aggregation problem" renders the study of group member attributes an unfruitful path for inquiry. Our argument instead emphasizes that deliberation serves as a crucial and under-explored conduit through which individual-level dispositions, such as hawkishness, affect foreign policy outcomes. Aggregation does not forge, *ex machina*, a tabula rasa within the group. Knowing the dispositions of the advisers who dominate policy debates has substantial explanatory power for state behavior.

The basic logic of our adviser model suggests that leaders depend on advisers for psychological and informational reasons, which suggests broad applicability across countries beyond the United States. Yet three core elements of the model also imply corresponding scope conditions. First, because advisers must be able to provide counsel congruent with their disposition without fear of leader retribution, our model may be less applicable in authoritarian regimes, particularly personalist ones (Weeks, 2014), in which leaders can severely and arbitrarily punish advisers who disagree. Second, because advisers must have access to the leader, the adviser model may offer limited insight into countries with bureaucratic institutions designed to exclude advisers from decision-making, such as China during Deng Xiaoping's early years (Jost, 2023). Third, since leaders must be at least somewhat receptive to the counsel that advisers provide, the adviser model may not apply under leaders who are extremely closed-minded either for situational (Reiter and Stam, 2002) or dispositional (Preston, 2001) reasons.

Our adviser model suggests a broad agenda for future research. Most broadly, it calls for more

scholarly attention to be paid to leaders, advisers, and the institutions that connect them. Yet this study intentionally sets aside several questions that subsequent scholarship could explore. First, it does not differentiate between cases where advisers are successful in shaping decision-making because leaders rely on their counsel to form beliefs on issues they have not fully considered, versus cases where advisers successfully persuade leaders to change their views. Future scholarship could leverage the data we introduce to tease apart these two mechanisms.

Second, while we provide evidence suggesting that what happens during deliberations matters to a leader's decision, it is also possible that adviser influence depends on expending effort by setting the agenda and building bureaucratic coalitions *prior to* deliberations. It is also possible that some advisers may be more influential than others—and that past deliberations shape future ones in intriguing, path-dependent ways.

Third, while our primary aim in this manuscript is to show how advisers matter systematically for outcomes of broad concern to the field of IR, these dynamics are clearly the final set of steps in a long causal chain. A more comprehensive approach would systematically study the antecedent stages prior to entering the "room where it happens"—from institutional design, to adviser appointment, to adviser attendance, to adviser behavior, to the decision.

Finally, our method of inquiry suggests how scholars might connect the study of decision-making with the study of numerous other international behaviors, ranging from interstate signaling and threat perception, to alliance formation and crisis escalation. Traditionally, the field of IR has studied these questions either through rich qualitative examination of archival documents or through quantitative methods that focus on state behavior rather than decision-making. The approach developed here offers a middle path: to study state behavior by quantitatively analyzing archival documents that span an extended period of time (Min, 2022), but in a way that still directly observes the decision-making process.

References

Allison, Graham T. and Philip Zelikow. 1999. *Essence of Decision: Explaining the Cuban Missile Crisis*. 2 ed. New York: Longman.

Aristotle. 1920. *Politics*. Oxford: Clarendon Press.

Brutger, Ryan and Joshua D. Kertzer. 2018. "A Dispositional Theory of Reputation Costs." *International Organization* 72(3):693–724.

Burke, John P. and Fred L. Greenstein. 1989. *How Presidents Test Reality: Decisions on Vietnam, 1954 and 1965*. Russell Sage Foundation.

Byman, Daniel L. and Kenneth M. Pollack. 2001. "Let Us Now Praise Great Men: Bringing the Statesman Back In." *International Security* 25(4):107–146.

Carter, Jeff and Charles E. Smith. 2020. "A Framework for Measuring Leaders' Willingness to Use Force." *American Political Science Review* 114(4):1352–1358.

Chiozza, Giacomo and H.E. Goemans. 2011. *Leaders and International Conflict*. Cambridge, UK: Cambridge University Press.

Croco, Sarah E. 2011. "The Decider's Dilemma: Leader Culpability, War Outcomes, and Domestic Punishment." *American Political Science Review* 105(3):457–477.

Destler, I.M. 1972. *Presidents, Bureaucrats, and Foreign Policy: The Politics of Organizational Reform*. Princeton University Press.

Feaver, Peter D. 2003. *Armed Servants: Agency, Oversight, and Civil-Military Relations*. Cambridge: Harvard University Press.

Fordham, Benjamin O. 2007. "The Evolution of Republican and Democratic Positions on Cold War Military Spending: A Historical Puzzle." *Social Science History* 31(4):603–636.

Fuhrmann, Matthew. 2020. "When Do Leaders Free-Ride? Business Experience and Contributions to Collective Defense." *American Journal of Political Science* 64(2):416–431.

Garrison, Jean A. 1999. *Games Advisors Play*. College Station: Texas A&M University Press.

George, Alexander. 1980. *Presidential Decision Making in Foreign Policy*. Boulder: Westview Press.

George, Alexander L. 1969. "The 'Operational Code': A Neglected Approach to the Study of Political Leaders and Decision-Making." *International Studies Quarterly* 13(2):190–222.

Goldhamer, Herbert. 1978. *The Adviser*. New York: Elsevier.

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

Hafner-Burton, Emilie M., Stephan Haggard, David A. Lake and David G. Victor. 2017. "The Behavioral Revolution and International Relations." *International Organization* 71(S1):S1–S31.

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.

Herrmann, Richard K. and Michael P. Fischerkeller. 1995. "Beyond the Enemy Image and Spiral Model: Cognitive-Strategic Research After the Cold War." *International Organization* 49(3):415–50.

Herrmann, Richard K., Philip E. Tetlock and Penny S. Visser. 1999. "Mass Public Decisions to Go to War: A Cognitive-Interactionist Framework." *American Political Science Review* 93(3):553–573.

Holmes, Marcus. 2018. *Face-to-Face Diplomacy: Social Neuroscience and International Relations*. Cambridge, UK: Cambridge University Press.

Holsti, Ole R. 2004. *Public Opinion and American Foreign Policy*. Revised edition ed. Ann Arbor, MI: University of Michigan Press.

Holsti, Ole R. and James N. Rosenau. 1984. *American Leadership in World Affairs: Vietnam and the Breakdown of Consensus*. Boston: Allen and Unwin.

Horowitz, Michael C., Allan C. Stam and Cali M. Ellis. 2015. *Why Leaders Fight*. New York: Cambridge University Press.

Hurwitz, Jon and Mark Peffley. 1987. "How Are Foreign Policy Attitudes Structured?" *American Political Science Review* 81(4):1099–1120.

Jeong, Gyung-Ho. 2018. "Measuring Foreign Policy Positions of Members of the US Congress." *Political Science Research and Methods* 6(1):181–196.

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

Jervis, Robert. 1998. *System Effects: Complexity in Political and Social Life*. Princeton: Princeton University Press.

Jervis, Robert. 2013. "Do Leaders Matter and How Would We Know?" *Security Studies* 22(2):153–179.

Jost, Tyler. 2023. *Bureaucracy and War: The Institutional Origins of Miscalculation in International Politics*. New York: Cambridge University Press.

Jost, Tyler and Joshua D. Kertzer. 2023. "Armies and Influence: Public Deference to Foreign Policy Elites." *Journal of Conflict Resolution* Forthcoming.

Kaarbo, Juliet. 1998. "Power Politics in Foreign Policy: The Influence of Bureaucratic Minorities." *European Journal of International Relations* 4(1):67–97.

Katagiri, Azusa and Eric Min. 2019. "The Credibility of Public and Private Signals: A Document-Based Approach." *American Political Science Review* 113(1):156–172.

Kennedy, Andrew. 2011. *The International Ambitions of Mao and Nehru: National Efficacy Beliefs and the Making of Foreign Policy*. New York: Cambridge University Press.

Kenwick, Michael R and Sarah Maxey. 2022. "You and Whose Army? How Civilian Leaders Leverage the Military's Prestige to Shape Public Opinion." *The Journal of Politics* 84(4).

Kertzer, Joshua D. 2016. *Resolve in International Politics*. Princeton: Princeton University Press.

Kertzer, Joshua D., Brian C. Rathbun and Nina Srinivasan Rathbun. 2020. "The Price of Peace: Motivated Reasoning and Costly Signaling in International Relations." *International Organization* 74(1):95–118.

Kertzer, Joshua D., Marcus Holmes, Brad L. LeVeck and Carly Wayne. 2022. "Hawkish Biases and Group Decision Making." *International Organization* 76(3):513–548.

Khong, Yuen Foong. 1992. *Analogies at War: Korea, Munich, Dien Bien Phu, and the Vietnam Decisions of 1965*. Princeton: Princeton University Press.

Krasner, Stephen D. 1972. "Are Bureaucracies Important? (Or Allison Wonderland)." *Foreign Policy* (7):159–179.

LeVeck, Brad L. and Neil Narang. 2017. "The Democratic Peace and the Wisdom of Crowds." *International Studies Quarterly* 61(4):867–880.

Lin-Greenberg, Erik. 2019. "Backing Up, Not Backing Down: Mitigating Audience Costs through Policy Substitution." *Journal of Peace Research* 56(4):559–574.

Maoz, Zeev. 1990. "Framing the National Interest: The Manipulation of Foreign Policy Decisions in Group Settings." *World Politics* 43(1):77–110.

Mattes, Michaela and Jessica L. P. Weeks. 2019. "Hawks, Doves and Peace: An Experimental Approach." *American Journal of Political Science* 63(1):53–66.

McManus, Roseanne W. 2021. "Crazy Like a Fox? Are Leaders with Reputations for Madness More Successful at International Coercion?" *British Journal of Political Science* 51(1):275–293.

Mearsheimer, John. 2001. *The Tragedy of Great Power Politics*. W.W. Norton & Company.

Min, Eric. 2022. "Painful Words: The Effect of Battlefield Activity on Conflict Negotiation Behav-

ior." *Journal of Conflict Resolution* 66(4-5):595–622.

Mintz, Alex and Carly Wayne. 2016. "The Polythink Syndrome and Elite Group Decision-Making." *Political Psychology* 37:3–21.

Mitchell, Sara McLaughlin and Will H. Moore. 2002. "Presidential Uses of Force During the Cold War: Aggregation, Truncation, and Temporal Dynamics." *American Journal of Political Science* 46(2):438–452.

Murray, Shoon Kathleen. 2002. *Anchors against Change: American Opinion Leaders' Beliefs After the Cold War*. University of Michigan Press.

Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71(S1):S265–S277.

Preston, Thomas. 2001. *The President and His Inner Circle: Leadership Style and the Advisory Process in Foreign Policy Making*. New York: Columbia University Press.

Rathbun, Brian C. 2007. "Hierarchy and Community at Home and Abroad: Evidence of a Common Structure of Domestic and Foreign Policy Beliefs in American Elites." *Journal of Conflict Resolution* 51(3):379–407.

Rathbun, Brian C. 2014. *Diplomacy's Value: Creating Security in 1920s Europe and the Contemporary Middle East*. Ithaca: Cornell University Press.

Reiter, Dan and Allan C. Stam. 2002. *Democracies at War*. Princeton: Princeton University Press.

Renshon, Jonathan. 2017. *Fighting for Status: Hierarchy and Conflict in World Politics*. Princeton: Princeton University Press.

Russett, Bruce. 1990. "Doves, Hawks, and US Public Opinion." *Political Science Quarterly* 105(4):515–538.

Saunders, Elizabeth N. 2011. *Leaders at War: How Presidents Shape Military Interventions*. Ithaca: 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.

Saunders, Elizabeth N. 2018. "Leaders, Advisers, and the Political Origins of Elite Support for War." *Journal of Conflict Resolution* 62(10):2118–2149.

Schafer, Mark and Scott Crichlow. 2010. *Groupthink versus High-Quality Decision Making in International Relations*. New York: Columbia University Press.

Schub, Robert. 2022. "Informing the Leader: Bureaucracies and International Crises." *American Political Science Review* 116(4):1460–1476.

Schultz, Kenneth A. 2005. "The Politics of Risking Peace: Do Hawks or Doves Deliver the Olive Branch?" *International Organization* 59(1):1–38.

Snyder, Richard C., H.W. Bruck and Burton Sapin. 1962. Decision-Making as an Approach to the Study of International Politics. In *Foreign Policy Decision-Making*, ed. Richard C. Snyder, H.W. Bruck and Burton Sapin. Glencoe, IL: The Free Press.

Sylvan, Donald A. and James F. Voss, eds. 1998. *Problem Representation in Foreign Policy Decision-Making*. New York: Cambridge University Press.

't Hart, Paul, Eric Stern and Bengt Sundelius. 1997. *Beyond Groupthink: Political Group Dynamics and Foreign Policy-Making*. Ann Arbor: University of Michigan Press.

Taber, Charles S. and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50(3):755–69.

Weeks, Jessica L.P. 2014. *Dictators at War and Peace*. Ithaca: Cornell University Press.

Wendt, Alexander. 2004. "The State as Person in International Theory." *Review of International Studies* 30(2):289–316.

Whitlark, Rachel Elizabeth. 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, and Assessment of Intentions in International Relations*. Princeton: Princeton University Press.

Yarhi-Milo, Keren. 2018. *Who Fights for Reputation in International Politics? Leaders, Resolve, and the Use of Force*. Princeton, NJ: Princeton University Press.

Yarhi-Milo, Keren, Joshua D. Kertzer and Jonathan Renshon. 2018. "Tying Hands, Sinking Costs, and Leader Attributes." *Journal of Conflict Resolution* 62(10):2150–2179.