

Foreign Policy Attitudes As Networks

Joshua D. Kertzer^{*} and Kathleen E. Powers[†]

Forthcoming in the *Oxford Handbook of Behavioral Political Science*,
Alex Mintz and Lesley Terris, Editors.

Last revised: May 2019

ABSTRACT: For the past fifty years, public opinion scholars have searched for signs of "constraint" in the American public's foreign policy attitudes. We review these attempts here, suggesting that the ensuing work has ultimately fallen into two research traditions that have largely been conducted in isolation of one another: horizontal models that portray attitudes as being sorted along multiple dimensions on the same plane, and vertical models that imply a hierarchical organization in which abstract values determine specific policy positions. We then offer a new – networked – paradigm for political attitudes in foreign affairs, leveraging tools from network analysis to show that both camps make unrealistically strict assumptions about the directionality and uniformity of attitude structure. We show that specific policy attitudes play more central roles than our existing theories give them credit for, and the topology of attitude networks varies substantially with individual characteristics like political sophistication.

ACKNOWLEDGMENTS: This research was supported by the Behavioral Decision Making initiative at Ohio State, and benefited from presentations at Ohio State, the 2013 NYU Experiments Conference, and the 2013 MPSA and APSA meetings. A special thanks to Robin Baidya, Jan Box-Steffensmeier, Greg Caldeira, Chris Gelpi, Brian Greenhill, Matt Hitt, George Ibrahim, Luke Keele, Samara Klar, Vlad Kogan, Neil Malhotra, Kathleen McGraw, Joanne Miller, William Minozzi, Irfan Nooruddin, Dave Peterson, Brian Rathbun, Jason Reifler, Jonathan Renshon, Molly Roberts, Chris Skovron, Ben Valentino, and Ben Wagner for their helpful assistance and feedback on previous versions of this project, and to Briggs DeLoach for helpful research assistance. Authors are listed in alphabetical order.

^{*}Paul Sack Associate Professor of Political Economy, Department of Government, Harvard University. 1737 Cambridge St, Cambridge, MA 02138. Email: jkertzer@gov.harvard.edu. Web: <http://people.fas.harvard.edu/~jkertzer/>.

[†]Assistant Professor of Government, Dartmouth College. 223 Silsby Hall, Hanover, NH 03755. Email: kathleen.e.powers@dartmouth.edu. Web: www.kepowers.com.

1 Introduction

How do ordinary citizens organize their foreign policy views? Ever since [Converse \(1964\)](#) first asked whether the American public was “innocent of ideology,” public opinion scholars have searched for signs of constraint in the public’s beliefs, in both domestic ([Peffley and Hurwitz, 1985](#); [Feldman, 2003](#)) and foreign policy ([Gravelle, Reifler and Scotto, 2017](#); [Wittkopf, 1986](#); [Hurwitz and Peffley, 1987](#)) domains. Pointing to structural patterns that characterize how attitudes and beliefs relate to one another as part of a broader system — what psychologists call *interattitudinal structure* ([Judd and Krosnick, 1989](#); [Lavine, Thomsen and Gonzales, 1997](#); [Dinauer and Fink, 2005](#)) — political scientists have demonstrated that attitudes are indeed better organized “than a bowl of cornflakes” ([McGuire, 1989](#), 50), and that the public perhaps deserves more credit for their belief systems than originally thought. Against the Almond-Lippmann consensus, the past forty years have produced a plethora of models examining the many different ways in which the public’s attitudes are connected to one another.

Our existing models of how people organize their foreign policy attitudes were built in two waves. Political scientists first turned to correlations and factor analysis to show meaningful patterns in how people think about world affairs ([Modigliani, 1972](#)), and then borrowed insights from social psychology to theorize about the underlying properties and processes of belief systems ([Hurwitz and Peffley, 1987](#)). These investigations have provided considerable insight into how people think about the world, demonstrating piece by piece how different attitudes predict one another: for example, beliefs about freedom are associated with support for free trade ([Rathbun, 2016](#)), internationalism with support for war ([Herrmann, Tetlock and Visser, 1999](#)), and globalization with attitudes towards rising powers like China ([Scotto and Reifler, 2017](#)).

In this chapter, we point to two challenges faced by the foreign policy attitude literature. The first is a lack of theoretical integration. Surveying the landscape of scholarship on the structure of foreign policy attitudes, we show that much of this work falls into two largely

distinct camps: horizontal models that portray attitudes as being sorted along multiple dimensions on the same plane, and vertical models that imply a hierarchical organization in which abstract values determine specific policy positions. Indeed, a network analysis of citation patterns among the past five decades of scholarship on foreign policy attitudes reveals four clusters separated by discipline, substance, and theory. Despite an increase in communication across traditions, scholars remain divided by their horizontal or vertical assumptions.

This divide thus points to a second, methodological challenge. Contemporary foreign policy research agrees that attitudes are constrained, but disagrees about how, often using causal language without employing techniques capable of adjudicating between competing claims about the direction and nature of this constraint.

As a way of bridging this divide, we introduce political scientists to a new — networked — paradigm for political attitudes. If one of the overarching goals of behavioral approaches to political science is to incorporate more realistic assumptions about human cognition into our theoretical models, the flexibility of this networked paradigm is well suited to the task, enabling us to integrate some of the rival expectations of the existing foreign policy public opinion literature while asking new questions about the properties of the system. To understand belief systems, we need to focus not just on beliefs, but also on systems; not only relationships between attitudes, but also on the holistic properties that characterize how they work at a higher level of analysis. What makes belief systems “hang together”? How are they systemic in the first place? To answer these questions, we apply tools from network analysis that are frequently used to yield insights about interpersonal ([Sokhey and Djupe, 2011](#)) and international ([Hafner-Burton, Kahler and Montgomery, 2009](#)) networks, but as far as we are aware, have not been applied to the study of interattitudinal networks.

After introducing our networked paradigm, we present a novel experimental design — a *directed multigraph experiment* — that disentangles the directional nature of attitude constraint while leveraging network analysis tools to study belief systems systemically. As a

plausibility probe for our theoretical framework, we present our experimental results, suggesting that studying foreign policy attitudes as networks leads to three important findings.

First, we find evidence that belief systems can be driven from the bottom-up rather than the top-down: specific policy attitudes are more central in belief systems than existing work would lead us to expect. Second, although IR scholars often model the public as a monolithic entity, we find that the topology of attitude networks varies dramatically across individuals. Third, we show that political knowledge and interest play major roles in foreign policy attitude networks: knowledge is critical in shaping the degree to which attitudes are structured, while foreign policy interest determines whether specific policy positions or abstract attitudes are more central to belief systems. Finally, we conclude by reviewing implications for political psychologists, IR scholars, and political scientists more broadly.


2 The Network of Foreign Policy Attitude Research

Ever since [Converse \(1964\)](#) claimed that many citizens in the mass public are innocent of ideology, political scientists have sought to show that foreign policy belief systems are nevertheless coherent, searching for signs of “constraint” among political attitudes.³ Foreign policy public opinion scholars have proposed a plethora of models to explain the structure of foreign policy attitudes — by our count, there have been at least 112 different foreign policy attitude structure pieces published since 1970 — but we can understand them as falling into two different camps.

The first wave of scholarship consisted of *horizontal* models that posit foreign policy attitudes are organized along multiple (usually independent) dimensions that exist on the same plane ([Wittkopf, 1986](#); [Holsti and Rosenau, 1988](#)). Work in this tradition tends to search for clusters of related policy attitudes along one or several dimensions, typically recovered inductively through factor analysis. For example, in a theoretical framework that still exerts

³We follow [Eagly and Chaiken \(2007, 583\)](#) in defining attitudes as “an individual’s propensity to evaluate a particular entity with some degree of favorability or unfavorability.”

Figure 1: The citation network of foreign policy attitude structure research, 1970-2017



By examining the citation patterns between the 112 foreign policy attitude network pieces, a latent position cluster model (LPCM) classifies the literature into four clusters. An arrow from node i to node j shows that article j cites article i : thus, for example, in the top left-hand corner, article 54 ([Mayton, Peters and Owens, 1999](#)) is cited by article 77 ([Vail and Motyl, 2010](#)). The pie chart for each article represents the posterior probability (estimated using 5000 MCMC samples) that each article falls into a particular cluster; here, the light grey represents work on trade attitudes, the dark grey represents a cluster of psychological work on foreign policy attitudes, white shows an older horizontal tradition, and black represents a newer wave that combines both vertical and horizontal models. The x and y axes represent the first two principal components in a four-dimensional Euclidean latent space as measured by the Minimum Kullback-Leibler (MKL) estimates. The LPCM is estimated using the `latentnet` package in R ([Krivitsky and Handcock, 2008](#)). See Table 1 in the online appendix for a legend identifying each article in the network.

enormous influence in how foreign policy scholars think about American public opinion, [Witkopf \(1990\)](#) argued that foreign policy attitudes throughout the Cold War could be sorted along two dimensions: militant internationalism (MI) and cooperative internationalism (CI). Of the 29 pieces published before 1990, 27 adhere to the assumptions of horizontal models.

The second wave consists primarily of *vertical* models that understand attitudes as being sorted hierarchically according to their degree of specificity, with attitudes towards specific policy issues determined by general postures and/or underlying values ([Hurwitz and Peffley, 1987](#); [Jenkins-Smith, Mitchell and Herron, 2004](#); [Lieberman, 2006](#)).⁴ These pieces look for evidence of top-down constraint, often using regression-based or structural equation modeling approaches. For example, [Kertzer et al. \(2014\)](#) argue that foreign policy attitudes are driven by moral values: individuals who understand morality in terms of fairness and avoiding harm tend to be opposed to militarist foreign policies, while individuals who understand morality in terms of respecting authority and tradition, and loyalty to one’s ingroup are more likely to favor the use of force. A full 43 of the 56 works published after 2000 posit top-down constraint, reflecting the growth of interest in the role that core values play in shaping political attitudes more generally (e.g. [Feldman, 2003](#)).⁵

⁴In the public opinion literature more broadly, see [Peffley and Hurwitz \(1985\)](#); [Feldman \(2003\)](#); [Goren \(2005\)](#); [Jacoby \(2006\)](#).

⁵To illustrate the divisions in the field, we conducted a search of the literature for all scholarship making arguments about the structural basis of foreign policy attitudes at the individual level published from 1970 to 2017. Of 112 pieces that met our criteria, 57 adopt horizontal assumptions, while 55 put forward vertical models. To empirically evaluate the extent to which different parts of the foreign policy attitude literature engage one another, we generate a citation network by encoding citations between each of the 112 pieces in an adjacency matrix, analyzing clustering in the network using a 4-dimensional latent position cluster model (LPCM) presented in Figure 1 ([Handcock, Raftery and Tantrum, 2007](#)). A full discussion of the methodology and results is available in the online appendix, but we briefly note four key findings here: First, there is a clear disciplinary divide – work by psychologists on foreign policy issues forms a single cluster in the network that is relatively disconnected from developments in political science. Second, there is a substantive divide between security and economics: 15 of 16 pieces on trade attitudes comprise a distinct cluster in the citation network. Third, the data confirm our observation that the first wave of attitude research contains primarily horizontal models, with one cluster that contains almost exclusively horizontal models, all published in 1998 or earlier. Fourth, the rise in vertical models over the past 20 years has accompanied more acknowledgement across opposing camps, with a cluster of contemporary pieces in which horizontal and vertical pieces cite each other.

3 Deconstructing Attitude Structure

As a whole, the voluminous literature on public opinion about foreign policy has taught us at least three things about how foreign policy attitudes are structured.

First, horizontal models show that foreign policy attitudes “go together”: evaluations of one foreign policy issue are related to attitudes towards other foreign policy issues. We know from Wittkopf’s (1990, 33) early factor analyses that the same Americans who, e.g., favored arms control agreements with the Soviet Union tended to also support cultural exchanges with the Soviets. Similarly, citizens who believe that the U.S. should prioritize its national interests also call for a strong military (Gravelle, Reifler and Scotto, in press). These examples demonstrate that foreign policy attitudes are not entirely unorganized, but rather cluster in predictable ways. Holding one type of foreign policy attitude makes people more likely to hold another, related, position. Since the dimensions in horizontal models are typically assumed to be independent of one another (but see Reifler, Scotto and Clarke, 2011), where people stand on one dimension (e.g. militant internationalism) will not be affected by where they stand on another: both hardliners and internationalists, for example, are perched at the same hawkish place on Wittkopf’s (1986) militant internationalism axis.

Second, vertical models show us that specific foreign policy attitudes depend not only upon each other, but also on more general dispositions and values: they have *abstract antecedents*. Vertical models thus depict a top-down version of attitude constraint whereby a change in a position on one attitude element requires compensatory changes on other attitude elements within the system. Scholars in this tradition typically identify the relevant values or orientations a priori and use regression or structural equation models to demonstrate that the variables of interest influence foreign policy attitudes. Conservation values predict support for the use of force (Rathbun et al., 2016), for example; general attitudes towards globalization predict foreign policy preferences towards China (Scotto and Reifler, 2017). Drawing from an array of potentially influential dispositions, others show that beliefs about retributive justice (Lieberman, 2013), general political orientations (Hurwitz and

[Peffley, 1987](#); [Scotto and Reifler, 2017](#)), or commitments to hierarchy/community ([Rathbun, 2007](#)) predict a host of foreign policy attitudes. For vertical models, then, the centrality of an element in the attitude system depends on its position in the hierarchy: values, at the top of the hierarchy, are central because they directly or indirectly shape all elements below them.

Third, both horizontal and vertical models suggest that connections between attitudes are not coincidental, but tell us something about how people form and update political attitudes; structure both shapes and reveals process. Models of social cognition suggest not only that how we respond to new information in the world depends heavily on information already in our heads, but that we form associational links when we consider multiple attitudes simultaneously ([Judd and Krosnick, 1989](#); [Lodge and Taber, 2013](#)). In other words, attitudes are neither formed nor updated in a vacuum. Individuals rely on existing ideas to shape their attitudes toward global politics, and related attitudes will interact with one another. For [Wittkopf \(1990\)](#) and others in the horizontal tradition, this means that a person's attitude toward the Iraq war should inform her opinion toward a proposed intervention in Iran. Or as vertical modelers like [Hurwitz and Peffley \(1987\)](#) suggest, people already carry around ideas about the morality of warfare, which then inform their degree of militarism. When a person forms her position on a potential intervention in Iran, these background beliefs shape it.

Although horizontal and vertical models make different assumptions about how attitudes are structured, the tension between these assumptions is rarely discussed. Perhaps this is because proponents of each type of model are focused on rejecting the null hypothesis of no constraint, rather than adjudicating between competing claims about the directionality. We propose to address this theoretical divide by building on the insights reviewed above to reconsider foreign policy attitudes as part of a network marked by both horizontal and vertical dynamics.

Attitudes as networks

These properties of foreign policy attitudes — that they are interconnected, have abstract antecedents, and rely on existing structure to give them context — reflect a broader understanding of attitudes as systems. Like other systems in social science, attitudes are an assemblage of interacting units in which the overall structure is important in its own right (Converse, 1964; Boutyline and Vaisey, in press). Our approach studies the systemic properties of belief systems by explicitly conceptualizing attitudes as networks. This networked paradigm understands each attitude — cooperative internationalism, the Iraq war, fairness — as an individual node. Nodes form various, multidirectional connections to one another when they are brought to mind simultaneously, and the system’s structure emerges from these patterns of connections. System level properties, in turn, affect our expectations about how attitudes in the network form or update. This broadens the scope of our research questions: rather than ask whether one attitude correlates with another, we ask how stimulating attitudes at different levels of abstraction induces change in connected attitudes throughout the network, and whether the system is tightly connected or relatively sparse. Moreover, this enables us to explicitly explore, rather than assume, the directionality of connections between attitudes and thus test horizontal and vertical dynamics simultaneously. Below, we briefly review the psychological and political science foundations of our approach before describing how an explicit networked paradigm advances our understanding of foreign policy attitudes.

Network models have played an important role in psychological research on the organization of memory (Loftus and Collins, 1975) and attitudes (Judd and Krosnick, 1989). While political scientists have done much to connect foreign policy public opinion with social cognition research, which explores how social information is stored in memory (Fiske and Taylor, 1984, McGraw, 2000), the network implications have to this point been largely under-appreciated. Psychologists understand attitudes as “networks of associated cognitions” (Monroe and Read, 2008, 734), whether through associative (Anderson, 1983) or

connectionist (Feldman and Ballard, 1982; Smith, 2009) models. Network models envision attitude objects (Democrats, gun control, war in Afghanistan) as *nodes* stored in a system and connected to one another through *associational links*, the structure of which has implications for how attitudes are accessed and formed (Lodge and McGraw, 1995; Lavine, Thomsen and Gonzales, 1997). Network models posit that links emerge or strengthen when multiple attitudes are activated simultaneously, and assuming that attitudes are linked necessarily implies that “the representational structure becomes a network” (Judd and Krosnick, 1989, 109).

As Sinclair (2012, 14) argues regarding the influence of social networks on political behavior, little political science work explicitly models attitudes as networks, but it has been nevertheless suggestive of the paradigm. Whenever we talk about “spreading activation” in priming effects (Valentino, 1999), memory-based processing in theories of survey responses (Zaller, 1992), or schemas in political belief systems (Conover and Feldman, 1984), we are implicitly studying the consequences of networked structure. Despite invoking networked metaphors, neither our theories of how attitudes interact nor the methods we use to test them fully embrace the implications of treating attitudes as networks. A networked paradigm challenges our models of foreign policy attitudes in at least three ways.

First, if attitudes are systems, much work in political psychology has been what systems theorists would call “reductionist,” identifying relationships between individual attitudes but telling us little about the system as a whole. Systems are greater than the sum of their parts, and network theory directly incorporates system-level properties: the location of a node in a network has consequences (Borgatti et al., 2009), for example. Highly central nodes, those with many connections, are structurally situated to exert more influence. The overall degree to which a network is connected – its density – also matters for how people form and update their attitudes. Denser attitude networks are subject to more widespread changes than sparse networks in which there are few pathways between different types of attitudes. This has implications for attitude formation that cannot be gleaned from individual relationships.

In this respect, although our methods differ, our interests are similar to those of scholars who study elite beliefs. Operational code research values the whole as much as its parts (Holsti, 1970; Schafer and Walker, 2006; Renshon, 2009). A leader’s operational code — her “philosophical and instrumental beliefs about the nature and use of power in the international system” — comprises multiple components (Renshon, 2009, 650). Knowing a leader’s stance on one component is valuable, but less informative than analyzing how the components combine as a system. Taber (1992), for example, emphasizes this complexity with a computational model that accommodates dynamic cognitive environments. Individual attitudes and bivariate associations are useful objects of study, but like operational code scholars, we argue that their systemic properties are also theoretically important.

Second, most recent research on foreign policy attitude structure assumes that attitudes are organized from the top-down (vertical constraint): general attitudes like ethnocentrism shape specific attitudes towards the Iraq War, but not the reverse (Hurwitz and Peffley, 1987; Rathbun, 2007; Rathbun et al., 2016; Scotto and Reifler, 2017). Yet we know that even core attitudes are subject to change — John Foster Dulles’ operational code shifted after WWII (Holsti, 1970), and in public opinion, the link between values and partisanship changes with one’s social network (Lupton, Singh and Thornton, 2014). Moreover, psychological research shows that directional links between attitudes are not necessarily organized by abstract schemas like ideology, and may instead be based on other concrete positions — attitudes about Iraq might shape judgments about future interventions (Lavine, Thomsen and Gonzales, 1997). Thus, both political scientists and psychologists raise the possibility of bottom-up attitude constraint — something we can test by analyzing structural positions in the network. Directed ties (which flow from one attitude element to another) indicate whether stimulating one attitude element produces a change in another. If priming core values leads to widespread attitude changes, top-down processes are at work. If instead an individual’s position on the Iraq War primes other attitudes, we learn that foreign policy events can provoke broader attitude change.

Similarly, while most horizontal models treat militant and cooperative internationalism as independent dimensions that do not interact with the policy positions they organize (Wittkopf, 1986; Rathbun, 2007), studying the network allows us to adjudicate this claim. We can determine whether an individual’s position on one dimension stimulates changes in the other or on attitudes throughout the system. If so, they are more central sources of constraint than current research suggests.

Third, past work implies an important degree of homogeneity in foreign policy belief systems. While individual attitudes vary, the manner in which they are connected—the underlying structure—remains the same (e.g. Wittkopf, 1990; Rathbun, 2007; Rathbun et al., 2016). For example, in Hurwitz and Peffley’s (1987) model, individuals can be more or less militaristic, but militarism always shapes attitudes towards defense spending. We argue that systemic properties vary across subsets of the population. In a representational network, links are formed and strengthened over time as individuals “[come] to believe that one object implies, favors, contradicts, or opposes the other object” (Judd and Krosnick, 1989, 109), or generally when two or more elements are simultaneously represented, as individuals must have formed an association between two attitude objects in order for one to influence the other (Lavine, Thomsen and Gonzales, 1997). Given the diversity of political information and experience in the public, however, we expect that political sophistication shapes how foreign policy attitudes are organized — and we can use tools from network analysis to assess how densely foreign policy attitudes are connected in individuals with high and low levels of political sophistication. Much of the public opinion literature has sought to show that foreign policy elites and the mass public structure their foreign policy views in similar ways (e.g. Holsti and Rosenau, 1988; Chittick, Billingsley and Travis, 1995; Rathbun, 2007) but our networked paradigm implies otherwise: individuals with relatively high levels of political knowledge should more readily connect abstract and specific attitudes (Zaller, 1992). Further, building upon Judd and Krosnick’s (1989) finding that political attitudes are more consistent among people interested in politics, we expect that those most interested in foreign

policy will have denser foreign policy attitude networks.

Thinking about attitudes as networks thus suggests an understanding of foreign policy attitude structure that differs from previous work in three substantive ways. First, it shifts the conversation from isolated relationships toward a holistic discussion of how the system operates. Second, while scholars have largely sought to exonerate the public by showing that constraint flows from the top-down, it can also flow from the bottom-up: if attitudes are networks, then attitudes at lower levels of abstraction (e.g. specific policy positions) may play a larger role in the system than often assumed. Finally, past work may overestimate the uniformity of attitude structure: if the particular pattern of connections varies with experience, then the topology of attitude networks should vary with individual characteristics.

Methods and materials

Observational data, network analysis, and directionality

Treating attitudes as networks carries two important methodological implications. First, if one attitude can simultaneously affect multiple others, we need a research design that can leverage this directional question: which attitudes stimulate change, and which are stimulated? Existing models have made considerable advancements in understanding how people think about world politics, but political scientists traditionally approach attitude structure as a measurement problem to be resolved with structural equation modeling (SEM), in which structure is a product of prior theoretical assumptions, rather than something to be tested empirically.⁶ Because SEM with static observational data is poorly suited to testing causal claims, this literature adopts the directional language of causal inference without employing the methodological techniques conducive to carrying it out.⁷ We could forgo claims of

⁶This has produced spirited debate in the public opinion literature about the specific criteria used to establish the topology of foreign policy beliefs (e.g. [Kegley 1986](#); [Chittick, Billingsley and Travis 1995](#), [Gravelle, Reifler and Scotto, 2017](#), in press).

⁷Notable exceptions include uses of panel data (e.g. [Peffley and Hurwitz 1993](#); [Murray 1992](#); [Goren 2005](#)), but these studies stand out in the foreign policy public opinion literature because their dynamic setup is rare.

directionality altogether, but whenever we estimate a regression model where we treat one attitude as a dependent variable, and another as an exogenous independent variable, we are make a directional claim about constraint. “The meaning of a strong relationship between one attitude and another is far from obvious” in this setup, given that switching the IV and DV to opposite sides of an equation tends to produce a similar coefficient (Fordham and Kleinberg, 2012, 316). Instead, we tackle the directional question head-on. Our experimental approach follows the advice of Judd et al. (1991, 193), who call for scholars to explore the “dynamic implications” of attitude structure, and Peffley and Hurwitz (1993, 83), who are “unable to say much about the exact reasoning processes that give rise to top-down... patterns of attitude constraint, a task better left to future, experimental studies.”

Second, if attitudes are networks, they should be studied as such. Past work deals with attitude complexity using computational models (Taber, 1992), analyzing how media consumption shapes connections between attitudes (Althaus and Kim, 2006), or by asking individuals to categorize concepts (Berinsky and Kinder, 2006) — but only recently has research begun to explore political attitudes with network analysis (Boutyline and Vaisey, in press). Given that scholars use network analysis to study everything from connections between organizations to ties between neurons (Butts, 2009), and that our theories already imply a network structure, nothing precludes its use here. Indeed, influential works in public opinion already invoke network properties like degree centrality as ways to think about belief systems (Luskin, 1987, 859). We develop an experiment and employ network analysis, which allows us to derive structure from relationships that emerge in the data rather than imposing it *ex-ante*. Our intent is to explore a novel paradigm for studying belief systems, which requires far more empirical testing than would be possible in a single article (Lakatos, 1970); in this sense, the empirical tests we outline below should be understood as a proof of concept that illustrates the value of studying attitude systems holistically.

Study sample

We conducted an online experiment in 2013 to investigate dynamic attitude networks. 1,307 participants were recruited using [Amazon Mechanical Turk](#).⁸ Participants, 59.6% of whom identified as male, and 78.4% as white/Caucasian, ranged in age from 18-75 (median: 28).

Measures, design, and instrumentation

As [Sinclair \(2012, 27-28\)](#) notes, there are at least two ways to experimentally study network structure. One is to directly manipulate the network by altering its structure; the other is to “stimulate some part of the network directly and then observe to what extent these manipulations spill over into other parts of the network.” We adopt this latter strategy, employing an experimental design where we prime one foreign policy attitude at a time and trace the network structure by measuring how these effects spill over onto other attitudes.

Network models of memory are often tested using simulation methods, which allow for networks with a relatively large number of nodes ([Monroe and Read, 2008](#)). In order to maintain a tractable experimental design, the foreign policy attitude network we test here is far smaller, consisting of 11 nodes at three levels of abstraction: four core values (authority/respect, harm/care, ingroup/loyalty, and fairness/reciprocity), three foreign policy orientations (militant internationalism, cooperative internationalism, and isolationism), and four policy attitudes (towards the war in Iraq, NATO-led intervention in Libya, potential intervention in Syria, and renewing the Kyoto protocol). The elements of the network tested here thus represent a theoretically-motivated sample of the potentially massive number of relevant nodes — an often necessary step when analyzing large network populations.⁹

Following terminology from graph theory, we test foreign policy attitude networks with what we call a *directed multigraph experiment*, because it allows us to simultaneously adjudicate the direction of connections between multiple attitude nodes as well as make claims

⁸See Appendix §2.1 for an additional discussion about the experimental sample.

⁹See the Appendix §2.2 for a detailed discussion of how we selected these 11 attitudes for the experimental network.

about system-level properties such as network density.¹⁰ The experiment proceeds in four steps: Participants 1) answer a series of demographic questions, 2) are exposed to a single stimulus paragraph that primes one of 11 possible nodes (or a control message), 3) respond to 11 question blocks that measure each value, orientation, and policy position from the network, and 4) respond to a series of questions on political knowledge, ideology, and partisanship.

The experiment employs a persuasive-message stimulus (Dinauer and Fink, 2005) in which participants are exposed to either one of eleven paragraphs designed to prime their position on the attitude, or to an unrelated control. This approach is designed to activate the targeted attitude object together with the positive associations offered in the message and any additional positive or negative associations that they already bear. The experiment thus primes a single attitude object — a node in the system — to activate existing associational links between that attitude and others in the network. Connected attitude objects become targets of indirect change, as participants change or entrench these attitudes in response to the original, overtly unrelated priming paragraph. Past work provides evidence for similar indirect pathways, in which a prime about one attitude affects related positions. Judd et al. (1991) show that when a survey question activates one attitude, participants’ responses to subsequent, linked attitudes become more polarized due to spreading activation processes. Blankenship, Wegener and Murray (2012, 609) show that policy attitudes shift through an “indirect route” — priming participants with equality changes their positions on affirmative action, for example. We expect these similar effects throughout participants’ foreign policy attitude networks where directed links have already formed. If priming attitudes about the war in Iraq influences how participants’ respond to militaristic foreign policy generally, they will bring the activated associations about Iraq to bear on questions about militant internationalism — despite the fact that MI questions are cast at a general level.

The 11 stimulus paragraphs — presented in full in the appendix §2.1 — were similar

¹⁰In graph theory, attitude constraint implies a *directed* graph; given distinct attitudes A and B , if A shapes B , we draw a directed edge (A, B) . If both (A, B) and (B, A) exist, we have a multigraph.

in length and composition, designed to prime positive attitudes toward the manipulated attitude element; for example, we prime authority/respect by noting that “authority figures represent the wisdom of their position and can help people have better lives.” These short messages were targeted one attitude each and avoid direct references to other nodes in the network.¹¹ The manipulations thus directly stimulate only one attitude, such that position changes on other elements, relative to participants who received the control message, could be attributed to the network association/link. In other words, the manipulations prime individual nodes in the network in order to observe the indirect effect on other nodes — mapping the *ex ante* structure of the system, and not priming the network itself.

This experimental approach differs dramatically from most political science research on attitude structure, which infers structure by calculating correlations across survey items. We contend that our approach has three advantages. First, to suggest that one attitude influences another posits a dynamic, directed process, and experiments are better suited to determining whether *A* constrains *B* or *B* constrains *A* (Fordham and Kleinberg, 2012). Second, it ties into Converse’s (1964) original understanding of dynamic attitude constraint: “when new information changes the status of one idea-element in a belief system, by postulate some other change must occur as well” (208). Our experiment targets new information at each node in the network to mimic this process, and measures change with respect to other nodes. Third, by mapping changes across the network rather than from one attitude to another, we can investigate the system holistically.

Participants were first exposed to the persuasive message stimulus, and then answered a series of questions about each of the eleven attitudes — which were viewed in randomly presented blocks. To measure the four abstract moral values, we relied on twenty-four items from the moral foundations questionnaire (Haidt and Graham, 2007). Participants record both the relevance of particular considerations to their judgments of right and wrong

¹¹Because they relate to salient foreign policy issues, some of the messages may be more likely to trigger affective reactions while others may be read as more cognitively-laden. Consistent with the psychology literature on automatic attitude activation, where both affect and cognition are implicated in spreading activation (Fazio, 2001), we expect that either process will facilitate change in associated attitudes.

on a 5-point scale ranging from “not at all relevant” to “extremely relevant,” and their (dis)agreement with statements about whether each value is important on a 6-point scale.

While widely used scales exist to measure militant and cooperative internationalism (Wittkopf, 1986), many typical items gauge support for specific policies such as Cold War containment or humanitarian interventions. In order to maintain a distinction between abstract orientations and policy attitudes, we chose subsets of the standard MI and CI items. These questions tap participants’ general ideas about the use of force in global politics (MI) and about the need for the U.S. to cooperate with other nations and organizations to solve transnational problems (CI). Isolationism is measured using a scale that taps the extent to which individuals think that the U.S. would be better served by focusing on national issues as opposed to becoming involved in the world.

The final four blocks of foreign policy questions measured attitudes toward the 2003 U.S. intervention in Iraq, the 2011 airstrikes conducted by NATO in Libya, the potential for U.S. military involvement in the Syrian conflict, and the possibility of a new agreement to replace the expired Kyoto Protocol.¹² Participants then reported their political ideology and party identification, indicated their personal interest in foreign policy, and completed a 9-item political knowledge quiz. This included four questions recommended by Delli Carpini and Keeter (1993), and five supplementary questions, three of which specifically tested knowledge on international political issues.¹³

Creating the directed multigraphs

A conventional approach would be to simply analyze the magnitude and statistical significance of treatment effects on each dependent variable in sequence. This sequential approach is problematic for two reasons. First, with eleven treatments and eleven dependent variables, the results are too complex to clearly convey and thus obscure more than they reveal. Sec-

¹²Measurement scales for each item are presented in full in appendix §2.2; scale reliabilities in appendix §3.

¹³Questions are presented in the appendix §2.4.

ond, our quantities of interest are not individual treatment effects, but overall configuration; our hypotheses concern the network-level structure of foreign policy attitudes, rather than whether one particular attitude influences another.

We thus present our results by analyzing the experimental data using tools from network analysis. In our novel setup, each attitude object (authority, the Iraq war, etc.) constitutes a *node* or *vertex* in a network. If manipulating one node causes a change in another node above a certain threshold, we deem the two nodes to be connected to one another (or, in network terminology, to share an *edge*). For example, if stimulating authority causes a change in attitudes towards the Iraq war, we draw an edge from authority to Iraq. Because the experimental design lets us calculate treatment effects bidirectionally — that is, we can tell both whether manipulating node i causes a change in node j , and whether manipulating node j causes a change in node i — the network is a *directed* graph. We can not only display the network graphically, but also use simple statistical techniques to analyze the networks’ topology, thereby highlighting the “big picture” even while comparing across subgroups. This approach has four steps: (i) calculate the treatment effects, (ii) “thin” or “threshold” the results by retaining only the largest treatment effects, which are then entered into an adjacency matrix, (iii) plot the adjacency matrix as a network, and (iv) employ permutation and randomization tests to calculate the probability that the observed network topologies are due to chance. We discuss each step in turn.

The first step to generating any network is to define its edges (Butts, 2009). The conventional approach would be to estimate each of our treatment effects using t-tests, drawing an edge from node i to node j if i ’s treatment effect on j is statistically significant according to threshold α . However, the sheer number of hypothesis tests produces an obvious multiple comparisons problem, for which even relatively powerful empirical Bayes techniques like local fdr will produce overly sparse networks. We propose two alternatives. First, we could simply abandon thresholding altogether (Thomas and Blitzstein, 2011), drawing edges between all nodes and weighting them according to the magnitude of each respective treatment

effect. To facilitate the use of simpler binary network statistics, however, we instead adopt a thresholding approach, in which we threshold edges based upon the distribution of potential effect sizes in the data, rather than by making reference to a broader population from which our sample is assumed to be drawn.¹⁴

To produce the edges, we first calculate all of the treatment effects at once, both for the full sample of participants, and for subsamples that divide participants along two measures of political sophistication.¹⁵ We represent our treatment effects with t-statistics, since they capture both the difference in means between the treatment and control groups and also their sample sizes and standard deviations. For the analyses reported below, we take the absolute value of the test statistics and threshold them by declaring the top 10% of effects to merit consideration and discarding the remaining 90%; in the Appendix we show that our results are robust to the selection of different cut points.

Next, effects that meet this 10% threshold are then used to produce an 11×11 adjacency matrix that depicts the relationships between the 11 nodes in a given sample, where entry (i, j) refers to an edge from treatment i to dependent variable j . Since the matrix is binarized and thresholded, those treatment effects that meet the 10% test are represented with a 1, and those that fail to meet the test are represented with a 0. We then plot this adjacency matrix, and calculate a variety of network statistics to characterize the network topology.

Finally, we engage in network inference to produce uncertainty estimates around these network statistics, comparing our observed networks with randomly generated networks of similar size and density. To evaluate the likelihood that the distributions of edges amongst values, orientations, and policy attitudes are due to chance, we perform a series of permutation tests. We randomly generate 10000 networks of an identical size and density as our empirically observed network, to test the probability that observed centrality statistics (the number of edges connected to a node or category of nodes) are due to network size and

¹⁴In using the observed data to determine whether a particular effect merits consideration, our analyses are similar in spirit to other forms of local inference (Efron, 2010; Keele, McCannaghay and White, 2012).

¹⁵By calculating the distribution of treatment effects for all of our analyses at once, it prevents the analyses from being ad-hoc.


density alone (Butts, 2008).¹⁶

Results

We present our results in two stages. First, we analyze the attitude network across our full sample. Second, because we expect heterogeneous treatment effects, we generate separate attitude networks for different subgroups to show how foreign policy attitude structure varies as a function of political knowledge and foreign policy interest. We assess the topologies of observed networks with reference to whether the results support assumptions implicit in extant models of attitude structure, alongside a discussion of what we learn by “zooming out” to the system level.

The full sample

Figure 2: Foreign policy attitude network: full sample



Note: Arrows depict the thresholded treatment effects, and thus, patterns of attitude constraint; nodes are scaled by their degree centrality, so larger nodes are connected to a higher number of other attitudes. Core values are represented as circles, foreign policy orientations as diamonds, and policy attitudes as triangles. Thus, core values play a smaller role in the results from the full sample — and specific policy attitudes play a larger role — than typical models of foreign policy attitude structure would suggest.

Figure 2 plots the 11-node network for the full sample.¹⁷ Core values are depicted with

¹⁶See appendix §5 for permutation tests to evaluate network density.

¹⁷All network plots and statistics are generated using the **statnet** suite of packages in R. (Handcock et al., 2008)

circles, foreign policy orientations with diamonds, and specific policy attitudes with triangles. Arrows indicate the direction of attitude constraint, and nodes are scaled by their degree centrality, which measures each nodes' connectedness to other nodes in the sample: the larger the node, the more important a role it plays in the network.¹⁸

Additionally, we differentiate between *outdegree* centrality — how many nodes a particular attitude influences — versus *indegree* centrality — how many nodes influence a particular attitude, to assess directional hypotheses from hierarchical models. Models assuming that abstract attitudes shape more concrete positions from the top-down suggest two observable implications: 1) values will have only outgoing connections (as abstract values shape orientations and policy positions below them) and thus higher out-degree centrality, but very low levels of in-degree centrality. Policy positions should display the opposite pattern, with higher levels of in-degree centrality implying that they are constrained by higher-order concepts, and an absence of out-degree ties that would otherwise indicate bottom-up change in belief systems.

Contrary to many assumptions about attitude constraint, values play a minimal role in shaping orientations and policy positions: indeed, fairness/reciprocity, ingroup/loyalty, and harm/care are isolates disconnected from the network. Indeed, permutation tests demonstrate that moral values play a *smaller* role ($p = 0.024$), and policy positions play a *larger* role ($p = 0.046$), as measured by degree centrality, than we would expect by chance.¹⁹ Constraint occurs both from the bottom-up — positions on the Iraq war shape isolationism and authority — and within a level — attitudes toward Libya influence support for Syrian intervention. Iraq war attitudes are most central, with more overall and outgoing connections than expected by chance ($p = 0.012$ and $p = 0.001$, respectively). This corroborates our

¹⁸We estimated additional measures of centrality, such as Eigenvector centrality, but because of the small number of nodes in the network, many of these measures failed to converge. As Butts (2008, 22) notes, degree centrality is highly correlated “with most other measures of centrality, making it a powerful summary index,” and thus a satisfactory measure of a node’s connectedness in attitude networks.

¹⁹See Table 1 for results from the permuted analyses. Cell entries show the number of total edges (degree centrality), outgoing edges (outdegree centrality), and incoming edges (indegree centrality) by node type (values, orientations, policy positions). They demonstrate whether specific node types display greater levels of centrality than we would see in a random network of equal size and density.

expectation that salient foreign policy events can link to and influence more general ideas about how to conduct foreign policy — and builds on conclusions from operational code research.

Table 1: Permuted Centrality Statistics

		Degree Centrality		Out-degree Centrality		In-degree Centrality	
		Observed [95% CI]	p-value	Observed [95% CI]	p-value	Observed [95% CI]	p-value
<i>Full Network</i>							
	Values	1 [2,8]	0.024	0 [0,5]	0.057	1 [0,5]	0.261
	Orientations	3 [1,7]	0.531	1 [0,4]	0.450	2 [0,4]	0.549
	Policies	8 [2,8]	0.046	5 [0,5]	0.038	3 [0,5]	0.419
<i>Low Interest</i>							
	Values	11 [6,15]	0.503	5 [2,9]	0.560	6 [2,9]	0.444
	Orientations	4 [4,12]	0.054	2 [1,7]	0.180	2 [1,7]	0.185
	Policies	17 [6,16]	0.010	9 [2,9]	0.039	8 [2,9]	0.109
<i>Medium Interest</i>							
	Values	4 [3,10]	0.221	2 [1,6]	0.380	2 [1,6]	0.367
	Orientations	4 [1,8]	0.526	1 [0,5]	0.294	3 [0,5]	0.395
	Policies	6 [3,10]	0.578	4 [1,6]	0.343	2 [1,6]	0.387
<i>High Interest</i>							
	Values	0 [4,13]	0.000	0 [1,8]	0.004	0 [1,7]	0.004
	Orientations	13 [3,10]	0.002	6 [1,6]	0.064	7 [1,6]	0.016
	Policies	9 [5,13]	0.505	5 [1,8]	0.441	4 [1,8]	0.564
<i>Low Knowledge</i>							
	Values	6 [4,12]	0.255	2 [1,7]	0.172	4 [1,7]	0.609
	Orientations	7 [2,10]	0.378	5 [0,6]	0.138	2 [0,6]	0.400
	Policies	9 [4,12]	0.392	4 [1,7]	0.611	5 [1,7]	0.357
<i>Medium Knowledge</i>							
	Values	10 [9,19]	0.106	6 [3,11]	0.415	4 [3,11]	0.103
	Orientations	13 [6,15]	0.196	5 [2,9]	0.578	8 [2,9]	0.104
	Policies	13 [9,19]	0.452	7 [3,11]	0.582	6 [3,11]	0.416
<i>High Knowledge</i>							
	Values	9 [5, 15]	0.445	5 [2,8]	0.599	4 [2,8]	0.413
	Orientations	11 [3,12]	0.081	5 [1,7]	0.295	6 [1,7]	0.132
	Policies	10 [5,15]	0.559	5 [2,8]	0.592	5 [2,8]	0.603

Note: Table displays the number of observed connections — total, outgoing, and incoming — for a given node type in each network [along with 95% CIs derived from 10000 permuted networks], and the corresponding p-values. Statistically significant p-values indicate that the observed centrality score differs from a randomly generated network of equal size and density.

Figure 2 offers an initial indication of what we can glean from a network approach, as it displays patterns that are inconsistent with our current understanding of attitude structure. At the same time, drawing on psychological research on interattitudinal structure (Judd and

Krosnick, 1989), we emphasize that properties of attitude networks will vary with individual-level characteristics like political sophistication. The relatively sparse attitude network derived from average treatment effects across the full sample might thus belie considerable treatment heterogeneity — a hypothesis we explore below.

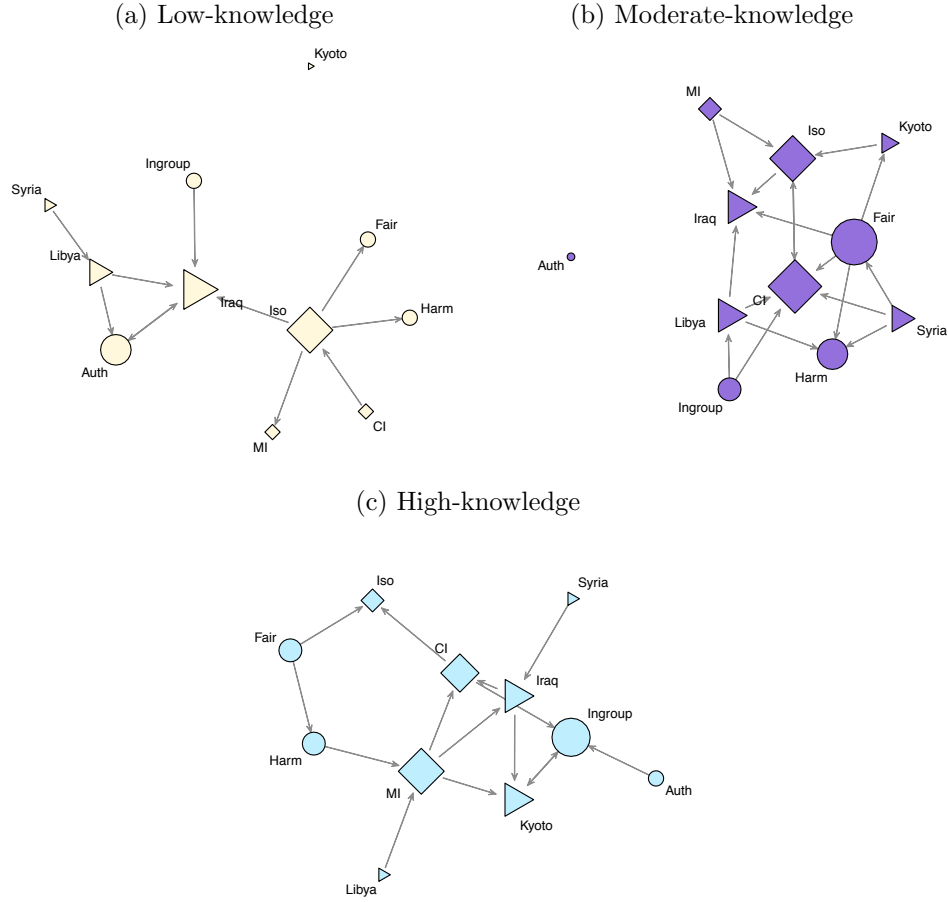
Subgroup analyses

Political sophistication

We next investigate the appropriateness of “one-size fits all” models of attitude structure. Although much of the foreign policy attitude literature has sought to show that experts and the general public structure their attitudes similarly (e.g. Holsti and Rosenau, 1988; Rathbun, 2007), political knowledge affects both people’s likelihood of receiving political information (Zaller, 1992) and how they process the information they do receive (Fiske, Lau and Smith, 1990). Given that the reception and processing of information affect how associational links form, we expect that participants with different levels of political sophistication will exhibit distinct system-level constraint patterns. Since conclusions about political sophistication depend on how it is measured (Luskin, 1987), we investigate it first, as political knowledge (Delli Carpini and Keeter, 1993), and second as foreign policy interest (Krosnick, 1990).

Figure 3 displays networks for three subsamples based on political knowledge, and we note three interesting results. First, network density statistics — which measure how tightly connected the network is organized — reveal a curvilinear effect. Moderately knowledgeable participants have more densely connected foreign policy attitude networks than their high and low-knowledge counterparts – the middle group’s network is 1.63 and 1.2 times more dense than the low- and high-knowledge networks, respectively. This comports with Zaller’s (1992) conclusion that the most sophisticated individuals have attitudes that are more resistant to change, and that the least sophisticated are unable to form connections between abstract and specific concepts like those between values and policy positions. Instead, mod-

Figure 3: Foreign policy attitude network by political knowledge



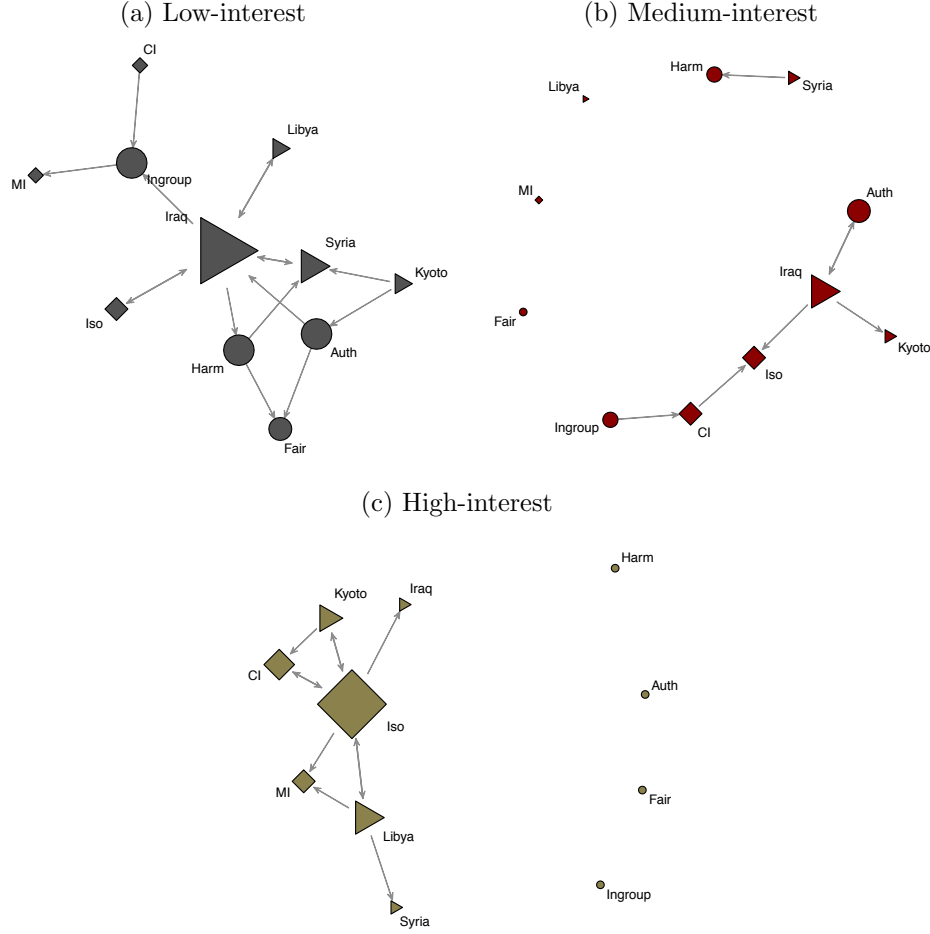
Note: Arrows depict the thresholded treatment effects, and thus, patterns of attitude constraint; nodes are scaled by their degree centrality. Core values are represented as circles, foreign policy orientations as diamonds, and policy attitudes as triangles. Consistent with Zaller (1992), knowledge displays a curvilinear effect here, as the medium-knowledge network has a greater density than either the low-knowledge or high-knowledge counterparts.

erate sophisticates, who hold multiple considerations but link diverse political attitudes, are most open to indirect attitude change. Second, each subgroup network is more dense than the aggregate network in Figure 2, likely due to the fact that aggregating attitude structure masks important structural heterogeneity. Third, only the network for the most knowledgeable subgroup includes a role for MI and CI that is statistically significant in the permuted analyses ($p = 0.081$). It appears that only the most sophisticated individuals engage the abstract views about foreign policy so familiar to academics when forming their attitudes.

Figure 4 operationalizes political sophistication differently by looking at the extent to which participants expressed interest in foreign policy, since interest is often understood as both a cause and consequence of political sophistication (Fiske, Lau and Smith, 1990). Compared to political knowledge, the patterns are strikingly different: with political knowledge, the moderate knowledge network was the most dense, whereas here we see that individuals moderately interested in foreign policy have attitude networks that are less dense than those of the high- and low-interest subgroups. The network for the least interested group is driven primarily by policy positions, especially the war in Iraq — policy attitudes have higher degree centrality generally ($p = 0.010$) and out-degree centrality specifically ($p = 0.039$) than expected by chance. The direction of influence thus flows primarily from specific policy attitudes to other parts of the network, revealing a bottom-up structure for the uninterested public. Such patterns imply that widespread attitude change might occur after a single salient foreign policy event — the chief concern of Almond, Lippmann, and other postwar realists. When the uninterested engage foreign policy information, they might use it to form their broader ideas about how the US should conduct itself internationally as they come to believe that it implicates other pre-existing but newly associated attitude objects.

In contrast, general orientations play a greater role in the work for our highly interested subsample ($p = 0.002$). Moral values, however, play no role — which is telling given the many edges in this dense network ($p = 0.000$). Judd and Krosnick (1989) find that attitudes are more reliably consistent among the group of individuals who ascribe importance to the political attitudes under investigation. Similarly, we find that those most interested in foreign policy closely connect their general orientations toward foreign policy to specific issues. Abstract values, not obviously associated with foreign policy orientations and positions, are less relevant to their foreign policy attitudes. With constraint that flows from orientations to policy positions, the high interest network comes closest to matching top-down assumptions about directionality. Bottom-up effects such as the influence of Libya on both MI and Isolationism, bidirectional edges indicating within-level relationships between CI and Isolationism,

Figure 4: Foreign policy attitude network by interest in foreign policy issues



Note: Arrows depict the thresholded treatment effects, and thus, patterns of attitude constraint; nodes are scaled by their degree centrality. Core values are represented as circles, foreign policy orientations as diamonds, and policy attitudes as triangles. Here, we see the inverse curvilinear effect to that from Figure 3, in that the medium-interest network displays *less* constraint than its low-interest and high-interest counterparts, with specific policy events playing a greater role influencing attitudes throughout the network for low-interest participants, and the more abstract foreign policy orientations playing a greater role influencing other attitudes for high-interest participants.

and core values' disconnectedness all suggest, though, that to assume that constraint flows only in one direction will obscure much of the story.

The results from two political sophistication analyses reveal two important findings. First, both political knowledge and foreign policy interest affect how foreign policy attitudes are structured and the degree of structure itself — something that we can only understand by examining system properties. Against many existing models of foreign policy attitude

structure (e.g., [Holsti and Rosenau, 1988](#); [Rathbun, 2016](#)), we find that expertise matters. Second, regardless of how one operationalizes political sophistication, we detect violations of top-down constraint patterns assumed by past work, including bottom-up influence, bidirectional edges that cross levels, and constraint among attitudes within levels of abstraction.

Conclusion

Political scientists have long sought evidence of constraint in foreign policy attitudes, producing a plethora of models with divergent assumptions about how attitudes are organized. Our comprehensive review of the literature demonstrated that the field has made considerable progress in rebutting the claim that the mass public holds unorganized, unprincipled views about foreign policy. At the same time, it reveals room for progress in at least two areas. First, some bodies of work have been conducted in relative isolation: psychologists and political scientists who study the same phenomenon stay in their disciplinary lanes, and support for free trade policies are treated as distinct from “foreign policy attitudes” writ large, rather than as part of the whole foreign policy belief system. Future research should integrate these areas to provide a more holistic understanding of how, for example, protectionism might fit into an isolationist orientation. Second, the early wave of research on horizontal models was relatively closed off from work that relied on vertical assumptions. Though citation-based engagement has increased over time, and some hint at fruitful ways to combine orientations with abstract antecedents ([Kertzer et al., 2014](#); [Rathbun et al., 2016](#); [Gravelle, Reifler and Scotto, 2017](#)), the key assumptions that divide these traditions tend to be acknowledged in passing and rarely tested against one another.

Beyond this review, we take the first step on the path toward reconciling the theoretical division in research on foreign policy attitudes. We propose a new way of studying political attitudes: as networks. We argued that interattitudinal structure can be conceptualized as a flexible network, where attitudes interact in a system, and developed a novel experimental design that illuminates system properties rather than component parts — while studying

attitudes as network data.

The results from our experimental “plausibility probe” suggest four key conclusions. First, conceptualizing attitudes as networks offers a promising theoretical framework. With our design, we can visualize interactions in a system of attitudes, and directly observe which types of attitudes stimulate more change throughout the system. Consistent with behavioral political science’s emphasis on empirically verifying the micro-level assumptions that animate our theories, we bring together findings from both the horizontal and vertical traditions in foreign policy attitude research to determine how various relationships work within an attitude network. The networked paradigm also allows us to study system-level properties like network density — which suggests important differences in attitude structure based on political sophistication — that would be inaccessible in traditional frameworks.

Second, [Converse’s \(1964\)](#) pessimistic assessment about the paucity of top-down constraint in the American public appears more accurate than revisionist accounts suggest: while regression and SEM-based research designs assume that values and orientations determine policy positions, we find that policy positions — not values — generally play a larger role in the system. Given the extent to which IR scholars have long complained about the moralistic nature of public opinion — Hans Morgenthau claimed that “intoxication with moral abstractions” was “one of the great sources of weakness and failure in American foreign policy” ([Morgenthau, 1951](#), 4) — it is noteworthy that we find more evidence of policy issues shaping moral values than the reverse. The outsize role played by bottom-up constraint suggests that the public is more likely to tie leaders’ hands in response to specific policy issues rather than abstract principles per se. Moreover, this adds to recent evidence that moral values do not reflect the stable predispositions assumed by [Graham, Haidt and Nosek \(2009\)](#), but are open to short- and long-term shifts ([Smith et al., 2017](#)).

Third, one size does not fit all: the oft-dramatic subgroup differences demonstrate that there is no single attitude structure that can be applied to a whole population. Instead, we find sensibly divergent patterns based on political sophistication. This suggests that it is

misleading to speak of such a thing as “the” structure of foreign policy attitudes — but also that the IR literature, which frequently views the public as a monolithic unit responding to international events, would benefit from taking this heterogeneity into account.

Finally, we show that political sophistication helps explain heterogeneous belief systems. Moderately knowledgeable individuals have more densely connected attitude networks than their high and low-knowledge counterparts, and foreign policy interest shapes whether constraint flows mainly from more specific or abstract attitude objects in the network. The central role of expertise in our findings thus raises questions about whether it is realistic to expect the mass public to organize their foreign policy attitudes in the same manner as foreign policy elites.

Importantly, the results also have implications far beyond foreign policy attitudes. Although political psychologists are well aware of order effects in survey responses, the network patterns of interdependence detected here suggest that order effects may be more problematic than we may currently realize. Moreover, directional debates in attitude structure are relatively common in political science: do policy preferences cause party ID, or does party ID cause policy preferences ([Jackson, 1975](#))? Does racial prejudice cause attitudes towards policies like welfare and busing ([Sears, 1988](#))? Traditionally, these hypotheses have been tested either using path analysis of cross-sectional data (whereupon assumptions about the direction of causality are governed entirely by theoretical assumptions), and panel surveys, which offer more leverage on causal inference. The directed multigraph experimental design fielded here offers another way to test directional theories of attitude interdependence by matching method to theory.

References

- Althaus, Scott L and Young Mie Kim. 2006. "Priming Effects in Complex Information Environments: Reassessing the Impact of News Discourse on Presidential Approval." *Journal of Politics* 68(4):960–976.
- Anderson, John R. 1983. *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.
- Berinsky, Adam J and Donald R Kinder. 2006. "Making Sense of Issues Through Media Frames: Understanding the Kosovo Crisis." *Journal of Politics* 68(3):640–656.
- Blankenship, Kevin L, Duane T Wegener and Renee A Murray. 2012. "Circumventing Resistance: Using Values to Indirectly Change Attitudes." *Journal of Personality and Social Psychology* 103(4):606.
- Borgatti, Stephen P, Ajay Mehra, Daniel J Brass and Giuseppe Labianca. 2009. "Network Analysis in the Social Sciences." *science* 323(5916):892–895.
- Boutyline, Andrei and Stephen Vaisey. 2017. "Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes." *American Journal of Sociology* 122(5):1371–1447.
- Butts, Carter T. 2008. "Social Network Analysis: A Methodological Introduction." *Asian Journal of Social Psychology* 11:13–41.
- Butts, Carter T. 2009. "Revisiting the Foundations of Network Analysis." *Science* 325:414–416.
- Chittick, William O., Keith R. Billingsley and Rick Travis. 1995. "A Three-Dimensional Model of American Foreign Policy Beliefs." *International Studies Quarterly* 39(3):313–331.
- Conover, Pamela J. and Stanley Feldman. 1984. "How people organize the political world: a schematic model." *American Journal of Political Science* 28(1):95–126.
- Converse, Philip E. 1964. The nature and origin of belief systems in mass publics. In *Ideology and Discontent*, ed. David E. Apter. New York: Free Press pp. 206–261.
- Delli Carpini, Michael X. and Scott Keeter. 1993. "Measuring Political Knowledge: Putting First Things First." *American Journal of Political Science* 37(4):1179–1206.

- Dinauer, Leslie D. and Edward L. Fink. 2005. "Interattitude Structure and Attitude Dynamics: A Comparison of the Hierarchical and Galileo Spatial-Linkage Models." *Human Communication Research* 31(1):1–32.
- Eagly, Alice H. and Shelly Chaiken. 2007. "The Advantages of an Inclusive Definition of Attitude." *Social Cognition* 25(5):582–602.
- Efron, Bradley. 2010. *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.
- Fazio, Russell H. 2001. "On the Automatic Activation of Associated Evaluations: An Overview." *Cognition & Emotion* 15(2):115–141.
- Feldman, J.A. and D.H. Ballard. 1982. "Connectionist Models and Their Properties." *Cognitive Science* 6(3):205–254.
- Feldman, Stanley. 2003. Values, Ideology, and the Structure of Political Attitudes. In *Oxford Handbook of Political Psychology*, ed. David O. Sears, Leonie Huddy and Robert Jervis. Oxford: Oxford University Press pp. 477–508.
- Fiske, Susan T., Richard R. Lau and Richard A. Smith. 1990. "On the varieties and utilities of political expertise." *Social Cognition* 8(1):31–48.
- Fiske, Susan T. and Shelley E. Taylor. 1984. *Social Cognition*. New York: McGraw-Hill.
- Fordham, Benjamin O and Katja B Kleinberg. 2012. "How Can Economic Interests Influence Support for Free Trade?" *International Organization* 66(02):311–328.
- Goren, Paul. 2005. "Party Identification and Core Political Values." *American Journal of Political Science* 49(4):881–896.
- Graham, Jesse, Jonathan Haidt and Brian A. Nosek. 2009. "Liberals and Conservatives Rely on Different Sets of Moral Foundations." *Journal of Personality and Social Psychology* 96(5):1029–1046.
- Gravelle, Timothy B, Jason Reifler and Thomas J Scotto. 2017. "The structure of foreign policy attitudes in transatlantic perspective: Comparing the United States, United Kingdom, France and Germany." *European Journal of Political Research* .
- Hafner-Burton, Emilie M., Miles Kahler and Alexander H. Montgomery. 2009. "Network Analysis for International Relations." *International Organization* 63:559–592.
- Haidt, Jonathan and Jesse Graham. 2007. "When Morality Opposes Justice: Conservatives

- Have Moral Intuitions that Liberals may not Recognize.” *Social Justice Research* 20(1):98–116.
- Handcock, Mark S., Adrian E. Raftery and Jeremy M. Tantrum. 2007. “Model-based clustering for social networks.” *Journal of the Royal Statistical Society: Series A* 170(2):301–354.
- Handcock, Mark S., David R. Hunter, Carter T. Butts, Steven M. Goodreau and Martina Morris. 2008. “statnet: Software Tools for the Representation, Visualization, Analysis and Simulation of Network Data.” *Journal of Statistical Software* 24(1):1–11.
- Herrmann, Richard K., Philip E. Tetlock and Penny S. Visser. 1999. “Mass Public Decisions to Go to War: A Cognitive-Interactionist Framework.” *American Political Science Review* 93(3):553–573.
- Holsti, Ole R. 1970. “The ‘Operational Code’ Approach to the Study of Political Leaders: John Foster Dulles’ Philosophical and Instrumental Beliefs.” *Canadian Journal of Political Science* 3(1):123–157.
- Holsti, Ole R and James N. Rosenau. 1988. “The Domestic and Foreign Policy Beliefs of American Leaders.” *Journal of Conflict Resolution* 32(2):248–294.
- Hurwitz, Jon and Mark Peffley. 1987. “How Are Foreign Policy Attitudes Structured?” *American Political Science Review* 81(4):1099–1120.
- Jackson, John E. 1975. “Issues, Party Choices, and Presidential Votes.” *American Journal of Political Science* 19(2):161–185.
- Jacoby, William G. 2006. “Value Choices and American Public Opinion.” *American Journal of Political Science* 50(3):706–723.
- Jenkins-Smith, Hank C., Neil J. Mitchell and Kerry G. Herron. 2004. “Foreign and Domestic Policy Belief Structures in the U.S. and British Publics.” *Journal of Conflict Resolution* 48(3):287–309.
- Judd, Charles M. and Jon A. Krosnick. 1989. The structural bases of consistency among political attitudes: Effects of political expertise and attitude importance. In *Attitude structure and function*, ed. Anthony R. Pratkanis, Steven J. Breckler and Anthony G. Greenwald. Hillsdale, NJ: Erlbaum pp. 99–128.
- Judd, Charles M., Roger A. Drake, James W. Downing and Jon A. Krosnick. 1991. “Some Dynamic Properties of Attitude Structures: Context-Induced Response Facilitation and

- Polarization.” *Journal of Personality and Social Psychology* 60(2):193–202.
- Keele, Luke, Corrine McConaughy and Ismail White. 2012. “Strengthening the Experimenter’s Toolbox: Statistical Estimation of Internal Validity.” *American Journal of Political Science* 56(2):484–499.
- Kegley, Charles W., Jr. 1986. “Assumptions and Dilemmas in the Study of American’s Foreign Policy Beliefs: A Caveat.” *International Studies Quarterly* 30(4):447–471.
- Kertzer, Joshua D, Kathleen E Powers, Brian C Rathbun and Ravi Iyer. 2014. “Moral support: How moral values shape foreign policy attitudes.” *The Journal of Politics* 76(3):825–840.
- Krivitsky, Pavel N. and Mark S. Handcock. 2008. “Fitting Position Latent Cluster Models for Social Networks with `latentnet`.” *Journal of Statistical Software* 24(5).
- Krosnick, Jon A. 1990. “Government Policy and Citizen Passion: A Study of Issue Publics in Contemporary America.” *Political Behavior* 12(1):59–92.
- Lakatos, Imre. 1970. Falsification and the Methodology of Scientific Research Programmes. In *Criticism and the Growth of Knowledge*, ed. Imre Lakatos and Alan Musgrave. Cambridge: Cambridge University Press pp. 91–196.
- Lavine, Howard, Cynthia J. Thomsen and Marti Hope Gonzales. 1997. “The Development of Interattitudinal Consistency: The Shared-Consequences Model.” *Journal of Personality and Social Psychology* 72(4):735–749.
- Liberman, Peter. 2006. “An Eye for an Eye: Public Support for War Against Evildoers.” *International Organization* 60(3):687–722.
- Liberman, Peter. 2013. “Retributive support for international punishment and torture.” *Journal of Conflict Resolution* 57(2):285–306.
- Lodge, Milton and Charles S Taber. 2013. *The Rationalizing Voter*. Cambridge University Press.
- Lodge, Milton and Kathleen M. McGraw, eds. 1995. *Political Judgment: Structure and Process*. University of Michigan Press.
- Loftus, Elizabeth and Allan M. Collins. 1975. “A Spreading Activation Theory of Semantic Processing.” *Psychological Review* 82(6):407–428.
- Lupton, Robert N, Shane P Singh and Judd R Thornton. 2014. “The Moderating Impact of

- Social Networks on the Relationships Among Core Values, Partisanship, and Candidate Evaluations.” *Political Psychology* .
- Luskin, Robert C. 1987. “Measuring Political Sophistication.” *American Journal of Political Science* 31(4):856–899.
- Mayton, Daniel M. II, Danya J Peters and Rocky W Owens. 1999. “Values, militarism, and nonviolent predispositions.” *Peace and Conflict* 5(1):69–77.
- McGraw, Kathleen M. 2000. “Contributions of the Cognitive Approach to Political Psychology.” *Political Psychology* 21(4):805–832.
- McGuire, William J. 1989. The Structure of Individual Attitudes and Attitude Systems. In *Attitude structure and function*, ed. Anthony R. Pratkanis, Steven J. Breckler and Anthony G. Greenwald. Erlbaum pp. 37–69.
- Modigliani, Andre. 1972. “Hawks and doves, isolationism and political distrust: An analysis of public opinion on military policy.” *The American Political Science Review* 66(3):960–978.
- Monroe, Brian M. and Stephen J. Read. 2008. “A General Connectionist Model of Attitude Structure and Change: The ACS (Attitudes as Constraint Satisfaction) Model.” *Psychological Review* 115(3):733–759.
- Morgenthau, Hans J. 1951. *In Defense of the National Interest*. New York: Alfred A. Knopf.
- Murray, Shoon. 1992. “Turning an Elite Cross-Sectional Survey into a Panel Study while Protecting Anonymity.” *Journal of Conflict Resolution* 36(3):586–595.
- Peffley, Mark and Jon Hurwitz. 1985. “A Hierarchical Model of Attitude Constraint.” *American Journal of Political Science* 29(4):871–890.
- Peffley, Mark and Jon Hurwitz. 1993. “Models of Attitude Constraint in Foreign Affairs.” *Political Behavior* 15(1):61–90.
- 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. 2016. “Wedges and Widgets: The Ideological Origins of Trade Attitudes.” *Foreign Policy Analysis* 12(1):85–108.
- Rathbun, Brian C., Joshua D. Kertzer, Jason Reifler, Paul Goren and Thomas J. Scotto.

2016. "Taking Foreign Policy Personally: Personal Values and Foreign Policy Attitudes." *International Studies Quarterly* 60(1):124–137.
- Reifler, Jason, Thomas J. Scotto and Harold D. Clarke. 2011. "Foreign Policy Beliefs in Contemporary Britain: Structure and Relevance." *International Studies Quarterly* 55(1):245–266.
- Renshon, Jonathan. 2009. "When Public Statements Reveal Private Beliefs: Assessing Operational Codes at a Distance." *Political Psychology* 30(4):649–661.
- Schafer, Mark and Stephen G Walker. 2006. "Democratic Leaders and the Democratic Peace: The Operational Codes of Tony Blair and Bill Clinton." *International Studies Quarterly* 50(3):561–583.
- Scotto, Thomas J. and Jason Reifler. 2017. "Getting tough with the dragon? The comparative correlates of foreign policy attitudes toward China in the United States and UK." *International Relations of the Asia-Pacific* 17(2):265–299.
- Sears, David O. 1988. Symbolic Racism. In *Eliminating Racism: Profiles in Controversy*, ed. Phyllis A. Katz and Dalmis A. Taylor. Plenum Press pp. 53–84.
- Sinclair, Betsy. 2012. *The Social Citizen: Peer Networks and Political Behavior*. Chicago: University of Chicago Press.
- Smith, Eliot R. 2009. "Distributed Connectionist Models in Social Psychology." *Social and Personality Psychology Compass* 3(1):64–76.
- Smith, Kevin B, John R Alford, John R Hibbing, Nicholas G Martin and Peter K Hatemi. 2017. "Intuitive Ethics and Political Orientations: Testing Moral Foundations as a Theory of Political Ideology." *American Journal of Political Science* 61(2):424–437.
- Sokhey, Anand E. and Paul A. Djupe. 2011. "Interpersonal Networks and Democratic Politics." *PS: Political Science & Politics* 44(1):55–59.
- Taber, Charles S. 1992. "POLI: An Expert System Model of US Foreign Policy Belief Systems." *American Political Science Review* 86(04):888–904.
- Thomas, A. C. and J. K. Blitzstein. 2011. "Valued Ties Tell Fewer Lies: Why Not To Dichotomize Network Edges With Thresholds." *ArXiv e-prints* .
URL: <http://adsabs.harvard.edu/abs/2011arXiv1101.0788T>
- Vail, Kenneth E., III and Matt Motyl. 2010. "Support for Diplomacy: Peacemaking and Mil-

- itarism as a Unidimensional Correlate of Social, Environmental, and Political Attitudes.” *Peace and Conflict* 16(1):29–57.
- Valentino, Nicholas A. 1999. “Crime News and the Priming of Racial Attitudes During Evaluations of the President.” *Public Opinion Quarterly* 63(3):293–320.
- Wittkopf, Eugene R. 1986. “On the Foreign Policy Beliefs of the American People: A Critique and Some Evidence.” *International Studies Quarterly* 30(4):425–445.
- Wittkopf, Eugene R. 1990. *Faces of internationalism: Public Opinion and American foreign policy*. Duke University Press.
- Zaller, John R. 1992. *The Nature and Origins of Mass Public Opinion*. Cambridge: Cambridge University Press.