Title
Theories and Methods for a Cognitive Macro-Sociology of Culture

Author
Boutyline, Andrei Grigoryevich

Publication Date
2017

Peer reviewed|Thesis/dissertation

Theories and Methods for a Cognitive Macro-Sociology of Culture

By

Andrei Grigoryevich Boutyline

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Sociology

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Neil D. Fligstein, Co-Chair
Professor Robb Willer, Co-Chair
Professor Ann Swidler
Professor Laura Stoker

Summer 2017

Theories and Methods for a Cognitive Macro-Sociology of Culture

Abstract

Theories and Methods for a Cognitive Macro-Sociology of Culture

by

Andrei Boutyline

Doctor of Philosophy in Sociology

University of California, Berkeley

Professor Neil D. Fligstein, Co-Chair

Professor Robb Willer, Co-Chair

This dissertation emerges out of an effort to understand the supra-individual aspects of attitudes and tastes, and especially those political attitudes that make up "public opinion." Various theoretical accounts conceive of large-scale attitude systems in terms of fields, shared schemas, cultural logics, or partisan ideologies. Though diverse, these accounts all depict attitudes as *structural* phenomena, defined by patterns of relations between cultural, cognitive, or social elements. In the three substantive chapters, I draw on network analysis, statistics, information theory, and computer science to create original methods for such structural analyses. I use them to provide new insights on culture as both individual cognition and macro-scale social organization. The first project examines networks of political attitudes. I find that, across subgroups, attitudes either follow the dominant liberal-conservative logic, or lack systemic organization. In the second project, I clarify and extend existing theories of cultural schemas to develop a greatly improved approach to detecting them in surveys. In the final project, I approach political attitudes as a field of competition. If public opinion is a debate between competing ideological camps, do the camps at least agree on which issues they are debating? My analyses of the skill with which individuals at different positions in the opinion space report their attitudes suggests that no such agreement exists.

To Mom, Dad, and Fabi

Contents

Acknowledgements

It takes a village to write a dissertation, and the number of those who helped me along the way is large enough that I am certain I will not remember everyone. I was immensely lucky to be mentored by Robb Willer and Neil Fligstein, who co-chaired this dissertation. They were exceptionally engaged, generous, and caring advisers. I am forever indebted to them both. Robb's mentorship on crafting a research paper is also evident throughout these pages.

I benefitted greatly from working with Stephen Vaisey, who first turned my attention towards the kinds of questions I address in the following chapters. Ann Swidler, Claude Fischer, Michael Hout, Laura Stoker, and John Levi Martin have all given crucial advice at various turns in this dissertation. The National Science Foundation Graduate Research Fellowship Program and National Science Foundation Integrated Graduate Education Research Traineeship in Politics, Economics, Psychology, and Public Policy have provided necessary financial support.

I also want to thank all of the members of the Berkeley Mathematical, Analytical and Experimental Sociology (MAXSoc) workshop, who have given me feedback and vital camaraderie at every stage of my doctoral work. Thank you Lindsay Bayham, Eric Giannella, Adam Goldstein, Jaren Haber, Jacob Habinek, Pat Hastings, Casey Homan, Katherine Hood, Joe LaBriola, Monica Lee, Kinga Makovi, Sanaz Mobasseri, Laura Nelson, Liana Prescott, Matt Stimpson, and William Welsh. A finer workshop has never existed. I also want to thank my friend John Flournoy for always reminding me what good research and good friendship are.

I owe thanks to my grandma Antonina Vasilyevna and my aunt Olga, who have always inspired me with their deep love of the written word. Any shred of eloquence in my writing is due to them.

Finally, I want to thank my mom, my dad, and my partner Fabiana Silva. Without their love, encouragement, and emotional, material, and intellectual help, none of this would have been possible. My dad's insistence that I learn to program computers at age eight, and the years he spent tutoring me in mathematics afterwards, gave me some of the most important tools I used to write this dissertation. The help he gave me in figuring out some of this dissertation's trickiest formal reasoning was also generous and irreplaceable. My mom's stories about the materials science lab where she worked, on the other hand, are what first gave me the desire to lead a life of the mind that has persisted to this day. I am still so happy that she encouraged me pursue this degree, even though it is not in mathematics, physics, or chemistry. Finally, my partner Fabiana has read, debated, re-read, edited, and re-re-read this dissertation's three substantive chapters. Every sentence in them has benefitted immensely from her brilliant clarity of thought.

Chapter 1

Introduction

This dissertation emerges out of an effort to understand the supra-individual aspects of attitudes and tastes, and especially those political attitudes that make up "public opinion." Various theoretical accounts conceive of large-scale cultural systems in terms of fields, shared schemas, cultural logics, or partisan ideologies. Though diverse, these accounts all depict attitudes as *structural* phenomena, defined by patterns of relations between cultural, cognitive, or social elements. However, because these accounts cannot be readily examined with standard statistical tools, they remain empirically understudied. In the three substantive chapters, I draw on network analysis, statistics, and computer science to create original methods for such structural analyses. I use them to provide new insights on culture as both individual cognition and macro-scale social organization.

Object of Investigation

The overarching interest of this dissertation is in the large-scale patterns of shared cultural elements. In the following chapters, I examine how ideological identity acts as a heuristic to organize political beliefs in the population; how simple transformations of latent cultural schemas structure popular musical tastes; and how ideological competition in the political field influences the stability with which individuals hold their political views. These projects are motivated by a view of culture that differs from those most commonly articulated in sociology. In short, I think of "culture" as any non-independence between individuals' cognitive contents that is not attributable to physiology or immediate effects of an external stimulus. In this introduction, I provide an initial sketch of this model of culture, beginning with the meta-theoretical orientation behind it.

The meta-theoretical orientation of my project comes from foundational works of cognitive science. The junction of sociology of culture and cognitive science has been an active area of investigation since at least the publication of DiMaggio's (1997) and Zerubavel's (1999) well-known programmatic pieces. The work of these two scholars engendered two distinct scholarly traditions. DiMaggio and his followers used findings from cognitive science to question and elaborate the micro-mechanisms behind the cultural processes documented by sociologists. This sociology of "culture and cognition" thus used cognitive science as a lens to study sociology of culture. Zerubavel and his students, on the other hand, engaged in a "culturalist cognitive sociology" (Brekhus 2007), which instead began with the presupposition that cognition is largely social in character. Work in this tradition has focused on the cultural contingency of seemingly universal cognitive processes such as categorization, attention, and perception. In short, DiMaggio's work drew the causal arrow from cognition to culture, and Zerubavel's pointed it the other way around.

The perspective I begin to develop here has much more in common with DiMaggio's approach than it does with Zerubavel's. However, rather than drawing on cognitive science for any particular set of findings or middle-range theories, I organize my thinking around its meta-theory and founding metaphor, which was first prominently articulated by von Neumann (1958). This metaphor conceives of the of brain as computer, mind as software, and thinking as information processing (Gardner 1987).

The metaphor of brain as computer owes much of its power to the work of Claude Shannon. In “A Mathematical Theory of Communication” (1948), Shannon demonstrated that information is a measurable quantity on its own right, and one that exhibits systematic properties that are independent of the physical medium that encodes it—whether that medium is electrical signal, spoken language, non-verbal interaction, or anything else that can serve as an input to the system. Cognitive science as a discipline began with applying this thinking to the study of the mind (Gardner 1987).

One way of seeing the relevance of information theory to the study of cognition begins by noting that an individual’s senses provide the sum-total of the information entering the mind, and the behaviors represent the sum-total of the information at the output. The entirety of an individual’s mental activity then definitionally becomes information processing or cognition—a category that can thus include explicit and implicit thought, and rational calculation as well as emotion.[1]

Within this perspective, if the sensory inputs encountered by a subject are recordable with sufficient accuracy (e.g., as written language or sets of codable behaviors), as are his or her subsequent behavioral outputs, it is possible to track, for example, how much information from the input is contained in the output—and thus to ask what algorithmic properties these cognitive processes exhibit. It is also possible to ask what properties any algorithm that produces the observed outputs from the observed inputs would need to have. Furthermore, such questions can be studied in many of the same ways whether or not the subject is a human, an animal, or a machine, including direct observation of the subject, examinations of problem-solving machinery, and formal analyses of the problem itself.

The concept of information can thus divorces the study of mental processes from the study of brains. In the 1950’s, its introduction thereby created a common interdisciplinary conversation in which psychologists, linguists, computer scientists, and later anthropologists and neuroscientists could partake.[2] Indeed, it is no coincidence that

---

[1] In some usages, the term has been used to draw a boundary between the emotionless and emotionally laden aspects of human mental life, with an “affective science” of emotion sometimes defined as a contrast to “cognitive science” of reasoning. This boundary rests in part on outdated views of the role of rationality in thought, and is now rarely drawn when defining cognitive science as a discipline. More importantly, the distinction is unnecessary and counterproductive in the context of an information-theoretic definition of cognition. Both emotion and rationality are ways of applying existing mental structures to sensory inputs in order to produce behavioral outputs. Therefore, both are forms of information processing and thus cognition.

[2] For example, Chomsky’s famous “poverty of the stimulus” argument (1980) held that human must be born with many aspects of language present in our brains from birth, because the amounts of information that babies are exposed to before they begin speaking are insufficient to learn language without such biological priors. Note that, in arguing for a unique biological feature of human brains, the argument itself makes no recourse to physiology: if it were robots rather than children that were able to learn to speak in the same setting, exactly the same

the field of cognitive science was founded at the Symposium on Information Theory (1958), where psychologist George Miller, linguist Noam Chomsky, and computer scientists Allen Newell and Herbert Simon presented their work (Gardner 1987). As recent sociological work has demonstrated, information theory can also serve as a powerful basis for formal theorizing about culture (Vilhena et al. 2014). Below, I propose a definition of "culture" that is compatible with this approach to cognition.

Definition of "Culture"

To improve the clarity of theoretical work within the sociology of culture, and ease translation between sociology of culture and the cognitive and computational sciences, I propose the following three-part definition of culture[3]:

1) Cognitive contents (i.e. information) can be "trivially" or "non-trivially" shared between individuals.
2) Cognitive contents are "trivially shared" when the incidence of the same cognitive contents across individuals can be effectively attributed to either shared physiology, immediate effects of the same stimulus; or chance. They are "non-trivially shared" when their shared incidence cannot be so attributed.
3) "Culture" is any and all cognitive contents that are non-trivially shared between individuals.

Some examples of cognitive contents that would *not* count as culture under this definition include idiosyncratic thoughts, because they are not shared; the inborn priors that allow humans to learn language (Chomsky 1980), because they are shared wholly due to physiology; or similar patterns of activation in occipital lobe areas V1-V4 of two individuals who are being shown the same image, because they are shared wholly due the immediate effects of the same stimulus.

On the other hand, by this definition, a musical melody or story known by two or more individuals would count as culture, as would a shared emotional experiences or pattern of associations attached to this melody or story. So would shared attitudes, beliefs, values, and internalized norms. The same goes for words, symbols, and categories that individuals acquire through similar acculturation or experience.

While neither events nor objects can themselves count as shared culture[4], the semiotic codes, emotional associations, learned skills, and remembered formative experiences that allow individuals to partake in the same religious rituals or public celebrations would count as culture. So would those cognitive contents that allow one to produce a piece of writing interpretable by others. Conversely, a group of people

argument would imply that such robots arrived with some features of the language pre-programmed.

[3] My focus here is specifically on what Lizardo (2017) calls "personal culture." I do not intend for this definition to cover the many other meanings of the word culture—e.g., all artistic products produced by a social group, or "the cultivation of soil; tillage" (The Free Dictionary).

[4] I.e., by virtue of not being cognitive contents.

witnessing the same dramatic car crash would not be itself an example of culture. However, the shared memory of the crash they retain afterwards would.

Advantages and Implications

This definition of culture brings a number of advantages that resemble those of the informational definition of cognition. First, it greatly improves on the clarity of most existing definitions of culture in sociology, which tend to define culture either entirely by referring to things that count as culture but not the *principles* that make them count as culture; or else by reference to other undefined terms, such as "meaning."[5] The incomplete nature of these definitions makes it so the meanings of the terms need to instead be gleaned inductively through repeated exposure. This can leave scholarship in the sociology of culture unnecessarily opaque to scholars from other fields. It can also make theoretical claims harder to pin down to measurable properties or testable propositions. My definition, in contrast, can more readily accommodate measurable and testable claims.[6]

Second, although my definition draws clear boundaries around culture, these remain broad enough to encompass most aspects of culture currently studied by sociologists. Indeed, one of my primary goals is to delineate culture as a common object of investigation that can be approached by survey research, interviews, ethnographic observation, experiments, formal modeling, and comparative analysis of texts, among many other methods. For example, unlike definitions of culture based around meaning and meaningful action, this definition applies equally well to both deliberative aspects of culture like Swidler's (1986) "cultural toolkit," and automatic aspects of culture such as those that may underlie moral values (Vaisey 2009). It also covers narrative and discourse as well as skill and habitus (Bourdieu 1984; Lizardo 2017). Additionally, it can apply to both very large scale cultural phenomena such as cross-national differences in values (Inglehart and Baker 2000), and very small-scale ones, such as local cultures that emerge within small groups (Fine 2012).

This definition is perhaps least compatible with those treatments of culture that focus specifically on external process, such as work on "meaning-making" (e.g., Lamont 2000). This may be unsurprising, as the "meaning-making" approach to culture intentionally draws a boundary between sociological and social-psychological investigations. My definition instead aims to be amenable to both approaches.

---

[5] Many influential definitions of culture do both, with "meaning" as the key undefined concept. So, for example, in her foundational work, Swidler defines culture as "symbolic vehicles of meaning, including beliefs, ritual practices, art forms, and ceremonies, as well as informal cultural practices such as language, gossip, stories, and rituals of daily life" (Swidler 1986:273). Sewell writes about cultural sociologists as "studying the place of meaning in social life" (Sewell 2005:37). And Pachucki and Breiger define "culture" as referring to "meanings, local practices, discourse, repertoires, and norms" (Pachucki and Breiger 2010:207). None of these treatments define the term "meanings."

[6] Perhaps most prominently, it makes the amount and the complexity of shared culture into concrete quantities that can be measured via mutual information and other entropic measures.

Nonetheless, I note again that, while processes and events cannot themselves count as culture under this definition, they can both still be seen as strong *indicators* of culture. Thus, many key aspects of such “externalist” treatments of culture may be translatable into the theoretical language I introduced here.

Third, because biology only enters the definition in the negative, as something that shared culture cannot be reducible to, this definition creates a layer of abstraction that lets it apply irrespective of whether the “individuals” in question are humans or machines. As with the definition of cognition as information processing, this layer of abstraction opens up the study of culture to formal theoretical analyses, thus expanding the interdisciplinary conversation to include formal disciplines like computer science and theoretical electrical engineering. This same advantage should ease the development of new quantitative methodologies.

In addition to the above-listed advantages, this definition also carries some more concrete theoretical consequences. Most prominently, the focus of *shared* cognitive contents is central to the definition. In order for cognitive contents to become non-trivially shared among large groups of people, any contents not arising from widely shared experiences need to diffuse between people without being lost or distorted. Since, without specific precautions, every act of social information transmission risks informational damage, this is a difficult achievement.

This conception thus suggests an empirical investigation into the social and cognitive mechanisms that enable pieces of culture to become widely shared. The second and third chapter of this dissertation both explore such mechanisms—the identity heuristic in chapter two, and simple transformations of shared schemas in chapter three.[7] The conception also suggests an investigation into the ways that cultural elements fail to be widely shared. This is the focus of my fourth dissertation chapter, which examines how characteristics of the public opinion field create asymmetries in how individuals wield their political views. Detailed summaries of these three substantive chapters follow.

Summaries of Three Substantive Chapters

Chapter 2, Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes (co-authored with Stephen Vaisey).

Many accounts of political belief systems conceive of them as networks of interrelated opinions, in which some beliefs are central and others peripheral. Working together with Stephen Vaisey, we formally show how such structural features can be used to construct direct measures of belief centrality in a network of correlations. We apply this method to the 2000 ANES data, which have been used to argue that political beliefs are organized around parenting schemas. Our structural approach instead yields results

[7] My definition implies that the complexity of shared culture can be measured via information-theoretic techniques. This is one of the approaches I take in the third chapter of this dissertation, where I show that the shared cultural logics organizing popular political attitudes can be accounted for by elementary entailments between pairs of cultural elements.

consistent with the central role of political identity, which individuals may use as the organizing heuristic to filter information from the political field. We search for population heterogeneity in this organizing logic first by comparing 44 demographic subpopulations, and then with inductive techniques. Contra recent accounts of belief system heterogeneity, we find that belief systems of different groups vary in the *amount* of organization, but not in the *logic* which organizes them.

Chapter 3, Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method.
The measurement of shared cultural schemas is a central methodological challenge for the sociology of culture. Relational Class Analysis (RCA) is a recently developed technique for identifying such schemas in survey data. However, existing work lacks a clear definition of such schemas, leaving RCA's accuracy largely unknown. In this chapter, I build on the theoretical intuitions behind RCA to arrive at this definition. I demonstrate that shared schemas should result in linear dependencies between survey rows—the relationship usually measured with Pearson's correlation. I thus modify RCA into a "Correlational Class Analysis" (CCA). When I compare the two methods using a broad set of simulations, results show that CCA is reliably more accurate at detecting shared schemas than RCA, even in scenarios that substantially violate the assumptions behind CCA. I find no evidence of theoretical settings where RCA is more accurate. I then revisit a prior RCA analysis of the 1993 GSS musical tastes module. While RCA had partitioned these data into three schematic classes, CCA partitions them into four. I compare these results with a multiple groups analysis in SEM, finding that CCA's partition yields greatly improved model fit over RCA. I conclude the chapter with a parsimonious framework to guide future work.

Chapter 4, Holding a Position: Public Opinion as Cognition in a Disorganized Field.
Different research traditions have approached political attitudes as either individual-level (psychological) or macro-scale (social) phenomena. In the fourth chapter, I develop "public opinion fields" as a theoretical model for integrating these levels of analysis. This model conceives of political attitudes in the population as a field of competition, where positions correspond to stances on issues—a metaphor encapsulated by the phrase "to hold a position," which can mean both "to have an opinion" and "to defend a location." To demonstrate the utility of this model, I focus on a concrete question: if public opinion is a debate between competing ideological camps, do the distributions of mass attitudes suggest that these competitors agree on what issues the debate is about?

Both sociological practice theories and research on the cognitive demands of survey response suggest that reliably answering survey questions is an acquired cultural skill that requires substantial training to achieve. Existing work furthermore suggests that this training likely comes via the messaging disseminated by political actors and social movements to their followers. Therefore, it follows that respondents' differential skills at answering various questions should reveal which issues their camp prepares them to debate. To explore this, I develop a formal model of position stability, and estimate it with General Social Survey panel data. I find that opposing positions on a topic are rarely held with the same amount of stability: that is, the terms of public debate appear overwhelmingly in dispute. Moreover, majority positions are generally held more stably than minority ones, as would be expected if ideological camps predominantly focused on issues where they were winning.

Chapter 2.

Belief Network Analysis:
A Relational Approach to Understanding the Structure of Attitudes

Theories of the structure of political beliefs typically conceive of them as networks of interrelated opinions, in which some beliefs are central and others are derived from these more fundamental positions (Converse 1964; Jost, Federico, and Napier 2009).[1] There are many such *center-periphery* theories of political ideology, each of which places something different at the center (e.g., political identity, authoritarianism, moral relativism). Research has established the plausibility of such accounts using such distinct quantities as the reliability of survey responses or the ability of "central" opinions to predict peripheral items (e.g., Converse 1964; Barker and Tinnick 2006). Though the empirical examinations of these theories have yielded valuable findings, they have generally not made use of the rich *structural* features of the theoretical accounts they test. Following the intuition of sociological network analysis (Breiger 1974; Wellman 1988; Freeman 2004; Pachucki and Breiger 2010) and building on recent work in the sociology of culture (Baldassarri and Goldberg 2014), we demonstrate how such structural features can be used to construct direct measures of belief centrality in the network of correlations. These centrality measures, together with other network metrics, enable intuitive comparisons between many theories of political belief structure. In this paper, we use this Belief Network Analysis (BNA) approach to contrast several prominent accounts of belief structure and to further elaborate the account most supported by the comparisons. We further demonstrate that these results are robust to sampling error and model selection. We then conduct additional analyses to support key assumptions behind the model against competing claims about heterogeneity of belief structure.

As an orienting case, we first focus on Lakoff's (2002) theory of "moral politics." This theory posits that people reason about the complex domain of policy by metaphorically mapping it onto the domain of family and parenting. By this process, cultural schemas describing two common parenting styles, nurturant and strict, become the "deep structures" underlying the liberal and conservative political worldviews. Though Lakoff's model is frequently cited in sociology (e.g., DellaPosta, Shi, and Macy 2015; Edgell 2012; Gross, Medvetz, and Russell 2011; Hitlin and Vaisey 2013; Jacobs and Carmichael 2002; Somers and Block 2005; Vaisey 2009; Wuthnow 2007), and has also has been deeply influential outside the academy,[2] we know of only a single peer-reviewed work that has lent it full support. This paper (Barker and Tinnick 2006), published in the *American Political Science Review*, interpreted the ability of parenting variables to predict other opinions in the 2000 American National Elections Study data as evidence of their central role in the belief system. To enable comparisons between our method and existing techniques, we revisit this study with our BNA methodology.

Our analysis shows that parenting values in fact occupy a peripheral position in the observed network. We find that the center is instead occupied by ideological identity, which is

---

[1] In this paper, we use the term "beliefs" as shorthand for what Converse (1964) calls "idea-elements" and Zaller (1992) calls "considerations": that is, the various kinds of persistent mental content that make up a person's political ideology, including both information and moral values. This is also the same usage as in Borhek and Curtis's *A Sociology of Belief* (1975).

[2] A 2014 *Salon* profile called *Moral Politics* "a book that should have utterly transformed our understanding of politics", and continued with "[a]nd for many who read it, it certainly did" (Rosenberg 2014). Lakoff's guides to political framing, which are based on this theory, carry endorsements from prominent figures like Howard Dean, Anthony Romero and George Soros (e.g., Lakoff 2014; Lakoff and Rockridge Institute 2006). In 2005, *New York Times* reported that one of these guides was "as ubiquitous among Democrats in the Capitol as Mao's *Little Red Book* once was in the Forbidden City," and quoted Nancy Pelosi describing his framing advice as "perfect for us, because we were just arriving in an unscientific way at what Lakoff was arriving at in a scientific way" (Bai 2005). In the 2016 election cycle, Lakoff's work has appeared in the popular media to explain Donald Trump's rise within the Republican party (DeVega 2016; Williamson 2016).

broadly consistent with theories of *social constraint* (e.g., Campbell et al. 1960; Converse 1964; Zaller 1992; Mondak 1993; Goren, Federico, and Kittilson 2009). In such theories, actors begin with a rudimentary understanding of the institutional field of politics (Sniderman and Stiglitz 2012), and a political identity within this field, which they then use as a heuristic for selecting political views from mass media and other information sources. We find that this centrality of ideological identity is remarkably robust to both statistical noise and the specific choice of variables to include in the model. The low centrality of parenting attitudes is equally robust. Our results thus provide strong evidence in favor of the social constraint account and against the theory of moral politics.

Both of these theories—and our BNA technique—make the assumption that the organization of attitudes is driven primarily by a single dominant process in the population. This assumption is shared by many theories of belief organization (e.g, Jost et al. 2003). However, a number of recent sociological treatments (Baldassarri and Goldberg 2014; Boltanski and Thévenot 1999; Thornton, Ocasio, and Lounsbury 2012; van Eijck 1999; Achterberg and Houtman 2009) assume substantial heterogeneity in how attitudes are organized, arguing for the existence of many different "logics" of constraint. If overall patterns mask substantial heterogeneity, our method could lead to invalid conclusions. Thus, in the second part of the paper, we partition the population along 16 key demographic dimensions, and examine heterogeneity in belief organization between the 44 resulting subpopulations. In the appendices, we also investigate potential heterogeneity with a novel information-theoretic method we introduce here, and using model comparison techniques from structural equation modeling. Contrary to Baldassarri and Goldberg's (2014) high-profile work and other recent sociological accounts, we find that heterogeneity in the organizational logic of political beliefs is the exception rather than the rule.

Among these 44 subpopulations, we find that all belief networks with substantial levels of organization are centered on political identity, and feature similar patterns of pairwise constraint. For those groups further removed from the field of "mainstream" U.S. politics, we find that belief systems instead generally *lack* organization—a result in line with a substantial volume of older work that showed the belief systems of such populations to be low in constraint (e.g., Converse 1964). In a few key subpopulations, however, we find some tentative evidence of a different belief system—one centered on religious identity rather than political identity. But any potential alternate scheme of political belief organization would appear highly limited in scope, suggesting that religious identity may not generally provide a heuristic sufficient to organize the full range of beliefs usually deemed "political."

The rest of the paper proceeds as follows. First, we develop a formal model of belief formation and introduce an empirical method that can adjudicate between competing center-periphery models of beliefs. Second, we use this method to contrast existing theories of belief formation with survey data from the 2000 American National Election Study (ANES). These analyses lead us to reject the theory of moral politics, and to offer an elaboration of the theory of social constraint. Third, we empirically examine a key assumption about population heterogeneity made by our method through demographic comparisons between subgroups, as well as with information-theoretic and psychometric techniques. We again find results consistent with the social constraint account. More broadly, our overarching goal in this investigation of belief structures is to help students of culture better understand how some cultural elements can organize and structure others within a cultural system, which Swidler (2001: 206) has identified as "the biggest unanswered question in the sociology of culture."

BELIEF STRUCTURES AS NETWORKS

Most prominent accounts define ideology as "a learned knowledge structure consisting of an interrelated network of beliefs, opinions and values" (Jost et al. 2009:310; but see Martin 2000). The network metaphor for belief systems fits well with both the definitions and the questions posed by the literature on ideology. A network, after all, is simply a system consisting of a finite set of identifiable entities called "nodes," as well as a set of defined relationships between them called "edges" or "ties." Converse's (1964) classic description of a system of belief elements held together by pairwise constraint or functional interdependence fits this definition of a network. Lakoff's (2002) account of belief generation, in which models of parenting are extended through metaphor and logical inference into full-fledged political belief systems, can also be rendered in network terms.

The benefit of the network language comes from the leverage it provides in succinctly describing the relational properties of such systems. Two key ideas which recur in many accounts of belief systems—the structural positions occupied by different beliefs and the degree of organization of the belief system—make use of network thinking to evoke such properties. We draw on social network analysis to show that the network understanding of belief systems need not stop at evocative metaphor (Breiger 1974; Wellman 1988; Freeman 2004; Pachucki and Breiger 2010). Like Baldassarri and Goldberg (2014), we interpret a set of survey responses as an empirical manifestation of the belief network, where the belief items are nodes and the associations between the beliefs are weighted ties. Going beyond existing work, we construct a formal network model of belief structure, and use it to demonstrate that a particular set of network indices—shortest-path betweenness centrality and centralization—provide theoretically relevant measurements for such a system. We develop this approach, which we term "belief network analysis" (BNA), in the context of comparing two prominent theories of political belief structure: the theory of moral politics and the theory of social constraint. Both lend themselves well to a network conceptualization, and rest on core concepts for which measures are readily available.

Moral Politics

We begin by describing the structural features of Lakoff's theory of moral politics (Lakoff 2002). Lakoff's model has roots in his earlier work on metaphor theory (Lakoff 1990; Lakoff and Johnson 1980). Lakoff proposes that conceptual systems are structured largely through metaphorical inference: metaphors project complex cognitive domains onto simpler ones. For example, two common metaphors used to understand anger are "anger is an opponent" (e.g., "he was struggling with his anger"; "his anger overpowered him") and "anger is a fluid heated up in a container" (e.g., "he was boiling with rage"; "simmer down"). He argues that this word usage reflects deeper differences in conceptual structure: the person who speaks of anger as an opponent may thus decide that he should try his best to fight it, while someone who thinks of anger as a boiling fluid concludes he should let some of it out lest he explode (Lakoff 1990).

In *Moral Politics*, Lakoff (2002) argues that political cognition is also fundamentally metaphorical. He points to terms like "fatherland," "Uncle Sam," "founding fathers" and "big brother" to argue that the "nation is a family" metaphor is the key to understanding political differences. This metaphor maps the complex domain of government onto the more familiar domain of family, allowing people to use their intuitions about parenting to make judgments in the otherwise opaque domain of policy. Lakoff concludes that ideological divisions stem from the fact that “liberals and conservatives have different models of how to raise children” (2002:337). The “strict father” model used by conservatives emphasizes authority, strict discipline and “tough love” as ways to lead the child to self-reliance. The “nurturant parent” model used by liberals emphasizes caring, protection, and respect as the best ways to help children grow up to be fulfilled and happy adults. Liberals thus support environmental protection and generous welfare policies because they are metaphorically understood as forms of parental caring. Conservatives oppose abortion and support mandatory sentencing for drug possession because their morality stresses personal accountability.

The derivation of political beliefs in Lakoff’s (2002) account can be roughly broken down into three phases. A person starts with a simple, unelaborated model of parenting—strict or nurturant. In the first derivation phase, beliefs about how to parent lead to broader moral judgments involving parenting and family. In the second phase, these expanded moral claims are applied to the domain of government. Via the “nation is a family” metaphor, the government becomes the parent, citizens become children, and proper governance becomes proper parenting. Intuitions about family thus yield intuitions about government. In the final phase, these intuitions are used to develop specific policy stances. Parenting beliefs thus become “central” to the system of political views in the sense that they serve as the initial basis from which this system is formed.

Under this model, debates over parenting philosophy are an important form of partisan conflict. Lakoff argues that liberals benefit from the popularity of advice books promoting nurturant parenting, as it “means that there are plenty of parents and children who have an intuitive understanding of the basis of Nurturant Parent morality and liberal politics” (2002:364). However, he proposes that the greater prominence of conservative parenting groups like Focus on the Family may give them an advantage in the longer term, as “the more children brought up with Strict Father values, the more future conservatives we will have” (2002:424). These arguments are consistent with a dynamic model in which peripheral beliefs are recursively generated from a central belief.

To date, only two studies have attempted to provide empirical support for Lakoff’s model.[3] McAdams and colleagues (2008) tested Lakoff’s assertions by examining the authority figures appearing in 128 life-narrative interviews. They found only mixed support: while conservatives were more likely to have strict authority figures, liberals were not more likely to have nurturant ones. Another study (Barker and Tinnick 2006), published in the *American Political Science Review*, used data from the 2000 ANES to show that respondents’ parenting values can predict many policy positions net of a large number of controls. Although the authors interpreted this as evidence for Lakoff’s theory, the existence of net associations is not sufficient to make the *structural* claim that parenting morality is the element which “unifies the collections of liberal and conservative issue positions” (Lakoff 2002:12). Below, we will use our network-analytic methodology to test this claim directly.

Social Constraint

The main alternative we consider to Lakoff’s (2002) account comes from theories of social constraint. Beginning with the classic works of Campbell et al. (1960) and Converse (1964), this diverse body of theories has been unified by the claim that people use political identity as a heuristic for acquiring further political beliefs via the flow of information from

---

[3] *Moral Politics* itself contains little systematic support for Lakoff’s argument beyond its intuitive plausibility and consistency with some anecdotal evidence (see Lakoff 2002:158).

opinion leaders including politicians, journalists and activists (Zaller 1992:6). We draw our term for this theory from Converse (1964:209), who characterized the belief systems produced by this process as "much less logical in the classical sense that they are psychological—and less psychological than social." Research on social constraint has also appeared under other titles, including "source cues," "elite theory," "partisan information processing," and "psychological" (as opposed to "rational") theories of partisan behavior (e.g., Goren 2005; Lee 2002; Mondak 1993; Zaller 1992). Since our primary interest is in the structural features of the belief systems described by these theories, we focus on prominent structural statements (Converse 1964; Zaller 1992) and elaborations of relevant mechanisms (e.g., Goren et al. 2009; Mondak 1993).

Social constraint theorists begin by highlighting the cognitive complexity of ideological reasoning (e.g., Converse 1964; Zaller 1992) and ask how one can come to a coherent or "constrained" political worldview. The difficulties in this task are multiple. Policy positions are not so well-defined that consistent positions across multiple domains could be logically derived from some bounded set of principles. On the other hand, the empirical makeup of most policy issues is complex enough that fully considered judgments would require a prohibitive amount of information. Part of the solution may come from broadly applicable psychological principles such as cognitive heuristics or moral values, which can make it possible to arrive at judgments on such issues based on only a partial understanding (Lau and Redlawsk 2001; Mondak 1993). However, to be mutually consistent, these principles themselves require systematization. Moreover, it may still be far from apparent *which* principle to apply to each issue, as most issues have many aspects and can often be judged using multiple conflicting principles (see e.g., Feinberg and Willer 2013).

For these reasons, adopting an existing system of belief organization is vastly easier than creating such a system from scratch. Various kinds of political elites, such as politicians or television pundits, have much to gain by becoming a "cognitive authority" (Martin 2002) for their audiences. However, adopted views would likely be consistent only if they are received from elites that agree with each other. Ideological and partisan identity can solve this coordination task. Once a person acquires such an identity—by, e.g., imitating their parents or following widely known cultural stereotypes (Green, Palmquist, and Schickler 2002)—he or she can replace the abstract question of "what should I believe?" with the social question "which team am I on?" Humans appear to be highly adept at this kind of social reasoning (Goren et al. 2009; Mondak 1993; Sniderman and Stiglitz 2012).

Ideological identity allows people to tune in to ideological information streams that contain relatively consistent stances on policy issues. They also convey broadly applicable ideological heuristics and stereotypical beliefs about the social world (Hurwitz and Peffley 1997; Kinder 1998; Martin and Desmond 2010; Petersen 2009; Zaller 1992)—for example, heuristics about which social groups require help and which punishment (e.g., drug addicts or the homeless); which potential threats are real and which overblown (e.g., global warming or voting by non-citizens); and which social domains are or are not "the government's business" to regulate (e.g., gun control or abortion restrictions). Individuals can then deploy these principles to pass original judgments on newly encountered issues, or to fill in the gaps in their knowledge about an issue with ideologically consistent stereotypes (Martin and Desmond 2010; Zaller 1992). Such beliefs can also give people positions within specific issue domains (e.g., "anti-war" or "pro-life"), enabling segments of the population to know who is "on their side" even without direct references to ideological identity, and thus easing their acquisition of more finely differentiated domain-specific knowledge. This process yields an expanding and branching belief network that structurally resembles the one suggested by the moral politics account.

MODEL OF BELIEF FORMATION AND NETWORK STRUCTURE

Though the accounts reviewed above have obvious differences, the generative processes they depict share key structural features. We will now develop a formal model of belief formation and network structure that is consistent with these features, which we summarize as follows:

- Individuals start with a single central belief (parenting model or political identity).
- This central belief is used to produce a number of broad stances (moral views or political heuristics), which are then used to stochastically produce further beliefs.
- Newly added sets of beliefs then form the basis for yet newer and more specific beliefs, repeating recursively to yield a center-periphery structure.

This summary is clearly a simplification of both accounts, and is not intended to capture the full psychological or social complexity of belief formation. For example, the model does not leave room for individuals to update their beliefs after they are created. Though prior work indeed suggests that individuals do not often change their political beliefs (Zaller 1992), even in the face of disconfirming information (Taber and Lodge 2006), it is unlikely that proponents of either account would argue that belief change never occurs. Other potential complexities are similarly elided, leaving a minimal model that captures only the main features of both accounts while remaining parsimonious enough to formalize and examine mathematically. We use the analytical leverage provided by this model to answer the following questions: given a correlation network of survey responses from people who formed their beliefs in this way, is it possible to identify the original, central, belief? And if so, how?

We begin with the central belief, designated x_0. New beliefs are recursively produced from older ones, beginning with x_0's direct descendants. When one belief is created from another, each position on the older belief corresponds to some position on the new belief created from it. However, since the inference process is imperfect, the newer variables may assume values other than the ones implied by the central belief. We can formalize this generation process as:

$$x_i = x_h + \phi_i$$

That is, belief i is produced from belief h, with exogenous error ϕ_i which represents imperfections in the inference process. For example, if x_0 produces x_1, and x_1 produces x_2 and x_3, then $x_1 = x_0 + \phi_1$, $x_2 = x_1 + \phi_2$, and $x_3 = x_1 + \phi_3$ (and thus $x_2 = x_0 + \phi_1 + \phi_2$). We will refer to x_2 and x_3 as the "descendants" of x_1, and to all three of those beliefs as descendants of x_0. We will assume[4] that all the ϕ terms have a variance of ϵ and are independent of each other and x_0 , and that x_0 has a variance of 1.

Let us now imagine a very simple belief system consisting of only the central belief x_0 and two derivative beliefs $x_1$and x_2, so that $x_1 = x_0 + \phi_1$ and $x_2 = x_0 + \phi_2$ (see first diagram in Figure 1). All three of these beliefs are positively correlated. Furthermore, it can be shown[5] that $|cor(x_0, x_1)| = |cor(x_0, x_2)| = 1/\sqrt{1+\epsilon}$, while $cor(x_1, x_2) = 1/(1+\epsilon)$. Since $\epsilon >$

[4] For a complete formal statement of these assumptions and other details of the model, see Appendix A.

[5] Both of these quantities can be derived from equation (1.6) in Appendix A by substitution ($P = 0, K = 0, S = 1$ in the first case, and $P = 0, K = 1, S = 1$ in the second).

0, $|cor(x_0, x_1)|$ and $|cor(x_0, x_2)|$ are both greater than $|cor(x_1, x_2)|$. In this simple case then, the central belief can be discovered by simply examining the sum of all pairwise absolute correlations for each variable, which we will call the *total constraint* of those variables:

$$totalcons(x_i) = \sum_{j \neq i} |cor(x_i, x_j)|$$

The central belief will have the highest value of $totalcons(x_i)$. In empirical problems where the derivative beliefs are hypothesized to be closely related to the central belief, total constraint is the simplest and most intuitive centrality measure to examine. In fact, if the theory under examination proposes a simple one-step belief derivation from origin to outcome, this simple sum of absolute correlations would provide a tight methodological fit.

[Figure 1 about here]

Now consider a slightly more complex case. Let us imagine that six more beliefs are added to the system: x_{11}, x_{12} , x_{13} and x_{21}, x_{22} , x_{23} are derived from x_1 and x_2, respectively, so that $x_{1i} = x_1 + \phi_{1i} = x_0 + \phi_1 + \phi_{1i}$ and $x_{2i} = x_2 + \phi_{2i} = x_0 + \phi_2 + \phi_{2i}$ for $i = 1,2,3$ (see second diagram in Figure 1). Perhaps surprisingly, the central belief may no longer be the variable with the greatest total constraint. Some straightforward (if tedious) algebra can be used to show that $totalcons(x_0) > totalcons(x_1)$ only if ϵ is less than approximately 0.98. Thus, even with only two generations of derivative beliefs, the central belief may already not be the most highly correlated belief in the sample. The accumulation of error variance introduced by imperfect inference can "swamp out" the variance of the central belief. Put another way, total constraint is too local a feature of the belief network to correctly identify the central belief.

Fortunately however, this same accumulation of error variance can be used to locate the central belief even in very spread out belief systems. Our method proceeds from a simple intuition. Many center-periphery accounts of ideology describe the central belief as being the "glue" (Converse 1964) that holds together the disparate parts of the belief system. That is, the central belief is what enables coherent stances to exist across the relatively disconnected domains like environmental protection and gay rights (Converse 1964; Lakoff 2002). By this logic, the center may not be the most constrained belief, but it should be the "broker" (Burt 2004) possessing the most unique and valuable pattern of constraint. Below we formally demonstrate that this intuition can be used to find the center of such a belief system.

First, we need to introduce the notions of tie length and path length. We define the *length* of tie T_{ij} to equal 0 if $i = j$, and otherwise be:

$$|T_{ij}| = |T_{ji}| = \frac{1}{cor(x_i, x_j)^2}$$

The network defined by such ties is a symmetric, complete weighted network of the kind that can be analyzed by many software packages.[6] We will be interested in analyzing the *paths* between pairs of nodes in this network. A path Λ_{kl} between beliefs $x_k, x_{l \neq k}$ is as an ordered set of connected ties. The length of the path is the sum of the tie lengths it contains: $|\Lambda| = \sum_{T_{ij} \in \Lambda} |T_{ij}|$. The shortest path between any two nodes x_i and x_j (that is, the path with the lowest value of $|\Lambda|$) is termed their *geodesic*.[7] We use the term *transverse* to describe ties or paths between two nodes

[6] The exponentiation of the correlation coefficient is referred to as "soft thresholding" (Zhang and Horvath 2005). It is a standard technique in the analysis of correlation networks, and is used to dampen the effects of statistically insignificant correlations.

[7] Note that the length of a path is the sum of tie lengths that compose it, as opposed to the simple *count* of ties as would usually be the case in networks with unweighted ties. Since the observed tie lengths are continuous random

which have x_0 as their only common ancestor. We assume that most geodesics in the belief system are transverse.[8]

Let us return to the belief structure depicted in diagram 2 of Figure 1 and consider the transverse tie $T_{11,21}$. Due to accumulation of error, this tie will be relatively long: $|T_{11,21}| = (1 + 2\epsilon)^2$. The tie lengths $|T_{11,0}|$ and $|T_{0,21}|$, however, will be significantly shorter: in fact, $|T_{11,0}| = |T_{0,21}| = \sqrt{|T_{11,21}|}$. Thus, as long as $\epsilon > 0.5$, the direct path $\Lambda_1 = (T_{11,21})$ will be *longer* than the indirect path $\Lambda_2 = (T_{11,0}\ , T_{0,21})$,[9] which may come as a surprise to those of us used to spending our lives in Euclidian space. This indirect path is indeed their geodesic.

In Appendix A, we derive a general formal model of geodesics in such belief systems. In brief, we prove that, in general, these geodesics share a uniform structure (see Theorems 3 and 4). We use this structure to derive an algebraic formula for their length, which lets us apply standard calculus optimization techniques to find the nodes these geodesics pass through (see Theorem 5). We show that, in general, the transverse geodesics connecting any two k^{th} generation beliefs will consist of more than one tie (i.e., be "*non-trivial*") whenever $k > 1/\epsilon$ (see Corollary 6A). In other words, after the number of generations in the belief system exceeds $1/\epsilon$, no generation of newly added transverse nodes will be connected by single-tie geodesics.

At the end of Appendix A, we arrive at the key finding of this reasoning, which is the strong and persistent "central pull" of the belief system. Whenever transverse geodesics increase from a single tie to two or more ties, these intervening ties take them closer to the center of the system. Even though further generations of beliefs added to the system will grow less and less correlated with x_0, the shortest paths connecting them will still pass close to the center. As we prove in Theorem 6, every single non-trivial transverse geodesic passes through either x_0 or through a node highly correlated with x_0 ($|r| \geq 0.74$). Thus, in the absence of interference from highly correlated nodes (collinearity), every one of such geodesics will pass through the center. And, even if such highly correlated nodes exist, no single one of them will generally lay on enough geodesics to be confused for the center.[10]

The formal proof thus confirms our informal intuitions. The center of a large belief system (see Figure 1C) may be identified by finding which belief lies on the greatest portion of geodesics in the network of squared correlations—i.e., the node with the highest "*shortest-path betweenness*" (Freeman 1978). If *M* is the total number of beliefs in the system and $\mathbb{1}(x) = 1$ if condition x is true and $\mathbb{1}(x) = 0$ otherwise, then the betweenness of node x_k is:

$$betweenness(x_k) = \frac{\sum_{i \neq j} \sum_{T_{ab} \in \Gamma_{ij}} [\mathbb{1}(a = k) * \mathbb{1}(a \neq i)]}{(M-1) * (M-2)/2}$$

---

variables, it is practically impossible that two distinct observed paths should have exactly equal lengths. This allows us to use a simpler notion of geodesic than is usual for networks of unweighted ties, which need to deal with the possibility that multiple equally short paths may exist between the same two nodes.

[8] This is always the case unless a single first-generation node counts half or more of all the descendants of x_0 as its own descendants. Intuitively, this assumption can be understood as a prohibition against highly "lopsided" networks.

[9] See Corollary 2A in Appendix A for proof of this statement.

[10] The geodesics that bypass x_0 can only do so via a transverse tie T_{ab}, where *a* and/or *b* are highly correlated with x_0. However, two geodesics cannot use the same tie for their "shortcut" unless the endpoints of one geodesic are related to the endpoints of the other (see Corollary 6B). In other words, while $x_0$lays on geodesics between nodes in all branches of the system, any given "shortcut" node will lie only on geodesics between *particular* branches of the system. Thus, though the possibility of such shortcuts means that users of this method should exercise caution in the presence of multicollinearity, it does not appear likely to interfere with finding x_0.

The numerator is the count of geodesics which pass through x_k, while the denominator is the number of pairs of beliefs not including x_k (see Wasserman and Faust 1994:189–191). As expected of a proportion, this quantity varies from 0 to 1.

By the same logic, Freeman's (1978) index of *betweenness centralization* can be used to measure the extent to which the belief network as a whole possesses a single, well-defined center. The centralization of a network is the sum of pairwise differences between the centrality of the most central node and the centrality of each other node, all normalized by the maximum possible value such a sum could obtain in any network of M nodes (Wasserman and Faust 1994:176). This index achieves its maximum of 1 when one belief in the network has a betweenness of 1, and every other belief has a betweenness of 0. It achieves its minimum of 0 when every belief has exactly the same betweenness centrality.

The centrality and centralization indexes introduced above can be used determine which belief lies at the center of the system, and how much more central it is than the rest of the network. However, they provide no basis to judge whether the difference in centrality is robust to statistical variation—i.e., whether or not its position at the center of the network is "statistically significant." We will use a non-parametric bootstrap to produce estimates of this statistical robustness. In each iteration of the bootstrap, we will resample the respondents (rows) in the survey dataset with replacement, construct a correlation network for the resample, and finally recalculate the betweenness indexes for this network. We use these results to estimate the confidence intervals for each variable's centrality.

Our primary reason for constructing these confidence intervals is to determine whether any given node is reliably more central than others in the network. Since confidence intervals constructed from raw values can yield a misleading picture of this comparison, we will compare the distributions of *relative* rather than absolute betweenness centrality scores.[11] To calculate the relative centrality of the nodes in each bootstrap sample, we first calculate their betweenness centralities, and then divide the centrality of each node by the maximum centrality for each sample. Thus, if a network of K beliefs contained three beliefs a, b, c with betweenness centralities of 0.80, 0.40, 0.08, and no node in the network had a centrality higher than 0.80, their relative centralities would equal 1, 0.5, and 0.1. A node with an absolute betweenness of 0 also has a relative betweenness of 0. To avoid confusion, we will adopt a convention of reporting absolute betweenness scores as proportions, and relative betweenness scores as percentages.

DATA AND ANALYTIC STRATEGY

With this model in mind, we can test competing center-periphery accounts of ideology by applying shortest-path betweenness to a correlation network of survey responses. In order to ensure comparability with previous work, we construct the network from the same 2000 American National Election Study dataset Barker and Tinnick (2006) used to argue for the

[11] Imagine, for example, a belief network of K beliefs contained nodes a and b and that this network was resampled 200 times, with a having the highest centrality in the network in all 200 resamples. Further imagine that, in the first 100 resamples, a has a betweenness of 0.8 and b has a betweenness of 0.7. In the second 100, a has a betweenness of 0.9, and b has a betweenness of 0.81. Thus, in all 200 resamples, a is more central than b. However, the 95% confidence range for the raw centrality of a would then be [0.8,0.9], and for b would be [0.7,0.81], indicating that their centralities are not significantly different from each other. On the other hand, the relative centrality measure we introduce here would produce confidence intervals of [1,1] for a and [0.875,0.9] for b, thus capturing the fact that a has a reliably higher relative centrality than b.

central role of parenting values using linear regression. This dataset is also a good fit for our analyses because of the wide diversity of political attitude items it contains.

Barker and Tinnick use three items to measure parenting values. All of these items start with the stem "Although there are a number of qualities that people feel that children should have, every person thinks that some are more important than others. I am going to read you pairs of desirable qualities. Please tell me which one you think is more important for a child to have." The first one offers the response options "independence" or "respect for elders", the second "curiosity" or "good manners", and the third "being considerate" or "being well-behaved." These items let Barker and Tinnick distinguish between respondents who prefer the independent, curious and considerate child of a nurturant parent, and those who favor the respectful, well-mannered and well-behaved child of a strict parent. In their analyses, they demonstrate with fifteen regression equations that these parenting values predict a variety of political attitudes (values and issue positions) net of roughly two dozen demographic and attitudinal covariates.

Because Barker and Tinnick use regression analysis, they must make some arbitrary decisions about which beliefs to treat as dependent variables and which to treat as controls. They ultimately classify 15 variables as outcomes, and predict each of them using a separate regression model. Although the BNA method of course requires us to decide which variables to include in the network, we do not need to decide which beliefs are causes and which are effects. BNA also allows us to collapse the 15 distinct models used by Barker and Tinnick into one. To decide which of the resulting variables to retain for our model of subjective political beliefs, we followed the rule of thumb established by Alwin (2007): a question is "factual" (non-subjective) if the answer can be verified against objective records. Thus demographic and behavioral questions tend to be factual, while beliefs, attitudes, values and many self-descriptions are not (Alwin 2007). In addition to removing non-subjective questions, we also dropped the three-variable "need for cognition" scale because the questions it contained did not pertain to politics. The remaining 46 variables are summarized in Table 1.

[Table 1 about here]

We also made a number of different methodological choices from Barker and Tinnick. First, as Layman et al. (2007) also point out in their unpublished critique, roughly a third of the respondents volunteered the answer "both" to at least one of the parenting questions. Barker and Tinnick treat these responses as missing data, and drop the respondents from the sample. Since doing so biases the sample towards those who have a strong opinion on parenting, this procedure may artificially inflate its apparent importance. Thus, instead of dropping these respondents, we treat these variables as ordinal instead of binary, and code "both" as the middle of three values. We follow the same procedure with other variables where respondents frequently volunteer responses indicative of ambivalence or indifference. Finally, since our method uses only the pairwise correlations between attitudes, we deal with other missing values by pairwise deletion.

The second pertinent issue with Barker and Tinnick's analysis is their use of scales with low reliability scores. We found that many of their scales have Cronbach's alpha values below 0.6, and some far below 0.5 (estimated via polychoric correlations),[12] which is substantially

[12] Throughout this paper, we always compute polychoric correlations between ordinal variables, polyserial correlations between numeric and ordinal, and Pearson's correlations between numeric variables. This includes the correlations we use to estimate all factor loadings. Note that correlation measures implicitly assume that the pairwise relationships between the latent and/or manifest variables in our data are predominantly linear in character. In Appendix C, we compare these correlations with non-parametric measures of non-independence between pairs of variables based on entropy and mutual information. Our results confirm that this assumption is justified.

below accepted levels. For example, their scale for gun control / crime policy contains a question about gun ownership rights, a question about federal spending on crime prevention, and a question about support for the death penalty. This scale has a Cronbach's alpha of 0.41, suggesting that these items may not be closely related to a single underlying concept. The scales representing equal rights, equal opportunity, abortion, and defense spending also have alphas below 0.6.

Scale construction can be used to remove the variance that stems from the response error associated with individual items (Ansolabehere, Rodden, and Snyder 2008). However, scales constructed of weakly related items could instead remove large amounts of non-error variance, producing composite variables that are instead more weakly correlated than the component variables. We thus retain only those scales where all the items have pairwise polychoric correlations of 0.6 or above. [13] In those scales where some items are correlated above this threshold while others are correlated below, we construct the scale using only the strongly correlated items. We include all the remaining items as separate variables. This includes the three parenting items, which have pairwise polychoric correlations of 0.47, 0.10 and 0.44. We also replicate our primary analysis with these three items joined into one scale, and find that it does not affect the substance of our results (see Appendix D).

We proceed with the analysis as follows. First, we examine the belief correlation network we constructed from the full 2000 ANES dataset. To test the moral politics and social constraint accounts, we compare the relative centralities of parenting and ideological identity. We then investigate the robustness of our findings using bootstrapping. We resample the respondents of the survey dataset to demonstrate the reliability our findings to sampling error. As an extra robustness check, we then resample the *variables* to examine the robustness of our findings to the specifics of variable selection. After addressing these methodological concerns, we turn to the theoretical challenge presented by potential population heterogeneity. We create 16 partitions of the survey population along major demographic and cultural variables, yielding 44 subsamples corresponding to various social groups (e.g., women, middle-income respondents, etc.). We compare these across different dimensions of belief structure. Finally, we analyze three subgroup belief networks in more detail to search for exceptions to the primary pattern.

RESULTS

Belief Network for full ANES sample

We depict the belief network we constructed from the 2000 ANES dataset in Figure 2. We used darker and thicker lines to represent stronger correlations, omitting correlations below $|r| = 0.15$ for legibility.[14] The force-directed plot reveals a network structure with a sparse periphery and relatively densely connected groups of nodes near the core. Visually, the densely connected core appears to contain two groups of variables, which we label "A" and "B" on the diagram.[15] The bulk of the items that make up group A have connections to the social welfare agenda that has divided the two major U.S. parties since the New Deal (Carmines and Layman 1997), including items on affirmative action and government efforts to redress inequality. It also contains

[13] We replicated our primary analysis with scales constructed at other thresholds. The substantive findings remained the same when other thresholds were used.

[14] The omission is for visual purposes only. All analyses are based on the full network with no ties omitted.

[15] The node groupings we discuss here have a conceptual resemblance to partitions produced by modularity maximization. Modularity analyses also suggest that the belief groups we label A and B likely belong to two different modules. However, we found the modularity results for this network to be unreliable (see Appendix B). To avoid creating an undue impression of accuracy, we report this informal visual analysis of group structure instead.

questions on the size and scope of government activity (e.g., regulation of the environment, gun control.) The group labelled "B," on the other hand, contains many items that correspond to the "New Left" issue agenda which became part of the mainstream political discourse in the late 1960s and early 1970s (Carmines and Layman 1997), such as items on abortion, gender equality, and gay rights. It also contains items concerning religious identity and moral worldview.

[Figure 2 about here]

The first column of Table 2 shows the centrality estimates for the individual variables. The "absolute" column shows that the centralities range from 0 for the least-central nodes to 0.35 for the most central (overall centralization: 0.33). The variables Parenting 1 and Parenting 3 both have centralities of 0, as do eight other nodes located near the periphery. The remaining parenting variable, Parenting 2, has a centrality of 0.07. Thus, the three parenting variables combined lie on only 7% of the geodesics.[16] The network plot in Figure 2 provides context for this low centrality. Except their ties to each other, their strongest ties are to items concerning religious identity and abortion that make part of group B. All of their ties to the social welfare or limits of government items that make up group A were too weak to depict. They thus appear to occupy a peripheral position near the edge of group B, far from the network's center.

[Table 2 about here]

Ideological identity, on the other hand, can be found near the middle of the plot in Figure 2, between the groups we labelled A and B. Its position between the two relatively densely connected clusters appears intuitively central. The betweenness scores confirm this visual impression. Its betweenness centrality of 0.35 makes it the most central node in the network. It lies on five times as many geodesics as all three parenting variables combined.

The second- and third- most central nodes are Limited Government and Gay Rights. They are positioned near the visual centers of A and B, respectively. Their respective centralities of 0.12 and 0.10 indicate that they are only roughly one third as central as ideological identity. They thus do not appear to occupy the same brokerage position as ideological identity. We therefore find that ideological identity is the clear center of this belief network. These results are consistent with the theory of social constraint, but not with the theory of moral politics.

Potential Sources of Error

<u>Sampling Variability.</u> Thus far, we have found that ideological identity is the most central node in this network. We now use the non-parametric bootstrap to establish the statistical significance of this finding. In each of the 1000 iterations of the bootstrap, we drew a sample of N=1543 ANES respondents with replacement, and followed the same belief network analysis procedure as above to create a set of betweenness estimates. We then used these scores to estimate the 95% confidence intervals for the relative betweenness centralities of each belief, which we report in the "Relative" column of table 2. The leftmost and rightmost endpoints of each error bar correspond to the 2.5th and 97.5th percentiles of the estimate, respectively, and the black circles represents the means (also printed numerically next to each error bar). Twenty five of the forty beliefs in this column have confidence intervals starting at 0%, indicating that their centralities are not statistically distinct from zero. Two of the parenting variables (1 and 3) are among these twenty-five low-centrality nodes. Their mean relative centralities are 1% and 0%, respectively. The remaining parenting node (2) has a mean relative centrality of 20%, indicating that, in an average iteration of the bootstrap, it is roughly 1/5th as central as the most central node.

---

[16] An enumeration of these geodesics reveals that they all end in either Parenting 1 or Parenting 3, which is consistent with the visual intuition that Parenting 2 serves as a "gatekeeper" (Freeman 1980) for these further removed nodes and nothing else.

The mean relative centrality of Limited Government is 34%, which indicates that, on the average, it lay on roughly 1/3rd as many geodesics as the most central node. It is again the second most central node in the network. Gay Rights and Equal Rights 1 have relative centralities of 27% and 25%, respectively, which places both at roughly 1/4th of the centrality of the most central node, and make them the 3rd and 4th most central nodes. The confidence intervals for these three nodes begin at 11%, 6% and 7%, and extend up to 67%, 55% and 54%, respectively. These confidence intervals are wide and overlapping, as are those belonging to most other nodes.

The reliably central position of Ideological Identity stands in stark contrast with the largely undifferentiated centralities of the other variables. Its mean relative centrality of 100% shows that, in the average run, it is the most central node. Moreover, the lower and upper ends of its confidence interval are also at 100%, which indicates that its relative centrality does not significantly deviate from the maximum. In fact, when we examined the full range of the 1000 bootstrap iterations, we found that it was the most central node in *every* iteration. The dominant role played by ideological identity thus appears both substantively significant and remarkably robust to statistical variation ($p \approx 0$). Our analyses are thus consistent with the social constraint view that ideological identity lies at the core of the system of political attitudes.[17] In contrast, they offer no support to the moral politics view that parenting attitudes play a central role in the system of political views.

Variable Selection. The analyses we describe above made use of a set of 40 belief variables constructed from the 2000 American National Election Study dataset. While this collection of political attitude items is large and seemingly comprehensive, its representativeness of the domain of politics as a whole cannot, of course, be guaranteed statistically. Given that it is impossible to enumerate the full set of political attitudes, we cannot rule out the possibility that the belief network we examine may feature too many beliefs from some parts of the unobserved belief network, and too few from others. Since prior work has demonstrated that betweenness centrality estimates may be unstable to changes in the node set (Zemljic and Hlebec 2005), this raises the possibility that our findings may be skewed by particular details of the set of variables we included in our model.

To rule out this possibility, we again turn to resampling. In the preceding analyses, we resampled the *rows* (respondents) of the 2000 ANES dataset to demonstrate that ideological identity is reliably central in the face of fluctuations in pairwise correlations (i.e., tie strengths). To examine reliability to fluctuations in the set of beliefs included in the analysis (i.e., the node set), we now resample its *columns* (items). Since including two copies of the same variable in one betweenness analysis would not yield meaningful results, we resampled the item set using an "*m* out of *n*" resampling scheme (Bickel, Götze, and Zwet 2012). In 2000 additional resamples, we dropped different 12-belief subsets from the network,[18] and analyzed the networks consisting of the remaining 28 beliefs. Since the number of ties in a fully connected network is roughly proportional to the square of the number of nodes, each of these resamples retained only 351 out of the 780 ties (49%) that composed the original network. These resamples were thus highly

[17] As an additional robustness check, we repeated this bootstrapping analysis with the three parenting variables joined into a single scale. In 1000 bootstraps of the resulting 38-variable network, we found that this parenting scale was on average no more central than the parenting variables were individually (see Appendix Table D2). The average absolute centrality of the parenting scale was 0.001, which is *lower* than the absolute centrality Parenting 2 had in the main sample (0.06). The centrality of ideological identity remained unchanged.

[18] We also replicated this analysis with k = 2,3,7,10,13 and 15 nodes dropped from the network. As could be expected, lower values of k result in smaller confidence intervals for our results, and vice versa. However, the substantive content of our findings was unaffected by these changes.

distinct from the original network as well as from each other. To distinguish this procedure from the first set of bootstraps, we will call the first set "row bootstraps" and this second set "column bootstraps."

We considered two distinct ways in which details of the node set could affect the results of our analysis. First, the apparent centrality of ideological identity could be exaggerated by eccentric features of the node set, such as the inclusion of structurally redundant nodes that dilute each other's centrality. Second, the presence of ideological identity may mask the centrality of another node that, in its absence, would have occupied an equally central position. We used two different column resampling procedures to rule out these possibilities.

We first examined whether the centrality of ideological identity drops when subsets of variables are removed from the analysis. To do this, we drew 1000 column resamples by first selecting ideological identity, and then adding a uniform random sample of 27 other beliefs drawn without replacement out of the remaining 39 beliefs. This yielded 1000 networks of 28 beliefs each. The relative betweenness centralities for this set of resamples can be found on the left side of table 3, under the title "Ideological Id. Retained."[19] These centrality results strongly resemble the ones previously depicted in table 2. As before, the 95% confidence interval belonging to ideological identity never declines from the 100% mark, indicating that it reliably remained the most central node throughout the full range of belief subsets analyzed. All other beliefs have lower relative centralities with confidence intervals that overlap one another, and remain reliably below that of ideological identity. Thus, the central position occupied by ideological identity appears robust even to dramatic changes in the set of beliefs we include in the analysis.

[Table 3 about here]

In the second column-resampling procedure, we instead omitted ideological identity from our belief set. We then drew 1000 uniform random samples of 28 beliefs each, again sampling from the remaining 39 beliefs without replacement. We examined each of these resampled 28-belief networks to determine whether any other belief node comes to occupy a reliably central position when ideological identity is dropped. We report these results in the right column of table 3. The 95% confidence intervals in this set are noticeably wider than previously, indicating that the nodes' relative centralities become substantially less stable when ideological identity is omitted. Moreover, in contrast with results in the left column of table 3, five distinct beliefs—limited government, gay rights, equal rights 1, party identification, and moral relativism—now have relative centralities statistically indistinguishable from 100%. Thus, in the absence of ideological identity, no other belief appears to occupy a robustly central position.

We also used our row (respondent) resampling procedure to examine the belief network that excludes ideological identity but includes all remaining 39 beliefs. To construct each of these additional resamples, we drew a uniform random sample of 1543 rows with replacement from a dataset that excluded ideological identity but retained all other columns. We then performed betweenness analyses on each of the 1000 resulting 39-belief networks. We found that the resulting centrality distribution again consisted of overlapping confidence intervals with no clear central variable, and overall resembled the results of the second set of column bootstraps

[19] When we calculated the centrality distributions for any one node, we simply omitted all the cases that this node was dropped from the analysis. Thus, those samples where a variable was absent have no effect its centrality scores.

we report above (see left column of Appendix Table D2). These results are again consistent with ideological identity occupying a uniquely central position in the belief network.[20]

Population Heterogeneity

Our analyses thus far have produced substantial evidence consistent with the theory of social constraint and inconsistent with theories (like Lakoff's) that put a different concept at the center of a belief system. The method we used, however, rests on the assumption that the population-wide organization of attitudes is produced largely through a single dominant process that does not vary systematically for subgroups of the population. That is, it assumes a single network of which each person's beliefs are a noisy realization.

This single-network view is shared by both moral politics theory and social constraint theory, as well as many other center-periphery accounts of ideology (e.g., Jost et al. 2003). However, a number of recent sociological treatments (Baldassarri and Goldberg 2014; Boltanski and Thévenot 1999; Thornton et al. 2012; Achterberg and Houtman 2009; van Eijck 1999) instead assume substantial heterogeneity by arguing for the existence of many different "logics" that organize beliefs differently for different subgroups. We develop these contrasting views of heterogeneity in more detail below, and then test them empirically on the ANES dataset.

In the social constraint view, individuals construct their belief systems by acquiring pieces of political content from attitude producers, which they select using their political identity as a heuristic. However, individuals vary greatly in the extent to which they care about politics and attend to political information flows (Delli Carpini and Keeter 1996). Moreover, the taste for political communication, much as tastes for other cultural products, is highly socially patterned, with some social groups systematically further removed from the institutional field of organized politics than others (Bourdieu 1984). Those who consume enough informational flows may learn the partisan pseudo-logics which make certain positions on far-flung issues such as global warming and gay rights entail views on, e.g., military policy and healthcare spending—attitudes which may otherwise be mostly unconstrained. By this logic, different social groups should then vary in the *amount* of belief system organization they exhibit—a point that has been demonstrated in much empirical work (Converse 1964, 2000)—but they should not vary in the *logic* of organization of their beliefs.

Baldassari and Goldberg propose a contrasting view of population heterogeneity, arguing that "the heterogeneity of political belief systems does not simply derive from differences in levels of political sophistication", i.e., amount of belief organization, "but in in individuals' social identities: people with different sociodemographic profiles structure their political preferences in systematically different ways." (2014:78). They thus see demographic positions as laying at the root of differences in both the amount and the logic of belief organization. Under this view, sets of political positions that are perceived to be coherent from the perspective of one population may appear to be contradictory from the perspective of another. For example, though support for environmental regulations is generally negatively correlated with support for gun ownership in the population as a whole, we can imagine, say, a sub-population of hunting enthusiasts where environmental protection and gun ownership rights go hand in hand. The practical implication of this argument is that attitudes that are positively correlated in one

[20] We additionally reanalyzed all 2000 28-node column bootstraps using multivariate linear regression, with individual simulation runs as observations, network centralization as the outcome variable, and 39 dummy variables indicating which variables were dropped as the predictors. As expected, we found that the presence or absence of ideological identity was by far the strongest predictor of centralization.

subgroup may be negatively correlated in another. And, if two such evenly-sized populations are mixed together in a single sample, the two patterns may simply cancel out, yielding two variables that appear uncorrelated in the full sample.

To compare these views of heterogeneity empirically, we constructed separate belief networks for 44 different subpopulations, which we produced by partitioning the population 16 times along different demographic dimensions (see Table 4 for descriptive statistics and Appendix E for details of variable coding). Prior work has found that various forms of social status are predictive of average belief constraint (i.e., mean absolute correlation between beliefs), with higher-status groups generally exhibiting a higher level of constraint than lower-status groups (see review in Gordon and Segura 1997). For this reason, we examine 9 dimensions that are associated with major social and economic cleavages in contemporary American society. These dimensions are respondent's income bracket, occupational category, class self-identification, education, gender, age, race (black, not black), Hispanic status, and religious denomination.[21]

We include 6 further dimensions to tap cultural cleavages. These measure whether the respondent's parents are foreign- or US-born, whether the respondent attends church, the type of populated place where the respondent resides (large city, rural area, etc.), and whether or not this location is in the South-Eastern United States. Because of the importance of child-rearing to the theory of moral politics, we also partition respondents by the number of children they have (zero, one or more). We also include an index of the respondent's factual knowledge about politics,[22] which is frequently used as a measure of the respondent's involvement with the field of organized mainstream politics (Delli Carpini and Keeter 1996).

In their recent study of belief heterogeneity, Baldassarri and Goldberg (2014) use an inductive partitioning approach that does not require subpopulations with different belief structures to lie on different sides of a demographic divide. The demographic positions of the belief systems they locate, however, are central their interpretation of these results, as well as to arguing against their possible spuriousness.[23] They state that "sociodemographic

[21] Each respondent is assigned to exactly one subpopulation of each demographic dimension. Aside from the interaction between religious attendance and income we discuss below, we do not intersect these dimensions. The respondents with missing demographic data along a dimension are not included in any of the groups in that dimension. (They are only omitted from the analyses of dimensions where they have missing data, and are present for analyses of all other dimensions.)

[22] To measure the information levels of the respondents, we made use of the factual political information quiz included on the ANES (8 questions, 0 or 1 points each). See Appendix E for quiz questions. Since this quiz-based measure can confound knowledge with personality traits such as confidence and competitiveness (Mondak 2001, 2000), we also make use of the interviewers' subjective assessments how informed the respondent appeared (2 items, 0 to 4 points each). We summed these scores into an index that ranges from 0 to 16, and labelled the bottom third of the range (0-5) "low information", and the top third (11-16) "high information."

[23] Inductive heterogeneity detection can yield positive results even in the absence of any true heterogeneity of belief structure. Consider two unrelated attitudes A (Yes/No) and B (Yes/No) with no logics connecting any position on A to a position on B, and with all response pairs (A=Yes, B=Yes), (No, No), (Yes, No), (No, Yes) equally likely. This population can then be partitioned into two groups, with (Yes, Yes) and (No, No) respondents assigned to one group, and (Yes, No) and (No, Yes) to the other. A and B would then be positively correlated in the first group, and negatively in the second. Since we know, however, that A and B are simply unconstrained, it is incorrect to interpret this as evidence of alternate systems of belief organization; the result is completely artefactual. Inductively located heterogeneity thus requires external validation. Baldassarri and Goldberg seek it partly in demographic position: "while RCA allows us to identify groups of respondents that exhibit distinctive patterns of opinion, we cannot, with survey data alone, determine the underlying psychological processes that generate these patterns. Nevertheless, we can make reasonable assumptions about these causes and how they relate to people's location in sociodemographic

characteristics—particularly class and religiosity—account for this divergence in political belief systems" (47), and "nonreligious high earners and religious low earners […] occupy social positions that push them to take ideological stances that are seemingly contradictory" (69). They thus claim that, "if the overlap between people's class and religiosity has a bearing on how they combine their political preferences, then we should find that the interaction between the two explains how respondents combine their political beliefs" (69). We use our demographic heterogeneity analysis to test the validity of these claims.

Empirically, Baldassarri and Goldberg operationalize this social position via church attendance and income, arguing that, while high-income church attendees and lower-income non-attendees experience economic and moral pressures which are aligned, lower-income church attendees and high-income non-attendees face pressures that are at odds (69). Thus, they conclude that the former two groups would have traditional belief systems with consistently conservative or liberal attitudes, while the latter would hold liberal positions on economic issues and conservative positions on moral ones, or vice versa. They do not, however, test this hypothesis directly. To carry out this test, we constructed an "Economic and Moral Pressures" stratifying variable. We labelled higher-income church attendees and lower-income non-attendees as "Aligned Pressures," lower-income attendees and higher-income non-attendees as "Cross Pressures," and all middle-income respondents as "Neither".

Heterogeneous Logics. To examine whether different demographic groups display distinctive logics of beliefs organization, we contrasted the way the same pairs of beliefs are correlated with each other in different groups. For each subpopulation, we constructed a matrix consisting of polychoric, polyserial and Pearson's correlations, as appropriate. Since statistical noise can make weak correlations fluctuate around zero, we first took each belief network and removed from it all the correlations that were not statistically different from 0 at $p < 0.05$. Out of the 780 correlations in each network, this left a median of 577, or roughly three quarters (see first numerical column in Table 4).

We first compared each pair of mutually exclusive subpopulations. These are the populations that come from partitions along the same demographic dimension. For example, for the dimension "age", such groups are "under 40", "40 to 55", and "over 55", this yields three unique comparisons: "under 40" versus "40 to 55"; "under 40" versus "over 55"; and "40 to 55" versus "over 55." Within each of these 45 subpopulation pairs, we compared the directions (signs) of all the correlations which were significant for both groups. For example, out of the 780 unique belief pairs, 649 attained statistical significance in the male sample, 587 in the female sample, and 535 in both the male and the female subsample. If different groups indeed use substantially different logics to organize their political views, we should observe that these signs frequently point in different directions.[24]

The results of these comparisons can be found in Table 4. For example, out of the 535 correlations that were significant in both the male and female samples, only 2 correlations (0.4%) had different signs for males and females, with the remaining 533 pointing in the same direction. The table shows that the same basic finding occurs for all comparisons on class, parents' nativity, age, education, income, region, religion, occupation, and type of place. It also holds for "Economic and Moral Pressures," where 98.4% of the significant correlations retained the same

space" (2014:59). Below, we examine whether the stated demographics actually correspond to substantial differences in attitude structure, and find that they do not. This raises questions about potential spuriousness.

[24] We discuss the theoretical meaning of these sign comparisons in more detail in Appendix F.

sign between the "Cross Pressures" and "Aligned Pressures" groups. Overall, in 43 of these 45 group comparisons, 95% or more of the significant correlations pointed in the same direction. Even in the two comparisons which recorded the most extreme differences—between blacks and non-blacks and between high and low-information respondents—89.5% and 87.4% of the correlations still retained the same sign.[25] Overall, among all the 45 comparisons we carried out, 98.7% (median) of the correlations retained the same sign for both groups, with only 1.3% switching directions (IQR: 0.5% to 1.6%).

[Table 4 about here]

We then extended this analysis to each unique pairing of the 44 demographic subgroups, independent of the dimension used to create them.[26] There are 946 such group pairs. In 901 of the resulting comparisons (95.2%), 95% or more of the correlations had the same sign, and in 941 comparisons (99.5%), this same-sign proportion exceeded 90%.[27] Among all the group pairs, 99.6% (median) of the correlations had the same sign, with 0.4% switching directions (IQR: 0% to 1.2%). Given the much-bemoaned statistical noisiness of survey data, the constancy of sign is strikingly robust.

Our results therefore provide no evidence to support the assertion that different groups typically organize their beliefs according to different logics. Rather, we find that, even in the most contrasting of groups, the overwhelming majority of political attitudes "go together" in the same way. The issue positions that go together for one group—at least to a statistically significant extent—are very rarely opposed for another group. Heterogeneity in the organizing logic of political beliefs thus appears to be the exception rather than the rule.

In addition to this analysis of heterogeneity across demographic dimensions, we also conducted a more general heterogeneity analysis using mutual information (see Appendix C). Normalized mutual information ($\hat{I}_{ij}$) and squared correlation yield similar estimates of pairwise relationships between variables when these relationships are strong and linear. However, while correlation captures only linear relationships, mutual information is a general non-parametric measure of non-independence. As we demonstrate in Appendix C, $\hat{I}_{ij}$ can detect relationships between variables even in the presence of two subpopulations where the variables obtain opposite linear relationships. The same heterogeneity would cause the linear relationships to cancel each other out, yielding an overall correlation of zero. Thus, in the presence of such heterogeneity, $\hat{I}_{ij}$ and squared correlation should diverge. However, when we apply both $\hat{I}_{ij}$ and squared correlation to our data, we find that the two measures are instead mutually correlated at $r = 0.91$, indicating that they overwhelmingly vary in unison. These supplementary analyses thus also find no evidence of heterogeneous logics of belief organization.

Amount of Organization. We next examine whether the belief networks exhibit a heterogeneity in their *amount* of organization, as suggested by the theory of social constraint. There are two senses in which beliefs can vary in the degree of organization between groups. First, in different groups, pairs of beliefs (node dyads) can "hold together" to a greater or lesser extent. We have previously referred to this quantity as "mean constraint," which we operationalize as the mean of the absolute correlations between pairs of beliefs. Second, whole networks can vary in how much

---

[25] We examine these groups in detail in in the following section.

[26] E.g., while the previous analysis compared the category "Male" only to "Female", this analysis also compares "Male" to "Black," "Under 40," etc.

[27] All 5 of the remaining group pairs, for which between 87.3% and 89.8% of the correlations retained the same sign, again involved either low-information or African American respondents.

structure they exhibit. For social constraint and other center-periphery accounts, this network-wide property can be captured by betweenness centralization.[28]

We present an overview of these measures in Figure 3. Here, each of the 44 subgroup belief networks is represented with a point, the coordinates of which correspond to its centralization (x) and mean constraint (y). The figure reveals significant differences in the *amount* of belief system organization between different populations. The subgroups range from 0.02 to 0.09 in constraint ($\sigma = 0.01$), a difference of a factor of 4. They also range from 0.10 to 0.48 in centralization ($\sigma = 0.09$), or roughly a factor of 5. A detailed analysis of which properties of demographic groups predict greater or lesser belief constraint is outside the scope of this paper. However, we note that, consistent with prior work on group differences in political knowledge (Delli Carpini and Keeter 1996), higher status groups appear to generally possess higher belief constraint than the lower status groups on the same dimension. Network centralization follows a similar pattern.

[Figure 3 about here]

We previously drew on the social constraint account to predict that belief networks would either be centered on ideological identity, or have no discernible center at all. In figure 3, the networks where ideological identity is the most central node are marked with filled circles, while those with a different central variable are indicated with hollow circles. All of the hollow circles are clustered near the lower-left corner of the plot, indicating that all networks with high constraint or high centralization have ideological identity at their center. The relationship between network centralization and the centrality of ideological identity is plotted in figure 4. The Pearson's correlation between these two quantities among the 44 networks is $r = 0.93$. Thus, as expected, the bulk of the variance in network centralization ($R^2 = 0.86$) can be explained by the centrality of ideological identity. Taken together, the results reported in this section show that although demographic groups vary in the *degree to which* they are organized, they do not vary in the *way in which* they are organized.

[Figure 4 about here]

Comparison to Existing Work. Baldassarri and Goldberg (2014) have previously analyzed ANES data to claim support for the existence of different logics of belief organization—a claim contrary to ours. In Appendix F, we reexamine some key evidence they offer in support of their argument to clarify this disagreement. Their heterogeneity analyses of the eight ANES years each partitioned respondents into three groups, which they termed "ideologues," "agnostics," and "alternatives." They propose that agnostics follow the same logic of belief organization as ideologues, albeit to a lesser extent, whereas alternatives employ a wholly different logic. Their results show that, for ideologues, the average cross-domain beliefs correlations are strong and positive, while for agnostics they are indeed either positive but weaker, or are insignificant. However, even for alternatives, who are supposed to follow a different logic, only 3 out of the 48 average cross-domain correlations are actually negative—and, even in those rare cases, the average negative correlations are weak. The remaining correlations are generally either still positive but weaker than for ideologues, or are insignificant. The belief system of alternatives thus appears to overwhelmingly be a subset of the belief system of ideologues. Much like agnostics, alternatives follow some of this system's logics to their full extent and weaken or omit other logics, but very rarely actually introduce unique logics of their own. Thus, contra their interpretation, we argue that Baldassarri and Goldberg's results support our view that belief

[28] We can also measure it using row bootstrapping, as the portion of resamples in which the node with the highest overall centrality across the resamples occupies the most central position. We take this approach later in the analysis.

systems generally differ in the extent of organization, but not in its logic (see Appendix F for more details).

Given that agnostics, alternatives and ideologues thus appear to overwhelmingly follow the same logic of belief organization, it may seem surprising that the RCA algorithm used by Baldassarri and Goldberg identified them as separate groups. RCA is, after all, designed to find distinct patterns of correlation within the population. However, as we point out in Appendix G, RCA does not currently provide a goodness-of-fit statistic to indicate whether the groups it located actually differ significantly from one another in their belief organization. Furthermore, existing work indicates that the modularity maximization partitioning technique used by RCA may have a substantial bias towards detecting heterogeneity in the data even when none exists (see discussion in Appendix B), making this absence of a goodness-of-fit statistic especially problematic. In Appendix G, we show how multiple group analysis in structural equation modeling can be adapted to provide such a goodness-of-fit statistic for RCA. We applied this technique to the RCA results for the 2000 ANES, comparing the heterogeneity model with respondents partitioned into the three RCA-detected groups to a no-heterogeneity model where all respondents kept in a single group. Both the AIC and BIC goodness-of-fit indices greatly preferred the no-heterogeneity model over the model with RCA-based partitions. This fits with our argument that the RCA-identified classes do not actually follow different logics of organization (Appendix F), and again supports our view of heterogeneity over Baldassarri and Goldberg's (see Appendix G for more details.).

Belief Systems by Political Information and Race

When we contrasted the pairwise relationships between beliefs for different demographic subpopulations, we found that the direction of association does not typically vary, with 95% or more of the significant correlations having the same sign in 43 of the 45 subgroup comparisons we conducted. The remaining two comparisons were those between black and non-black respondents, and between respondents with high and low levels of political information. Black and non-black respondents differed on 10.5% of the correlation signs, and high- and low-information respondents differed on 12.6% of the signs. Though these differences are not large, these subgroups are nonetheless the most likely to contain evidence of an alternate basis of political belief organization from the mainstream. To explore this possibility, we examine whether an alternate belief occupies the central position in any of these subgroups.

Table 5 shows the relative betweenness centralities for high information, low information, and black subsamples. We omit the non-black subsample from this table because it contains 90% of all ANES respondents, and its pattern of relative centralities is nearly identical to the full population's. Following our row bootstrapping procedure, we resampled each of these three subgroups 1000 times to determine the distributions of relative centralities in each.

[Table 5 about here]

We plotted the belief network for the high information subsample in figure 5. The average constraint of this network (0.26) is visibly higher than that of the full population sample (0.16). Ideological identity again has a mean relative centrality of 100%, with the other beliefs occupying significantly less central positions (left column of table 5). In fact, the gap between ideological identity and the other variables has grown. The runner-up belief now has a mean centrality of 19%, as compared to 40% in the full sample. Of the remaining 39 beliefs, only 10 have confidence intervals that do not include zero, as compared to 15 in the full sample. The belief network for the high information sample thus in many ways appears to be an exaggerated version of the same structure we observed in the full sample, with ideological identity as an even more clearly defined central node.

[Figure 5 about here]

Since the high-information subsample follows the same general pattern of organization as the overall sample, we turn to the low-information group to search for evidence of an alternate structure (figure 6). Indeed, religiosity (75%) and Biblical literalism (69%) have the highest centralities in the network, whereas ideological identity (18%) occupies a relatively peripheral role. This leads to the intriguing suggestion that religion may play a more important structuring role among respondents further removed from the institutional field of mainstream politics—a topic we return to in the discussion. However, the central column of Table 5 is composed almost entirely of wide confidence intervals, indicating that this network exhibits little of the way of stable centrality structure. Every one of the seven confidence intervals extending to 100% also dips to 25% or below, indicating that the same nodes which occupy central positions in some iterations also occupy peripheral positions in others. This leaves the centralities of religiosity and Biblical literalism statistically indistinct from those of 37 of the remaining 38 nodes.

[Figure 6 about here]

The results for the black subsample follow a similar pattern (last column of Table 5), with religiosity (88%) occupying the most central position. As with the low-information network, however, the centrality distribution is characterized by wide confidence intervals that leave this centrality statistically indistinct from that of most other nodes. Thus, our analyses of the subpopulations that appeared the most likely to provide evidence for alternate systems of organization instead better support the conclusion that the political attitudes in these populations have no reliable center. To examine whether this result holds more broadly, we extended this analysis to encompass all 44 subpopulations.[29] We found that 11 of the 44 networks had statistically reliable centers at $p < 0.05$. In all 11 of these networks, the central position was occupied by ideological identity. Relaxing the reliability cutoff to $p < 0.10$ or even an unusually lax $p < 0.25$ increased the number of qualifying networks to 15 and 22, respectively. Nonetheless, all of these networks still had ideological identity as their most central node. We thus found no evidence of subgroup networks reliably centered on any node other than ideological identity.

DISCUSSION

In this paper, we developed Belief Network Analysis, a novel correlation network-based method for examining the structure of beliefs. We used this method to compare different theoretical accounts of belief structure. To focus our analysis, we used the same National Election Study data as Barker and Tinnick (2006), who employed regression analyses to argue that parenting models play a central role structuring political beliefs (Lakoff 2002). We found no evidence to support this claim, and found instead that ideological identity (liberal/conservative) is most likely to provide that organization. Since regression analyses of these data have, to our knowledge, been the only quantitative work to offer support for Lakoff's theory of moral politics, our results suggest significant skepticism towards this popular theory.

Our results are also not consistent with accounts that emphasize the heterogeneity of belief structures across different social groups (e.g., Baldassarri and Goldberg 2014). An analysis of 44 demographic subgroups showed that, at least in the domain of politics, there appears to be

[29] We estimated the relative centrality confidence intervals for each of the 41 remaining subgroups by drawing 250 row resamples of each (10250 resamples total), and then performing the same betweenness analyses as above.

a single dominant logic of belief organization. On different sides of the major demographic divisions we examined—even those which past work suggested delineate alternative ways of organizing political attitudes—we found little reliable evidence of beliefs fitting together in opposite ways. If "support issue A" implied "support issue B" in one subpopulation, it generally either implied "support B" in the other population, or else support for A had little relationship with support for B. Across all the subpopulations we examined, the cases where "support A" implied "oppose B" in the other subpopulation were very rare, with only 1.3% (median) of the significant correlations between beliefs switching signs in a typical comparison. We thus concluded that groups generally differ in the *extent to which* their attitudes are organized at all, but not in the *logic* around which they are organized. Additional information-theoretic analyses and reexaminations of results from prior work also supported this conclusion.

Our centrality and heterogeneity analyses were therefore both consistent with the view that political identity serves as a key heuristic in structuring political beliefs (Converse 1964; Mondak 1993; Zaller 1992; Goren et al. 2009). Such social constraint accounts hold that individuals acquire their attitudes via attention to information flows from political elites, which they select by using their political identity as a filter. In support of the social constraint view, we found that ideological identity occupied the most central position in the overall population—a result that was extremely robust to a variety of changes in the model. Our resampling analyses also showed that, in the absence of ideological identity from the model, the belief network simply appears uncentralized, which further highlights the unique structural position of ideological identity and is consistent with its role as the dominant organizing heuristic.

In our heterogeneity analyses, ideological identity also occupied the central position in every sub-population whose belief system had a stable center, leaving us with no reliable evidence of belief systems organized around anything but ideological identity. While low constraint between beliefs and high noise in attitude measurement mean that practically any possible combination of attitudes can be empirically observed, we found no reason to believe that these combinations represent an alternate organized belief system. Across subpopulations, the centrality of ideological identity and the overall level of organization in the belief system were closely correlated—that is, belief systems appeared organized to the extent that ideological identity served as their center. Consistent with prior work, it was the demographic groups with greater participation in the field of organized partisan politics (Delli Caprini and Keeter 1996) that had belief systems organized around ideological identity. However, for groups further away, no other organizing principle appeared to step in and fill the gap. Taken together, these diverse findings fit with the argument that that the logic of the political field may be the dominant organizing principle holding together the attitudes we term "political," and that political identity—that is, a position within this field—is the main conduit via which individuals acquire this organization.

To search for possible exceptions, we closely examined low-information and African-American respondents, who showed some evidence of using religious beliefs to structure their political views. This examination, however, produced no statistically reliable evidence of alternate organization. Though in these populations, religiosity-related variables were somewhat more central, and political identities were less central, the overall belief networks lacked any clear center. And, even in these two groups, more than 85% of the statistically significant belief correlations still retained the same sign as in the comparison group. Thus, even if these results were to be interpreted as evidence of an alternate, religion-based system of organization, such a system would appear to provide alternate organization for only a small subset of political attitudes.

One possible explanation for this finding comes from the partial endogeneity of the concept of "politics" to the political field. As some critical scholars have argued, attitudes may come to be classified as "political" when they concern issues that have become the subject of competition between political parties, social movements, or other recognizably political actors (Lee 2002). By this reasoning, political attitudes may be exactly those attitudes which actors in the political field wish to influence in the general population, and thus weave into their competing belief systems. By contrast, since issues of special interest to other fields (like religion) do not automatically come to be classified as political, the belief systems they produce may then systematize political attitudes only to the extent that their interests intersect with those of the political field. Viewed as political belief systems, they would likely appear incomplete.

Although our results agree with the basic tenets of the social constraint view, they deviate from existing accounts on the relative importance of different political identities. All social constraint accounts posit that individuals acquire their political views from elite opinion leaders. To most scholars of American politics, this has specifically meant political *parties* (e.g., Campbell et al. 1960; Carmines and Wagner 2006; Goren et al. 2009; Sniderman and Stiglitz 2012).[30] However, there may be theoretical reasons to doubt this focus on parties and politicians to the exclusion of other political identities. Many popular political commentators appear to flaunt their independence from the major political parties, and to affiliate themselves with ideological rather than partisan labels (e.g., "Although I am a 'conservative,' I'm not a 'Republican,' and there's a big difference" [Glenn Beck (2008)]). If identity is indeed primarily useful as the heuristic individuals use to evaluate political communication flows, then, *ceteris paribus*, the most relevant identities should be those used by the most visible communicators. In our analyses, we repeatedly found ideological identity to be more central than party identity, which supports this alternate account. It thus merits investigating whether existing work underestimates the importance of political identification in attitude formation.

Limitations and Future Directions

BNA allowed us to detect interesting patterns in survey data and to use these patterns to compare competing theories. The technique produced centrality and centralization scores that lent themselves to intuitively clear interpretations in terms of brokerage within belief structure, and this interpretation was further aided by network diagrams. The bootstrapping-based confidence intervals also provided clear measures of statistical significance and sensitivity to variable selection, and in our case indicated that the primary centrality result is extremely robust to both sampling error and changes in the model. In addition to political beliefs, our approach can be applied directly to other cultural domains that can be reasonably approximated by the center-periphery model. The development of cultural tastes may be one such area for investigation.

However, like all data analysis techniques, BNA has limitations. Most important, it gains its leverage from making some simplifying assumptions about its target domain. We focus our investigation on theoretical accounts of belief structure in which some beliefs are central, and others are produced from them through a noisy inferential process. The logic of our method derives from the fact that this process of belief derivation should leave behind a correlation

[30] Thus, for example, Goren and colleagues (2009) open their work with the claim that "party identification represents the most stable and influential political predisposition in the belief systems of ordinary citizens" (2009:805), while Leege and colleagues (2009) argue that voter preferences are "largely the products of ambitious politicians seeking issues that will carry them to victory."

network where central items act as brokers uniting otherwise disparate parts of the system—a consequence which we prove formally in Appendix A. Our proof, however, rests on a stylized model of the belief acquisition process which strips away complexity in order to leave a parsimonious structure suitable for formal investigation. Future work should examine the behavior of the model when assumptions are relaxed and complexities re-introduced. Such work could potentially greatly increase the applicability of the method we described here.

While the belief generation process we capture in our model is in line with many theoretical accounts of belief structure, other accounts disagree. For example, some accounts assume broad heterogeneity in the belief structures of different subpopulations, while others envision deeply non-linear relationships between the belief variables. We examined both of these conflicting accounts in this paper, and provided evidence that supported our model over these alternate accounts in the political domain (see also Appendixes C and G). We hope that, by clearly laying out our assumptions and formally deriving our method from them, we ease the task of potential challengers wishing to dispute or extend this model of belief structure. We include our full formal reasoning in Appendix A so that others can build on this work, perhaps extending it to domains where the current model's assumptions do not hold.[31]

Another current limitation of our method is that it is a test of structure rather than causality. Though we present many diverse pieces of evidence consistent with belief system generation by social constraint, we do not uniquely identify this causal process. More broadly, while many other approaches to belief structure aim to ascertain the causal precedence of some beliefs over others, our analysis uses a cross-sectional dataset and is thus unsuitable for such questions. Instead, our analysis focuses on the *structural significance* of beliefs within the system. Since theoretical accounts of belief systems make both causal and structural claims, both approaches are necessary. We show that, among the beliefs we analyze, ideological identity is unique in occupying a structural position at the center of the network; even in its absence, no other belief comes to occupy a reliably central network position. We thus rule out accounts that would place parental values, limited government, symbolic and/or explicit racism, or any of the culture wars issues at the center of the belief structure. In our heterogeneity analyses, we also show that it is not the case that different populations have belief systems centered on different attitudes, which points against accounts that conceive of different sub-populations as achieving attitude organization through drastically different processes.

Our structural results provide a complement rather than a replacement to those analyses which have used experimental or longitudinal data to show that, for example, changes in political identification can cause changes to other items in the political belief system, that partisan source cues play an outsize role in how individuals reason about political information, and that identity changes lead to changes in political values but not vice versa (e.g., Bartels 2002; Cohen 2003; Goren 2005; Goren, Federico, and Kittilson 2009; Mondak 1993; Mondak et al. 2004; Zaller 1992; but see Johnston 2006). We rest on these existing findings to argue that political identity is not simply a post-hoc label attached to constellations of beliefs acquired through another unknown process—an alternate mechanism that could, under some circumstances, place

[31] Future work can also extend BNA to examine other structural features of belief systems. For example, while we have focused on centrality and centralization, many of the theoretical accounts we reviewed here also suggest that some beliefs may cluster into densely-tied subgroups or "communities" that are relatively weakly connected to the rest of the network. These structural features can be examined using a network partitioning algorithm. While Newman's (2006) modularity maximization is the conventional approach to detecting this kind of community structure, our analyses documented some apparent problems with the method that preclude us from recommending its use here (see Appendix B).

ideological identity at the center of the belief network, and that we cannot ourselves rule out. In the future, it may be possible to combine some of the strengths of these approaches by using panel data to construct a network based on within-person belief changes. Since BNA is based on the well-developed foundation of correlation and network analysis, it may be able to benefit from the wealth of knowledge developed in these domains to quickly make these and other methodological advancements.

CONCLUSION

We believe that Belief Network Analysis represents a novel and important contribution to the study of belief systems. Building a belief network out of correlations enabled us to draw on the rich methodological and theoretical toolkit of network analysis to construct intuitive measures of structural features theorized in the literature on belief systems. Our specific empirical results were broadly consistent with the conception of political identity as the dominant heuristic for acquiring attitudes. They also provided considerable evidence against Lakoff's theory of moral politics, and suggest the need for a degree of skepticism toward sociological accounts which assume substantial social heterogeneity in logics of belief organization, at least in the political domain. While this analysis focused on political attitudes, the techniques offered here are general and can be applied to other domains. Since no single methodological approach alone can provide sufficient understanding of culture's complex structures, our hope is that BNA will be joined by other inventive methodologies in a renewed effort to tackle "the biggest unanswered question in the sociology of culture" (Swidler 2001:206)—how some cultural elements structure others.

TABLES FOR CHAPTER 2: —

TABLE 1

VARIABLE MNEMONICS, SURVEY ITEMS (RIGHT-ALIGNED IN ITALICS), AND DESCRIPTIVE STATISTICS

BELIEF VARIABLE	SHORT NAME	TYPE AND RANGE†	MEDIAN OR MEAN‡	1ST - 3RD QUART. OR STDEV	*ITEM (ITALICS)*
Abortion Legal					*In what cases should abortion be legal: never, if clear need, if rape/incest/risk to life, or always?*
	Abort. 1	Ordinal (4)	If there is clear need	If rape/incest/life risk - Always legal	
Abortion for Teens					*Favor a law to require "girls under age 18 to receive her parent's permission before she could obtain an abortion?"*
	Abort. 2	Ordinal (4)	Favor strongly	Favor strongly - Favor not strongly	
Moral Relativism					*"The world is always changing and we should adjust our view of moral behavior to those changes"*
	Moral rel.	Ordinal (5)	Disagree somewhat	Agree somewhat - Disagree strongly	
Affirmative Action					*"Should companies that have discriminated against blacks have to have an affirmative action program?"*
	Affirm. act.	Ordinal (4)	Weak Yes	Strong Yes - Strong No	
Military Spending					*"Should the government decrease/increase defense spending"?*
	Milit. spend	Ordinal (5)	Keep the same	Keep the same - Increase Spending	
Military					*Feelings about the military, from not favorable/"cold" (0) to favorable/"warm" (100)*
	Military	Numeric [0,100]	$\mu = 72.7$	$\sigma = 20.4$	
Environmentalism 1					*"Toughen regulations to protect the environment" or regulations are "too much of a burden on business".*
	Envir. 1	Ordinal (5)	Neither	Toughen Regulations - Neither	
Environmentalism 2					*More important to protect the environment or maintain jobs and standard of living?*
	Envir. 2	Ordinal (5)	Equally Important	Protect Environment - Equally Important	
Environmentalism 3					*Increase / decrease "federal spending on environmental protection"?*
	Envir. 3	Ordinal (4)	Increase	Increase - Keep about the same	
Equal Rights 1					*Agreement with "we have gone too far in pushing equal rights in this country"*
	Eq. rights 1	Ordinal (5)	Neither ag. nor disag.	Agree somewhat - Disagree somewhat	
Equal Rights 2					*Agreement with "This country would be better off if we were less worried about how equal people are"*
	Eq. rights 2	Ordinal (5)	Neither ag. nor disag.	Agree somewhat - Disagree somewhat	
Equal Treatment					*Agreement with "If people were treated more equally in this country we would have many fewer problems"*
	Equal treat.	Ordinal (5)	Agree somewhat	Agree strongly - Disagree somewhat	
Foreign Aid					*Increase / decrease "federal spending on foreign aid"?*
	For'gn aid	Ordinal (4)	Keep the same	Keep the same - Decrease	
Biblical Literalism					*Is the Bible the word of God? Yes, literally / Yes, but not literally / No, not the word of God.*
	Biblical lit.	Ordinal (3)	Yes but not literally	Yes, literally - Yes but not literally	

TABLE 1 (CONTINUED)

BELIEF VARIABLE	SHORT NAME	TYPE AND RANGE†	MEDIAN OR MEAN‡	1ST - 3RD QUART. OR STDEV	*ITEM (ITALICS)*
Gay Rights (two-item scale)					*(1) Should "homosexuals [...] be allowed to serve in the United States Armed Forces"? (strong no - strong yes)* *(2) Should "gay"/"lesbian"/"homosexual couples" "be legally permitted to adopt children?" (no - yes)*
	Gay rights	Numeric [1,4]	$\mu = 2.14$	$\sigma = 1.06$	
Newer Lifestyles					*Agreement with "The newer lifestyles are contributing to the breakdown of our society"*
	New lifest.	Ordinal (5)	Agree somewhat	Agree strongly - Neither agree nor disag.	
AIDS Spending					*Increase / decrease "federal spending on AIDS research"?*
	AIDS spend	Ordinal (4)	Increase	Increase - Keep about the same	
Buying Guns					*Should federal government make it "more difficult" or "easier for people to buy a gun"?*
	Gun rights	Ordinal (5)	Somewhat more diffic.	A lot more difficult - Keep about the same	
Death Penalty					*"Do you favor or oppose the death penalty for persons convicted of murder?"*
	Death penal.	Ordinal (5)	Favor strongly	Favor strongly - Oppose not strongly	
Crime Spending					*Increase / decrease "federal spending on dealing with crime"?*
	Crime spend	Ordinal (4)	Increase	Increase - Keep about the same	
Ideological Identity					*Self-placement on ideological seven-point scale from "extremely liberal" to "extremely conservative".*
	Ideol. id	Ordinal (7)	Moderate	Moderate - Slightly conservative	
Immigration 1					*Increase / decrease "number of immigrants from foreign countries who are permitted to come to the [U.S.] to live"*
	Immig. 1	Ordinal (5)	Leave the same	Leave the same - Decrease a lot	
Immigration 2					*Increase / decrease "federal spending on tightening border security to prevent illegal immigration"*
	Immig. 2	Ordinal (4)	Increase	Increase - Keep about the same	
Immigration 3					*"Favor a law making English the official language of the United States"?*
	Immig. 3	Ordinal (3)	Favor	Favor - Neither favor nor oppose	
Individualism					*Should be "cooperative person who works well with others" or "a self-reliant person able to take care of oneself"?*
	Indiv.	Ordinal (2)	Cooperative	Self-reliant - Cooperative	
Limited Government (three-item scale)					*(1) "There are more things that the government should be doing" vs. "the less government the better"* *(2) "Need a strong gov't to handle today's complex economic problems" vs. "the free market can handle [them]"* *(3) Gov't bigger because it does things people should do for themselves vs. because we face bigger problems*
	Limit. gov't	Numeric [-1.1,1.5]	$\mu = 0$	$\sigma = 1.03$	
Equal Opportunity					*"Our society should do whatever is necessary to make sure that everyone has an equal opportunity to succeed"*
	Eq. oppor.	Ordinal (5)	Agree strongly	Agree somewhat - Agree strongly	
Equal Chance					*"One of the big problems in this country is that we don't give everyone an equal chance"*
	Eq. chance	Ordinal (5)	Neither ag. nor disagr.	Disagree somewhat - Agree somewhat	

TABLE 1 (CONTINUED)

BELIEF VARIABLE	SHORT NAME	TYPE AND RANGE†	MEDIAN OR MEAN‡	1ST - 3RD QUART. OR STDEV	*ITEM (ITALICS)*
Inequality					*Agreement with "It is not really that big a problem if some people have more of a chance in life than others"*
	Ineq.	Ordinal (5)	Disagree somewhat	Agree somewhat - Disagree somewhat	
Parenting 1					*"Please tell me which one you think is more important for a child to have: independence, or respect for elders?"*
	Parent. 1	Ordinal (3)	Respect for elders	Both - Respect for elders	
Parenting 2					*"Please tell me which one you think is more important for a child to have: curiosity, or good manners?"*
	Parent. 2	Ordinal (3)	Good manners	Curiosity - Good manners	
Parenting 3					*"Please tell me which one you think is more important for a child: being considerate, or well behaved?"*
	Parent. 3	Ordinal (3)	Considerate	Considerate - Well behaved	
Party Identity					*Self-placement on party-identification seven point scale from "Strong Democrat" to "Strong Republican"*
	Party id	Ordinal (7)	Independent	Weak Democrat - Weak Republican	
Anti-Black Racism					*Rate blacks from (1)"hard-working" to "lazy"(2)"intelligent" to "unintelligent"(3)"trustworthy" to "untrustworthy"*
(three-item scale)	Racism	Numeric [-2.5,2.7]	$\mu = 0$	$\sigma = 0.9$	
Gender Equality					*Men and women should "have equal roles" or "a woman's place is in the home."*
	Gender equal.	Ordinal (5)	Feel strongly: equal	Feel strongly: equal - Feel not strongly: equal	
Religiosity					*"Do you consider religion to be an important part of your life, or not?"*
	Relig.	Ordinal (2)	Important	Important - Important	
Surplus Taxes					*Approve / disapprove : "most of the expected federal budget surplus should be used to cut taxes"*
	Taxes	Ordinal (2)	Approve	Approve - Disapprove	
Tolerance					*Should be "tolerant of people who choose to live according to their own moral standards" even if different from ours*
	Toler.	Ordinal (5)	Agree somewhat	Agree somewhat - Neither agree nor disagr.	
Welfare Recipients					*Feelings about "people on welfare", from not favorable/"cold" (=0) to favorable/"warm"(=100)*
	Welf're recip.	Numeric [0,100]	$\mu = 51.8$	$\sigma = 20.3$	
Welfare Spending					*Increase / decrease : (1) "federal spending on welfare programs?" and (2) "federal spending on food stamps?"*
(two-item scale)	Welf're spend	Numeric [1,4]	$\mu = 2.22$	$\sigma = 0.64$	

NOTE.—The "short name" column contains the abbreviations we use in the network diagrams. Numeric variables are coded so that lower values correspond to respondents providing the first response option. E.g., for Anti-Black Racism, lower values corresponds to "hard-working", "intelligent", and "trustworthy"; for Welfare Spending, they correspond to "increase" for both types of funding.

†—Ordinal variable (with number of categories), or numeric variable with range [from, to].

‡—For ordinals, these columns show median and the 1st and 3rd quartile values. For numerics, they show mean and standard deviation.

TABLE 2

CENTRALITY OF NODES IN BELIEF NETWORK FROM FULL 2000 ANES SAMPLE

ATTITUDE	BETWEENNESS CENTRALITY		
	Original sample	Bootstrapped (1000 row resamples)	
	Absolute	Relative (mean & 95% CI)	Absolute (mean)
Ideological Identity	0.352	100	0.319
Limited Government	0.119	34	0.105
Gay Rights	0.104	27	0.082
Equal Rights 1	0.074	25	0.076
Welfare Spending	0.089	21	0.065
Parenting 2	0.070	20	0.061
Party Identity	0.065	16	0.048
Biblical Literalism	0.020	16	0.050
Environmentalism 3	0.023	13	0.041
Environmentalism 1	0.046	12	0.039
Immigration 2	0.035	12	0.037
Immigration 3	0.036	12	0.037
Abortion Legal	0.030	12	0.038
Affirmative Action	0.024	10	0.032
Immigration 1	0.024	9	0.029
Equal Rights 2	0.023	8	0.026
Newer Lifestyles	0.016	5	0.017
Buying Guns	0.008	5	0.016
Religiosity	0.011	5	0.017
Foreign Aid	0.007	4	0.011
Equal Chance	0.008	3	0.009
Moral Relativism	0.005	3	0.010
Welfare Recipients	0.024	5	0.015
Equal Treatment	0.004	2	0.006
Military Spending	0.003	2	0.006
Crime Spending	0.004	2	0.005
Tolerance	0.001	1	0.004
AIDS Spending	0.003	2	0.005
Parenting 1	0	1	0.004
Equal Opportunity	0	1	0.004
Abortion for Teens	0	1	0.003
Environmentalism 2	0	1	0.002
Death Penalty	0	0	0.001
Military	0	0	0.001
Surplus Taxes	0	0	0.001
Inequality	0.001	0	0.001
Gender Equality	0	0	0.001
Parenting 3	0	0	0
Individualism	0	0	0
Anti-Black Racism	0	0	0

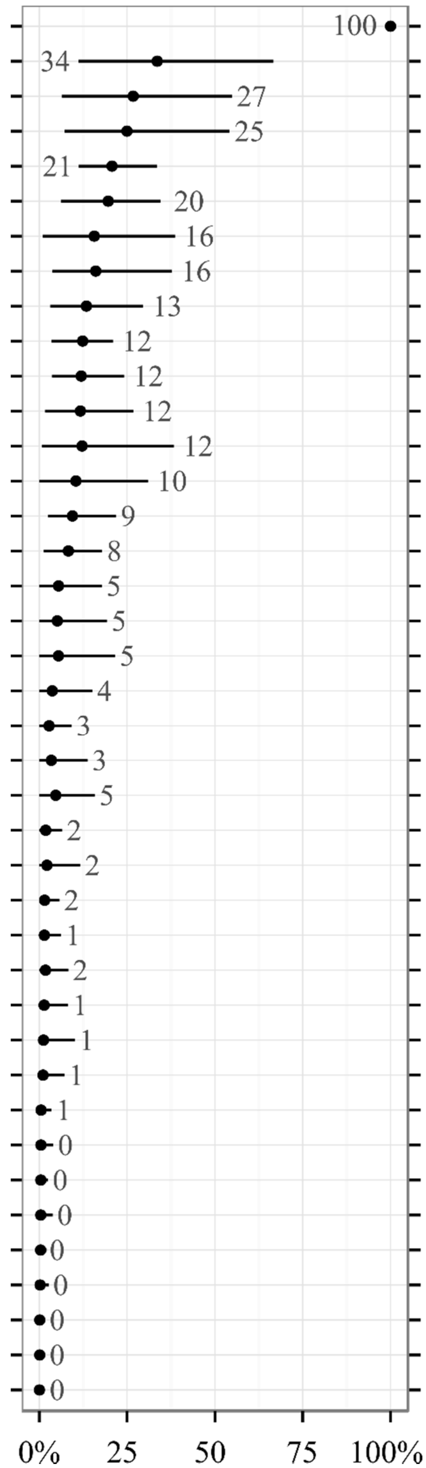

TABLE 3

STABILITY TO CHANGES IN VARIABLE SET: BETWEENNESS CENTRALITY IN 1000 COLUMN RESAMPLES

ATTITUDE	IDEOLOGICAL ID. RETAINED: Relative (mean and 95% CI)	IDEOLOGICAL ID. RETAINED: Absolute (mean)	IDEOLOGICAL ID. DROPPED: Relative (mean and 95% CI)	IDEOLOGICAL ID. DROPPED: Absolute (mean)
Ideological Identity	100	0.388	n/a	—
Limited Government	40	0.153	65	0.137
Gay Rights	31	0.114	85	0.185
Welfare Spending	23	0.089	32	0.067
Equal Rights 1	25	0.090	96	0.218
Party Identity	20	0.078	56	0.111
Parenting 2	19	0.070	42	0.089
Environmentalism 1	12	0.047	20	0.040
Immigration 3	13	0.049	25	0.052
Abortion Legal	12	0.047	18	0.036
Environmentalism 3	12	0.046	34	0.067
Biblical Literalism	14	0.054	37	0.076
Equal Rights 2	9	0.035	21	0.041
Immigration 2	10	0.037	26	0.051
Affirmative Action	10	0.037	22	0.043
Immigration 1	8	0.030	26	0.049
Newer Lifestyles	7	0.027	27	0.057
Welfare Recipients	5	0.021	7	0.015
Buying Guns	5	0.018	23	0.046
Religiosity	2	0.010	14	0.028
Foreign Aid	3	0.011	8	0.015
Equal Chance	2	0.008	5	0.010
Moral Relativism	2	0.007	32	0.061
Equal Treatment	2	0.006	8	0.016
Crime Spending	1	0.005	3	0.006
AIDS Spending	2	0.006	4	0.009
Military Spending	1	0.003	15	0.030
Tolerance	1	0.002	9	0.017
Death Penalty	1	0.002	1	0.003
Equal Opportunity	0	0.001	1	0.002
Inequality	0	0.001	1	0.002
Parenting 1	0	0.001	1	0.001
Gender Equality	0	0	5	0.010
Environmentalism 2	0	0	1	0.002
Anti-Black Racism	0	0	0	0.000
Individualism	0	0	0	0.000
Military	0	0	2	0.004
Abortion for Teens	0	0	2	0.004
Surplus Taxes	0	0	2	0.003
Parenting 3	0	0	0	0
	0% 25 50 75 100%		0% 25 50 75 100%	

TABLE 4
SIGNIFICANT CORRELATIONS BETWEEN PAIRS OF ATTITUDES WHICH HAVE OPPOSITE SIGNS WITHIN DIFFERENT SUBGROUPS (PERCENTAGE)

DEMOGRAPHIC SUBGROUPS (WITH GROUP SIZE)	COR'S WITH $p < 0.05$	% OPPOSITE				
		1	2	3	4	5
Gender:						
1. Male (671) ………………………	649	—	0.4			
2. Female (872) ……………………	587	0.4	—			
Class (self-identified) †:						
1. Working class (699) …………..	582	—	1.3			
2. Middle class (844) ………………	639	1.3	—			
Parents foreign-born:						
1. Parents Foreign-born (229) ……..	449	—	1.0			
2. Not foreign-born (1310) ………..	663	1.0	—			
Number of children:						
1. One or more children (1124) …..	628	—	0.2			
2. No children (403) ………………	577	0.2	—			
Black:						
1. Black (159) …………………….	226	—	10.5			
2. Not Black (1344) ………………	683	10.5	—			
Hispanic:						
1. Hispanic (101) …………………	243	—	2.6			
2. Not Hispanic (1432) ……………	676	2.6	—			
Age group:						
1. Under 40 (545) …………………	577	—	0	1.4		
2. 40 to 55 (491) …...………………	589	0	—	1.3		
3. Over 55 (499) …………………...	517	1.4	1.3	—		
Education:						
1. High School or below (578) …….	464	—	1.8	4.0		
2. Associate or some college (467) ..	561	1.8	—	0.2		
3. Bachelor's or above (494) …….	636	4.0	0.2	—		
Income:						
1. Under \$35,000 / year (513) …….	510	—	0.5	1.4		
2. \$35,000 to \$65,000 / year (385) ..	586	0.5	—	0.4		
3. Over \$65,000 / year (408) ………	576	1.4	0.4	—		
South-Eastern United States:						
1. Southern (550) …………………	584	—	0.4			
2. Not Southern (993) …………….	640	0.4	—			

TABLE 4 (CONTINUED)

DEMOGRAPHIC SUBGROUPS (WITH GROUP SIZE)	COR'S WITH P <0.05	% OPPOSITE SIGNS IF COMPARED TO SUBGROUP NUMBER 1	2	3	4	5
Religion:						
1. Catholic (396) ……………….	410	—	1.1	1.4	1.9	
2. Mainline Protestant (420) ……	574	1.1	—	2.0	0.8	
3. Other Protestant (287) ……….	423	1.4	2.0	—	2.5	
4. Other Religion or none (323) ..	588	1.9	0.8	2.5	—	
Occupational category:						
1. Manager (203) ………………….	445	—	0.5	0.3	0	1.4
2. Professional (326) …………...	587	0.5	—	0.5	0.5	2.4
3. Routine non-manual (354)……	474	0.3	0.5	—	0.3	0.9
4. Skilled or semi-skilled (241) ….	435	0	0.5	0.3	—	1.6
5. Unskilled or farm (276) ……….	440	1.4	2.4	0.9	1.6	—
Type of place:						
1. Larger city (282) ………………	436	—	1.5	1.3	1.6	
2. Rural area (431) ………………	488	1.5	—	1.2	1.8	
3. Town or smaller city (507) ……	592	1.3	1.2	—	0.2	
4. Suburb (310) …………………..	569	1.6	1.8	0.2	—	
Church Attendance:						
1. Attends Church (1057)	641	—	0.6			
2. Does Not Attend Church (478)	604	0.6	—			
Cross-Pressures:						
1. Pressures aligned (471)	591	—	0	1.6		
2. Neither (385)	586	0	—	0.9		
3. Pressures crossed (448)	499	1.6	0.9	—		
Political knowledge:						
1. Low (398) ………………………	372	—	4.2	12.6		
2. Medium (667) ………………….	543	4.2	—	1.6		
3. High (473) …………………….	700	12.6	1.6	—		

NOTE.—Each category is compared to each other category within the same demographic dimensions. The column numbers (1 – 5) index the categories in the current dimension (E.g., for the categories within Religion, subgroup number 1 is Catholics, 2 is mainline Protestants, etc.) The counts of statistically significant correlations are out of 780 possible. See Appendix E for detailed descriptions of the variables.

† Identification with other classes is rare. 97% of respondents either "thought of" themselves as middle or working class, or identified with one of the two when prompted to do so explicitly.

TABLE 5

BELIEF CENTRALITIES FOR DIFFERENT SUBPOPULATIONS

ATTITUDE	RELATIVE BETWEENNESS CENTRALITY (MEAN AND 95% CI)		
	High Information Respondents	Low Information Respondents	African American Respondents
Ideological Identity	100	18	18
Immigration 3	19	11	3
Welfare Spending	17	19	15
Equal Rights 2	16	20	22
Surplus Taxes	15	5	18
Environmentalism 1	14	37	42
Abortion Legal	11	29	27
Equal Rights 1	11	38	12
Limited Government	11	17	5
Parenting 2	11	30	36
Biblical Literalism	11	69	33
Immigration 2	10	11	8
Gay Rights	10	21	14
Abortion for Teens	10	6	11
Buying Guns	9	12	14
Affirmative Action	6	43	36
Party Identity	6	8	12
Environmentalism 3	5	30	29
Newer Lifestyles	5	5	7
Welfare Recipients	4	15	8
Tolerance	4	11	12
Military Spending	4	0	4
Parenting 1	3	9	18
Immigration 1	3	28	8
Religiosity	3	75	88
Individualism	3	23	13
Foreign Aid	3	43	6
Equal Chance	2	14	16
Moral Relativism	2	4	18
Death Penalty	1	4	5
AIDS Spending	1	14	49
Environmentalism 2	1	13	20
Parenting 3	1	13	44
Military	0	4	7
Gender Equality	0	18	30
Equal Opportunity	0	66	36
Inequality	0	4	12
Equal Treatment	0	35	20
Crime Spending	0	15	16
Anti-Black Racism	0	6	4

0% 25 50 75 100 | 0% 25 50 75 100 | 0% 25 50 75 100

NOTE.—Each of the three sets of betweenness estimates is based on 1000 row-wise bootstraps.

FIGURES FOR CHAPTER 2

FIG. 1.—Belief network at different stages.

(1) After 1st generation

(2) After 2nd generation

(3) Snippet of a large belief network.

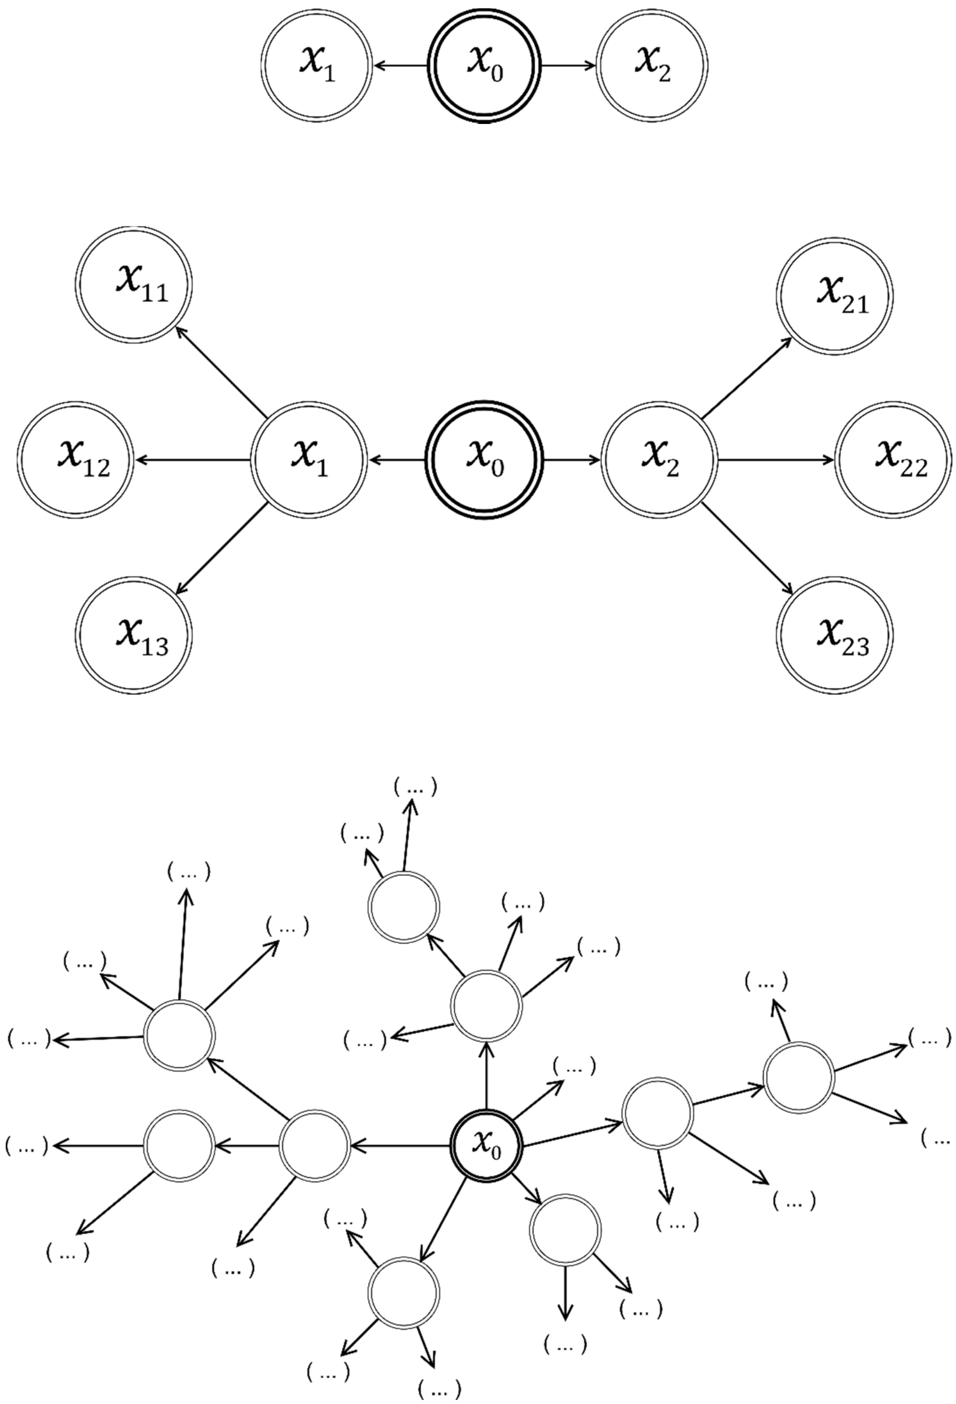

FIG. 2.— Correlation network for the full population sample (see Table 1 for full node names). Tie strength $|T_{ij}| = cor^2(x_i, x_j)$ is represented by thickness and boldness (see inset). Correlations below $|r| = 0.15$ are not depicted. The force-directed layout places strongly connected nodes closer together, and weakly connected and unconnected nodes further apart.

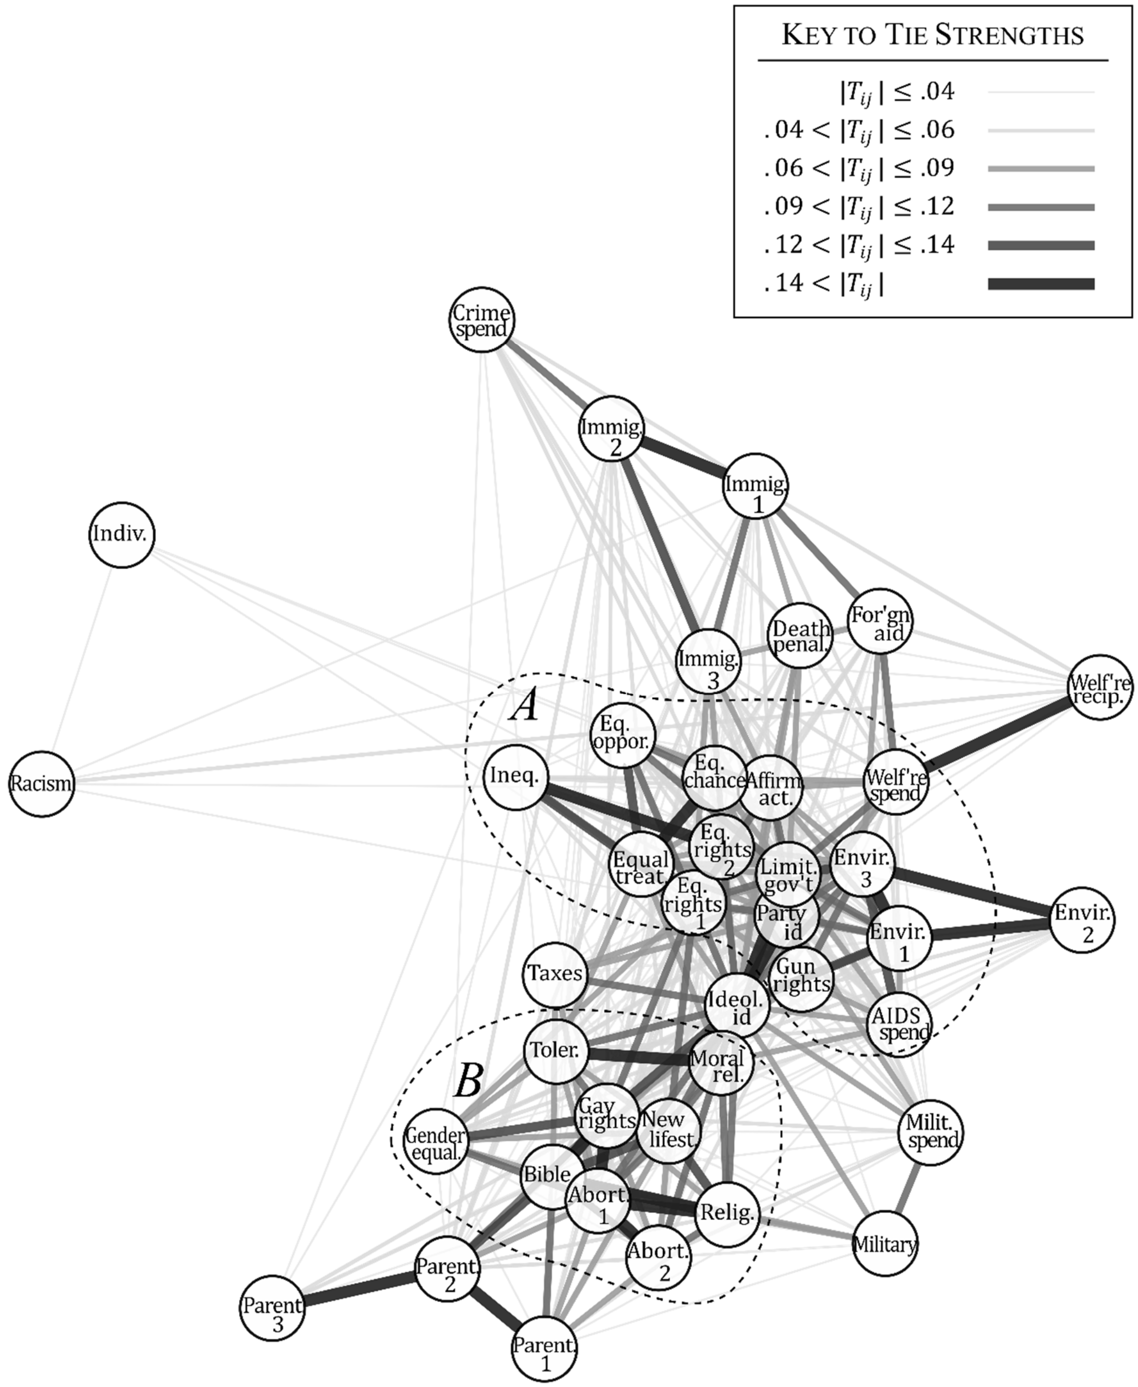

FIG. 3.— Degree of organization in belief networks of 44 demographic subgroups. Filled circles indicate networks where the most central node was ideological identity.

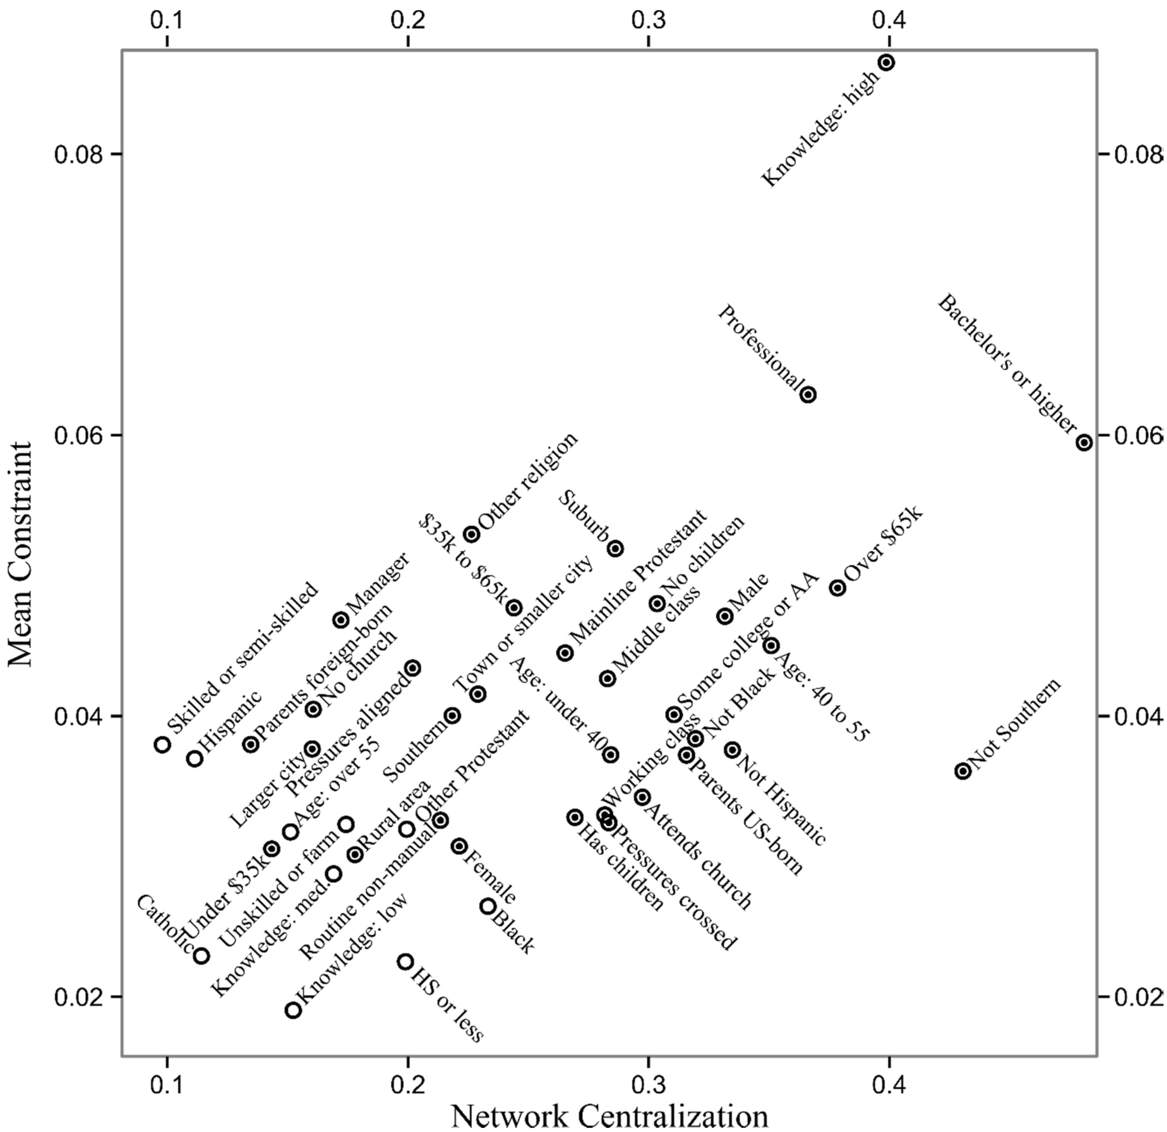

FIG. 4.— Network centralization by centrality of ideological identity in 44 demographic subgroups, with line of best fit.

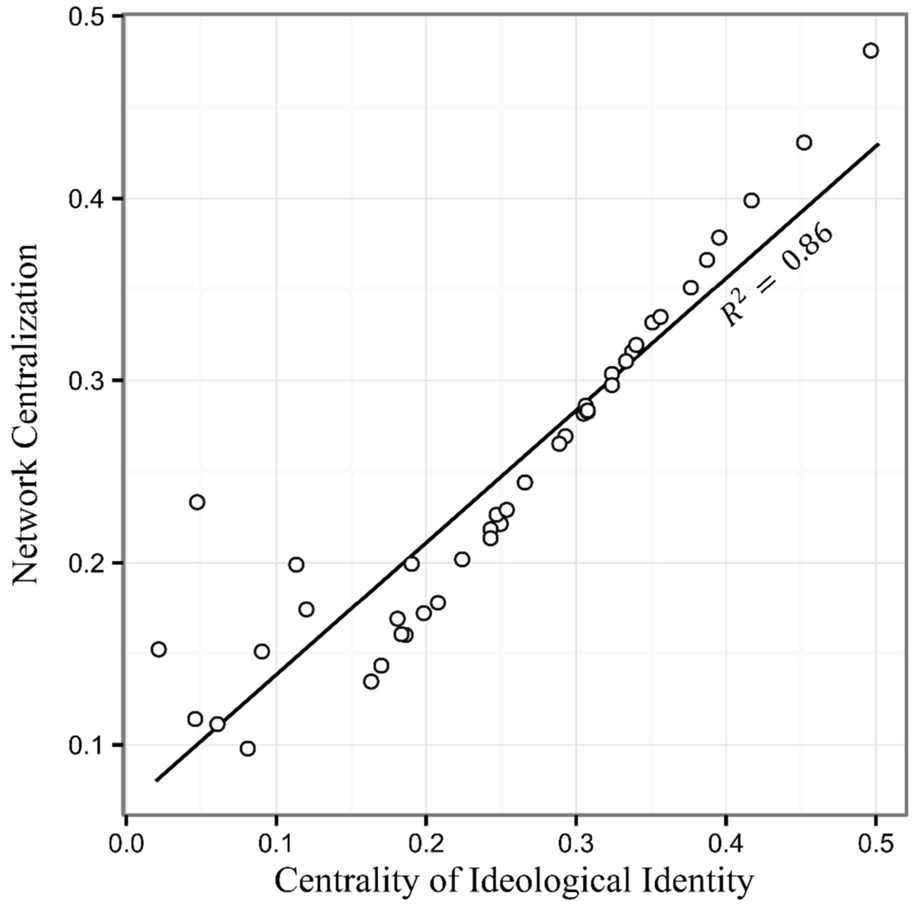

FIG. 5. Correlation network for the high information subsample (see Table 1 for full node names). Tie strength is represented by thickness and boldness (see inset). Correlations below $|r| = 0.15$ are not depicted. Force-directed layout.

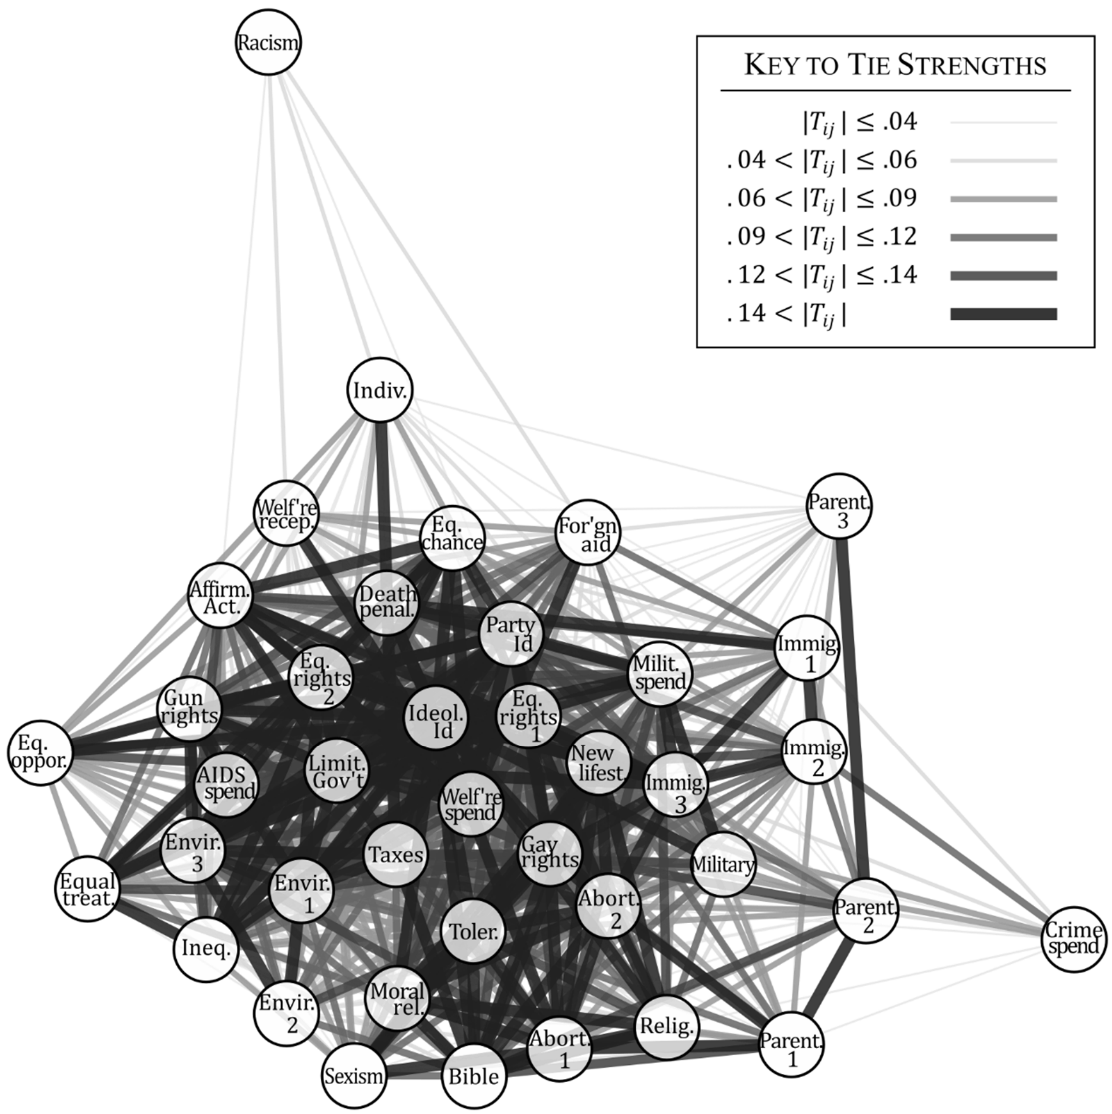

FIG. 6. Correlation network for the low information subsample (see Table 1 for full node names). Tie strength is represented by thickness and boldness (see inset). Correlations below $|r| = 0.15$ are not depicted. Force-directed layout.

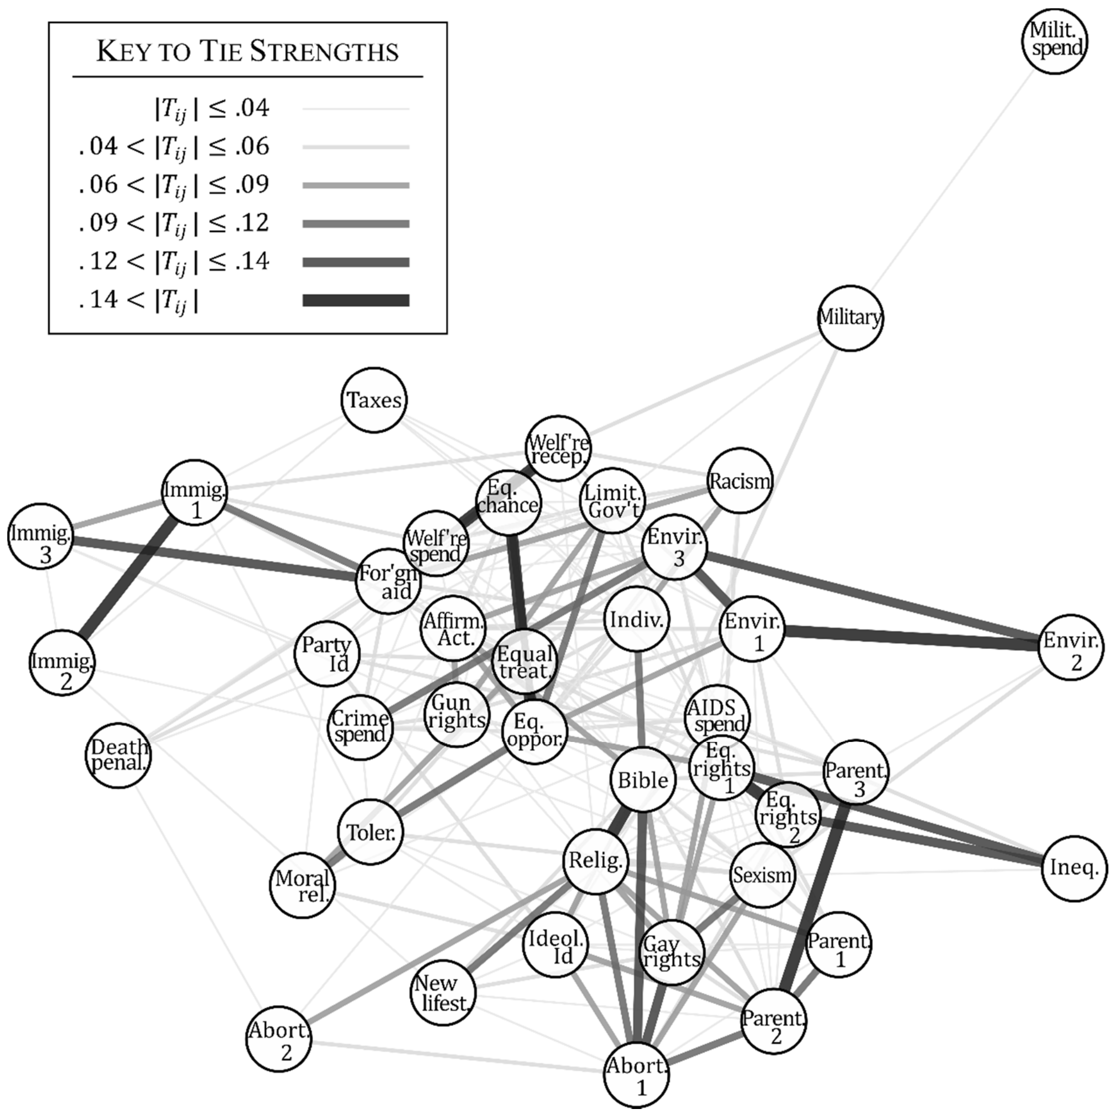

Chapter 3.

Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method

The task of revealing intelligible structures of meaning beneath complex collections of cultural data is among the most central methodological challenges posed by the sociology of culture (Mohr 1998; Mohr and Rawlings 2012). From the perspective of culture and cognition, this task is a search for shared "cultural schemas"—abstract cognitive structures which specify relationships between cultural elements. In a high-profile recent work, Goldberg (2011) proposes an innovative methodology for identifying groups of survey respondents who share such cultural schemas, which he terms Relational Class Analysis (RCA). RCA has rightfully garnered a substantial amount of attention across diverse domains of study including cultural tastes (Goldberg 2011; Daenekindt 2016), public opinion (Baldassarri and Goldberg 2014; Wu 2014), organizational behavior (Miranda, Summers, and Kim 2012), and economic sociology (DiMaggio and Goldberg 2010). However, existing work has not yet provided a clear definition of shared cultural schemas, the central concept under investigation. This is a crucial limitation as, without such a definition, RCA's accuracy at locating such schemas cannot be convincingly demonstrated.

Goldberg (2011) introduces RCA using a survey of musical tastes as his case study. Taking a cue from relational theories of meaning (e.g., Saussure 1916 [2013]; Lévi-Strauss 1963), RCA searches the data for schemas that define not the musical tastes themselves, but the *relationships* between these tastes—that is, which genre tastes are perceived as similar and which as opposed. This kind of schema can be found in the implicit agreement between an individual who likes musical genres A and B but dislikes genre C, and another who dislikes A and B but likes C: though the two hold no tastes in common, they nonetheless agree that A "goes with" B, whereas C is "opposed to" A and B. On the other hand, an individual who likes all three of the genres A, B and C has two tastes in common with the first individual, but does not agree that C is the opposite of A and B. Thus, under this conception, the first two individuals arrange their tastes according to the same cultural schema, while the third one does not.[1] The goal of RCA is to partition the survey population into classes of respondents that share such cultural schemas—a novel methodological task which is, in itself, a bold conceptual innovation.

At the core of RCA's approach is a novel similarity measure termed "relationality." Goldberg contends that relationality can quantify the extent to which two respondents organize their attitudes according to a shared cultural schema. However, he does not provide a clear definition of such shared schemas. With this key link between theory and measurement missing, relationality's ability to successfully capture such schemas cannot be convincingly shown. In this paper, I examine the theoretical intuitions implicit in current work to arrive at this missing definition. I demonstrate that, in order to detect shared cultural schemas like those in Goldberg's (2011:1404-1405) motivating example, relationality must measure the degree of linear dependency between two individuals' vectors of responses. This lends itself to a simple, intuitively plausible formal model of schematic similarity as linear dependence between response vectors, and suggests that Pearson's correlation may already provide a solution for the task that relationality sets out to solve. When I reexamine Goldberg's motivating example with correlation, I find that it indeed yields more accurate results than relationality. Its results match Goldberg's own description of the data, while relationality's do not.

---

[1] In Martin's (2002) terminology, such schemas thus underlie the "tightness" rather than the "consensus" of a system of attitudes. If one imagines attitudes as an abstract space where each dimension represents a like/dislike of a given musical genre, such schemas would specify an axis or plane along which culturally valid tastes can be arranged rather than a specific point in space at which tastes should be located.

I then use simulations to verify that this difference in accuracy generalizes more broadly. First, I simulate 10,000 test cases where shared schemas result in linear dependencies between responses, as they do in my model and in Goldberg's example. I analyze each simulated dataset with both relationality and correlation, and compare their results to the true schematic structure of the simulated data, which is known by design. The results confirm that this switch from "Relational" to "Correlational" Class Analysis (CCA) reliably increases the accuracy of the technique, so much so that, when the two disagree, the odds that CCA's results are more accurate approach 17:1. I also document that RCA introduces a strong distributional assumption which has not previously been noted in published work. When it is violated, the odds in favor of CCA's results further increase to 23:1.

Both RCA and CCA seek to measure the extent to which each pair of respondents employs the same cultural schema, and use these measures to create groupings. But what if the socio-cognitive processes giving rise to shared cultural schemas are actually non-linear in character, contrary the theoretical intuitions formalized in this paper? Though the motivating examples behind relationality do not entail such a possibility, Goldberg nonetheless designed RCA to function even in the presence of complex nonlinearities in the data (Goldberg 2011:1433). Since the model elaborated here is only one theoretical possibility of how schematic cultural cognition may function, this kind of robustness may be a prudent goal. I thus conduct further simulations to determine whether RCA is preferable over CCA when shared cultural schemas depart from the theorized linearity.

In the first three sets of simulations, which I term "sub-linear," I investigate the possibility that some of the basic theorized relationships between schemas—scaling, shift or inversion—do not take place. In the second three sets, which I call "super-linear," I examine the possibility of schematic transformations that are more complex than those envisioned by the theory: high-degree polynomial transformations, multi-way interactions, and independence between parts of the schema. The simulations show that CCA remains reliably more accurate than RCA under every scenario examined. While substantial departures from linearity cause CCA's average accuracy to decrease, RCA's average accuracy decreases to an even lower level. I find no evidence of settings where RCA is more accurate than CCA.

Finally, to explore the substantive consequences of switching from relationality to correlation, I revisit Goldberg's (2011) RCA analysis of the 1993 General Social Survey (GSS) musical taste data. The RCA analyses had detected three schematic classes: "Omnivore – Univore," "Contemporary – Traditional," and "Highbrow – Lowbrow." The CCA results confirm the first two of these schemas, but contain two other schemas in place of RCA's third. I term these additional schemas "Anything (but) Country" and "Anything (but) Heavy Metal." These schemas more closely resemble the patterns of exclusionary omnivorousness documented by Bryson (1996) than the hierarchical "Highbrow – Lowbrow" logic described by Goldberg. Thus, the CCA analyses confirm broad outlines of RCA findings, but also contain potentially important differences. I further compare these results via a widely-known multiple groups analysis technique from structural equation modeling (SEM), which can be used to provide a goodness-of-fit measure for inductively located heterogeneity (Author et al., under review). The multiple groups analysis indicates that CCA's classes have a far better fit to the GSS data than RCA's. I conclude by discussing how future work can further advance the methodological and theoretical project stemming from Goldberg's (2011) deeply innovative contribution.

SHARED SCHEMAS

RCA's conception of shared cultural schemas builds on relational theories of meaning in the tradition of Saussure and Lévi-Strauss (e.g., Cerulo 1993, Emirbayer 1997). Such theories hold that the meaning of symbols in a cultural system rests not in the signs themselves, but in the relationships among them. So, for example, the concept "hot" acquires its meaning from its opposition to "cold" rather than from any innate properties of the word itself (Saussure 1916 [2013]; Ritzer and Stepnisky 2013). The concepts that make up a single cultural domain then exist in semantic networks, where each element gains its significance from its relationships to the other elements. For example, each of the concepts "mother," "father," "son," and "daughter" is defined in part through its distinction from the other three concepts, and cannot be fully understood without understanding them as well (Lévi-Strauss 1963). Such a configurations of concepts within a single cultural domain can be thought of a *shared cultural schema*—an abstract cognitive structure that individuals acquire through experience or acculturation.

As I noted above, Goldberg (2011) does not formally define what it means for a set of survey respondents to share such a cultural schema.[2] He instead illustrates this relationship by way of a motivating example and accompanying diagram, which I recreate as Figure 1. Describing this figure, Goldberg states that A and B have "identical" patterns of musical tastes, C's pattern is "almost a mirror image" of A's and B's, and D's is "different but not antithetical." He thus concludes that "respondents A, B, and C exhibit the same logic of musical taste construction ... as they all exhibit the same structure of relevance and opposition" (Goldberg 2011:1405), whereas respondent D does not.

[2] Goldberg describes cultural schemas as "complex structures of mental representation […] that are built up incrementally through interaction with the environment [that] embody our taken-for-granted assumptions about the world," and further states that "schemas are not clear sets of behavioral rules but rather implicit recognition procedures that emerge from intricate associational links among salient aspects of our cognitively represented experiences." (2011:1401). This description provides some intuitions about the functional role schemas play in culture and cognition. However, it is not the clear definition needed to assess whether relationality can accurately measure such schemas.

To arrive at a formal definition of schematic similarity, I expand Goldberg's discussion of this introductory example by writing out the implied algebraic operations. Respondent A likes pop, blues and rock, strongly likes classical and opera, and is indifferent towards bluegrass and country: $A = [4,4,4,5,5,3,3]$. Respondent B, on the other hand, dislikes pop, blues and rock, is indifferent towards classical and opera, and strongly dislikes bluegrass and country: $B = [2,2,2,3,3,1,1]$. Except for an overall downward shift in the appraisal of all the genres, this pattern of tastes is identical to that of the first respondent: $B = A - 2$.

In contrast to A and B, respondent C is indifferent towards pop, blues and rock, strongly dislikes classical and opera, and strongly likes bluegrass and country: $C = [3,3,3,1,1,5,5]$. These tastes again follow the same relative pattern as A and B, except all tastes are vertically shifted, inverted and amplified: $C = 2 * (-1) * A + 11$, or, equivalently, $C = 2 * (-1) * B + 7$. Finally, respondent D strongly dislikes pop and rock, strongly likes blues, likes classical, opera and bluegrass, and dislikes country: $D = [1,5,1,4,4,4,2]$. Unlike A, B and C, this respondent construes an opposition between bluegrass and country, but not between bluegrass and opera. No series of inversions, multiplications or shifts of this pattern can transform it into the one exhibited by A, B and C. We thus conclude that, while respondents A, B and C follow the same schema, respondent D does not.

From this example, we can surmise that two respondents follow exactly the same schema if (i) their attitudes are identical, (ii) their attitudes are exact inverses of each other's, (iii) the attitudes of either respondent are uniformly more extreme than those of the other, (iv) the attitudes of either respondent are uniformly more positive than of the other, or (v) any combination of (ii), (iii) and (iv). These conditions specify the mathematical operations of identity ($Y = X$), inversion ($Y = -X$), scaling ($Y = kX$) and vertical shift ($Y = X + b$) (see Figure 2). They can thus be captured by a single algebraic statement: two respondents X and Y

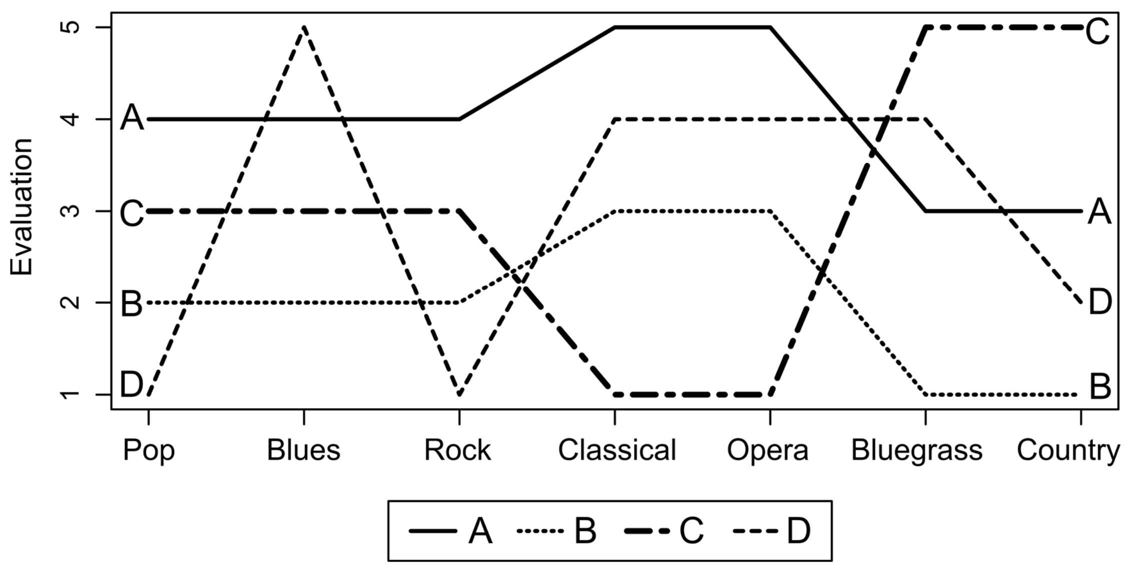

Figure 1. Musical tastes of four respondents, with evaluations ranging from 1 (strongly dislike) to 5 (strongly like) for each genre. Respondents A, B and C follow the same taste schema, while D does not. This figure recreates the contents of Goldberg's Figure 1A (2011:1405).

follow exactly the same schema *if and only if there exists a linear transformation that can produce one vector of responses from the other one*, or, more formally, if there exist such constants b and $k \neq 0$ that $Y = kX + b$. It is therefore intuitively clear that any measure of schematic similarity between two respondents should obtain its maximum value when such *k* and b exist, and should otherwise capture the degree to which one pattern can be approximated by linear transformations of the other.

This model of schematic cognition thus formalizes the implied operations that give Goldberg's illustrative example of schematic similarity its intuitive plausibility. Because of the simplicity of these linear operations, it is easy to imagine the social processes which could bring about such schematic transformations. For example, a highly opinionated person may follow the same pattern of genre tastes as a less opinionated one, but turn all "likes" into "strong likes" and all "dislikes" into "strong dislikes," thus yielding situation (iii); a music lover may begin with a pattern of musical tastes common to her social group, but shift all of them upward, yielding situation (iv); a rebellious teenager could begin with his or her parents' musical tastes and invert all of them, yielding situation (ii), etc.

Since many aspects of physical reality are linear (or approximately linear) in character, similar algebraic transformations can also be found in other cognitive domains. For example, in

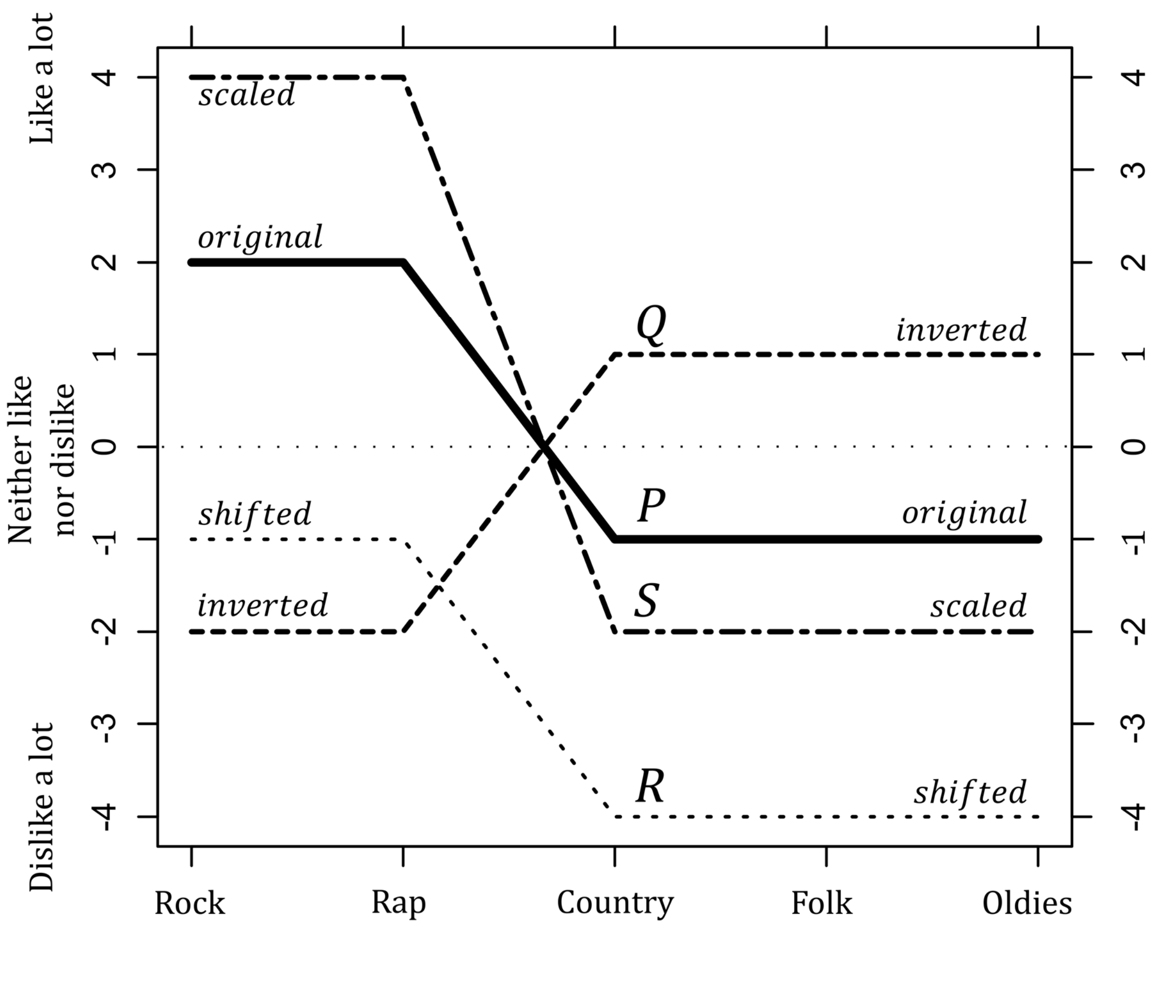

Figure 2. Original schema and three basic linear schematic transformations.

visual object perception, individuals are capable of recognizing the same stimulus placed in different positions and orientations (e.g., an apple that is rotated onto its side and moved further away, making it appear smaller). This process has been theorized to involve the perceiver effectively carrying out matrix-algebraic transformations of the perceived object so as to match it with representations of objects stored in memory (Ullman 1989; Hummel 2013). The simplicity of the three basic cognitive transformations also fits with the "cognitive miser" model of mental processing, which proposes that humans minds seek to minimize cognitive effort (e.g., Kool et al. 2010; Taylor 1981). The theorized transformations of taste schemas captured in this model thus appear both socially and cognitively plausible.

Relationality

Following his motivating example, Goldberg (2011) offers relationality R_{ij} as a measure of this schematic similarity. It is computed by first taking the row vector containing the attitudes belonging to a respondent, and calculating the differences between each pair of that respondent's attitudes by subtracting them from one another.[3] Each survey row i is thus transformed into a square matrix X_i of pairwise arithmetic differences between variables in that row. Then, to calculate the relationality between a pair of respondents i and j, the absolute values of their respective difference matrices X_i and X_j are element-wise subtracted from each other. Each element of the resulting matrix D_{ij} is assigned a sign based on whether the corresponding entries of X_i and X_j were in the same or in opposite directions. Finally, the elements of matrix D_{ij} are summed together to yield the relationality R_{ij}, which is rescaled to range from 1 to -1.

The distinction between positive and negative relationalities is not useful for RCA, as either extreme indicates that respondents i and j follow the same schema. Thus, following a bias-reduction step that I examine later, RCA uses only the absolute values of relationality $|R_{ij}|$, which range from 1 (same schema) to 0 (unrelated schemas). RCA interprets the absolute relationalities as an adjacency matrix for a weighted network, with respondents as nodes and their pairwise absolute relationalities as ties. Finally, it uses a modularity maximization algorithm (Newman 2006) to partition this network into groups of respondents who have relatively high absolute relationalities $|R_{ij}|$.

Goldberg's (2011) argument that relationality measures schematic similarity is discursive rather than formal. Drawing on the structuralist idea that symbols acquire their meaning from their distinctions with other symbols, he argues that "[c]omparing how two individuals organize meaning therefore requires examining the associations between their attitudes. This calls for a method that looks at the extent of dissimilarity between the pairwise differences between their individual opinions" (1403). He therefore constructs relationality around an across-respondent comparison of within-respondent arithmetic differences in genre ratings, apparently reasoning that, since the shared schemas specify the *relative* value individuals assign to musical genres, they would be captured by such subtractions. But that reasoning is based on a conflation of *distinction* (semantic difference) and *subtraction* (arithmetic difference)—two concepts that are related in some contexts, but substantially different in others. To illustrate the problem with this

[3] Formally, Goldberg (2011) defines the relationality between two respondents i and j to equal $R_{ij} = \frac{2}{K(K-1)} \sum_{k=1}^{K-1} \sum_{l=k+1}^{K} (\lambda_{ij}^{kl} * \delta_{ij}^{kl})$, where $\delta_{ij}^{kl} = 1 - ||\Delta X_i^{kl}| - |\Delta X_j^{kl}||$ and where $\Delta X_i^{kl} = X_i^k - X_i^l$ is the difference between the values of the variables k and l for respondent i, and $\lambda_{ij}^{kl} = 1$ if ΔX_i^{kl} and ΔX_j^{kl} have the same sign or are both zero, and $\lambda_{ij}^{kl} = -1$ otherwise.

conflation, consider an analogy to melody perception, which is another cultural domain where individual elements are defined relative to one another. Instead of the question "are two individuals who report different musical tastes nonetheless following the same schema in different ways?" the question here becomes "are two individuals who play different notes nonetheless performing the same melody in different keys?"

As is commonly known, most people perceive only relative rather than absolute pitch: that is, rather than discerning the absolute sound frequency of a musical note, they perceive its frequency relative to other notes played. Thus, if one were to change the key a melody is performed in—i.e., to multiply all the tone frequencies in the melody by the same constant—most listeners would perceive the result to be identifiably the same melody (Radocy and Boyle 2012).[4] Algebraically, if an audible frequency k is taken to be the note *A* in an equal-tempered music scale, then the following two notes in the scale ($B\flat$ and *B)* would have frequencies $k \times \sqrt[12]{2}$, and $k \times \sqrt[12]{2} \times \sqrt[12]{2} = k \times \sqrt[6]{2}$, respectively. For any melody, changing the value of $k > 0$ would alter the scale's key, but keep the melody recognizably the same (since, e.g., for any $k \neq 0, B/A = \sqrt[6]{2}$). Therefore, within the constraints of human hearing, two tone sequences (x_1, y_1, z_1) and (x_2, y_2, z_2) would be recognizable as the same melody as long as $y_1/x_1 = y_2/x_2$ and $z_1/y_1 = z_2/y_2$. An algorithm that determines whether two individuals are performing the same melody in different keys would thus need to effectively carry out an across-performer comparison of within-performer *ratios* (geometric differences) between the tone frequencies. If it were to instead compare the melodies by subtracting the tones from one another, as RCA does with musical tastes, it would generally get the comparison wrong.

The goal of this analogy is not to argue that subtraction-based relationality is necessarily the wrong measure for schematic similarity. Rather, it is to point out that distinctions between elements in a relative system may not be subtractive in character. We thus cannot conclude that relationality $|R_{ij}|$ is a valid measure of schematic similarity simply because it compares within-respondent arithmetic differences in tastes. Its ability to correctly detect shared schemas must instead be tested directly, by applying it to sets of response patterns that either do or do not follow the same schema, and examining whether it can correctly determine which are which.[5] Goldberg's introductory example discussed above can be used as such a test case.

Recall that, in the introductory example, respondents A, B and C follow exactly the same schematic pattern, while D follows a different one (Goldberg 2011:1405). The values $|R_{ij}|$ obtains in the introductory example should thus clearly identify that A, B and C share a schema, but D does not. However, relationality's difficulties at this task are evident in Figure 1B of the

[4] Similarly, multiplying the timing sequence that makes up a melody's tempo by the same constant also keeps the melody recognizably the same (Radocy and Boyle 2012).

[5] Goldberg (2011) reports some simulation tests of RCA's accuracy in his online Appendix B. However, his simulations consist entirely of individual taste vectors $Y = X$ and their exact inverses $Y = -X$, plus noise. He reports an RCA accuracy of 100% in test cases without noise, and 97.8% with noise. But these simulations are too easy: they cover only one out of the three kinds of schematic transformation, missing $Y = kX$ and $Y = X + b$. They also grant RCA's strong assumption that inversion probability equals 50% (see "Distributional Assumptions" below). To test performance under the full range of schematic similarity patterns in the motivating example, simulations would need to cover a substantially broader variety of scenarios. The simulations I report in the following sections cover this full range. On these more realistic tests, neither RCA nor CCA can reach such a high accuracy.

original paper (Goldberg 2011:1405). I depict the relevant parts of this diagram[6] in Figure 3, where diagonally shaded bars represent $|R_{ij}|$. Relationality achieves its maximum value for the respondent pair A and B (AB = 1.00), thus clearly indicating that the two follow the same schema. Conversely, the relationality between the pair A and D is approximately 0.2, which is appropriately low as A and D follow different schemas. Since C follows the same schema as A and B, the absolute relationalities AC and BC should optimally be equal to the same value as AB

[6] Respondents are assigned to classes based on the absolute values of the relationalities $|R_{ij}|$, so I plot the values $|R_{ij}|$ instead of R_{ij} to enable easier visual comparison of their magnitudes. I also omit the Euclidean distances, which are not relevant to the present discussion. These changes make the limited dynamic range of relationality values more visually apparent. Since Goldberg does not present the numerical relationality scores for this example, I reproduce these values by measuring the bar lengths in his figure.

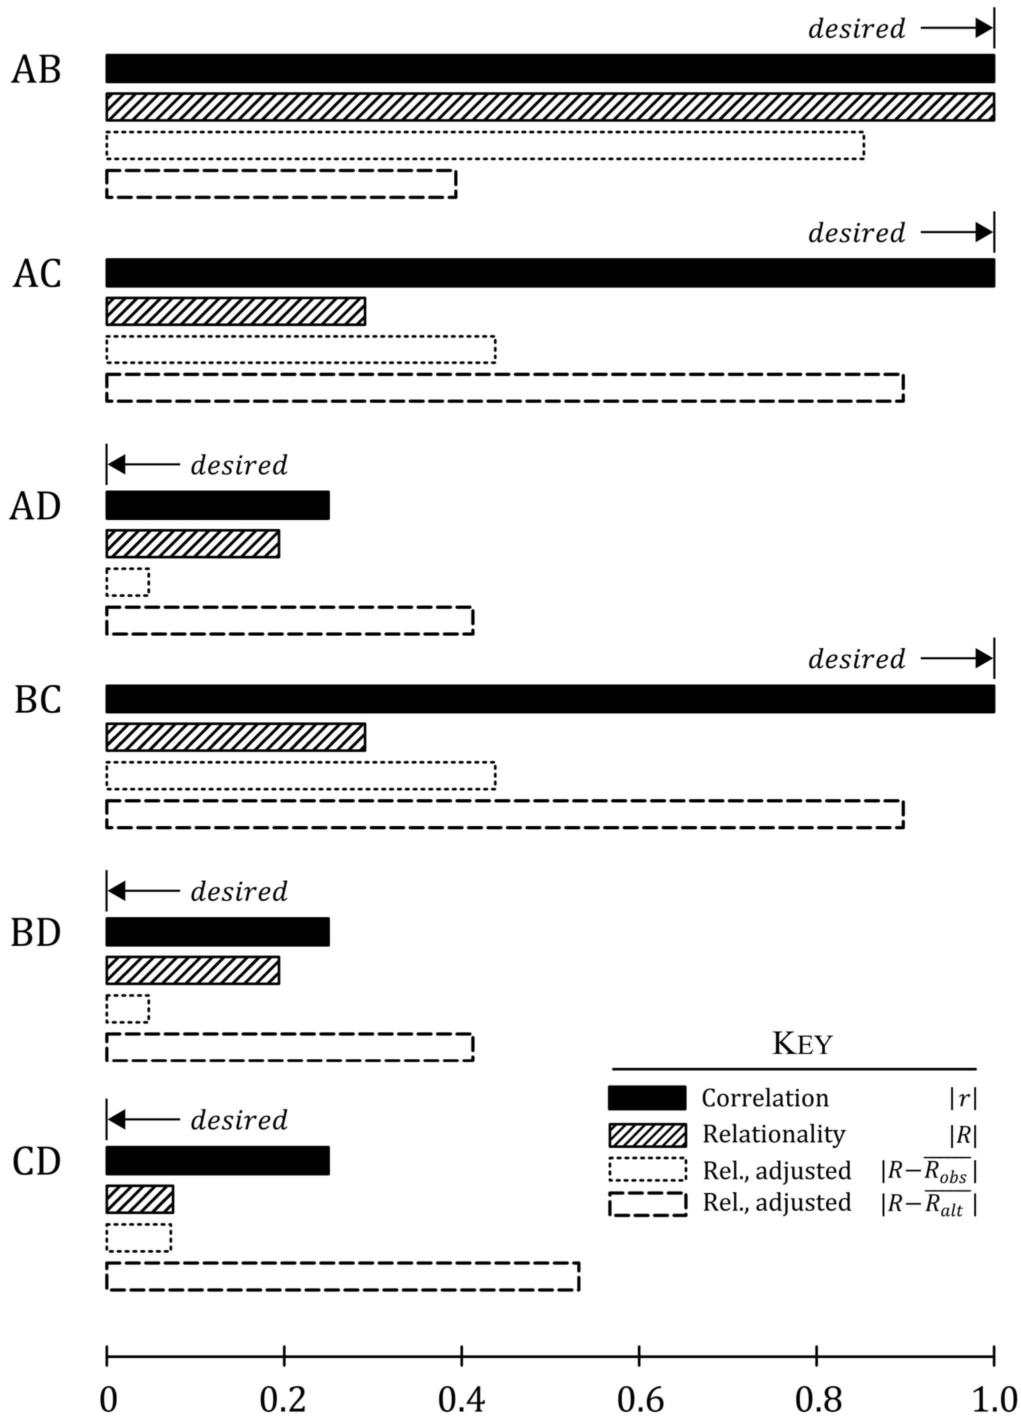

Figure 3. Absolute values of pairwise correlations (solid black) and relationalities (diagonal hatch) between the four patterns from Figure 1. The goal is to correctly determine that A, B and C belong to one schematic class, but D does not. The "desired" arrows indicate the values of each comparison that would best lend themselves to this correct answer. Remaining bars depict $|R_{ij}|$ after bias adjustment if each pattern occurs once (dotted), and if patterns A and B occur three additional times (dashed).

(1.00). Unfortunately, this is not the case: both AC and BC have absolute relationalities of approximately 0.3, which is far closer to the relationality of the unrelated pair AD (0.2). Thus, relationality appears to grossly understate the schematic similarity between respondent C and respondents A and B.

To determine whether this inaccuracy affects the solution yielded by RCA, I created a dataset consisting entirely of rows A, B, C and D, each repeated 200 times for a total of 800 rows. I then analyzed it with the RCA software provided by Goldberg (Goldberg and Zhai 2013). To produce the correct solution, RCA would have to partition this population in two classes, the first containing all the copies of rows A, B and C (600 rows total), and the second all copies of D (200 rows). However, RCA instead produced an erroneous solution consisting of three distinct classes, with the copies of C incorrectly assigned to their own class, separate from copies of A and B[7]. Since Goldberg uses this example to introduce relationality as a tool for detecting shared schemas, RCA's inability to do so here is troubling.[8]

Correlation

This shortcoming means that there is merit in trying out a different measure of schematic similarity. Recall that two respondents X and Y exactly follow the same schema if there exist such constants $k \neq 0$ and b such that $Y = kX + b$. Thus, provided that X and Y have a finite non-zero variance, it can be easily shown that the absolute Pearson's correlation $|r|$ between X and Y equals 1 if and only if they follow exactly the same schema. As the two responses become more and more linearly independent of one another—that is, as the best possible linear transformation of X leaves an ever larger percentage of Y's variance unexplained—the value of $|r|$ decreases monotonically towards 0. Finally, $|r|$ will be equal to 0 if and only if $k = 0$ gives the best linear approximation of Y, or, in other words, if the best linear approximation of Y ignores the contents of X altogether. This is why Pearson's correlation is often interpreted as the "measure of the degree of linear relationship between two variables" (Stockburger 2007; see Rodgers and Nicewander 1988 for a detailed treatment). Thus, Pearson's correlation appears to be a perfect candidate for this task.

The solid bars in Figure 3 demonstrate the results obtained by applying Pearson's correlation to the same problem. The absolute correlations AB, AC, and BC all equal 1, whereas AD, BD, and CD equal 0.25. The absolute correlations between responses that follow the same schema are thus at their theoretical maximum, while the ones between members of different schematic classes are closer to their minimum. Thus, correlation appears to produce a far clearer depiction of the schematic relationships between these respondents than relationality. To examine whether this improvement results in a correct partition into classes, I adapted

[7] When I ran RCA with default parameters, it partitioned the population into 800 separate classes, thus assigning even identical rows to different classes. This obviously faulty solution appears to be due to the pseudo-significance testing RCA uses to filter weak relationalities, which is based on strong assumptions about how relationalities are distributed in the data. Disabling it produced the substantially more realistic solution I report above. (As I discuss below, this filter appears to generally decrease the average accuracy of RCA.)

[8] Because of its complexity, I have thus far omitted discussion of the bias adjustment procedure that RCA performs on the relationality scores before taking their absolute value. I will return to this introductory example later in the paper to describe this procedure in due detail. In particular, I will show that the bias reduction procedure relies on strong distributional assumptions that may be violated in much empirical data, and can cause substantial problems when these assumptions do not hold. See "Distributional Assumptions" section below. Please also note that the RCA software did not offer an option to disable the bias adjustment step. All the RCA analyses I report in this paper thus include this bias adjustment.

Goldberg's technique to use row correlations in place of relationalities. I term the resulting algorithm Correlational Class Analysis (CCA; see Appendix A for details). And indeed, when I applied CCA to the same 800-row dataset, it correctly assigned all the rows into the two schematic classes present in the data. Thus, while RCA failed to correctly recover the schematic classes in Goldberg's example, CCA produced a perfect answer.

SIMULATION

The above analysis suggests that CCA is a more accurate tool than RCA for detecting groups of respondents whose tastes follow the same cultural schema. However, one may rightly object that a single example does not provide a sufficient basis for drawing such a broad conclusion. To rule out the possibility that CCA's apparently superior performance is due to features specific to Goldberg's introductory example, I turn to simulation to carry out a more thorough analysis of the accuracy of both methods. In each of the 10,000 simulation runs reported in this section, I create a dataset where simulated respondents arrange their tastes according to one of a set of randomly generated shared cultural schemas by a process consistent with the theoretical model described above (in later sections, I examine simulations that violate the assumptions of this model). Since the schematic class membership of each respondent is known by design, these simulations enable me to test how accurately CCA and RCA assign respondents to schematic classes under the conditions described by the theory.

The formal definition of shared schemas can serve as the basis for such simulations. Since two response vectors exactly follow the same schema if and only if they are linear transformations of one another, the schema specifying relationships between *N* tastes can itself be specified with a vector $\rho = [\rho_1, \rho_2, \dots, \rho_N]$.[9] Such a schema-specifying vector (hereafter "schema") can be randomly generated by drawing a vector of integers from an appropriate probability distribution. In turn, each response vector X that exactly follows this schema can be generated by randomly drawing a pair of linear transformation constants k and b, which can invert, rescale, or shift the pattern in ways consistent with the theory discussed earlier. Since real survey respondents do not perfectly reproduce cultural schemas, a simulated response must also include substantial stochastic deviations from the schema. These can be introduced via an independent error vector ϵ of the same length as the original pattern, so that $X = k\rho + b + \epsilon$. This is the basic formula behind the first two sets of simulations I examine.

To ensure that the simulations cover a wide range of potential cases, the simulation procedure contains three randomization steps. The first step of each run randomizes the broad characteristics of the simulated dataset, such as the ranges and variances that will be used to generate the values of ρ, k, b and ϵ, as well as the number of distinct taste schemas behind the responses. The second step generates these schemas using the variance parameters produced in the first step. The final step generates a random number of respondents following each of these schemas by applying random linear transformation and adding random noise, both generated using the ranges and variances set in the first step. I repeat the entire procedure 5000 times,[10] creating simulated datasets that widely differ in the ranges of simulated variables, variance of individual responses, signal to noise ratio, and many other parameters. A more detailed

[9] For example, the schema shared by A, B and C in Figure 1 can be specified as $\rho = (0,0,0,1,1,-1,-1)$, so that $A = \rho + 4$, $B = \rho + 2$, and $C = -2\rho + 3$. The schemas are thus also themselves defined up to a linear transformation.

[10] Because of an apparent bug, the RCA software repeatedly crashed for 69 (1.3%) of these simulated datasets, producing no results. I excluded these cases from the analysis.

description of the simulation procedure can be found in Appendix B, and an *R* implementation is also available online.

Figure 4 illustrates a single simulation run. The randomly determined parameters from the first step of the run set the schema variance to 0.51, number of schemas to 3, and the maximum error variance, shift and scaling to $1.02, 1$ and 2, respectively. Thus, to make the schemas A, B and C, the simulation drew vectors from the Normal distribution with $\mu = 0$ and $\sigma^2 = 0.51$, and rounded to them the nearest integer. The resulting schemas are depicted with solid black lines, one per plot. For each schema, the simulation created a set of followers by randomly picking a value of shift *b* from $\{-1,1\}$, scaling and inversion factor *k* from $\{-2,-1,1,2\}$, and a noise vector ϵ drawn from the normal distribution with variance of no more than 1.02. A small sample of such respondents is depicted in dashed lines behind the appropriate schema.

Measuring Accuracy
Each of the 5000 simulated datasets generated by this procedure consists of randomly generated respondents who arrange their tastes according to one of a set of randomly generated taste schemas in a fashion consistent with the formal model I developed here. Thus, as in Goldberg's (2011) introductory example, the true schematic class membership for each simulated respondent is known by design. The simulation's goal is to assess the accuracy with which the group assignments made by the two algorithms correspond to this known membership. If two

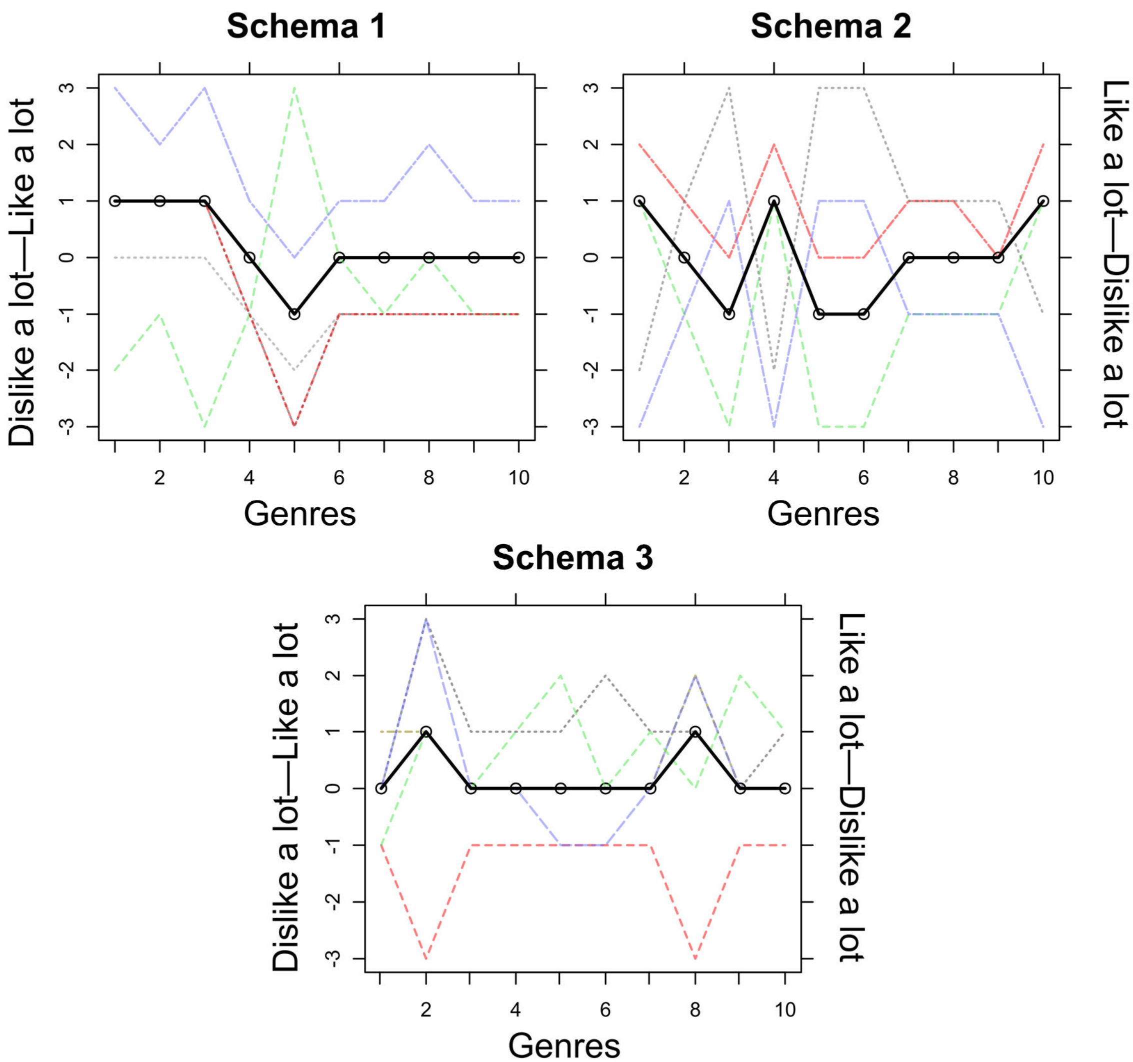

Figure 4. Three simulated schemas (solid black) for one simulation run, with a small sample of responses derived from each schema depicted behind each one in the same plot (dashed). In this run, the pattern variance was 0.51, and the maximum noise variance for individual responses was 1.05. This moderately high ratio of noise to schema variance creates a relatively difficult classification task.

respondents were generated from the same schema, they should be assigned to the same group; if they were created from different schemas, they should belong to different groups.

For each run, I measured this classification accuracy of each algorithm with Normalized Mutual Information (NMI):

$$NMI(\Omega, C) = \frac{2 * I(\Omega; C)}{H(\Omega) + H(C)},$$

where C is the vector of true class memberships for every respondent, Ω contains the corresponding class assignments made by the algorithm, I is mutual information, and H is Shannon entropy. NMI is an established criterion for measuring the accuracy of network partitioning algorithms (e.g., Danon et al. 2005; Lancichinetti, Fortunato, and Kertész 2009). It ranges from 1 when the estimate Ω perfectly recreates the true membership structure C, to 0 when Ω is independent with respect to these true schematic classes (Manning, Raghavan, and Schütze 2008).[11]

Table 1: Comparison of RCA and CCA accuracy in 10,000 simulation runs

Measure	Simulation 1 (5000 runs) Relationality (RCA)	Simulation 1 (5000 runs) Correlation (CCA)	Simulation 2 (5000 runs) Relationality (RCA)	Simulation 2 (5000 runs) Correlation (CCA)
Overall accuracy (median NMI)	0.74	0.87	0.67	0.87
Accuracy, interquartile range (25% to 75%)	(0.54, 0.88)	(0.69, 0.97)	(0.46, 0.84)	(0.69, 0.97)
Runs with near-perfect accuracy (NMI > 0.95)	13.2%	30.5%	9.2%	30.3%
Runs with near-complete inaccuracy (NMI < 0.05)	2.8%	0.1%	3.1%	0.1%
Runs with higher accuracy than other method	5.2%	88.1%	3.8%	91.6%
Odds of higher CCA accuracy in a given run	16.9 : 1		23.8 : 1	

Notes: Results from simulations with randomly varying schema variances, noise amounts, and ranges of linear transformations. In Simulation 2, inversion odds also randomly varied.

[11] In spite of its popularity, this measure may at times produce unexpected results. In particular, NMI can be surprisingly forgiving of solutions that return too many classes: e.g., if the true data contains 4 classes, we would intuitively consider a 100-class solution to be almost completely wrong; however, NMI may still assign it an accuracy score above 0.5. In supplementary analyses, I found that RCA returned such many-class solutions much more frequently than CCA; thus, any potential bias introduced by this measurement property appears to be conservative with regards to my argument here. I am indebted to John Levi Martin for alerting me to potential unexpected behaviors of this measure.

Simulation Results

The results of these 5000 tests are presented in Table 1 under the heading “Simulation 1”. The median accuracy of CCA (0.87) is higher than that of RCA[12] (0.74), a difference that is highly significant statistically (Wilcoxon W = 8585271, p < 0.0001). The interquartile range (IQR) of CCA’s accuracy extends from 0.69 to 0.97, while RCA’s extends 0.54 to 0.88. Thus, while CCA’s 75th percentile is just shy of a perfect accuracy, RCA’s 75th percentile barely surpasses CCA’s median. The substantive significance of these differences is clearer when the CCA accuracies (Y) are plotted against the RCA accuracies (X) in Figure 5. While the accuracies are strongly associated ($R^2 = 0.79$), CCA is more accurate than RCA in the vast majority of cases (88.1%). In contrast, RCA is more accurate than CCA in only 5.2% of the cases. Thus, when RCA and CCA disagree, which they do in 93.3% of the cases, the odds that CCA’s result is more accurate than RCA’s equal 17:1.

To determine if the results point to any classes of data where RCA would nonetheless be preferable to CCA, I disaggregated them by schema variance and noise variance, which are the parameters most responsible for the difficulty of the classification task. Lower schema variances or higher noise variances result in more challenging signal to noise ratios, which should make the performance of both algorithms poorer. The loess curves demonstrating this effect are presented in Figure 6.

On both plots, the accuracies of the two algorithms overwhelmingly rise and fall together, thus suggesting that the algorithms find the same cases challenging and the same cases easy. In spite of this similarity, CCA’s accuracy remains reliably above RCA’s throughout the full ranges of both variances, with a median gap of roughly 0.10 in favor of CCA. This accuracy gap remains remarkably stable throughout most of the two ranges. For example, in simulations within the top 5% of noise variance, the median accuracies equal 0.61 for CCA and 0.51 for RCA. Within the bottom 5% of noise variance, they equal 1.00 and 0.91, respectively. CCA thus retains the same advantage over RCA under both the least and the most challenging noise conditions examined.

[12] RCA software contains a user-configurable filtering step based on pseudo-significance testing, where weak relationalities are dropped prior to partitioning to reduce potential noise. I examined how filtering affected RCA’s performance using 250 simulation runs. Filtering increased accuracy in 56% of the cases but and decreased it in 43%. However, the average decrease (-0.33) was three times greater than the average increase (0.11). Overall, disabling the filter substantially raised RCA’s median accuracy, from 0.56 when enabled to 0.69 when disabled. Additionally, in 10% of the cases with filtering, RCA encountered an error and yielded no solution at all (as compared to 1% without filtering). Thus, to increase RCA’s accuracy and avoid potential bias from substantial missing results, I disabled the filter for all the simulations reported in this paper.

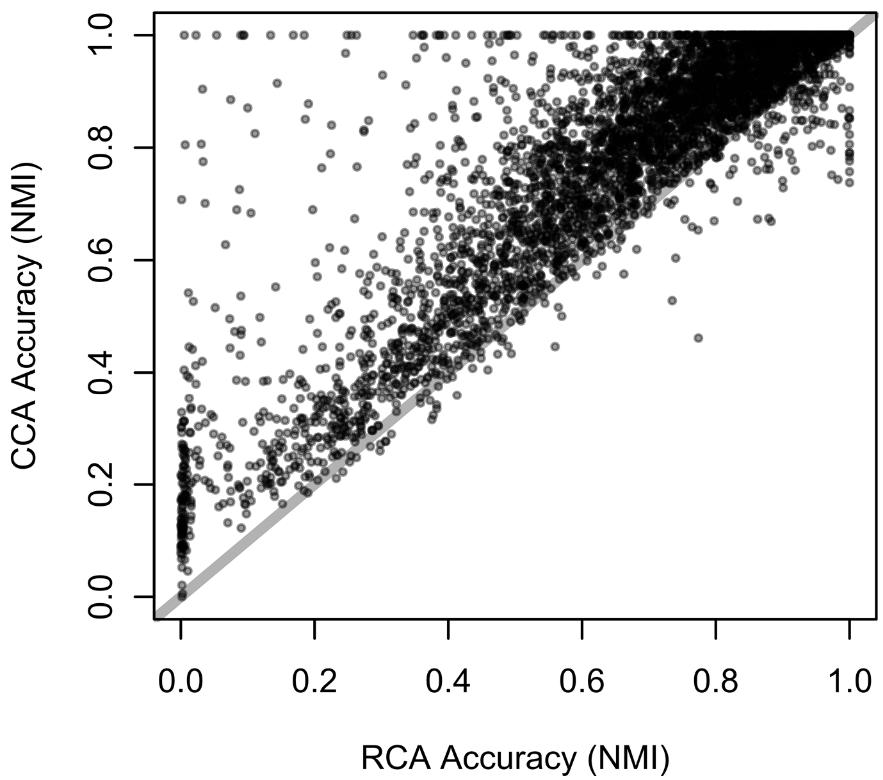

Figure 5. CCA accuracy compared to RCA accuracy for 5000 simulation runs. Each point is a single run. Runs where CCA was more accurate are above the $Y = X$ diagonal (in gray), while those where RCA was more accurate are below. Note the absence of points in the bottom-right corner.

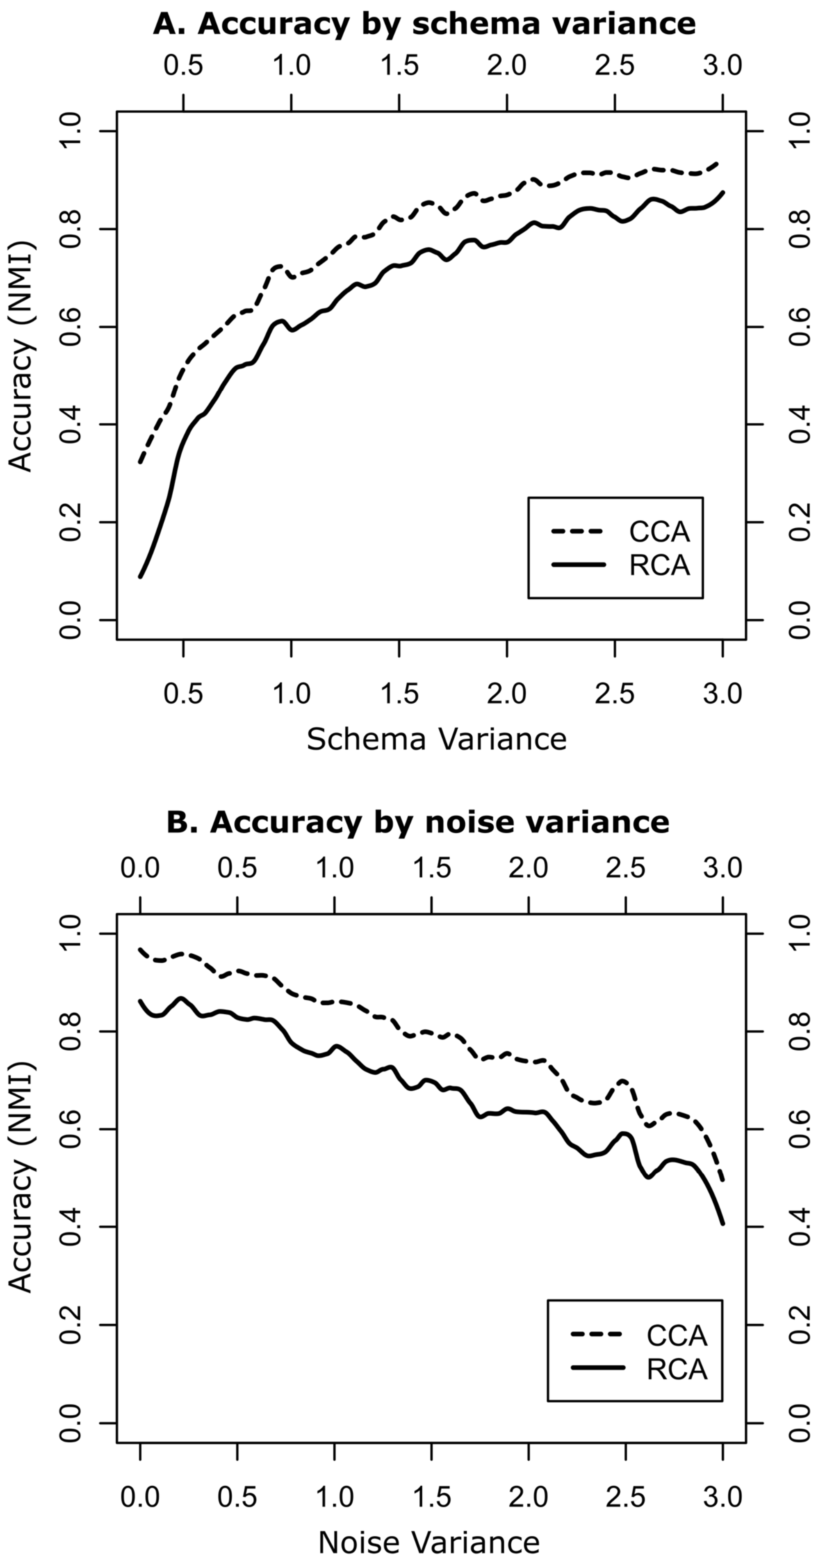

Figure 6. Loess curves comparing RCA and CCA accuracy, based on 5000 simulation runs.

The only instance where the two accuracy curves substantially deviate from each other occurs when the schema variance is low. In such situations, the pairwise correlations between responses tend towards zero, which is a well-known bias, while their relationalities approach one. Therefore, as Goldberg points out, the relationality between low-variance respondents is systematically higher than the correlation. When Goldberg briefly considers using correlation in his online appendix, he dismisses it on the basis of this difference. However, Goldberg incorrectly interprets this difference to mean "relationality does a better job at examining relationships between respondents whose responses have relatively low variance" (Goldberg 2011:Appendix A). The fact that relationality produces *higher* values than correlation does not imply that it produces *more accurate* values. And indeed, as can be seen on the left side of Figure 6A, the opposite appears to be true. When schema variance is in the lowest 5% of its simulated range, CCA's median accuracy decreases to 0.31. However, RCA's accuracy experiences a disproportionally large drop, decreasing to 0.07. This low accuracy indicates that RCA results for low-variance schemas contain almost no information about the true membership structure of the data. It further suggests that relationality may exhibit an upward bias for low-variance observations that is substantially more damaging to its performance than correlation's downward bias.[13]

Distributional Assumptions

Relationality also introduces a strong distributional assumption which may further degrade its accuracy when violated. While a correlation of zero always indicates an absence of a linear relationship, the equivalent "null value" of relationality differs from dataset to dataset and is generally skewed above zero (Goldberg 2011:Appendix A). RCA attempts to compensate for this bias by re-centering the matrix of relationalities by its mean. This bias adjustment procedure crucially rests on the assumption that the true mean relationality between all the rows in the data is zero, or, equivalently, that the relationality values are distributed symmetrically around their null value. This would normally be the case only if the proportion of respondents following a schema without inverting it equals the proportion following its inverse[14]—a quantity I call *inversion probability*. For example, when this probability is 50%, the number of highbrow respondents following the schema "like classical, like opera, dislike rock, dislike country" would equal the number of lowbrow respondents following its inverse, "dislike classical, dislike opera, like rock, like country". However, since in reality the number of highbrow respondents can differ greatly from the number of lowbrow respondents, there is no reason to expect that the inversion probability generally equals 50%. This suggests that RCA's symmetry assumption may be frequently violated by empirical data.

To illustrate the magnitude of potential error that occurs when this assumption is violated, I return to the introductory example. In Figures 1 and 3, the mean relationality between the four respondents equals $\overline{R_{obs}} = (AB + AC + AD + BC + BD + CD)/6 \approx 0.15$. Subtracting this value from the other relationalities yields the adjusted relationalities $|R - \overline{R_{obs}}|$, depicted in dotted lines in Figure 3. In this adjusted result, the related pairs AC and BC are assigned higher relationalities than before, whereas the unrelated pairs AD, BD and CD are decreased. Aside

[13] Those rare respondents with a variance of absolute zero require special treatment. See Appendix D for detailed discussion.

[14] It can also hold in a number of degenerate or improbable cases—e.g., when there are roughly as many taste schemas as there are respondents.

from the relationality of the related pair AB, which is also decreased by the bias adjustment, the bias-adjusted scores $|R - \overline{R_{obs}}|$ appear a clear improvement over the unadjusted values of $|R|$.

However, the bias adjustment procedure can easily have the opposite effect. Consider an alternate dataset which is just like the one in the introductory example, but with the addition of three more copies each of patterns A and B. In this alternate dataset, the mean relationality would then be $\overline{R_{alt}} \approx 0.6$. The relationality matrix that is adjusted for this new mean,$|R - \overline{R_{alt}}|$, is depicted in dashed lines in Figure 3. This adjustment correctly brings the relationalities between AC and BC closer to 1. However, it also causes a large decrease in the relationalities between the related pair A and B, which are identical except for a vertical shift. It further substantially increases the estimated relationalities between unrelated pairs AD, BD and CD. These completely unrelated patterns are now erroneously assigned a *higher* relationality than the closely related pair AB. As this example illustrates, RCA's bias adjustment procedure crucially depends on the relative number of times each pattern appears in the observed dataset—a quantity that is fundamentally arbitrary.

All the simulations presented above have granted RCA's assumption of symmetric distribution of relationalities by keeping the inversion probability fixed at 50%. To examine the effects of relaxing this assumption, I created a second simulation where the inversion probability is instead drawn from a uniform distribution over its full range, and varies between each of the 5000 simulation runs.[15] The results of this second simulation are reported on the right side of Table 1. As expected, both the median (0.86) and the interquartile range (0.69 to 0.97) of CCA accuracies remain unchanged from the first simulation. On the other hand, the median accuracy of RCA drops to 0.67, significantly lower than its prior median of 0.74 (Wilcoxon W = 13802245, $p < 0.0001$), and below CCA's 25th percentile of accuracy. CCA is now more than 3 times as likely as RCA to produce a nearly perfect answer (NMI > 0.95), while RCA is more than 20 times as likely to produce an almost completely incorrect one (NMI < 0.05). Thus, RCA's accuracy appears to suffer a significant further drop when its assumption of symmetrical distribution is violated, as may generally happen in empirical applications. CCA remains unaffected by this change.[16]

The results of both simulations thus reinforce my earlier suppositions. Though the accuracies of RCA and CCA were highly correlated, they nonetheless differed in almost 95% of the 10,000 combined simulation runs. In simulation 1, which obeyed RCA's symmetry assumption, the odds that CCA's result was more accurate equaled 17:1. In simulation 2, where this assumption was relaxed, these odds further rose to 23:1. CCA was more accurate over the full range of schema and noise variances examined. The only major deviation from this otherwise stable gap in accuracies occurred in simulations with low schema variance. In those

---

[15] See Appendix B ("Random Inversion Probability") for details of procedure.

[16] For a point of comparison, note that Pearson's correlation is only guaranteed to be unbiased and asymptotically efficient if the vectors are bivariate normal in distribution. To examine whether this is the case with my primary simulations, I applied the Royston H test for multivariate normality (Royston 1983) to 425 datasets randomly generated by the full linear simulation procedure from Appendix B. In each one of the 425 datasets, the hypothesis of multivariate normality was rejected at $p < 0.0001$. I then partitioned each simulated dataset by schematic class membership, so that respondents in each resulting dataset would come from exactly one schematic class. In each one of these 1728 single-class datasets, the hypothesis of multivariate normality was again rejected at $p < 0.0001$. The data produced by my linear simulation procedure thus violate correlation's bivariate normality assumption. Note that CCA was nonetheless able to analyze many of the simulated linear datasets with near-perfect accuracy.

simulations, RCA performance suffered a significant additional drop relative to CCA's. Overall, the simulations provided no evidence of any use cases where RCA would be preferable to CCA.

DEVIATIONS FROM THE MODEL

The 10,000 simulations I described above demonstrate that CCA yields substantially more accurate results than RCA when the shared cultural schemas result in linear dependencies between rows in the survey dataset. Such linear dependencies are consistent with the motivating examples of schematic similarity in Goldberg's (2011) work, as well as the formal reasoning behind CCA I laid out here. As I argue above, they are also socially and cognitively plausible. Nonetheless, schematic cultural cognition is still little understood, and whether this model of schematic similarity as linearity is correct is not known. For this reason, the performance of both algorithms when this model is *violated* is also of interest. Another reason for this interest stems from the close conceptual relationship between correlation and linearity, which could potentially lead correlation-based CCA to perform disproportionally well in simulations that obey the linearity assumption. Since Goldberg may not have been aware that the examples he presented in his work were linear in character, he did not specifically design RCA to detect linear transformations. Simulations that violate the model would thus yield a more conservative test of CCA's accuracy, and could potentially point to use cases where RCA has an advantage over CCA. For these reasons, I now turn to simulations where respondents' patterns are produced through processes that differ from the linear transformations theorized here.

There are two fundamentally distinct ways that schematic transformations could deviate from the theorized model. First, schematic transformations could in reality be based on only a subset of the possible linear operations (shift, scaling and inversion). Though such cases would technically remain linear, the two algorithms may perform differently on these "edge cases" of linearity than they do with all three operations present.[17] And second, schematic transformations could be based on operations which are non-linear in character, such as in the presence of multi-way interactions between variables. I will use the shorthands "sub-linear" and "super-linear" to refer to these two broad kinds of deviations from the model. In this section, I examine six distinct sets of simulations to investigate both of these possibilities.

Sub-linear Transformations

Recall that musical taste schemas are theorized to consist of relative evaluations of different musical genres. Because these evaluations are relative rather than absolute, individuals can use one schema to produce a variety of contrasting taste patterns. Specifically, they can invert the schema, turning likes into dislikes; they can scale it, making all the appraisals uniformly more or less extreme; and they can vertically shift it, making all the appraisals more positive or negative.[18] There is, to my knowledge, no alternate theory that clearly articulates a different set of cognitive processes that may underlie this kind of schematic cognition. However, there are plausible scenarios in which each of the three theorized transformations would not exist. And,

[17] Among other reasons, such differences in accuracy may arise if two measures have different amounts of sensitivity to inversion, scaling and shift. *Ceteris paribus*, the measure that is more likely to identify that X and its noisy inverted copy $Y = -1 * X + \epsilon$ follow the same schema would probably yield a more accurate solution in the linear simulations examined above. However, the same property would then make it more likely to yield false positives in the "no inversion" scenario, where two patterns that appear to be inverses of one another would actually be produced by two different schemas.

[18] I describe these operations in terms of individual action largely for the sake of readability. The actual socio-cognitive processes of schematic transformation could well be supra-individual.

given the noisiness of much survey data, a method that searches for rare or nonexistent schematic transformations could generate substantial amounts of false positives, thus yielding inaccurate results. In the absence of empirical work documenting that the three theorized processes do indeed take place, it appears useful to investigate the robustness of both methods to their potential absence. In this section, I use three sets of simulations to do so.

I begin by briefly sketching some potential reasons to question the prevalence of the basic linear operations of scaling, inversion, or shift. My goal here is not to argue that any of these operations are likely to be absent, but rather to point out that such an absence is plausible enough to warrant further attention. To do so, I return to Figure 2, which illustrates each of the three basic linear transformations in isolation.

The patterns in this figure are all derived from one shared schema (labeled "original"). Respondent *P* follows this schema without any transformations, and thus likes both rock and rap ($+2$) and mildly dislikes country, folks and oldies (-1). Respondent Q uses the same schema as *P*, but inverts it to create the opposite pattern of tastes ($Q = -1 * P$). The cognitive demands of such an inversion are minimal, as it simply requires taking a pattern of tastes and replacing the "likes" with "dislikes". However, whether two social actors with very different taste patterns like "highbrow" and "lowbrow" would actually employ exactly the same cultural schema is uncertain. Since cultural taste is intertwined with social position, individuals with dramatic cultural differences are also likely to live in different kinds of places, consume media from different sources, and have network ties to distinct kinds of alters (Bourdieu 1984; Fararo and Skvoretz 1987; McPherson 2004; McPherson, Smith-Lovin, and Cook 2001). This social distance suggests that, rather than inverting the same taste schema, two groups with contrasting tastes may employ separate cultural schemas that maintain only an indirect, imprecise, symbolic or accidental opposition to one another.

It is also possible to raise questions regarding the empirical prevalence of vertical shifts like the one employed by respondent *R* ($= P - 3$). While the original schema consisted of a pattern of contrasting positive and negative appraisals, *R*'s pattern contains only genres she dislikes a little (rock and rap $= -1$) and those she dislikes a great deal (country, folk and oldies $= -4$). Such a transformation could occur if respondent *R* had internalized the same set of distinctions as *P* but simply did not enjoy listening to music, thus giving lower appraisal to all the genres. However, internalizing cultural distinctions takes effort and time (e.g., Bourdieu 1984). It is thus reasonable to wonder if respondents who dislike music would *generally* put in the effort to internalize the same complex set of genre contrasts as music lovers, who presumably have a far greater intrinsic motivation to expose themselves to music.

The case of respondents with multiplicatively scaled response patterns like S ($= P * 2$) could raise similar questions. Unlike *P*, who perceives relatively mild distinctions between musical styles, *S*'s tastes feature dramatic contrasts between very strong positive assessments of some genres and very negative assessments of others. This again invites the same concern regarding generally unequally levels of motivation.

These questions about the empirical prevalence of the three kinds of linear transformations will need to be settled by future empirical work. In the meantime, they merit investigating how CCA and RCA would perform in the scenario that each of these transformations does not take place. Among the second set of 5000 simulations I reported earlier in the paper, 1185 featured no vertical shift ($b \in \{0\}$) and 1700 featured no scaling ($k \in \{1\}$). To examine the performance of CCA and RCA in the absence of these transformations, I reanalyzed

Table 2: Comparison of RCA and CCA accuracy in six sets of simulations with departures from theorized model

Simulation	RCA accuracy (median & IQR)	CCA accuracy (median & IQR)	Methods equally accurate (% of runs)	CCA more accurate (% of runs)
Sub-linear:				
No shift	0.68 (0.46, 0.86)	0.84 (0.65, 0.96)	5.4%	88.0%
No scaling	0.66 (0.44, 0.84)	0.82 (0.61, 0.95)	5.1%	88.5%
No inversion	0.41 (0.21, 0.57)	0.87 (0.69, 0.97)	0.3%	98.7%
Super-linear:				
Random polynomial functional form	0.27 (0.16, 0.38)	0.42 (0.31, 0.57)	0.1%	94.2%
Independent subschemas	0.49 (0.31,0.69)	0.69 (0.49,0.85)	1.1%	76.1%
Multi-way interactions	0.23 (0.12, 0.42)	0.52 (0.26, 0.79)	0.1%	95.8%

these cases in isolation. I also performed 1000 further simulations to examine the performance of the two algorithms in the absence of schematic inversion ($\zeta = 0$).

I report the accuracy of CCA and RCA in the three sub-linear scenarios in Table 2. Among the "no shift", "no scaling" and "no inversion" simulations, CCA had median accuracies of 0.84, 0.82 and 0.87, respectively. These are slightly lower on the average than its prior accuracy of 0.87 (Table 1, simulation 2), though the difference is minor. With RCA, the "no shift," "no scaling" and "no inversion" simulations yielded accuracies of 0.68, 0.66 and 0.41, respectively. While the first two of these scores are close to RCA's prior accuracy of 0.67, its score in the "no inversion" scenario is almost 40% lower. RCA was thus capable of retaining its prior accuracy in the "no shift" and "no scaling" scenarios, but not the "no inversion" scenario. This large drop in accuracy may again originate with RCA's strong assumption of symmetric distribution discussed above.

Comparing the CCA and RCA accuracies on each individual run showed that CCA was more accurate in 88%, 88.5% and 99% of the scenarios that omitted shift, scaling and inversion, respectively. Thus, CCA remains the preferred choice of method if there is doubt regarding the empirical prevalence of some of the theorized linear transformations, and especially so if these doubts concern schematic inversion. In Appendix C, I also show how Pearson's formula can be altered to produce modified coefficients that specifically reflect each of these three sub-linear scenarios. The three modified coefficients treat inversion, shift, or scaling as schematic difference rather than schematic similarity. Though CCA performed well in each of these settings, it may be possible to use such "specially tuned" correlation coefficients to increase its accuracy even further while retaining the method's speed and simplicity.

Super-linear Transformations

It is also conceivable that shared cultural schemas result in schematic transformation that are more complex than scaling, shift, and inversion. The theory of schematic similarity elucidated here does not predict the existence of such transformations. However, Goldberg intended relationality to function even in the presence of such non-linear relationships between variables (e.g., Goldberg 2011:1433), arguing that the complex way that relationality is computed makes it "more sensitive to interdependencies" between variables than correlation (Goldberg 2011:Appendix A). In this section, I use further simulations to investigate whether relationality's complexity in fact makes it more accurate than correlation in the presence of such "super-linear" transformations.

There are at least three broadly distinct kinds of conceivable super-linear transformations: polynomial (rather than linear) functions, independence between parts of the taste schema, and multi-way interactions between tastes. All three of these scenarios contradict the theory of schematic similarity developed here. All three also violate different assumptions made by Pearson's correlation, thus creating potential difficulties for CCA with respect to RCA. To be clear, these three super-linear scenarios are not backed by well-articulated theoretical models of shared cultural schemas like the one developed in this paper. These simulations are thus general tests of robustness to potential unforeseen deviations from the theorized model rather than examinations of concrete alternate use cases for either method.[19]

Functional form. I begin by using polynomials to examine deviations from the model that can come from more complex functional forms. High-degree polynomials are straightforward to generate randomly and can be used to approximate a wide variety of other continuous functions, with the maximum degree of the polynomial providing the limit on the number of inflection points and thus the complexity of the transformation. In contrast to the linear transformations, which always express clearly understandable processes of scaling, inversion, and shift, most of these random polynomial transformations have no clear theoretical justification.[20] They are instead informative because of the diversity of transformations they produce. This scenario thus provides insight on the methods' performance under a broad range of unspecified deviations from the theorized model.

In these simulations, I replace linear transformations with randomly generated polynomials of the form:

$$X = g(\rho) = \sum_{i=0}^{10} a_i b_i \rho^i = a_0 b_0 + a_1 b_1 \rho + a_2 b_2 \rho^2 + \cdots + a_{10} b_{10} \rho^{10},$$

where $a_i \in \{-1,0,1\}$ and $b_i > 0$. The a_i coefficients determine whether a polynomial term is present as well as whether it is inverted, and the b_i coefficients determine its relative weight. I adapt the three-step randomization process described earlier to generate such polynomials.[21]

---

[19] Any such examination would instead begin with an alternate model of how such schemas are represented and transformed (see "Limitations and Future Directions" section below).

[20] One example of a polynomial transformation that does have an interpretation is a "center-versus-extremes" transformation that retains the extremity of the tastes in the schema but ignores the valence, thus treating "like a lot" and "dislike a lot" as equivalent (e.g., $X = \rho^2, -10 \leq \rho \leq 10$).

[21] See online source code supplement for details.

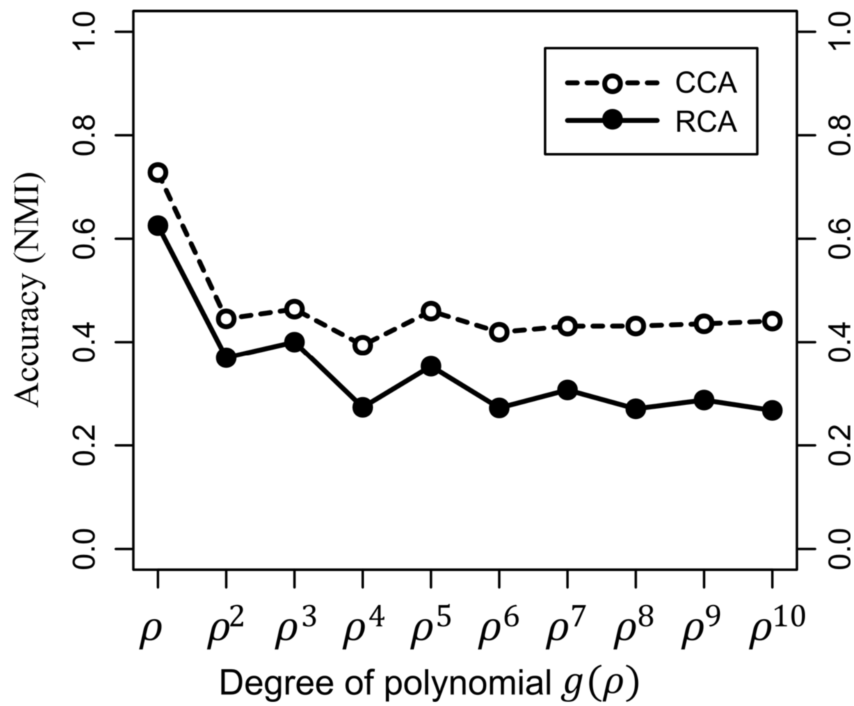

Figure 7. Accuracy for 3000 simulations with polynomial transformations.

The results of 2000 such simulations are reported in Table 2 with the label "Random functional form." In these simulations, CCA obtained a median accuracy of 0.42, which is roughly half of its prior accuracy of 0.87. RCA obtained a median accuracy of 0.27, or roughly 40% of its prior accuracy of 0.67. CCA yielded more accurate results in 94.2% of the runs. Thus, while polynomial transformations proved challenging for correlation, they appeared no less challenging for relationality.

In Figure 7, I disaggregate these results by the maximum degree of the random polynomial transformation in the simulation run, which tracks the complexity of the transformation. Some simulations were randomly assigned first-degree polynomial transformations, which are analogous to the linear transformations used before (albeit with different signal-to-noise ratios). The figure shows that CCA as well as RCA had their highest accuracies with these first-degree transformations. The accuracy of both algorithms then dropped as the degree of polynomial increased. However, CCA remained reliably more accurate than RCA throughout the entire range. These results thus yield no evidence that RCA's computational complexity improves its ability to detect more complex functional forms of schematic similarity.

Independent Subschemas. All the reported simulations have thus far assumed that each cultural schema specifies the relationships between all the attitude objects. However, it is conceivable that some respondents may use distinct cultural "subschemas" to organize different subsets of their tastes. For example, they may use one subschema to specify the relationship between genres to which most Americans are frequently exposed on radio and television, and a different subschema for genres encountered only in more specialized settings. Else, they could use one subschema for genres they were exposed to by their parents, and another subschema for genres they learned about through friends. Such respondents could then apply different transformations to different subsets of tastes, e.g., by inverting the subschema for popular genres without inverting the one for specialized genres.

When different transformations are applied to different parts of a schema, they combine into overall schematic transformations that cannot be captured in terms of simultaneous deviations from each vector's single common mean. While the correlation coefficient is computed from these deviations, relationality instead compares the deviations of variables from one another (Goldberg 2011: Appendix A), which could potentially give it an advantage under this kind of non-linearity. Any such advantage, however, would be limited in its usefulness, as both CCA and RCA assign each respondent to exactly one class that is assumed to correspond to exactly one schema, and thus cannot possibly produce the correct solution in most situations with independent subschemas.[22] Extending CCA/RCA to cover such cases would require substantial further theoretical and methodological development.

In their present form, the methods can still be applied to an edge case of this scenario where different subschemas cannot be arbitrarily recombined to produce an individual's pattern of tastes, but rather always occur in combination with exactly the same subschemas.[23] This is the scenario I simulate here. As this limitation greatly reduces its realism, this set of simulations is again presented only as a test of robustness rather than a theoretical alternative to the model of schematic similarity as linearity.

---

[22] Consider a situation where there are three schemas *a*, *b* and *c*, each of which can describe any musical genre. Imagine then that there are four groups of respondents: group 1 uses *a* to arrange all of their tastes; group 2 uses *a* to arrange tastes for popular genres only, and *b* for the remaining specialized genres; group 3 instead uses *a* for the specialized genres and *b* for popular ones; and group 4 uses *c* to evaluate jazz and rap, while using *a* for all remaining genres. It is not possible to partition this population in such a way that each respondent is assigned to a single class that corresponds to a shared schema.

[23] That is, if *a, b, c* and *d* are subschemas, and group *X* uses the combination (*a,b*) to arrange their tastes while group *Y* uses combination (*c,d*), we must assume that there is no group *Z* that uses any other combination of these subschemas—(*a,c*), (*b,c*), *(a,d*) or (*b,d*). We must also assume that *a* defines no genres defined by *b,* and *c* no genres defined by *d.*

To examine this scenario, I simulated 2000 datasets in which individual responses were derived from the schemas that contained up to four different subschemas inside of them.[24] I report the results of these "independent subschemas" analyses in Table 2. In these simulations, the median accuracy of CCA was 0.68, and the median accuracy of RCA was 0.49. Thus, for both CCA and RCA, this scenario proved substantially more challenging than the linear ones, where their accuracies were 0.87 and 0.67, respectively. Overall, CCA produced more accurate results in 76% of the runs.

In Figure 8, I disaggregate these results by the average count of independent subschemas within the schemas used in each simulation run, which tracks the extent to which a run deviates from the assumption of variance around a single common mean. The runs where this count equals 1 are analogous to the linear simulations analyzed above (albeit again with different signal-to-noise ratios). CCA's median accuracy in these single-subschema runs was 0.81, and RCA's was 0.55. The accuracy of both algorithms degraded as the number of subschemas grew, falling to 0.44 for CCA and 0.31 for RCA when this quantity reached the simulation maximum of 4. CCA again remained consistently more accurate than RCA throughout the full range of the simulation.

Multi-way interactions. I now turn to potential departures from the model caused by interactions between variables in the taste schema. In the examples presented by Goldberg (2011) and the simulations reported above, each elements of a respondent's vector R was always produced from transformations of exactly one element of the schema ρ. For example, consider a respondent R who begins with a taste schema $S = [\mathrm{S}_{\mathrm{rock}} = 2, \mathrm{S}_{\mathrm{rap}} = x, S_{classical} = y, S_{opera} = z]$, then inverts it, scales it by 2 and finally shifts it by +1. This respondent is, in effect, applying a

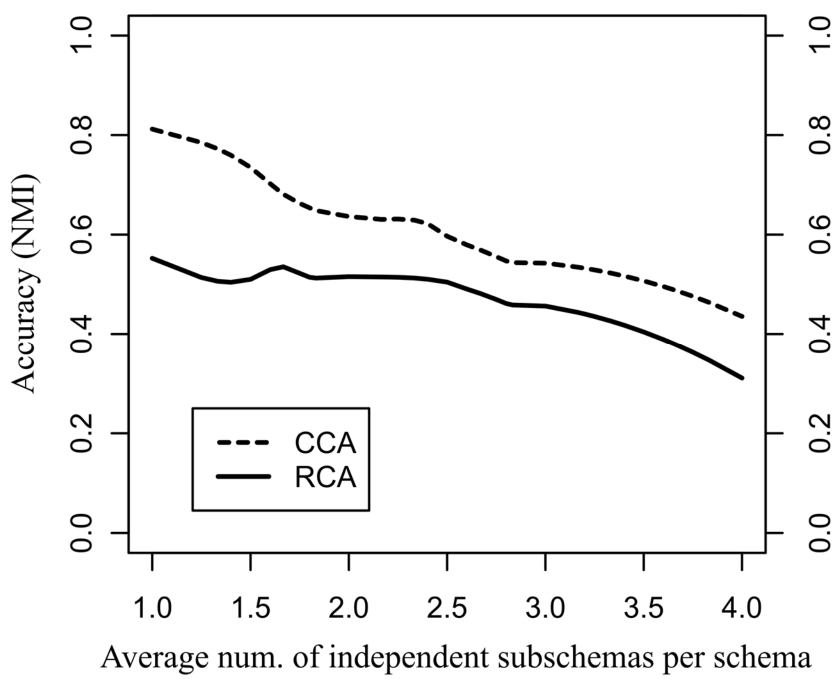

Figure 8. Accuracy for 2000 simulations with independent subschemas.

[24] Source code available in online supplement.

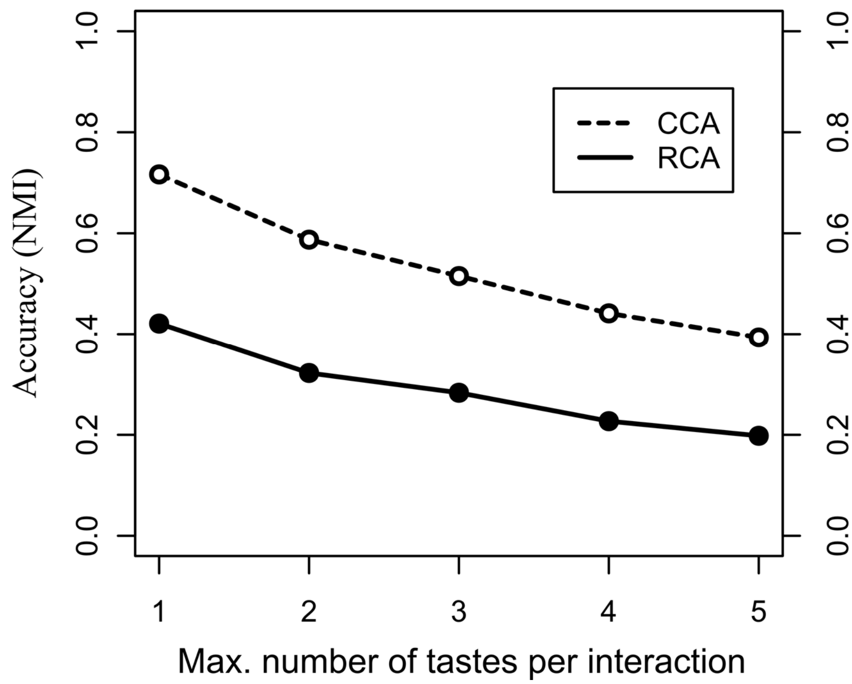

Figure 9. Accuracy for 2000 simulations with multi-way interactions.

univariate function $R_k = f(S_k) = (-1) * 2 * S_k + 1$ to each element of *S*. No matter what values *x, y* and *z* assume, *R*'s attitude to rock would equal $f(2) = -3$.

In contrast, in the presence of interactions between attitudes, a given respondent's attitude regarding a single genre may be a function of two or more tastes in the schema, e.g., $R_{rock} = g(S_{rock}, S_{rap}, S_{classical}) = (S_{rock}^2 * S_{rap} * S_{classical})^{\frac{1}{4}} = \sqrt[4]{4xy}$. The presence of such multi-way interactions would entail complex and perhaps unintuitive relationships between cultural schemas and the responses they generate[25]. In contrast to the linear transformations that are the focus of this paper, it is difficult to imagine any plausible socio-cognitive mechanism that would consistently yield such intricate transformations. However, in the interest of examining the robustness of both algorithms under a maximally broad range of unexpected departures from linearity, I use further simulations to study their performance in this setting.

In this procedure, each simulated respondent partitions a shared schema into groups of tastes I term "interaction blocks", and then generates each genre taste by calculating its randomly weighted geometric mean with other tastes in its block[26]. Each respondent thus uses a random

[25] A population of respondents who use transformations such as these would be introducing cultural logics that are not part of the schema. Consider, for example, a schema like $S = [\mathrm{S_{rock}} = 1, \mathrm{S_{rap}} = 1, S_{classical} = 4, S_{opera} = 4]$ on a scale from 1 = like to 4 = dislike, and a respondent Z using the transformations $Z_{rock} = (S_{rock}^2 * S_{opera} * S_{classical})^{.25} = 2$ and $Z_{classical} = (S_{classical}^2 * S_{rock} * S_{rap})^{.25} = 2$. Though the lowbrow/highbrow schema *S* placed rock in opposition to classical, and *Z* produced her tastes without any error, she nonetheless ended up with an omnivorous taste pattern.

[26] The general functional form of the transformation that produces a respondent's taste R_k in interaction block $\{R_k, R_m, R_n\}$ from schema ρ is the weighted geometric mean $R_k = g(\rho_k, \rho_m, \rho_n) = \sqrt[a+2]{\rho_k^a \rho_m \rho_n}$. I attach

transformation of the form similar to g above, e.g., $R_k = g(k, m, n, p) = \sqrt[6]{k^3 mnp}$. Since each respondent is assigned idiosyncratic interaction blocks, each response vector consists of multiplicative interactions between different sets of tastes. To create a range of complexities, the maximum number of tastes per interaction block differs at random from run to run.

I report the results of 2000 runs of this simulation in Table 2 ("multi-way interactions"). The median accuracy was 0.52 for CCA and 0.23 for RCA, indicating that correlation again yielded more accurate results than relationality. To examine how the complexity of interactions affected accuracy, I disaggregated the simulation runs by maximum number of tastes per interaction block in the given run (see Figure 9). When this number equals 1, the transformations add noise but do not actually feature any interactions between terms. As expected, this is the easiest scenario for both CCA and RCA, which achieve accuracies of 0.72 and 0.42, respectively. These simulations become increasingly more challenging for both CCA and RCA as the number of tastes per interaction increases. By the time the interaction size reaches its maximum of 5, both CCA and RCA see their initial accuracies drop by roughly half, to 0.39 and 0.20, respectively. Throughout the 2000 runs of this simulation, CCA yielded a more accurate result than RCA on 96% of the runs, indicating that it is the preferable technique even in the presence of complex interactions.

The simulations I examined in this section compared three broad ways in which relationships between taste vectors can violate the assumptions of linearity. I examined situations where linear transformation functions were replaced by polynomials; where some parts of the taste schema were transformed independently from others; and where transformations of the taste schema involved multi-way interactions between separate tastes. While Goldberg (2011) contended that relationality's computational complexity makes is better equipped than Pearson's correlation to handle complex relationships between variables, these simulations indicate that this is likely not the case. CCA yielded more accurate results than RCA in 94% of simulations with random polynomial functional forms, 96% of simulations with multi-way interactions, and 76% of simulations with independent subschemas (see Table 2). When I disaggregated each scenario by difficulty, I found that RCA's accuracy consistently trailed RCA's. The simulations thus yielded no evidence of cases—linear or nonlinear—where RCA would be preferable to CCA. Moreover, it is interesting to note that, in Figure 7, where the relationship between polynomial degree and classification accuracy was highly non-monotonic, RCA's accuracy still repeatedly rose and fell in close synchrony with CCA's. The drops in accuracy again closely resembled each other in the presence of interactions (figure 9). And, though the two methods responded differently to small numbers of independent subschemas, the drops to RCA's accuracy also closely paralleled CCA's through most of that simulation's range (right side of figure 8). This remarkably persistent pattern suggests that, minor differences aside, relationality may generally be detecting the same substantive relationships between respondents as correlation, albeit at a consistently lower accuracy.

EMPIRICAL EXAMPLE: MUSICAL TASTES

To compare the results produced by the two methods in an empirical setting, I applied CCA to the 1993 GSS music tastes module previously analyzed with RCA (Goldberg 2011).

---

weight a to the focal element ρ_k because without weights the respondent's tastes for all genres in the interaction block would be identical: $\sqrt[3]{\rho_k \rho_m \rho_n} = \sqrt[3]{\rho_m \rho_k \rho_n} = \sqrt[3]{\rho_n \rho_k \rho_m}$. I override this unintuitive behavior in roughly 75% of the runs by randomly drawing $a \sim U\{1,2,3,4\}$. When $a > 1$, ρ_k has a stronger effect on R_k than do ρ_m or ρ_n. See online source code supplement for detailed simulation procedure.

This dataset contains 1532 respondents' evaluations of 17 musical genres. Each respondent rated each genre using a five-point Likert scale that ranges from "like very much" to "dislike very much." For comparability, I followed exactly the same coding procedures as Goldberg (2011). Below, I contrast the substantive contents of the results. I then compare their fit to the data using a multiple groups analysis technique from structural equation modeling.

The RCA analyses partitioned the survey population into three schematic classes, which Goldberg labeled "Omnivore – Univore", "Highbrow – Lowbrow", and "Contemporary – Traditional." For respondents in the "Omnivore – Univore" class, most genre tastes were positively correlated among each other. Goldberg interpreted this as evidence of a culturally omnivorous taste schema, in which no genres are perceived as opposites, but rather a high appraisal of most genres is opposed to a low appraisal of most genres. In the "Highbrow – Lowbrow" class, tastes for "elitist" genres such as opera and classical music were positively correlated among each other, but negatively correlated with most tastes for popular genres. Finally, Goldberg characterized the third schematic class as "Contemporary – Traditional." Here, a cluster of positively correlated tastes for well-established musical genres including gospel, bluegrass, and country is negatively correlated to tastes for arguably more contemporary genres, including heavy metal, pop and rap, as well as oldies and jazz.

The persistence of the Highbrow – Lowbrow taste schema was perhaps Goldberg's most surprising finding, as much contemporary work has argued that omnivorousness has replaced highbrow tastes as a marker of high status in the contemporary United States (Peterson and Kern 1996; Peterson 1997, 2005; see Goldberg 2011 for a more detailed discussion). The CCA analyses of these data, however, do not replicate this finding. While RCA identified three classes in these data, CCA identified four, which are presented in Figure 10. The first two of these closely resemble those located by RCA. The first class features practically no negative correlations between the genres, suggesting that respondents in this class perceive little opposition between different musical styles. In this population, positive appraisal of any one genre generally "goes with" positive appraisal of any other genre, suggesting an undiscriminating logic of taste that ranges between near-uniformly positive appraisals of all genres on one extreme, and a near-uniformly negative appraisal of all genres on the other. This is the same omnivorous logic as behind the Omnivore – Univore class identified by RCA.

The second class located by CCA appears to be defined by an opposition between rock, rap, and metal on one extreme, and gospel, country, folk, and bluegrass on the other. This suggests a bifurcation of respondents into those who prefer newer musical genres and those who prefer more established ones, which closely resembles the logic of the Contemporary – Traditional class identified by RCA. The two methods, however, deviate in their classification of blues and jazz, which RCA had categorized as "contemporary" rather than "traditional." In contrast, CCA analyses instead suggest that blues belongs to the "traditional" side of the divide, while jazz, along with latin, straddles the two sides without clearly belonging to either. Since both blues and jazz were already well-established genres by the mid-twentieth century, the Contemporary – Traditional class identified by CCA fits better with the intuitive chronological understanding of those terms, and may thus have greater face validity.

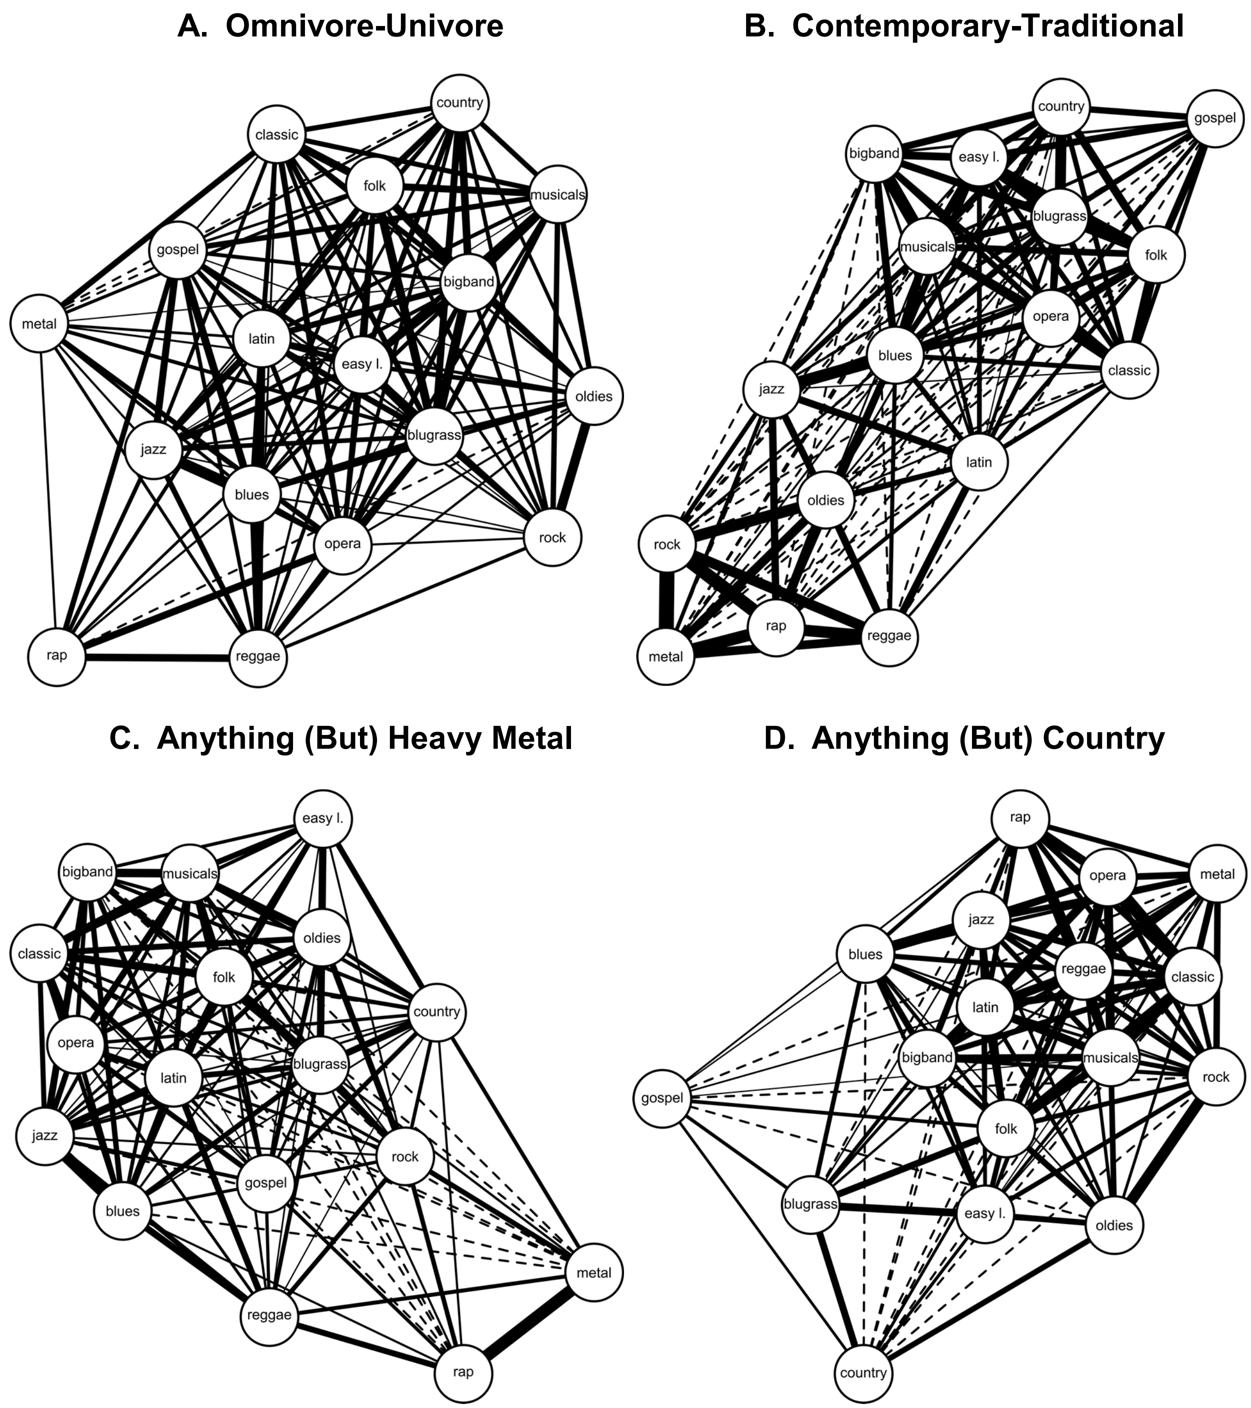

Figure 10. Networks illustrating the four schematic classes identified by CCA. Ties represent correlations between genre tastes, with stronger correlations indicated by thicker lines. Negative correlations are depicted with dashed lines. Weak correlations ($|r| < 0.05$) not plotted.

Table 3: Cross-tabulation of Estimated Schematic Class Memberships in 1993 GSS Music Tastes Data

	CCA class			
RCA class	Omnivore – Univore	Contemporary – Traditional	Country – Anything But	Heavy Metal – Anything But
Omnivore – Univore (673)	281	92	167	133
Contemporary – Traditional (394)	60	241	35	58
Highbrow – Lowbrow (461)	142	36	159	124
Total (N=1528)	483	369	361	315

Notes: Cross-tabulation of class memberships estimated by CCA (columns) and RCA (rows). RCA class sizes are indicated in parentheses. Four out of 1532 respondents had no response variance and were omitted from this analysis.

The biggest apparent difference between the CCA and RCA analyses concerns the remaining group of respondents. RCA had placed the remaining respondents into a single class, which appeared to follow a traditional hierarchical logic with "highbrow" genres such as opera and classical music on the one extreme, and popular "lowbrow" genres on the other. In contrast, CCA separated the remaining population into two further schematic classes which appear to follow a different set of logics (see panels C and D of Figure 10). In both CCA-identified classes, the majority of genres are tied together in a dense cluster of positive correlations, thus suggesting that both are variants of an omnivorous taste schema. These results resemble previous findings that documented the existence of multiple distinct logics of omnivorousness (e.g., Tampubolon 2008).

However, the omnivorousness of respondents in these last two classes features clear exceptions. In the class depicted on the left, a higher appraisal of most genres is generally accompanied by a *lower* appraisal of heavy metal (and frequently also rap music.) The class on the right exhibits a nearly identical structure, except country and gospel music occupy the same position of exclusion as metal and rap did in the class on the left. These patterns closely echo the analyses of Bryson (1996), who famously showed that omnivores may retain a symbolic boundary against genres most closely associated with low education: heavy metal, rap, country and gospel music. Thus, drawing on the title of Bryson's work, I term these latter two classes "Anything (but) Heavy Metal" and "Anything (but) Country."[27]

Table 3 cross-tabulates the group assignments made by the two algorithms. A plurality (though not a majority) of respondents that RCA grouped into its first two classes remain grouped together in the CCA results. The third class identified by RCA, however, does not have an analogous relationship to any one of CCA's classes. Such a divergence between CCA and RCA class assignments is consistent with the simulation analyses above, where the results of the two methods were correlated but rarely the same. While the simulations had demonstrated a

[27] The parentheses around "but" in "Anything (but) Country" are meant to signify that the class could equally be named "Anything But Country" or "Anything Country"—i.e., that the logic opposes "any music that *isn't* country" to "any music that *is* country." The same logic applies to "Anything (but) Heavy Metal."

reliable difference in accuracy between the two methods, these empirical analyses confirm that the two methods can point to substantively different conclusions.

Multiple Groups Analysis

Both RCA and CCA are theoretically understood as methods for partitioning the population into groups of respondents who follow different cultural schemas. However, from a more practical perspective, their task is to detect unobserved heterogeneity in how attitudes or tastes correlate for different groups. This heterogeneity is exemplified by the hypothetical situation where a pair of attitudes are correlated positively in one subgroup and negatively in another, leading them to appear uncorrelated in the overall sample (e.g., Baldassarri and Goldberg 2014:56). RCA and CCA reveal these latent logics of taste organization by inductively partitioning the population into groups where the same pairs of tastes correlate differently with one another. The sets of logics these methods locate are thus subsample correlation matrixes which were mixed together to yield the overall sample.

Their success at locating this kind of unobserved heterogeneity can be measured by how well these resulting subsample correlation matrixes describe the observed distribution of tastes in the population. A well-known multiple groups testing technique from structural equation modeling can provide such a measure (Bollen 1989; review in MacCallum and Austin 2000). Multiple groups analysis is widely used in psychology and public health to ascertain whether survey scales "have the same meaning across groups" (Gregorich 2006:S78), such as across age cohorts or education levels. Here, we are interested instead in ascertaining that the subgroups located by one of the methods show evidence of substantial *differences* in meaning.

To test the extent of this difference, I use SEM to evaluate whether using separate correlation matrixes for each subgroup meaningfully improves model fit over a single correlation matrix for the whole population, taking parsimony into account. The differences in AIC and BIC between the single-matrix and multiple-matrix models can be used to quantify the "statistical significance" of the heterogeneity detected by RCA or CCA. Roughly speaking, these indicators answer the question "are the groups located by these methods distinct enough in their logics of taste organization to make partitioning the population worth it?" Since neither AIC nor BIC require the models to be nested, the same approach can also be used to compare the fit of RCA and CCA partitions to one another.

I use SEM to fit three models to the musical tastes data: a homogeneity model where a single correlation matrix describes the whole population; a model with three correlation matrixes corresponding to RCA's partition of the population; and a model with four correlation matrixes corresponding to CCA's partition. The log likelihoods are -35376.85 (df=153) for the single-matrix model; -34054.2 (df=459) for the RCA partition model; and -33381.32 (df=612) for the CCA partition model. Comparing their model fit via AIC and BIC, I find that both the RCA (Δ AIC = -2033.3; Δ BIC = -401.8) and CCA (Δ AIC = -3073.07; Δ BIC = -625.81) partitions are strongly preferred to the single-matrix model. This is consistent with the existence of significant heterogeneity in the schemas that different subgroups use to organize their musical tastes. Both indicators also show that the CCA-based partition fits the data far better than the RCA-based partition (Δ AIC = -1039.76 and Δ BIC = -224.01). The multiple groups analysis thus suggests that CCA offers a better description of the heterogeneity present in this data than RCA.

DISCUSSION

In this paper, I introduced Correlational Class Analysis (CCA), a technique that builds on Goldberg's (2011) Relational Class Analysis (RCA) methodology. Both methods aim to partition a survey population into groups of respondents who arrange their tastes according to shared cultural schemas. RCA uses an eponymous relationality measure to quantify the extent that two respondents appear to use such a “shared schema”—a central concept that has lacked a clear definition. By formalizing and making explicit the intuitions about such schemas implicit in Goldberg’s work, I was able to substantially advance the clarity of this approach. Furthermore, my formalization showed that these kinds of shared schemas should manifest themselves in linear dependencies between pairs of response vectors—the same measurement task for which Pearson’s correlation has long been an established solution. When I applied Pearson’s correlation to the same example that Goldberg (2011) used to introduce relationality, I found that correlation yielded substantially more accurate results. I thus proposed CCA as a correlation-based alternative to RCA.

I then used simulations to generalize and broaden the comparison between the two methods. In the first 10,000 simulations, I generated random taste schemas and then used random noisy linear transformations to simulate respondents following these schemas. I measured the ability of both methods to correctly determine which sets of respondents were produced from the same schema, and which from different ones. I found that the accuracy of correlation-based CCA remained reliably higher than that of relationality-based RCA across the full range of simulation parameters. I also demonstrated that RCA relies on a strong distributional assumption. When the two methods produced different results, the odds that CCA was more accurate than RCA were 17:1 in runs that obeyed RCA’s assumption, and rose to 23:1 in the runs that violated it.

Because of the theoretical uncertainty surrounding schematic cognition, and because Goldberg meant RCA to capture non-linear dependencies between tastes, I then investigated the performance of both methods when the model of schematic similarity as linearity was violated. I used simulations to examine six broad ways that schematic transformations could deviate from the theorized model. In the first three sets of simulations, I examined how the algorithms would perform if some basic linear operation (scaling, inversion, or shift) did not take place. In the second three sets, I examined performance in the presence of polynomial schematic transformations, independent subschemas, and multi-way interactions between genre preferences. Even though these simulations violated the assumptions of my model, CCA remained consistently more accurate than RCA throughout the six sets of simulations, indicating that relationality was not better than correlation at detecting any of these alternate patterns of schematic similarity. Overall, I found no evidence of any use cases where RCA would be methodologically preferable over CCA.

I concluded with a re-analysis of the 1993 GSS musical tastes module previously analyzed by Goldberg (2011). While RCA had partitioned this population into three classes, CCA partitioned it into four. Two of the four classes resembled those found by RCA, while the other two did not. This confirmed that the two methods can yield substantively different conclusions in empirical settings. I then examined their partitions of the GSS population with a multiple groups analysis in SEM, which indicated that CCA’s results yielded far better model fit to the GSS data than did RCA’s.

In the above analyses, I found no evidence of any settings in which RCA would be methodologically justified over CCA. I am also not presently aware of any compelling theoretical reasons that could suggest the use of RCA over CCA. While the theoretical justifications behind RCA and CCA have drawn on different rhetorical languages, these linguistic differences are not reflected in the actual substance of the two methods. One such linguistic difference that deserves special emphasis concerns the “relational” nature of the methods. Though CCA does not use Goldberg's relationality measure, CCA remains as relational a technique as RCA in both senses that Goldberg (2011) uses that term. First, both methods "simultaneously examine the relationships between variables and individuals" (Goldberg 2011:1404) by constructing networks of schematic similarities between respondents, and then using these to partition them into classes. And second, both methods consider "the relationships between [respondents'] positions on a variety of issues that make up a certain social domain" (Goldberg 2011:1402) rather than examining these positions separately. Indeed, it appears appropriate that Galton (1888) originally referred to the measure he invented as "co-relation" rather than "correlation"—an alternate orthography which would render CCA as "Co-Relational Class Analysis," thus highlighting CCA's equally relational nature.

Limitations and Future Directions

To highlight further avenues for methodological improvement, it is useful to locate RCA and CCA as two instances of a more general theoretical and methodological framework of *schematic class analysis*. This framework begins with three theoretical propositions:

1) Latent *shared cultural schemas* specify relative evaluations of tastes vis-à-vis each other in a cultural domain.

2) Individuals' taste patterns are manifestations of such schemas, created from them via one of a finite set of *schematic transformations*.

3) The individuals who created their taste patterns from the same schema make up a *schematic class*.

Then, methodologically,

4) The extent to which two respondents *i* and *j* appear to follow the same latent schema can be measured via some *schematic similarity measure* S_{ij}, which yields a similarity matrix *S* when applied to all pairs of respondents.

5) Finally, it is possible to partition the matrix *S* into zones of greater and lesser similarity via some *partitioning method* $\Phi(S)$, thus yielding an estimate of schematic classes present in the data.

6) These can then be used to estimate the contents of the shared cultural schemas.

Building on the intuitions behind Goldberg's (2011) work, I proposed here that shared cultural schemas (1) can be formalized as vectors of real numbers, and schematic transformations (2) by linear functions. I therefore proposed to use Pearson's correlation as schematic similarity measure S_{ij} (4). Like Goldberg, I employed modularity maximization for the partitioning method (5), which yields a partition of the survey population into schematic classes (3). I then showed how multiple groups analysis can be used to quantify the fit of such a partition.

A theoretical argument for a different approach to detecting shared schemas would need to provide alternate answers to points (1) through (3). That is, an alternate model would need to either specify a different way that shared schemas can be captured (e.g., matrixes, functions, etc.); or a different way that a schema can manifest itself in individuals' responses (i.e., other than through inversion, scaling and shift); or offer an alternate conception of schematic class membership. Such theoretical challenges would then require different methodological answers to points (4)-(6).

To maximize analytical clarity, correlation can be used as a starting point for many such alterations. In the 129 years since Galton proposed an index of "co-relation", a substantial number of other measures of correlation, such as Spearman's ρ or point-biserial correlation, have been developed to capture this relationship under different special cases of data (Rodgers and Nicewander 1988). Many of these can be used as "drop-in replacements" for Pearson's correlation, thus making CCA easily adaptable to situations for which specialized alternatives to correlation already exist.

Moreover, the simple algebraic form of Pearson's *r* makes it possible to derive "custom" correlation coefficients for many other potential situations. For example, above, I examined three "sub-linear" scenarios where the basic linear schematic transformations of inversion, scaling, or shift were theorized not to occur. In Appendix C, I demonstrate how Pearson's r can be modified to yield three new indexes $r_{\sim}, r_{\times}$ and r_{+}, each of which makes the coefficient account for these theoretical changes.

Datasets. Other avenues for methodological development lie in adapting schematic class analysis to new types of data, such as the datasets from review websites like *Yelp!*, *Epinions.com* and *Amazon.com*, which individuals use to publicly share their ratings of businesses, services or products. Such data, which are private but regularly made available to researchers, greatly exceed traditional survey data in scale and richness of detail, and thus provide important new opportunities for students of culture. Since Pearson's correlation is easy to compute efficiently and fast modularity maximization algorithms are available, scaling CCA to a large *N* is straightforward.

Schematic classes. While this paper has formally defined the concept of shared cultural schemas, the concept of schematic classes remains a second black box to be opened. How far can an individual's responses deviate from a schema for him to still count as a member of a schematic class, and how similar do two schemas have to be to actually count as one and the same schema? The answers to these questions fall to the partitioning method $\Phi(S)$. Presently, CCA, like RCA, uses modularity maximization as its partitioning method. As the simulations in this paper demonstrate, it does its task with acceptable accuracy. But, while modularity maximization is among the most widely used network analytic techniques across many disciplines (e.g., Neal 2014; Shwed and Bearman 2010; Porter, Onnela, and Mucha 2009), its use may implicitly introduce undesired theoretical assumptions. A substantial literature notes that modularity maximization suffers from a "resolution limit" that biases it against detecting both very small and very large modules in many empirical settings, joining even very different modules together when they are too small in proportion to the whole network, and conversely breaking apart modules that are too large (Fortunato and Barthélemy 2007; Lancichinetti and Fortunato 2011; for a thorough overview, see Good, de Montjoye, and Clauset 2010).

By relying on modularity maximization, RCA and CCA may thus unintentionally introduce an assumption that each member of the survey population belongs to one of a moderate number of schematic classes, and may produce misleading results when this is not the case—e.g., in situations where many respondents follow their own idiosyncratic response patterns, or where the whole population belongs to only one schematic class. Future methodological work should develop a theoretically-informed partitioning method $\Phi(S)$ that avoids these biases.[28] Before this is done, RCA and CCA results should be examined to verify that they indeed contain a significant amount of heterogeneity, much as the results of other methods are tested for statistical significance or goodness of fit. In the GSS case study above, I show how this can be accomplished via a multiple groups analysis in SEM.

[28] For an example of a partitioning technique with a domain-specific fix to problems stemming from modularity's resolution limit, see Sohn and colleagues (2011).

CONCLUSION

In this paper, I demonstrated that Pearson's correlation can be used to accurately measure schematic similarity between survey responses. Across the broad range of possible theoretical scenarios I examined, correlation-based CCA proved reliably more accurate at partitioning survey populations by shared cultural schema than relationality-based RCA. The switch from relationality to correlation brings a number of further benefits. Relationality is substantially more computationally costly to calculate than correlation, and also requires bias correction and extensive bootstrapping for significance testing. Correlation obviates the need for these further steps. This leaves an algorithm that is clear, fast, and easy to implement (see Appendix A). It also clarifies and standardizes the method, thus placing it in closer conversation with other methodological work. Future improvements to CCA can draw insights from these existing literatures, thus helping further advance the methodological project that began with RCA.

Chapter 4.

Holding a Position:
Public Opinion as Cognition in a Disorganized Field

How firmly do people on the opposite sides of some political issue hold their positions? Our intuitions may suggest that the answer depends on the mainstream popularity of the positions in question. For example, picture someone with a position that most would consider extreme, such as opposition to abortion even in cases of risk to life of mother. Both existing academic literature and popular culture frequently imagine individuals with such extreme views as "true believers" (Hoffer 1951) who stick to their extreme positions all the more doggedly because they place them in a clearly delineated principled minority (Hogg 2000, 2004). This line of thinking is consistent with a view that places political issues into two categories: those pitting two more-or-less mainstream popular positions against one another, where both sides hold their positions with the same stability; and those pitting a popular position against an extreme one, where the extreme positions are held more firmly. I argue theoretically and demonstrate empirically that this implicit model is exactly wrong. First, equal stability on both sides of an issue is an exception rather than the rule; and second, there is a significant and robust *positive* relationship between the popularity of a position and the stability with which it is held.

The social contours of public opinion have long been an area of interest across the social sciences. Much research on political attitudes has investigated two related questions about this landscape: first, how individuals' demographic and cultural characteristics relate to their stances on political issues; and second, how these characteristics relate to the manner in which individuals hold and report their attitudes. So, for example, we know much about how individuals' income or religion affects their political views (e.g., Hout, Brooks, and Manza 1995; Hout and Fischer 2002; Weakliem 1993); and about how economic class, gender, age, and political knowledge influence their ability and willingness to report political attitudes (Alwin and Krosnick 1991; Delli Carpini and Keeter 1996; Laurison 2012, 2015; Mondak and Anderson 2004). I opened the present paper with a reference to a principled minorities because it points to a third question, which is rarely asked explicitly: how is the relative standing of competing issue positions within a public debate related to the way that these positions are held? While the former two questions examine how the field of public opinion is inscribed within other sociocultural fields, this third question can lead us to understand the field of public opinion as a sociocultural entity of its own.

I use the language of field theory to define this object of investigation. The greatest virtue of "fields" as an organizing concept for public opinion comes from their dual character as both socio-structural and cognitive entities (Bourdieu 1984; Fligstein and McAdam 2012; Martin 2003). Fields are abstract arenas of social action where actors hold one of a set of socially recognizable positions, and experience *field effects* corresponding to these positions (Martin 2003). The *public opinion fields* I examine here are formed by disagreement over political issues, with field positions corresponding to the different socially recognizable stances on these issues (for example, "environmentalist", "pro-choice", or "pro-gun"). Their field effects can include social pressures to become more like those on "their side" of the field, and feelings of animosity towards those on the other side (Martin 2000; Nicholson 2012; Smith and Hogg 2008). Most relevantly for the topic I describe here, field effects also include exposure to ideological information streams from political parties and the mass media, which play a key role in influencing the content of individuals' political beliefs and the manner in which these beliefs are manifested in behavior (Zaller 1992). This relationship between structural positions and the motivations and competences of their occupants is a central aspect of fields (Bourdieu 1984; Savage and Silva 2013).

As I discuss below, public opinion fields generally have substantially less structure and less powerful field effects than many conventional objects of field-theoretic investigation. They may be best characterized as weak or inchoate fields, with an existence somewhere between fully formed fields and what Fligstein and McAdam call "unorganized social space" (2012:5). Nonetheless, as I illustrate here, the concept of fields is useful for analyzing public opinion both in spite of and *because* of these divergences. Public opinion fields share enough properties with stronger fields to make field-theoretic language applicable. This language enables parsimonious descriptions relating the macro-scale political, cultural, and media "context" in which attitudes are formed and maintained to the micro-scale cognitions and behaviors of "having an opinion." But the fields metaphor also brings out the ways that public opinion differs from more settled fields: specifically, the deep lack of consensus between competitors over what issues the public debate is about, and the absence of strong "governance units" (Fligstein and McAdam 2012) that could enforce any agreement about the nature of disagreement underlying the field.

In this paper, I use the language of fields to bring together existing findings from across diverse traditions for studying attitudes and political culture, including those in sociology, political science, psychology, and cognitive science. This kind of systematization is useful because it can point to new hypotheses at the junctions of these fields of study, both by highlighting connections between different findings and by revealing tensions latent in present work. I then make use of the parsimonious structural features of the fields metaphor to formally develop a latent class model of position-holding within public opinion fields, which I then estimate with data from the 2008-2010-2012 General Social Survey (GSS) panel. Finally, by linking public opinion to the burgeoning study of social fields (e.g., Fligstein and McAdam 2011, 2012; Green 2013; Martin 2003), I also hope that this perspective may help reintegrate the study of public opinion—long a "missing concept" in sociology (Manza and Brooks 2012)—into closer dialog with contemporary sociological theory.

Field Structure of Public Opinion

Public opinion is a concept with multiple competing meanings, and thus requires disambiguation. By *public opinion*, I mean the distribution of attitudes about politically-relevant issues within a population—or, equivalently, the distribution of a population within an abstract social space formed by these attitude objects (Martin 2000), where an individual actor's field position corresponds to their stance on the issues defining the space (see, e.g., Figure 1). The public opinion field around a given issue can become *settled* or *institutionalized* when the opposing camps of actors reach a consensus about what exact issue is being debated, about the possible stances one could have on that issue, and about the importance of debating that issue. Definitionally, I treat each political issue as forming its own field, though empirically some of these fields may be so deeply interrelated as to function as one. Since, as Fligstein and McAdam (2012) noted, fields have a Russian stacking-doll character, I also will also talk of the overall "public opinion field" composed of these smaller fields.

[Figure 1 about here]

Whereas public opinion, under this definition, is a feature of how individuals are distributed across issue positions, public opinion *fields* also contain organizational actors who compete with one another over the positions, attentions, and allegiances of these individuals. It is these organizational actors that are responsible for many of the broad characteristics of the field. Such actors include political parties, media companies, and social movement organizations. The primary focus of my analysis is on individual people rather than on these competing organizations. However, the outsize role played by these organizational actors means that they must feature prominently in this model. The field model of public opinion relates features of their competition to the ways that individuals hold their opinions.

In order to mobilize supporters and expand coalitions, the parties and other organizational actors work to frame issue positions as consistent with their ideology, and then champion them to their audiences (Carmines and Stimson 1989; Sniderman 2000; Sniderman and Stiglitz 2012). To select information flows that correspond to their current position, individuals use their ideological or political identities—that is, their positions in the field—to determine which of these competing organizations are "on their side" (Boutyline and Vaisey 2017; Broockman and Butler 2015; Cohen 2003; Sniderman and Stiglitz 2012). Individuals then attend to these ideological information streams, and use them to develop and maintain their political attitudes. In addition to reinforcing the individuals' existing field positions, these political communications can also lead them to acquire positions in the public opinion fields formed around other issues. The amount of information that is successfully transmitted from the ideological producers to individual opinion holders can be thought of as the *field strength* at that position.

Through this process, these major organizations become the cultural producers behind many (or perhaps most) of the "logics" uniting disparate topics into seemingly coherent ideologies (see, e.g., Boutyline and Vaisey 2017; Converse 1964; Green, Palmquist, and Schickler 2002). They also provide the public with the vocabulary of terms and stereotyped arguments that individuals then deploy to describe and justify their own political views (Strauss 2012), and many of the background assumptions about how the social world operates that allow individuals to reason ideologically (Martin and Desmond 2010). In short, these ideological actors provide their audiences with the cultural skill of holding a political opinion—a skill individuals put to use when reporting their attitudes during a survey interview.

The challenge of stability

Scholars across different intellectual traditions have stressed that answering survey questions is an acquired cultural competency (e.g., Bourdieu 1984; Chang and Krosnick 2009; Zaller 1992). One aspect of it is the respondent's "sense that [he or she] is a legitimate producer of political opinions" (Laurison 2015:925)—that is, the confidence necessary to offer an answer to a survey question as opposed to a "Don't know" response (but see Luskin and Bullock 2011). My focus, however, is on the *skill* of producing attitude reports themselves. It is now well-accepted that individuals rarely possess coherent political ideologies which they could use as the bases for their responses (e.g., Converse 1964; Martin 2010). An individual's process of deciding between the answer choices offered by the interviewer instead appears resemble a search through the variously internalized cognitive contents which could potentially serve as the bases for such a choice (Krosnick 1991, 1999; Zaller and Feldman 1992). These "considerations" (Zaller 1992) include beliefs, identities, narratives, "source cues" (Goren, Federico, and Kittilson 2009), and non-verbal intuitions which individuals can effectively weigh against one another in deciding on the response. This process appears to often be guided by "satisficing," with individuals stopping their search for relevant considerations once an apparently "good enough" set has been found (Krosnick 1991).

The more unfamiliar the topic at hand, the more the process of finding relevant considerations is stochastic. Absent habitual search paths, the most likely considerations to be retrieved are those that were retrieved recently, or that were primed by recent exposure to related environmental stimuli. Since most individuals have at least some opposing considerations on most issues (Zaller and Feldman 1992), this process often results in the same individual choosing different responses at different times. This difficulty of achieving response stability across panel waves is why Zaller and Feldman (1992) argue that it should be a central topic of investigation—a position that echoes Bourdieu's (1987) call to focus on "don't know" responses.

As with most skills, practice with reasoning about a given political topic can increase the skill with which individuals retrieve relevant intuitions. Research on the effects of repeated exposure on verbal and social reasoning tasks shows that repetition may make the relevant cognitive schemas more readily retrievable, and may also improve the speed and accuracy of the retrieval process itself (e.g., Smith 1994; Smith, Branscombe, and Bormann 1988). Those who frequently encounter discourses about a topic may thus develop habits that let them reliably pick out the same intuitions when asked to do so "on the spot" in the context of a survey interview. Indeed, survey researchers have frequently documented this effect of repeated exposure in the form of panel conditioning, where a respondent's exposure to a survey question on one panel wave increases the reliability with which he or she answers questions on that topic in future waves (e.g., Kroh, Winter, and Schupp 2016).

As I argued above, field position ends up exposing individuals to different streams of ideological information. Thus, individuals in different field positions likely come to acquire cultural competencies for answering different kinds of attitude questions. This makes *the stability with which individuals who hold a given position can report attitudes consistent with this position* a key measure of field strength at a given point in the space of attitudes: positions frequently reiterated by political messaging should have far higher stability than those less frequently rehearsed. This link between field position and attitude stability captured by the phrase "to hold a position," which can mean both "to have an opinion" and "to stand without wavering."

Degree of Institutionalization

Recall that the institutionalization of a public opinion field involves the emergence of consensus about which issue is being debated, the positions that one could conceivably take on that issue, and the importance of debating that issue. In more concrete terms, this kind of institutionalization would mean that opposing positions on that issue are actively discussed by individuals and promoted by organizational actors, and that the opposing camps dedicate similar portions of their resources to messaging about this issue. Then, *ceteris paribus*, individuals on both sides of the field would be exposed to equal amounts of relevant cultural materials, and would thus hold their preferred positions with equal amounts of stability. Returning to the theme of relative stability with which I opened the paper, the question of whether opposing positions in the field are generally held with the same stability can thus be rendered as a question about the degree of institutionalization of the field of public opinion.

In these terms, the American public opinion field appears to have a very low degree of institutionalization. The root cause of this disorganization may be that, in championing political positions, political actors generally aim to win supporters rather than arguments. They thus attempt to sway the public debate towards topics on which they have majority support (Carmines and Stimson 1989). Though news organizations at times try to force political actors to focus on the same questions—perhaps most iconically so during the Presidential debates—there are no "governance units" (Fligstein and McAdam 2012) within the public opinion field that are strong enough to force the political actors to switch their general focus to a neutral set of issues.

Indeed, a key insight of research on "issue ownership" is that, in the minds of voters, many major issue categories are firmly divided up between the two major parties. For example, Democratic candidates are trusted to provide better solutions for social welfare issues such as public education and aid to the poor, whereas Republicans are trusted to make better decisions with regards defense and the size of government (Petrocik 1996). As Petrocik and colleagues demonstrated with a content analysis of presidential campaign rhetoric from 1952 to 2000, candidates from both parties then emphasize the issues their party is perceived to own. This adherence to topic was not perfect: Democrats, in particular, still spent a substantial amount of time discussing Republicans topics. However, these trespasses were then ignored by the media, with Democratic discussion of Republican issues receiving almost no news coverage (Petrocik, Benoit, and Hansen 2003). Thus, popular audiences were largely exposed to only the Republican framings of Republican-owned issues, and the Democratic framings of Democratic-owned ones.

Even when opposing sets of politicians and social movements do focus on the same broad issue category, they engage in framing contests (Snow et al. 1986) which involve, in part, shifting focus to those sub-issues that are likely to win greater levels of support from the general public. For example, in her classic study of abortion politics, Luker's (1984) invoked the division of possible reasons for abortion into "hard" ones (e.g., the woman's life is endangered or the pregnancy is a result of rape), where public support for the pro-choice position was high, and "soft" ones (e.g., when the woman does not want a child or cannot afford to raise another child), where it was lower. The pro-choice movement thus worked to focus the public debate specifically on these "hard" cases.

Luker further observed that, rather than asserting the pro-life position when addressing "hard" cases, the pro-life camp took a stand against the *pertinence* of these questions. For example, with regards to threat to life of mother, they made the case that modern medicine had virtually eliminated the cases where the life of the mother is pitted against that of the child. With regards to pregnancies from rape, they argued that "something biological happens to rape victims that precludes the possibility of pregnancy" (Luker 1984:235)—a widely discredited argument that nonetheless was still reiterated by pro-life candidates for the Senate as recently as 2010.

The competition between political and ideological actors thus touches not only on answers to political questions, but also on which questions should be asked. Political actors focus their communications on those positions where they have an advantage. Even when focused on the same broad issue category, they focus their messaging on specific issues where their stance is more popular, and avoid communicating about those where it is less popular. Since, as I argued above, such political communication is the major vehicle by which individuals acquire the cultural competency necessary for answering survey questions reliably, I predict that the following two patterns should hold across issues:

H_1: Adherents of opposite positions on any given issue generally exhibit different amounts of stability when reporting their position on the issue.

H_2: All things being equal, adherents of more popular positions exhibit more stability when reporting their positions on the issue than adherents of less popular positions.

METHOD

Assumptions and Notation

As I argued above, the relative stability of survey responses can be used as a tool for tracing the contours of the public opinion field around a given topic. To arrive at an estimator for the relative stability of field positions, I begin by laying out and formalizing a "position-holding model" of survey response. The basic idea behind the model is that the public opinion field around any issue contains a number of competing positions, and that these positions differ from each other in how stably their occupants report their positions on the survey. I thus assume that the set of meaningfully different response options to a given question corresponds to the set of possible positions. While the respondents' answers choices across *t* waves of a survey are observed, the true position of any given respondent is unobservable. This is thus a latent class model, with latent classes representing positions. However, this model differs from many more traditional applications of latent class analysis (LCA) in that the number and meaning of the classes is known by design rather than determined inductively. This feature allows for a simpler and more intuitive formal notation than the one usually used for LCA (e.g., Goodman 1974).

To make this derivation easier to follow, I will use a concrete example of a single attitude question that has two response options, "Yes" and "No"—e.g., "do you favor or oppose the death penalty?" I will also assume that each of *K* respondents responded to this question once in each of the three waves $t \in \{1,2,3\}$. I will refer to the set of all possible single-wave answers as $\mathcal{A} = \{y, n\}$, and the set of all possible three-wave answer sequences as

$$s = \{yyy, nyy, yny, yyn, ynn, nyn, nny, nnn\}.$$

I will refer to the total number of elements in s as $S = |\mathcal{A}|^{\max(t)}$, which in this example equals 8. I will use s_i to refer to the i^{th} element of s, and s_{it} to the t^{th} element of s_i, so that, e.g., if i = 2 and t = 1, then $s_i = nyy$ and $s_{it} = n$. I will designate a respondent's answer sequence across all three waves as ω, and her answer from wave t as ω_t.

I will first assume that each respondent has some unobserved true position τ in the field—in this case, either "yes" or "no." (I relax this assumption in the generalization below.) These positions are the "classes" of the latent class model. I will call the set of possible positions $\Omega = \{\mathrm{Y}, \mathrm{N}\}$. Since the process of question answering is stochastic, a respondent's answers do not have to correspond to their latent position, so it is possible that someone whose true position is "yes" would have answered "no, no, no" on the three waves of the survey. To provide a link between true and observed positions, I assume only that the respondent whose true attitude is Y is more likely to respond with a "yes" than with a "no", and vice versa. I formally state this "linking assumption" (1.1) below.

Since I will frequently reference the probabilities of individual answers (e.g., $P(\omega_t = y)$ or $P(\omega_t = s_i)$), three-wave answer sequences (e.g., $P(\omega = ynn)$), and true positions (e.g., $P(\tau = \mathrm{Y})$), I will generally omit the ω_t, ω, and τ (yielding, e.g., $P(y), P(s_i), P(ynn)$, and $P(Y)$.). So, for example, $P(yny|\mathrm{Y})$ stands for $P(\omega = yny|\tau = \mathrm{Y})$, which is the probability that a respondent whose true position is "Yes" answers "yes" on the first and third waves and "no" on the second wave. Using this notation, the linking assumption can be stated as:

$$P(x \mid X) > P(z \mid X) \; \forall \, z \neq x, \tag{1.1}$$

where *x* is the answer option that corresponds to true position *X*.

I further make the common latent class assumption that a respondent's answers on two different waves are independent of one another conditional on the respondent's true position,

$$P(\omega_t | \tau, \omega_q) = P(\omega_t | \tau) \; \forall \, q \neq t \tag{1.2}$$

Using this notation, the null hypothesis of equal position stability can be phrased as:

$$H_0: P(y \mid \mathrm{Y}) = P(n \mid \mathrm{N}) \tag{1.3}$$

Conversely, the alternate hypothesis of unequal position stability becomes:

$$H_1: P(y \mid \mathrm{Y}) \neq P(n \mid \mathrm{N}) \tag{1.4}$$

Finally, the hypothesis that more popular positions are held more stably can be rendered (with some simplification) as:

$$H_2: cor[P(X), P(x|X)] > 0 \tag{1.4}$$

Likelihood Function
My goal is to estimate the above parameters from a contingency table T, which cross-tabulates responses across the three waves. Each $T_i = |s_i|$ contains the count of respondents whose three-wave sequence of answers was s_i, e.g., $T_1 = |yyy|$ is the count of respondents answering "yes" on each wave. Assuming independence between respondents, T follows a multinomial distribution with S categories and $K = \Sigma_i T_i$ trials. Its likelihood function is thus given by

$$\mathcal{L}(T) = \binom{K}{|yyy|, |nyy|, |yny|, |yyn|, |ynn|, |nyn|, |nny|, |nnn|} * \\ P(yyy)^{|yyy|} * P(nyy)^{|nyy|} * P(yny)^{|yny|} * P(yyn)^{|yyn|} * \\ P(ynn)^{|ynn|} * P(nyn)^{|nyn|} * P(nny)^{|nny|} * P(nnn)^{|nnn|}$$

This can be rewritten more compactly as

$$\mathcal{L}(T) = \binom{K}{T_1, \dots, T_S} * \prod_{i=1}^{S} P(s_i)^{T_i}, \tag{1.5}$$

where $P(s_i)$ is the probability of a three-wave sequence in *s* (e.g.,yny). By law of total probability and assumption (1.2),

$$P(s_i) = P(\mathrm{Y}) * P(s_{i1}|\mathrm{Y}) * P(s_{i2}|\mathrm{Y}) * P(s_{i3}|\mathrm{Y}) + P(\mathrm{N}) * P(s_{i1}|\mathrm{N}) * P(s_{i2}|\mathrm{N}) * P(s_{i3}|\mathrm{N})$$

or, more compactly,

$$P(s_i) = P(\mathrm{Y}) \prod_{t=1}^{\max(t)} P(s_{it}|\,\mathrm{Y}) + P(\mathrm{N}) \prod_{t=1}^{\max(t)} P(s_{it}|\mathrm{N}) \tag{1.6}$$

Finally, combining (1.5) and (1.6) produces

$$\mathcal{L}(T) = \binom{K}{T_1, \dots, T_S} * \prod_{i=1}^{S} \left[P(\mathrm{Y}) \prod_{t=1}^{\max(t)} P(s_{it}|\,\mathrm{Y}) + P(\mathrm{N}) \prod_{t=1}^{\max(t)} P(s_{it}|\mathrm{N}) \right]^{T_i} \tag{1.7}$$

Estimation
The likelihood function $\mathcal{L}(T)$ can be used as a basis for testing hypotheses about conditional stabilities $P(y|\mathrm{Y})$ and $P(n|\mathrm{N})$. To examine their relationship, I expand and rearrange (1.7) as

$$\begin{aligned} \mathcal{L}(T) = \binom{\Sigma_i T_i}{T_1, \dots, T_8} &* [P(y|\mathrm{Y})^3 P(\mathrm{Y}) + P(y|\mathrm{N})^3 P(\mathrm{N})]^{T_1} * \\ &* (3 * [P(y|\mathrm{Y})^2 P(n|\mathrm{Y}) P(\mathrm{Y}) + P(y|\mathrm{N})^2 P(n|\mathrm{N}) P(\mathrm{N})])^{T_2+T_3+T_4} * \\ &* (3 * [P(y|\mathrm{Y}) P(n|\mathrm{Y})^2 P(\mathrm{Y}) + P(y|\mathrm{N}) P(n|\mathrm{N})^2 P(\mathrm{N})])^{T_5+T_6+T_7} * \\ &* [P(n|\mathrm{Y})^3 P(\mathrm{Y}) + P(n|\mathrm{N})^3 P(\mathrm{N})]^{T_8}. \end{aligned} \tag{2.1}$$

The values of T_i are observed. The remaining parameters have the constraint

$$P(\mathrm{Y}) + P(\mathrm{N}) = P(y|\mathrm{Y}) + P(n|\mathrm{Y}) = P(y|\mathrm{N}) + P(n|\mathrm{N}) = 1$$

This leaves three free parameters: $P(\mathrm{Y}), P(y|\mathrm{Y})$ and $P(n|\mathrm{N})$, which are the parameters needed to test the hypotheses stated above. This model can be estimated from an observed contingency table via approximate maximum likelihood. Since constraints like (1.1), (4.1), and (4.2) do not appear to be readily supported by any freely available LCA software, I instead implemented my own estimator for this class of models in *Python*. I then used a series of simulations to verify that it can indeed accurately estimate the parameters of interest.

Note that, because the ordering of an individual's responses make no difference to the model (e.g., $P(yyn) = P(yny) = P(nyy)$), the observation counts T_2 through T_7 occur in equation (2.1) only as the sums $T_2 + T_3 + T_4$ and $T_5 + T_6 + T_7$. As a consequence, the degrees of freedom *df* available in the data is smaller than the number of cells in *T*. It instead equals the number ways of drawing $max(t)$ unordered samples with replacement from a set of $|A|$ elements, or

$$table\ df = \binom{\max(t) + |\mathcal{A}| - 1}{\max(t)} \tag{2.2}$$

When max(t) = 3, this can be simplified to

$$table\ df = \frac{1}{6}|\mathcal{A}|(|\mathcal{A}| + 1)(|\mathcal{A}| + 2) \tag{2.2}$$

With two response options $|\mathcal{A}| = |\{y, n\}| = 2$, *table df* equals 4, which is exactly enough to estimate the three free parameters in equation (2.1).

Generalization

The above approach can be easily generalized to questions with more answer choices, positions, or panel waves. It is sufficient to substitute $\mathcal{A} = \{a, b, \dots, z\}$ and $\Omega = \{A, B, \dots, Z\}$ into equation (1.4), and alter *s* and max(*t*) as appropriate. This transforms the likelihood function (1.6) to

$$\mathcal{L}(T) = \binom{K}{T_1, \dots, T_S} \prod_{i=1}^{S} \left[\sum_{j=1}^{|\Omega|} \left(P(\tau_j) \prod_{t=1}^{\max(t)} P(s_{it}|\tau_j) \right) \right]^{T_i} \tag{3.1}$$

When the number of panel waves $\max(t) = 3$, this likelihood function has $|\Omega||\mathcal{A}| - 1$ free parameters, making the overall model degrees of freedom equal:

$$model\ df = |\Omega| * |\mathcal{A}| \tag{3.2}$$

This generalized model can also be estimated from an observed contingency table via maximum likelihood.

The hypotheses H_0, H_1, and H_2 concern positions on the opposite ends of an issue, and so now relate to the probabilities $P(A), P(Z), P(a|A)$, and $P(z|Z)$. The labels A and Z here refer to two opposite positions on an issue. For our purposes, it is not important which position gets called A and which Z. The model also contains parameters for other classes and probabilities, though these will be of less interest to us. Using this more general notation, the two-position model I laid out in the previous sections can be summarized as:

Model 1. $\mathcal{A} = \{a, z\}, \Omega = \{A, Z\}$

Respondents without a position

Another generalization concerns the assumption that each respondent has some true position in the space of attitudes. This assumption, which is present in various forms in many methods of analyzing survey data, has been frequently criticized as untenable (e.g., Bourdieu 1972; Krosnick 1991; Martin 2000; Perrin and McFarland 2011). Some respondents may have no opinion at all on the issue they are asked about, or may not wish to exercise the cognitive effort necessary to determine whether they have a position. Research on the cognitive bases of survey response identifies four forms of "strong satisficing" (Krosnick 1991; 1999) that could allow a respondent to answer a survey item without having recourse to any position on issue in question: (i) choosing a neutral response option, such as the middle of a rating scale (e.g., "neither" or "both") or the "no change" alternative; (ii) "coin-flipping", or randomly choosing between the options offered; (iii) "non-differentiating," or giving the same answer to all questions that have the same response options; (iii) and saying "don't know" or refusing to answer rather than choosing one of the options offered. Here, I develop two alternatives to the position-holding model 1 to account for the first two of these strategies, giving neutral responses and coin-flipping. (I discuss the other two forms of strong satisficing in later parts of the paper.)

Neutral responses. The first form of strong satisficing pertains to questions that offer a neutral or non-contentious response option, which can let a respondent avoid asserting an ideological position on an issue. In the GSS data I use here, the middle response category on attitude questions with an odd number of response options can serve this purpose. For example, items asking the respondent to report agreement or disagreement with an ideological statement may also provide the option "neither agree nor disagree." Questions asking a respondent to choose between a conservative and a liberal explanation a social problem may also offer the choice of "neither" or "both." Numerical thermometer-type scales can offer a middle position that serves the same purpose.

Another common question format posits two different alternatives to the status quo—for example, whether the number of immigrants admitted into the country should be increased or decreased, or whether the government is spending too much or too little on some social policy. In such cases, answer choices like "the number of immigrants should remain the same as is" or "the current spending levels are about right" can also function as ways for respondents to seemingly avoid taking a stand (Krosnick 1991). These kinds of survey items can be modelled by including a neutral respondent class $\emptyset$ and neutral response option ν:

Model 2a. $\mathcal{A} = \{a, \nu, z\}, \Omega = \{A, \emptyset, Z\}$

Unlike the classes A and Z, which correspond to actual positions in a public opinion debate, the defining characteristic of the neutral class $\emptyset$ is its unwillingness or inability to choose between the ideological choices offered. For this reason, I do not link class $\emptyset$ to response ν via the linking assumption (1.1). I thus let the response probabilities for members of the neutral class vary freely, requiring only that $P(a \mid \emptyset) + P(\nu \mid \emptyset) + P(z \mid \emptyset) = 1$.[1] Per (2.2), three-wave tables with such three-category questions have a $table\ df = 10$. Per (3.2), the Model 2a's *df* = 9.

<u>*Coin Flipping.*</u> Though model 2a allows respondents opt out of taking a position on an issue by picking the neutral position, such respondents must still make enough distinctions between the available answer categories to recognize the position as neutral or non-contentious. A population of respondents who are even more disengaged from the political field (or uncaring about the survey interview process) could go further and engage in "mental coin flipping" (Krosnick 1991:220, see also Krosnick 1999; Barge and Gehlbach 2012). That is, they may pick one of the answer choices at random, without any attention to its contents. If such "coin flipping" is frequent enough, the assumption that every respondent has a position may be a source of error.

To account for the possibility of such respondents, model 2a can be expanded to include both the neutral class $\emptyset$ and a second neutral class $\emptyset_{\mathrm{U}}$:

Model 2b: $\mathcal{A} = \{a, \nu, z\}, \Omega = \{A, \emptyset, \emptyset_{\mathrm{U}}, Z\}$

[1] Note that, due to this lack of link, the parameter $P(\nu|\emptyset)$ should not be interpreted as the stability of class $\emptyset$ (unlike, e.g., $P(a|A)$ and $P(z|Z)$, which represent the stability of classes *A* and *Z,* respectively).

Rather than having a tendency to prefer any one answer choice over the others, "coin-flipping" respondents $\emptyset_{\mathrm{U}}$ pick a response from the available ones at uniform random:

$$P(a \mid \emptyset_{\mathrm{U}}) = \cdots = P(z \mid \emptyset_{\mathrm{U}}) = \frac{1}{|\mathcal{A}|} \tag{4.1}$$

The overall likelihood function remains the same as in (2.4). Because restriction (4.1) reduces the degrees of freedom required to estimate this model by 2, its model df equals 10. It can thus be estimated from the same three-category items as model 2a (table *df* = 10). Note that it is not possible to include class $\emptyset_{\mathrm{U}}$ in analyses of two-category items due to insufficient table *df*.[2]

Alternately, in an effort to give a safer answer, respondents choosing a response at random could be proportionally more likely to pick popular responses than unpopular ones (the same behavior might happen if, lacking a position of their own, they decided to imitate the position held by one of their friends). Such a class of respondents $\emptyset_P$ would yield the following position-holding model:

Model 2c: $\mathcal{A} = \{a, v, z\}, \Omega = \{A, \emptyset, \emptyset_P, Z\}$

$$P(x \mid \emptyset_P) = P(x)\ \forall\ x \in \mathcal{A} \tag{4.2}$$

The marginal probability $P(x)$ can be estimated by plugging in its observed empirical frequency of x. This model again has 10 *df*. I will use BIC and likelihood ratio tests to select between models 2a, 2b, and 2c.

Regression analyses

Applying the above models to a panel dataset yields a new set of data describing each position on each survey question asked of the panel. This new dataset has positions rather than individuals as the units of analysis. For example, the question "Should abortion be legal in cases of risk to life of mother" yields two observations, one for the pro-choice position on this question, and one for the pro-life position. The resulting dataset would contain the estimates of the popularity of these two positions, P(A) and P(Z), and their stability, P(a|A) and P(z|Z). Because the stability estimates for items with two and three response categories are not directly comparable to one another[3], I will examine the two sets of items with separate analyses. For each set, I will first contrast the stability of opposite positions visually and via bivariate analyses. I will then examine the relationship between the stability of each position and its popularity with multivariate regression.

There are three important kinds of confounds these analyses will need to rule out: differences between respondents, differences between survey items, and population-level attitude trends. An extensive literature in the study of survey response has focused on the effects of question and respondent characteristics on the reliability of respondents' answers (e.g., Alwin

[2] While items without a third position (e.g., $\mathcal{A} = \{y, n\}$) do not provide sufficient degrees of freedom to estimate this model from three-wave data, a two-position model with coin flippers could be estimated from four-wave panel data.

[3] Perhaps the most immediate reason for this non-comparability comes from their differences in range. By equation (1.1), position stability $P(x \mid X) > P(z \mid X)\ \forall\ z \neq x$. For two-category questions, this makes the range of stability equal $1 \geq P(x \mid X) > 0.5$. For three categories, the stability range is instead $1 \geq P(x \mid X) > 0.3$. It is also reasonable to expect that the quantity of response categories would affect the frequency with which any one response category is picked (e.g., due to random guessing), which would affect estimates of both *P(X)* and *P(x | X)*.

and Krosnick 1991; Holbrook, Cho, and Johnson 2006). Since my hypotheses concern only the characteristics of positions, such effects of question and respondent characteristics would present themselves as confounds to be eliminated.

First, some questions place fewer cognitive demands on the respondent, and are thus easier to answer reliably than other questions. Question topic, clarity of phrasing, number of response categories, and many other question characteristics have been linked to reliability of survey response (e.g., Alwin and Krosnick 1991). The position of the question within the survey also makes a difference, with questions towards the end answered less reliably than those in the beginning. If more unpopular positions corresponded to questions that also happen to place the greatest cognitive demands on the respondent because of their phrasing or position in the survey, these item-level differences could create a confound. Because each question in the survey corresponds to two different positions in these analyses, I can use question-level fixed effects to eliminate such question-level confounds.[4]

Second, individuals occupying different attitude positions can be expected to differ across demographic and cultural characteristics in ways relevant to this investigation. For example, respondents who are higher-educated, more politically involved, have a better memory or verbal ability, and are less tired or more motivated can be expected to provide more stable responses to the survey questions than those who are not (e.g., Converse 1964; Kaminska, McCutcheon, and Billiet 2010; Krosnick, Narayan, and Smith 1996; Schuman and Presser 1996). Similarly, respondents also differ in their tendency to engage in various forms of strong satisficing such as offering the same answer to a series of questions sharing the same response choices. It is thus necessary to rule out the possibility that the more popular positions are more stable simply because they attract respondents who, in general, hold their positions more stably than other respondents.

One approach to ruling out this confound would involve estimating position-level expected values of various demographic control variables, which can be accomplished via repeated application of the Bayes rule to estimates produced by the position-holding model (I demonstrate this approach in Appendix A). Note, however, that demographic characteristics are theorized to influence a respondent's *general* ability to stably answer survey questions, and it is this general ability which in turn affects the respondent's stability on a particular issue position. Thus, rather than controlling for demographic characteristics, a more straightforward solution to this confound is to control for respondent's general ability to answer question stably.

To arrive at such a position-level control, I make use of the fact that each respondent in my data answered at least 66 questions on each of the panel waves.[5] For the occupants of position X on question q, I can thus estimate their average stability across the other 65+ questions they answered.[6] Using these individual-level estimates, I can then use the approach I describe in Appendix A to create a position-level measure of the average stability with which occupants of a given position $X_{(q)}$ held their positions on all the *other* questions they answered,

[4] This is possible because each question has two positions associated with it in the position-level dataset.

[5] The structure of these data do not make it possible to use individual-level fixed effects; this approach, however, approximates the logic of such fixed effects by focusing the analysis on the variance in stability around an individual respondent's mean.

[6] Recall that the position X of a respondent on any given question is a random variable $P(X \mid \omega)$, where ω is the respondent's three-wave response. Their stability on the question is the probability $P(x|X)$ that their response matched up with this random variable, and thus is a random variable itself. Both of these probabilities can be calculated by applying Bayes rule to the position-holding model estimates.

$$general\ stability\ of\ occupants\ of\ position\ X_{(q)} = E_{Y(r)}\left[P\left(y\middle|Y_{(r)}\right)\middle|X_{(q)}\right]$$

Here, $X_{(q)}$ is position X on question q, $Y_{(r)}$ is any position on question $r \neq q$, and y is the answer option corresponding to $Y_{(r)}$.

This position-level variable lets me control for the average differences in stability between the occupants of different positions, which eliminates the need for including any other controls for individual characteristics of occupants. Because the stability estimates for two-and three- category questions are not on the same scale, I estimate the individual-level stability separately for two- and three-category questions.

The third confound to my analysis can come from broad shifts in public opinion. For example, the first decade of the 21st century saw a general liberalization of attitudes regarding gay rights. As a result of this trend, many of the respondents who reported anti-gay marriage views earlier in the decade could be expected to change their responses to later waves of the survey. This is a different kind of instability than I theorized in this paper. Since the position-holding model does not itself take time trends into account, it cannot distinguish position instability that is due to unreliable reporting from position instability due to secular shifts in public opinion. To compensate for this potential confound, I include a control for attitude trends in the regression analyses.

DATA

I tested my hypotheses with data from the 2008-2010-2012 General Social Survey Panel. I examined all the survey questions that (a) concerned a subjective attitude on a political issue, broadly construed, and (b) were asked on all three waves of the study. There were 107 distinct questions fitting these criteria, covering abortion attitudes (7 questions), government spending (28 questions), confidence in major institutions such as the press or the Supreme Court (12 questions), crime and punishment (4 questions), freedom of various forms of unpopular speech (15 questions), gender and marriage (9 questions), race and racism (8 questions), police use of force (5 questions), sexuality (5 questions), limits of government (5 questions), the right to suicide (5 questions), and a single question each regarding prayer in public schools, immigration restrictions, and possibility of a third world war.[7]

For each item, Table 1 contains the question mnemonic, GSS variable name, and effective sample size. Since the position-holding model is estimated separately for each survey question, I dropped respondents with missing data only from the analyses that involved the actual questions they did not answer. Estimating the position-level control variables for the regression analyses similarly required only pairwise deletion. There were a total of N=1294 distinct respondents across all the analyses.

[Table 1 around here]

All of the items were either binary or ordinal, but some featured far more answer choices than others. For items with large numbers of answer choices, and especially long Likert scales, the difference between two proximate positions may be small. The difference between answer choices like “strongly support” and “support” is minor enough that a shift from one to the other may not indicate a meaningful instability in the respondent’s reported attitude. These minor movements may then obscure more meaningful instabilities. To prevent this from happening, I

[7] This count excludes one question about sex education, which I dropped from the analyses because its skewed response distribution caused problems with estimation. Across the three waves, 92.4% of all non-missing responses were in favor of sex education in public schools.

collapsed all answer choices with 5 or more categories to 3 categories if the number of categories was odd, and to 2 if it was even. Of the 107 items, 52 were thus coded as two-category, and 55 as three-category. Table 1 also contains mnemonics for the answer choices on the opposite ends of each issue—that is, the response options corresponding to the first ("A") and last ("Z") positions for each item (to reiterate, the choice of which response assigned which position label is arbitrary; labels only indicate that the responses are on opposite ends of the same issue).

Some respondents avoid taking a stand on various survey topics by either volunteering a "I don't know" response when one is not offered, or otherwise refusing to pick between the response categories provided. While I could have approached these refusals in a similar fashion to how I dealt with neutral and status quo responses (models 2a, 2b, and 2c above), I chose not to do so for two reasons. First, "don't know" responses and refusals to answer may be especially confounded with socially-patterned personality traits such as self-confidence, risk-taking, or willingness to state an opinion (e.g., Laurison 2015; Mondak 2001; Mondak and Anderson 2004). And second, refusing to choose an answer is relatively rare compared to choosing a neutral or status quo option when one is offered. For these reasons, I instead treated "don't know" and "refused to answer" responses as missing data.

RESULTS

Position-Holding Model Estimates

Model 1. I began by applying model 1 to the 52 two-category items in the GSS data. The results produced by the model are plotted in Figure 2. In the figure, each of these 52 items is represented with a point. The horizontal coordinate of each point is the estimated stability of the first position (A) for the item, and the vertical coordinate is the estimated stability of the last position (Z). The letter over each point indicates the general topic of the question. For example, the item asking whether the respondent thinks it is acceptable (A) or not (Z) for police to strike a murder suspect is depicted as a point with label "P" (police). The coordinates of this point correspond to the stability of the "yes" $\mathrm{P(a|A)}$ and "no" $\mathrm{P(z|Z)}$ positions on this question, respectively (0.57 and 0.95). Across these 104 positions, the average stability equaled 0.84. These results can also be found in table 1.

[Figure 2 about here]

I now turn to the difference in stability between opposite positions. Dotted lines connect each point $(\mathrm{P(a|A)}, \mathrm{P(z|Z)})$ in the figure to the point $(\mathrm{P(a|A)}, \mathrm{P(a|A)})$ that lies on the diagonal. The length of each dotted line thus equals $|\mathrm{P(z|Z)} - \mathrm{P(a|A)}|$, which is exactly the difference in how stable the two opposite positions A and Z are on that issue. For example, for the police violence question, $|\mathrm{P(z|Z)} - \mathrm{P(a|A)}| = 0.38$. Across all the items, the average length of the asymmetry $|\mathrm{P(z|Z)} - \mathrm{P(a|A)}| = 0.15$, which is significantly different from 0 ($t = 10.4, df = 51, p < 0.001$), or 18% of the average position's stability. As is readily apparent from the figure, this difference is substantively large. These results are thus in line with Hypothesis 1, which states that the two sides of an issue do not generally have the same stability.

To examine a subset of these items more closely, I plot the relative stabilities for the seven abortion items in Figure 3 below, together with the question mnemonics for each one. These items correspond to the different abortion cases addressed by Luker (1984). The first four items are the "soft" abortion cases: elective abortion, abortion for unmarried women, and abortion for women who do not want or cannot afford to have any more children (A-D). The remaining three are the "hard" cases of abortion when there is risk to life of mother, rape, or chance of birth defect (E-G). As I noted above, Luker argued that pro-life activists wanted the debate to focus on question regarding the hard cases, where the pro-life position had overwhelming public support, while pro-choice activists refused to engage these questions for similar reasons. The stability results are consistent with this avoidance: while the pro-choice position on the hard cases is exceptionally stable, the pro-life position on them is exceptionally unstable.

[Figure 3 about here]

Model 2. The above analyses do not account for the fact that some respondents may not have a position on the issues in question. To rule out the possibility that the results I report are due to such respondents, I turn to results that the position-holding models produced from the 55 three-category items in the data. I proposed three different models for such items. Model 2a adds a class of positionless respondents who are drawn to the neutral or status-quo response category in the middle of the scale. Model 2b extends model 2a by adding a further class of positionless respondents who act as "mental coin-flippers," picking one response option at uniform random. Model 2c instead extends model 2a by adding a class of positionless respondents who are proportionally more likely to pick popular responses over unpopular ones.

Since models 2b and 2c are nested within model 2a, they can be compared to it with likelihood ratio tests ($\chi^2\ df = 1$). Across the 55 three-category items, likelihood ratio tests for model 2b produced an average p-value of $p = 0.79$, with the lowest p-value equaling $p = 0.196$. Tests for model 2c produced an average p-value of $p = 0.77$, with the lowest equaling $p = 0.152$. This indicates that adding either class of randomly-answering respondents to the model never brought any substantial improvements to model fit. Accordingly, model 2b's and 2c's *BIC* scores were always greater than model 2a's *BIC* score. Thus, for all of the 55 three-category items examined, the model (2a) without either class of randomly-answering respondents appears uniformly preferable to the models with them (2b and 2c). I will therefore base my further analyses on the results of model 2a.

The stability estimates for the A and Z positions produced by model 2a from the 55 three-category items are plotted in Figure 4. Across these 110 positions, the average stability equaled 0.74. As in figure 3, the dotted lines correspond to the absolute differences between the stability of the opposite positions on each of these questions, $|P(a|A) - P(z|Z)|$. The mean value of this difference equals 0.162, which is 22% of the average position stability. This difference is substantively and statistically different from 0 ($t = 11.6$, $df = 54$, $p < 0.001$). This again offers support for Hypothesis 1.

[Figure 4 about here]

Regression estimates

I will now use regression analyses with results from position-holding models 1 and 2a to test my second hypothesis, which states that occupants of more popular positions should hold them more stably than occupants of less popular ones. Because the stability estimates from models 1 (two-category items) and 2a (three-category items) are on different scales, I again analyze two- and three-category items separately from one another.

Model 1. I again begin these analyses with estimates from model 1 (two-category items, no positionless respondents). The univariate regression of position stability on position popularity is reported as Model I in Table 2, and depicted in the top panel of Figure 5 (throughout the text, I use Arabic numerals to indicate position-holding models, and Roman numerals to indicate regression models). The regression coefficient for popularity is 0.357 (SE = 0.022, $p < 0.001$). Thus, when the popularity of a position increases by one percentage point, the probability that its occupants correctly report their position increases by 1/3rd of a percentage point. This relationship is both statistically and substantively significant, and thus provides support for Hypothesis 2.[8]

[Table 2 about here]

[Figure 5 about here]

As I discussed above, the zero-order correlation between the popularities and stabilities of positions in the data may be confounded by differences in the general stability of the positions' occupants, as well as by general trends in the popularity of different positions. To compensate for these confounds, I add controls for general stability of occupants and change in position popularity to the regression. These results are reported as Model II in table 2. The coefficient for general occupant stability equals 0.017 (SE = 0.006, $p < 0.001$), indicating that, when the occupants of a position were one standard deviation more stable in their stances on the other 2-category items, the average stability of that position increased by roughly 0.02 of a point. The coefficient for popularity change equals 0.018 (SE = 0.007, $p < 0.05$), indicating that, when a position experienced one standard deviation more change in popularity, its stability decreased by roughly 0.02 of a point. Both effects are thus significant and in the predicted direction, although neither is substantively large. With these controls included, the coefficient for the popularity of a position decreased to 0.270 (SE = 0.031, $p < 0.001$), but still remained substantively large and highly significant statistically.

Finally, different survey items differ from one another in how easy they are to respond to reliably. To compensate for item-level differences in stability, I include item-level fixed effects in the regression. These results are reported as Model III in table 2. With the fixed effects in place, the coefficient for popularity again decreased in magnitude but remained highly significant, equaling 0.215 (SE = 0.040, $p < 0.001$). This means that, within each survey item, a one-point increase in the popularity of a position was associated with a 0.22 point increase in stability. This result thus again provides evidence in favor of Hypothesis 2.

Model 2. The above results do not account for the possibility that some respondents may not occupy a meaningful position on an item. To account for such respondents, I now turn to regression analyses of the estimates produced by position-holding model 2a, which are based on items with three categories and feature a neutral class for positionless respondents. These results are depicted in the bottom panel of figure 5, and reported in Table 3. As with the previous set of analyses, Model I contains the univariate regression results of position stability on position popularity. The coefficient for popularity equaled 0.429 (SE = 0.041, $p < 0.001$), indicating that a one-point increase in the popularity of a position was associated with a significant 0.43 point increase in its stability.

[Table 3 about here]

In Model II, I added controls for the general stability of a position's occupants on all other three-category items ($\beta = 0.041, SE = 0.007, p < 0.001$) and the population-level popularity change of a position ($\beta = -0.028, SE = 0.007, p < 0.001$). The key coefficient for position popularity still remained substantively large and statistically significant ($\beta = 0.378, SE = 0.037, p < 0.001$) in the presence of these controls.

[8] I used simulations to rule out the possibility that the positive correlation between popularity and stability may be an artifact of the model. Simulations instead indicated that the model contains a slight bias in the *opposite* direction—i.e., more popular positions are likely to appear *less* stable than they truly are. These results thus do not appear to be produced by model bias.

Finally, Model III includes item-level fixed effects. In this model, the coefficient for popularity retained both its magnitude and its significance ($\beta = 0.386, SE = 0.40, p < 0.001$). Thus, when my analysis accounts for neutral respondents, position-level differences in occupant stability, trends in position popularity, and item-level fixed effects, I find that a one point increase in position popularity is associated with a 0.39 point increase in position stability. This again offers support for my second hypothesis.

CONCLUDING DISCUSSION

In this paper, I developed "public opinion fields" as an organizing model for tying together different strands of research regarding political attitudes. I defined public opinion fields as abstract social spaces around political issues, where competing views on an issue correspond to field positions. Field positions are occupied by both major political actors and individual opinion holders. These major political actors produce ideological information streams as part of their effort to grow coalitions and mobilize individual attitude holders. Individuals use their field positions to select among these streams, and this attention in turn stabilizes their field positions and provides them with positions in new fields. The information streams contain not just political facts, but also the cultural logics uniting together disparate attitudes, schemas describing the political world, frames providing ways of perceiving it, and stereotyped arguments that individuals can use to define, defend, and justify their positions. Different field positions thus provide their occupants with distinct sets of cultural competencies. In a survey setting, these competencies are what allow individuals to stably report their attitudes across multiple waves. Thus, I argued that the stability with which respondents report a given issue position can be used as a key indicator of the strength of the public opinion field at that location.

I observed that public opinion fields appear to have a very low degree of settlement or institutionalization—that is, rather than having an agreed-upon set of issues to debate, these actors appear to compete over which issues should be debated. I thus arrived at the prediction that the stability of positions on opposite ends of one issue should generally not be equal. Moreover, since political actors tend to be far more concerned with enrolling supporters and winning elections than they are with convincing individuals of the truth of particular arguments, they appear to focus their communications on the positions where they are winning. I thus also predicted that the stability of a position should be positively related to its popularity.

To test these suppositions, I developed a formal latent class model of position-holding, and estimated it with data from the 2008-2010-2012 GSS panel. I found that, on the average, opposite positions on issues differed significantly from each other in stability, with the average magnitude of this difference equaling roughly $1/5^{th}$ of the stability of an average position. Furthermore, across all positions, the popularity of a position had a significant positive association with its stability. A one percentage point increase in the popularity of a position was associated with at least a 0.22% increase in its stability. The significant positive association between popularity and stability remained in place in models that accounted for neutral respondents. It was also robust to controls for the average characteristics of position occupants, for population-level shifts in public opinion, and to the addition of item-level fixed effects.

The analyses thus offered consistent support for both of my hypotheses. Ideologically-opposed individuals appear to be best prepared to respond to different questions. Moreover, in general, popular positions are held more firmly, as would be expected if they were the focus of much political communication; unpopular positions are instead held with more uncertainty, as would happen if they were relatively underemphasized. This fits with an image of a public debate where the opposing camps within the mass public do not actually agree on which issues are being debated, with each camp focusing its attention on topics where it is already winning. It may be the case then that the informational dynamics within the public opinion field are not as good at sustaining a substantive debate between ideologically opposed individuals as they are at creating the *illusion* that such a debate exists—a competition that is less like combatants jousting, and more like them simultaneously tilting at different windmills.

From the perspective of the public opinion fields model, this finding turns attention to a second, broader kind of field settlement: the emergence of institutionalized alignments across different issue domains, via which topics as logically distant from one another as economic redistribution and traditional moral values come to exhibit persistent constraint or correlation (Converse 1964; see, e.g., see Figure 1). Recent work in the sociology of culture has explained such broad alignments via the existence of latent shared cultural schemas that specify which positions are consistent with which other ones, and which are opposed (Boutyline 2017; Goldberg 2011; for application of this thinking to political attitudes, see Baldassarri and Goldberg 2014; Daenekindt, Koster, and Waal 2017). For example, if one group wants less economic redistribution and endorses traditional values, and another wants more redistribution and opposes traditional values, they may hold no positions in common, but may nonetheless share a latent schema wherein support for economic distribution is considered the "opposite of" support for traditional moral values. However, the existence of such an agreement about the nature of the disagreement between competing ideological camps appears at odds with the general asymmetry of political position-holding I document in the present work. Future work on this topic should thus examine whether the cultural schemas held by competing ideological camps differ from one another to a greater extent than is assumed by the present work.

The measurement approach I followed in this paper let me estimate both the popularity and the stability of positions from any attitude question with three or more panel waves. This enabled me to demonstrate that the relationship between position popularity and stability is strong, and that it holds robustly across a broad range of issues. It did not, however, let me isolate the exact mechanism behind this relationship. To get at this mechanism, future work could complement this project by focusing on a small number of issues and tracking the volume of partisan communication around them. This would enable the researcher to test whether differences in partisan communication can indeed explain the association between issue popularity and the reliability of survey response.

A different way of building on the present project would be through an investigation of position holding in a panel dataset with four or more waves. For example, my analyses accounted for the existence of "strong satisficing," which can let respondents answer a survey question without any true latent position on an issue. The literature on cognitive bases of survey response, however, has also documented the existence of "weak satisficing." Respondents following this strategy *do* have opinions on issues and draw on them to answer survey questions, but use cognitive effort-saving heuristics that leave their responses unduly influenced by irrelevant features of question phrasing and survey setting. If repeated exposure to a political topic enables well-prepared respondents to reliably answer questions without much effort, then such respondents also appear to have little reason to use weak satisficing. Differences in the rates of weak satisficing may thus be the mechanism through which the less popular positions lose response stability relative to more popular ones. The position-holding model could be adapted to investigate this possibility with a four-wave panel dataset.[9] The extra modelling room provided

[9] Model 2a could be expanded to ask the following research question: if A and Z are two groups on opposite sides of some issue, can they be be partitioned by response strategy into subgroups of those who use weak satisficing

by such a dataset would also enable further elaborations to the position-holding model, making it a promising avenue for further work.

(A_w, Z_w) and those who do not (A', Z'), where the probability of stable response is the same by response strategy, $P(a|A_w) = P(z|Z_w)$ and $P(a|A') = P(z|Z')$? If so, then the difference in stability between A and Z can be explained by difference in the relative frequency of A_w within A and Z_w within Z.

FIGURES FOR CHAPTER 3

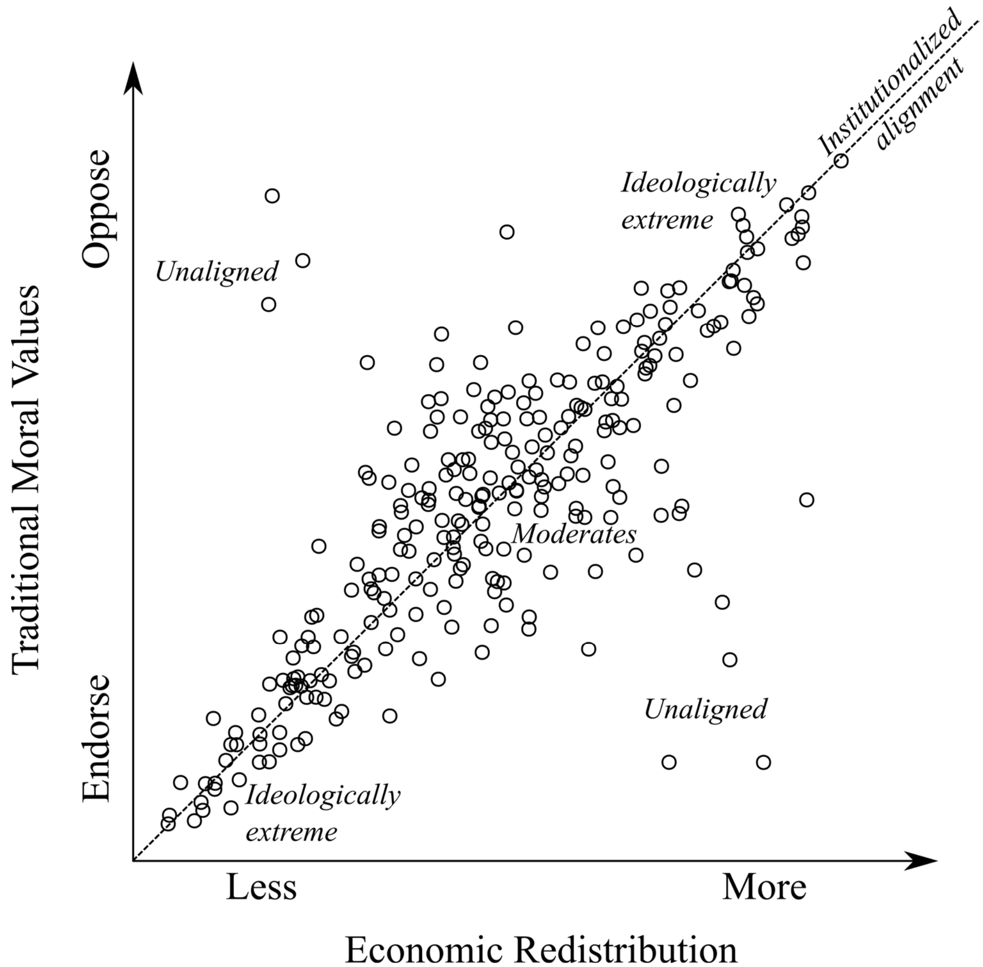

Figure 1. Example of a possible public opinion field around two attitudes.

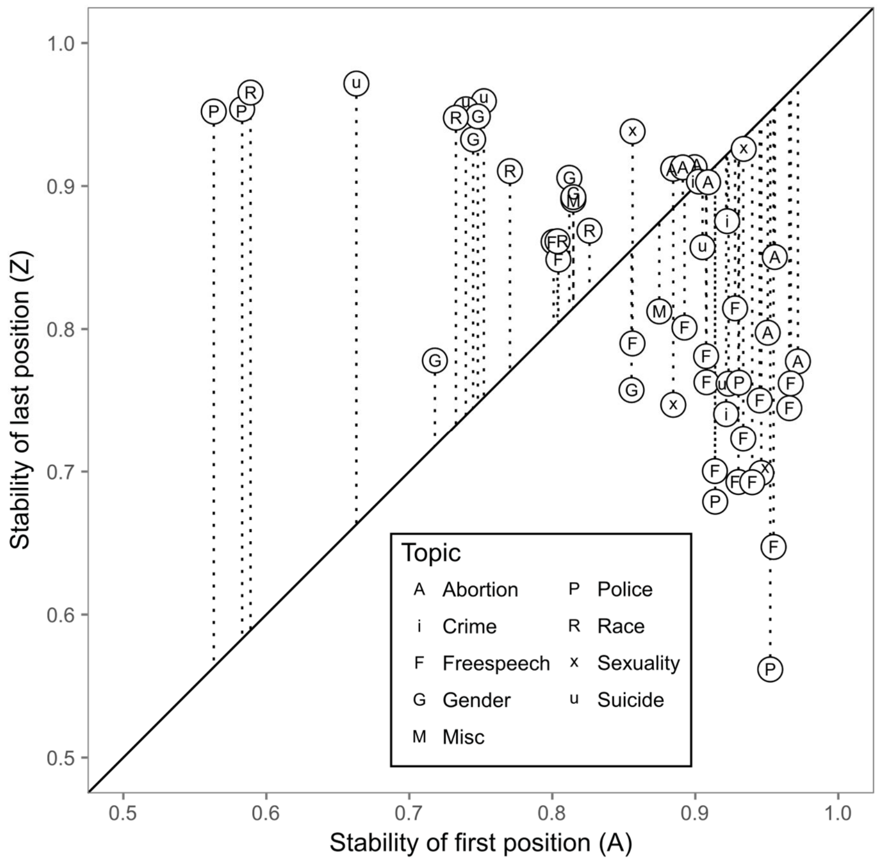

Figure 2. Relative stability of positions at opposite sides of 52 two-position issues from the GSS data. The X coordinate of each point indicates stability of first answer choice ("A"), while the Y coordinate indicates stability of the last answer choice ("Z"). The lengths of the dotted lines indicating the distance from each point to the X = Y diagonal thus correspond to the difference between the stability of the opposite positions on the issue, $|P(a|A) - P(z|Z)|$.

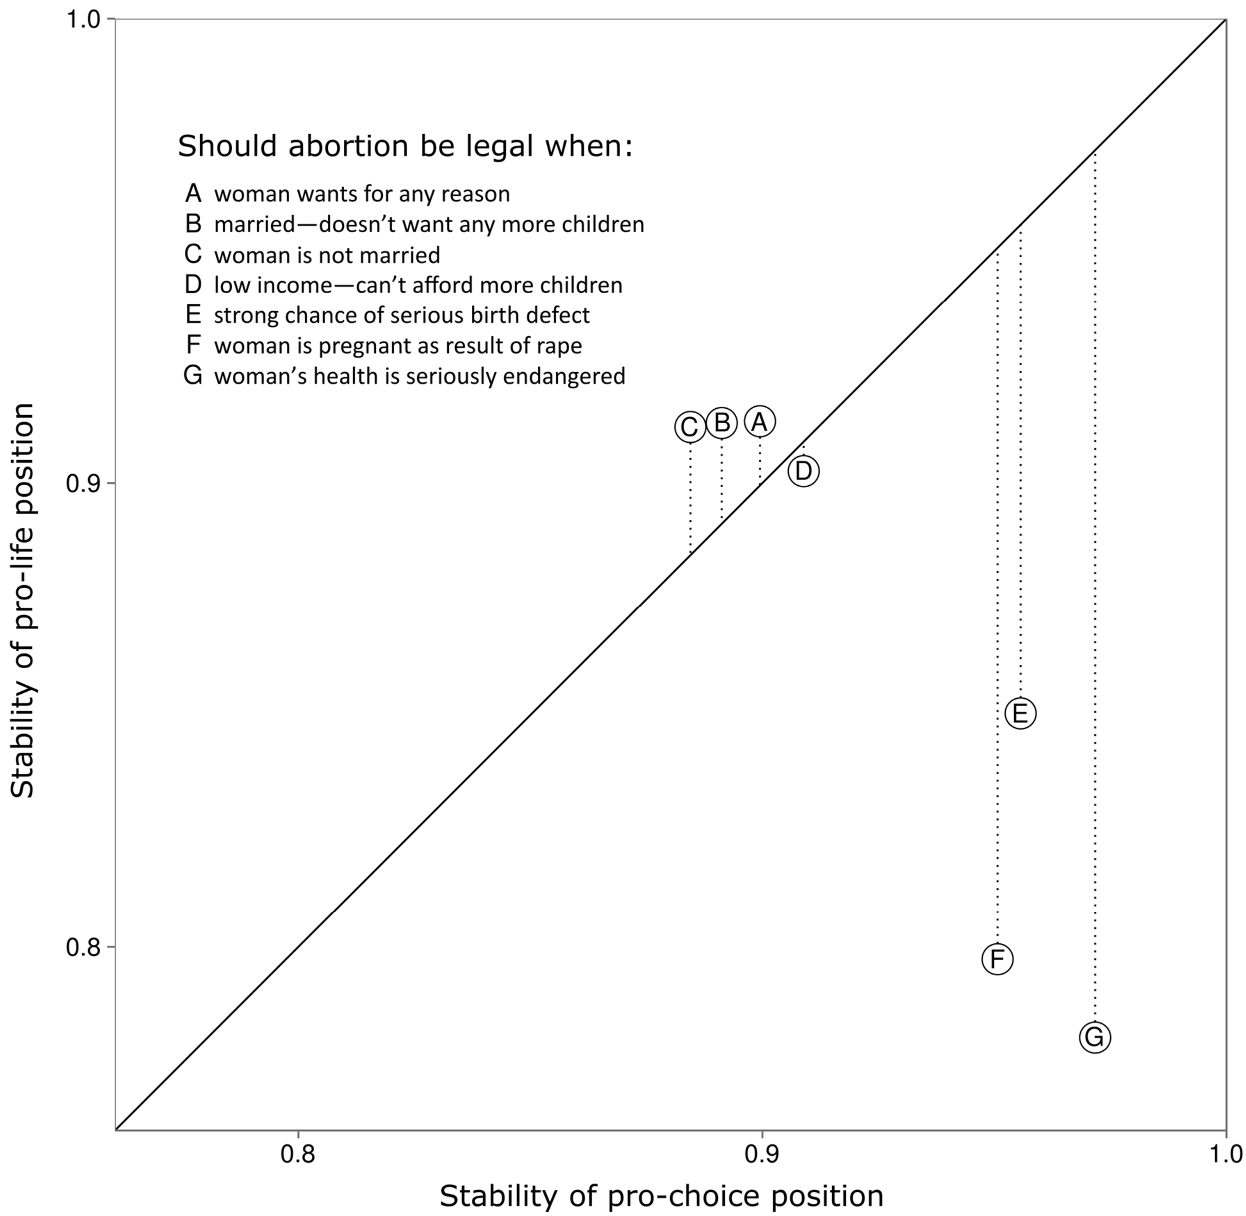

Figure 3. Relative stability of abortion attitudes. For items below the diagonal, the pro-choice position is more stable. For those above, the pro-life position is more stable.

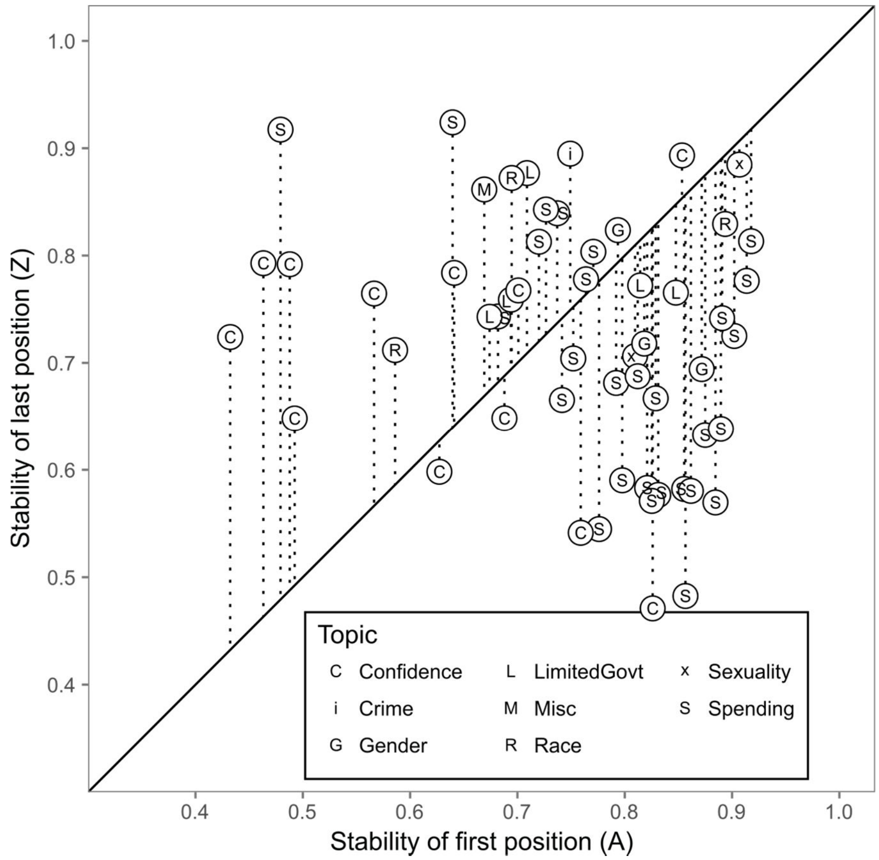

Figure 4. Relative stability of positions at opposite sides of 55 three-position issues from the GSS data. For each point, the X coordinate correspondents to stability of first answer choice, and the Y coordinate to the stability of the last answer choice. The lengths of the dotted lines indicating the distance from each point to the X = Y diagonal correspond to the difference between the stability of this pair of positions, $|P(a|A) - P(z|Z)|$.

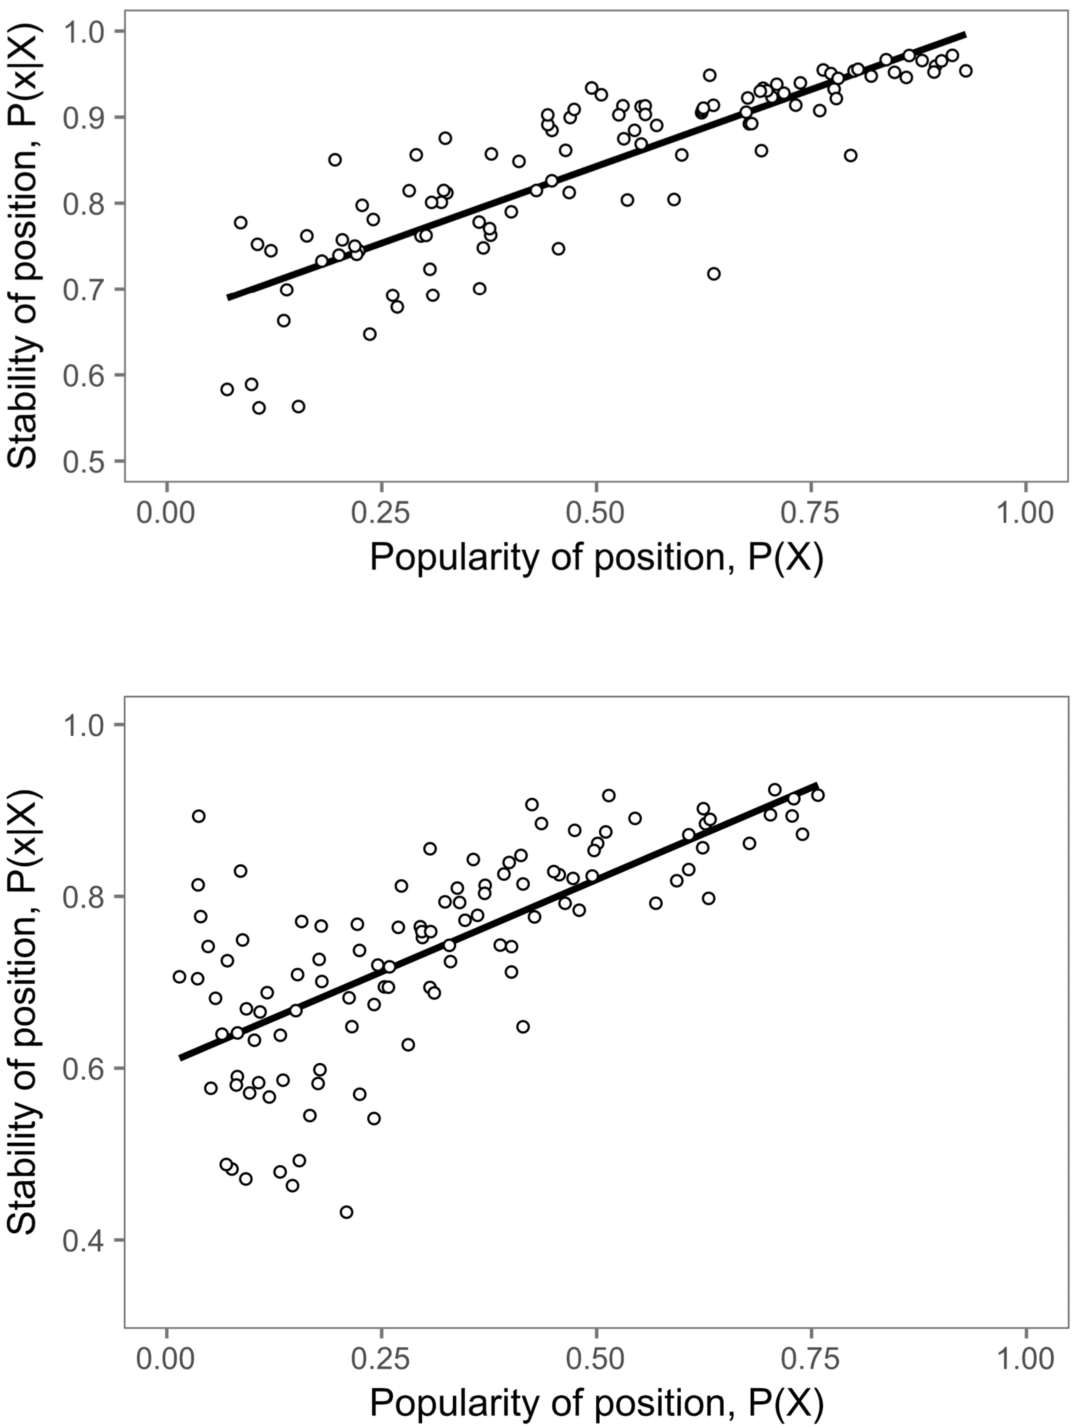

Figure 5. Relationship between the popularity (x-axis) and stability (y-axis) of leftmost and rightmost positions for 107 items in the GSS data, with lines of best fit. The top plot depicts the 104 positions for the 52 two-position items. The bottom plot shows the 110 positions for the 55 three-position items (the neutral position is not depicted).

TABLES FOR CHAPTER 4

Table 1. Results from applying position-holding model to the 2008 – 2010 – 2012 GSS panel (107 attitudes).

GSS		Positions						Stability		
variable		First (A)	(mid)	Last (Z)	N	P(A)	P(Z)	P(a|A)	P(z|Z)	Diff.
	Abortion should be legal when:									
absingle	woman is not married	yes		no	799	0.44	0.56	0.89	0.91	0.02
abany	woman wants for any reason	yes		no	789	0.45	0.55	0.88	0.91	0.03
abpoor	low income—can't afford more children	yes		no	781	0.47	0.53	0.90	0.91	0.01
abnomore	married—doesn't want any more children	yes		no	780	0.47	0.53	0.91	0.90	-0.01
abdefect	strong chance of serious birth defect	yes		no	763	0.77	0.23	0.95	0.80	-0.15***
abrape	woman is pregnant as result of rape	yes		no	781	0.80	0.20	0.96	0.85	-0.11***
abhlth	woman's health is seriously endangered	yes		no	767	0.91	0.09	0.97	0.78	-0.19***
	Opinion regarding government spending on:									
natrace	improving the conditions of blacks	too little	✓	too much	475	0.31	0.18	0.85	0.58	-0.27***
natfare	welfare	too little	✓	too much	552	0.22	0.40	0.74	0.84	0.10*
natdrug	dealing with drug addiction	too little	✓	too much	566	0.64	0.09	0.79	0.58	-0.21#
natroad	highways and bridges	too little	✓	too much	1220	0.47	0.11	0.82	0.57	-0.25***
natsci	supporting scientific research	too little	✓	too much	1124	0.40	0.11	0.74	0.66	-0.08
natcrime	halting the rising crime rate	too little	✓	too much	573	0.61	0.05	0.83	0.57	-0.26*
natcity	solving the problems of the big cities	too little	✓	too much	486	0.43	0.17	0.78	0.55	-0.23#
natsoc	social security	too little	✓	too much	1173	0.62	0.07	0.86	0.48	-0.37***
natmass	mass transportation	too little	✓	too much	1127	0.46	0.10	0.83	0.56	-0.26***
natheal	improving and protecting nat's health	too little	✓	too much	583	0.69	0.08	0.86	0.58	-0.28**
natpark	parks and recreation	too little	✓	too much	1231	0.30	0.04	0.75	0.68	-0.08
natspac	space exploration program	too little	✓	too much	554	0.18	0.36	0.73	0.84	0.11*
nateduc	improving the nation's education system	too little	✓	too much	601	0.74	0.04	0.91	0.76	-0.15#
natarms	the military, armaments and defense	too little	✓	too much	574	0.25	0.37	0.72	0.81	0.09#
natchld	assistance for childcare	too little	✓	too much	1091	0.57	0.06	0.79	0.68	-0.11#
natenvir	improving and protect the environment	too little	✓	too much	596	0.54	0.05	0.89	0.73	-0.16*
nataid	foreign aid	too little	✓	too much	546	0.14	0.51	0.47	0.92	0.45***
natcityy	assistance to big cities	too little	✓	too much	503	0.21	0.39	0.69	0.74	0.06
nathealy	health	too little	✓	too much	640	0.63	0.22	0.88	0.57	-0.31***
natcrimy	law enforcement	too little	✓	too much	648	0.51	0.10	0.87	0.63	-0.24***
nataidy	assistance to other countries	too little	✓	too much	625	0.06	0.70	0.65	0.93	0.27**

GSS variable		Positions: First (A)	(mid)	Last (Z)	N	P(A)	P(Z)	Stability: P(a|A)	P(z|Z)	Diff.
natfarey	assistance to the poor	too little	✓	too much	641	0.63	0.07	0.90	0.71	-0.19*
natarmsy	national defense	too little	✓	too much	620	0.27	0.36	0.76	0.78	0.01
natenviy	the environment	too little	✓	too much	636	0.64	0.13	0.89	0.64	-0.25***
nateducy	education	too little	✓	too much	659	0.76	0.04	0.92	0.79	-0.12
natracey	assistance to blacks	too little	✓	too much	504	0.27	0.31	0.81	0.69	-0.12*
natdrugy	drug rehabilitation	too little	✓	too much	589	0.45	0.15	0.83	0.66	-0.16*
natspacy	space exploration	too little	✓	too much	603	0.16	0.37	0.77	0.80	0.03
	Confidence in people running instutitions:									
confinan	banks & financial institutions	a lot	✓	hardly any	853	0.15	0.41	0.49	0.65	0.16
conbus	major companies	a lot	✓	hardly any	836	0.12	0.22	0.69	0.64	-0.05
conclerg	organized religion	a lot	✓	hardly any	835	0.18	0.22	0.70	0.77	0.06
coneduc	education	a lot	✓	hardly any	857	0.28	0.18	0.63	0.60	-0.03
confed	executive branch	a lot	✓	hardly any	838	0.21	0.33	0.43	0.72	0.29***
conlabor	organized labor	a lot	✓	hardly any	786	0.12	0.29	0.58	0.77	0.19*
conpress	press	a lot	✓	hardly any	845	0.08	0.48	0.65	0.78	0.14*
contv	television	a lot	✓	hardly any	845	0.15	0.34	0.46	0.79	0.33**
conjudge	Supreme Court	a lot	✓	hardly any	821	0.31	0.24	0.76	0.55	-0.21**
consci	scientific community	a lot	✓	hardly any	796	0.39	0.09	0.83	0.47	-0.36***
conlegis	Congress	a lot	✓	hardly any	842	0.07	0.46	0.49	0.79	0.30**
conarmy	military	a lot	✓	hardly any	839	0.50	0.04	0.85	0.89	0.04
	Crime and punishment									
courts	courts dealing with criminals should be:	less hrsh	✓	harsher	1078	0.09	0.70	0.75	0.90	0.15***
cappun	favor or oppose death penalty for murder	favor	✓	oppose	1154	0.68	0.32	0.92	0.88	-0.05**
gunlaw	favor or oppose gun permits	favor		oppose	841	0.78	0.22	0.92	0.74	-0.18***
grass	should marijuana be made legal	legal		not legal	711	0.44	0.56	0.90	0.90	0.00
	Free speech									
spkath	allow anti-religionist to speak	allow		not allow	851	0.76	0.24	0.95	0.65	-0.31***
colath	allow anti-religionist to teach	allow		not allow	818	0.68	0.32	0.89	0.80	-0.09**
spkrac	allow racist to speak	allow		not allow	840	0.62	0.38	0.91	0.76	-0.15***
colrac	allow racist to teach	allow		not allow	801	0.59	0.41	0.80	0.85	0.04
librac	allow racists book in library	allow		remove	820	0.64	0.36	0.91	0.70	-0.21***
spkcom	allow communist to speak	allow		not allow	832	0.72	0.28	0.93	0.81	-0.11***
colcom	should communist teacher be fired	fired		not fired	781	0.31	0.69	0.80	0.86	0.06

GSS variable		Positions: First (A)	(mid)	Last (Z)	N	P(A)	P(Z)	Stability: P(a|A)	P(z|Z)	Diff.
libcom	allow communists book in library	allow		remove	817	0.76	0.24	0.91	0.78	-0.13***
spkmil	allow militarist to speak	allow		not allow	828	0.69	0.31	0.93	0.72	-0.21***
colmil	allow militarist to teach	allow		not allow	803	0.60	0.40	0.86	0.79	-0.07#
spkhomo	allow homosexual to speak	allow		not allow	835	0.88	0.12	0.97	0.74	-0.22***
colhomo	allow homosexual to teach	allow		not allow	824	0.84	0.16	0.97	0.76	-0.20***
libath	allow anti-religious book in library	allow		remove	820	0.74	0.26	0.94	0.69	-0.25***
libmil	allow militarists book in library	allow		remove	823	0.69	0.31	0.93	0.69	-0.24***
libhomo	allow homosexual's book in library	allow		remove	831	0.78	0.22	0.95	0.75	-0.19***
	Gender									
fejobaff	support preferential hiring of women	support	✓	oppose	350	0.32	0.30	0.61	0.82	0.21**
discaffm	a man won't get a job or promotion	likely	✓	unlikely	372	0.09	0.22	0.47	0.39	-0.09
discaffw	a woman won't get a job or promotion	likely	✓	unlikely	439	0.17	0.02	0.72	0.71	0.00
fepol	women not suited for politics	agree		disagree	756	0.22	0.78	0.74	0.93	0.19***
fehire	should hire and promote women	agree	✓	disagree	450	0.61	0.26	0.87	0.69	-0.18**
meovrwrk	men hurt family when work too much	agree	✓	disagree	836	0.60	0.26	0.82	0.72	-0.10
fepresch	preschool kids suffer if mother works	agree	✓	disagree	823	0.07	0.12	0.44	0.65	0.20
fefam	better for man to work, woman tend home	agree	✓	disagree	823	0.02	0.19	0.80	0.70	-0.10
divlaw	make divorce easier, or more difficult?	easier		more diff.	751	0.32	0.49	0.79	0.82	0.03
	Race									
discaff	are whites hurt by aff. action?	yes	✓	no	794	0.13	0.40	0.60	0.71	0.12
affrmact	favor preference in hiring blacks	support	✓	oppose	756	0.09	0.58	0.72	0.79	0.07
racopen	allow owner to choose race of homebuyer	allow	✓	not allow	852	0.25	0.73	0.69	0.88	0.18***
wrkwayup	blacks overcome prejudice without favors	agree	✓	disagree	814	0.73	0.09	0.89	0.83	-0.06
racdif1	differences due to discrimination	yes		no	778	0.38	0.62	0.77	0.91	0.14***
racdif2	differences due to inborn disability	yes		no	814	0.10	0.90	0.58	0.97	0.38***
racdif3	differences due to lack of education	yes		no	804	0.45	0.55	0.83	0.87	0.04
racdif4	differences due to lack of will	yes		no	762	0.54	0.46	0.80	0.86	0.06
	Police striking citizen:									
polhitok	ever approve of police striking citizen	yes		no	358	0.70	0.30	0.93	0.76	-0.17***
polabuse	if citizen said vulgar or obscene things	yes		no	815	0.07	0.93	0.57	0.95	0.38***
polmurdr	if citizen questioned as murder suspect	yes		no	820	0.15	0.85	0.57	0.95	0.38***
polescap	if citizen attempting to escape custody	yes		no	793	0.73	0.27	0.91	0.68	-0.23***
polattak	if citizen attacking policeman with fists	yes		no	838	0.89	0.11	0.95	0.56	-0.39***
	Sexuality									

GSS variable		Positions First (A)	(mid)	Last (Z)	N	P(A)	P(Z)	Stability P(a|A)	P(z|Z)	Diff.
premarsx	is unmarried sex wrong	wrong	✓	not wrong	807	0.22	0.49	0.85	0.90	0.05
teensex	is unmarried sex wrong for teens (14-16)	wrong	✓	not wrong	828	0.70	0.02	0.90	0.69	-0.21
homosex	homosexual sex relations	wrong	✓	not wrong	772	0.46	0.43	0.93	0.89	-0.04#
pillok	teens shld alwys be able to get birth ctrl	agree	✓	disagree	367	0.14	0.21	0.79	0.60	-0.19
pornlaw	pornography should be illegal	illegal	✓	legal	859	0.34	0.02	0.81	0.66	-0.15
marhomo	homosexuals should have right to marry	agree	✓	disagree	844	0.42	0.44	0.91	0.88	-0.02
	Should government:									0.00
helpnot	do more, or leave to indiv's & business	yes	✓	no	802	0.24	0.33	0.67	0.74	0.07
eqwlth	should govt reduce income differences	yes	✓	no	850	0.41	0.35	0.81	0.77	-0.04
helppoor	should govt improve standard of living	yes	✓	no	813	0.31	0.30	0.69	0.76	0.06
helpsick	should govt help pay for medical care	yes	✓	no	813	0.41	0.18	0.85	0.76	-0.08
helpblk	should govt aid blacks	yes		no	799	0.15	0.47	0.71	0.88	0.16**
	Right to suicide:									0.00
suicide1	if incurable disease	yes		no	776	0.62	0.38	0.91	0.86	-0.05*
suicide2	if bankrupt	yes		no	814	0.11	0.89	0.75	0.96	0.21***
suicide3	if dishonored family	yes		no	815	0.14	0.86	0.66	0.97	0.31***
suicide4	if tired of living	yes		no	799	0.20	0.80	0.74	0.95	0.21***
letdie1	allow incurable patients to die	yes		no	351	0.70	0.30	0.92	0.76	-0.16***
	Misc.									
prayer	bible prayer in public schools	approve		disapprove	780	0.43	0.57	0.81	0.89	0.08**
letin1	number of immigrants to U.S. should be:	increased		reduced	811	0.09	0.50	0.68	0.86	0.18**
uswary	expect U.S. in a world war in 10 years	yes		no	373	0.53	0.47	0.87	0.81	-0.06

Note: Check mark (✓) between first and last positions indicates that the variable was coded as 3-position. All variables with 4 or more levels were collapsed to 4 levels if the number of levels was even, and to 3 if it was odd. The “Diff.” column equals P(z|Z) - P(a|A). Its significance is *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, # $p < 0.10$, as determined by likelihood ratio test versus H_0: P(a|A) = P(z|Z). The mean absolute difference across the whole table equals 0.14.

Table 2. OLS models predicting the stability of first and last positions for two-category GSS items. The estimates being analyzed come from position-holding model 1.

	Model I	Model II	Model III
Popularity	0.357***	0.270***	0.215***
	(0.022)	(0.031)	(0.040)
Occupant stability†		0.017**	0.016**
		(0.006)	(0.005)
Popularity change†		-0.018*	-0.038**
		(0.007)	(0.012)
(Intercept)	0.664***	0.708***	0.803***
	(0.012)	(0.016)	(0.039)
Item fixed effects?	—	—	✓
R^2	0.715	0.752	0.911
Adjusted R^2	0.712	0.745	0.813
N‡	104	104	104

Note: Standard errors in parentheses.
† Occupant stability and popularity change were standardized ($\sigma = 1$).
‡ The *N* for these regressions is the total number of first (A) and last (Z) item positions (i.e., double the number of survey items.)
*p < .05; **p < .01; ***p < .001

Table 3. OLS models predicting the stability of first and last positions for three-category GSS items. The estimates being analyzed come from position-holding model 2a.

	Model I	Model II	Model III
Popularity	0.429^{***}	0.378^{***}	0.386^{***}
	(0.041)	(0.037)	(0.040)
Occupant stability†		0.041^{***}	0.044^{***}
		(0.007)	(0.008)
Popularity change†		-0.028^{***}	-0.024^{*}
		(0.007)	(0.011)
(Intercept)	0.605^{***}	0.621^{***}	0.743^{***}
	(0.015)	(0.013)	(0.047)
Item fixed effects?	—	—	✓
R^2	0.505	0.656	0.864
Adjusted R^2	0.500	0.647	0.714
N‡	110	110	110

Note: Standard errors in parentheses.
† Occupant stability and popularity change were standardized ($\sigma = 1$).
‡ The *N* for these regressions is the total number of first (A) and last (Z) item positions (i.e., double the number of survey items.)
$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

CHAPTER 5:

CONCLUSION

The preceding three chapters approached shared culture from the perspective of a cognitive sociology. Chapter 2 examined political ideologies. I began by noting that many accounts of ideologies conceive of them as networks of interrelated opinions, in which some beliefs are central and others are derived from these more fundamental positions. I formally showed how such structural features can be used to construct direct measures of belief centrality in a network of correlations. With my coauthor Stephen Vaisey, we applied this method to the 2000 ANES, finding that ideological identity serves as the primary attitude-organizing heuristic. We searched for possible heterogeneity by contrasting 44 demographic groups, and then via two novel techniques for detecting latent complexity. Contrary to existing theories of how such belief systems are structured, we found that groups' belief systems vary in the *amount* of organization, but rarely in its *logic*. Across all groups, attitudes either follow the dominant liberal-conservative structure, or lack systemic organization. This paper has been conditionally accepted for publication by the *American Journal of Sociology*.

This investigation raises a broader question about how culture is organized: if culture provides the logics by which individuals arrange their attitudes, beliefs and tastes, how complex and varied do these logics tend to be? While scholars have often assumed that cognitively internalized culture is intricately organized, recent theoretical work in culture and cognition has argued that various simplifying properties of human memory make such complexity unlikely. I am currently developing a novel information-theoretic technique that will allow me to test this proposition empirically. Using this technique, I will examine survey data corresponding to a broad range of different cultural domains, including religious beliefs, collective memories, cultural tastes, and models of family and social relationships. Within each domain, I will assess the complexity of cultural logics necessary to account for its systematic character.

In Chapter 3 of the dissertation, I developed a method for detecting shared cultural schemas in survey data, which is a central methodological challenge in the sociology of culture. Such schemas define which attitude positions "go together" and which are opposed. For example, two individuals may share such a schema if one supports environmental regulation and welfare spending, and another opposes both policies: though the two hold opposing views, they implicitly agree on which stance on the environment goes with which stance on welfare. Such examples aside, existing theoretical reasoning has left the central concept of shared schemas loosely defined. I extended and clarified this reasoning to arrive at the missing definition. Surprisingly, my analyses demonstrated that such schemas are simply patterns of linear dependency between survey rows—the relationship usually measured by correlation. I then created a correlation-based alternative to the existing computationally intensive method, which both simplified it and greatly improved its accuracy, even in scenarios that greatly violate the assumption of linearity.

My clarified theory also yields testable predictions about schemas as vehicles for shared culture which I intend to explore in subsequent projects. Specifically, it highlights the potential role of the three linear schematic transformations—inversion, scaling, and shift—in establishing patterns of shared tastes and attitudes across the populations. The existence and prevalence of such transformations is an empirical question I will examine in future work.

Finally, in Chapter 4, I approached the distribution of political attitudes in the population as a field of competition. If public opinion is a debate between competing ideological camps, do the distributions of mass attitudes suggest that these competitors agree on what issues the debate is actually about? Both sociological practice theories and research on the cognitive demands of survey response suggest that reliably answering survey questions is an acquired cultural skill that requires substantial training to achieve. Much existing work suggests that such training most likely comes from the political, organizational, and social movement actors competing over public support for political issues. Therefore, it follows that respondents' differential skills at answering various questions should reveal which issues their ideological camp prepares them to debate. Drawing on my field-theoretic understanding of public opinion, I further observed that it is a highly unstable and un-institutionalized field, with little agreement between competing actors over the details of the competition. I thus predicted that two people with opposing views would rarely hold them with the same certainty.

I then formally developed a latent class model of conditional response reliability, and implemented software to estimate it from the 2008-2010-2012 panel data from the General Social Survey. I found that two people who have opposing views on a topic rarely hold them with the same amount of stability: that is, the terms of public debate appear overwhelmingly in dispute. Moreover, popular positions are generally held more stably than unpopular ones, which is consistent with a political field where competitors predominantly focused on issues where they were winning.

My next step building on this project involves the formal model of position holding I developed to test the hypotheses. This model can be easily extended to test hypotheses about the different kinds of "weak satisficing" respondents use to answer questions probabilistically. Such an investigation of weak satisficing should be able to reveal the immediate mechanisms behind the response instability I observed in Chapter 4. I further intend to extend the analyses to the other two 3-wave GSS panels (2006-2008-2010 and 2010-2012-2014), and thus to gain a better understanding of the across-time dynamics of position holding within the field of public opinion.

References

Achterberg, Peter, and Dick Houtman. 2009. "Ideologically Illogical? Why Do the Lower-Educated Dutch Display so Little Value Coherence?" *Social Forces* 87(3):1649–70.
Alwin, Duane F., and Jon A. Krosnick. 1991. "The Reliability of Survey Attitude Measurement The Influence of Question and Respondent Attributes." *Sociological Methods & Research* 20(1):139–81.
Alwin, Duane. 2007. *Margins of Error: A Study of Reliability in Survey Measurement.* 1st ed. Wiley-Interscience.
Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder. 2008. "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting." *American Political Science Review* 102(02):215–32.
Bai, Matt. 2005. "The Framing Wars." *The New York Times*, July 17.
Baldassarri, Delia, and Amir Goldberg. 2014. "Neither Ideologues nor Agnostics: Alternative Voters' Belief System in an Age of Partisan Politics." *American Journal of Sociology* 120(1):45–95.
Barge, Scott, and Hunter Gehlbach. 2012. "Using the Theory of Satisficing to Evaluate the Quality of Survey Data." *Research in Higher Education* 53(2):182–200.
Barker, David C., and James D. Tinnick. 2006. "Competing Visions of Parental Roles and Ideological Constraint." *American Political Science Review* 100(02):249–63.
Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24(2):117–50.
Beck, Glenn. 2008. "Commentary: Obama No, McCain Maybe." *CNN.com*, June 25.
Bickel, Peter, Friedrich. Götze, and Willem van Zwet. 2012. "Resampling Fewer Than N Observations: Gains, Losses, and Remedies for Losses." Pp. 267–97 In *Selected Works of Willem van Zwet, Selected Works in Probability and Statistics*, edited by Sara van de Geer and Marten Wegkamp. Springer New York.
Bollen, Kenneth A. 1989. *Structural Equations with Latent Variables.* New York: John Wiley & Sons.
Boltanski, Luc, and Laurent Thévenot. 1999. "The Sociology of Critical Capacity." *European Journal of Social Theory* 2(3):359–77.
Borhek, James, and Richard Curtis. 1975. *A Sociology of Belief.* New York: Wiley.
Bourdieu, Pierre. 1972. "Public Opinion Does Not Exist." In *Communication and Class Struggle.* New York: International General.
Bourdieu, Pierre. 1984 [1987]. *Distinction: A Social Critique of the Judgement of Taste.* Cambridge, Mass.: Harvard University Press.
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.
Boutyline, Andrei. 2017. "Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method." *Sociological Science* 4:353–93.
Breiger, Ronald L. 1974. "The Duality of Persons and Groups." *Social Forces* 53(2):181–90.

Brekhus, Wayne. 2007. "The Rutgers School A Zerubavelian Culturalist Cognitive Sociology." *European Journal of Social Theory* 10(3):448–64.
Broockman, David E., and Daniel M. Butler. 2015. "The Causal Effects of Elite Position-Taking on Voter Attitudes: Field Experiments with Elite Communication." *American Journal of Political Science* n/a-n/a.
Bryson, Bethany. 1996. "'Anything But Heavy Metal': Symbolic Exclusion and Musical Dislikes." *American Sociological Review* 61(5):884–99.
Burt, Ronald S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110(2):349–99.
Campbell, Angus, Philip Converse, Warren Miller, and Donald Stokes. 1960. *The American Voter*. Oxford, England: John Wiley.
Carmines, Edward G., and James A. Stimson. 1989. *Issue Evolution: Race and the Transformation of American Politics*. Princeton University Press.
Carmines, Edward, and Geoffrey Layman. 1997. "Issue Evolution in Postwar American Politics." In *Present Discontents: American Politics in the Very Late Twentieth Century*, edited by Byron E. Shafer. Chatham House Publishers.
Carmines, Edward, and Michael Wagner. 2006. "Political Issues and Party Alignments: Assessing the Issue Evolution Perspective." *Annual Review of Political Science* 9:67–81.
Cerulo, Karen A. 1993. "Symbols and the World System: National Anthems and Flags." *Sociological Forum* 8(2):243–71.
Cerulo, Karen A. 2010. "Mining the Intersections of Cognitive Sociology and Neuroscience." *Poetics* 38(2):115–32.
Chang, Linchiat, and Jon A. Krosnick. 2009. "National Surveys Via Rdd Telephone Interviewing Versus the Internet Comparing Sample Representativeness and Response Quality." *Public Opinion Quarterly* 73(4):641–78.
Chomsky, Noam. 1980. "A Review of BF Skinner's Verbal Behavior." *Readings in the Philosophy of Psychology* 1:48–63.
Cohen, Geoffrey L. 2003. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85(5):808–22.
Converse, Philip. 1964. "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, *International yearbook of political behavior research ; v. 5*, edited by David Apter. Free Press.
Converse, Philip. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3(1):331–53.
Daenekindt, Stijn, Willem de Koster, and Jeroen van der Waal. 2017. "How People Organise Cultural Attitudes: Cultural Belief Systems and the Populist Radical Right." *West European Politics* 0(0):1–21.
Danon, Leon, Albert Díaz-Guilera, Jordi Duch, and Alex Arenas. 2005. "Comparing Community Structure Identification." *Journal of Statistical Mechanics: Theory and Experiment* 2005(09):P09008.
DellaPosta, Daniel, Yongren Shi, and Michael Macy. 2015. "Why Do Liberals Drink Lattes?" *American Journal of Sociology* 120(5):1473–1511.
Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.

DeVega, Chauncey. 2016. "The GOP's Gross Adam Sandler Primary." *Salon*, March 9.
DiMaggio, Paul, and Amir Goldberg. 2010. "Searching for Homo Economicus: Variation in the Structure of Americans' Moral Evaluations of Markets." Atlanta, GA.
DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23:263–87.
Edgell, Penny. 2012. "A Cultural Sociology of Religion: New Directions." *Annual Review of Sociology* 38(1):247–65.
Emirbayer, Mustafa. 1997. "Manifesto for a Relational Sociology." *American Journal of Sociology* 103(2):281–317.
Fararo, Thomas J., and John Skvoretz. 1987. "Unification Research Programs: Integrating Two Structural Theories." *American Journal of Sociology* 92(5):1183–1209.
Feinberg, Matthew, and Robb Willer. 2013. "The Moral Roots of Environmental Attitudes." *Psychological Science* 24(1):56–62.
Fligstein, Neil, and Doug McAdam. 2011. "Toward a General Theory of Strategic Action Fields." *Sociological Theory* 29(1):1–26.
Fligstein, Neil, and Doug McAdam. 2012. *A Theory of Fields*. Oxford University Press.
Fortunato, Santo, and Marc Barthélemy. 2007. "Resolution Limit in Community Detection." *Proceedings of the National Academy of Sciences* 104(1):36–41.
Freeman, Linton C. 1978. "Centrality in Social Networks Conceptual Clarification." *Social Networks* 1(3):215–39.
Freeman, Linton C. 1980. "The Gatekeeper, Pair-Dependency and Structural Centrality." *Quality & Quantity* 14(4):585–92.
Freeman, Linton C. 2004. *The Development of Social Network Analysis: A Study in the Sociology of Science*. Vancouver, BC; North Charleston, S.C.: Empirical Press ; BookSurge.
Galton, Francis. 1888. "Co-Relations and Their Measurement, Chiefly from Anthropometric Data." *Proceedings of the Royal Society of London* 45:135–45.
Gardner, Howard E. 1987. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books.
Goldberg, Amir, and Jinjian Zhai. 2013. *RCA R Package*. Retrieved from CRAN repository.
Goldberg, Amir, and Sarah K. Stein. 2016. *RCA: Relational Class Analysis*. Retrieved on June 15, 2016. (https://cran.r-project.org/web/packages/RCA/index.html)
Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116(5):1397–1436.
Good, Benjamin H., Yves-Alexandre de Montjoye, and Aaron Clauset. 2010. "Performance of Modularity Maximization in Practical Contexts." *Physical Review E* 81(4):046106.
Goodman, Leo A. 1974. "The Analysis of Systems of Qualitative Variables When Some of the Variables Are Unobservable. Part I-A Modified Latent Structure Approach." *American Journal of Sociology* 79(5):1179–1259.
Gordon, Stacy, and Gary Segura. 1997. "Cross–National Variation in the Political Sophistication of Individuals: Capability or Choice?" *The Journal of Politics* 59(01):126–47.

Goren, Paul, Christopher Federico, and Miki Kittilson. 2009. “Source Cues, Partisan Identities, and Political Value Expression.” *American Journal of Political Science* 53(4):805–20.
Goren, Paul. 2005. “Party Identification and Core Political Values.” *American Journal of Political Science* 49(4):881–96.
Green, Adam Isaiah. 2013. *Sexual Fields: Toward a Sociology of Collective Sexual Life.* University of Chicago Press.
Green, Donald, Bradley Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds: Political Parties and the Social Identity of Voters.* 1st ed. Yale University Press.
Gregorich, Steven E. 2006. “Do Self-Report Instruments Allow Meaningful Comparisons Across Diverse Population Groups? Testing Measurement Invariance Using the Confirmatory Factor Analysis Framework.” *Medical Care* 44(11 Suppl 3):S78–94.
Gross, Neil, Thomas Medvetz, and Rupert Russell. 2011. “The Contemporary American Conservative Movement.” *Annual Review of Sociology* 37(1):325–54.
Henderson, Lol, and Lee Stacey. 2014. *Encyclopedia of Music in the 20th Century.* Routledge.
Hitlin, Steven, and Stephen Vaisey. 2013. “The New Sociology of Morality.” *Annual Review of Sociology* 39(1).
Hoffer, Eric. 1951. *The True Believer: Thoughts on the Nature of Mass Movements.* HarperCollins.
Hogg, Michael A. 2000. “Subjective Uncertainty Reduction through Self-Categorization: A Motivational Theory of Social Identity Processes.” *European Review of Social Psychology* 11(1):223–55.
Hogg, Michael A. 2004. “Uncertainty and Extremism: Identification with High Entitativity Groups under Conditions of Uncertainty.” In *The Psychology of Group Perception: Perceived Variability, Entitativity, and Essentialism.* Psychology Press.
Holbrook, Allyson, Young IK Cho, and Timothy Johnson. 2006. “The Impact of Question and Respondent Characteristics on Comprehension and Mapping Difficulties.” *Public Opinion Quarterly* 70(4):565–95.
Hout, Michael, and Claude S. Fischer. 2002. “Why More Americans Have No Religious Preference: Politics and Generations.” *American Sociological Review* 67(2):165–90.
Hout, Michael, Clem Brooks, and Jeff Manza. 1995. “The Democratic Class Struggle in the United States, 1948-1992.” *American Sociological Review* 60(6):805–28.
Huckfeldt, Robert, Jeffrey Levine, William Morgan, and John Sprague. 1999. “Accessibility and the Political Utility of Partisan and Ideological Orientations.” *American Journal of Political Science* 43(3):888–911.
Hummel, John. 2013. “Object Recognition.” In *The Oxford Handbook of Cognitive Psychology.* Oxford University Press, USA.
Hurwitz, Jon, and Mark Peffley. 1997. “Public Perceptions of Race and Crime: The Role of Racial Stereotypes.” *American Journal of Political Science* 41(2):375–401.
Jacobs, David, and Jason T. Carmichael. 2002. “The Political Sociology of the Death Penalty: A Pooled Time-Series Analysis.” *American Sociological Review* 67(1):109–31.
Johnston, Richard. 2006. “Party Identification: Unmoved Mover or Sum of Preferences?” *Annual Review of Political Science* 9(1):329–51.

Jost, John, Christopher Federico, and Jaime Napier. 2009. "Political Ideology: Its Structure, Functions, and Elective Affinities." *Annual Review of Psychology* 60(1):307–37.

Jost, John, Jack Glaser, Arie Kruglanski, and Frank Sulloway. 2003. "Political Conservatism as Motivated Social Cognition." *Psychological Bulletin* 129(3):339–75.

Kaminska, Olena, Allan L. McCutcheon, and Jaak Billiet. 2010. "Satisficing Among Reluctant Respondents in a Cross-National Context." *Public Opinion Quarterly* 74(5):956–84.

Kinder, Donald R. 1998. "Communication and Opinion." *Annual Review of Political Science* 1(1):167–97.

Kool, Wouter, Joseph T. McGuire, Zev B. Rosen, and Matthew M. Botvinick. 2010. "Decision Making and the Avoidance of Cognitive Demand." *Journal of Experimental Psychology: General* 139(4):665–82.

Kroh, Martin, Florin Winter, and Jürgen Schupp. 2016. "Using Person-Fit Measures to Assess the Impact of Panel Conditioning on Reliability." *Public Opinion Quarterly* 80(4):914–42.

Krosnick, Jon A. 1991. "Response Strategies for Coping with the Cognitive Demands of Attitude Measures in Surveys." *Applied Cognitive Psychology* 5(3):213–236.

Krosnick, Jon A. 1999. "Survey Research." *Annual Review of Psychology* 50(1):537–67.

Krosnick, Jon A., Sowmya Narayan, and Wendy R. Smith. 1996. "Satisficing in Surveys: Initial Evidence." *New Directions for Evaluation* 1996(70):29–44.

Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*. University Of Chicago Press.

Lakoff, George, and Rockridge Institute. 2006. *Thinking Points: Communicating Our American Values and Vision : A Progressive's Handbook*. New York: Farrar, Straus and Giroux.

Lakoff, George. 1990. *Women, Fire, and Dangerous Things*. University Of Chicago Press.

Lakoff, George. 2002. *Moral Politics: How Liberals and Conservatives Think*. University of Chicago Press.

Lakoff, George. 2014. *The ALL NEW Don't Think of an Elephant!: Know Your Values and Frame the Debate*. 10th Anniversary Edition edition. Chelsea Green Publishing.

Lamont, Michele. 2000. "Meaning-Making in Cultural Sociology: Broadening Our Agenda." *Contemporary Sociology* 29(4):602–7.

Lancichinetti, Andrea, and Santo Fortunato. 2011. "Limits of Modularity Maximization in Community Detection." *Physical Review E* 84(6):066122.

Lancichinetti, Andrea, Santo Fortunato, and János Kertész. 2009. "Detecting the Overlapping and Hierarchical Community Structure in Complex Networks." *New Journal of Physics* 11(3):033015.

Langfelder, Peter, and Steve Horvath. 2007. "Eigengene Networks for Studying the Relationships between Co-Expression Modules." *BMC Systems Biology* 1(1):54.

Lau, Richard, and David Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45(4):951–71.

Laurison, Daniel. 2012. "Political Competence in the United States." *Praktiske Grunde* 2012(2):43–56.

Laurison, Daniel. 2015. “The Willingness to State an Opinion: Inequality, Don’t Know Responses, and Political Participation.” *Sociological Forum* 30(4):925–48.
Layman, Geoffrey, John McTague, Shanna Pearson-Merkowitz, and Michael Spivey. 2007. “Which Values Divide?” New Orleans, LA.
Lee, Taeku. 2002. “The Sovereign Status of Survey Data.” In *Navigating Public Opinion : Polls, Policy, and the Future of American Democracy: Polls, Policy, and the Future of American Democracy*. Oxford University Press.
Leege, David C., Kenneth D. Wald, Brian S. Krueger, and Paul D. Mueller. 2009. *The Politics of Cultural Differences: Social Change and Voter Mobilization Strategies in the Post-New Deal Period*. Princeton University Press.
Lévi-Strauss, Claude. 1963. *Structural Anthropology*. Basic Books.
Lizardo, Omar. 2014. “Beyond the Comtean Schema: The Sociology of Culture and Cognition Versus Cognitive Social Science.” *Sociological Forum* 29(4):983–89.
Lizardo, Omar. 2014. “Omnivorousness as the Bridging of Cultural Holes: A Measurement Strategy.” *Theory and Society* 43(3-4):395–419.
Lizardo, Omar. 2017. “Improving Cultural Analysis: Considering Personal Culture in Its Declarative and Nondeclarative Modes.” *American Sociological Review* 82(1):88–115.
Luker, Kristin. 1984. *Abortion and the Politics of Motherhood*. Duke University Press.
Luskin, Robert, and John Bullock. 2011. “‘Don’t Know’ Means ‘Don’t Know’: DK Responses and the Public’s Level of Political Knowledge.” *The Journal of Politics* 73(02):547–57.
MacCallum, Robert, and James Austin. 2000. “Applications of Structural Equation Modeling in Psychological Research.” *Annual Review of Psychology* 51(1):201–26.
Manning, Christopher D., Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. New York: Cambridge University Press.
Manza, Jeff, and Clem Brooks. 2012. “How Sociology Lost Public Opinion A Genealogy of a Missing Concept in the Study of the Political.” *Sociological Theory* 30(2):89–113.
Martin, John Levi, and Matthew Desmond. 2010. “Political Position and Social Knowledge.” *Sociological Forum* 25(1):1–26.
Martin, John Levi. 1999. “Entropic Measures of Belief System Constraint.” *Social Science Research* 28(1):111–34.
Martin, John Levi. 2000. “The Relation of Aggregate Statistics on Beliefs to Culture and Cognition.” *Poetics* 28(1):5–20.
Martin, John Levi. 2002. “Power, Authority, and the Constraint of Belief Systems.” *American Journal of Sociology* 107(4):861–904.
Martin, John Levi. 2003. “What Is Field Theory?” *American Journal of Sociology* 1–49.
Martin, John Levi. 2010. “Life’s a Beach but You’re an Ant, and Other Unwelcome News for the Sociology of Culture.” *Poetics* 38(2):229–44.
McAdams, Dan et al. 2008. “Family Metaphors and Moral Intuitions: How Conservatives and Liberals Narrate Their Lives.” *Journal of Personality and Social Psychology* 95(4):978–90.
McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27(1):415–44.

McPherson, Miller. 2004. “A Blau Space Primer: Prolegomenon to an Ecology of Affiliation.” *Industrial and Corporate Change* 13(1):263–80.
Miranda, Shaila, Jama Summers, and Inchan Kim. 2012. “Visions of Social Media: Surfacing Schemas from Firms’ Informational Engagements.” *ICIS 2012 Proceedings.*
Mohr, John W., and Craig Rawlings. 2012. “Four Ways to Measure Culture: Social Science, Hermeneutics and the Cultural Turn.” In *The Oxford Handbook of Cultural Sociology*. Oxford University Press, USA.
Mohr, John. 1998. “Measuring Meaning Structures.” *Annual Review of Sociology* 24:345–70.
Mondak, Jeffery J., and Mary R. Anderson. 2004. “The Knowledge Gap: A Reexamination of Gender-Based Differences in Political Knowledge.” *Journal of Politics* 66(2):492–512.
Mondak, Jeffery, Christopher Lewis, Jason Sides, Joohyun Kang, and J. Olyn Long. 2004. “Presidential Source Cues and Policy Appraisals, 1981-2000.” *American Politics Research* 32(2):219–35.
Mondak, Jeffery. 1993. “Public Opinion and Heuristic Processing of Source Cues.” *Political Behavior* 15(2):167–92.
Mondak, Jeffery. 2000. “Reconsidering the Measurement of Political Knowledge.” *Political Analysis* 8(1):57–82.
Mondak, Jeffrey J. 2001. “Developing Valid Knowledge Scales.” *American Journal of Political Science* 45(1):224–38.
Neal, Zachary. 2014. “The Devil Is in the Details: Differences in Air Traffic Networks by Scale, Species, and Season.” *Social Networks* 38:63–73.
Newman, Mark E. J. 2006. “Modularity and Community Structure in Networks.” *Proceedings of the National Academy of Sciences* 103(23):8577–82.
Nicholson, Stephen P. 2012. “Polarizing Cues.” *American Journal of Political Science* 56(1):52–66.
Pachucki, Mark, and Ronald Breiger. 2010. “Cultural Holes: Beyond Relationality in Social Networks and Culture.” *Annual Review of Sociology* 36(1):205–24.
Perrin, Andrew J., and Katherine McFarland. 2011. “Social Theory and Public Opinion.” *Annual Review of Sociology* 37(1):87–107.
Peterson, Richard A. 1997. “The Rise and Fall of Highbrow Snobbery as a Status Marker.” *Poetics* 25(2–3):75–92.
Peterson, Richard A. 2005. “Problems in Comparative Research: The Example of Omnivorousness.” *Poetics* 33(5–6):257–82.
Peterson, Richard A., and Roger M. Kern. 1996. “Changing Highbrow Taste: From Snob to Omnivore.” *American Sociological Review* 61(5):900–907.
Petrocik, John R. 1996. “Issue Ownership in Presidential Elections, with a 1980 Case Study.” *American Journal of Political Science* 40(3):825–50.
Petrocik, John R., William L. Benoit, and Glenn J. Hansen. 2003. “Issue Ownership and Presidential Campaigning, 1952–2000.” *Political Science Quarterly* 118(4):599–626.
Porter, Mason, Jukka-Pekka Onnela, and Peter Mucha. 2009. “Communities in Networks.” *Notices of the American Mathematical Society* 56(9).
Radocy, Rudolf, and J. David Boyle. 2012. *Psychological Foundations of Musical Behavior.* 5th ed. Charles C Thomas Publisher.

Ritzer, George, and Jeff Stepnisky. 2013. *Sociological Theory*. 9th edition. New York, NY: McGraw-Hill.
Rodgers, Joseph Lee, and W. Alan Nicewander. 1988. "Thirteen Ways to Look at the Correlation Coefficient." *The American Statistician* 42(1):59–66.
Rosenberg, Paul. 2014. "This Is Why Conservatives Win: George Lakoff Explains the Importance of Framing — and What Democrats Need to Learn." *Salon*. Retrieved April 10, 2015 (http://www.salon.com/2014/11/22/this_is_why_conservatives_win_george_lakoff_explains_the_importance_of_framing_and_what_democrats_need_to_learn/).
Royston, J. P. 1983. "Some Techniques for Assessing Multivariate Normality Based on the Shapiro- Wilk W." *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 32(2):121–33.
Rubinov, Mikail, and Olaf Sporns. 2011. "Weight-Conserving Characterization of Complex Functional Brain Networks." *NeuroImage* 56(4):2068–79.
Saussure, Ferdinand de. 1916. *Course in General Linguistics*. Columbia University Press.
Savage, Mike, and Elizabeth B. Silva. 2013. "Field Analysis in Cultural Sociology." *Cultural Sociology* 7(2):111–26.
Schuman, Howard, and Stanley Presser. 1996. *Questions and Answers in Attitude Surveys: Experiments on Question Form, Wording, and Context*. SAGE.
Sewell, William H. 2005. "The Concept (S) of Culture." *Practicing History: New Directions in Historical Writing after the Linguistic Turn* 76–95.
Shannon, C. E. 1948. "A Mathematical Theory of Communication." *The Bell System Technical Journal* 27:379–423, 623–656.
Shwed, Uri, and Peter S. Bearman. 2010. "The Temporal Structure of Scientific Consensus Formation." *American Sociological Review* 75(6):817–40.
Smith, Eliot R. 1994. "Procedural Knowledge and Processing Strategies in Social Cognition." Pp. 1–99 In *Handbook of Social Cognition*, edited by R. Wyer and T. Srull. Lawrence Erlbaum.
Smith, Eliot R., Nyla R. Branscombe, and Carol Bormann. 1988. "Generality of the Effects of Practice on Social Judgment Tasks." *Journal of Personality and Social Psychology* 54(3):385–95.
Smith, Joanne R., and Michael A. Hogg. 2008. "Social Identity and Attitudes." Pp. 337–360 In *Attitudes and Attitude Change*.
Sniderman, Paul M., and Edward H. Stiglitz. 2012. *The Reputational Premium: A Theory of Party Identification and Policy Reasoning*. Princeton University Press.
Sniderman, Paul. 2000. "Taking Sides: A Fixed Choice Theory of Political Reasoning." In *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*. Cambridge University Press.
Snow, David A., E.Burke Rochford, Steven K. Worden, and Robert D. Benford. 1986. "Frame Alignment Processes, Micromobilization, and Movement Participation." *American Sociological Review* 51(4):464–81.
Sohn, Yunkyu, Myung-Kyu Choi, Yong-Yeol Ahn, Junho Lee, and Jaeseung Jeong. 2011. "Topological Cluster Analysis Reveals the Systemic Organization of the Caenorhabditis Elegans Connectome." *PLoS Comput Biol* 7(5):e1001139.

Somers, Margaret R., and Fred Block. 2005. "From Poverty to Perversity: Ideas, Markets, and Institutions over 200 Years of Welfare Debate." *American Sociological Review* 70(2):260–87.
Stockburger, David. 2007. *Introductory Statistics: Concepts, Models, and Applications.* Cengage Learning.
Strauss, Claudia. 2012. *Making Sense of Public Opinion: American Discourses about Immigration and Social Programs.* New York: Cambridge University Press.
Swidler, Ann. 1986. "Culture in Action: Symbols and Strategies." *American Sociological Review* 51(2):273–86.
Swidler, Ann. 2001. *Talk of Love: How Culture Matters.* University Of Chicago 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.
Tampubolon, Gindo. 2008. "Revisiting Omnivores in America circa 1990s: The Exclusiveness of Omnivores?" *Poetics* 36(2–3):243–64.
Taylor, Shelley E. 1981. "The Interface of Cognitive and Social Psychology." *Cognition, Social Behavior, and the Environment* 189–211.
Thornton, Patricia H., William Ocasio, and Michael Lounsbury. 2012. *The Institutional Logics Perspective: A New Approach to Culture, Structure and Process.* 1st ed. Oxford University Press.
Ullman, Shimon. 1989. "Aligning Pictorial Descriptions: An Approach to Object Recognition." *Cognition* 32(3):193–254.
Vaisey, Stephen. 2009. "Motivation and Justification: Toward a Dual-Process Theory of Culture in Action." *American Journal of Sociology* 114(6):1675–1715.
Van Eijck, Koen. 1999. "Socialization, Education, and Lifestyle: How Social Mobility Increases the Cultural Heterogeneity of Status Groups." *Poetics* 26(5–6):309–28.
Vilhena, Daril et al. 2014. "Finding Cultural Holes: How Structure and Culture Diverge in Networks of Scholarly Communication." *Sociological Science* 1:221–38.
von Neumann, John. 1958. The Computer and the Brain. New Haven, CT, USA: Yale University Press.
Wasserman, Stanley, and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications.* 1st ed. Cambridge University Press.
Watts, D. J., and S. H. Strogatz. 1998. "Collective Dynamics of 'small-World'networks." *Nature* 393(6684):440–442.
Weakliem, David. 1993. "Class Consciousness and Political Change: Voting and Political Attitudes in the British Working Class, 1964 to 1970." *American Sociological Review* 58(3):382–97.
Wellman, Barry. 1988. "Structural Analysis: From Method and Metaphor to Theory and Substance." Pp. 19–61 In *Social Structures: A Network Approach, Structural analysis in the social sciences, Vol. 2.*, edited by B. Wellman and S. D. Berkowitz. New York, NY, US: Cambridge University Press.
Williamson, Kevin D. 2016. "Chaos in the Family, Chaos in the State: The White Working Class's Dysfunction." *National Review*, March 17.
Wu, Angela Xiao. 2014. "Ideological Polarization Over a China-as-Superpower Mindset: An Exploratory Charting of Belief Systems Among Chinese Internet Users, 2008-2011." *International Journal of Communication* 8: 2243–2272.

Wuthnow, Robert. 2007. "Cognition and Religion." *Sociology of Religion* 68(4):341–60.
Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. Cambridge University Press.
Zaller, John, and Stanley Feldman. 1992. "A Simple Theory of the Survey Response: Answering Questions versus Revealing Preferences." *American Journal of Political Science* 36(3):579–616.
Zemljic, Barbara, and Valentina Hlebec. 2005. "Reliability of Measures of Centrality and Prominence." *Social Networks* 27(1):73–88.
Zerubavel, Eviatar. 1999. *Social Mindscapes: An Invitation to Cognitive Sociology*. Harvard University Press.
Zhang, Bin, and Steve Horvath. 2005. "A General Framework for Weighted Gene Co-Expression Network Analysis." *Statistical Applications in Genetics and Molecular Biology* 4:Article17.

APPENDIXES

Appendixes for Chapter 2

APPENDIX A. Formal Proofs

In this Appendix, we will prove a number of theorems about center-periphery belief networks produced by the derivation scheme described in the Model of Belief Formation and Network Structure section of the paper. In Theorems 1 and 2, we derive basic formulas that serve as the foundation for the other proofs. In Theorems 3 and 4, we show that all geodesics in such networks follow a simple topology. In Theorem 5, we use this to derive a formula for lengths of transverse geodesics, which compose the majority of geodesics in the system. Finally, in Theorem 6, we draw on the preceding theorems to demonstrate what we term the "central pull" of the belief network, which is the key result of this formal study. We prove that non-trivial transverse geodesics always tend toward the origin of the system, passing either directly through it, or through a highly correlated node. This leads us to recommend shortest-path betweenness as the tool for identifying the belief at the origin of such networks.

Model, Definitions and Notation

All the previously defined terms retain their original definitions, though we restate some of these below with greater precision. We also adopt a more expressive variable naming scheme. We use ■ to indicate ends of proofs; $A \equiv B$ to mean "A is equal by definition to B." We also occasionally use the notation $X =$ (reason) $= Y$ or $X \Leftrightarrow$ (reason) $\Leftrightarrow Y$ to add concise explanations of why X is equivalent to Y.

<u>*Nodes.*</u> We will use the letters *a, b, c, d, f, g* and x to refer to nodes (beliefs). Node a_K will be called the *ancestor* of node a_S, and a_S the *descendant* of node a_K, if and only if $a_S = a_K + \sum_{i=1}^{n} \phi_i$. Node a_K will be called the *parent* of node a_S, and a_S the *child* of a_K, if and only if $a_S = a_K + \phi_i$. The central belief x_0 is the ancestor of all other nodes, so any node b can be expressed as $b = x_0 + \sum_{i=1}^{m} \phi_i$. We assume that all ϕ_i are independent of x_0, and of all ϕ_j unless $i = j$. We also assume x_0 and ϕ_i have variances $Var(x_0) = 1$ and $Var(\phi_i) = \epsilon$, and finite means.[1]

Node b's *generation* $\eta(b)$ is then equal to *m*, which is the number of ϕ_i terms added to x_0 to produce b. Node x_0 is the only zeroth generation node in the system. We will use subscript indexes to refer to a node's generation.[2] For example, a_1 is a first-generation node; and a_K and b_K are two nodes of the same generation $\eta(a_K) = \eta(b_K) = K$. If a_K is the parent of a_S, then $S = K + 1$. In cases when we need to use a node's index to indicate anything other than the node's generation, we will use a superscript instead: e.g., a^K or a_3^K.

We will refer to a node and its full set of ancestors as the node's *ancestry* $\pi(a)$. The node c_i in $\pi(a) \cap \pi(b)$ that has the largest value of $\eta(c_i)$ will be called the *lowest common ancestor* of a and b. We will call two nodes a and b *strangers* if and only if (*i*) neither node is x_0, and (*ii*) x_0 is their only common ancestor. $T(a, b)$ would then be a *stranger* tie. We assume that the majority of node pairs in the network are strangers. We make no further assumptions about the network's topology.

The belief network $\mathcal{G} = \{\mathcal{N}, \mathcal{T}\}$ is the set of all nodes $\mathcal{N}$ and all ties $\mathcal{T}$, which contains one tie for *every* pair of nodes in $\mathcal{N}$ (i.e., $\mathcal{G}$ is fully connected). We define the length of tie connecting any pair of nodes a and b as

[1] The means must be finite for correlation to be defined. However, the means do not have to be the same, and the variables do not have to be identically distributed.

[2] Please note that this indexing convention applies to nodes only.

$$|T(a,b)| \equiv 1/cor^2(a,b). \tag{1.1}$$

Paths. We will use capital Greek letters to refer to network paths. A path is an ordered sequence of adjacent ties, or, equivalently, of the nodes that are the endpoints of these ties. We will use summation of ties or paths to indicate concatenation into paths. The length of a path is the sum of all the tie lengths composing it: e.g., if $\Lambda = T(a,b) + T(b,c)$, then $|\Lambda| = |T(a,b)| + |T(b,c)|$.

If a path between a and b consists of S ties, we will call it an *S-path*. The distinction between number of ties in a path and path length is important: since path length is the sum of tie lengths as opposed to the count of ties, different S-paths will generally have different lengths. We will use superscript indexes and the function $\tau(\Lambda)$ to specify the number of ties in a path: $(\tau(\Lambda) = S) \equiv (\Lambda$ is a path with S ties$) \equiv (\Lambda = \Lambda^S)$. Since there is only one 1-path connecting any pair of nodes a and b, $\Lambda^1(a,b) = T(a,b)$. We will refer to 1-paths as *trivial* paths. *Central* paths are paths containing x_0. *Transverse* paths are those between strangers.

We will use the letter Γ to refer to geodesics. $\Gamma(a,b)$ is the shortest path connecting *a* and *b*. Note again that this refers to the path with the lowest sum of tie lengths, which is frequently not the same as the path with the fewest ties. We will also occasionally refer to shortest paths between pairs of nodes as *absolute* geodesics to distinguish them from *S-geodesics*, which are the shortest among all S-paths connecting the same pair of nodes. We denote S-geodesics as $\Gamma^S(a,b)$. If Θ^S is the set of all S-paths connecting a and b, then

$$\Lambda(a,b) = \Gamma^S(a,b) \text{ if and only if } \Lambda \in \Theta^S \text{ and } |\Lambda| \leq |\theta^S(a,b)| \ \forall\ \theta^S \in \Theta^S. \tag{1.2}$$

Theorem 1. General formula for length of ties

Theorem. For any two nodes $a = a_K$ and $b = b_S$,

$$|T(a,b)| = \frac{(1 + K * \epsilon)(1 + S * \epsilon)}{(1 + P * \epsilon)^2} \tag{1.3}$$

where P is the generation of the lowest common ancestor of a and b.

Proof. We begin with the variances of a and b. $Var(a) = Var(x_0 + \sum_{i=1}^{K} \phi_i)$ = (due to independence of x_0 and ϕ_i) = $Var(x_0) + \sum_{i=1}^{K} Var(\phi_i) = 1 + K * \epsilon \equiv \sigma^2(a)$. By analogy, $Var(b) = 1 + S * \epsilon \equiv \sigma^2(b)$. (1.4)

We now calculate $Cov(a,b) = Cov\left(x_0 + \sum_{i=1}^{K} \phi_i, x_0 + \sum_{j=1}^{S} \phi_j\right) = Cov(x_0, x_0) + Cov\left(x_0, \sum_{j=1}^{S} \phi_j\right) + Cov\left(\sum_{i=1}^{K} \phi_i, x_0\right) + Cov\left(\sum_{i=1}^{K} \phi_i, \sum_{j=1}^{S} \phi_j\right) = 1 + \sum_{j=1}^{S} Cov(x_0, \phi_j) + \sum_{i=1}^{K} Cov(\phi_i, x_0) + \sum_{i=1}^{K} \sum_{j=1}^{S} Cov(\phi_i, \phi_j)$ = (due to independence of x_0 and ϕ_i, and independence of ϕ_i and ϕ_j unless $i = j$) $= 1 + P * \epsilon$. (1.5)

Using (1.4), we now get

$$Cor(a,b) = [1 + P * \epsilon]/\sqrt{(1 + K * \epsilon) * (1 + S * \epsilon)}, \tag{1.6}$$

and finally,

$$|T(a,b)| = \frac{1}{Cor^2(a,b)} = \frac{(1 + K * \epsilon)(1 + S * \epsilon)}{(1 + P * \epsilon)^2} \ \blacksquare$$

Theorem 2. Lengths of ties to strangers, ancestors and the center

Theorem. If $a = a_K$ and $b = b_S$ are strangers, then:

$$|T(a,b)| = (1 + K * \epsilon) * (1 + S * \epsilon) \tag{2.1}$$

If $b = b_S$ is an ancestor of $a = a_K$, then:

$$|T(a,b)| = \frac{1 + K * \epsilon}{1 + S * \epsilon} \tag{2.2}$$

If $b = x_0$, then:

$$|T(a,b)| = (1 + K * \epsilon) \tag{2.3}$$

Proof. To prove each of these three statements, substitute appropriate values into equation (1.3): $P = 0$ (for 2.1), $P = S$ (for 2.2), and $S = 0$ (for 2.3). ■

Corollary 2A. For any two strangers of the same generation, a_K and b_K, the central 2-path $\Lambda^2 = T(a_K, x_0) + T(x_0, b_K)$ is longer than the 1-path $\Lambda^1 = T(a_K, b_K)$ if and only if $K * \epsilon < 1$.

Proof of corollary. The path Λ^2 has length $|\Lambda^2| = |T(a_K, x_0)| + |T(x_0, b_K)|$. By (2.3), this equals $(1 + K * \epsilon) + (1 + K * \epsilon) = 2 * (1 + K * \epsilon)$. On the other hand, by (2.1), Λ^1 has length $|\Lambda^1| = (1 + K * \epsilon)^2$. Therefore $|\Lambda^2| > |\Lambda^1| \Leftrightarrow 2 * (1 + K * \epsilon) > (1 + K * \epsilon)^2 \Leftrightarrow 2 > (1 + K * \epsilon) \Leftrightarrow 1 > K * \epsilon$. ■

Application to Figure 1B. In our discussion of Figure 1B, we claimed that, if $\epsilon > 0.5$, the path $\Lambda^1 = (T_{11,21})$ will be longer than the path $\Lambda^2 = (T_{11,0}, T_{0,21})$. Translating this to the more detailed notation we use in this appendix, path $\Lambda^1 = T(x^{11}, x^{21})$, whereas $\Lambda^2 = T(x^{11}, x_0) + T(x_0, x^{21})$. The generation K of x^{11} and x^{21} is 2. Substituting this into Corollary 2A, we see that $|\Lambda^2| > |\Lambda^1|$ if and only if $2 * \epsilon < 1$, or equivalently, if $\epsilon < 0.5$.

Theorem 3. S-geodesics contain at most one stranger tie

For the remainder of this appendix, we will focus on S-geodesics (see 1.2), which we use as an analytical tool for studying absolute geodesics. Each absolute geodesic $\Gamma(a,b)$ can also be seen as an S-geodesic $\Gamma^S(a,b)$, with $S = \tau(\Gamma)$. Additionally, $|\Gamma(a,b)| = \min\{|\Gamma^S(a,b)|: 1 \leq S < \infty\}$. The crucial advantage of S-geodesics is that we can express their length analytically, and thus study them with optimization via partial derivatives. Any statement proven to hold for all S-geodesics by extension also holds for absolute geodesics.

Theorem. If an S-path Ω between g and f contains two or more stranger ties, Ω is not the S-geodesic between g and f. (3.1)

Proof. Let us traverse path Ω from g to f. Let $T(a,b)$ be the first of two stranger ties we encounter, and $T(c,d)$ the second of these ties. There are only three possibilities about how $T(a,b)$ and $T(c,d)$ are related (see figure A1):

i. $T(a,b)$ and $T(c,d)$ are immediately adjacent, i.e., $b = c$.
ii. c is a descendant of b.
iii. c is an ancestor of b.

[Figure A1 around here]

In each of these cases we will show that some 2-tie segment $\Theta \subset \Omega$ can be replaced by a shorter 2-tie segment Θ_0, which passes through x_0. This contradicts the definition of an S-geodesic.

In case (i), we will show that $\Theta = T(a,b) + T(b,d)$ is longer than $\Theta_0 = T(a,x_0) + T(x_0,d)$. By (2.1) and (2.3), $|\Theta_0| = (1 + \eta(a) * \epsilon) + (1 + \eta(d) * \epsilon)$ and $|\Theta| =$

$(1+\eta(a)*\epsilon)*(1+\eta(b)*\epsilon)+\ (1+\eta(c)*\epsilon)*(1+\eta(d)*\epsilon)$. Since $\eta(b), \eta(d)$ and $\epsilon > 0, |\Theta| > |\Theta_0|$.

In case (ii), we will assume for simplicity that *b* is the parent of *c* (otherwise, the same proof holds if $T(b,c)$ is replaced with the tie from *c*'s parent to *c*). We will show that $\Theta = T(b,c) + T(c,d)$ is longer than $\Theta_0 = T(b,x_0) + T(x_0,d)$. By (2.1) – (2.3), $|\Theta| = \frac{1+\eta(c)*\epsilon}{1+\eta(b)*\epsilon} + (1+\eta(c)*\epsilon)*(1+\eta(d)*\epsilon) > 1 + (1+\eta(c)*\epsilon)*(1+\eta(d)*\epsilon) > 1 + (1+\eta(b)*\epsilon)*(1+\eta(d)*\epsilon) = 1 + (1+\eta(b)*\epsilon) + \eta(d)*\epsilon + \eta(b)*\eta(d)*\epsilon^2 > (1+\eta(b)*\epsilon) + (1+\eta(d)*\epsilon) \Leftrightarrow |\Theta| > |\Theta_0|$.

Since the tie lengths are symmetric, case (iii) can be made identical to (ii) by simply reversing the order in which the nodes in Ω are traversed, and renaming them accordingly. ■

Theorem 4. S-geodesics pass only through ancestors

Theorem. An S-geodesic $\Lambda = \Gamma^S(a,b)$ passes only through nodes that are ancestors of *a* or *b*:

$$\Gamma^S(a,b) \subseteq \pi(a) \cup \pi(b) \tag{4.1}$$

Proof. Let us assume that (4.1) is not the case, and Λ contains nodes not in $\pi(a)$ or $\pi(b)$. The ties connecting such nodes to those in $\pi(a)$ and $\pi(b)$ will then be stranger ties. Thus, Λ then contains at least two stranger ties, which contradicts Theorem 3. ■

Theorem 5. Length of transverse S-geodesics

In this theorem, we derive a formula for the length of transverse S-geodesics by addressing it as a minimization problem over all possible S-paths. In order to perform this optimization task using partial derivatives, from this point on we let all N_i and K_j assume continuous values. It can be shown that this solution is asymptotically equivalent to the discrete case when $max(N,K) \to \infty$.[3]

Theorem: For any two strangers $a = a_K$ and $b = b_N$, if the path Λ is an S-geodesic $\Gamma^S(a,b)$, then its length satisfies the following equation:

$$|\Lambda(a,b)| = S * [(1+K*\epsilon)*(1+N*\epsilon)]^{1/S} \tag{5.1}$$

Additionally, path Λ either passes through the center x_0 or through a pair of nodes a^1 and b^1 which satisfy the following equation:

$$(1\ +\ \mathrm{K}_1*\epsilon)*(1\ +\ \mathrm{N}_1*\epsilon) = \big((1\ +\ \mathrm{K}*\epsilon)*(1\ +\ \mathrm{N}*\epsilon)\big)^{1/S} \tag{5.2}$$

where $K_1 = \eta(a^1)$ and $N_1 = \eta(b^1)$.

Proof. First let $\Lambda(a,b)$ be a non-central S-geodesic. According to Theorems 3 and 4, all nodes in Λ belong to $\pi(a) \cup \pi(b)$, and Λ contains a single stranger tie $T(a^1,b^1)$ which connects a node in $\pi(a)$ to a node in $\pi(b)$. We will refer to $T(a^1,b^1)$ as the *bridge*. We will number the nodes beginning with a^1 and b^1 and counting outwards toward *a* and *b* (see right side of figure A2). Then, by Theorem 4 and definition of non-central S-geodesics, Λ can be represented as:

$$\Lambda = (a = a^T, a^{T-1}, \dots, a^2, a^1, b^1, b^2, \dots, b^{S-T+1} = b). \tag{5.3}$$

[3] Their equivalence is in fact stronger than asymptotic, as when both the discrete and continuous solutions to the minimization problem approach ∞, the absolute difference between the two remains less than a fixed value.

There are $T-1$ ties between nodes in $\pi(a)$ and $S-T$ ties between nodes in $\pi(b)$, so that the total number of ties in Λ is $(T-1)+1+(S-T)=S$.

[Figure A2 about here]

We will denote $K_i = \eta\left(a^i\right)$ and $N_j = \eta(b^j)$. The following are true by definition:

$$0 < K_1 < K_2 < \cdots < K_{T-1} < K_T \text{ and } 0 < N_1 < N_2 < \cdots < N_{S-T} < N_{S-T+1} \tag{5.4}$$

$$K_T = K = \eta(a_K) \tag{5.5}$$

$$N_{S-T+1} = N = \eta(b_N) \tag{5.6}$$

By (2.1) and (2.2), the length of Λ can be expressed as the following function L, which adds up the lengths of all the ties that lie within $\pi(a)$ and $\pi(b)$, as well as of the bridge:

$$L = \sum_{i=1}^{T-1} \frac{1 + K_{i+1} * \epsilon}{1 + K_i * \epsilon} + (1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) + \sum_{j=1}^{S-T} \frac{1 + N_{j+1} * \epsilon}{1 + N_j * \epsilon} \tag{5.7}$$

Since *a* and *b* are fixed, L is a function of $S-1$ independent variables:

$$L = L(K_1, \dots, K_{T-1}, N_1, \dots, N_{S-T}) \tag{5.8}$$

By the definition of *S*-geodesics, $|\Lambda(a,b)|$ is the minimum value of L over all the possible combinations of intermediate nodes from within the appropriate ancestries. To find the length of the *S*-geodesic, we thus minimize L using the partial derivatives $\frac{\partial L}{\partial K_i}(1 \le i \le T-1)$ and $\frac{\partial L}{\partial N_j}(1 \le j \le S-T)$. Note that each K_i term except K_1 appears in exactly two elements of the first summation in (5.7), and different K_i terms are not functions of each other. The same situation also holds for N_j. Thus:

$$\frac{\partial L}{\partial K_i} = \frac{\partial}{\partial K_i}\left[\frac{1 + K_{i+1} * \epsilon}{1 + K_i * \epsilon} + \frac{1 + K_i * \epsilon}{1 + K_{i-1} * \epsilon}\right] \text{for } 2 \le i \le (T-1) \tag{5.9}$$

$$\frac{\partial L}{\partial N_j} = \frac{\partial}{\partial N_j}\left[\frac{1 + N_{j+1} * \epsilon}{1 + N_j * \epsilon} + \frac{1 + N_j * \epsilon}{1 + N_{j-1} * \epsilon}\right] \text{for } 2 \le j \le (S-T) \tag{5.10}$$

Terms K_1 and N_1 appear once in their respective summations and once in the bridge term. Thus:

$$\frac{\partial L}{\partial K_1} = \frac{\partial}{\partial K_1}\left[\frac{1 + K_2 * \epsilon}{1 + K_1 * \epsilon} + (1 + K_1 * \epsilon) * (1 + N_1 * \epsilon)\right] \tag{5.11}$$

$$\frac{\partial L}{\partial N_1} = \frac{\partial}{\partial N_1}\left[\frac{1 + N_2 * \epsilon}{1 + N_1 * \epsilon} + (1 + K_1 * \epsilon) * (1 + N_1 * \epsilon)\right] \tag{5.12}$$

Performing the differentiations in (5.9) – (5.12) and setting each partial derivative to zero, we get the following system of equations:

$$\begin{cases} \text{from 5.9:} & (1 + K_{i+1} * \epsilon) * (1 + K_{i-1} * \epsilon) = (1 + K_i * \epsilon)^2 \text{ for } 2 \le i \le (T-1) & (5.13) \\ \text{from 5.10:} & \left(1 + N_{j+1} * \epsilon\right) * \left(1 + N_{j-1} * \epsilon\right) = \left(1 + N_j * \epsilon\right)^2 \text{ for } 2 \le j \le (S-T) & (5.14) \\ \text{from 5.11:} & (1 + K_2 * \epsilon) = (1 + K_1 * \epsilon)^2 * (1 + N_1 * \epsilon) & (5.15) \\ \text{from 5.12:} & (1 + N_2 * \epsilon) = (1 + N_1 * \epsilon)^2 * (1 + K_1 * \epsilon) & (5.16) \end{cases}$$

We note that equations (5.13) and (5.14) are geometric progressions. Thus we will search for a solution to (5.13) – (5.16), as well as (5.5) – (5.6), that satisfies the following two equations with five unknown parameters $r, \alpha, \beta, \gamma, \delta$, which define such geometric progressions:

$$\begin{cases} 1 + K_i * \epsilon = r^{\alpha * i + \beta}, 1 \le i \le T & (5.17) \\ 1 + N_j * \epsilon = r^{\gamma * j + \delta}, 1 \le j \le S - T + 1 & (5.18) \end{cases}$$

As can be shown by substitution, (5.17) – (5.18) satisfy (5.13) and (5.14) for any values of $r, \alpha, \beta, \gamma,$ and δ. To determine which values satisfy the remaining equations, we substitute the following into (5.15): $(1 + K_2 * \epsilon) = r^{\alpha * 2 + \beta}$, $(1 + K_1 * \epsilon)^2 = r^{2*(\alpha*1+\beta)} = r^{2\alpha+2\beta}$, and $(1 + N_1 * \epsilon) = r^{\gamma*1+\delta}$. This yields

$$r^{2\alpha+\beta} = r^{2\alpha+2\beta} * r^{\gamma+d} \Leftrightarrow 2\alpha + \beta = 2\alpha + 2\beta + \gamma + \delta \Leftrightarrow 0 = \beta + \gamma + \delta \quad (5.19)$$

Analogous substitutions into (5.16) yield

$$r^{2\gamma+\delta} = r^{2\gamma+2\delta} * r^{\alpha+\beta} \Leftrightarrow 0 = \beta + \alpha + \delta \quad (5.20)$$

Subtracting (5.19) from (5.20) produces $\beta + \alpha + \delta - (\beta + \gamma + \delta) = 0 \Leftrightarrow \alpha = \gamma$.
Since the values of r, β and δ can be adjusted to make α and γ equal any constant, we assume that $\alpha = \gamma = 1$. (5.21)
Substituting this into (5.19) or (5.20) yields $\beta = -1 - \delta$. (5.22)
Substituting (5.21) and (5.22) into (5.17) and (5.18) now yields:

$$\begin{cases} 1 + K_i * \epsilon = r^{i-1-\delta}, 1 \le i \le T & (5.23) \\ 1 + N_j * \epsilon = r^{j+\delta}, 1 \le j \le S - T + 1 & (5.24) \end{cases}$$

To find the value of r, we now examine the end nodes A_K and B_N. By (5.5) and (5.6),

$$1 + K * \epsilon = 1 + K_T * \epsilon = (\text{by } 5.23) = r^{T-1-\delta} \quad (5.25)$$

$$1 + N * \epsilon = 1 + N_{S-T+1} * \epsilon = (\text{by } 5.24) = r^{S-T+1+\delta} \quad (5.26)$$

Multiplying both sides of (5.25) by (5.26) produces:

$$\begin{gathered} (1 + K * \epsilon) * (1 + N * \epsilon) = r^{T-1-\delta} * r^{S-T+1+\delta} = r^S \Leftrightarrow \\ r = [(1 + K * \epsilon) * (1 + N * \epsilon)]^{\frac{1}{S}} \end{gathered} \quad (5.27)$$

Now, we substitute (5.23) – (5.24) into each term of (5.7), beginning with the first summation:

$$\sum_{i=1}^{T-1} \frac{1 + K_{i+1} * \epsilon}{1 + K_i * \epsilon} = \sum_{i=1}^{T-1} \frac{r^{i+1-1-\delta}}{r^{i-1-\delta}} = \sum_{i=1}^{T-1} r = (T-1) * r$$

$$(1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) = r^{1-1-\delta} * r^{1+\delta} = r$$

$$\sum_{j=1}^{S-T} \frac{1 + N_{j+1} * \epsilon}{1 + N_j * \epsilon} = \sum_{i=1}^{T-1} \frac{r^{j+1+\delta}}{r^{j+\delta}} = \sum_{j=1}^{S-T} r = (S-T) * r$$

Therefore, the minimum value of L over all real values of $K_1, \dots, K_{T-1}, N_1, \dots, N_{S-T}$ is:
$(T-1) * r + r + (S-T) * r = S * r = (\text{by } 5.27) = S * [(1 + K * \epsilon) * (1 + N * \epsilon)]^{1/S}$, which proves (5.1) for non-central S-geodesics.

If $\Lambda^*(a, b)$ is a central geodesic, it will have the form $\Lambda^* = (a = a^T, a^{T-1}, \dots, a^2, a^1, x_0, b^1, b^2, \dots, b^{S-T+1} = b)$, as depicted on the left side of figure A2. Since Λ^* does not contain a bridge tie, its length is simply $L^* = \sum_{i=1}^{T} \frac{1 + K_{i+1} * \epsilon}{1 + K_i * \epsilon} + \sum_{j=1}^{S-T} \frac{1 + N_{j+1} * \epsilon}{1 + N_j * \epsilon}$. The proof of (5.1) for the central case is otherwise nearly identical to the non-central case above, and we omit it for reasons of space.

Finally, to prove statement (5.2), we combine (5.27), (5.23) and (5.24) when $i = j = 1$:

$$(1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) = r^{1-1-\delta} * r^{1+\delta} = r = \left((1 + K * \epsilon) * (1 + N * \epsilon)\right)^{1/S}$$

This concludes the proof of Theorem 5.■

Theorem 6. The "central pull" of the system

Finally, we will prove that non-trivial transverse geodesics—which are the majority of geodesics in the system—pass through x_0 directly or through a closely correlated node x'. This "central pull" of the system suggests shortest-path betweenness as the tool for identifying the center of such systems. The two corollaries further support the use of betweenness centrality. Corollary 6A shows that trivial geodesics occur only near the center of the system—or, in other words, that a large-enough system will have many non-trivial geodesics. This is a useful property, as betweenness centrality relies on non-trivial geodesics to identify the center. Since some geodesics may pass through a closely correlated node x' instead of x_0, empirical researchers should exercise caution in the presence of multicollinearity. However, Corollary 6B shows that, if two transverse geodesics connecting unrelated pairs of nodes both bypass x_0, they will do so via different nodes. In other words, while there is only one x_0, there many possible nodes can play the role of x'. This lessens the inferential threat presented by collinearity, as it makes it less likely that any one x' should lie on enough geodesics to be mistaken for the center.

Theorem. All non-trivial transverse geodesics $\Gamma(a_K, b_N)$ pass through either x_0 or through node x' whose absolute Pearson's correlation with x_0 exceeds $\left(2/3\right)^{3/4} \approx 0.74$.

Corollary 6A. If a transverse geodesic $\Gamma(a_K, b_N)$ is trivial, then $\eta(a) \leq 1/\epsilon$ or $\eta(b) \leq 1/\epsilon$.

Corollary 6B. Assume $\Gamma(a, b)$ and $\Omega(c, d)$ are two non-central, non-trivial transverse geodesics. Let x' be the node in $\Gamma(a, b)$ that is closest to x_0 and x'' be the node in $\Omega(c, d)$ that is closest to x_0. If a, b, c and d are strangers, then $x' \neq x''$.

Proof of theorem. By Theorem 3, $\Gamma(a,b)$ can contain either no stranger ties or one stranger tie. In either case, $\Gamma(a,b)$ must contain a tie $T(f, g)$ uniting a node in $\pi(a)$ with a node in $\pi(b)$. (6.1)

If $\Gamma(a, b)$ contains no stranger ties, f or g must be in $\pi(a) \cap \pi(b) = \{x_0\}$. Thus, without stranger ties, $\Gamma(a, b)$ must pass through the origin, which satisfies the condition of this theorem.

We now turn to cases where Γ contains one stranger tie. If $S = \tau(\Gamma)$, by (5.1),

$$|\Gamma(a, b)| = S * [(1 + K * \epsilon) * (1 + N * \epsilon)]^{\frac{1}{S}} \tag{6.2}$$

Substituting $C \equiv (1 + K * \epsilon) * (1 + N * \epsilon)$ into (6.2) transforms it to

$$|\Gamma(a, b)| = S * C^{1/S} \tag{6.3}$$

Furthermore, since Γ is the absolute geodesic between *a* and *b*, it by definition cannot be longer than the S-geodesic $\Gamma^{S+1}(a, b)$. Thus, by (5.1),

$$|\Gamma(a, b)| \leq (S + 1) * C^{1/(S+1)} \tag{6.4}$$

Combining (6.3) and (6.4) and transforming the result yields $S * C^{\frac{1}{S}} \leq (S+1) * C^{\frac{1}{S+1}} \Leftrightarrow$

$$\Leftrightarrow C^{\frac{1}{S}-\frac{1}{S+1}} \leq \frac{S+1}{S} \Leftrightarrow C^{\frac{1}{S*(S+1)}} \leq \frac{S+1}{S} \Leftrightarrow C \leq \left(\frac{S+1}{S}\right)^{S*(S+1)} \Leftrightarrow$$

$$C^{-1} \geq \left(\frac{S}{S+1}\right)^{S*(S+1)} \tag{6.6}$$

By (5.2), there are two nodes in $\Gamma = \Gamma^S(a,b)$ with generations K_1 and N_1, respectively, which satisfy the equation $(1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) = [(1 + K * \epsilon) * (1 + N * \epsilon)]^{\frac{1}{S}}$. Thus, by definition of C, $(1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) = C^{\frac{1}{S}}$. (6.7)

We will consider the case when $K_1 \leq N_1$ (the opposite case is nearly identical):

$(1 + K_1 * \epsilon)^2 \leq (1 + K_1 * \epsilon) * (1 + N_1 * \epsilon) \Leftrightarrow (1 + K_1 * \epsilon)^2 \leq C^{\frac{1}{S}} \Leftrightarrow (1 + K_1 * \epsilon) \leq C^{1/(2*S)} \Leftrightarrow$

$$1/(1 + K_1 * \epsilon)^{1/2} \geq C^{-1/(4*S)} \tag{6.8}$$

By (1.1), $Cor^2(K_1, x_0) = 1/(1 + K_1 * \epsilon) \Leftrightarrow Cor(K_1, x_0) = 1/(1 + K_1 * \epsilon)^{1/2}$. Thus, by (6.8), $Cor(K_1, x_0) \geq C^{-1/4*S}$, and by (6.6),

$$Cor(K_1, x_0) \geq \left(\frac{S}{S+1}\right)^{(S+1)/4} \tag{6.9}$$

We will now consider the natural logarithm of the function $R(S) = \left(\frac{S}{S+1}\right)^{(S+1)/4}$: $\ln R(S) = \frac{S+1}{4} * \ln\left(\frac{S}{S+1}\right) = \frac{S+1}{4} * \ln\left(1 - \frac{1}{S+1}\right)$. Since both factors increase with S, $\ln R(S)$ increases monotonically, and thus so does $R(S)$. Since Γ is non-trivial, $\min(S) = 2$. Substituting this into (6.9) yields $Cor(K_1, x_0) \geq \left(\frac{2}{3}\right)^{3/4} \approx 0.74$. ■

__Proof of Corollary 6A.__ For trivial geodesics, $S = 1$. Substituting $S = 1$ and $Cor(K_1, x_0) = 1/(1 + K_1 * \epsilon)^{1/2}$ into (6.9) yields

$$\frac{1}{\sqrt{1 + K_1 * \epsilon}} > \left(\frac{1}{2}\right)^{\frac{1}{2}} \Leftrightarrow 1 + K_1 * \epsilon < 2 \Leftrightarrow K_1 * \epsilon < 1 \Leftrightarrow K_1 < 1/\epsilon$$

Since in a trivial geodesic $K_1 = \min(\eta(a), \eta(b))$, this implies $\eta(a) < 1/\epsilon$ or $\eta(b) < 1/\epsilon$. ■

__Proof of Corollary 6B__. Let us assume that 6B is not the case, and that $x' = x''$. Then, by Theorem 4, $x' \in (\pi(a) \cup \pi(b))$ and $x' = x'' \in (\pi(c) \cup \pi(d))$. Thus $x' \in [\pi(a) \cup \pi(b)] \cap [\pi(c) \cup \pi(d)] = ([\pi(a) \cup \pi(b)] \cap \pi(c)) \cup ([\pi(a) \cup \pi(b)] \cap \pi(d)) \Leftrightarrow$

$$x' \in [\pi(a) \cap \pi(c)] \cup [\pi(b) \cap \pi(c)] \cup [\pi(a) \cap \pi(d)] \cup [\pi(b) \cap \pi(d)]. \tag{6.10}$$

If two nodes *f* and *g* are strangers, then by definition $\pi(f) \cap \pi(g) = \{x_0\}$. Thus (6.10) equals

$$x' \in \{x_0\} \cup \{x_0\} \cup \{x_0\} \cup \{x_0\} \Leftrightarrow x' = x_0.$$

This contradicts the definition of non-central geodesics. ■

APPENDIX B. NETWORK PARTITIONING.

The moral politics and social constraint perspectives both suggest that belief networks may contain subgroups of beliefs that exist in relative independence from the rest of the network. For example, in Zaller's (1992) account, anti-war beliefs exhibit a partial decoupling because over-time changes in party positions on this issue cause different followers of the same party to acquire receive messages. In Lakoff's (2002) account, beliefs about gender roles are similarly decoupled because they share a common foundation that varies independently from other beliefs. Both accounts suggest that beliefs may form subgroups with ties that are stronger within groups than they are between them. Since the average strength of belief correlations varies between different populations (Converse 1964) and depends on properties of the survey instrument (Martin 1999), the specific criteria for within- and between-group tie strengths cannot be determined a-priori. A method for locating community structure in belief networks should thus instead do so based on the observed distribution of tie strengths. Newman's (2006) *modularity maximization* technique is commonly used for this kind of partitioning. However, recent methodological work raises concerns about that method's accuracy and validity in many empirical settings. In this appendix, we review this work, and examine whether the accuracy problems apply to our empirical case.

Newman's approach uses an objective function called "modularity." The modularity of any partition of a given network into mutually exclusive subgroups measures the degree to which that partition results in stronger-than-chance ties within each subgroup and weaker-than-chance ties between subgroups. Modularity maximization algorithms search the space of possible partitions of a given network for the one that yields the greatest modularity. This technique thus tries to locate partitions defined by "statistically surprising" (Newman 2006) arrangements of ties. However, two well-documented problems often compromise its ability to do so.

First, the modularity function suffers from a "resolution limit" that can bias it against *(i)* detecting small modules in large networks, as well as *(ii)* large modules in small networks, even when the existence of such modules appears intuitively clear and can be easily determined by other methods (Fortunato and Barthélemy 2007; Good, de Montjoye, and Clauset 2010; Lancichinetti and Fortunato 2011). Though more recent versions of the algorithm introduce a tuning parameter that can make either bias *(i)* or bias *(ii)* less likely, it is practically impossible to avoid both biases simultaneously (Lancichinetti and Fortunato 2011). The two biases make it less likely that modularity maximization would determine that all the nodes in the network belongs to only a single module, or that every node belongs in its own single-node module—i.e., that there is no subgroup structure. This renders it a poor fit for testing structural hypotheses like those suggested by the moral politics and social constraint accounts. As we discuss in Appendix G, the RCA technique used by Baldassarri and Goldberg (2014) to claim evidence for heterogeneous logics of belief organization also relies on modularity maximization. We demonstrate problems with RCA results that are consistent with this bias (see Appendix G.)

Second, the modularity function exhibits degeneracies that greatly compromise its reliability in many empirical settings (Rubinov and Sporns 2011; Good et al. 2010). Rather than a unique maximum that clearly recommends one optimal partition of the network, the modularity function often has multiple near-maxima corresponding to distinct partitions with approximately equal modularities. This leaves its results unstable to small changes in the network. To determine whether this occurs with our data, we analyzed the modularity function for the full population, low- and high-information, and African American belief networks. For each, we compared the modularity scores across different partitions produced by the hierarchical "*fastgreedy*" maximization algorithm (in separate analyses, we compared *fastgreedy* to an exhaustive search and found that *fastgreedy* generally located the optimal or nearly-optimal belief network partitions). Since these are only a small portion of all possible partitions of each network, this search for degeneracies was conservative. Nonetheless, all four distributions exhibited these near-maxima, with 5 to 8 distinct partitions having modularity scores within 5% of the maximum in each network. A bootstrapping analysis confirmed that the resulting subgroup assignments were unreliable.

Because of these problems, we do not report the modularity maximization results in our analyses, and cannot recommend the algorithm's use with belief networks. Future work should explore the reliability and theoretical fitness of other network partitioning methods, and may instead need to develop a new method that is better tailored to this kind of data. Such work should also compare these partitioning schemes to more traditional factor-analytic techniques for grouping variables, both in terms of reliability and of theoretical fit.[4]

APPENDIX C. BIVARIATE LINEARITY.

The BNA methodology we develop in this paper analyzes weighted networks constructed from squared correlations between pairs of belief variables. We use polychoric correlations between ordinal variables, polyserial correlations between numeric and ordinal, and Pearson's correlations between numeric variables. Our model and method thus carry an assumption that correlation can fully capture the pairwise relationships between beliefs—or, equivalently, that the relationships between these beliefs are linear in character.[5] This assumption is commonly made in the literature on belief structures (e.g., Converse 1964; Jennings 1992). Nonetheless, it may be possible that relationships between beliefs are more complex, due to either non-linear functional form or to unobserved heterogeneity. For example, if two beliefs could be captured by integer-valued variables $x \in \{-2, \dots, 2\}$ and y, the true relationship between them may take the form of a parabola $y = x^2$. Even though x and y are deterministically interrelated, their Pearson's correlation $r(x, y = x^2)$ would equal 0. In the presence of such non-linearity, our reliance on correlation could thus lead us to overlook relationships between variables.

To examine this possibility, we build on Martin's (1999, 2002) work on entropic measures of constraint to develop an information-theoretic test for whether such non-linearity is present in our data. We base this test on *mutual information*, which is a general-purpose non-parametric measure of association between discrete variables. Unlike linear relationship measures like Pearson's and polychoric correlation, mutual information quantifies the amount of non-independence between x and y without any assumptions about the functional form of their relationship, or about the relative ordering of each variable's levels. The mutual information $I(x, y) \in [0, +\infty)$ equals zero if and only if x and y are fully independent. It otherwise quantifies the extent to which knowing the value of one variable would reduce uncertainty regarding the value of the other.

The relationship between mutual information and marginal (univariate) entropies $H(x)$ and $H(y)$ is roughly analogous to that between covariance and variance. To yield a correlation-like measure of pairwise association between x and y, mutual information can be normalized to the [0,1] range via $\hat{I}(x, y) = I(x, y)/\min\{H(x), H(y)\}$ (Kvalseth 1987; Yao 2003). If either of the variables fully describes all the variation in the other, $\hat{I}(x, y)$ achieves its maximum of 1.

[4] Older work has in fact used factor analyses of the adjacency matrix to partition social networks into subgroups, though Wasserman and Faust (1994:290) note that this approach may not yield subgroups with theoretically desirable properties.

[5] In the case of our formal model, this is a consequence of our assumed belief derivation process (Appendix A) rather than a separate assumption in the formal sense of the term. For ordinal variables, the assumption concerns latent beliefs rather than manifest indicators.

Since in the above example, each value of *y* maps onto exactly one value of *x*, $\hat{I}(x, y) = 1$. In the presence of random noise, this relationship will decrease, but will remain above zero as long as the noise does not "overwhelm" *y* to the point that *x* and *y* are independent (we return to issues of noise later in the appendix).

Mutual information can also be used to detect a relationship between a pair of variables *x* and *z* even in the presence of unobserved heterogeneity, as long as the heterogeneity does not render *x* and *z* independent (in the latter case, no method could possibly detect the latent relationship between *x* and *z* without introducing further assumptions). For example, consider a further pair of integer-valued belief variables *x* and *z*, where $x \in \{-2, -1, 1, 2\}$. Belief *x* may have a positive linear relationship with belief $z = (z_1; z_2)$ in one half of the population, where $z = z_1 = 2x$, and a symmetrical negative relationship in the other half, where $z = z_2 = -2x$. The resulting relationship between *x* and *z* is an "×" shape centered at the origin. Due to the absence of a linear relationship, their Pearson's correlation is $r(x, z) = 0$. However, *x* and *z* are not independent, as each value of *z* maps onto only two values of *x*, and vice versa. Correspondingly, their normalized mutual information is above zero: $\hat{I}(x, z) = 0.5$.

In general, if the population consists of two subgroups where $z = z_1 = ax + b$ in the first subgroup and $z = z_2 = cx + d$ in the second, and the variable *x* assumes at least 3 different values,[6] each value of *x* would correspond to only a proper subset of values of *z*. Therefore, even in the presence of such heterogeneity, *x* and *z* would still not be independent, and their $\hat{I}$ would be above zero. The same can clearly not be said of correlation, which equals zero if $a = -c$.

With this consideration in mind, we calculated the mutual informations $\hat{I}_{ij}$ between each pair (i, j) of the 46 variables in our ANES data, placing them in the matrix $\hat{\boldsymbol{I}} = [\hat{I}_{ij}]$. To calculate $\hat{I}$, we used the empirical estimators in the *R* package "*entropy*" (Hausser and Strimmer 2013). Since mutual information only exhibits the properties described above when the variables are discrete, we used the original survey items from the Limited Government, Gay Rights, Welfare and Racism scales in place of the factored scores. This yielded 46 belief variables in total.[7] For comparison, we also constructed a correlation matrix $\boldsymbol{C} = [C_{ij}]$, where each cell C_{ij} contains the squared polychoric correlation between the same beliefs *i* and *j*. We plot the values of C_{ij} and $\hat{I}_{ij}$ for all variable pairs $i \neq j$ in the data (figure C1).

[Figure C1 about here]

Consider the element-wise Pearson's correlation between these two matrixes, $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$. This correlation would only be high if *N* and *C* generally varied together, suggesting that mutual information and correlation are registering the same relative rates of interrelatedness between each of the pairs of variables.[8] On the other hand, $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ would be low if $\boldsymbol{C}$ and $\hat{\boldsymbol{I}}$

[6] 41 of the 46 variables in our data (89%) assume 3 or more values

[7] We make this matrix available in the online appendix to this paper.

[8] Of course, since correlation reflects only simultaneous deviations from each matrix's mean, *r(C,N)* could also be high if *C* always missed the same amount of the pairwise relationships between all pairs of variables, i.e., if the degree of non-linearity in the pairwise relationships between variables was constant. If we include the diagonal elements of the matrixes in this comparison, this becomes impossible: the diagonal cells of each matrix measure the relationship of each belief to itself, i.e., $C_{ii} = N_{ii} = 1$, and thus cannot possibly contain any non-linearity. On the other hand, since the diagonal elements of the two matrixes are both at their maximum, their presence can lead us to overestimate the similarity between the two matrixes. We thus report *r(C,N)* with and without the diagonal elements. Additionally, we note that high pairwise correlations between the variables that formerly made up

frequently varied independently of one other, or if their simultaneous variances were related in a non-linear fashion. This makes $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ a relatively conservative measure of whether or not polychoric correlation is missing any substantial non-independence between the variables.

When we calculated $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ for our observed ANES data, we found that $r(\boldsymbol{C}, \hat{\boldsymbol{I}}) = 0.91$, or 0.97 if the matrix diagonals are included in the comparison. This high value of $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ shows that polychoric correlations can account for the majority of the pairwise non-independence between variables in this dataset. It is thus inconsistent with the presence of substantial non-linear relationships in our data.

It is possible, however, that measurement error in survey data may conceal some heterogeneity or complexity that is present in the population. Since survey data are not perfectly reliable, this error could conceivably distort a truly non-linear or heterogeneous relationship between variables to the point that only a single linear trend remains observable. This could lead $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ to underestimate the true amount of heterogeneity or non-linearity in the relationship. Conversely, it is also possible that measurement error may distort linear trends, creating the appearance of complexity when the true relationship between two variables is simply linear. If this were the case, $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ could instead overestimate the amount of heterogeneity or non-linearity, thus leading to an even more conservative test.

To adjudicate between these two possibilities, we make use of the reliability estimates Alwin (2007) calculated for the 1990 ANES panel, which included 18 of the same questions we analyzed here. We subset our data to these 18 variables, and examine the corresponding 18×18 matrixes (excluding diagonals). Within this subset, $r(\boldsymbol{C}, \hat{\boldsymbol{I}})$ equals 0.90, which is close to the value it obtained for the full set of 46 belief variables we examined above. With all the matrix cells $(i, j \neq i)$ as the units of analysis, we regress the normalized mutual informations $\hat{I}_{ij}$ on the reliabilities of the two variables *i* and *j*, controlling for their squared polychoric correlations C_{ij}:

$$\hat{I}_{ij} = \beta_0 + \beta_1 \times reliability_i + \beta_2 \times reliability_j + \beta_3 \times C_{ij} + \varepsilon_{ij}$$

The standardized coefficient for C_{ij} equals 0.93 ($t(302) = 40.39, p < 0.001$), which indicates that correlation retains a strong relationship with mutual information in the presence of these controls. Turning to the reliability variables, the standardized coefficients for both equal -0.13 ($t(302) = -5.75, p < 0.001$), indicating that an increase in reliability is accompanied by a moderate but highly significant *decrease* in the mutual information between the two variables, net of the amount of mutual information predicted by correlation. This result is not consistent with the idea that noise in the survey data leads us to underestimate the potential non-linearity or heterogeneity in the bivariate relationships. Rather, it shows that noise may lead us to overestimate these possible complexities, and that the true relationships between the variables may thus be even more homogeneous and linear than the already-high correlation $r(\boldsymbol{C}, \hat{\boldsymbol{I}}) = 0.91$ suggests. We therefore conclude that our choice of using correlation as the building block for the networks we examine in this paper is justified.

APPENDIX D. ADDITIONAL CENTRALITY ANALYSES.

As further robustness checks, we used row bootstrapping to examine belief centralities in the 2000 ANES network with some key variables omitted or differently coded. The first column of

Limited Government, Gay Rights, Welfare and Racism scales also leave little room for substantial deviations from linearity.

Table D1 contains the results from 1000 bootstraps of the 39-belief network that excludes ideological identity. These results, like those in the right column of Table 3, show that the network without ideological identity has no clear center. The second column of Table D1 contains the results of 1000 further bootstraps of the 39-belief network that retains ideological identity (Conservative – Liberal) but omits party identity (Democrat – Republican). These results retain much the same centrality distribution as the first set of relative centrality scores we reported in Table 1, with ideological identity as the reliably most central variable. Finally, the last column of Table D1 contains results from 1000 bootstraps of the network that omits all broad identities and moral stances, and retains only the more narrow domain-specific beliefs. Like other analyses that excluded ideological identity, these yielded wide overlapping confidence intervals with no clear center.

[Table D1 about here]

We also investigated whether our results are affected by the ways the key variables were coded. Our ideological identity variable contains 7 levels, while each parenting variable contains only 3. It may thus be possible that the greater observed centrality of ideological identity is due simply to better variable quality. To rule out this possibility, we examined the results of joining the three parenting items into a single variable via summation (which yields a 7-category scale). We also examined the results of reducing the number of levels in the ideological identity variable by collapsing the top two, middle three and bottom two levels of this variable into single levels. This yielded a 3-category ideological identity variable. As Table D2 shows, neither change affected our results.

[Table D2 about here]

APPENDIX E. STRATIFYING VARIABLES.

This appendix describes how we constructed the demographic variables we used in our heterogeneity analyses.

<u>Parents foreign-born</u>: Respondent answered ‘No’ to “Were both of your parents born in this country?”

<u>Class (self-identified)</u>: Coded from branching question: “Most people say they belong either to the middle class or the working class. Do you ever think of yourself as belonging in one of these classes? (IF YES:) Which one? (IF NO:) Well, if you had to make a choice, would you call yourself middle class or working class?”

<u>Black</u>: Coded from the primary racial group self-description variable. Respondents who identified as “Black” or its synonyms were the only ones coded as “Black.” Responses without a clear racial designation (e.g, ‘American’ or ‘None’) were coded as missing. The remaining respondents were coded as non-black.

<u>Hispanic</u>: Respondent coded as “Hispanic” if *(i)* the respondent’s self-identified ethnicity was Mexican, Central American, South American, Cuban, Puerto Rican, or Spanish, or if *(ii)* respondent answered “yes” to the question “Are you of Spanish or Hispanic origin or descent?”

<u>South-Eastern US</u>: Interview location in the “South” region of the US census.

<u>Religion</u>: Coded from the religion/denomination variable. “Catholics” were self-identified. We identified “Mainline Protestants” using the list of Mainline Protestant denominations in Steensland et al. (2000). We coded remaining Protestants as “Protestant (Other)”. Due to small

N's, we merged the remaining religious and non-religious respondents into "Other Religion or none".

Occupational Category: Coded from 14-level occupational category variable. We coded "Executive, Administrative and Managerial" as "Managerial"; "Professional Specialty Occupations" as "Professional"; "Technicians and Related Support Occupations," "Precision Production, Craft and Repair" and "Machine Operators, Assemblers and Inspectors" as "Skilled or Semi-Skilled"; "Sales Occupation" and "Administrative Support" as "Routine Non-manual"; and "Private Household", "Protective Service", "Service except Protective and Household", "Farming, Forestry and Fishing", "Transportation and Material Moving" and "Handlers, Equipment Cleaners, Helpers" as "Unskilled or Farm" (the remaining occupational categories had no respondents.)

Type of place: Coded from response to "Please tell me which category best describes where you were mostly brought up?" We coded "on a farm" and "in the country, not on a farm" as "Rural"; city/town of up to 100,000 residents as "Smaller City"; city of over 100,000 residents as "Larger City"; and suburb of any size of city as "Suburban."

Cross Pressures: Constructed by interacting the three-category income variable with the two-category church attendance variable. Lower-income respondents (under $35,000 a year) were coded as "Pressures Crossed" if they attended church, and "Pressures Aligned" if they did not. Higher-income respondents (over $65,000 a year) were coded as "Pressures Crossed" if they did not attend church, and "Pressures Aligned" if they did. Middle-income respondents ($35,000 to $65,000 a year) were coded as "Neither."

Political information: constructed from an 8-item quiz and two interviewer assessments. The 8-item quiz began with the stem "Now we have a set of questions concerning various public figures. We want to see how much information about them gets out to the public from television, newspapers and the like." The first question was (1) "The first name is TRENT LOTT. What job or political office does he NOW hold?" The next three questions had the same format and asked about (2) "WILLIAM REHNQUIST", (3) "TONY BLAIR", (4) "JANET RENO". The remaining four questions were (5) "What U.S. state does George W. Bush live in now?" (6) "What U.S. state is Al Gore from originally?" (7) "Do you happen to know which party had the most members in the House of Representatives in Washington BEFORE the election (this/last) month?" and (8) "Do you happen to know which party had the most members in the U.S. Senate BEFORE the election (this/last) month?". The quiz score was the count of responses that were identified as correct in ANES data. Interviewers were also asked to subjectively assess each respondent's "general level of information about politics and public affairs" both before and after the November 2000 elections. These responses were on 5-point Likert scales, which we coded as "Very High" (4) to "Very Low" (0), and summed together with the quiz score. Finally, we coded total scores from 0 to 5 points as "Low", 6 to 10 points as "Medium", and 11 to 16 as "High".

The remaining variables (Gender, Number of children, Age group, Education, Income, and Church attendance) were coded from the appropriate ANES variables and are unambiguous.

APPENDIX F. COMPARISON TO BALDASSARRI & GOLDBERG (2014).

We have argued that social groups differ in the extent to which their belief systems are organized, but rarely in the logic which organizes them. Baldassarri and Goldberg (2014; hereafter B&G) instead claim that differences in both the amount and the logic of organization are frequently encountered. In this appendix, we examine some key evidence they present in favor of their argument, and argue that their results support our view of heterogeneity and contradict theirs.

In their theoretical reasoning, B&G note that "it is important to make an analytical distinction between differences that are the result of weak opinion constraint and those that present alterative, internally coherent belief system" (B&G:55). In our terminology, the differences that result from constraint strength are differences in the *amount* of organization. If X and Y are two subgroups that differ in amount of belief organization, and in group X the belief "P is true" strongly implies "Q is true", then in group Y "P is true" should imply "Q is true" more weakly. So, if A and B are a pair of belief variables, and $cor_X(A,B)$ and $cor_Y(A,B)$ are the correlations observed between these beliefs in subgroups X and Y, respectively, then we can express this kind of difference formally as:

$$\text{if } cor_X(A,B) > 0, \text{ then } cor_X(A,B) > cor_Y(A,B) \geq 0; \quad (1)$$
$$\text{if } cor_X(A,B) < 0, \text{ then } cor_X(A,B) < cor_Y(A,B) \leq 0.$$

B&G propose that the difference between *ideologues* and *agnostics* can be accounted for by amount of organization: for ideologues, all belief domains have strong positive mutual associations, while for agnostics, many of these associations are weakly positive or nonexistent. Thus, if we compare pairwise correlations between ideologues and agnostics, we should frequently see relationships of type (1).

We refer to differences that stem from the presence of "alterative, internally coherent" belief systems as differences in the *logic* of belief organization. If there are two internally coherent belief systems, the differences between them should take the form of contrasting entailments: if in population X the belief "A is true" implies "B is true", in population Y it should imply "B is false." Thus, the same pairs of beliefs must often exhibit *opposite* correlations in the two groups, so that there exist a substantial number of beliefs A and B such that:

$$\text{if } cor_X(A,B) > 0, \text{ then } cor_Y(A,B) < 0; \quad (2)$$
$$\text{if } cor_X(A,B) < 0, \text{ then } cor_Y(A,B) > 0.$$

This appears to be the distinction that B&G propose between ideologues and *alternatives*. As B&G note, "if alternatives are inherently different from ideologues and agnostics, we should find that issue domains correlate differently with one another in this group" (B&G:64). The term "correlate differently" is ambiguous. However, in order for the contrast between alternatives and ideologues to be "inherently different" from the contrast between agnostics and ideologues, we reason that the differences in correlations cannot be again due to weak opinion constraint, and should thus be the differences of entailments as in situation (2) above.

Our disagreement with B&G concerns how to interpret the situation where $|cor_X(A,B)| > 0$, and $cor_Y(A,B)$ is not significantly different from zero (i.e., $cor_Y(A,B) \approx 0$.) This situation occurs when "P is true" implies "Q is true" for population X, but "P is true" does not imply anything about Q for population Y. In the typology we offer above, this closely resembles condition (1): "if $cor_X(A,B) > 0$, then $cor_X(A,B) > cor_Y(A,B) \geq 0$." The weaker the overall level of belief constraint in population Y, the more often we will see pairs of beliefs A and B where $cor_Y(A,B) \approx 0$. At the most extreme end, if population Y displays absolutely no belief organization, this situation would apply to all belief pairs. We thus interpret it as difference in the amount of belief organization. On the other hand, B&G interpret such non-significant relationships as a tendency to "dissociate" (62) or "decouple" (66) between pairs of attitudes, which they appear to treat as evidence of belief system difference that is *not* simply due to "weak opinion constraint." Since insignificant correlations correspond to the *weakest* of the weak opinion constraints, this choice appears indefensible to us.

With this in mind, we examine the results B&G report for the 1984, 1986, 1988, 1992, 1994, 1996, 2000 and 2004 ANES in their figure 3 (B&G:63). For each of the survey years, the authors report the average correlation between economic and moral attitudes for members of the ideologue, agnostic and alternative groups. Of these 8 average correlations between economic and moral attitudes presented for ideologues, all 8 are positive and significant. Of the 8 presented for agnostics, 2 are not significant, and the remainder are weaker than for ideologues. This is a clear example of situation (1), and is consistent with their argument that the difference between ideologues and agnostics can be attributed to difference in constraint. We agree with both this interpretation of the data and the broader theoretical claim.

We now turn to the results they present for alternatives. Since the average correlations between economic and moral attitudes for ideologues were positive, the correlations for alternatives would need to frequently be negative to provide evidence for alternate belief systems, as in situation (2) above. However, the results indicate that, for alternatives, 5 correlations are not significant, 2 are weakly positive, and only 1 is negative. Thus only a single ANES year analyzed by B&G actually fits under condition (2)—the same year (2004) they highlight in detail in their paper. For the other 7 years, alternatives exhibit weakly positively constrained or unconstrained attitudes between these domains. As we argued above, this is a difference in the *amount* of belief organization, not in its *logic*. This is consistent with our claims, but not with those made by B&G.

Overall, out of all the 48 average cross-domain correlations between economic, civil, moral and foreign attitudes reported by B&G in their figure 3, all 48 are positive and significant for ideologues. For agnostics, 35 are positive and significant and the rest are insignificant. For alternatives, 26 are positive and significant, 19 are insignificant, and only 3 are negative (all of them weakly.)[9] Thus, for both alternatives and agnostics, beliefs overwhelmingly follow pattern (1), but not pattern (2): if support for A implies support for B for ideologues, it only very rarely implies opposition to B in either of the two groups. In fact, though agnostics are supposedly the group defined by "weak associations among political beliefs" (B&G:60), alternatives appear to actually have *weaker* cross-domain belief associations than agnostics. Across all the years, the "alternatives" identified by B&G thus appear to actually be primarily "agnostics," who generally lack belief constraint as compared to the "ideologues". This is consistent with our argument that groups overwhelmingly vary in the amount of their belief organization rather than in its logic.

As we noted above, B&G provide a different interpretation of these results by treating the alternatives' insignificant correlations between economic and moral attitudes as evidence of an alternate belief system. We think this choice is theoretically unjustified. Moreover, it contradicts how B&G reason about this population elsewhere in the paper. In the abstract, they describe alternatives as "morally conservative but economically liberal or vice versa"—a group which they later propose consists of "free-market supporters [who are] culturally and socially progressive," (B&G:77) and free-market opponents who are culturally and socially conservative. But since the average correlations between economic and moral attitudes for this group are generally insignificant rather than negative, this group likely also contains as many individuals who are both morally and economically conservative, or both morally and economically liberal. And, given the well-documented relationship between low constraint and political moderation,

[9] Given the large number of comparisons, these 3 average correlations may be attributable to noise alone.

it may also consist of those who are simply moderate across all the issues. Contra B&G's description, many of these would not be individuals for whom "selecting one party over the other necessarily entails suppressing one ideological orientation in favor of another" (69). The latter would be the case if and only if their moral and economic attitudes were correlated *negatively*—that is, if the evidence for alternatives having an alternate belief system fulfilled condition (2) we define above.

APPENDIX G. GOODNESS-OF-FIT TEST FOR INDUCTIVELY-DETECTED HETEROGENEITY

In this paper, we have argued that social groups generally vary in the extent of their belief organization, and not in their organizing logic. This contradicts Baldassarri and Goldberg's (2014; hereafter, B&G's) claim that substantial portions of the population use conflicting logics to organize their attitudes. We provided evidence for our claim by comparing belief networks between different demographic groups, and demonstrating that they overwhelmingly arrange their beliefs according to the same logic. We also used an information-theoretic technique to demonstrate that linearity accounts for most of the relationship between variables, thus making heterogeneity unlikely (Appendix C). Additionally, we revisited the ANES time series results reported by B&G, and argued that they are more consistent with our view of heterogeneity than with theirs (Appendix F). In light of this evidence, it may seem surprising that the Relational Class Analysis (RCA) technique used by B&G detected multiple logics of belief organization in the population. In this appendix, we apply RCA to our 2000 ANES data, and use structural equation modeling to examine whether the groups it identified follow different logics of belief organization.

RCA (Goldberg 2011) is a technique for identifying multiple logics of belief organization. Its goal is to locate respondent subgroups where the same pairs of beliefs correlate differently with one another. Thus, in essence, it is a technique for using individual-level data to partition a single sample correlation matrix into multiple subsample correlation matrices that are mixed to generate the overall sample matrix. The method represents the survey dataset as a network, with individuals as nodes and their estimated response pattern similarities as ties, and then uses a modularity maximization algorithm to partition the network into groups. However, recent work demonstrates that modularity maximization can be strongly biased against correctly detecting situations where the network contains no partitions, dividing the population into groups even when a single-group solution would appears intuitively more correct (see discussion in Appendix B). RCA likely inherits this potential for bias, which means that it may indicate the existence of multiple logics of belief organization even when none truly exist. Since RCA provides no goodness-of-fit statistic to describe how distinct the logics it locates truly are, it is difficult to determine whether the partitions it produced are spurious. RCA is thus best thought of as an exploratory technique that detects multiple logics of organization *assuming they exist.*

In order to test whether multiple logics are needed to interpret the ANES data or whether, as we argue, there is only one dominant logic, we use a well-known multiple group testing technique from structural equation modeling (SEM). This technique tests whether using the separate correlation matrices of a proposed partition improves model fit over using one matrix alone. This approach is widely used in psychology to test whether, say, factor loadings in a confirmatory factor analysis are the same for different genders or ethnic groups (Bollen 1989; for a review, see MacCallum and Austin 2000). We hope that this approach will be more widely used in future research to evaluate the partitions suggested by RCA and related techniques.[10]

We used RCA to examine the 2000 ANES items we studied in our primary analyses, which are largely the same items that B&G used in their analyses of the 2000 data. Like B&G (2014:87), we list-wise deleted all respondents with one or more missing responses, leaving N=727 respondents. We then rescaled each of the belief variables to the [0, 1] range. The RCA software (Goldberg and Stein 2016) detected three groups of respondents, as it did in the B&G (2014) analyses. It also left four respondents unassigned to any group, perhaps because their dissimilarity from other respondents left them as isolates in the network. We dropped these isolates from our analyses.

We use SEM to examine the fit of the three-group partition detected by RCA by comparing models in which the full dataset is produced by a single correlation matrix to a model in which the data are produced separately for each RCA-detected group via its own correlation matrix. Of course, whenever any extra parameters are added to a model, the likelihood of the model will almost always increase, and can never decrease. To assess whether this increase in model fit is meaningful, AIC and BIC can be used to examine it in light of the added complexity of the model (Raftery 1995). The single-matrix model has a log-likelihood of -45,836 (1127 df) and the three-matrix model has a log-likelihood of -43,697 (3381 df). Using both AIC and BIC as measures of model fit, the single-matrix model is overwhelmingly preferred (Δ AIC = 235; Δ BIC = 10,578). Consistent with our reasoning, the population homogeneity model thus offers a better description of the data than the model of heterogeneity detected by RCA. Thus even when we apply RCA directly to examine heterogeneity, we find no evidence for it in our data.

[10] It is worth noting here that this confirmatory SEM approach does provide evidence for different logics in other datasets (see Author 1, in preparation).

APPENDIX TABLE D1

BETWEENNESS CENTRALITIES FOR ALTERNATE SELECTIONS OF BELIEFS

ATTITUDE	RELATIVE BETWEENNESS CENTRALITY (MEAN AND 95% CI) No Ideological Id.	No Party Id.	No Broad Stances
Ideological Identity	N/A	100	N/A
Equal Rights 1	89	27	N/A
Gay Rights	78	28	89
Limited Government	59	36	64
Party Identity	49	N/A	N/A
Parenting 2	41	21	55
Biblical Literalism	39	17	N/A
Environmentalism 3	37	14	61
Welfare Spending	35	21	42
Immigration 2	29	13	45
Moral Relativism	26	4	N/A
Immigration 1	26	10	47
Affirmative Action	25	12	34
Immigration 3	24	12	25
Environmentalism 1	22	13	35
Buying Guns	22	6	60
Newer Lifestyles	21	6	7
Abortion Legal	19	12	24
Equal Rights 2	18	9	N/A
Religiosity	16	6	N/A
Military Spending	16	2	22
Foreign Aid	9	4	17
Welfare Recipients	8	5	13
Tolerance	7	2	N/A
Equal Treatment	7	2	N/A
AIDS Spending	7	2	9
Equal Chance	6	4	4
Abortion for Teens	5	1	10
Parenting 1	4	2	6
Separate Spheres	4	0	8
Crime Spending	4	2	6
Equal Opportunity	3	1	0
Surplus Taxes	3	0	3
Environmentalism 2	2	1	6
Military	2	0	1
Death Penalty	1	0	2
Inequality	1	0	0
Parenting 3	0	0	1
Individualism	0	0	N/A
Anti-Black Racism	0	0	0

NOTE.—Each of the three sets of betweenness estimates is based on 1000 row-wise bootstraps.

APPENDIX TABLE D2
BETWEENNESS CENTRALITIES FOR ALTERNATE VARIABLE FORMATS

ATTITUDE	7-CATEGORY PARENTING		3-CATEGORY IDEOLOGY	
	Relative(μ and 95% CI)	Abs.	Relative(μ and 95% CI)	Abs.
Ideological Identity (7)†	100	0.32	N/A	
Ideological Identity (3)	N/A		100	0.34
Limited Government	35	0.11	34	0.11
Equal Rights 1	24	0.08	25	0.08
Gay Rights	21	0.07	22	0.07
Welfare Spending	22	0.07	20	0.07
Parenting 2	N/A		18	0.06
Party Identity	16	0.05	17	0.06
Biblical Literalism	11	0.03	17	0.06
Environmentalism 3	15	0.05	13	0.04
Environmentalism 1	13	0.04	13	0.04
Immigration 2	13	0.04	12	0.04
Affirmative Action	11	0.04	10	0.03
Immigration 3	13	0.04	10	0.03
Immigration 1	8	0.02	9	0.03
Abortion Legal	10	0.03	8	0.03
Equal Rights 2	9	0.03	7	0.02
Religiosity	5	0.01	7	0.02
Buying Guns	5	0.02	5	0.02
Welfare Recipients	5	0.02	4	0.02
Newer Lifestyles	5	0.02	4	0.01
Foreign Aid	4	0.01	3	0.01
Moral Relativism	3	0.01	3	0.01
Equal Chance	3	0.01	3	0.01
Military Spending	2	0.01	2	0.01
Equal Treatment	2	0.01	2	0.01
AIDS Spending	2	0.01	2	0.01
Parenting 1	N/A		1	0.00
Crime Spending	2	0.00	1	0.00
Tolerance	1	0.00	1	0.00
Equal Opportunity	1	0.00	1	0.00
Parenting (7)	1	0.00	N/A	
Abortion for Teens	1	0.00	1	0.00
Environmentalism 2	0	0.00	0	0.00
Death Penalty	0	0.00	0	0.00
Inequality	0	0.00	0	0.00
Surplus Taxes	0	0.00	0	0.00
Military	0	0.00	0	0.00
Separate Spheres	0	0.00	0	0.00
Parenting 3	N/A		0	0.00
Individualism	0	0.00	0	0.00
Anti-Black Racism	0	0.00	0	0.00
	0% 25 50 75 100		0% 25 50 75 100	

NOTE.—Each of the three sets of betweenness estimates is based on 1000 row-wise bootstraps.
†—The parentheticals (7) and (3) indicate the number of categories in the variable.

APPENDIX FIG. A1.—There are three possible relative locations for stranger ties $T(a, b)$ and $T(c, d)$ in a path $\Lambda(f, g)$: *(i)* the two ties could be immediately adjacent, i.e., $b = c$ *(ii)* c could be a descendant of b *(iii)* c could be an ancestor of b. Solid lines indicate ties, whereas dashed lines indicate ancestry. Possible omitted network segments. Note that the proof does not require f and g to be strangers.

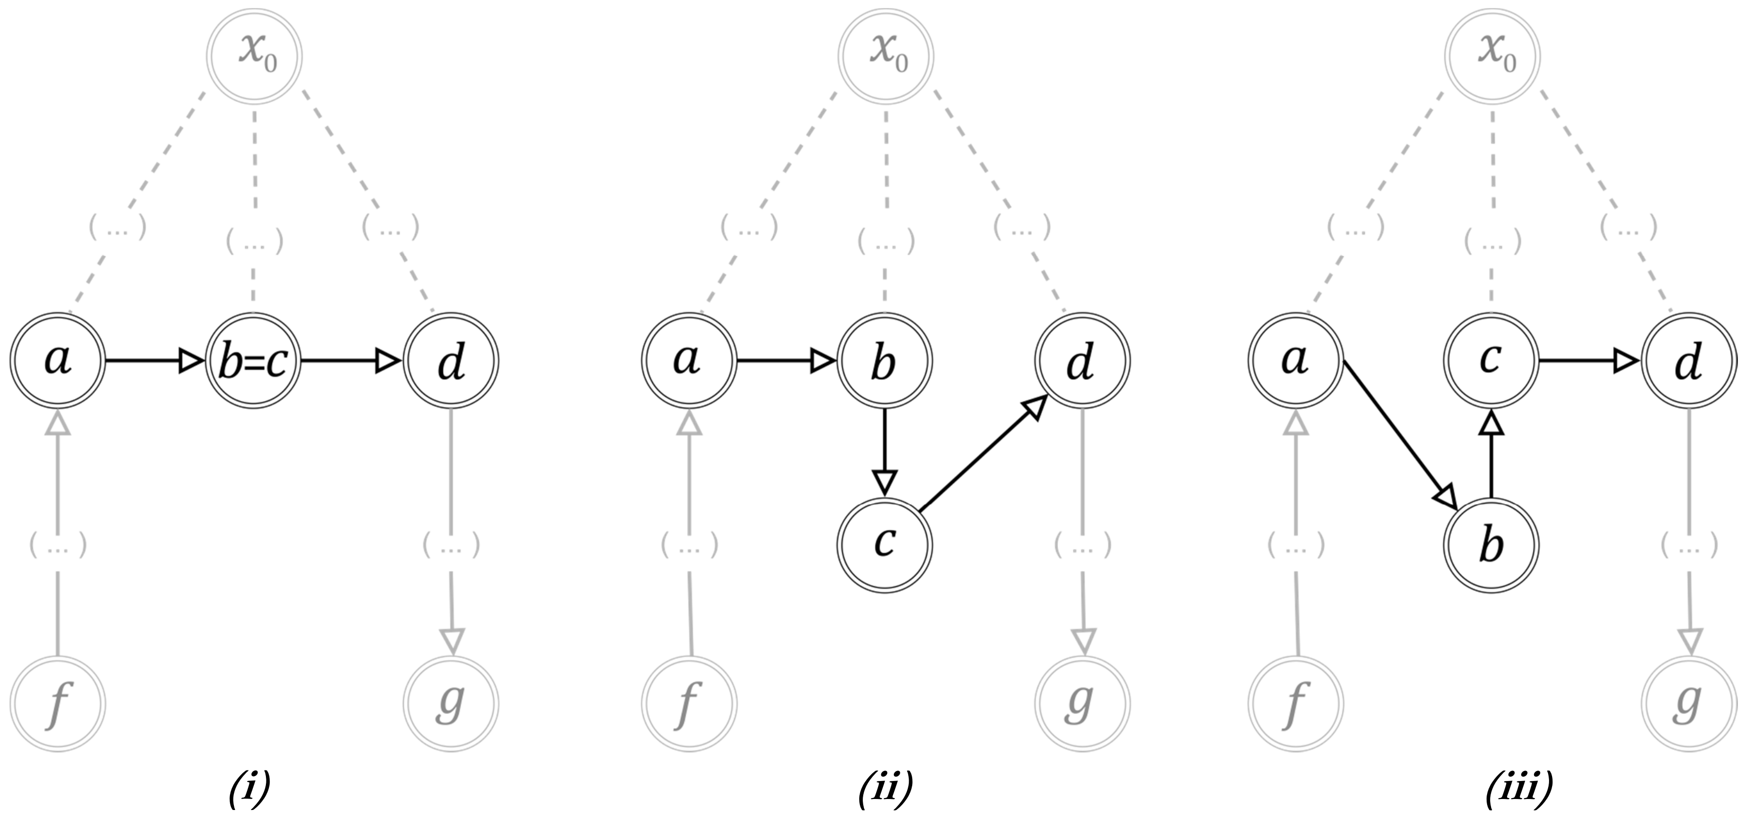

Appendix Fig. A2.—Geodesics between stranger nodes can follow one of two topologies: *(i)* central and *(ii)* non-central.

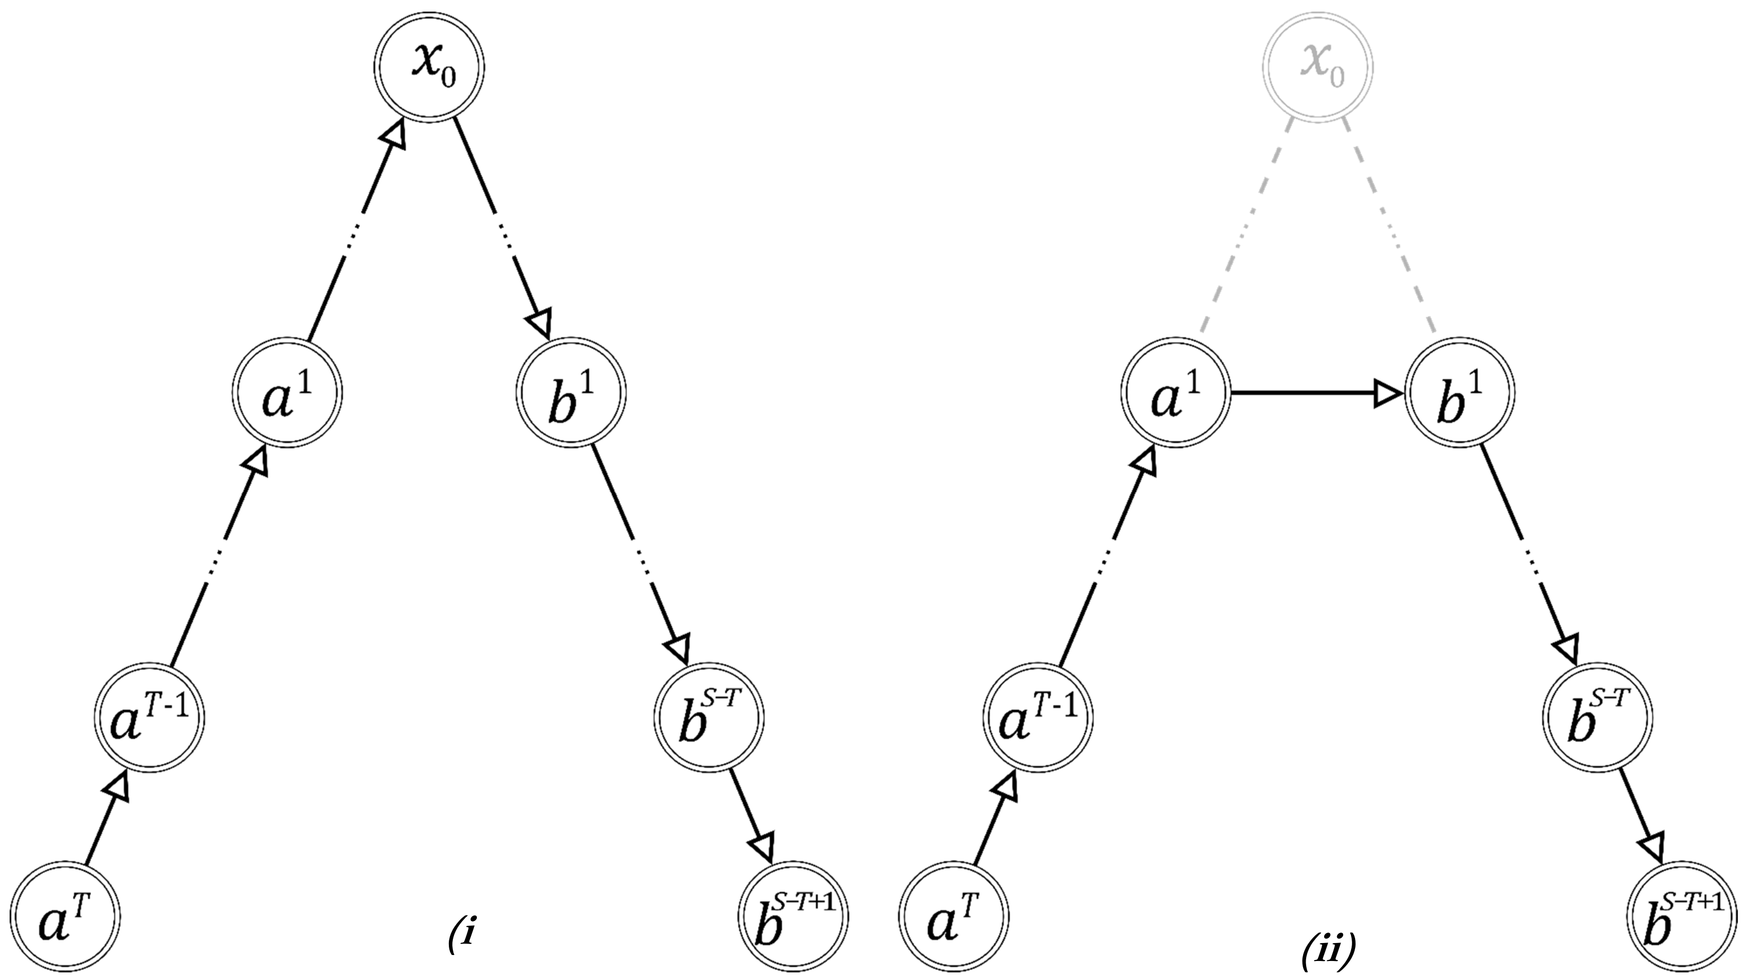

Appendixes for Chapter 3

APPENDIX A. CCA Algorithm

Correlational class analysis can be easily implemented in any programming environment which supports network partitioning by modularity maximization. It consists of four steps:

1. Create a matrix *G* of absolute row correlations between survey respondents.
2. Set statistically insignificant correlations to 0 to reduce noise (e.g., using *t*-tests[11]).
3. Import *G* into a network analysis package, treating it as an adjacency matrix.
4. Use the existing network partitioning routines to produce the class assignments.

In the *R* statistical environment with the *igraph 0.7* library, this can be implemented as:

```
CCA <- function (dataset, min.significant.row.cor = 0.60) {
   C <- abs(cor(t(dataset)))                       # 1st step
   C[C < min.significant.row.cor] <- 0             # 2nd step
   G <- graph.adjacency(C, mode="undirected",
            weighted = TRUE, diag = FALSE)         # 3rd step
   leading.eigenvector.community(G)$membership # 4th step
}
```

A more full-featured implementation of the method is available on *CRAN*, and can be installed in *R* with `install.packages("corclass")`.

See Appendix D for discussion of how to treat respondents with zero variance.

APPENDIX B. Simulating the Theorized Model

This appendix contains the simulation procedures for the first two sets of 5000 simulations I report in the paper.

Procedure 1—Initial linear simulation procedure (fixed inversion probability)

<u>Step 1 (parameters):</u> I first randomly set the maximum ranges of various broad simulation parameters by drawing them from the uniform distribution: schema variance $\upsilon \sim U[0.3,3]$, noise variance $\epsilon \sim U[0,3]$, maximum shift $\delta \sim U\{0,\ldots,3\}$ maximum scaling $\gamma \sim U\{1,\ldots,3\}$, and number of schematic classes $c \sim U\{2,\ldots,6\}$.

<u>Step 2 (schemas):</u> Then, for each of the c classes, I randomly generate a schema vector $\rho = [\rho_1, \ldots \rho_{10}]$ by drawing from the Normal distribution, $\rho_i \sim N(\mu = 0, \sigma^2 = \upsilon)$, and rounding to the nearest integer. Any duplicate vectors are discarded, and new vectors generated in their place, until I have c unique vectors. Then, I randomly set the counts $n_1, \ldots, n_c$ of respondents in each schematic class, $n_i \sim U\{100,101,\ldots,500\}$.

<u>Step 2b (range limits):</u> The range of the 10 taste variables is then limited to $z_i = \pm[\max(|\rho_i|) * \gamma + \delta] \cap \mathbb{Z}$ (this limit is enforced at the end of Step 3).

<u>Step 3 (responses):</u> Finally, for each respondent $f \in [1, n = n_\rho]$ following schema ρ, I generate the 10-element response vector $X_f = (k_f * \rho) + \delta_f + \epsilon_f$ by first drawing the values of the

[11] My exploratory results suggest that more stringent cutoffs may produce more accurate results as long as they are not so extreme as to turn some nodes into isolates. I used $\alpha = 0.05$ as the cutoff for the simulations reported above, and $\alpha = 0.01$ for the GSS analyses. A `min.significant.row.cor` of 0.60 approximates a *t*-test at $\alpha = 0.01$ for rows of 17 variables.

vertical shift $\delta_f \sim U\{-\delta, \dots, \delta\}$ and the scaling and inversion factor $k_f \sim U\{[-\gamma, \dots, -1] \cup [1, \dots, \gamma]\}$. I then generate each respondent's noise vector ϵ_f by first determining *f*'s individual noise variance $\mathrm{E}_f \sim U(0, \epsilon)$, and then drawing each *i*th element $\epsilon_{fi} \sim N(\mu = 0, \sigma^2 = E_f)$, $i \in [1, \dots, 10]$, rounded to the nearest integer. If any $X_{fi} \notin z_i$, where X_{fi} is the *i*th value of X_f, set it to the nearest value in z_i to enforce the range of the variable.

Procedure 2—Full linear simulation procedure (Random inversion probability)
For the full linear simulation procedure, I extends Procedure 1 as follows:

Step 1: I now also draw a random inversion probability: $\zeta \sim U[0, 0.5]$.
Step 3: I now draw a random inversion factor $z_f \in \{1, -1\}$, with $P(z_f = -1) = \zeta$. Since factor k_f now controls the scaling but not the inversion, I now restrict it to positive values: $k_f \sim U[1, \gamma]$. Each respondent f following schema ρ is generated by $X_f = (z_f * k_f * \rho) + \delta_f + \epsilon_f$.

APPENDIX C: Theory-Driven Changes to Pearson's Correlation Coefficient

Consider the absolute value of Pearson's correlation coefficient $|r(X, Y)| =$

$$\left|\frac{Cov(X, Y)}{\sqrt{Var[X]}\,\sqrt{Var[Y]}}\right| =$$

$$\left|\frac{E[(X - \bar{X})(Y - \bar{Y})]}{\sqrt{E[(X - \bar{X})^2]} * \sqrt{E[(Y - \bar{Y})^2]}}\right|$$

Different components of this formula make the coefficient invariant to inversion, scaling, and shift in the vector. Each of the three "sub-linear" scenarios I described earlier can be specifically accommodated by altering the relevant component (and rescaling the result to the $[0,1]$ range if needed). While I leave a fuller methodological treatment of this topic for future work, I derive some basic formulas below as an example of this approach.
No inversion. Most obviously, $|r(X, Y)|$ is invariant to inversion because of the absolute value operator. If inversion is to be interpreted as maximum schematic difference rather than schematic similarity, i.e., $r_{\sim}(X, Y) = 1 \rightarrow r_{\sim}(X, -Y) = 0$, the absolute value operator can simply be removed, with the resulting formula shifted and rescaled to [0,1]:

$$r_{\sim}(X, Y) = 0.5 * \left(\frac{Cov(X, Y)}{\sqrt{Var[X]}\,\sqrt{Var[Y]}} + 1\right)$$

No scaling. To create a version of the coefficient that is sensitive to scaling, it is useful to note that correlation between a variable X and its multiple $Y = kX$ equals 1 because both the numerator and denominator of $r(X, Y)$ scale along with *k*:

$$r(X, kX) = \frac{Cov(X, kX)}{\sqrt{Var[X]}\,\sqrt{Var[kX]}} = \frac{k * Var[X]}{\sqrt{Var[X]}\,\sqrt{k^2 Var[X]}} = \frac{k}{k} = 1$$

It is possible to transform this mechanism into one that penalizes differences in scaling in proportion to the multiplier *k*:

$$r_{\times}(X, Y) = \left|\frac{Cov(X, Y)}{\max(Var[X], Var[Y])}\right|$$

If the variances of X and Y are equal, $r_{\times}(X,Y) = |r(X,Y)|$. However, if $Var[X] > Var[Y]$, $r_{\times}(X,Y) = \left|\frac{Cov(X,Y)}{Var[X]}\right| = |r(X,Y)| * \frac{\sqrt{Var[Y]}}{\sqrt{Var[X]}} = |r(X,Y)| * \sigma_y/\sigma_x$. Thus, as desired, $r_{\times}(X,kX) = |\frac{1}{k}|$ if $|k| > 1$, and $|k|$ otherwise. More broadly, for any $X \neq Y$, $r_{\times}(X,Y)$ will penalize their correlation in proportion to the ratio of their standard deviations.

<u>No shift.</u> Finally, it is possible to modify the correlation coefficient to penalize vertical shifts by replacing the variances in the denominator with the variables' second moments around the grand mean $\bar{M} = 0.5(\bar{X} + \bar{Y})$, yielding

$$r_{+}(X,Y) = \left|\frac{Cov(X,Y)}{\sqrt{E[(X-\bar{M})^2]} * \sqrt{E[(Y-\bar{M})^2]}}\right|$$

Since $Var[X] = E[(X-\bar{X})^2] < E[(X-\bar{M})^2], \forall\ \bar{M} \neq \bar{X} \in R$, this quantity will equal $|r(X,Y)|$ only if the two variables have the same mean, and will otherwise penalize it towards 0.

APPENDIX D: Zero-variance responses

Since correlation is normalized by the product of variances, it is undefined when the variance of either respondent is at absolute zero. The minimal implementation in Appendix A requires that such respondents be dropped from the analysis. Since zero-variance respondents are relatively rare in empirical survey data (e.g., out of the 1532 respondents to the 1993 GSS musical tastes module, there was a total of only 4 with zero variance), dropping them will often be the most pragmatic solution. There are, however, empirical settings where zero-variance respondents can be common, such as when a group of respondents is reliably drawn to one extreme of the scale over the other (e.g., party-line voters on ballots). In such situations, this property of correlation can become a serious limitation necessitating a more considered solution.

The theory of cultural schemas does not appear to suggest any clear way of dealing with zero-variance respondents. Since such respondents express the same attitude towards all musical styles, their tastes literally contain no distinctions between any pair of genres. This lack of cultural judgment means that they are, in a sense, sitting out the game of distinction altogether. Thus, depending on theoretical context, an "undefined" schematic class membership may be justified. On the other hand, this lack of distinctions could be interpreted as a kind of "null schema" that specifies no contrasts between any genres. Following this logic, correlations between two zero-variance respondents could be set to 1, and their correlations with others to 0. (To keep CCA's accuracy directly comparable to RCA's, the algorithm I use to analyze the simulated datasets took this latter approach.)

In other theoretical settings, other solutions to the classification problem may be preferable. For example, in the absence of substantial differences in variance between respondents (e.g., when the survey items feature small numbers of response categories), it may be possible to use covariance instead of correlation, since the covariance between two zero-variance respondents equals zero. However, as long as multiplicative scaling $(Y = kX)$ is one of the theorized schematic transformations, many intuitively appealing solutions may prove conceptually problematic. For example, in some theoretical settings, it may make sense to treat zero-variance responses as extreme cases of a low-variance schema, and thus to assign them to the schematic class containing respondents with the lowest average variance. But variance, unlike correlation, is not invariant to multiplication: if $Y = kX,\ Var[Y] = k^2 * Var[X]$. Therefore, in the presence of multiplicative scaling and a sufficiently broad range of *k*, any non-null schema could produce both low- and high-variance responses. This would then make "amount of variance" a problematic way of distinguishing between schemas.

Appendixes for Chapter 4

Appendix A. Estimating position-level demographic covariates.

Respondents occupying different positions can differ in their average age, education, gender, and other characteristics. It may thus be useful to create position-level demographic variables of the form E[attribute | position]. If the survey question asked respondents whether they support abortion in cases of risk to life of mother, and "Y" is the "Yes" position, then E[age | Y] or E[male | Y] would be the expected age and percentage male of those who support abortion in in cases of risk to life of mother.

These kinds of position-level variables can be estimated by repeated application of Bayes rule. To demonstrate how this can be done, I use expected sex of a respondent in the pro-choice position E[male|Y] as an example. Since "male" is a binary variable, $E[male|Y] = 1 \times P(male|Y)$. By Bayes rule,

$$P(male|Y) = \frac{P(Y|male)P(male)}{P(Y)} \tag{A1}$$

The probability P(Y) is the chance that a respondent's latent position equals Yes, which can be estimated via the position-holding model. P(male) can be estimated directly from the observed frequencies. The remaining probability $P(Y|male)$ is the chance that any male respondent's latent position equals Yes. This can be estimated by averaging the estimated probabilities of a Yes position across all male respondents, using the observed response ω_i of each respondent *i* to estimate the probability $\tau_i = Y$:

$$P(Y|male) = \frac{\sum_i P(\tau_i = Y|\omega_i) \times \mathrm{m}_i}{\sum_i \mathrm{m}_i} \tag{A2}$$

Here, $\mathrm{m}_i = 1$ if *i* is male, and 0 otherwise.

The values of $P(\tau_i = Y|\omega_i)$ can be calculated by again applying the Bayes rule and expanding the equation, which differs slightly based on the different values of ω_i. For example, if respondent *j* has sequence $\omega_j = yyn$, the probability that *j*'s position $\tau_j = Y$ is given by:

$$P\left(\tau_j = Y \middle| \omega_j = yyn\right) = \frac{P(yyn|Y)P(Y)}{P(yyn)} = \frac{P(y|Y)^2 P(n|Y)P(Y)}{P(y)^2 P(n)} \tag{$A3$}$$

For respondent *k* whose sequence $\omega_k = yyy$, this probability instead equals:

$$P(\tau_k = Y|\omega_k = yyy) = \frac{P(yyy|Y)P(Y)}{P(yyy)} = \frac{P(y|Y)^3 P(Y)}{P(y)^3} \tag{$A4$}$$

The $P(\tau_i = Y|\omega_i)$ equations for the remaining response sequences follow analogous forms.

I have already estimated P(Y), P(y|Y), P(n|Y) via the position-holding model, while P(y) and P(n) can be estimated directly from observed frequencies. Substituting these probabilities into (A3) yields the probability that the true position of respondent *j*, whose response sequence equals *yyn*, is Y. Substituting them into (A4) yields this probability for respondent *k* with $\omega_k = yyy$. I repeat the same procedure for each respondent *i*. I then substitute all the resulting values of $P(\tau_i = Y|\omega_i)$ into equation (A2) to get $P(Y|male)$, which can finally be substituted into (A1) to calculate the desired quantity $P(male|Y)$. This yields the expected proportion of male respondents among supporters of the pro-life position in cases of risk to life of mother.

This procedure is straightforward to generalize to integer- or real-valued demographic variables by simply summing over all the unique observed values of the variable. For example, the expected age of respondents in the pro-choice position Y is given by

$$E[age \mid Y] = \sum_i age_i \times P(age = age_i \mid Y).$$

Each element in this summation can be estimated via the same steps I used to estimate (A1).