

Three Papers on Culture, Time, and Attitudes

by

Kevin Kiley

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Stephen Vaisey, Advisor

\_\_\_\_\_  
Craig Rawlings

\_\_\_\_\_  
Kieran Healy

\_\_\_\_\_  
Christopher Johnston

Dissertation submitted in partial fulfillment of  
the requirements for the degree of Doctor  
of Philosophy in the Department of  
Sociology in the Graduate School  
of Duke University

2021

ABSTRACT

Three Papers on Culture, Time, and Attitudes

by

Kevin Kiley

Department of Sociology  
Duke University

Date: \_\_\_\_\_

Approved:

\_\_\_\_\_  
Stephen Vaisey, Advisor

\_\_\_\_\_  
Craig Rawlings

\_\_\_\_\_  
Kieran Healy

\_\_\_\_\_  
Christopher Johnston

An abstract of a dissertation submitted in partial  
fulfillment of the requirements for the degree  
of Doctor of Philosophy in the Department of  
Sociology in the Graduate School of  
Duke University

2021

Copyright by  
Kevin Robert Kiley  
2021

## **Abstract**

This dissertation uses the lens of cultural sociology to understand variance in people's attitude reports over time. Across three studies, I use a variety of panel surveys and statistical approaches to understand how and why people change their attitudes and to adjudicate theories of culture's role in shaping attitudes. Study 1 uses data from the 2006 to 2014 rotating panels of the General Social Survey to adjudicate between a settled dispositions model, in which changes in attitudes are temporary and people return to a settled baseline, and an active updating model, where changes persist. Study 2 explores heterogeneity within the settled dispositions group, asking whether people's attitude reports should be thought of as temporary constructs drawn from stored considerations or whether they represent durable opinions. It quantifies the prevalence of these opinion behaviors for 544 items from 10 panel data sets. Study 3 seeks to predict variance in attitude responses over time. Using data on religious, moral, and family structure beliefs in the National Study of Youth and Religion, I use Latent Class Analysis to deduce a set of constraints that should shape people's response patterns over time. I test these constraints on people's subsequent attitude reports. Taken together, results of these studies suggest 1) people's attitudes are stable, on average, over the long term; 2) this average stability often masks high levels of instability in the short term, though some proportion of the population is stable on any issue; and 3) both this stability and

instability are somewhat predictable based on a person's pattern of beliefs. Findings suggest a model of culture where people internalize a diverse set of considerations when they are young and are shaped, in the short-term, by environmental influences. But durable cognitive structures, likely formed when people are young, limit the power of changing social circumstances to induce durable change.

## **Dedication**

For Adam, the best idea I had in graduate school.

# Contents

Abstract.....	iv
List of Tables.....	xii
List of Figures.....	xiii
Acknowledgements.....	xv
1. Introduction.....	1
2. Study 1: Measuring Stability and Change in Personal Culture Using Panel Data.....	7
2.1 Introduction.....	7
2.2 Theoretical Background.....	11
2.2.1 Belief Formation.....	11
2.2.2 Active Updating Model (AUM).....	12
2.2.3 Settled Dispositions Model (SDM).....	15
2.2.4 Consistency in Response.....	18
2.2.5 Toward Theoretical Synthesis.....	19
2.3 Research Questions and Expectations.....	24
2.4 Analytical Strategy.....	25
2.4.1 Statistical Models.....	25
2.4.2 Combined Model.....	26
2.4.3 Comparison to Other Approaches.....	28
2.4.4 Limitations of the Method.....	31
2.4.5 Analysis Steps.....	33

2.4.6 Item Selection .....	35
2.5 Results .....	36
2.5.1 Evidence for Active Updating .....	37
2.5.2 Age Heterogeneity .....	40
2.5.3 Item Heterogeneity.....	43
2.5.4 Political Beliefs .....	51
2.6 Discussion and Conclusion .....	54
2.6.1 Implications for Cultural Sociology.....	57
2.6.2 Implications for Political Sociology .....	59
2.6.3 Methodological Implications .....	62
3. Study 2: Ambivalence Is Everywhere: Quantifying Opinion Behavior Across Topic Domains.....	66
3.1 Introduction.....	66
3.2 Theoretical Framework.....	70
3.2.1 Stable Opinions and ‘Non-Attitudes’ .....	70
3.2.1 Some Do, Some Do Not .....	74
3.3 Expectations and Hypotheses.....	77
3.4 Data.....	79
3.5 Methods .....	82
3.5.1 Measuring Opinion Behavior .....	82
3.5.2 Opinion Behaviors.....	84
3.5.3 Mixture Model .....	88

3.5.4 Missing Data .....	92
3.5.5 Model Modifications .....	93
3.5.6 Comparison with Measurement Error Model.....	95
3.5.7 Testing Issue Publics .....	95
3.6 Results .....	96
3.6.1 Gay Marriage .....	96
3.6.2 Full Data Results.....	100
3.6.3 Issue Publics .....	108
3.7 Discussion.....	111
3.7.1 Implications for Cultural Sociology.....	113
3.7.2 Implications for Measuring Attitudes .....	116
3.7.3 General Conclusions .....	117
4. Study 3: Predictably Unpredictable: The Dynamic Constraint of Cultural Belief Systems .....	119
4.1 Introduction.....	119
4.2 Theoretical Framework.....	123
4.2.1 The ‘Problem’ of Instability.....	123
4.2.2 Culture, Cognition, and Survey Response.....	127
4.2.3 Consequences for Measurement .....	129
4.2.4 Latent Class Analysis and Belief Systems.....	133
4.2.5 Culture and Social Structures .....	134
4.3 Hypotheses .....	137

4.3.1 Latent Class Analysis as a Measure of Belief Systems .....	137
4.3.2 Adjudicating Culture and Structure .....	138
4.4 Data and Measures .....	140
4.4.1 The National Study of Youth and Religion .....	140
4.4.2 Measures .....	142
4.4.3 Belief Systems.....	146
4.4.4 Change Over Time .....	147
4.4.5 Pattern Prediction.....	149
4.4.6 Changing Social Structures .....	153
4.5 Results .....	155
4.5.1 Belief Systems.....	155
4.5.2 Change over Time .....	161
4.5.3 Predicting Response Patterns .....	165
4.5.4 Changing Social Structures .....	169
4.6 Discussion.....	171
4.6.1 Implications for Cultural Sociology.....	173
4.6.2 Implications for Methods .....	175
4.6.2 General Conclusions .....	176
5. Conclusions.....	177
Appendix A: Chapter 2 Phi Estimates .....	182
Appendix B: Chapter 3 Finite Mixture Model Estimation .....	186
B.1 Re-Expressing the Data .....	186

B.2 Drawing Group Indicators Given Parameters.....	187
B.3 Drawing Parameters Given Group Indicators.....	189
B.4 Convergence .....	191
Appendix C: Chapter 3 Regression Model.....	192
C.1 Vacillation.....	194
C.2 Durable Change .....	196
References.....	199
Biography .....	208

## List of Tables

Table 1: Finite mixture model parameter estimates for gay marriage question (2006 to 2014 GSS). .....	97
Table 2: Average within-group standard deviation, by opinion behavior group.....	99
Table 3: Belief measures, by wave. ....	143
Table 4: Expected response pattern probability for two hypothetical belief systems.....	150
Table 5: Covariates of class assignment.....	159
Table 6: Results of fixed effects linear regressions of within-person, within-question standard deviation and mean on within-group standard deviation and mean. ....	164
Table 7: Summary statistics for covariates of belief prediction, times 1 and 2.....	169

## List of Figures

Figure 1: Graphical representation of active updating and settled dispositions models over time.....	11
Figure 2: Distribution of phi estimates and probabilities that items show evidence of active updating.....	38
Figure 3: Comparison of phi estimate for individuals over and equal to or less than 30 years old, with 95 percent confidence intervals. ....	41
Figure 4: Distribution of beta and phi estimates for GSS items by whether model prefers phi = 0.5 constraint. ....	45
Figure 5: Summary of phi estimates for all items, by topical group. Numbers in parentheses indicate the number of items in each topical group. ....	48
Figure 6: Phi (updating) and beta (consistency) estimates for 43 political attitudes and identities. ....	52
Figure 7: Mean proportion of respondents in each behavioral group, by question and data set. ....	100
Figure 8: Mean estimated proportion of "strong" responses in option holder and vacillating changer behavior groups, by question. ....	103
Figure 9: Average within-person standard deviation, by behavioral group and question response options.....	105
Figure 10: Network graph of pairwise correlations of co-occurring stability greater than 0.2.....	109
Figure 11: Over-time response probabilities for hypothetical belief systems. ....	131
Figure 12: Religious, Moral, and Family Beliefs by Latent Class. ....	156
Figure 13: Scatterplot of within-group standard deviation at time 1 against average within-person standard deviation, times 1 through 3. ....	162
Figure 14: Boxplots of deviation of observed from expected counts of response patterns, 10,000 iterations each, religious beliefs. ....	166

Figure 15: Boxplots of deviation of observed from expected counts of response patterns, 10,000 iterations each, moral and family beliefs .....	167
Figure 16: Phi (updating) estimates for items about religious activity and beliefs, social life, subjective socioeconomic status, and suicide. ....	182
Figure 17: Phi (updating) estimates for items about guns, law, crime, policing, government, and public spending.....	183
Figure 18: Phi (updating) estimates for items about civil liberties, confidence in institutional leadership, health, morale, and trust. ....	184
Figure 19: Phi (updating) estimates for items about race, gender, sex, sexuality, and abortion.....	185
Figure 20: Coefficient estimates of regression of proportion of respondents giving a vacillating response on question content and structure indicators. ....	195
Figure 21: Coefficient estimates of regression of proportion of respondents giving a durable change response pattern on content and structure indicators.....	197

## Acknowledgements

This dissertation would not exist without the support and guidance of Stephen Vaisey. As a dissertation chair, collaborator, statistics and culture teacher, director of graduate studies, lab leader, and friend, Steve shaped both how I asked questions and the kinds of questions I found interesting. The ideas presented in this dissertation grew out of an extended collaboration with Steve that I hope continues well into the future. I will be eternally grateful that he asked me to think about measuring culture over time. Steve's enthusiasm, support, feedback, suggestions, and questions touch every part of this dissertation.

In defiance of what this dissertation argues, my academic interests and work changed substantially in graduate school through the addition of two people to my social network. When my department hired Craig Rawlings in my third year, I found another person who asked the kinds of questions that made sociology interesting to me in the first place. Craig's feedback helped the third chapter cohere from a wandering mess to a real sociological paper. Chris Johnston's suggestions to connect chapter 2 to public opinion work shifted the direction of my own work, and his offhand comment to look into a political science paper gave me the tool I need to answer the question underlying chapter 3. Chris and Craig's feedback on a very early draft of chapter 4 pushed that paper in a much better direction. Kieran Healy did not at all care that my interests vacillated widely.

I owe a significant debt of gratitude to Jennifer Hill, who I have never met, but whose work on the finite mixture model of opinion behavior inspired chapter 3. The papers that inspired that chapter, and her own dissertation work, are not nearly as influential as they should be. The anonymous reviewers at the *American Sociological Review* who evaluated the paper that constitutes chapter 2, in particular Reviewer 2, made thoughtful, constructive comments that dramatically improved that paper and inspired other portions of this dissertation.

I am also deeply indebted to my fellow graduate students at Duke University. Danielle Montagne, as a friend and collaborator on a project that ultimately went nowhere, helped rekindle an interest in sociology that waned at significant moments in graduate school. She would make a great sociologist, and I wish she still wanted to be one. Nicolas Restrepo Ochoa asks great questions. He provided great feedback on much of the work in this dissertation. I am eager to read everything he produces. Lauren Valentino, Mary Beth Fallin Hunzaker, and Joshua Doyle, who I view as intellectual older siblings, were inspirations throughout graduate school, and their own research motivates much of what is in here. Colin Birkhead, Nicholas Bloom, Joshua Bruce, Christina Kamis, and Jessie West were all sources of social and emotional support at different points graduate school.

I had the very rare opportunity of completing a Ph.D. only 30 minutes from where my family lives. My parents, Edward and Patricia Kiley, have been a source of

unwavering emotional support and love even as they long ago stopped understanding what I was doing or why I was still in school. They structured my cognition early in life and are the principal reason I love learning. My brother, Brian Kiley, and his partner Cassie Irwin, have been true friends over the years. My mother-in-law, Kathe Baer, has been an endless source of support and joy.

Most importantly, I could not have accomplished this dissertation without the support of my wife and best friend, Annie Baer Kiley. Every idea in this dissertation was discussed with her and is better as a result. Her consistent belief in my ability to have good ideas, even when I doubted it, has been invaluable during my Ph.D. experience. In the past year, as the pandemic confined her to spending almost every minute of the day in the same room as someone writing a dissertation, she has helped temper the wild emotional swings of someone engaged in intellectual work. She is the smartest person I know, and the best partner a person could have.

# 1. Introduction

In his seminal paper on belief systems in the general public, Philip Converse concluded that “large portions of an electorate ... simply do not have meaningful beliefs, even on issues that have formed the basis for intense political controversy among elites for substantial periods of time” (Converse 1964: 245). This finding spawned decades of debate in public opinion research about the nature of stability and change in beliefs and attitudes, with different theoretical traditions highlighting or downplaying the importance of Converse’s findings (Achen 1975; Alwin 2007; Ansolabehere, Rodden, and Snyder 2008; Zaller 1992). But the central empirical fact of Converse’s study is undeniable: on average, people are not consistent in their attitude reports over time.

How do we make sense of this variation over time? The most common interpretation of this variation is that it amounts to measurement error, deviations attributable to “vague question wordings, vague response categories or categories that do not reflect the individual’s actual attitude, inattentiveness on the part of the respondent, and even typographical errors” (Ansolabehere, Rodden, and Snyder 2008: 216). Under this approach, people have a “true score” that can be deduced when we sift through their responses over time or average their responses to similar questions (Alwin 2007; Lord & Novick 1968). Other approaches ascribe more meaning to this change, suggesting that it reflects a real change of attitudes over time as people move across

social spaces, encountering new challenges that force them to rethink their dispositions (Gross 2009).

This dissertation brings the lens of sociology, in particular cultural sociology, to bear on the question of attitude change and stability over time. A central finding of cultural sociology in the last half-century is that rather than internalize a consistent cultural worldview through a single socialization process, people are exposed to and internalize a diverse, contradictory cultural repertoire of beliefs, practices, and assessments (DiMaggio 1997; Swidler 1986). A long line of research documents this cultural contradiction across domains. In America, love is both a choice entered into freely and a special commitment that transcends people's desires (Swidler 2001). American culture is individualist and collectivist; managerial and therapeutic; biblical and republican (Bellah et al. 1985). Morality is fixed and relative (Baker 2004), orthodox and progressive (Hunter 2000).

It is not just that culture is diverse and contradictory at the public level and that concepts compete in the public sphere while people maintain distinct worldviews. This contradiction is apparent in personal culture, the declarative and non-declarative attitudes, worldviews, values, and dispositions that manifest at the individual level (Lizardo 2017). Because of their cognitive limitations, people consume a broad array of cultural information and cannot or tend not to engage in the effort to reconcile these contradictions (Martin 2010; Zaller 1992). As a result, "our heads are full of images,

opinions, and information, untagged as to truth value, to which we are inclined to attribute accuracy and plausibility” (DiMaggio 1997: 267).

Given the heterogeneity both around and inside of people’s heads, it is not surprising that people contradict themselves in interviews (Swidler 2001) and vacillate widely in their responses to the same attitude questions over time (Hout and Hastings 2016; Zaller 1992). At the same time, some people display remarkable consistency over time in their beliefs and behaviors, and cultural beliefs, even in their contradictions, shape behavior, suggesting they are not just noise (Harding 2007; Miles 2015; Vaisey 2009).

The three papers that constitute the remainder of this dissertation seek to make sense of these contradictions.

Chapter 2 adjudicates two models of individual change that underlie theories of population-level cultural change. The first, an active updating model, emphasizes the role of changing discourses, environments, and interactions on attitude formation. At the individual level, this model is associated with pragmatist models of change and connects to social influence and diffusion models of population-level change (Gross 2009; Joas 1996). The second, a settled dispositions model, emphasizes the durable influence of dispositions formed early in life. This model has connections to cohort replacement theories of population-level change and Bourdieu’s concept of the habitus at the individual level (Bourdieu 1990; Ryder 1965). I formalize and test these competing

models on 183 attitude and behavior items from the 2006 to 2014 rotating panels of the General Social Survey. The pattern of results is complex but more consistent with the settled dispositions model than with the active updating model. Most of the observed change in the GSS appears to be short-term attitude change or measurement error rather than persisting changes. When persistent change occurs, it is somewhat more likely to occur in younger people and for public behaviors and beliefs about high-profile issues than for private attitudes.

Chapter 3 addresses a key debate elided in Chapter 2. The settled dispositions model is broadly consistent with two very different models of opinion formation. The first, what some researchers call “ambivalence” and others call a “toolkit” model of culture, posits that beliefs and attitudes are temporary constructs with no durability and that people construct new dispositions each time they are asked from considerations in their heads (Zaller 1992; Swidler 2001). The second suggests that people, constrained by underlying cultural dispositions or social structures, have real, durable opinions (Vaisey 2009). I formalize these opinion behaviors and estimate their prevalence across more than 500 attitude questions in 10 panel data sets. The results are broadly consistent with Zaller’s model of “ambivalence:” people internalize heterogeneous considerations and struggle to maintain consistent attitudes over time as social influences push them in different directions. At the same time, on any question, some people demonstrate consistent attitudes. And this stability is patterned with people showing stability on

related topics. The pattern of results is consistent with an “issue publics” model and models in cultural sociology that root stability in institutions, social structures, and durable cultural beliefs (Lizardo and Strand 2010; Martin 2010).

Chapter 4 attempts to make sense of the dual finding that people have stable dispositions over the long term while being highly inconsistent in the short term. Drawing on models of culture’s role in shaping people’s cognition (DiMaggio 1997; Strauss and Quinn 1998), I posit that heterogeneity of responses over time is principally a function of the competing considerations people internalize while young. Because social influences shift people’s responses in each survey wave more or less at random, and because these influences are temporary, instability should be predictable. I use Latent Class Analysis to deduce five belief systems regarding religious, moral, and family structure beliefs in the National Study of Youth and Religion’s second wave. These systems suggest distinct constraints on people’s beliefs over time. The constraints evident across people in these belief systems at one time strongly predict people’s attitude behavior over time. A model that roots this variation in these belief systems outperforms a model that roots changing attitudes in changing social structures in predicting attitudes over time.

Taken together, the papers presented in chapters 2 through 4 suggest 1) that people’s attitudes are stable, on average, over the long term; 2) that this average stability often masks high levels of short-term instability, though some proportion of the

population is remarkably stable on any issue; and 3) that both this stability and instability are socially patterned and somewhat predictable based on people's patterns of beliefs. Combined, the findings suggest a model of culture where people internalize a diverse set of considerations when they are young and are shaped in the short-term by environmental influences. But the constraints and heterogeneity inherent in their belief systems, likely formed when people are young, limits the ability of social change to facilitate durable changes in attitudes. Chapter 5 discusses these conclusions and their implications for our model of how culture works in people in more depth.

## **2. Study 1: Measuring Stability and Change in Personal Culture Using Panel Data**

### ***2.1 Introduction***

Culture is an important part of social life, but cultures are continuously evolving. In 1972, for example, over 40 percent of US adults supported a law outlawing interracial marriage. Three decades later, this opinion had become so uncommon that the question was removed from the U.S. General Social Survey. How does this kind of cultural change happen?

Attempts to account for opinion changes in society have produced conflicting theories about the process of opinion formation and individuals' ability to maintain consistent attitudes. Some models suggest that people lack cognitive tools to maintain consistent beliefs on social and political issues. As a result, people construct responses on the fly in interview settings, drawing on ideas from opinion leaders and changing their attitudes as elite discourse changes (Converse 1964; Perrin and McFarland 2011; Zaller 1992). Cohort replacement theories, in contrast, posit that people do hold opinions and are unwilling to alter them in the face of societal change; public opinion thus changes only with generational turnover (Mannheim 1952; Ryder 1965). Another set of models claim that people hold "a number of real, stable, and sensible opinions about public policy," and they change their opinions in response to new information (Page and Shapiro 1992: xi; see also Achen 1992). More recently, sociologists have suggested that

people attempt to align malleable peripheral beliefs with relatively fixed “core” beliefs using social cues (Boutyline and Vaisey 2017; Goldberg and Stein 2018; Lakoff 1996).

Despite their diversity, models of aggregate attitude and behavior change often implicitly invoke one of two broad models of individual change with strong connections to theories in cultural sociology. The first, an *active updating* model, emphasizes the role of changing discourses, environments, and interactions on attitude formation. This model is rooted in a broadly pragmatist approach to action, which claims that contemporary social environments and problems provoke people to adapt their views and make new meaning (e.g., Joas 1996, Gross 2009, Swidler 2001). The second is a *settled dispositions* model, which emphasizes the continuing influence of durable dispositions acquired early in life. This model has affinities with the Bourdieusian tradition, which deemphasizes (but does not ignore) the current environment in favor of the “past conditions of production” (e.g., Bourdieu 1990; Vaisey and Lizardo 2016). These two models represent different approaches to understanding how people come to hold diverse forms of *personal culture*, or the declarative and nondeclarative attitudes, worldviews, values, dispositions, and associations that manifest at the individual level (Lizardo 2017).

In this paper, we make these two models of personal culture change explicit, deduce some of their empirical implications, and derive a statistical model for estimating the prevalence of active updating using panel data. Previous models that

tend to assume one of these data-generating processes to measure stability and reliability of estimates over time. Our approach, in contrast, separates persisting change from non-persisting change to estimate whether people make persistent changes in their attitudes and behaviors. We then apply this method to 183 items from the 2006 to 2014 General Social Surveys (GSS). By classifying the pattern of change in personal culture observed across GSS items, we can clarify when different accounts of aggregate change are more likely to apply. Due to data limitations, we cannot speak to all types of cultural objects (e.g., music styles, baby names). We do, however, investigate a wide variety of opinions, including views on politics, free speech, race, and gender roles, and practices including socializing at bars or attending church, that are important in contemporary U.S. society.

Our analysis yields several results. First, we find that the majority of what appears to be individual-level change in attitudes or practice probably reflects short-term (i.e., non-persisting) change or measurement error rather than actual persistent change. Simply put, there is little evidence that large numbers of U.S. adults changed their beliefs or practices in lasting ways over this period. Second, settled opinions vary in how consistently individuals report the same answer. Consistent with theories arguing that people lack clear opinions, some survey items appear to elicit inconsistent or random responses. At the same time, people are consistent on both high- and low-profile public policy items, suggesting a greater degree of “real” attitudes than these

theories suggest. Third, the persistent change that we *do* see in the data is somewhat more concentrated among younger respondents. On several items, it appears that younger adults are still in the process of acquiring dispositions and habits they will take into later life. Fourth, we find that changes in social behavior (e.g., church attendance, political party membership, socializing) are more likely to persist than changes in private attitudes (e.g., political ideology), and people are more likely to report these attitudes and behaviors consistently. This suggests that interactional and institutional mechanisms may provide stronger support for lasting change than do pressures for intrapsychic consistency.

Our findings offer broad support for theories claiming that cultural change comes through generational turnover rather than contemporaneous persuasion and social influence. However, a pattern of exceptions and caveats that can help us understand how institutions and events shape the process of cultural change; these patterns challenge the idea that change in all attitudes follows a similar trajectory over the life course. The pattern of results also supports models of attitude change that put ideological identification at the center of a network of political beliefs and suggest that individuals are more likely to make lasting changes in their partisan identification than to their general political beliefs. We argue that there is a place for both the active updating and settled dispositions models in accounting for cultural change, but we need more research on the circumstances under which each is more likely to apply.

Active Updating Model

Settled Dispositions Model




Figure 1: Graphical representation of active updating and settled dispositions models over time.

## 2.2 Theoretical Background

### 2.2.1 Belief Formation

How does cultural change happen at the individual level? Let us make the question concrete by imagining a person who answers the same question each year for several years. The question could be anything, but assume it is this GSS question: “Please tell me whether you strongly agree, agree, disagree, or strongly disagree [with this statement]: ‘a working mother can establish just as warm and secure a relationship with her children as a mother who does not work.’”

How does the respondent formulate a response to that question each time, year after year? To make things as explicit as possible, we can write the data-generating process formally. Although this presentation may make it seem like we are assuming rationality or conscious deliberation, this way of writing the models makes no particular cognitive assumptions. We will explain this in greater detail below. For now, consider the following two simple models:

$$y_{it} = y_{it-1} + v_{it} \quad (1)$$

$$y_{it} = U_i + v_{it} \quad (2)$$

These models may seem similar at first glance, but they have different implications for the pattern of individuals' responses we would observe over time. Equation (1) represents an *active updating* model and Equation (2) represents a *settled dispositions* model, as we explain below. Figure 1 shows these models and helps highlight their differences. In the next two sections, we consider each of these models and briefly discuss their links to influential sociological theories.

### **2.2.2 Active Updating Model (AUM)**

Equation (1) represents the active updating model. Respondents form their answers by starting with what they said last time ( $y_{it-1}$ ) and then incorporating any new considerations ( $v_{it}$ ). There is no need to remember responses from earlier time points (e.g.,  $y_{it-2}$ ) because this information gets folded into the updated response each time. Formally, Equation (1) is a Markov model, where future states depend entirely on the

current state. This formal property is often assumed to underlie the data-generating process in studies of change in and reliability of repeated survey measures (Alwin 2007; Krosnick and Alwin 1989). In this framework, we use the term “updating” to refer to the change to a new baseline for whatever reason, not necessarily requiring the incorporation of new or better information.

More informally, this AUM posits people who are updating their views in the face of social experience. There are formal Bayesian ways of modeling updating, and this model underlies theories of rational updating in the face of new information (Achen 1992; Bartels and Jackman 2014; Page and Shapiro 1992), but we need not rely on any assumptions of rationality, optimality, or conscious thought for this basic process to apply. Following Gross’s (2009: 367) pragmatist account, we could instead regard this model as consistent with an “active and creative relation to the world” that “lead[s] actors to see themselves in new ways, to value different kinds of goods, and to become attached to problem solutions that they could not have imagined previously.”

Active updating lies at the heart of most theories suggesting contemporaneous social environments influence individual attitudes and behaviors. When sociologists posit that adults adopt the beliefs and behaviors of their friends (Centola 2018; Christakis and Fowler 2010; DellaPosta, Shi, and Macy 2015), or that individuals adopt behaviors that they view as consistent with their cultural dispositions (Goldberg and Stein 2018), they invoke this model. When social scientists argue people adapt their

partisan affiliations to conform to their social groupings, change their ideological commitments to conform to their partisan identifications (Green, Palmquist, and Schickler 2002), change their partisan affiliations to match their ideological commitments (Abramowitz and Saunders 2006; Baldassarri and Gelman 2008), or respond to exposure to diverse opinions by becoming more extreme in their viewpoints (Bail et al. 2018), the claim is rarely that these changes are temporary shocks that will revert to a stable baseline over time. Instead, the claim is that the changing social environment changes dispositions, which will then remain stable until the environment changes again.

Regardless of whether the exact process is rational, heuristic, discursive, embodied, or any combination of these, the key notion is that people are continually influenced by the contemporaneous (social) environment in ways that might lead to long-term, persistent change in beliefs, practices, and identities (DeGloma 2014; Gross 2009). As the person encounters new considerations throughout her life (e.g., adding working mothers to one's social network, seeing changing media representations of working mothers), she can continue to revise her views.

The AUM makes no assumption about the distribution of  $v_{it}$ . Specifically, it does not assume that  $v_{it}$  has an expected value of 0, either for any time  $t$  or for any person  $i$ . This leaves open the possibility of a population-wide shift in responses as many people react to the same changes in the environment.

Several theories of cultural change at the population level require evidence of active updating at the individual level, such as arguments that the individuals in the population are becoming more polarized on political issues (DiMaggio, Evans, and Bryson 1996; Iyengar and Westwood 2015), that individuals adjust their political preferences based on the performance of the governing party (Bartels and Jackman 2014; Page and Shapiro 1992), that new ideas diffuse across the population via person-to-person contact (Centola n.d.), and that changes in elite discourse drive change in mass opinion (Zaller 1992). In the absence of evidence that people change their minds in persisting ways, these theories cannot be valid at the aggregate level, and we must seek other explanations for aggregate changes.

### **2.2.3 Settled Dispositions Model (SDM)**

Equation (2) represents the settled dispositions model. Here each respondent begins the study period with a settled tendency to respond to the question in a particular way ( $U_i$ ). Temporary considerations, like current events, can play a role in what response a person gives at each time (part of the  $v_{it}$ ), but these considerations have no lasting impact beyond time  $t$ . As the right panel of Figure 1 makes clear, there is no mechanism by which a particular consideration can “propagate up” into the settled disposition and change the baseline for future responses. Such considerations are thus temporary influences only. At its core, this model reflects the Bourdieusian model of action that emphasizes “the past conditions of production” (Bourdieu 1990: 66ff) and not

the contemporaneous environment. In other words, updating happened, but it happened in the past, prior to the time of the study. In this sense,  $U_i$  reflects the *habitus*.

The SDM underlies theories that suggesting people's dispositions are stable throughout their lives after early enculturation or, if they do change, tend to revert back to a relatively stable baseline in short time; this includes "control" theories such as Identity Control Theory and Affect Control Theory (Robinson 2007; Smith-Lovin and Heise 1988). The notion that individuals' dispositions are stable is commonly associated with cohort replacement theories of attitude change at the population level. These models posit that people hold relatively stable opinions, that few people change their attitudes as they age, and most change in public opinion must come from older people dying and being replaced in the population with young people who hold different opinions (Mannheim 1952; Ryder 1965). Vaisey and Lizardo (2016), looking the empirical balance between period and cohort effects on a range of opinion items in the GSS, find evidence suggesting that population-wide cultural change most likely comes in the form of cohort replacement.

On its face, the settled dispositions model might seem to imply people hold strong, consistent opinions, but this need not be the case. The settled dispositions model is also consistent with theories suggesting that people lack stable opinions and are unable to report them as such in a survey context (Bourdieu 1979; Perrin and McFarland 2011). These theories tend to focus on questions of public policy, because the act of

opinion creation requires more information than individuals typically have (Converse 1964; Zaller 1992). Such accounts imply that peoples' responses at any particular time are a deviation from a baseline, with this deviation often assumed to be random and commonly referred to as a form of "measurement error" (Page and Shapiro 1992). As long as these deviations are random, there is no reason to perceive these changes as a form of updating, even if wave-to-wave differences are quite large.

Like the AUM, the SDM does not require  $v_{it}$  to be 0 in expectation for the population at time  $t$ , but it does assume that  $v_{it}$  has an expected value of 0 within an individual over time. That is, current considerations can move people from their baseline temporarily, but there is a tendency to "bounce back" to that baseline over time. The SDM thus allows for population-wide shifts in beliefs, practices, or identities at a particular time (temporary period effects), but it assumes that within individuals these shifts will be erased over time as people return to their baselines.

Like with the active updating model, there are theories of attitude development and change at the societal level that would require evidence of the settled dispositions model at the individual level. For example, if people's attitudes change in a way that reflects updating, then it would be wrong to consider their deviations from their average to be "measurement error"; theories suggesting that changes in belief are not real (Page and Shapiro 1992) would thus be unsupported. Similarly, an item would have to be

relatively stable for most adults for aggregate cultural change to come primarily through cohort replacement (Ryder 1965).

The two general models thus differ fundamentally in their emphasis on the character of personal change. In the AUM, changes tend to persist because they propagate into future responses, shifting the baseline over time. In the SDM, changes tend to revert because each person returns to their baseline. We can therefore summarize the models' implications for individual change as follows: the AUM predicts *persisting change*, whereas the SDM predicts *non-persisting change*.

#### **2.2.4 Consistency in Response**

Neither model makes any assumptions about the magnitude of the variance of  $v_{it}$ , meaning that both models allow for the possibility of either high or low levels of change in responses from wave to wave. This means we must consider separately the overall *amount* of change in responses from wave to wave and the *persistence* of this change. Knowing that there is more wave-to-wave change in one item than another tells us little about how these two items relate to underlying theories of attitude development. However, knowing that two attitudes have the same rate of wave-to-wave change, but one follows an active updating pattern with little random fluctuation and the other follows a settled disposition pattern with a high amount of random fluctuation, allows us to make more precise inferences about the applicability of different theories to these items. As such, we designate a second term, *consistency*, to

reflect this second dimension of change. *Consistency* here simply refers to the degree that the attitude development process departs from pure randomness.

### **2.2.5 Toward Theoretical Synthesis**

We do not suggest that all attitudes and behaviors captured by survey questions follow a settled dispositions model or that all display active updating. Instead, we suggest that by classifying the pattern of change at the individual level, we can use that as a tool for adjudicating debates of change at the population level.

The clearest contrast is between models that assume beliefs are settled during adulthood, such as cohort replacement theories, which would produce a pattern of low active updating, and those that posit individuals update beliefs throughout their lives, which would produce an active updating pattern. This disagreement thus provides the primary framework of our analysis. However, it is possible that some items display active updating while others do not. There are at least three ways to understand how both models could be present in a population without contradiction: age-based differences, core-periphery models, and public-private differences.

*Age-based differences:* The concept of a “cohort effect” is that the environment influences the baseline response tendency when a person is young and then stabilizes for the remainder of one’s life (Elder 1974; Bartels and Jackman 2014). To put it differently, cohort formation requires an early period of updating followed by a later

period characterized by either absolute stability (where the variance of  $V_{it}$  is low) or by temporary, non-persisting changes that disappear as people revert to baseline.

Items reflecting a cohort replacement pattern should display low active updating at the individual level for most adults but high active updating for young people, as this is the period most commonly found to be susceptible to attitude change (see Danigelis, Hardy, and Cutler 2007; Inglehart and Baker 2000; P. S. Visser and Krosnick 2004).

*Core-periphery models:* A number of theories of attitude development posit that people hold “core” beliefs, which they use to adjudicate peripheral attitudes (Boutyline and Vaisey 2017). This work has predominantly focuses on political beliefs and adjudicating whether people use ideological identity, policy positions, moral beliefs, or partisanship to form their positions on other issues (Green, Palmquist, and Schickler 2002; Kinder and Kalmoe 2017; Lakoff 1996). In their analysis of cross-sectional data, Boutyline and Vaisey (2017) find support for ideological identification – assessment of oneself as a conservative or a liberal, not as a Republican or Democrat – as the “core” belief in the political space. Similarly, Baldassarri and Gelman (2008) find increased alignment between partisan identification and policy positions but little increase in the correlation between specific policy positions over time. They interpret this finding as indicating that “voters are splitting along party lines according to the issues that are most salient to them, while they do not bother to adjust their (weak) preferences on the remaining issues” (Baldassarri and Gelman 2008: 442).

The picture emerging from this work is that ideological identity should be a core, if poorly reported, disposition. Changes in ideological identity should be non-persisting, occurring either because people either do not fully understand these terms and report their views with error (Converse 1964) or because people bounce around a stable mean, meaning this item should show low updating and medium-to-low consistency. On the other hand, we should see evidence that individuals adjust their partisan identification to match their ideological identification as they come to find that one party better represents their core positions (Baldassarri and Gelman 2008). In other words, liberal Republicans should become liberal Democrats, not conservative Republicans. This means partisan identification should display high levels of updating. Again, this does not mean many people change, only that changes tend to persist.

Alongside ideological identification, moral issues and general political sentiments might serve as core organizing principles, but what these beliefs are will not be consistent across people (Baldassarri and Gelman 2008). For example, some people might hold strong, consistent views on abortion whereas others care about the amount of government intervention in the economy. Due to this heterogeneity, these beliefs should have a moderate level of consistency, because people who do not value them should report them inconsistently. There should be no evidence of updating in these beliefs, as previous work suggests people do not bring these attitudes in line with other beliefs, even if they are discordant.

Peripheral beliefs that “grow” out of core beliefs, such as preferences for government spending on various priorities and specific policy positions, should show some evidence of updating, as individuals attempt to align these with their core beliefs. Yet, because of the cognitive effort required to map peripheral beliefs onto core beliefs, we generally expect the former to be inconsistently reported (Converse 1964; Zaller 1992). However, there are conditions that might decrease that cognitive effort and therefore increase the amount of updating and consistency with which peripheral issues are reported, to which we now turn.

*Public versus Private Culture:* A third dimension likely influencing the degree of active updating and consistency of attitude and behavior reports is the publicness of the item, or whether it is in some way externalized in public symbols, discourses, and institutions. The act of opinion construction often taps what Lizardo (2017: 91) calls “personal declarative culture,” the “explicit, symbolically mediated culture” such as language that people use to reason, evaluate, judge, and categorize objects. This kind of knowledge can be contradictory and flexible, meaning a person can reach and justify a range of different opinions in response to the same prompt, depending on the circumstances. However, when public culture provides a strong “scaffold” through clear signals of how identities and social locations should influence opinions, it becomes easier for people to maintain consistency (Lizardo and Strand 2010).

Theories that argue individuals do not hold consistent opinions also posit that issues receiving significant media attention can be reported more consistently than other issues (Zaller 1992). This means a small number of high-profile issues could show a higher rate of consistency, but we do not expect high rates of active updating for most of these, because signals are consistent over time. However, we do highlight one issue that saw significant change in elite opinion, which should have led to active updating in the general population. In the time frame we observe (2006 to 2014), there was a major shift in elite opinion on gay rights. Leaders of the Democratic Party (including President Barack Obama) came out in favor of legalizing gay marriage, and a number of Republican Party leaders also began to express support for gay marriage. For this issue, we expect to find evidence of active updating throughout the population, not just among young people.

Items tapping some public dimension of behavior and attitudes are also likely to show higher levels of active updating because social mechanisms, such as publicly joining a group, can help maintain changes in ways that intrapsychic forces cannot. Any item where the burden of maintaining consistency is externalized should show both higher rates of active updating and greater consistency.

As noted previously, we do not expect that all items will align with just one model. If we find a diverse array of patterns, with items displaying a range of active updating and consistency, then the above theories provide a window through which to

interpret these distinctions. Our goal here is neither to declare victory for one of the theoretical perspectives nor to simply say that all “matter.” Rather, our objective is to improve sociological models of cultural evolution by more precisely specifying when and where different types of processes are at work. We believe achieving a better understanding of these processes will be relevant for many subfields of sociology as well as for other social science fields that study changes in beliefs and behaviors.

### **2.3 Research Questions and Expectations**

With these considerations in mind, we ask three questions. First, *to what extent are patterns of personal cultural change generally better described by an active updating model or a settled dispositions model?* Previous work on cultural change using cross-sectional data suggests that cohort effects are generally more important than period effects in explaining broad cultural change at the population level (Vaisey and Lizardo 2016). This implies that, in the repeated measures data on adults we will use here, we should find that the settled dispositions model performs better on most items because cohort formation should be (mostly) complete.

Second, *is there evidence that younger respondents are doing more active updating than older respondents?* In a sample of adults, the possibility exists that cohort formation may be complete for most beliefs and behaviors before the study period (i.e., before age 18). However, if some cohort formation is still occurring among younger respondents, we

should find evidence consistent with the AUM is disproportionately located among younger respondents.

Third, *are there systematic differences in item content among questions exhibiting different levels of active updating and consistency?* The preceding sections have made some predictions based on existing literature, but we cannot enumerate predictions for all reported beliefs and behaviors. As noted, we use these theories as a lens through which to interpret the overall pattern of results.

## **2.4 Analytical Strategy**

To investigate these ideas empirically, we examine 183 survey items from the 2006 to 2014 General Social Survey (GSS) panels in search of evidence in favor of an active updating model. This period of the GSS contains three different three-wave panels, each of which surveys a sample of adults three times over a four-year period (e.g., 2008–2010–2012). Three waves of data are the minimum amount needed to compare the predictions of the active updating model and the settled dispositions model. We discuss item selection below. A broader range of years would be preferable, but the panel component of the GSS began in 2006 and was discontinued in 2014, so these years represent the full range of what we can analyze using the GSS.

### **2.4.1 Statistical Models**

Our main goal is to obtain separate estimates of the amount of active updating and the amount of non-persisting change in responses over time. We first consider

measuring the amount of active updating. Our two different models make different predictions in the three-wave panel context. The AUM makes the following prediction for wave 3:

$$E(y_{i3}) = y_{i2} \quad (3)$$

That is, the AUM predicts that a respondent's most recent response is the best available predictor of her next response, and that wave 1 carries no additional information in predicting wave 3 once we control for wave 2. If change is persisting, then our best guess is that a person's response will be close to what they said last time, and previous responses will provide no additional predictive power.

The SDM makes the following prediction:

$$E(y_{i3}) = U_i \quad (4)$$

Because the best estimate of  $U_i$  is the mean of the respondent's two previous answers, we can rewrite the SDM prediction as follows:

$$E(y_{i3}) = \frac{y_{i2} + y_{i1}}{2} \quad (5)$$

That is, the SDM predicts that the average response of previous waves is the best predictor of the next response. If change is non-persisting, then taking the average of the last two responses will be our best guess about a person's underlying position.

### **2.4.2 Combined Model**

Both of the models above include  $y_{i2}$  as a predictor of  $y_{i3}$ , but only the settled dispositions model includes  $y_{i1}$  as a predictor. If the SDM is correct,  $y_{i1}$  should be just as predictive as  $y_{i2}$  because both are (on average) equally informative about the

respondent's stable disposition. Therefore, to test for evidence of active updating we use a model that evaluates whether  $y_{i2}$  carries any additional predictive power over  $y_{i1}$ . If the two previous estimates are equally predictive, then we can be relatively confident that the data we observe came from a settled dispositions model. However, if  $y_{i2}$  is a better predictor of  $y_{i3}$  than  $y_{i1}$ , this provides evidence that some respondents engage in active updating. We use the following nonlinear model to estimate the relative influence of  $y_{i1}$  and  $y_{i2}$ :

$$E(y_{i3}) = \alpha + \phi\beta y_{i2} + (1 - \phi)\beta y_{i1} \quad (6)$$

Rather than generate separate coefficient estimates for  $y_{i2}$  and  $y_{i1}$ , this model generates two parameter estimates of interest:  $\beta$ , which captures how well any combination of previous waves predicts a person's response at wave 3, and  $\phi$ , the relative proportion of wave 3 explained by wave 2 compared to wave 1. If the settled dispositions model is the preferred data-generating process for an item, then both  $y_{i2}$  and  $y_{i1}$  should be equally predictive of  $y_{i3}$ , and  $\phi$  will equal .5, meaning the best estimate of wave 3 is a function of the mean of previous waves, consistent with Equation (5). If the active updating model is present in at least some respondents and wave 1 provides no additional predictive power when we control for wave 2, then  $\phi$  will increase toward 1 to converge with Equation (3) in certain circumstances.

Our estimates of  $\beta$  provide a measure of the consistency of individuals' responses, contingent on the degree of active updating. We can think of this parameter

as analogous to an  $R^2$  measure in a traditional linear model, capturing the total “predictiveness” of the model. If individuals pick a random response at each wave, the best predictor for a person at wave 3 will be the sample average, and  $\beta = 0$ . If there is little random fluctuation between waves, once the amount of active updating is accounted for,  $\beta$  will approach 1.

### **2.4.3 Comparison to Other Approaches**

Our model is not the first to measure stability and change in panel data, but existing models make assumptions that eliminate the distinction between data-generating processes we seek to test. Hout and Hastings (2016), for example, use a hierarchical model to measure reliability in GSS responses. This model assumes that there is no change in the underlying latent item other than wave-specific period effects (akin to our SDM), so the design precludes the possibility of quantifying the level of active updating in an item over time, assuming that this change is just measurement error. Hout and Hastings also test a structural-equation model designed by Alwin (2007) and Heise (1969) that assumes the process that generates the data is the Markovian active updating process we outlined earlier. This approach gets closer to our AUM by generating a parameter for stability and reliability, but it would require us to make an assumption that the amount of change is consistent across waves. This approach also combines persisting and non-persisting change into two similar but distinct kinds of change: “structural” and “non-structural,” both of which can be persisting and non-

persisting. This distinction, although important for some theoretical questions, is not our focus.

A number of other approaches seek to understand the consistency of latent beliefs by combining and scaling responses to questions representing the same latent concept (Ansolabehere, Rodden, and Snyder 2008), assuming that wave-to-wave changes in responses represent measurement errors around a “true” latent belief. This raises the distinction between the *stability of a belief* and the *stability of a survey question response*. Because we at times invoke both these models, we include in our analysis several composite scales of related items in our analysis. If wave-to-wave changes in survey responses are non-persisting measurement errors, then scales should have higher consistency than the measures they comprise, but we should see no difference in their levels of active updating.

Another possibility is that participation in the survey itself produces change or stability, a phenomenon referred to as panel conditioning bias (Warren and Halpern-Manners 2012; Oh, Yeatman, and Trinitapoli 2019). Warren and Halpern-Manners (2012) outline several forms of panel conditioning, and we can group these into two broad patterns.

On the one hand, people’s responses might become more consistent over time as participation in the survey forces them to crystallize their beliefs, seek out new information that helps them form beliefs, realize their beliefs are out of sync with the

general population, or learn to “game” the survey to get through it faster. If this were taking place in the GSS, it would result in a pattern of high active updating and high consistency, as respondents would change between waves 1 and 2, and wave 2 would become a better predictor of wave 3.

We do not view this as a problem for our theoretical models. If people change their attitudes or behavior as a result of participating in a survey, they are conforming to the AUM, being open to change throughout their life course, and the source of that change is irrelevant. This might lead to a higher estimate of active updating than we might observe in a population that did not take the survey, which would hinder our ability to extrapolate our findings, but it would still provide evidence that people update beliefs over time.

A second form of panel conditioning posits that people exhibit low levels of updating because of commitment bias, or an attempt to maintain consistency in their responses over time, even if they actually change. In this scenario, individuals respond to a question at wave 1 and give the same response in subsequent waves, even if their true beliefs or positions change. This would be problematic for our study, as it would under-estimate the amount of real change in the population. However, if this is the case, we should observe no active updating in responses and high levels of consistency, as it would be illogical for individuals to report random changes if they were attempting to maintain consistency.

Finally, a number of approaches exist for evaluating theoretical process of belief formation and change for the population, such as examining the association between theoretically related values (Baldassarri and Gelman 2008; Boutyline and Vaisey 2017; DellaPosta 2020), or by looking at changes in the distribution of responses over time (DiMaggio, Evans, and Bryson 1996). These tools are well designed to address the questions they set out to answer. However, because our theoretical questions focus on the process of belief change within individuals, these do not speak to our core concerns.

#### **2.4.4 Limitations of the Method**

Three challenges limit our ability to evaluate the presence of settled dispositions and active updating models using our approach and therefore limit the conclusions we can reach. First, our model is designed to allocate variance explained to each of the prior waves rather than to assign probabilities to each data-generating process. Because of this, a few individuals making large persisting changes can inflate the  $\phi$  estimate even if most individuals make small non-persisting changes.

The second is measurement error, which is a form of non-persisting change. For some researchers, measurement error represents the inconsistency that results from constructing responses anew each wave (Zaller 1992), and in that case it should not be considered “error” so much as an indicator of that process at work, since there is no “true” item to measure. For other theories, measurement error reflects individuals’ inability to accurately report their responses. It is also possible that measurement error

reflects errors of selection and interpretation, such as misunderstanding the question or incorrect coding.

Because measurement error looks the same as non-persisting real change (and because the latter is sometimes interpreted as the former),  $\phi$  estimates will be biased toward .5, because responses with error would be random departures from the baseline. There is evidence that many of the items explored in our analysis are measured with significant error (Alwin 2007; Hout and Hastings 2016). On the other hand, previous studies of reliability tend to conflate measurement error and non-persisting real change in attitudes, meaning that although we might have good estimates for the combination of these two processes, we cannot separate them. Because of measurement error, it is unlikely that  $\phi$  and  $\beta$  will reach 1 for any item, even if the underlying process is fully based on active updating.

Our third challenge is that we focus on predicting wave 3. If individuals have a high likelihood of changing between waves 2 and 3, our ability to predict responses at wave 3 will be limited and  $\beta$  will be low. Our model relies on the assumption that “persisting” change is relatively rare and that individuals who change between waves 1 and 2 do not also make persisting changes in the opposite direction between waves 2 and 3. If the rate of active updating is so high that individuals make changes between each wave, then the model becomes indistinguishable from the settled dispositions

model with high measurement error, and it may not be reasonable to consider this sort of change “persistent.”

In addition to these three challenges, there are two forms of change our model is not well designed to account for. The first is a unidirectional shock to the population. Because our model includes an intercept, changes that shift all responses toward one end of the scale are absorbed into that term and not accounted for in our  $\phi$  estimates. The second is change in the variance of responses. If all individuals shift outward or inward toward the population mean but maintain their relative position in the overall distribution, this change will be absorbed into  $\beta$  but not enter into  $\phi$ .

Despite these limitations, the model is capable of detecting the presence of persisting change even in the presence of high levels of measurement error. Because of this, it is best to think of our approach as seeking any evidence in favor of active updating, rather than allocating probabilities to each model. We can only detect whether there is any evidence of persisting belief changes, and therefore whether there is any evidence that active updating is taking place in the population.

#### **2.4.5 Analysis Steps**

Our analysis proceeds in three steps to answer our three research questions. First, we evaluate the overall evidence in support of the active updating model. To do this, we compare for each item the Bayesian Information Criteria (BIC) of a model estimated using Equation (6) with a free estimate of  $\phi$  to a model that constrains  $\phi = .5$ .

We calculate the posterior probability that the model with the free parameter fits the data better. If the model with the constraint is preferred, then we conclude that both wave 1 and wave 2 are equally good predictors of wave 3, meaning there is no evidence that respondents are actively updating on that item.<sup>1</sup>

Second, for variables showing at least some evidence of active updating ( $\phi > 0.5$ ), we ask whether the persistent change is concentrated among younger respondents. To test this, we re-estimate our original model and allow  $\phi$  to have different values above and below a given age cutoff. Rather than test a single age cutoff, we again use BIC comparisons to evaluate whether including the dummy variable improves the model fit using a cutoff of every age between 20 and 45. We test a range of cutoffs to ensure robustness of the overall pattern to specific ages.

Finally, we consider whether there are any meaningful patterns in the relative distribution of evidence for active updating across variables as suggested by existing theories. Although, as we discussed earlier, previous work gives some indications about what we might expect, the approach here will necessarily be inductive.

---

<sup>1</sup> All data used in this analysis is publicly available through the General Social Survey at the National Opinion Research Center at <http://gss.norc.org/>. R code for how we cleaned the GSS panels, all analyses performed in this paper, and all figures created in this paper can be found at [http://github.com/krkiley/panel\\_change](http://github.com/krkiley/panel_change).

## 2.4.6 Item Selection

To test our model on as broad a range of items as possible, we sought measures of attitudes, beliefs, self-assessments, self-perceptions, and social behaviors that were asked in three waves of the GSS panels. We excluded from our analysis questions that focused on demographic characteristics (marital status, household size, region, gender, race, ethnicity), work activity (employment status, income, hours worked, size of workplace), objective socioeconomic status (years of education and highest degree, home ownership), and an interviewer's evaluation of a respondent. We follow Hout and Hastings (2016) and group questions into 15 categories based on subject material. Questions in the same category tend to be asked in the same block during the survey and have the same structure, such as questions about confidence in institutions, questions about government spending, and questions about social life.

We also follow Hout and Hastings (2016) in recreating common scales about gender roles, access to abortion, and social trust. This includes a six-question scale of support for abortion and a seven-question scale that includes the question asking about abortion under any circumstances ("abany"). We use Smith's (1997) scale of "misanthropy" by combining questions about how helpful, fair, and trustworthy people are. We use four questions to create a scale of gender role attitudes (Cotter, Hermsen, and Vanneman 2011). Like Hout and Hastings, we combine civil liberties items into six scales about the freedom of atheists, communists, militarists, racists, and, in the 2010 to

2014 panel, Muslim clergy. We combine four parallel questions about how frequently individuals socialize to create a “social life” scale. We combine four questions about support for suicide under different circumstances. We also created a scale of support for police use of violence against criminal suspects by combining five binary questions about the conditions under which individuals support police use of violence.

In total, we test the model on 183 GSS items, including the composite scales. For each question, we use all cases for which the respondent gave responses in all three waves. Models are estimated using survey weights that account for the GSS’s sampling design as well as non-response adjustment.

## **2.5 Results**


Our model estimates two parameters of interest for each GSS item:  $\beta$ , our measure of consistency, captures how well any combination of previous waves predicts a person’s response at wave 3. High values of  $\beta$  indicate individuals are relatively consistent in their responses, once we control for the amount of active updating. Our measure of active updating,  $\phi$ , captures the relative proportion of wave 3 variance predicted by wave 2. If responses are generated through a true settled dispositions model, then  $\phi$  will be .5 (i.e., both wave 1 and wave 2 are equally good predictors of wave 3). As the evidence of active updating increases,  $\phi$  will increase toward 1. Both  $\phi$  and  $\beta$  equaling 1 would indicate that all individuals who changed between waves 1 and

2 persisted in their change, that an item was measured with no measurement error, and that there was no additional change between waves 2 and 3.

### **2.5.1 Evidence for Active Updating**

To evaluate the evidence in favor of the active updating model, we compare for all 183 items the Bayesian Information Criteria (BIC) of a model with a free estimate of  $\phi$  to a model that constrains  $\phi = .5$ . If the model with the constraint is preferred, then there is no evidence that respondents engage in an active updating process with respect to that item.

Figure 2 plots the distribution of  $\phi$  estimates for the 183 items evaluated in this analysis and the posterior probabilities that the model without the constraint fits the data better, generated by comparing the BIC from models with and without the constraint. On the left side of the figure, we see that the majority of  $\phi$  estimates fall between .5 and .6, meaning that wave 2 is only a slightly better predictor than wave 1 for most items. This suggests that if active updating is happening in these responses, it is relatively infrequent or small compared to temporary change and measurement error.



**Figure 2: Distribution of phi estimates and probabilities that items show evidence of active updating.**

To provide a concrete example, consider the GSS question asking respondents whether they think it should be possible for a woman to receive a legal abortion if she became pregnant as a result of rape, to which individuals can respond either “yes” or “no.” This produces a  $\phi$  estimate of 0.62, above the 75th percentile of all  $\phi$  estimates. Of the 2,259 people who responded to the question in three waves, 257 changed between waves 1 and 2. Under the settled dispositions model, these responses would reflect either measurement error or a temporary shift at either wave 1 or wave 2, and we would expect that about 50 percent, or 129 individuals, would maintain the same response into

wave 3. Only 147 of the 257 people (57 percent) who changed between waves 1 and 2 maintained the same response at wave 3. Therefore, we only have evidence that about 18 individuals (less than 1 percent of the sample) showed evidence of persisting change.

Considered this way, the majority of GSS items demonstrate persisting change at a rate of less than 1 percent of the total sample, and none shows evidence of persisting change greater than 5 percent, with confidence in the leadership of the executive branch of the federal government, confidence in banks and the financial system, and the respondent's belief about whether they will be able to find a good job topping the list. In other words, even for items that show strong evidence of some active updating, the overall amount of attitude change in the population is likely small.

The right side of Figure 2 shows that although the majority of items prefer the free parameter, 75 items (about 40 percent of the total) prefer the constraint, meaning these items show no evidence of active updating over this period. That is, although respondents might give different answers to these items in any particular wave because of measurement error or a transient change of opinion, they tend to revert to their previous positions. This group includes many items about abortion, civil liberties, confidence in institutions, and views on race and gender.


We will discuss in more detail below how different items perform. To answer our first question, however, we need only focus on the overall distribution. Forty percent of items show no evidence for active updating, and among items that do show

some evidence of persistent change, very few come close to approaching 1. This means that for almost all items, measurement error or non-persistent change tend to be much more common than persistent change. We can only be really confident in detecting substantial amounts of persistent change (greater than 2 percent of the population) among a small minority of items, perhaps 1 in 5. This means most of the “change” that shows up in the GSS panels reflects some combination of measurement error or non-persistent change.

### **2.5.2 Age Heterogeneity**

Our second research question asks whether there is evidence that younger respondents update their views more than older respondents. Although it is impossible to determine what proportion of people are following each data-generating process, it is possible to compare the age distribution of evidence for updating in each item.

Of the 108 items that showing any evidence for active updating in the last section, 22 showed differential effects of age for over 50 percent of cutoff ages we tested, meaning the majority did not. Figure 3 plots the estimates of these 22 items for people older than 30 and people aged 30 and younger to get a sense of the magnitude of difference between older and younger individuals on these items.



**Figure 3: Comparison of phi estimate for individuals over and equal to or less than 30 years old, with 95 percent confidence intervals.**

The majority of items showing evidence for age concentration indicate active updating is more prevalent among younger respondents than among older respondents. These items include views on affirmative action, women in the workforce, and politics; several civil liberties items; general views of whether people can be trusted; and views on whether doctors should let terminal patients die. These items tend to be in subject areas where a large proportion showed no evidence of active updating, which suggests an overall trend of these views being formed earlier in life (i.e., prior to becoming eligible for the GSS at age 18) and remaining relatively stable over time.

For some items, such as whether individuals can be trusted, political views, whether physicians should allow terminal patients to die, and whether companies should make special efforts to hire and promote women to address past discrimination, all evidence of active updating disappears for people over age 30. This suggests that these items follow an “impressionable years” pattern: these opinions malleable during early adulthood but quickly harden into “durable dispositions” (Alwin and Krosnick 1991; Krosnick and Alwin 1989; Vaisey and Lizardo 2016). For other items, such as how important people believe it is for children to be popular and views on how much the government should spend on health care, there is still evidence of active updating in older individuals even though it is substantially less than for younger individuals. This is consistent with the “increasing persistence” hypothesis, where attitude change gradually becomes less likely as individuals age (Glenn 1974; Inglehart and Baker 2000).

Eight items show a negative effect of being below the age cutoff on the  $\phi$  value, meaning younger people showed *less* evidence of active updating than did older individuals. These items include how often individuals were active in religious activities, views on suicide in the case of bankruptcy, and views on whether aging parents should live with their children. Some of these items might be things people are not forced to consider until later in life, and as a result people do not form clear opinions while young. This pattern of older individuals changing their attitudes and behaviors at higher rates than younger individuals is somewhat unanticipated in the attitude change literature (P. S. Visser and Krosnick 2004; Danigelis, Hardy, and Cutler 2007), and it suggests greater heterogeneity in the relationship between age and attitude change than previously theorized.


The remaining 86 items (just under half of all items explored here) show evidence for some active updating but do not show consistent evidence for age heterogeneity, suggesting a more complicated relationship between age and attitude change than previously theorized. However, this does not mean these items show strong evidence of an active updating model. These items may simply be susceptible to updating for a small proportion of the population.

### **2.5.3 Item Heterogeneity**

Although 40 percent of items show no evidence of active updating, and those that do show evidence tend to show only weak support for active updating, it is difficult

with just these findings to draw any broad conclusions about how these results speak to theories of attitude development and change. Here we bring in our second dimension of attitude change, consistency in responses, to clarify the overall pattern.


Figure 4 plots the  $\phi$  and  $\beta$  estimates for items shaded by whether they preferred the  $\phi = .5$  constraint or not; a few items that stand out are labeled. Items tend to prefer the  $\phi = 0.5$  constraint for a combination of two reasons: because wave 1 and wave 2 have equal predictive power ( $\phi$  is close to .5) or because the measure is so unpredictable ( $\beta$  is low) that neither wave 1 nor wave 2 has much predictive power, making any observed active updating close to meaningless.



**Figure 4: Distribution of beta and phi estimates for GSS items by whether model prefers  $\phi = 0.5$  constraint.**

Items showing evidence of active updating tend to have  $\phi$  estimates greater than .55, and most have  $\beta$  estimates greater than .6. A small group of items, including confidence in the leadership of the executive branch of the federal government have low  $\beta$  estimates, meaning prediction at wave 3 is difficult, but have large  $\phi$  estimates, meaning wave 2 is still a better predictor than wave 1.

Which items show the strongest evidence for active updating? There is no way we can discuss all 183 items in detail without the discussion becoming tedious. Figure 5 summarizes the distributions of  $\phi$  by the content of the question, and Appendix A includes  $\phi$  estimates for all items. We constrain items showing no evidence of active updating to  $\phi = .5$ . In addition to showing the median and interquartile range of each distribution, Figure 5 also highlights the item in each group that shows the greatest degree of evidence for active updating.



**Figure 5: Summary of phi estimates for all items, by topical group. Numbers in parentheses indicate the number of items in each topical group.**

There is a lot to process, even in this summary figure. Nevertheless, the main takeaway is that, even for items showing some evidence for active updating, the values of  $\phi$  are still quite low. Only two groups of items have median  $\phi$  values above 0.6: public spending and religious activity. In general, it is accurate to say that most of the “change” measured by the GSS is not persistent but some combination of measurement error and short-term fluctuations.

Consistent with the findings of Vaisey and Lizardo (2016), we see that more than half of items in the gender, family, race, sex, civil liberties, and confidence in institutions groups show no evidence for updating. These categories also contain several items that show evidence of active updating in younger cohorts, suggesting that these items became settled by the time respondents entered the GSS sample. Views on these issues are likely shaped by early socialization experiences and mostly settled by the time a respondent reaches adulthood. This means that, for most of these items, population-level change must occur through cohort succession rather than through individual change.

There are several exceptions to this general pattern, however, even in categories with otherwise low  $\phi$  values. Items with the largest values generally have one or more attributes in common. We consider these attributes to give some general impressions of the pattern.

Some of the high- $\phi$  items rely on *external mechanisms* to help maintain them. If a person starts going to church or starts socializing with friends at a bar, she builds social networks that make this behavior more likely to continue. This is clear when contrasted with how often individuals socialize with friends, relatives, or neighbors, which are more nebulous questions that display less active updating. Switching political parties (which involves changing public registration) is a more persistent change than changing political ideology (which can happen privately in the mind). Owning a gun has a high  $\phi$  value because a new physical object either enters or leaves the person's possession.

Other high- $\phi$  items have a *changing referent*. That is, although the item wording is the same, the object to which the question refers may change between survey waves. The most obvious example of this is the item about confidence in the executive branch of the federal government (which has the highest  $\phi$  value of all items in the analysis). The president changed between the 2008 and 2010 waves of the GSS, meaning the question no longer referred to the same administration. If we generate estimates for this item for each of the three panels (2006 to 2010, 2008 to 2012, and 2010 to 2014) for the "confidence in the leadership of the executive branch of the federal government" item, it is only for the middle panel (2008 to 2012), where the president changed between waves 1 and 2, that shows significant evidence of persistent change ( $\phi = .95$ ,  $\beta = .51$ ). In the 2006-10 panel, waves 1 and 2 have almost no predictive power ( $\beta = .10$ ). For the 2010-14 panel,

which takes place entirely during the Obama administration,  $\phi$  moves much closer to .5 and predictive power increases ( $\phi = .57, \beta = .65$ ).

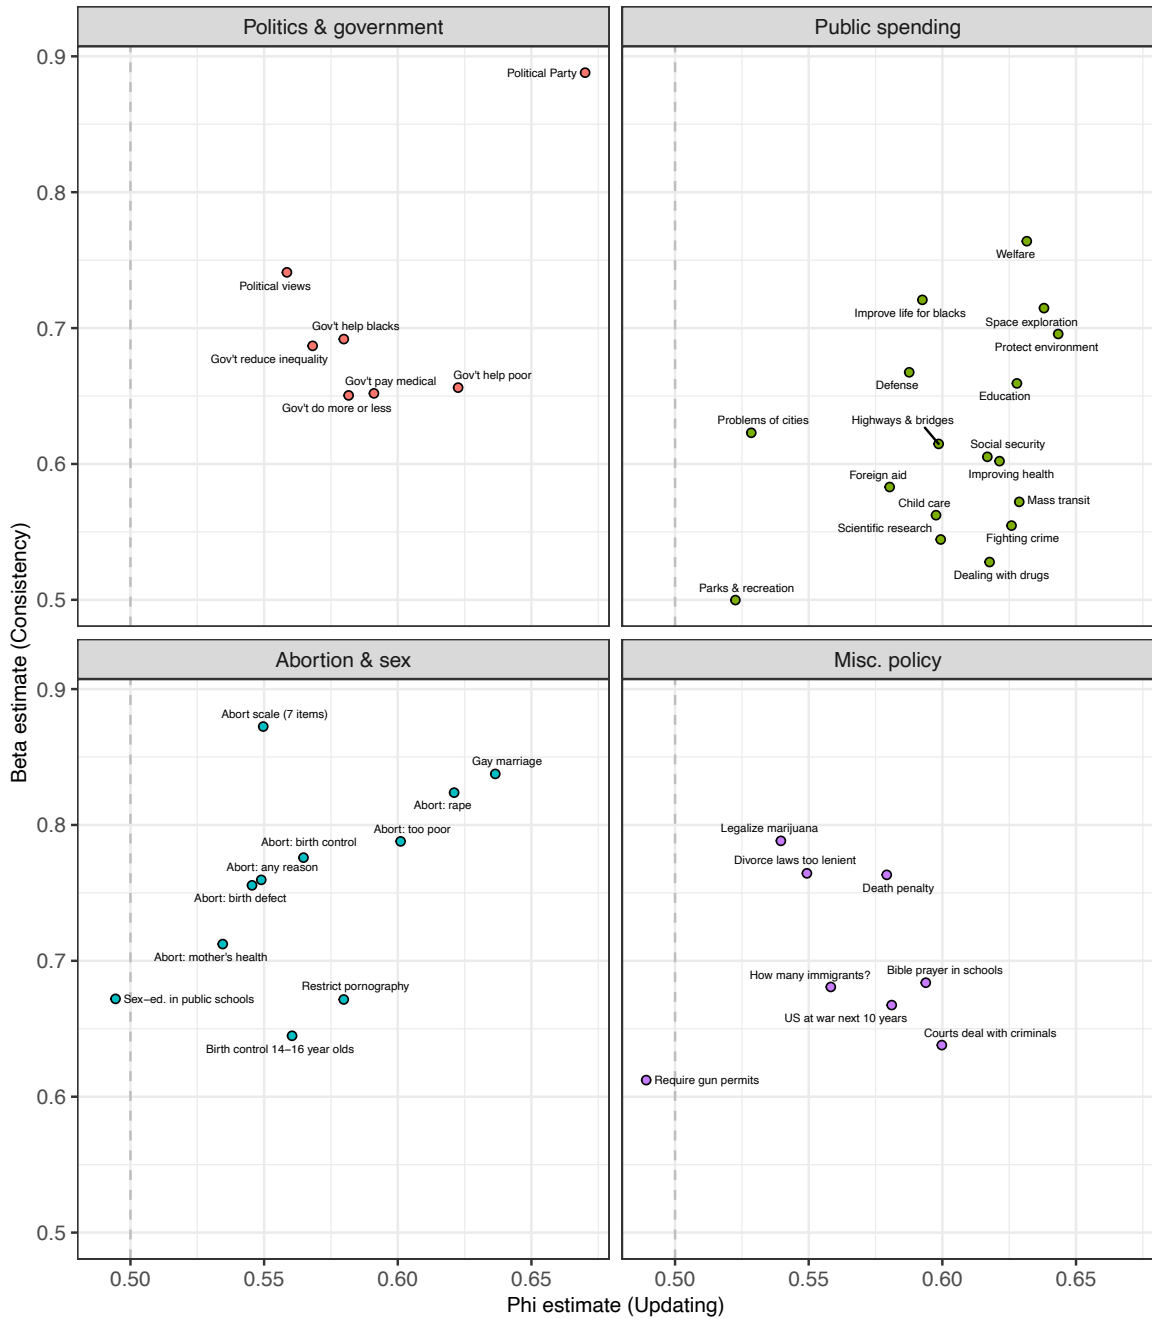
Likewise, all public spending items refer to whether the government is spending “too much, too little, or the right amount” on different areas. The change of administration and changing federal spending policies likely affected these items. The same applies to most questions about subjective SES, where respondents are asked questions about their personal financial or work situations (which changed for many Americans during the time of the study due to the Great Recession). If the environment is changing, we should see exactly this sort of pattern.

Perhaps the most striking pattern in our findings pertains to questions about *gay rights*. Six items ask about some aspect of gay rights, and all show evidence for active updating. Questions about civil liberties for gays and about gay marriage are the highest in their categories. The huge public and political salience of this issue throughout the study period likely made this issue one where more people than usual were open to revising their views. This pattern is consistent with Zaller’s (1992) argument that highly salient issues where elite opinion shifts can lead to large changes in public opinion.

A few items lack external mechanisms, represent forms of private culture, and did not achieve salience during this period, but still display active updating, including some views on abortion and the morality of different forms of sex. We address these items in greater detail in the discussion section.

### 2.5.4 Political Beliefs

Because theories of public opinion formation and the development of political beliefs form the bulk of the theoretical tension that frames this analysis, and because many of these beliefs display evidence of active updating, we examine these beliefs in greater detail than other items. Figure 6 plots the  $\phi$  and  $\beta$  estimates for questions about political identity, the role of government, and specific policies. We break these items up into categories and remove a few items for ease of viewing.



**Figure 6: Phi (updating) and beta (consistency) estimates for 43 political attitudes and identities.**

The figure displays several notable features. First, political party affiliation is a clear outlier, with greater consistency and greater active updating than other items. In contrast, questions asking about specific public spending priorities tend to have low consistency compared to other items. Political views—including ideological identification (named “Political views” in the figure) and general views of the role of government (“Gov’t do more or less,” “Gov’t reduce inequality”)—display higher consistency than specific policy questions but show weak evidence of active updating. It is important to restate that for political views, all evidence of active updating disappears by the time individuals reach age 30, and individuals become more consistent in reporting their ideological identity as they age.

People tend to be inconsistent in their responses to spending priorities, but views on other policy questions with both high and low salience are reported quite consistently. Views on abortion, especially when considered as a scale, are consistently reported. As with all other scales, aggregating the composite items increases consistency but does not affect the estimate of the amount of active updating, validating our method’s ability to separate persisting and non-persisting change. As discussed previously, views on gay marriage display high levels of active updating. In contrast, the item asking about support for legalizing marijuana use is reported with consistency but weak active updating. This is notable, as public opinion and policies regarding both have shifted considerably in recent years. Previous work suggests that change in both

could be driven by the same underlying process (Schnabel and Sevell 2017). Our findings suggest that, at least in recent years, many individuals have changed their views on gay marriage, but changing views on marijuana have been driven primarily by cohort replacement.

## **2.6 Discussion and Conclusion**

This study was motivated by a theoretical contrast in cultural sociology between the settled dispositions model (SDM), which emphasizes the power of the past, and the active updating model (AUM), which emphasizes contemporary responses to external conditions. These two models of individual change are implicit in analyses of cultural change at the population level, but they are rarely compared empirically. We asked three general questions in light of this distinction and developed an empirical approach to adjudicating (albeit imperfectly) between these models. We now revisit our research questions to summarize what we have learned.

*First, to what extent are patterns of cultural change generally better described by an active updating model or a settled dispositions model?* Generally, and consistent with previous research (e.g., Vaisey and Lizardo 2016), we see a greater degree of evidence in support of the settled dispositions model. Around 40 percent of all items show no evidence for updating. For items that do show evidence of active updating, the overall rate of persisting change in the population is likely low. For most items measured by the GSS, less than 1 percent of the population appears to make any persisting change in their

views in a two-year period, and most changes appear to be short term or random deviations. In practice, this means that knowing what a person said two years ago said two years ago does not help us predict their current views more accurately than knowing what they said four years ago. Unfortunately, because of measurement error, we cannot be sure exactly how much updating there is. But the average level of updating for most views appears to be low.

*Second, is there evidence that younger respondents are doing more active updating than older respondents?* For a limited subset of items, there seems to be evidence that younger respondents are updating their views, but older respondents are not. This is consistent with a “cohortization” model that views young respondents as susceptible to updating shocks and older respondents as relatively insensitive to such shocks (see e.g., Bartels and Jackman 2014). Because the youngest respondents in the GSS are 18 years old, we may lack the ability to detect updating occurring earlier, so this is probably an underestimate. This may be the case for the roughly 70 items that show no evidence of persistent change over time. Overall, however, we find significant heterogeneity in the relationship between age and attitude and behavior change.

*Third, are there systematic differences in item content between questions that are better described by each model?* With 183 items tapping very different kinds of opinions and demonstrating a range of active updating and consistency, there is no one overall pattern of attitude change in the population. We find patterns of responses that provide

a range of support for diverse theories of attitude formation, which we discuss below. At the same time, our results are consistent with some general expectations about updating. Questions with public (or otherwise changing) referents and questions tapping high-salience topics over the study period (e.g., gay rights) showed the most evidence for active updating. Items assumed to be “peripheral” beliefs, such as specific policy questions, also showed evidence for updating and low consistency. Most items about gender, family, race, and institutions showed the least evidence for updating, suggesting most have settled on views on these topics by the time they turn 18.

A handful of items display active updating and enough consistency to merit further consideration but are not easily explained by the theoretical mechanisms outlined so far. These include questions about abortion in the case of rape or poverty, the morality of premarital and teen sex, the ideal number of children in a family, and whether children should obey adults or think for themselves. These items did not achieve high salience during the study period and there were not, to the best of our knowledge, clear changes in elite opinion on these issues. This suggests additional mechanisms can drive persisting change at the individual level, such as exposure to these ideas via mass media or direct interactions with people in these groups. In the case of questions about children and family structures, the experience of having children might lead to changes in these attitudes. It is important to note that even though these

items demonstrate active updating, only a small proportion of the population that are changing their views during the study window.

The overall pattern of results, while diverse and at times hard to reconcile, does shed light on several theoretical debates. We now consider these implications.

### **2.6.1 Implications for Cultural Sociology**

In the domain of cultural theory, our findings support the view that a great deal of cultural change happens slowly through the mechanism of cohort succession. Most beliefs about gender roles, sexual morality, and abortion appear to be settled by early adulthood. The settled nature of these beliefs is often coupled with a high rate of consistency, suggesting individuals truly hold these beliefs or at least have sufficient external support to consistently report them over time. In contrast, many views about race were so inconsistent that it would be difficult to call them either settled or updating. Even items that did display strong evidence for active updating (e.g., “whites rich or poor” and “whites work hard”) were still hard to predict wave after wave.

Our results suggest that one reason attitudes are largely stable is because most issues simply do not reach the level of salience necessary to shift opinions. In contrast to other beliefs, the pattern of findings for gay rights shows that a high degree of public salience and social movement activity can accelerate change by encouraging people to update their views. By definition, salience is a limited resource, meaning only a few beliefs and behaviors could change at this rate during any given period. The baseline

process of attitude change appears to be more consistent with a model that shows people do not really change; rather, they die and are replaced by cohorts with different views. This general model is more consistent with a Bourdieusian theory that emphasizes the “conditions of past production” rather than processes of active meaning construction with little long-term memory.

The dominant pattern is stability during adulthood, but the results suggest greater heterogeneity in the relationship between age and attitude change than what is emphasized in existing theories. The most prominent and well supported theories of attitude change suggest a peak of susceptibility to attitude change early in adulthood and either a rapid or gradual decline in attitude change with age (Alwin and Krosnick 1991; P. S. Visser and Krosnick 2004). An additional view, the “life stages” hypothesis, suggests a similar pattern with a surge in attitude change late in life (P. S. Visser and Krosnick 2004). This age-related decline appears to be true for many political attitudes, yet the pattern is not nearly as consistent for other kinds of attitudes. Some items, such as views on the Bible, suggest equal openness to attitude change as individuals age. Others, such as views on most abortion questions, suggest early adult socialization is so strong that lifetime opinions are settled by the time most people reach 18 years old.

For some beliefs, such as whether aging parents should live with their adult children and whether divorce laws are too lenient, persistent change becomes more common with age, a pattern not accounted for in any major theoretical account of the

relationship between age and attitude change. This pattern suggests the salience of an issue can matter at the individual level, as well as the societal level. Rather than supporting a single theory linking age to attitude change, our results call for more work linking attitude content to social factors that encourage openness to change at different ages. Rather than assuming that “attitudes” in general are more or less likely to change at particular ages, we should explore the relationship between age and stability for a range of attitudes. Doing so will expand our understanding of the institutional and developmental factors that that give rise to stable or variable attitudes (Howe and Krosnick 2017). Recent work that has attempted to provide social explanations for (mostly political) attitude stability in middle age is an important step in that direction (Eaton et al. 2009; Visser and Mirabile 2004).

### **2.6.2 Implications for Political Sociology**

A major takeaway of our analysis is that ideological identity (identity as a liberal or conservative, and the extremity of this identification) is in all practical terms stable for individuals over age 30. Respondents might express different positions from wave to wave, but in guessing what a person will say in the future, we are better off guessing the mean of their previous responses than their most recent response. Our sample does not cover a large enough window of time or the life course to say for certain whether this represents a regular pattern (perhaps there is some period-specific reason why younger individuals in our data changed while older individuals did not), but the pattern is

consistent with theories and previous findings that political dispositions become settled by age 30. With the exception of some low-profile government spending questions, most policy questions showed greater evidence of active updating than did the question of ideological identification, although the overall level of active updating was still limited. Partisan identification, in contrast, showed the highest degree of active updating of all political questions, as well as some of the highest consistency, and this updating was active across all age ranges.

These results are consistent with theories positing that, at least in the current era, ideological identification (view of oneself as liberal or conservative), rather than some other belief such as partisan identification, moral views, religion, or particular policy positions, is the “central” political belief (Boutyline and Vaisey 2017; Converse 1964).

It is important not to overstate the role of this belief in forming other beliefs. Our findings suggest individuals bring their partisan identification in line with their ideological identification, but there is limited evidence of adjusting other beliefs. This pattern is most consistent with Baldassarri and Gelman’s (2008) model of “partisans without constraint,” which suggests that individuals hold a few strong beliefs and align their partisan identification with these, rather than adopting beliefs as a function of their partisan identification.

Because there is no evidence of ideological change for the majority of the sample, and only very weak evidence of changes in specific policy positions, even high-profile

ones, the pattern of results is not consistent with the popular conception of political polarization in which individuals become more extreme in their views over time. This is notable given the time frame of our study, which covered the Obama administration, a time commonly assumed to have seen a conservative shift for Republicans and a liberal shift for Democrats. Similarly, our findings are not consistent with the idea that individuals adopt a partisan affiliation based on their social groupings and subsequently adjust their ideological commitments to conform to that (Green, Palmquist, and Schickler 2002). For this to be true, ideological identification and other political beliefs would have to demonstrate at least as much updating as partisan identification, which is not true in our data.

The finding that ideological identification is comparatively stable should not be taken to imply that most people in the population have the kind of tightly knit belief structures that political scientists typically call “ideologies.” Outside of a handful of high-profile items such as partisan identity, abortion, and gay marriage, individuals appear to lack clear opinions on most specific policy questions. The low degree of consistency in beliefs is consistent with the idea that people are “ideologically innocent” (Kinder and Kalmoe 2017). What we can say is that people over age 30 in the United States during the period we studied did not make lasting changes in whether and how strongly they thought of themselves as “liberal” or “conservative.”

The overall picture that emerges from evaluating the active updating and stability of political items in the GSS is one in which the majority of respondents, especially those over age 30 years old, hold a general political identity and a few clear views on issues like abortion, attempt to align their partisan identification to their views, and respond to elite opinion change when it provides clear signals. It is not a picture of a rapidly polarizing society or one wholly ignorant of public debates.

### **2.6.3 Methodological Implications**

Methodologically, our results highlight the challenges of evaluating population-wide attitude change using short-term panel studies. The evidence strongly suggests that most of what might be interpreted as “change” in the GSS panels is some combination of measurement error or non-persistent change. It does not matter whether measurement error or short-term change is the predominant driver behind this pattern; what matters is that substantive change is too rare in a sample of adults to measure accurately on the vast majority of items. This strongly argues against using two-wave panels to measure attitude change, which do not allow researchers to separate persisting from non-persisting changes.

The fact that persistent change is practically nonexistent for many items bolsters the case for using repeated survey responses to measure the reliability of survey items (Alwin 2007; Hout and Hastings 2016), since it is often a valid assumption that the underlying view is unchanging. At the same time, our results call for greater focus on

methodological tools that can separate short-term attitude change from measurement error. Although we generally assume that lasting changes in attitudes are more likely to influence behavior, this is not necessarily true. Short-term attitude changes might be meaningful in shaping short-term behaviors but identifying this is difficult.

We said at the outset that patterns of change might be the result of panel conditioning, or that the process of participating in the survey could lead to more active updating or stability than might be expected in the absence of survey participation. One could view our results through this prism and claim that items that exhibiting high active updating and high consistency (e.g., views on gay marriage or partisan identification) do so because of panel conditioning bias, or that items that exhibiting low active updating and high consistency (e.g., views on abortion or the legalization of marijuana) do so because of commitment bias, but it becomes difficult to explain why these biases operate for specific questions and not others.

We believe the overall pattern of results we observe is more consistent with other theoretical models of belief change than with those outlined by panel conditioning. We see too much inconsistency in responses for commitment bias to be a major explanatory factor. Items that we have theoretical reasons to suspect might succumb to the updating form of panel conditioning bias, such as questions where “respondents’ initial attitudes are less crystallized” (Warren and Halpern-Manners 2012: 499), questions that “increase respondents’ knowledge of the behavior and/or their motivation to engage in it”

(Warren and Halpern-Manners 2012: 500), or questions that “induce respondents to provide socially non-normative or stigmatized responses” (Warren and Halpern-Manners 2012: 501), tend to show low active updating and low consistency.

We cannot (and would not want to) rule out the possibility that panel conditioning is taking place in the GSS. We believe it is worthwhile to explore these same GSS panels for evidence of panel conditioning. However, we do not believe that panel conditioning bias is the principal driver of the overall pattern of change and consistency we observe.

Because many attitudes, including views on abortion, race, gender roles, social trust, and institutional confidence, have mostly stabilized by the time individuals enter the GSS, our results also call for greater emphasis on surveying the attitudes of adolescents and children to understand how these attitudes are formed. Panel studies tracing the political socialization of adolescents are rare but could be highly fruitful. In a similar vein, it does not seem worthwhile to ask certain GSS questions repeatedly. Questions about racial stereotypes, which show almost no consistency from wave to wave but have been asked every wave since 1996, strike us as particularly problematic. Repeated questions should be specifically targeted to topics that seem to be changing broadly (e.g., politics, gay rights).

Our results ultimately suggest that real, persistent attitude change is an uncommon phenomenon among adults. Understanding the social origins of individuals’

attitudes requires greater focus on the “conditions of past production” — childhood and adolescence—that give rise to persistent beliefs in adulthood.

## **3. Study 2: Ambivalence Is Everywhere: Quantifying Opinion Behavior Across Topic Domains**

### ***3.1 Introduction***

In their day-to-day lives, people are exposed to an incredible range of heterogeneous, conflicting culture (DiMaggio 1997; Martin 2010; Swidler 1986). They hear arguments for and against abortion, government spending, affirmative action, and science; religious leaders and others debate the proper role of religion in the public sphere; and cultural critics debate how good heavy metal music is (Bryson 1996; Ferree 2003; Hunter 1992; Mann and Schleifer 2020). People also come to these debates with a relatively impoverished cognitive structure that struggles to make sense of this “screeching, honking, throbbing mess,” and they cannot help but internalize most of these considerations to some degree (Martin 2010: 229). Because of this diversity of public and private culture, the question of whether people are able to maintain consistent attitudes has been a key debate in the social sciences.

On one hand, researchers point to low wave-to-wave correlations in people’s responses to the same question over time, question-order and -wording effects, the effects of psychological primes, the complexity of culture, and the general limitations of human cognition to argue that people do not carry around fixed beliefs (Converse 1964; Martin 2010; Perrin and McFarland 2011). Instead, they argue, people carry around diverse considerations on issues of public interest and construct opinions in the interview or survey setting, drawing on recent considerations from their social

environment (Swidler 2001; Zaller 1992). Other researchers point to the stability of attitudes in the aggregate, both aggregating across questions within people and aggregating the same question across a population; the infrequency of durable changes in attitudes over time; and the ability of specific attitudes to predict a range of behaviors over time as evidence that people carry at least some stable latent attitudes or dispositions, even if these positions are obscured by measurement error (Ansolabehere, Rodden, and Snyder 2008; Inglehart 1985; Kiley and Vaisey 2020; Vaisey 2009).

Recent work argues dispositions are mostly stable by the time people reach adulthood (Kiley and Vaisey 2020; Vaisey and Lizardo 2016). But these findings elide the “non-attitudes” debate by grouping together people who hold stable, unchanging opinions with people who change randomly (but not durably) from wave to wave. In other words, while finding that people tend not to make durable changes in their beliefs over time, these approaches cannot say whether this is because people hold real, stable attitudes or whether it is because they hold no attitudes at all. This is, to put it mildly, an important distinction.

A middle ground – going back to Converse’s development of the “non-attitudes” thesis – suggests that on any particular issue (and in any particular window of time) the population comprises three groups: people who hold stable opinions they can report consistently, people who hold weak opinions subject to temporary influences, and a small group of people who make durable change (Freder, Lenz, and Turney 2019; Hill

and Kriesi 2001a; Zaller 1992). In other words, rather than assuming that all members of the population either have or lack opinions, this approach says, “some do, and some do not,” a position consistent with advances in psychology distinguishing between strong attitudes, which govern behavior and thought, and weak attitudes, which do not (Howe and Krosnick 2017). But for any particular question, these groups will vary in size and composition.

The debate about whether people hold stable attitudes has implications for the study of attitudes and beliefs in social behavior. If people’s dispositions and behaviors are swayed by temporary influences, there is substantial room for contemporaneous social structures, opinion leaders, and situations to shape attitudes, and explanations for attitudes should be rooted in these social structures, as pragmatist theories suggest (Gross 2009; Joas 1996). If people’s attitudes are relatively durable and impervious to social influence, these attitudes would be more likely to shape patterns of behavior, affiliation, and belief over time, and explanations for attitudes should be located in people’s backgrounds, rather than contemporaneous social structures. A world where some people hold stable attitudes and others do not directs attention to “institutions and contexts and other forms of objectified cultural structure” that facilitate attitude stability (Lizardo and Strand 2010: 206; Martin 2010) and suggests different attitudes might matter in explaining different peoples’ behavior.

Research on variation in attitudes over time tends to assume people have latent beliefs and variation over time is measurement error (Alwin 2007; Hout and Hastings 2016). As a result, these approaches do not distinguish between people to demonstrate stable attitudes and those who appear to be constructing a new response each time. Research attempting to separate these two types of behaviors tends to focus on a small number of attitudes (Converse 1964; Hill and Kriesi 2001a). While these works have produced significant insights into attitude structuring, their limited scope prevents the comparison of attitudes across topic domains and limits the ability to make general inferences about attitude stability and instability over time.

In this paper, I formalize three models of attitude behavior rooted in cultural sociology and use a finite mixture model approach to estimate for more than 500 attitude<sup>1</sup> questions across 10 panel data sets the proportion of respondents who report stable opinions, the proportion of respondents who report vacillating attitudes or no opinion at all, and the proportion of respondents who demonstrate durable changes of opinion. These questions include topics typically addressed by political scientists, as well sociologically relevant questions about religion, gender roles, race relations, morality, institutional and social trust, and more.

---

<sup>1</sup> The term “attitude” has diverse meanings across the social sciences. For example, Alwin uses the term to indicate “Evaluative responses to particular objects or actors, assumed to exist along a positive/negative continuum of acceptance, favorability, or agreement” (2007: 123). I use the term more broadly to encompass not just this kind of question but all subjective evaluations, including values, beliefs, and self assessments.

I find wide variation across these questions in the proportion of people who demonstrate stable opinions, with some questions demonstrating widespread response stability (with 70 percent of respondents or more giving what could be considered “stable” opinions) and others demonstrating rampant inconsistency, with more than 70 percent of respondents appearing to construct a new opinion each wave. Overall, rates of vacillating change exceed rates of stable attitudes, and high rates of ambivalence are found in attitudes about political issues, religious and moral beliefs, self-assessments, sentiment toward groups, and more. People disproportionately demonstrate stability on related issues (e.g., stability on one political issue predicts stability on other political issues, but not on religious issues), suggesting different forces give rise to stability in different domains. Rather than argue people in general “have” or “lack” opinions, these results suggest both are generally true for any particular question, and they reinforce the argument that stable opinions are principally a function of attention and social structure.

## ***3.2 Theoretical Framework***

### **3.2.1 Stable Opinions and ‘Non-Attitudes’**

Broad streams of social science research assume people have relatively stable attitudes or dispositions, which is why Converse’s original “non-attitudes” findings generated such a robust debate. The most sustained challenge to the “non-attitudes” model are what are called “measurement-error” models, which posit people holding (latent) attitudes but reporting them with some error (Achen 1975; Ansolabehere,

Rodden, and Snyder 2008; Inglehart 1985). Error arises for a variety of reasons, but when responses to the same question are aggregated at the population level, when responses to related questions are aggregated within people, or when certain statistical methods are employed, respondents display much higher levels of stability than when we look at their individual responses over time (Ansolabehere, Rodden, and Snyder 2008; Inglehart 1985; Judd and Milburn 1980; Page and Shapiro 1992). Researchers interpret this as evidence of stable “latent” attitudes.

In cultural sociology, this line of thinking takes the form of suggesting people’s implicit dispositions shape their attitude behavior over time, even if they cannot articulate these dispositions in interview contexts (Miles 2015; Vaisey 2009). Vaisey argues interviews and surveys tap distinct cognitive processes. Interviews tap discursive reasoning, bringing to mind the breadth of contradictory considerations people have internalized. Fixed-choice survey questions invoke practical reasoning, gut feelings about which answer is correct. In support of his argument, he and others show that responses to questions about worldviews and values predict a range of behaviors across contexts (Miles 2015; Vaisey 2009; Vaisey and Lizardo 2010).

These “stable attitude” theories suggest people are consistent in their cultural commitments over time, especially in survey responses. When asked to give an opinion about an issue, such as whether they agree or disagree that “Morality is a personal matter and society should not force everyone to follow one standard,” a moral relativist

might vacillate between “agree” or “strongly agree,” but when they look at the “lineup” of options, they are not going to say society should enforce one standard, and they are not going to say they lack a position. These results will feel wrong, and the person will avoid them, even if they are not conscious of why.

At the same time, a key finding of cultural sociology in the last half-century is that rather than internalize a consistent cultural worldview, people are exposed to a diverse, contradictory cultural repertoire of beliefs, practices, and assessments (Swidler 1986; DiMaggio 1997). Research documents this contradiction across domains. In America, love is a choice entered into freely and a unique and irreplaceable commitment fundamental to a person’s identity (Swidler 2001). Morality is fixed and relative (Baker 2004), orthodox and progressive (Hunter 2000). American culture is individualist and collectivist; managerial and therapeutic; biblical and republican (Bellah et al. 1985).

This contradiction is present in the public sphere and in personal culture, the declarative and non-declarative attitudes, worldviews, values, and dispositions that manifest at the individual level (Lizardo 2017). Because of their cognitive limitations, people consume a broad array of cultural information and cannot or tend not to engage in the effort to reconcile these contradictions (Martin 2010; Zaller 1992). As a result, “our heads are full of images, opinions, and information, untagged as to truth value, to which we are inclined to attribute accuracy and plausibility” (DiMaggio 1997: 267). Because they have these contradictory considerations, people struggle to maintain a single line of

cultural reasoning over time (Swidler 1986). In successive interviews (and often in the same interview) people demonstrate different opinions on the same issue without recognizing these contradictions (Swidler 2001). People draw on different cultural resources to justify institutionally constrained behavior (Mills 1940; Scott and Lyman 1968). And people often demonstrate high levels of inconsistency in response to the same questions over time, (Alwin 2007; Converse 1964; Hout and Hastings 2016).

This work parallels a line of public opinion scholarship. In his Receive-Accept-Sample model, Zaller (1992; Zaller and Feldman 1992) argues that people are characterized by ambivalence toward political issues and tend to be uncritical toward messages they receive. They store conflicting arguments, and “most of the time, there is no need to reconcile or even recognize their contradictory reactions to events and issues. Each can represent a genuine feeling, capable of coexisting with opposing feelings and, depending on momentary salience in the person’s mind, controlling responses to survey questions” (Zaller 1992: 93). When people are called to present an attitude, they conjure up some or all of these considerations and construct an answer. But because considerations are called up in a haphazard way, which manifest at any time can vary, and people’s assessments can be influenced by question wording, question order, psychological primes, and changes in their information environment.

Neither the RAS model nor the heterogeneous culture model suggest people lack considerations. Both argue conflicting considerations make people express different,

often strongly felt, opinions over time. Neither model prohibits people from consistently reporting strong contradictory beliefs – support for more government spending, lower taxes, and a balanced budget, for example – over time. But these models suggest that questions calling forth these conflicts will lead people to give inconsistent responses as different considerations are pushed to the foreground of cognition each time.

These “inconsistent attitude” theories expect people to vacillate in their cultural or political beliefs when measured over time. On the morality question raised earlier, these approaches expect people to sometimes say they lack an opinion as they feel torn between their considerations. Other times, because they recently heard a compelling argument in favor of moral absolutism, they will say they disagree. Still other times, after hearing a counter argument, they will say they agree.

### **3.2.1 Some Do, Some Do Not**

The debate about the stability of attitudes persists because at different times both seem true. People sometimes behave as if they have clear, consistent political or cultural beliefs informing their behavior. At other times they behave as if they do not. This suggests a path forward.

Researchers invoke Converse’s work (1964; Converse and Markus 1979) to suggest the American public lacks opinions on political matters, but his “black and white” model got its name by suggesting on any particular issue the population could be divided into “a ‘hard core’ of opinion on a given issue, which is well crystallized and

perfectly stable over time” and a group of people whose responses are “statistically random” (Converse 1964: 242). The population did not comprise ideologically distinct camps but “issue publics” that care about different issues. Converse’s model has received support over time and across data sets (Converse and Markus 1979; Taylor 1983; Hill and Kriesi 2001a).

Similarly, Zaller’s RAS model supports a model in which some people hold stable attitudes and others do not, lodging attitude stability in social structures that either help people avoid heterogeneous messages or provide motivation to reconcile contradictions (Zaller 1992). It is often not the case that people who report consistent attitudes lack considerations on the “other side” of an issue. These people are often the most engaged in a subject and are often better than others at making cases with which they disagree. But something about their social situations motivates them to reconcile conflicting considerations.

Lending support to this model, psychological research distinguishes between strong attitudes, which are “resistant to change, stable over time, influential on cognition, and influential on action” and weak attitudes, which are not (Howe and Krosnick 2017: 329). This line of research conceptualizes of attitude strength as a multi-dimensional concept with diverse antecedents and suggests people vary in the strength they attach to different attitudes.

These approaches contrast with measurement error models that assume changes over time are a function of the question being asked, not the people answering questions. While initial assessments argued nothing predicted stability (Achen 1975; Erikson 1979), subsequent research finds individual characteristics such as political awareness are predictive of attitude stability, suggesting the source of instability lies in people (Freeder, Lenz, and Turney 2019; Zaller 1992).

Converse and Zaller's theories focus principally on political issues, but the public contradictions, individual cognitive limitations, and social features that facilitate both stability and instability are not confined to the political domain. There is no reason to assume a model in which some people hold stable opinions while others express ambivalence in the form of vacillating attitudes does not describe the behavior of a range of beliefs across topics.

A "some do, some do not" model of attitude stability closely aligns with work in cultural sociology suggesting social structures – institutions, contexts, and public culture – are principally responsible for shaping attitude structure and constraint (Lizardo and Strand 2010; Martin 2002, 2010; Rawlings 2020). If social reinforcement and scaffolding through public culture is a necessary condition for holding consistent beliefs, then we would not expect everybody to be equally exposed to these processes. Some people are going to be in environments where they hear contradictory messages and

considerations, while other people are in environments where they hear a single line of reasoning that makes holding beliefs relatively easy (Zaller 1992).

### **3.3 Expectations and Hypotheses**

The preceding discussion suggests the question “do people have stable beliefs or do they lack stable beliefs?” is a false dichotomy, rooted in an outdated understanding of attitudes. The central features of Zaller’s political ambivalence model – heterogeneous messages, internalization of conflicts, and weak motivation to reconcile inconsistencies – are present across domains. At the same time, some people will be in structures that facilitate stable attitudes. As a result, I expect on any particular issue measured in a survey, some people will hold strong, stable attitudes, while other people will demonstrate weak, ambivalent attitudes. This generates my central hypothesis.

*Hypothesis 1: On any particular issue measured in surveys, some people will demonstrate stable attitudes and others will demonstrate vacillating responses.*

Stable opinions and vacillating attitudes are not the only potential response patterns. People might make real, durable changes in attitudes over time as they move into new social structures. If the strong/weak attitude distinction underlies attitude behavior, durable changes in attitudes should be much less common than temporary changes. As Howe and Krosnick note, “Attitudes that can be easily changed are weak and unlikely to shape behavior. The attitudes that most powerfully shape behavior are the hardest to change” (2017: 328). Previous work suggests people form dispositions

early in life and carry them over time, rather than update attitudes over time (Kiley and Vaisey 2020; Vaisey and Lizardo 2016). This means we should expect very few people to durably change their opinions.

*Hypothesis 2: Durable change will be substantially less common than either vacillating change or stable attitudes.*

An exception to Hypothesis 2 should be made for younger people. Existing work suggests that people under 30 are likely still forming durable dispositions they will carry with them over time, and as a result should be susceptible to environmental influences (Kiley and Vaisey 2020). If people are still in the process of forming dispositions, they might be more likely to demonstrate durable change than older peers.

There are a handful of other general expectations that emerge from the preceding discussion. Questions asked in social science surveys tend to be asked because people disagree on these issues, meaning people will have consumed heterogeneous arguments in favor of and against whatever issue is being asked. Because screening out competing considerations takes cognitive work, time, and attention, which are limited resources, I expect that people can only be stable on a few issues at a time. This leads to the third hypothesis:

*Hypothesis 3: Across survey questions that researchers tend to ask, vacillating attitudes will be more common than stable attitudes.*

If attitudes are characterized by the strong and weak attitude dichotomy (Howe and Krosnick 2017), with strong attitudes being relatively impervious to temporary influences and weak attitudes being relatively susceptible to these influences, we should see this distinction manifest in responses, with people who give “strong” responses being less likely to change their attitudes over time.

If the measurement error model is correct, the probability of being stable over time on one question should not be related to stability on other questions. The issue publics model posits that people should demonstrate stability on related issues. Stable attitudes on one question of general morality should be correlated with stable opinions on other questions of general morality. Stable opinions on one political issue (opposition to abortion) might span to general political ideology (conservatism), but not necessarily to other political issues, such as opposition to government spending (Baldassarri and Gelman 2008). This suggests the fourth hypothesis:

*Hypothesis 4: People are will be more likely to demonstrate stability on related issues than unrelated issues.*

### **3.4 Data**

The preceding hypotheses make claims about attitudes in general. While it is possible to take a random sample to generalize to a population of people, it is unclear how to randomly sample to draw inferences to a larger population of attitudes. In the absence of a random sample, I follow Alwin (2007) and cast a wide net to capture a

broad range of attitudes and beliefs. In searching for questions to test, I employed two major criteria. First, the method outlined below requires observing the same person responding to the same question a minimum of three time. Second, questions need to tap some form of belief, attitude, value, or self-assessment, rather than statements of fact (Alwin 2007).

A search of multiple panel data sets produced 544 questions that met these criteria. For this analysis I use data from the General Social Survey's rotating three-wave panel (166 questions)<sup>2</sup>; the American National Election Studies panels from 1956 to 1960 (10 questions), 1972 to 1976 (71 questions), 1992 to 1997 (66 questions), and 2000 to 2004 (57 questions); and the Cooperative Congressional Election Study (36 questions). I also use data from the National Study of Youth and Religion (45 questions), two panels – one of mothers, one of their children – from the Intergenerational Study of Parents and Children (17 questions in the mother panel, 18 in the child panel), and the Longitudinal Study of Generations (58 questions). These data sources have distinct advantages when compared with each other. The GSS, ANES, and CCES panels are representative samples of U.S. adults over the age of 18. The NSYR is a nationally representative sample of young people and includes questions about religious and moral issues the adult panels lack. The LSG and ISPC are not nationally representative samples, so it is hard to make

---

<sup>2</sup> The GSS panels were collected over an eight year period as a series of three rotating three-wave panels. Each panel extended over a four-year period (2006-2008-2010; 2008-2010-2012; 2010-2012-2014). Because of the significant overlap in time periods, I treat the GSS panels as a single panel.

inferences about the larger population from them, but they include distinct sets of attitude items. The GSS, CCES, ANES panels, and ISPC all have only three waves. The NSYR has four waves, though many questions are only asked in three. The LSG has a total of eight waves, with some questions asked in three. In these latter two data sets, I limit my attention to three waves for comparability.

The GSS comprises the widest range of topics, including attitudes about politics, race relations, gender roles, religion, morality, social trust, and civil liberties, among other topics. The ANES and CCES panels focus on political issues but also include assessments of people and groups, measures of social trust, and some measures of religious beliefs. The NSYR focuses on religious beliefs, moral beliefs, and self-assessments. The ISPC focuses on gender and family structure. The LSG contains questions about religious beliefs, family relationships, gender roles, and values.

The panels also cover different time intervals. The GSS, ANES, and CCES panels cover about four years. The ISPC panel covers 13 years. Other panels fall between these. In general, I suspect stable opinions will be easier to hold over shorter time frames, so panels covering a longer time frame will see more vacillation and durable change.

It is impossible to say whether the attitudes captured in the 544 questions analyzed here are “representative” of attitudes in general. Because of the nature of the surveys and the kinds of topics get covered in panel studies that extend at least three waves, political views make up a large proportion of questions in this analysis.

However, the data sets – particularly the GSS – contain a broad array of questions on other topics including beliefs about race relations, gender roles, family structures, moral worldviews, intergroup relations, and social and institutional trust. Presenting a full list of all 544 questions and their wording would be cumbersome. A full list of questions, their wording, any modifications made, and coefficient estimates from the model outlined below is available in an online supplement.

### **3.5 Methods**

#### **3.5.1 Measuring Opinion Behavior**

Hypotheses 1 through 3 require estimates of the proportion of survey respondents who demonstrate each of three opinion behaviors: stable opinions, vacillating change, and durable change. Hypothesis 4 requires an individual-level probability of attitude stability. Most methods for measuring opinion behavior conflate two of these groups or avoid quantifying their prevalence, especially at the individual level. Kiley and Vaisey's (2020) methods compare the relative prevalence of vacillating changers and durable changers without quantifying these groups or the proportion giving stable responses. Measures of how many people change their responses over time, the average level of change in responses, or how many people remain stable over time tend to conflate people who make durable change and people who vacillate. Methods measuring correlations over time assume members of the population follow the same process and treat questions as reflecting continuous scales (Alwin 2007; Hout

and Hastings 2016). These methods also assume a real position exists underneath measurement error (Achen 1975; Ansolabehere, Rodden, and Snyder 2008), which may not be a valid assumption.

Latent class methods can deductively group cases such that their manifest variables are statistically independent (McCutcheon 1987). Using such a method, we might be able to quantify how many people consistently give “agree” responses, “disagree” responses, and responses that change from one side to the other (similar to Taylor 1983). The approach outlined below is similar to such approaches but groups people on the basis of an underlying data-generating process, rather than on manifest responses.

An ideal approach would formally specify the three opinion behaviors and quantify the probability a person’s set of responses came from each of these models. Such a model should respect the nominal structure of the data, especially the unique nature of “no opinion” or “don’t know” response options and response options that explicitly express ambivalence. And it should allow for a comparison with the measurement error model. Hill and Kreisi outline a finite mixture model congruent with these requirements and use it to estimate the prevalence of these three kinds of response behavior to questions about pollution-abatement policies in a sample of Swiss residents (Hill and Kriesi 2001a, 2001b; Hill 2001). This model formalizes the three distinct theoretical data-generating processes and estimates the probability each one generated a

set of responses. I outline this model below, starting by formalizing the three behavioral groups.

### 3.5.2 Opinion Behaviors

Consider the question from the General Social Survey's three-wave panels: "Do you agree or disagree with the following statement: Homosexual couples should have the right to marry one another." Respondents are given five options to select from: "strongly agree", "agree", "neither agree nor disagree", "disagree", "strongly disagree." Volunteered responses of "can't choose" and "no answer" are also recorded. Survey respondents answered this question three times over four years. How would the theoretical models outlined above expect people to behave?

**Opinion Holders:** Stable attitude theories expect people to report being on the same side of an issue over time as internalized dispositions motivate their responses or social structures constrain them. With regard to the question above, a person who supports gay marriage might switch between "agree" and "strongly agree" based on recent considerations. But barring actual error in coding the response,<sup>3</sup> it would be

---

<sup>3</sup> Measurement error arises from "vague question wordings, vague response categories or categories that do not reflect the individual's actual attitude, inattentiveness on the part of the respondent, and even typographical errors" (Ansolabehere, Rodden, and Snyder 2008: p 216). There will likely be cases of error where the response is coded incorrectly, but it is unlikely that these cases will be so prevalent as to produce the level of vacillation that we see in some questions.

unreasonable to assume a person who supports gay marriage would declare that they oppose it or “neither agree nor disagree.”<sup>4</sup>

We could summarize this group’s responses over time with only two parameters. First, we would need to know what proportion of people agreed ( $\alpha_1$ ) or disagreed ( $1 - \alpha_1$ ) with a statement, an unchanging disposition for each person, and then, in a particular wave, draw whether they give a strong ( $\delta_1$ ) or weak response.<sup>5</sup> Theoretically,  $\alpha_1$  would represent the balance of forces in society that create stable opinions in favor of an issue, such as socializing and constraining institutions.

**Vacillating Changers:** Inconsistent attitude theories posit people constructing new opinions in each wave on the basis of indiscriminately consumed considerations (Swidler 1986; Zaller 1992). In our example, these people have likely heard arguments both in favor of and opposed to gay marriage but have not taken time to reconcile them. When asked to give an opinion, they haphazardly call forth some of these considerations. Because these people are assumed to be disconnected from public debates, these influences on responses are local and assumed to be random.

---

<sup>4</sup> It is debatable whether “neither agree nor disagree” represents ambivalence, as I assume here, or a nuanced position where a person thinks that homosexuals should have the right to marry in some circumstances and not others. If it is the latter, then people should maintain this position over time. However, of the 301 people in the GSS panel who report this position in wave 1 and respond in each wave, only 31 maintain this position across all three waves. Results below further argue against this position.

<sup>5</sup> For simplicity, the model makes the assumption that people expressing opinions are equally likely (within opinion behavior groups) to give them strongly. This assumption can be relaxed.

What matters in these theories is not people's dispositions but the balance of considerations in the social environment they might have internalized over time. We would only need to know the probability a person gives a "no opinion" response ( $\varphi_2$ ), the probability they give an opinion that is some version of "agree" ( $\alpha_2$ ), and the probability that, if they give an opinion, they give it strongly ( $\delta_2$ ), which should be roughly the same across all people. Under this model, these parameters reflect the distribution of considerations in society that people can draw on to give responses, including the social factors that might make admitting to no opinion acceptable. Because their opinions are statistically independent across waves, each person's response in each wave would be a multinomial draw of the five response options ("strongly agree", "agree", "neither agree nor disagree", "disagree", "strongly disagree") defined by those three parameters.

**Durable Changers:** The final group comprises two sets of opinion behaviors that fill in gaps between the other two groups. The first comprises people who truly hold opinions but are influenced by events in the public sphere to change their attitude. On the gay marriage example, this might include the Democratic party's evolution on gay marriage. These people should go from one side of the issue to the other over the course of the survey – saying "disagree" or "strongly disagree" in wave 1 or waves 1 and 2 but

“agree” or “strongly agree” for the remainder.<sup>6</sup> Second, some people might lack an opinion on the issue and say “neither agree nor disagree” in the first wave, and, upon being asked to consider the question, develop an opinion for the next wave, what is commonly called “panel conditioning” (Oh, Yeatman, and Trinitapoli 2019; Warren and Halpern-Manners 2012).

The defining feature of this group is that they change opinions once. Response patterns that switch twice or do not switch at all would be incompatible with this model. Because this group is defined by this switch, parameters summarizing their behavior revolve around where they start, where they end, and when during the survey they change, which might also affect how people shift. The first parameter,  $\tau_3$ , captures the probability of change after the first wave instead of after the second wave. The second and third parameters,  $\varphi_3^{(Pre1)}$  and  $\varphi_3^{(Pre2)}$ , quantify the probability an opinion changer starts by saying they lack an opinion and changes after the first and second waves, respectively. A fourth parameter,  $\alpha_3^{(post)}$  defines the probability these people change to agreeing with the statement. Two more parameters,  $\alpha_3^{(Pre1)}$  and  $\alpha_3^{(Pre2)}$ , define the probability someone who changes from one side of the scale to the other starts by

---

<sup>6</sup> It is plausible that in a long enough survey interval, a person might truly change their opinion twice. The surveys explored here tend to look at intervals of about four years, reducing my concern for this possibility. A change that reverts in two years does not seem very “durable,” and seems to reflect more local influences than real change.

agreeing with the statement. Finally, a single parameter,  $\delta_3$ , quantifies the probability someone gives a strong opinion, contingent on them giving an opinion.

These three models are simplifications of diverse causal pathways, but they represent the distinct theoretical processes that cultural sociologists and others have argued produce attitudes in a survey context.

### **3.5.3 Mixture Model**

Given these groups, how do we estimate their prevalence? Let  $\pi_j, j = 1,2,3$  represent the marginal probability of belonging to each of these three groups, with 1 indicating Opinion Holders, 2 indicating Vacillating Changers, and 3 indicating Durable Changers.

The fundamental challenge of quantifying response patterns is that these labels are unobserved. Some patterns could only be generated by one of the models (e.g., “agree”-“disagree”-“agree” could only be generated by a vacillating changer model), but most could be generated by the vacillating changer model and one of the other groups. This is true even for people who remain on the same “side” of a question over time. Over three waves, a vacillating changer might, by chance, give the same response three times.

If we knew which group each person was in, we could count the number of opinion holders who say they agree with a statement to get an estimate of  $\alpha_1$ . If these parameters were known, we could calculate which theoretical process was most likely to

generate each response. In the absence of directly observing these parameters, I can estimate a range of plausible values of these parameters  $(\alpha_1, \delta_1, \dots, \tau_3)$ , as well as estimates about the proportion of the sample belonging to each opinion behavior group  $(\pi_1, \pi_2, \pi_3)$ , that are compatible with the data we observe and the assumptions made above about these groups.

To do this, I use a Bayesian estimation technique known as the Data Augmentation algorithm, a form of Gibbs sampler that generates this range of plausible parameter estimates. The algorithm has two basic steps. First, it samples the “missing data” – group membership – given a set of randomly drawn parameters.<sup>7</sup> It asks: “given this person’s responses, how common each group is, and how often members of each group give certain response patterns, how likely is it that this pattern came from group 1, group 2, and group 3?” It then samples from those three probabilities and assigns a group to that person. It then uses these group memberships to re-estimate parameters, asking, “given that these people are in group 1, what’s the probability of someone in group 1 saying they strongly agree or strongly disagree ( $\delta_1$ )?” The model draws a plausible value for each parameter from a specified posterior distribution.

---

<sup>7</sup> At the first iteration, there are no group assignments with which to generate parameter estimates, so random values are drawn from the parameter space. These starting values have minimal, logical constraints (for example, the proportion of people in each behavioral group must sum to 1). To ensure that these starting draws do not unduly influence results, I wait for multiple chains with different starting values to converge to the same posterior distribution before using draws to summarize the distribution.

Posterior distributions combine the distribution of the data, given the unknown parameter values, with a prior distribution of the parameters, which quantifies belief about the parameters before seeing any data. I use minimally informative priors outlined in Hill and Kreisi (2001b), the equivalent of adding two additional people to each opinion behavior group and splitting them up among behaviors within these groups.<sup>8</sup> The non-informative priors have the effect of generating slightly conservative parameter estimates. This effect is most clearly illustrated with the durable changer group. Even in iterations where nobody in the sample is assigned to the “durable change” group, which happens quite often, the model still assumes two people display this behavior, inflating the proportion estimates away from 0.

The key to the estimation procedure is that this process – assign groups based on parameters and estimate parameters based on group assignments – iterates until it converges on a stationary distribution summarizing plausible values of the parameters given the observed data and the model assumptions. I can then draw values from repeated iterations in this distribution to summarize it.

Variation in parameter estimates comes from ambiguous patterns (e.g., “agree”-“agree”-“disagree”) being alternately assigned to different behavioral groups (e.g., durable changer or vacillating changer). In practice, the range of plausible parameters

---

<sup>8</sup> Hill and Kreisi test a range of non-informative priors – comparable to between one and six additional people per behavioral group – and find little sensitivity to different non-informative priors.

for each question tends to be quite small. People who never change can never be assigned to the durable change group, and people who change cannot be stable opinion holders. For any question, several response patterns can *only* be classified as vacillating changers, specifically people who “revert” – go from one side of the scale to the other and back again during the three waves – and people who switch to “no opinion” responses in later waves. These people constrain the range of parameters for the vacillating changer group, which shapes the likelihood that ambiguous patterns get classified as vacillating changers. Because responses are assumed to be independent across waves in the vacillating changer group, each “agree”-“disagree”-“agree” reversion in the data set increases the probability that an “agree”-“agree”-“disagree” was also generated under the vacillating change model rather than the durable changer model. If these restricted patterns are very common, the probability that ambiguous patterns were also generated by the vacillating changer group increases. If they are rare, the probability that ambiguous patterns will be classified as opinion holders or durable changers increases. Full details of the estimation procedure are outlined in Appendix B. More details can be found in Hill and Kreisi (2001b).

The model Hill and Kreisi (2001b) outline is designed to estimate parameters on four waves of data, but the model is compatible with three waves of data with minor modifications. The Longitudinal Study of Generations and the National Study of Youth and Religion include more than three waves of data for several questions. Comparing

models estimated on four waves of data to the modified model estimated on only three of waves dropping an internal wave (wave 2 or 3) produces almost identical parameter estimates, especially for group size parameters ( $\pi_1$ ,  $\pi_2$ , and  $\pi_3$ ), affirming the model's ability to generate valid inferences using three waves.

### **3.5.4 Missing Data**

Across the panels, people often lack responses to questions in at least one wave. I assume these responses are Missing at Random, meaning once I condition on a person's observed responses, responses are assumed to be uncorrelated with missingness. Hill (2001) finds the distribution of parameter estimates is similar under a variety of assumptions, with the proportion of vacillating changers slightly higher under less restrictive assumptions than the one made here. Accommodating this assumption adds a step to each iteration: filling in missing data on the basis of group assignment and the parameter estimates. If a person is assigned to the "stable opinion" group and expresses either agree or strongly agree in their non-missing waves, their missing response is filled in from a draw of "strongly agree" and "agree" with probabilities  $\delta_1$  and  $1 - \delta_1$ .

The model iterates through these three steps – assign group membership on the basis of responses and estimated parameters; estimate parameters on the basis of group membership and responses; and fill in missing data on the basis of parameters, group membership, and observed data – until iterations reach a stable posterior distribution.

### 3.5.5 Model Modifications

The model outlined above assumes a five-point response scale. I modify the model and the questions being examined to accommodate additional question structures. Several response scales lack a clear “no opinion” midpoint. I assume that in the absence of an explicit “no opinion” option people to select a second-best option. This would lead to a decreased value of the  $\varphi$  parameter, but it should not affect the distribution of the opinion group proportions. For some question, midpoints are not labeled “no opinion”, but rather some other phrase. In cases where the substantive interpretation of this midpoint option is equivalent to “no opinion” or ambivalence, I employ the model as it is outlined above. I collapse the response options into a binary response structure, discussed below, in cases where the midpoint has a substantive interpretation other than “no opinion.”

I remove the  $\delta$  parameters from the model to accommodate common two- and three-point scales. The lack of “strong” response options has the consequence of greater uncertainty about whether a response pattern falls into one behavioral group or another, but the model still generates stable group proportion estimates.

To handle seven-point scales by I add an additional parameter ( $\gamma_j$ ) to account for “weakly” agree and “weakly” disagree or comparable points (response 3 and response 5 on a seven-point scale). For most questions, I maintain the assumption that the scale midpoint reflects a lack of opinion rather than a substantive “moderate” position. If

people hold substantive “moderate” positions, it will show up as lower average within-person variance for vacillating changers responding to these questions, which I discuss below.<sup>9</sup> Some questions have much larger scale ranges, including 10-, 12-, or 100-point feeling thermometers. When these questions have symmetrical structures, I collapse them questions into five-point scales.<sup>10</sup>

Finally, I collapse non-symmetrical scales to a binary response structure. For example, the GSS includes a question about people’s interpretation of the Bible: “Which of these statements comes closest to describing your feelings about the Bible? a. The Bible is the actual word of God and is to be taken literally, word for word. b. The Bible is the inspired word of God but not everything in it should be taken literally, word for word. c. The Bible is an ancient book of fables, legends, history, and moral precepts recorded by men.” I collapse this into a binary question of whether a person chooses Biblical literalism (choice a) or not. This approach sacrifices nuance for direct comparability with other questions. In these cases, I try to focus on the most substantive division in the structure.

---

<sup>9</sup> Kinder and Kalmoe suggest that on seven-point political ideology scales, moderates are “indistinguishable from those who remove themselves from ideological analysis” (2017: 160) and that we should regard these people as people who lack opinions, not people who hold sophisticated positions in the middle of the distribution. This strengthens my confidence on using the midpoint as a non-attitude point, rather than a moderate position.

<sup>10</sup> On 100-point scales, I consider respondents giving responses of less than or equal to 20 and greater than or equal to 80 as giving “strong” responses. Respondents who say 50 are considered to give an ambivalent or no opinion response.

### **3.5.6 Comparison with Measurement Error Model**

The model outlined above cannot explicitly rule out the measurement error model, which assumes the population comprises one group with different latent positions and responding with normally distributed errors (Alwin 2007). If this were true, people in each group should demonstrate similar levels of within-person change. Opinion holders would have a latent position near the ends of the scale (between 1 and 2 or between 4 and 5), vacillating changers and durable changers would have latent positions in the middle of the scale (between 2 and 3 or between 3 and 4), and all groups would respond with similar average errors. People with (random) large variation would have to be in the vacillating changer or durable changer group, but they should be relatively rare so that, on average, the groups' errors are comparable.

I follow Hill and Kreisi (2001a) and estimate the average within-person standard deviation for each behavioral group at each iteration of the model. If groups have similar average within-person standard deviations, then I cannot rule out the measurement error model. If these groups have substantially different average within-person standard deviations, it is less likely the data was generated through a measurement error model.

### **3.5.7 Testing Issue Publics**

To adjudicate whether patterns of stability are consistent with the "issue publics" theory, as Hypothesis 5 suggests, I calculate the individual-level probabilities of belonging to the opinion holder group over the model's 10,000 iterations. I then calculate

the pairwise correlation of these probabilities across the questions in the GSS panels. The GSS covers the widest range of questions, giving it the greatest room to detect clusters of questions that demonstrate related stability. If the measurement error model is correct, there should be little correlation between stability on different questions. If stability is principally a function of individual characteristics (education, socioeconomic status, intelligence), these correlations should be relatively high. If the issue publics model is correct, I should observe higher correlations on related issues than unrelated issues.

### **3.6 Results**

My discussion of results proceeds in three stages. I first demonstrate the model and the interpretation of its estimates using the GSS question about gay marriage discussed above. I then explore parameter estimates across all 544 questions. I then explore the correlations of stability within the GSS questions to evaluate whether the pattern of stability can be thought of as representing “issues publics.”

#### **3.6.1 Gay Marriage**

Table 1 summarizes the parameter distribution of the gay marriage question. It presents the mean and the central 95 percent of values most compatible with the observed data. The mean and distribution of values are drawn from 10,000 iterations of the finite mixture model once the model has converged to a stationary parameter distribution (see Appendix B).

**Table 1: Finite mixture model parameter estimates for gay marriage question (2006 to 2014 GSS).**

Parameter Group	Meaning	Parameter	Mean	Interval
Behavior Groups	Opinion holders	$\pi_1$	0.51	(0.49, 0.53)
	Vacillating changers	$\pi_2$	0.44	(0.42, 0.47)
	Durable changers	$\pi_3$	0.05	(0.03, 0.06)
Opinion Holders	P(agree)	$\alpha_1$	0.40	(0.37, 0.42)
	P(strong)	$\delta_1$	0.73	(0.71, 0.76)
Vacillating Changers	P(agree)	$\alpha_2$	0.44	(0.41, 0.47)
	P(strong)	$\delta_2$	0.25	(0.22, 0.28)
	P(no opinion)	$\varphi_2$	0.26	(0.23, 0.28)
Durable Changers	P(no opinion -> agree)	$\alpha_3^{(post)}$	0.40	(0.19, 0.60)
	P(agree -> disagree)	$\alpha_3^{(pre1)}$	0.33	(0.20, 0.46)
	P(agree -> disagree)	$\alpha_3^{(pre2)}$	0.09	(0.00, 0.28)
	P(strong)	$\delta_3$	0.68	(0.54, 0.82)
	P(no opinion)	$\varphi_3^{(pre1)}$	0.57	(0.43, 0.71)
	P(no opinion)	$\varphi_3^{(pre2)}$	0.35	(0.02, 0.60)
	P(change t1 -> t2)	$\tau_3$	0.68	(0.56, 0.81)

As a top-line result, the model estimates that between 49 and 53 percent of the sample give responses consistent with the “opinion holder” model ( $\pi_1$ ), either agreeing ( $\alpha_1 = .40$ ) or disagreeing with the question in all waves, while between 42 and 47 percent of the sample give responses consistent with the “vacillating changers” model ( $\pi_2$ ). Between 3.28 and 6.45 percent of the sample made durable changes in attitudes during the survey window of 2006-2014 ( $\pi_3$ ).

Is about half the population holding a stable view on the gay marriage issue a surprising number? It is significantly higher than the 20 percent of people Converse suggested had stable opinions on the issue of government control of utilities (Converse 1964), and it is roughly comparable to what Hill and Kreisi found when they explored

opinions on pollution-abatement policies (Hill and Kriesi 2001a). That being said, it was a dominant political issue of the time, and the fact that more than four in ten people demonstrated response patterns suggesting they lacked durable opinions on the issue is notable.

Other estimates in Table 1 are consistent with expectations about attitude change and stability. People who hold stable attitudes are more likely to report “strong” attitudes than people who hold vacillating attitudes (Howe and Krosnick 2017). Nothing in the model requires this; it is a function of people who do not give strong responses changing between sides of the scale over time. Vacillating changers said “neither agree nor disagree” about a quarter of the time. When they gave an opinion, they were about evenly split between agree and disagree ( $\alpha_2 = .44$ ).

Consistent with the historical record, people who made durable change tended to change from opposing gay marriage ( $\alpha_3^{(pre1)} = 0.318$  and  $\alpha_3^{(pre2)} = 0.081$ ) to supporting it. They were also more likely to change between the first and second wave ( $\tau_3 = .68$ ) than between the second and third wave, consistent with theories of panel conditioning.<sup>11</sup> Durable changers more closely resemble opinion holders than vacillating changers in opinion strength, though there is more uncertainty in estimates for the durable change group, given its small size.

---

<sup>11</sup> The full panel collapses three three-wave panels, so wave 1 actually represents three two-year periods (2006-08, 2008-10, and 2010-12). This makes it difficult to draw conclusions about timing of change, though analyzing the panels separately can provide clarification.

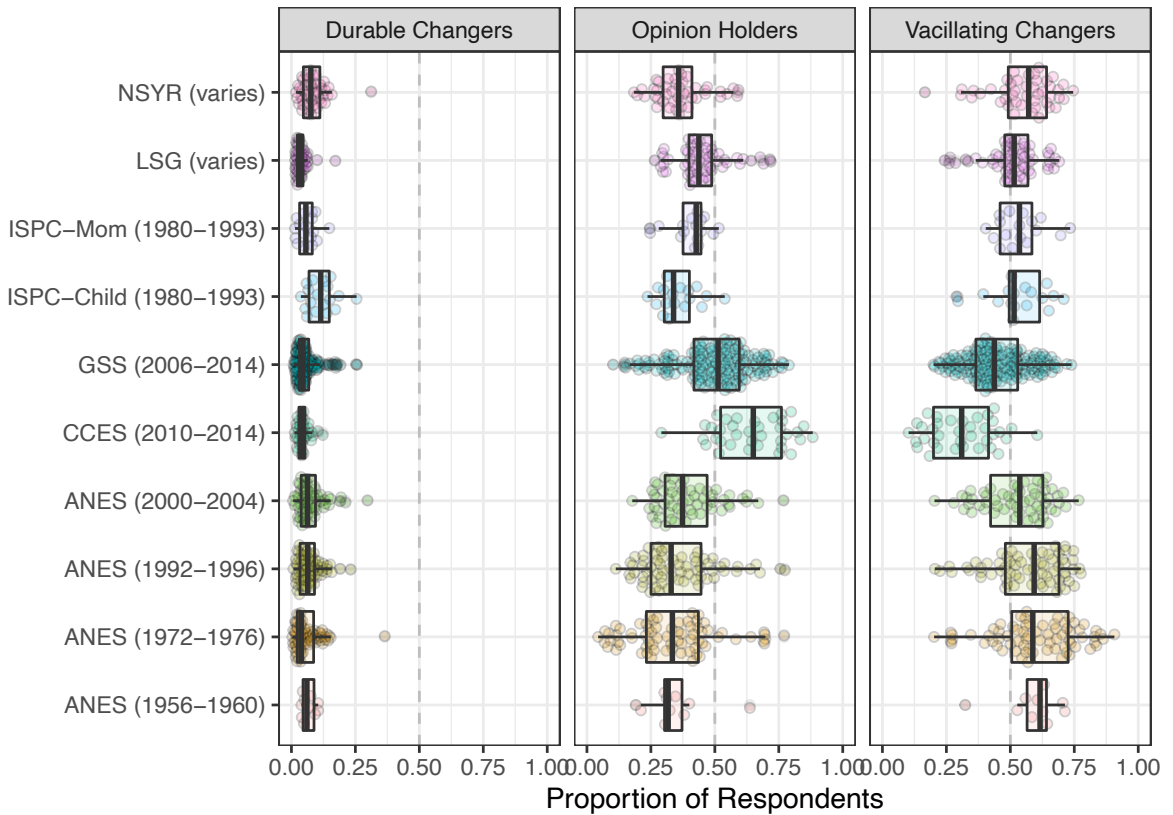
Table 2 shows that average within-person standard deviations are different across groups. Opinion holders have a small average standard deviation ( $\sigma_{G1} = 0.303$ ), vacillating changers have average within-person standard deviations about three times as large as opinion holders ( $\sigma_{G2} = 0.909$ ), and durable changers have an average standard deviation larger still ( $\sigma_{G3} = 1.469$ ). Vacillating changers and opinion holders do not reflect similar behavior at different points in the scale. Vacillating changers demonstrate much more wave-to-wave change than opinion holders, choosing strong responses more often than we would anticipate if they had similar error ranges to opinion holders. People who make durable change tend to go from one end of the spectrum to the other, a behavior very different than vacillating changers. As Hill and Kreisi note in their examination of the pollution-abatement data, “Given the implausibility of assuming vastly different measurement errors for different people, it seems rather more likely that this variation is composed of both measurement error and true opinion instability” (2001a: 408).

**Table 2: Average within-group standard deviation, by opinion behavior group.**

Behavioral Group	$\sigma$	Plausibility Interval (95 percent)
Opinion Holder	0.30	(0.295, 0.312)
Vacillating Changer	0.91	(0.885, 0.932)
Durable Changer	1.47	(1.308, 1.631)

### 3.6.2 Full Data Results

I now turn to examining patterns across the 544 questions outlined above. Figure 7 presents a series of boxplots, three for each data set, plotting the distribution of mean values from the parameter distribution for the proportion of durable changers, vacillating changers, and opinion holders for each question. The figure contains a lot of information, but it highlights clear patterns across the data sources that speak to the hypotheses outlined above.



**Figure 7: Mean proportion of respondents in each behavioral group, by question and data set.**

Starting at the left, durable change is rare across all data sets and almost all questions. Consistent with Hypothesis 2, for about half of all questions, the mean estimated proportion of respondents demonstrating durable change is less than 5 percent, with many of these questions showing rates indistinguishable from 0.<sup>12</sup> Only 17 percent of questions have mean rates of durable change above 10 percent. Attitudes that exhibit high levels of durable change typically pertain to high-profile public issues, and they are exceptions that prove the rule that durable change is, in general, rare.

There does not appear to be any substantial effect of panel length on rates of durable change. While the mother panel of the ISPC covers a 13-year span, and some questions from the LSG cover windows of more than 10 years, these panels do not show any higher rates of durable change than the panels covering only four years.

Most of the data sets analyzed here include questions that demonstrate both high (greater than 70 percent) and low (less than 30 percent) levels of stability and vacillation. Since within-survey comparisons reflect the same sample, vacillation is not a feature of the population (the kinds of people being surveyed) but an interaction between people taking the survey and the kind of question being asked. In other words, there are some questions that most of the sample gives clear opinions over time, and there are questions on which very few people in the sample articulate clear opinions over time. However,


---

<sup>12</sup> Because of the conservative nature of the estimation process, the estimate for the proportion of durable changers can never be 0. But most of these questions have confidence intervals that extend below 1 percent and many iterations for these questions assign no respondents to the durable change group.

most questions fall somewhere in the middle, with some people answering consistently and others not. This provides broad support for Hypothesis 1.

Third, consistent with Hypothesis 3, the vacillating changer group tends to be larger than the opinion holder group for most of the questions analyzed. The 1956-60 ANES panel, which Converse analyzed to generate his original insights about non-attitudes, displays the highest average level of vacillation, but it is not an outlier. Almost all other surveys include questions that surpass that average estimate of ambivalence. The GSS and the CCES are the only data sets where more questions show higher levels of stability than vacillation.

Figure 8 plots the proportion of opinion holders and vacillating changes who give strong responses, conditional on them expressing an opinion (not saying “no opinion”/ “neither agree nor disagree”), for questions with four, five, or seven response options. Questions with two and three response options are excluded here as they lack “strong” response options. Points above the diagonal line indicate questions where opinion holders give “strong” responses at a higher rate than vacillating changers, conditional on giving an opinion.

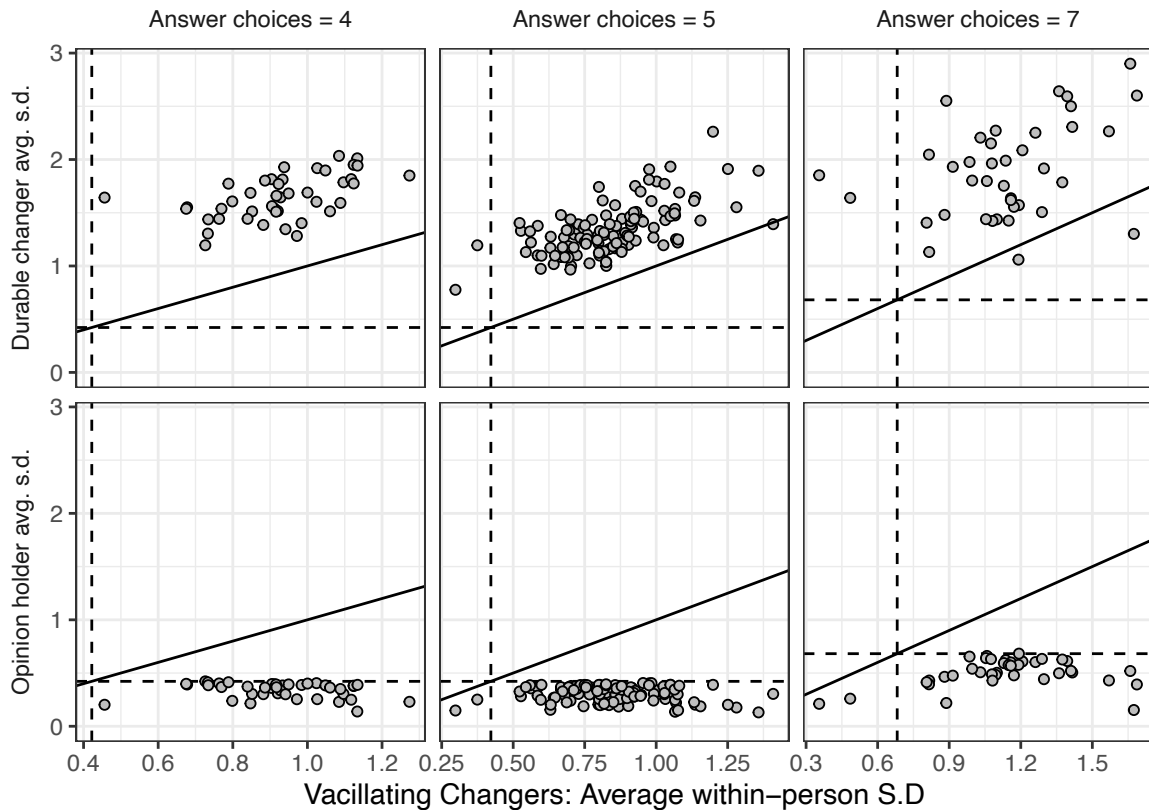


**Figure 8: Mean estimated proportion of "strong" responses in option holder and vacillating changer behavior groups, by question.**

Opinion holders are more likely than vacillating changers attitudes to respond with a strong position, often much more likely. Even on questions where few people overall give strong responses, opinion holders are still more likely to give them than vacillating changers. Vacillating changers tend to say that their opinions are weak and, as we would expect of weak opinions, switch sides of the scale. At the same time, vacillating changers still select the "strong" option, often quite frequently, suggesting that they tend to have a larger response range than opinion holders.

To compare the three-class model to the measurement error model, Figure 9 presents two sets of plots that compare the mean estimate of average within-person

standard deviations for different behavioral groups. The top row compares vacillating changers to durable changers, and the bottom row compares vacillating changers to opinion holders. Because opinion holders always have a within-person standard deviation of 0 on questions with 2 or 3 response options, these questions are omitted. The solid black line indicates equal within-group standard deviations. Points below the solid diagonal line indicate questions where vacillating changers have greater within-person variance than the comparison group (durable changers in the top row, opinion holders in the bottom). Dashed lines indicate the largest within-person standard deviation for stable opinion holders.



**Figure 9: Average within-person standard deviation, by behavioral group and question response options.**

Consistent with the gay marriage example presented above, vacillating changers uniformly have higher within-person standard deviations than opinion holders for questions with four or more response options, regardless of question structure, and almost uniformly lower within-person standard deviations than durable changers.<sup>13</sup>

Durable changers appear to move from one strong attitude to another while vacillating

---

<sup>13</sup> The three questions where durable changes have lower average within-person standard deviation than vacillating changers ask respondents if the government should bus students to ensure racial equality (72-76 ANES), whether the government should do more to improve the conditions of blacks (92-96 ANES), and whether they are happy with their bodies (NSYR).

changers cluster in the middle of the scale. In other words, we can be fairly confident the model is separating three different groups with distinct sets of opinion behavior.

As an overall conclusion, questions follow a similar pattern. Vacillating changers and opinion holders are present for all issues, consistent with Converse's original black and white model, and there is often, though not always, a small number of people making durable change. As expected, opinion holders exhibit strong opinions at the ends of scales. Vacillating changers have a much wider range of responses that vacillate on either side of the issue. Durable changers move between ends of the scale, often from one extreme to another. We can say that overall, durable change is rare, but there is large variance in stability and vacillating change that remains to be explained.

It is difficult just looking at the overall patterns to say whether question content or structure drives overall rates of stability or vacillation. The GSS has more questions with fewer response options (questions with two, three, or four choices make up most of the GSS questions), but also more questions about non-political content. ANES questions tend to have more response options but focus almost exclusively on politics. Appendix C outlines a regression approach to evaluate the competing influences of structure and content on attitude stability and change.

Structure and content explain about 50 percent of the variation in rates of vacillation. Questions with fewer response options show higher stability, net of question content. Questions about sex and religion show slightly higher rates of stability than

questions about politics, and questions about general morality and civil liberties show higher rates of vacillation than other questions, net of question structure. Questions about race, foreign relations, work, crime, and family and gender are not more or less stable than general political questions once I account for question structure.

There are some questions whose proportions of vacillating changers and opinion holders are poorly explained by these models. Across data sets, partisan identification is more stable than other measures of general political ideology (position on liberal-conservative scale; role of government in economy; role of government in equal opportunity, etc.). Questions about specifics – specific policies and specific behaviors – tend to be more stable than predicted while questions about general concepts (ideology, broad morality, etc.) tend to vacillate more than predicted.

Content explains less variation in rates of durable change, and there is no clear pattern in terms of question structure. The NSYR and the child panel of the ISPC, the two panels with the youngest average age, show higher rates of durable change than other panels, even controlling for questions structure and content. The largest residual in the durable change model pertains to people's evaluation of Richard Nixon in the 1972 to 1976 ANES panel, in which a large proportion of people shifted from a positive opinion or no opinion to a negative opinion. Given that this window covers from Nixon's victory in the 1972 election through the Watergate scandal and his eventual resignation, this shift is not surprising.

### 3.6.3 Issue Publics

The 166 questions in the GSS panels examined here produce 13,695 pairwise correlations.<sup>14</sup> In general, pairwise correlations between rates of stability are very low, with an average correlation of .04, suggesting stability in one attitude is not a strong predictor of stability in another attitude and that stability is not a function of the person answering the question. At the same time, some pairwise correlations are quite strong. To better understand the distribution of these strong correlations, I plot them as a network diagram in Figure 10. For parsimony, I focus on the 541 correlations where  $\rho > .2$ , about 4 percent of all correlations. Ties between issues indicate that people who are stable on one issue are stable on the other, and that people who vacillate on one issue vacillate on the other. The 58 questions with no  $\rho > .2$  are held out of the figure. The isolates include a heterogeneous mix of questions about gun laws, whether there is an afterlife, whether premarital sex is acceptable, and more.

---

<sup>14</sup> There are seven federal spending questions where different versions are tested for different people, examined separately above. The two versions show minimal differences, so I consolidate each pair into a single item for this analysis.



questions of police use of force (polhitok), abortion (abhlth), and the values that should be demonstrated by children (obey, helpoth). These questions tend to reflect the application of general standards across specific instances, and preliminary analysis suggests that higher levels of education are associated with stability on issues in this cluster. This component is disconnected from a second large cluster about political views, including ideology (polviews), partisan identification (partyid), and views on government and race (helpblk, affact). Other “issue publics” connect issues on government spending priorities (nacity, natrace, natdrug), confidence in institutions (conlegis, coneduc), beliefs about gender (fefam, fepresch), and beliefs about science (advfront, natci).

Notable in the figure are the lack of connections between certain issues. Despite having a similar structure and appearing next to each other in the survey, two sets of questions about abortion are disconnected – holding a stable belief about abortion in some circumstances (child defect, women’s health, and rape) is not related to a stable position on abortion in other circumstances (single mother, as a method of birth control, if the mother is poor, and any circumstance the woman wants). The same is true of beliefs about euthanasia and suicide. Stability in general political beliefs (liberal-conservative ideology, partisan identification) is disconnected from a broad set of questions about government spending. Questions about generalized trust are disconnected from specific questions about confidence in institutional leadership, and

questions about specific morality are disconnected from questions about general morality.

### **3.7 Discussion**

The results presented above suggest several conclusions about the behavior of attitudes, at least the attitudes that get measured in social science surveys. First, it is wrong to say that people do not hold durable opinions on these kinds of issues. On most questions explored here, there is a sizeable proportion of the sample (more than 30 percent and sometimes as high as 70 percent) that consistently locates itself on one side of an issue or another in ways that cannot be attributed to chance. When people hold opinions stably, they are disproportionately likely to say they hold these opinions strongly, a finding consistent with psychological models that distinguish between strong and weak attitudes (Howe and Krosnick 2017). People are more likely to give consistent answers to simpler questions – those with fewer answer choices or about specific rather than general issues – but some proportion maintain stable attitudes even on complicated, abstract, conceptual questions. And when people give consistent responses to a question, they tend to also give consistent responses to issues we expect to be related.

At the same time, it is wrong to say that people in general “have” durable opinions. There is typically a large group of people who do not maintain consistent opinions on any particular issue, vacillating between ends of the scale and in their

extremity in ways that suggest they are subject to short-term considerations. Differences in the average within-person standard deviations for opinion behavior groups suggests people who vacillate are not just people at the middle of the scale who report real attitudes with error. They often span the scale more or less at random. These patterns are strongly consistent with Zaller's argument that on any particular issue, "people are likely to internalize many contradictory arguments" and "form considerations that induce them both to favor and oppose the same issues" (1992: 59), as well as "toolkit" theories of culture that suggest that people store heterogeneous considerations and can be shifted to either side of an issue through local influences, including their own desire to justify their behavior (Scott and Lyman 1968; Swidler 2001).

This inconsistency is not just present when people are asked about politics, where ambivalence has most frequently been demonstrated, but it is present at comparable rates when people are asked about general morality, racial issues, evaluations of public figures and groups, and more. Many of these issues, especially morality, are believed to be "core" dispositions that are relatively invariant over time and shape other beliefs. This seems unlikely if people change these views frequently. With the exception of questions about religious beliefs and sex, people were not more likely to be consistent on non-political issues than political issues. In fact, the regression results suggest that when issues in other domains – race, social trust, medicine, etc. – intersect with the political space, people are more likely to report the same opinion

consistently. This suggests that public debates help crystallize people's views on issues, regardless of domain (Martin 2002; Rawlings 2020; Zaller 1992). In other words, people demonstrate stability when they get clear signals about what to say.

While stability is not strongly patterned by topic domain, there are general patterns that cut across domains. Across topic domains, attitudes about specific issues, policies, and people tend to be more stable than general attitudes. People are more likely to be consistent on specific moral prohibitions (e.g., extramarital sexual relations) than on general questions about morality. They are more stable on specific questions about government spending on various priorities (e.g., health care) than on general questions about the role of government and political ideology. This suggests people work inductively from specifics to general principles and have a hard time reconciling disparate specifics into general principles. They do not seem to apply general principles to specific outcomes. In the political domain, this is consistent with the notion that people are "ideologically innocent" (Kinder and Kalmoe 2017) and further argues against general or latent beliefs as begetting specific beliefs.

### **3.7.1 Implications for Cultural Sociology**

The results presented above are consonant with a "toolkit" theory of culture (DiMaggio 1997; Swidler 1986), in which people internalize a broad array of culture and draw on different aspects of it to develop attitudes in the moment to justify institutional constraints. In general, people do not appear to internalize one culture that leads them to

consistently report particular beliefs. They have the cultural resources to generate a range of beliefs across circumstances.

Instead of resting stability in the internalization of a single worldview, results presented here reinforce findings that stable belief systems are social products that require external scaffolding – institutions, contexts, and social networks (Lizardo and Strand 2010; Martin 2002). Because people tend to indiscriminately consume information, and because they have little incentive to reconcile conflicts, they should not be thought of as carrying “attitudes” through time. They become stable when some external force ties these disparate beliefs together. Exactly what kinds of social scaffolding, and where in the life course these influences matter most, is unclear from these results. Young people appear to be making durable change at higher rates than older people, suggesting they are more susceptible to environmental influences. It could be the case that people with stable dispositions acquire these early in life, whereas people who do not acquire these early never do. It could also be the case that people who hold stable attitudes do so because they are embedded in contemporaneous social structures (networks, organizations, institutions, etc.) that facilitate them, but changing these structures could make these people to vacillate.

Adjudicating the mechanisms that facilitate stability requires further work. A key benefit of the finite mixture model approach presented here is that it can be extended by adding group-predicting parameters. Researchers could then adjudicate the sources of

attitude stability, regardless of the specific attitudes people hold. For example, it might be the case that people embedded in closed networks of dense ties might be more likely to present stable opinions over time while people bridging structural holes are particularly likely to vacillate. Or it could be the case that people who regularly participate in social institutions such as churches or volunteer organizations receive social reinforcement that produces stability. Or it could be that people who pay attention to charismatic leaders or opinion leaders take their cues from these leaders to develop coherent beliefs.

The low proportions of people demonstrating durable change reinforce the recent finding that people do not commonly make durable changes in their attitudes after they reach adulthood (Kiley and Vaisey 2020; Vaisey and Lizardo 2016). At the same time, high rates of vacillation or ambivalence in many questions challenges the finding that adults have “settled dispositions.” Many survey respondents agree with a statement in one wave and disagree in another, but these changes are not durable. If people tend to be ambivalent toward issues of public importance, then it is possible for their attitudes to be influenced by a consolidation of elite opinion around a topic. However, so long as culture remains heterogeneous, it seems likely that large groups of the population will vacillate.

### 3.7.2 Implications for Measuring Attitudes

Measurement error theories argue that over-time variance in survey responses is principally a function of question structure, but these theories struggle to find features of attitude questions that explain of this variation (Alwin 2007). The results presented here reinforce findings in public opinion research that instability is an interaction of the question and the person answering the question, not just the question itself (Freder, Lenz, and Turney 2019; Zaller 1992). In this sense, people should not be thought of as having a “true score” with random error. They likely have a range of answers they could give under a variety of conditions, a reflection of the considerations they have internalized and call up in the survey context. Some people will demonstrate more or less variance over time. This undermines the central assumption of measurement error theories.

Understood from this perspective, however, variance itself is an interesting property of attitudes. Understanding why some people give consistent attitudes and others do not, and how this variation relates to which attitudes people give, becomes an interesting phenomenon worthy of explanation, rather than “noise” to be discarded. It suggests that repeated measurement of the same attitude over time provides a distinct way of conceptualizing attitudes that can be used to adjudicate theories about how people internalize considerations from the social world.

Researchers invoke the weak but present predictive power of attitudes to argue both for and against a link between attitudes and behaviors (Jerolmack and Khan 2014; Vaisey 2014). The prevalence of weak attitudes – inconsistency that seemingly cannot be attributed to measurement error – should caution researchers away from assuming that cross-sectional measures of attitudes are a good proxy for an underlying belief. For many questions, more than half a sample might be responding with a temporary attitude construct, while others report real attitudes that matter to them. Given this heterogeneity, we should not expect strong predictive ability for an attitude in general, but that tells us nothing about how it influences the people for whom the attitude does matter. “Attitude effects” are going to be shaped by what proportion of the population finds an attitude to be meaningful, the distribution of stable attitude holders in the population, and the actual effect of the attitude on behavior. At the same time, it is not clear how many respondents need to be making stable responses for an attitude to be a good predictor for behavior (a coarsened version of the question Vaisey (2009) used to predict behavior was only stable for about 27 percent of respondents).

### **3.7.3 General Conclusions**

This paper sought to clarify a divide in both cultural sociology and public opinion scholarship between theories suggesting that people lack clear opinions and those suggesting that people hold stable, real opinions. Going back to Converse’s original formulation of the “black and white” model, I suggested that for any particular

question, both were likely true: some people hold real, stable attitudes while other people draw on disparate considerations to construct a new opinion each time they are asked. Using a finite mixture model with three theoretically grounded models of opinion behavior, I quantified the proportion of survey respondents who held stable opinions, those who vacillated or gave ambivalent responses, and those who formed an opinion or changed their opinion during the course of the survey. The vacillating group tended to outnumber opinion holders, but this varied by question. This set of results is consistent with Zaller's model of ambivalence and a "toolkit" theory of culture, but also suggests that some people in the population do form stable, real opinions in certain circumstances.

Ultimately, these results suggest that researchers should devote attention to the social conditions that facilitate stable attitudes and dispositions, whether they are located in people's past or in their contemporary environment. While we know that political awareness tends to facilitate stability in political beliefs, we have no expectation that political awareness should facilitate stability beliefs about religion, morality, or more, but it does suggest that domain-specific awareness might predict domain-specific stability. Attention to the mechanisms that create attitudes can help researchers develop a better understanding of the role of culture in behavior.

## **4. Study 3: Predictably Unpredictable: The Dynamic Constraint of Cultural Belief Systems**

### ***4.1 Introduction***

Theories of culture's role in shaping social behavior posit that culture, internalized in people through durable cognitive structures that facilitate interpretation of the social world, competes with and interacts with other social influences to shape people's attitudes and behaviors over time (DiMaggio 1997; Lizardo and Strand 2010; Martin 2010). Within a society, much of culture is shared, but groups often develop distinct cultural meanings that shape how members understand objects, events, and situations, which lead them to engage in divergent lines of behavior as a result (Frye and Trinitapoli 2015; Harding 2007; Vaisey 2009). Because of culture's theoretical role in shaping dispositions, finding ways to measure these distinct cultural understandings in people, particularly a measure that can be compared against other influences, has been a critical question in sociology (Hunzaker and Valentino 2019; Mohr 1998; Mohr et al. 2020).

When looking for cultural differences in observational data, sociological researchers have tended to take one of two approaches. The first uses single attitude reports or scales designed to probabilistically tap underlying cognitive structures (Vaisey 2009; Miles 2015). The second approach measures the relationships between attitude indicators to tap schematic logics that organize and constrain people's understanding of belief space (Baldassarri and Gelman 2008; Goldberg 2011; Boutyline

2017). The principal assumption of these approaches is that people with similar culturally shared schemas – the underlying networks of associations between abstract concepts that structure people’s cognition (DiMaggio 1997) – will produce similar attitudes or relationships between attitudes.

These cultural measurement approaches struggle to account for the dynamics of attitudes over time. A consistent finding from studies in cultural sociology and public opinion is that people change their attitudes frequently and, for lack of a better term, randomly (Alwin 2007; Converse 1964; Hout and Hastings 2016; Zaller 1992). If culture only manifests only as restrictions to one set of responses, or if culture only manifests as designated relationships between attitudes, then these methods imply that people’s cognition is not durably shaped by cultural systems. This has been the conclusion from studies of ideology in the political realm (Kinder and Kalmoe 2017).

But unlike models of political ideology that make strong assumptions about how attitudes should be related in people’s heads, cultural sociology suggests that heterogeneity and contradiction across social settings are features built into shared cultural systems (Swidler 1986, 2001). To the extent we can measure people’s cultural schemas, people with similar underlying cognitive structures still give different responses to the same questions (Hunzaker and Valentino 2019). This means to the extent that culture can be conceptualized as belief systems that place constraints on people’s attitude change (Martin 2002), these constraints might not be as restrictive as

existing approaches assume. While they likely restrict some beliefs, they should leave room for people to vary on other dimensions. To measure culture's durability in people, account for the role it plays in structuring attitudes, and assess its influence on behavior, sociologists need a measure of cultural belief systems that accounts for this instability.

This paper seeks to reconcile these conceptual and methodological issues by rethinking the empirical signature of cultural belief systems in observational data. I make three principal contributions. First, drawing on insights from sociology of culture and cognition and public opinion research, I argue that the influence of culture on attitudes is not well demonstrated by attitude clustering at a single time, measures of the relationships between attitudes at a single time, or even pairwise change over time. Because culture exists in people as schematized networks of concepts that shape cognition well below the level of attitudes, and because the link between these schemas and attitudes is subject to other cognitive biases, shared cultures do not deterministically produce similar attitude reports. Instead, they *probabilistically* produce survey responses subject to local influences (DiMaggio 1997; Hunzaker and Valentino 2019). If two people have similar (culturally shaped) cognitive structures that connect diverse and contradictory considerations, they might produce different attitudes at a single point in time. But that does not make their cognitive structures any less cultural, shared, or potentially influential on behavior.

Second, I argue that Latent Class Analysis – a method of data reduction that groups people into classes with similar probabilities of giving responses to particular questions – reflects the theoretical tenets of this kind of cultural-cognitive structuring better than measures such as pairwise correlation, relational class analysis, and correlational class analysis (Converse 1964; Goldberg 2011; Boutyline 2017). This generates a clear testable prediction: the constraints evident within groups identified by LCA at a single point in time should reflect the range of responses an individual could give under different circumstances. Therefore, these constraints should predict changes and stability in attitudes over time. I test this proposition using data on religious, moral, and family-structure beliefs from the National Study of Youth and Religion. LCA identifies five belief systems. These systems vary in the degree to which they constrain different beliefs and the portions of belief space to which they constrain respondents. Constraints evident in cross-sectional data at a single time point predict which attitudes people change between waves and how they change them better than models that assume attitudes are largely independent.

Third, I adjudicate the relative importance of these belief systems and structural influences on the pattern of changes. Belief systems deduced at a single time better predict the pattern of attitude changes over time than models rooting beliefs in social circumstances. This suggests cultural belief systems exist independent of these

structures and continue to constrain beliefs over time even as people move across social contexts.

These results have two principal implications. First, they suggest culturally patterned cognition often manifests as shared instability rather than shared attitudes. Culture is messy, but just because it is messy does not mean it is idiosyncratic. Second, they suggest cultural background plays a large role in the structuring of attitudes over time. As people move across contexts, the web of considerations they draw from in constructing attitudes appears more stable than this movement implies. This further directs attention to the circumstances of socialization early in life if researchers want to understand why people believe what they believe and think how they think.

## ***4.2 Theoretical Framework***

### **4.2.1 The 'Problem' of Instability**

Culture and cognition theories argue that people's interpretation of the world is heavily shaped by schemas, "knowledge structures that represent objects or events and provide default assumptions about their characteristics, relationships, and entailments under conditions of incomplete information" through patterned exposure to concepts (DiMaggio 1997: 269). These schemas are information-processing mechanisms, shaping which features of a social situation we attend to, how we internalize and store new information, and how we recall information when prompted (DiMaggio 1997; Hunzaker and Valentino 2019; Strauss and Quinn 1998). Because they form through repeated

exposure, these schematic structures tend to reflect institutionalized and recurring social structures and the cultural environment (Strauss and Quinn 1998). Within larger societies, differences in exposure to these influences produce variation in these cognitive structures that facilitate difference, but potentially predictable, lines of behavior over time (Vaisey 2009).

A standard assumption of cultural measurement is that similar schematic structures will tend to produce similar attitude reports or will tend to produce similar relationships between questions across people (Boutyline 2017; Boutyline and Vaisey 2017; DiMaggio et al. 2018; Goldberg 2011). Under this assumption, researchers work backward from a person's observed pattern of responses to infer that two people who oppose abortion and support government spending on services – and a person who supports abortion access but opposes government spending on services – share a similar “alternative” view of the political world (Baldassarri and Goldberg 2014).

But this inference breaks down when people's attitudes change over time, which public opinion research finds happens frequently (Converse 1964; Zaller 1992). People who oppose abortion and support government spending on services in one survey wave frequently change one of these answers in a subsequent wave. While over-time instability has been most clearly documented in political beliefs, high over-time instability in attitudes has been found across domains (Alwin 2007; Hout and Hastings 2016; Kiley and Vaisey 2020). Because they rely on strong assumptions about the

connections between attitudes, cultural measurement approaches do not provide a way to understand these changes. Did this person change their understanding of the political space? Are the constraints observed at the first time not as durable as assumed? Is one of these responses erroneous?

Given that many people are inconsistent over time, a potential conclusion is that they do not have strongly constrained cognition and that their cultural beliefs are not influential for action (Converse 1964; Jerolmack and Khan 2014; Swidler 1986). But nothing about culturally patterned schematic structures requires that they produce similar responses to a survey or interview question. Schematic cognition exists well below the level of attitudes in the networks of abstract concepts in people's heads (D'Andrade 1995; Hunzaker and Valentino 2019; Strauss and Quinn 1998). These networks frequently connect conflicting concepts in the same domain.

Decades of work in cultural sociology and public opinion document how people consume diverse and contradictory bits of culture, often storing this heterogeneous mixture without taking time to reconcile its contradictions (DiMaggio 1997; Martin 2010; Swidler 1986; Zaller 1992). As a result, "our heads are full of images, opinions, and information, untagged as to truth value, to which we are inclined to attribute accuracy and plausibility" (DiMaggio 1997: 267). People might connect marriage to pragmatic considerations of social support and shared management of household responsibilities but also to strong feelings of infatuation and sexual desire (Swidler 2001).

Just because cognition is conflicting does not mean it is idiosyncratic. Swidler's (2001) romantic/prosaic model of love is instructive because it is explicitly not an idiosyncratic belief structure. Its recurrence across people is a key reason it is interesting. Across a range of interviews with middle class professionals, Swidler found that respondents had no trouble believing that "love is (1) a clear, all-or-nothing choice; (2) of a unique other; (3) made in defiance of social forces; and (4) permanently resolving the individual's destiny" while simultaneously believing that "(1) Real love is not sudden or certain ... (2) There is no 'one true love' ... (3) The kind of love that leads to marriage should not depend on irrational feeling in defiance of social convention ... [and] (4) Love does not necessarily last forever," despite the inherent contradictions in these sentiments (Swidler 2001: 113-114). This connection of the abstract notion of love to competing notions of romance and prosaic considerations is a durable cognitive structure, repeated across people, produced through repeated exposure to cultural objects and social behavior, that connects the concept of love to heterogeneous and conflicting considerations. It is, to the best of our understanding, schematized culture (this conflicting romantic/prosaic model of love is also one of the central "schemas" that Strauss and Quinn [1997] use to illustrate the concept).

When asked, "Do you think that, in general, a couple without children should end their marriage if it is empty and unfulfilling, or should they stick with it even if they are not happy?" people with this model of love might struggle to be consistent over time

as competing notions of love pull them in different directions. These over-time changes are a feature of the interaction between the cultural-cognitive structure and the social situation. In a different situation – the question, “is love a complicated mixture of prosaic and romantic sentiments?” – this cultural belief system might produce stability where another belief system of love, such as one rooted in Biblical notions of love, might struggle. It is easy to identify other places where accepted models of culturally schematized cognition produce inconsistency in the face of misaligned social structures. A question asking whether Jesus Christ was a man or God might prove problematic for the most structured Christian belief system (Martin 2002), but not very challenging for an atheist who is relatively unconstrained on other beliefs. In other words, instability is not necessarily a reflection of the durability or sharedness of the belief system, only the interaction between the schema and the social environment.

#### **4.2.2 Culture, Cognition, and Survey Response**

Schemas are principally *reactive* structures, shaping people’s interpretation of situations, internalization of new considerations, and recall of information in response to prompts (Hunzaker and Valentino 2019; Strauss and Quinn 1998). They might prevent people from internalizing conflicting information, which would help them facilitate stable attitude responses over time (Hunzaker 2016). When people internalize conflicting information, schematic structures might facilitate the recall of some bits over others. These two factors help explain why on any particular question, some proportion of the

population clearly articulates the same opinions over time, with people differing on which issues they are stable (Converse 1964; Hill and Kriesi 2001).

But people's cognition is subject to a range of other influences. Public opinion research suggests that when prompted to give an opinion on an issue, such as in response to a survey or interview question, people bring up a range of considerations stored in their heads, shaped by cognitive biases and local influences such as question structure, recent interactions with peers, and institutional constraints. They then use these considerations to generate an opinion (Zaller 1992; Perrin and McFarland 2011). People with conflicting considerations stored in their cognition do not simply average all considerations and pick scale midpoints (though they do this occasionally). They can range widely in their stated beliefs over time as local influences shift, shaping which considerations are most prominent in their minds at any time.

Importantly, considerations that have been recently called up through discussions with friends, stories on the local news, or other social experiences are more likely to get called up in the survey context even if people, in the abstract, have the same underlying cultural-cognitive structure (Kahneman 2011). This does not mean that people's cognition will be unimportant for attitude behavior over time, or that any schematic structure could produce any pattern of attitudes. The underlying belief system might restrict people to relatively narrow bounds on some questions. But on

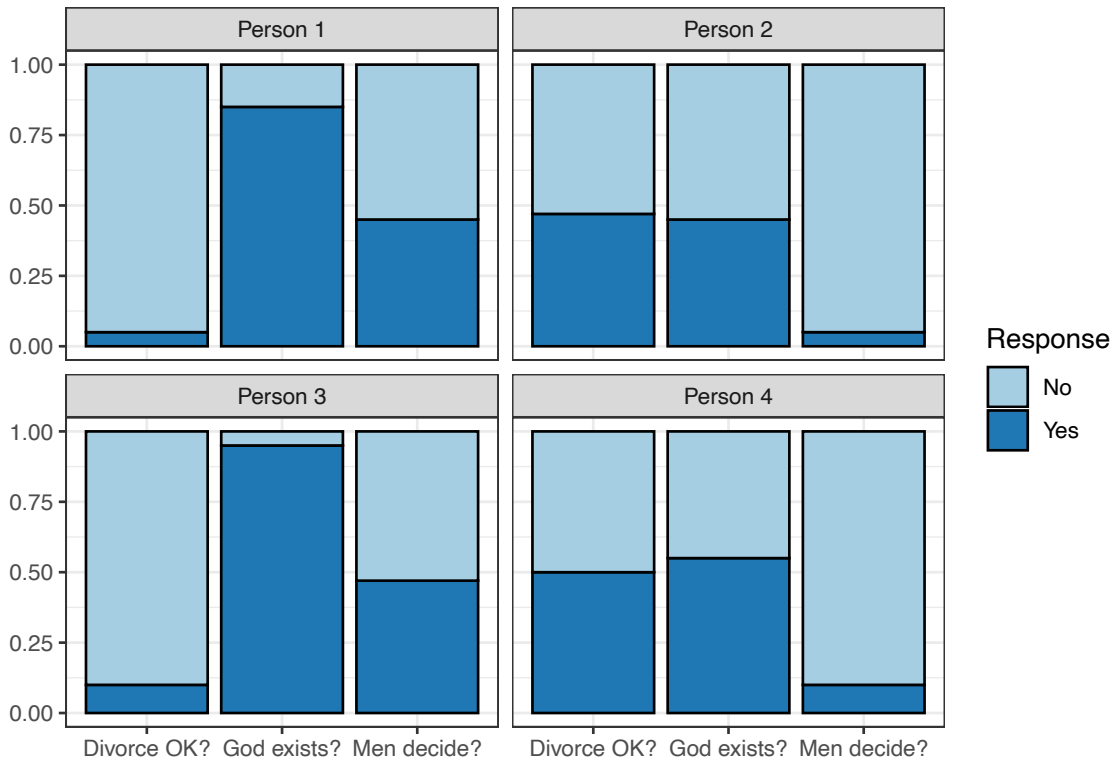
relatively unconstrained dimensions, these cognitive biases might produce a range of responses.

The preceding discussion suggests that shared cultures, to the extent that they are internalized as similar cognitive structures across people, do not have to manifest as shared attitudes or networks of attitudes at a single point in time. They exist below attitudes and, in conflict with other biases, shape the internalization and recall of considerations that drive people to give different responses when prompted. In other words, cultural schemas are networks of cognitive connections that *probabilistically* produce survey responses under a variety of conditions, and thereby limit (or not limit) people's responses to certain portions of the belief space over time.

#### **4.2.3 Consequences for Measurement**

Because local circumstances might shape the interaction between schematic cognition and attitudes, to truly map someone's cultural belief system, and to understand whether two people's attitudes were governed by similar cultural belief systems, we would ideally observe them across a variety of social circumstances as they were exposed to a range of stimuli and either allow or do not allow these influences to shape their responses over time. It is only through such a process that we would observe the effect of culture on belief systems, or how the "arbitrary movement of individuals in [belief] space has been reined in" (Martin 2002: 865).

Imagine four people are asked the same three binary questions across an infinite variety of social situations carefully controlled by a researcher. One question asks whether God exists, one question asks whether divorce is an acceptable option for unhappy couples, and one that asks whether men should make most decisions in relationships. We poke, prod, and prime these respondents, asking them questions in their house or their church, with their significant other present or absent, surrounded by friends or parents, while they are in their car on the way to work or while they are relaxing on their couch, wiping their mind of any memory of the question before the next asking. We change the order of the questions, change the gender and race of the interviewer, change the time of year, and on and on and on. Figure 11 presents the results of this hypothetical process, with each panel representing one of these people.



**Figure 11: Over-time response probabilities for hypothetical belief systems.**

When we look at the distribution of responses, we would see that the two people on the left have a similar over-time distribution of response probabilities across the three questions, possibly shaped by an underlying Christian worldview that creates schematic connections between religious concepts and relationships. Because the system links marriage to the existence of a deity and eternal salvation, people with this belief system have a relatively easy time rejecting competing considerations and influences as they move across social environments and a relatively constant recall process in response to the questions about divorce and God. As a result, they generate the same answer in most circumstances. A crisis of faith or actual error in filling out the survey at one time point

might lead these people to say that God does not exist or that divorce is acceptable, but these instances are rare. At the same time, conflicting influences from their church and from the broader society cause them to be inconsistent on the question of gender roles as social circumstances push different considerations to the forefront of their cognition.

The two people on the right have internalized the conflicting messages present in the population about love and religion and vacillate on those two questions over time as local influences bring different considerations to the forefront of their cognition. When something triggers the prosaic model of love, such as a discussion with a friend about marital difficulties, these people say people should get divorced if they are unhappy. When something triggers the romantic model of love, such as recently watching a romantic movie, they oppose divorce. But because they lack the connection of marriage to traditional conceptions of gender roles, they have a relatively straightforward time giving the same answer to that question over time.

The key point here is that the lack of constraint on questions in each system is itself a reflection of culturally patterned cognition. If we could get into the heads of the people on the right, we would see similar schematic connections from love to prosaic and romantic notions, but maybe not the religious notions of love that we would observe in the heads of the people on the left. We could trace this structure to exposure to cultural artifacts, such as movies, books, and music, and patterned observations of how people behave.

In this framework, we observe a shared belief system if we see groups of people who share probabilities of answering questions in particular ways over time and across a variety of (reasonable) social circumstances. The obvious challenge of this definition is that we do not frequently observe people's responses to the same question repeatedly over time. And if we use beliefs over time to deduce belief systems, we have nothing on which to test whether these assumptions are true. Fortunately, a method for detecting such cultural systems in cross-sectional data exists and has been used in sociological studies of attitude structuring before (Bonikowski and DiMaggio 2016; DiMaggio et al. 2018).

#### **4.2.4 Latent Class Analysis and Belief Systems**

Latent Class Analysis is a data-reduction method that seeks to group people into unobserved categories where, within these categories, the probability of giving a particular response to a question is independent from the probability of giving responses to other questions (Goodman 1974; Lazarsfeld and Henry 1968; McCutcheon 1987). This fundamental assumption, the conditional independence assumption, posits that once the latent class is identified, each person's response on a particular question is an independent draw from the probabilities of the different responses observed within that group.

This statistical model aligns closely with the theoretical model outlined above: people's patterned exposure to concepts gives them similar probabilities of answering a

question with a particular response. LCA allows for the presence of competing sets of these belief systems in the population, reflected in the different latent classes. In contrast to other methods, LCA allows people to give a range of responses over time to remain in the same “system.” And it allows beliefs to be constrained to different degrees in different systems. In cultural sociology, LCA has been used to deduce belief systems about religion and science from survey responses in the General Social Survey (DiMaggio et al. 2018) and to deduce varieties of popular nationalism, also using the GSS (Bonikowski and DiMaggio 2016). However, these works do not make the strong assumptions about underlying cognition made here or, importantly, test these assumptions over time.

If, as the preceding discussion argues, a cultural belief system is a shared set of considerations and the connections between them that probabilistically produce certain responses over time across a range of questions, LCA should be able to detect these systems in a single wave of data. Assuming that belief systems do not undergo major revisions over time, the features of these systems observed at one time point should predict the behavior of attitudes over time.

#### **4.2.5 Culture and Social Structures**

To this point, I have assumed that cultural cognitive structures operate independently from other social considerations in forming beliefs over time. This is obviously not true. Attitudes reflect a dynamic interplay between cultural-cognitive

structures, social structures, and the physical environment (DiMaggio 1997; Goldberg and Stein 2018; Lizardo and Strand 2010; Martin 2010). And adjudicating the relative influence of these processes is difficult because people's cultural beliefs and preferences appear to shape their social networks and their organizational participation (Lewis and Kaufman 2018; Lizardo 2006; Vaisey and Lizardo 2010).

If differences in cognitive structures exist, they have to come from somewhere, but sociological theories disagree on how much social structure matters at different points in the life course. Some approaches suggest that people form durable dispositions when they are young, that they carry these dispositions with them across social contexts, and that these dispositions are relatively impervious to change over time (Bourdieu 1990; Ryder 1965). This might be because of biological processes that make early experiences more formative for cognition, it might be because early experiences shape people's movement into subsequent social environments, or it might be because schematized cognition helps filter out competing considerations over time. Either way, this line of reasoning suggests that constrained beliefs will likely be unaffected by changing influences as people move across social structures.

Other approaches suggest cognitive structures are less durable and that social structures reshape cultural belief systems over the life course. These models posit that people generally internalize a range of conflicting considerations. As people move into new social environments, they get exposed to new considerations that rework their

underlying cognition. Cognitive authorities, leaders endowed with the social responsibility for shaping beliefs, provide clear guidelines for which attitudes go together, and organizational hierarchies make certain belief structures appear impossible (Martin 2002). Affect-laden social interactions can make holding some attitudes feel uncomfortable, leading people to change their attitudes and producing constrained attitudes within communities (DellaPosta, Shi, and Macy 2015; Rawlings 2020). New considerations diffuse across social networks (Centola 2018). These approaches would ascribe little durability to the cognitive structures that people form and lodge a lot more explanatory power in the social environments in which people find themselves (Martin 2010).

Rethinking the hypothetical example above, it is plausible that all four people share similar underlying cognitive structures and what varies between the two pairs is their social circumstances, which shape which considerations get called up. However, if durable cognition played no role in attitude stability over time, we would expect to observe durable changes in attitudes as people move across social settings. But we tend to find the opposite, that people's dispositions are often more durable than these changing circumstances would imply (Alwin and Krosnick 1991; Kiley and Vaisey 2020; Sears and Funk 1999). This suggests that some form of durable cognitive structuring does play a role in attitude stability.

### **4.3 Hypotheses**

The preceding discussion suggests a set of hypotheses about the behavior of attitudes over time. The first three hypotheses are methodological, revolving around the ability of LCA at one time point to identify belief systems that predict attitude behavior over time. The second set of hypotheses pertain to the relative role of cognitive structures in shaping attitude behavior compared to structural influences such as social networks and organizational participation.

#### **4.3.1 Latent Class Analysis as a Measure of Belief Systems**

The central assertion of the methodological argument is that Latent Class Analysis, using data across people at a single time point, should be able to deduce belief systems that predict within-person attitude behavior over time. As constraints – restrictions of movement of beliefs over time – are taken to be the central signature of a belief system (Martin 2002), I focus on those first. Specifically, the constraints on attitudes identified within a latent class at one time should predict the degree to which people change their attitudes over time. This produces the first two hypotheses:

*Hypothesis 1: Within cultural belief systems, beliefs that are more constrained will demonstrate less change over time than less constrained beliefs.*

*Hypothesis 2: Across cultural belief systems, the same belief will show less movement over time if it is in a belief system that constrains it more.*

The theoretical model outlined above makes a stronger prediction. It says that the belief system at time 1 should not only predict the degree to which attitudes in any particular system will change, but the probabilities that people in these belief systems will give certain responses, assuming changing social circumstances do not severely disrupt belief systems (a point I discuss below). If the theoretical model is correct, responses at any time should be conceptualized as independent draws from the distribution deduced at time 1. These probabilities are shaped by broad cultural forces that influence people's considerations, but which specific response a person gives at any wave will be shaped by (random) local influences. While it will be hard to predict what any particular person will say in each wave, especially if we are looking at an unconstrained belief, assuming these draws are independent can give us strong predictions for the overall count of observed patterns over time. This produces the third hypothesis:

*Hypothesis 3: Over time, the aggregate response patterns of the sample should reflect independent draws from the deduced belief systems.*

#### **4.3.2 Adjudicating Culture and Structure**

The preceding hypotheses assume either that social circumstances are invariant or that changing social circumstances do not significantly affect people's belief systems. I test this assumption and, in doing so, compare the relative influence of social structures and cultural belief structures on attitudes over time.

If people have durable schematic thinking, then they should be less susceptible to the influence of alternative considerations that come from changing social environments (Hunzaker 2016). This might in part be because they select into social settings, but it should principally be because these schematic structures influence how they process new information.

The most prominent alternative model is that socially patterned collections of attitudes are largely a product of people being located in similar social structures, and that cultural background, shaping ongoing cognition through durable cognitive structures, is largely irrelevant. Under this hypothesis, the reason people exhibit similar sets of attitudes is not because they have similar cultural belief systems shaping their attitude responses over time, but because they exist in similar social environments that produce similar independent attitudes. As people move across social environments, they hear different sets of considerations that reshape their consideration sets and the attitudes they report in surveys. If this is the case, measures of social structure at the time of the survey should outperform previously deduced cultural belief systems in explaining current beliefs.

As noted previously, attitude structuring is undoubtedly an interaction between cultural schemas and social structures (DiMaggio 1997; Martin 2010; Lizardo and Strand 2010). But given a lack of evidence for durable change in attitudes over time and the fact that we observe substantially less durable change in attitudes than changing social

circumstances would predict (Kiley and Vaisey 2020), I expect that belief systems play a much larger role in shaping people's attitudes over time than this structural reasoning suggests. This leads to my final hypothesis:

*Hypothesis 4: Belief systems will better predict people's attitude reports over time than models that account for changing social circumstances.*

#### **4.4 Data and Measures**

Testing the propositions outlined above requires data on the same beliefs and relevant social structures measured over time. The National Study of Youth and Religion meets these criteria. Below I outline the data set and measures used to test the hypotheses outlined above.

##### **4.4.1 The National Study of Youth and Religion**

The National Study of Youth and Religion is a four-wave panel survey of adolescents that began in 2002 when respondents were between the ages of 13 and 17 and surveyed them every three or four years for four waves. The survey began with a sample of 3,370 adolescents and is designed to be a probabilistic sample of adolescents in the United States at the time the survey started. The same respondents were asked to complete subsequent waves with response rates varying over time: 2,604 at wave 2 (72 percent); 2,182 at wave 3 (65 percent); and 2,144 at wave 4 (64 percent). In wave 2, respondents were between ages 16 and 20, in wave 3 respondents were between ages 17 and 24, and in wave 4 respondents were between ages 20 and 32.

The age range of the NSYR is important to the theoretical argument outlined above as it pertains to the structuring of cognition and the movement across organizational and social contexts. Adolescence and early adulthood, the periods covered in the NSYR, are assumed to be particularly formative for cultural beliefs (Kiley and Vaisey 2020; Ghitza, Gelman, and Auerbach, n.d.; Bartels and Jackman 2014). At the same time, these periods also represent a significant time of transition for young people in the United States as they move out of their parents' homes and into college and the workforce, begin to form long-term romantic attachments, and generally transition from adolescence to adulthood. There is likely more movement across social contexts at this period than most other periods in the life course. As such, this provides a good window in which to test the competing influences of cultural belief structures, organizational settings, and social change.

My analysis focuses principally on Waves 2 through 4 of the NSYR.<sup>15</sup> Because I only draw a few background variables from the first wave of the NSYR, and because time matters significantly in the testing of the theoretical model outlined above, for clarity I will refer to waves 2, 3, and 4 of the NSYR as times 1, 2, and 3 for the rest of this paper.

---

<sup>15</sup> Many of the attitude measures used in this analysis were not added to the survey until wave 2. Focusing on attitudes present in all four waves severely curtails the number of beliefs that can be explored.

Deducing belief systems at time 1 requires people to have responses on all belief measures and covariates. Because the relationships among beliefs are a central question of this analysis, any missing data imputation strategy might bias the deduction of classes in unforeseen ways. To measure belief systems, I only use people who are observed on all covariates and beliefs at time 1, leaving a sample of 2,521 respondents. Testing hypotheses 1 and 2 requires people to have beliefs observed at all three time points, but it does not necessarily require covariates in later waves, and the test does not require people to be observed on all beliefs. Because of this, I use data on people whose classes are deduced at time 1 and who were observed on any beliefs at times 2 and 3. The sample size for each question varies between 1,668 (divorceok) and 1,699 (demons).

Testing hypothesis 3 and 4 requires comparable social structure and belief variables over time. Because some key social structural questions were dropped at time 3 (specifically questions about social networks), I focus on change between times 1 and 2. Again, it is not necessary that respondents report all beliefs at time 2, but respondents must have full beliefs at time 1 and full covariates at times 1 and 2. The analytical samples range from 2,153 (divorceok) to 2,175 (demons).

#### **4.4.2 Measures**

**Beliefs:** In times 1 through 3, NSYR respondents were asked a set of questions about their religious, moral, and family-structure beliefs. Seven questions ask about religious beliefs, four ask about morality and the role of religion in daily life, and six ask

about gender relations and family structures.<sup>16</sup> These questions are reported on either a three-point scales of “yes,” “maybe,” and “no,” or five-point scales of “strongly agree,” “agree,” “undecided/don’t know,”<sup>17</sup> “disagree,” “strongly disagree.” These variables are summarized in Table 3.

**Table 3: Belief measures, by wave.**

Religion	Do you believe...	Time	1. Yes	3. Maybe	5. No	N. Obs
afterlife	...that there is life after death?	1	1188	1078	255	2521
		2	1075	886	207	2168
		3	862	640	197	1699
angels	...in the existence of angels?	1	1403	866	252	2521
		2	1187	734	254	2175
		3	910	488	298	1696
astrology	...in astrology, that stars and planets affect people’s lives?	1	301	579	1641	2175
		2	225	451	1498	2174
		3	140	521	1036	1697
demons	...in the existence of demons and evil spirits?	1	1118	834	569	2521
		2	980	639	556	2175
		3	754	575	370	1699
god	...in God, or not, or are you unsure?	1	1950	448	123	2521
		2	1662	378	134	2174
		3	1206	285	204	1695
godworld	...god created the world?	1	1731	125	665	2521
		2	1414	119	630	2163
		3	706	582	408	1696
heaven	...in heaven as a place where some people go after death?	1	2121	35	365	2521
		2	1778	18	376	2172
		3	1106	311	281	1698
miracles	...in the possibility of divine miracles?	1	1531	696	294	2521
		2	1317	557	300	2174
		3	922	505	272	1699
reincar	...in reincarnation, that people have lived previous lives?	1	436	766	1319	2521
		2	312	655	1206	2173
		3	213	701	779	1693

<sup>16</sup> An obvious omission from this list of questions is the one Vaisey (2009) uses to predict adolescent behavior and social networks over time, which he argues represents people’s “moral typology.” Because of a coding error, responses to that question were lost for almost all respondents at Wave 3 of the NSYR, making it hard to compare across waves.

<sup>17</sup> Responses of “don’t know” were not provided in the survey for questions asked on the five point scales, but if respondents gave them, they were coded as such. As such, they are relatively rare. When responses to questions that include a “maybe” option are coded as “don’t know,” I recode these responses to the “maybe” response.

		1.	2.	3.	4.	5.	N.	
		Strongly	Agree	Unsure/	Disagree	Strongly	Obs	
Morality	Time	Agree	Agree	DK	Disagree	Disagree		
brkmoral	It is sometimes okay to break moral rules if it works to your advantage	1	28	352	7	1531	603	2521
		2	35	321	5	1129	678	2168
		3	19	111	97	873	593	1693
moralchg	The world is always changing, and we should adjust our views of what is morally right.	1	272	1272	19	702	256	2521
		2	215	1130	17	610	195	2167
		3	149	606	161	491	282	1689
moralrel	Morals are relative, there are no definite rights and wrongs for everybody.	1	221	1224	33	744	299	2521
		2	146	876	10	817	316	2165
		3	125	489	117	637	322	1690
relprvte	Religion is a private matter that should be kept out of public debates	1	368	1076	13	829	235	2521
		2	400	919	10	661	177	2167
		3	433	574	139	381	162	1689

		1.	2.	3.	4.	5.	N.	
		Strongly	Agree	Unsure/	Disagree	Strongly	Obs	
Family	Time	Agree	Agree	DK	Disagree	Disagree		
divrceok	A couple without children should end their marriage if it is empty and unfulfilling.	1	1755		37		729	2521
		2	1483		26		644	2153
		3	1183		0		499	1682
mandecid	Important decisions in the life of the family should be made by the man of the house.	1	68	380	6	1429	638	2521
		2	54	322	16	1101	674	2167
		3	47	181	57	759	644	1688
manmar	A man can have a fully satisfying life without getting married.	1	403	1646	8	404	60	2521
		2	412	1392	8	287	71	2170
		3	526	863	70	187	47	1693
unmarsex	It is alright for two married people who are not in love to have sex.	1	157	1177	9	862	316	2521
		2	175	1104	13	663	211	2166
		3	398	723	80	309	180	1690
wommar	A woman can have a fully satisfying life without getting married.	1	459	1632	16	371	43	2521
		2	485	1362	10	269	44	2170
		3	531	850	89	184	38	1692
wrkngmom	Working mom can establish as warm relationship with child as one who does not work.	1	489	1581	12	401	38	2521
		2	440	1303	9	374	40	2166
		3	594	836	70	166	23	1689

These responses are treated as nominal in some models and continuous in others.

To make the range of responses to each question comparable, I scale all attitude measures to five-point scales between 1 and 5 by converting questions on three-point scales: “yes” to 1, “maybe” to 3, and “no” to 5.

**Covariates:** Given the belief domains explored here, I examine three principal sources of attitude structuring: sociodemographic backgrounds, participation in organizations, and social networks. Sociodemographic background variables include

respondent gender (indicator for female), and whether at least one parent has a bachelor's degree.<sup>18</sup>

A second set of covariates is designed to tap organizational participation, which is expected to change to some extent between waves. Given the role of religious organizations in shaping attitudes explored here, I include a set of indicator variables for religious tradition (Steensland et al. 2000) and a measure of church attendance, measured on a scale from never (0) to more than once a week (6).<sup>19</sup> Because participation in formal education might introduce an alternative set of considerations into people's belief systems, I include a variable measuring the number of years of education a person has received above ninth grade. I also include an indicator for whether the respondent lives in the South census region to capture movement across geographic contexts. Other census regions were uncorrelated with beliefs.

---

<sup>18</sup> These demographic variables are measured at the survey's first wave and are assumed to be invariant during the survey. Preliminary analyses found that other demographic variables, including race/ethnicity, parent's income, and whether people grew up in a two-parent household were unrelated to beliefs, net of other predictors.

<sup>19</sup> The NSYR includes a separate indicator for members of the Church of Jesus Christ of Latter-day Saints. The religious tradition measure places members of the Church of Jesus Christ of Latter-day Saints in the "other" category. I group these respondents with evangelical protestants, based on preliminary analysis of their response patterns. While there are deep theological differences between the groups, these are not reflected in the questions analyzed here, and the two groups behave similarly in terms of the beliefs measured in this analysis. The NSYR also includes a code for respondents with an indeterminate religion, separated by whether they attend services or not. I group these into a distinct category of indeterminate religious tradition. Finally, the NSYR includes separate codes for "black Evangelical protestants" and "black Mainline protestants." I group these two categories together.

Finally, to measure social network influence, I include the proportion of a respondents' friends who share that person's religious orientation, including no religious orientation for people who do not express one.<sup>20</sup> As a second measure of social networks, I include the highest level of closeness a respondent reported with either parent. These covariates are measured at times 1 and 2.

#### 4.4.3 Belief Systems

I use Latent Class Analysis to deduce a set of belief systems using the 19 attitude items asked at time 1. LCA attempts to assign a class to each respondent such that their responses are independent from each other within classes. The model uses the Newton-Raphson method to maximize the log-likelihood of multiple parameters under the assumption that indicators are independent conditional on class membership.

A model that predicts the observed response pattern  $y$  using a latent class variable  $X$  with  $L$  values has the following probability structure:

$$P(Y = y) = \sum_l P(X = l)P(Y = y|X = l)$$

Because of the conditional independence assumption that  $K$  indicators are independent within each latent class  $l$ , the joint probability of a given response pattern

---

<sup>20</sup> Almost all respondents at time 1 (about 94 percent) said they had five close friends, making a count of friends as an additional measure superfluous.

can be rewritten as the product of individual item response probabilities, and  $P(Y = y|X = l)$  can be rewritten as:

$$\prod_k P(Y_k = y_k|X = l)$$

LCA estimates a set of parameters that includes the relative class proportions and the probabilities of each response for each class on each question. For the purposes of the LCA, responses to the belief questions are treated as nominal.

The model also includes the covariates outlined above as predictors of class assignment. The model simultaneously estimates two conditional probabilities: the probability of response conditional on group assignment and the probability of group assignment conditional on covariates. I use the Bayesian information criterion to select the best-fitting number of classes.

#### **4.4.4 Change Over Time**

The first two hypotheses make predictions about how much attitudes should change over time as a function of the constraint within a belief system at a single point in time. Within a system, more constrained beliefs should change less over time than less constrained beliefs, and across systems, a belief should change less in a more constrained system.

I measure constraint of a particular attitude,  $j$ , for members of a designated belief system,  $k$ , by calculating the within-class standard deviation of responses to that question over all people,  $i$ , in at time 1. LCA assigns each person a probability of

belonging to each class. For this analysis, I assign each person to the class with the highest probability. I then calculate the standard deviation of responses within that group, treating responses as continuous, rather than nominal as the LCA does. A group where people tend to give the same response or cluster in adjacent responses will have a low standard deviation and therefore high constraint. A group where people tend to give answers across the scale will have a high standard deviation and therefore low constraint.

$$\sigma_{jk} = \sqrt{\frac{\sum(x_{ijk} - \mu_{jk})^2}{N_k - 1}}$$

The outcome of interest is the amount of variation people express in their attitudes over time. I measure within-person variance over time using the within-person standard deviation of responses given at times 1, 2, and 3.

$$\sigma_{ij} = \sqrt{\frac{\sum(x_{ijt} - \mu_{ij})^2}{N_i - 1}}$$

I test the first and second hypothesis using a single linear regression of within-person variance on the within-class standard deviation at time 1, with fixed effects for question ( $J$ ) and for person ( $I$ ). This amounts to simultaneously testing whether people exhibit more variation in their less constrained beliefs over time than their more-constrained beliefs and whether a belief demonstrates more over-time variation when it is in a less constrained belief system than when it is in a more constrained belief system.

$$\sigma_{ij} = \sigma_{jk} + I + J + \epsilon_{ij}$$

I also conduct a similar fixed-effects regression of each person's mean for each question at times 2 and 3 ( $\mu_{ij,t>1}$ ) on the group mean at time 1 ( $\mu_{jk}$ ) for each question at time 1. This tests whether constraints about where group members fall on average continue to constrain people over time.

$$\mu_{ij,t>1} = \mu_{jk} + I + J + \epsilon_{ij}$$

#### 4.4.5 Pattern Prediction

The remaining hypotheses posit that modeling people's responses over time as independent draws from the belief system outperforms other potential data-generating processes. To assess this proposition, I take a predictive approach (Hofman, Sharma, and Watts 2017; Salganik et al. 2020).

Hypotheses 3 and 4 focus on the observed counts of change patterns over time. To illustrate this approach, assume two belief systems that differently constrain people's views on the following question: "Do you think that, in general, a couple without children should end their marriage if it is empty and unfulfilling, or should they stick with it even if they are not happy?" In one belief system, people are constrained to oppose divorce quite strongly ( $Pr(yes) = .9$ ). These people have many considerations against divorce, but there is a small chance that a local event could make them support it. In the second belief system, people have roughly equal considerations in favor of and opposed to divorce ( $Pr(yes) = .5$ ). Which response they give at a particular wave will be affected by the balance of considerations on their mind at any time.

If responses over time are independent draws from this distribution, then if we asked people the question two times, people in the first group should say “yes” in both waves about 81 percent of the time ( $.9 * .9 = .81$ ). People in the second group should say “yes” in both waves about 25 percent of the time ( $.5 * .5 = .25$ ). We can calculate the probability of each of the four possible two-wave response patterns, presented in Table 4 below.

**Table 4: Expected response pattern probability for two hypothetical belief systems.**

Pattern	Pr(yes) = 0.90	Pr(yes) = 0.50
Yes -> Yes	0.81	0.25
Yes -> No	0.09	0.25
No -> Yes	0.09	0.25
No -> No	0.01	0.25

A key assumption of the belief systems model is that predicting any person’s response at any particular time will be difficult, especially if it is deduced that beliefs are relatively unconstrained in a particular system, such as in the rightmost column. But since it assumes local influences are statistically random, the theoretical model can generate strong predictions of counts of response patterns in the aggregate. I can use the distribution of these two belief systems in the population, as well as the distribution of responses observed at time 1, to generate a range of plausible predictions for the count of each pattern.

Predicting response patterns in the latent class model requires two steps: sampling class identification and sampling responses. The LCA model assigns each observation a probability of belonging to each class based on their covariate profile. I sample class assignment from these probabilities. Then, using these class assignments, I sample responses from the probabilities assigned to members of that class. I count the number of people who demonstrate each response patterns (“Agree” at time 1 to “Disagree” at time 2) and compare that to the observed count of response patterns. While the theoretical framework makes within-class predictions, because people are probabilistically assigned to different classes, and to make comparisons to other theoretical processes, I aggregate counts of response patterns at the question level.

To measure the predictive accuracy of a model for each question,  $j$ , I square the difference between the expected number of response patterns generated by the model for all combinations of responses at time 1 and time 2,  $E_j(t_1, t_2)$ , and the observed number of cases that had that response pattern,  $O_j(t_1, t_2)$ , penalizing larger deviations. I then sum across all potential patterns and take the square root of this value to get a measure of predictive accuracy.

$$\lambda_j = \sqrt{\sum_{t_1} \sum_{t_2} (E_j(t_1, t_2) - O_j(t_1, t_2))^2}$$

Because both class assignment and response probabilities reflect sources of uncertainty, I iterate this process 10,000 times to generate a distribution of accuracy that reflects the probabilities of class assignment and response probabilities.

This range of numbers provides a quantification of how good the model predicts response patterns, with 0 being a perfect prediction of the count of response patterns. But this range is meaningless on its own, since there is no clear alternative expectation for how many counts we might observe. It is unlikely that any model would perfectly predict responses over time. But I can compare this theoretical process to other theoretical models. To generate a plausible upper-bound of prediction, I sample from the marginal distribution at wave 3 for each respondent.

The first theoretical alternative to the belief system model posits that beliefs are independent from each other. In this framework, belief in the existence of God is shaped by exposure to one set of considerations, and belief about divorce is shaped by another. While patterned exposure to social settings might lead people who believe in God to also oppose divorce, there is no cognitive link between these issues that constrains people's responses. Under this model, people have more or less idiosyncratic belief systems (or sets of considerations) as a function of their social experiences. People would receive separate influences on each belief from their social environments – churches, schools, families, friends, etc. – and these would shape their responses at each wave.

To estimate these idiosyncratic responses, I conduct a multinomial logistic regression for each individual attitude at time 1 using the same set of covariates

included in the latent class analysis.<sup>21</sup> This produces a set of individual-specific probabilities of giving each response to a question. I then use those probabilities to simulate responses over time and similarly quantify predictive accuracy.

Model comparisons typically penalize models for complexity, as complexity leads to greater predictive accuracy within a sample. The latent class model, while complex, is simpler than estimating separate models for each belief. If the latent class model makes better predictions, there is no reason to prefer estimating separate probabilities for each response on the grounds of parsimony. There are obvious ways to simplify both models by removing parameters that do not aid in prediction, or by treating responses as ordinal rather than multinomial. However, the goal of using the same predictors and same outcome scale is to design two models that reflect two similar but distinct theoretical processes: one where beliefs influence and constrain each other, and one where they do not.

#### **4.4.6 Changing Social Structures**

Hypothesis testing to this point has been oriented toward establishing LCA as a good methodological fit for the theoretical concept of a belief system and the predictions

---

<sup>21</sup> This approach assumes that each person's response at each wave is a draw from a multinomial distribution. An alternative is to assume that each person's response is a latent variable observed with error. This would model the outcome not as a set of independent categories (multinomial logit/probit), but as manifestation of a latent variable (ordinal logit/probit). In practice, the multinomial logit is a less constrained version of the ordinal logit model. If attitudes do reflect an underlying latent construct, the multinomial logit will reflect this structure, but the reverse is not true. Since I am not principally concerned with model parsimony, but rather on adjudicating theoretical processes, I use the multinomial logit model.

it makes over time. If that is established, then I can use the deduced belief systems to compare the relative influence of the belief system with social structural changes that might produce changes in beliefs over time.

Both the LCA model and the idiosyncratic beliefs model outlined above assume that belief systems are shaped early in life and endure, but Hypothesis 4 suggests that changing social circumstances will reshape beliefs over time. To test the influence of organizational and social network change, I use the coefficients derived from the latent class and multinomial logit model at time 1 to predict class assignment and responses at time 2 using social structural and social network variables observed at time 2. If changing circumstances – increased church attendance or a more diverse friend group, for example – have the effect of producing changes in attitudes, then using information about social change between waves will produce better estimates of the patterns of change over time.

This approach makes a rather strong assumption that the coefficients estimated at time 1 hold at time 2, but that assumption is inherent in the strong structural argument. Under these approaches, beliefs at time  $t$  are a reflection of structures at time  $t$ . If coefficients derived at time 1 are a bad predictor of beliefs at time 2, it suggests that different influences matter at each time point or that past structures carry influence forward, which is the theoretical argument of the belief systems model.

To ensure comparability across prediction models, each prediction model uses all people with full beliefs and covariates at time 1 to generate coefficients (2,521 respondents), and people with beliefs at time 1 and all observed covariates at times 1 and 2. A handful of people with covariates at time 2 failed to answer some of the belief questions. They are evaluated on the questions they did respond to.

## **4.5 Results**

The results proceed in three parts. First, I deduce and explain the belief systems identified through LCA. Second, I test the proposition that the constraints implicit in each belief system are good predictors of over-time change. Third, I adjudicate the competing influences of the belief system and social structures in predicting responses over time.

### **4.5.1 Belief Systems**

Based on goodness of fit measures and substantive interpretation, I selected and present a five-class model to summarize the belief systems across the three domains outlined above. Figure 12 presents the expected probability of each response option for all 19 questions for all of the classes. Table 5 summarizes the distribution of covariates across classes. I briefly summarize each belief system, giving a substantive interpretation based on response probabilities and covariates, as well as their implications for expected over-time change.

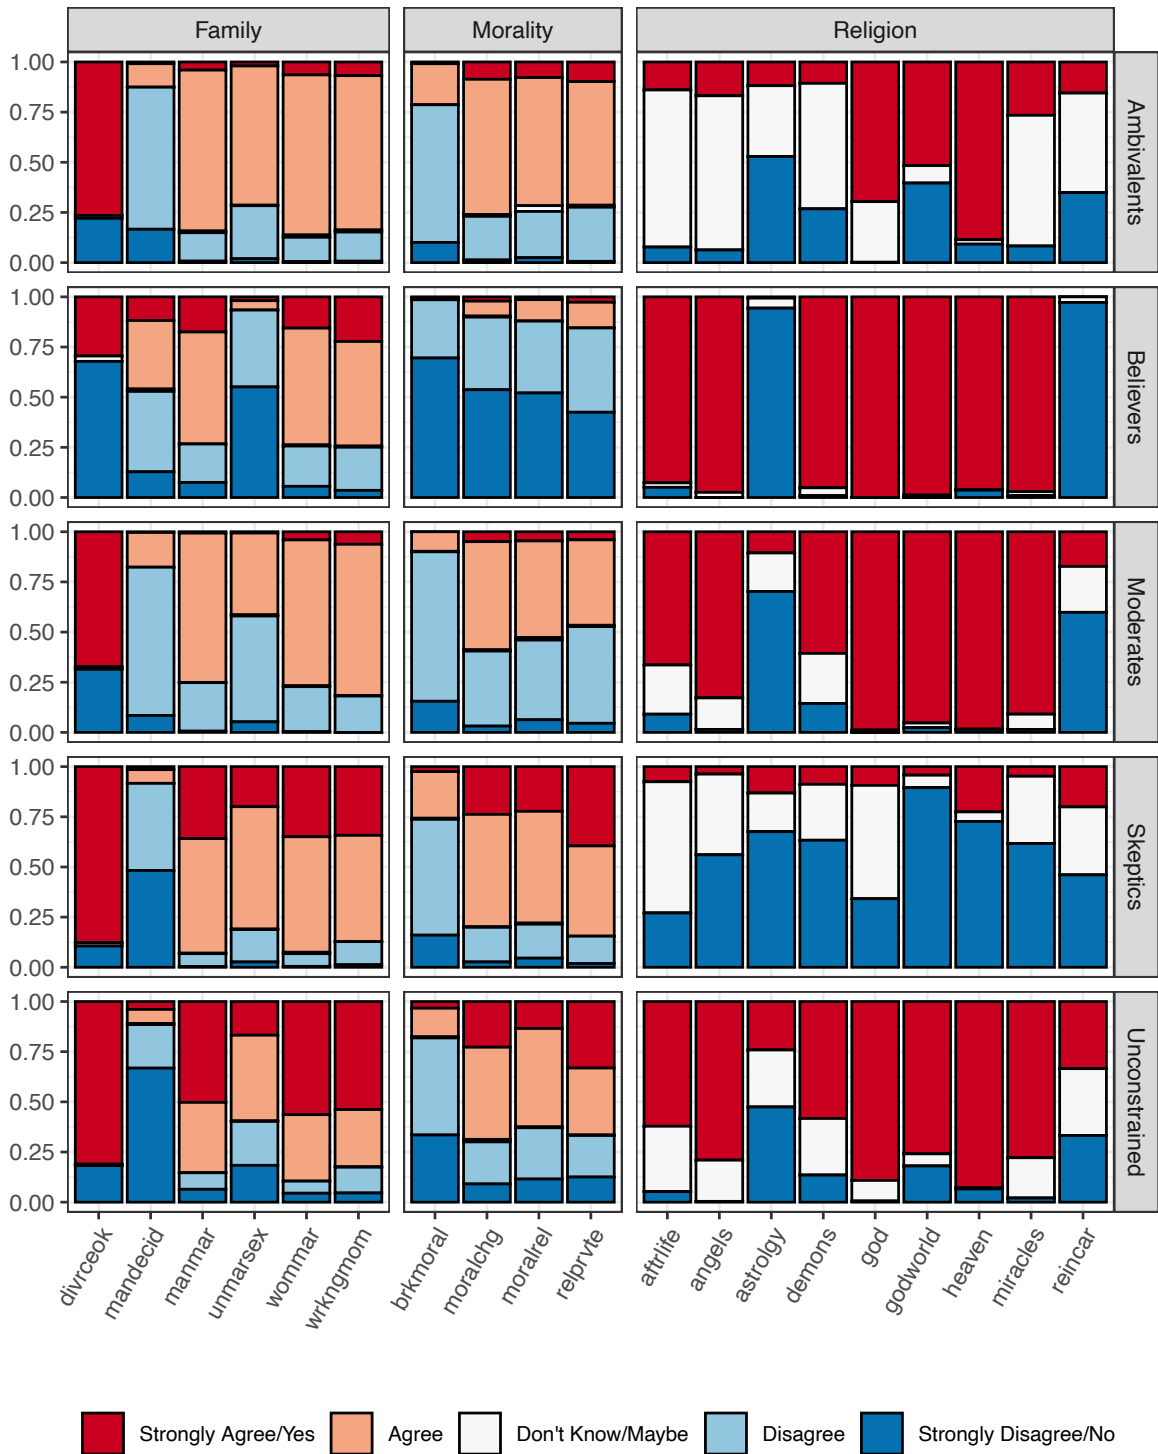


Figure 12: Religious, Moral, and Family Beliefs by Latent Class.

*Believers:* The first group, about 13 percent of survey respondents, has strongly constrained religious and moral beliefs. Almost everybody in this class expresses a belief in the major tenets of Christian theology, and they reject non-Christian beliefs (reincarnation and astrology). They strongly contrast with other classes in being much more likely to say they disagree and strongly disagree with moral relativism (moralchg; moralrel) and the notion that religion is a private matter (relprvte). They are also strongly constrained to believe that sex before marriage is wrong. Identification as an Evangelical Protestant/LDS is a strong predictor of being in this class, and members of this group attend religious services substantially more often than members of other classes.

A key feature of this class is that they are less constrained in their beliefs about family and gender than many of the other classes. This lack of constraint arises because their belief system appears to present them with considerations that are at odds with the prevailing culture. While most classes tend to think that men and women can live a “fully satisfying life without getting married,” members of this group appear to have considerations that push them to occasionally disagree. Similarly, while most other groups are constrained to the “disagree” side of the scale on whether “Most of the important decisions in the life of the family should be made by the man of the house,” members of this group occasionally agree or strongly agree. They are also the group most likely to say that sex before marriage is not acceptable.

Under the belief system framework outlined above, I expect members of this group to be highly unlikely to make changes in their beliefs about religious phenomena, both relative to their other beliefs and relative to other groups. They will also be less likely to change their views of morality than members of other groups. At the same time, because they have conflicting considerations about family structures, they should be more likely to change those beliefs – both more likely to change them than other groups and more likely to change them than other beliefs.

*Skeptics:* The second class, about 14 percent of respondents, is characterized by a rejection or questioning of religious beliefs. They also reject or question astrology and reincarnation. In fact, they look more similar to Believers on these two issues than other classes do. They are the most constrained to the “relative” side of the moral relativism-moral absolutism scales. In terms of covariates, they tend not to identify with a religious denomination or attend religious services. However, people who identify as Jewish cluster in this group.

This group is also constrained on some religious beliefs, some family structure questions, and views on morality, though in the opposite direction from Believers. They should change little on these beliefs over time. They are the least constrained on the question of god’s existence, about equally split between saying “no” and “maybe.” I expect members of this group to vacillate on this question.

**Table 5: Covariates of class assignment.**

	Skeptics	Ambivalents	Believers	Moderates	Unconstrained
Class Prevalence	0.14	0.27	0.13	0.32	0.15
Black Protestant	0.04	0.24	0.07	0.49	0.16
Catholic	0.05	0.45	0.02	0.28	0.19
Indeterminate Religion	0.06	0.23	0.10	0.45	0.16
Evangelical Protestant/LDS	0.01	0.11	0.34	0.44	0.10
Jewish	0.46	0.41	0.01	0.02	0.10
Mainline Protestant	0.07	0.28	0.10	0.38	0.17
No Affiliation	0.48	0.34	0.00	0.06	0.12
Other	0.30	0.15	0.15	0.22	0.18
Female	0.11	0.22	0.14	0.34	0.19
Male	0.16	0.32	0.12	0.30	0.10
South	0.09	0.20	0.16	0.40	0.15
Not South	0.17	0.32	0.12	0.30	0.10
Parents have B.A.	0.19	0.26	0.17	0.26	0.12
Parents no B.A.	0.10	0.27	0.10	0.37	0.16
Mean age	17.65	17.66	17.76	17.67	17.85
Mean service attendance (0-6)	0.53	1.74	5.12	3.26	2.10
Mean years of education	11.53	11.29	11.65	11.41	11.56
Mean parental closeness (1-6)	4.80	4.99	5.26	5.09	5.22
Mean prop. friends share beliefs	0.48	0.54	0.79	0.70	0.63

*Moderates:* This group, the largest in the sample with about 32 percent of respondents, closely resembles the Believers in their responses to questions about religious beliefs, but their constrained religious beliefs do not appear to spill over into other domains. Members of this group appear torn between their religious commitments and the culture of contemporary American society, or at least have not taken the time to reconcile these contradictions, producing relatively high levels of ambivalence on issues of family structure and morality, rarely giving “strong” responses to either. This is the

largest class in the data set, drawing members from all religious groups, principally people who do attend religious services but do not attend them frequently.

*Ambivalent:* The third group, about 27 percent of respondents, is characterized by a high degree of uncertainty on religious and moral beliefs. They are the most likely to say “maybe” in response to questions about the existence of angels, demons, and God, as well as the non-Christian belief questions such as astrology and reincarnation. These respondents tend to be Catholic or unaffiliated with a religious tradition and infrequent service attenders. Members of this group are the least likely to give strong opinions on questions of morality and family structures.

*Unconstrained:* The final group, about 15 percent of the sample, demonstrates little constraint across the board. While it is tempting to interpret this group as displaying idiosyncratic belief systems, the theoretical model suggests that these people have internalized a broad range of considerations that make them highly subject to local influences. As a result, they should vacillate over time on these questions, especially questions about morality, where they give both moderate and strong opinions on both sides of the issue.

Contrasting this Unconstrained group with the Ambivalent group helps draw out the theoretical implications of the belief system model. Members of the Ambivalent group explicitly pick scale midpoints, while members of the unconstrained group range widely. This suggests the former demonstrate a higher level of constraint in their

cognition, recognizing their conflicting considerations and averaging them. Members of the latter group appear not to recognize these conflicts and appear to be more subject to the whims of temporary influences. Where a member of the Ambivalent group might say that morality is a complicated mix of relativism and absolutism, members of the Unconstrained group might strongly agree that morality is relative at one moment and strongly disagree at another.

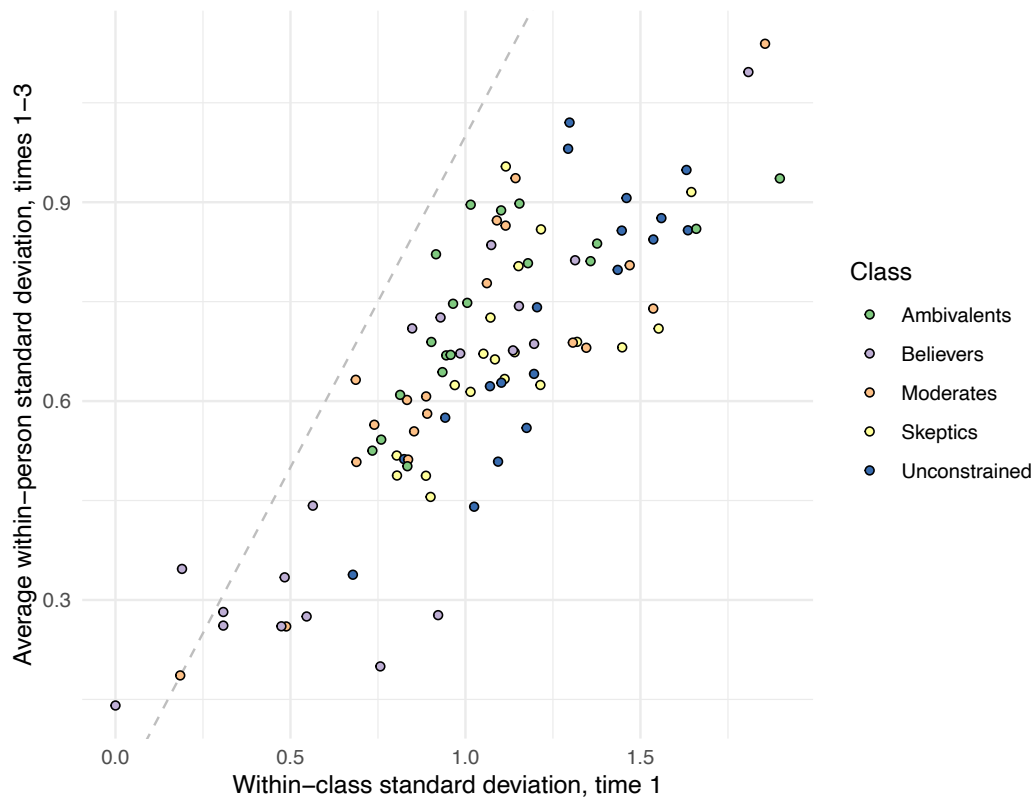
While I call these five groups “belief systems,” it is not necessary that members of these groups see these domains or beliefs as connected. It could be the case that the group I have deemed Moderates do not see connections between family life, morality, and religious behavior. These groups should be conceptualized as socially patterned sets of influences – schooling, parent’s education, social networks, and religious participation – that shape the range of considerations members hold over time.

With the exception of the Unconstrained group, most of these groups present some level of constraint in a portion of the belief space. Both Believers and Skeptics are constrained in their religious, moral, and (most) family beliefs, though in different ways. Moderates have constraints on religious beliefs. Ambivalents are somewhat uniquely constrained to the middle of scales.

#### **4.5.2 Change over Time**

Figure 13 plots the average within-question, within-class standard deviation at time 1 against the average within-question, within-person standard deviation over time

for that class. If within-group constraint of a particular question is a good proxy for within-person considerations, there should be a positive correlation between these two measures.



**Figure 13: Scatterplot of within-group standard deviation at time 1 against average within-person standard deviation, times 1 through 3.**

There is an incredibly strong relationship between the amount that a particular question varies within a group at time 1 and the average within-person standard deviation that members of the group demonstrate over time ( $\rho = 0.829$ ). This relationship holds across questions within groups (lowest correlation is 0.650 for Skeptics; highest correlation is 0.876 for the Believers) and within questions across

groups, ranging from 0.449 for belief in the afterlife to 0.969 for whether divorce is acceptable.

This does not just reflect different degrees of constraint across groups. While Believers are highly constrained on several questions, with a cluster of points in the lower left of the figure, they also have one of the two least constrained beliefs of any group – their view on divorce – in the top right of the figure. Most groups have some questions on which they are relatively constrained and relatively unconstrained, and this appears to equate to levels of stability and change over time.

The diagonal line in Figure 13 represents a 1:1 prediction of within-group variance at time 1 and average within-person variance over time, which is what we would expect if the variation observed in the classes was a perfect prediction of people's internal belief systems. Points generally fall below the line, suggesting that people are more constrained than their belief systems would suggest. This is not surprising, as idiosyncratic cognitive and social forces likely further constrain beliefs over time. But many attitudes fall very close to this line, suggesting a good fit with the theoretical model. A similar comparison of within-group means at time 1 to the groups' average means at times 2 and 3 also produces a very strong positive correlation, close to the 1:1 prediction line ( $\rho = 0.978$ ).

To test hypotheses 1 and 2, I estimate a regression of within-person change between times 2 and 3 on within-group variance at time 2 with fixed effects for question

and person. Table 6 presents the results of that regression. I also present a regression of the group mean at time 1 on within-person means for times 2 and 3 to examine whether people are, in general, constrained to the same portion of the belief space.

**Table 6: Results of fixed effects linear regressions of within-person, within-question standard deviation and mean on within-group standard deviation and mean.**

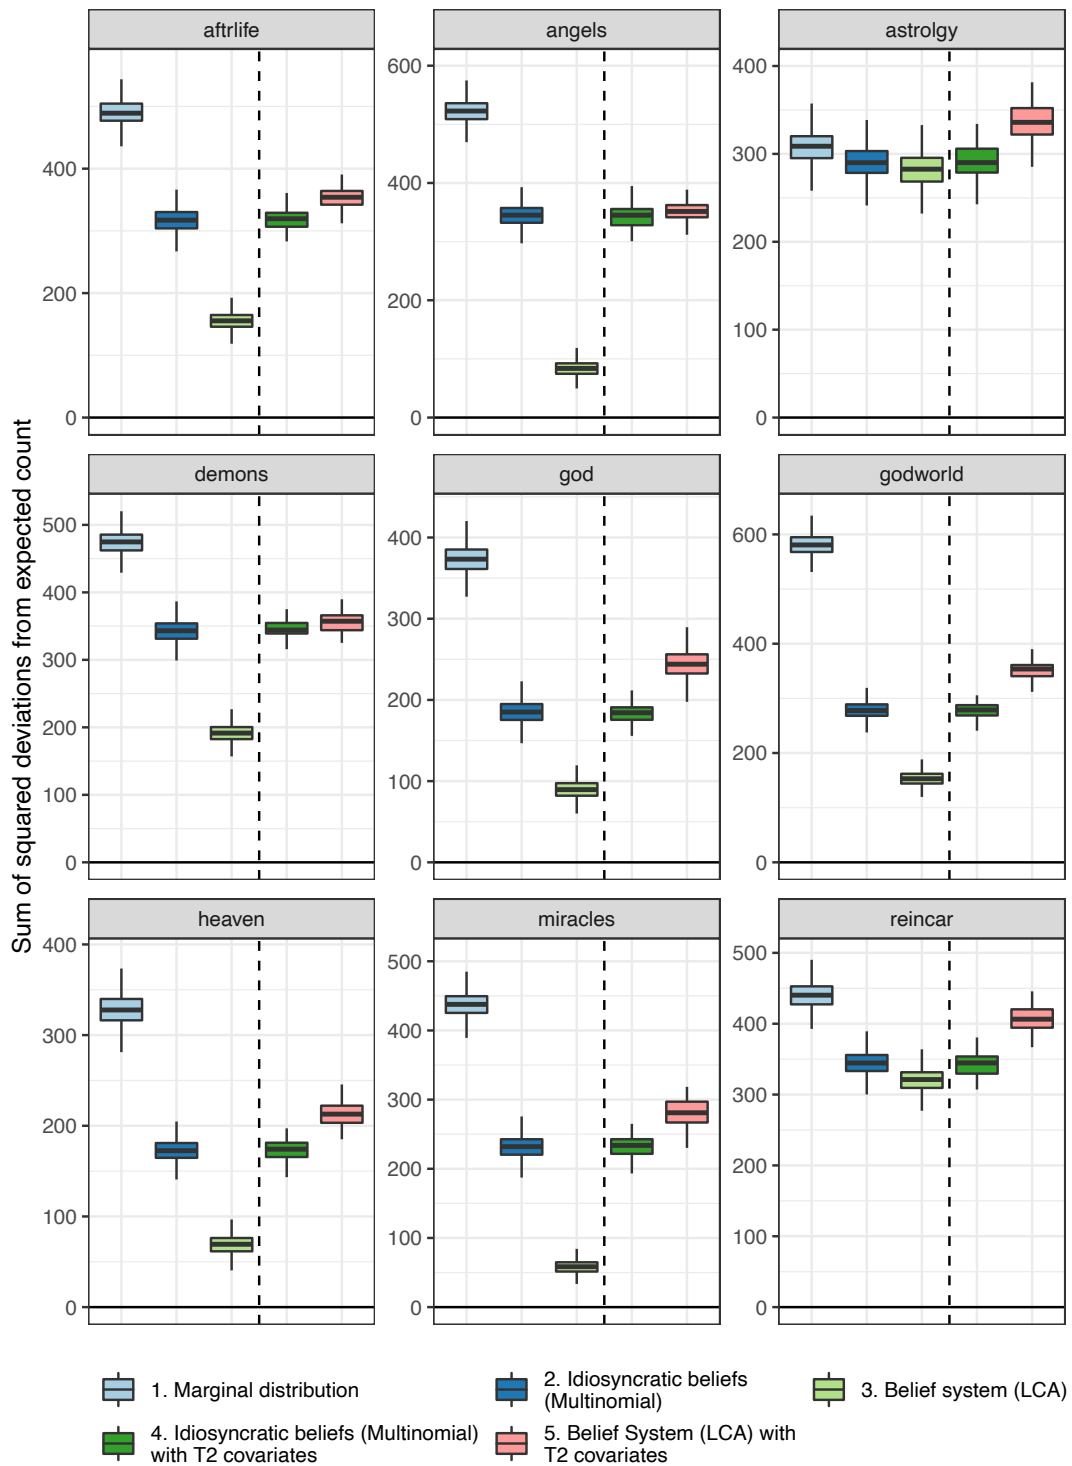
	Model 1: Individual standard deviation (time 1-3)	Model 2: Individual mean (time 2-3)
Within-group, within-question standard deviation (time 1)	0.55 (0.01)	
Within-group, within-question mean (time 1)		0.80 (0.01)

As expected, Table 6 shows a strong positive association between within-group variance at time 1 and within-person change over time. In other words, consistent with Hypothesis 1, people are more likely to change attitudes that are less constrained in their group at time 1. And consistent with Hypothesis 2, within questions, groups that are less constrained exhibit more change their answers over time. The second columns show that people are, on average, highly constrained to parts of the belief space indicated by their class at wave 1, even if they sometimes have wide variance within that.

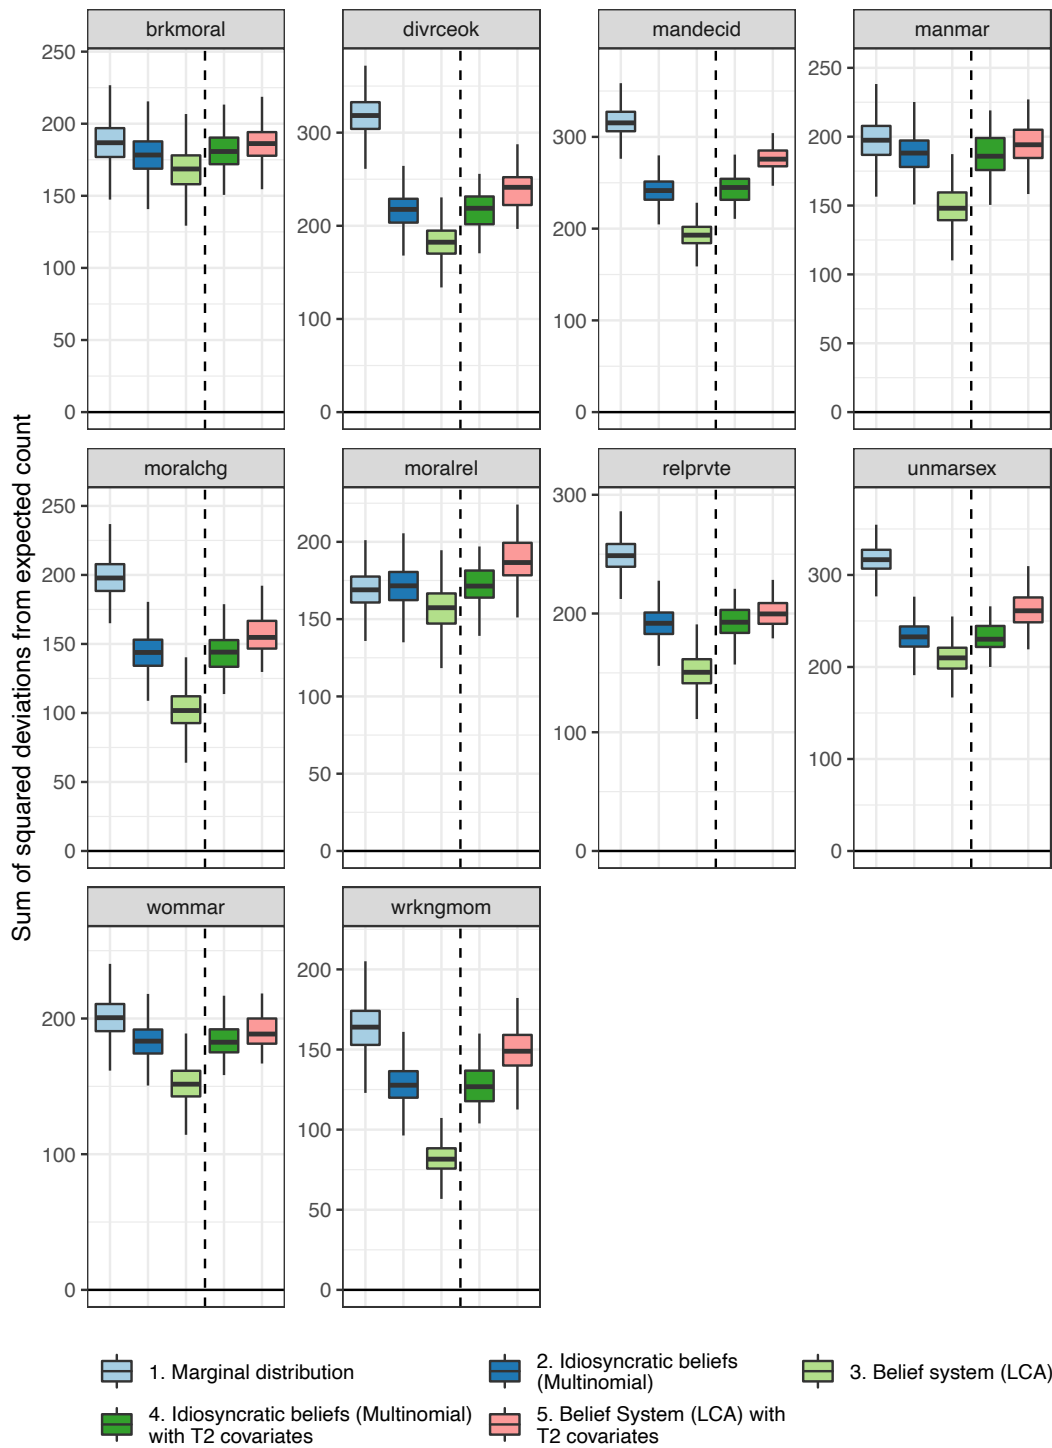
To summarize, the takeaway from these results is that the Latent Class Analysis produces a set of classes that not just group people with similar attitudes but similar propensities to give different responses over time. The Unconstrained group is not people with a set of idiosyncratic beliefs that do not fit into any of the other categories. They are a set of people with uniquely unstable views that can range widely over time. The Moderates are not people with similar religious beliefs but different family structure beliefs. They have stable religious beliefs and unstable moral and family beliefs.

### **4.5.3 Predicting Response Patterns**

Hypothesis 3 predicts responses people give over time will reflect a multinomial draw from the probabilities derived from the belief system at time 1. Figure 14 and Figure 15 compare the accuracy of predicted counts of response patterns with predictions generated using the marginal distribution at wave 3 and predictions generated from the multinomial logit model reflecting independent belief formation, for the religion variables and non-religion variables, respectively. While an overall count of accuracy across all beliefs can be estimated, beliefs are plotted separately to highlight where the belief system improves predictive accuracy and where it does not. Because the marginal distributions shape the total amount of deviance possible, each plot has a different y-axis, though all begin at 0.



**Figure 14: Boxplots of deviation of observed from expected counts of response patterns, 10,000 iterations each, religious beliefs.**



**Figure 15: Boxplots of deviation of observed from expected counts of response patterns, 10,000 iterations each, moral and family beliefs.**

The predictions generated through the latent class model consistently outperform the predictions made through the multinomial logit model. For some questions, especially the religious beliefs, differences between the two models are stark, with the belief system model improving prediction substantially over the marginal distribution and the independent beliefs model. On questions about the existence of belief in miracles, angels, and heaven, the latent class belief system model is very accurate. Other questions are less conclusive, but the latent class model still tends to outperform the independent beliefs model on average. Unsurprisingly, the model appears to perform best on questions when the five belief systems impose different constraints (e.g., belief in miracles) and worst when belief systems present similar constraints (e.g., whether morals are relative). Because there are not stark differences in the constraints the systems impose around whether it is acceptable to break moral rules for personal gain or belief in astrology, there is not much room for either model to improve upon the marginal distribution.

The key interpretation of this comparison is that information about people's beliefs tell us something about their other beliefs, suggesting an underlying structure to the belief system. The assumption that responses at any wave are a multinomial draw from the distributions deduced at time 1 is well supported, especially for the Christian religious beliefs.

Both the multinomial logit (independent beliefs) and the latent class (belief system) model assume the belief systems that people have, whether idiosyncratic or culturally shared, are durable over time. They cannot adjudicate whether beliefs are principally a result of stable social circumstances or durable cognition. I now turn to adjudicating that question.

#### 4.5.4 Changing Social Structures

The final hypothesis explores whether changes organizational participation and social networks better account for changing beliefs than the belief systems explored above. Table 7 shows how the organizational participation and social structure variables change between times 1 and 2. There are two notable changes that are likely to produce shifts in beliefs. First, people decrease the amount they attend religious services. Second, there are shifts in religious affiliation. People move out of the indeterminate religion group and the “no affiliation” group grows. Since both affiliation and frequency of attendance strongly differentiated the five classes, changes in these variables should produce changes in beliefs under a non-durable cognition model. Parental closeness, and the proportion of one’s close friends who share their beliefs remain relatively unchanged.

**Table 7: Summary statistics for covariates of belief prediction, times 1 and 2.**

	Time 1 Mean	Time 1 S.D.	Time 2 Mean	Time 2 S.D.
Age	17.68	1.34	20.00	1.43
Freq. of service attendance (0-6)	2.62	2.22	2.01	2.10

Black Protestant	0.06		0.07	
Catholic	0.20		0.17	
Indeterminate Religion	0.14		0.04	
Evangelical Protestant/LDS	0.29		0.30	
Jewish	0.04		0.04	
Mainline Protestant	0.08		0.11	
No Affiliation	0.16		0.24	
Other	0.02		0.03	
South	0.40		0.40	
Years of education	11.46	1.41	12.83	1.54
Mean parental closeness (1-6)	5.07	0.86	5.12	0.91
Mean prop. friends share beliefs	0.63	0.36	0.61	0.35

To the right side of the dashed line in each plot in Figure 14 and Figure 15 are predictions generated using time 2 covariates. If changing social circumstances lead to changes in attitudes, these models should outperform those that do not.

In general, the time 1 belief system predictions outperform both the idiosyncratic influence model (the multinomial model) that uses time 2 covariates and a model that uses covariates at time 2 to revise membership in belief systems. There is little change in the predictions for the multinomial logit model compared to just using wave 1 covariates, but the latent class model performs worse when we account for changes in social structure.

In other words, a model that assumes attitudes are principally a reflection of current social structures makes worse predictions for people's attitudes than one that

assumes no influence of contemporaneous social structures other than what is caused by prior experiences.

#### **4.6 Discussion**

This paper had two related goals. First, it sought to rethink how researchers interested in the measurement of culturally constrained cognition conceptualize and measure the constraints of these belief systems in observational data. Because the cultural structuring of cognition happens well below the level of a survey response, measuring culture as the relationship between responses can produce misleading conclusions about culturally patterned cognition. Existing measures of schematic cognition do not fully take into account how culturally patterned cognition might produce variation in survey responses over time that reflect, rather than undermine, those cognitive structures. I argued that Latent Class Analysis could be used to deduce belief systems that manifest as shared probabilities of giving certain responses to different questions over time.

Using LCA, I deduced five belief systems in the population of adolescents surveyed in the NSYR regarding family structure, morality, and religious beliefs. These systems differ in the constraints (and lack of constraints) they place on attitudes. Results from the first three hypothesis tests find that the within-group constraints implied by the model at time 1 are reflected in individuals' change over time. The key finding of this paper is that the response probabilities derived across people at a single point in

time provide a remarkably good prediction of how people changed over time. In other words, the population contains groups that are stable on some issues but unstable on others. Tools designed to capture culture by looking at just whether people give the same responses will fail to capture these distinctions.

The five belief systems identified above should not be interpreted as sharply divided “cultures” on these issues. They are somewhat porous classifications that represent common belief constraints created by patterned exposure to different cultural influences – churches, education, family structures, etc. – and some people have roughly equal probabilities of being assigned to multiple classes.

As a second goal, the paper sought to adjudicate the relative influence of these cultural-cognitive structures on attitude behavior over time compared to social-structural influences such as social networks and organizational participation. The fourth hypothesis test finds that the constraints (or lack of constraints) present at time 1 better predict the aggregate pattern of responses over time than models that directly root attitudes in contemporaneous social circumstances. Despite competing expectations that changes in organizational participation and social influences should produce changes in beliefs, the constraints deduced at time 1 appear to strongly limit the effect of these changes, especially on religious beliefs and some family beliefs. This suggests that these underlying dispositions are quite durable, even as they produce heterogeneous responses.

#### **4.6.1 Implications for Cultural Sociology**

The assumption that durable cultural cognition has to manifest as durable attitudes reflects an implausible assumption about the relationship between cultural schemas and attitudes. The results presented here suggest that people's cognition is more culturally structured and more durable than lots of work supposes, but that durable cognitive structures often produce unstable "attitudes." Rather than constraining people to certain responses, culture shapes how people internalize and recall various considerations over time. On some dimensions, they tend to be highly stable. On others, they vacillate. But these variations in stability and instability are patterned. In the analysis presented above, people who are stable in their Christian religious and absolutist moral views tend to be unstable in their views on family structures. People who are stable on their progressive family views tend to be unstable on their religious views.

Because these variations in attitudes are durable over time, the findings presented here suggest that research focused on explaining attitude behavior over time might put too much emphasis on contemporary social structures at the expense of cognitive structures formed early in life.

Because culture makes room for heterogeneity, it has been difficult to understand when people make meaningful changes in their beliefs. The method above deduces plausible limits on attitudes within a belief system. This provides a way to understand

when people might be making real, lasting changes in their cultural beliefs and when these changes are temporary. When a Moderate in the belief systems deduced above changes from moral relativism to moral absolutism, it likely reflects nothing more than temporary influences. If a Believer makes the same shift, especially if he also shifts his view of religion as a private matter, it likely points to something more durable.

Understanding the limits of belief systems and identifying when people transgress them can point to the social structures that do induce substantive changes in beliefs.

The belief systems deduced above also suggest that some people are more open to (temporary) social influences than other people. A change in the social environment – moving from the south to the northeast, for example – might lead an Unconstrained person to report different religious beliefs, but that same change should not affect a Believer. This means that the average effect of some structural change on an attitude is going to be shaped by how these constraints vary across the population. Without deducing cultural constraints first, however, we might not be able to make sense of this kind of treatment effect.

Because belief systems appear to be quite durable in the face of changing circumstances, there is no reason to expect that inconsistency in survey responses should be thought of as undermining the causal power of cultural beliefs (Jerolmack and Khan 2014). Inconsistency in survey response should be thought of as reflecting a misalignment between the structures of cognition and the survey question, not reflecting

unstructured thought. Studies that find a strong predictive link between survey questions and attitudes might simply get lucky in finding both a question and behaviors that are strongly predicted by a belief system. But better measures of belief systems could potentially show stronger links between these underlying cognitive structures and social action. As Harding (2007) shows, external cultural heterogeneity predicts social behaviors. Internal cultural heterogeneity might also produce specific lines of behavior. A schematic structure that produces inconsistency on the divorce question potentially causes more divorces than one that constrains beliefs over time to oppose divorce.

To be clear, these results should not downplay the role of organizational structures and social networks in shaping people's belief patterns. Organizational participation, especially religious tradition and frequency of religious service attendance, was a strong predictor of belief systems at time 1 (and, therefore, beliefs in subsequent waves). These "past conditions of production" should be considered particularly important in shaping people's cognition over time (Bourdieu 1990).

#### **4.6.2 Implications for Methods**

There are undoubtedly more direct measures getting at people's underlying schemas (Hunzaker and Valentino 2019) or the unconscious processes that shape people's cognition (Miles, Charron-Chénier, and Schleifer 2019). But to the extent that researchers want to understand the limits that culture places on people's cognition, LCA appears to do a reasonable job.

The results pose a challenge to how we understand measurement error. The central assumption of measurement error arguments is that people have an underlying position, and that inconsistency is inherent in questions, not in people (Alwin 2007; Ansolabehere, Rodden, and Snyder 2008). The theoretical model outlined here builds on previous arguments suggesting people do not have “attitudes” or “beliefs” as measured by surveys. They have considerations that probabilistically produce responses under a variety of conditions (Zaller 1992). What the latent class analysis shows is that inconsistency is a relationship between people’s range of beliefs and the question they answer. Some belief structures strongly constrain people to present consistent responses, but others free people to give a range of attitude reports over time. Moderates, Unconstrained, and Ambivalents have similar average responses on a number of questions. Where they differ is in their variance, but that variance is predicted by their other beliefs, not random error.

#### **4.6.2 General Conclusions**

Measuring culture is challenging, and measuring culture in the morass that is people’s cognition is even harder. For a long time, researchers looked for the signature of shared culture in shared attitudes and stability. But shared culture is often heterogeneous and conflicting, and the product can be shared instability. But that does not make this any less “culture.” Taking seriously the measurement of culture means accounting for this instability.

## 5. Conclusions

This dissertation set out to understand the variance in people's attitude reports over time. The three studies detailed above present a range of findings about attitude reports over time and the conditions under which people form and change their attitudes, attempting to draw general conclusions across diverse topic domains. The first study suggested that people update their views throughout adolescence and early adulthood. But following these stages of the life course, attitudes stop durably updating, with rare, instructive exceptions. The second study suggests that this over-time consistency masks short-term heterogeneity within people. While some people give stable responses to questions over time, the majority are characterized by ambivalence – giving responses on different sides of an issue over time as local (random) conditions shift. However, these shifts (in adults) are not durable. It also found that when people are stable over time, they are stable on related issues, but which issues people are stable on differs significantly. The third paper suggests that this over-time stability and instability itself patterned and shaped by exposure to socializing institutions when people are young. For the moral, religious, and family structure beliefs explored here, these institutions include churches, schools, and social networks.

What do these findings mean for the sociology of culture? In this chapter, I attempt to use these findings to draw general conclusions about the social and cultural origins of attitudes and opinions and how they evolve and change over the life course.

As noted at several points throughout the preceding chapters, the findings presented here are strongly consistent with a toolkit model of culture in which people indiscriminately internalize the heterogeneous mixture of considerations present in the cultural environment (DiMaggio 1997; Martin 2010; Swidler 2001). Because people are, in general, unconcerned with most issues that are asked about in a survey, they rarely take the time to reconcile these conflicts, and they can be pushed by local circumstances to give a variety of responses to the same question. However, when something about the social environment drives people to reflect on their beliefs, they are quite capable of harmonizing and stabilizing beliefs.

Just because people are ambivalent, however, does not mean that all people have the same consideration sets. The fourth chapter of this dissertation suggests that people's variance in attitude reports is also socially patterned. Some peoples' belief systems constrain them to narrow portions of the belief space while others free them to express a wide range of beliefs, even if these groups have the same central tendency. Importantly, two groups of people can be constrained and unconstrained on opposite beliefs. These beliefs also act on each other. If a person is constrained on two dimensions, they tend to be ambivalent on a dimension that asks for them to reconcile these competing beliefs. In Chapter 4, for example, people with traditional religious beliefs and progressive beliefs about family structure struggle to maintain consistency on questions of morality, which put these domains in conflict.

Together, the three chapters also suggest that people's belief systems – the range of considerations they can bring with them to a survey interview – are formed relatively early in life. People do not make durable changes in their attitudes over time, even as they vary in their responses. This suggests that if we want to understand why people believe what they believe, we are often better off looking at their cultural background – church participation, educational attainment, neighborhoods, etc. – than their contemporaneous social environments.

A key contribution of this dissertation is to draw parallels between this model of culture and work on public opinion, particularly Zaller's model of ambivalence and opinion formation, and suggest places where they can inform each other. Cultural sociology and cultural anthropology provide an empirically grounded a model of human cognition that can underlie the actor at the heart of Zaller's theory. The life course patterns of openness to new considerations highlighted here likely interact with elite messaging, forming generational effects in political dispositions. Similarly, extending Zaller's theory to other domains gives us expectations for when and why we observe inconsistency over time in some areas but not others. At the macro-level, domains such as morality tend to resemble the political field where people regularly hear conflicting considerations and have little incentive to reconcile these conflicts. This helps explain why, despite the frequent assumption that moral worldviews are "core" beliefs that shape our ability to hold other beliefs, I observe high levels of inconsistency

in them. In other domains, such as specific sexual prohibitions, there is little public disagreement among opinion elites that married people should not have sexual relations with others and that teens should wait until they are at least 16 to have sex. As a result, people are able to give highly consistent answers to these questions.

Analyzing the correlation of stability measures in Chapter 3 and observing the effects of organizational participation identified in Chapter 4 suggest that meso-level structures can also stabilize views even if there is significant disagreement in the general public. This is another area where cultural sociology expands Zaller's model (Hardin 2007, 2011; Frye and Trinitapoli 2015). Zaller suggests that when there is macro-level disagreement, stability is principally driven by individual attention to the issues in question. However, while people might hear diverse considerations at the macro-level, their local circumstances might induce them reject some of these considerations, at least at the stage of recall, even if they do not pay significant attention to these debates. People embedded in conservative religious communities are surely aware of arguments for moral relativism (they are likely invoked frequently to argue against), but these people have no problem disagreeing with the idea over time. Sociological research directs attention to the features of organizations, social networks, and institutions that might facilitate stability beyond macro processes (Lizardo and Strand 2010; Martin 2002, 2010; Rawlings 2020).

As a general conclusion, the theoretical framework advanced here suggests that people's inconsistency arises not because they are ignorant or uninformed, but because they have too many considerations and do not prioritize them in their day to day lives. I posit that people maintain this cultural inconsistency (and might be cognitively structured to do so) because day-to-day social life is inherently complex. The buzzing mass of social interactions gives people little incentive to develop totalizing worldviews and every incentive to keep their cultural repertoire diverse and conflicting. Rigid ideologies close off social relationships, limit social opportunities, and hamper the acquisition of new considerations. Totalizing worldviews might keep people from developing new ideas or achieving desired ends. People need diverse cultural tools to maintain marriages (Swidler 2001) and friendships, to code-switch between different components of their social network, and to harmonize life across institutional domains of the economy, politics, family, and religion (Eliasoph 1997). Taken from this perspective, inconsistency and change in survey responses should not be surprising. What is surprising, and what deserves increased theoretical attention in the social sciences, are the circumstances that lead people to stabilize their views over time.

## Appendix A: Chapter 2 Phi Estimates







Figure 16: Phi (updating) estimates for items about religious activity and beliefs, social life, subjective socioeconomic status, and suicide.



**Figure 17: Phi (updating) estimates for items about guns, law, crime, policing, government, and public spending.**



**Figure 18: Phi (updating) estimates for items about civil liberties, confidence in institutional leadership, health, morale, and trust.**



**Figure 19: Phi (updating) estimates for items about race, gender, sex, sexuality, and abortion.**

## Appendix B: Chapter 3 Finite Mixture Model Estimation

The following is a summary of the model used here, outlined in Hill and Kreisi (Hill and Kriesi 2001b), with modifications for three waves rather than the four they use in their analysis. Comparing a four-wave model to a three-wave model using data from the NSYR and LSG, which have more than three waves, to a model with the same data, dropping an interior wave (wave 2 or 3), produces comparable estimates for most coefficients, especially behavioral group proportions ( $\pi_1$ ,  $\pi_2$ , and  $\pi_3$ ), suggesting the model behaves just as well with three waves.

The goal of the Data Augmentation algorithm is to draw a plausible set of parameters,  $p(\theta|Z)$ , compatible with the observed data ( $Z$ ) and the assumptions made in the model. It iterates through two main steps: Drawing group membership indicators, given the parameter, and drawing parameters given the group membership indicators. I first outline the variables that go into the model and then outline those steps below.

### B.1 Re-Expressing the Data

Rather than estimate parameters on the basis of full patterns, such as “agree”-“disagree”-“strongly disagree,” I characterize a respondent by a set of general variables. These variables retain the features of the data without having to generate a probability for each pattern separately. For example, by re-expressing the data in this manner I collapse the distinction between someone who says “agree”-“strongly agree”-“agree”

and someone who says “agree”-“agree”-“strongly agree.” Both patterns become someone who remains on the same side of an issue and gives one strong response.

$$\begin{aligned}
A_i &= \begin{cases} 1 & \text{is the } i\text{th person's initial response is a 4 or 5} \\ 0 & \text{otherwise} \end{cases} \\
B_i &= \text{number of the } i\text{th individual's responses that are either 1 or 5 across all } t \\
C_i &= \text{number of the } i\text{th individual's responses that are 3} \\
D_i &= \text{number of times the } i\text{th individual crosses an opinion boundary} \\
E_i &= \begin{cases} 0 & \text{if } D_i \neq 1 \\ t_i^* & \text{otherwise} \end{cases} \\
F_i &= \begin{cases} 0 & \text{if the } i\text{th individual's preswitch response is a 1, 2, or 3, or } D_i \neq 1 \\ 1 & \text{if the } i\text{th individual's preswitch response is a 4 or 5} \end{cases} \\
H_i &= \begin{cases} 0 & \text{if the } i\text{th individual's preswitch response is a 1, 2, 4, or 5, or } D_i \neq 1 \\ 1 & \text{if the } i\text{th individual's preswitch response is a 3} \end{cases} \\
M_i &= \begin{cases} 0 & \text{if the } i\text{th individual's postswitch response is a 1, 2, or 3, or } D_i \neq 1 \\ 1 & \text{if the } i\text{th individual's postswitch response is a 4 or 5} \end{cases} \\
Q_i &= \begin{cases} 0 & \text{if the } i\text{th individual's postswitch response is a 1, 2, 4, or 5, or } D_i \neq 1 \\ 1 & \text{if the } i\text{th individual's postswitch response is a 3} \end{cases} \\
R_i &= \text{number of the } i\text{th individual's responses that are either 1 or 2 across all } t
\end{aligned}$$

I denote the vector of these random variables for individual  $i$  as  $Z_i$ . For the seven-point scale, I include an additional variable that counts “weak” responses.

## B.2 Drawing Group Indicators Given Parameters

I use the observed data  $Z_i$  and parameter estimates to determine the probability that a person falls into each behavioral group. The conditional probability of belonging to each behavior group given the observed responses and estimated parameters is expressed by the following functions:

$$\begin{aligned}
p_{1,i}^Z &= (\alpha_i^{(1-A_i)}(1-\alpha_1)^{A_i}(1-\delta_1)^{(3-B_i)}\delta_1^{B_i})I(C_i=0)I(D_i=0) \\
p_{2,i}^Z &= \varphi_2^C \alpha_2^{R_i} (1-\varphi_2-\alpha_2)^{(3-C_i-R_i)} \delta_2^{B_i} (1-\delta_2)^{(3-C_i-B_i)} \\
p_{3,i}^Z &= [\tau_3^{I(E_i=1)}(1-\tau_3)^{I(E_i=2)}(\varphi_3^{(pre1)H_i} \alpha_3^{(pre1)(1-F_i)(1-H_i)} (1-\alpha_3^{(pre1)} - \varphi_3^{(pre1)})^{F_i})^{I(E_i=1)} \\
&\quad \times (\varphi_3^{(pre2)H_i} \alpha_3^{(pre2)(1-F_i)(1-H_i)} (1-\alpha_3^{(pre2)} - \varphi_3^{(pre2)})^{F_i})^{I(E_i=2)} \\
&\quad \times ((1-\alpha_3^{(post)})^{M_i} (\alpha_3^{(post)})^{(1-M_i)})^{H_i} \delta_3^{B_i} (1-\delta_3)^{(3-C_i-B_i)}] I(D_i=1) I(Q_i=0)
\end{aligned}$$

where  $I(\cdot)$  equals 1 if the condition in the parentheses is met and 0 otherwise. These probabilities represent the expected frequency of a particular pattern given the parameters that describe behavior in that opinion behavior group. For example, if a person responds “no opinion” or “neither agree nor disagree” in one wave, then  $D_i \neq 0$ , and  $p_{1,i}^z$  will equal 0 for that person, since people who express “no opinion” cannot be opinion holders. If  $\alpha_1 = .5$  and  $\delta_1 = .75$ , we expect patterns where a person gives a “strongly disagree” response in all three waves to make up about 21 percent of responses in group 1.

Generating these probabilities requires parameter estimates. In the first iteration of the model, these parameters are drawn at random from the parameter space. These random draws are only bound by logical constraints, such as proportions summing to 1.

With these probabilities and draws for  $\pi_1$ ,  $\pi_2$ , and  $\pi_3$  (sampled randomly from the parameter space in the first iteration), I sample from the following trinomial distribution for each person to generate group membership indicators for each iteration.

$$p(G_i|\theta, Z_i) = \text{Mult}\left(\frac{\pi_1 p_{1,i}^z}{\sum_{j=1}^3 \pi_j p_{j,i}^z}, \frac{\pi_2 p_{2,i}^z}{\sum_{j=1}^3 \pi_j p_{j,i}^z}, \frac{\pi_3 p_{3,i}^z}{\sum_{j=1}^3 \pi_j p_{j,i}^z}\right)$$

These draws generate group membership indicators for a single iteration of the model. The probability that a person’s responses are generated from a particular pattern are a function of the likelihood that a pattern was generated by a particular behavioral group,  $p_{i,j}$ , and the overall prevalence of that behavioral group in the population,  $\pi_j$ . While a pattern “strongly agree”-“agree”-“strongly agree” might be most consistent

with a stable opinion, as the proportion of the population vacillating changers increases, the probability that the vacillating group produced this pattern increases as well. While it might be unlikely (probability = 0.001) that someone behaving consistent with the vacillating changer group in the gay marriage example would give three “strongly agree” responses, the probability we observe someone in the sample doing that would be much larger if the group is 1,000 people rather than 10.

### B.3 Drawing Parameters Given Group Indicators

With group memberships temporarily assigned for one iteration, I then draw parameter estimates from their distribution conditioning on the data ( $Z_i$ ) and the group indicators. In other words, I fit separate models for each (temporarily assigned) behavioral group to generate parameter estimates for that iteration.

I use Beta and Dirichlet distributions for prior distributions for each parameter. Assuming a priori independence of the appropriate parameters, I factor  $p(\theta)$  into six independent Beta distributions –  $p(\alpha_1, 1 - \alpha_1)$ ,  $p(\delta_1, 1 - \delta_1)$ ,  $p(\delta_2, 1 - \delta_2)$ ,  $p(\alpha_3^{(post)}, 1 - \alpha_3^{(post)})$ ,  $p(\delta_3, 1 - \delta_3)$ ,  $p(\tau_3, 1 - \tau_3)$  – and four independent Dirichlet distributions –  $p(\varphi_2, \alpha_2, 1 - \varphi_2 - \alpha_2)$ ,  $p(\varphi_3^{(pre1)}, \alpha_3^{(pre1)}, 1 - \varphi_3^{(pre1)} - \alpha_3^{(pre1)})$ ,  $p(\varphi_3^{(pre2)}, \alpha_3^{(pre2)}, 1 - \varphi_3^{(pre2)} - \alpha_3^{(pre2)})$ ,  $p(\pi_1, \pi_2, \pi_3)$ .<sup>22</sup> Parameters are then drawn from the appropriate

---

<sup>22</sup> In the seven-point-scale model,  $\delta$  and  $\gamma$  estimates are drawn from a Dirichlet distribution (one for each behavioral group), rather than the Beta distribution for the  $\delta$  estimate in the five-point-scale model. In the three-point-scale model, no  $\delta$  parameter is estimated.

posterior distributions, using only people in that particular behavioral group. These posterior distributions amount to summing up people or responses that demonstrate that behavior and adding the relevant prior to the count. For example, the posterior distributions for the three  $\pi_j$  parameters and the two group 1 behaviors are:

$$\begin{aligned}
 p(\pi_1, \pi_2, \pi_3) &= \text{Dirichlet}\left[\left(\sum_{i=1}^N G_i = 1 + 2\right), \left(\sum_{i=1}^N G_i = 2 + 2\right), \left(\sum_{i=1}^N G_i = 3 + 2\right)\right] \\
 p(\alpha_1, 1 - \alpha_1) &= \text{Beta}\left[\left(N - \sum_{i=1}^N A_i + 1\right), \left(\sum_{i=1}^N A_i + 1\right)\right] \\
 p(\delta_1, 1 - \delta_1) &= \text{Beta}\left[\left(\frac{\sum_{i=1}^N B_i}{3} + 1\right), \left(N - \frac{\sum_{i=1}^N B_i}{3} + 1\right)\right]
 \end{aligned}$$

The formulas look complicated, but they amount to counting the number of response patterns that follow the particular rule the parameter pertains to. This is most obvious in the case of the  $\pi_j$  estimates, which are a Dirichlet draw from the count of each person assigned to that group for the current iteration, plus two additional cases in each group (the prior). Similarly, the estimate for  $\alpha_1$ ,  $(N - \sum_{i=1}^N A_i + 1)$ , is simply the total number of people in Group 1 (opinion holders),  $N$ , minus the number of people who disagree,  $\sum_{i=1}^N A_i$ , plus 1 (the prior). Since opinion holders can only agree or disagree, this gives the count of people who agree, which is what  $\alpha_1$  indicates.

As noted above the priors are functionally equivalent to adding two people to each behavior group and splitting up their behavior within those groups. Hill and Kreisi find that changes in these priors have minimal effects on the actual estimates. The

greatest instability is observed in the durable changers group, which is often small and therefore more susceptible to the influence of priors.

## **B.4 Convergence**

To prevent the randomly drawn starting values to unduly influence the parameter estimates, I draw five different sets of starting parameters for each question. The model iterates until these five chains reach a stable distribution, from which draws were taken to capture the posterior distribution. Functionally, I used five chains, each with 2,500 iterations. I discarded the first 500 iterations to eliminate estimates generated prior to convergence. Diagnostics suggest that, for most questions, parameter estimates converged well before the 500th iteration. These posterior distributions are then summarized.

## Appendix C: Chapter 3 Regression Model

To better assess the influence of question content and structure, I tagged questions with broad indicators for content areas: government and politics; civil liberties; medicine and science; race; sex; family and gender;<sup>23</sup> socioeconomic position; work; morality; foreign relations; self-evaluations; religion; crime; social trust; and other. These categories are not mutually exclusive, and many questions are asked in a survey because feature the intersection of multiple domains. Because of this, I focus on an explicit reading of the question and code items parsimoniously. Because more than half of all questions analyzed here touch on government and politics in some capacity, I instead use an indicator to signal that a question does *not* explicitly deal with politics. I also include indicators for whether a question asks people to assess another person's position, and indicators for whether the question refers to a clear group of people or a particular person.

I also create indicator variables for question structure. First, I include five indicators about question format. The first format asks whether people agree with a statement. These include traditional five- and four-point Likert scales and two- and three-point agree/disagree or yes/no statements. The second format asks people to select a position select between two poles or options. The third format, common in the ANES

---

<sup>23</sup> Ideally I would include separate indicators for questions relating to gender and questions relating to family structures. However, almost all questions about gender touch on some other subject material, with family structures being the largest overlap.

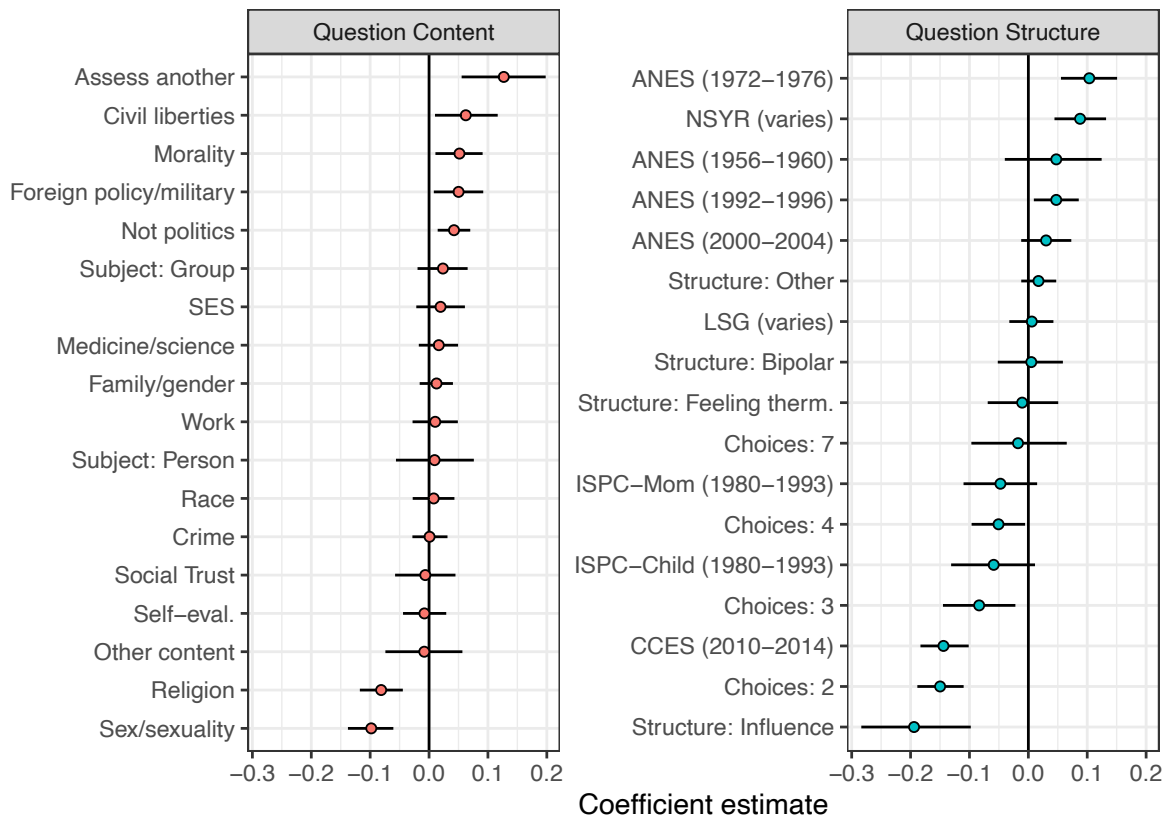
panels, asks respondents to place a respondent on a “feeling thermometer.” A fourth format asks people whether groups have “too much” influence in politics. The final category includes other question structures. Most questions in this last category evaluate the intensity of something, such as whether an amount should increase, stay the same, or decrease. While many questions with “other” structures were collapsed to reflect a “yes/no” structure, I code based on the original wording. I also include indicator variables for the number of response options (2, 3, 4, 5, 7) and indicator variables for data set.

I then regress the proportion of respondents falling into each behavioral group on these indicators. As noted above, the questions explored here do not represent a random sample of questions, so standard confidence intervals and statistical inferences are not meaningful. However, I incorporate uncertainty in two ways. First, I bootstrap coefficient estimates, randomly sampling questions with replacement 10,000 times. Second, to incorporate the uncertainty in each question’s proportion estimate, each time a parameter is sampled in the bootstrap, I draw a random value from that parameter’s posterior distribution. By bootstrapping the estimates, I can be relatively confident that whatever patterns I observe are not driven by a single question or small group of questions with unusual behavior, even if I cannot generalize to all attitudes.

In all regressions, the reference group consists of general questions about government (particularly government's role in the domestic economy) on a five-point agree/disagree scale, with the General Social Survey serving as the reference data set.

### **C.1 Vacillation**

Figure 20 presents coefficient estimates for a regression of the proportion of respondents giving vacillating response patterns on question content and structure. An unpictured regression of the proportion of respondents giving durable change response patterns on the same parameters is almost a mirror-image of the figure presented here, with the variables predicting higher vacillation predicting lower stability. As a result, I will discuss both stability and vacillation in reference to the Figure 20.



**Figure 20: Coefficient estimates of regression of proportion of respondents giving a vacillating response on question content and structure indicators.**

A general five-point agree/disagree question about domestic government and politics from the GSS has an expected proportion of people vacillating of about .58 with a standard error of .033. Two content areas demonstrate notably less vacillation than other kinds of questions: questions about religious beliefs and questions about sex and homosexuality. On the other end of the scale, questions about general morality, questions about civil liberties, and questions about medicine and science show more vacillation than other kinds of questions. When questions fall outside of the political space (“Not politics” in Figure 20), they tend to show higher levels of vacillation.

Notably, questions about race, foreign relations, work, crime, and family and gender do not appear substantially different from questions about politics once I control for question structure.

Questions that ask people to assess others – where politicians or parties fall on particular issue scales – demonstrate much higher vacillation than questions about people’s own opinions. This is consistent with previous work (Alwin 2007) and suggests a lack of knowledge about other peoples' beliefs.

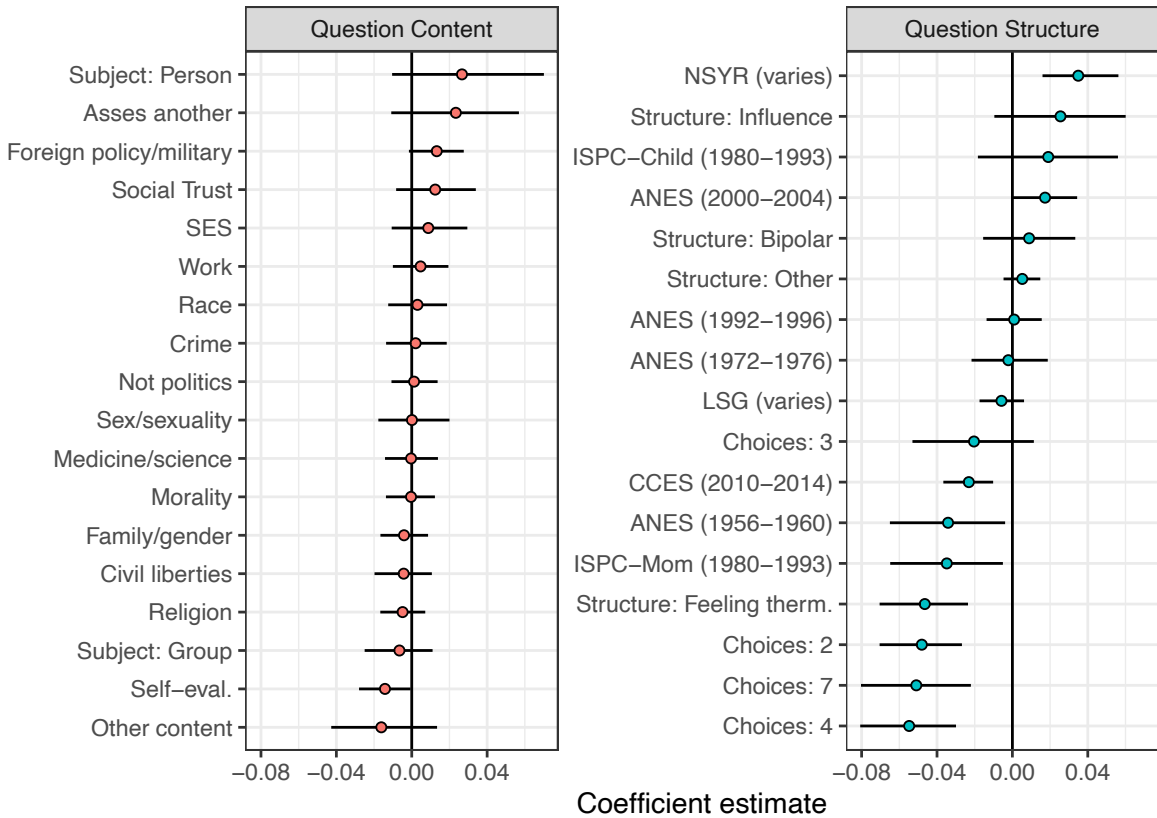
Stability decreases and vacillation increases as the number of response options increases. This helps explain the higher levels of stability for GSS and CCES questions, which have fewer response options on average, relative to the other data sets. Questions about groups’ political influence show more stability than other kinds of questions, but beyond that, question structure does not seem to matter much.

Finally, even controlling for question structure, data sets demonstrate different rates of vacillation. Notably, the 1972-1976 ANES and the NSYR demonstrate more vacillation than other data sets. The CCES demonstrates lower rates of vacillation.

## **C.2 Durable Change**

Figure 21 plots coefficient estimates of the proportion of respondents demonstrating durable change. The coefficients associated with question content and structure are much smaller on durable change than on vacillating change, reflecting its truncation in low end of the proportion range, and the parameters here do not explain

nearly as much variance as the model for vacillating change. There is almost no appreciable effect of question content on the rate of durable change.



**Figure 21: Coefficient estimates of regression of proportion of respondents giving a durable change response pattern on content and structure indicators.**

Two data sets demonstrate substantially higher levels of durable change: the NSYR and the ISPC-Child, the two panels with the lowest average age.

There is a less clear pattern on question structure for durable change than for vacillating change. Questions with five response options (the reference category) demonstrate the most durable change, followed by questions with three responses.

Questions with two, four, and seven responses demonstrate comparable, low levels of durable change.

In addition to the Nixon question discussed above, other questions with high rates of durable change include a shift in the NSYR in which many adolescents go from saying that people should wait to have sex until they are married to saying, “not necessarily,” with very few shifting the other way, and a large change of opinion around whether the presidential election in 2000 was fair.

## References

- Abramowitz, Alan I., and Kyle L. Saunders. 2006. "Exploring the Bases of Partisanship in the American Electorate: Social Identity Vs. Ideology." *Political Research Quarterly* 59 (2): 175–87.
- Achen, Christopher H. 1975. "Mass Political Attitudes and the Survey Response." *The American Political Science Review* 69 (4). [American Political Science Association, Cambridge University Press]: 1218–31.
- Alwin, Duane F. 2007. *Margins of Error: A Study of Reliability in Survey Measurement*. Hoboken, N.J.: John Wiley & Sons.
- Alwin, Duane F., and Jon A. Krosnick. 1991. "Aging, Cohorts, and the Stability of Sociopolitical Orientations over the Life Span." *American Journal of Sociology* 97 (1). University of Chicago Press: 169–95.
- Ansolahehere, 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 (2): 215–32.
- Bail, Christopher A., Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. "Exposure to Opposing Views on Social Media Can Increase Political Polarization." *Proceedings of the National Academy of Sciences* 115 (37): 9216–21.
- Baker, Wayne. 2004. *America's Crisis of Values: Reality and Perception*. Princeton, New Jersey: Princeton University Press.
- Baldassarri, Delia, and Andrew Gelman. 2008. "Partisans Without Constraint: Political Polarization and Trends in American Public Opinion." *American Journal of Sociology* 114 (2): 408–46.
- 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). The University of Chicago Press: 45–95.
- Bartels, Larry M., and Simon Jackman. 2014. "A Generational Model of Political Learning." *Electoral Studies* 33 (March): 7–18.

- Bellah, Robert N., Richard Madsen, William M. Sullivan, Ann Swidler, and Steven M. Tipton. 1985. *Habits of the Heart, with a New Preface*. Berkeley: University of California Press.
- Bryson, Bethany. "'Anything But Heavy Metal': Symbolic Exclusion and Musical Dislikes." *American Sociological Review* 61 (5): 844-899.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics (1964)." In *Ideology and Discontent*, edited by D. E. Apter, 18:206–61. New York: Free Press.
- Bonikowski, Bart, and Paul DiMaggio. 2016. "Varieties of American Popular Nationalism." *American Sociological Review* 81 (5). SAGE Publications Inc: 949–80.
- Bourdieu, Pierre. 1979. "Public Opinion Does Not Exist." In *Communication and Class Struggle*, edited by A. Mattelart and S. Siegelau, 124–30. Bagnolet, France: International General.
- — —. 1990. *The Logic of Practice*. Stanford, Calif. : Stanford University Press, 1990.
- Boutyline, Andrei, and Stephen Vaisey. 2017. "Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes." *American Journal of Sociology* 122 (5). The University of Chicago Press: 1371–1447.
- Centola, Damon. 2019. *How Behavior Spreads*. Princeton: Princeton University Press.
- Christakis, Nicholas, and James Fowler. 2010. *Connected: The Amazing Power of Social Networks and How They Shape Our Lives*. HarperCollins Publishers.
- Converse, Philip E., and Gregory B. Markus. 1979. "Plus ca Change...: the New CPS Election Study Panel." *The American Political Science Review* 73 (1). [American Political Science Association, Cambridge University Press]: 32–49.
- Cotter, David, Joan M. Hermsen, and Reeve Vanneman. 2011. "The End of the Gender Revolution? Gender Role Attitudes from 1977 to 2008." *American Journal of Sociology* 117 (1): 259–89.
- D'Andrade, Roy G. 1995. *The Development of Cognitive Anthropology*. Cambridge: Cambridge University Press.
- Danigelis, Nicholas L., Melissa Hardy, and Stephen J. Cutler. 2007. "Population Aging, Intracohort Aging, and Sociopolitical Attitudes." *American Sociological Review* 72 (5). SAGE Publications Inc: 812–30.

- DeGloma, Thomas. 2014. *Seeing the Light: The Social Logic of Personal Discovery*. Illustrated edition. Chicago ; London: University of Chicago Press.
- DellaPosta, Daniel. 2020. "Pluralistic Collapse: The 'Oil Spill' Model of Mass Opinion Polarization." *American Sociological Review* 85 (3). SAGE Publications Inc: 507–36.
- DellaPosta, Daniel, Yongren Shi, and Michael Macy. 2015. "Why Do Liberals Drink Lattes?" *American Journal of Sociology* 120 (5). The University of Chicago Press: 1473–1511.
- DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23: 263–87.
- DiMaggio, Paul, John Evans, and Bethany Bryson. 1996. "Have American's Social Attitudes Become More Polarized?" *American Journal of Sociology* 102 (3): 690–755.
- DiMaggio, Paul, Ramina Sotoudeh, Amir Goldberg, and Hana Shepherd. 2018. "Culture Out of Attitudes: Relationality, Population Heterogeneity and Attitudes Toward Science and Religion in the U.S." *Poetics* 68 (June): 31–51.
- Eaton, Asia A., Penny S. Visser, Jon A. Krosnick, and Sowmya Anand. 2009. "Social Power and Attitude Strength over the Life Course." *Personality and Social Psychology Bulletin* 35 (12): 1646–60.
- Elder, Glen H. 1974. *Children of the Great Depression: Social Change in Life Experience*. Chicago: University of Chicago Press.
- Eliasoph, Nina. 1997. "'Close to Home': The Work of Avoiding Politics." *Theory and Society* 26 (5): 605-47.
- Erikson, Robert S. 1979. "The SRC Panel Data and Mass Political Attitudes." *British Journal of Political Science* 9 (1). Cambridge University Press: 89–114.
- Ferree, Myra Marx. 2003. "Resonance and Radicalism: Feminist Framing in the Abortion Debates of the United States and Germany." *American Journal of Sociology* 109 (2): 304-44.
- Freeder, Sean, Gabriel S. Lenz, and Shad Turney. 2019. "The Importance of Knowing 'What Goes with What': Reinterpreting the Evidence on Policy Attitude Stability." *The Journal of Politics* 81 (1): 274–90.
- Frye, Margaret, and Jenny Trinitapoli. 2015. "Ideals as Anchors for Relationship Experiences." *American Sociological Review* 80 (3): 496–525.

- Ghitza, Yair, Andrew Gelman, and Jonathan Auerbach. n.d. "The Great Society, Reagan's Revolution, and Generations of Presidential Voting," 21.
- Glenn, Norval D. 1974. "Aging and Conservatism." *The Annals of the American Academy of Political and Social Science* 415. [Sage Publications, Inc., American Academy of Political and Social Science]: 176–86.
- Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116 (5). The University of Chicago Press: 1397–1436.
- Goldberg, Amir, and Sarah K. Stein. 2018. "Beyond Social Contagion: Associative Diffusion and the Emergence of Cultural Variation." *American Sociological Review* 83 (5). SAGE Publications Inc: 897–932.
- Goodman, Leo A. 1974. "Exploratory Latent Structure Analysis Using Both Identifiable and Unidentifiable Models." *Biometrika* 61 (2): 215–31.
- Green, Donald P., Bradley Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. Yale University Press.
- Gross, Neil. 2009. "A Pragmatist Theory of Social Mechanisms." *American Sociological Review* 74 (3). SAGE Publications Inc: 358–79.
- Halpern-Manners, Andrew, and John Robert Warren. 2012. "Panel Conditioning in Longitudinal Studies: Evidence from Labor Force Items in the Current Population Survey." *Demography* 49 (4): 1499–1519.
- Harding, David J. 2007. "Cultural Context, Sexual Behavior, and Romantic Relationships in Disadvantaged Neighborhoods." *American Sociological Review* 72 (3): 341–64.
- Heise, David R. 1969. "Separating Reliability and Stability in Test-Retest Correlation." *American Sociological Review* 34 (1). [American Sociological Association, Sage Publications, Inc.]: 93–101.
- Hill, Jennifer L. 2001. "Accommodating Missing Data in Mixture Models for Classification by Opinion-Changing Behavior." *Journal of Educational and Behavioral Statistics* 26 (2). American Educational Research Association: 233–68.
- Hill, Jennifer L., and Hanspeter Kriesi. 2001a. "An Extension and Test of Converse's "Black-and-White" Model of Response Stability." *The American Political Science*

- Review* 95 (2). [American Political Science Association, Cambridge University Press]: 397–413.
- — —. 2001b. “Classification by Opinion-Changing Behavior: A Mixture Model Approach.” *Political Analysis* 9 (4). [Oxford University Press, Society for Political Methodology]: 301–24.
- Hofman, Jake M., Amit Sharma, and Duncan J. Watts. 2017. “Prediction and Explanation in Social Systems.” *Science* 355 (6324). American Association for the Advancement of Science: 486–88.
- Hout, Michael, and Orestes P. Hastings. 2016. “Reliability of the Core Items in the General Social Survey: Estimates from the Three-Wave Panels, 2006–2014.” *Sociological Science* 3 (November): 971–1002.
- Howe, Lauren C., and Jon A. Krosnick. 2017. “Attitude Strength.” *Annual Review of Psychology* 68 (1): 327–51.
- Hunter, James Davison. 1992. *Culture Wars The Struggle To Control The Family, Art, Education, Law, And Politics In America*. New York: Basic Books.
- — —. 2000. *The Death of Character : Moral Education in an Age Without Good or Evil*. 1st ed. New York : Basic Books.
- Hunzaker, M. B. Fallin. 2016. “Cultural Sentiments and Schema-Consistency Bias in Information Transmission.” *American Sociological Review* 81 (6): 1223–50.
- Hunzaker, M.B. Fallin, and Lauren Valentino. 2019. “Mapping Cultural Schemas: From Theory to Method.” *American Sociological Review* 84 (5). SAGE Publications Inc: 950–81.
- Inglehart, Ronald. 1985. “Aggregate Stability and Individual-Level Flux in Mass Belief Systems: The Level of Analysis Paradox.” *American Political Science Review* 79 (1): 97–116.
- Inglehart, Ronald, and Wayne E. Baker. 2000. “Modernization, Cultural Change, and the Persistence of Traditional Values.” *American Sociological Review* 65 (1): 19.
- Iyengar, Shanto, and Sean J. Westwood. 2015. “Fear and Loathing Across Party Lines: New Evidence on Group Polarization.” *American Journal of Political Science* 59 (3): 690–707.

- Jerolmack, Colin, and Shamus Khan. 2014. "Talk Is Cheap: Ethnography and the Attitudinal Fallacy." *Sociological Methods & Research* 43 (2): 178–209.
- Joas, Hans. 1996. *The Creativity of Action*. Chicago: University of Chicago Press.
- Judd, Charles M., and Michael A. Milburn. 1980. "The Structure of Attitude Systems in the General Public: Comparisons of a Structural Equation Model." *American Sociological Review* 45 (4): 627.
- Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kiley, Kevin, and Stephen Vaisey. 2020. "Measuring Stability and Change in Personal Culture Using Panel Data." *American Sociological Review* 85 (3). SAGE Publications Inc: 477–506.
- Kinder, Donald R., and Nathan P. Kalmoe. 2017. *Neither Liberal nor Conservative: Ideological Innocence in the American Public*. Chicago: University of Chicago Press.
- Krosnick, Jon, and Duane Alwin. 1989. "Aging and Susceptibility to Attitude Change." *Journal of Personality and Social Psychology* 57 (October): 416–25.
- Lakoff, George. 1996. *Moral Politics*. Chicago: University of Chicago Press.
- Lazarsfeld, Paul F., and Neil W. Henry. 1968. *Latent Structure Analysis*. 1st edition. Houghton Mifflin Company.
- Lewis, Kevin, and Jason Kaufman. 2018. "The Conversion of Cultural Tastes into Social Network Ties." *American Journal of Sociology* 123 (6): 1684–1742.
- Lizardo, Omar. 2006. "How Cultural Tastes Shape Personal Networks." *American Sociological Review* 71 (5): 778–807.
- — —. "Improving Cultural Analysis: Considering Personal Culture in Its Declarative and Nondeclarative Modes." *American Sociological Review* 82 (1). SAGE Publications Inc: 88–115.
- Lizardo, Omar, and Michael Strand. 2010. "Skills, Toolkits, Contexts and Institutions: Clarifying the Relationship Between Different Approaches to Cognition in Cultural Sociology." *Poetics* 38 (2): 205–28.
- Lord, Frederic M., and Melvin R. Novick. 1968. *Statistical Theories of Mental Test Scores*. Reading, Mass.: Addison-Wesley Publishing Co.

- Mann, Marcus, and Cyrus Schleifer. 2020. "Love the Science, Hate the Scientist: Conservative Identity Protects Belief in Science and Undermines Trust in Scientists." *Social Forces* 99 (1): 305-32.
- Mannheim, Karl. 1952. "The Problem of Generations." In *Essays on the Sociology of Knowledge*, 276-322. New York: Oxford University Press.
- Martin, John Levi. 2002. "Power, Authority, and the Constraint of Belief Systems." *American Journal of Sociology* 107 (4): 861-904.
- — —. 2010. "Life's a Beach but You're an Ant, and Other Unwelcome News for the Sociology of Culture." *Poetics* 38 (2): 229-44.
- McCutcheon, Allan L. 1987. *Latent Class Analysis*. SAGE.
- Miles, Andrew. 2015. "The (Re)Genesis of Values: Examining the Importance of Values for Action." *American Sociological Review* 80 (4): 680-704.
- Miles, Andrew, Raphaël Charron-Chénier, and Cyrus Schleifer. 2019. "Measuring Automatic Cognition: Advancing Dual-Process Research in Sociology." *American Sociological Review* 84 (2). SAGE Publications Inc: 308-33.
- Mills, C. Wright. 1940. "Situated Actions and Vocabularies of Motive." *American Sociological Review* 5 (6): 904.
- Mohr, John W. 1998. "Measuring Meaning Structures." *Annual Review of Sociology* 24 (1): 345-70.
- Mohr, John W., Christopher A. Bail, Margaret Frye, Jennifer C. Lena, Omar Lizardo, Terence E. McDonnell, Ann Mische, Iddo Tavory, and Frederick F. Wherry. 2020. *Measuring Culture*. Columbia University Press.
- Oh, Jeong Hyun, Sara Yeatman, and Jenny Trinitapoli. 2019. "Data Collection as Disruption: Insights from a Longitudinal Study of Young Adulthood." *American Sociological Review* 84 (4): 634-63.
- Page, Benjamin I., and Robert Y. Shapiro. 1992. *The Rational Public*. Chicago: University of Chicago Press.
- Perrin, Andrew J., and Katherine McFarland. 2011. "Social Theory and Public Opinion." *Annual Review of Sociology* 37 (1): 87-107.

- Rawlings, Craig M. 2020. "Cognitive Authority and the Constraint of Attitude Change in Groups." *American Sociological Review* 85 (6). SAGE Publications Inc: 992–1021.
- Robinson, Dawn T. 2007. "Control Theories in Sociology." *Annual Review of Sociology* 33 (1): 157–74.
- Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30 (6). [American Sociological Association, Sage Publications, Inc.]: 843–61.
- Salganik, Matthew J., Ian Lundberg, Alexander T. Kindel, Caitlin E. Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M. Altschul, et al. 2020. "Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration." *Proceedings of the National Academy of Sciences* 117 (15). National Academy of Sciences: 8398–8403.
- Schnabel, Landon, and Eric Sevell. 2017. "Should Mary and Jane Be Legal? Americans' Attitudes Toward Marijuana and Same-Sex Marriage Legalization, 1988–2014." *Public Opinion Quarterly* 81 (1): 157–72.
- Scott, Marvin B., and Stanford M. Lyman. 1968. "Accounts." *American Sociological Review* 33 (1): 46.
- Sears, David O., and Carolyn L. Funk. 1999. "Evidence of the Long-Term Persistence of Adults' Political Predispositions." *The Journal of Politics* 61 (1): 1–28.
- Smith, Tom W. 1997. "Factors Relating to Misanthropy in Contemporary American Society." *Social Science Research* 26 (2): 170–96.
- Smith-Lovin, Lynn, and David R. Heise. 1988. *Analyzing Social Interaction: Advances in Affect Control Theory*. New York: Gordon and Breach Sciences Publishers.
- Steensland, Brian, Jerry Z. Park, Mark D. Regnerus, Lynn D. Robinson, W. Bradford Wilcox, and Robert D. Woodberry. 2000. "The Measure of American Religion: Toward Improving the State of the Art." *Social Forces* 79 (1). Oxford University Press: 291–318.
- Strauss, Claudia, and Naomi Quinn. 1998. *A Cognitive Theory of Cultural Meaning*. Publications of the Society for Psychological Anthropology. Cambridge: Cambridge University Press.

- Swidler, Ann. 1986. "Culture in Action: Symbols and Strategies." *American Sociological Review* 51 (2): 273.
- . 2001. *Talk of Love: How Culture Matters*. Chicago: University of Chicago Press.
- Taylor, Marylee C. 1983. "The Black-and-White Model of Attitude Stability: A Latent Class Examination of Opinion and Nonopinion in the American Public." *American Journal of Sociology* 89 (2): 373–401.
- Vaisey, S., and O. Lizardo. 2010. "Can Cultural Worldviews Influence Network Composition?" *Social Forces* 88 (4): 1595–1618.
- Vaisey, Stephen. 2009. "Motivation and Justification: A Dual-Process Model of Culture in Action." *American Journal of Sociology* 114 (6): 1675–1715.
- . 2014. "The 'Attitudinal Fallacy' Is a Fallacy: Why We Need Many Methods to Study Culture." *Sociological Methods & Research* 43 (2): 227–31.
- Vaisey, Stephen, and Omar Lizardo. 2010. "Can Cultural Worldviews Influence Network Composition?" *Social Forces* 88 (4): 1595–1618.
- . 2016. "Cultural Fragmentation or Acquired Dispositions? A New Approach to Accounting for Patterns of Cultural Change." *Socius* 2 (January).
- Visser, Penny S, and Jon A Krosnick. 2004. "Development of Attitude Strength over the Life Cycle: Surge and Decline." *Journal of Personality and Social Psychology* 87 (6): 779–95.
- Visser, Penny S., and Robert R. Mirabile. 2004. "Attitudes in the Social Context: The Impact of Social Network Composition on Individual-Level Attitude Strength." *Journal of Personality and Social Psychology*, 779–95.
- Warren, John Robert, and Andrew Halpern-Manners. 2012. "Panel Conditioning in Longitudinal Social Science Surveys." *Sociological Methods & Research* 41 (4): 491–534.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge Studies in Public Opinion and Political Psychology. Cambridge: 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). [Midwest Political Science Association, Wiley]: 579–616.

## Biography

Kevin Kiley graduated with a B.A. degree in Journalism and Political Science with highest honors and distinction from the University of North Carolina at Chapel Hill in 2010. At UNC, he was a Morehead-Cain Scholar. He earned a Master of Arts in Social Science from the University of Chicago in 2014. His Master's thesis paper at the University of Chicago won the Earl S. and Esther Johnson Prize.

He started his doctoral studies at Duke University in 2015. While at Duke he co-authored two publications: "Measuring Change and Stability in Personal Culture Using Panel Data" in the *American Sociological Review* and "A Model-Based Method for Detecting Persistent Cultural Change in Panel Data" in *Sociological Science*. At Duke, Kevin was the recipient of the James B. Duke Fellowship, the University Scholars Fellowship, the Vorsanger-Smith Scholar Award, and the Kenan Institute for Ethics Graduate Fellowship.