

A Mixed-Membership Approach to the Assessment of Political Ideology from Survey Responses

Justin H. Gross, *University of North Carolina at Chapel Hill*

Daniel Manrique-Vallier, *Duke University*

July 13, 2012

Abstract

We employ mixed-membership (or grade-of-membership) techniques—of growing popularity in medical diagnostics, psychology, genetics, and machine learning—in order to identify prototypical profiles of survey respondents based on their answers to questions aimed at uncovering their basic orientations or ideological predispositions. In contrast with factor analytic techniques and IRT approaches, we treat both manifest and latent variables as categorical. A mixed membership model may be thought of as a generalization of latent class modeling, in which individuals act as members of more than one class. This notion is well-aligned with earlier theoretical work of Zaller, Feldman, Stimson, and others, who at times envision respondents to be internally complex, answering survey questions probabilistically according to what Zaller calls varying “considerations.” Reanalyzing data in this way, we develop new insights into the sorts of constraints that may structure mass belief systems.

1 Introduction

A rich and important tradition in political science involves the analysis of patterns of political ideology.¹ Initially, the identification and characterization of such patterns had been performed in a purely ad-hoc manner, based on philosophical and qualitative considerations within a theoretical framework. This approach reached its peak with what many consider its best and most influential example, *The American Voter*, by Campbell et al. (1960). Beginning with Converse (1964), a wave of critical rethinking on the subject emerged. Along the way, various researchers have applied modern empirical tools of survey analysis and statistical inference to revisit previously held assumptions about the structure of American political attitudes and beliefs (e.g. Marcus et al. (1974); Achen (1975); Stimson (1975); Feldman (1988); Conover and Feldman (1984); Zaller (1992); Pew Research Center (2011); Ellis and Stimson (2012)) . In essence this constitutes a measurement problem: we treat data, e.g. responses to survey questions, as manifest indicators of respondents’ latent dispositions on politics and policy. A number of different analytical tools have been utilized in approaching the problem, with factor analysis being the most frequently employed and item response models gaining in popularity. The basic goal of these endeavors is typically to understand the structure of—or constraints on—beliefs, values, and attitudes at two levels: the individual and the population at large. Thus, the objects of interest include configurations of views that may be expected to coexist within a single person and the relative frequency with which these particular sets of views are held (or called upon in responding to survey items).

A basic task in the study of ideology is the construction of typologies. Simple examples of typologies are the well known distinctions between “left” and “right” or “liberal”

¹Except where noted, we used the term *ideology* somewhat generically to include any patterns of political and policy-oriented belief, values, and attitudes. Within political science, psychology, and the scholarly study of public opinion, the term is more narrowly defined as a highly constrained special case of this.

and “conservative”. Types such as “liberal” or “conservative” correspond to particular configurations of attitudes, values and beliefs that hypothetical adherents to these ideological types are supposed to exhibit. Of course nothing precludes the construction of typologies with more than two classes, as long as they offer meaningful analytical distinctions. For example, in the most recent (as of 2012) of a series of reports from the Pew Center for People and the Press, survey response data were grouped using clustering techniques, revealed two distinct groups of Republican-leaning respondents (labeled post facto as *Staunch Conservatives* and *Main Street Republicans*); three categories of people inclined toward the Democratic party (*New Coalition Democrats*, *Hard-Pressed Democrats* and *Solid Liberals*); and three so-called “Middle Groups” (*Libertarians*, *Disaffecteds* and *Post-Moderns*) (Pew Research Center, 2011). In previous reports in the Pew series, issued in 1987, 1994, 1999, and 2005, somewhat different typologies were identified.²

Even though systems that describe a few reference types are of analytic value, pretending that real-life individuals actually fully conform to one and only one type would clearly be an oversimplification. Essentially that would force us to accept that we can partition the population into a few disjoint subsets (e.g. “liberals” and “conservatives” or perhaps a larger set as in the Pew report), and that within each subset all individuals share the exact same configuration of attitudes, beliefs and values.

²Worth noting is that those identified as middle-groups were not necessarily neatly placed on a scale from liberal to conservative, as is typically done for “moderates” not clearly identifiable as liberal or conservative in contemporary terms. For example, Libertarians were those who tended to express strong views in favor of reduced government in all aspects of life, social and economic, leading to positions more associated with Republicans on economic issues and with Democrats on a number of social issues including support for political secularism. Meanwhile, Disaffecteds would not be fruitfully described in terms of a continuum between liberal and conservative as they were typified more by their cynicism regarding politics and voting generally, yet did tend to be somewhat more likely to consider themselves Republicans. While the Disaffected category had been identified in all three previous reports, Post-Moderns, on the other hand were a newly emergent type. Young and heavily Democratic in party membership, they agreed with staunch liberals on such issues as the environment, immigration, and separation of church and state, yet were more wary of New Deal and Great Society policies. Distinctions among Democrats included ethnic and economic diversity, with Solid Liberals as mostly white and liberal on most issue, while Hard-Pressed Democrats included more minorities and people with lower incomes and more conservative attitudes on social matters (Pew Research Center, 2011, pp. 20–21).

A more realistic way of describing actual individual’s ideology, while keeping the analytic advantages advantage of specifying a reduced number of reference types is through the use of a Mixed Membership (MM) framework. Mixed Membership allows us to specify partial membership into multiple reference types, and to quantify that membership. This enables us to describe ideology in terms of a few prototypical configurations, such as liberal and conservative, while allowing these configurations to coexist within a single individual. This way we could, for instance, describe an individual’s ideology as a combination of “24% conservative and 86% liberal”.

In the remainder of this introductory section, we consider the challenge of measuring individuals’ ideologies or political belief systems, typical methods for handling the task, and what a mixed-membership modeling approach may offer political scientists wishing to answer key problems in the study of ideology. We also briefly reflect upon the notion of individuals as partial adherents to more than one ideological profile and look a bit more closely at why MM models provide such a suitable empirical counterpart to this analytical framework. The data with which we illustrate an application of this measurement model are described in section 2. Next, in section 3, we present a general mixed-membership model for ideology, treating survey respondents as drawing upon partial memberships in certain latent ideological prototypes (or *extreme profiles*), in order to determine their responses. Some details regarding model fit are offered in section 4, after which we discuss our results in 5. Finally, we conclude with a brief examination of how scholars of political psychology and public opinion stand to benefit more broadly from a mixed-membership approach to their investigations, and offer a candid assessment of the limitations of the current model and possible remedies to be pursued in future work.

1.1 Understanding Ideology and the Structure of Politically-Oriented Beliefs, Values, and Attitudes

Among scholars who wish to understand how members of the public reach evaluations about parties, policies, and candidates, or simply about what they see on the evening news, there are a variety of approaches that may be taken. What is common to most of these is an assumption that people have certain dispositions, outlooks, or “basic orientations” (Feldman, 1988) upon which they rely in making such evaluations. Important debates have revolved around the question of whether ordinary people seem to apply “abstract ideological principles, sweeping ideas about how government and society should be organized” (Kinder, 1983, p. 390) in order to reach opinions on a variety of issues. For some, the notion of *ideology* itself is inherently unidimensional, a “general left-right scheme . . . organizing a wide range of fairly disparate concerns” Zaller (1992, p. 26). In this strict sense of “ideology,” as the term is employed by political scientists, the observation that most people do not rely on such a unified structuring of political views and perceptions has long been of great interest (Campbell et al., 1960; Converse, 1964). And yet, even if the members of the mass public do not think about most political issues using a left-right scheme to the extent that many political elites appear to do so, they do not approach each new object of evaluation independently, nor do they think or care enough about politics and policy to be able to do so (Page and Shapiro, 1992; Lupia and McCubbins, 1998). Thus, the basic orientations people use to make sense of the political world may be somewhat varied and not well captured by a single—or possibly even a small number—of dimensions.

We start from the widely shared assumption that such latent structure does exist in individuals and that it drives, probabilistically, their responses to survey questions Zaller (1992). It is this latent structure that we refer to here as *ideology*. As a prominent example of fundamental latent structures that do not meet the classical notion of a left-right overarching

scale, some have suggested that particular nations or cultures have a few prominent core beliefs and values that may have a high degree of popularity, but which may be of more or less importance in individuals' psyches. Converse (1964, p. 211) posits that "psychological constraints" may be at play, whereby "a few crowning postures—like premises survival of the fittest in the spirit of Social Darwinism—serve as a sort of glue to bind together many more specific attitudes and beliefs, and these postures are of prime centrality in the belief system as a whole."

Feldman (1988) and others follow up on this by examining the *core beliefs* and *core values* that may provide just this sort of psychological constraint. More recently, Ellis and Stimson (2012) make similar ideas the centerpiece of their "alternative conception of 'ideology,' ... defined by citizens' specific beliefs and values regarding what governments should and should not be doing." This "operational ideology," is distinguished from a "symbolic ideology" based on a person's self-identification or one based on more vague sentiments about " 'government' or 'government programs' broadly framed."

Our own conceptualization of ideology here follows that of Ellis and Stimson in its reliance on fundamental values and beliefs, especially as related to the appropriate role and obligations of government. Note that there is nothing inherent in such a definition that requires ideology to be unidimensional, although left-right orientation can certainly be a useful heuristic and is the focus of these authors' own discussion. To the extent that ideology may be thought of as the degree to which various of these widely embraced postures are salient for individuals in their evaluation of political objects, it will be useful to have options other than a continuous, unbounded interval for measuring this. For instance, we might expect that the vast majority of Americans will embrace values of reward for hard work, equality of opportunity, or freedom from governmental interference. Different sorts of people may find certain such considerations more compelling than others, but it would be surprising to find, for example, a group of Americans openly hostile to the notion of equal opportunity. Thus

we would like our measurement tools to be able to reflect this, by allowing us to distinguish groups of respondents not only on the items with respect to which they are diametrically opposed, but also those items that are widely embraced across groups but more consistently by one group than another.

1.2 Measuring Ideology with Survey Data

Regardless of the details of a latent structure approach to ideology (whether, for example, we treat latent and observed variables as continuously varying, ordinal, or measurable in terms of unordered levels), a key assumption is that all variation in survey responses can in fact be explained by an underlying latent variable. Survey responses will thus be conditionally independent given one's ideology (i.e. belief/value structure). The matter of how to conceptualize the latent space, whether as a multidimensional continuum or a typology with multiple possible latent classes, is largely a pragmatic question about what best reveals an otherwise invisible structure to the researcher in a manner appropriate to the questions being asked. It may be that certain renderings of this space (as perhaps, a unidimensional continuum) are rather limited in what they can tell us about how opinions are generated, but a choice among different representations may properly hinge upon what best allows lucid communication of findings. A discrete multivariate approach to the survey responses themselves makes sense, since such a treatment reflects the actual structure of Likert scale items typically found in public opinion surveys, allowing the researcher to avoid the false assumption of a continuous scale and comparable units between response levels.

The main approaches one may take in measuring ideology as a latent construct, inferred from responses to carefully selected survey questions, can be divided into heuristic or purely descriptive techniques on one hand and principled, model-based approaches on the other. Among the former are basic principal components analysis (PCA)³, multidimen-

³Social scientists regularly use the term “principal components analysis” interchangeably with (ex-

sional scaling (MDS) (Marcus et al., 1974), Q-analysis (Conover and Feldman, 1984) and correspondence analysis (CORA), a categorical analogue to PCA. The latter include Factor Analysis (FA) (Feldman, 1988) and other forms of latent structure analysis, such as latent trait analysis/item-response theory (IRT) (Treier and Hillygus, 2009) and latent class analysis (LCA) (Taylor, 1983; Feldman and Johnston, 2009), as well as MM/GoM, which may be thought of as either a sort of discrete factor analysis (Erosheva, 2002, pp. 16–20) or a sort of extension of latent class analysis.

Although there are a number of different options for handling the measurement of ideology (and beliefs, attitudes, values, etc.), the most common approach is some form of factor analysis. As the oldest latent variable measurement technique, and the most deeply ingrained in the habits of social scientists, it has the advantage of being easily related to ordinary regression techniques, and dominates the early literature on mass belief structures. Converse sets the precedent of actually equating the factors discovered or confirmed via FA with dimensions of belief structure, generating political evaluations much as Spearman’s general intelligence quotient g generates responses to IQ test items (Spearman, 1904). “Factor analysis is *the* statistical technique designed to reduce a number of correlated variables to a more limited set of *organizing dimensions*” (Converse, 1964, our emphasis). One reason that factor analysis became the dominant approach to measuring latent ideological structure was that it was the earliest to be implemented in standard statistical computing packages. The representation of individuals’ ideals and beliefs located in a low-dimensional continuous space also conformed well with evocative metaphors, borrowed from economics, by which voters were considered to have in mind ideal *amounts* of different characteristics desired of their world and of their candidates (Downs, 1957). Although the application of factor analysis in

ploratory) factor analysis—and PCA is treated as a special case of FA in statistical computer packages—but we are referring to its standard statistical meaning, a process by which orthogonal basis vectors of the reduced-dimensional space are chosen to maximize variance accounted for with each additional dimension included. The goal of PCA, as with the other descriptive approaches, is simply dimension reduction, not modeling or inference regarding the data generation process.

such situations is a deeply entrenched tradition in the study of ideology and public opinion, and is not an unreasonable approach, it is more appropriate for continuous multivariate data than discrete multivariate data typically found in survey responses.

1.3 Citizens as Partial Adherents to Distinct Ideologies

Converse (1964) set forth a research agenda, carried out in various forms over the decades since, aimed at understanding the “constraints” on patterns of belief which people may simultaneously hold. He refers to the “combinations” and “permutations” of “idea-elements” actually observed for individuals. When the constraints are severe enough, the resulting packet of beliefs to which a set of people subscribes is considered an *ideology*, in its strict sense as a psychological term of art. As Kinder (1983, p. 390) puts it, the notion of ideology upon which scholars once focused, but which guides the political evaluations of few actual citizens, consisted of “abstract ideological principles, sweeping ideas about how government and society should be organized”. In this previously dominant understanding of ideology, answers to a variety of public opinion questions could be thought to follow logically from a highly rigid, overarching outlook. Of course, if ideology operated as a purely deductive process among adherents, survey responses would be generated deterministically, and we should see certain beliefs and opinions always occurring together and others never co-occurring.⁴

Given such a narrow view of ideology, it is easy to look at actual patterns of response as evidence that people are haphazard in their thinking about politics and policy. Zaller (1992) countered this by developing a highly influential theory of how individuals formulate their responses to public opinion polls by randomly sampling from a number of privately held “considerations” relevant to the question at hand. Such a formulation helps account for

⁴Empirically, not only do certain beliefs and values fail to logically entail others, but neither may logically inconsistent sets be ruled out; the commonly held trio of preferences for more government spending, lower taxes, and a reduced deficit is but one prominent example.

a number of puzzling observations, such as the tendency for particular individuals to give different answers to the same question on different occasions.

Latent class statistical modeling (Goodman, 1974) corresponds fairly well to Zaller’s theoretical model, as each member of a distinct class responds by drawing a particular response from a distribution associated with that type. However, such an approach has the intrinsic limitation of assuming that individuals belong exclusively to just one ideological group and that each such group is homogeneous. This characterization leaves out the possibility of individuals who do not fully conform to any of the categories of a typology, but rather lie somewhere in between them.

Mixed membership models offer a conceptually attractive way of overcoming this limitation. Under Mixed Membership analysis, we still try to identify and characterize typical ideological classes. However, we regard individuals not as full members of those classes, but as *partial members*. This way, we take individuals’ responses as arising from all the distributions associated with the classes, weighted according to individually specified membership on all of them.

Mixed Membership models help formalize the idea of people as being partial adherents to different recognizable ideologies. Some—especially political leaders or “elites”—may adhere vigorously to a particular ideology and this would be reflected by full or nearly full membership on one to the exclusion of the others. Others, perhaps the vast majority of the mass public, will draw on more widely dispersed vectors of partial membership in each. This offers a nice compromise between a continuous Euclidean latent space on one hand and a categorical latent space on the other; patterns in the population and in individuals themselves are described in terms of easily understood prototypical distributions over categorical responses, and yet individuals are treated as combinations of the various prototypes, with their partial memberships allowed to vary continuously. The generating process of responses may indeed be thought of hierarchically: in encountering a survey item, the respondent first

randomly draws a ideological profile based upon his or her relative degree of membership in each extreme profile, and then randomly draws a response from that profile's response distribution.

2 Application: The American National Election Survey

The data we analyze here come from a pilot study for the 1984 American National Election Study (NES), conducted by the Center for Political Studies of the Institute of Social Research at the University of Michigan during the summer of 1983. The study's purpose was to introduce and test new survey items, including a number of questions on core values that we will be using to illustrate our mixed membership modeling approach to measuring ideology. The complete data consist of reinterviews with 314 randomly selected respondents to the earlier 1982 National Election Study. We reanalyze the same nineteen items investigated by Feldman (1988),⁵ using only the 279 complete responses.

The initial national sample, obtained for the 1982 NES and from which the individuals in the 1983 pilot study considered here were subsampled, consisted of 1418 respondents living within the primary areas of the survey's county-based sampling frame. These areas were all located within the forty-eight contiguous states (not including military bases). They include twelve major metropolitan areas, thirty-two other standard metropolitan statistical areas, and thirty counties or county-groups representing the rural subpopulation. Stratification was implemented independently within in each of the four major geographical regions of the United States, as recognized at the time: northeast, north central (loosely, the midwest), south, and west, with each represented in proportion to population. The population under

⁵Feldman served on the ANES planning committee and was apparently directly involved with formulation of these questions, intended to measure three particular core values and beliefs of Americans.

study included only United States citizens eighteen years or older on Election Day, 1982.

According to Feldman’s review of the existing literature at the time, three political attitudes dominate the American political psyche: “belief in equality of opportunity, support for economic individualism, and support for the free enterprise system,” and the nineteen items to be analyzed are all intended to measure these components of ideology. We have labeled all of these with our own variable names, which are intended to capture the spirit of the questions and distinguish similar items from one another based on the subtle wording differences. (See the Appendix for a complete list of the wording of questions.) Seven of the nineteen items are intended to measure support for or belief in what Feldman calls “equal opportunity,” including statements that attribute inequality to inherent individual differences (*natural inequality 1*, *natural inequality 2* and *equality goal misguided*), one that claims a key role for society—perhaps understood as “government” by some respondents—in ensuring equal opportunity for success (*equal opportunity–society’s responsibility*), assessments of whether inequality is a serious problem (*equal treatment* and *inequality big problem*), and one on a belief in shared governance. The items dealing with “economic individualism” are closely related to one another and differ mostly in subtle ways, as indicated in our choice of variable names: *hard work optimism*, *hard work realism*, *hard work idealism*, *ambition pessimism*, *effort pessimism*, and *individual responsibility for failure*. Finally, the “free enterprise” items (*less intervention is better*, *intervention populism*, *laissez-faire capitalism*, *regulations not a threat to freedom*, *intervention causes problems*, and *free enterprise not intrinsic feature of gov’t*) allow respondents to weigh possible tradeoffs between positive and negative consequences of governmental regulations. All items consist of statements to which the respondents may say that they “agree strongly,” “agree but not strongly,” “can’t decide,” “disagree but not strongly,” or “disagree strongly.” In order to avoid overparametrization for such a small sample, yet capture the main qualitative differences in responses, we collapse these responses into three categories: agree, disagree, or can’t decide.

3 Methods

For application to the study of political ideology we will use a technique known as the Grade of Membership model (GoM). GoM models (Woodbury et al., 1978; Manton et al., 1991; Erosheva et al., 2007) are a sub-family of Mixed Membership models (Erosheva and Fienberg, 2005). They are well-suited to obtaining low-dimensional representations of high-dimensional multivariate unordered categorical data, such as that that arises from opinion surveys. Similar to other MM techniques, GoM models represent individuals as individually weighted combinations of a small number of “ideal individuals” or “extreme profiles” and uses the data to estimate both the extreme profiles and the each subject’s membership structure. The Bayesian version of the GoM model that we introduce here and use throughout this chapter was introduced by Erosheva et al. (2007) for an application to the study of disability in elders.

3.1 Grade of Membership Models

We consider a sample of N individuals. Each individual $i = 1, 2, \dots, N$ has a corresponding J dimensional manifest variables vector, $X_i = (X_{i1}, \dots, X_{iJ})$, that collects the outcomes of interest. We assume that the components of the outcomes vector are unordered categorical variables with n_j levels each ($j = 1, \dots, J$). In our application, these outcomes are the answers to each of the J questions of the survey. For convenience, we label component’s levels using consecutive numbers, $X_{ij} \in \{1, 2, \dots, n_j\}$. We assume that there is only one response vector per individual.

GoM models assume the existence of a specific number, K , of “extreme profiles” or “pure types”. These are idealized versions of individuals that we use as reference types for specifying the response’s distribution for actual individuals. We assume that real individuals are combinations of these extremes types. To formalize this, we endow each individual with

its own *membership vector*, $g_i = (g_{i1}, \dots, g_{ik}, \dots, g_{iK})$. Each component of g_i , g_{ik} for $k = 1, \dots, K$ specifies the degree of membership of individual i in each of the K extreme profiles. We restrict membership vectors so that $g_i \in \Delta_{K-1} = \{(g_1, \dots, g_K) : g_k \geq 0, \sum_{k=1}^K g_k = 1\}$, where Δ_{K-1} is the $K - 1$ -dimensional simplex. Ideal individuals of the k th extreme profile have a membership vector whose k th component is $g_{ik} = 1$ and the rest, zeros.

The next step is to characterize extreme profiles. For any individual that is a *full member* of the k th extreme class (i.e. such that its membership vector has $g_{ik} = 1$ and $g_{ik'} = 0$ for $k' \neq k$), we assume that the response distribution of the j th entry of the manifest variables vector is a simple discrete distribution:

$$\Pr(X_{ij} = l | g_{ik} = 1) = \lambda_{jk}(l), \tag{1}$$

where $l \in \{1, 2, \dots, n_j\}$, and $\lambda_{jk} = (\lambda_{jk}(1), \dots, \lambda_{jk}(n_j)) \in \Delta_{n_j-1}$.

For generic individuals, with membership vector g_i , we characterize their component-wise response distribution as the convex combination

$$\Pr(X_{ij} = l | g_i) = \sum_{k=1}^K g_{ik} \lambda_{jk}(l).$$

Geometrically, this specification means that the individual response distributions are located within the convex hull defined by the extreme profiles.

We further assume that the item responses j are conditionally independent given membership vectors. This local independence assumption (Holland and Rosenbaum, 1986), expresses the idea that the membership vector g completely explains the dependence structure among the J binary manifest variables. By making this assumption, we can construct the conditional joint distribution of responses:

$$\Pr(X_i = x_i | g_i) = \prod_{j=1}^J \sum_{k=1}^K g_{ik} \lambda_{jk}(x_{ij}).$$

Further assuming that the individuals are randomly sampled from the population we finally get,

$$\Pr(X = x | g) = \prod_{i=1}^N \prod_{j=1}^J \sum_{k=1}^K g_{ik} \lambda_{jk}(x_{ij}). \quad (2)$$

Membership vectors are unobserved latent quantities. In order to get an unconditional expression for the joint distribution of observed responses, we assume that membership vectors are sampled from a common distribution, G_α , with support in Δ_{K-1} ; whereby

$$\Pr(X = x) = \prod_{i=1}^N \int_{\Delta_{K-1}} \prod_{j=1}^J \sum_{k=1}^K g_k \lambda_{jk}(x_{ij}) G(dg). \quad (3)$$

An interesting perspective on the GoM model results from considering the following equivalent data generation process, which generates N variates $x_i = (x_{i1}, \dots, x_{iJ})$, for $i = 1, \dots, N$, according to a GoM model with K extreme profiles (Haberman, 1995; Erosheva et al., 2007):

According to this process, we can understand the generation of individual GoM variates as arising from a two-step procedure: (1) Given a membership vector g_i , we obtain the components of the response vector one by one. (2) For each of the J components, we determine an effective extreme profile—which is allowed to vary from component to component—by sampling it with probabilities given by g_i . Next, we sample the actual response as if the

GoM data generation process

For each $i = 1, 2, \dots, N$

Sample $g_i \sim G$

For each $j = 1, 2, \dots, J$

Sample $z_{ij} \sim \mathbf{Discrete}_{1:K}(g_1, g_2, \dots, g_K)$

Sample $y_{ij} \sim \mathbf{Discrete}_{1:n_j}(\lambda_{jk}(1), \dots, \lambda_{jk}(n_j))$

individual were a full member of that extreme profile for that question. The multiple membership is reflected by the fact that the individual answers to each question are generated according to different extreme profiles.

3.2 Full Bayesian Specification

For this application we closely follow Erosheva et al. (2007). We complete the specification of the GoM in a Full Bayesian fashion by choosing the distribution of membership vectors, G_α , and a prior distribution for all parameters of our model.

For the membership vectors distribution G_α , we specify their common distribution as

$$g_i \stackrel{iid}{\sim} \mathbf{Dirichlet}(\alpha),$$

with $\alpha = (\alpha_0 \cdot \xi_1, \dots, \alpha_0 \cdot \xi_K)$, $\alpha_0 > 0$ and $\xi = (\xi_1, \dots, \xi_K) \in \Delta_{K-1}$. Parameter ξ is the expected value of distribution G_α . Using the generative process interpretation from the previous section, each component of ξ represents the expected proportion of item responses from each of the extreme profiles; thus we can informally understand it as the relative importance of each extreme profile in the population. Parameter α_0 is a concentration parameter that expresses how concentrated the probability distribution is about its expected value (as α_0 increases) or near the vertices of the simplex Δ_{K-1} (as α_0 decreases). When

$\alpha_0 = K$ and $\xi = (1/K, \dots, 1/K)$, so that $\alpha = \mathbf{1}_K$, distribution G_α becomes uniform over Δ_{K-1} .

We specify the hyper priors of G_α , as $\alpha_0 \sim \mathbf{Gamma}(1, 2)$ (in shape/rate parametrization) and $\xi \sim \mathbf{Dirichlet}(\mathbf{1}_K)$. These choices specify an a priori ignorance about the relative importance of each extreme profile in the population and a slight preference (although not so strong) for small values of α_0 , with individuals likely to be relatively pure adherents to one or another extreme profile unless the data provide suggest otherwise.

Each conditional response distribution for item j and extreme profile k , $\lambda_{jk}(\cdot)$, consists of n_j scalar parameters restricted to the simplex Δ_{n_j} . For these parameters, we chose the prior distribution

$$\lambda_{jk} = (\lambda_{jk}(1), \dots, \lambda_{jk}(n_j)) \stackrel{iid}{\sim} \mathbf{Dirichlet}(\mathbf{1}_{n_j}),$$

or a uniform distribution over Δ_{n_j-1} .

4 Fitting the Models

We have employed an MCMC algorithm to obtain samples from the posterior distribution of parameters given the data. The algorithm is an extension for multilevel variables of the sampler presented in Erosheva et al. (2007), originally developed for binary variables. This sampler is based on a data augmentation strategy, using the equivalent generative process outlined in Section 3. We have fitted models with $K = 2, 3, 4$ and 5 extreme profiles using the prior distributions described in Section 3.2.

Similar to other latent structure models, GoM models are invariant to permutations of the extreme profile labels. For this reason we have re-labeled extreme profiles according to the decreasing sequence of the posterior estimates (posterior means) of the components of ξ . This ordering makes comparisons easier.

Table 4 show posterior estimates (posterior means and standard deviations) for the

K	α_0	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5
2	0.510 (0.236)	0.971 (0.009)	0.029 (0.009)			
3	0.765 (0.233)	0.604 (0.129)	0.373 (0.129)	0.023 (0.007)		
4	0.772 (0.238)	0.591 (0.131)	0.384 (0.131)	0.016 (0.010)	0.010 (0.010)	
5	0.827 (0.229)	0.613 (0.122)	0.359 (0.121)	0.013 (0.010)	0.011 (0.011)	0.003 (0.003)

Table 1: Posterior estimates of parameters α_0 and ξ for models with $K = 2, 3, 4$ and 5 extreme profiles. Numbers between parenthesis are posterior standard deviations.

population-level distribution of membership vectors, α_0 and ξ for models with $K = 2, 3, 4, 5$ extreme profiles. In all cases, posterior estimates of α_0 are relatively small. This makes most membership vectors in the population to be dominated by a single extreme profile. However, α_0 is large enough so that the mixed membership becomes an important structural feature. Not surprisingly, given the scarcity of data ($n = 279$), posterior dispersions are rather large.

Investigating the posterior estimates of ξ we see that, for $K \geq 3$, all models feature two dominant extreme profiles, with $\xi_1 \approx 0.60$ and $\xi_2 \approx 0.36$ and $K - 2$ profiles with very small values of ξ_k . Closer inspection reveals that those two dominant extreme profiles ($k = 1$ and $k = 2$) are very similar for all models with $K \geq 3$. Plots in Figure 4 show the posterior estimates of $\lambda_{j1}(l)$ (parameters of the first extreme profile) for the model with $K = 3$ extreme profiles vs. their counterparts for models with $K = 4$ and $K = 5$ extreme profiles. We see that all points lie almost perfectly over the main diagonal. The situation is similar for the second extreme profile ($k = 2$, not shown).

Based on these observations and the qualitative inspection of the estimates, have selected the model with $K = 3$ for our inferences. This was the smallest model for which the two dominant extreme profiles appear, and any more complex model basically gives us the same information, supplemented only by additional extreme profiles with very small values of ξ_k . Interestingly, the dominant extreme profile for model with $K = 3$ is almost numerically equal to the weighted sum (by ξ_1 and ξ_2) of the estimates of two dominant extreme profiles of our selected model (and also for $K = 4$ and 5). This indicates that the two dominant extreme




Figure 1: Comparison of posterior estimates of the first extreme profile, $\lambda_{j1}(l)$, for model with $K = 3$ vs. models with $K = 4$ and $K = 5$ extreme profiles.

profiles in our chosen model are basically a split of the first dominant profile of model $K = 2$. Our attempts to perform more formal evaluations failed to produce anything illuminating. Posterior predictive counts were difficult to produce and analyze, due to the small sample size and large number of variables. We also evaluated the AICM index (Raftery et al., 2007; Erosheva et al., 2007), which selected a model with $K = 2$ extreme profiles.

5 Results and Discussion

5.1 Results

The multinomial conditional probabilities of responses given full membership, λ_{jk} , define the extreme profiles. Two of the three extreme profiles, $k = 1, 2$, account for 98% of the item responses, while profile $k = 3$ generates around two percent, and seems to be associated

with a high probability of responding “can’t decide” ($l = 2$) to the survey items (estimated as anywhere from around 10% to 83% for different items when $k = 3$).

Analyzing the posterior membership estimates of the 279 respondents, the vast majority have partial membership of less than .01 in $k = 3$. Just five members of the population would answer survey questions as primarily a member of this class (from .71 to .91 membership), and only twenty-three of the 279 have greater than 2% membership in this “neutral” prototype.

In order to better characterize the two dominant ideal types identified, let us examine each simply in terms of the probability of agreeing or disagreeing with each item when answers are based on considerations⁶ rooted in one or the other dominant profile. For convenience—and in order to connect our findings with standard writings on American ideology—we use the term *conservative* as shorthand for $k = 1$ and *liberal* for $k = 2$. For reasons that will become apparent shortly, we might more accurately refer to these as something like *Individualist/Believers in Realized American Ideals* and *Social Responsibility-Oriented/Still Waiting for American Ideals to be Realized*. Given the clumsiness of such labels, we will stick with the more common ideological identifiers, but consider them to be best understood in terms of response distributions on survey items, which we are about to examine.

5.2 Analyzing the Extreme Profiles: Americans’ Core Values vs. Core Beliefs

In considering the estimated response distributions λ_k for the “conservative” ($k = 1$) and “liberal” ($k = 2$) ideal types, one important thing to notice is the presence of certain high-valence items, enjoying the consensus one might expect of core values shared by most

⁶We intentionally use the term *considerations*, from Zaller (1992), here in order to emphasize the connection between our measurement strategy and Zaller’s theoretical framework. Just as Zaller depicts respondents drawing at random from an unobserved distribution of considerations in order to answer each question, we model such individuals as drawing an ideal type at random in proportion to their own latent membership vector, and then generating a response according to the distribution associated with the selected ideal type on the particular item.

members of a society. For such items, the distinction between liberals and conservatives is not especially stark, but to the extent that one type is more predictably supportive of a statement than the other, the differences are in the direction that would be expected. For instance, both prototypical respondents would be unlikely to agree that our inherent differences should lead us to give up on the goal of equality ($j = 2$, *equality goal misguided*), but the prototypical conservative may have a greater probability of breaking with the norm: $\hat{\lambda}_{21}(1) = .272$ ($sd. = .045$) as opposed to $\hat{\lambda}_{22}(1) = .136$ ($sd. = .063$) for the prototypical liberal. Whether responding as a liberal or conservative, an individual would very likely support the democratic ideal of governance by all sorts of people—not only the most successful—($\sim .94$ or $.86$, respectively) as well as the notion that society has a responsibility to ensure equal opportunity of success for all ($\sim .89$ or $.82$, respectively). Yet, while a commitment to the ideal of equal opportunity in personal and public life is widely embraced, so too is the recognition that people are not equally well suited to leadership positions (*natural inequality 1* ($\sim .87$ and $.76$) and *natural inequality 2* ($\sim .95$ and $.85$) among prototypical conservatives and liberals, respectively.)

In order to clarify which items are most important in defining each dominant extreme profile, we consider the quantity

$$Cohes_{jk} = \frac{\max_{l=1, \dots, n_j} \{\lambda_{jk}(l)\}}{\min_{l=1, \dots, n_j} \{\lambda_{jk}(l)\}}, \quad (4)$$

or the *cohesion* of extreme profile k with respect to item j . The cohesion scores reflect the reliability with which each extreme type responds to an item. In Zaller’s (1992) theory of survey response to opinion polling this might correspond to individuals tending to answer a question predictably, perhaps because nearly all relevant considerations lead to the same response. This may alternatively be thought to measure the cohesiveness of hypothetical adherents to each extreme profile.

Additionally, we consider the hypotheses

$$DR_j : \arg \max_{l=1, \dots, n_j} \{\lambda_{j1}(l)\} \neq \arg \max_{l=1, \dots, n_j} \{\lambda_{j2}(l)\}, \quad (5)$$

for $j = 1, \dots, J$. Hypothesis DR_j states that full adherents to the two dominant profiles have different modal responses to a given item j . Obtaining posterior estimates of the probability of DR_j enables us to draw inferences about how well different items distinguish the extreme profiles from one another. For example, a posterior probability value 0.01 for item $j = 7$, *inequality big problem*, $Pr[DR_7|Data]$, means that, given our data on 279 respondents, we find only a one-percent chance that the most likely response for the two types of pure respondent are the same—it is highly probable that the top response for *liberals* is recognition of inequality as a big problem while *conservatives* are more likely than not to deny inequality as a persistent issue.

Table 3 shows our posterior estimates (posterior means) of $Cohes_{jk}$ and our estimated posterior probability of hypotheses DR_j , for extreme profiles $k = 1$ (conservative) and $k = 2$ (liberal), and for every item in the survey ($j = 1, \dots, 19$). Analyzing the cohesion scores we see that for certain items, both dominant extreme profiles are predictable and give identical responses (e.g., *intervention causes problems*); for others they reliably give opposite responses, i.e., one type is expected to agree and the other to disagree with an item (e.g., *inequality big problem*); for still other items, prototypical adherents to one profile are highly likely to give their modal response while prototypical adherents of the other are far less predictable (e.g., *laissez-faire capitalism*).

For the first set of five items listed in Table 3, there is a greater than .50 posterior probability that pure liberals and pure conservatives will disagree on their favored responses. We can be highly confident ($> .99$) that prototypical liberals and conservatives—at least to the extent that these labels are appropriately applied to $k = 1$ and $k = 2$ —will tend to

disagree when it comes to their reactions to *hard work idealism* (“If people work hard, they almost always get what they want”), *individual responsibility for failure* (“Most people who don’t get ahead should not blame the system; they really have only themselves to blame”), and *inequality big problem* (“One of the big problems in this country is that we don’t give everyone an equal chance”). Conservatives are also likely to disagree with liberals when it comes to the item most closely associated with contemporary definitions of American liberalism and conservatism: *less intervention is better* (“The less government gets involved with business and the economy, the better off this country will be”). Similarly, we expect them to disagree on an item that perhaps best captures a quintessential American belief that hard work pays off: *hard work optimism* (“Any person who is willing to work hard has a good chance of succeeding”). For the remaining 14 items, the bulk of our posterior probability is placed on matching modal responses for liberals and conservatives, though the two ideal types may differ substantially in how predictable they are in choosing this modal response. For example, while both are more likely to agree with the statement that “There are many goods and services that would never be available to ordinary people without governmental intervention” (*intervention populism*), the mean posterior cohesion score for a prototypical liberal is 59, in contrast to around 2 for a prototypical conservative. For no fewer than eight items, we can be virtually certain that both dominant extreme profiles share a modal response.

If we look closely at the three items that most clearly distinguish our prototypical liberals from conservatives, two have been identified by Feldman (1988) as measures of the core belief in *Economic Individualism* and one as a measure of the belief in *Equal Opportunity*. All three, however, tap into beliefs about what *is* rather than what *should be*:

- *hard work idealism* ($j = 12$): If people work hard, they almost always get what they want.

- *individual responsibility for failure* ($j = 10$): Most people who don't get ahead should not blame the system; they really have only themselves to blame.
- *inequality big problem* ($j = 7$): One of the big problems in this country is that we don't give everyone an equal chance.

Indeed, much of what seems to separate the response distributions for the two dominant ideal types has to do with how well respondents view the United States as actually living up to the ideals shared by many in both camps. In order to appreciate this, a distinction should be drawn between *beliefs* and *values*.

According to Glynn et al. (1999), "Values are ideals. Beliefs represent our understanding of the way things are, but values represent our understanding of the way things should be" (p. 105). The difference between beliefs and values is not always well delineated and, in fact, some survey questions may capture aspects of both. In certain cases, what is presented as a value may imply some belief about the way things actually are, and this may affect the responses of some individuals surveyed. For example, while a majority of individuals answering from either principal extreme profile claim a belief that equal treatment leads to fewer problems (*equal treatment*), pure liberals are nearly in uniform agreement with the statement, but conservatives have around a 38% chance of disagreeing with the sentiment. Why might there be resistance among conservatives, who otherwise generally embrace the goal of equality and the notion that society has a responsibility to ensure equal opportunity, according to their responses other survey questions? Hidden within the question is an implied belief about the way things actually are: "If people were treated more equally in this country, we would have many fewer problems" ... *than we have now*, with the italicized words as implied subtext. So if one believes that inequality leads to problems, but also that people already are treated equally and perhaps that commonly advocated programs aimed at the issue (e.g., affirmative action) are misguided, one might disagree with the survey item. Thus,

one's national pride and a tendency to view the nation as having already realized the ideals of equal opportunity are considerations of prototypical conservatives that have a non-trivial probability of being primed by the choice of wording here.

Of the five items on which pure liberals and conservatives are expected to differ on their most likely responses, four deal with the locus of responsibility for individual success and failure. For all four, conservatives lean in the direction of individual responsibility for success and failure, while liberals are less convinced. Conservatives are unified in their belief that hard work has a “good chance” of yielding success, while liberals tend to disagree (albeit only at 3:2 ratio). On a similar question, phrased differently, liberals largely reject an idealistic view of hard work, with a .79 probability of disagreeing that it will “almost always” lead to satisfying results, while conservatives have a .63 of *embracing* such idealism. On one of the most divisive questions, prototypical conservatives agree three to one that individuals should blame themselves if they “don't get ahead,” while liberals find “the system” more at fault, by more than four to one! When it comes to the explicit assertion that inequality—specifically a lack of equal opportunity—remains a “big problem” in the United States, there is again a clear distinction between the two dominant types of respondent; liberals identify inequality as a big problem at over seven to one, while conservatives are nearly two to one in the opposite direction.

In short, the results of our grade-of-membership analysis reveal hidden structure in the beliefs and values of survey respondents missing from the original factor analytic results in Feldman (1988). While several included values are widely embraced across extreme profiles, prototypical liberals tend to be more cohesive in their support of those ideals typically associated with them (equality and democratic principles), while prototypical conservatives tend to be more consistent in embracing values tied to their own central narratives (rewards of hard work, individual responsibility and self-reliance, and antipathy towards government intervention). Only a few survey items serve to starkly contrast the two dominant extreme

profiles, and those that most clearly distinguish them involve beliefs about the United States in which they live rather than simply ideals about their nation as it could be.

6 Conclusion

Typologies are ubiquitous in political science, providing useful frameworks for understanding variation in ideology, beliefs and values, as well as conceptual development of many other areas (e.g., comparative political systems, conflicts, and policy areas). Typically, the manner in which such typologies are developed is ad hoc, reflecting a priori qualitative judgments about how the political world is naturally partitioned. As we demonstrate, it is also possible to use a mixed-membership approach to construct typologies from data in a principled, model-based way, with qualitative interpretation taking place only after model-based estimates have been drawn, and without imposing a possibly artificially crisp partition on the population of interest. In some cases, the resulting typology will correspond well to what we expect and in others it might offer surprises. In our illustration, we see a bit of both: on one hand, two dominant extreme profiles emerge, which correspond roughly to what might be identified as conservatives and liberals, but on the other hand, the nature of these profiles presents us with a more nuanced view of what these ideal types look like. Analyzing survey response data with reference to such prototypes, we maintain a measure of simplicity that promotes understanding, while accommodating the heterogeneity actually present in real-world populations.

Where some have previously analyzed public opinion in terms of types, they have used ad hoc clustering techniques (Pew Research Center, 2011), without justification for the particular algorithm or verification that the results are robust to other choices of clustering routine. Latent class analysis is a surprisingly underutilized technique in political science that improves upon this by assuming that data are generated from distributions associated with

j	Question	Level Profile	$\lambda_{jk}(l)$					
			$l = 1$ (Agree)		$l = 3$ (Disagree)			
			$k = 1$	$k = 2$	$k = 1$	$k = 2$	$k = 1$	$k = 2$
1	<i>Equal treatment</i>		0.61 (0.10)	0.92 (0.05)	0.37 (0.10)	0.07 (0.05)		
2	<i>Equality goal misguided</i>		0.27 (0.05)	0.14 (0.06)	0.7 (0.05)	0.83 (0.06)		
3	<i>Equal opportunity societys responsibility</i>		0.82 (0.05)	0.89 (0.05)	0.17 (0.05)	0.10 (0.05)		
4	<i>Natural inequality 1</i>		0.87 (0.04)	0.76 (0.07)	0.12 (0.04)	0.22 (0.07)		
5	<i>Natural inequality 2</i>		0.95 (0.02)	0.85 (0.06)	0.05 (0.02)	0.14 (0.05)		
6	<i>Democracy</i>		0.86 (0.04)	0.94 (0.04)	0.14 (0.04)	0.05 (0.04)		
7	<i>Inequality big problem</i>		0.30 (0.14)	0.88 (0.07)	0.69 (0.14)	0.10 (0.07)		
8	<i>Hard work optimism</i>		0.97 (0.02)	0.45 (0.17)	0.02 (0.02)	0.54 (0.17)		
9	<i>Hard work realism</i>		0.12 (0.06)	0.47 (0.09)	0.87 (0.06)	0.52 (0.09)		
10	<i>Individual responsibility for failure</i>		0.77 (0.06)	0.19 (0.12)	0.22 (0.06)	0.77 (0.12)		
11	<i>Ambition pessimism</i>		0.76 (0.05)	0.88 (0.05)	0.23 (0.05)	0.11 (0.05)		
12	<i>Hard work idealism</i>		0.64 (0.06)	0.21 (0.12)	0.35 (0.06)	0.78 (0.11)		
13	<i>Effort pessimism</i>		0.75 (0.07)	0.95 (0.03)	0.25 (0.07)	0.04 (0.03)		
14	<i>Less intervention is better</i>		0.81 (0.05)	0.42 (0.13)	0.17 (0.05)	0.55 (0.13)		
15	<i>Intervention populism</i>		0.62 (0.06)	0.83 (0.06)	0.36 (0.05)	0.11 (0.06)		
16	<i>Laissez-faire capitalism</i>		0.36 (0.05)	0.07 (0.07)	0.63 (0.05)	0.91 (0.07)		
17	<i>Regulations not a threat to freedom</i>		0.33 (0.05)	0.49 (0.08)	0.66 (0.05)	0.49 (0.08)		
18	<i>Intervention causes problems</i>		0.94 (0.04)	0.58 (0.13)	0.05 (0.03)	0.40 (0.13)		
19	<i>Free enterprise not intrinsic feature of govt</i>		0.12 (0.07)	0.41 (0.08)	0.87 (0.07)	0.58 (0.08)		

Table 2: The two dominant extreme profiles for $K = 3$: Profile $k = 1$ (60.4% of responses) vs. Profile $k = 2$ (37.3% of responses). Numbers between parenthesis are posterior standard deviations. The grouping of items is based on Feldman (1988) and the original intent of the survey questionnaire design: the first concern *Equal Opportunity*; the second, *Economic Individualism*; and the third, *Free Enterprise*. The variable names, generic in the original, are our own.

j	Question	$Cohes_{jk}$		DR_j
		$k = 1$ (cons)	$k = 2$ (lib)	
12	<i>Hard work idealism</i>	1.95	22.29	1.00
10	<i>Individual responsibility for failure</i>	4.14	16.08	1.00
7	<i>Inequality big problem</i>	11.29	14.56	0.99
14	<i>Less intervention is better</i>	5.79	19.11	0.66
8	<i>Hard work optimism</i>	330.92	8.62	0.61
17	<i>Regulations not a threat to freedom</i>	2.13	14.48	0.46
9	<i>Hard work realism</i>	29.55	7.43	0.36
18	<i>Intervention causes problems</i>	67.89	11.28	0.27
1	<i>Equal treatment</i>	1.89	7.35	0.16
19	<i>Free enterprise not intrinsic feature of govt</i>	70.77	58.83	0.14
15	<i>Intervention populism</i>	1.79	58.46	0.02
16	<i>Laissez-faire capitalism</i>	1.82	14.02	0.01
13	<i>Effort pessimism</i>	3.31	21.71	0.00
11	<i>Ambition pessimism</i>	3.53	5.60	0.00
6	<i>Democracy</i>	6.90	5.89	0.00
5	<i>Natural inequality 2</i>	49.45	7.53	0.00
4	<i>Natural inequality 1</i>	8.18	32.55	0.00
3	<i>Equal opportunity society's responsibility</i>	5.15	55.54	0.00
2	<i>Equality goal misguided</i>	2.66	17.90	0.00

Table 3: Extreme profiles for $K = 3$: Profile $k = 1$ (60.4% of responses) vs. Profile $k = 2$ (37.3% of responses). Cohesion scores represent posterior means for the odds that a prototypical adherent will give the top response to the question. * indicates that a full member of one extreme profile is clearly more predictable (i.e., displays higher cohesion) in its response than a full member of the other.

distinct latent types of respondent (Feldman and Johnston, 2009). Our mixed-membership approach may be thought of as a generalization of LCA, which combines advantages of categorical data inference available in LCA with advantages of continuity assumptions from factor analysis. In fact, one way of thinking about what we are doing here is that we have taken what would be allocated to error in the case of LCA and have integrated it with the structural component of our model. In a latent class analysis, good model fit must reduce the “error” associated with individuals who mostly answer as if they belonged to one type, but respond anomalously on certain questions. Allowing for grades of membership in these classes (reconceptualized as ideal types) takes what would otherwise be considered noise and attribute it to a person’s internal complexity.

We see great potential for mixed-membership modeling in the study of political psychology, behavior and public opinion. It offers a sort of compromise between the concreteness of classification by types and the flexibility of multidimensional continuous latent variables as in factor analysis. Informal and qualitative accounts of political behavior, for example, rely heavily on such classifications as the “likely voter” and the “independent voter,” the “alienated working class voter”; for the most part, evocative labels such as these are replaced in quantitative analysis by measures further removed from the familiar and useful prototypes, but which utilize interval-level scores to reflect diversity across the population. Assuming individuals to manifest partial membership in multiple recognizable types lets researchers use prominent response patterns as familiar reference points without reducing people to stereotypes. Furthermore, it allows us to discover new ideal types, patterns that we might not have otherwise noticed.

Mixed membership is a general idea that can be implemented and exploited in many ways. The particular technique that we employed in this application, the Grade of Membership model, has a fairly simple structure and is an appropriate tool for the basic soft-clustering that we presented here. However, in order to investigate a wider array of political science

research questions and to better use the available data, we need to develop more tools. First, we would like to investigate the relationship between individual ideology and other relevant variables, like cohort or income. This can be achieved by incorporating covariates into the model. One possible approach, introduced by Manrique-Vallier (2010), is to specify the population-level membership distribution conditional on the covariates, keeping the extreme profiles common to the whole population. Such an extension would allow to estimate the effect of the covariates into the membership composition of the individuals, enabling us to answer questions like “are younger generations more conservative than younger ones?”

Another useful direction would be to use individual estimates of membership as predictors for some dependent variable of interest, for example one’s position on a particular policy issue, or reaction to an experimental intervention. This is a common approach with factor analysis, where we use scores on estimated dimension-reduced scales as inputs into a regression-style model of substantive interest. We can also understand Mixed Membership analysis as a dimension reduction technique: in our example we reduced the 19-dimensional response vector into a two-dimensional membership vector. Thus performing a similar analysis with membership vectors instead of factor loadings as inputs would achieve a similar aim, but make for a more intuitive analysis of results. For example, we could replace statements like “for every standard deviation increase on the economic individualism factor, we expect . . .”—with the meaning of this factor obscured—with more appealing statements of the form “an extra 25% conservatism leads to . . .” The actual implementation of this idea carries some difficulties, though. While it might be tempting to perform regular regression analysis conditional on posterior point estimates of the individual membership scores, we have to make sure of reflecting the inherent posterior uncertainty of these estimates into the regression. One possible approach is to set up comprehensive hierarchical models that include the mixed membership and the regression parts. Another possible approach is to obtain samples from the posterior distribution of individual membership vectors and use multiple imputation

techniques (Rubin, 1987) to perform the combined analysis.

One limitation of GoM models, which stems from its simple local independence structure, is that given membership all answers to questions are taken to be essentially equivalent. This basically attributes all dependence between item responses to the underlying membership structure. However, researchers usually design and organize surveys so that questions belong to specific domains, such as “economic issues” or “social issues”, and therefore illuminate different (often known) aspects of ideology. One can envision a hierarchical extension in which, in addition to the mixed membership structure, questions are organized into domains and interact with the membership in different ways. The structure could be such that we take the original classification of questions as prior information with some degree of uncertainty, and learn the rest from the data.

If we or others are to effectively extend mixed-membership analysis in any of these potential directions, we would be well-advised to keep in mind the simple observation of Achen (1975), who reminds us: “The greater the distance from data to conclusions, the more opportunity for errors.” While latent variable modeling techniques grant us a principled way to measure underlying hidden concepts only indirectly revealed through survey responses, this typically comes at the expense of transparency; the connection between abstractions such as factor loadings and the observed data is often obscured in the minds of researchers and their audience. Among the various advantages of the MM/GoM approach to survey data in seeking to better understand the structure of mass attitudes, one of the most compelling is that it allows us the luxury of abstraction while preserving the close connection to data.

References

- Achen, C. (1975), “Mass political attitudes and the survey response,” *The American Political Science Review*, 69, 1218–1231.
- Campbell, A., Converse, P., Miller, W., and Donald, E. (1960), *The American Voter*, New York: John Wiley & Sons.

- Conover, P. and Feldman, S. (1984), “How people organize the political world: A schematic model,” *American Journal of Political Science*, 95–126.
- Converse, P. (1964), “The nature of belief systems in mass publics,” in *Ideology and its Discontents*, ed. Apter, D., New York: Free Press, chap. 6, pp. 207–261.
- Downs, A. (1957), *An Economic Theory of Democracy*, Harper.
- Ellis, C. and Stimson, J. A. (2012), *Ideology in America*, Cambridge University Press.
- Erosheva, E. (2002), “Grade of membership and latent structure models with application to disability survey data,” Ph.D. thesis, Department of Statistics, Carnegie Mellon University.
- Erosheva, E. and Fienberg, S. (2005), “Bayesian mixed membership models for soft clustering and classification,” *Classification—the Ubiquitous Challenge: Proceedings of the 28th Annual Conference of the Gesellschaft Für Klassifikation EV, University of Dortmund, March 9-11, 2004*.
- Erosheva, E., Fienberg, S., and Joutard, C. (2007), “Describing disability through individual-level mixture models for multivariate binary data,” *Annals of Applied Statistics*, 1, 502–537.
- Feldman, S. (1988), “Structure and Consistency in Public Opinion: the Role of Core Beliefs and Values,” *American Journal of Political Science*, 32, 416–440.
- Feldman, S. and Johnston, C. (2009), “Understanding the determinants of political ideology: The Necessity of a Multi-Dimensional Conceptualization,” in *2010 Annual Meeting of the Midwest Political Science Association*, Chicago, IL.
- Glynn, C., Herbst, S., O’Keefe, G., Shapiro, R., and Lindeman, M. (1999), *Public opinion*, Boulder, CO: Westview Press.
- Goodman, L. A. (1974), “Exploratory latent structure analysis using both identifiable and unidentifiable models,” *Biometrika*, 61, 215–231.
- Haberman, S. J. (1995), “Review: Statistical applications using fuzzy sets, by K. Manton, M. Woodbury and H. Tolley,” *Journal of the American Statistical Association*, 90, 1131–1133.
- Holland, P. W. and Rosenbaum, P. R. (1986), “Conditional association and unidimensionality in monotone latent variable models,” *Annals of Statistics*, 14, 1523–1543.
- Kinder, D. (1983), “Diversity and complexity in American public opinion,” in *Political science: The state of the discipline*, ed. Finifter, A., Washington, D.C.: American Political Science Association, pp. 389–425.
- Lupia, A. and McCubbins, M. (1998), *The democratic dilemma: Can citizens learn what they need to know?*, Chicago: Cambridge Univ Press.
- Manrique-Vallier, D. (2010), “Longitudinal Mixed Membership Models with Applications to Disability Survey Data,” Ph.D. thesis, Department of Statistics, Carnegie Mellon University.
- Manton, K. G., Stallard, E., and Woodbury, M. (1991), “A multivariate event history model based upon fuzzy states: Estimation from longitudinal surveys with informative nonre-

- sponse,” *Journal of Official Statistics*, 7, 261–293.
- Marcus, G., Tabb, D., and Sullivan, J. (1974), “The application of individual differences scaling to the measurement of political ideologies,” *American Journal of Political Science*, 18, 405–420.
- Page, B. and Shapiro, R. (1992), *The rational public: Fifty years of trends in Americans’ policy preferences*, New York: University of Chicago Press.
- Pew Research Center (2011), “Beyond Red vs. Blue: Political Typology,” Tech. rep., Pew Research Center for the People and the Press, Washington, D.C. <http://www.peoplepress.org/files/legacy-pdf/Beyond-Red-vs-Blue-The-Political-Typology.pdf>.
- Raftery, A., Newton, M., Satagopan, J., and Krivitsky, P. (2007), “Estimating the Integrated Likelihood via Posterior Simulation Using the Harmonic Mean Identity,” in *Bayesian Statistics*, eds. Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., Oxford University Press, vol. 8.
- Rubin, D. (1987), *Multiple imputation for nonresponse in surveys*, John Wiley & Sons.
- Spearman, C. (1904), ““General Intelligence,” Objectively Determined and Measured,” *The American Journal of Psychology*, 15, 201–292.
- Stimson, J. (1975), “Belief systems: Constraint, complexity, and the 1972 election,” *American Journal of Political Science*, 19, 393–417.
- Taylor, M. (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, 373–401.
- Treier, S. and Hillygus, D. (2009), “The nature of political ideology in the contemporary electorate,” *Public Opinion Quarterly*, 73, 679–703.
- Woodbury, M., Clive, J., and Garson Jr, A. (1978), “Mathematical typology: A grade of membership technique for obtaining disease definition.” *Computers in Biomedical Research*, 11, 277–98.
- Zaller, J. (1992), *The nature and origins of mass opinion*, Cambridge Univ Pr.

Appendix—Survey Items

Equal Opportunity

- *equal treatment* If people were treated more equally in this country, we would have many fewer problems. (V2169/V3120)
- *equality goal misguided* We should give up on the goal of equality, since people are so different to begin with. (V2172/V3122)
- *equal opportunity-society’s responsibility* Our society should do whatever is necessary to make sure that everyone has an equal opportunity to succeed. (V2175/V3123)

- *natural inequality 1* Some people are just better cut out than others for important positions in society. (V2178/V3121)
- *natural inequality 2* Some people are better at running things and should be allowed to do so. (V2250)
- *democracy* All kinds of people should have an equal say in running this country, not just those who are successful. (V2253, Not in wave 2)
- *inequality big problem* One of the big problems in this country is that we don't give everyone an equal chance. (V2256, V3125)

Economic Individualism

- *hard work optimism* Any person who is willing to work hard has a good chance of succeeding. (V2170)
- *hard work realism* Hard work offers little guarantee of success. (V2173)
- *individual responsibility for failure* Most people who don't get ahead should not blame the system; they really have only themselves to blame. (V2176)
- *ambition pessimism* Even if people are ambitious, they often cannot succeed. (V2251)
- *hard work idealism* If people work hard, they almost always get what they want. (V2254)
- *effort pessimism* Even if people try hard, they often cannot reach their goals. (V2257)

Free Enterprise

- *less intervention is better* The less government gets involved with business and the economy, the better off this country will be. (V2171)
- *intervention populism* There are many goods and services that would never be available to ordinary people without governmental intervention. (2174)
- *laissez-faire capitalism* There should be no government interference with business and trade. (V2177)
- *regulations not a threat to freedom* Putting government regulations on business does *not* endanger personal freedom. (V2252)
- *intervention causes problems* Government intervention leads to too much red tape and too many problems. (V2255)
- *free enterprise not intrinsic feature of gov't* Contrary to what some people think, a free enterprise system is not necessary for our form of government to survive. (V2258)