Assessing the Structure of Policy Preferences: A Hard Test of the Low-Dimensionality Hypothesis

Christopher Hare, University of California, Davis
Ben Highton, University of California, Davis
Bradford Jones, University of California, Davis

This article analyzes belief systems in a novel way, modeling relational patterns of policy disagreements using nonmetric multidimensional scaling. Because of its flexible assumptions, the approach enables us to conduct a notably harder test of the "low-dimensionality" hypothesis than is found in previous work. The results support the proposition that a basic space (consisting of a small number of interwoven issue domains) anchors the policy dimension of public opinion. Among our findings, we show that voters—especially those meeting a minimum threshold of political sophistication—neither lack meaningful attitudes nor hold distinct preferences across a wide range of issues. Rather, their policy attitudes are organized alongside relevant core values and affective evaluations in a common, low-dimensional cognitive space. A unidimensional approximation of these belief structures often exhausts the explanatory power of vote choice models.

For the real environment is altogether too big, too complex, and too fleeting for direct acquaintance. We are not equipped to deal with so much subtlety, so much variety, so many permutations and combinations. And although we have to act in that environment, we have to reconstruct it on a simpler model before we can manage with it. To traverse the world men must have maps of the world.
—Walter Lippmann (1922, 16)

V. O. Key (1966, 8) famously characterized the analysis of public opinion as a "task not unlike coming to grips with the Holy Ghost." Less colorfully, Philip Converse (1964, 206) wrote that "belief systems have never surrendered easily to empirical study or quantification." Despite the challenges, the importance of understanding, explaining, and interpreting public opinion has kept scholarly interest alive for decades.

In this article, we focus on the dimensional structure of public opinion. One important aspect of the quality of mass-elite representative linkages depends on citizens' ability to navigate the low-dimensional political choice space organized by elite actors.[1] In this view—articulated in the "basic space" theory (Hinich and Munger 1994)—democratic accountability requires low-dimensional policy space because it makes it possible for voters to organize political information, reach decisions in line with their underlying policy preferences, and ultimately make choices between competing parties, especially in a system like the United States with just two major parties (e.g., Lupton, Smallpage, and Enders 2020; Sniderman 2017).

However, an ongoing debate among scholars relates to the ability of and whether citizens structure their political preferences in a basic space. One side reasserts the "ideological innocence" thesis, contending that voters' belief systems (to the extent they can be characterized as "systems") remain at most weakly structured (e.g., Achen and Bartels 2016; Kinder and Kalmoe 2017; Lenz 2012). According to this view, voters remain generally nonideological and continue to hold

Christopher Hare (cdhare@ucdavis.edu) is an associate professor in the Department of Political Science, University of California, Davis. Ben Highton (bhighton@ucdavis.edu) is a professor in the Department of Political Science, University of California, Davis. Bradford Jones (bsjjones@ucdavis.edu) is a professor in the Department of Political Science, University of California, Davis.

Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/726961.

1. To be sure, there are other important nonpolicy models of democratic accountability (e.g., Achen and Bartels 2016; Duggan and Martinelli 2017).

Published online March 22, 2024.

The Journal of Politics, volume 86, number 2, April 2024. Published by The University of Chicago Press for the Southern Political Science Association. https://doi.org/10.1086/726961

"nonattitudes" on many policy matters. In contrast, others are more sanguine about the existence of meaningful policy attitudes in the mass electorate and argue or imply that a larger policy web structures and constrains them (e.g., Ansolabehere, Rodden, and Snyder 2008; Broockman 2016; Lewis-Beck et al. 2008).

This article makes an important contribution to scholarly understanding of policy preferences, their dimensionality, and voting behavior by using a novel approach to analyze mass preferences. In contrast to existing methods employed to estimate dimensionality, ours organizes traditional survey response data in the form of a respondent-by-respondent "disagreement matrix." We then use nonmetric multidimensional scaling (MDS) (Borg and Groenen 2005) to estimate a geometric configuration of respondents in which the distances between voters are proportional to their observed level of policy disagreement. Voters may differ along a single policy dimension, a small number of policy domains (e.g., economic, social/cultural, etc.), or on a greater number of policy dimensions.

Our use of nonmetric MDS provides a much harder test of the low-dimensionality hypothesis than found in existing empirical work because it imposes less restrictive assumptions on the data. For example, most existing scaling methods conflate respondents who hold moderate views with those whose opinions include extreme liberal and conservative preferences (Broockman 2016). Our approach addresses this observational equivalence and can separate those whose preferences are truly moderate from those who have a set of mixed, but extreme, preferences.

The article's most important findings are twofold. From a substantive perspective, the results provide strong evidence that a low-dimensional basic space structures the expanse of voters' specific policy attitudes and, in particular, their voting decisions. This is most pronounced for the politically sophisticated, but not to the exclusion of less politically sophisticated voters. The article's findings also have important methodological implications concerning the measurement of mass political attitudes. They imply a "basic space" measurement of policy preferences that reflects the mapping of multiple issues into a lower-dimensional space of domain-specific preferences. Because the economic and social/cultural domains are so tightly intertwined in contemporary American public opinion, a unidimensional measure of preferences will be a good approximation to both latent dimensions. However, these dimensions remain conceptually distinct and empirically recoverable.

THEORY AND BACKGROUND

We begin by describing how we employ a number of concepts in this article. We refer to "policy preferences" as the most specific form of preferences. In the context of survey questions, it is often the case that respondents are asked to agree or disagree with a particular policy statement, or to select their most preferred option from a range of possibilities. Questions like these are asking about policy preferences.[2]

Often, surveys include multiple policy preference questions about the same "issue." For example, a survey might include a series of questions about abortion designed to assess the conditions under which people believe the procedure should be legal. Researchers often consider issues—and the set of policy alternatives within them—as falling within particular "domains." The two most common domains are economic (sometimes called social welfare) and social (sometimes referred to as cultural or moral) (e.g., Feldman and Johnston 2014; Layman and Carsey 2002; Layman et al. 2010; Miller and Schofield 2003; Treier and Hillygus 2009). If there are close links between preferences across issue domains, then we will refer to people as having preferences that are "constrained" (Converse 1964).[3]

One body of research views the public as similar to political elites in that preferences track a single dimension conforming to the traditional left-right or liberal-conservative continuum. For our purposes, the key finding regarding political elites is that in contemporary American politics, their preferences exist on a single dimension.[4] In the mass public, a host of researchers find evidence of a similar, single underlying dimension (Bafumi and Herron 2010; Fowler et al. 2022; Jessee 2012; Shor and Rogowski 2018).[5]

In contrast, other work develops the idea that issue and policy preferences follow from more general values and beliefs that serve to constrain preferences within issue domains but not across domains (e.g., Feldman 1988; Goren 2013; Peffley

2. Respondents' answers may or may not reveal their preferences, depending on a variety of factors like the quality of the question, measurement error, etc. We return to the question of measurement below.

3. In our analysis we investigate the dimensionality of preferences. While in theory it is possible that there are multiple dimensions within issues, our primary concern is with dimensionality across issues and issue domains. For example, we ask whether there are meaningful differences in the preferences people hold on different economic issues and whether there are meaningful differences in the preferences people hold in the economic and social issue domains.

4. For instance, allowing for additional dimensions adds little if anything to our understanding of the policy positions endorsed by members of Congress (Poole and Rosenthal 2007). Likewise, other political elites—party activists and campaign contributors, in particular—exhibit similar levels of unidimensional constraint (Bonica 2018; Jennings 1992; Layman et al. 2010; Lupton et al. 2015).

5. Stimson (2012, 2015) offers a somewhat weaker view, arguing for "one-plus" dimensions of mass public opinion in which the two-party system serves to collapse political conflict on all but a handful of new issues into a unidimensional configuration. "While we can think of economic and cultural domains as clearly separable . . . they are far from completely distinct in the view of the American electorate. . . . The two dimensions are correlated, not independent" (Stimson 2012, 29–31).

and Hurwitz 1985). As Feldman and Johnston (2014, 339) write, "Citizens possess abstract beliefs which constrain specific policy preferences, but they do not necessarily see a higher-order connection between these political values." Empirically, a host of studies report results consistent with the view that public opinion is best understood and represented in multidimensional terms (e.g., Ansolabehere et al. 2008; Carmines, Ensley, and Wagner 2012; Klar 2014; Layman and Carsey 2002; Lupton, Myers, and Thornton 2015; Treier and Hillygus 2009).

The final perspective casts doubt on the existence of any meaningful dimensional structure underlying public opinion. This view finds its modern roots in work by theorists such as Walter Lippmann and Joseph Schumpeter, both of whom questioned whether voters live up to the democratic ideal of the "omnicompetent citizen" (Bennett 2006). Converse's (1964) landmark essay serves as the empirical genesis of what remains this highly influential view (Achen and Bartels 2016) that most citizens are uninterested and uninformed about politics, lack consistent views on most policy issues, and do not organize their political attitudes using higher-order schema—including the liberal-conservative ideological continuum used by political elites. Mass policy preferences are rather like "scattered croutons floating in the undifferentiated cognitive soup" (Luskin 1987, 860).

Of course, one can take a more qualified view of Converse's (and subsequent) results and conclude that citizens are capable of holding meaningful policy stances under certain circumstances, but that these attitudes exist in relative isolation. That is, their policy preferences (on at least some issues) are genuine but uncorrelated. In this case, public opinion could be modeled in a high-dimensional policy space with each issue constituting its own disjoint dimension. Representing preferences like these in any space with dimensionality lower than the number of unique issues has the potential to conflate a mixture of extreme left and right preferences with consistently moderate preferences (Broockman 2016; but see Fowler et al. 2022). Taken to its limit, this perspective could radically change how we represent policy attitudes in models of voting behavior and representation.

Theoretically, there are plausible reasons to expect that voters' policy preferences—more specifically, the systematic component of their preferences—reduce to a low-dimensional space. The first derives from a long line of work from psychology and behavioral economics on bounded rationality (Simon 1955). Taking account of the complexity of the information environment relative to the limits of cognitive resources and attention, people adopt economizing heuristic strategies to manage their decision-making processes (Kahneman 1973). While the quality and effectiveness of these heuristics are debatable, it is nonetheless clear that they are ubiquitous components of voters' information processing and preference formation processes (Lau and Redlawsk 2006). Clearly, all people (including the more sophisticated, aware, and motivated) simplify the world in order to make sense of it.[6] There is also ample evidence that core human values (including political values) are organized along a small number of dimensions, though there remains disagreement over what the dimensions represent and the precise structure of the space (Inglehart 1997; Jacoby 2014; Rokeach 1973; Schwartz 1992).

Political and intellectual elites serve to facilitate the process by grouping controversies together. Often there is no necessary or logical connection between many policy positions. However, elites "package" them together and define "what goes with what" (Converse 1964; Noel 2013; Poole 2005). To the extent that these ideological packages diffuse through the mass public, the political choice space will exist along a small number of abstract spatial dimensions.[7]

Basic space theory (Enelow and Hinich 1982; Hinich and Munger 1994; Ordeshook 1976) unifies elite and individual-centered (or top-down/bottom-up) explanations for the existence of low-dimensional structure of mass policy preferences. The theory posits the existence of two spaces: a complex "action space" where issues are represented by separate, orthogonal dimensions and a predictive "basic space" that condenses the information in the action space. Bonica (2018, 832) portrays this as a "holographic interpretation," one in which "issue preferences are understood as a higher-dimensional representation of information existing in a low-dimensional ideological space." This approach provides the theoretical foundation for ideal point models to account for behavior such as congressional roll call voting using only one or two dimensions. In public opinion, the theory implies that the basic space—not the action space—defines the structure of citizens' preferences and choices.

The central mechanism in the basic space theory is the set of mappings between the action space (i.e., the observed survey responses) and the basic space (i.e., the latent policy space). In reality, these mappings are unobservable "black

6. This perspective is empirically bolstered by multidimensional scaling results showing that human judgments across a wide set of domains can be represented in low-dimensional conceptual spaces (Gärdenfors 2000). These include emotions, linguistics, and perceptions of colors, sounds, shapes, and facial expressions. Shepard (1987) provides an overview of relevant work. See also Croft and Poole (2008) and Russell (1980).

7. Ample evidence indicates that citizens also use group identities and an "affective calculus" to learn about similarities and policy linkages between the components of partisan-ideological coalitions (Brady and Sniderman 1985; Goggin, Henderson, and Theodoridis 2020; Orr and Huber 2020; Sniderman 2017).

boxes" unique to each individual (Bonica 2018; Poole 2005). They may be derived from abstract ideological principles (as in Converse 1964), but they may also be induced by constraints on the choice set in a way that allows even unsophisticated citizens to rely on the political culture to employ ideology as a heuristic (Downs 1957; Popkin 1994). Both forms of constraint will be observationally equivalent even though their sources differ.[8]

The challenge for scaling analysis is to approximate the mapping patterns in a parsimonious manner while accounting for the most important sources of heterogeneity. First, a long line of literature rooted in Converse (1964) and expanded by Freeder, Lenz, and Turney (2019), Jacoby (1995), Stimson (1975), and Zaller (1992) argues for the central importance of political sophistication in influencing the structure and effects of policy preferences.[9] In addition to political sophistication, ideological context likely plays an important role in structuring citizens' policy attitudes by defining issue linkages. Specifically, polarized environments in which political competition is presented in stark, clearly differentiated terms between the parties promotes the level of attitudinal constraint in the mass public (Layman and Carsey 2002; Levendusky 2009; Nie, Verba, and Petrocik 1979; Pomper 1972). As noted, in such an environment—or culture—the less sophisticated, by following more easily distinguishable party cues, may also evince a low-dimensional structure but for different reasons than the more sophisticated (Popkin 1994; Smidt 2017). Whether they do, of course, is an empirical question.

METHODS

Complicating the dimensionality question is the tension in existing research regarding the role of measurement error. In light of the possibility that an individual's response to a survey question about public policy reflects at least some random measurement error due to questionnaire design, wording, or administration, it becomes important to distinguish answers to survey questions (the indicator variables) from an individual's underlying actual preferences (Achen 1975; Ansolabehere et al. 2008; Erikson 1979; Norpoth and Lodge 1985). Given the dominant view of the survey response that it represents a probabilistic draw from individuals' underlying attitudinal distributions (Zaller and Feldman 1992), basic statistical theory and the theory of errors tell us that with sufficient draws there will be asymptotical convergence on the truth. This helps to explain the widespread success of summated scales in the social sciences.

However, the proper way to aggregate stated policy preferences—answers to survey questions—into dimensional scales is not obvious and may introduce other problems. For example, Broockman (2016) explains how some favored techniques—like simple averages across items or more sophisticated unfolding or item response theory (IRT) models—might erroneously categorize people as holding moderate views, when in fact they hold a set of extreme views that "cancel out" when reducing high-level multidimensionality to low-level or unidimensional scales. Moreover, researchers usually lack firm a priori theoretical footing about the number and substantive meaning of dimensions encapsulated by a set of survey items. Questions concerning a particular policy domain may be over- or underrepresented in a survey, and it may be ambiguous as to whether a given issue should be used to construct a domain scale.[10] Further, the popularity of the two-parameter Bayesian IRT model (Martin and Quin 2002) to measure attitude structures may have inadvertently stymied research into the multidimensional structure of mass policy preferences. This is not to say that Bayesian IRT methods cannot estimate multidimensional solutions (e.g., Treier and Hillygus 2009). But, because identification is tricky and requires an increasing number of fixed constraints on the subject and/or item parameters in higher-dimensional configurations, researchers usually opt for a unidimensional result a priori. This may obscure interesting multidimensional structure in public opinion.

In this article, we use a novel approach that simultaneously addresses both methodological problems—the presence of mixed extreme preferences and the indeterminacy of the issue mappings onto one or more latent policy dimensions—in measuring the dimensionality of policy conflict in the mass public. Specifically, we use multidimensional scaling (MDS) to recover how people's constellations of policy preferences are organized. MDS procedures produce a geometric representation of observed (dis)similarities data. Observations are represented as points, and the distances between points are proportional to

8. See also Ordeshook (1976). We thank an anonymous reviewer for raising this point.

9. We use the terms "awareness," "sophistication," and "knowledge" interchangeably to refer to how much people know, understand, and are engaged with the political world. Others suggest a less consequential role for political sophistication, arguing that accounting for measurement error (Ansolabehere et al. 2008) or focusing on the structure-inducing effects of values (Goren 2013) serves to level the playing field between low- and high-sophistication citizens.

10. For instance, environmental policy touches on both economic and cultural/postmaterialist concerns—should it be included alongside abortion and LGBTQ rights questions when estimating a social/cultural policy dimension? Is immigration a social issue or a racial one? These kinds of questions cannot be answered with simple averaging. Of course, the same indeterminacy problem exists in factor analysis and is addressed through rotation.

the observed level of dissimilarity between observations.[11] That is, similar observations will be placed in closer proximity while dissimilar observations will be placed further apart in latent Euclidean space of given dimensionality.

MDS has been used in a wide array of fields (including psychology and psychometry, marketing, and physics) to visualize and measure the underlying structure of a dataset (Borg and Groenen 2005). The standard example is that of a matrix of driving distances between cities. These data naturally represent the level of dissimilarity between each pair of cities, and MDS will produce a two-dimensional "map" (where the dimensions represent geographic North-South and East-West differences) that makes clear the relative proximity between all of the cities. We can also think of this process in terms of modeling the data-generating process with use of a latent geometric space, in which differences along two underlying dimensions (i.e., latitude and longitude) give rise to the observed dissimilarities between observations.

For instance, imagine five respondents: a progressive (**P**), a conservative (**C**), a libertarian (**L**), a moderate (**M**), and an authoritarian (**A**, sometimes referred to as a communitarian). These respondents provide a mix of left (0), right (1), and centrist (0.5) responses to economic and cultural policy questions, as shown in figure 1*A*. Averaging their responses produces identical scores for the last three respondents (**L**, **M**, and **A**) in spite of their differences (e.g., Broockman 2016). However, we can also organize the data in a symmetric disimilarities matrix, in which cells represent the level of disagreement between each pair of respondents. Doing so yields figure 1*B*. MDS translates these dissimilarities into spatial (geometric) distances, as in figure 1*C*.[12] MDS perfectly reproduces the sources of observed dissimilarities between respondents in two-dimensional space, recovering the underlying sources of disagreement (i.e., the economic and cultural dimensions) even though it operates on the combined matrix in figure 1*B*. As we discuss below, it is this flexibility that allows us to more rigorously test the claim that citizens' patterns of policy disagreements are organized in low-dimensional space.

MDS methods have a long lineage in political science dating back to Weisberg and Rusk (1970), though never (to our knowledge) using policy preference data (see Poole [2008] for a review). Here, we use survey respondents' answers to a series of issue scales to generate an $n \times n$ dissimilarities matrix

A Issue scale responses

	Economic	Cultural	**Average**
Respondent P	0	0	0
Respondent C	1	1	1
Respondent L	1	0	0.5
Respondent M	0.5	0.5	0.5
Respondent A	0	1	0.5

↓

B Disagreement matrix

	P	C	L	M	A
P	0.0				
C	1.0	0.0			
L	0.5	0.5	0.0		
M	0.5	0.5	0.5	0.0	
A	0.5	0.5	1.0	0.5	0.0

↓

C MDS configuration

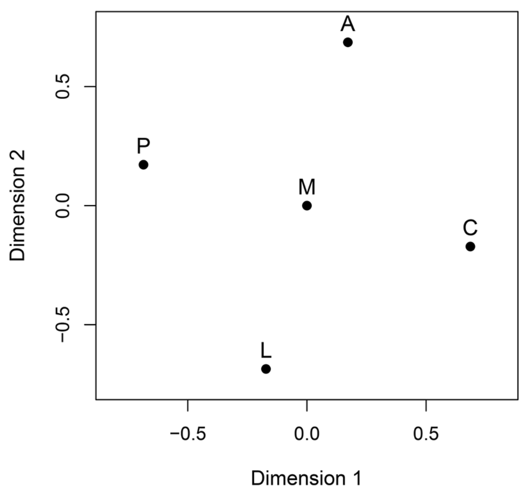

Figure 1. Illustration of the nonmetric multidimensional scaling (MDS) method

(called D), where n is the number of respondents. The cell entries (i.e., the dissimilarities) are the root sum of squared differences (i.e., Euclidian distances) between each pair of respondents over a series of q policy questions (with responses scaled to run between 0 and 1).[13]

11. The observed (dis)similarities are treated cardinally by metric MDS and ordinally by nonmetric MDS.

12. Note that in two dimensions, the diagonal distances (e.g., between **C** and **L**) will be $\sqrt{2}$ if the interior distances (e.g., between **C** and **M**) are 1. Nonmetric MDS accounts for this by relaxing the interval-level assumption between observed and reproduced data.

13. We use squared (Euclidian) distances rather than absolute (city block or "Manhattan") distances in eq. (1) both as a matter of convention and because the Euclidian metric is more appropriate when individual judgments are nonseparable (as the literature suggests is true of voters' political preferences; e.g., Lacy 2001; Stoetzer and Zittlau 2020). However, the estimated point configurations using squared and absolute distances appear to be highly correlated, especially the first and second dimensions ($r \geq 0.98$).

That is, the entries of D are calculated as:

$$\delta_{i,i'} = \left(\sum_{j=1}^{q} (p_{ij} - p_{i'j})^2 \right)^{1/2}, \quad (1)$$

where p_{ij} is the stated preference of respondent i on item j and $p_{i'j}$ is the stated preference of respondent i' on item j.

Explicitly arranging and analyzing the data in terms of the degree to which respondents disagree over policy matters solves the problem of extreme responses canceling each other out. As in the above example (fig. 1), scaling methods that average across responses will place voters with mixed left-right preferences (such as libertarians and authoritarians/communitarians) at the same, middle position. With MDS, however, the entry in D will indicate maximum disagreement, and the method will attempt to maximize the distance between these respondents.

Another strength of MDS is that it is agnostic about the specific item-dimension mappings. Its only objective is to capture the largest sources of variance in the observed dissimilarities data. The estimation process usually proceeds by estimating configurations in $1, 2, \ldots, S$ dimensions and selecting an s-dimensional solution that balances model parsimony and explanatory power.

The fit of MDS models is usually assessed in terms of Stress, a badness-of-fit measure that represents the amount of distortion between the observed dissimilarities δ and approximated distances $\hat{d}$. Hence, higher Stress values correspond to more ill-fitting point configurations. Over all pairs of respondents i and i', Stress is calculated as:

$$\sigma_1 = \left(\frac{\sum_{i<i'} (\delta_{i,i'} - \hat{d}_{i,i'})^2}{\sum_{i<i'} (\delta_{i,i'}^2)} \right)^{1/2}, \quad (2)$$

which avoids scale dependency by dividing the raw Stress by the sum of the squared observed dissimilarities (i.e., normalizing the raw Stress value in the numerator).[14] Though Stress values lack a natural metric for interpretation, a crude rule of thumb is that values of 0.2 constitute "poor" solutions, 0.1 "fair" solutions, and 0.05 "good" solutions. However, it is usually more appropriate to assess Stress values by comparing them to baseline values obtained from random permutations of the data (Borg and Groenen 2005; Mair, Borg, and Rusch 2016).[15]

Organizing issue scale responses in this particular format offers a final advantage. By explicitly representing these data in terms of inter-respondent differences, we better capture the relational nature of political competition. That is, preferences may be viewed as the product of a process in which policy linkages are learned by experience: ongoing evaluations of the political world in terms of similarity and dissimilarity between actors and objects (Sniderman 2017). By fitting an MDS model to citizens' patterns of disagreements, we more directly test the characterization of the basic space in capturing "structural descriptions of political differences between groups of positions" (Hinich and Munger 1994, 128).

Overall, then, compared to existing methods employed to assess, scale, and measure policy preferences, nonmetric MDS provides a notably stronger test of the low-dimensionality hypothesis. It avoids the common assumption that responses to policy items are cumulatively structured, which is problematic when locating respondents who hold a mix of extreme left- and right-wing preferences (Broockman 2016). Because the nonmetric MDS procedure operates directly on the respondent-by-respondent disagreement matrix, it is able to avoid making such cumulative assertions on the data. This allows us to obtain a more faithful representation of the complex webs of policy disagreements in the mass public—one that implicitly accommodates higher dimensions if they are manifest in the response patterns.

DATA

We apply nonmetric MDS to analyze policy disagreement patterns between respondents in the 2012 American National Election Study (ANES).[16] The 2012 ANES includes an especially large number (80) of issue questions that tap into a diverse array of policy controversies in American politics. We later leverage this feature to compare the predictive power of nonmetric MDS results based on random subsets of issues in models of vote choice.

In the section "Model Fit by Dimensionality," we use the Stress values to assess overall and subgroup-specific model fit in spaces of different dimensionalities. In addition to estimating nonmetric MDS on the complete ANES dataset, we also estimate separate models for those with low, medium, and high levels of political sophistication.[17] In order to provide a

14. Because values of the quotient are usually quite small, eq. (2) takes its square root. This is the first (and most widely used) version of the Stress calculation, hence the subscript in σ_1 (Borg and Groenen 2005, 42).

15. Another advantage of MDS in this context is that it provides individual-level fit statistics, as we can easily decompose eq. (2) to calculate each respondent's contribution to overall Stress. As detailed in app. L (apps. A–L are available online), these Stress per point (SPP) values provide a measure of each respondent's fit in the estimated configuration (i.e., how well their configuration of policy disagreements can be geometrically represented in an s-dimensional abstract space).

16. We also analyze data from the Broockman (2016) data and the 2018 Cooperative Congressional Election Study (CCES). The results, presented in app. E, closely mirror those from the 2012 ANES.

17. Political sophistication is measured using a summated scale of 24 correct/incorrect political knowledge items (Cronbach's $\alpha = 0.9$) to

comparable null or baseline measure of fit, we perform nonmetric MDS on randomly permuted versions of the data.[18] This provides a null hypothesis (H_0) against which to assess the fit (Stress) of our model (Mair et al. 2016). Finally, we use a series of Monte Carlo simulations to generate disagreement matrices under known conditions (varying the level of error, the dimensionality, and the salience weights attached to each dimension) and estimate nonmetric MDS configurations for each simulated matrix in 1 through 10 dimensions.[19] This allows us to compare the patterns of Stress values recovered from the survey data with those obtained from a series of hypothetical data-generating processes with known properties (Spence and Graef 1974).

Following Kruskal and Wish (1978, 48), we should not rely entirely on statistical (Stress-based) criteria to evaluate the dimensionality of voters' policy preferences. When deciding on the appropriate number of latent dimensions, we should also consider external validity by examining relevant, out-of-sample political attitudes and choices.[20] A meaningful cognitive dimension is not just one that satisfies some statistical criterion of improvement in model fit but is also interwoven with other politically relevant attitudes.

To evaluate the dimensionality of voters' policy preferences on a substantive basis, in the section "Mappings between the MDS Dimensions and Political Variables," we use a technique known as property vector fitting (PVF) to examine the relationship between the recovered dimension(s) and external political variables (such as core value measures and feeling thermometer ratings). PVF uses multiple linear regression to project each external measure onto the respondent ideal point configurations and use the mapping parameters (i.e., the β regression coefficients) to assess the correspondence between each external attitude or disposition and dimension from nonmetric MDS (Jones and Koehly 1993). Although MDS does not estimate item-specific parameters (for instance, the discrimination parameter in IRT models), PVF provides an alternate strategy to substantively interpret the dimensions of the latent space.

Finally, in the section "Simulations of Ideal Points as Predictors of Vote Choice," we perform two sets of simulations designed to test the power of a reduced set of basic space dimensions from nonmetric MDS as predictors of vote choice.[21] In the first set of simulations, we vary the number of dimensions estimated from one to eight, using all 80 issue questions in the 2012 ANES. In the second set of simulations, we vary both the number of items used to construct the disagreement matrix and the number of dimensions estimated. Our goal is to assess how well composite information from the set of policy items—that is, estimates of respondent positions on the latent dimensions—explains presidential vote choice. If voters are indeed operating in a reduced policy space in their decision-making processes, then only a small number of dimensions should be needed to explain citizens' voting choices. Moreover, if the relevant policy information contained in the responses is largely redundant, we should be able to estimate those dimensions using only a reduced set of items.

RESULTS

Model fit by dimensionality

We begin by examining model fit (using Stress values as our criterion) for each of the estimated nonmetric MDS configurations in the 2012 ANES. These values are displayed in the scree plots in figure 2. The quantities of greatest interest to us are the location of the "elbow" (the point at which the inclusion of additional dimension provides only marginal improvement in fit) and the difference in Stress values between the observed and randomly permuted data.

What is immediately apparent in figure 2 is that the MDS configurations of the observed policy disagreement data (both overall and subset by level of political knowledge) have smaller Stress values than those estimated from the randomly permuted versions of the same data. In other words, MDS provides a significant improvement in fit over the null hypothesis that policy disagreements among citizens exhibit no latent structure. We also find a steady gradation in Stress values among people based on political knowledge, with low-knowledge respondents possessing the highest Stress configurations followed by medium- and then high-knowledge respondents.[22] In all three cases, the elbow occurs at either two or three dimensions, though it is sharpest among respondents

bin respondents into three equality sized (nonweighted) groups. Additional details are provided in app. B.

18. Specifically, we randomly shuffle values within each of the original columns (i.e., survey items) and use these to construct the permuted disagreement matrices (see, e.g., Jacoby and Armstrong 2014).

19. We also simulated datasets in which responses are generated under the assumptions of normal (correlated and uncorrelated) and bimodal utility functions across the q policy items. Details and results for these simulations are provided in app. F.

20. "Out-of-sample" here simply refers to variables (such as feeling thermometer ratings and voting preferences) that were not used to construct the disagreement matrix for nonmetric MDS.

21. In each of 100 trials, we draw with replacement a random set of 500 respondents with probability of selection equal to the survey weights.

22. In app. G, we show results from our Monte Carlo experiments that generate issue responses under different assumptions about voter utility functions and degree of correlation between issues. The MDS fit statistics for the three observed datasets are most in line with data simulated under the assumption that respondents possess utility functions that are normally distributed and correlated across issues.

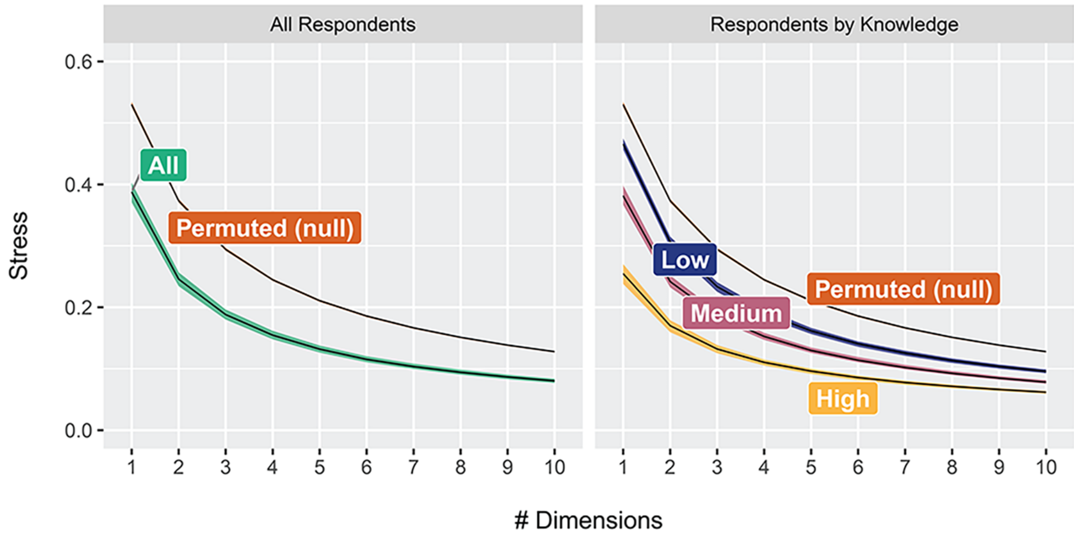

Figure 2. Scree plots from nonmetric multidimensional scaling (MDS) of respondent policy disagreement matrices from the 2012 American National Election Study (ANES). Shaded regions show 95% bootstrapped confidence intervals.

with high levels of political sophistication and becomes flatter as we move to medium- and low-sophistication respondents. Hence, the observed patterns suggest that we can characterize voters' policy disagreements as being organized in low- (≤ 3) dimensional space.[23] High-sophistication respondents' policy differences exhibit the most structure, but even those with low sophistication show a significant difference from the permuted data (i.e., null structure).[24]

We also take a reverse engineering approach to test the dimensionality of the data by using a simulation strategy introduced by Spence and Graef (1974). As discussed above, this involves randomly generating respondent-by-respondent disagreement matrices under a variety of known conditions, estimating nonmetric MDS Stress values for those configurations, and comparing the simulated Stress patterns to those from the observed data. We vary three conditions in our simulations: (1) the number of latent generating dimensions (between 1 and 10), (2) the salience or importance weights attached to each dimension (equal, reciprocal, or exponential weights), and (3) the level of error (between zero and ∞).[25]

Figure 3 highlights the two patterns of simulated Stress values that most closely track (based on sum of squared differences) the observed Stress values from the 2012 ANES. In both cases, the latent dimensionality of the simulated space is high (six or eight dimensions). But critically, these simulations also feature exponential weights (in the six-dimensional configuration) and reciprocal weights (in the eight-dimensional configuration). Both of these conditions greatly diminish the importance of higher dimensions: for instance, the fourth dimension is four (20) times less influential than the first dimension in the reciprocal (exponential) conditions. If we instead limit ourselves to the equal weight condition (i.e., conceive of the dimensions as representing equally important sources of attitudinal constraint), we find that the closest simulated patterns are generated from the two-dimension ($\varepsilon = 0.25$) and three-dimension ($\varepsilon = 0.1225$) conditions. Hence, the results provide further evidence that the low-dimensional interpretation best characterizes public opinion. At most, higher dimensions explain only a small fraction of voters' patterns of policy disagreements.

Mappings between the MDS dimensions and political variables

The recovered low-dimensional point configurations appear to be good fits to the policy disagreement data, but to what extent do the estimated dimensions reflect contemporary political

23. Appendix L disaggregates Stress values by respondent, which reveals very few instances where individual fit is meaningfully improved by the inclusion of four or more dimensions.

24. Appendix J presents scree plots obtained when permuting the data within each knowledge group. Particularly among respondents to the 2012 ANES, permutation produces virtually identical Stress values across groups.

25. Additional details are provided in app. G. The reciprocal and exponential weight conditions reduce the importance of higher dimensions by contracting respondents along those dimensions. Under reciprocal weights, distances between respondents on the sth dimension are multiplied by $1/s$. Under exponential weights, distances are multiplied by $\exp(-s+1)$. Error levels (ε) represent the variance of the Gaussian shocks added to the distances. Infinite ε corresponds to a random permutation of the data (see Spence and Graef 1974).

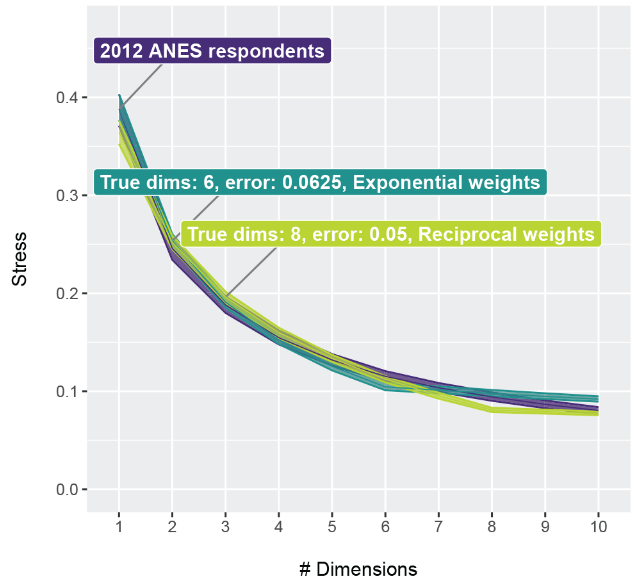

Figure 3. Selected scree plots from nonmetric multidimensional scaling (MDS) of the main 2012 American National Election Study (ANES) respondent policy disagreement matrix and the two closest simulated disagreement matrices. Shaded regions show 95% bootstrapped confidence intervals.

divisions? That is, do voters organize their policy differences in the same space as they make other political choices? Figure 4 suggests that they do. It shows the estimated two-dimensional MDS configurations of respondents divided by level of political sophistication and colored by presidential vote choice. Across levels of political sophistication, we see that the first dimension separates Obama and Romney voters most acutely and that these divisions are sharpest among politically sophisticated respondents. The second dimension appears to offer minimal explanatory power in terms of vote choice for any knowledge group.

To further assess the meaning of the estimated MDS configurations of voter positions, we next use property vector fitting (PVF) to analyze the mapping of several external political variables in the MDS space. These items were not used to construct the policy disagreement matrices, but PVF allows us to project them into the same space via linear regression. In this case, we regress each external variable on respondents' MDS coordinates and use the estimated coefficients to measure the partial associations between the dimensions and external items.[26]

Specifically, we use PVF to project two types of responses—those to core value questions and feeling thermometer ratings—into the 10-dimensional MDS space.[27] Figure 5 shows the distribution of ordinary least squares (OLS) coefficients for each MDS dimension across the core values and feeling thermometers. If the recovered scores reflect underlying and meaningful attitudes (rather than simply serving as useful summary measures of overall policy dispositions), they should exhibit a strong association with value structures and affective evaluations. Indeed, core values—as universal and central components of human behavior—serve as a crucial "bottom-up" constraining mechanism in the basic space theory and provide a further test of the veracity of our theoretical framework.[28]

The results in figure 5 demonstrate that the forces structuring citizens' policy disagreements in low-dimensional space share a close connection to major value and affective differences. Respondents' first-dimension MDS scores are consistently predictive of both their core value dispositions (especially egalitarianism and moral traditionalism) and feeling thermometer ratings (including ideological and nonideological groups). The second MDS dimension is generally less relevant than the first but nonetheless contributes additional explanatory power across the three core values (especially authoritarianism) and is most important in explaining affect toward (ir)religious groups such as Christian fundamentalists and atheists.[29] The role of the third dimension is muted, though it appears to be weakly related to authoritarianism and feelings about illegal immigrants.

Critically though, in assessing the dimensionality of MDS space, figure 5 uncovers no linkages between the external variables

26. Additional details and results are provided in app. H.

27. We repeat these regressions over 100 random samples of 500 respondents from the 2012 ANES. In each trial, we store the absolute values of the regression coefficients and set $\beta = 0$ if $p > .05$.

28. We create standard indices of moral traditionalism (four items, $\alpha = 0.70$), egalitarianism (six items, Cronbach's $\alpha = 0.78$), and authoritarianism (four items, $\alpha = 0.61$) for the core value scales.

29. We subset the analysis by political knowledge groups and provide results for additional feeling thermometers in app. H.

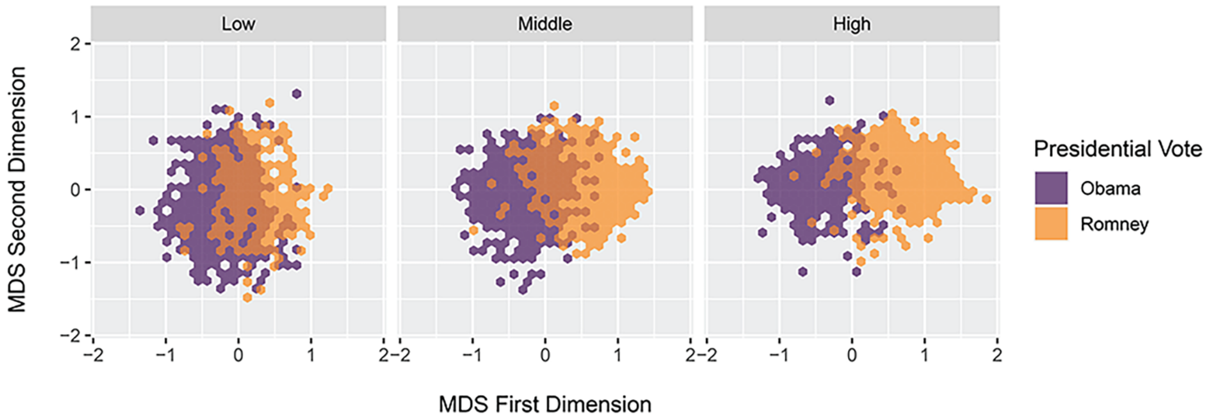

Figure 4. Respondent coordinates from nonmetric multidimensional scaling (MDS) of policy disagreement matrices; 2012 American National Election Study (ANES) respondents, by political knowledge.

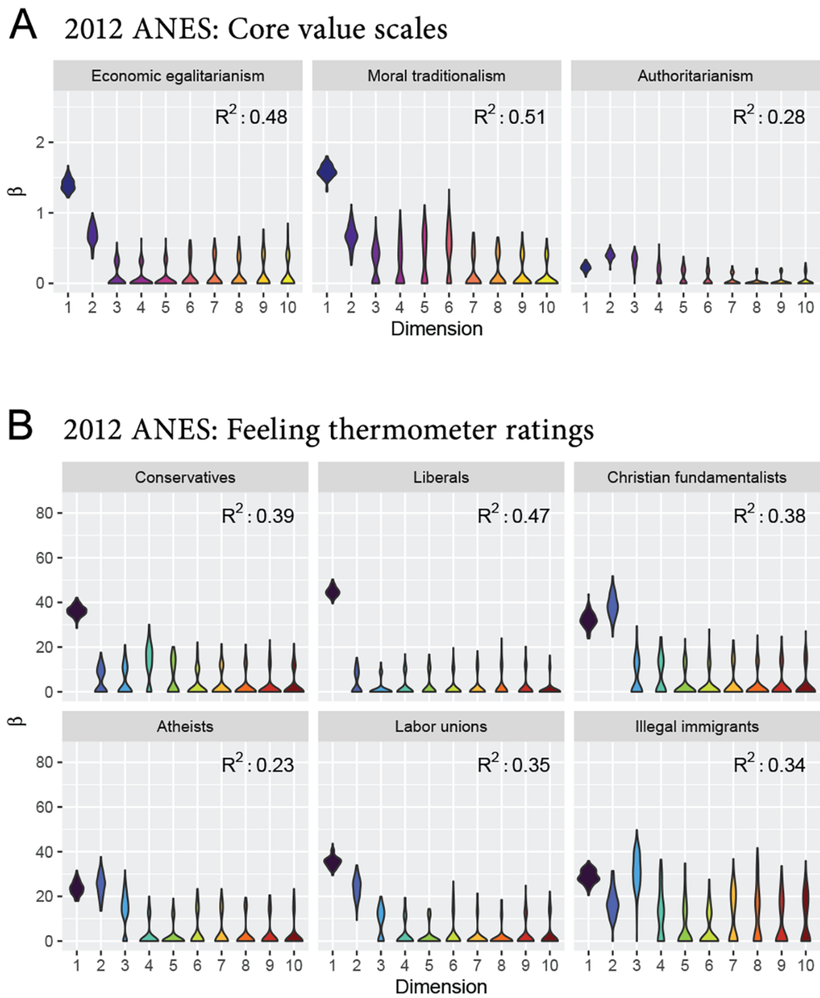

Figure 5. Property vector mappings of external core values (*A*) and group feeling thermometer ratings (*B*) onto 10-dimensional nonmetric multidimensional scaling (MDS) configurations. The β coefficients estimated with OLS, using 100 bootstrap samples of 500 respondents. Mean R^2 values shown.

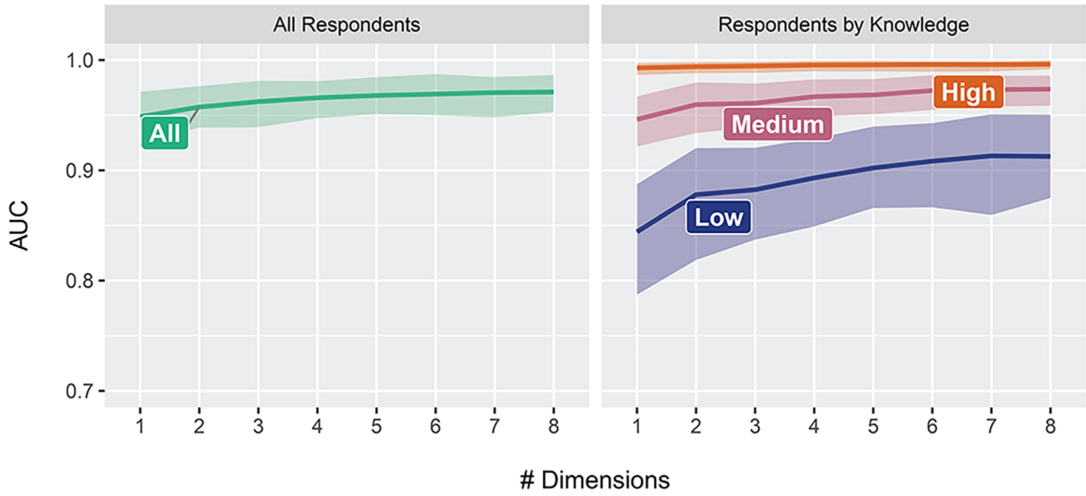

Figure 6. Presidential vote model fit by dimensionality of nonmetric multidimensional scaling (MDS) configuration, 2012 American National Election Study (ANES). Shaded regions show 95% bootstrapped confidence intervals. Models include controls for race, gender, church attendance, age, education, and income. AUC (area under the receiver-operator characteristic curve) value for controls-only model is 0.8.

and any of the higher dimensions from MDS. Given that the core value and affective measures represent a diverse array of politically relevant dispositions, these results provide additional evidence that our MDS-based approach is recovering a small number of meaningful attitudinal dimensions structuring citizens' patterns of policy disagreements.

Simulations of ideal points as predictors of vote choice

As a final test of our theoretical expectations, we perform two sets of Monte Carlo simulations that extract policy dimensions from nonmetric MDS and use them to predict presidential vote choice. Above, we found that a policy space of three or fewer dimensions proved sufficient to model the basic structure of most citizens' policy disagreements. In these tests, we further consider the possibility that higher dimensions are nonetheless meaningful drivers of individual voting behavior.

In the first set of simulations, we use nonmetric MDS to estimate one to eight dimensions on policy disagreement matrices constructed from bootstrapped samples of 500 respondents. We then include the MDS scores (alongside standard demographic controls; though omitting party identification) to model presidential vote choice with a probit regression model, using the AUC value as our measure of model fit.[30] Here, we are interested in locating the point at which additional dimensions provide only marginal improvement in voting behavior model fit. The second test uses this same design but also varies the number of policy items. That is, we randomly select a subset of issues from the full set of issue questions to estimate the MDS dimensions. If all or most issues tap into the same basic space, then only a subset should be needed to explain the policy-based component of vote choice.

The results are displayed in figures 6 and 7. Both show that a single policy dimension is nearly sufficient to explain the voting behavior of low- and medium-sophistication respondents and entirely sufficient to explain the voting behavior of high-sophistication respondents. In no case do more than two or three dimensions meaningfully improve model fit.[31] Of course, this does not mean that the inclusion of one or two dimensions will perfectly predict vote choice—policy considerations are, after all, tangential for many voters. Rather, our emphasis is on the finding that additional dimensions that extend the basic space are not capturing variance meaningfully related to voting behavior.

Figure 7 provides further support for this argument. Here, we find that using only 10 random issue questions is usually adequate to estimate behaviorally consequential dimensions.[32] Regardless of the number of issue responses included, two-dimensional MDS configurations modestly outperform one-dimensional configurations, while the improvement provided

30. AUC (area under the receiver-operator characteristic curve) is a widely used fit statistic in discrete choice models. Unlike accuracy, it measures the model's discriminatory power in predicting outcomes (in this case, Obama and Romney voters) across a range of classification thresholds other than 0.5. Additional details are provided in app. A.

31. We do find a gradual increase in AUC values offered by estimating higher-dimensional MDS configuration among low-sophistication respondents, though this improvement is marginal and not statistically significant at conventional levels.

32. Appendix J disaggregates the results by political sophistication. While adding issues provides the most dramatic improvement in fit among medium sophistication respondents, the increase is negligible.

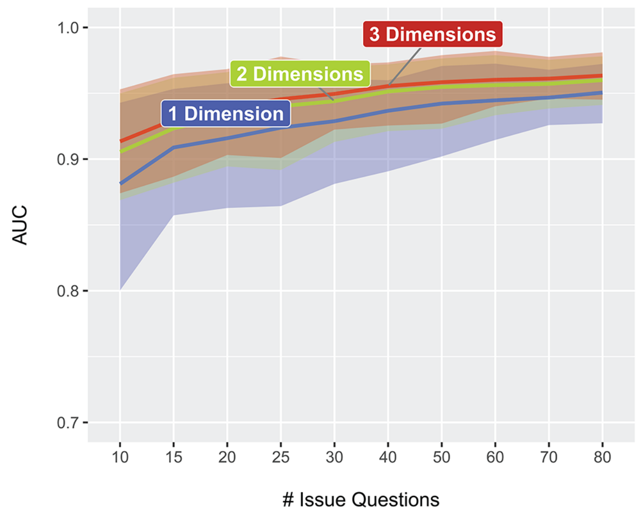

Figure 7. Presidential vote model fit by number of issues included in nonmetric multidimensional scaling (MDS), 2012 American National Election Study (ANES). Shaded regions show 95% bootstrapped confidence intervals. Models include controls for race, gender, church attendance, age, education, and income. AUC (area under the receiver-operator characteristic curve) value for controls-only model is 0.8.

by a third dimension is marginal. There is a faint indication that the difference in fit between the one- and two-dimensional models increases alongside the number of issue questions selected (presumably, as the samples capture a more diverse set of policy attitudes). But, as in figure 6, we find scant evidence for exotic high-dimensional policy influences on vote choice.[33]

DISCUSSION

> When voters can expertly judge every detail of every stand taken and relate it directly to their own views of the good society, they are interested only in issues, not in philosophies. . . . Uncertainty alters this whole situation by removing the voters' perfect competence at relating every party decision to their own ideologies. Voters do not know in great detail what the decisions of the government are, and they cannot find out except at a significant cost. . . . Under these conditions, many a voter finds party ideologies useful because they remove the necessity of his relating every issue to his own philosophy. (Downs 1957, 98)

Downs's seminal work stands like a monolith in studies of policy voting, but one of its key insights about the role of issues in voters' political universes has nonetheless gone underappreciated. His spatial model makes clear we should not expect citizens to vote on the basis of exhaustive, issue-by-issue comparisons. Indeed, as confirmed by countless empirical studies of public opinion since, we should not even expect voters to hold coherent attitudes on most specific issues. Rather, the central mechanism of the Downsian model is an underlying dimension that binds parties to abstract societal goals. This is the space that most citizens use to understand politics and evaluate parties and candidates. The basic space theory expands and formalizes this understanding of how citizens engage in policy voting. Specifically, it emphasizes the role of political elites in creating mappings between the two spaces. These linkages are defined and reinforced by political competition in a path-dependent process: "the product of the accumulated experience citizens have with the political system" (Hinich and Munger 1994, 165). The basic space theory is consistent with work from behavioral economics and cognitive psychology emphasizing our proclivity to store and process information using heuristics and spatial organization.

We assess the basic space theory's applicability to the structure of public opinion with a novel approach based on nonmetric multidimensional scaling (MDS). Our methodology flexibly considers the way citizens organize their disagreements

33. These experiments are designed to test the validity of the low-dimensional hypothesis in explaining voting behavior and get a practical sense of the number of survey items necessary to capture the behaviorally consequential basic dimension(s). However, they also highlight the challenge of substantively interpreting the dimensions recovered by MDS. Besides applying the property vector fitting (PVF) technique (as above), researchers might consider using MDS in tandem with scaling methods that directly estimate item-level parameters (such as item response theory and factor analytic models) to supplement interpretation of the latent space.

about policy matters, thereby providing a more rigorous test of the low-dimensionality hypothesis. In doing so, it also addresses the important critique of the now-standard dimensional analyses of policy preferences raised by Broockman (2016) concerning voters with conflicting, extreme views.

Two influential approaches to the dimensionality question—one that rejects the existence of widespread structure in mass policy preferences and the other that conceives of such preferences as positions in high-dimensional space—provide unrealistic models of how voters navigate the political world. To the extent that citizens' policy disagreements are structured, it is along a small number of intertwined basic dimensions. The inclusion of more than three dimensions provides only trivial improvements in fit.[34] Even though the MDS procedure is agnostic about the source(s) of disagreements across issues, we nonetheless recover a clear left-right dimension.

For most of the electorate, a basic policy space is behaviorally predictive and sufficient to capture the systematic component of their policy attitudes. This is especially true among citizens with moderate and high levels of political sophistication. To clarify, we do not mean to suggest that voters' policy attitudes exhibit perfect or even high levels of constraint but, rather, we simply fail to find patterns that would seriously invalidate the use of aggregate policy scales. From a methodological perspective, this implies that scales and indices are useful measurement tools for the vast majority of citizens. Scaling-based estimates—whether from factor analysis, IRT models, or simple summated scales—are successful because they combine insights from measurement theory as well as the basic space theory to derive estimates of voter positions in a simplified cognitive space. This basic space consists of a small number of correlated dimensions and it—rather than the action space defined by the full universe of policy conflicts—structures citizens' issue attitudes, drives the policy component of voting behavior, and reflects affective and value dispositions.

At the same time, to the extent that an underlying structure is evident in the policy preferences of ordinary citizens, the method for uncovering it used here (and those used elsewhere) does not reveal the explanation for that structure. Especially in a period when party differences have grown considerably, similarities across people with different levels of political sophistication may be evident. The less attentive to politics may nevertheless receive sufficiently clear party signals that result in a lower dimensional structure than would otherwise be evident (Popkin 1994; Smidt 2017). This may have had the effect of producing greater similarity with the more informed whose structure may be based less on party cues and more on actual ideology.

Looking forward, there are some promising avenues for future research. For one, in this study we have not analyzed the Stress per point (SPP) values as a measure of individual voters' fits to the MDS model. Future work could make greater use of SPP values, including validating their measurement properties and leveraging them to examine heterogeneity in mass attitude structures.[35] In this vein, individual differences scaling could also be used to identify variation in the relative importance of the recovered dimensions among subgroups of voters.[36] Additionally, Broockman's (2016) argument that standard scaling procedures may conflate moderation with mixed-extreme views does not only apply to the mass public. Given the use of these models for locating a host of actors—including members of Congress, judges, the president, and political parties and candidates—a key extension of the work we have done here would be to employ MDS to analyze these and other political stimuli. Because MDS allows attitude structures to be observed without imposing any assumptions that make it more or less likely, it is an extremely valuable method for assessing whether a host of political institutions and actors are organized in the way we have come to think they are.

ACKNOWLEDGMENTS

This article was previously titled "The Space of American Public Opinion: Ideological Dimensionality in Models of Political Behavior." An earlier version of this article was presented at Measurement Meets Politics and Policy: Honoring the Scholarship and Mentorship of William G. Jacoby and Saundra K. Schneider, April 3, 2019, Chicago. Thanks to Ryan Bakker, Raymond Brandt, Adam Enders, Bob Erikson, Bill Jacoby, Paul Sniderman, Walt Stone, members of the University of California Davis American Politics Reading Group, and especially Bob Lupton for helpful conversations and suggestions over previous drafts.

REFERENCES

Achen, Christopher H. 1975. "Mass Political Attitudes and the Survey Response." *American Political Science Review* 69 (4): 1218–31.

Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton, NJ: Princeton University Press.

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder Jr. 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.

34. The results also suggest that individuals of the type Broockman (2016) theorizes are generally uncommon.

35. We provide a preliminary analysis of the SPP values in app. L.

36. Individual differences scaling is also referred to as weighted or three-way multidimensional scaling (Borg and Groenen 2005, 449–94).

Bafumi, Joseph, and Michael C. Herron. 2010. "Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress." *American Political Science Review* 104 (3): 519–42.

Bennett, Stephen Earl. 2006. "Democratic Competence, Before Converse and After." *Critical Review* 18 (1–3): 105–41.

Bonica, Adam. 2018. "Inferring Roll-Call Scores from Campaign Contributions Using Supervised Machine Learning." *American Journal of Political Science* 62 (4): 830–48.

Borg, Ingwer, and Patrick J. F. Groenen. 2005. *Modern Multidimensional Scaling: Theory and Applications*. 2nd ed. New York: Springer.

Brady, Henry E., and Paul M. Sniderman. 1985. "Attitude Attribution: A Group Basis for Political Reasoning." *American Political Science Review* 79 (4): 1061–78.

Broockman, David E. 2016. "Approaches to Studying Policy Representation." *Legislative Studies Quarterly* 41 (1): 181–215.

Carmines, Edward G., Michael J. Ensley, and Michael W. Wagner. 2012. "Who Fits the Left-Right Divide? Partisan Polarization in the American Electorate." *American Behavioral Scientist* 56 (12): 1631–53.

Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." In David E. Apter, ed., *Ideology and Discontent*. New York: Free Press, 206–61.

Croft, William, and Keith T. Poole. 2008. "Inferring Universals from Grammatical Variation: Multidimensional Scaling for Typological Analysis." *Theoretical Linguistics* 34 (1): 1–37.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper & Row.

Duggan, John, and César Martinelli. 2017. "The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness." *Journal of Economic Literature* 55 (3): 916–84.

Enelow, James M., and Melvin J. Hinich. 1982. "Ideology, Issues, and the Spatial Theory of Elections." *American Political Science Review* 76 (3): 493–501.

Erikson, Robert S. 1979. "The SRC Panel Data and Mass Political Attitudes." *British Journal of Political Science* 9 (1): 89–114.

Feldman, Stanley. 1988. "Structure and Consistency in Public Opinion: The Role of Core Beliefs and Values." *American Journal of Political Science* 32 (2): 416–40.

Feldman, Stanley, and Christopher Johnston. 2014. "Understanding the Determinants of Political Ideology: Implications of Structural Complexity." *Political Psychology* 35 (3): 337–58.

Fowler, Anthony, Seth J. Hill, Jeffrey B. Lewis, Chris Tausanovitch, Lynn Vavreck, and Christopher Warshaw. 2022. "Moderates." *American Political Science Review* 117 (2): 643–60.

Freeder, Sean, Gabriel S. Lenz, and Shad Turney. 2019. "The Importance of Knowing 'What Goes with What': Reinterpreting the Evidence on Policy Attitude Stability." *Journal of Politics* 81 (1): 274–90.

Gärdenfors, Peter. 2000. *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.

Goggin, Stephen N., John A. Henderson, and Alexander G. Theodoridis. 2020. "What Goes with Red and Blue? Mapping Partisan and Ideological Associations in the Minds of Voters." *Political Behavior* 42 (4): 985–1013.

Goren, Paul. 2013. *On Voter Competence*. New York: Oxford University Press.

Hinich, Melvin J., and Michael C. Munger. 1994. *Ideology and the Theory of Political Choice*. Ann Arbor: University of Michigan Press.

Inglehart, Ronald. 1997. *Modernization and Post-modernization: Cultural, Economic, and Political Change in 43 Societies*. Princeton, NJ: Princeton University Press.

Jacoby, William G. 1995. "The Structure of Ideological Thinking in the American Electorate." *American Journal of Political Science* 39 (2): 314–35.

Jacoby, William G. 2014. "Is There a Culture War? Conflicting Value Structures in American Public Opinion." *American Political Science Review* 108 (4): 754–71.

Jacoby, William G., and David A. Armstrong. 2014. "Bootstrap Confidence Regions for Multidimensional Scaling Solutions." *American Journal of Political Science* 58 (1): 264–78.

Jennings, M. Kent. 1992. "Ideological Thinking among Mass Publics and Political Elites." *Public Opinion Quarterly* 56 (4): 419–41.

Jessee, Stephen A. 2012. *Ideology and Spatial Voting in American Elections*. Cambridge: Cambridge University Press.

Jones, Lawrence E., and Laura M. Koehly. 1993. "Multidimensional Scaling." In Gideon Keren and Charles Lewis, eds., *A Handbook for Data Analysis in the Behavioral Sciences: Methodological Issues*. Hillsdale, NJ: Erlbaum, 95–163.

Kahneman, Daniel. 1973. *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.

Key, V. O., Jr. 1966. *The Responsible Electorate*. New York: Vintage.

Kinder, Donald R., and Nathan P. Kalmoe. 2017. *Neither Liberal nor Conservative: Ideological Innocence in the American Public*. Chicago: University of Chicago Press.

Klar, Samara. 2014. "A Multidimensional Study of Ideological Preferences and Priorities among the American Public." *Public Opinion Quarterly* 78 (S1): 344–59.

Kruskal, Joseph B., and Myron Wish. 1978. *Multidimensional Scaling*. Newbury Park, CA: Sage.

Lacy, Dean. 2001. "A Theory of Nonseparable Preferences in Survey Responses." *American Journal of Political Science* 45 (2): 239–58.

Lau, Richard R., and David P. Redlawsk. 2006. *How Voters Decide: Information Processing during Election Campaigns*. New York: Cambridge University Press.

Layman, Geoffrey C., and Thomas M. Carsey. 2002. "Party Polarization and 'Conflict Extension' in the American Electorate." *American Journal of Political Science* 46 (4): 786–802.

Layman, Geoffrey C., Thomas M. Carsey, John C. Green, Richard Herrera, and Rosalyn Cooperman. 2010. "Activists and Conflict Extension in American Party Politics." *American Political Science Review* 104 (2): 324–46.

Lenz, Gabriel S. 2012. *Follow the Leader? How Voters Respond to Politicians' Policies and Performance*. Chicago: University of Chicago Press.

Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. Chicago: University of Chicago Press.

Lewis-Beck, Michael S., William G. Jacoby, Helmut Norpoth, and Herbert F. Weisberg. 2008. *The American Voter Revisited*. Ann Arbor: University of Michigan Press.

Lippmann, Walter. 1922. *Public Opinion*. New York: Macmillan.

Lupton, Robert N., William M. Myers, and Judd R. Thornton. 2015. "Political Sophistication and the Dimensionality of Elite and Mass Attitudes, 1980–2004." *Journal of Politics* 77 (2): 368–80.

Lupton, Robert N., Steven M. Smallpage, and Adam M. Enders. 2020. "Values and Political Predispositions in the Age of Polarization: Examining the Relationship between Partisanship and Ideology in the United States." *British Journal of Political Science* 50 (1): 241–60.

Luskin, Robert C. 1987. "Measuring Political Sophistication." *American Journal of Political Science* 31 (4): 856–99.

Mair, Patrick, Ingwer Borg, and Thomas Rusch. 2016. "Goodness-of-Fit Assessment in Multidimensional Scaling and Unfolding." *Multivariate Behavioral Research* 51 (6): 772–89.

Martin, Andrew D., and Kevin M. Quinn. 2002. "Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953–1999." *Political Analysis* 10 (2): 134–53.

Miller, Gary, and Norman Schofield. 2003. "Activists and Partisan Realignment in the United States." *American Political Science Review* 97 (2): 245–60.

Nie, Norman H., Sidney Verba, and John R. Petrocik. 1979. *The Changing American Voter*. Cambridge, MA: Harvard University Press.

Noel, Hans. 2013. *Political Ideologies and Political Parties in America*. Cambridge: Cambridge University Press.

Norpoth, Helmut, and Milton Lodge. 1985. "The Difference between Attitudes and Nonattitudes in the Mass Public: Just Measurements." *American Journal of Political Science* 29 (2): 291–307.

Ordeshook, Peter C. 1976. "The Spatial Theory of Elections: A Review and a Critique." In Ian Budge, Ivor Crewe and Dennis J. Farlie. eds., *Party Identification and Beyond: Representations of Voting and Party Competition*. New York: Wiley, 285–313.

Orr, Lilla V., and Gregory A. Huber. 2020. "The Policy Basis of Measured Partisan Animosity in the United States." *American Journal of Political Science* 64 (3): 569–86.

Peffley, Mark A., and Jon Hurwitz. 1985. "A Hierarchical Model of Attitude Constraint." *American Journal of Political Science* 29 (4): 871–90.

Pomper, Gerald M. 1972. "From Confusion to Clarity: Issues and American Voters, 1956–1968." *American Political Science Review* 66 (2): 415–28.

Poole, Keith T. 2005. *Spatial Models of Parliamentary Voting*. New York: Cambridge University Press.

Poole, Keith T. 2008. "The Evolving Influence of Psychometrics in Political Science." In Janet M. Box-Steffensmeier, Henry E. Brady, and David Collier, eds., *The Oxford Handbook of Political Methodology*. New York: Oxford University Press, 199–213.

Poole, Keith T., and Howard Rosenthal. 2007. *Ideology and Congress*. New Brunswick, NJ: Transaction.

Popkin, Samuel L. 1994. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. 2nd ed. Chicago: University of Chicago Press.

Rokeach, Milton. 1973. *The Nature of Human Values*. New York: Free Press.

Russell, James A. 1980. "A Circumplex Model of Affect." *Journal of Personality and Social Psychology* 39 (6): 1161–78.

Schwartz, Shalom H. 1992. "Universals in the Content and Structure of Values: Theoretical Advances and Empirical Tests in 20 Countries." *Advances in Experimental Social Psychology* 25:1–65.

Shepard, Roger N. 1987. "Toward a Universal Law of Generalization for Psychological Science." *Science* 237 (4820): 1317–23.

Shor, Boris, and Jon C. Rogowski. 2018. "Ideology and the US Congressional Vote." *Political Science Research and Methods* 6 (2): 323–41.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69 (1): 99–118.

Smidt, Corwin D. 2017. "Polarization and the Decline of the American Floating Voter." *American Journal of Political Science* 61 (2): 365–81.

Sniderman, Paul M. 2017. *The Democratic Faith: Essays on Democratic Citizenship*. New Haven, CT: Yale University Press.

Spence, Ian, and Jed Graef. 1974. "The Determination of the Underlying Dimensionality of an Empirically Obtained Matrix of Proximities." *Multivariate Behavioral Research* 9 (3): 331–41.

Stimson, James A. 1975. "Belief Systems: Constraint, Complexity, and the 1972 Election." *American Journal of Political Science* 19 (3): 393–417.

Stimson, James A. 2012. "On the Meaning and Measurement of Mood." *Daedalus* 141 (4): 23–34.

Stimson, James A. 2015. *Tides of Consent: How Public Opinion Shapes American Politics*. 2nd ed. Cambridge: Cambridge University Press.

Stoetzer, Lukas F., and Stephen Zittlau. 2020. "Vote Choice and the Nonseparability of Economic and Social Issues." *Public Opinion Quarterly* 84 (1): 158–70.

Treier, Shawn, and D. Sunshine Hillygus. 2009. "The Nature of Political Ideology in the Contemporary Electorate." *Public Opinion Quarterly* 73 (4): 679–703.

Weisberg, Herbert F., and Jerrold G. Rusk. 1970. "Dimensions of Candidate Evaluation." *American Political Science Review* 64 (4): 1167–85.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. 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): 579–616.