

What Ordered Optimal Classification reveals about ideological structure, cleavages, and polarization in the American mass public

Christopher Hare¹ · Tzu-Ping Liu¹ · Robert N. Lupton²

Received: 19 December 2017 / Accepted: 28 March 2018 / Published online: 31 March 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract This paper develops an extension of Poole’s (Polit Anal 8(3):211–237, 2000) Optimal Classification (OC) scaling procedure to the analysis of polytomous or ordered choice data. This type of data is regularly encountered in public opinion and expert surveys, legislative and judicial bodies where abstention is relevant, and measures of policy that are coded along ordinal scales. OC is nonparametric and requires only minimal assumptions about voters’ utility functions and the error term. As such, Ordered Optimal Classification (OOC) provides a flexible modeling strategy to estimate latent ideological spaces from ordinal choice data. OOC is also easily estimated in multidimensional space without identifying restrictions. After describing the OOC procedure, we perform a series of Monte Carlo experiments and apply the method to analyze survey data from the 2015 Cooperative Congressional Election Study. We then conclude with a discussion of how scholars can utilize OOC in future work involving multidimensional spatial models of choice.

Keywords Ideal point estimation · Ideology · Public opinion · Optimal Classification

An earlier version of this paper was presented at the Keith T. Poole Career Retrospective Conference, May 2017, Athens, GA. Thanks to Keith Poole for his patient and invaluable guidance throughout this project and to Bob Erickson, Howard Rosenthal, and an anonymous review for their insightful comments and suggestions for improvement. Thanks also to the University of Georgia, the Albert B. Saye Fund, and the School of Public and International Affairs for their generous financial support for conference travel. An R package and accompanying code to estimate Ordered Optimal Classification is available at <http://github.com/tzuliou/ooc>.

✉ Christopher Hare
cdhare@ucdavis.edu

Tzu-Ping Liu
tpliu@ucdavis.edu

Robert N. Lupton
robert.lupton@uconn.edu

¹ Department of Political Science, University of California, Davis, Davis, USA

² Department of Political Science, University of Connecticut, Storrs, USA

1 Introduction

In recent years, scholars have employed a range of methodological tools to examine an ideologically sorted, if not polarized, American electorate (Ansolabehere et al. 2008; Treier and Hillygus 2009; Jessee 2009; Lupton et al. 2015). Their findings add nuance to the orthodox view that most voters lack ideological “constraint” and hold poorly structured political attitudes (Converse 1964). In line with the spatial theory of voting (Enelow and Hinich 1984; Hinich and Munger 1994, 1997), these results indicate that mass political preferences—particularly by politically sophisticated voters—exhibit structure in low-dimensional space, often with separate dimensions for major issue domains (e.g., economic, social/cultural, foreign policy issues). This research raises not only substantive questions about the structure and distribution of mass political attitudes, but methodological ones as well; in particular, what is the most appropriate method to recover the underlying ideological structure of public opinion?

In this paper, we demonstrate how Keith Poole’s (2000, 2005) Optimal Classification (OC) procedure can be used to analyze public opinion data and recover spatial “maps” of mass ideology that parallel the iconic NOMINATE-based plots of Poole and Rosenthal (Poole and Rosenthal 1997) for Congress. OC is a nonparametric unfolding procedure that was developed to study voting behavior in legislative bodies, but has a number of statistical properties that make it a useful tool for the study of public opinion. OC, unlike other scaling methods such as factor analysis and item response theory, is nonparametric and makes only weak assumptions about the functional forms of the respondents’ preferences and the error distributions. This is especially important in scaling mass political attitudes, which include more “noise” than legislative roll call data. Owing to its flexibility as a nonparametric method and the efficiency of its search algorithms, OC has been shown to achieve superior classification performance on a range of political and non-political (e.g., Croft and Poole 2008) datasets.

We extend the fundamental geometry and algorithms underlying the (binary) OC method to accommodate polytomous or ordered choice data—particularly, the standard Likert and issue scale formats regularly used in public opinion survey research. Aldrich and McKelvey (1977) developed a revolutionary scaling method to account for differential item functioning in respondent usage of these scales and recover bias-corrected estimates of respondents and political stimuli (parties and candidates) positions in latent policy space. Poole (1998) later extended the Aldrich–McKelvey model to allow for the presence of missing data and the estimation of multiple ideological dimensions from issue scale data with his basic space (blackbox) procedure.

The ordered OC (OOC) procedure progresses from the OC and basic space lineages, allowing researchers to estimate multidimensional scaling results from ordered choice data without requiring strong parametric assumptions concerning individual utility or error processes. OOC, unlike IRT methods, also does not require the imposition of a priori identification restrictions when estimating multidimensional configurations. In addition to public opinion applications, OOC can be used to analyze ordinal choice data arising from expert surveys; roll call voting in legislative and judicial bodies where abstention is relevant; and public policies and political systems with categorical attributes.¹

¹ This includes, for instance, the Chapel Hill Expert Survey (Bakker et al. 2015), the Congressional Election Study (Stone and Simas 2010), the Convention Delegate Study (Layman et al. 2010), the United Nations (Bailey et al. 2017), the IMF (Thacker 1999), human rights policies (Fariss and Schnakenberg 2014), and legal systems (Rosenthal and Voeten 2007).

The paper proceeds as follows. The next section provides the theoretical motivation behind the application of OC to analyze ordinal choice data generally and public opinion survey data specifically. We then develop and describe the OOC procedure and reports the results of a series of Monte Carlo experiments that demonstrate its effectiveness in recovering the latent positions of voters and the orientation of issues through the latent space. Finally, we apply OOC to analyze public opinion survey data from the 2015 Cooperative Congressional Election Study and outline possible directions for future research.

2 The challenge of scaling public opinion and the Optimal Classification method

Generally stated, scaling methods are concerned with the measurement of unobservable, latent quantities. Scaling techniques recover the latent dimensions of the data and produce estimates of individuals' ideal points along those dimensions (e.g., Weisberg 1974). The low-dimensional space recovered from political choice data “constrains” a more complex set of attitudes on a multitude of political issues. The latent, organizing dimensions have been termed “basic” or “ideological” dimensions since they are closely related to how political ideologies weave together political attitudes into a consistent (if not necessarily logically coherent) whole (Converse 1964; Hinich and Munger 1994). For example, a political conservative is likely to oppose nationalized health care and environmental regulations, as well as to support tax cuts and abortion restrictions. Hinich and Munger (1994, 1997) emphasize the role of elite “packaging” of positions on what often seem to be unrelated issues (e.g., social welfare spending and gay marriage) in determining precisely how the latent, ideological space maps onto the issue dimensions.

Scaling methods have produced important findings about the underlying organization of political belief systems. For example, Poole and Rosenthal's (2007) NOMINATE scaling procedure shows that no more than two dimensions are needed to explain the vote choices of members of Congress, with the first dimension representing the familiar economic liberal-conservative continuum and the second dimension accounting for sources of party cleavages (e.g., racial issues in the mid-twentieth century). The use of scaling methods has advanced furthest in studies of legislative and judicial behavior and institutions (Poole and Rosenthal 2007; Clinton et al. 2004; Martin and Quinn 2002), but the use of public opinion data to generate spatial models of electoral competition (the locations of voters and/or candidates) has a long lineage itself (Weisberg and Rusk 1970; Poole and Rosenthal 1984; Palfrey and Poole 1987; Jacoby 1994; Treier and Hillygus 2009; Jessee 2012).

Indeed, there are some relative advantages to scaling public opinion data over sources of elite preferences. For one, the issue of separating the sincere and strategic elements of observed choice data is less of an issue than with roll call data (Poole and Rosenthal 1997) or campaign contributions (Bonica 2013).² Survey respondents have some motivation to disguise their true preferences on sensitive topics (i.e., social desirability bias) and clearly

² OOC can also be used to analyze the strategic and sincere components of roll call voting on a series of related dichotomous votes (particularly amendment voting) by coding legislators' voting patterns (e.g., YY, NY, NN) categorically (Silberman and Durden 1976; Nunez and Rosenthal 2004; Ladha 1991). As discussed in the next section, OOC uses a constrained normal vector to model the dichotomous components of a single ordinal scale. This means that all of the votes that comprise a given voting pattern will have an identical orientation in the recovered ideological space. We thank Howard Rosenthal for raising this point.

exhibit partisan and ideological bias in their responses to perceptual questions such as the perceived state of the economy (though monetary incentives have been shown to reduce these kinds of perceptual biases; Bullock et al. 2015). Otherwise, however, respondents have little incentive to answer preferential questions in a strategic manner. Public opinion surveys are also less constrained by the forces of agenda control and freer to gauge respondent preferences on a diverse set of issues and policy alternatives than is the case in professional legislative and judicial bodies.³ Finally, respondents can select from a wider array of choices than legislators (e.g., a seven-point issue scale vs. a binary yea/nay vote) and register their opinions with greater nuance.

However, the foremost challenge to scaling public opinion is the noisy and idiosyncratic nature of the data. Measurement error is present at high rates in even the most sophisticated public opinion surveys (Ansolabehere et al. 2008). Moreover, political attitudes have long been found to be less crystallized in the minds of voters than for political elites (Converse 1964; Zaller and Feldman 1992). As Lewis (2001, p. 276) notes: “Not only do we tend to observe far fewer bits of data from which to infer each voter’s preferences, but voter behavior appears to be more stochastic than legislative behavior.”

A nonparametric approach avoids these problems by not making strong assumptions about the functional form of respondents’ utility functions and the distribution of error. Parametric procedures like Poole and Rosenthal’s (2007) NOMINATE and Clinton, Jackman, and Rivers’s (2004) IDEAL Bayesian item response theory (IRT) model both make strong assumptions about the utility function (normal and quadratic, respectively) and the error distribution (logit for W-NOMINATE, and normal for DW-NOMINATE and IDEAL). Empirical evidence indicates that in the case of analyzing voting data in professional legislatures, assumptions about the nature of utility functions and error distributions are relatively benign (Poole 2000, p. 215). For example, the analysis of roll call voting in the European Parliament is unaffected by the choice to use parametric or nonparametric ideal point estimators (Hix et al. 2006).

This may not be the case for public opinion data. In particular, parametric models impose the unrealistic constraint that mass political attitudes are uniformly structured across the electorate. Specifically, this entails choosing which functional forms to model voters’ utility functions and error distributions. However, since Shepsle’s (1972) pioneering discussion of the influence of disposition towards risk on the shape of individual preference functions, no scholarly consensus has emerged as to which functional form (quadratic/concave, normal/convex, or linear/absolute distance) best characterizes voters’ utility functions. Brady and Ansolabehere (1989) found that voters possess utility functions which are well-structured but often contorted by the influence of factors like risk, uncertainty, and indifference. In electoral situations, evaluations about candidates’ intangible qualities or valence factors may also influence utility curves (Stokes 1963; Groseclose 2001). It is not surprising, then, that the literature has featured each of the competing models of quadratic utility (Alvarez 1997; Clinton and Jackman 2009), normal utility (Poole and Rosenthal 2007; Carroll et al. 2013), and linear/absolute distance utility (Berinsky and Lewis 2007). It seems likely that voters employ a wide variety of convex, concave, and linear functions in their utility formulations.

³ However, an underappreciated aspect of measuring mass ideology concerns the researcher’s reliance on which issues and policy alternatives are and are not included in public opinion surveys. This is a subtle—though still consequential—form of agenda control. We thank Bob Erickson for this observation.

The same idiosyncratic characteristics which introduce heterogeneity into respondents' utility functions also impede our ability to make parametric assumptions about voting errors. While mass political preferences are likely better structured (particularly in the contemporary, polarized political environment) than conventionally held (Sniderman and Bullock 2004), survey response data almost certainly remains heteroskedastic due to differences in political information and the frequency of competing predispositions and values among respondents (Alvarez and Brehm 1995; Kellstedt et al. 2017). Some voters—and indeed, political elites (Lauderdale 2010)—are simply more unpredictable when making political choices. In addition, Palfrey and Poole (1987) show that the heterogeneity of political information in the American electorate means that errors in voters' perceptions of ideological space are heteroscedastic. Thus, those with low political information are more likely to make voting errors than their highly informed counterparts. Severe violations of assumptions about the error process can have tremendous substantive consequences for scaling results (Rosenthal and Voeten 2004).

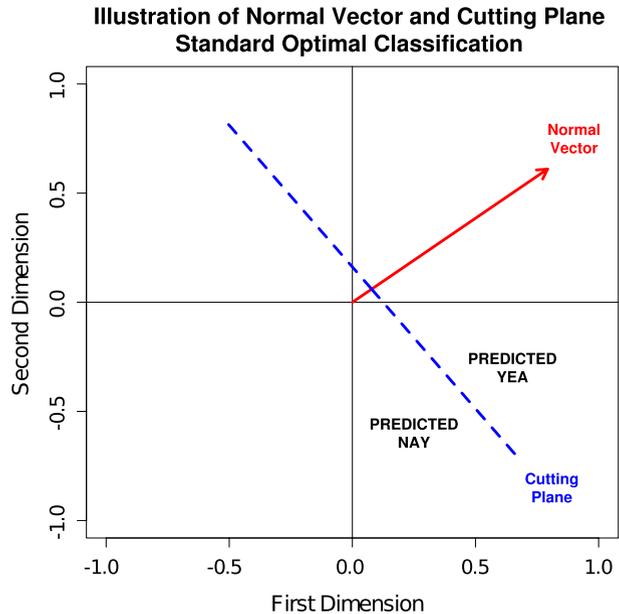
Consistent with Jacoby (1985) and Zaller and Feldman (1992), if we understand survey responses as arising from a data generating process in which survey respondents sample from their preference distributions, then it is critical to appreciate the nature of those distributions. Respondent preference distributions vary, likely considerably, within the electorate, and this makes the reliance of parametric assumptions problematic. The danger presented by the idiosyncratic characteristics of public opinion data is that scaling methods that either (a) seek to reproduce metric information from a correlation/covariance matrix of responses or (b) impose uniform parametric assumptions about the functional forms of respondents' utility functions and error distributions may corrupt the analysis.

For example, factor analysis is one popular parametric scaling method that is prone to exaggerating the dimensionality of the latent space (Coombs and Kao 1960; van Schuur and Kiers 1994; Brazill and Grofman 2002). This can have important substantive implications for the study of mass policy attitudes. For example, Jacoby (2008) shows that while confirmatory factor analysis of government spending attitudes produces a two-dimensional solution, Mokken scaling (a nonparametric IRT method) indicates that spending attitudes exhibit unidimensional structure. In this case, the unidimensional result is both parsimonious and exhibits broad explanatory power.

Many situations, though, suggest the presence of an underlying multidimensional policy space. Some scaling methods, though, otherwise constrain the dimensionality of the estimated configuration. For instance, Tahk (2018), develops a nonparametric method for ideal point estimation and inference that requires the assumption of unidimensionality. Dimensionality is also a concern with Bayesian item response theory (IRT) models. Multidimensional IRT models require an increasing number of constraints on the individual and/or item parameters to identify higher-dimensional solutions (Rivers 2003). Moreover, these constraints must be set a priori, requiring the researcher to make decisions about which issues correspond to the latent dimensions. Accordingly, the vast majority of IRT models in political science are unidimensional (but see Treier and Hillygus 2009; Sohn 2017), even in cases where multidimensional configurations would be of at least exploratory value.

We contend that OC is an ideal method for scaling public opinion data because it specifically addresses these issues. OC is a flexible nonparametric unfolding method that was built on the fundamental geometry of the spatial (geometric) model of voting (Poole 2000, 2005). Given a set of binary choice data (such as Yea and Nay votes by legislators along a series of roll call votes), OC produces a configuration of legislator ideal point coordinates and roll call cutting planes (which divide predicted Yeas from predicted Nays) that

Fig. 1 Illustration of standard (binary) Optimal Classification. The cutting plane (the dashed blue line) defines the prediction surface, while the normal vector (the solid red line) is perpendicular by construction and indicates the direction or orientation of the issue in the latent ideological space. (Color figure online)



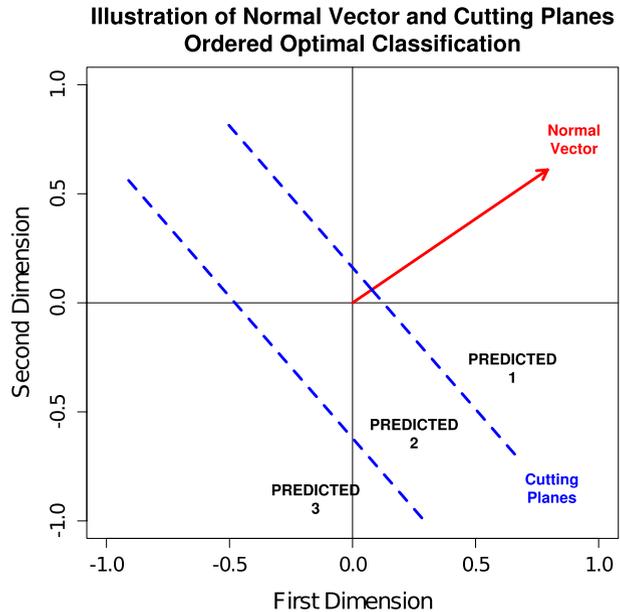
maximizes correct classification of the choices.⁴ Each roll call also has an estimated normal vector that is perpendicular to the cutting plane and indicates the direction of the policy alternatives through the latent space. Figure 1 illustrates the relationship between normal vectors and cutting planes in OC. Optimal Classification iteratively adjusts the cutting planes and normal vectors to maximize correct classification (or, equivalently, minimize the number of classification errors).

Cutting planes are almost certain to make classification errors on any given vote (e.g., incorrectly classifying “Yea” voters who are on the “Nay” side of the cutting plane). OC’s *cutting plane procedure* works from a starting configuration of voter coordinates and uses an iterative process to find cutting planes on each vote that maximize the number of voters who are correctly classified.⁵ OC creates polytopes from the intersection of the cutting planes from multiple votes (sometimes referred to as the Coombs mesh). The *legislator procedure* then searches through the grid to locate the available polytope for each voter which maximizes the correct classification. This produces the best available configuration of voters (points) and roll call votes (cutting lines) in a space of specified dimensionality.

⁴ Although OC is not guaranteed to find the global maximum, it regularly does so or gets very close to it. Poole (2000) reports the results of extensive Monte Carlo tests in one to ten dimensions which show that, at worst, only about 43 misclassifications per 50,000 total choices occur. Such a figure indicates that OC is very closely approximating the global classification maximum.

⁵ The starting values for the ideal points are obtained from an eigenvalue–eigenvector decomposition of the double-centered voter agreement score matrix.

Fig. 2 Illustration of Ordered Optimal Classification. A single normal vector is estimated for each survey item, and the $c - 1$ cutting planes (where c is the number of response categories in the corresponding item) continue to be restricted to be perpendicular to the normal vector



3 Ordered Optimal Classification

We expect OC to be an effective tool for estimating the ideological space of public opinion from survey data. However, as OC was developed with the goal of scaling binary choice (particularly legislative roll call) data, standard OC is not ideal for analyzing the ordered choice format encountered when analyzing survey data, roll call votes in which abstentions are relevant, and policies that are coded in an ordinal fashion. In voter and expert surveys, these data often come in the form of four or five-point Likert scales (in which respondents states how strongly they agree or disagree with a given statement or position) and seven or eleven-point issue scales (in which respondents place their most preferred position along a continuum of extreme left to extreme right positions on a given issue). Such responses can of course be coded in binary format, but this requires choosing a point at which to split the scale alongside losing information when collapsing the data to a binary scheme.

This consideration motivates our development of an ordinal version of the OC procedure. To extend OC to analyze polytomous or ordered choice data, we continue to estimate a single normal vector for each issue scale (as with binary OC), but wish to estimate multiple cutting planes that divide each pair of alternatives along the scale. For instance, we wish to separate predicted “Strongly Agree (1)” responses from “Somewhat Agree (2)” responses, and from “Somewhat Agree (2)” responses from “Somewhat Disagree (3)” responses, and so forth. An issue scale with c categories will require $c - 1$ cutting planes to classify the choices. As with binary OC, this configuration of cutting planes will form a Coombs mesh and OOC employs the same legislator algorithm in OC to search through the Coombs mesh for the polytope that maximizes each respondent’s correct classification. In addition, given a single normal vector for each issue scale, we also use the OC algorithm that searches for the point along the normal vector to locate the cutting plane so as to maximize that correct classification of the competing choices

Table 1 Binary coding of responses on a four-point issue scale

Response	1 versus higher	2 versus higher	3 versus higher
1 (strongly disagree)	Yea	Yea	Yea
2 (somewhat disagree)	Nay	Yea	Yea
3 (somewhat agree)	Nay	Nay	Yea
4 (strongly agree)	Nay	Nay	Nay

(1 vs. 2, 2 vs. 3, etc.). Figure 2 illustrates OOC's extension of normal vectors and cutting planes to model polytomous choices.

Because we wish to continue employing the specific search algorithms underlying the OC procedure, we need to organize the polytomous data in binary fashion. To do so, we adopt van Schuur's (2011, p. 74) coding scheme that dichotomizes polytomous choices with a series of $c - 1$ binary choices (where c is the number of categories in the scale). In this scheme, respondents first choose between the lowest category and higher categories. For this choice, all respondents who chose 1 are coded as Yeas, and all others are coded as Nays. The second binary choice is between the second-lowest category and higher categories. For this choice, respondents who chose 1 or 2 are coded as Yeas, and all others are coded as Nays. This sequence proceeds until the choice between the second-highest category and the highest category is reached. Table 1 provides an example of how responses to a 4-point scale are coded using $c - 1$ (or 3) binary choices.

The reason that the new binary choices cannot be analyzed by standard OC is that a separate normal vector would be estimated for each of the $c - 1$ choices, even though they are generated from the same issue scale. The normal vector represents the direction of a given issue through the latent space, and of course each issue scale (and the choices along it) represent the same issue.

Accordingly, we add a new step to OOC that uses a regression model to recalculate the normal vectors at each iteration of the procedure. As explained in Poole (2005, pp. 37–40), the coefficients from a regression of the respondent choices onto the ideal point coordinates x can be used to calculate the normal vector of item j on the k th dimension ($k = 1, \dots, s$) using Eq. 1:

$$NV_{jk} = \frac{\beta_{jk}}{\sqrt{\beta_{j1}^2 + \dots + \beta_{js}^2}} \quad (1)$$

The most straightforward way to estimate Eq. 1 is to use a generalized linear model with an appropriate link function (such as ordered probit) to regress the survey responses to item j on the s -dimensional ideal point coordinates and substitute the regression coefficients into $\beta_{j1}, \beta_{j2}, \dots, \beta_{js}$. This approach is computationally efficient, but it also changes OOC from a nonparametric to a semiparametric estimator. For those who wish to preserve the nonparametric character of the standard OC procedure, we implement two alternative kernel-based methods to ordered probit regression in estimating the normal vector routine.

The first nonparametric alternative is a support vector machine (SVM) routine with a linear kernel function.⁶ In classification problems, SVMs estimate the separating hyperplane between two classes of data that maximizes the margin M : the distance between the separating hyperplane and a subset of observations (termed “support vectors”) that are closest to the decision boundary.⁷ In the regression setting, SVMs estimate a generic prediction function \hat{f} in the predictor space that considers only errors greater than a certain magnitude (ϵ) in calculating the loss function. Hence, in both classification and regression applications, SVMs ignore “low error” observations that are distant from the separating hyperplane (classification problems) or have small residuals (regression problems) and instead focus on optimizing prediction for “high error” points (Hastie et al. 2009). In both cases, we can transpose the coefficients that define the separating hyperplanes for each issue (β'_{jk}) and insert them into Eq. 1 to estimate the normal vector routine.⁸

The second alternative method we implement for estimating the normal vector routine is Hainmueller and Hazlett’s (2014) kernel-based regularized least squares (KRLS) procedure.⁹ KRLS is a nonparametric analogue to regression that avoids the standard linearity and/or additivity assumptions implicit in regression-based models. It instead uses a kernel function to measure the covariate similarity between each pair of observations. Observations with greater similarity (i.e., closer proximity in the covariate space) to a particular point x^* exert greater influence in defining the prediction surface at that point ($f(x^*)$).¹⁰ We use the average partial derivative of y_j (the responses to item j) with respect to each x_k (the ideal point coordinates on the k th dimension) as our estimate of β_{jk} , allowing us to estimate Eq. 1 in a nonparametric fashion.

We make all three options for estimating the normal routine—ordered probit, support vector classification/regression, and KRLS—available in the `ooc` package in R.¹¹ To estimate the OOC procedure, we first obtain starting values for the respondent ideal points (by running standard OC on the generated binary choices) and the issue normal vectors (by running our normal vector routine of choice on the initial values of the ideal point coordinates). OOC then successively loops through a series of iterations, at each iteration estimating (1) the legislator procedure, (2) the cutting plane procedure, and (3) the normal vector routine. We have found that OOC converges on the solution fairly quickly—by

⁶ As Adam Bonica has pointed out, OC and SVMs are closely related in their pursuit of a separation hyperplane that optimally divides two classes of data. The major differences between the two methods concerns their loss functions (OC uses overall correct classification rates while SVMs assess classification performance using both the correct classification rate and the robustness of the derived hyperplane), constraints on the separating hyperplane (SVMs allow for nonlinear separating hyperplanes, while OC uses strictly linear separating hyperplanes), and their treatment of the predictor variables X (standard SVMs requires the predictor variables to be observed, while in OC the predictor variables [i.e., the ideal point coordinates of the observations] are treated as latent variables to be estimated) (Hastie et al. 2009; Bonica 2018).

⁷ The cost parameter C controls the width of the margin—i.e., the number of observations that are allowed to violate the margin and constitute the support vectors.

⁸ We use a linear kernel to simplify analysis, although one of the attractive properties of SVMs is that alternative kernel functions can be used to estimate nonlinear decision boundaries and prediction functions. We note that the adaptation of such kernels provides one possible avenue for future development of the OOC procedure.

⁹ Implemented in the `krls` package in R (Ferwerda et al. 2017).

¹⁰ We follow Hainmueller and Hazlett (2014) and use the Gaussian kernel function.

¹¹ We have found that support vector regression (SVR) tends to slightly outperform its competitors in terms of classification performance, and does so with reasonable computation efficiency. Hence, we set SVR as the default method for the normal vector routine and use it in the analyses presented in this paper.

around five iterations, the estimated parameters from successive iterations are correlated at > 0.99 . By default, we run OOC for 25 iterations, but users may increase this number or establish a stopping criterion to terminate the procedure.

3.1 Monte Carlo experiments

We next report the results of a series of Monte Carlo experiments in which ordered survey response data is generated under a variety of specifications and used to assess the performance of OOC in recovering the true respondent and issue locations in ideological space. Each simulation uses 1500 survey respondents (i in $1, \dots, n$) and 40 issue scales (j in $1, \dots, q$). All of the issue scales have five points or response categories (c in $1, \dots, C$). Respondent ideal points (x_i) in two and three-dimensional ideological space are drawn from a multivariate normal distribution with mean 0 and interdimension correlations (i.e., the off-diagonal of the Σ matrix) randomly drawn from a uniform distribution between -0.1 and 0.7 . The issue normal vectors N_j are randomly drawn from the edges of a unit hypersphere in the corresponding two or three-dimensional space, with the outcome locations (O_1, O_2, \dots, O_5) randomly selected and projected onto their respective issue normal vectors. Respondent ideal points are also projected onto each of the q normal vectors ($\kappa_i = x_i N_j'$), which simplifies the math by allowing us to work with relative distances along a single dimension.

We calculate respondent utilities for each point on the issue scales using the random utility model $U_{ijc} = F(\|x_i - O_{jc}\|) + \epsilon_j$, with issue and alternative-specific values for O_{jc} , random issue-specific shocks ϵ_j , and where $\|\cdot\|$ denotes Euclidian distance (McFadden 1976). Our choice of F comes from three standard functional forms: linear ($1 - |x_i - O_j|$), normal/Gaussian ($\exp(-\frac{1}{2}(x_i - O_j))$), and quadratic ($1 - \frac{1}{2}(x_i - O_j)^2$). Each simulation randomly selects probability weights for each of the three utility functions and then uses these weights to randomly assign respondents to one of the three functional forms in calculating their utilities. As a result, the simulations will reflect a diverse set of utility rules.

Finally, in order to allow for the presence of heteroskedastic error, we randomly sample respondent-specific error variances (σ_i^2) from a uniform distribution between 0 and 0.75 and issue-specific error variances (σ_j^2) from a gamma distribution with a scale parameter of 0.5 and a shape parameter randomly selected between 0 and 3. Following Lauderdale (2010), we calculate the probability of a given response c on issue j by respondent i as:

$$Pr_{ijc} = \Phi \left(\frac{\left(\frac{U_{ijc}}{\exp(\sigma_i^2 \sigma_j^2)} \right)}{\sum_{c=1}^C \left(\frac{U_{ijc}}{\exp(\sigma_i^2 \sigma_j^2)} \right)} \right) \quad (2)$$

where Φ is the standard normal CDF. Responses are then generated using this probability matrix. Missing values are randomly inserted into 10, 25, and 40% of the entries in the final response matrix.

Figures 3 and 4 summarizes the results of the Monte Carlo experiments, varying the levels of error and missingness across 100 simulations. The level of error is determined by the proportion of spatially incorrect choices by respondents: between 0.26 and 0.51 for “low,” 0.49 and 0.59 for “medium,” and 0.47 and 0.72 for “high.” We measure performance using three sets of correlations: between the interpoint distances of the true

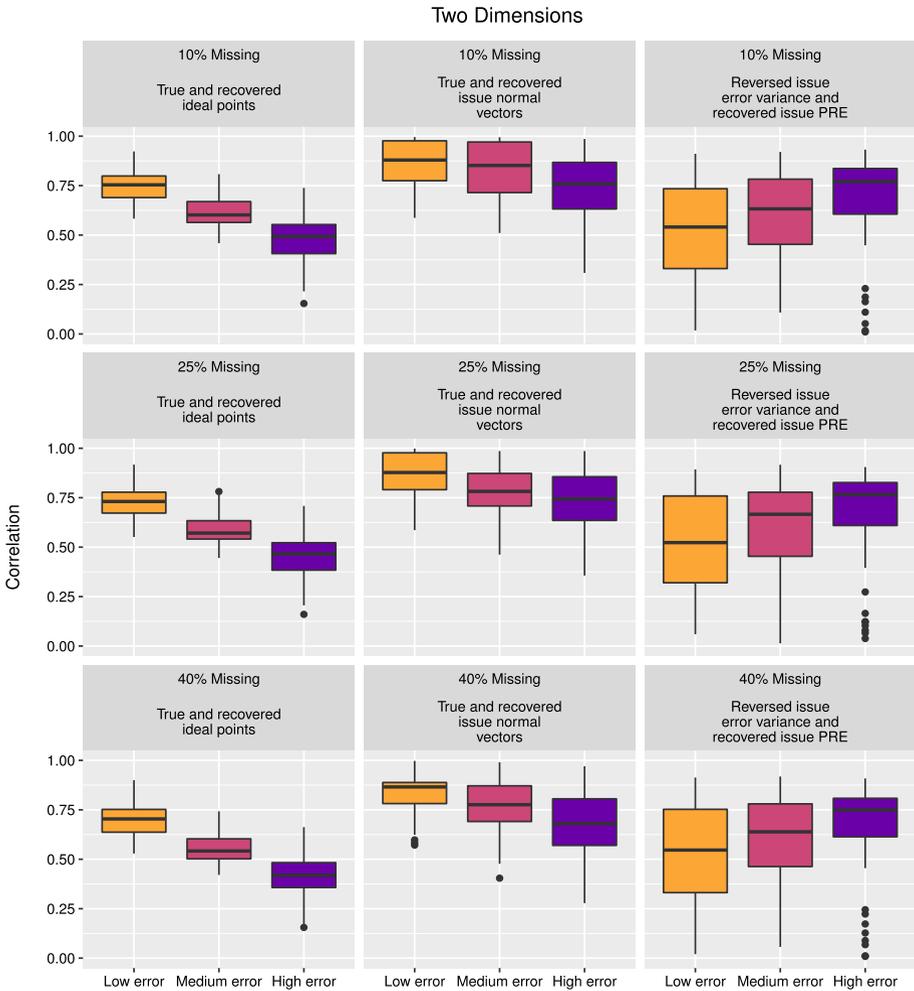


Fig. 3 Monte Carlo tests of Ordered Optimal Classification performance in two dimensions

and recovered respondent ideal point configurations, between the interpoint distances of the true and recovered normal vector configurations, and the simulated issue error variances and the recovered issue fit statistic (for which we use the proportional reduction in error, or PRE).

The results indicate that even with highly noisy, error-laden voting data (in which between half and three-quarters of respondents’ simulated choices are spatially incorrect), the correlations between the true and recovered normal vector configurations seldom fall below 0.7. In cases with more reasonable levels of error, the correlations are usually about 0.8 or higher. Hence, OOC appears to perform exceptionally well in its recovery of the relative ideological orientation of the individual issue scales. Its recovery of the respondent ideal points is less impressive, but nonetheless respectable in situations where the rate of spatially incorrect choices is less than one-half (i.e., the “low

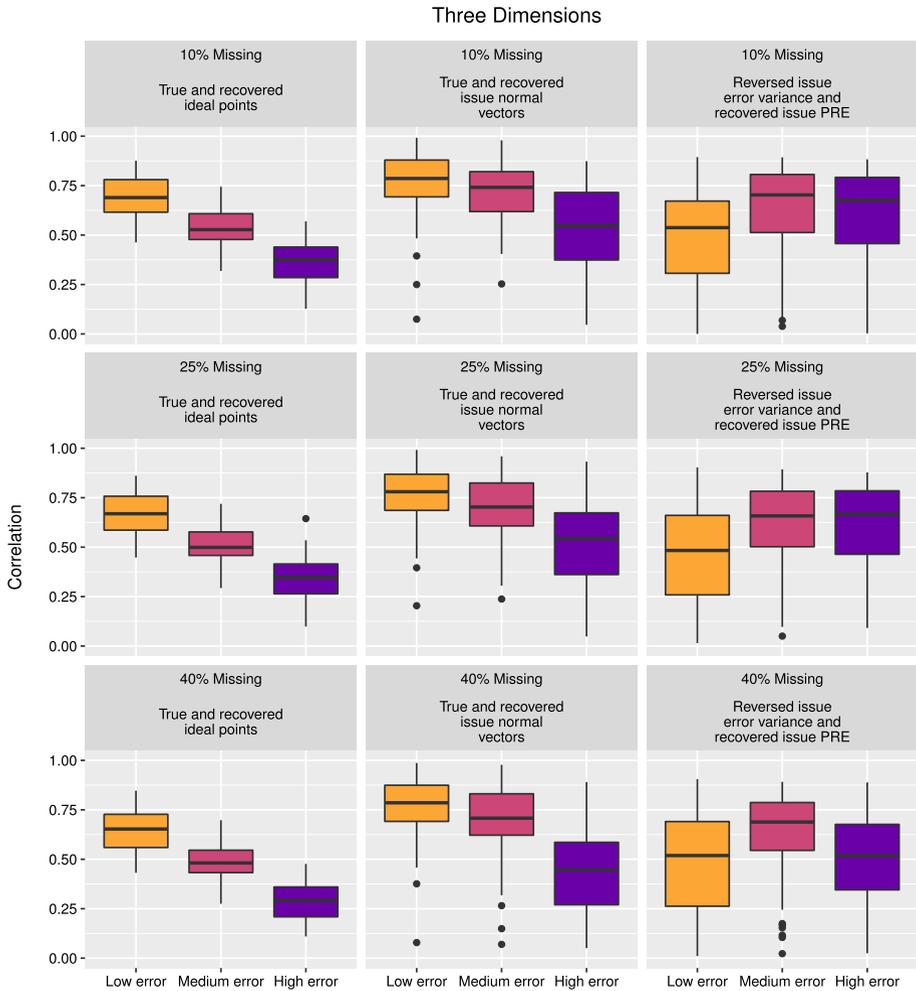


Fig. 4 Monte Carlo tests of Ordered Optimal Classification performance in three dimensions

error” category). In these cases, the interpoint distances between the true and recovered ideal point configurations are correlated around $r = 0.7$.

In addition, the OOC procedure is mostly unaffected by the level of missingness in the data and the number of dimensions estimates, both results that are consistent with Monte Carlo experiments on the original OC procedure (Poole 2000). The issue-specific PRE fit statistics better reflect the underlying level of error variance in conditions where the overall level of incorrect spatial voting is higher. While these correlations are relatively low, they are nonetheless large enough as to provide a good indication as to how well specific attitudes are structured in ideological space.

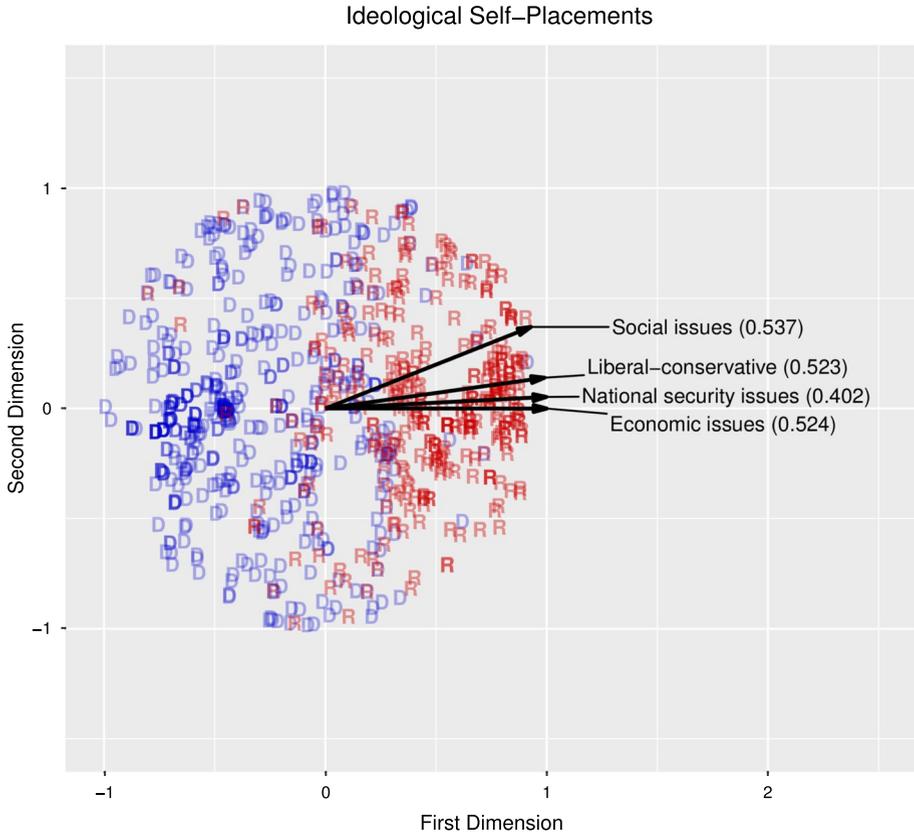


Fig. 5 Ordered Optimal Classification scaling of the 2015 Cooperative Congressional Election Study data. Democratic (D) and Republican (R) respondents shown alongside the selected ideological question normal vectors (with PRE fit statistics in parentheses)

4 Application: assessing ideological structure in the Contemporary American Electorate

We next fit OOC to 52 survey questions from the University of Georgia module of the 2015 Cooperative Congressional Election Study. The questions cover both policy issues and three core value batteries: economic egalitarianism (six questions), moral traditionalism (four questions), and militarism (two questions). The issue questions cover abortion, the environment, spending preferences, LGBT rights, gun control, health care, immigration, social welfare programs, religion and morality, free trade, the military, and foreign intervention. We also include respondents' liberal-conservative self-placements alongside three similar seven-point scales from Klar (2014) that specifically ask respondents about their broad economic, social/cultural, and national security ideological orientations.

In two dimensions, OOC correctly classifies 73.1% of respondent choices with an APRE (aggregate proportional reduction in error) value of 0.415, indicating a good fit of these data to a low-dimensional ideological model (particularly considering the ordinal nature of many of the survey items). Figure 5 displays the respondent ideal points (denoted by *D* for

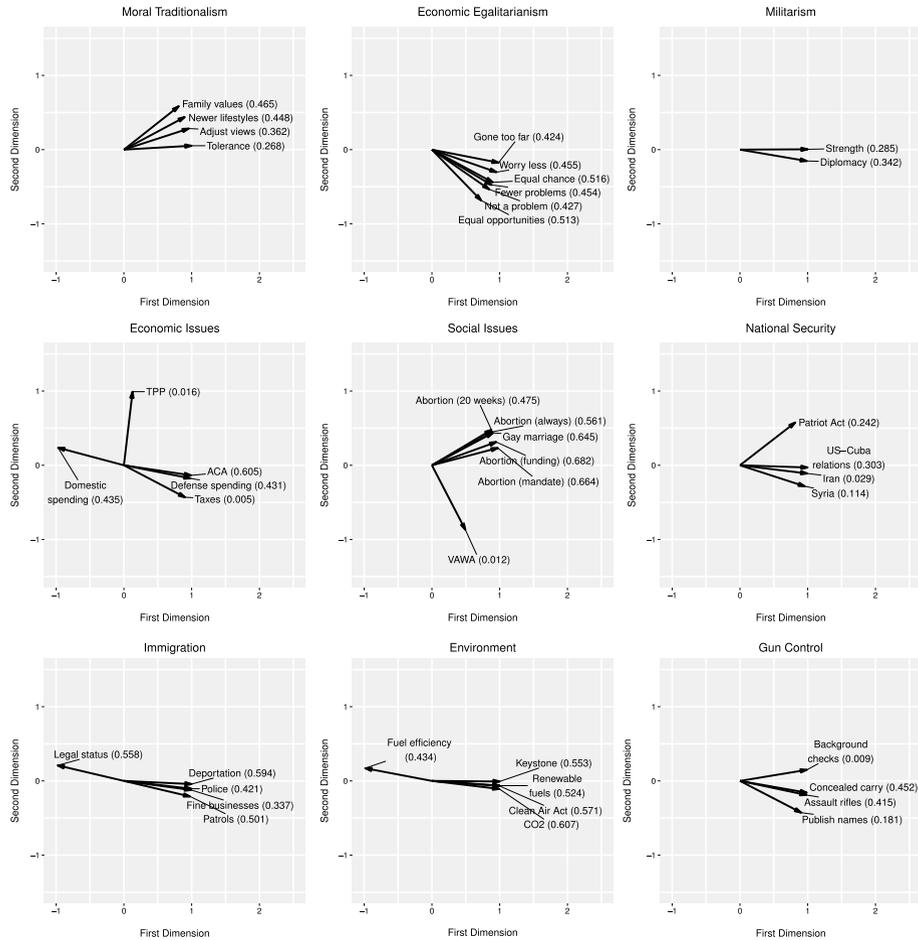


Fig. 6 Selected normal vectors from Ordered Optimal Classification scaling of the 2015 Cooperative Congressional Election Study. Arrows denote the relative orientation of policy preferences and core values in the latent two-dimensional ideological space

Democratic identifiers/leaners and *R* for Republican identifiers/leaners), and this configuration shows a clear split between partisans along the first dimension. If the first dimension recovered OOC indeed represents the familiar partisan/ideological conflict dimension that has become increasingly salient in American politics in the polarization context (e.g., Abramowitz 2010), then we should be able to map individuals’ ideological attachments and identifications onto this space.

We thus add to Fig. 5 the normal vector projections (with PRE fit statistics in parentheses) for individuals’ general symbolic identification as ideological liberals or conservatives alongside their responses to the more operational economic, social, and national security ideology scales. Although the concepts can be disjointed—for example, self-identified conservatives routinely espouse support for element of the social welfare state (Stimson 2004)—evidence shows that voters have become more aware of party differences and more likely to identify ideologically as elite partisanship has

intensified and partisan and ideological cues have become clearer (Bafumi and Shapiro 2009; Jacoby 2002; Smidt 2017). Therefore, we expect these items to map cleanly onto the first dimension in the CCES data. We observe precisely this result, which highlights voters' reasonable degree of "vertical constraint," Converse's (1964) term for the higher-order linkages between individuals' ideological labels and their policy preferences. Though some separation is evident between the items representing different economic, social, and national security domains, there is nonetheless a clean constraining effect of a single liberal-conservative ideological dimension that bisects the economic and social ideology vectors.

Figure 6 expands this analysis by showing the normal vectors for selected groups of survey items. Here, we find greater separation between the different economic, social, and national security policy domains. Crucially, though, attitudes are still well-constrained in two-dimensional space: both issues and values that tap into similar domains (for instance, moral traditionalism and specific social/cultural issue attitudes) show a considerable degree of overlap in the same regions of the ideological space. The economic issues and values tend to project into the southeast quadrant of the space, while the social normal vectors are located mostly in the northeast quadrant. The national security normal vectors are generally positioned between the economic and social normal vector clusters. Interestingly, attitudes on gun control (which is often classified as social/cultural issue) shows greater proximity to the economic-focused normal vectors.

Especially revealing are the results concerning the ideological orientation of core values representing individuals' orientation toward bedrock principles in American life. The core value normal vectors all have moderate-to-high PRE values, indicating a good spatial fit, and suggest these values provide a deeper moral or belief-based foundation to respondents' ideological preferences. Egalitarianism is a longstanding cultural value (McClosky and Zaller 1984), and postures toward equality and the redistribution of economic resources is a preeminent source of elite political contestation and mass public preoccupation (Jacoby 2014; Lane 1959; Layman and Carsey 2002). Evidence shows that egalitarianism shapes attitudes to an array of issue attitudes involving government transfers and questions of fairness—including government spending (Feldman 1988; Goren 2008)—and the value also relates strongly to partisanship and ideology (Keele and Wolak 2006; Jacoby 2006, 2014; Lupton et al. 2017). Moral traditionalism, or support for traditional family arrangement and social strictures, relates to hot-button social attitudes ranging from abortion (Alvarez and Brehm 1995) to gay marriage (Brewer 2003) and transgender rights. And, as with egalitarianism, the value is associated with partisan and ideological identities (Layman and Green 2006; Weisberg 2005). Militarism, measured with a two-item scale concerning relative preferences for diplomacy versus military force and strength versus understanding in foreign policy, has been less studied (but see Rathbun et al. 2016), but also taps into fundamental personality and authoritarian dispositions (Hetherington and Weiler 2009) and also corresponds to the immigration issue normal vectors in our results.

Taken together, our previous results and the significant relationship observed between the ideological mappings of core values and policy preferences estimated by our OOC shows that the principles, predispositions and policy attitudes that most help citizens navigate an often-bewildering political environment and are most closely associated with electoral choice all occupy a well-structured low-dimensional space in voters' minds. We conclude that despite substantial heterogeneity in mass public opinion, the American electorate brings to bear coherent orientations when they confront the political world in the age of polarization.

5 Discussion

Optimal Classification is a novel and useful tool for researchers who wish to empirically assess the latent space of political actors' policy preferences in a nonparametric manner. The spatial "maps" produced by such scaling methods produce have supplied the discipline with geometric models of political cleavages and ideological divides. These models constitute meaningful scientific progress (Poole 2017) and have become "iconic and revelatory" tools for conducting and presenting research in political science (Brady 2011).

The subfields of public opinion and voting behavior have been among the many beneficiaries of empirical spatial voting models. In an age of polarization, contemporary American politics has increasingly become defined by the intersection of issue preferences and value dispositions in a combined multidimensional partisan-ideological space (Jacoby 2014; Gibson and Hare 2016). The examination of such phenomena naturally lend themselves to spatial-based analyses, but the application of scaling methods to public opinion data have been hampered by violations of parametric assumptions concerning individual utility and the error term and/or limitations in estimating multiple latent dimensions without strict identifying assumptions.

The Optimal Classification framework addresses both of these problems and offers a flexible exploratory tool to uncover new insights into policy spaces in mass electorates. This paper's extension of standard OC to accommodate ordered choice data allows public opinion researchers to analyze survey data within this framework. We also wish to emphasize the potential applications of OOC in other contexts in which political actors' preferences and judgments are measured in an ordinal fashion: expert surveys, voting behavior in institutions where abstentions are meaningful, and public policy and legal systems with categorical features.

An important next step in this research program is to develop uncertainty estimates for inferential purposes, either using the nonparametric bootstrap to randomly sample respondents and/or issues or, following Bonica (2014), a jackknifing scheme to estimate standard errors for these quantities. Particularly for the issue normal vectors, this can serve to substantiate the claim that a diverse set of issue attitudes are collapsing onto a single ideological dimension for the American electorate. Of course, the inclusion of survey data from multiple years is also needed to elucidate the dynamics of the dimensionality of American public opinion, particularly as concerns the mapping of core values and beliefs onto the ideological space. Finally, we plan to break down the item mappings in Fig. 6 by party, as both the direction and fit of issues and attributes may vary across partisan lines.

Appendix

Below we perform two additional sets of Monte Carlo experiments on the statistical properties of the Ordered Optimal Classification estimator. The first assesses OOC's recovery of the true ideal points of ideologically moderate and ideologically extreme respondents. The second provides an informal test of the consistency of the OOC estimator, replicating

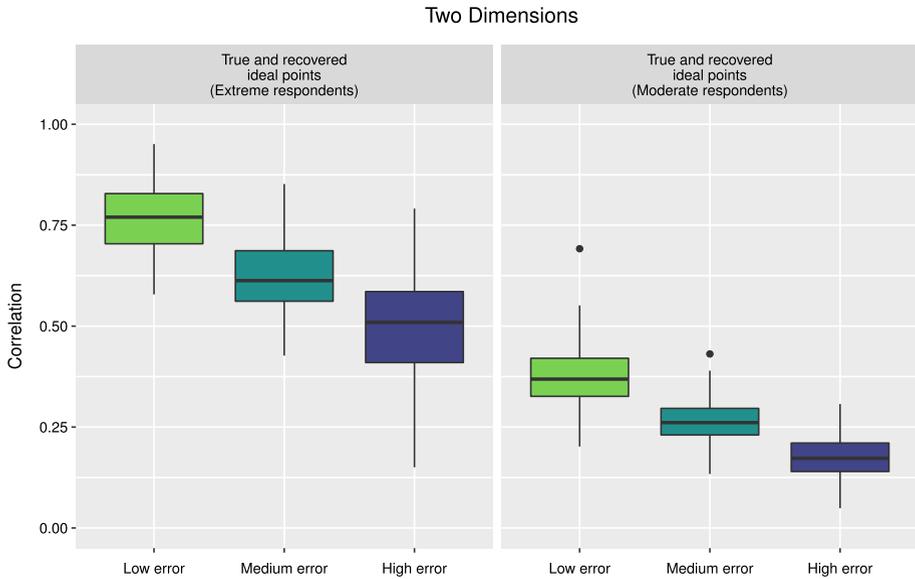


Fig. 7 Monte Carlo tests of Ordered Optimal Classification recovery of ideologically moderate and extreme respondents' ideal points

the analysis in the main text (specifically, Figs. 3, 4) across an increasing number (25, 50, and 100) of issue questions.

Figure 7 shows the correlations between the true and recovered ideal points separately for ideologically moderate and extreme respondents in two dimensions. Ideologically moderate respondents are defined as those with ideal points in the interquartile range on both dimensions, while ideologically extreme respondents are those with ideal points outside of the interquartile range on both dimensions. OOC clearly performs better in its recovery of extremists' ideal points, though some of this is an artifact of the wider (more polarized) range of ideal point values for ideologically extreme respondents relative to moderate respondents.

Figures 8 and 9 replicate the analysis in Figs. 3 and 4 while increasing the number of simulated issue scales from 25 to 50 to 100. If OOC is a consistent estimator, the correlations between true and estimated parameters should increase alongside the number of issues for a given level of error and dimensionality. This is precisely what we observe, with the improvements most apparent in the three-dimensional case in Fig. 9.

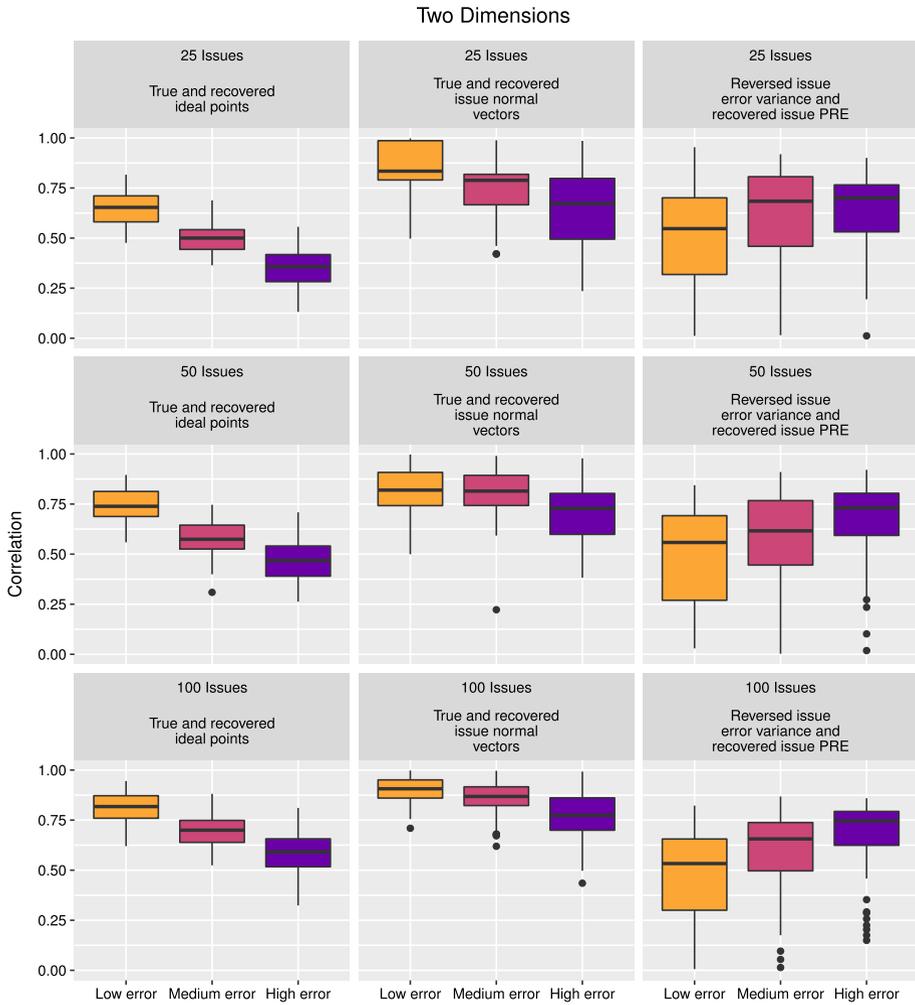


Fig. 8 Monte Carlo tests of Ordered Optimal Classification performance in two dimensions with 25, 50, and 100 issues

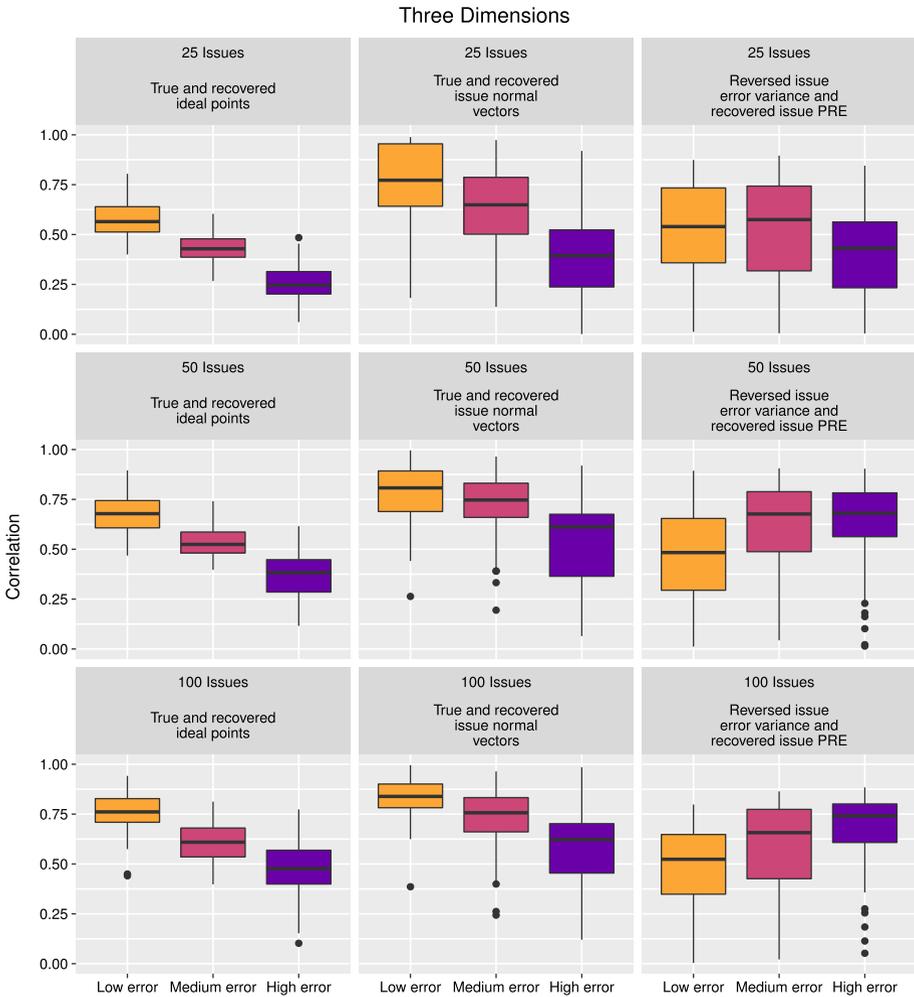


Fig. 9 Monte Carlo tests of Ordered Optimal Classification performance in three dimensions with 25, 50, and 100 issues

References

- Abramowitz, A. (2010). *The disappearing center: Engaged citizens, polarization, and American democracy*. New Haven: Yale University Press.
- Aldrich, J. H., & McKelvey, R. D. (1977). A method of scaling with applications to the 1968 and 1972 presidential elections. *American Political Science Review*, *71*(1), 111–130.
- Alvarez, R. M. (1997). *Information and elections*. Ann Arbor: University of Michigan Press.
- Alvarez, R. M., & Brehm, J. (1995). American ambivalence towards abortion policy: Development of a heteroskedastic probit model of competing values. *American Journal of Political Science*, *39*(4), 1055–1082.
- Ansolabehere, S., Rodden, J., & Snyder, J. M, Jr. (2008). The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting. *American Political Science Review*, *102*(2), 215–232.
- Bafumi, J., & Shapiro, R. Y. (2009). A new partisan voter. *Journal of Politics*, *71*(1), 1–24.
- Bailey, M. A., Strezhnev, A., & Voeten, E. (2017). Estimating dynamic state preferences from United Nations voting data. *Journal of Conflict Resolution*, *61*(2), 430–456.
- Bakker, R., Vries, Cd, Edwards, E., Hooghe, L., Jolly, S., Marks, G., et al. (2015). Measuring party positions in Europe: The Chapel Hill Expert Survey trend file, 1999–2010. *Party Politics*, *21*(1), 143–152.
- Berinsky, A. J., & Lewis, J. B. (2007). An estimate of risk aversion in the U.S. electorate. *Quarterly Journal of Political Science*, *2*(2), 139–154.
- Bonica, A. (2013). Ideology and interests in the political marketplace. *American Journal of Political Science*, *57*(2), 294–311.
- Bonica, A. (2014). The punctuated origins of Senate polarization. *Legislative Studies Quarterly*, *39*(1), 5–26.
- Bonica, A. (2018). Inferring roll call scores from campaign contributions using supervised machine learning. *American Journal of Political Science*. <https://doi.org/10.2139/ssrn.2732913>.
- Brady, H. E. (2011). The art of political science: Spatial diagrams as iconic and revelatory. *Perspectives on Politics*, *9*(2), 311–331.
- Brady, H. E., & Ansolabehere, S. (1989). The nature of utility functions in mass publics. *American Political Science Review*, *83*(1), 143–163.
- Brazill, T. J., & Grofman, B. (2002). Factor analysis versus multi-dimensional scaling: Binary choice roll-call voting and the US Supreme Court. *Social Networks*, *24*(3), 201–229.
- Brewer, P. R. (2003). The shifting foundations of public opinion about gay rights. *Journal of Politics*, *65*(4), 1208–1220.
- Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, *10*(4), 519–578.
- Carroll, R., Lewis, J. B., Lo, J., Poole, K. T., & Rosenthal, H. (2013). The structure of utility in spatial models of voting. *American Journal of Political Science*, *57*(4), 1008–1028.
- Clinton, J. D., & Jackman, S. (2009). To simulate or NOMINATE? *Legislative Studies Quarterly*, *34*(4), 593–621.
- Clinton, J., Jackman, S., & Rivers, D. (2004). The statistical analysis of roll call data. *American Political Science Review*, *98*(2), 355–370.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In D. E. Apter (Ed.), *Ideology and discontent* (pp. 206–261). New York: Free Press.
- Coombs, C., & Kao, R. (1960). On a connection between factor analysis and multidimensional unfolding. *Psychometrika*, *25*(3), 219–231.
- Croft, W., & Poole, K. T. (2008). Inferring universals from grammatical variation: Multidimensional scaling for typological analysis. *Theoretical Linguistics*, *34*(1), 1–37.
- Enelow, J. M., & Hinich, M. J. (1984). *The spatial theory of voting: An introduction*. New York: Cambridge University Press.
- Fariss, C. J., & Schnakenberg, K. (2014). Measuring mutual dependence between state repressive actions. *Journal of Conflict Resolution*, *58*(6), 1003–1032.
- Feldman, S. (1988). Structure and consistency in public opinion: The role of core beliefs and values. *American Journal of Political Science*, *32*(2), 416–440.
- Ferwerda, J., Hainmueller, J., & Hazlett, C. J. (2017). Kernel-based regularized least squares in R (KRLS) and Stata (krls). *Journal of Statistical Software*, *79*(3), 1–26.
- Gibson, T., & Hare, C. (2016). Moral epistemology and ideological conflict in American political behavior. *Social Science Quarterly*, *97*(5), 1157–1173.
- Goren, P. (2008). The two faces of government spending. *Political Research Quarterly*, *61*(1), 147–157.

- Groseclose, T. (2001). A model of candidate location when one candidate has a valence advantage. *American Journal of Political Science*, 45(4), 862–886.
- Hainmueller, J., & Hazlett, C. (2014). Kernel regularized least squares: Reducing misspecification bias with a flexible and interpretable machine learning approach. *Political Analysis*, 22(2), 143–168.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). New York: Springer.
- Hetherington, M. J., & Weiler, J. D. (2009). *Authoritarianism and polarization in American politics*. Cambridge: Cambridge University Press.
- Hinich, M. J., & Munger, M. C. (1994). *Ideology and the theory of political choice*. Ann Arbor: University of Michigan Press.
- Hinich, M. J., & Munger, M. C. (1997). *Analytical politics*. Cambridge: Cambridge University Press.
- Hix, S., Noury, A., & Roland, G. (2006). Dimensions of politics in the European Parliament. *American Journal of Political Science*, 50(2), 494–511.
- Jacoby, W. G. (1985). Inconsistent preferences and the multidimensional unfolding model. *Political Methodology*, 11(3/4), 201–220.
- Jacoby, W. G. (1994). Public attitudes toward government spending. *American Journal of Political Science*, 38(2), 336–361.
- Jacoby, W. G. (2002). Core values and political attitudes. In B. Norrander & C. Wilcox (Eds.), *Understanding public opinion* (2nd ed., pp. 177–201). Washington, DC: CQ Press.
- Jacoby, W. G. (2006). Value choices and American public opinion. *American Journal of Political Science*, 50(3), 706–723.
- Jacoby, W. G. (2008). Comment: The dimensionality of public attitudes toward government spending. *Political Research Quarterly*, 61(1), 158–161.
- Jacoby, W. G. (2014). Is there a culture war? Conflicting value structures in American public opinion. *American Political Science Review*, 108(4), 754–771.
- Jessee, S. A. (2009). Spatial voting in the 2004 presidential election. *American Political Science Review*, 103(1), 59–81.
- Jessee, S. A. (2012). *Ideology and spatial voting in American elections*. Cambridge: Cambridge University Press.
- Keele, L., & Wolak, J. (2006). Value conflict and volatility in party identification. *British Journal of Political Science*, 36(4), 671–690.
- Kellstedt, P. M., Ramirez, M. D., Vedlitz, A., & Zahran, S. (2017). Does political sophistication minimize value conflict? Evidence from a Heteroskedastic graded IRT model of opinions toward climate change. *British Journal of Political Science*. <https://doi.org/10.1017/S0007123417000369>
- Klar, S. (2014). A multidimensional study of ideological preferences and priorities among the American public. *Public Opinion Quarterly*, 78(S1), 344–359.
- Ladha, K. K. (1991). A spatial model of legislative voting with perceptual error. *Public Choice*, 68(1–3), 151–174.
- Lane, R. (1959). The fear of equality. *American Political Science Review*, 53(1), 35–51.
- Lauderdale, B. E. (2010). Unpredictable voters in ideal point estimation. *Political Analysis*, 18(2), 151–171.
- Layman, G. C., & Carsey, T. M. (2002). Party polarization and “conflict extension” in the American electorate. *American Journal of Political Science*, 46(4), 786–802.
- Layman, G. C., Carsey, T. M., Green, J. C., Herrera, R., & Cooperman, R. (2010). Activists and conflict extension in American party politics. *American Political Science Review*, 104(2), 324–346.
- Layman, G. C., & Green, J. C. (2006). Wars and rumours of wars: The contexts of cultural conflict in American political behaviour. *British Journal of Political Science*, 36(1), 61–89.
- Lewis, J. B. (2001). Estimating voter preference distributions from individual-level voting data. *Political Analysis*, 9(3), 275–297.
- Lupton, R. N., Myers, W. M., & Thornton, J. R. (2015). Political sophistication and the dimensionality of elite and mass attitudes, 1980–2004. *Journal of Politics*, 77(2), 368–380.
- Lupton, R. N., Smallpage, S. M., & Enders, A. M. (2017). Values and political predispositions in the age of polarization: Examining the relationship between partisanship and ideology in the United States, 1988–2012. *British Journal of Political Science*. <https://doi.org/10.1017/S0007123417000370>
- Martin, A. D., & Quinn, K. M. (2002). Dynamic ideal point estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953–1999. *Political Analysis*, 10(2), 134–153.
- McClosky, H., & Zaller, J. (1984). *The American ethos: Public attitudes toward capitalism and democracy*. Cambridge, MA: Harvard University Press.
- McFadden, D. L. (1976). Quantal choice analysis: A survey. *Annals of Economic and Social Measurement*, 5(4), 363–390.

- Nunez, S., & Rosenthal, H. (2004). Bankruptcy “reform” in Congress: Creditors, committees, ideology, and floor voting in the legislative process. *Journal of Law, Economics, and Organization*, 20(2), 527–557.
- Palfrey, T. R., & Poole, K. T. (1987). The relationship between information, ideology, and voting behavior. *American Journal of Political Science*, 31(3), 511–530.
- Poole, K. T. (1998). Recovering a basic space from a set of issue scales. *American Journal of Political Science*, 42(3), 954–993.
- Poole, K. T. (2000). Nonparametric unfolding of binary choice data. *Political Analysis*, 8(3), 211–237.
- Poole, K. T. (2005). *Spatial models of parliamentary voting*. New York: Cambridge University Press.
- Poole, K. T. (2017). The scientific status of geometric models of choice and similarities judgment. *Public Choice*, 171(3/4), 245–256.
- Poole, K. T., & Rosenthal, H. (1984). U.S. presidential elections 1968–80: A spatial analysis. *American Journal of Political Science*, 28(2), 282–312.
- Poole, K. T., & Rosenthal, H. (1997). *Congress: A political-economic history of roll call voting*. New York: Oxford University Press.
- Poole, K. T., & Rosenthal, H. (2007). *Ideology and Congress*. New Brunswick, NJ: Transaction.
- Rathbun, B. C., Kertzer, J. D., Reifler, J., Goren, P., & Scotto, T. J. (2016). Taking foreign policy personally: Personal values and foreign policy attitudes. *International Studies Quarterly*, 60(1), 124–137.
- Rivers, D. (2003). *Identification of multidimensional spatial voting models*. Working paper.
- Rosenthal, H., & Voeten, E. (2004). Analyzing roll calls with perfect spatial voting: France 1946–1958. *American Journal of Political Science*, 48(3), 620–632.
- Rosenthal, H., & Voeten, E. (2007). Measuring legal systems. *Journal of Comparative Economics*, 35(4), 711–728.
- Shepsle, K. A. (1972). The strategy of ambiguity: Uncertainty and electoral competition. *American Political Science Review*, 66(2), 555–568.
- Silberman, J. I., & Durden, G. C. (1976). Determining legislative preferences on the minimum wage: An economic approach. *Journal of Political Economy*, 84(2), 317–329.
- Smidt, C. D. (2017). Polarization and the decline of the American floating voter. *American Journal of Political Science*, 61(2), 365–381.
- Sniderman, P. M., & Bullock, J. (2004). A consistency theory of public opinion and political choice: The hypothesis of menu dependence. In W. E. Saris & P. M. Sniderman (Eds.), *Studies in public opinion: Attitudes, nonattitudes, measurement error, and change* (pp. 337–357). Princeton, NJ: Princeton University Press.
- Sohn, Y. (2017). Identification and estimation for multidimensional item response theory: An analysis of roll call votes in the United States Congress. Presented at the 2017 Annual Meeting of the Annual Political Science Association, San Francisco, CA.
- Stimson, J. A. (2004). *Tides of consent: How public opinion shapes American politics*. Cambridge: Cambridge University Press.
- Stokes, D. E. (1963). Spatial models of party competition. *American Political Science Review*, 57(2), 368–377.
- Stone, W. J., & Simas, E. N. (2010). Candidate valence and ideological positions in U.S. House elections. *American Journal of Political Science*, 54(2), 371–388.
- Tahk, A. (2018). Nonparametric ideal-point estimation and inference. *Political Analysis*. <https://doi.org/10.1017/pan.2017.38>.
- Thacker, S. C. (1999). The high politics of IMF lending. *World Politics*, 52(1), 38–75.
- Treier, S., & Hillygus, D. S. (2009). The nature of political ideology in the contemporary electorate. *Public Opinion Quarterly*, 73(4), 679–703.
- van Schuur, W. H. (2011). *Ordinal item response theory: Mokken Scale analysis*. Thousand Oaks, CA: Sage.
- van Schuur, W. H., & Kiers, H. A. (1994). Why factor analysis often is the incorrect model for analyzing bipolar concepts, and what model to use instead. *Applied Psychological Measurement*, 18(2), 97–110.
- Weisberg, H. F. (1974). Dimensionland: An excursion into spaces. *American Journal of Political Science*, 18(4), 743–776.
- Weisberg, H. F. (2005). The structure and effects of moral predispositions in contemporary American politics. *Journal of Politics*, 67(3), 646–668.
- Weisberg, H. F., & Rusk, J. G. (1970). Dimensions of candidate evaluation. *American Political Science Review*, 64(4), 1167–1185.
- Zaller, J., & Feldman, S. (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American Journal of Political Science*, 36(3), 579–616.