

A Nonparametric Bayesian Model for Detecting Differential Item Functioning: An Application to Political Representation in the US

Yuki Shiraito¹, James Lo² and Santiago Olivella³

¹Department of Political Science, University of Michigan, Ann Arbor, MI, USA. E-mail: shiraito@umich.edu

²Department of Political Science and International Relations, University of Southern California, Los Angeles, CA, USA. E-mail: lojames@usc.edu

³Department of Political Science, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. E-mail: olivella@usc.edu

Abstract

A common approach when studying the quality of representation involves comparing the latent preferences of voters and legislators, commonly obtained by fitting an item response theory (IRT) model to a common set of stimuli. Despite being exposed to the same stimuli, voters and legislators may not share a common understanding of how these stimuli map onto their latent preferences, leading to differential item functioning (DIF) and incomparability of estimates. We explore the presence of DIF and incomparability of latent preferences obtained through IRT models by reanalyzing an influential survey dataset, where survey respondents expressed their preferences on roll call votes that U.S. legislators had previously voted on. To do so, we propose defining a Dirichlet process prior over item response functions in standard IRT models. In contrast to typical multistep approaches to detecting DIF, our strategy allows researchers to fit a single model, automatically identifying incomparable subgroups with different mappings from latent traits onto observed responses. We find that although there is a group of voters whose estimated positions can be safely compared to those of legislators, a sizeable share of surveyed voters understand stimuli in fundamentally different ways. Ignoring these issues can lead to incorrect conclusions about the quality of representation.

Keywords: item response theory, nonparametric Bayes, Dirichlet process, differential item functioning, joint scaling

1 Introduction

Measurement models, such as the popular two-parameter item response theory (IRT) model, are commonly used to measure latent social-scientific constructs like political ideology. Such models use observed responses to a common set of stimuli (e.g., congressional bills to be voted on) in order to estimate underlying traits of respondents and mappings from those traits to the responses given (e.g., a “yea” or “nay” vote). Standard applications of these models typically proceed on the assumption that the set of stimuli used to measure constructs of interest are understood equally by all respondents, thus making their answers (and anything we learn from them) comparable. This assumption is commonly known as *measurement invariance*, or *measurement equivalence* (King *et al.* 2004; Stegmüller 2011).

As early as 1980, however, researchers were aware that violations of this assumption were possible. Today, violations of this assumption are commonly referred to as differential item functioning (DIF). In the language of the time, Lord (1980, 212) defined DIF by stating that “if an item has a different item response function for one group than for another, it is clear that the item is biased.”

Since Lord’s description of the problem that DIF poses to measurement, a number of researchers have developed and adopted various techniques to mitigate its effects. Lord (1980, 1977) proposed a general test of joint difference between the item parameters estimates for two groups of respondents in the data. Thissen, Steinberg, and Wainer (1993) build on this work,

Political Analysis (2023)

DOI: [10.1017/pan.2023.1](https://doi.org/10.1017/pan.2023.1)

Corresponding author
Yuki Shiraito

Edited by
Lonna Atkeson

© The Author(s), 2023. Published by Cambridge University Press on behalf of the Society for Political Methodology. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

proposing additional methods for fitting IRT models to a known reference and focal group and then testing for the statistical differences in item parameters between the two groups. This work in *identifying* DIF is complemented by work that attempts to *correct* DIF under very specific circumstances and assumptions, including Aldrich and McKelvey (1977), Hare *et al.* (2015), Jessee (2021), King *et al.* (2004), Poole (1998), and Stegmüller (2011).

In this paper, we propose a model designed to improve measurement when DIF is present. To do so, we rely on Bayesian nonparametrics to flexibly estimate differences in the mappings used by respondents when presented with a common set of items. While we are not the first scholars to combine Bayesian nonparametric techniques (and specifically the Dirichlet process) with IRT models (see, e.g., Jara *et al.* 2011; Miyazaki and Hoshino 2009), to the best of our knowledge, we are the first to do so explicitly with the goal of diagnosing DIF. Our model—which we refer to as the multiple policy space (MPS) model—addresses one specific violation of measurement invariance that is of particular importance in political methodology.

Our model identifies subgroups of respondents who share common item parameter values, and whose positions in a shared latent space can thus safely be compared. Thus, while subgroups in our model will not necessarily be distinct from each other, our model can estimate group-specific latent traits by first learning a sorting of observations across unobserved groups of respondents who share a common understanding of items, and conditioning on these group memberships to carry out the measurement exercise. This is similar in spirit to work done by Lord (1980) and Thissen *et al.* (1993), but a crucial difference in our work is that we do not require researchers to *a priori* specify a set of group memberships of members before testing. Rather, our work offers an automated, model-based approach to discover these group memberships from response patterns alone, which in turn also identifies groups of respondents for whom common latent trait mappings can and cannot be validly compared. In discovering these latent group memberships, we can also distinguish the set of respondents in our data that are comparable on a common latent score (i.e., a liberal-conservative ideological spectrum) from those who think on a different dimension (i.e., a libertarian-authoritarian spectrum).¹

To empirically illustrate our model, we apply it to the estimation of political ideology using a dataset that contains both legislators and voters. Our application is based on the dataset analyzed by Jessee (2016), which contains 32,800 respondents in a survey conducted in 2008 and 550 U.S. Congress members who served in the same year. As we discussed above and will elaborate in the next section, the aim of the MPS model in this application is to identify subsets of the voters and legislators within which IRFs are shared and to measure latent traits within each subset, rather than jointly scaling the actors into a common ideology space or determining whether joint scaling disrupts ideal point estimates or not. In our analysis, we find that the 73% of the voters in the dataset share item parameters with the legislators, whereas the 27% of the voters do not.

Our paper proceeds as follows. First, we introduce the substantive context and dataset of our application, focusing on the work of Jessee (2016). Second, we discuss and motivate the details of our IRT model for dealing with measurement heterogeneity, discussing the role of the Dirichlet process prior—the underlying technology that our proposed model uses to nonparametrically separate respondents into groups. Third, we offer Monte Carlo simulation evidence demonstrating the ability of our model to recover the key parameters of interest. Fourth, we present a substantive application of our model to the debate on the joint scaling of legislators and voters. This debate focuses on the extent to which we can reasonably scale legislators and voters into the

1 Ideal points of people belonging to the same substantive cluster are comparable, assuming that we take the spatial model of voting as our preferred model of political preferences. While it is possible to compare preferences on individual issues of individuals in separate clusters (i.e., opinions on tax cuts), we see no straightforward way to standardize ideal points of individuals in different ideological clusters (i.e., a 0.5 on a liberal-conservative scale versus a -0.5 on a libertarian-authoritarian scale).

same ideological space, which effectively can be reframed as a question regarding the extent to which voters share the same item parameters as legislators. We conclude with some thoughts on potential applications of our approach to dealing with heterogeneity in measurement.

2 Application: Scaling Legislators and Voters

In recent years, a literature extending the canonical two-parameter IRT model to jointly scale legislators and voters using bridging items has emerged (Bafumi and Herron 2010; Hirano *et al.* 2011; Jessee 2012; Saiegh 2015). In such applications, researchers begin with a set of items that legislators have already provided responses to, such as a set of pre-existing roll call votes. Voters on a survey are then provided with the same items and asked for their responses. The responses of the voters and legislators are grouped together and jointly scaled into a common space, providing estimated ideal points of voters and legislators that in theory can then be compared to one another.

In an influential critique of this work, Jessee (2016) argued that this approach did not necessarily guarantee that legislators and voters could jointly be scaled into a common space.² Jessee's core critique was that legislators and voters potentially saw the items and the ideological space differently, even if they were expressing preferences on the same items. Joint scaling effectively constrains the item parameters for those items to be identical for both groups, but does not guarantee that they are actually identical in reality. In the language of the MPS model, Jessee claimed that there were potentially two separate clusters—one for legislators and another for voters—through which DIF can occur.

For Jessee, the question of whether voters and legislators could be jointly scaled was essentially a question of sensitivity analysis. He conceptualized the answer to this question as a binary one—that is, either all voters and legislators could be jointly scaled together, or they could not be. His proposed solution to answer this question was to estimate two separate models for legislators and voters. Jessee then used the legislator item parameters to scale voters in “legislator space,” and the voter item parameters to scale legislators into “voter space.” If these estimates were similar to those obtained via joint scaling, then the results were robust and legislators and voters could be scaled together. The Jessee approach essentially adopts Lord (1980) and Thissen *et al.* (1993) approach for testing for DIF, and adds an extra step by reestimating latent traits for the reference and focal groups conditional on the item parameters of the other group.

Our approach to answering this question differs substantially from Jessee, but it is worth noting that his conception of the problem is a special case of our approach. To answer this question using our model, we can estimate an MPS model where we constrain all of the legislators to share a common set of item parameters, but allow voters to move between clusters. Voters can thus be estimated to share membership in the legislator cluster, or they can split-off into other separate clusters occupied only by voters. This highlights the principal difference between the MPS model and Jessee's approach. Jessee's approach is a sensitivity analysis in the spirit of Lord (1980) that provides a binary Yes/No answer to the question of whether jointly scaling legislators and voters together will change the ideal points estimates meaningfully—that is, it scales voters using the item parameters of the legislators, and legislators using the item parameters of the voters. Substantial deviation in the estimated ideal points between these approaches suggests that voters and legislators cannot be scaled together in a common space. In contrast, the MPS model identifies the subset of voters that can be jointly scaled with legislators, which the Jessee model does not. While two special cases of the MPS model (i.e., either all voters lie share item parameters with the legislators, or none of them do) correspond to potential answers that Jessee's model can provide,

2 A critique of joint scaling by Lewis and Tausanovitch (2013) is conceptually similar to Jessee's critique in sharing concern that parameter values for different groups of respondents differ, but employs a different methodology.

our model can provide intermediate answers—notably, we can identify the number and identity of the voters who share an ideological space with legislators, and voters need not all share a common ideological space with one another.

3 Model Description

Our modeling approach adopts the same group-based definition of DIF previously described by Lord (1980) and Thissen *et al.* (1993). Specifically, we assume that there are subsets of respondents who share the same IRFs, which in turn are different from those used by members of other subsets.

If we knew *a priori* what these groups were (e.g., gender of legislators in legislative voting), correcting/accounting for DIF would be relatively easy, and would amount to conditioning on group membership during the scaling exercise. However, the subsets of respondents for whom items are expected to function in different ways is often not immediately obvious. In such cases, we can use response patterns across items to *estimate* membership into groups of respondents defined by clusters of item parameter values (i.e., of the parameters that define different IRFs). This is the key insight behind our approach, which relies on a Dirichlet process prior for item parameters that allows us to identify collections of individuals for whom IRFs operate similarly without the need to fix memberships or the number of such groups *a priori*.

To this end, we propose a model that addresses DIF violations occurring across groups of respondents. When group membership is held constant across items, we are able to identify sets of respondents who are effectively mapped onto different spaces, but who are guaranteed to be comparable *within* group assignment. Our approach, which we call the *MPS* model, is a latent-variable generalization of the standard nonparametric Dirichlet process mixture regression model (e.g., Hannah, Blei, and Powell 2011).³

With these intuitions in place, we now present our DP-enhanced IRT model, including a discussion of how the Dirichlet process prior can help us address the issue of heterogeneous IRFs, but leave the details of our Bayesian simulation algorithm to the Appendix.

3.1 The Multiple Policy Space Model

Let $y_{i,j} \in \{0, 1\}$ be respondent i 's ($i \in 1, \dots, N$) response on item $j \in 1, \dots, J$. Our two-parameter IRT model defines

$$\begin{aligned}
 y_{i,j} \mid \boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma} &\stackrel{\text{i.i.d.}}{\sim} \mathcal{B} \left(\Phi \left(\boldsymbol{\beta}_{k[i]j}^\top \boldsymbol{\theta}_i - \gamma_{k[i]j} \right) \right), \forall i, j \\
 \boldsymbol{\theta}_i &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_D \left(\mathbf{0}, \boldsymbol{\Lambda}^{-1} \right), \forall i \\
 (\boldsymbol{\beta}_{k,j}, \gamma_k) &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_{D+1} \left(\mathbf{0}, \boldsymbol{\Omega}^{-1} \right), \forall k, j,
 \end{aligned} \tag{1}$$

where $k[i] \in 1, \dots$ is a latent cluster to which respondent i belongs; $\boldsymbol{\theta}_i$ is a vector of latent respondent positions on D -dimensional space; $\boldsymbol{\beta}_{k,j}$ is a vector of cluster-specific item-discrimination parameters; $\gamma_{k,j}$ is a cluster-specific item-difficulty parameter.⁴ Substantively, cluster-specific item parameters reflect the possibility that the IRF is shared by respondents belonging to the same group k but heterogeneous across groups.

To aid in the substantive interpretation of this model, it is helpful to consider the case where we only keep respondents in group $k = k'$, and discard respondents belonging to all other groups. Thus, we are only using the item parameters from the cluster k' , which are common to all respondents in that cluster. Since this is the case, we can discard the cluster indexing altogether,

3 As such, it differs from other uses of the DP prior (DPP), such as that of Kyung, Gill, and Casella (2009) or Trauttmüller, Murr, and Gill (2015), where a DPP is defined as part of a semiparametric model.
 4 $\boldsymbol{\Lambda}$ and $\boldsymbol{\Omega}$ are prior precisions of ideal points and item parameters, respectively, with $\boldsymbol{\Lambda} \equiv \mathbf{I}_D$ for identification purposes.

and the first line of Equation (1) reduces to

$$y_{i,j} | \theta, \beta, \gamma \stackrel{\text{i.i.d.}}{\sim} \mathcal{B} \left(\Phi \left(\beta_j^T \theta_i - \gamma_j \right) \right), \forall i \text{ s.t. } k[i] = k'.$$

This is the standard two-parameter IRT model. Thus, we can summarize our model as follows: if cluster memberships were known, the MPS model is equivalent to taking subsets of respondents by cluster, and scaling each cluster separately using the standard two-parameter IRT model. This implies that even though they are expressing preferences on the same items, respondents in different clusters are mapping the same items onto different latent spaces. Thus, comparisons of θ_i are only meaningful when those θ_i belong to the same cluster (i.e., would have been scaled together in the same IRT model).⁵

Given that we do not observe which observations belong to which clusters, however, we need to define a probabilistic model for the cluster memberships that does not require *a priori* specifying how many clusters respondents can be sorted into. For this, we rely on the Dirichlet process prior.

3.2 Sampling Cluster Memberships Using a Dirichlet Process Mixture

The Dirichlet process is a popular nonparametric Bayesian prior (Ferguson 1973; see also Teh 2010). The basic idea of the Dirichlet process is that any sample of data for which one typically estimates a set of parameters can be split into subgroups of units, letting the data guide discovery of those groups instead of requiring users to pre-specify their number *a priori*. Technically, the Dirichlet process prior allows mixture models to have a potentially infinite number of mixture components, but in general it allows a small number of components to be occupied by observations by penalizing the total number of occupied components. It is known that the number of mixture components is not consistently estimated. Nevertheless, when used for density estimation (Ghosal, Ghosh, and Ramamoorthi 1999) and nonparametric generalized (mixed) linear models (Hannah *et al.* 2011; Kyung *et al.* 2009), Dirichlet process mixture models consistently estimate the density and the mean function, respectively.

We now describe the Dirichlet process mixture of our MPS model.⁶ Let $p_{k'}$ denote the probability that each observation is assigned to cluster k' , for $k' = 1, 2, \dots$, that is, $p_{k'} \equiv \Pr(k[i] = k')$, and let the last line of Equation (1) be the base distribution from which cluster-specific item parameters are drawn. Then, under a DP-mixture model of cluster-specific IRT likelihoods, we have

$$k[i] \stackrel{\text{i.i.d.}}{\sim} \text{Categorical}(\{p_{k'}\}_{k'=1}^{\infty}), \tag{2}$$

$$p_{k'} = \pi_{k'} \prod_{l=1}^{k'-1} (1 - \pi_l), \tag{3}$$

$$\pi_{k'} \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(1, \alpha). \tag{4}$$

Equations (2)–(4) are the key to understanding how the Dirichlet process mixture makes non-parametric estimation possible. At the first step in the data generating process, we assign each observation to one of clusters $k' = 1, 2, \dots$. The assignment probabilities are determined by Equations (3) and (4), which is called the “stick-breaking” process. The origin of the name sheds light on how this process works. When deciding the probability of the first cluster ($k' = 1$), a stick of length 1 is broken at the location determined by the Beta random variable (π_1). The probability that each observation is assigned to the first cluster is set to be the length of the broken stick, π_1 . Next,

5 Item parameters follow a similar logic in the sense that they are only comparable within the same cluster, but not across clusters.
 6 The description of the Dirichlet process here is based on the stick-breaking construction developed by Sethuraman (1994).

we break the remaining stick of length $1 - \pi_1$ again at the place π_2 within the remaining stick. The length of the second broken stick ($\pi_2(1 - \pi_1)$) is used as the probability of each observation being assigned to the second cluster. After setting the assignment probability of the second cluster, we continue to break the remaining stick following the same procedure an infinite number of times. The probabilities produced by the stochastic process vanish as the cluster index increases because the remaining stick becomes shorter every time it is broken. Although we do not fix the maximum number of clusters and allow the number to diverge in theory, the property of the stick-breaking process that causes the probability to quickly shrink toward zero prevents the number of clusters from diverging in practice.⁷

Accordingly, when clusters over which DIF occurs are unobserved (both in membership and in number), we can rely on this probabilistic clustering process over a potentially infinite number of groups. In this context, each cluster k' effectively defines a (potentially) different IRF, which in turn allows us to automatically sort observations into equivalence classes within which measurement invariance is expected to hold, without guaranteeing that observations sorted into *different* clusters will be comparable. Hence, our model partitions respondents across a (potentially infinite) set of multiple policy spaces.

In general, the substantive interpretation of estimated clusters needs to be approached cautiously. While our model is useful for identifying which respondents perceive a common latent space with each other, it will generally *overestimate* the total number of actual (i.e., substantively distinct) clusters in the data (Kyung *et al.* 2009; Womack, Gill, and Casella 2014).⁸ In the MPS model, multiple DP clusters can be thought of as being part of the same substantive group—even if their corresponding item parameters are not exactly the same. What is more, this sub-clustering phenomenon can exacerbate known pathologies of mixture modeling and IRT modeling, such as *label switching* (i.e., invariance with respect to component label permutations) and *additive and multiplicative aliasing* (i.e., invariance with respect to affine transformations of item parameters and ideal points).

Thus, even if all respondents actually belonged to the same cluster k' , we could estimate more than one cluster (denoted here by k'') with the other clusters recovering the transformed set of item parameters $\beta_{k'',j} = (\beta_{k',j}^\top K)$ (where K is an arbitrary rotation matrix). However, we would still be able to see that clusters k' and k'' were similar by examining the correlation between $\beta_{k'}$ and $\beta_{k''}$, as well as the patterns of correlation between these and the item parameters associated with other clusters. When sub-clustering is an issue, two sub-clusters can be thought of as being part of the same substantive cluster if their items are highly correlated, or if they share similar correlation patterns with parameters in other sub-clusters.⁹

Having presented the details of our model, we now present the results of a Monte Carlo simulation that illustrates its ability to accurately partition respondents across clusters and recover the associated item parameters within each cluster.

7 The value of the prior parameter α determines how quickly the probabilities to form a new cluster vanish. For $\alpha = 1$, the Beta distribution in Equation (4) turns out to be the uniform distribution. This is the standard choice in the literature (and is our default option in all results presented here), whereas a smaller (larger) value of α leads to a faster (slower) decrease in the cluster probabilities, depending on the total number of respondents in each cluster. Rather than experiment with defining different values for this hyper-parameter for problems of different sizes, we adopt a fully Bayesian approach and define an Gamma hyper-prior over α ,

$$\alpha \sim \text{Gamma}(a_0, b_0)$$

and learn a posterior distribution over α supported by the data.

8 In the context of DP *mixtures*, this issue arises as a result of multiple components having very similar (though not exactly equal) item parameters. Accordingly, and in contrast to models that rely on DPPs to approximate arbitrary densities (as is the case for DP random-effects models), clusters in DP mixtures can be thought of as proper sub-clusters—partitions that are nested within actual, substantive groupings in the data.

9 Correlations, not being a proper metric, can violate the triangle inequality. Thus, high correlations between any two sets of item parameters do not always guarantee similar patterns of association to the parameters of other clusters.

4 Monte Carlo Simulations

As an initial test of our MPS model, we conduct a Monte Carlo simulation to test the ability of our model to correctly recover our parameters of interest. We simulate a dataset in which $N = 1,000$ respondents provide responses to $J = 200$ binary items. Respondents are randomly assigned to one of three separate clusters with probabilities 0.5, 0.2, and 0.3, respectively. In each cluster, respondent ability parameters and item difficulty and discrimination parameters are all drawn from a standard normal distribution. For starting values, we use k -means clustering to generate initial cluster assignments, and principal components analysis on subsets of the data matrix defined by those cluster assignments for starting ability starting values. Item difficulty and discrimination starting values were generated for each cluster and item by running probit regressions of the observed data on the starting ability parameter values by cluster. We run 1,000 Markov Chain Monte Carlo (MCMC) iterations, discarding the first 500 as burn-in, and keeping only the sample that produces the highest posterior density as the maximum *a posteriori* (MAP) estimate of all parameters and latent variables, to avoid issues associated with label switching.¹⁰

Table 1 shows a cross-tabulation of the simulated versus estimated cluster assignments. The estimation procedure is able to separate the simulated clusters well, in the sense that none of the estimated clusters span multiple simulated clusters. However, we see evidence of the sub-clustering phenomenon discussed earlier. Members of simulated cluster 1, for instance, were split into estimated clusters 3, 7, 9, and 10. Since members of simulated cluster 1 were all generated using the same item parameters, the four estimated clusters that partition them are effectively noisy affine transformations of each other. Thus, we expect that the four sets of estimated item parameters for clusters 3, 7, 9, and 10 will be correlated. Simulated clusters 2 and 3 are similarly split between multiple estimated clusters, and we could expect these parameters to be similarly correlated.

In a real-case application, of course, access to the true underlying cluster memberships is not available. And as we discussed earlier, Dirichlet process mixtures are ideal for capturing the *distribution* of parameters by discretizing their support into an infinite number of sub-clusters. As a result, many of these Dirichlet sub-clusters may share very similar parameter values, effectively representing the same substantive groupings in terms of item functionings. Accordingly, using DP mixtures for diagnosing DIF requires a formal procedure for establishing which sub-clusters belong together by virtue of sharing similar item parameters, and which contain observations that truly differ in their item functionings.

The practical issue of establishing equivalence across groups can be approached from a number of perspectives. For example, researchers could employ pair-wise equivalence tests on the item parameters (see, e.g., Hartman and Hidalgo 2018; Rainey 2014, for illustrations in Political Science), being careful to account for the problems raised by conducting multiple comparisons (e.g., using a Bonferroni-style correction, or the Benjamini–Hochberg procedure to control the false discovery rate). Given the potentially large number of pairings, however, we rely on an alternative approach that studies the second and third order information contained in the item parameter correlation matrix. Specifically, we study the graph induced by correlations across entire vectors of estimated item parameters to reconstruct substantive clusters from the

10 An anonymous reviewer pointed out that it would be useful to obtain information about uncertainty by keeping MCMC iterations rather than using one iteration as the MAP estimate. While we agree with this, there are some technical difficulties. First, keeping cluster assignments for all MCMC iterations requires a large memory size, especially because recent datasets for ideal point estimation contain a massive number of respondents (Imai, Lo, and Olmsted 2016). Moreover, since our model is a mixture model, label switching across iterations may mislead uncertainty measures. By contrast, since cluster assignments are discrete random variables, getting uncertainty might not provide much additional information. However, we are aware of the possibility that for some summary statistics of cluster assignments, for example, the minimum proportion of voters who are in the same cluster as the legislators, uncertainty can be computed as an interval on the continuous scale. While we did not use such measures in our application, these could be of interest in other contexts.

Table 1. Simulated versus estimated clusters, MPS model. The estimated clusters recover the simulated clusters, but the sub-clustering phenomenon results in multiple estimated versions of the same cluster. For example, estimated clusters 2 and 4 represent two different ways to identify the simulated cluster 2.

Estimated cluster	Simulated cluster		
	1	2	3
1	0	0	74
2	0	110	0
3	99	0	0
4	0	99	0
5	0	0	79
6	0	0	63
7	139	0	0
8	0	93	0
9	118	0	0
10	126	0	0

sub-clusters identified through the DP mixture, and encourage applied researchers to follow the same approach.

To do so, we treat correlations among parameters as the adjacency matrix of a weighted, undirected graph defined on the set of sub-clusters. The problem of finding substantive clusters can then be cast as the problem of finding the optimal number of *communities* of sub-clusters on this graph—a problem for which a number of approximate solutions exist (for a succinct review, see Sinclair 2016).

For instance, a simple tool for identifying the optimal number of communities in a network is given by the *Gap Statistic* (Tibshirani, Walther, and Hastie 2001), which compares an average measure of dissimilarity among community members relative to the dissimilarity that would be expected under a null distribution of edge weights emerging from a no-heterogeneity scenario¹¹

$$\text{Gap}(k) = \mathbb{E}_{H_0} [\log(\bar{D}_k)] - \log(\bar{D}_k).$$

The optimal number of communities (i.e., of substantive clusters) can then be established by finding the k^* that maximizes $\text{Gap}(k)$. Figure 1 shows the value of gap statistic for different values of k , suggesting that the correct number of substantive clusters is 3 or 4.

Indeed, Figure 2 shows the result of applying a simple community detection algorithm¹² to the graphs formed by using correlations across discriminations (left panel) and correlations across difficulties (right panel). In both instances, the true simulated clusters are denoted using shapes for the graph nodes, and the substantive groupings discovered by the community detection algorithm are denoted using shaded areas. In all instances, the communities identified map perfectly onto the known simulation clusters.

While our previous analyses tested the correspondence between the true and estimated clusters, they say little about the recovery of the correct item parameters. In Figure 3, we explore the item discrimination parameters in a series of plots, where each panel plots two sets of item

¹¹ Implementations can vary with respect to the way dissimilarity is operationalized and with respect to how the null distribution is defined.

¹² Given the small number of sub-clusters in our estimation, we use a greedy procedure that starts by assigning each sub-cluster to its own community, and then proceeds to bind them together while locally optimizing a measure of *modularity*—the extent to which edge density is higher within communities than it is between them (Newman 2003).




Figure 1. Gap statistic over different numbers of substantive clusters, defined as communities in a graph of item parameter correlations. High values of the gap statistic indicate a grouping with high within-cluster similarity relative to a null model (in which edges are drawn uniformly at random) with no heterogeneity. Thus, the k that maximizes the gap statistic is a reasonable estimate for the number of substantive clusters in the data.




Figure 2. Graphs defined on nodes given by DP mixture sub-clusters. The graph has weighted edges defined using pair-wise correlations between discrimination parameters (left panel) and difficulty parameters (right panel). True simulation clusters are denoted with different node shapes, and communities detected by a modularity-maximizing algorithm are denoted with shaded regions. Recovery of simulated clusters is exact in both instances.

discrimination parameters against each other. Along the main diagonal, we plot combinations of the simulated item discrimination parameters (columns) for each cluster against the estimated parameters (rows) for the corresponding known cluster. In all three cases, the item parameters are well recovered and estimates are highly correlated with truth, with correlations of $r = 0.99$, $r = 0.97$, and $r = 0.97$ for the three plots.¹³

¹³ In all cases, and because of the identification problems discussed earlier, estimates are only identified to an affine transformation of the true parameters. We therefore rotate all estimated parameters so that they match their known signs under the correspondence in Table 1.




Figure 3. Correlation of item discrimination parameters. Main diagonal plots estimated versus simulated parameters for each cluster and show that the item discrimination parameters are correctly recovered to an affine transformation. Off-diagonal plots show cross-cluster correlation between estimated and true item parameters, which is expected (under the simulation) to be zero.

In turn, the off-diagonal panels present each combination of the *simulated* item discrimination parameters versus their (mis-matched) counterparts in other clusters. Since parameters in each cluster were generated from independent draws, the items are uncorrelated in reality. As expected, this independence is reflected in the estimated item parameters, which appear similarly uncorrelated with one another and with parameters in other known clusters.

We repeat the same exercise in Figure 4, but this time for the latent traits. In all cases, the latent traits are highly correlated, again demonstrating correct recovery of the traits of interest. The figures also highlight the fact that, in the MPS model, estimated latent traits are only comparable to other respondents belonging to the same cluster. If the MPS model facilitated comparisons across clusters, then at a minimum all of the figures shown here would consistently either be positively or negatively correlated with the simulated true ideal point. However, this is not the case. This is of course not surprising—the MPS model effectively estimates a separate two-parameter IRT model for each cluster of legislators, allowing the same items to assume different item parameters for each group. Thus, ideal points across groups would not be comparable, any more than ideal points from separate IRT models would be comparable. Of course, the MPS model makes a significant innovation in this regard—it allows us to use the data itself to sort respondents into clusters, rather than forcing the researcher to split the sample *a priori*.

Notably, standard measures of model fit also suggests that the MPS model fits the data better in the Monte Carlo. The MPS model produced a log-likelihood of $-85,776.71$, but when we fit the

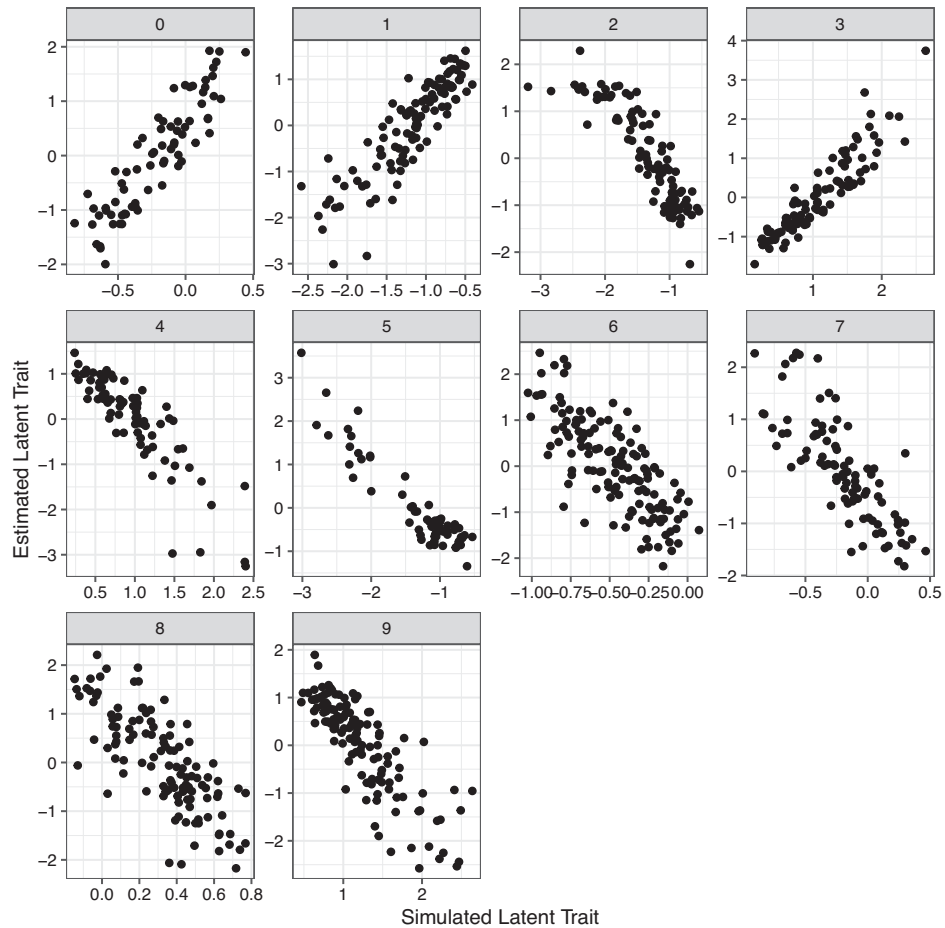


Figure 4. Correlation of latent traits parameters. Plots show simulated against estimated latent traits for all 10 estimated clusters.

standard IRT model on the data that constrains all legislators to share the same single cluster, the log-likelihood drops significantly to $-117,477.2$. This improvement in fit is not surprising—compared to standard two-parameter IRT, MPS fits a much more flexible model. Whereas the standard, single cluster model involves estimating 1,000 respondent and 400 item parameters for a total of 1,400 parameters, the MPS model estimates 1,000 respondent parameters and 400 item parameters *per cluster*. Since the maximum number of clusters in the estimation is set to 10, effectively the MPS model estimates 5,000 total parameters. Thus, a better measure of fit would penalize MPS for the added flexibility afforded by the substantial increase in parameters. The Bayesian Information Criterion (BIC) offers one such measure. It is equal to 252,043 for the single cluster model and for 232,604.7 the MPS model, which confirms that the MPS model fits the data better—even after accounting for the substantial increase in model flexibility. Note that this BIC test is essentially a test of DIF across the identified clusters using methods similar in spirit to those proposed by Lord (1980) and Thissen *et al.* (1993).

Finally, it is important to note that while MPS will partition observations into sub-clusters even when there is no underlying heterogeneity (i.e., even when the standard IRT model is correct), the similarity of item parameters across sub-clusters will immediately suggest that the resulting partition is substantively spurious. To see this, consider Figure 5, which depicts the values of the gap statistic as computed on a graph defined as those in Figure 3, but resulting from a model estimated on data that has no underlying heterogeneity in IRFs. The gap statistic correctly suggests that the correct number of substantive clusters is, in fact, 1. The idea that there is no heterogeneity is further supported by the fact, under such a data-generating process, the standard IRT model




Figure 5. Gap statistic. Statistic defined over different numbers of substantive clusters, when true Data Generating Process (DGP) has no heterogeneity. In this case, the gap statistic again recommends the correct number of clusters—one, in this case.

with a single cluster fits the data better, with $BIC_{IDEAL} = 168,430.8$ versus $BIC_{MPS} = 173,686.3$. Thus, there is little evidence that MPS will overfit data when there is no heterogeneity to be identified.

We now turn to our original motivating application: evaluating whether (or rather *which*) U.S. voters can be scaled on the same space as their legislators.

5 Empirical Results

We apply the MPS model to one of the main examples used in Jessee (2016)—the 2008 Cooperative Congressional Election Study (CCES). This is an online sample of 32,800 survey respondents from the YouGov/Polimetrix panel administered during October and November 2008. In total, the CCES included eight bridging items that directly corresponded to votes taken during the 110th House and Senate, which can be matched to 550 legislators.¹⁴ The policy items included withdrawing troops from Iraq within 180 days, increasing the minimum wage, federal funding of stem cell research, warrantless eavesdropping of terrorist suspects, health insurance for low earners, foreclosure assistance, extension of free trade to Peru and Colombia, and the 2008 bank bailout bill.¹⁵ In this example, Jessee found that joint scaling appeared to work relatively well for this dataset—that is, the ideal points from the grouped model look relatively similar regardless of whether one uses item parameters derived from respondents, the House, or the Senate.

We run 110,000 MCMC iterations, discarding the first 10,000 as burn-in, and keeping only the MAP estimate of the parameters of interest. The maximum number of clusters is constrained to be 10. Similar to the Monte Carlo, we generate starting ideal point values using principal components analysis within each cluster, and probit regression for starting item parameter values. However, rather than generating initial cluster assignments using *k*-means clustering, we instead start all legislators in one cluster, and all voters in a second cluster. Legislators are constrained to remain in the same cluster throughout each iteration, but voters are permitted to change cluster memberships.¹⁶

Table 2 shows a cross-tabulation of the final estimated clusters on the rows against the two separate starting clusters for the legislators and voters. All 550 legislators start in the same cluster, and are constrained to remain so (although their ideal points within the cluster are permitted to change). In turn, the 32,800 surveyed voters divide themselves across six different clusters, with 15,732 respondents remaining in the same cluster as the legislators.

¹⁴ We lose two legislators who recorded no votes on any of the items under study.

¹⁵ The example here makes the same assumption that all joint scaling papers make—that legislators and voters understand the roll call item in a consistent manner. This is known to not literally be true—see Hill and Huber (2019) for research on how legislators and voters may understand even the same roll call vote differently. However, for the purpose of detecting DIF with our model, we adopt and focus on the “common understanding of items” assumption that is prevalent throughout this literature.

¹⁶ This constraint fits the substantive question (i.e., identifying which voters move into a cluster occupied by all legislators), but we acknowledge that for other substantive questions, it may be appropriate to set other constraints. For example, one could separate Southern and Northern Democrat legislators into separate fixed clusters and allow voters to move into those clusters.

Table 2. Estimated versus starting clusters. Legislators all started in cluster 1, and remained there throughout estimation.

Estimated cluster	Legislator starting cluster	Voter starting cluster
1	550	15,732
2	0	8,256
3	0	7,469
4	0	17
5	0	114
6	0	964

Table 3. Correlations of item discrimination parameters between estimated CCES 2008 clusters. Standard errors in parenthesis.

Estimated cluster	Estimated cluster					
	1	2	3	4	5	6
1						
2	0.76 (0.27)					
3	-0.43 (0.37)	-0.14 (0.40)				
4	0.13 (0.40)	-0.10 (0.41)	-0.80 (0.25)			
5	-0.75 (0.27)	-0.62 (0.32)	0.37 (0.38)	-0.41 (0.37)		
6	-0.13 (0.40)	-0.00 (0.41)	-0.49 (0.36)	0.32 (0.39)	0.33 (0.39)	

The 15,732 respondents estimated to share the same cluster with the legislators are almost certainly underestimated, due to the fact that different clusters in DP-prior models may nevertheless share similar parameter values. Table 3 explores this further, tabulating the correlations of the item discrimination parameters between each of the six populated estimated clusters. From examining this table, we see that estimated clusters 2 and 5 have item parameters that are highly correlated with those in the constrained legislator cluster. Combining respondents from clusters 1, 2, and 5 together, 24,102 of the 32,800 respondents in the CCES sample, or approximately 73% of the sample, lie in the same ideological space as legislators.

With this large number of observations falling in a single cluster, it is not surprising that different model selection criteria provide different indications as to whether a standard IRT or MPS fits the data better. For instance, while the comparison between the BIC produced by our model (viz., 408,016.4) and the BIC produced by a standard IRT model (viz., 407,033.7) would suggest the latter offers a better fit to these data, the evidence is reversed when we consider Akaike Information Criterion (AIC) as a selection criterion (with values of 355,419.4 and 370,214.8 for MPS and the regular IRT, respectively.). Nevertheless, an evaluation of the extent to which communities of sub-clusters emerge from these pair-wise correlations suggests the importance of separating between two sets of voters.

The right panel of Figure 6 depicts this correlation-weighted graph, along with the substantive clusters identified by the same greedy algorithm used in the previous section (indicated using gray shaded areas). In this case, both the greedy community-detection procedure and the gap statistic (depicted on the left panel of Figure 6) identify two communities—one containing all legislators and a large number of voters, and another composed of the remaining voters who *do not* share the same policy space as legislators.




Figure 6. (Left) Gap statistic. (Right) Graph on nodes given by DP mixture sub-cluster. Left panel shows two substantive clusters appear to fit the data best. Right panel graph has weighted edges defined using pairwise correlations between discrimination parameters in a model estimated on the 2008 CCES data. Shaded regions denote communities detected by a modularity-maximizing algorithm. Again, two substantive clusters appear summarize the data best, with a “legislator cluster” formed by sub-clusters 1, 2, and 5.

To further validate this sorting, we study the extent to which a model that forces all voters in sub-clusters 1, 2, and 5 to remain fixed in the cluster containing all legislators results in a better fit to the observed responses. Such a model results in an unequivocally better fit versus a model that allows all voters to be freely allocated to clusters, with a BIC of 407,426.8 and an AIC of 365,820.8.¹⁷

In addition, and to explore the question of what characterizes the 24,102 survey respondents who “think like a legislator” (i.e., who are sorted into estimated clusters 1, 2, and 5), we group these respondents together and predict membership in this pseudo-legislator group with a Bayesian binomial probit regression (with vague, uniform priors), using a range of standard covariates—including education, gender, age, income, race, party identification, political interest, and church attendance. We report these results in Figure 7.¹⁸

We find that older voters and people who express more interest in politics tend to map their latent traits onto observed responses similarly to the way legislators do, while Black and Hispanic voters are less likely than their white counterparts to share an ideological space with legislators. And while the coefficients associated with education, income and gender all fail to attain our chosen level of significance, their signs do indicate that more educated and richer voters also tend to think more like legislators, while women appear less likely to share the policy space of their (mostly male) legislative counterparts.

Overall, our findings are largely consistent with Jessee, who found that latent trait estimates from this dataset were consistent regardless of whether one used the item parameters estimated from legislators or voters. However, the key difference from our approach is that we not only identify the 73% of survey respondents who follow this pattern, but also the 27% of survey respondents that do not share an ideological space with legislators. Furthermore, our improved fit statistics suggests that the improvement in model fit for this subset of respondents is quite significant, even for a dataset where the recovered ideal points would be somewhat similar regardless of whether one used only the voter, House, or Senate item parameters to generate ideal points.

¹⁷ In turn, a model that only fixes the membership of the 15,732 voters who are estimated to be in cluster 1 results in a BIC of 407,623.5 and an AIC of 368,304.6, again indicating a worse fit than a model in which everyone in clusters 1, 2, and 5 are fixed from the beginning.

¹⁸ We fit our model using R function `MCMCPack::BayesProbit()`, in package version 1.6-3. We take 9,000 samples from the posterior, having discarded the first 1,000 samples as burn-in.




Figure 7. Point estimates and 90% credible intervals for coefficients in Bayesian probit regression of membership into estimated legislator cluster. A reference line is added at zero. We find that “political interest,” “race,” and “age” are likely to be characteristic of voters in the legislator cluster.

6 Conclusion

When implementing commonly used measurement models, most researchers implicitly subscribe to the idea that all individuals share a common understanding of how their latent traits map onto the set of observed responses: legislators are believed to have shared sense of where the cut-point between voting alternatives lies, survey respondents are assumed to ascribe a common meaning to the scales presented in the questions they confront, and voters are understood to perceive the same candidates and parties as taking on similar ideological positions.

When this assumption is violated by the real data-generating process, however, adopting this widespread strategy can be a costly over-simplification that results in invalid measures of the characteristics of interest. By assuming that units can be separated into groups for whom comparable item functioning holds, we propose a modeling strategy that relaxes the stringent measurement invariance assumption, allowing researchers to identify sets of incomparable units who can be mapped onto multiple latent spaces. The distinctive feature of our proposed approach is that it does not require *a priori* identification of group memberships—or even a prior specification of the number of heterogeneous groups present in the sample.

On this note, it is important to reiterate that the clusters we obtain from our Dirichlet process prior models are not distinct groups, in the sense that they may share parameters that are similar enough to be considered part of the same sub-population. Our models, therefore, are designed to account for the existence of these heterogeneous groups without directly identifying *a posteriori* memberships into them. In so doing, our models assume that the target of inference is the latent traits, rather than the group memberships. And while it is sometimes possible to tease out sub-populations from estimated Dirichlet process clusters, we generally discourage users from trying to ascribe substantive meaning to the clusters directly identified by our nonparametric model—except to say that observations that are estimated to be in the same Dirichlet process cluster have latent traits that can be safely compared to one another. If a more thorough interpretation of which sub-clusters are, in fact, substantively equivalent is of interest, we encourage researchers

to post-process the Dirichlet mixture clusters in order to identify the more substantive groupings defined by item parameters that are similar enough, as we did through the use of the gap statistic on the graph of item parameter correlations in our illustration of the MPS model.¹⁹ Having done so, researchers can then make data-driven decisions about the presence and pervasiveness of DIF in their data. Alternatively, design-based solutions (such as anchoring vignettes) can help ascribe meaning to different subgroups, while other model-based approaches—such as the product partition DP-prior model proposed by Womack *et al.* (2014), or the repulsive DP-mixture model proposed by Xie and Xu (2020)—may offer potential analytical avenues, if adapted to the IRT framework. We leave these possibilities for future research.

Despite these caveats, we believe our proposed model can offer researchers a simple alternative to the standard modeling approach and its strong invariance assumptions. If heterogeneity in item functioning is a possibility—as we suspect is often the case in the social science contexts in which probabilistic measurement tools are usually deployed—our approach offers applied researchers the opportunity to assess that possibility and identify differences across units if said differences are supported by the data, rather than simply assuming those differences across sub-populations away.

A broader substantive question that this paper does not address directly is whether our empirical results hold for joint scaling of legislators and voters using different datasets and/or in other contexts. While we found that most voters share an ideological space with legislators in the CCES dataset, it is still an open question whether most voters and legislators can be jointly scaled particularly when there are a greater number of bridging items that provide more information about how similar their IRFs are. Having presented the methodology that allows researchers to address this question, we leave it for future research.

A. Computational Details

Gibbs Sampler.

Truncate the stick-breaking process at some constant K . Define

1. Update the stick-breaking weight $\pi_{k'}$ for $k' = 1, \dots, K - 1$ by sampling from a Beta distribution s.t.

$$\pi_{k'} \sim \text{Beta}\left(1 + N_{k'}, \alpha + \sum_{l=k'+1}^K N_l\right),$$

where N_k is the number of observations assigned to cluster k under the current state.

2. Update $k[i] \in \{1, \dots, K\}$ for $i = 1, \dots, N$ by multinomial sampling with

$$\Pr(k[i] = k' \mid \mathbf{y}_i, \boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}) \propto \rho_{k'} \Pr(\mathbf{y}_i \mid \boldsymbol{\theta}_i, \boldsymbol{\beta}_{k'}, \boldsymbol{\gamma}_{k'}),$$

where

$$\rho_{k'} \equiv \pi_{k'} \prod_{l=1}^{k'-1} (1 - \pi_l),$$

$$\Pr(\mathbf{y}_i \mid \boldsymbol{\theta}_i, \boldsymbol{\beta}_{k'}, \boldsymbol{\gamma}_{k'}) = \left(\Phi\left(\boldsymbol{\beta}_{k',j} \boldsymbol{\theta}_i - \boldsymbol{\gamma}_{k',j}\right) \right)^{y_{ij}} \left(1 - \Phi\left(\boldsymbol{\beta}_{k',j} \boldsymbol{\theta}_i - \boldsymbol{\gamma}_{k',j}\right) \right)^{1-y_{ij}}.$$

¹⁹ The meaning of “similar enough” is, of course, a matter of researcher discretion. In our illustration, we relied on the gap statistic and community detection tools defined on the correlation graph of item discriminations. Alternative approaches that make the notion of sufficient similarity more explicit could rely on equivalence tests, as they require the definition of a clear equivalence range.

In practice, we augment the latent variable $y_{i,j}^*$, so that we have

$$\Pr(k[i] = k' \mid \mathbf{y}_i^*, \boldsymbol{\theta}_i, \boldsymbol{\beta}_{k'}, \boldsymbol{\gamma}_{k'}) \propto p_{k'} \mathcal{N}(y_{i,j}^* \mid \boldsymbol{\beta}_{k',j}^\top \boldsymbol{\theta}_i - \gamma_{k',j}, 1).$$

3. Conditional on $\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}$, and \mathbf{k} , sample

$$y_{i,j}^* \sim \begin{cases} \mathcal{N}(\theta_i \beta_{k',j} - \gamma_{k',j}, 1) \mathcal{I}(y_{i,j}^* < 0), & \text{if } y_{i,j} = 0, \\ \mathcal{N}(\theta_i \beta_{k',j} - \gamma_{k',j}, 1) \mathcal{I}(y_{i,j}^* \geq 0), & \text{if } y_{i,j} = 1, \end{cases}$$

which can be parallelized over respondents and items, for dramatic speedups.

4. Conditional on $\boldsymbol{\theta}, \mathbf{y}^*$, and \mathbf{k} , sample

$$(\boldsymbol{\beta}_{k',j}, \gamma_{k',j}) \sim \mathcal{N}_{D+1}(\boldsymbol{\mu}_{k',j}, \mathbf{M}_{k',j}^{-1}),$$

where $\mathbf{M}_{k',j} = (\mathbf{X}_{k'}^\top \mathbf{X}_{k'} + \boldsymbol{\Omega})$; $\boldsymbol{\mu}_{k',j} = \mathbf{M}_{k',j}^{-1} \mathbf{X}_{k'}^\top \mathbf{y}_{k',j}^*$; $\mathbf{X}_{k'}$ is a matrix with typical row given by $\mathbf{x}_i = [\boldsymbol{\theta}_i, -1]$ for i s.t. $k[i] = k'$, and $\mathbf{y}_{k',j}^*$ is a vector with typical element $y_{i,j}^*$, again restricted to i s.t. $k[i] = k'$.

Once again, this can be parallelized over items and clusters, reducing user computation times.

5. Conditional on $\boldsymbol{\beta}, \boldsymbol{\gamma}$, and \mathbf{k} , and for each i s.t. $k[i] = k'$, sample

$$\boldsymbol{\theta}_i \sim \mathcal{N}_D(\boldsymbol{\nu}_{k'}, \mathbf{N}_{k'}^{-1}),$$

where $\mathbf{N}_{k'} = (\mathbf{B}_{k'}^\top \mathbf{B}_{k'} + \boldsymbol{\Lambda})$; $\boldsymbol{\nu}_{k'} = \mathbf{N}_{k'}^{-1} \mathbf{B}_{k'}^\top \mathbf{w}_i$; $\mathbf{B}_{k'} = [\boldsymbol{\beta}_{k',1}, \dots, \boldsymbol{\beta}_{k',j}]^\top$ is a $J \times D$ matrix, and $\mathbf{w}_i = \mathbf{y}_i^* + \boldsymbol{\gamma}_{k'}$ is a $J \times 1$ vector. We parallelize these computations over respondents.

6. Finally, conditional on cluster assignments and stick-breaking weights, sample

$$\alpha \sim \text{Gamma}(a_0 + N - 1, b_0 - \sum_{k'=1}^{N-1} \log(1 - \pi_{k'})).$$

Acknowledgment

We would like to thank Kevin Quinn, Iain Osgood, participants in the 2019 Asian PolMeth conference and at the UCLA Political Science Methods Workshop, and two anonymous reviewers for their useful feedback.

Data Availability Statement

Replication materials are available in Shiraito, Lo, and Olivella (2022).

Conflict of Interest

The authors have no conflicts of interest to declare. All co-authors have seen and agree with the contents of the manuscript and there is no financial interest to report.

References

- Aldrich, J. H., and R. D. McKelvey. 1977. "A Method of Scaling with Applications to the 1968 and 1972 Presidential Elections." *American Political Science Review* 71 (1): 111–130.
- Bafumi, J., and M. C. Herron. 2010. "Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress." *American Political Science Review* 104 (3): 519–542.
- Ferguson, T. S. 1973. "A Bayesian Analysis of Some Nonparametric Problems." *The Annals of Statistics* 1 (2): 209–230.

- Ghosal, S., J. K. Ghosh, and R. Ramamoorthi. 1999. "Posterior Consistency of Dirichlet Mixtures in Density Estimation." *The Annals of Statistics* 27 (1): 143–158.
- Hannah, L. A., D. M. Blei, and W. B. Powell. 2011. "Dirichlet Process Mixtures of Generalized Linear Models." *Journal of Machine Learning Research* 12: 1923–1953.
- Hare, C., D. A. Armstrong, R. Bakker, R. Carroll, and K. T. Poole. 2015. "Using Bayesian Aldrich–McKelvey Scaling to Study Citizens' Ideological Preferences and Perceptions." *American Journal of Political Science* 59 (3): 759–774.
- Hartman, E., and F. D. Hidalgo. 2018. "An Equivalence Approach to Balance and Placebo Tests." *American Journal of Political Science* 62 (4): 1000–1013.
- Hill, S. J., and G. A. Huber. 2019. "On the Meaning of Survey Reports of Roll-Call 'Votes?'" *American Journal of Political Science* 63 (3): 611–625.
- Hirano, S., K. Imai, Y. Shiraito, and M. Taniguchi. 2011. "Policy Positions in Mixed Member Electoral Systems: Evidence from Japan." Unpublished manuscript. <https://imai.fas.harvard.edu/research/files/japan.pdf>.
- Imai, K., J. Lo, and J. Olmsted. 2016. "Fast Estimation of Ideal Points with Massive Data." *American Political Science Review* 110 (4): 631–656.
- Jara, A., T. E. Hanson, F. A. Quintana, P. Müller, and G. L. Rosner. 2011. "DPpackage: Bayesian Semi- and Nonparametric Modeling in R." *Journal of Statistical Software* 40 (5): 1–30.
- Jessee, S. A. 2012. *Ideology and Spatial Voting in American Elections*. Cambridge: Cambridge University Press.
- Jessee, S. A. 2016. "(How) Can We Estimate the Ideology of Citizens and Political Elites on the Same Scale?" *American Journal of Political Science* 60 (4): 1108–1124.
- Jessee, S. A. 2021. "Estimating Individuals' Political Perceptions while Adjusting for Differential Item Function." *Political Analysis* 29: 1–18.
- King, G., C. J. Murray, J. A. Salomon, and A. Tandon. 2004. "Enhancing the Validity and Cross-Cultural Comparability of Measurement in Survey Research." *American Political Science Review* 98 (1): 191–207.
- Kyung, M., J. Gill, and G. Casella. 2009. "Characterizing the Variance Improvement in Linear Dirichlet Random Effects Models." *Statistics & Probability Letters* 79 (22): 2343–2350.
- Lewis, J., and C. Tausanovitch. 2013. "Has Joint Scaling Solved the Achen Objective to Miller and Stokes?" Unpublished manuscript.
- Lord, F. M. 1977. "A Study of Item Bias, Using Item Characteristic Curve Theory." In *Basic Problems in Cross-Cultural Psychology*, edited by Y. H. Poortinga, 19–29, Swets & Zeitlinger Lisse, The Netherlands.
- Lord, F. M. 1980. *Applications of Item Response Theory to Practical Testing Problems*. Routledge, New York, NY.
- Miyazaki, K., and T. Hoshino. 2009. "A Bayesian Semiparametric Item Response Model with Dirichlet Process Priors." *Psychometrika* 74 (3): 375–393.
- Newman, M. E. 2003. "The Structure and Function of Complex Networks." *SIAM Review* 45 (2): 167–256.
- Poole, K. T. 1998. "Recovering a Basic Space from a Set of Issue Scales." *American Journal of Political Science* 42 (3): 954–993.
- Rainey, C. 2014. "Arguing for a Negligible Effect." *American Journal of Political Science* 58 (4): 1083–1091.
- Saiegh, S. M. 2015. "Using Joint Scaling Methods to Study Ideology and Representation: Evidence from Latin America." *Political Analysis* 23 (3): 363–384.
- Sethuraman, J. 1994. "A Constructive Definition of Dirichlet Priors." *Statistica Sinica* 4 (2): 639–650.
- Shiraito, Y., J. Lo, and S. Olivella. 2022. "Replication Data for: A Non-Parametric Bayesian Model for Detecting Differential Item Functioning: An Application to Political Representation in the US." Harvard Dataverse. <https://doi.org/10.7910/DVN/BCDALU>
- Sinclair, B. 2016. "Network Structure and Social Outcomes: Network Analysis for Social Science." In *Computational Social Science: Discovery and Prediction*, edited by R. Michael Alvarez, 121–139. Analytical Methods for Social Research. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781316257340.006>
- Stegmüller, D. 2011. "Apples and Oranges? The Problem of Equivalence in Comparative Research." *Political Analysis* 19 (4): 471–487.
- Teh, Y. W. 2010. "Dirichlet Process." In *Encyclopedia of Machine Learning*, edited by Claude Sammut and Geoffrey I, 280–287. Springer, New York, NY.
- Thissen, D., L. Steinberg, and H. Wainer. 1993. "Detection of Differential Item Functioning Using the Parameters of Item Response Models." In *Differential Item Functioning*, edited by P. W. Holland, and H. Wainer, 67–113. Lawrence Erlbaum Associates, Inc, New York, NY.
- Tibshirani, R., G. Walther, and T. Hastie. 2001. "Estimating the Number of Clusters in a Data Set via the Gap Statistic." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63 (2): 411–423.
- Trautmüller, R., A. Murr, and J. Gill. 2015. "Modeling Latent Information in Voting Data with Dirichlet Process Priors." *Political Analysis* 23 (1): 1–20.
- Womack, A., J. Gill, and G. Casella. 2014. "Product Partitioned Dirichlet Process Prior Models for Identifying Substantive Clusters and Fitted Subclusters in Social Science Data." Unpublished manuscript.
- Xie, F., and Y. Xu. 2020. "Bayesian Repulsive Gaussian Mixture Model." *Journal of the American Statistical Association* 115 (529): 187–203.