Beyond the Assembly: Estimating Multidimensional Foreign Policy Preferences from Multi-Modal Data

Evan A. Jones*

Dept. of Government and Politics,
University of Maryland, College Park

Prepared for the 2021 Lussier Conference for Graduate Research on International Conflict,
University of Maryland–College Park

Preliminary draft: comments welcome

Word count: 10177

Abstract

Conventional approaches to estimating latent preferences face numerous constraints. They are limited not only by the dearth of data from which preferences can be learned, usually roll-call votes, but also by the assumptions and limitations built into the statistical models they use to estimate these preferences. Commonly employed models such as DW-Nominate from the legislative studies literature or Bailey, Strezhnev, and Voeten's (2017) dynamic IRT from IR are no exception. Both are limited to roll-call votes and require *a priori* assumptions about the number of latent dimensions that must be validated *a posteriori*. I develop a flexible, non-parametric model that allows researchers to combine multi-modal data, such as speeches and votes, and extract ideal points along as many dimensions as are stably present. I validate the model by examining it's performance on the UNGA voting dating, using Bailey and company's ideal points as a baseline, and find evidence for upwards of three dimension. I also apply the model to combined UNGA votes and debates as well as UNGA votes and Universal Period Review statements, illustrating the model's utility for estimating topic-specific preferences that are still anchored to well-established latent constructs.

*eajones3@umd.edu

1 Introduction

Measures of foreign policy preferences based on United Nations General Assembly (UNGA) votes are ubiquitous in international relations research. Since their early appearance in the literature (Ball 1951; Lijphart 1963; Moon 1985; Vengroff 1976; Russett 1966), UNGA votes have taken on an increasingly important role as proxies for countries' latent preferences. Illustrating this, the current state-of-the-art in this area, Bailey, Strezhnev, and Voeten's (2017) dynamic IRT ideal points, appear in 367 manuscripts based on Google scholar search. This is unsurprising—the UNGA votes are a consistent data source and offer unparalleled temporal and country coverage. However, they are not without their drawbacks.

Our ability to infer latent preferences are constrained by the data and model we use to estimate them. The UNGA is but a single institution through which countries express preferences, and so we struggle to learn about dimensions of foreign policy preferences on issues that do not appear on the agenda, or do so infrequently. While useful for some analysts and research questions, these may be limited or inappropriate for others.

In addition, latent preferences are very much a function of the models used to estimate them. Though Bailey, Strezhnev, and Voeten (2017) overcome the longstanding challenge of discerning agenda changes from preference changes, they do so by imposing significant model constraints on the data. Their estimated ideal points are unidimensional and heavily regularized to ensure smoothness or "stability" over time.[1] The former ignores potential higher-order dimensions of country's preferences, and the latter limits our ability to detect all but large shifts or realignments in preferences such as the end of the Cold War or regime change.

To overcome these issues, I develop a flexible, data and dimension agnostic model for estimating ideal points. The model utilizes Bayesian non-parametric priors and extended rank likelihood (Hoff 2007) in a way that allows researchers to combine multi-modal data, such as speeches and votes, and extract ideal points along as many dimensions as are stably present. I thus call it multi-modal beta process factor analysis (`mmBPFA`). Although here I apply `mmBPFA` to roll call votes (binary) and debate speech (count) data, the model is theoretically unconstrained in the variety of data it can model. So long as the sources of data share at least one stable latent dimension, `mmBPFA` will detect it and provide ideal point estimates.

I validate it against UNGA vote-based ideal points of Bailey, Strezhnev, and Voeten (2017), and then extend it to combinations of votes and speech and votes and Universal Periodic Review (UPR) statements. The model captures the well-established preference dimensions such as the West versus the rest axis identified by 2017 and a major-minor power split similar to the North-South rupture of the 1960-80s identified by Voeten (2000). However, it also finds additional dimensions related to issue framing and agenda-setting when speeches are included as well as human rights-specific dimensions when UPR data is included.

1. The author is only aware of one exception: Bailey and Voeten's 2018 extension of their dynamic IRT model to two dimensions. However, this model is identical to 2017 and is unnecessarily constrained to two dimensions.

In the conclusion, I summarize the strengths and weaknesses of `mmBPFA` for estimating countries' foreign policy ideal points and discuss how it opens up new avenues for IR scholars to extract valuable information from hitherto underutilized data sources. I also highlight avenues for future improvements to the model.

2 Identifying Countries' Foreign Policy Preferences with UNGA Votes

The use of UNGA voting behavior as proxies for foreign policy preferences is for good reason. The United Nations general assembly is one the few international forums in which most countries express their preferences on the same items. The institution's longevity is an additional boon for researchers. As such, estimates of state preference based on this data are ubiquitous—appearing as both independent and dependent variables across a wide array of studies.

Voting behavior has been shown to predict interstate conflict onset (Gartzke 1998; Reed et al. 2008) and dynamics (Wolford 2014; Sweeney 2003), terrorism (Dreher and Gassebner 2008), and the provision of peacekeeping troops (Ward and Dorussen 2016). In the study of IOs, voting patterns correlate with lending in the IMF and World Bank (Thacker 1999; Dreher and Jensen 2007), joining the WTO (Davis and Wilf 2017), and the design of treaties (Koremenos 2005). Unsurprisingly, state's activity at the UN is also informative of their bilateral, predicting diplomatic missions (Neumayer 2008), foreign aid distribution (Alesina and Dollar 2000), compliance with aid agreements (Girod and Tobin 2016), and inclusion of the Chinese renminbi as a reserve currency (Liao and McDowell 2016).

Using changes in UNGA voting patterns as a measure of foreign policy proximity, scholars have shown changes in domestic leadership and regime types influence foreign policy orientations (Dreher and Jensen 2013), the European Union's foreign policy is becoming more consistent (Drieskens 2010), and how Chinese and US aid can buy foreign policy deference (Flores-Macías and Kreps 2013). They have also been used to explore whether the US becoming isolated on foreign policy (Voeten 2004) and convergence in UN member states' interests, more generally (Bearce and Bondanella 2007).

The prevailing measure state preferences using this data are Bailey, Strezhnev, and Voeten's (2017) (henceforth, BSV) dynamic item response theory ideal points. By successfully disentangling changes in the UN agenda from changes in preferences, BSV superseded the S-score (Signorino and Ritter 1999).

2.1 Dynamic IRT Ideal Points

BSV utilize a straight-forward IRT model to estimate country's ideal points but introduce a couple important constraints to overcome the agenda-preference change dilemma. The basic models is given by:

$$\begin{aligned} Z_{it\nu} &= \beta_\nu \theta_{it} + \epsilon_{i\nu} \\ \epsilon_{i\nu} &\sim \mathcal{N}(0,1) \end{aligned} \tag{1}$$

where $Z_{it\nu}$ is a normally-distributed latent variable representing country i's preference at time t on vote ν. The discrimination parameter, or "polarity" of each vote, is given by

β_ν, and θ_{it} denotes a country's ideological location. If β_ν is large, it the vote strongly delineates countries. And the signs of β_ν and θ_{it} tell us the direction of splits. When large, positive θ_{it} countries vote yea, the discrimination parameter will be positive for these votes. When they vote nay, it is negative.

Since the observed vote choices are ordinal—yea, abstain, nay—but latent preferences are assumed to be continuous, BSV utilize Johnson and Albert's (1999, 166) cut-point approach to handle ordinal responses. Thus, letting $Y_{it\nu}$ be the observed choice for a given vote in year t, we have:

$$Y_{it\nu} = \begin{cases} \text{yea} & \text{if } Z_{it\nu} < \gamma_{1\nu} \\ \text{abstain} & \text{if } \gamma_{1\nu} < Z_{it\nu} < \gamma_{2\nu} \\ \text{nay} & \text{if } Z_{it\nu} > \gamma_{2\nu} \end{cases} \tag{2}$$

In other words, when the latent preference is less than the lower cut point $\gamma_{1\nu}$ we observe a yea vote, abstain when it falls between the two cut points, and nay when it is above the upper cut point $\gamma_{1\nu}$.

BSV impose an important restriction on these cut points, known as across-time bridging, whereby resolutions with the same content are forced to have the same cut points within five year rolling windows. This links the latent spaces over time and ensures agenda changes do not manifest as altered preferences. The rolling window relaxes the constraint, allowing for the possibility that political context can change and, thus, the interpretation of bills with the same content can also change.

They also use Bayesian priors on the ideal points so that estimates from the previous year $\theta_{i,t-1}$ inform those in the current period θ_{it}. By restricting the variance on this prior, they enforce a significant degree of smoothing. Estimates do vary across periods, but by a small amount. They set the parameter to detect discrete shifts in preferences such as regime changes, but the degree of smoothing and, hence, the sensitivity of estimates to preference changes is an arbitrary researcher decision.

These constraints are a blessing and a curse—they allow us to extract well-defined, stable ideal point estimates from votes without being tainted by shifts in the agenda, but do so at the expense of expressive power. First, generalizing the model to higher dimensions is problematic because the second dimension may not be stable over time and so you cannot impose consistent priors across time. Second, bridging resolutions cannot be used for higher dimensions since a resolution may have concurrent interpretations that lead it to load on multiple dimensions. Though they can be applied to stable dimensions, stability usually cannot be ensured beyond the first dimension and may not even be desirable if we want to allow higher dimensions to evolve along with the broader IR landscape.

Finally, the question of how much smoothing in estimates is arbitrary and depends on the level of granularity desired by the research. More smoothing ensures stable ideal point estimates that are correlated with macro-level shifts in world politics such as the end of the Cold War or major regime shifts whereas less smoothing will allow ideal points to reflect micro-level shifts in foreign policy preferences. To illustrate the significance of this smoothing, Figure 1 plots the level of autocorrelation in the BSV ideal points.

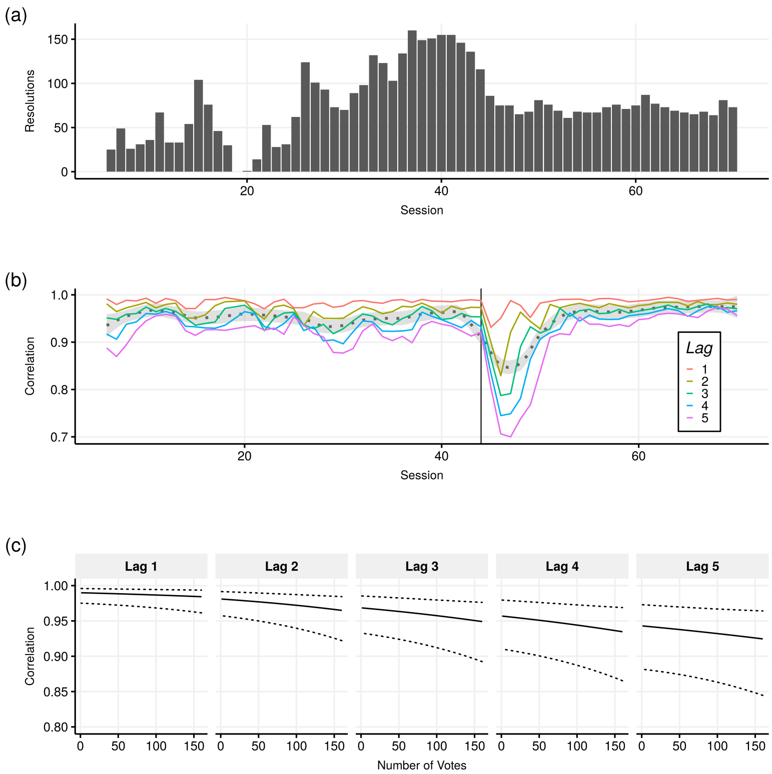

Figure 1: **Influence of Model versus Data in BSV Ideal Points.** (a) Number of resolutions voted on by session. (b) UNGA ideal point autocorrelation. (c) Predicted autocorrelation by number of votes.

The top panel shows the number of votes per session, the middle panel the level of aggregated autocorrelation for up to 5 lag years, and the bottom panel shows the predicted autocorrelation by lag period as a function of the number of votes. The predictions are based on an OLS regression of the Fisher-transformed ideal point autocorrelation contained in Appendix A. In general, as the amount of effective data increases the IRT model should place less emphasis on the prior ideal points and rely more on information contained in the votes from the current period.[2] We see there is no visible nor statistical relationship between the amount of data and autocorrelation. The predicted slopes are nearly flat and only begin to show some decrease as a function of data at 5 years. The year-to-year autocorrelation never drops below 0.9 and the only briefly drops below 0.85 among the longer lags during the post Cold-War transition period.

While there is correct answer as to the right amount of smoothing, the amount induced by BSV's constraints could hide interesting and informative shifts in preferences if they do not cause a country to significantly diverge from its location on the primary

2. By effective data, I mean resolutions that contain unique information pertaining to relevant IR issue areas as well as more members. The underlying assumptions are: 1) the general assembly agenda dynamically changes over time to address emergent issues and, therefore, does not vote on (nearly) identical sets of resolutions year-over-year; 2) increases in the number of resolutions are not due to surges in superfluous votes related to the functioning of the body such as budget or functions resolutions; and 3) larger membership should mean that the weight of the likelihood (current session ideal points) grows relative to the prior (previous session ideal points).

Liberal West-versus-the Rest dimension. Depending on researcher needs, detecting micro-level shifts or those along higher order axes might be desirable.

These constraints and the trade-offs they entail are not inherently problematic. They help solve the problem Bailey, Strezhnev, and Voeten set out to solve. Yet, this paradigm for estimating ideal points does not scale well to higher dimensions, nor does it accommodate the fusion of multiple data sources. The proliferation of data available to social science researchers, especially text (Wilkerson and Casas 2017), far outpaces the tools we have to analyze those sources in a consistent, unified way. Just consider the vast amount of unstructured and semi-structured data—statements, meeting minutes, debates, budgets, etc.—produced by IOs and member states that partially reveal country preferences. They either go unused or, more commonly, are analyzed in isolation and the insights they furnish are thereafter synthesized by the researcher. Such approaches are inherently *ad hoc* and inconsistent. A more principled approach is necessary to allow researchers to move beyond merely using UNGA votes.

2.2 Towards Data Fusion and Multidimensionality

The political methodology literature is rich with methods for estimating ideological preferences from different sources of data (for example, Benoit and Laver 2003; Lauderdale and Herzog 2016; Baturo, Dasandi, and Mikhaylov 2017; Clinton, Jackman, and Rivers 2004). Most of these offer bespoke solutions to specific data types—the two most common being roll call voting and text—and vary in the degree to which they correspond to a formal theoretic foundation.[3] The model I introduce below unites and extends this vein of research by providing a single framework for estimating dimension and data agnostic ideal points. By freeing the researcher from the constraints of a single data type—or certain data configurations compatible with established approaches—and a pre-specified number dimensions, my model provides specific advantages over the dynamic IRT UNGA ideal points.

In terms of data, using UNGA votes is often a case of satisficing: they provide the most expansive temporal and country coverage but are limited by the agenda and mandate of the UN. First, the resolutions that appear on the UNGA docket, like all other legislative bodies, are a function of agenda-setting which is influenced by the distribution of resources and power across member states (Panke 2017). Second, by virtue of the UNGA mandate combined with the increasing fragmentation IOs (Morse and Keohane 2014; Greenhill and Lupu 2017; Alter and Raustiala 2018) and resulting specialization, only certain issues have remained fixtures of the UN agenda such as the Middle East, nuclear disarmament, and human rights whereas many others have been delegated away to the purview of different UN bodies or other IO's altogether. Ideal points estimated from votes in the UNGA are constrained to reflects preferences over this set of issues. In short, variation in learned preferences is only as rich as the data used to estimate those preferences.

Richer data means richer sets of preferences. By enabling researchers to potentially combine a diverse set of sources such as UNGA speeches, bi- and multilateral treaties,

3. See 2004, 356 for a canonical example of the formal link between a theoretical model of respondents' expected utility (ideal point) and a Bayesian IRT model.

and official statements, to name a few, the model broadens the possibilities for discovering state preferences across different domains of international relations. Importantly, these may correspond more closely to their research goals.

In terms of dimensionality, Bailey and Voeten (2018b, 2) note that constraining the dynamic IRT comes at the expense of missing out on "valuable information contained in a higher dimensional policy space." As such, they expand the BSV model to two dimensions and find a second dimension representing a North-South conflict between developed and developing countries. However, the dimension is only stable and present throughout the 1970s and 80s. They also allude to other issue-specific dimensions that are resolution dependent and suggest researchers interested in these should cherry-pick resolution-specific ideal points. While this is an important acknowledgment of multidimensionality, it still unnecessarily assumes a specific number of dimensions *a priori*. The flexibility of my model detects not only the stable dimensions but also higher order, emergent dimensions where present.

It should be noted that, unlike BSV, my model is not dynamic. Country ideal points from previous years are not incorporated as prior information nor do I use bridging resolutions.[4] Nonetheless, as shown later, I discovered a first dimension that is stable, corresponds closely to their first dimension, and is fairly robust to agenda change. As such, I do not posit my model as a direct substitute for BSV's ideal point estimates but rather as a complementary approach that may better suit the needs of some researchers.

3 Multimodal Beta Process Factor Analysis

I estimate multidimensional state preferences by developing a multimodal beta process factor analysis model. At its core, the model is parameterized to be functionally equivalent to the familiar IRT approach. However, some specific modifications greatly enhance its flexibility.

First, as the name implies, the model utilizes a beta process prior (Hjort 1990; Paisley and Carin 2009) to learn the number of active dimensions K^+ in a given session. Second, it foregoes the traditional "cutpoint" approach (Albert and Chib 1993; Quinn 2004) altogether and draws on extended rank likelihood (Hoff 2007) to combine discrete and continuous margins.

The model is presented below, and each of the modifications are discussed in turn. Full mathematical and implementation details for available in Appendix B. Let the observed data be given by:

$$y_{ij} = F_j^{-1}[\Phi(x_{ij})] \tag{3}$$

where $F_j^{-1}(\cdot)$ and $\Phi(\cdot)$ denote the (pseudo) inverse of an unknown univariate CDF and Gaussian CDF, respectively. The data are assumed to be transformations from a latent

4. Theoretically, there is nothing restricting one from making my model dynamic, and this is a potential avenue for future research.

space x_{ij} that is modeled by a Gaussian copula[5]:

$$x_{ij} \sim \begin{cases} \mathcal{TN}_{x_{ij}^l, x_{ij}^u}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is observed} \\ \mathcal{N}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is missing} \end{cases} \tag{4}$$

Three parameters—λ, ω, and α—correspond to the standard IRT: $\boldsymbol{\lambda}_j$ is a vector of factor loadings for variable j and is equivalent to the discrimination parameter in an IRT model; $\boldsymbol{\omega_i}$ is a vector of K^+ ideal points for country i; and α_j is an item-difficulty parameter.

The $\boldsymbol{z}_j$ parameter, which is a binary vector of length K^+, reflects the beta process prior (more below). The $\odot$ denotes element-wise multiplication, and so one way of conceptualizing this parameter is as a regularizer that induces sparsity, i.e. sending some some of the factor loadings to zero. A more accurate description, however, is to consider the entire matrix $\boldsymbol{Z}$ as having an infinite number of columns K of which only some finite number of them are ever "active" K^+. If a variable j is active on dimension k then z_{jk} takes a value of one and λ_{jk} will be sampled.

The parameters x_{ij}^l and x_{ij}^u are lower an upper bounds on the truncated normal distribution. Although they function as cutpoints, they are not estimated as separate parameters but inferred indirectly from the ordering on the observed data or the extended rank likelihood (Hoff 2007). This simple modification turns out to be an elegant way to model discrete and continuous margins without assigning any probability distribution to them.[6]

Implementing a Gaussian copula factor model[7] is not without it's challenges, especially for complex combinations of continuous and discrete variables. First, conditional indpendence of the factors cannot be guaranteed in such cases (Murray et al. 2013) and so sparisity inducing priors need to be carefully implemented and interpreted (Pitt, Chan, and Kohn 2009; Dobra and Lenkoski 2011). Second, the above specification assumes you know the marginal distributions of the observed data F_j. While this may be true in some instances like binary roll call data, this likely does not hold for many applications especially if we aim to model complex combinations. Though the former issue cannot be easily solved in a mathematically principled way, the latter can be.

5. A copula is a multivariate cumulative distribution function where the marginal probability distributions of each variable are treated as uniform over the interal [0, 1]. Though simple in concept, copulas can be used to describe any multivariate joint distribution (Sklar 1959). As a result, they are powerful tools for statistical tasks that require learning latent dependence structures over, potentially, complicated multivariate distributions and find applications in modeling financial risk (Genest, Gendron, and Bourdeau-Brien 2009), missing data imputation (Hollenbach et al. 2018), and factor analysis (Murray et al. 2013).

6. It should be noted that `mmBPFA` is similar to Kim, Londregan, and Ratkovic (2018) since they also employ the extended rank likelihood to place vote and speech data in a latent Gaussian space. Though there are notable differences. They further tailor their model to deal with zero inflation in text data, whereas I do not. Although this improves performance for text data, it constrains the generalizability to mixed data that do not fit the specific votes-text paradigm. Moreover, they employ a Laplacian (LASSO) prior (Park and Casella 2008) to induce *ex post* sparsity on the prior. Whereas `mmBPFA` induces sparsity throughout sampling.

7. The Gaussian copula factor model is extremely general and subsumes many other common multidimensional scaling models such as Quinn (2004) and Albert and Chib's (1993) "cutpoint" probit factor model.

3.1 Combining Arbitrary Multimodal Data

The solution lies in Hoff's (2007) extended rank likelihood. Noticing that since the inverse transformation F_j^{-1} is monotonically increasing, there also exists a weak partial ordering on the latent variable $\boldsymbol{x}_j$. In other words, if an observed value y_{ij} is greater than another observation $y_{i'j}$, then the associated latent variables x_{ij} and $x_{i'j}$ also maintain this ordering. Using these partial orderings, Hoff (2007) and Murray et al. (2013) show can ignore the observed marginal distributions F_j altogether and define the observed $\boldsymbol{Y}$ in terms of partial orderings $D(\boldsymbol{Y})$ and correlation matrix $\boldsymbol{C}$ only:

$$P(\boldsymbol{Y}|\boldsymbol{C}, F_1, \ldots, F_p) = P(\boldsymbol{Z} \in D(\boldsymbol{Y})|\boldsymbol{C}) \tag{5}$$

and still maintain consistency in the dependence structure for most cases.

The set of partial orderings is fixed function of the observed data, and thus needs only be computed once. Let

$$\begin{aligned} D_{lower}(y_{ij}) &= \{y_{i\prime j} : y_{i\prime j} < y_{ij}\} \\ D_{upper}(y_{ij}) &= \{y_{i\prime j} : y_{i\prime j} > y_{ij}\} \end{aligned} \tag{6}$$

then we can use this ordering to find the bounds x_{ij}^l and x_{ij}^u and, thereby, a mapping between the latent and observed spaces.

Beyond the enabling the combination of multimodal margins, the extended rank likelihood also provides computational gains because one no longer needs to the additional step of estimating cutpoints. This can lead to considerable speed ups, especially for discrete data[8]

3.2 Beta Process Prior

The perennial question of the true number of latent clusters plagues not only multidimensional scaling, but all latent hierarchical models such as topic model (Grimmer and Stewart 2013). The conventional answer usually rests on some combination of domain knowledge, statistical knowledge, and guess-and-checking (i.e. fitting and re-fitting models). Bayesian non-parametric priors offer a solution to this issue that is fairly robust to "research degrees of freedom" (Simmons, Nelson, and Simonsohn 2011) and hyperparameter tuning. The beta process prior (Hjort 1990), in particular, is well-suited to learning the most appropriate number of dimensions in a factor analysis setting (Knowles and Ghahramani 2011; Paisley and Carin 2009; McAlister 2020).

The beta process prior follows a beta-Bernoulli generative structure:

$$\begin{aligned} P(\pi_k) &\sim Beta\left(\frac{a}{K}, \frac{b(K-1)}{K}\right) \\ P(z_{jk}|\pi_k) &\sim Bern(\pi_k) \end{aligned} \tag{7}$$

8. Note that although we avoid estimating cutpoints, this still does not guarantee *fast* convergence of our MCMC sampler. For instance, though my model converges quite quickly for discrete margins, for continuous margins there is a 1-to-1 correspondence between the dimensionality of the data and the number of samples that must be drawn. This constrains the model's scalability to some degree.

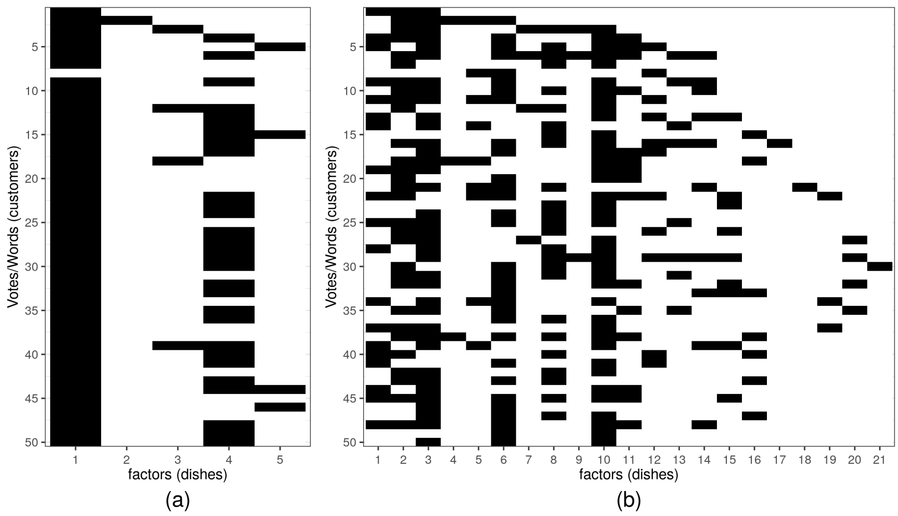

Figure 2: **Draws from a One Parameter IBP for Two Different Values of α.** (a) $\alpha = 2$. (b) $\alpha = 4$.

where a and b are hyper-parameters, and the number of dimensions $K \to \infty$.

There are finite (Doshi-Velez et al. 2009) and infinite (Knowles and Ghahramani 2011) derivations of the beta process in the literature. In the former, you assume a far greater number of dimensions than believed present in the data and the space shrinks over the course of sampling. The latter differs only in that at each step new potential dimensions (up to an infinite amount) are also sampled.

The latter, which I implement, happens to be equivalent to a *two parameter Indian Buffet Process* (IBP) (Ghahramani and Griffiths 2006; Thibaux and Jordan 2007). The IBP earns its quirky name from the analogy use to explain the stochastic process. Imagine customers eating at an Indian buffet.

1. The first customer enters and Indian buffet with an infinite number of dishes.
2. She helps herself to the first Pois(α) dishes.
3. The $j^{th} \in \{1, \ldots, P\}$ customer helps herself to each dish with probability $\frac{m_k}{\beta+P-1}$, where m_k is the number of times dish $k \in \{1, \ldots, \infty\}$ was previously sampled.
4. The j^{th} customer tries $Pois(\frac{\alpha\beta}{\beta+j-1})$ new dishes.

In the factor analysis setting, customers correspond to observed variables and dishes correspond to the latent factors they load on. Two hyper-parameters influence the IBP's behavior: α determines how many dishes the first customer tries, and β sets the *a priori* probability a customer will try a given dish. Figure 2 provides two hypothetical examples of the binary matrix $\boldsymbol{Z}$ for different values of the α parameter. As you can see, larger α values lead to more dimensions sampled, yet a "rich get richer" property is present in that factors which are loaded on initially are more likely to accrue further loadings as the process continues. This results in a factor loading matrix where a most of the variance is explained by a (relatively) small number of "universal" factors followed by an array of more "idiosyncratic" factors.

The beta process offers the advantage of allowing the dimensionality to grow in complexity with the data. Smaller datasets are inherently constrained to lower dimensionality—there is less information available to learn from. When you combine related but not perfectly congruent datasets, such as UNGA votes and speeches, then votes and speeches can be allowed to simultaneously load on shared factors as well as separate factors that are unique to them. In this way, the beta process is a principled approach not only to estimating multidimensionality but also to combining multiple sources of data.

Nonetheless, this approach is not without its downsides. The beta process is very demanding insofar as it behaves better on larger datasets. For instance, the number of active dimensions can collapse to zero or diverge to infinity in generate cases. However, I have only encountered this when the number of variables in the data set is quite small, e.g. $p < 30$. If one has such a small data set, then these methods are likely too demanding.

4 Model Performance

Before moving to more complex, multi-modal forms of data, I first submit `mmBPFA` to numerous tests to check it's fidelity. First, I ensure the ideal points do not change in response to agenda changes since the model does not employ the same constraints as BSV. Second, I examine the utility of moving beyond a single dimension. Third, I check the predictive performance of the model. And finally, I compare it to BSV's single-dimension model and Bailey and Voeten's (2018) later extension to two dimensions.

I limit the tests to votes-only data from the 25^{th} through 72^{nd} sessions since this is the range for which UNGA debate data is also available which I also test later. Additionally, as mentioned previously, because mmBPFA makes fewer assumptions it requires relatively more data to retrieve good results and can become degenerate on very low rank datasets. In general, performance is best when the matrix has rank 100 or greater. Prior to the 25^{th} session, the average number of votes is less than 50. While mmBPFA successfully ran for some of these sessions, in others the model diverged towards an infinite number of dimensions. I therefore am reluctant to accept the validity of any results on so few votes.

4.1 Validity

The key breakthrough made by BSV was their model's ability to discern preference from agenda changes. For mmBPFA to be of any utility in estimating preferences, it must, at a minimum, maintain this capability. Since mmBPFA is neither uses bridging resolutions nor year-to-year Bayesian priors, there is a very real possibility it could be prone to conflating agenda and preference shifts. Fortunately, this appears not to be the case.

I test the dependency of mmBPFA's first dimension ideal points on agenda changes by regressing the ideal points on session-wise issue proportions (Voeten 2013). Higher

Table 1: Lagged-Ideal Point versus Issue Proportion Regressions

	(1) Issue Prop.	(2) Issue Prop. Δ
Lagged-Dependent Variable	0.804***	0.804***
	(0.020)	(0.020)
Middle East	-0.026	-0.062
	(0.090)	(0.112)
Nuclear	-0.005	0.024
	(0.153)	(0.117)
Disarmament	0.015	0.018
	(0.175)	(0.172)
Human Rights	0.011	0.003
	(0.107)	(0.206)
Colonialism	-0.007	-0.005
	(0.201)	(0.120)
Economic	-0.016	-0.013
	(0.108)	(0.214)
Observations	8,052	8,052
R^2	0.6417	0.6418
Adjusted R^2	0.6414	0.6414

One-way (country) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

dimensions are not tested due to their instability, and thus incomparability, over time.[9] The first model uses the static, yearly issue proportions of all votes while the second model uses the first difference, or year-to-year changes, in issue proportions. I cluster standard errors by country in both.

Table 1 shows the results. In neither model do the ideal points respond to observed agenda changes. Indeed, only do the lagged-ideal points show significant correlation (coefficient of 0.804) with the current ideal points. Both models also have an R^2 of 0.64, indicating that although a majority of the variation in current preferences is explained by past preferences, there is still a fair amount of unexplained variation when also controlling for agenda shifts. For comparison, in a similar regression Bailey, Strezhnev, and Voeten (2017) find a coefficient of 0.982 on their lagged-ideal points and an R^2 of 0.969. This again highlights the different degree of constraints imposed by each of these models.

In some institutions—US Congress, for instance—multidimensionality is well-established (McCarty, Poole, and Rosenthal 2001; Kim, Londregan, and Ratkovic 2018; McAlister 2020). In under-explored institutions or novel research areas, however, dimensionality

9. For instance, if there are three stable dimensions year y_t but only two in the following year y_{t+1}, the second dimension in y_{t+1} could correspond to either the second or third dimension from y_t. There is nothing fixing the dimension labels over time. Moreover, the more the higher dimensions are correlated, the more difficult it is to disentangle them *a posteriori*.

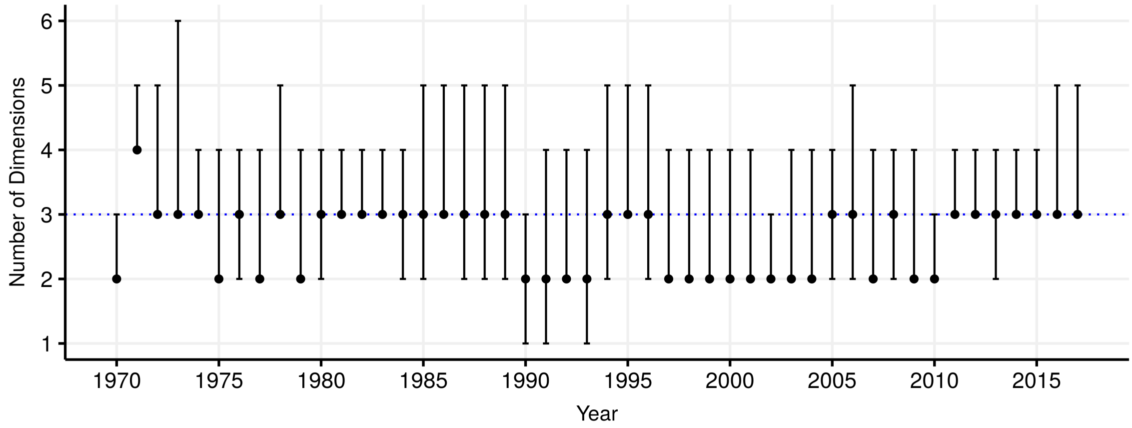

Figure 3: **Dimensionality of the 25th–72nd UN General Assembly.** Points indicate the median posterior number of dimensions mmBPFA found after burn-in. Bars represent 95% HPD intervals. The median number of dimensions across sessions is 3, indicated by the blue dotted-line.

is impossible to know *a priori* as one of the model's benefits is allowing the researcher to learn about preference dimensionality in an agnostic way. In the UNGA, Bailey and Voeten (2018b) find evidence for a plausible, second dimension. I provide further support for the existence of at least two dimensions, and likely more.

Figure 3 plots the median number of preference dimensions found by mmBPFA between 1970 and 2017 in the UNGA votes. The median values are based on 1000 posterior draws. I also plot 95% highest posterior density intervals. The median number of dimensions found across all years is three. In the 1970s there is one session for which mmBPFA finds six possible dimensions in some draws. Notably, in none of the sessions does mmBPFA find only a single dimension as the median value. Contrast this with 2018a who find the second dimension only to be stable from the mid 1960s until the mid 1980s, only to collapse in the post Cold War period, and potentially reemerge after 2010. While I find similar patterns changes in the number of dimensions, I consistently estimate at least one more dimension.

Nonetheless, the second and third dimensions can be quite difficult to distinguish from one another, both substantively (more below) and mathematically. One issue with the discrete Bayesian mixture models, a class of which mmBPFA is a member, is that the dimensions have no inherent, consistent ordering across draws. This "label switching" is a well-known issue (Rodriguez and Walker 2014). I solve it by reordering the dimensions based on their Euclidean distance from a representative "pivot" draw. In practice, the distance between the second and third dimensions is so small as to render them indistinguishable in terms of the resolutions loading on them. When this happens, I retain the dimension with a higher average number of resolutions loaded on it over all draws and drop the other dimension.

This approach has practical implications, namely that I ignore retain and analyze what can be considered a hard, lower bound on the number of dimensions found by mmBPFA. So although it finds three consistent dimensions, in practice I often retain

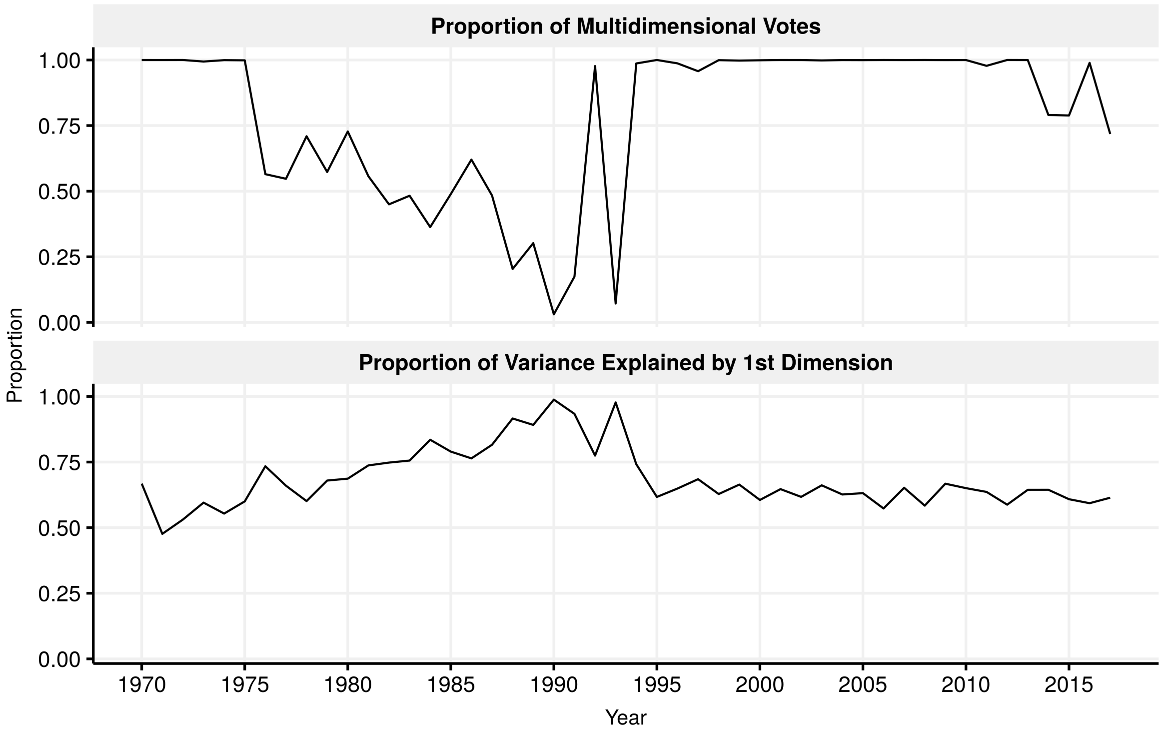

Figure 4: **Importance of Multidimensionality in UNGA Voting.** (top) The proportion of votes loading on more than one dimension per session. (bottom) Proportion of variance in voting behavior captured by the first dimension. *Note: plots are based on posterior means.*

only two and keep a third or fourth when they are distinguishable from each other.

Given that mmBPFA's additional modeling power comes with corresponding increases in computational costs and data constraints, an important question is whether the extensibility to higher dimensions is necessary. Answering this requires some measure of the substantive importance of dimensions beyond the first (Roberts, Smith, and Haptonstahl 2016; Smith 2007). One useful metric in making this determination is the proportion of variables (here votes) that load on more than one dimension. The higher the proportion, the more additional dimensions matter. While this metric provides a baseline heuristic for the relevance of additional dimensions, it suffers some deficiencies, namely it does not tell us the amount of explanatory power captured by dimensions beyond the first. If every vote in a given session loads on two dimensions but the first dimension explains 90% of the variation, then looking at higher dimensions may not be of much utility. Thus, one should also consider the proportion of variation explained (PVE) by each dimension. For a specific variable j, the PVE is given by:

$$PVE_{jk} = \frac{z_{jk}\lambda_{jk}^2}{\sum_{h=1}^{K} z_{jh}\lambda_{jh}^2} \tag{8}$$

PVE equals one for a given dimension if it explains all the variance in the variable and decreases as it explains less variation relative to the other K-1 dimensions.

I calculate each of these metrics for each session and plot the results in Figure 4. Both measures establish the utility in considering higher dimensions when estimating country preferences from the UNGA voting data as a whole, but also show there is considerable variation in the relevance of other dimensions over time. Looking at the top panel, we see that between 1970 and 1975 all votes load on more than one dimension, but this steadily declines as the Cold War approaches its denouement in 1991. After 1994, there is another roughly 20 year period where nearly every vote in every session is multidimensional in nature which slightly tapers off in 2013. The bottom panel plots the PVE of the first dimension over time which is inversely, but not perfectly, correlated to the proportion of multidimensional votes. In the 1970s, the PVE of the first dimension hits a floor of about 50% and steadily increases throughout the Cold War. After 1995, the PVE falls to 60% and hovers there until 2017, the final year in the data. Taken together, these two metrics support the relevance of additional dimensions in a majority of UNGA sessions, in which they explain anywhere from 25 to 50 percent of variation beyond that captured by the first dimension.

Finally, there is the question of the predictive performance of mmBPFA. Since more dimensions explain more variation in the observed data, in theory this should translate to better predictive performance compared with a single-dimensional model. A measure of model performance commonly employed in the roll call scaling literature is the geometric mean probability (GMP) (Carroll et al. 2009) which rewards correct predictions and penalizes incorrect predictions. The GMP is calculated as follows:

$$GMP = \left(\prod_{i=1}^{N} \prod_{j=1}^{P} P(\hat{y}_{ij} = y_{ij}) \right)^{\frac{1}{N*P}} \tag{9}$$

where $\hat{y}_{ij}$ is the predicted vote, y_{ij} is the observed vote, N is the number of voters and P is the number of votes. In addition to the GMP, I calculate the root mean-squared error of predictions:

$$RMSE = \sqrt{(\hat{y}_{ij} - y_{ij})^2} \tag{10}$$

Figure 5 plots the predictive performance results. The GMP and RMSE are displayed in the top and bottom panels, respectively. The predictive performance of mmBPFA varies considerably over time. In 1970, the GMP performs quite poorly at about 0.5, indicating the mean probability of correct classification is only 50%. On sessions thereafter, performance improves considerably to the 0.65-0.7 range. Lower RMSE indicates better performance, and thus shows a similar, but opposite, trend to GMP. The GMP is objectively low, especially since the GMP is calculated for the same data the model is trained on. By comparison, Bailey and Voeten's (2018) two-dimensional model exhibits a GMP that is consistently better by about 0.2 to 0.3 over the same period. Such significant performance differences give reason for pause.

There are a couple possible explanations for this. One is due to the differences in how mmBPFA estimates cut points. Since it does not directly estimate them, but infers them from the extended rank likelihood, they are treated as a nuisance. This could lead to poor inferences for cutpoints on extremely imbalanced votes. The second, and more likely culprit, is a bug somewhere in the code. In simulations, I have found predictions to be systematically biased towards zeros, suggesting the model is not allowing enough mass

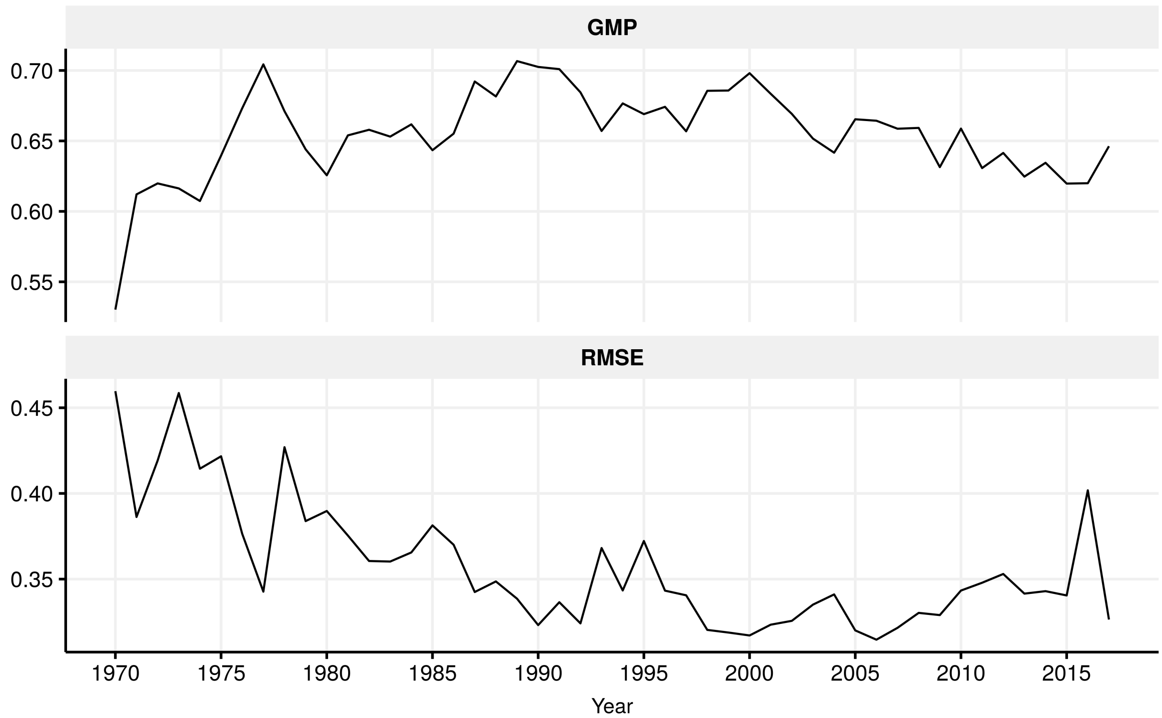

Figure 5: **Accuracy Measures of mmBPFA Predictions.** (top) The geometric mean probability of the estimates. (bottom) The root mean squared error of predictions per session.

to be placed in the tails of the distribution which points to an issue with the variance parameter on either Λ, Ω or both.

4.2 Model Interpretation

Having established the model's validity on the UNGA data, I now turn to the issue of substantive interpretation. My model identifies upwards of three dimensions in country foreign policy preferences. Yet, if those dimensions have no discernible or meaningful substantive interpretation, then they are of little use. I treat BSV's dynamic IRT as a baseline for substantive interpretation since their estimates are the standard within the literature.

Recall that BSV's first-dimension reflects contestation over the Western liberal international order, the primary ideological division within the United Nations. As a first gauge of how the dimensions found by mmBPFA map onto this division, I plot the the correlation between ideal points for each of these dimension and BSV's first-dimension ideal points in Figure 6. The first dimensions between both models are a nearly identical. Apart from a period between the 27^{th} and 40^{th} session where correlation between the two sets of ideal points vacillates between 0.8 and 0.9, the two are almost perfectly correlated. The second and third dimensions never exceed a correlation of 0.6 and, in fact, tend to show little correlation with the first dimension after the 43^{rd} session, indicating they are picking up elements of country preferences that are largely orthogonal to their location

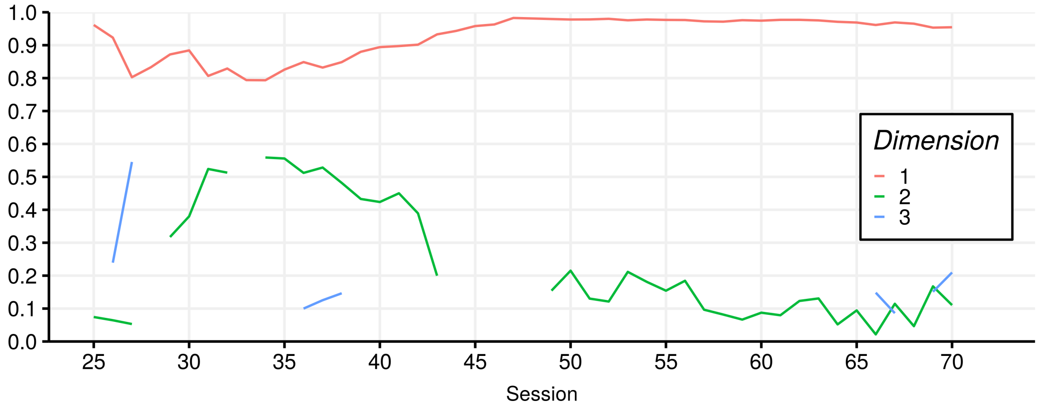

Figure 6: **Correlation between mmBPFA Ideal Points and 1st Dimension of BSV**

along the Liberal order dimension.

To help identify these dimensions, I ran structural topic models (Roberts, Stewart, and Tingley 2019) on the resolution descriptions with the dimension factor loadings as covariates for all sessions where given dimension was present. I also included B-splines for each session to account for temporal variation. Rather than selecting a specific number of topics which is known to be context dependent, I let the model choose a number of dimensions via spectral decomposition (2019, 9) and then filtered the topics to only those that were correlated with a dimension at the 95% confidence level. Vignettes of example resolutions for each topic are available in Appendix C.

The vignettes point to a second dimension largely defined by distributional concerns over the global commons, nuclear weapons, and UN reform (questions of contributions and refugee aid); colonialism; and human rights issues that overlap with sovereignty. This closely matches the "North-South" second dimension found by Bailey and Voeten (2018b) as well as Voeten (2000). The North-South division is historically associated with the Non-Aligned Movement (NAM) that began with the 1956 Bandung Conference in which smaller, developing countries utilized their strength in numbers to form an alternative to the superpower-led, bloc-driven geopolitics of the Cold War. By 1964, the Group of 77 (G-77) also emerged as a buffer against the capitalist international economic order (Doyle 1983). Issued raised by the G-77 also appear in the second dimension.

These two entities drive much of the second dimension until the end of the Cold War. The dimension disappears and re-emerges in 1994, reflecting contention over human rights and sovereignty concerns. Prior to this shift, the second dimension is more strongly correlated with the first (see figure 6); however, afterwards, the correlation attenuates. While the "North-South" label is apt during the 1960s-80s, this division does not hold up as well after the Cold War despite the dimension remaining stable. For instance, the second dimension produces an odd group of bedfellows in the post-Cold War Period. As the top panel of Figure 7 which plots the first versus second dimension ideal points in the 70^{th} session shows, Russia anchors the far negative pole and is followed closely by North Korea, Syria, China, Iran, Pakistan, and India. Slightly less far out countries such

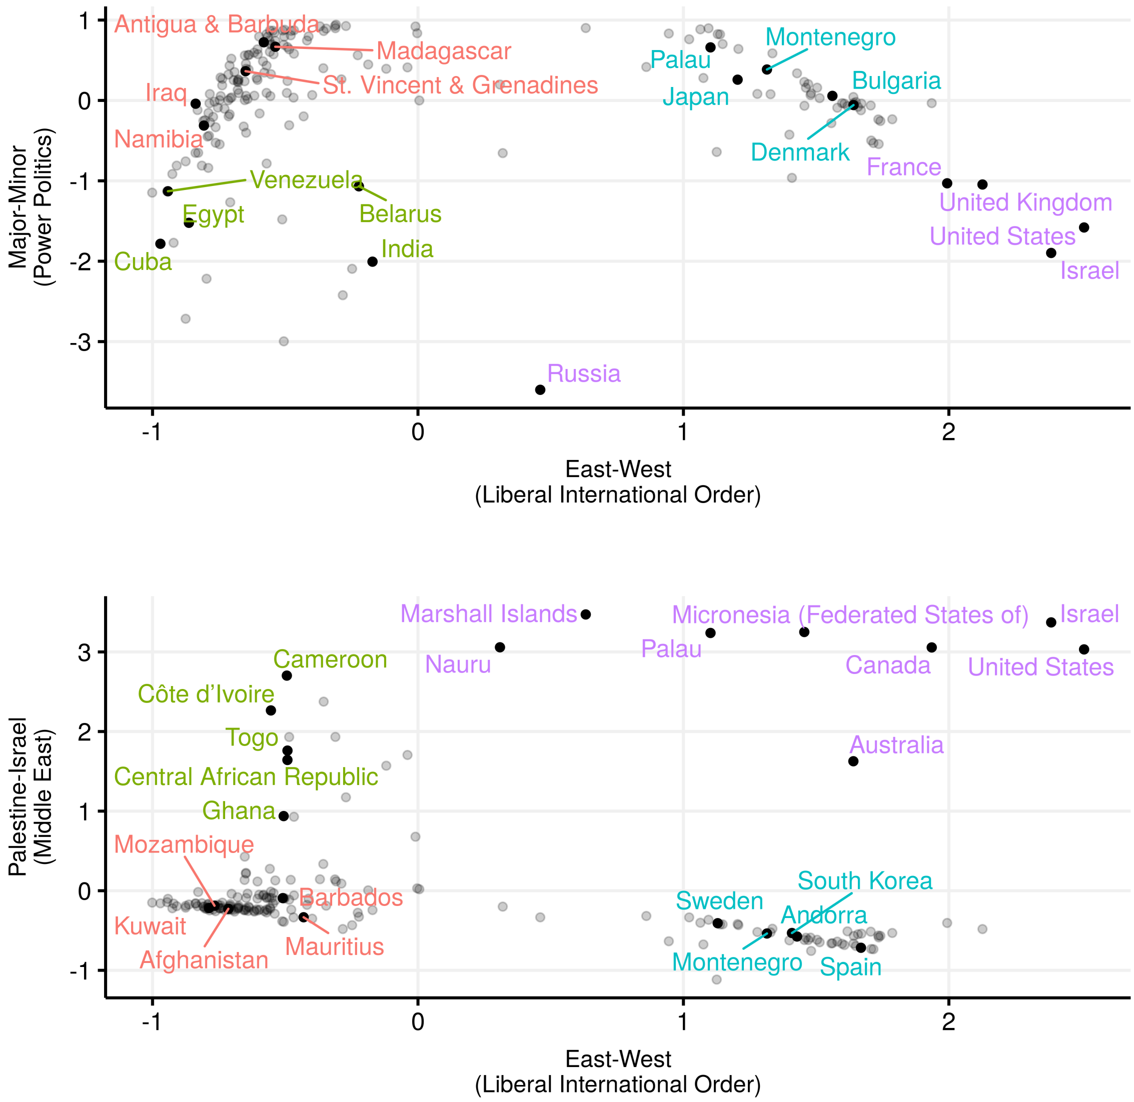

Figure 7: **Ideal Point-Dimension Scatter Plots for 70th Session of UNGA (2015–2016).** (top) Dimension 1 (x-axis) versus Dimension 2 (y-axis). (bottom) Dimension 1 (x-axis) versus Dimension 3 (y-axis).

as Israel, Cuba, Zimbabwe, the USA, Egypt, Uzbekistan, Myanmar, Sudan, Nicaragua, Venezuela, Belarus, Great Britain, and France all vote similarly on human rights and sovereignty issues. I therefore consider the second dimension as a Major-Minor Power divide. While not all of the aforementioned countries would rightly be considered major powers, many of them 1) have nuclear weapons or nuclear aspirations; 2) could be classified as attaining at least middle power status at some point; 3) have sizable militaries; or 4) want to insulate themselves from sovereignty incursions by international bodies.

The third dimension represents the Israeli-Palestine dispute and very instable. It tends to wax and wane depending on how prominently this issue appears on the agenda. While the Israeli-Palestine dispute also occasionally appears in the topic vignettes for the second dimension, this is likely due to label switching since the dimensions are not

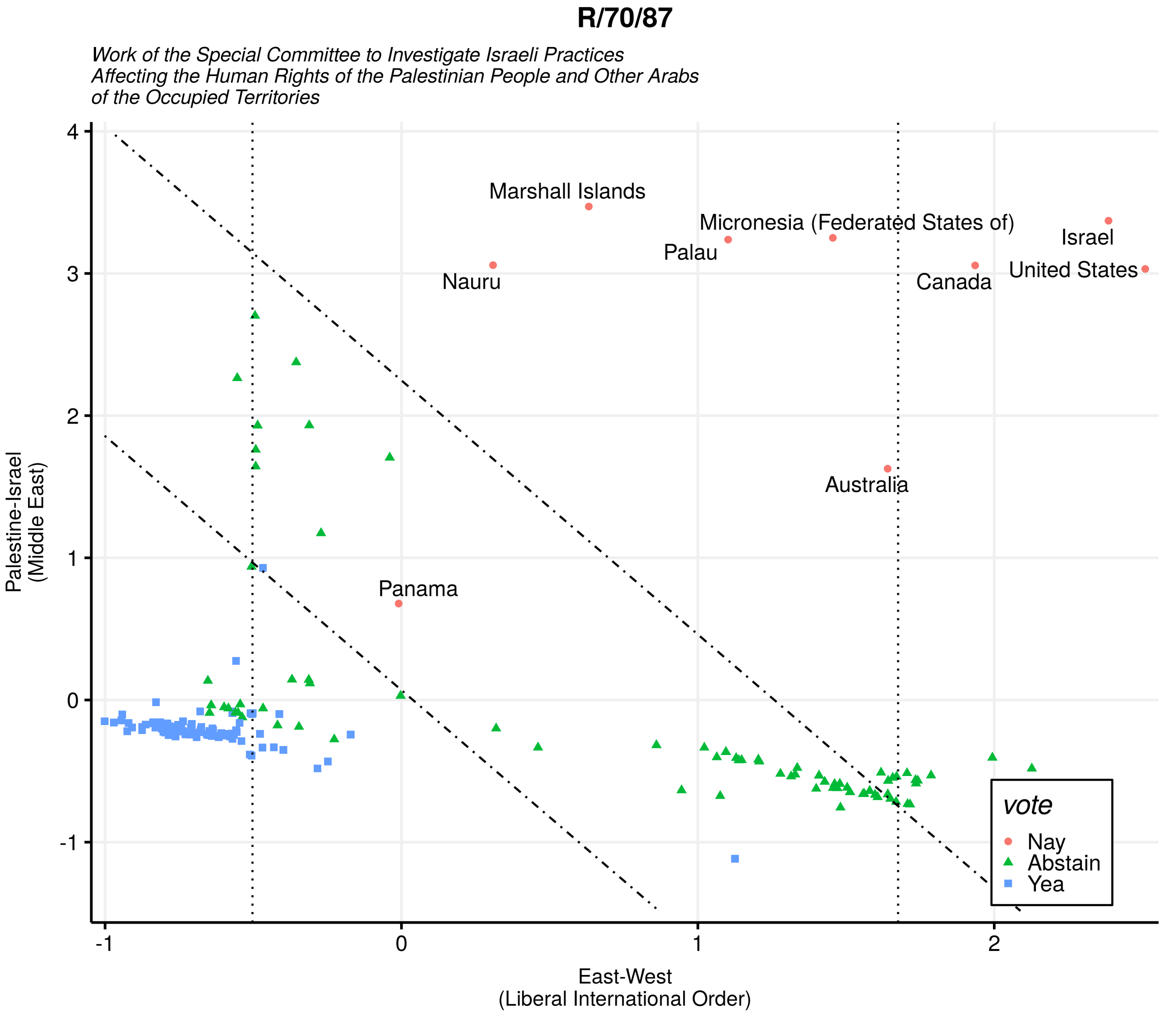

Figure 8: **One- and Two-Dimensional Cut Lines for Vote on Investigation of Israeli Human Rights Violations.** Resolution 87 of the 70^{th} Session of UNGA (2015–2016). The dot-dash lines indicate two-dimensional cuts and the dotted lines mark one-dimensional cuts.

temporally fixed.[10] The bottom panel in Figure 7 plots the distribution of first versus third dimension ideal points in the 70^{th} Session (2015), one of the most recent sessions in which the third dimension makes an appearance. Almost all the variation is driven by Israel and its closest diplomatic allies (purple). There are also a handful of African countries (green) that lie opposite from Israel on the first dimension, but who vote very similarly on Israel-Palestine resolutions. Many of these also have long-standing diplomatic relations with Israel. The remaining countries show very little variation on this dimension.

10. The appearance of Colonialism/apartheid topics in the third dimension vignettes further substantiate label switching as the culprit. These should be in the second dimension and likely are due to mis-labeling in Sessions 25–27. This will be addressed in the future.

Although the third dimension is highly unstable, it is extremely important in explaining country preferences on the Middle East Peace issue. To illustrate this, I plot the estimated cut lines for Resolution 87 in Session 70 in Figure 8. This resolution heavily criticized Israel's military operations in the Gaza strip in 2008-9 and 2014 and called for investigations into Israeli human rights violations in occupied Palestine. The dot-dash lines split the space into yea, abstain, and nay predicted voting blocs. While the probabilistic prediction is not perfect, too narrowly predicting abstentions, it does markedly better than the first dimension (dotted lines). More importantly, we see that although the lines are diagonal—indicating the first dimension contributes to explaining the vote as well—they would likely be even more orthogonal to the first dimension if not for some of mmBPFA's under-fitting.[11]

5 Moving Beyond Votes

To highlight mmBPFA's flexibility in terms of data, I also estimate ideal points based on two different extensions beyond votes: UNGA opening speeches (Baturo, Dasandi, and Mikhaylov 2017) and countries' Universal Period Review (UPR) statements (Terman and Voeten 2018). When moving beyond votes, whether in combination with votes or in isolation, the interpretation of dimensions becomes more complex, as illustrated below, since the data generating process is different. The data generating process informs not only what the dimensions represent, but how they should or could be used as independent variables in other analyses.

5.1 UNGA Votes and Debates

Due to space constraints and for interpretation's sake, I only analyze a recent subset of UNGA sessions 65 through 72 covering the 2010-2017 period. To generate the data matrices, I performed typical text pre-processing on the speeches—lemmatization; removing stop words, punctuation, special characters, and numbers. I used the `spaCy` software for named entity and nounphrase recognition so as to preserve important phrases during tokenization. I then trimmed words that appear in more than 80 or less than 50 percent of documents. This excludes extremely rare words that are likely to be unique to very countries as well as words that are so common as to provide little variation across countries. Finally, I joined the resulting document-term matrices of word counts with the votes by country.

In the typical IRT model, we assume countries do not vote strategically, but rather their votes represent genuine, expressed preferences. I assume speeches are also manifestations of a country's latent preferences. However, the two reflect preferences expressed at two different stages in the policymaking timeline and likely are not perfectly congruent. Voting on resolutions occurs at the final step in a multistage process, in which each stage offers countries different degrees of freedom in preference expression. By the time a resolution proceeds to a vote, many disagreements have already been worked out. Opening speeches within the UN occur at the outset of each session. Countries are completely unconstrained as to which agenda issues they want to advocate for and how. Thus, one might speculate that there are three different types of latent dimensions: those specific to votes, those to speeches, and one common to both.

11. See the concerns over the predictive validity in Figure 5.

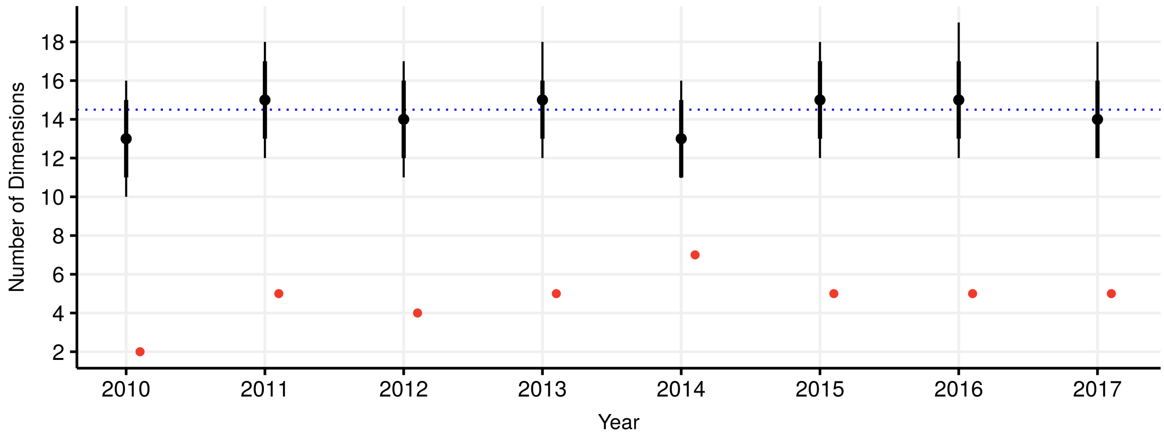

Figure 9: **Dimensionality of the 65$^{\text{th}}$–72$^{\text{nd}}$ UNGA Sessions based on Votes and Speeches.** Points indicate the median posterior number of dimensions. Thick and thin bars represent 80% and 95% credible intervals, respectively. The median number of dimensions across these session is 15, indicated by the blue-dotted line. The red dot indicates the number of dimensions retained after trimming highly similar dimensions based on pairwise Euclidean distances.

Compared with vote-only ideal points, the vote-speech ideal point estimates from mmBPFA exhibit a couple notable differences. The first is dimensionality. Figure 9 plots the number of dimensions found by mmBPFA for the vote-speech data. In all the sessions explored, the estimated number of dimensions for the mixed-mode data is considerably higher. The average across these sessions, indicated by the dotted blue line, is 15 dimensions. However, the number of dimensions retained drops significantly after I prune highly similar dimensions. Though this step introduces an additional researcher degree of freedom (Simmons, Nelson, and Simonsohn 2011), I find it aids greatly in interpretation of the higher order dimensions which can be difficult to distinguish from one another when they load on highly similar votes and words.[12]

In most sessions, post-trimming removes around 10 dimensions. The only exception being the 69th session where far fewer are removed. Nonetheless, even after trimming, the number of retained dimensions exceeds the number found by examining only votes. As I show below, these dimensions also differ substantively from the vote-only dimensions, reflecting instead contestation and disagreement over issue-specific framing.

The second notable difference is that when we move to the mixed-mode data, the dimensions become far less hierarchical in terms of their explanatory power. The spider plots in Figure 10 displays the PVE for each dimension across the sessions analyzed. Rather than a first dimension that drives most of the explanation, the first dimension

12. Whether additional post-trimming should be done is up for debate. On the one hand, it may seem counter-intuitive to remove dimensions since one key advantage of mmBPFA is it's ability to estimate higher-order dimensions. On the other hand, interpretability is key and higher order dimensions, at least in the case of votes and speeches, often reflect contestation over the framing of issues that are so nuanced as to elude easy identification.

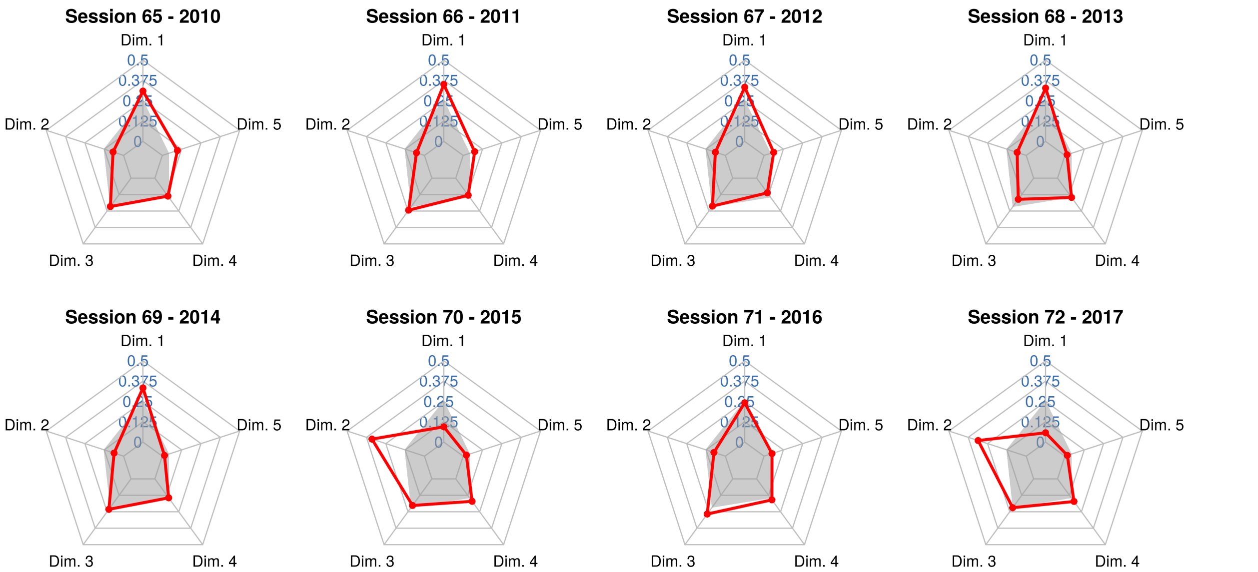

Figure 10: **Proportion of Variance Explained (PVE) by the First Five Dimensions in the 69th–72nd UNGA Sessions.** The spider plots are zoomed in to the [0,0.5] interval. Each rung represents 12.5% of variance explained. The red dots indicate the proportion of variance explained by that dimension. The light gray pentagon illustrates the average PVE across the sessions to help highlight session-wise deviations.

explains about 30 percent of the variation in the data on average and the second most important dimension explains about 25 percent. Thereafter, the remaining dimensions are very issue-specific, explaining little variation. The shape of the red lines show that the predominant dimension is fairly stable over time and occasionally switches as in sessions 70 and 72. This is likely due to session-specific agenda shifts and suggests these ideals points should not be thought of as stable ideological preferences.

To identify what each dimension represents, I zoom in on the 69^{th} session. In this session, the number of dimensions retained after trimming (7) was the closest to the median number identified by mmBPFA during sampling (13), and is the most dimensions kept after additional scrutiny out of all the sessions. The 69^{th} session also took place in 2014-15, following Russia's annexation of Crimea and during the United States' intervention in the Syrian civil war against ISIL, offering much for debate.

Figure 11 plots the ideal points of the top ten countries occupying each pole of the seven dimensions. The seven dimensions correspond to six different issues: post-colonial and de-nuclearization issues, security threats, development challenges, human rights, Western-led multilateralism, and middle east peace. There are two dimensions dedicated to security frames (discussed below). I identified the dimensions by examining 40 of the most polarizing resolutions and phrases from each dimension. Some dimensions are dominated by votes such as the first one on post-colonial and nuclear issues. Others are solely determined by speeches, such as the two on security threats, development challenges and multilateralism. Human rights and Middle East Peace were associated with both votes and speeches.

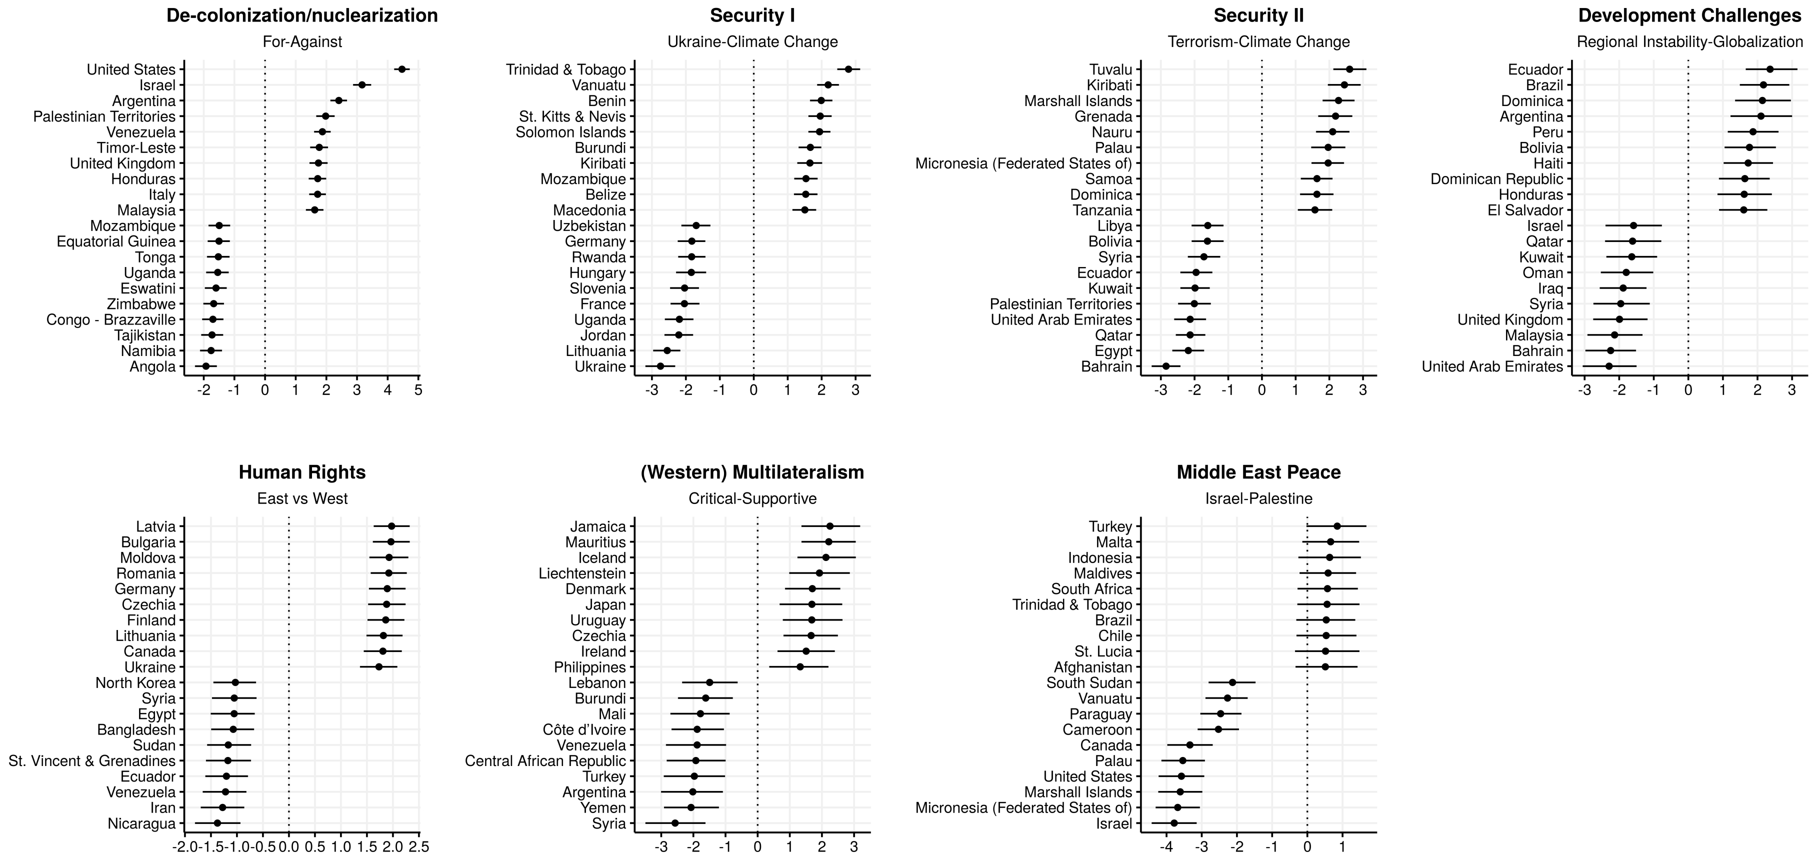

Figure 11: **UNGA 69$^{\text{th}}$ Session Ideal Points by Dimension.** Each panel displays the ten countries occupying each pole of the given dimension. Ideal points are estimated based on the median posterior draw, and black bars represent 95% credible intervals.

Topically, the de-colonization/nuclearization dimension aligns closely to the North-South and major-minor power divide. Issues of nuclear non-proliferation and the rights of the Non-Self Governing territories usually pit former colonizers and nuclear powers against de-colonized countries. However, in terms of ideal points, the positive pole is occupied by odd bedfellows. While countries such as the United States, Argentina, the United Kingdom and Italy understandably have an interest in the Non-Self Governing territories[13], and Israel cares about nuclear issues in the Middle East, the presence of other countries is less easily explained. Palestine sits only one position removed from Israel. Other notable outliers include Venezuela, Timor-Leste and Honduras.

The second and third dimensions represent divisions over how states frame the dominant security threats that should be addressed by the UN. One of them pits countries—many European or Post-Soviet—that emphasized the situation in Ukraine as the predominant security threat requiring UN attention against a coterie of small-island developing states and African nations that highlight climate change as the primary existential, and security, threat. The second pits countries primarily emphasizing terrorism, the treat posed by ISIL, and civil wars in Syria and Yemen against the climate change coalition. These two dimensions stand out since one might expect the same issue should be contained in a single dimension. The most plausible explanation is the the climate change group represents a unified ideological, but the Ukraine and Terrorism groups are too ideologically distinct in terms of their expressed voting preferences to be placed on a single dimension. Moreover, each of these represents competing frames. Given the number of intransigent security challenges facing the world in 2014, having two distinct dimensions dedicated to the issue is not perplexing.

The fourth dimension, development challenges, is purely speech-related and reflects two frames commonly employed by states when discussing the primary development challenges their country faces. Middle Eastern states heavily regional instability as inhibiting economic development whereas many Latin American countries—with their rocky history with international financial institutions such as the IMF—maintain that globalization and its associated inequities are the key hindrance to development.

A similar, yet slightly different, issue-framing dimension is that addressing multilateralism. Here, the two poles are occupied by states that are either critical or supportive of multilateralism. Though the two groups are united in their antipathy or approval, the reasons for criticism are not uniform. For instance, Yemen, Turkey and Syria are critical of intervention in the their countries under the guise of Western, especially American, led multilateralism; whereas Argentina and Venezuela criticize the economic and financial elements of Western multilateralism. Among the supporters, the messaging is less specific and more uniform in its championing of multilateralism as the best means to solve economic, political, and security challenges.

Each of the issue-frames should not be thought of as ideological dimensions since their poles need not be mutually exclusive and are driven by issues contemporaneous to a given session. They primarily reflect agenda-setting preferences on issue areas associated with the sub-committees of the UNGA. They will thus shift over time with the agenda,

13. Under Chapter XI of the Charter of the United Nations, the Non-Self-Governing Territories are defined as "territories whose people have not yet attained a full measure of self-government". For a full list, see: https://www.un.org/dppa/decolonization/en/nsgt

but are likely linked with contestation within issue areas anchored to the subcommittees.

The remaining two dimensions, human rights and Middle East peace, draw on both votes and speeches. They also hold ideological interpretations associated with the three votes only dimensions. In the case of Middle East peace, there is a one-to-one correspondence with the Israel-Palestine dimension. Human Rights is a unique sub-dimension of the East-West divide identified by BSV and also found by mmBPFA. It pits the Western definition of human rights which is centered on political, civil, and individual liberties against the various re-interpretations often offered by autocratic regimes and other human rights violators. Thus, we find countries such as Iran, Venezuela, Sudan, Egypt, Syria and North Korea occupying one pole, with primarily European countries situated at the other end. Whereas the United States anchors the pole in the votes-only first dimension, here it is absent. Instead, the Baltic states along with Germany and Canada act as anchors instead. The appearance of Ukraine, Bulgaria, Moldova, and Romania is driven by their emphasis on the situation in Ukraine and Crimea in their opening speeches.

5.2 UNGA Votes and Universal Period Review

The UPR is a mechanism conducted by the UN Human Rights Council to regularly analyze the human rights practices of member states. The mechanism relies on self- and peer-review to overcome participatory and political shortcomings of it's predecessor, the UN Human Rights Commission. The 47 member UPR Working Group meets three times a year to review 14-16 member states according to a lot. Each country is thus reviewed approximately once every 4.5 years. In each review, states commit a self-assessment followed by a 140 minute dialogue in which member states and permanent observers make recommendations for improvements. Recommendations are categorized into 47 different topics and recorded by Reviewer-Reviewee dyad.

I use the raw recommendation data from Terman and Voeten (2018) covering the first 20 sessions of the UPR and combine this with UNGA votes. Since UPR rounds and UNGA sessions are on differently schedules, I code rounds according to their contemporaneous UNGA session and then aggregate the reviews by reviewer country and UNGA session. This creates a set of topic counts for each reviewer covering UNGA sessions 62–69 that reflects their human rights priorities and ideology. Not all member countries participate in reviews each year. As Terman and Voeten note, the review processed is heavily colored by strategic foreign policy considerations. Countries decide whom to review, what criticisms to raise, and the severity of criticisms depending on reviewer-reviewee bilateral ties. Thus, the topic counts in any given session will be influenced by the agenda (i.e. which countries are reviewed in that period) and not necessarily reflect unbiased human rights preferences. However, the topic counts should help ameliorate this since about 45 countries are reviewed in each UNGA session, meaning the counts likely capture reviews of allies and adversaries alike.

The results show there are four dimensions of preferences on average contained in the Votes-UPR data and three after post-processing. The top panel in Figure 12 plots the number of dimensions found in each session and the bottom panel shows the PVE for each dimension. While the first dimension, which still corresponds to the liberal West versus the rest axis, dominates in most sessions, there is also a viable second dimension that

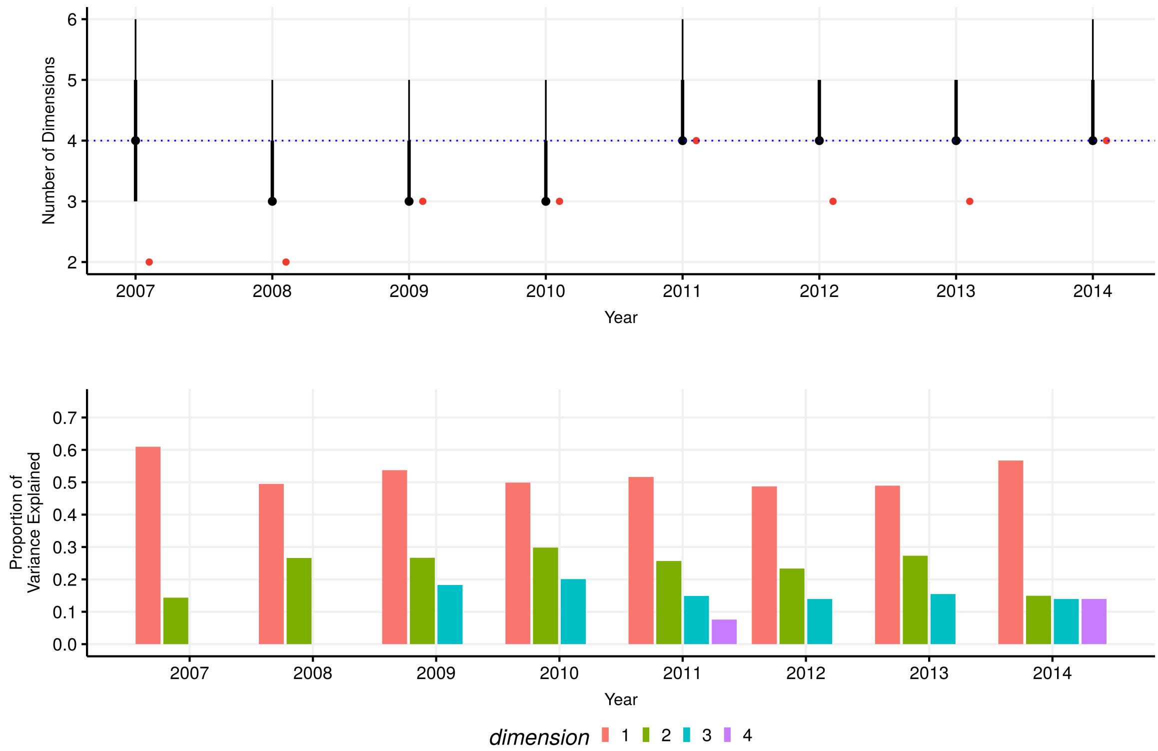

Figure 12: **Dimensionality and Proportion of Variance Explained in the 62nd–69th UNGA Sessions based on Votes and Universal Period Review Statements.** (top) Points indicate the median posterior number of dimensions. Thick and thin bars represent 80% and 95% credible intervals, respectively. The median number of dimensions across these session is 4, indicated by the blue-dotted line. The red dot indicates the number of dimensions retained after trimming highly similar dimensions based on pairwise Euclidean distances. (bottom) Proportion of Variance Explained by Dimension.

explains anywhere form about 12-30% of variation in depending on the session. Higher dimensions are also present beginning in 2009 and fluctuate in relevance depending on the session which is likely a function of the UNGA and UPR agendas.

Since these models are still anchored by votes, they share some commonalities with the the votes-only and votes-speeches but also exhibit differences. Table 2 provides examples of the most discriminating votes and UPR topics for each dimension along with their factor loadings. There are four distinct dimensions. Some dimensions are uni-directional in that variation is driven by a cohort of countries voting for (against) / raising a specific set of issues but no cohesive antipodal group or set of issues.

The first dimension corresponds to a socioeconomic vs political/civil liberty human rights conception axis. Although nearly identical to the liberal West versus the rest dimension (correlation of 0.98), note that the most informative items differ once UPR data is incorporated. The liberal west pole is driven by criticism of human rights in Belarus

Table 2: Sample of Most Discriminating Resolutions and UPR Topics, Sessions 62-69

Item	Factor Loading
Socioeconomic vs Political/Civil Liberty Human Rights	
R/62/27: Promotion of multilateralism in the area of disarmament and non-polifer...	-0.692
R/62/166: Respect for the purposes and principles contained in the Charter of th...	-0.691
R/64/156: Combating defamation of religions	-0.688
R/69/178: Promotion of a democratic and equitable international order	-0.687
R/68/184: Situation of human rights in the Islamic Republic of Iran	0.612
R/62/169: Situation of the human rights in Belarus	0.640
Race, Social Justice, and Minorities	
R/65/73: The Hague Code of Conduct against Ballistic Missile Proliferation	-0.197
R/67/202: Entrepreneurship for development	-0.121
R/64/97: Information from Non-Self-Governing Territories transmitted under Artic...	-0.090
R/64/106: Implementation of the Declaration on the Granting of Independence to C...	-0.087
R/64/105: Dissemination of information on decolonization	-0.087
UPR: International instruments	0.583
UPR: Racial discrimination	0.585
UPR: Justice	0.588
UPR: Minorities	0.613
Israel vs Palestine	
R/69/92: Israeli settlements in the Occupied Palestinian Territory including Eas...	-0.563
R/69/93: Israeli practices affecting the human rights of the Palestinian people ...	-0.556
R/69/91: Applicability of the Geneva Convention relative to the Protection of Ci...	-0.549
R/69/24: Jerusalem	-0.541
R/69/22: Special information programme on the question of Palestine of the Depar...	-0.539
R/69/94: The occupied Syrian Golan	-0.538
UPR: Elections	0.076
R/65/142: International trade and development	0.084
UPR: Public security	0.085
UPR: Disabilities	0.086
R/64/99: Implementation of the Declaration on the Granting of Independence to Co...	0.089
R/67/217: Towards a new international economic order	0.099
Economic and Military Development	
UPR: Development	-0.280
UPR: Disabilities	-0.234
UPR: Other	-0.233
UPR: Public security	-0.216
UPR: Poverty	-0.211
UPR: Right to health	-0.203
R/69/49: The Arms Trade Treaty	0.378
R/66/176: Situation of human rights in the Syrian Arab Republic	0.378
R/66/29: Implementation of the Convention on the Prohibition of the Use Stockpil...	0.389
R/69/59: Compliance with non-proliferation arms limitation and disarmament agree...	0.392
R/69/52: United action towards the total elimination of nuclear weapons	0.414

and Iran in each session. The other pole is more diverse, informed by votes on a more equitable international order, religious protections, and multilateralism in disarmament. These differences map onto larger debates over universal versus relative conceptions of human rights wherein the former is usually associated with Western notions of civil and political liberties inherent to all humans and the latter with economic, social and cultural rights (Renteln 2013).

The second dimension is unidirectional, capturing countries whose human rights ideologies focus on issues of racial, ethnic, and religious discrimination. These concerns are often raised by African and Asian nations when criticizing Western nations, and particularly the United States. On the surface, this may appear to be similar to the first dimension since it is anchored by the US at one pole and non-Western, illiberal African and Asian countries at the other. However, the two are almost completely uncorrelated, indicating there is more than strategic criticism of adversaries driving this dimension.[14] This is clear from Figure 13 which plots the first versus second dimension ideal points across the 62^{nd} through 69^{th} sessions. While there is clear separation into two clusters along the first dimension, positions along the second dimension are more evenly distributed. We can also see that attention to racial, ethnic and religious discrimination are not solely the tools of African and Asian nations, but central to the ideology of many Western countries as well, especially within Europe.

The third dimension is also unidirectional and corresponds to the Israel-Palestine axis that appears in the votes-only and votes-speeches analyses as well, providing further evidence of the staying power of this dimension within the United Nations. While human rights concerns are inherent drivers of countries' positions on this dimension, especially for European nations, adding the UPR data does not fundamentally alter ideal points nor the most informative items.

The fourth dimension is harder to clearly identify since it is not unidirectional, but consists of strong factor loadings in both directions on the seemingly unrelated issues of small arms/nuclear weapon non-proliferation and issues related to economic (under)development such as poverty, right to health, situation of disabled persons, and public security. Rather than representing an ideological spectrum, this dimension likely captures strategic behavior which would explain why it appears sporadically. Developing countries pushing for weapons non-proliferation such as those impacted by conflict are also likely to struggle with issues of poverty, development, and public security. To decrease the likelihood of being criticized on these issues when they are under review, they de-emphasize these in their reviews of others according to a tit-for-tat logic.

6 Conclusion

This paper contributes to the literature on multidimensional scaling and estimating countries' foreign policy preferences from IO data. First, it overcomes traditional model constraints that limit our inference to a single form of data, namely roll call votes. The use of roll call votes to estimate latent preferences is ubiquitous in the American and IR literatures. Only recently have scholars started to develop methods to estimate pref-

14. Indeed, all of the dimensions are essentially orthogonal to the first dimension and each other. Pairwise scatter plots for higher dimensions are available in Appendix D.

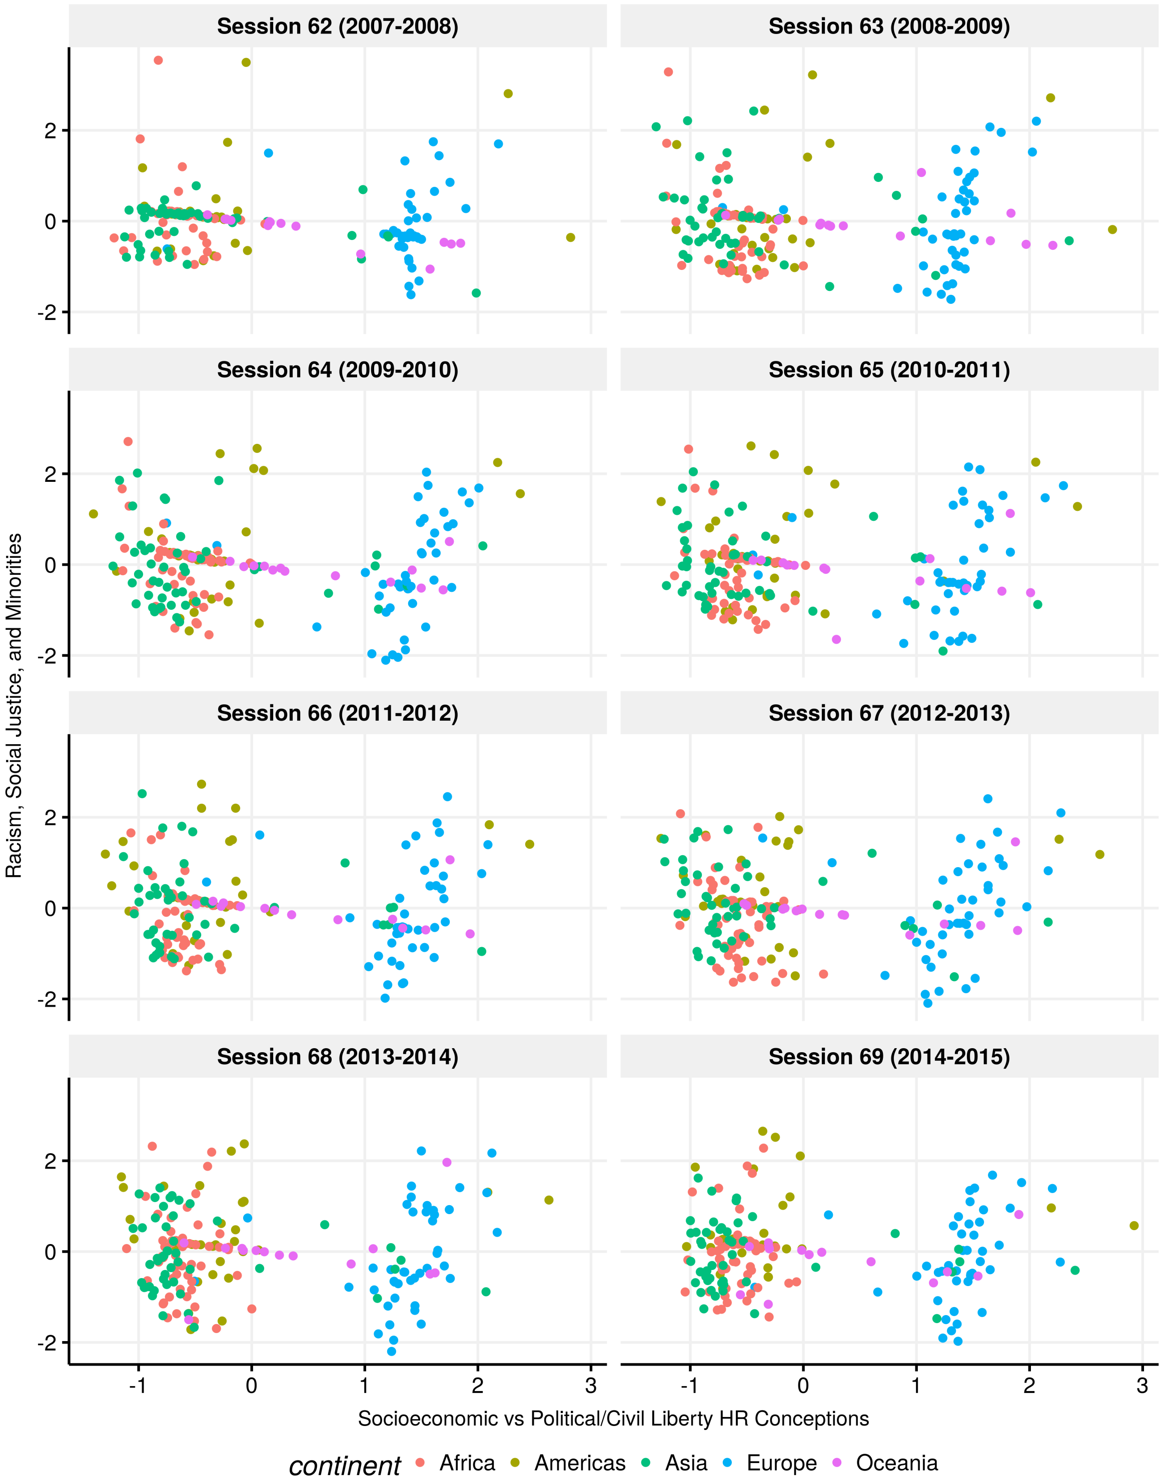

Figure 13: **First versus Second Dimension Ideal Points, 62nd–69th UNGA Sessions based on Votes and Universal Period Review Statements.**

erences from richer combinations of data such as votes and speech, but these methods are bespoke and not generalizable. `mmBPFA` allows for arbitrary combinations of data, opening the door for researchers to draw on novel, domain-specific datasets to estimate latent preferences. This has important implications for IR applications where countries regularly express their foreign policy preferences, but the information is unstructured, piecemeal and scattered across various institutions and formats.

Second, I draw on Bayesian nonparametric priors allowing researchers to escape the guess-and-check approach to identifying dimensions common in standard multidimensional models. Instead, `mmBPFA` will identify the as many dimensions as stably present in the data. As data sets grow (shrink), so too will the dimensionality. Additional post-processing further ensures that only the most stable dimensions are kept, leaving researchers confident the results represent a lower bound on the true dimensionality.

Finally, by applying `mmBPFA` to UNGA votes, votes and speeches, and combination of votes and UPR statements I find that foreign policy preferences exist within a macro-micro hierarchy. In addition to the first two macro dimensions—West versus the Rest and North versus South—identified by Bailey, Strezhnev, and Voeten (2017) and Bailey and Voeten (2018b), I also find significant variation in agenda-setting preferences within the UNGA speeches associated with questions of security, human rights, multilateralism, development, non-proliferation, and Middle East peace. Although these preferences capture divergent preferences over prevailing issues of the day, and so are not stable over long periods of time, many of them persist over shorter spans and may be of interest to researchers focusing on specific periods or issues. I also show that contestation over human rights conceptions plays out in the UNGA voting and UPR statements of member states which cannot be explained purely by strategic considerations. This opens up the door to future research to examine the evolution of these human rights preferences over time and the degree to which they are independent of countries' strategic interests.

Despite these contributions, the model has shortcomings. First, it is not dynamic which leads to label switching on dimensions across time and increased responsiveness to agenda-shifts. The former is a hindrance to inter-session inference while the latter allows for potentially undesirable session-to-session volatility in ideal points. Future iterations should include a dynamic element to overcome both of these issues. Second, as the analyses illustrate, the ability to discover higher dimensionality is a blessing and a curse. Higher order preferences can be difficult to substantively validate, at best, or spurious, at worst. Researchers interested in utilizing higher-order dimensions as independent or dependent variables in a regression should carefully validate them against other data sources to ensure they have a clear substantive interpretation that matches the desired ideological construct.

References

Albert, James H., and Siddharta Chib. 1993. "Bayesian Analysis of Binary and Polychotomous Response Data." *Journal of the American Statistical Association* 88:669–79.

Alesina, Alberto, and David Dollar. 2000. "Who Gives Foreign Aid to Whom and Why?" *Journal of Economic Growth* 5 (1): 33–63.

Alter, Karen, and Kal Raustiala. 2018. "The Rise of International Regime Complexity." *Annual Review of Law and Social Science* 14:329–249.

Bailey, Michael A., Anton Strezhnev, and Erik Voeten. 2017. "Estimating Dynamic State Preferences from United Nations Voting Data." *Journal of Conflict Resolution* 61 (2): 430–536.

Bailey, Michael A., and Erik Voeten. 2018a. "A Two-Dimensional Analysis of 70 Years of United Nations Voting." *Public Choice* 176:33–55.

———. 2018b. "A Two-Dimensional Analysis of Seventy Years of United Nations Voting." *Public Choice* 176 (1-2): 33–55.

Ball, Margaret M. 1951. "Bloc Voting in the General Assembly." *International Organization* 5 (1): 3–31.

Baturo, Alexander, Niheer Dasandi, and Slava Mikhaylov. 2017. "Understanding State Preferences With Text as Data: Introducing the UN General Debate Corpus." *Research and Politics* 4 (2). https://doi.org/10.1177/2053168017712821.

Bearce, David H., and Stacy Bondanella. 2007. "Intergovernmental Organizations, Socialization, and Member-state Interest Convergence." *International Organzation* 61:703–733.

Benoit, Kenneth, and Michael Laver. 2003. "Estimating Irish Party Positions Using Computer Wordscoring: The 2002 Elections." *Irish Political Studies* 17:97–107.

Carroll, Royce, Jeffrey B. Lewis, James Lo, Keith Poole, and Howard Rosenthal. 2009. "Measuring Bias and Uncertainty in DW-NOMINATE Ideal Point Estimates via the Parametric Bootstrap." *Political Analysis* 17 (3): 261–275.

Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. "The Statistical Analysis of Roll Call Data." *American Political Science Review* 98:355–370.

Davis, Christina L., and Meredith Wilf. 2017. "Joining the Club: Accession to the GATT/WTO." *Journal of Politics* 79 (3): 964–978.

Dobra, Adrian, and Alex Lenkoski. 2011. "Copula Gaussian Graphical Models and Their Application to Modeling Functional Disability Data." *The Annals of Applied Statistics* 5 (9): 969–993.

Doshi-Velez, Finale, Kurt T. Miller, Jurgen Van Gael, and Yee W. Yeh. 2009. "Variational Inference for the Indian Buffet Process." *International Conference on Artificial Intelligence and Statistics,* 137–144. http://proceedings.mlr.press/v5/doshi09a/doshi09a.pdf.

Doyle, Michael W. 1983. "Stalemate in the North-South Debate: Strategies and the New International Economic Order." *World Politics* 35 (3): 426–464.

Dreher, Axel, and Martin Gassebner. 2008. "Does Political Proximity to the U.S. Cause Terror?" *Economic Letters* 99 (1): 27–29.

Dreher, Axel, and Nathan Jensen. 2007. "Independent Actor or Agent? An Empirical Analysis of the Impact of U.S. Interests on International Monetary Fund Conditions." *Journal of Law and Economics* 50:105–124.

———. 2013. "Country or Leader? Political Change and UN General Assembly Voting." *European Journal of Political Economy* 29:183–196.

Drieskens, Edith. 2010. “Towards a Systematic Analysis of the EU as an Actor in the UN System.” *Journal of International Organizations Studies* 1 (1): 105–108.

Flores-Macías, Gustavo, and Sarah E. Kreps. 2013. “The Foreign Policy Consequences of Trade: China’s Commercial Relations with Africa and Latin America, 1992-2006.” *Journal of Politics* 75 (2): 325–371.

Gartzke, Erik. 1998. “Kant We All Just Get Along? Opportunity, Willingness, and the Origins of the Democratic Peace.” *American Journal of Political Science* 42 (1): 1–27.

Genest, Christian, Michel Gendron, and Michaël Bourdeau-Brien. 2009. “The Advent of Copulas in Finance.” *The European Journal of Finance* 15 (7-8): 609–618.

Ghahramani, Zoubin, and Thomas L. Griffiths. 2006. “Infinite Latent Feature Models and the Indian Buffet Process.” In *Advances in Neural Information Processing Systems,* edited by Y. Weiss, B. Schölkopf, and J. C. Platt, 475–482. MIT Press.

Girod, Desha M., and Jennifer L. Tobin. 2016. “Take the Money and Run: The Determinants of Compliance with Aid Agreements.” *International Organization* 70 (1): 209–239.

Greenhill, Brian, and Yonatan Lupu. 2017. “Clubs of Clubs: Fragmentation in the Network of Intergovernmental Organizations.” *International Studies Quarterly* 61:181–195.

Grimmer, Justin, and Brendan M. Stewart. 2013. “Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts.” *Political Analysis* 21 (2): 267–297.

Hjort, Nils L. 1990. “Nonparametric Bayes Estimators Based on Beta Processes in Models for Life History Data.” *The Annals of Statistics* 18 (3): 1259–1295.

Hoff, Peter D. 2007. “Extending the Rank Likelihood for Semiparametric Copula Estimation.” *The Annals of Applied Statistics* 1 (1): 265–283.

Hollenbach, Florian M., Iavor Bojinov, Shahryar Minhas, Nils W. Metternich, Michael D. Ward, and Alexander Volfovsky. 2018. “Multiple Imputation Using Gaussian Copulas.” *Sociological Methods & Research,* https://doi.org/10.1177%2F0049124118799381.

Johnson, Valen E., and James H. Albert. 1999. *Ordinal Data Modeling.* New York, NY: Springer.

Kim, In Song, John Londregan, and Marc Ratkovic. 2018. “Estimating Spatial Preferences from Votes and Text.” *Political Analysis* 26 (2): 210–229.

Knowles, David, and Zoubin Ghahramani. 2011. “Nonparametric Bayesian Sparse Factor Models with Application to Gene Expression Modeling.” *Annals of Applied Statistics* 5 (2B): 1534–1552.

Koremenos, Barbara. 2005. “Contracting around International Uncertainty.” *American Political Science Review* 99 (4): 549–565.

Lauderdale, Benjamin E., and Alexander Herzog. 2016. “Measuring Political Positions from Legislative Speech.” *Political Analysis* 24 (3): 374–394.

Liao, Steven, and Daniel McDowell. 2016. “No Reservations: International Order and Demand for the Renminbi as a Reserve Currency.” *International Studies Quarterly* 60 (2): 272–293.

Lijphart, Arend. 1963. “The Analysis of Bloc Voting in the General Assembly: A Critique and Proposal.” *The American Political Science Review* 57 (4): 902.

McAlister, Kevin. 2020. *Disagreement and Dimensionality: A Beta Process Approach to Estimating Dimensionality of Ideal Points in the US Congress.* Working Paper. University of Michigan. http://www.kevinmcalister.org/uploads/1/2/0/2/120250241/bpirt_draft_version.pdf.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2001. “The Hunt for Party Discipline in Congress.” *American Political Science Review* 95 (3): 673–687.

Moon, Bruce E. 1985. “Consensus or Compliance? Foreign-policy Change and External Dependence.” *International Organization* 39 (2): 297–329.

Morse, Julia C., and Robert O. Keohane. 2014. “Contested Multilateralism.” *Review of International Organizations* 9 (4): 385–412.

Murray, Jared S., David B. Dunson, Lawrence Carin, and Joseph E. Lucas. 2013. “Bayesian Gaussian Copula Factor Models for Mixed Data.” *Journal of the American Statistical Association* 108 (502): 656–665.

Neumayer, Eric. 2008. “Distance, Power, and Ideology: Diplomatic Representation in a World of Nation-States.” *Area* 40 (2): 228–236.

Paisley, John, and Lawrence Carin. 2009. “Nonparamteric factor analysis with beta process priors.” In *Proceedings of the 26th Annual International Conference on Machine Learning,* 777–784. ACM.

Panke, Diana. 2017. “The Institutional Deisgn of The United Nations General Assembly: An Effective Equalizer?” *International Relations* 31 (1): 3–20.

Park, Trevor, and George Casella. 2008. “The Bayesian Lasso.” *Journal of the American Statistical Association* 103 (482): 681–686.

Pitt, Michael, David Chan, and Robert Kohn. 2009. “Efficient Bayesian Inference for Gaussian Copula Regression Models.” *Biometrika* 93:537–554.

Quinn, Kevin M. 2004. “Bayesian Factor Analysis for Mixed Ordinal and Continuous Responses.” *Political Analysis* 12 (4): 338–353.

Reed, William, David H. Clark, Timothy Nordstrom, and Wonjae Hwang. 2008. “War, Power, and Bargaining.” *Journal of Politics* 70 (4): 1203–1216.

Renteln, Allison D. 2013. *International Human Rights: Universalism Versus Relativism.* Quid Pro Books.

Roberts, Jason M., Steven S. Smith, and Stephen R. Haptonstahl. 2016. “The Dimensionality of Congressional Voting Reconsidered.” *American Politics Research* 44 (5): 794–815.

Roberts, Margaret E., Brandon M. Stewart, and Dustin Tingley. 2019. "Stm: An R package for structural topic models." *Journal of Statistical Software* 91 (1): 1–40.

Rodriguez, Carlos E., and Stephen G. Walker. 2014. "Label Switching in Bayesian Mixture Models: Deterministic Relabeling Strategies." *Bayesian Methodology* 23 (1): 23–45.

Russett, Bruce M. 1966. "Discovering Voting Groups in the United Nations." *American Political Science Review* 60:327–339.

Signorino, Curtis S., and Jeffrey M. Ritter. 1999. "Tau-b or Not Tau-b: Measuring the Similarity of Foreign Policy Preferences." *International Studies Quarterly* 43 (1): 115–144.

Simmons, Joseph P., Leif D. Nelson, and Uri Simonsohn. 2011. "False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant." *Psychological Science* 22 (11): 1359–1366.

Sklar, Abe. 1959. "Fonctions de répartition à n dimensions et leurs marges [Distribution Functions in n Dimensions and their Margins]." *Publications de l'Institut de Statistique de l'Université de Paris* 8:229–231.

Smith, Steven S. 2007. *Party Influence in Congress.* Cambridge University Press.

Sweeney, Kevin J. 2003. "The Severity of Interstate Disputes." *Journal of Conflict Resolution* 47 (6): 728–750.

Terman, Rochelle, and Erik Voeten. 2018. "The Relational Politics of Shame: Evidence from the Universal Periodic Review." *The Review of International Organizations* 13 (1): 1–23.

Thacker, Strom. 1999. "The High Politics of IMF Lending." *World Politics* 52 (1): 38–75.

Thibaux, Romain, and Michael I. Jordan. 2007. "Hierarchical Beta Processes and the Indian Buffet Process." In *International Conference on Artificial Intelligence and Statistics,* 564–571.

Vengroff, Richard. 1976. "Instability and Foreign Policy Behavior: Black Africa in the U.N." *American Journal of Political Science* 20 (3): 425–438.

Voeten, Erik. 2000. "Clashes in the Assembly." *International Organization* 54 (2): 185–215.

———. 2004. "Resisting the Lonely Superpower: Responses of States in the United Nations to U.S. Dominance." *Journal of Politics* 66 (3): 729–754.

———. 2013. "Data and Analyses of Voting in the UN General Assembly." Chap. 4 in *Routledge Handbook of International Organization,* edited by Bob Reinalda, 54–66. London, UK: Routledge.

Ward, Hugh, and Han Dorussen. 2016. "Standing Alongside Your Friends: Network Centrality and Providing Troops to UN Peacekeeping Operations." *Journal of Peace Research* 53 (3): 392–408.

Wilkerson, John, and Andreu Casas. 2017. "Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges." *Annual Review of Political Science* 20:529–544.

Wolford, Scott. 2014. "Power, Preferences, and Balancing: The Durability of Coalitions and the Expansion of Conflict." *International Studies Quarterly* 58 (1): 146–157.

Appendix A Autocorrelation Regression

Table 3: Linear Regression of Fisher Transformed Ideal Point Autocorrelation

	Lag 1	Lag 2	Lag 3	Lag 4	Lag 5
Num. Votes	$-0.22^{\dagger}$	-0.31^{**}	-0.24^{*}	-0.21^{*}	-0.14
	(0.13)	(0.11)	(0.11)	(0.10)	(0.11)
Num. Members	0.24^{*}	0.24^{**}	0.23^{**}	0.26^{**}	0.26^{**}
	(0.10)	(0.09)	(0.08)	(0.08)	(0.08)
Transition	-0.42^{***}	-0.53^{***}	-0.56^{***}	-0.61^{***}	-0.65^{***}
	(0.09)	(0.08)	(0.08)	(0.07)	(0.07)
N	64	64	64	64	64
R^2	0.30	0.45	0.49	0.54	0.56
adj. R^2	0.27	0.42	0.46	0.52	0.54
Resid. sd	0.22	0.20	0.18	0.18	0.18

Standard errors in parentheses

† significant at $p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

Appendix B Mathematical and Implementation Details

Subsection B.1 Generative Model

The generative model for mmBPFA is given by:

$$
\begin{aligned}
&P(x_{ij}|—) \sim \begin{cases} \mathcal{TN}_{x_{ij}^l, x_{ij}^u}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is observed} \\ \mathcal{N}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is missing} \end{cases} \\
&P(\boldsymbol{\Omega}_i) \sim \mathcal{N}_K(\boldsymbol{0}, \boldsymbol{I}_K) \\
&P(\lambda_{jk}|z_{jk}, \gamma_k^{-1}) \sim z_{jk}\mathcal{N}_p(0, \gamma_k^{-1}) + (1 - z_{jk})\delta_0 \\
&P(z_{jk}|\pi_{jk}) \sim \text{Bern}(\pi_{jk}) \\
&P(\pi_{jk}|a, b, K^+) \sim \text{Beta}(\tfrac{a}{K^+}, b(1 - \tfrac{1}{K^+})) \\
&P(a|b, K^+) \sim \text{Gamma}\Big(e + K^+, f + b\sum_{j=1}^{P} \tfrac{1}{b+j-1}\Big) \\
&P(b) \sim \text{Gamma}(2, 1) \\
&P(\gamma_k) \sim \text{Gamma}\Big(c + \tfrac{m_k}{2}, d + \sum_{j=1}^{p} \lambda_{jk}^2\Big)
\end{aligned}
\tag{11}
$$

where x_{ij}^l and x_{ij}^u are the lower and upper bounds for the ij^{th} value based on the extended rank likelihood partial ordering, K^+ is the number of active features, and $c, d, e \text{and} f$ are tunable hyperparameters.

Subsection B.2 Gibbs Sampler

The detailed sampling steps for each of the parameters are listed below. The exact Gibbs sampling procedure implemented in the `mmBPFA` package is outlined in algorithm 1.

Sample X.

For each $i \in [1, \ldots, N]$ and $j \in [1, \ldots, P]$ sample x_{ij} from a truncated normal distribution:

$$
x_{ij} \sim \begin{cases} \mathcal{TN}_{x_{ij}^l, x_{ij}^u}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is observed} \\ \mathcal{N}\big((\boldsymbol{z}_j \odot \boldsymbol{\lambda}_j)\boldsymbol{\omega}_i^\top - \boldsymbol{\alpha}_j, 1\big) & \text{if } y_{ij} \text{ is missing} \end{cases}
$$

where the lower bound $x_{ij}^l = \max\{x_{i'j} : y_{i'j} < y_{ij}\}$ and $x_{ij}^u = \min\{x_{i'j} : y_{i'j} > y_{ij}\}$. The notation $i'j$ denotes all observations in column j excluding observation i. We can precalculate and store the set of all observations less than and greater than y_{ij} for each unique value of y_{ij} in column j. Let this mask be called $\mathcal{D}$:

$$
\begin{aligned}
\mathcal{D}_{lower}(y_{ij}) &= \{y_{i\prime j} : y_{i\prime j} < y_{ij}\} \\
\mathcal{D}_{upper}(y_{ij}) &= \{y_{i\prime j} : y_{i\prime j} > y_{ij}\}
\end{aligned}
$$

This allows for faster calculation of x_{ij}^l and x_{ij}^u at each step.

Sample $\boldsymbol{Z}$ and $\boldsymbol{\Lambda}$.

First, sample currently observed features. For each $j \in [1, \ldots, P]$ and each $k \in [1, \ldots, K^+]$ first sample the element of the IBP matrix, π_{jk}, to determine when factor k is active. Let r_{jk} be the product of likelihood ratio:

$$\begin{aligned} r_{jk} &= \frac{P(z_{jk} = 1|\boldsymbol{X}, -)}{P(z_{jk} = 0|\boldsymbol{X}, -)} \\ &= \frac{P(\boldsymbol{X}|z_{jk} = 1, -)}{P(\boldsymbol{X}|z_{jk} = 0, -)} \frac{P(z_{jk} = 1)}{P(z_{jk} = 0)} \end{aligned}$$

The likelihood ratio of $\boldsymbol{X}$ is given by:

$$\frac{P(\boldsymbol{X}|z_{jk} = 1, -)}{P(\boldsymbol{X}|z_{jk} = 0, -)} = \frac{\int P(\boldsymbol{X}|\lambda_{jk}, -)\mathcal{N}(0, \gamma_k^{-1})d\lambda_{jk}}{P(\boldsymbol{X}|\lambda_{jk} = 0, -)} = \sqrt{\frac{\gamma_k}{\gamma}} \exp\Big(\frac{1}{2}\gamma\mu^2\Big)$$

where $\gamma = \boldsymbol{\omega}_k^\top \boldsymbol{\omega}_k + \gamma_k$ and $\mu = \frac{1}{\gamma}\boldsymbol{\omega}_k^\top \hat{\boldsymbol{E}}_j$ with the matrix of residuals $\hat{\boldsymbol{E}} = \boldsymbol{X} - \boldsymbol{\Lambda}\boldsymbol{\Omega} + \boldsymbol{\alpha}$ evaluated with $\lambda_{jk} = 0$. The ratio of priors is:

$$\frac{P(z_{jk} = 1)}{P(z_{jk} = 0)} = \frac{m_{-jk}}{P - m_{-jk} - 1}$$

where $m_{-jk} = -z_{jk} + \sum_{h=1}^P z_{hk}$ is the number of columns for which factor k is active.

The probability of an IBP matrix element z_{jk} being active, then, is:

$$\pi_{jk} = \frac{r_{jk}}{1 + r_{jk}}$$

Then sample $z_{jk} \sim \text{Bern}(\pi_{jk})$. If $z_{jk} = 1$, sample $\lambda_{jk} \sim \mathcal{N}(\mu, \gamma^{-1})$. Otherwise, set $\lambda_{jk} = 0$.

Second, sample new features K^*. To do so, we must first sample IBP parameters a and b. Sample a as follows:

$$a \sim \text{Gamma}\Big(e + K^+, f + H_p(b)\Big)$$

where $H_p(b) = b\sum_{j=1}^P \frac{1}{b+j-1}$. To sample b, we must use a metropolis-Hastings step. First, draw a new proposal:

$$b^* \sim \text{Gamma}(2, 1)$$

Accept b^* with probability $\min(1, t_{b\to b^*})$:

$$t_{b\to b^*} = \frac{(ab^*)^{K^+}\exp[-aH_p(b^*)]\prod_{k=1}^{K^*}\mathbb{B}(m_k, P - m_k + b^*)}{(ab)^{K^+}\exp[-aH_p(b)]\prod_{k=1}^{K^*}\mathbb{B}(m_k, P - m_k + b)}$$

where $\mathbb{B}(\cdot, \cdot)$ is the Beta function and $m_k = \sum_{j=1}^P z_{jk}$.

New features are proposed by sampling κ_j with a metropolis-Hastings step. For each $j \in [1, \ldots, P]$, sample a new dimensions potential dimensions to load on:

$$\kappa_j \sim \text{Pois}(\kappa_j; \tau\nu)$$

where $\nu = \frac{ab}{P+b-1}$ and τ is a tunable scaling parameter the Poisson rate. Knowles and Ghahramani (2011, 1541) establish a way to improve the mixing for this update that I follow. For each of the new potential dimensions, $c \in [1, \dots, \kappa_j]$ draw values for the factor loading λ_{jc}:

$$\lambda_{jc} \sim \mathcal{N}(0, 1)$$

Given this new proposal ξ^*, new dimensions are accepted with probability $\min(1, v_{\xi \to \xi^*})$:

$$\begin{aligned} v_{\xi \to \xi^*} &= v_{\mathcal{L}} \cdot v_{Pr} \\ v_{\mathcal{L}} &= (2\pi)^{N\kappa_j/2} |\boldsymbol{M}|^{-N/2} \exp\Big(\tfrac{1}{2} \textstyle\sum_i \boldsymbol{m}_i^\top \boldsymbol{M} \boldsymbol{m}_i\Big) \\ v_{Pr} &= \frac{\text{Pois}(\kappa_j; \nu)}{\text{Pois}(\kappa_j, \tau\nu)} \end{aligned}$$

where we have defined $\boldsymbol{M} = \boldsymbol{\lambda}\boldsymbol{\lambda}^\top + \boldsymbol{I}_{\kappa_j}$ and $\boldsymbol{m}_i = \boldsymbol{M}^{-1}\boldsymbol{\lambda}\hat{\boldsymbol{E}}_{\boldsymbol{ij}}$ and $\hat{\boldsymbol{E}} = \boldsymbol{X} - \boldsymbol{\Lambda}\boldsymbol{\Omega} + \boldsymbol{\alpha}$.

Remove inactive features and Normalize $\boldsymbol{\Lambda}$. For each $k \in [1, \dots, K^+]$ if $z_{jk} = 0 \forall j \in [1, \dots, P]$, remove column k from z_{jk}. Update K^+ accordingly. Now, normalize the variance of $\boldsymbol{\Lambda}$. For each $j \in [1, \dots, P]$ and $k \in [1, \dots, K^+]$:

$$\lambda_{jk} = \frac{\lambda_{jk}}{\sqrt{1 + \sum_{h=1}^{K+} \lambda_{jh}^2}}$$

Sample $\boldsymbol{\Omega}$. For each $i \in [1, \dots, N]$, sample $\boldsymbol{\omega_i}$ from:

$$\begin{aligned} \boldsymbol{\omega_i} &\sim \mathcal{N}_{K^+}(\boldsymbol{\mu}_\omega, \boldsymbol{\sigma}_\omega^2) \\ \boldsymbol{\mu}_\omega &= \boldsymbol{\sigma}_\omega^2 \boldsymbol{\Lambda}^\top \boldsymbol{x}_i \\ \boldsymbol{\sigma}_\omega^2 &= (\boldsymbol{\Lambda}^\top \boldsymbol{\Lambda} + \boldsymbol{I}_{K^+})^{-1} \end{aligned}$$

Sample $\boldsymbol{\alpha}$. For each $j \in [1, \dots, P]$, sample the item-difficulty parameter α_j from:

$$\alpha_j \sim \mathcal{N}\Big(\bar{\mu}_j, \tfrac{1}{N^2} \sum_{i=1}^{N} (\mu_{ij} - \bar{\mu}_j)^2\Big)$$

where $\mu_{ij} = \boldsymbol{\lambda}_j^\top \boldsymbol{\omega}_i - x_{ij}$ and $\bar{\mu}_j = \frac{1}{N} \sum_{i=1}^{N} \mu_{ij}$.

Sample γ_k. For each $k \in [1, \dots, K^+]$, sample the factor precision γ_k from:

$$\gamma_k \sim \text{Gamma}\Big(c + \tfrac{m_k}{2}, d + \sum_{j=1}^{p} \lambda_{jk}^2\Big)$$

where c and d are tunable hyperparameters. However, if one would like to share power across factors, then we can also sample d. Where the prior is given by $d \sim \text{Gamma}(c_0, d_0)$ and the posterior update is $d|\gamma_k \sim \text{Gamma}(c_0 + cK, d_0 + \Sigma_{k=1}^{K} \gamma_k)$.

In practice, I split the sampler up into two-phases. In the pre-sparse phase, the beta-process prior is not used and thus the sampler is allowed to mix without inducing any sparsity or sampling new features, so the number of dimensions K remains fixed. In the main phase, the beta-process prior is instantiated and the features are added or dropped as determined by by the data. McAlister (2020) finds that scheduling the sampler in this way improves mixing and decreases the chances of the MCMC becoming degenerate (converging to 0 or infinite dimensions) on smaller datasets.

Algorithm 1: mmBPFA Gibbs Sampler

```
initialize parameters:
    K⁺ ← 100
    Z ← Bern(0.9)
    Λ ← Z ⊙ N_{K⁺}(0, 1)
    α_j ← 0
    Ω ← N_{K⁺}(0, 1)
    γ_k ← 1
    d ← 100
    a, b ← 1
for t ∈ [1, . . . , T] do
    sample X
    if pre-sparse phase then
        sample Λ, Z without beta-process prior
    else
        sample Λ, Z
        Remove inactive K
        sample a, b
    end
    sample Ω, α_j, γ_k
    if share factor power then
        sample d
    end
end
```

Appendix C Structural Topic Models

Figure 14: **Resolution-Topics Associated with the First Dimension**

Figure 15: **Resolution-Topics Associated with the Second Dimension**

Topic 5

resol ia authorizing syg continue construction project extension palais des nations geneva within total new revised estimated cost 31186000

---

resol ia authorizing syg continue program major maintenance improvements palais des nations geneva provisions appropriation funds

Topic 12

res noting satisfaction addl_protocol ratified uk neths urging states treaty may become parties addl protocol sign ratify soonest order people area ques may derive benefit

---

resol_ia declaring indian_ocean designated time zone peace eliminating indian_ocean bases military manifestations calling upon littoral hinterland states indian_ocean sc mbrs others

Topic 17

res inviting admin power determine prodecures holding referendum terr permit return political exiles terr req syg send special mission spanish sahara assist impl relevant ga esp regard un parti

---

res expressing grave indignation concern maltreatment torture opponents apartheid africa req special comte follow situation order take appropriate steps promote intl compaign purpose putting end tre

Topic 43

res_req sc resume consid ques race conflict africa resulting fr policies apartheid govt afr view taking action chapt vii charter ensure complete cessation states supply arms ammunit

---

requests states concerned pending decision the_security_council effective measures andor legislation broaden scope oil embargo order ensure complete cessation supply shipping oil pe

Topic 46

retain operative para section ii condemning countries encourage incite africa persist racial policy

---

retain operative para section ii reaffirming legitimacy struggles oppressed peoples southern africa particularly colonial racial alien domination

Topic 58

situation human rights the_islamic_republic iran

---

situation human rights iraq

Topic 62

determines legislative administrative measures actions taken israel occupying power altered purported alter character status holy city jerusalem particular socalled basic law

---

determines legislative administrative measures actions taken israel occupying power altered purported alter character status holy city jerusalem particular socalled basic law

Topic 63

resol ia authorizing syg enter commitments meet unforeseen extraordinary expenses financial_year_1972 within prescribed limitations

---

res authorizing syg enter commitments meet unforseen extraordinary expenses financial_year_1973 within prescribed limitations

Figure 16: **Resolution-Topics Associated with the Third Dimension**

Appendix D UPR and Votes Scatter Plots

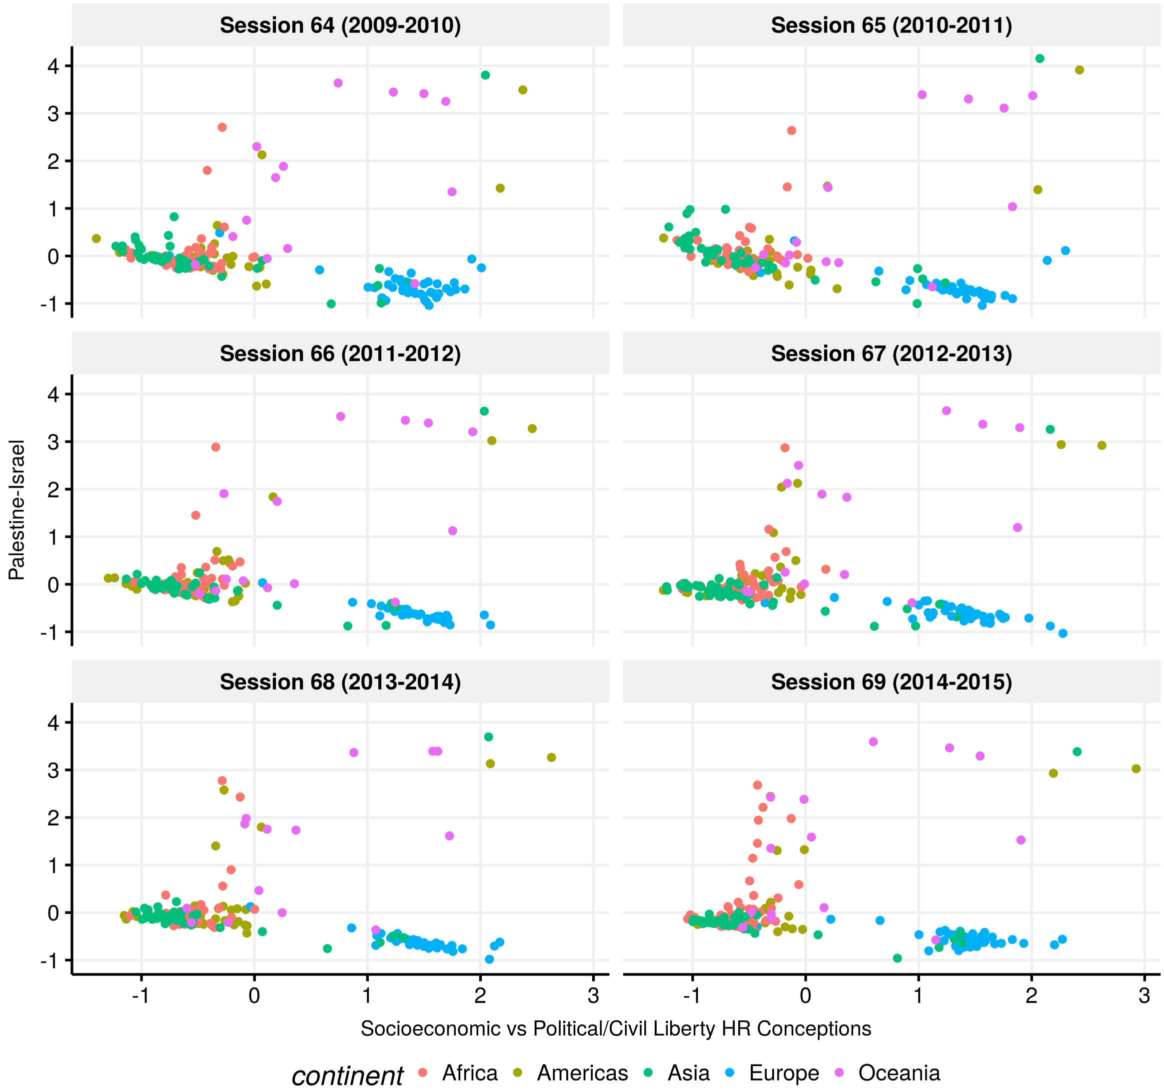

Figure 17: **First versus Third Dimension Ideal Points, 62^{nd}–69^{th} UNGA Sessions based on Votes and Universal Period Review Statements.**

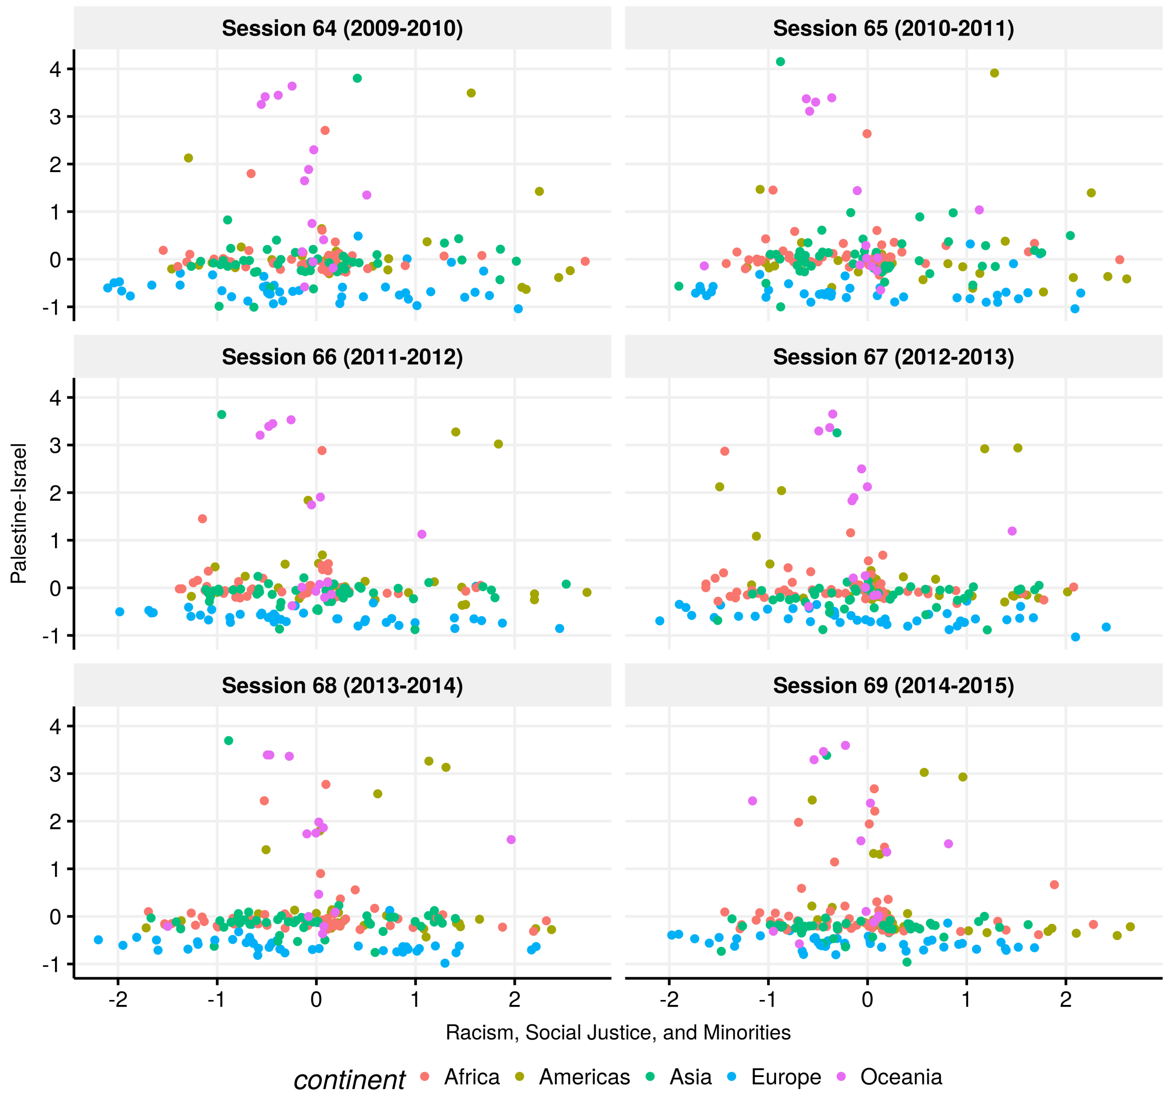

Figure 18: **Third vs Fourth Dimension Ideal Points, 62nd–69th UNGA Sessions based on Votes and Universal Period Review Statements.**

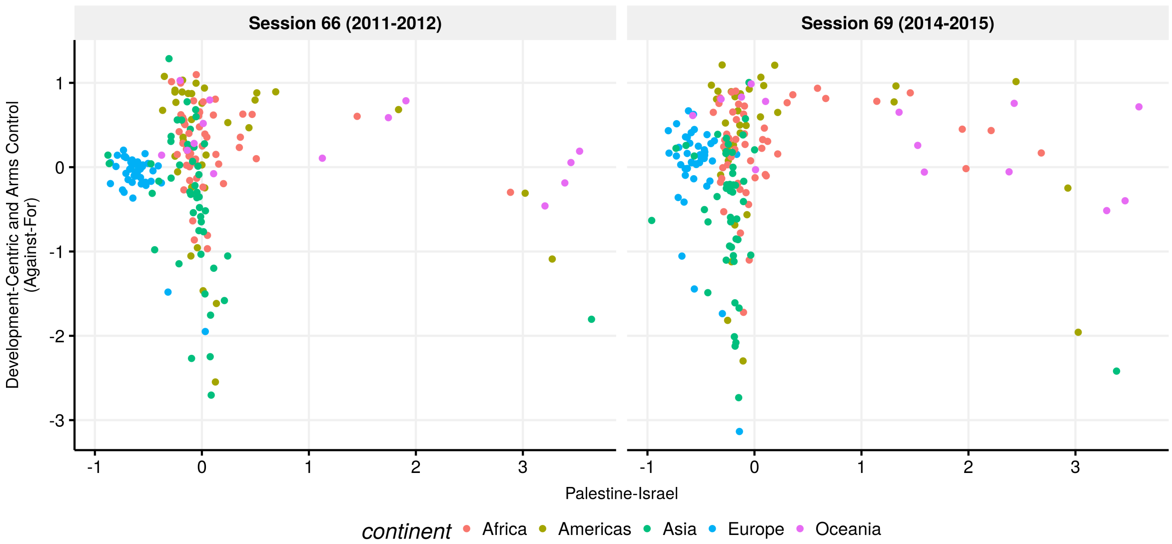

Figure 19: **Third vs Fourth Dimension Ideal Points, 62$^{\text{nd}}$–69$^{\text{th}}$ UNGA Sessions based on Votes and Universal Period Review Statements.**