Bayesian Latent Variable Models for Discrete Choice Data

by

Qiushi Yu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Political Science and Scientific Computing)
in the University of Michigan
2021

Doctoral Committee:

Professor Kevin Quinn, Chair
Professor Walter Mebane
Professor Iain Osgood
Professor Yuki Shiraito

Qiushi Yu
yuqiushi@umich.edu
ORCID iD: 0000-0003-3011-4765

DEDICATION

To Rebecca, my wife, counsel, and best friend.

ACKNOWLEDGMENTS

This dissertation would have never come into being without the help of many important people in my life. First, I would like to express my deepest gratitude to my committee chair, Professor Kevin Quinn. Professor Quinn ushered me into the challenging but fun subfield of political methodology, and he has been the lodestar in my academic journey ever since. He helped me accumulate a wide range of skills, such as programming in R, C++, and Python, machine learning, survey design, and academic writing and presentation. Throughout the process of my dissertation research and writing, he provided the imperative support by helping me identify appropriate research topics, advising me to operationalize research ideas, supporting my data collection, and aiding me in revising my dissertation. Under his mentorship, I have transcended myself and become more equipped to embrace challenges and opportunities in the future. For this reason, I am eternally indebted to him.

Next, I would like to thank the other three faculties on my committee, Professor Iain Osgood, Professor Walter Mebane, and Professor Yuki Shiraito. Professor Osgood has taught me critical mathematical skills in his class. He has also provided lots of detailed and constructive feedback on my multiple research projects. Professor Mebane encouraged me to take advantage of training opportunities on statistics and scientific computing outside the political science PhD program, and he has also provided thorough feedback on my dissertation. Professor Shiraito has served as a resourceful mentor to me and provided insightful suggestions for my research projects.

Third, I owe a tremendous debt of gratitude to Professor Qingmin Zhang at Peking University, who was my advisor throughout the undergraduate and the master's program. Since I entered college, he has served as a role model to me, and has always inspired me to an independent, innovative and hardworking scholar. He strongly encouraged me to apply to PhD programs in the United States, and has always been a source of encouragement and support throughout my PhD program.

Last, I want to thank my wife, Rebecca, for her unwavering love, staunch support, and comprehensive assistance during my PhD program. She has always held strong conviction in my potential, and has spared no effort in helping me to succeed. She has been an integral part of my life outside PhD study, which helps keep me sane and balanced. I would also like to thank my parents for their steadfast support in my pursuit of PhD study.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

ABSTRACT

This dissertation consists of three research projects that center on measuring voters' and political representatives' multi-dimensional latent preferences based on their binary opinions on public policy issues. This dissertation engages with the broad political methodology literature on pairwise comparison models, Dirichlet process mixture models, and Bayesian IRT models. I innovate statistical models to uncover partisan perceptions of information, to identify voting coalitions, and to estimate multi-dimensional latent preferences. I analyze survey data and roll call vote data, including both existing and newly-collected data. These studies advance the understanding of important political science topics, such as political biases in people's perceptions of COVID-19 related statements, politicization of human rights in the United Nations, and multiple issue dimensions of legislators' ideal points. Summaries of the three projects are detailed below.

In the first project, I propose a new multidimensional pairwise comparison model to improve the existing models in the pairwise comparison literature. Two new model specifications are proposed. The first version has a uniform prior for respondent attributes, and the second version has a Dirichlet process prior for respondent attributes. The new model allows for a richer structure of the latent attributes of the objects being compared than standard models as well as respondent-specific perceptual differences. I use the new multidimensional pairwise comparison model to analyze the survey data that I collected in the summer of 2020. This survey asked respondents to compare the truthfulness of pairs of statements about COVID-19. These statements were taken from the fact-checked statements on https://www.politifact.com. I thus have an independent measure of the truthfulness of each statement. I find that the actual truthfulness of a statement explains very little of the variability in individuals' perceptions of truthfulness. Instead, I find that the partisanship of the speaker and the partisanship of the respondent account for the majority of

the variation in perceived truthfulness, with statements made by co-partisans being viewed as more truthful.

In the second project, I advance the scholarship on politicization of humans rights within the United Nations. Previous research typically looks at simple associations between voting coalitions and observable variables, such as geographic location or membership in international organizations. My study is the first attempt at estimating the latent coalition structure based on the voting data. I propose a Bayesian Dynamic Dirichlet Process Mixture (DDPM) model to identify voting coalitions based on roll call vote data across multiple time periods. I also propose post-processing methods for analyzing the outputs of the DDPM model. I apply these methods to the United Nations General Assembly (UNGA) human rights roll call vote data from 1992 to 2017. I identify human rights voting coalitions in the UNGA after the Cold War, and polarizing resolutions that divide countries into different coalitions.

In the third project, I propose an innovative penalized EM algorithm for estimating sparse item-dimension loading structures in multidimensional IRT models. The new penalized EM algorithm identifies a sparse item-dimension loading structure for multidimensional IRT models by applying a L1 penalty on discrimination parameters in model estimation. The sparse item-dimension loading structure aids in identifying the anchoring items for each dimension and provides information on what each dimension means. In addition, the penalized EM algorithm has the flexibility to allow a discrimination parameter to be exactly zero. Hence, the penalized EM algorithm can consistently estimate the multidimensional ideal points when the data generation process is truly based on a sparse item-dimension loading structure. I first use simulation data to demonstrate how the penalized EM algorithm can accurately recover both the true item-dimension loading structure and ideal points. Then, I replicate a previous study on the 105th US senate roll call vote data to show how the penalized EM algorithm can help infer latent dimension meanings.

CHAPTER 1

Introduction

This dissertation contains three research projects that center on measuring voters' and political representatives' multi-dimensional latent preferences based on their binary opinions on public policy issues. The theme of identifying the ordinal or clustering patterns of political actors' latent preferences links the three projects into one coherent academic endeavor. We work with survey data and roll call vote data, and we develop statistical models to uncover partisan perceptions of information, identify voting coalitions, or estimate multi-dimensional latent preferences. These three studies advance our understanding of important political science topics, such as political biases in people's perceptions of COVID-19 related statements, politicization of human rights in the United Nations, and multiple issue dimensions of legislators' ideal points.

In addition to the shared substantive theme, the three projects all speak to the political methodology literature on latent variable models. Latent variable models are important empirical analysis tools in various social science fields, such as political science, psychology, educational studies, sociology, and economics. These models are critical for researchers to learn about the underlying patterns from behavioral and opinion data, and therefore make sense of actors' choices and behaviors. Specifically, the newly proposed models in this dissertation contribute to the literature on pairwise comparison models, Dirichlet process mixture models, and Bayesian IRT models. Compared with existing models, the newly proposed models relax unidimensional assumptions, account for temporal dependence of cluster numbers, and allow richer structures of latent variables.

The rest of the dissertation proceeds as follow. In the second chapter, we propose a 2-dimensional

pairwise comparison model to improve the unidimensional models in the current pairwise comparison literature. We show that the 2-dimensional pairwise comparison model provides more interpretable and informative latent estimates in a study on survey respondents' perceptions of the truthfulness of COVID-19 related statements. In the third chapter, we propose a Bayesian Dynamic Dirichlet Process Mixture (DDPM) model to identify voting coalitions based on roll call vote data across multiple time periods. We apply the new method to the United Nations General Assembly (UNGA) human rights roll call vote data from 1992 to 2017. In the fourth chapter, we propose an innovative penalized EM algorithm for estimating sparse item-dimension loading structures in multidimensional IRT models. The sparse item-dimension loading structure aids us in identifying the anchoring items for each dimension and provides information on what each dimension means. We replicate a previous study on the 105th US Senate roll call vote data to show how the penalized EM algorithm can help us infer latent dimension meanings. In the last chapter, we conclude by summarizing our findings and contributions.

CHAPTER 2

Multidimensional Pairwise Comparison Model for Heterogeneous Perceptions

(Joint Work with Kevin M. Quinn)

2.1 Introduction

Latent attribute measurement models have broad applications in political science, psychology, economics and other social science fields. The importance of latent measurement models lie in the fact that latent attributes are essential to explain actors' preferences, personalities, choices and behaviors. In political science, unobserved variables and concepts are at the core of many studies, such as ideology, preference, democracy, human rights score, and public opinions (Poole and Rosenthal, 1991; Martin and Quinn, 2002; Treier and Jackman, 2008; Fariss, 2014; Jessee, 2009). Scientific measurement of many latent variables greatly advance our understanding of the mechanisms deciding political actors' opinions, preferences and choices.

To have comparable measurement across objects or respondents, we have to assume that latent attributes are located on a common space, such as the left-right political spectrum. The ideal observed data for us to learn latent attributes from should share the same data generation process given the latent attributes. For example, if one research group assigns scores on observed domestic institutions for every country with the same benchmark, then we can believe that all the observed data are determined by a consistent latent attribute, such as democracy, on a common space (Treier

and Jackman, 2008). If we can collect data on incidents of different kinds of human rights violations, then we can also believe that countries' oppressive latent attributes (human rights score) on a common space determine these violation occurrences (Fariss, 2014). The observed data in the above two settings are "objective", because they represent either a professional agency's impartial and consistent evaluation, or historical records of incidents. Many popular models are ready for analyzing this kind of data, such as the factor analysis model, Bayesian IRT model, or multidimensional scaling model (Quinn, 2004; Clinton, Jackman and Rivers, 2004; Poole and Rosenthal, 1985).

Latent variable estimation becomes challenging when it's difficult to collect "objective" evaluations on objects' latent attributes. Perception or opinion data often fall in this category. For example, when respondents are asked to assign perceived blackness scores on a white-to-black ordinal basis to different photos of males, the resulting scores on the photos are inconsistent across respondents. This is because that a respondent's own race, life experience, and living environment would greatly affect how she perceives racial features. Therefore, it's difficult to assume that scores assigned by different respondents come from the same data generation process (Abrajano, Elmendorf and Quinn, 2018).

To avoid the inconsistent evaluations across respondents, researchers have proposed the pairwise comparison method for data collection and analysis (Thurstone, 1927; Bradley and Terry, 1952; David, 1963). It's a much easier task for respondents to compare two objects and pick the one having a greater latent attribute. In the above example, instead of asking respondents to score one photo at a time, researchers can instead ask respondents to compare a randomly selected photo pair and pick the one that in her opinion has a greater latent attribute. This way, researchers are able to collect perception data that are determined by attributes in a common space. The downside is that researchers can only collect dichotomous data with this method. Pairwise comparison models aid researchers in recovering the locations of objects on a continuous scale in a latent common space.

The pairwise comparison model was originally proposed in psychometrics. Researchers in

psychology, marketing studies, and information science have proposed various specifications of the pairwise comparison model to analyze different data, such as ranking data and pairwise comparison data by multiple respondents (Stark, Chernyshenko and Drasgow, 2005; Wang et al., 2017; Kim, Kim and Shim, 2017). The most common pairwise comparison model assumes a unidimensional latent attribute for objects (Thurstone, 1927; Bradley and Terry, 1952). Other researchers have proposed a unidimensional respondent attribute to account for respondents' different levels of sensitivity (Carlson and Montgomery, 2017). There are also attempts to generalize the pairwise comparison models to multidimensional latent spaces. The current multidimensional pairwise comparison models lack interpretabiltiy, because there are insufficient constraints in model specification. As a result, it's difficult to learn how different dimensions are aggregated in the data generation process. In addition, the existing multidimensional models also assume that every respondent must have unique respondent-specific parameters, and are thus unable to identify groups of respondents that share the same respondent-specific parameters. To fill this lacuna in the literature, we propose a new multidimensional pairwise comparison model in this chapter.

The proposed model assigns a unit-length positive vector to each respondent. This constraint guarantees that a respondent attribute is a weight vector, whose elements are non-negative and sum up to 1. This constraint makes it explicit to model how object attributes of different dimensions are aggregated and in turn determine respondents' choices. Compared with previous models, the proposed model has an advantage in model identification and interpretability due to the innovative constraint. Moreover, in the second version of the new model, we add a Dirichlet Process prior on respondent-specific parameters, so we can flexibly cluster respondents into groups that share similar perceptions. This innovation gives us leverage to learn the grouping structure among respondents, which can potentially advance our understanding of how latent attributes are perceived differently. Furthermore, this also allows us to make inferences on which characteristics of respondents affect their perceptions on the latent attribute in question.

We use the new model to analyze survey data that we collected in the summer of 2020. This survey asked respondents to compare the truthfulness of pairs of statements about COVID-19. Ap-

plying our new model to our survey data, We find that the objective truthfulness of a statement only weakly correlates with respondents' perceptions of truthfulness. On the other hand, we find strong correlations between the political valence of the statements and their perceived truthfulness, moderated by the political leaning of a respondent. Statements made by a co-partisan of a respondent tend to be viewed as more truthful by this respondent.

A sizable fraction of respondents gauge the truthfulness of COVID-19 statements through partisan lenses. For these respondents, partisanship has a stronger impact on their responses than does the actual truthfulness of the statements. Indeed, the responses from the most right-leaning respondents are negatively correlated with the objective truthfulness of the statements. That said, a plurality of respondents are relatively unswayed by partisanship but have a difficult time accurately gauging the truthfulness of COVID-19 statements. We also observe associations between the respondent-specific perceptual parameters and public-health-relevant behaviors, such as mask wearing and social distancing.

The rest of the chapter proceeds as follow. First, we review the assumptions and specifications of previous pairwise comparison models, before we detail the new multidimensional pairwise comparison model. Second, we specify the MCMC samplers for the new model in detail. Third, we show simulation studies on the new model to illustrate how the new model can accurately recover the true latent variable values. Fourth, we apply the new model to the pairwise comparison data collected in the survey. We report and compare the analysis results based on both existing unidimensional models and the newly-proposed multidimensional model. We conclude by discussing our findings and plans for future research.

2.2 A New Model for Pairwise Comparisons Data with Heterogeneous Perceptions

Pairwise comparisons provide an important method of measuring latent attributes in the social sciences. Pairwise, relative judgments are cognitively easier for human respondents than other

types of assessments, such as those based on Likert-type scales (Oishi et al., 2005; Phelps et al., 2015; Dittrich et al., 2007). Further, the focus on pairwise, relative judgments eliminates the possibility that respondents may use a Likert-type scale differently (see, for instance, Bachman and O'Malley (1984); Brady (1985); Suchman and Jordan (1990) and King et al. (2004)).

Traditional models for pairwise comparisons assume unidimensional latent attributes for objects (Thurstone, 1927; Bradley and Terry, 1952; David, 1963). Some researchers add a unidimensional respondent-specific parameter to account for respondents' different levels of ability or sensitivity (Carlson and Montgomery, 2017). There have also been attempts to generalize pairwise comparison models to multidimensional latent spaces (Carroll and De Soete, 1991; Yu and Chan, 2001; Balakrishnan and Chopra, 2012). In this section, we briefly review the existing pairwise comparison models, before we introduce our new model.

2.2.1 Existing Models

We start by reviewing unidimensional pairwise comparison models. Consider a set of J objects $\{o_j\}_{j=1}^{J}$. We assume that each o_j has a latent attribute $\theta_j \in \mathbb{R}$ that denotes an attribute of interest. While $\boldsymbol{\theta}$ is unobserved, we do observe $y_{ijj'}$, the result of a paired comparison of o_j and $o_{j'}$ by respondent i, in which i is asked to make a ranking judgment as to whether o_j or $o_{j'}$ has a larger value of the latent attribute. $y_{ijj'}$ is equal to 1, if respondent i judges o_j to have a larger value of the latent attribute in question than $o_{j'}$, 0 otherwise. More specifically, we assume

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = F(\theta_j - \theta_{j'})$$

where $F(\cdot)$ is a cumulative distribution function (CDF). If $F(\cdot)$ is the CDF of a standard normal distribution, the model above is the Thurstone model, the first pairwise comparison model ever proposed (Thurstone, 1927). If $F(\cdot)$ is the CDF of a logistic distribution, the model above is the Bradley-Terry model (Bradley and Terry, 1952).

If we assume that respondents make conditionally independent judgments of the objects, the likelihood for $\boldsymbol{\theta}$ can be written as the product of Bernoulli probability mass functions. A variant of the above model is to assume that respondents vary in their ability or sensitivity to discern the latent differences between objects, but the latent object attributes remain on the real line, i.e. $\theta_j \in \mathbb{R}$ for $j = 1, \ldots, J$. Here

$$p_{ijj'} = F(\beta_i [\theta_j - \theta_{j'}])$$

Typically, it is assumed that $\beta_i \in \mathbb{R}_+$ for $i = 1, \ldots, N$ (Carlson and Montgomery, 2017). An object pair is compared by multiple respondents, and a respondent does multiple tasks of pairwise comparisons. Therefore, we are able to estimate the two sets of latent variables, object attributes and respondent attributes, iteratively. The intuition is that we are able to update the object attributes conditioning the current values of respondent attributes, and vise versa. Fully Bayesian methods based on Markov chain Monte Carlo (MCMC) or the EM algorithm have both been proposed to estimate these unidimensional models (Böckenholt and Tsai, 2001; Johnson and Kuhn, 2013).

Political scientists have applied the above unidimensional pairwise comparison models to various projects, such as solving the coder inconsistency problem in text data coding tasks (Carlson and Montgomery, 2017), gauging legislators' degree of "grandstanding" in speech (Park, 2021), and measuring the sophistication of political text (Benoit, Munger and Spirling, 2019). Similar unidimensional pairwise comparison models are also popular in personality measurement research[1] (Stark, Chernyshenko and Drasgow, 2005; Wang et al., 2017). An alternative specification of the unidimensional pairwise comparison model is latent ranking estimation models, which focus on probabilistically estimating rankings of objects instead of specific object attributes. The assumption for these models is that there could be multiple latent rankings governing the pairwise comparison data generation process. The goal of these models is to learn the ranking that is the most likely. Similar latent ranking models are widely applied in the area of information sciences

[1]These tests are called multidimensional personality test. However, the word "multidimensional" does not refer to latent dimensions, but refer to an aspect of personality captured by a group of specially designed objects. A unidimensional pairwise comparison model is used to measure an individual "dimension" of personality with a group of pertinent objects. Therefore, the underling measurement model is still undimensional for an aspect of personality.

and recommendation systems. For example, Kim, Kim and Shim (2017) assume a categorical distribution among all the possible rankings among objects, and use a generative model to learn which overall ranking is the most likely (Kim, Kim and Shim, 2017).

As a variant of the unidimensional pairwise comparison model, the Thurstonian Factor Model further decomposes the J object attributes by the product of a loading matrix and a vector of multidimensional common factors (Maydeu-Olivares and Böckenholt, 2005; Maydeu-Olivares and Brown, 2010). We illustrate the model setup for only three unique objects, $J = 3$. Each unique object pair is compared only once, so the probability parameters for all the unique pairwise comparison tasks are modeled as below.

$$\underbrace{\boldsymbol{p}}_{\frac{J(J-1)}{2}\times 1} = F(\underbrace{\boldsymbol{A}}_{\frac{J(J-1)}{2}\times J}\underbrace{\boldsymbol{\theta}}_{J\times 1})$$

$$\underbrace{\boldsymbol{\theta}}_{J\times 1} = \underbrace{\boldsymbol{\mu}}_{J\times 1} + \underbrace{\boldsymbol{\Lambda}}_{J\times d}\underbrace{\boldsymbol{\xi}}_{d\times 1} + \underbrace{\boldsymbol{\epsilon}}_{J\times 1}$$

$$\boldsymbol{A} = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & -1 \end{bmatrix}, \text{ if } J = 3$$

where $\boldsymbol{p} \in [0,1]^{\frac{J(J-1)}{2}}$ is a $\frac{J(J-1)}{2}$-length probability vectors, representing the probability parameters for all the unique pairwise comparison tasks. $\boldsymbol{\theta}$ is the latent attribute vector for object $1, 2, \cdots, J$. $\boldsymbol{A}$ is a linear transformation matrix that helps to compute the differences between all object attribute pairs. We can easily construct $\boldsymbol{A}$ for other J values by enumerating all possible unordered unique object attribute pair differences by row. $F(\cdot)$ is the link function that maps an element in the $\boldsymbol{A\theta}$ vector to the corresponding element in vector $\boldsymbol{p}$. The first equation above is the matrix representation of the unidimensional pairwise comparison model. The second equation represents a factor analysis model to rewrite the vector of object attributes as the sum of the mean vector, $\boldsymbol{\mu}$, and the product of a loading matrix $\boldsymbol{\Lambda}$ and a common factor vector $\boldsymbol{\xi}$, conditioning on

$d < J$. The Thurstonian Factor Model serves as an exploratory method to show the common latent factors that generate a specific object's latent attribute.[2]

A limitation of the above unidimensional pairwise comparison models is that they either assume no perceptual differences between the individuals making the comparative judgments, or assume that individuals only vary in the ability to discern the signal from a common unidimensional comparison of latent attributes. However, the unidimensional assumption is too strong when respondents have to compare objects on more than one latent dimension, and they have to aggregate objects' attributes on different latent dimensions to make their choices. Moreover, when individuals are asked to compare objects on a general attribute that is itself made up of component attributes, they may vary in the way they construct the general attribute from the component attributes. For instance, it is common for individuals to be asked to rate the "racial stereotypicality" of photographs (Eberhardt et al., 2006). Yet one might reasonably think that racial stereotypicality is composed of multiple sub-attributes, such as skin color, hair texture, face shape, and so on. Further, one might also think that individuals differ in the weights they place on these sub-attributes when asked to make comparisons between photos on the general attribute of racial stereotypicality (Abrajano, Elmendorf and Quinn, 2018).

In addition to the unidimensional models, researchers have also proposed multidimensional pairwise comparison models (Cattelan, 2012). In multidimensional models, both objects and respondents are assumed to have locations on a common space. What determines a respondent's choice between objects is the relative distance between the respondent location to the two objects' locations. A respondent will prefer an object that's closer to her own latent location than an object that's far away.

Two methods for estimating multidimensional object and respondent attributes are originally proposed: the wandering vector method and the wandering ideal point method (Carroll and De Soete,

[2]For example, we ask employees to rank all unique pairs of nine work environment motivations: 1. Supportive Environment. 2. Challenging Work. 3. Career Progression. 4. Ethics. 5. Personal Impact. 6. Personal Development. 7. Social Interaction. 8. Competition. 9. Work Security. We use the Thurstonian Factor Model to learn a single common factor and a 9×1 loading matrix, the product of which gives us object attributes. This is equivalent to a unidimensional pairwise comparison model. We can also reduce the nine objects' attributes to two or three common factors. For more details, see Maydeu-Olivares and Brown (2010).

1991). These two methods use two common ways to compute distance in a multidimensional space. The wandering vector method uses the dot product to compute distance, whereas the wandering ideal point model uses a method similar to Euclidean distance. The wandering vector method provokes more innovations of multidimensional pairwise comparison models later, so we focus on this specification. The probability parameter for a pairwise comparison task in a d-dimensional wandering vector model is modeled as follows:

$$p_{ijj'} = F(\boldsymbol{\beta}_i \cdot \boldsymbol{\theta}_j - \boldsymbol{\beta}_i \cdot \boldsymbol{\theta}_{j'})$$

where $\boldsymbol{\beta}_i \in \mathbb{R}_+^d$ and $||\boldsymbol{\beta}_i|| = 1$ for $i = 1, \ldots, N$, and $\boldsymbol{\theta}_j \in \mathbb{R}^d$. A respondent attribute is a unit-length vector with non-negative elements. The distance between a respondent location and an object location is the dot product between the two vectors. No estimation method was provided for the wandering vector method when it was originally proposed.

A Bayesian sampler is proposed for the wandering vector method on analyzing full ranking data on a group of objects evaluated by multiple respondents (Yu and Chan, 2001). This Bayesian sampler uses a normal prior for all $\boldsymbol{\beta}_i$: $\boldsymbol{\beta}_i \sim N_d(\mathbf{1}, \boldsymbol{I}_d)$. A similar Bayesian model is proposed to analyze pairwise movie preference data by multiple users (Balakrishnan and Chopra, 2012). This model also uses a multivariate normal prior for respondent vectors. These Bayesian samplers do not constrain a respondent vector as a unit-length non-negative vector, but use a multivariate normal prior for all $\boldsymbol{\beta}_i$, such as $\boldsymbol{\beta}_i \sim N_d(\mathbf{1}, \boldsymbol{I}_d)$. These weaker constraints on respondent attributes greatly hurt the interpretabiltiy of the model, because the dot product of a respondent vector and an object vector is no longer the weighted average of attributes on different dimensions.

2.2.2 A New Multidimensional Model

The unidimensional pairwise comparison models discussed above have important limitations. They either assume no perceptual differences between respondents, or they assume that respondents only vary in the ability or sensitivity to discern the object attribute differences. Moreover, the unidimen-

sional attribute assumption is overly strong when respondents evaluate objects on more than one latent dimension. Further, respondents may differentially weight the attributes that correspond to different latent dimensions.

Existing multidimensional pairwise comparison models are difficult to intepret due to their lack of constraints on the respondent-specific parameters. When respondent-specific parameters are not constrained to be unit-length non-negative vectors, these parameters cannot be easily viewed as dimension-specific weights. In addition, the existing models do not allow for any clustering among the respondent-specific parameters that would represent shared perceptual frameworks among respondents. To address these issues, we propose a new multidimensional pairwise comparison model. We detail two versions of this model—each corresponding to a different prior distribution for the respondent-specific parameters.

In this new model, we operationalize a unit-length weight vector for each respondent with trigonometric functions. This allows us to model a respondent's perception of an object as the weighted average of the object's attributes on each latent dimension. The model therefore allows researchers to estimate how multiple latent sub-attributes are aggregated into a general latent attribute, and to assess the extent to which respondents differ in their construction of the general attribute from the sub-attributes.

In the first version of the model we assume a uniform prior for these respondent-specific parameters. In the second version, we assume a Dirichlet Process prior on the respondent-specific parameters. This second model allows researchers to learn how perceptual frameworks cluster among respondents and how various respondent characteristics relate to respondent perceptions of the latent attributes in interest.

We begin with the special case of a two-dimensional latent attribute space. Once again, consider a set of J objects $\{o_j\}_{j=1}^{J}$. However, we now assume each o_j has latent attributes that can be represented by a location in two-dimensional Euclidean space: $\boldsymbol{\theta}_j \in \mathbb{R}^2$. We assume that respondents differ in the weights they place on each of these two dimensions. More specifically, respondent i's judgment depends on a unit-vector $\mathbf{g}(\gamma_i) \equiv (\cos(\gamma_i), \sin(\gamma_i))^T$ with $\gamma_i \in [0, \frac{1}{2}\pi]$ in

the following way:

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi_1(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i))$$

where $\Phi_1(\cdot)$ is the CDF of a univariate standard normal distribution, and $\cdot$ denotes the dot product between two vectors. Intuitively, respondent i projects the latent attributes onto $\mathbf{g}(\gamma_i)$ and then uses the signed distance between the projected points to compare two objects. This is depicted graphically in Figure 2.1.

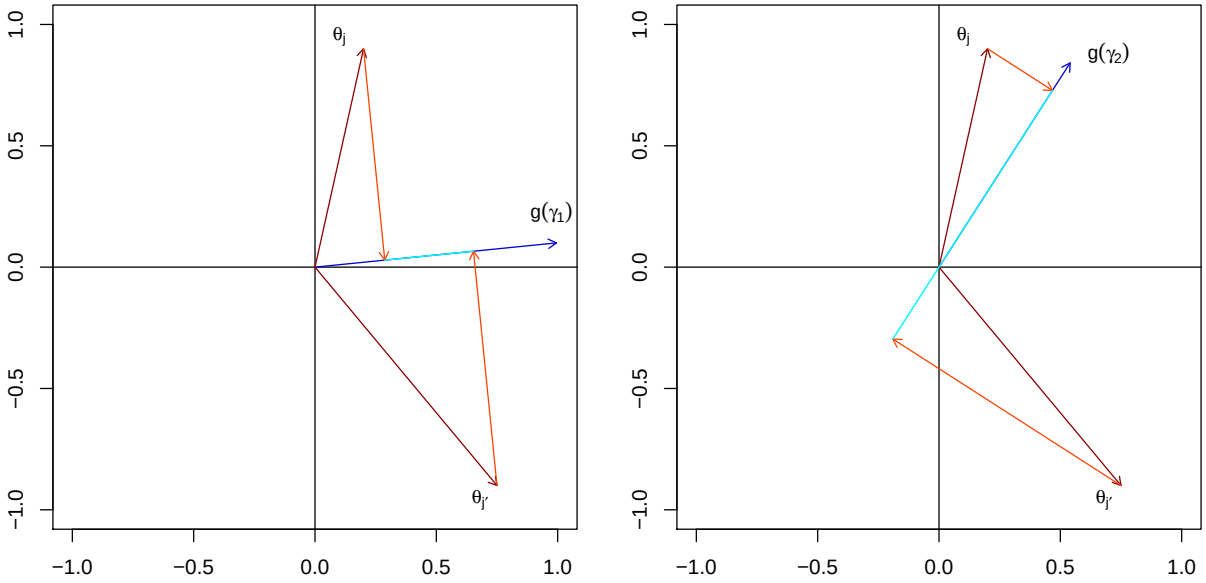

Figure 2.1: *Example of the Two-Dimensional Latent Attribute Model.* In the left panel, respondent 1 places much more emphasis on dimension 1 (the horizontal dimension). As a result, this individual is slightly more likely to evaluate $o_{j'}$ as being preferred to o_j. In the right panel, respondent 2 gives weight to both of the latent dimensions with slightly more weight placed on dimension 2 (the vertical dimension). As a result, this individual is much more likely to evaluate o_j as being preferred to $o_{j'}$.

A semi-conjugate prior distribution for $\boldsymbol{\theta}_j$ is:

$$\boldsymbol{\theta}_j \overset{iid}{\sim} \mathcal{N}_2(\mathbf{0}, \mathbf{I}_2) \quad j = 1, \ldots, J.$$

A number of priors for $\boldsymbol{\gamma}$ are reasonable. We consider two, and introduce the two versions of

the new model in the following subsections.

2.2.2.1 Uniform Prior

The most intuitive prior on γ_i is the uniform prior:

$$\gamma_i \overset{iid}{\sim} \mathcal{U}nif(0, \tfrac{1}{2}\pi) \quad i = 1, \ldots, N.$$

This specification has the advantage of simplicity but it does not allow for the possibility that γ_i is equal to $\gamma_{i'}$ for any two individuals i and i'. Such grouping may be desirable if we are interested in making inferences about the extent to which respondents share the same perceptual framework for evaluating the latent attributes in question. Further, allowing γ_i to equal $\gamma_{i'}$ with positive probability is also useful in situations where respondents only rate a small-to-moderate number of paired comparisons. In these situations, allowing some form of clustering among the γ parameters will lower the variance of the resulting estimates of the γ parameters.

2.2.2.2 Dirichlet Process Prior

An alternative is to assume that each γ_i is drawn from a distribution G that is itself drawn from a Dirichlet process. More formally,

$$\gamma_i \overset{iid}{\sim} G \quad i = 1, \ldots, N$$

$$G \sim \mathcal{DP}(\alpha G_0)$$

where $\alpha \in \mathbb{R}_+$ is a concentration parameter and G_0 is the centering distribution, which is specified as $\mathcal{U}nif(0, \frac{1}{2}\pi)$. α could be either fixed at a constant value or given a prior distribution and estimated. If α is to be estimated, then we assume a Gamma prior distribution for α:

$$\alpha \sim \mathcal{G}amma(a, b)$$

While the Dirichlet process prior for $\boldsymbol{\gamma}$ complicates estimation, it has the advantage of allowing for the possibility of perceptual clustering among respondents.

The new models can be generalized to $d > 2$ dimensions by assuming that each $\boldsymbol{\theta}_j \in \mathbb{R}^d$, and the perceptual unit-vectors $\mathbf{g}(\boldsymbol{\gamma}_i)$ are constrained to lie in the positive orthant of $\mathbb{R}^d$. $\boldsymbol{\gamma}_i$ would be $(d-1)$-dimensional in this case with each element being an angle in the positive orthant. For example, if $d = 3$, then we can use $[\gamma_{i1} \quad \gamma_{i2}]^T$ to represent a unit-length vector in the positive orthant, $[\cos(\gamma_{i1})\cos(\gamma_{i2}), \quad \cos(\gamma_{i1})\sin(\gamma_{i2}), \quad \sin(\gamma_{i1})]^T$. The MCMC algorithms we use for fitting these two versions of the model are discussed in the next section.

2.3 Markov Chain Monte Carlo Algorithm

In this section, we detail the samplers for the two versions of the new model.

2.3.1 Sampler for the First Version of the New Model

The sampler for the first version of the new model consists of a Gibbs sampler component and a random walk Metropolis-Hastings (MH) sampler component. We use the Gibbs sampler to sample $\boldsymbol{\theta}_j$'s and the augmented parameters, $y^*_{ijj'}$'s. We use the random walk MH sampler to sample γ_i's.

2.3.1.1 Sample $y^*_{ijj'}$

We use the data augmentation method in Bayesian statistics to replace the binary choice data points with continuous values (Albert and Chib, 1993, 1995). A binary choice data point has the following Bernoulli distribution:

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi_1(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i))$$

We define a continuous latent attribute difference, $y^*_{ijj'}$, to correspond every binary data point, $y_{ijj'}$.

We use $\epsilon_{ijj'} \sim N(0,1)$ to denote an i.i.d error term. Therefore, $y_{ijj'} = 1$ is equivalent to a positive attribute difference between object j and object j' for respondent i: $y^*_{ijj'} = \boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i) + \epsilon_{ijj'} > 0$. Likewise, $y_{ijj'} = 0$ is equivalent to a negative attribute difference between object j and object j' for respondent i: $y^*_{ijj'} = \boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i) + \epsilon_{ijj'} < 0$. Without loss of generality, we impose a truncated standard normal distribution on $y^*_{ijj'}$. The sign of $y^*_{ijj'}$ must be equal to the sign of the corresponding $y_{ijj'} - \frac{1}{2}$. Given the current values of $\boldsymbol{\theta}_j$,$\boldsymbol{\theta}'_j$ and γ_i, we can sample $y^*_{ijj'}$ from the following distribution:

$$y^*_{ijj'} \sim \begin{cases} \mathcal{N}\Big(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i), 1\Big)\mathbb{I}(y^*_{ijj'} > 0), & \text{if } y_{ijj'} = 1 \\ \mathcal{N}\Big(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i), 1\Big)\mathbb{I}(y^*_{ijj'} < 0), & \text{if } y_{ijj'} = 0 \end{cases}$$

where $\mathcal{N}\Big(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i), 1\Big)\mathbb{I}(y^*_{ijj'} > 0)$ is a univariate truncated normal distribution, which only takes positive values. Similarly, $\mathcal{N}\Big(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i), 1\Big)\mathbb{I}(y^*_{ijj'} < 0)$ is a univariate truncated normal distribution, which only takes negative values.

2.3.1.2 Sample $\boldsymbol{\theta}_j$

For sampling the values of $\boldsymbol{\theta}_j$, we need to do careful bookkeeping of all the pairwise comparison tasks that involve object j. Let's use M_j to denote the number of all the unique comparison tasks involving object j. Accordingly, we need to record the sign of $\boldsymbol{\theta}_j$, the counterpart object attribute $\boldsymbol{\theta}_{j'}$, the respondent attribute γ_i, and the augmented parameter $y^*_{ijj'}$ or $y^*_{ij'j}$ (depending on which side object j shows in this task) in each one of the M_j tasks.

We use the rows of matrix $\tilde{\boldsymbol{\Theta}}$ to store the counterpart object attribute $\boldsymbol{\theta}_{j'}$'s in all the comparison tasks involving object j. We use vector $\tilde{\boldsymbol{\gamma}}_j$ to store the respondent attribute γ_i's in all the comparison tasks involving object j. We use vector $\tilde{\boldsymbol{y}}^*_j$ to store the augmented parameter y^*'s in all the comparison tasks involving object j. When object j shows up in a comparison task, it either shows as the left-side choice or right-side choice. For a unique comparison task, we denote the sign of object j as $+1$ if it's the left-side option, or as -1 if it's the right-side option. We use vector $\boldsymbol{s}_j$ to

store the signs of object j in all the comparison tasks involving object j. These four containers are filled in a specific order so that the m'th element or row of the four containers correspond to the parameters for the m'th comparison task involving object j.

To illustrate the relationship between $\boldsymbol{\theta}_j$, $\tilde{\boldsymbol{\Theta}}_j$, $\tilde{\boldsymbol{\gamma}}_j$, and $\tilde{\boldsymbol{y}}_j^*$, we write out the equation representing the m'th comparison involving object j.

$$\boldsymbol{\theta}_j \cdot \left(\boldsymbol{s}_j[m] \times \mathrm{g}(\tilde{\boldsymbol{\gamma}}_j[m])\right) + \text{error term} = \tilde{\boldsymbol{y}}_j^*[m] + \left(\boldsymbol{s}_j[m] \times \left(\tilde{\boldsymbol{\Theta}}_j[m] \cdot \mathrm{g}(\tilde{\boldsymbol{\gamma}}_j[m])\right)\right)$$

where $[m]$ indicates the m's element of a vector or the m's row of a matrix, and the error term has an i.i.d standard normal distribution.

We can write a similar equation for every comparison task involving object j. The left hand side of the equation is a dot product of $\boldsymbol{\theta}_j$ and another vector, and the right hand side of the equation is a scalar. For $m = 1, 2, \cdots, M_j$, we repeat the same algebraic manipulation in the above equation. We compute $\boldsymbol{s}_j[m] \times \mathrm{g}(\tilde{\boldsymbol{\gamma}}_j[m])$, and store the resulting vector in the m'th row of matrix $\boldsymbol{X}_j$. We compute $\tilde{\boldsymbol{y}}_j^*[m] + \left(\boldsymbol{s}_j[m] \times \left(\tilde{\boldsymbol{\Theta}}_j[m] \cdot \mathrm{g}(\tilde{\boldsymbol{\gamma}}_j[m])\right)\right)$, and store the resulting value in the m'th element of vector $\boldsymbol{z}_j$. Then we can express the distribution of $\boldsymbol{z}_j$ with the multidimensional normal distribution below.

$$\underbrace{\boldsymbol{z}_j}_{M_j \times 1 \text{ vector}} \sim \mathcal{N}_{M_j}\Big(\underbrace{\boldsymbol{X}_j}_{M_j \times 2 \text{ matrix}} \times \underbrace{\boldsymbol{\theta}_j}_{2 \times 1 \text{ vector}}, \mathbf{I}_{M_j}\Big)$$

$\boldsymbol{\theta}_j$ has a semi-conjugate bivariate normal prior distribution.

$$\boldsymbol{\theta}_j \sim \mathcal{N}_2(\mathbf{0}, \mathbf{I}_2)$$

Then, we are able to derive the conditional posterior of $\boldsymbol{\theta}_j$ as follows:

$$\boldsymbol{\theta}_j | \boldsymbol{X}_j, \boldsymbol{z}_j \sim \mathcal{N}_2\left(\left(\boldsymbol{X}_j^T \boldsymbol{X}_j + \mathbf{I}_2\right)^{-1} \boldsymbol{X}_j^T \boldsymbol{z}_j, \quad \left(\boldsymbol{X}_j^T \boldsymbol{X}_j + \mathbf{I}_2\right)^{-1}\right)$$

2.3.1.3 Sample γ_i

For sampling the values of γ_i, we need to do careful bookkeeping of all the pairwise comparison tasks that involve respondent i. Let's use M_i to denote the number of all the unique comparison tasks involving respondent i. Accordingly, we need to record the left-side object attribute $\boldsymbol{\theta}_j$, the right-side object attribute $\boldsymbol{\theta}_{j'}$, and the augmented parameter $y^*_{ijj'}$ in each one of the M_i tasks.

We use the rows of matrix $\tilde{\boldsymbol{\Theta}}_i$ to store the left-side object attribute $\boldsymbol{\theta}_j$'s in all the comparison tasks involving respondent i. We use the rows of matrix $\tilde{\boldsymbol{\Theta}}'_i$ to store the right-side object attribute $\boldsymbol{\theta}_{j'}$'s in all the comparison tasks involving respondent i. We use vector $\tilde{\boldsymbol{y}}^*_i$ to store the augmented parameter y^*'s in all the comparison tasks involving respondent i. These three containers are filled in a specific order so that the m'th element or row of the three containers correspond to the parameters for the m'th comparison task involving respondent i.

Given the current values of $\gamma_i^{(t)}$, $\tilde{\boldsymbol{\Theta}}_i$, and $\tilde{\boldsymbol{\Theta}}'_i$, the density function of $\tilde{\boldsymbol{y}}^*_i$ is the product of M_i normal distribution densities:

$$\mathcal{L}(\tilde{\boldsymbol{y}}^*_i|\gamma_i^{(t)}, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}'_i) = \prod_{m=1}^{M_i} \phi_1(\tilde{\boldsymbol{y}}^*_i[m]; \quad \tilde{\boldsymbol{\Theta}}_i[m] \cdot \mathbf{g}(\gamma_i^{(t)}) - \tilde{\boldsymbol{\Theta}}'_i[m] \cdot \mathbf{g}(\gamma_i^{(t)}), 1)$$

where $[m]$ indicates the m'th element of a vector or the m'th row of a matrix, and $\phi_1(\cdot; \mu, \sigma^2)$ is the PDF of a univariate normal distribution with mean μ and variance σ^2.

We use a random walk MH sampler to sample γ_i, and we sample each γ_i separately for $i = 1, 2, \cdots, I$. We generate a random walk step, τ, from a uniform distribution, $\tau \sim \mathcal{U}nif(-\delta, \delta)$. δ is the positive tuning parameter that determines the accepting rate of the random walk MH sampler. (Chib and Greenberg, 1995*a*) The proposed new value of γ_i is $\gamma_i^{(t+1)} = \gamma_i^{(t)} + \tau$. We plug $\gamma_i^{(t+1)}$ in the density function, and get $\mathcal{L}(\tilde{\boldsymbol{y}}^*_i|\gamma_i^{(t+1)}, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}'_i)$.

Due to the uniform prior on γ_i, the acceptance ratio, r, is determined only by the ratio of the likelihoods.

$$r = \min\left(1, \frac{\mathcal{L}(\tilde{\boldsymbol{y}}^*_i|\gamma_i^{(t+1)}, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}'_i)}{\mathcal{L}(\tilde{\boldsymbol{y}}^*_i|\gamma_i^{(t)}, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}'_i)}\right)$$

With probability r, we accept the proposed new $\gamma_i^{(t+1)}$, and with probability $1-r$, we reject it.

2.3.2 Sampler for the Second Version of the New Model

The sampler for the second version of the new model shares the same steps for sampling $y^*_{ijj'}$ and $\boldsymbol{\theta}_j$ in the first version. We only introduce the rest of the steps in the sampler for the second version of the new model, given the current values of $y^*_{ijj'}$'s and $\boldsymbol{\theta}_j$'s. We assume a Dirichlet process prior on γ_i. Before specifying the sampler for γ_i, we compare two approaches to implement a Dirichlet process Mixture model: the collapsed sampler and the blocked Gibbs sampler (Müller, Rodriguez et al., 2013).

The collapsed sampler approach analytically computes the probability for assigning a unit to a cluster by integrating out the parameters characterizing each cluster (Ferguson, 1973; Escobar, 1994; MacEachern, 1994; Escobar and West, 1995; MacEachern and Müller, 1998; Neal, 2000). This approach works well with conjugate priors, and cleverly uses the integral trick to account for infinite values of the cluster parameters when deciding a cluster assignment probability. The collapsed sampler has wide applications in various fields. The limitation of the collapsed sampler lies in the relative difficulty for it to work with non-conjugate priors.

The blocked Gibbs sampler has its theoretical foundation in the stick-breaking process reparameterizaton of the Dirichlet process (Sethuraman, 1994; Ishwaran and Zarepour, 2000; Ishwaran and James, 2001). The blocked Gibbs sampler further simplifies the sampling procedure by assuming a finite number of candidate clusters to start with (Müller, Rodriguez et al., 2013). There are other augmented variables in the blocked Gibbs sampler to facilitate large clusters to grow larger and small clusters to disappear. Even if we assume a large number of finite candidate clusters at the beginning, the block Gibbs sampler will eventually converge to a small number of clusters as the MCMC mixes. Therefore, the block Gibbs sampler represents a close and efficient approximation to the original Dirichlet process with infinite candidate clusters.

Due to the non-conjugate prior employed on γ_i, we use the block Gibbs sampler for sampling γ_i in the second version of the new model. In this subsection, we first specify the Dirichlet process

prior on γ_i. Then we introduce the sampling steps for the Dirichlet process part of the second version of the new model.

We assume a finite maximum number of clusters K. We denote each cluster membership of respondent i as L_i, $L_i \in \{1, 2, \cdots, K\}$. Cluster k is characterized by parameter, γ_k. In contrast to the first version of the new model where each respondent has a unique γ_i, different respondents may share the same γ_k if they are in the same cluster k in the second version. All the γ_k's have the Dirichlet process prior.

$$\gamma_k \overset{iid}{\sim} G \quad k = 1, \ldots, K$$

$$G \sim \mathcal{DP}(\alpha G_0)$$

where $\alpha \in \mathbb{R}_+$ is a concentration parameter and G_0 is the centering distribution, which is specified as $\mathcal{U}nif(0, \frac{1}{2}\pi)$.

Without considering any density function, we devise an augmented cluster weight parameter, ω_k. The purpose of ω_k's is to induce sparsity in clustering, so that large clusters tend to grow larger and small clusters tend to disappear. The prior for the vector $\boldsymbol{\omega}$ is a generalized Dirichlet distribution (Ishwaran and James, 2001; Connor and Mosimann, 1969). ω_k is generated from a stick-breaking process, for $k = 1, 2, ..., K$, and is only determined by the current sizes of all the clusters.

To express the density function of the augmented parameters, $y^*_{ijj'}$'s, associated with respondent i, conditioning on respondent i being in cluster k, we need to use the notations elaborated in the last section, $\tilde{\boldsymbol{\Theta}}_i$, $\tilde{\boldsymbol{\Theta}}'_i$, and $\tilde{\boldsymbol{y}}^*_i$. Given respondent i being in cluster k with γ_k, the conditional density function of $\tilde{\boldsymbol{y}}^*_i$ is the product of M_i normal distribution densities:

$$\mathcal{L}(\tilde{\boldsymbol{y}}^*_i | \gamma_k, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}'_i) = \prod_{m=1}^{M_i} \phi_1(\tilde{\boldsymbol{y}}^*_i[m]; \quad \tilde{\boldsymbol{\Theta}}_i[m] \cdot \mathbf{g}(\gamma_i) - \tilde{\boldsymbol{\Theta}}'_i[m] \cdot \mathbf{g}(\gamma_k), 1)$$

Both the cluster weight, ω_k, and the conditional density of respondent i's augmented parameters, $\tilde{\boldsymbol{y}}^*_i$, given respondent i being in cluster k, contribute to the probability of assigning respondent i to cluster k. We denote the probability of assigning respondent i to cluster k as q_{ik}. For

$k = 1, 2, \cdots, K$, we compute q_{ik} as follows:

$$q_{ik} \propto \omega_k \mathcal{L}(\tilde{\boldsymbol{y}}_i^* | \gamma_k, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}_i')$$

$$\sum_{k=1}^{K} q_{ik} = 1$$

The cluster label for respondent i has a categorical distribution:

$$L_i \sim \mathcal{C}ategorical(q_{i1}, q_{i2}, \cdots, q_{iK})$$

We can supply a fixed value for the precision parameter, α. Or we can treat α as a parameter to estimate based on the data. In the latter case, we put a conjugate Gamma prior on α with shape a and rate b:

$$\alpha \sim \mathcal{G}amma(a, b)$$

In the rest of this subsection, we demonstrate the steps to sample the parameters above.

2.3.2.1 Sample γ_k

Given the current cluster memberships of all the respondents, we update each cluster's γ_k with either a mini random walk Metropolis-Hastings sampler or a simple draw from the prior. If cluster k is empty, then we don't have any empirical data for updating γ_k. We simply draw a new γ_k from $\mathcal{U}nif(0, \frac{\pi}{2})$. If cluster k has members, we treat these respondents and their associated $\tilde{\boldsymbol{\Theta}}_i$, $\tilde{\boldsymbol{\Theta}}_i'$, and $\tilde{\boldsymbol{y}}_i^*$ as belonging to cluster k. Then we use a random walk MH sampler to update γ_k. Theoretically, we can use a one-step MH sampler for each cluster, and the MCMC should eventually traverse to the mode of each γ_k. In order to improve the efficiency of MCMC, we do multiple steps of MH sampler and update γ_k with the last-step value. We need to specify the iteration number and tuning parameters for these mini MH samplers.

In each iteration within a mini MH sampler, we do the following steps. Given the current value of $\gamma_k^{(t)}$, $\{\tilde{\boldsymbol{\Theta}}_i\}_{i:L_i=k}$, and $\{\tilde{\boldsymbol{\Theta}}_i'\}_{i:L_i=k}$, the conditional density of $\{\tilde{\boldsymbol{y}}_i^*\}_{i:L_i=k}$ is the product of multiple

normal distribution densities:

$$\begin{aligned}
&\mathcal{L}(\{\tilde{\boldsymbol{y}}_i^*\}_{i:L_i=k}|\gamma_k^{(t)}, \{\tilde{\boldsymbol{\Theta}}_i\}_{i:L_i=k}, \{\tilde{\boldsymbol{\Theta}}_i'\}_{i:L_i=k}) \\
&= \prod_{i:L_i=k} \prod_{m=1}^{M_i} \phi_1(\tilde{\boldsymbol{y}}_i^*[m]; \quad \tilde{\boldsymbol{\Theta}}_i[m] \cdot \mathbf{g}(\gamma_k^{(t)}) - \tilde{\boldsymbol{\Theta}}_i'[m] \cdot \mathbf{g}(\gamma_k^{(t)}), 1)
\end{aligned}$$

where $[m]$ indicates the m'th element of a vector or the m'th row of a matrix, and $\phi_1(\cdot)$ is the PDF of a univariate normal distribution.

We generate a random walk step, τ, from a uniform distribution, $\tau \sim \mathcal{U}nif(-\delta, \delta)$. δ is the positive tuning parameter that determines the accepting rate of the mini random walk MH sampler. The proposed new value of γ_k is $\gamma_k^{(t+1)} = \gamma_k^{(t)} + \tau$. We plug $\gamma_k^{(t+1)}$ in the density function, and get $\mathcal{L}(\{\tilde{\boldsymbol{y}}_i^*\}_{i:L_i=k}|\gamma_k^{(t+1)}, \{\tilde{\boldsymbol{\Theta}}_i\}_{i:L_i=k}, \{\tilde{\boldsymbol{\Theta}}_i'\}_{i:L_i=k})$.

Due to the uniform prior on γ_k, the acceptance ratio, r, is determined only by the ratio of the density functions.

$$r = \min\left(1, \frac{\mathcal{L}(\{\tilde{\boldsymbol{y}}_i^*\}_{i:L_i=k}|\gamma_k^{(t+1)}, \{\tilde{\boldsymbol{\Theta}}_i\}_{i:L_i=k}, \{\tilde{\boldsymbol{\Theta}}_i'\}_{i:L_i=k})}{\mathcal{L}(\{\tilde{\boldsymbol{y}}_i^*\}_{i:L_i=k}|\gamma_k^{(t)}, \{\tilde{\boldsymbol{\Theta}}_i\}_{i:L_i=k}, \{\tilde{\boldsymbol{\Theta}}_i'\}_{i:L_i=k})}\right)$$

With probability r, we accept the proposed new $\gamma_k^{(t+1)}$, and with probability $1-r$, we reject it. We store the last-step value $\gamma_k^{(T)}$, and update the old γ_k with $\gamma_k^{(T)}$. We repeat this process for each cluster, until we finish updating all the γ_k's.

2.3.2.2 Sample ω_k

Given the current respondents' cluster memberships and each cluster's size, we use a stick-breaking process to update ω_k's. We denote the size of cluster k as ζ_k. To generate ω_k, we need to introduce auxiliary parameters, V_k, for $k = 1, 2, \cdots, K-1$. Given the current cluster sizes, ζ_k's, we first generate the auxiliary parameters, V_k as below.

$$V_k \sim \mathcal{B}eta(1 + \zeta_k, \alpha + \sum_{l=k+1}^{K} \zeta_l), \text{ for } k = 1, 2, \cdots, K-1$$

Then we update ω_k's according to the following formula:

$$\begin{aligned}
\omega_1 &= V_1 \\
\omega_k &= V_k \prod_{l=1}^{k-1}(1-V_l), \text{ for } k = 2, \cdots, K-1 \\
\omega_K &= \prod_{l=1}^{K-1}(1-V_l) = 1 - \sum_{l=1}^{K-1}\omega_l
\end{aligned}$$

2.3.2.3 Sample L_i

Given the current value of γ_k for cluster k, we compute the conditional density, $\mathcal{L}(\tilde{\boldsymbol{y}}_i^*|\gamma_k, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}_i')$, for respondent i to be in cluster k. We then take the product of ω_k and $\mathcal{L}(\tilde{\boldsymbol{y}}_i^*|\gamma_k, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}_i')$, and use it to form the categorical distribution below to draw the new cluster label, L_i, for respondent i, from the discrete cluster label set, $\{1, 2, \cdots, K\}$.

$$\begin{aligned}
&L_i \sim \mathcal{C}ategorical(q_{i1}, q_{i2}, \cdots, q_{iK}) \\
&q_{ik} \propto \omega_k \mathcal{L}(\tilde{\boldsymbol{y}}_i^*|\gamma_k, \tilde{\boldsymbol{\Theta}}_i, \tilde{\boldsymbol{\Theta}}_i') \\
&\sum_{k=1}^{K} q_{ik} = 1
\end{aligned}$$

2.3.2.4 Sample α

The vector $\boldsymbol{\omega}$ has a generalized Dirichlet distribution prior, and the concentration parameter α is a parameter in this prior. The conditional distribution of α, given the current $\boldsymbol{\omega}$, has the kernel of a Gamma distribution: $f(\alpha|\boldsymbol{\omega}) \propto \alpha^{K-1}\omega_K^{\alpha} = \alpha^{K-1}\exp\big(-(-\alpha\log\omega_K)\big)$ (Ishwaran and Zarepour, 2000). Given the conjugate Gamma prior on α, $\alpha \sim \mathcal{G}amma(a, b)$, we can express the conditional posterior for α as follows:

$$\alpha|\boldsymbol{\omega} \sim \mathcal{G}amma(a + K - 1, b - \log \omega_K)$$

2.4 Simulation Study

We conduct simulation studies to illustrate how the samplers for the two versions of the new model work. The experiments show that both versions of the new model are able to recover the true latent variable values from the observed binary choice data. For both versions of the new model, we specify four configurations of respondent number, I, and object number, J: $(I = 40, J = 40), (I = 40, J = 80), (I = 80, J = 40), (I = 80, J = 80)$. For each configuration, we repeat the simulation steps below for 50 times, so we end up with 50 simulated data sets for each configuration.

2.4.1 Simulation Results for the First Version of the New Model

For each simulation data set with respondent number I and object number J, we let each respondent compare $\texttt{round}(\frac{0.03J(J-1)}{2})$ object pairs. We generate the true $\boldsymbol{\theta}_j$'s from a bivariate standard normal distribution. We generate γ_i's from a uniform distribution, $\gamma_i \overset{i.i.d}{\sim} \mathcal{U}nif(0, \frac{\pi}{2})$. We generate respondent i's binary choice between object pair j and j', $y_{ijj'}$, in the following way:

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi_1(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i))$$

Before fitting the first version of the new model on a simulated data set, we generate the starting values of $\boldsymbol{\theta}_j$'s from a bivariate standard normal distribution. We generate the starting values of γ_i's from a uniform distribution, $\mathcal{U}nif(0, \frac{\pi}{2})$. After some trial runs, we set the random walk MH sampler tuning parameter as $\delta = 0.5$. We run 2500 iterations with the first 500 iterations as burnins. We thin the MCMC chain by 2.

The simulation results show that the sampler for the first version of the new model is able

to recover the true latent variable values based a simulated binary choice data set. Moreover, as either I or J grows, a simulation data set provides more information about the latent variables, and the sampler for the first version of the new model also achieves higher accuracy in latent variable estimation. We use two measures to gauge how well the model can uncover the true latent variables values: the correlations between the estimated parameters and the true values, and the Mean Squared Errors (MSE) of the estimated parameters. We compute the correlations and MSEs for the results from each simulation data set under each simulation configuration.

As shown by Figure 2.2 through Figure 2.5, the estimated γ parameters and θ parameters have consistently high correlations with the true values, and low MSE values across simulation data sets and simulation configurations. Under the simulation configuration with the least information ($I = 40$ and $J = 40$), the mode of the correlations between the estimated γ's and their true values is around 0.85, and the mode of the MSEs is around 0.09. As the information in simulation data sets grows, the correlation mode grows to about 0.95, and the MSE mode drops to about 0.03. Similarly, under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated θ's and their true values is around 0.9, and the mode of the MSEs is around 0.25. As the simulation data set size grows, the correlation mode grows to about 0.97, and the MSE mode drops to about 0.07.

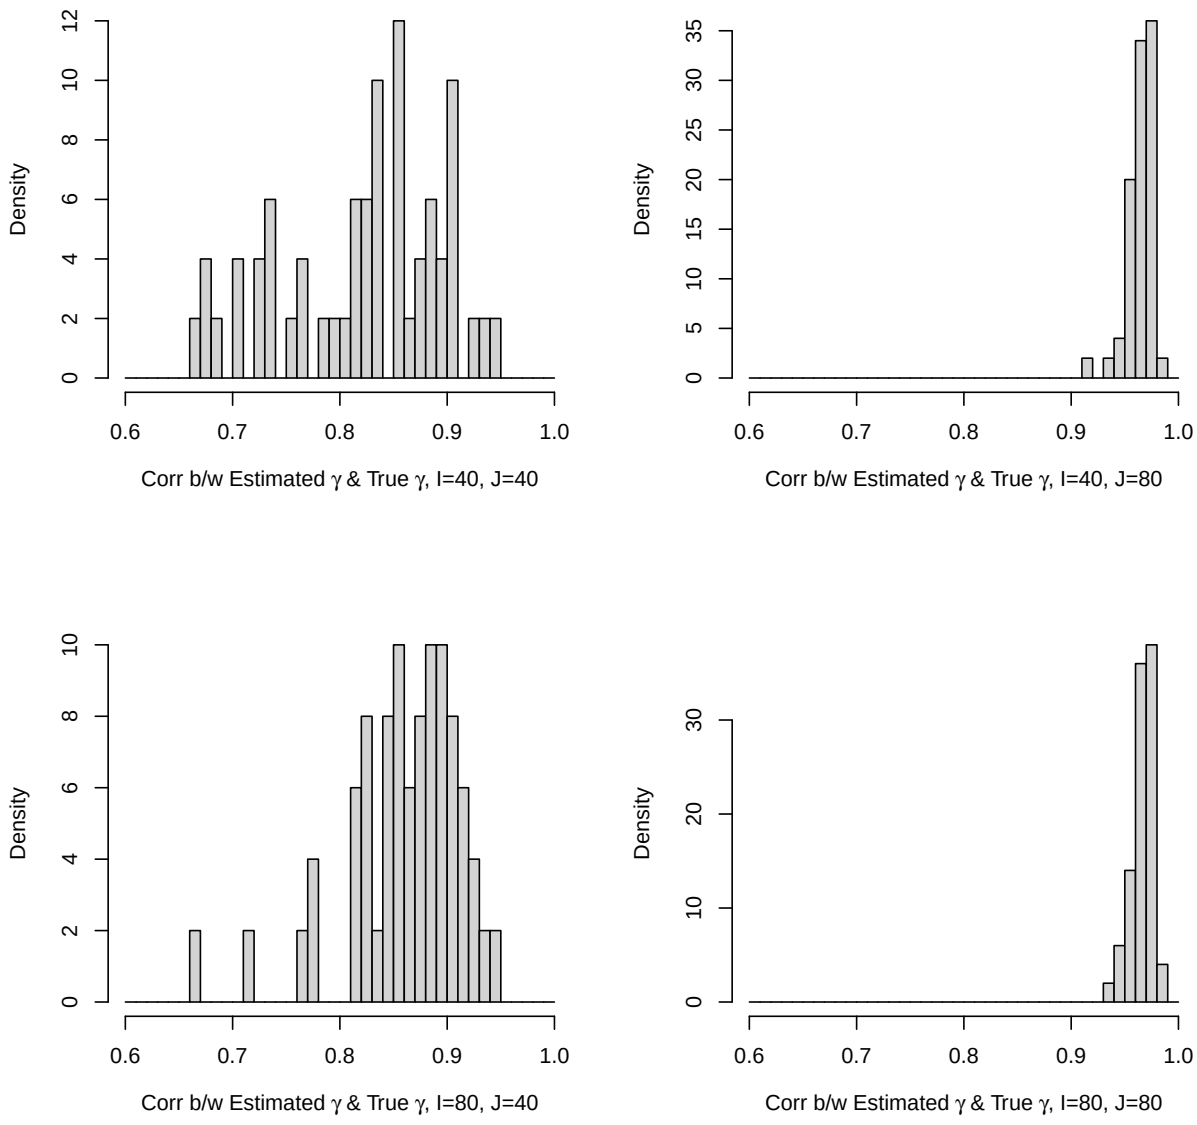

Figure 2.2: *Correlation between Estimated γ and True γ for the First Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated γ values and their true values is around 0.85. As the simulation data set size grows, the correlation mode grows to about 0.95.

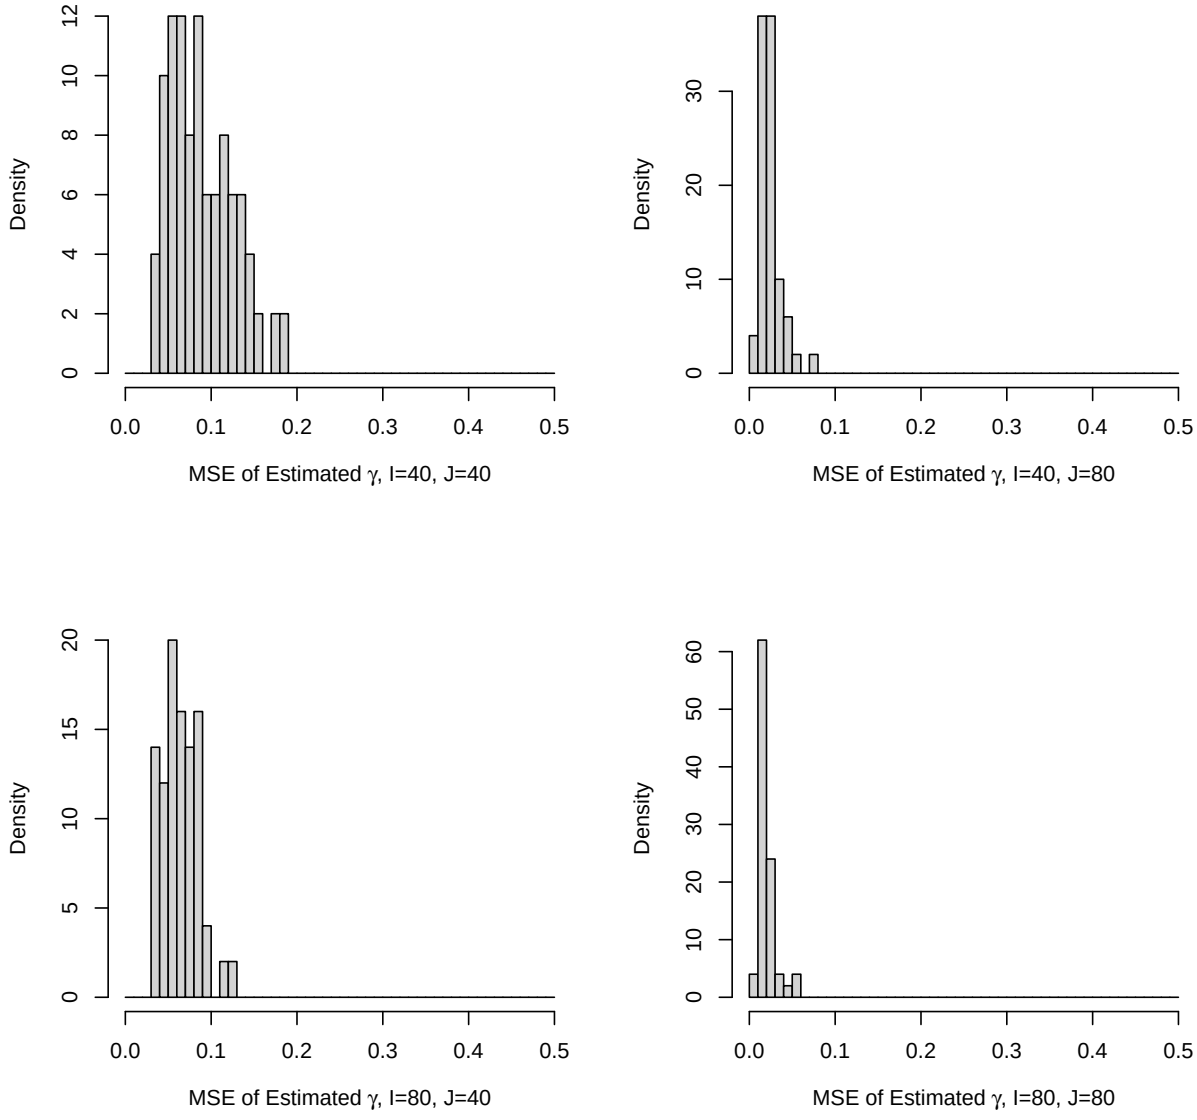

Figure 2.3: *MSE of Estimated γ for the First Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the MSEs of the estimated γ values is around 0.09. As the simulation data set size grows, and the MSE mode drops to about 0.03.

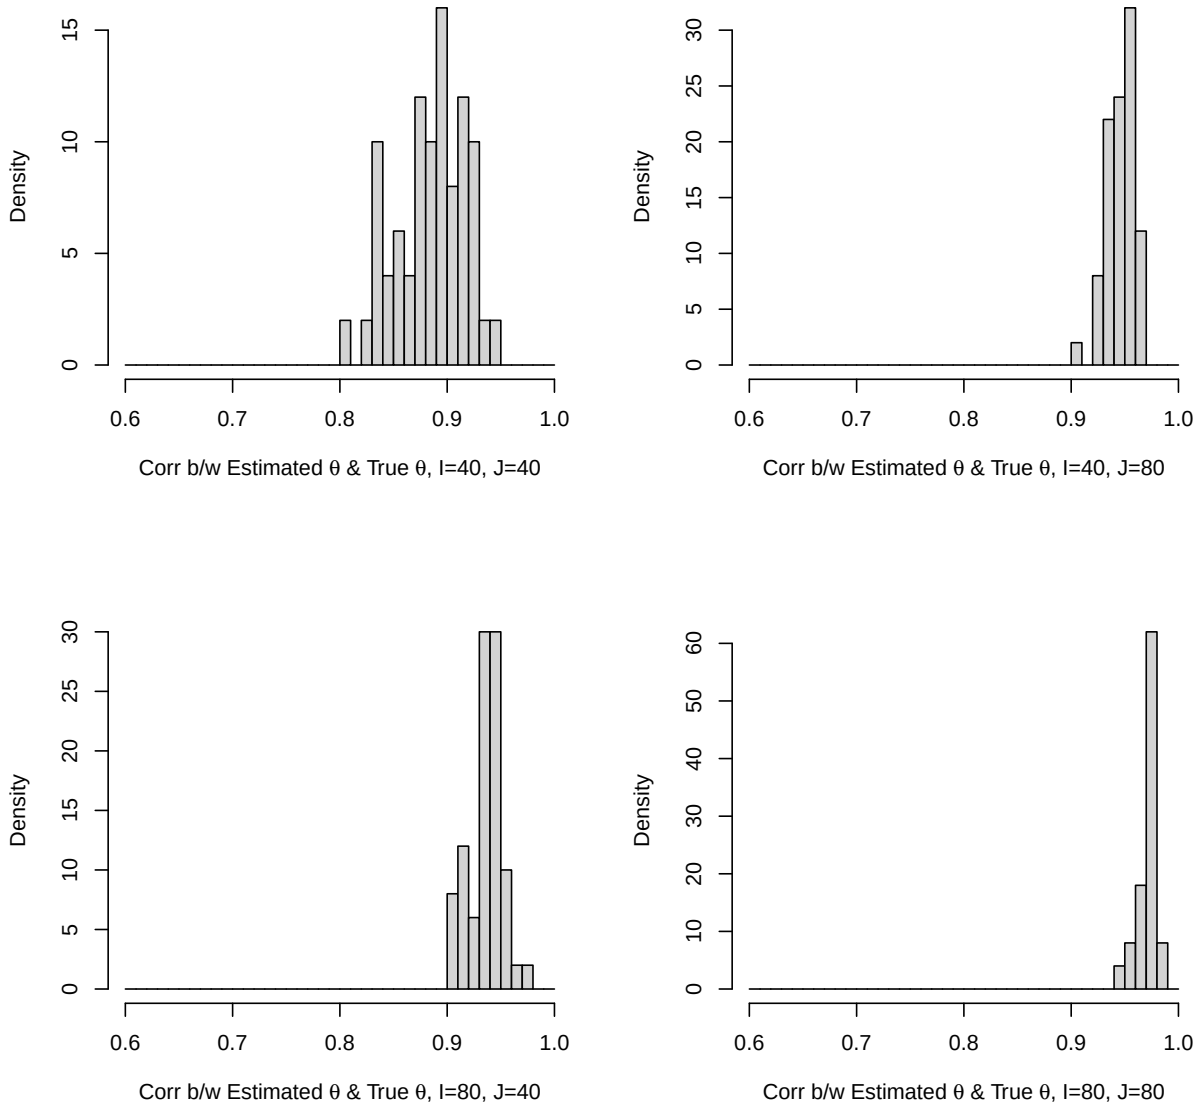

Figure 2.4: *Correlation between Estimated θ and True θ for the First Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated θ values and their true values is around 0.9. As the simulation data set size grows, the correlation mode grows to about 0.97.

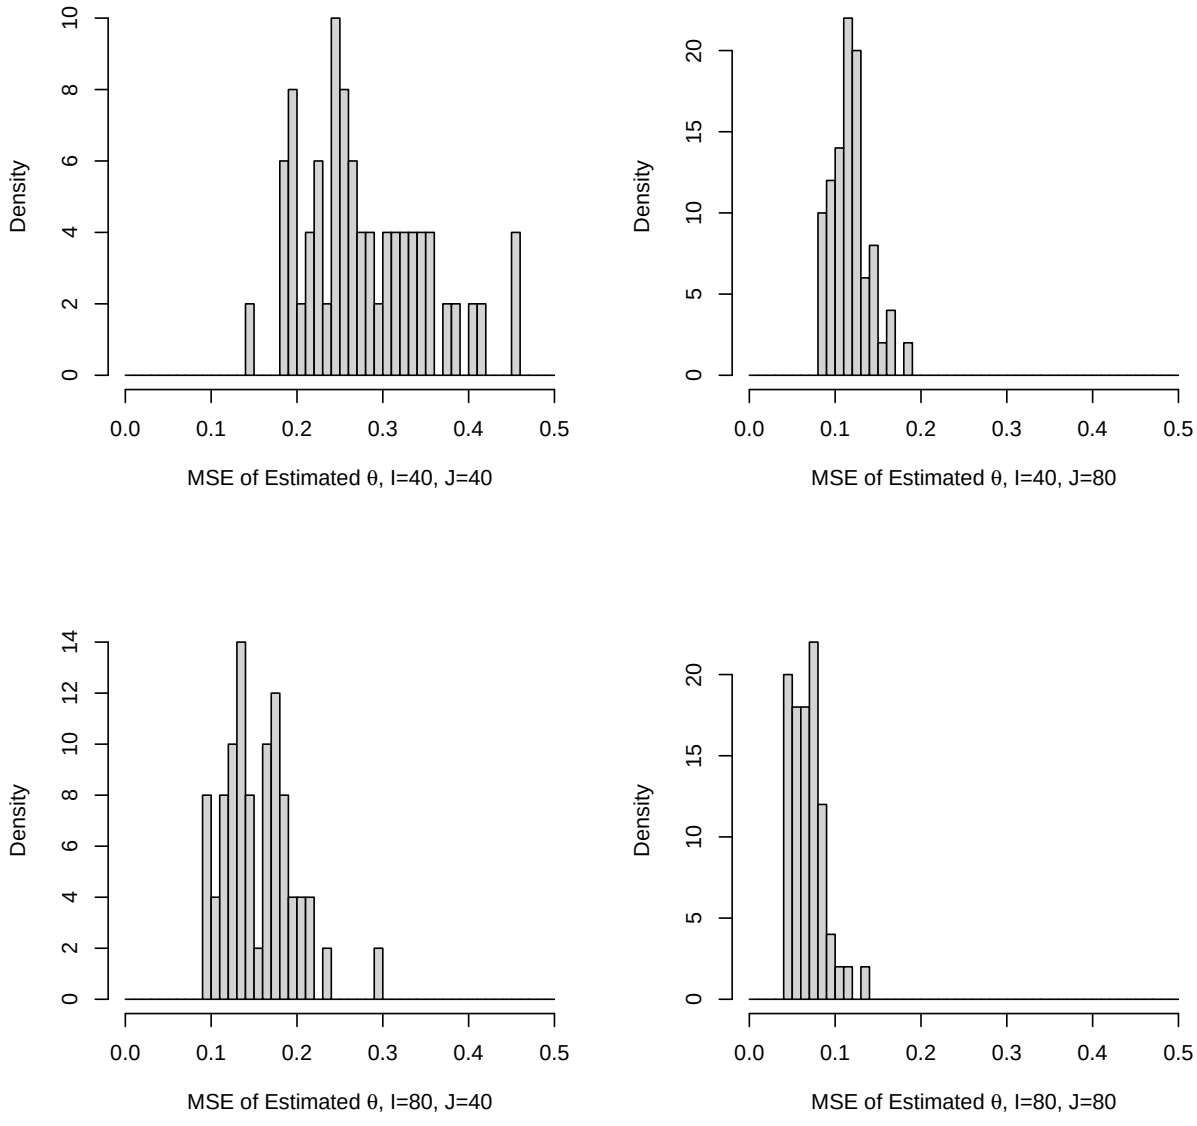

Figure 2.5: *MSE of Estimated θ for the First Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the MSEs of the estimated θ values is around 0.25. As the simulation data set size grows, the MSE mode drops to about 0.07.

2.4.2 Simulation Results for the Second Version of the New Model

For each simulation data set with respondent number I and object number J, we let each respondent compare $\texttt{round}(\frac{0.03J(J-1)}{2})$ object pairs. We generate the true $\boldsymbol{\theta}_j$'s from a bivariate standard normal distribution. We generate γ_i's from a Categorical distribution among three unique values, $\gamma_i \overset{i.i.d}{\sim} \mathcal{C}ategorical(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. We draw the three unique values from $\mathcal{U}nif(0, \frac{\pi}{2})$.[3] We generate respondent i's binary choice between object pair j and j', $y_{ijj'}$, in the following way:

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi_1(\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma_i) - \boldsymbol{\theta}_{j'} \cdot \mathbf{g}(\gamma_i))$$

Before fitting the second version of the new model on a simulated data set, we generate the starting values of $\boldsymbol{\theta}_j$'s from a bivariate standard normal distribution. We set the maximum cluster number at 20. We use the evenly distanced sequence of values from 0.03 to $\frac{\pi}{2} - 0.03$ as the starting values of γ_i's. After some trial runs, we set the random walk MH sampler tuning parameter as $\delta = 0.5$. We use 300 iterations for a mini-MCMC step for updating the γ value for an existing cluster. We don't fix the α value. Instead, we set $a = 0.01$ and $b = 1$ in the prior distribution for α, $\mathcal{G}amma(a, b)$. We run 2500 iterations with the first 500 iterations as burnins. We thin the MCMC chain by 2.

The simulation results show that the sampler for the second version of the new model is able to recover the true latent variable values. Moreover, as a simulation data set provides more information about the latent variables, the sampler for the second version of the new model also performs better at latent variable estimation. Similarly to the simulation study on the first version of the new model, we compute the correlations and MSEs for the results from each simulation data set under each simulation configuration.

As shown by Figure 2.6 through Figure 2.9, the estimated γ values and θ values have consistently high correlations with the true values, and low MSE values across simulation data sets and

[3] We make sure that any two unique values are not too close to each other. Namely, we make sure that the absolute difference between any unique value pair is greater than $\frac{\pi}{8}$.

simulation configurations. Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated γ values and their true values is around 0.88, and the mode of the MSEs is around 0.07. As the information in a simulation data set grows, the correlation mode grows to about 0.99, and the MSE mode drops to about 0.01. Similarly, under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated θ values and their true values is around 0.9, and the mode of the MSEs is around 0.25. As the simulation data set size grows, the correlation mode grows to about 0.98, and the MSE mode drops to about 0.05.

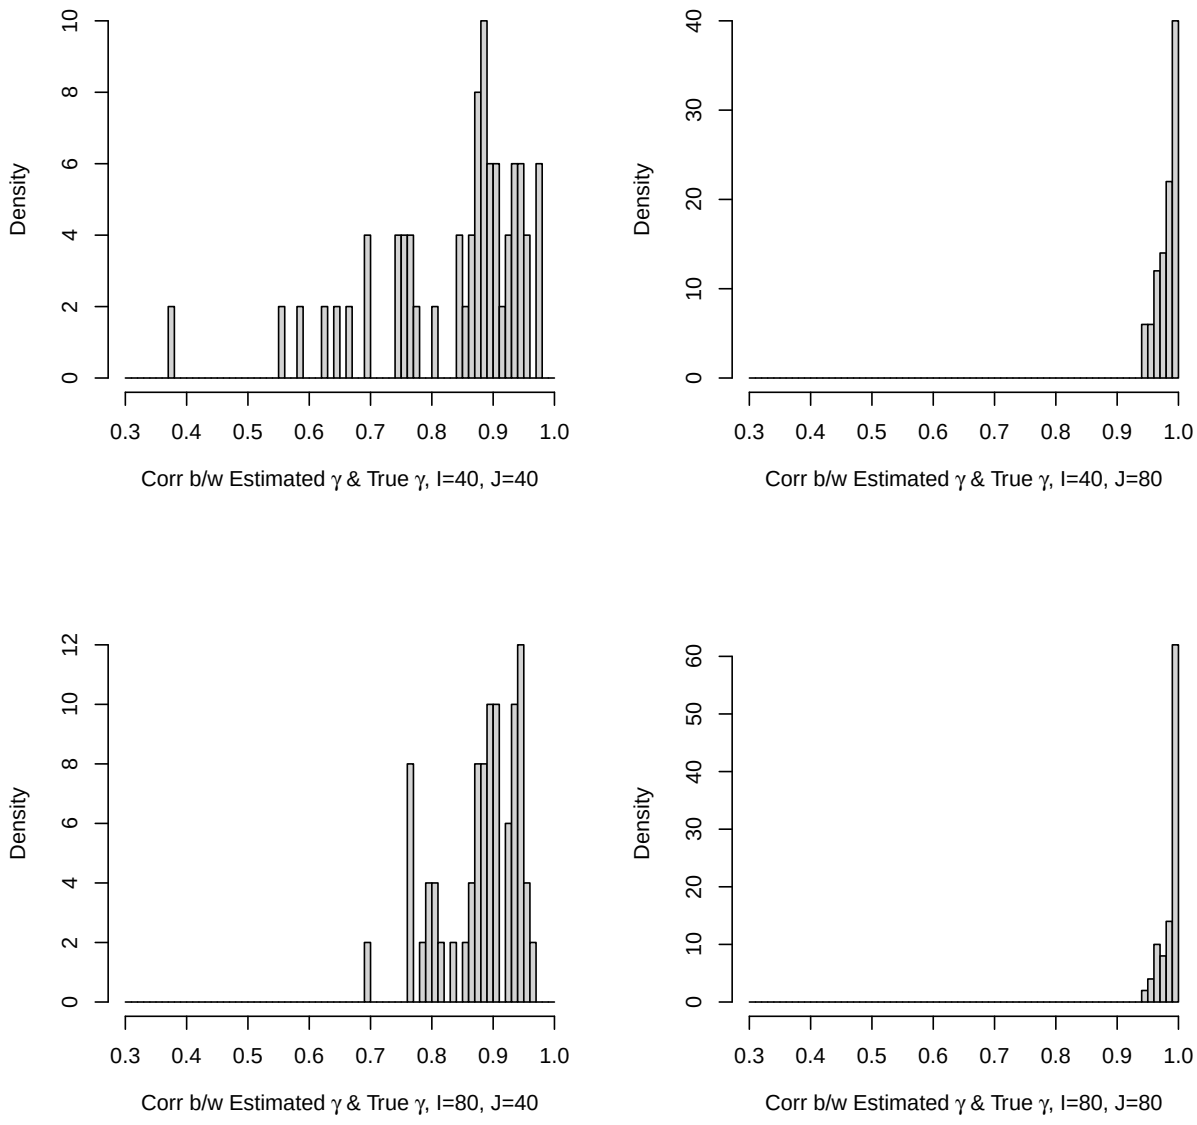

Figure 2.6: *Correlation between Estimated γ and True γ for the Second Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated γ values and their true values is around 0.88. As the simulation data set size grows, the correlation mode grows to about 0.99.

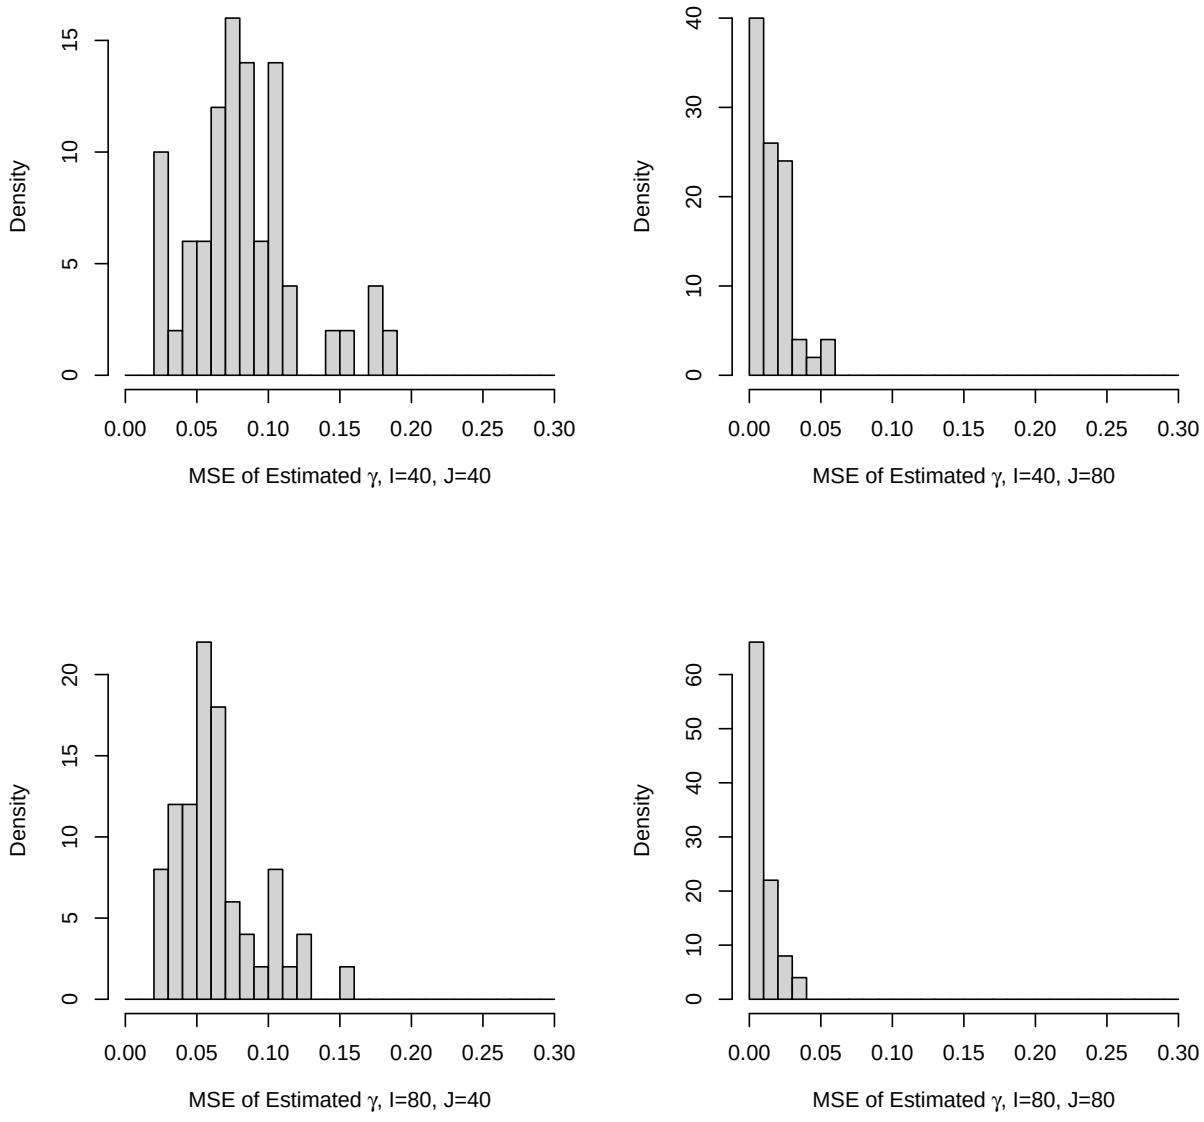

Figure 2.7: *MSE of Estimated γ for the Second Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the MSEs of the estimated γ values is around 0.07. As the simulation data set size grows, the MSE mode drops to about 0.01.

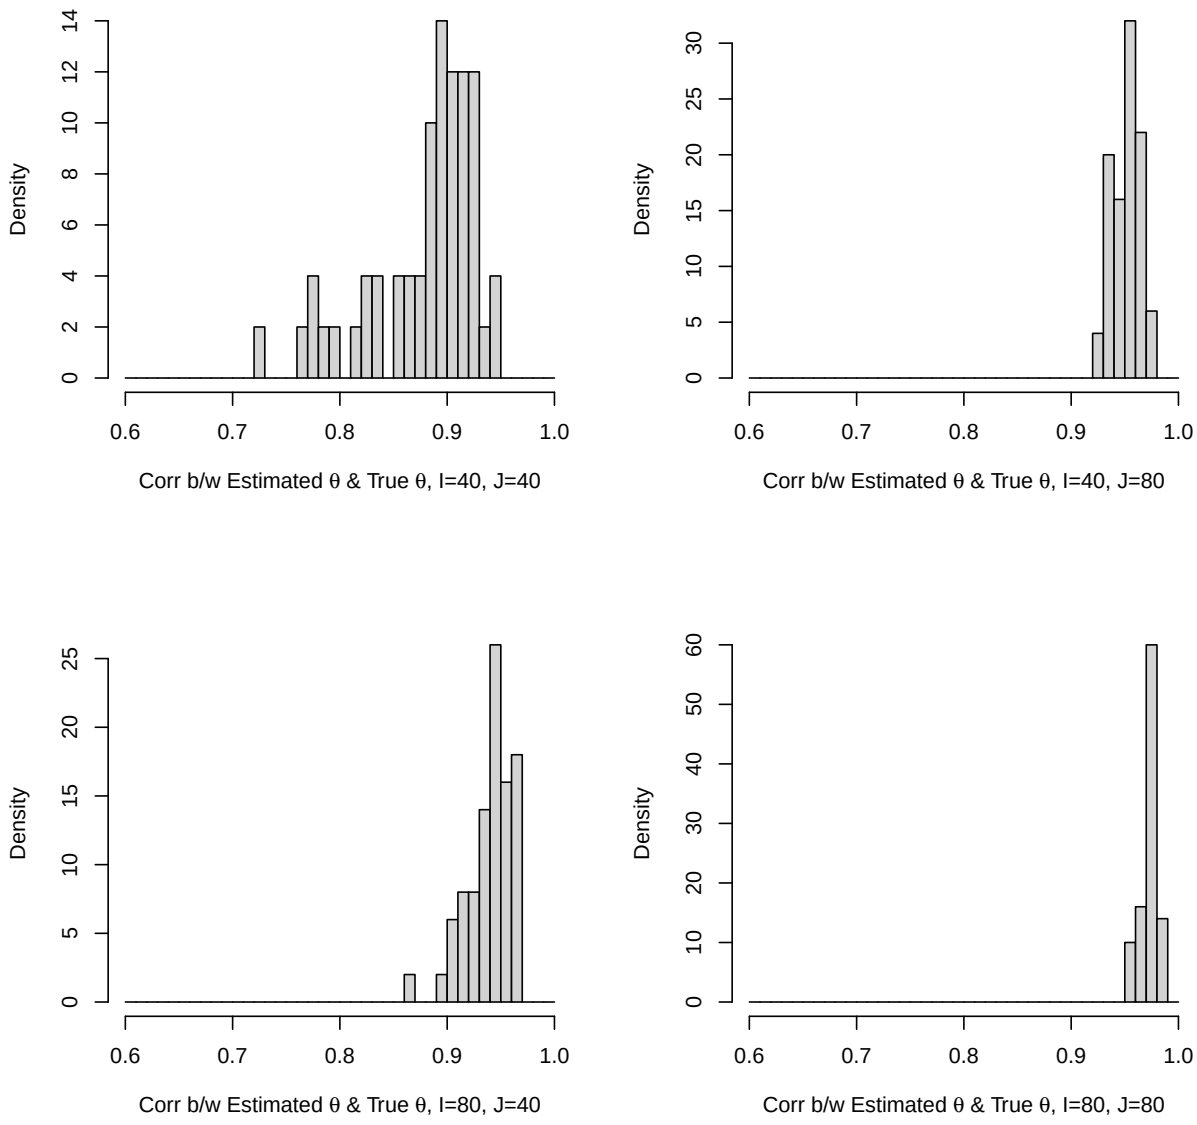

Figure 2.8: *Correlation between Estimated θ and True θ for the Second Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the correlations between the estimated θ values and their true values is around 0.9. As the simulation data set size grows, the correlation mode grows to about 0.98.

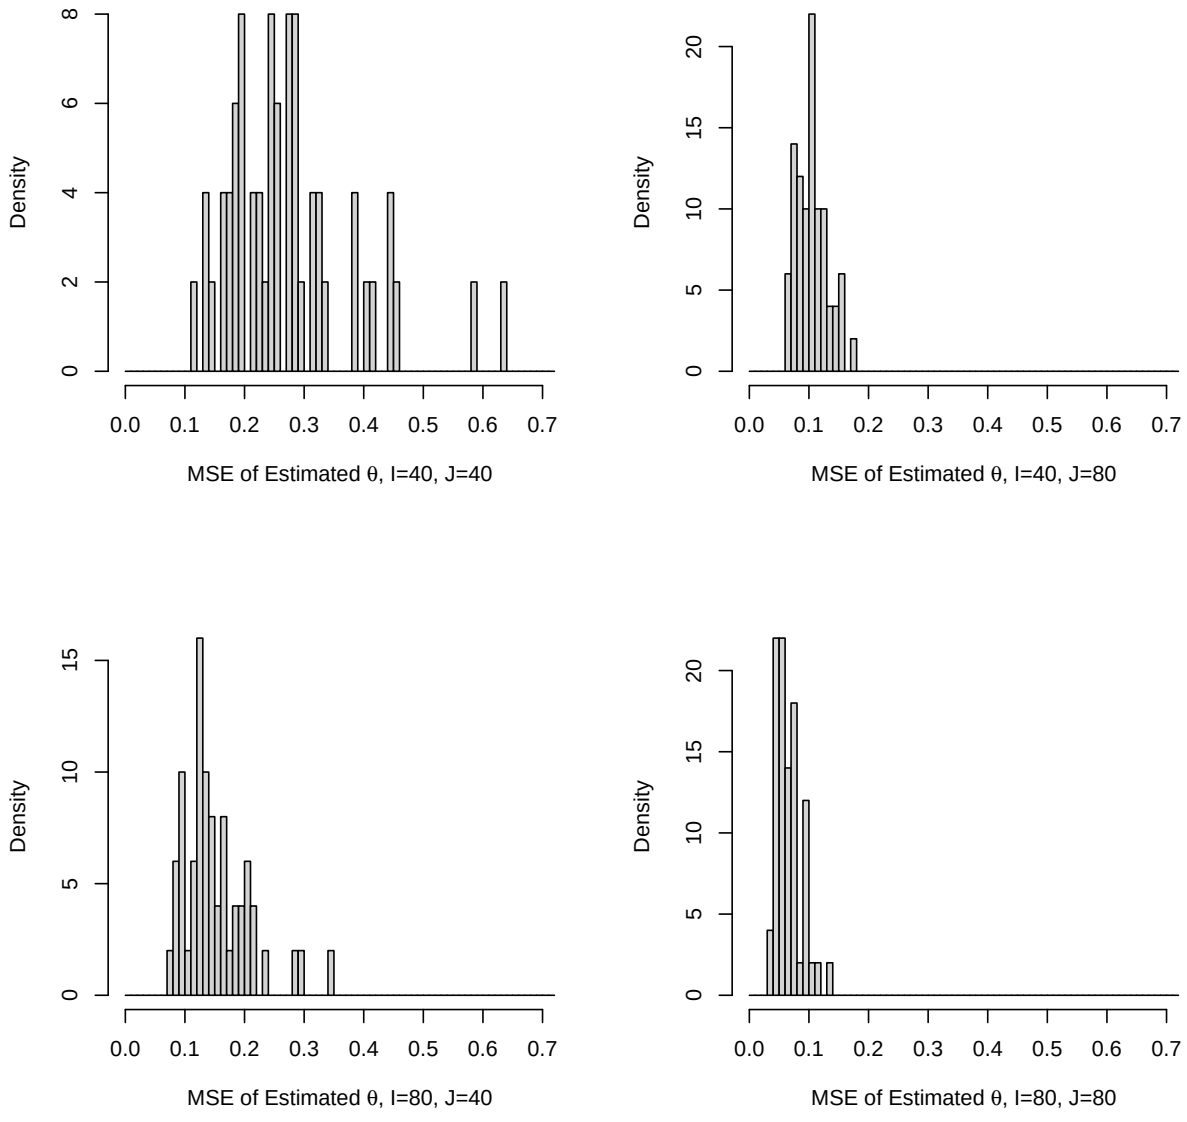

Figure 2.9: *MSE of Estimated θ for the Second Version of the New Model.* Under the simulation configuration with $I = 40$ and $J = 40$, the mode of the MSEs of the estimated θ values is around 0.25. As the simulation data set size grows, the MSE mode drops to about 0.05

2.5 Modeling the Perceived Truthfulness of Public Statements on COVID-19

One important challenge in combating the COVID-19 pandemic is to help the public identify scientific information and to rein in the spread of misinformation (Galvão, 2020). The COVID-19 "infodemic", which stands for the "overabundance of information—some accurate and some not—that occurs during an epidemic", has great implications for the global efforts in fighting the COVID-19 pandemic (Eysenbach, 2020). Rumors, stigma, and conspiracy theories about COVID-19 have caused damage across the world (Islam et al., 2020). To better understand how COVID-19 misinformation spreads, we need to measure and analyze how information about COVID-19 is perceived by individuals, and what individual characteristics impact the perception. In today's environment where COVID-19 is highly politicized (Hart, Chinn and Soroka, 2020), it is valuable to investigate how individuals perceive the truthfulness of public statements about COVID-19. To do this we collect original survey data and analyze the data with the new multidimensional pairwise comparison model proposed above.

More specifically, we focus our attention on the following question: Does the objective truthfulness of a statement on COVID-19 or the political valence of the statement better account for individual perceptions of the truthfulness of the statement? To answer this question, we implemented a nation-wide online survey that elicited perceptions of truthfulness on 42 public statements on the COVID-19 pandemic. To minimize cognitive demands on respondents and to eliminate the possibility of differential Likert-type scale usage, our survey asked respondents to compare the *relative* truthfulness of pairs of COVID-19 statements.

We hypothesize that perceptions of truthfulness are influenced by two distinct attributes of the statements: 1) the objective truthfulness of the statements, and 2) the political valence of the statements. Such multidimensional latent structure of the objects being rated cannot be represented by standard, unidimensional models for pairwise comparisons data (e.g. Thurstone (1927); Bradley and Terry (1952)). Therefore, we apply the new multidimensional pairwise comparison model

above to the survey data, which allows the objects being rated to have a rich structure of latent attributes and also allows for respondent-specific differences in perception. The remainder of this section proceeds as follow. First, we describe our data collection procedure and survey design. Next, we apply the new model to the pairwise comparison data collected in the survey. We report and compare the analysis results based on both existing unidimensional models and the newly-proposed multidimensional model.

2.5.1 Data and Survey Design

In this section we briefly discuss the source of the COVID-19 statements used in our survey, and the design and implementation of the survey.

2.5.1.1 COVID-19 Statements

Since we are interested in the extent to which members of the mass public accurately assess the truthfulness of statements about COVID-19, it is important that we use fact-checked statements so as to have an independent measure of the truthfulness of each statement. Our source of these fact-checked COVID-19 statements is the website `https://www.politifact.com`.[4]

PolitiFact catalogs a range of statements that have political content. According to PolitiFact's own website:

> Each day, PolitiFact journalists look for statements to fact-check. We read transcripts, speeches, news stories, press releases, and campaign brochures. We watch TV and scan social media. Readers send us suggestions via email to `truthometer@politifact.com`; we often fact-check statements submitted by readers. Because we can't feasibly check all claims, we select the most newsworthy and significant ones.[5]

PolitiFact journalists fact check these statements and categorize the truthfulness of each statement into one of six categories (from most truthful to least truthful): true, mostly true, half true,

[4] PolitiFact's Editor-in-Chief, Angie Drobnic Holan, gave us permission to use the PolitiFact data for this survey in an email on May, 11, 2020.

[5] https://www.politifact.com/article/2018/feb/12/principles-truth-o-meter-politifacts-methodology-i/

mostly false, false, pants on fire.[6]

We selected 42 statements with the intent of balancing the truthfulness of the statements and the slant of the statements (left, neutral, right). These statements were made between February 22, 2020 and May 8, 2020. Ideally, we would have used equal numbers statements of left, neutral, and right-leaning statements from all six truthfulness categories. However, some categories were sparsely populated and we were forced to dichotomize the truthfulness categories into high truth (true, mostly true, and half true) and low truth (pants on fire, false, mostly false). This gave us 7 statements in each of the 3×2 combinations of slant $\times$ truthfulness. The full set of 42 statements along with their truthfulness ratings and slant is presented in Appendix A.1.

2.5.1.2 The Survey

The survey was conducted on July 8, 2020.[7] Respondents were recruited from the Lucid Marketplace.[8] Quotas were used to make the sample approximate the U.S. voting age population. The survey was conducted online using the Qualtrics interface.

After a respondent provided their informed consent to continue with the survey, a short training page was provided to the respondent. This training page made it clear that the questions about the relative truthfulness of pairs of COVID-19 statements were eliciting the respondent's belief about which statement was more truthful when it was stated. The key language here was:

> Factual statements can be placed on a line. At one extreme end of the line are statements that are completely truthful and accurate. At the other extreme end are statements that are intentionally false. Between these two extremes we find statements that contain elements of truth and falsity and / or half-truths. For each task, you will see two statements on the coronavirus pandemic. Your task is to read both and to select the statement that you believe was **more truthful when it was stated.**

[6] See https://www.politifact.com/article/2018/feb/12/principles-truth-o-meter-politifacts-methodology-i/ for full description of PolitiFact's fact-checking process.

[7] This survey was judged exempt from review by our university's IRB (study ID HUM00184241).

[8] https://luc.id/marketplace/

After this brief training, each respondent was given a single attention check question that provided the respondent with two COVID-19 statements and asked them to select both statements.[9] Recent work on attention checks in online surveys suggests that eliminating respondents who fail attention checks may introduce demographic bias.[10] Consequently, we do not exclude respondents who fail this check. The purpose of including this attention check is solely to encourage respondents to read the following response prompts carefully.

After the attention check, respondents were asked to report their view of the relative truthfulness of COVID-19 statements given to them in randomly selected pairs of statements. Figure 2.10 depicts what this looks like for one randomly selected pair of statments.

After the paired comparisons of COVID-19 statements were given to respondents, the respondents were asked a sequence of demographic, attitudinal, and behavioral questions.

We removed a small number of respondents that Lucid flagged as having a high likelihood of being fraudulent. We received usable responses from 2,621 respondents. On average, each respondent gave us their view of the relative truthfulness of just less than 15 pairs of randomly selected statements. Appendix A.2 provides descriptive statistics on our respondent sample.

2.5.2 COVID-19 Statement Perception Data Analysis Results

Before presenting results from our new model, we present results from simple unidimensional models. As we show, the unidimensional models obscure the structure that underlies perceptions of truthfulness of COVID-19 statements. As we see in Section 2.5.2.2, our two-dimensional model more accurately represents the respondent-level heterogeneity: The responses from some respondents are more highly correlated with the objective truth of the statements, while the responses from other respondents are more strongly associated with the political valence of the statements.

[9]The text of the question was: "The following are two statements about the coronavirus pandemic. These statements were made between late February and early May, 2020. We are interested in which statement you believe was more truthful when it was made. However, for this question, we care more about whether you are paying attention. Please choose both the first and second statements to indicate you are paying attention."

[10]For example, see https://www.qualtrics.com/blog/using-attention-checks-in-your-surveys-may-harm-data-quality/.

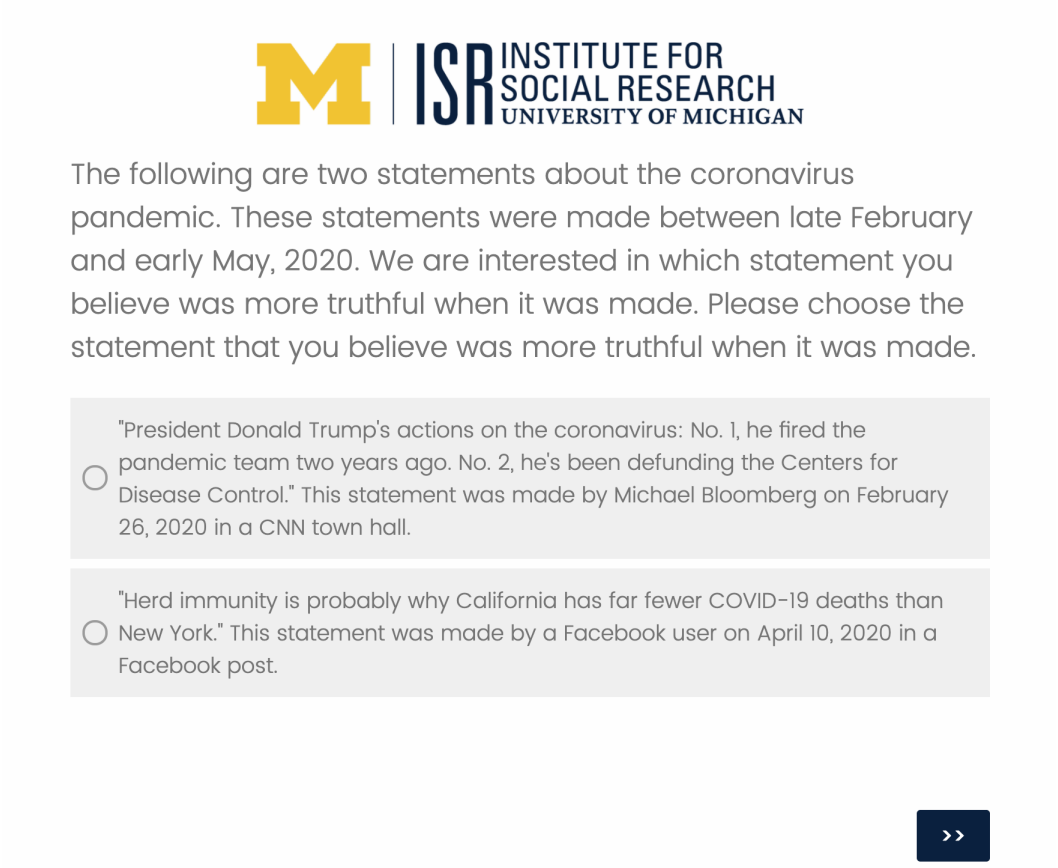

Figure 2.10: *Screen Shot of COVID-19 Statement Comparison Survey.*

The latter association varies with the strength of a respondent's partisanship and political ideology.

2.5.2.1 Results from Unidimensional Models

As a starting point, we fit the simple Thurstone model,

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi(\theta_j - \theta_{j'})$$

, to the pairwise comparison data from our survey.[11] Here $\Phi(\cdot)$ is the Gaussian cumulative distribution function.

Inspection of the output reveals that this simple model provides a poor fit to the observed data. For instance, for each observed $y_{ijj'}$ we calculate the in-sample posterior expectation of a correct classification:

$$\frac{1}{M}\sum_{m=1}^{M}\left(\mathbb{I}(y_{ijj'}=1)\Phi(\theta_j^{(m)}-\theta_{j'}^{(m)})+\mathbb{I}(y_{ijj'}=0)\Phi(\theta_{j'}^{(m)}-\theta_j^{(m)})\right)$$

where $m = 1, \ldots, M$ indexes the MCMC draws. Note that a "correct" classification is simply defined to be a classification equal to the observed response—it is not necessarily related to whether respondent i accurately perceived the true truthfulness of statement j relative to statement j'.

The average of these posterior expectations of a correct response, taken over all observed $y_{ijj'}$s, is 0.52. We can also aggregate to the statement by averaging over respondents. Doing this, we see that the average probability of a correct classification across all statements is also 0.52, and that no statement has a probability of being correctly classified greater than 0.56. If we aggregate to the respondent by averaging over the statement pairs seen by each respondent, we see that the average probability of a correctly classified response by a respondent is also 0.52. Further, we find that 26% of respondents have probabilities of a correctly classified statement less than 0.5 and only 0.3% of respondents (8 out of 2,621) have probabilities of a correctly classified statement greater than 0.6.

We also examine how the posterior means of the θ parameters correlate with the objective truthfulness and partisan valence of the statements. To do this we give a "pants-on-fire" statement a value of 0, a "false" statement a value of 1, a "mostly-false" statement a value of 2, a "half-true" statement a value of 3, a "mostly-true" statement a value of 4, and a "true" statement a value of 5. We then calculated the Spearman rank correlation between these truthfulness ratings and the posterior means of the θ parameters. This produced a rank correlation of 0.42.

[11] To identify the model, we constrained θ_{1014} to be negative and constrained $\theta_{1015} = 0.25$. The remaining θ parameters were assumed to have independent standard normal prior distributions. The MCMC sampler was run for 120,000 iterations with the first 20,000 discarded as burn-in iterations. Every 10th iteration was stored.

Similarly, we gave right-valence statements a value of 1, neutral-valence statements a value of 0, and left-valence statements a value of -1. Then we calculated the Spearman rank correlation between the partisan-valence of the statements and the posterior means of the θ parameters. This resulted in a rank correlaton of -0.17.

The simple unidimensional Thurstone model produces estimates of the statement-specific parameters that are only weakly correlated with objective truth and even more weakly correlated with the other factor that we expect to structure responses—the political valence of the statements.

As noted above, a natural extension of the basic Thurstone model is to introduce a respondent-specific parameter β_i that allows for differential ability to perceive differences between statements. This produces the model:

$$y_{ijj'} \sim \mathcal{B}ernoulli(p_{ijj'})$$

$$p_{ijj'} = \Phi(\beta_i[\theta_j - \theta_{j'}]).$$

We fit this model to the pairwise comparisons data from our survey.[12]

If we calculate the in-sample posterior expectation of a correct classification for this model in the analogous way that we did for the simple Thurstone model, we find that the average of these posterior expectations of a correct response, taken over all observed $y_{ijj'}$s, is 0.55. While the inclusion of the respondent-specific β parameters ensures that the respondent-level predictions match the observed data at least 50% of the time, it is still the case that 35% of the observed $y_{ijj'}$s have posterior probabilities of a correct classification less than 0.50. At the statement level, we see that, on average, statements are classified correctly 55% of the time with only 2 of 42 statements having a probability of correct classification greater than 0.6.

The Spearman rank correlation between the posterior means of the statement-specific θ parameters and the ojective truthfulness of the statements is 0.32, which is lower than in the simple Thurstone model. However, the rank order correlation between the posterior means of θ and the

[12]To identify the model, we constrained θ_{1014} to be negative and constrained $\theta_{1015} = 0.25$. The remaining θ parameters were assumed to have independent standard normal prior distributions. The β parameters were also assumed to have independent standard normal priors. The sign of the β parameters was not restricted. The MCMC sampler was run for 120,000 iterations with the first 20,000 discarded as burn-in iterations. Every 10th iteration was stored.

partisan valence of the statements is -0.73. Even though we constrained the model so that a neutral-valence, high-truth statement was to the right of 0 and a neutral valence, low-truth statement was to the left of 0, the resulting estimates of θ are more strongly correlated with the partisan valence of the statements than the objective truthfulness of the statements.

Indeed, this warping of truthfulness at the responent level can be seen in the posterior means of the respondent-specific β parameters. 38% of respondents have a β parameter with a posterior mean less than 0. In other words, 38% of respondents are, on average, viewing objective truth as subjective falsity, and vice versa.

2.5.2.2 Results from the Two-Dimensional Dirichlet Process Model

The results from the simple unidimensional models are not fully satisfying. The Thurstone model does a poor job of representing observed patterns in the data, and produces estimates of statement-specific parameters that only weakly correlate with objective truthfulness. The inclusion of a respondent-specific parameter slighly improves model fit at the expense of weakening the already weak correlation between the statement-specific parameter estimates and objective truth.

We fit the two-dimensional Dirichlet process model discussed in Section 2.2.2.2 to the data in the hope that this provides a better fit than the unidimensional models. [13]

Calculating the in-sample posterior expectation of a correct classification for this model in the analogous way that we did for the unidimensional models above, we find that the average of these posterior expectations of a correct response, taken over all observed $y_{ijj'}$s, is 0.57. The respondent-level predictions and statement-level predictions also match the observed data 57% of the time. These numbers are slightly better than the values of 0.52 and 0.55 we achieved with the unidimensional models.

However, this slight improvement in in-sample predictive accuracy is not the main advantage

[13] To identify the model, we constrained $\boldsymbol{\theta}_{1015}$ to be equal to 0.25 on the first dimension and greater than 0 on the second dimension; we constrained $\boldsymbol{\theta}_{1004}$ to be less than 0 on the first dimension; and we constrained $\boldsymbol{\theta}_{1042}$ to be greater than 0 on the first dimension. The remaining $\boldsymbol{\theta}$ parameters were assumed to have independent bivariate normal prior distributions with mean $\mathbf{0}$ and variance-covariance matrices equal to identity matrices. G_0 was set to $\mathcal{U}uni(0, \frac{1}{2}\pi)$ and the concentration parameter α was assumed to follow a $\mathcal{G}amma(1, 1)$ distribution. The MCMC sampler was run for 440,000 iterations with the first 40,000 discarded as burn-in iterations. Every 40th iteration was stored.

of the two-dimensional Dirichlet process model. The main advantage is that it allows us to uncover a more nuanced understanding of how certain types of respondents assess the truthfulness of statements. More specifically, in this application it allows us to see how some respondents make assessments of truthfulness based on their partisanship or political ideology, while other respondents seem to be more guided by the objective truthfulness of the statements.

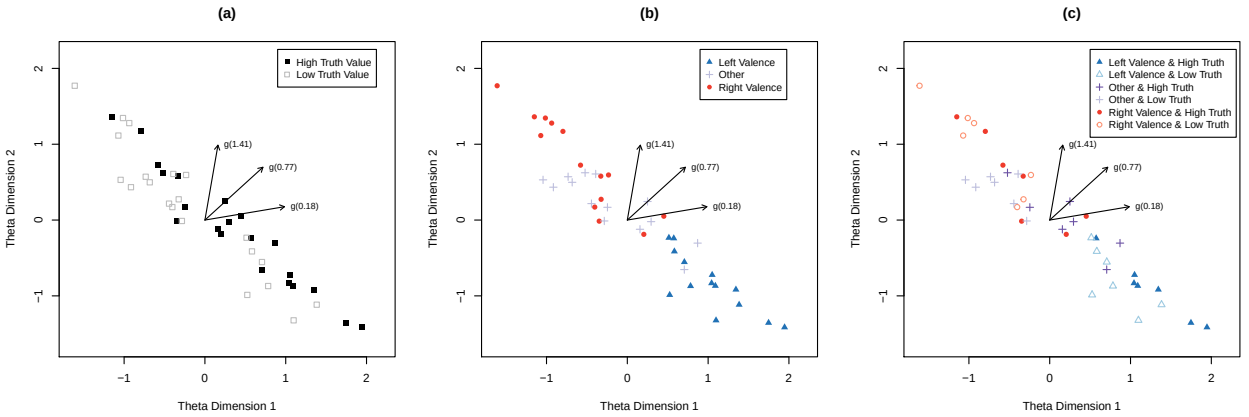

Figure 2.11: *Posterior Means of $\boldsymbol{\theta}$ and the Minimum, Maximum, and Median Posterior Means of* $\mathbf{g}(\gamma)$. In each panel, the points correspond to the posterior means of the $\boldsymbol{\theta}$ parameters for the 42 statements. The arrows correspond to the $\mathbf{g}(\gamma)$ vectors at the minimum posterior mean of γ_i, the maximum posterior mean of γ_i, and the median posterior mean of γ_i for $i = 1, \ldots, N$. In panel (a), the $\boldsymbol{\theta}$ points are shaded based on the objective truthfulness of the statements. In panel (b), the $\boldsymbol{\theta}$ points are color-coded based on the left-right valence of the statements. Finally, in panel (c), the $\boldsymbol{\theta}$ points are coded according to both the objective truthfulness of the statements and the left-right valence of the statements. Note that projecting the $\boldsymbol{\theta}$ points onto $\mathbf{g}(0.77)$ (0.77 is the median posterior mean of $\gamma_i \quad i, \ldots, N$) produces values associated with the objective truthfulness of the statements, albeit weakly. On the other hand, projecting the $\boldsymbol{\theta}$ points onto $\mathbf{g}(0.18)$ and $\mathbf{g}(1.41)$ results in points where higher values correspond to more left-leaning and right-leaning valence respectively.

As a starting point, consider Figure 2.11. This figure plots the posterior means of $\boldsymbol{\theta}_j$ for the $j = 1, \ldots, 42$ statements along with $\mathbf{g}(0.18)$, $\mathbf{g}(0.77)$, and $\mathbf{g}(1.41)$, where 0.18, 0.77, and 1.41 are the minimum, median, and maximum values of the posterior means of γ_i, for $i = 1, \ldots, N$ respectively. Figure 2.11 allows us to see how three types of respondents (those with $\gamma_i = 0.18, 0.77$ and 1.41) perceive the truthfulness of the statements.

Respondents with γ_i parameters near the median of 0.77 project the statement-specific $\boldsymbol{\theta}$ pa-

rameters onto a dimension that correlates with the objective truthfulness of the statements, albeit weakly. On the other hand, respondents with γ_i parameters near the minimum of 0.18 project the statement-specific $\boldsymbol{\theta}$ parameters onto a dimension which is positively correlated with the leftward valence of the statements. Finally, respondents with γ_i parameters near the maximum of 1.41 project the statement-specific $\boldsymbol{\theta}$ parameters onto a dimension which is positively correlated with the rightward valence of the statements.

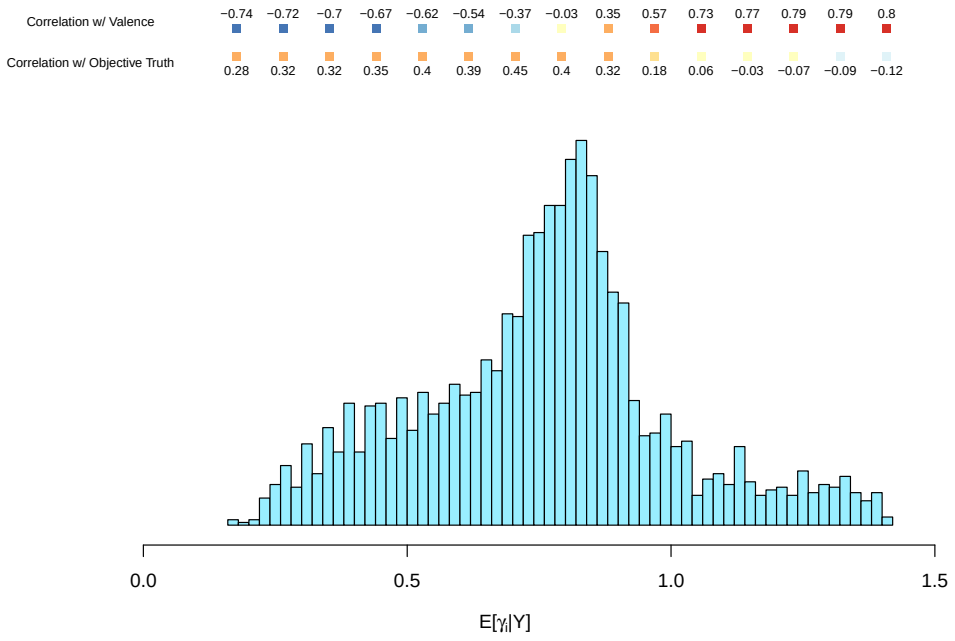

Figure 2.12: *Histogram of the Posterior Means of γ_i for $i = 1 \ldots, N$ Along with the Spearman Rank Correlations Between $\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma)$ and Objective Truthfulness and Left-Right Valence for Various Values of γ.* Note that respondents whose posterior mean γ parameter is near 0.77 tend to assess statements primarily based on the objective truthfulness of the statements but that this association is weak (correlation slightly greater than 0.4). Respondents with γ parameters that are closer to the extremes of 0.18 and 1.41 assess the truthfulness of the COVID-19 statements in ways that are strongly associated with the left-right valence of the statements. Further, respondents with γ parameters greater than approximately 1.1 not only assess the truthfulness of the COVID-19 statements such that right-valence statements are perceived as more truthful, they also assess the truthfulness of COVID-19 statements in ways that are negatively correlated with the objective truthfuness of the statements.

While the information in Figure 2.11 is useful, it doesn't provide information on three impor-

tant things: a) the distribution of respondent-specific γ_i parameters, b) the precise strength of the correlation between particular $\mathbf{g}(\gamma_i)$ projections and the objective truthfulness of statements, and c) the precise strength of the correlation between particular $\mathbf{g}(\gamma_i)$ projections and the left-right valence of the statements. This information is displayed in Figure 2.12.

Looking at Figure 2.12, three things are apparent. First, most respondents have estimated γ parameters near the median posterior mean value of 0.77 (the modal estimate of γ_i is just to the right of 0.77). There are a smaller number of respondents who have lower estimated γ parameters—about 16% of the respondents have estimated γ parameters less than 0.5—and there is a still smaller number of respondents with much larger estimated γ parameters—about 13% have estimated γ parameters greater than 1.0.

Second, Figure 2.12 presents the correlation between $\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma)$ and the objective truthfulness of the statements for various values of γ. This is the multidimensional analog to the correlation between $\boldsymbol{\theta}$ and objective truthfulness from the unidimensional models discussed in Section 2.5.2.1. Once again, we give a "pants-on-fire" statement a value of 0, a "false" statement a value of 1, a "mostly-false" statement a value of 2, a "half-true" statement a value of 3, a "mostly-true" statement a value of 4, and a "true" statement a value of 5. We then calculated the Spearman rank correlation between these truthfulness ratings and $\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma)$ for the posterior means of $\boldsymbol{\theta}$ for the 42 statements and 15 equally-spaced values of γ from 0.18 to 1.41. This produces the 15 color-coded correlations at the top of Figure 2.12.

What we see here is that the γ value that induces the highest correlation with the objective truth of the statements is $\gamma = 0.70$ which gives rise to a correlation of 0.45. γ values less than or equal to 0.88 give rise to correlations with objective truth that are greater than or equal to 0.28. However, respondents with γ values greater than or equal to 1.14 tend to rate COVID-19 statements in ways that are negatively correlated with the objective truth of the statements.

Third, Figure 2.12 presents the correlation between $\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma)$ and the left-right valence of the statements. As above, we gave right-valence statements a value of 1, neutral-valence statements a value of 0, and left-valence statements a value of -1. We calculated the Spearman rank correla-

tion between these left-right valence ratings and $\boldsymbol{\theta}_j \cdot \mathbf{g}(\gamma)$ for the posterior means of $\boldsymbol{\theta}$ for the 42 statements and 15 equally-spaced values of γ from 0.18 to 1.41. This produces the 15 color-coded correlations at the very top of Figure 2.12.

The resulting pattern of correlations with the left-right valence is stark. Respondents with the highest values of γ, say above or equal to 1.14, perceive the truthfulness of the COVID-19 statements in a way that positively correlates with the rightward valence of the statements with correlations of 0.77 or above. These same individuals' evaluations of the truthfulness of the COVID-19 statements are negatively correlated with objective truth. On the other hand, individuals with the lowest values of γ, say below or equal to 0.35, perceive the truthfulness of the COVID-19 statements in a way that negatively correlates with the rightward valence of the statements with correlations of -0.70 or below. These individuals' evaluations of the truthfulness of the COVID-19 statements are weakly positively correlated with the objective truth of the statements.

To summarize, respondents with the modal value of γ are primarily responding to the objective truthfulness of the COVID-19 statements when evaluating the truthfulness of pairs of statements. These respondents are not rating statements in ways that are correlated with the left-right valence of the statements. That said, there is only a weak correlation between the objective truth of the statements and their subjective perceptions. On the other hand, respondents with γ values at the two extremes are rating the truthfulness of statements in ways that are strongly associated with left-right valence of the statements—respondents with low values of γ tend to see left-leaning statements as more truthful while respondents with high values of γ tend to see right-leaning statements as more truthful. The objective truth of the statements is less relevant for these respondents than is the left-right valence of the statements. Indeed, those respondents who tend to see right-leaning statements as more truthful tend to perceive the truthfulness of the statements in ways that are slightly negatively correlated with the objective truth of the statements.

We also examine how respondent perceptions of COVID-19 statement truthfulness, as measured by their estimated γ parameters, correlate with respondent characteristics and behaviors. Figure 2.13 plots the relationship between the respondent-specific γ estimates and three measures

related to the political attitudes of respondents: partisanship (as operationalized by an indicator of whether a respondent self-identifies as a strong Republican), ideology (as operationalized by respondent self-placement on a 7-point Likert-type scale running from 1 = "very liberal" to 7 = "very conservative"), and the slant of news media consumption (as operationalized by respondent self-statement of their preferred news outlet combined with the media bias ratings from https://www.allsides.com/media-bias/media-bias-ratings). Each panel of Figure 2.13 plots a local regression estimate of the conditional expectation function of the variable in question on γ_i for respondents $i = 1, \ldots, N$.[14]

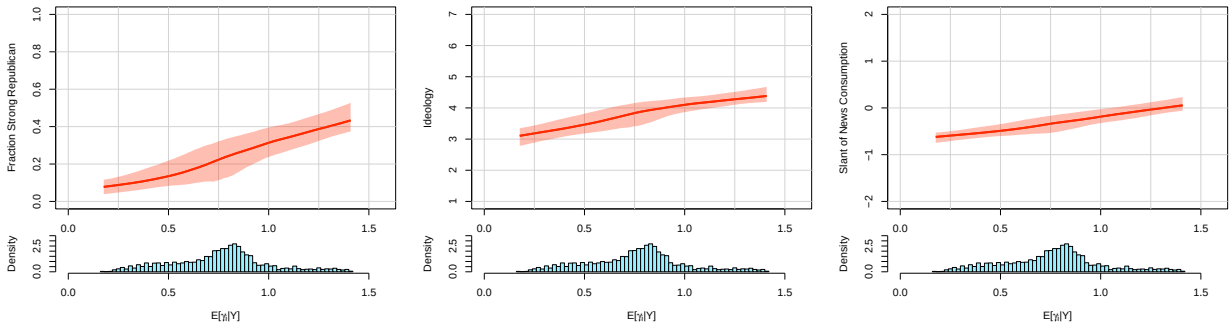

Figure 2.13: *Associations Between Posterior Means of γ_i for $i = 1 \ldots, N$ and Respondent Partisanship, Ideology, and Slant of News Media Consumption.* The dark orange lines are the posterior means of local regression predictions averaged over the posterior distribution of γ. The light orange band is the pointwise central 95% credible region for these local regression predictions, again averaged over the posterior distribution of γ.

Not suprisingly, inspection of Figure 2.13 reveals that rightwing partisanship, ideology, and news media consumption is increasing in γ. Respondents with the largest values of γ tend to be the respondents with the most right-leaning political views. Those with the lowest γ values tend to be the most left-leaning respondents.

We also examine whether respondent-specific γ values (and thus the perceptual framework that respondents use to evalute the truthfulness of COVID-19 statements) are associated with behaviors important for public health. More specifically, Figure 2.14 plots the relationship between the

[14]Each panel was constructed by fitting M local regressions of the variable in question on each of the M posterior samples of $\boldsymbol{\gamma}$. The pointwise average of these M estimated regression functions is the dark orange line in each panel. The light orange band in each panel is the pointwise central 95% credible region for these local regressions (the empirical 2.5th and 97.5th pointwise percentiles of the M estimated regression functions).

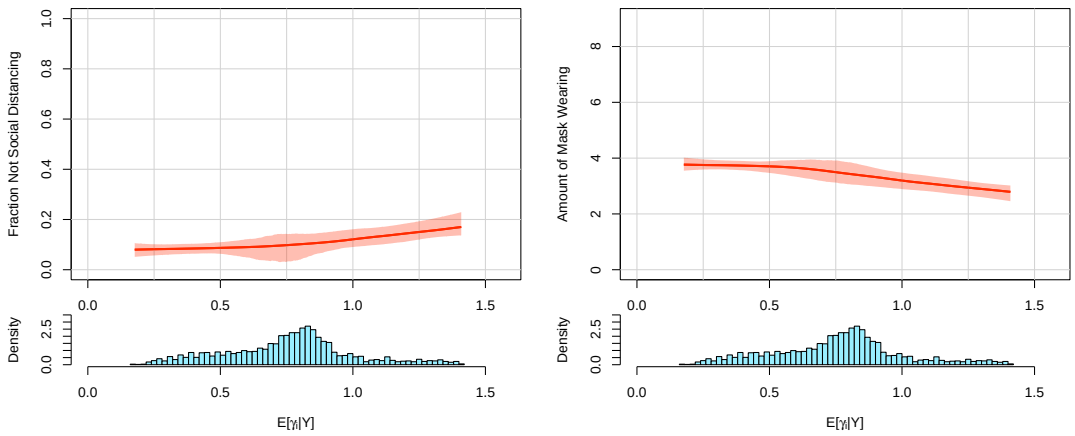

Figure 2.14: *Associations Between Posterior Means of γ_i for $i = 1 \ldots, N$ and Self-Reported Social-Distancing Behavior and Mask-Wearing Behavior.* The dark orange lines are the posterior means of local regression predictions averaged over the posterior distribution of γ. The light orange band is the pointwise central 95% credible region for these local regression predictions, again averaged over the posterior distribution of γ.

respondent-specific γ estimates and a) a measure of a lack of social distancing (operationalized as 0/1 indicator equal to 1 if a respondent said that 21 or more people were 6 feet or closer to them in the past week), and b) a measure of mask wearing (operationalized as the number of situations, out of nine possible, where the respondents said they wear a mask). The panels are constructed in the same way as Figure 2.13.

Figure 2.14 shows that the structure underlying how respondents judge the truthfulness of COVID-19 statements (as measured by their γ values) is associated with behaviors that have consequences for publich health. Specifically, lack of social distancing is increasing in γ, while mask wearing is decreasing in γ.

2.6 Conclusion

In this chapter, we have proposed a new multidimensional pairwise comparison model to measure multidimensional latent attributes and respondent-specific perceptual parameters. This model incorporates desirable constraints on respondent-specific parameters, and thereby has much better interpretabiltiy than previous multidimensional models. We derive a computationally efficient MCMC algorithm for estimating the new model and we will make the resulting code available to the public in the `MCMCpack` R package.

We apply this new model to original survey data where respondents are asked to judge the truthfulness of pairs of statements on COVID-19. Our analysis sheds light on how statements on COVID-19 are perceived by respondents and what respondent characteristics are associated with the perceptual frameworks used by respondents. Importantly, we find a weak correlation between the actual truthfulness of a statement and respondents' perceptions of truthfulness. More importantly, we find that the political valence of statements is largely responsible for the variation in perceived truthfulness. Co-partisanship between a respondent and the speaker of a statement predicts higher perceived truthfulness.

Additional findings directly speak to the puzzle about why and how misinformation on COVID-19 spreads. Our findings show that individuals generally have a hard time differentiating the scientific information on COVID-19 from the misinformation. Moreover, many respondents rely on partisanship as a cue to gauge the truthfulness of statements on COVID-19. Among these partisan respondents, the most rightward-leaning respondents' tend to view objectively truthful statements as subjectively false. Finally, we also observe associations between the respondent-specific perceptual parameters and respondents' practice of mask-wearing or social distancing. This shows that perceptions of information on COVID-19 have powerful consequences for what actions individuals take to cope with the pandemic.

CHAPTER 3

Dynamic Dirichlet Process Mixture Model for Identifying Voting Coalitions on Human Rights Roll Call Votes in the United Nations General Assembly

3.1 Introduction

In his 2018 address to the United Nations General Assembly (UNGA), President Trump stated that the United Nations' body overseeing human rights issues "had become a grave embarrassment to this institution, shielding egregious human rights abusers while bashing America and its many friends" (Trump, 2018). Indeed, others have lodged similar critiques of the United Nations' handling of human rights. Many critiques point out that human rights are highly politicized, making it difficult for the United Nations to fulfill its original mission of defending human rights (Habibi, 2007; Normand and Zaidi, 2008; Freedman, 2013; Hug, 2016).

One important avenue for countries to make their voices heard on human rights issues is to vote on human rights resolutions in the UNGA. These human rights resolutions concern both general human rights principles, such as the right to development, the use of mercenaries, children's and women's rights, and human rights violation reports on a specific country. Researchers have analyzed countries' human rights votes from two main perspectives. First, researchers believe that countries' votes on human rights resolutions reflect their individual preferences in terms of human rights protection, negligence, or violation (Boockmann and Dreher, 2011). Researchers found that

countries' domestic political regimes and domestic human rights records are predictive of their votes on human rights resolutions (Hug and Lukács, 2014). This approach is in line with many studies that apply multidimensional scaling models or Bayesian IRT models to the UNGA voting data to estimate countries' latent preferences in terms of global governance and world affairs (Voeten, 2000, 2004; Bailey, Strezhnev and Voeten, 2017; Bailey and Voeten, 2018).

Second, other researchers focus on the "peer group" effect. Studies of the United Nations have shown evidence for "voting blocs" in the UNGA vote data in the early days of the institution (Hovet, 1958). Namely, there are apparent indications of clustering of states' voting in the UNGA (Lijphart, 1963; Newcombe, Ross and Newcombe, 1970; Holloway, 1990). However, most of these studies take a qualitative approach, and explain this phenomenon in terms of the historical and political backgrounds. Specifically, researchers of human rights voting in the UNGA and the Human Rights Council of the UN also contend that "peer groups" of countries affect individual countries' votes on human rights resolutions (Boockmann and Dreher, 2011; Hug and Lukács, 2014; Hug, 2016). The "peer groups" may arise from geographic closeness, such as regional country groups, or cultural connections, such as countries sharing the same religion or language. Specifically, researchers contend that the European Union (EU) member countries tend to coalesce in human rights voting (Luif, 2003). Other researchers find that member countries of the Organization of Islamic Cooperation (OIC) are inclined to vote together on human rights issues (Besant and Malo, 2009).

Our study on the human rights votes in the UNGA engages with the "peer group" argument in the literature. The existing analyses rely solely on observable variables to group countries, such as geographic region or international organization membership. To the best of our knowledge, there has been no attempt to learn the latent clustering structure in the UNGA human rights votes data.[1] To better understand the politicization and the challenges of dealing with human rights issues in the UN, we propose a new statistical clustering model to analyze the UNGA human rights

[1]There are studies aimed at detecting communities based the entire UNGA votes in the network literature (Macon, Mucha and Porter, 2012; Pauls and Cranmer, 2017; Pomeroy, Dasandi and Mikhaylov, 2019). These studies illustrate some general community structures in the UNGA, but do not engage the specific topic of UNAG human rights votes.

vote data. The goal of this study is to inductively identify the voting coalitions in the UNGA human rights vote data without making assumptions about countries' political, historical, or cultural backgrounds.[2] We also identify the most polarizing resolutions that divide countries into different voting coalitions. The latent voting coalition structures and the polarizing resolutions advance our understanding of the lines of conflict on human rights and the driving force behind the politicization of human rights issues in the UNGA. In addition, our study also engages and contributes to the broad literature on model-based clustering methods for identifying voting blocs (Gormley and Murphy, 2008; Gormley, Murphy et al., 2008).

In this chapter, we propose a Dynamic Dirichlet Process Mixture (DDPM) model to identify voting coalitions based on roll call vote data across multiple time periods. We also propose post-processing methods for analyzing the outputs of the DDPM model. The strength of the Dirichlet Process Mixture (DPM) model lies in its weak assumptions and flexibility. In this study, the DPM model assumes an individual voter's vote on a specific bill as a Bernoulli trial.[3] We make less data generation assumptions than other common models applied to roll call vote data, such as multidimensional scaling models and Item Response Theory models. In addition, as a Bayesian non-parametrics clustering algorithm, the DPM model does not require users to pre-specify the number of clusters for data analyses. Theoretically, the method can identify as many clusters as there are individual voters. With proper priors, the DPM model lets the data decide how many clusters there should be. Moreover, as a strength of Bayesian modeling in general, the DPM model allows researchers to add hierarchical priors on appropriate parameters to model dynamic dependence.

Many existing roll call vote datasets cover decades of voting records in one institution. The dynamic nature of the datasets provide important information on the varying numbers of voting coalitions across time. Intentionally modeling the dynamic dependence of the numbers of voting coalitions allows the DDPM model to borrow information across time periods, resulting in more

[2] In this project, the term "voting coalition" is based on statistical patterns in the voting data, and it's different from what a "coalition" means in a game-theoretic model.

[3] In the extreme case where every voter is in her own coalition (cluster), DPM model jointly models NJ Bernoulli trials, which is equivalent to estimating NJ Bernoulli parameters, for a $N \times J$ voting matrix.

accurately identifying voting coalitions in each period. Aside from the purpose of borrowing information across time, the variation and trend of the numbers of voting coalitions across time could be the direct quantity of interest. For example, the number of voting coalitions in the UNGA human rights vote data bears on many important topics in international relations, such as polarity and politicization of human rights issues. To model the dynamic dependence among roll call votes across periods, we propose a Dynamic Linear Model (DLM) to model the precision parameters for each period's DPM. The precision parameter in a period directly reflects a researcher's prior belief about how many voting coalitions there should be in this period. The DLM on the precision parameters adds a dynamic dependence structure on modeling the numbers of voting coalitions across time.

Due to the label switching problem, interpreting the MCMC output from the DDPM model can be challenging. To aid applied researchers in analyzing the outputs of the DDPM models, we propose a Maximum A Posteriori estimation method to extract fixed cluster labels for each voter in a period based on the posterior samples. Moreover, we propose a visualization method to present the coalition structure in each time period, based on the posterior probability for any voter pair to be in the same cluster. We implement the DDPM model and post-processing methods in the `MCMCddpmbb` R package, and make it freely available to the public.

The rest of the chapter proceeds as follows: First, we illustrate the model setup for a single period in a roll call vote setting. Second, we specify the DLM structure to model time dependence in the data. Third, we outline the sampling scheme for the DDPM model. Fourth, we propose post-processing methods for analyzing the posterior samples of the DDPM model. Fifth, we show the DDPM model's effectiveness at uncovering the true latent variables and its robustness to various model specifications in multiple simulation studies. Last, we use the DDPM model to analyze the UNGA human rights roll call vote data after the Cold War.

3.2 Model Setup for a Single Period

We illustrate the proposed model for a single period in the setting of identifying voting coalitions (clusters) based on voters' roll call votes in a parliament-like institution across multiple periods. The number of time periods is T, i.e., $t = 1, ..., T$. The number of voters and bills can both vary across time. For any given time period t, there are I_t voters, i.e., $i = 1, ..., I_t$, and there are J_t bills (items), i.e., $j = 1, ..., J_t$. Let v_{ijt} be an indicator variable representing the observed vote of voter i on bill j at time t.

$$v_{ijt} = \begin{cases} 1 & \text{if voting yea} \\ 0 & \text{if voting nay} \\ NA & \text{if missing} \end{cases}$$

In each time period, a roll call vote matrix $\boldsymbol{V}_t$ stores all the data. Column i of $\boldsymbol{V}_t$ stores voter i's voting record vector $\boldsymbol{v}_{it}$ that consists of $1, 0$ or NA. We use J_t independent Bernoulli distributions to model each $\boldsymbol{v}_{it}$. For the voters belonging to the same cluster l, they share the same J_t Bernoulli parameters, which are stored in $\boldsymbol{\theta}_{lt}$. If voter i is in cluster l, then we know $v_{ijt} \sim$ Bernoulli(θ_{ljt}), for $j = 1, 2, \cdots, J_t$, and we express the likelihood of voter i's voting record, $\boldsymbol{v}_{it}$, as follows:

$$F(\boldsymbol{v}_{it}|\boldsymbol{\theta}_{lt}) = \prod_{j=1}^{J_t} \theta_{ljt}^{\mathbf{I}(v_{ijt}=1)} (1 - \theta_{ljt})^{\mathbf{I}(v_{ijt}=0)}$$

where $\mathbf{I}()$ is an indicator function.

We use a Beta-Binomial DPM model to cluster voters, which is equivalent to clustering all the $\boldsymbol{v}_{it}$ vectors. We put a Dirichlet Process prior on cluster l's Bernoulli parameter vector $\boldsymbol{\theta}_{lt}$:

$$\boldsymbol{\theta}_{lt}|G_t \sim G_t$$

$$G_t|\alpha_t, \boldsymbol{\lambda}_t \sim DP(\alpha_t G_{0t}(\cdot|\boldsymbol{\lambda}_t))$$

α_t is the precision parameter in period t. G_t does not have a specific parametric form but is assumed to be generated from a Dirichlet process, whose center distribution is G_{0t}. G_{0t} is the product of J_t independent Beta distributions, acting as the priors for an arbitrary cluster l's Bernoulli parameter vector $\boldsymbol{\theta}_{lt}$. $\boldsymbol{\lambda}_t$ is a J_t by 2 matrix, the j'th row of which stores the two parameters for the Beta prior distribution corresponding to bill j in time t. For an arbitrary new cluster l, the prior for its Bernoulli parameter vector, $G_{0t}(\boldsymbol{\theta}_{lt}|\boldsymbol{\lambda}_t)$, is expressed as follows:

$$G_{0t}(\boldsymbol{\theta}_{lt}|\boldsymbol{\lambda}_t) = \prod_{j=1}^{J_t} \{ \frac{\Gamma(\lambda_{1jt} + \lambda_{0jt})}{\Gamma(\lambda_{1jt})\Gamma(\lambda_{0jt})} \theta_{ljt}^{\lambda_{1jt}-1} (1 - \theta_{ljt})^{\lambda_{0jt}-1} \}$$

We introduce the parameter of direct interest, c_{it}, indicating the cluster affiliation of voter i in time t. Given $\boldsymbol{\lambda}_t$ and α_t, c_{it}'s are drawn separately across different time periods. When considering the distribution for c_{it}, we take the current cluster affiliations of all the other voters, $\boldsymbol{c}_{-i,t}$, as given. Assuming that there are currently L_t existing unique clusters in period t among all the voters except for voter i, then the cluster affiliation of voter i is determined by a categorical distribution. When drawing a new value of c_{it}, we know that the drawn value is either one of the existing cluster's label, or a new label representing a completely new cluster ($\boldsymbol{v}_{it}$ will be the first member of this new cluster). Therefore, the categorical distribution on c_{it} is defined on all the existing clusters' labels and a potentially newly generated cluster's label. We can express the probability for c_{it} as follows:

$$p(c_{it} = l, \text{ for } l \in \boldsymbol{c}_{-i,t} | \boldsymbol{c}_{-i,t}, \boldsymbol{V}_t) = b \frac{n_{-i,lt}}{I_t - 1 + \alpha_t} \int F(\boldsymbol{v}_{it}|\boldsymbol{\theta}_{lt}) dH_{-i,l}(\boldsymbol{\theta}_{lt})$$
$$p(c_{it} \neq c_{i't}, \forall i' \neq i | \boldsymbol{c}_{-i,t}, \boldsymbol{V}_t) = b \frac{\alpha_t}{I_t - 1 + \alpha_t} \int F(\boldsymbol{v}_{it}|\boldsymbol{\theta}_{0t}) dG_{0t}(\boldsymbol{\theta}_{0t})$$

$\boldsymbol{c}_{-i,t}$ represents all voters' cluster affiliations except for voter i's, c_{it}. b is an appropriate normalizing constant. $n_{-i,lt}$ is the size of cluster l, not considering voter i. $F(\boldsymbol{v}_{it}|\boldsymbol{\theta}_{lt})$ is the likelihood above, and $H_{-i,l}$ is the updated prior density of $\boldsymbol{\theta}_{lt}$ based on the Dirichlet Process prior G_{0t} and all the vote

vectors belonging to cluster l, except for voter i's. Note that we analytically integrate out $\boldsymbol{\theta}_{lt}$ and $\boldsymbol{\theta}_{0t}$ vectors by relying on the conjugate Beta prior, so we never sample $\boldsymbol{\theta}$ parameters in the MCMC sampler. The intuition to use the product of cluster l's size and the likelihood of voter i's vote vector is to induce sparsity of cluster assignment. Without considering cluster sizes, assigning each voter to her own cluster is guaranteed to achieve the largest likelihood for the data. However, identifying only single-member clusters provides no additional insight beyond the information presented by the raw data.

In a single period, we draw c_{it}'s with the Chinese Restaurant Process (Blackwell, MacQueen et al., 1973; Aldous, 1985; Pitman, 1996). The Chinese Restaurant Process induces large clusters to grow larger and small clusters to disappear. The Chinese Restaurant Process is not a very strong assumption in itself, because we can theoretically tune the algorithm with the precision parameter, α_t, to achieve a vast possibility of cluster configurations. α_t acts as the imaginary cluster size for a cluster yet to be generated. If $\alpha_t \to 0$, then we put all voters in one cluster, which is equivalent to a completely pooled model. If $\alpha_t \to +\infty$, then we put each voter in a single-member cluster, which is equivalent to separately modeling each voter's vote vector without borrowing information. We can add proper priors on α_t's to let the data tell how many clusters there should be, and this modeling strategy reflects the flexibility of DPM.

We assume that the values of $\boldsymbol{\lambda}_t$ are known a priori. As pointed out in Spirling and Quinn (2010), roll call vote data usually do not provide enough information for sampling $\boldsymbol{\lambda}_t$ parameters, so models that sample values of $\boldsymbol{\lambda}_t$ are not ideal. This finding is also confirmed by our exploration with models that sample $\boldsymbol{\lambda}_t$ with simulated data. Sampling $\boldsymbol{\lambda}_t$ prevents the posterior samples of cluster labels from converging to the true values. For this chapter, we assume that all values in each matrix $\boldsymbol{\lambda}_t$ are 1. This is equivalent to assuming a uniform$(0, 1)$ prior distribution for all θ_{ljt}'s. However, it's reasonable to assume other values for elements in $\boldsymbol{\lambda}_t$ based on alternative theories or prior beliefs.

We assume that a voter's voting record on a bill is missing at random (Rubin, 1976). The implication of this assumption is that a missing voting record does not affect estimating any voter's

cluster affiliation. For a single period, if voter i's voting record on bill j is missing ($v_{ijt} = NA$), then the likelihood of v_{ijt}, $F(v_{ijt})$, is a constant in voter i's cluster affiliation, c_{it}.[4] Therefore, v_{ijt} does not affect computing the weight parameters for the categorical distribution for drawing voter i's cluster affiliation, c_{it}, and only observed voting records for voter i have an impact. Similarly, suppose voter i' belongs to cluster l for the current sampling iteration, and we are about to draw the cluster affiliation parameter for voter i, c_{it}. Then voter i''s missing voting record on bill j does not affect computing $H_{-i,l}(\boldsymbol{\theta}_{lt})$, the updated prior density of $\boldsymbol{\theta}_{lt}$. [5]

3.3 Modeling Dynamic Dependence

For each period t, the precision parameter α_t is a very important parameter. The value of α_t reflects a researcher's prior belief about the number of clusters in period t. Let L_t be the number of clusters in period t. The prior distribution for L_t depends critically on α_t, and is stochastically increasing with α_t. As shown in the DPM model literature, the approximate conditional expectation of L_t, given α_t, the precision parameter, and I_t, the number of voters, can be expressed as follows: $E(L_t|\alpha_t, I_t) \approx \alpha_t \ln(1 + \frac{I_t}{\alpha_t})$ (Escobar and West, 1995). α_t is always positive, so for a single time period t, we put a log-normal prior on α_t.

$$\ln(\alpha_t) \sim N(\gamma_t, V)$$

We assume an underlining Bayesian Dynamic Linear Model (DLM) on $z_t = \log(\alpha_t)$ for the purpose of smoothing the values of α_t's by borrowing information from adjacent time periods. The reason for adding this dynamic smoothing structure on α_t's is that adjacent time periods should show similar degrees of division or cooperation in the voting records. Therefore, the prior belief

[4] If $v_{ijt} = NA$, then $F(v_{ijt}|\theta_{ljt}) = \theta_{ljt}^{\mathbf{I}(v_{ijt}=1)}(1-\theta_{ljt})^{\mathbf{I}(v_{ijt}=0)} =$ constant. It does not matter if we put voter i in any existing cluster or a newly generated cluster, because $F(v_{ijt}|\theta_{ljt})$ is the same constant. Therefore, the missing v_{ijt} does not affect voter i's cluster affiliation sample draws.

[5] The updated Beta prior, $H_{-i,l}(\boldsymbol{\theta}_{lt})$, is based on the Dirichlet Process prior G_{0t} and all the observed voting records of the voters belonging to cluster l. We put a conjugate Beta prior for θ_{ljt}, and we only use observed voting records of the voters in cluster l for updating the Beta prior of θ_{ljt}. A missing voting record for any voter belonging to cluster l does not affect this process.

of the number of clusters for one period should be correlated with those in adjacent periods. The sampler will be able to borrow information from draws of α_t in neighboring periods for sampling α_t for a specific period.

Specifically, we put a random walk DLM prior on z_t. The random walk DLM prior specification is extensively used in various Bayesian dynamic IRT models for modeling latent variables in political science research, such as the Supreme Court Justices' ideal points (Martin and Quinn, 2002; Bailey, 2013) and countries' human rights scores (Schnakenberg and Fariss, 2014). The popularity of the random walk DLM prior is due to its flexibility and the efficient Gibbs sampler developed for it. The random walk DLM prior on z_t is specified as follows:

$$
\begin{aligned}
z_t &\sim N(\gamma_t, V) \\
\gamma_t &\sim N(\gamma_{t-1}, W) \\
\gamma_0 &\sim N(m_0, H_0) \\
V &\sim \text{Inverse-Gamma}(\frac{r_0}{2}, \frac{s_0}{2}) \\
W &\sim \text{Gamma}(r_1, s_1)
\end{aligned}
$$

V is the variance of the Gaussian disturbance, e_2, in the observation equation, $z_t = \gamma_t + e_2, e_2 \sim N(0, V)$. V has a conjugate Inverse-Gamma prior with user-supplied parameters r_0 and s_0. W is the variance of the Gaussian disturbance, e_1, in the evolution equation, $\gamma_t = \gamma_{t-1} + e_1, e_1 \sim N(0, W)$. The smaller W is, the more smoothing of z_t is achieved across time, whereas large W value allows the value of z_t to change more drastically across time. W has a Gamma prior with user-supplied shape parameter r_1 and rate parameter s_1. We choose a Gamma prior for W over a conjugate Inverse-Gamma prior commonly used in the DLM literature (Petris, Petrone and Campagnoli, 2009). The Gamma prior is a better alternative, because of its potentially desirable shrinkage property (Bitto and Frühwirth-Schnatter, 2019). If a user believes that cluster numbers are very stable across time, then she can pull the estimated W toward a small value, by tuning r_1

and s_1 to reduce both the prior mean $\frac{r_1}{s_1}$ and the prior variance $\frac{r_1}{s_1^2}$. Conversely, if a user does not have a strong belief about the stability of cluster numbers across time, then she can use a more diffused Gamma prior on W, and little shrinkage effect will take place in this case. m_0 and H_0 are the mean and variance for Gaussian prior on γ_0 for the imaginary period before period 1. Except for the Gamma prior on W, the above specification is the standard setup of a unidimensional Gaussian random walk DLM (West and Harrison, 2006).

3.4 Sampling Scheme

The sampler for the entire model is a mixture of Gibbs sampler and Metropolis-Hastings sampler. Due to conjugate prior specifications, the DPM part of the model only relies on a Gibbs sampler for sampling c_{it}'s. The conjugate property breaks down for sampling α_t, so we use a random walk Metropolis-Hastings step to sample α_t in each period separately. Given the last draws of α_t's, the DLM part of the model mostly relies on a Gibbs sampler, except for the sampling step on W, which employs a random walk Metropolis-Hastings sampler. We lay out the sampling scheme in the following three parts:

1. For each period t, we sample cluster labels, c_{it}'s, given the current value of α_t, vote matrix, $\boldsymbol{V}_t$, and user-supplied values of $\boldsymbol{\lambda}_t$.

2. For each period t, we draw the precision parameter, α_t, given the current value of γ_t, the current number of clusters, L_t, and the constant number of voters, I_t.

3. Given the current values of α_t's and the user-supplied values of $m_0, H_0, r_0, s_0, r_1, s_1$, we sample the latent means, γ_t's, the observation equation variance, V, and the innovation equation variance, W, in the DLM model.

3.4.1 Sample $c_{it} \mid \boldsymbol{c}_{-i,t}, \alpha_t, \boldsymbol{V}_t, \boldsymbol{\lambda}_t$,

Cluster affiliation parameters, c_{it}, are sampled separately in each time period given $\boldsymbol{V}_t$ and $\boldsymbol{\lambda}_t$, so we drop subscript t for simplification in this subsection. Every cluster affiliation corresponds to either a unique existing cluster, or to a newly generated cluster. $\boldsymbol{c}_{-i}$ represents all voters' cluster affiliations except for voter i's, c_i. We use the Chinese Restaurant Process to draw c_i given $\boldsymbol{c}_{-i}$. Voter i is assigned to a new cluster or an existing cluster with a categorical distribution. Due to the conjugacy of the Beta-Binomial prior-likelihood specification, we can integrate out $\boldsymbol{\theta}_0$ or $\boldsymbol{\theta}_l$ when computing the weights for a new cluster or the existing clusters (Neal, 2000). Therefore, the sampler never directly samples or stores the values of $\boldsymbol{\theta}_0$ or $\boldsymbol{\theta}_l$.

Let $n_{-i,l}$ be the size for an existing cluster l, not considering voter i. For each existing cluster l (not considering voter i), we define the count of yea vote on bill j as $n_{-i,lj1}$, and the count of nay vote on bill j as $n_{-i,lj0}$, based on all the affiliating voters' observed voting records. A yet-to-be-generated cluster has α size and zero counts of yea or nay vote on any bill. Let b be an appropriate normalizing constant, and I be the number of voters. We express the conditional posterior probability of c_i as below.[6] We can thus directly draw c_i with a closed-form Gibbs sampler.

$$
\begin{aligned}
p(c_i = l, \text{ for } l \in \boldsymbol{c}_{-i} | \boldsymbol{c}_{-i}, \boldsymbol{V}) = b \frac{n_{-i,l}}{I - 1 + \alpha} \prod_{j=1}^{J} \Bigg(& \Big(\frac{n_{-i,lj1} + \lambda_{1j}}{n_{-i,lj1} + n_{-i,lj0} + \lambda_{0j} + \lambda_{1j}} \Big)^{\mathbf{I}(v_{ij}=1)} \\
& \times \Big(\frac{n_{-i,lj0} + \lambda_{0j}}{n_{-i,lj1} + n_{-i,lj0} + \lambda_{0j} + \lambda_{1j}} \Big)^{\mathbf{I}(v_{ij}=0)} \Bigg) \\
p(c_i \neq c_{i'}, \forall i' \neq i | \boldsymbol{c}_{-i}, \boldsymbol{V}) = b \frac{\alpha}{I - 1 + \alpha} \prod_{j=1}^{J} & \Bigg(\Big(\frac{\lambda_{1j}}{\lambda_{0j} + \lambda_{1j}} \Big)^{\mathbf{I}(v_{ij}=1)} \Big(\frac{\lambda_{0j}}{\lambda_{0j} + \lambda_{1j}} \Big)^{\mathbf{I}(v_{ij}=0)} \Bigg)
\end{aligned}
$$

[6] Please refer to section, "Example: Bernoulli Data with a Conjugate Beta Prior", in Jain and Neal (2004) for the derivations of the Gibbs sampler for c_i. We adapt the notations in Jain and Neal (2004) to express the formula in this chapter.

Researchers find that pure Gibbs samplers for DPM models have poor mixing, so it's necessary to attempt at merging and splitting existing clusters to improve mixing (Jain and Neal, 2004). Therefore, in addition to the Gibbs sampler for c_i, we also use the Sequentially-Allocated Merge-Split (SAMS) sampler for c_i. We implement the above Gibbs sampler and SAMS sampler in alternative turns. Interested readers should refer to Dahl (2003) for the details of SAMS sampler.

3.4.2 Sample α_t (z_t) $\mid \gamma_t, V, L_t, I_t$

For period t, the conditional likelihood of α_t, given the number of clusters L_t and the number of voters I_t, is $p(\alpha_t|L_t, I_t) \propto \alpha_t^{L_t} \frac{\Gamma(\alpha_t)}{\Gamma(\alpha_t+I_t)}$ (Antoniak, 1974; Escobar and West, 1995, 1998). We have a log-normal prior on α_t, $\ln(\alpha_t) \sim N(\gamma_t, V)$. We can express the conditional posterior of α_t as follows:

$$p(\alpha_t|\gamma_t, V, L_t, I_t) \propto \underbrace{\frac{1}{\alpha_t} \exp\Big(- \frac{(\ln \alpha_t - \gamma_t)^2}{2V} \Big)}_{\text{log-normal prior kernel}} \underbrace{\alpha_t^{L_t} \frac{\Gamma(\alpha_t)}{\Gamma(\alpha_t + I_t)}}_{\text{conditional likelihood}}$$

The conditional posterior of α_t does not have a kernel of any recognizable distribution. We use a unidimensional random walk Metropolis-Hastings sampler to sample α_t for each period separately, and users supply the positive value τ to generate the random walk step, δ , from $\delta \sim$ uniform$(-\tau, \tau)$ for updating α_t. We accept the proposed new $\alpha_{t,\text{ new}} = \alpha_t + \delta$ value, with the following acceptance rate.

$$\min \Big\{ 1, \frac{p(\alpha_{t,\text{ new}}|\gamma_t, V, L_t, I_t)}{p(\alpha_t|\gamma_t, V, L_t, I_t)} \Big\}$$

Users can monitor the acceptance rate for α_t and tune the τ parameter to achieve the desirable acceptance rate around 0.45 (Roberts et al., 1997; Chib and Greenberg, 1995*b*). After sampling α_t, we can use an logarithm transformation to directly get z_t by $z_t = \ln(\alpha_t)$.

3.4.3 Sample $\{\gamma_t\}_{t=0}^T$, V, $W \mid z_t, m_0, H_0, r_0, s_0$

Based on the random walk DLM prior imposed on z_t's, we use a Gibbs sampler to draw $\{\gamma_t\}_{t=0}^T$, V and W, given the latest draws of z_t. We use the efficient Forward Filtering Backward Sampling (FFBS) algorithm for this Gibbs sampler. Interested readers can refer to Carter and Kohn (1994), Frühwirth-Schnatter (1994) or West and Harrison (2006) for the detailed proof for the FFBS algorithm. We lay out the sampling steps below.

Let D_t denote all the information available up to period t, and we know $D_t = \{D_{t-1}, z_t\}$. For period 0, we know $D_0 = \{m_0, H_0, V, W\}$. We exploit the factorization of the joint probability of $\{\gamma_t\}_{t=1}^T$:

$$p(\{\gamma_t\}_{t=0}^T) = p(\gamma_T|D_T)p(\gamma_{T-1}|\gamma_T, D_T)\cdots p(\gamma_1|\gamma_2, D_2)p(\gamma_0|\gamma_1, D_1)$$

We can thus draw a joint sample of $\{\gamma_t\}_{t=0}^T$ by sampling $p(\gamma_T|D_T)$ first, then we can subsequently draw γ_t by $p(\gamma_t|\gamma_{t+1}, D_{t+1})$ for $t = T-1, T-2, \cdots, 1, 0$.

Given the current values of $\{z_t\}_{t=1}^T$, V and W, and the fixed values of m_0 and H_0, we start the forward filtering procedure. For $t = 1, 2, \cdots, T$, we compute the following quantities in sequence. First, we compute, a_t, the prior mean of γ_t, given D_{t-1} by $a_t = E[\gamma_t|D_{t-1}] = E[\gamma_{t-1}+e_1] = m_{t-1}$. Second, we compute, R_t, the prior variance of γ_t, given D_{t-1} by $R_t = Var[\gamma_t|D_{t-1}] = Var[\gamma_{t-1}+e_1] = H_{t-1} + W$. Third, we compute, f_t, the prior mean of the one-step forecast of z_t, given D_{t-1} by $f_t = E[z_t|D_{t-1}] = E[\gamma_t + e_2|D_{t-1}] = a_t$. Fourth, we compute, Q_t, the prior variance of the one-step forecast of z_t, given D_{t-1} by $Q_t = Var[z_t|D_{t-1}] = Var[\gamma_t + e_2|D_{t-1}] = R_t + V$. Fifth, we compute, m_t, the posterior mean of γ_t given D_t by $m_t = E[\gamma_t|D_t] = E[\gamma_t|D_{t-1}] + \frac{Var[\gamma_t|D_t]}{Var[z_t|D_{t-1}]}\Big(z_t - E[z_t|D_{t-1}]\Big) = a_t + \frac{R_t}{H_t}(z_t - f_t)$. Sixth, we compute, H_t, the posterior variance of γ_t given D_t by $H_t = Var[\gamma_t|D_t] = Var[\gamma_t|D_{t-1}] - \Big(\frac{Var[\gamma_t|D_{t-1}]}{Var[z_t|D_{t-1}]}\Big)^2 Var[z_t|D_{t-1}] = R_t - (\frac{R_t}{Q_t})^2 Q_t$.

After the above forward filtering procedure, we store the following values: $\{a_t\}_{t=1}^T$, $\{R_t\}_{t=1}^T$, $\{f_t\}_{t=1}^T$, $\{Q_t\}_{t=1}^T$, $\{m_t\}_{t=1}^T$, and $\{H_t\}_{t=1}^T$. We start the backward sampling procedure from $t = T$,

since in the factorization of $p(\{\gamma_t\}_{t=0}^T)$, the PDF of γ_T is only conditioning on D_T. We can directly draw γ_T from the conditional posterior below:

$$\gamma_T|D_T \sim N(m_T, H_T)$$

Sequentially, for $t = T-1, \cdots, 2, 1, 0$, we compute, u_t, the mean of the conditional distribution of γ_t, given γ_{t+1} and D_t, by $u_t = E[\gamma_t|\gamma_{t+1}, D_t] = E[\gamma_t|D_t] + \frac{Var[\gamma_t|D_t]}{Var[\gamma_{t+1}|D_t]}(\gamma_{t+1} - E[\gamma_{t+1}|D_t]) = m_t + \frac{H_t}{R_{t+1}}(\gamma_{t+1} - a_{t+1})$. Then we compute, U_t, the variance of the conditional distribution of γ_t, given γ_{t+1} and D_t, by $U_t = Var[\gamma_t|\gamma_{t+1}, D_t] = Var[\gamma_t|D_t] - \big(\frac{Var[\gamma_t|D_t]}{Var[\gamma_{t+1}|D_t]}\big)^2 Var[\gamma_{t+1}|D_t] = H_t - \big(\frac{H_t}{R_{t+1}}\big)^2 R_{t+1}$. Given the values of u_t and U_t, we update γ_t by its full conditional posterior below.

$$\gamma_t|\gamma_{t+1}, D_t \sim N(u_t, U_t)$$

Next, we update the value of V, given the updated values of $\{\gamma_t\}_{t=1}^T$ and $\{z_t\}_{t=1}^T$. V has a conjugate Inverse-Gamma prior distribution, $V \sim$ Inverse-Gamma$(\frac{r_0}{2}, \frac{s_0}{2})$. The conditional posterior distribution of V is expressed below.

$$V|\{\gamma_t\}_{t=1}^T, \{z_t\}_{t=1}^T \sim \text{Inverse-Gamma}\big(\frac{r_0 + T}{2}, \frac{s_0 + \sum_{t=1}^T (z_t - \gamma_t)^2}{2}\big)$$

Lastly, we sample the value of W, given the updated values of $\{\gamma_t\}_{t=0}^T$. W has a non-conjugate Gamma prior distribution, $W \sim$ Gamma(r_1, s_1). We can express the conditional posterior of W as follows:

$$p(W|\{\gamma_t\}_{t=0}^T, r_1, s_1) \propto \underbrace{W^{r_1 - 1}\exp(-s_1 W)}_{\text{Gamma prior kernel}} \underbrace{\prod_{t=1}^T \frac{1}{\sqrt{W}} \exp\big(-\frac{(\gamma_t - \gamma_{t-1})^2}{2W}\big)}_{\text{conditional likelihood kernel}}$$

The conditional posterior of W does not have a kernel of any recognizable distribution. We use a unidimensional random walk Metropolis-Hastings sampler to sample W. Users supply the

positive value τ_W to generate the random walk step, δ_W , from $\delta_W \sim \text{uniform}(-\tau_W, \tau_W)$ for updating W. We accept the proposed new $W_{\text{new}} = W + \delta_W$ value, with the following acceptance rate. Users can monitor the acceptance rate for W and tune the τ_W parameter to achieve the desirable acceptance rate around 0.45.

$$\min\left\{1, \frac{p(W_{\text{new}}|\{\gamma_t\}_{t=0}^{T}, r_1, s_1)}{p(W|\{\gamma_t\}_{t=0}^{T}, r_1, s_1)}\right\}$$

3.5 Post-Processing

Due to the label switching problem present in any DPM model, we are not able to use simple statistical summaries of the exact draws of cluster labels for analysis (Jasra, Holmes and Stephens, 2005). Instead, we propose a maximum a posteriori (MAP) estimation method to extract fixed cluster labels for voters. This method draws on the post-processing procedure proposed for both finite and infinite mixture models in the model-based clustering literature (Frühwirth-Schnatter, 2006; Frühwirth-Schnatter and Malsiner-Walli, 2019). We compute the MAP cluster labels for each period t separately, so we drop subcript t in the following illustration for simplification.

3.5.1 Compute the Cluster Number Mode and Subset the Posterior Samples

In the posterior sample of cluster labels, $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{\text{iter}=1}^{N}$, the number of unique clusters, L_{iter}, varies across iterations, which reflects DPM's strength in letting the data decide how many clusters there should be. We compute the mode of cluster numbers, L_{mode}, across iterations. Then, we extract a subset of the cluster label posterior sample iterations, $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{\text{iter}\in\{1,2,\cdots,N|L_{\text{iter}}=L_{\text{mode}}\}}$, where the number of clusters is equal to the mode. We denote $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{\text{iter}\in\{1,2,\cdots,N|L_{\text{iter}}=L_{\text{mode}}\}}$ as $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{L_{\text{iter}}=L_{\text{mode}}}$. By using the mode as the subseting criterion, we let the algorithm decide the number of clusters in our final result. Working with the above subset of the cluster label posterior samples, we have a fixed number of clusters, L_{mode}, across iterations. Therefore, we can use the likelihood of a finite mixture model as an approximate likelihood function, in order to get

the MAP cluster labels in the next step.

3.5.2 Identify the MAP Cluster Labels

We compute the likelihood for each iteration in $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{L_{\text{iter}}=L_{\text{mode}}}$. First, we compute, w_l, the marginal probability for a voter to belong to cluster, l, by the ratio of cluster l's size, n_l, to the number of voters, I: $w_l = \frac{n_l}{I}$. Second, for each cluster l, we can compute the posterior Bernoulli probability parameter vector, $\tilde{\boldsymbol{\theta}}_l$, based on the voters' voting records, who belong to cluster l. For a specific bill j, we compute the Bernoulli parameter by $\tilde{\theta}_{lj} = \dfrac{\lambda_{1j} + n_{lj1}}{\lambda_{1j} + \lambda_{0j} + n_{lj1} + n_{lj0}}$, where n_{lj1} is the count of yea vote on bill j from all the voters belonging to cluster l, n_{lj0} is the count of nay vote on bill j from all the voters belonging to cluster l, and λ_{1j} and λ_{0j} are the Beta prior parameters supplied by users.

We compute the approximate posterior likelihood for a cluster label posterior sample iteration as follows:

$$p(\{c_{i,\text{iter}}\}_{i=1}^{I}|\boldsymbol{V},\boldsymbol{\lambda}) = \prod_{i=1}^{I}\left(\prod_{l}\left[w_l\Big(\prod_{j=1}^{J}\tilde{\theta}_{lj}^{\mathbf{I}(v_{ij}=1)}(1-\tilde{\theta}_{lj})^{\mathbf{I}(v_{ij}=0)}\Big)\right]^{\mathbf{I}(c_{i,\text{iter}}=l)}\right)$$

We compute the above approximate posterior likelihood for each iteration in $\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{L_{\text{iter}}=L_{\text{mode}}}$, and use the iteration of cluster label draws, $\{c_{i,\text{iter}}\}_{i=1}^{I}$, corresponding to the largest approximate posterior likelihood value, as the MAP cluster labels. The MAP cluster labels are expressed as below.

$$\underset{\{c_{i,\text{iter}}\}_{i=1}^{I}\in\left\{\{c_{i,\text{iter}}\}_{i=1}^{I}\right\}_{L_{\text{iter}}=L_{\text{mode}}}}{\text{argmax}} \quad p(\{c_{i,\text{iter}}\}_{i=1}^{I}|\boldsymbol{V},\boldsymbol{\lambda})$$

3.5.3 Visualization of Voter-Pair Posterior Probability of Being in the Same Cluster

In addition to the MAP cluster labels, we propose a visualization method to present the clustering structure in a time period. Based on the posterior sample of cluster labels, we compute the posterior probability for any pair of voters, i_1 and i_2, to be in the same cluster across iterations in time t. These probabilities are stored in an $I_t \times I_t$ symmetric square matrix, M_t, for period t, and element (i_1, i_2) of M_t indicates the posterior probability for voters, i_1 and i_2, to be in the same cluster. We use Ward's Hierarchical Agglomerative Clustering Method to reorder the rows and columns of M_t (Murtagh and Legendre, 2014). We use the Ward's dissimilarity criterion to move the similar columns and rows to be adjacent to each other. `hclust()` function in `R` implements this algorithm with various distance criteria, and we use `ward.D2` distance criterion for this study. The resulting reordered matrix M_t' has anti-diagonal blocks storing larger probability values and off-anti-diagonal elements storing lesser probability values. We plot matrix M_t' like a heat map to visualize the clustering structure.

3.6 Simulation Study

In this section, we report findings from various simulation studies. First, we show how the DDPM model is able to achieve high accuracy at both identifying the correct cluster numbers and classifying the relationships of unique voter pairs in each period, based on different simulated data sets. Second, we simulate cluster numbers in a Monte Carlo experiment for an imaginary time period, and show how hyperparameter specifications affect the prior distribution for the number of clusters. Lastly, we run multiple robustness checks to show that the DDPM model is able to consistently identify the true cluster numbers under various model specifications.

Period	1	2	3	4	5	6	7	8	9	10
number of voters	45	51	52	57	40	49	46	54	59	58
number of items	75	80	73	55	49	76	52	74	69	72
number of clusters	4	4	4	2	2	3	2	3	2	2

Table 3.1: *True Data Structure in the Stable-Cluster-Number Example.*

3.6.1 Two Demonstrative Examples

To demonstrate the effectiveness of the DDPM model, we run the DDPM model on two simulated data sets of multiple periods. In one data set, the numbers of clusters are relatively stable across time, and in the other data set, the numbers of clusters fluctuate in a more volatile way. There are 10 periods in total, and in each period t we randomly sample (with equal likelihood and replacement) the number of voters, I_t, among the integers between 40 and 60, and the number of bills, J_t, among the integers between 40 and 80.

We use the following two DLM models to specify the true numbers of clusters across periods. For the stable-cluster-number example, we set $\gamma_0 = 0$, $W = 0.05$, $V = 0.05$. For the volatile-cluster-number example, we set $\gamma_0 = 0$, $W = 0.1$, $V = 0.05$. For each example respectively, We generate the values of $\gamma_1, \gamma_2, \cdots, \gamma_T$ by the evolution equation in DLM: $\gamma_t = \gamma_{t-1} + e_1, e_1 \sim N(0, W)$. After obtaining all the γ_t's, we generate all the z_t's through the observation equation in the DLM: $z_t = \gamma_t + e_2, e_2 \sim N(0, V)$. All the α_t's directly follow by the transformation, $\alpha_t = \exp(z_t)$. Given the number of voters, I_t, and α_t, we generate the true numbers of clusters by the function, $L_{t,true} = \text{round}(E(L_t|\alpha_t, I_t))$, where $E(L_t|\alpha_t, I_t) \approx \alpha_t \ln(1 + \frac{I_t}{\alpha_t})$ (Escobar and West, 1995). The true numbers of clusters, voters and items for each period in the stable-cluster-number example are reported in the Table 3.1, and those in the volatile-cluster-number example are reported in the Table 3.2.

In both examples, a roll call vote matrix is generated with underling Bernoulli distributions, for each period t. Given the number of clusters, L_t, we generate L_t Bernoulli parameter vectors of length J_t, $\boldsymbol{\theta}_{lt}$, by generating each element of $\boldsymbol{\theta}_{lt}$ from a Beta distribution, $\theta_{ljt} \sim \text{Beta}(1, 1)$. For

Period	1	2	3	4	5	6	7	8	9	10
number of voters	53	49	57	42	47	56	54	59	44	51
number of items	67	53	52	65	74	40	73	77	72	56
number of clusters	3	4	2	4	3	5	5	4	4	7

Table 3.2: *True Data Structure in the Volatile-Cluster-Number Example.*

period t, we distribute all voters into L_t clusters evenly.[7] For all the voters that belong to cluster l, we simulate their vote vectors, $\boldsymbol{v}_{it}$, with the Bernoulli parameter vector, $\boldsymbol{\theta}_{lt}$, through the following Bernoulli distributions, $v_{ijt} \sim \text{Bernoulli}(\theta_{ljt})$, for $j = 1, 2, \cdots, J_t$. The roll call vote matrices of ten periods are supplied as the data set to the DDPM model sampler.

The DDPM model specifications in the stable-cluster-number example and the volatile-cluster-number example are the same. We run the SAMS sampler every third of the scan. The tuning parameter for the random walk Metropolis-Hastings step for drawing α_t's, τ, is 1. We set the parameters in the DLM priors as follow: $m_0 = 0, H_0 = 0.1, r_0 = 10, s_0 = 0.5, r_1 = 1, s_1 = 50$. The tuning parameter for the random walk Metropolis-Hastings step for drawing W, τ_W, is 0.02. We use the following values as starting values for V, W and all α_t's: $V_{start} = 0.05$, $W_{start} = 0.02$, and $\alpha_{t,start} = 1$ for $t = 1, 2, \cdots, T$. We run 2000 iterations and discard the first 1000 iterations as burnin. In the stable-cluster-number example, the acceptance rates of the random walk Metropolis-Hastings steps for drawing α_t's lie in the interval, $[0.2430, 0.3250]$. The acceptance rate of the random walk Metropolis-Hastings step for drawing W is 0.5485. In the volatile-cluster-number example, the acceptance rates of the random walk Metropolis-Hastings steps for drawing α_t's lie in the interval, $[0.3320, 0.4730]$. The acceptance rate of the random walk Metropolis-Hastings steps for drawing W is 0.5215. These acceptance rates are all acceptable for a unidimensional Metropolis-Hastings sampler.

In Figure 3.1 and Figure 3.2, we report the cluster number histograms against the true values indicated by the dashed lines across periods in both examples. In both examples, the modes of cluster numbers are all equal to the true cluster numbers across all periods.

[7]For voter i, we get the modulo m from the ratio between $i - 1$ and L_t, and we assign voter i to cluster $m + 1 \in \{1, 2, \cdots, L_t\}$.

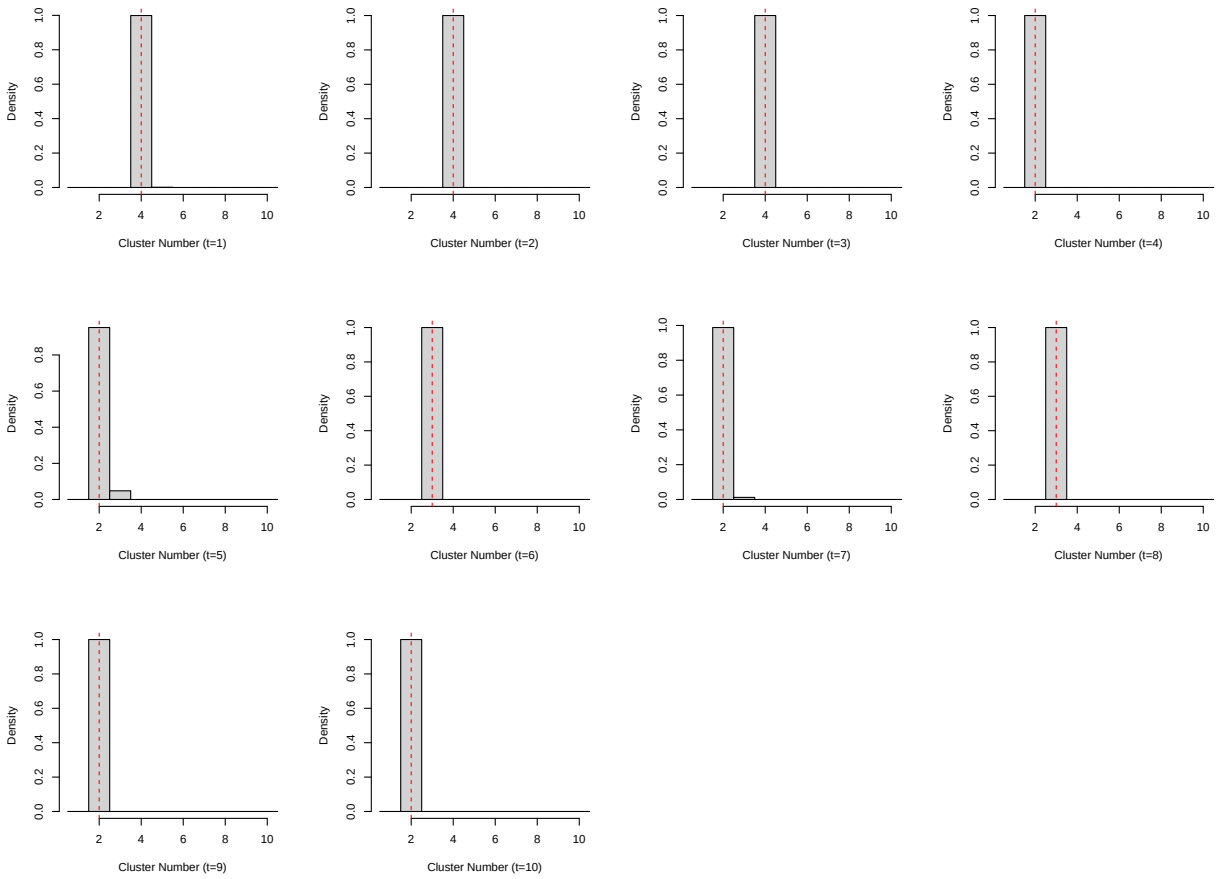

Figure 3.1: *Cluster Number Histograms in the Stable-Cluster-Number Example.* The figures show the histograms of the numbers of clusters across iterations in each period. The dashed lines indicate the true cluster numbers.

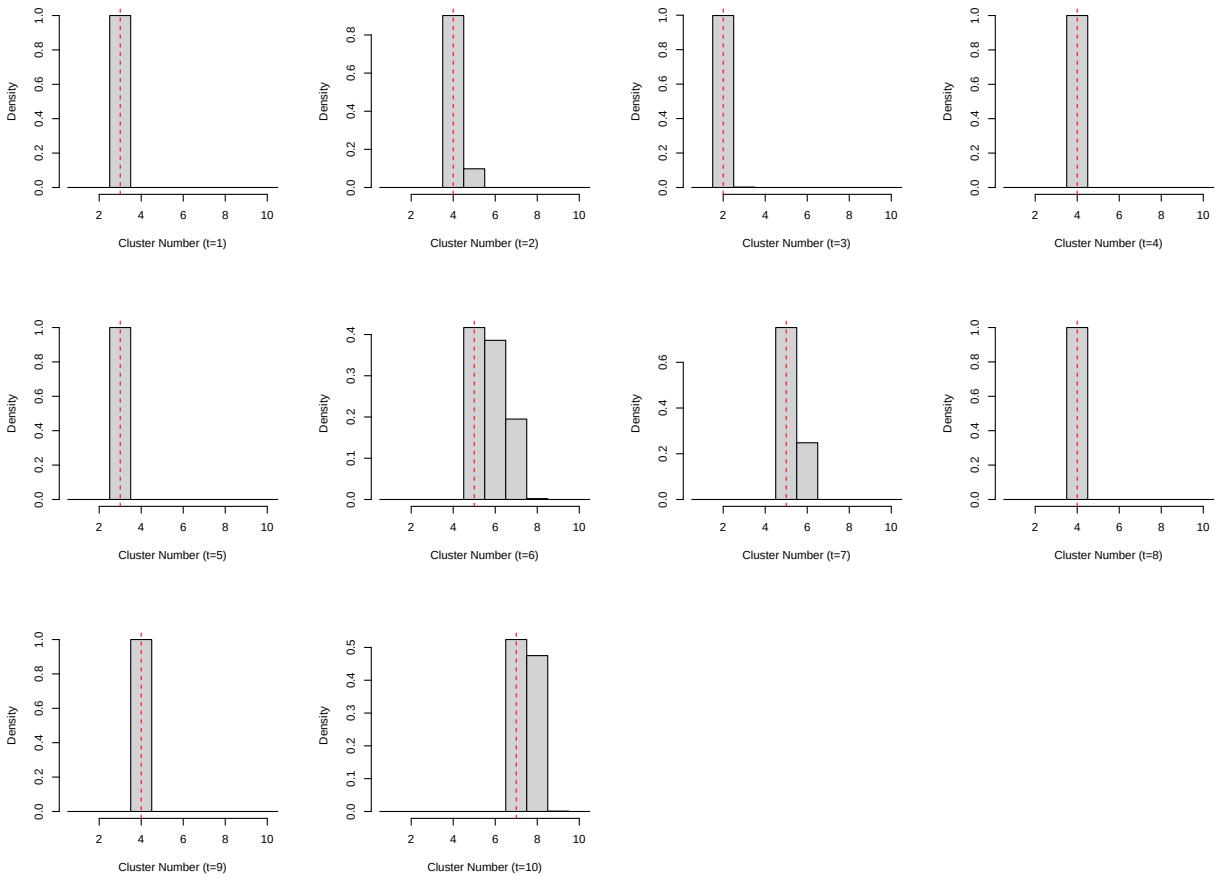

Figure 3.2: *Cluster Number Histograms in the Volatile-Cluster-Number Example.* The figures show the histograms of the numbers of clusters across iterations in each period. The dashed lines indicate the true cluster numbers.

Period	1	2	3	4	5	6	7	8	9	10
Sensitivity	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Specificity	1.00	1.00	0.96	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Table 3.3: *Classification Performance of the DDPM Model in the Stable-Cluster-Number Example.* In the stable-cluster-number example, the MAP cluster estimators correctly predict all the classification labels, except for a few incorrect predictions in period 3.

Period	1	2	3	4	5	6	7	8	9	10
Sensitivity	1.00	1.00	1.00	1.00	1.00	0.98	1.00	1.00	1.00	0.98
Specificity	1.00	1.00	1.00	1.00	1.00	0.92	1.00	1.00	1.00	0.87

Table 3.4: *Classification Performance of the DDPM Model in the Volatile-Cluster-Number Example.* In the volatile-cluster-number example, the MAP cluster estimators correctly predict all the classification labels, except for a few incorrect predictions in period 6 and period 10.

To evaluate the performance of the DDPM model on identifying unit-wise cluster affiliations, we frame the clustering task as a classification task. For period t, we have $\frac{I_t(I_t-1)}{2}$ unique voter pairs. If the (i_1, i_2) voter pair are in the same cluster, we code the label as 1, and 0 otherwise. We post-process the posterior samples from the above two examples with the MAP estimation method and identify the MAP cluster estimators for each period. Based on the MAP cluster estimators, we predict the voter pair label as 1 if they are in the same cluster, and 0 otherwise.

We report the sensitivity and specificity of the DDPM model's classification performance in the stable-cluster-number example in Table 3.3, and those in the volatile-cluster-number example in Table 3.4. Sensitivity indicates the success rate of correctly predicting the voter pairs that are truly in the same cluster. Specificity indicates the success rate of correctly predicting voter pairs that are truly not in the same cluster. In the stable-cluster-number example, the MAP cluster estimators correctly predict all the classification labels, except for a few incorrect predictions in period 3. In the volatile-cluster-number example, the MAP cluster estimators correctly predict all the classification labels, except for a few incorrect predictions in period 6 and period 10. Overall, the DDPM model has a decent performance at predicting whether voter pairs are in the same cluster in both examples.

To conserve space, we only visualize the pair-wise posterior probabilities of being in the same cluster for period 10 in Figure 3.3 and Figure 3.4, for the stable-cluster-number example and the volatile-cluster-number example, respectively. The two heat map plots show that the DDPM models identify the true clustering structures with great accuracy in both examples. In the stable-cluster-number example, all the units are estimated to be in their true clusters accurately. In the volatile-cluster-number example, only unit 41, unit 43, and unit 48 seem to switch between two clusters across iterations, and all the other units are estimated to be in their true clusters accurately.

3.6.2 Prior Distribution of the Cluster Number Simulation

In the DDPM setup, it's important to understand how the hyperprior parameter specifications affect the prior distribution of the cluster number in a period. We simulate cluster numbers in a Monte Carlo experiment for an imaginary time period $t = 0$, based on the following hyperprior specifications: $m_0 = 0, H_0 = 1, V = 0.05$. We simulate the cluster numbers under three scenarios, and the only difference across the scenarios is the number of voters. We specify the following voter numbers across the three scenarios: $I_0 \in \{40, 50, 60\}$.

We use the following steps to generate the cluster numbers under each scenario. Under each scenario, we simulate 10000 cluster numbers. We generate a value of γ_0 by $\gamma_0 \sim N(m_0, H_0)$. After obtaining a γ_0 value, we generate the corresponding z_0 through the observation equation in the DLM: $z_0 = \gamma_0 + e_2, e_2 \sim N(0, V)$. The corresponding α_0 directly follows by the transformation, $\alpha_0 = \exp(z_0)$. Given the number of voters, I_0, and α_0, we generate the cluster number by the function, $L_0 = \text{round}(E(L_0|\alpha_0, I_0))$, where $E(L_0|\alpha_0, I_0) \approx \alpha_0 \ln(1 + \frac{I_0}{\alpha_0})$ (Escobar and West, 1995). We plot the empirical densities of the simulated cluster numbers under each scenario in Figure 3.5. The modes of the three empirical densities fall within the interval, $[3, 4]$. The three empirical densities are all right-skewed.

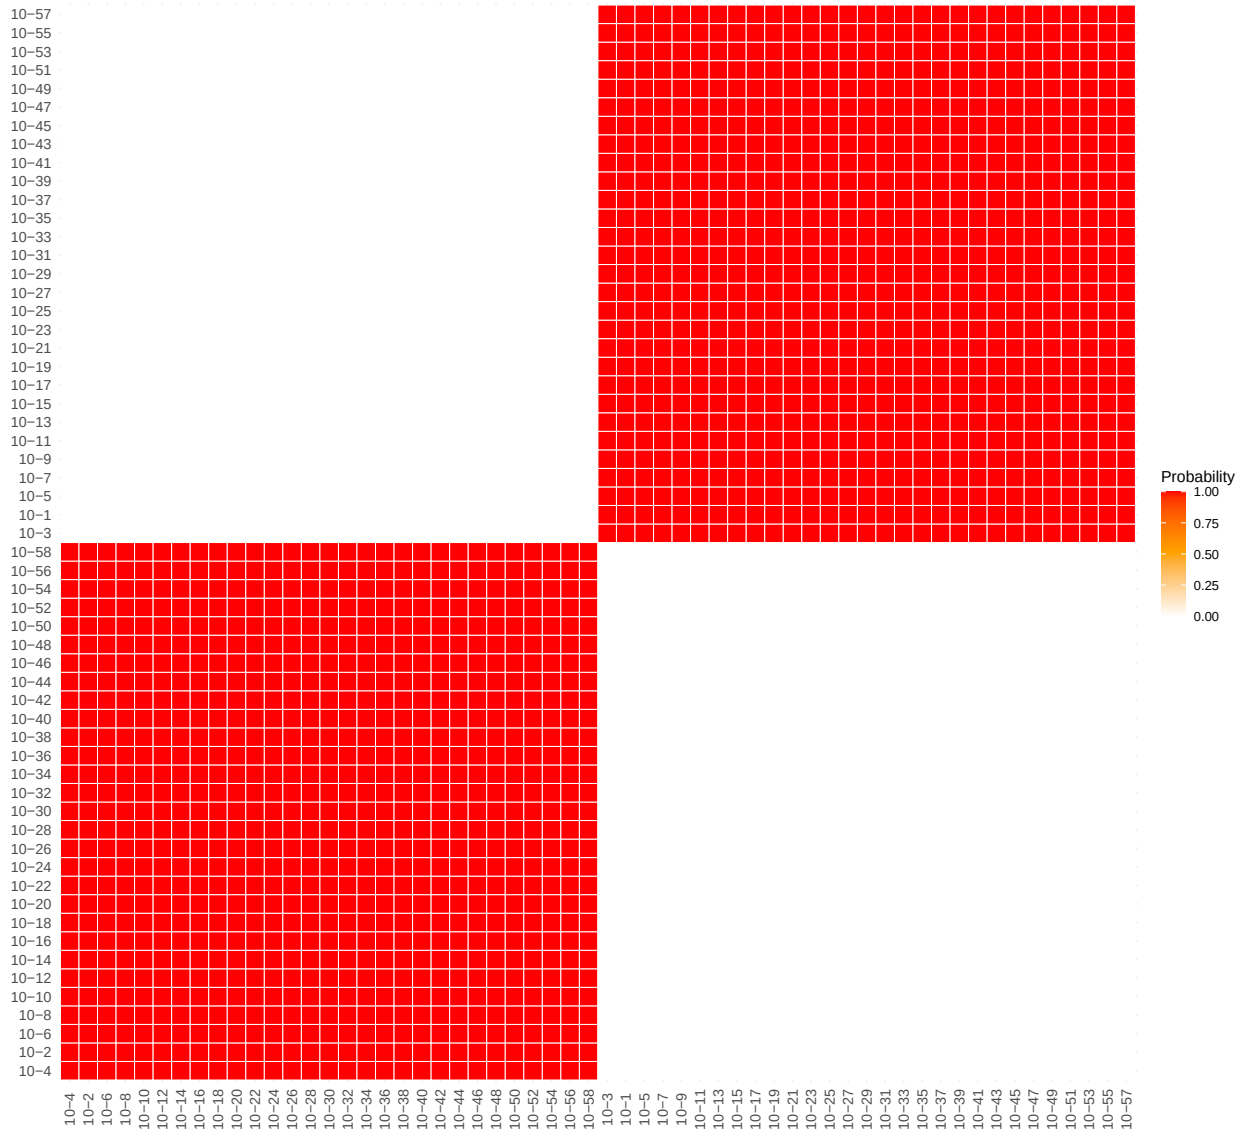

Figure 3.3: *Heat Map Plot for Period* 10 *in the Stable-Cluster-Number Example.* Each element in the heat map indicates the posterior probability for the unique voter pair to be in the same cluster in period 10.

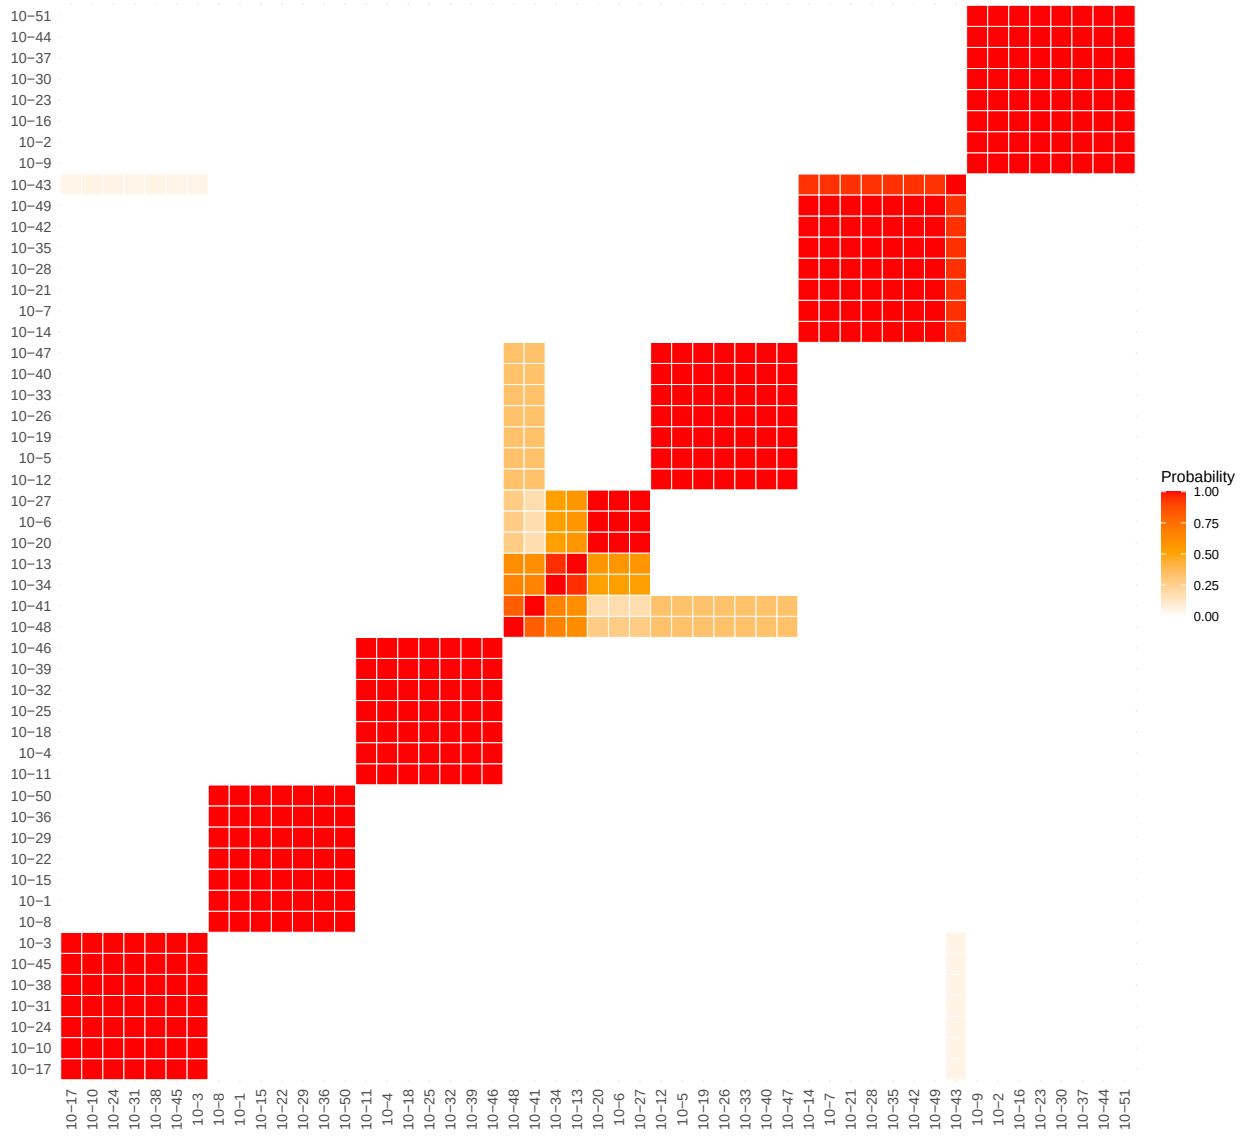

Figure 3.4: *Heat Map Plot for Period* 10 *in the Volatile-Cluster-Number Example.* Each element in the heat map indicates the posterior probability for the unique voter pair to be in the same cluster in period 10.

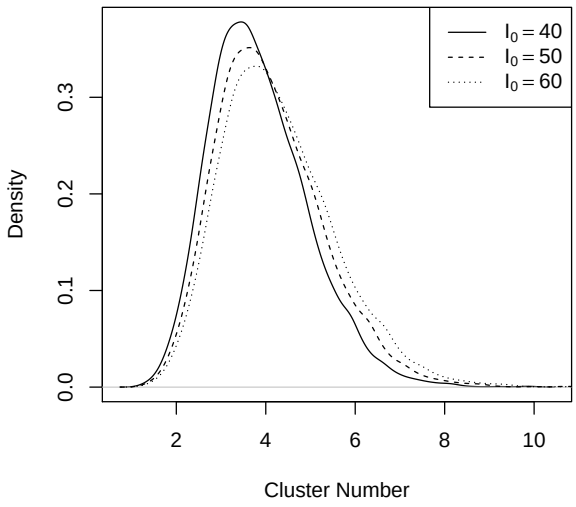

Figure 3.5: *Empirical Densities of the Simulated Cluster Numbers under Each Scenario.* The modes of the three empirical densities fall within the interval, $[3, 4]$. The three empirical densities are all right-skewed.

3.6.3 Robustness Check

Researchers have found that some Dirichlet Process Mixture model-based (DPM model-based) clustering methods have a hard time in consistently identifying the true cluster number, and they tend to estimate more clusters than the ground truth (Miller and Harrison, 2013). To solve this challenge, other researchers have developed alternative sparse finite mixture models to induce sparsity in cluster number estimation, leading to more accurate identification of cluster numbers (Malsiner-Walli, Frühwirth-Schnatter and Grün, 2016, 2017). More recently, researchers have found that a DPM model and a sparse finite mixture model yield similar estimation results with respect to the number of clusters, as long as the hyperparameter specifications in the two models induce a similar degree of sparsity (Frühwirth-Schnatter and Malsiner-Walli, 2019).

In response to the above general concern about DPM model-based clustering methods, we conduct various robustness checks. We run robuestness checks for a stable-cluster-number example and a volatile-cluster-number example. The robustness check results show that the DDPM model

is able to accurately identify the true cluster numbers under various model specifications.

In the subsection, "Two Demonstrative Examples", we have explained the details about simulating the data sets for the stable-cluster-number example and the volatile-cluster-number example. For robustness checks, we follow the exact same data simulation specifications and steps for both examples. Therefore, Table 3.1 shows the true data structure for the stable-cluster-number example, and Table 3.2 shows the true data structure for the volatile-cluster-number example. We simulate 100 roll call vote data sets for each example.

We fit three DDPM models on the 100 simulated data sets, for both examples respectively. In all of the three models, we run the SAMS sampler every third of the scan. In all of the three models, we run 2000 iterations and discard the first 1000 iterations as burnin. We specify different hyperparameters in the DLM parts of the three models. From Model One to Model Three, the hyperparameter specifications reflect increased degrees of restriction on sampling α_t's, and should therefore induce more sparsity in estimating cluster numbers.

In Model One, the tuning parameter for the random walk Metropolis-Hastings steps for drawing α_t's, τ, is 4. We set the parameters in the DLM priors as follow: $m_0 = 0, H_0 = 0.4, r_0 = 10, s_0 = 2, r_1 = 1, s_1 = 10$. The tuning parameter for the random walk Metropolis-Hastings step for drawing W, τ_W, is 0.1. We use the following values as starting values for V, W and all α_t's: $V_{start} = 0.2$, $W_{start} = 0.1$, and $\alpha_{t,start} = 1$ for $t = 1, 2, \cdots, T$. In Model Two, the tuning parameter for the random walk Metropolis-Hastings steps for drawing α_t's, τ, is 2. We set the parameters in the DLM priors as follow: $m_0 = 0, H_0 = 0.2, r_0 = 10, s_0 = 1, r_1 = 1, s_1 = 20$. The tuning parameter for the random walk Metropolis-Hastings step for drawing W, τ_W, is 0.05. We use the following values as starting values for V, W and all α_t's: $V_{start} = 0.1$, $W_{start} = 0.05$, and $\alpha_{t,start} = 1$ for $t = 1, 2, \cdots, T$. In Model Three, the tuning parameter for the random walk Metropolis-Hastings steps for drawing α_t's, τ, is 1. We set the parameters in the DLM priors as follow: $m_0 = 0, H_0 = 0.1, r_0 = 10, s_0 = 0.5, r_1 = 1, s_1 = 50$. The tuning parameter for the random walk Metropolis-Hastings step for drawing W, τ_W, is 0.02. We use the following values as starting values for V, W and all α_t's: $V_{start} = 0.05$, $W_{start} = 0.02$, and $\alpha_{t,start} = 1$ for

Period	1	2	3	4	5	6	7	8	9	10
Model One Average	4.00	4.00	4.00	2.00	2.31	3.00	2.31	3.00	2.00	2.00
Model Two Average	4.00	4.31	4.00	2.00	2.31	3.01	2.31	3.00	2.32	2.00
Model Three Average	4.00	4.00	4.00	2.00	2.31	3.00	2.31	3.00	2.00	2.00
True Cluster Number	4	4	4	2	2	3	2	3	2	2

Table 3.5: *Robustness Check Results for the Stable-Cluster-Number Example.* For the stable-cluster-number example, all three models are able to consistently identify the true cluster numbers on average across time periods, except for the slightly inflated cluster number estimation in period 5 and period 7.

Period	1	2	3	4	5	6	7	8	9	10
Model One Average	3.01	4.12	2.02	4.00	3.00	5.64	5.07	4.12	4.00	7.52
Model Two Average	3.00	4.10	2.05	4.03	3.00	5.72	5.05	4.08	4.06	7.49
Model Three Average	3.00	4.18	2.02	4.03	3.00	5.60	5.05	4.10	4.06	7.41
True Cluster Number	3	4	2	4	3	5	5	4	4	7

Table 3.6: *Robustness Check Results for the Volatile-Cluster-Number Example.* For the volatile-cluster-number example, all three models seem to have a decent performance at consistently identifying the correct cluster numbers on average across time periods, except for the slightly inflated cluster number estimation in period 6 and period 10.

$t = 1, 2, \cdots, T$.

Table 3.5 and Table 3.6 show the robustness check results for the stable-cluster-number example and the volatile-cluster-number example, respectively. For each Monte Carlo experiment, we use the mode of the posterior cluster number sample as the estimated cluster number. The average values reported in the above two tables are the simple averages of the estimated cluster numbers across 100 Monte Carlo experiments. We also report the true cluster numbers at the bottom of the tables. For the stable-cluster-number example, the estimated cluster number averages are very similar across the three model specifications. All three models are able to consistently identify the true cluster numbers on average across time periods, except for the slightly inflated cluster number estimation in period 5 and period 7. For the volatile-cluster-number example, all three models seem have a decent performance at consistently identifying the correct cluster numbers on average across time periods, except for the slightly inflated cluster number estimation in period

Table 3.7: *Summary Statistics of the UNGA Human Rights Vote Data.* The table reports the summary statistics for the numbers of countries and the numbers of resolutions in the UNGA human rights roll call vote data from 1992 to 2017.

Summaries	Number of Countries	Number of Resolutions
Min.	99.0	11.00
1st Qu.	119.2	17.25
Median	134.0	21.00
Mean	131.3	20.50
3rd Qu.	141.0	24.00
Max.	156.0	29.00

6 and period 10. However, there are slight differences in terms of the estimated cluster number averages across the three models. The most restrictive model, Model Three, seem to perform the best overall at consistently identifying the correct cluster numbers on average across time periods.

3.7 Human Rights Votes in the UNGA

We apply the proposed DDPM model and post-processing methods to the United Nations General Assembly (UNGA) human rights roll call vote data from 1992 to 2017. We rely on the publicly available "United Nations General Assembly Voting Data" for this study (Voeten, Strezhnev and Bailey, 2009). The original UNGA roll call vote data record "Yes", "No", "Abstain", or "Missing" for a voting record. Voeten argues that "UNGA resolutions are not binding, what really matters is whether or not a state is willing to go on the record for supporting a resolution." "No" votes and "Abstain" votes "both are essentially ways a state can express its unwillingness to comply with the text of a resolution" (Voeten, 2000). Therefore, in this study, we record "Yes" votes as 1, "Missing" as NA, and both "No" and "Abstain" votes as 0. Moreover, in each year, we only include countries that vote on at least 95% of the human rights resolutions in the year. Table 3.7 reports the summary statistics of the numbers of countries and resolutions throughout the years.

We use the following specifications for running the DDPM model on human rights vote data. We run the SAMS sampler every third iteration of the MCMC sampling scheme. We supply the

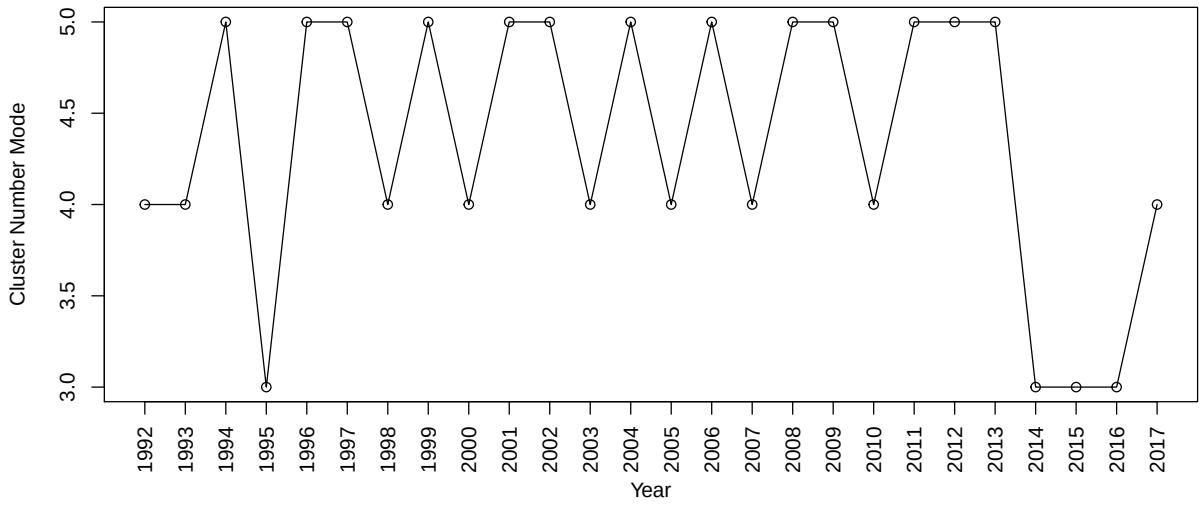

Figure 3.6: *Annual Cluster Number Mode.* For each year, we compute the mode of the numbers of unique clusters across iterations in the posterior sample. The annual cluster number modes represent our estimates of the annual voting coalition numbers, based on the DDPM model.

following fixed values. All elements in $\boldsymbol{\lambda}_t$'s are fixed at 1. The tuning parameter for the random walk Metropolis-Hastings steps for drawing α_t's, τ, is 0.5. We set the parameters in the DLM priors as follow: $m_0 = 0, H_0 = 0.1, r_0 = 10, s_0 = 0.5, r_1 = 1, s_1 = 50$. The tuning parameter for the random walk Metropolis-Hastings step for drawing W, τ_W, is 0.01. We use the following values as starting values for V, W and all α_t's: $V_{start} = 0.05$, $W_{start} = 0.02$, and $\alpha_{t,start} = 1$ for $t = 1, 2, \cdots, T$. We run 10000 iterations and discard the first 5000 iterations as burnin. The acceptance rates of the random walk Metropolis-Hastings steps for drawing α_t's lie in the interval, $[0.3871, 0.4661]$. The acceptance rate of the random walk Metropolis-Hastings step for drawing W is 0.344. These acceptance rates are all acceptable for a unidimensional Metropolis-Hastings sampler.

We post-process the posterior samples of cluster labels in each year, and report the year-wise cluster number modes in Figure 3.6. For all years, the numbers of coalitions vary between 3 and 5. There is no obvious time trend of the cluster numbers across years. The relative stability of

coalition numbers show preliminary evidence for "peer groups" in the UNGA human rights vote data. It also shows that there is no universal consensus on human rights issues in the UNGA. The persistent existence of coalitions lends support to the criticism of human rights politicization in the UNGA.

For each year, we estimate the MAP cluster labels for countries. In Appendix B.1, we report the voting coalition membership for each year from 1992 to 2017. To converse space, we only show the coalition membership in the most recent year, 2017, in Table 3.8. In 2017, the DDPM model identifies four coalitions. The most obvious pattern is the separation between developed and developing countries. There are two coalitions consisting of developed countries. Most western European countries and the developed countries in Asia-Pacific, such as Japan, South Korea, and New Zealand, form a coalition. The US, Israel, and two other close US allies form a coalition. The developing countries are also divided into two coalitions: a large coalition including most developing countries, and a small coalition joined mostly by a few Latin American countries.

Figure 3.7 shows the posterior probability for unique country pairs to be in the same coalition in 2017. The plot confirms the four coalitions reported in the Table 3.8. The heat map plot shows that Australia is likely to belong to both the smaller US-led coalition and the bigger EU-led coalition. Similarly, Russia and Argentina are likely to belong to the bigger developing country coalition and the smaller developing country coalition. However, there is no country that is likely to belong to both a developing country coalition and a developed country coalition. This shows that the differences between developing countries and developed countries are much starker than those among developing countries themselves or developed countries themselves.

After examining the coalition memberships from 1992 and 2017, we are able to find some consistent patterns. These findings help us reevaluate some of the arguments in the literature of human rights voting in the UNGA. First, across all the years, there are stable coalition structures in the voting data. This finding lends support to the long-lasting discussion of the politicization of human rights issues in the UNGA. Ideally, human rights debates in the UNGA should be based on principles and international laws, and consensus building on human rights issues should contribute

Table 3.8: *Voting Coalition Membership in 2017.* Four voting coalitions are identified in 2017. Coalition 3 is the EU-led big developed country coalition. Coalition 2 is the US-led small developed country coalition. Coalition 4 is the big developing country coalition including most countries in Asia and Africa. Coalition 1 is the small developing country coalition, mostly consisting of Latin American countries. Coalition $1, 2, 3, 4$ corresponds to the four anti-diagonal blocks from the bottom-left corner to the top-right corner in Figure 3.7.

Coalition	Member
1	Cameroon, Colombia, Guatemala, Honduras, Mexico, Panama, Paraguay, Peru, Solomon Islands, Togo
2	Australia, Canada, Israel, United States of America
3	Albania, Andorra, Austria, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Greece, Hungary, Iceland, Ireland, Italy, Japan, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, New Zealand, Norway, Poland, Portugal, Republic of Korea, Republic of Moldova, Romania, San Marino, Slovakia, Slovenia, Spain, Sweden, Switzerland, The former Yugoslav Republic of Macedonia, Turkey, Ukraine, United Kingdom of Great Britain and Northern Ireland
4	Algeria, Angola, Argentina, Bahamas, Bahrain, Bangladesh, Bhutan, Bolivia (Plurinational State of), Botswana, Brazil, Brunei Darussalam, Cabo Verde, Cambodia, Chile, China, Congo, Costa Rica, Côte D'Ivoire, Cuba, Djibouti, Dominican Republic, Ecuador, Egypt, El Salvador, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Guinea, Guyana, India, Indonesia, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kuwait, Kyrgyzstan, Lao People's Democratic Republic, Lesotho, Libya, Malaysia, Maldives, Mali, Mauritania, Mauritius, Morocco, Mozambique, Namibia, Nepal, Nicaragua, Oman, Pakistan, Philippines, Qatar, Russian Federation, Saint Kitts and Nevis, Saint Lucia, Saudi Arabia, Sierra Leone, Singapore, South Africa, Sri Lanka, Sudan, Syrian Arab Republic, Tajikistan, Thailand, Trinidad and Tobago, Tunisia, Uganda, United Arab Emirates, United Republic of Tanzania, Uruguay, Uzbekistan, Venezuela, Bolivarian Republic of, Viet Nam, Yemen, Zimbabwe

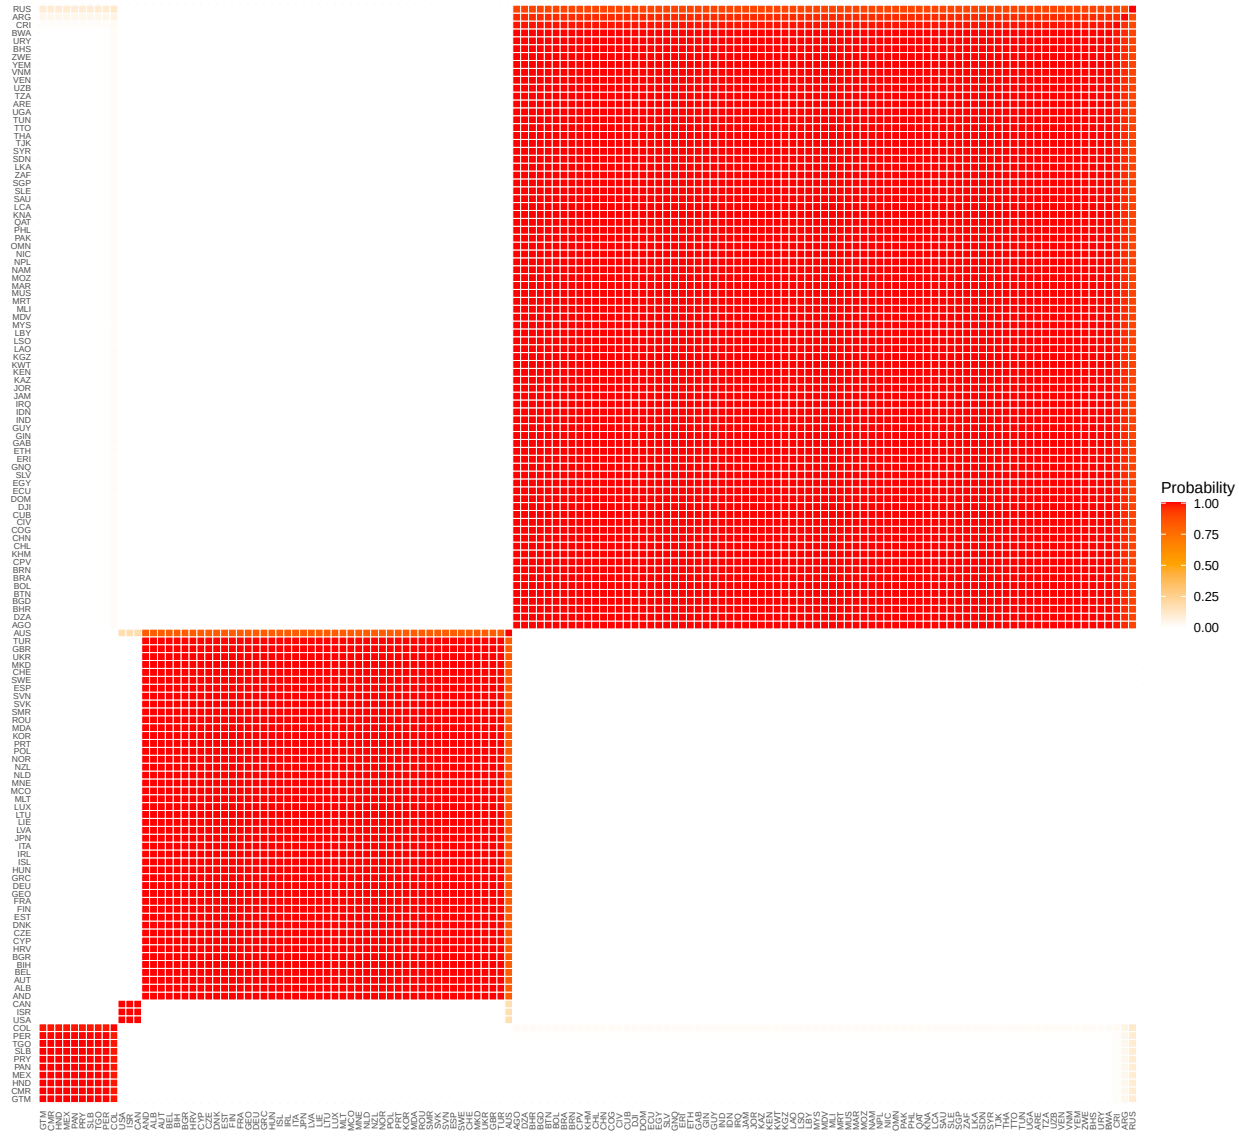

Figure 3.7: *Posterior Probability for Country Pairs Being in the Same Coalition in 2017.* Four voting coalitions emerge from the heat map plot. The four anti-diagonal blocks from the bottom-left corner to the top-right corner correspond to Coalition $1, 2, 3, 4$ in Table 3.8. The vast majority of countries belong to one coalition with very high probability. Australia is likely to belong to both Coalition 2 and Coalition 3. Russia and Argentina are likely to affiliate to both Coalition 1 and Coalition 4.

to global governance. However, the voting records show that different country coalitions take contradictory positions, and countries are divided in the political debates of human right issues.

Second, throughout the years, the EU member countries always coalesce in the same group. This finding lends support to the previous studies that treat the EU as a voting bloc (Luif, 2003). In addition to the EU member countries, a few developed countries in Asia-Pacific, including Japan, South Korea, and New Zealand, are also consistently a part of the EU-led group. Therefore, we should update the previous understanding of treating EU as a bloc, but redefine the EU-led group to include more developed countries. In some years, the US, Israel, Canada, and Australia are part of the EU-led group, but they form their own smaller US-led group in other years. This finding shows that US and Israel agree with the EU-led coalition in many aspects of human rights issues, but there is also some repeated disagreement between the US-led coalition and the EU-led coalition. Among the developing countries, there is always a big developing country coalition, including most developing countries from Asia and Africa. In addition, there also exists a smaller developing country coalition, including Latin American countries and former Soviet Union member countries. The coalition pattern among the developing countries contradicts previous studies that treat the Organization of Islamic Cooperation member countries as a bloc. As shown by the clustering results, most of the Muslim countries are a part of the big developing country coalition. Moreover, the coalition patterns among developing countries also corroborate some of the previous arguments that emphasize regional voting blocs. As shown by the smaller developing country coalition, we find that Latin American countries break with the big developing country coalition to form their own smaller coalition in some years. There is also evidence showing that Russia and other former Soviet Union member countries are apart from the big developing country coalition and end up in a smaller coalition in some years.

We examine what specific human rights issues divide countries into different coalitions. Based on the MAP cluster estimators and each country's voting records in year t, we are able to compute the probability vector, $\tilde{\boldsymbol{\theta}}_{lt}$, for each cluster l. There are J_t resolutions in year t. For each resolution $j \in \{1, 2, \cdots, J_t\}$, we compute the maximum probability difference for any unique cluster pairs,

$d_{ll'jt} = \max\left\{|\tilde{\theta}_{jlt} - \tilde{\theta}_{jl't}|\right\}$, for all $l, l' \in \{c_{it,\text{MAP}}\}_{i=1}^{I_t}$. For some resolutions, countries have almost unanimous votes, so $d_{ll'jt}$ is close to zero. However, other resolutions become the center of controversy, and cluster-wise probability parameters for this resolution can be very different. The numbers of resolutions vary in each year, so we identify the top 50% most polarizing resolutions in terms of $d_{ll'jt}$ values in each year. To conserve space, we only report the top 50% most divisive resolutions in 1992 in Table 3.9 . We report the top 50% most polarizing resolutions in all years in Appendix B.2.

We have a few important findings from showing these most polarizing resolutions. First, the Israeli-Palestinian issue is a persistent line of conflict in all years. This finding confirms previous studies that theorize the obsession of voting on the Israeli-Palestinian issue in the UNGA (Becker et al., 2015). Second, country-specific human rights reports tend to be polarizing, such as reports on Cuba and Iran. The reports targeting specific countries clearly show the political nature of human rights issues in the UNGA. Countries are divided in different coalitions by how they view the human rights record of a country, because they prioritize bilateral relations with the targeted country over objectively evaluating the facts. Third, the UNGA tends to vote on some of the exact same polarizing human rights resolutions for many years. For example, the resolution condemning "the use of mercenaries to violate human rights" has been a polarizing resolution in almost every year from 1992 until the late 2000s. This resolution-wise repetition also shows evidence for the persistent politicization of certain human rights issues. For the case of the "use of mercenaries" resolution, developing countries almost all vote for it, while developed countries almost all oppose it. This resolution has become a weapon for developing countries to embarrass developed countries on their human rights practice.

Last, we compare the DDPM model with a Bayesian Dynamic IRT model for analyzing the UNGA human rights vote data. A Bayesian IRT model is one of the most common choices for modeling roll call vote data. Bayesian IRT models assume that voters make voting decisions based on the relative utilities between voting yea and nay. In a latent policy space, if a voter's ideal point is closer to the new proposal's position, she is more likely to vote yea; if a voter's ideal point is

Table 3.9: *Polarizing Resolutions in 1992.* The six resolutions are the top 50% most polarizing resolutions in 1992. Resolution $1, 4, 5, 6$ are country-specific human rights resolutions about Iran or the Palestinian-Israeli issue. Resolution $2, 3$ are about general human rights principles, such as self-determination of nations and the use of mercenaries.

Count	Resolution Title
1	Programme of work, Committee on Palestinian rights
2	Importance of the universal realization of the rights of peoples to self determination and of the speedy granting of independence to colonial countries and peoples for the effective guarantee and observance of human rights
3	Use of mercernaries as a means to violate human rights and to impede the exercise of the right of peoples to self determination United Nations African Institute for the Prevention of Crime and the Treatment of Offenders
4	Situation of human rights in the Islamic Republic of Iran
5	Encourages all member states to lend assistance to the committee on Palestinian rights
6	Reaffirms the inalienable right of all displaced inhabitants to return to their homes or former places of residence in the territories occupied by Israel since 1967

closer to the status quo, she is inclined to vote nay. In each year, the Dynamic IRT model estimates one continuous ideal point variable for each country, and two continuous parameters for each resolution. Here, we use the same model specified in the study of the US Supreme Court justices' ideal points (Martin and Quinn, 2002). We implement the model with `MCMCdynamicIRT1d()` function in `MCMCpack` R package (Martin, Quinn and Park, 2011). We use the default priors and tuning parameters of the function.

To compare the performances of the DDPM model and the Dynamic IRT model, we first randomly draw one voting record per country-year to form a hold-out testing data set. The remaining data points become the training data set. We train the two models with the training data set, and predict the hold-out voting records with the estimated latent variables, respectively. We use sensitivity, specificity, and $F1$ score to evaluate prediction performance.[8] Table 3.10 shows the

[8]In the roll call vote context, sensitivity represents the share of the yea votes that are correctly predicted as a yea vote, and specificity represents the share of the nay votes that are correctly predicted as a nay vote. $F1$ score is a balanced prediction performance metric, considering both sensitivity and precision. $F1$ score is a function of sensitivity and precision: $F1 = \frac{2 \times \text{precision} \times \text{sensitivity}}{\text{precision} + \text{sensitivity}}$.

Table 3.10: *DDPM Model and Bayesian Dynamic IRT Model Comparison.* We train a DDPM model and a Bayesian Dynamic IRT model on the same training data set. We use the estimated latent variables to predict the unique country-pair relations (being in the same coalition or not) in the hold-out testing data set. The table shows that the DDPM model has a slight edge over the Bayesian Dynamic IRT model in terms of F_1 score.

Model	Sensitivity	Specificity	$F1$ Score
DDPM Model	0.9511	0.8501	0.8668
Dynamic IRT Model	0.9540	0.8127	0.8477

prediction performances of the two models. The Dynamic IRT model is slightly better than the DDPM model at predicting yea votes, while the DDPM model is much better than the Dynamic IRT model at predicting nay votes. In sum, the DDPM model has a larger $F1$ score than the Dynamic IRT model. The comparison of the DDPM model and the Dynamic IRT model shows that the DDPM is at least as a good of a modeling choice as the Dynamic IRT model. This provides further support for modeling voting coalitions instead of individual country's ideal points for the UNGA human rights vote data.

3.8 Conclusion

We have proposed the DDPM model to identify voting coalitions with roll call vote data across multiple periods, along with post-processing methods to analyze the posterior samples from the DDPM model. The proposed post-processing methods provide streamlined steps for applied researchers to make sense of the DDPM outputs and more easily interpret the results. We make the functions implementing the DDPM model and the post-processing methods available to the public in the `MCMCddpmbb` R package.

We have applied these methods to the UNGA human rights roll call vote data from 1992 to 2017. We have identified human rights voting coalitions in the UNGA after the Cold War, and the polarizing resolutions that divide countries into different coalitions. Through this study, we find a clear separation between developing and developed countries in human rights voting. Moreover,

we find the EU as a stable coalition among the developed countries, whereas some Latin American countries tend to form a small coalition, apart from a large developing country coalition. We also find that many polarizing resolutions repeatedly show up across years, and the lines of conflict lie in both debates on general human rights principles and human rights violation reports on specific countries.

Future research plans are in order. One future plan for improving the DDPM model is to investigate how to dynamically model voters' cluster affiliations across time directly. Another place for future improvement is to generalize the binary voting record assumption to accommodate ordinal voting records, such as "nay", "abstain", and "yea". Last, it's worth investigating how a sparse finite mixture model performs at producing interpretable results in the roll call vote context. A sparse finite mixture model option can be a potentially useful addition to the `MCMCddpmbb` R package.

CHAPTER 4

Penalized EM Algorithm for Multidimensional IRT Model

4.1 Introduction

Ideal point estimation models have extensive applications in political science. With ideal point estimation models, researchers measure national-level and state-level legislators' (Clinton, Jackman and Rivers, 2004; Shor, Berry and McCarty, 2010), the Supreme Court justices' (Martin and Quinn, 2002), or voters' (Jessee, 2009) ideal points, based on an actor's binary responses to bills, court cases, or survey questions. Researchers also use similar models to estimate countries' preferences according to the United Nations General Assembly roll call votes (Voeten, 2000; Bailey, Strezhnev and Voeten, 2017). There are various methods in the toolkit for ideal point estimation, such as NOMINATE Models (Poole and Rosenthal, 1985; Poole, 2001; Carroll et al., 2009), Bayesian IRT models (Martin and Quinn, 2002; Clinton, Jackman and Rivers, 2004; Imai, Lo and Olmsted, 2016), and non-parametric ideal point models (Poole, 2000; Tahk, 2018). Specifically, this chapter focuses on developing a penalized EM algorithm for IRT models.

The spatial voting model is the theoretical foundation for IRT models. The spatial voting model assumes that actors make voting decisions based on the relative utilities between voting yea and nay. In the policy space, if an actor's ideal point is closer to a new proposal's position, she is more likely to vote "Yea"; if an actor's ideal point is closer to the status quo, she is inclined to vote "Nay". The spatial voting model assumes a uni- or multi-dimensional policy space, and accordingly IRT

models are able to estimate uni- or multidimensional ideal points.

Unidimensional IRT models are commonly used. A unidimensional ideal point model is easy to estimate and interpret. Unidimensional IRT models simply aim at accounting for the maximum variation in the data with unidimensional ideal point estimates. Researchers do not have to investigate the dimension meaning since all items load onto the single dimension. Instead, researchers give a heuristic interpretation of the single dimension based on their understandings of the subject. Researchers show that unidimensional ideal point scores capture most of the variation and predict the observed data well for some binary choice data sets (Poole and Rosenthal, 1991). Therefore, we have a strong justification for estimating only unidimensional ideal points in these situations.

However, for other binary choice data, multidimensional models significantly improve the predictive power of ideal point estimates on the observed choices. The multidimensional ideal points are able to pick out different patterns of variation in the data that the unidimensional models are not able to. Furthermore, there are indeed concerns for the lack of fit or interpretability of unidimensional IRT models when estimating survey data in the literature (Broockman, 2016; Ahler and Broockman, 2018). New methods have also been proposed to model the number of dimensions itself, together with estimating ideal points, based on roll call vote data (McAlister, 2021). This paper does not directly speak to how to decide between uni- or multidimensional IRT models. Rather, we focus on how to improve upon the existing methods for estimating multidimensional IRT models.

In contrast to unidimensional IRT models, a multidimensional IRT model requires more rigorous investigation of the item-dimension loading structure and dimension meanings. The previous attempts at learning dimension meanings mostly rely on researchers' substantive understandings of the area or data. Poole and Rosenthal (1991) contend that for US legislators' two-dimensional ideal points, "one dimension ranges from strong loyalty to one party (Democrat-Republican or Democrat) to weak loyalty to either party to strong loyalty to a second, competing, party (Federalist, Whig, or Republican). Another dimension differentiates 'liberals' from 'conservatives' within the two competing parties." This observation is based on heuristic interpretations of the estimated

ideal points, and no specific "anchoring items (bills)" are identified to give concrete evidence.

Jackman (2001) takes a step further for learning dimension meanings, and he stresses the importance of learning the item-dimension loading structure in multidimensional IRT model estimation (Jackman, 2001). This is equivalent to estimating the sparse mapping relationship between items and dimensions, akin to estimating the factor loading structure in traditional factor analysis. Jackman (2001) estimates a multi-dimensional IRT model assuming all items load on all dimensions at the beginning, and examines whether a discrimination parameter is significantly different from zero or not ex post.[1] He identifies the bills that have a significant discrimination parameter on only one dimension as the "anchoring items" for this dimension. Then he infers the dimension meanings based on the contents of the anchoring bills. His approach gives inconsistent estimates if the loading structure is truly sparse in the data generation process (many discrimination parameters should be exactly zero). Moreover, he relies on an arbitrary threshold (significance level) to determine whether a discrimination parameter is zero.

Ideally we should integrate the estimation of the sparse item-dimension loading structure into the entire IRT model estimation process. This way, we are able to consistently estimate the item parameters and ideal points simultaneously. In turn, we will be able to identify items that only load onto a single dimension, and use these special items to identify the meanings of each dimension. Researchers have attempted to estimate the sparse item-dimension loading structure for a full Bayesian IRT model by employing spike-and-slab priors (Richard Hahn, Carvalho and Scott, 2012). This method is able to identify a stochastic sparse item-dimension loading structure but suffers from two potential drawbacks. First, the learnt loading structure is potentially different in every iteration of the posterior sample, and it's difficult to determine a fixed result of the item-dimension loading structure. Nor do we have a clear method to choose the optimal amount of regularization (the prior probability of the spike) in the Bayesian framework. Second, Bayesian simulations with spike-and-slab priors can be enormously costly in terms of computational time

[1] Jackman uses the 90% credible intervals of discrimination parameter posterior samples to decide whether the discrimination parameters are "significant" or not. This is an application of frequentist statistical tests to Bayesian posterior samples.

and memory.

An EM algorithm has the potential to avoid the above two drawbacks. To our best knowledge, there is one existing EM algorithm for estimating the sparse item-dimension loading structure for IRT models (Sun et al., 2016). Sun et al. (2016) develop this method for IRT models with a logit link and only positive discrimination parameters. This group of IRT models is commonly used in educational testing and psychology studies, where items are guaranteed to elicit subjects' latent attributes in a fixed direction.[2] However, in ideal point estimation, yea votes on some bills reveal an actor's conservative preference, whereas yea votes on other bills reflect an actor's liberal inclination. Therefore, we need to develop a method for estimating the IRT models suitable for ideal point estimation in the political science research.

In this paper, we propose a new penalized EM algorithm to estimate the sparse item-dimension loading structure in multidimensional IRT models. We hereby make two contributions. First, the proposed penalized EM algorithm is able to identify the sparse item-dimension loading structure for a multidimensional IRT model. The sparse item-dimension loading structure provides the evidence for identifying the anchoring items for each dimension and inferring the meaning of each dimension. Second, the penalized EM algorithm relaxes the assumption that all items load onto all dimensions. Therefore, the penalized EM algorithm is able to consistently estimate the multidimensional ideal points in the presence of a sparse item-dimension loading structure. Similar to all EM algorithms, the penalized EM algorithm for IRT model is sensitive to parameter starting values. Researchers should experiment with multiple starting values in exploration analysis, and then pick the most suitable starting values.

The rest of the chapter proceeds as follows. First, we revisit the spatial voting model and Bayesian IRT model for ideal point estimation. Second, we propose the penalized EM algorithm for multidimensional IRT models, and specify the estimation steps for the algorithm. Third, we use simulation data to illustrate that the penalized EM algorithm is able to successfully recover the true item-dimension loading structure and consistently estimate the true ideal points simultaneously.

[2]For example, all test items in SAT are designed to elicit test takers' aptitude, and a test taker's correct answer to a test item shows higher aptitude in expectation.

Last, we apply the proposed algorithm to the 105th US Senate roll call vote data to identify the anchoring items for a two-dimensional IRT model and infer dimension meanings.

4.2 Spatial Voting Model and Bayesian IRT Model

The data structure we study in this paper resembles a roll call vote matrix, where rows indicate voters and columns represent items. Let matrix $\boldsymbol{Y}$ $(N \times J)$ denote such a roll call vote matrix with N voters voting on J items. The vote outcome is binary in that each voter either votes "Yea" or "Nay" to a given item. An entry y_{ij} is recorded in the following way.

$$y_{ij} = \begin{cases} 0 & \text{if voter } i \text{ votes Yea on item } j \\ 1 & \text{if voter } i \text{ votes Nay on item } j \\ \text{NA} & \text{if voter } i\text{'s vote on item } j \text{ is missing} \end{cases}$$

The spatial voting model sets the theoretical foundation for inferring voters' ideal points based on the above roll call vote data (Enelow and Hinich, 1984; Jackman, 2009). The spatial voting model assumes that voters' ideal points, proposal positions and the status quo are projected onto a policy space. A voter' decision on a new proposal is defined as the function of her ideal point's relative distance to the proposal position and the status quo. Under the globally satiable and symmetric utility assumption, voters have a quadratic utility function over "Yea" or "Nay" choices. Let $\boldsymbol{\phi}_j$ denote the proposal position and $\boldsymbol{\psi_j}$ denote the status quo with respect to item j. The policy space has dimension K. Then the utility of voter i's "Yea" vote and "Nay" vote on item j are:

$$U_i(\boldsymbol{\phi}_j) = -||\boldsymbol{x}_i - \boldsymbol{\phi}_j||^2 + \eta_{ij}$$

$$U_i(\boldsymbol{\psi}_j) = -||\boldsymbol{x}_i - \boldsymbol{\psi}_j||^2 + \upsilon_{ij}$$

where $\boldsymbol{x}_i$ is the ideal point of voter i, and η_{ij} and υ_{ij} are both error terms. Voter i will vote "Yea" if $U_i(\boldsymbol{\phi}_j) > U_i(\boldsymbol{\psi}_j)$ and "Nay" otherwise. The errors are assumed to be independently and normally

distributed with mean, 0, and variance, σ_j^2. With this structural assumption on the errors, the probability of voter i voting "Yea" is modeled as

$$\begin{aligned} P(y_{ij}=1) &= P(U_i(\boldsymbol{\phi}_j) > U_i(\boldsymbol{\psi}_j)) \\ &= P(v_{ij} - \eta_{ij} < ||\boldsymbol{x}_i - \boldsymbol{\psi}_j||^2 - ||\boldsymbol{x}_i - \boldsymbol{\phi}_j||^2) \\ &= P(v_{ij} - \eta_{ij} < 2(\boldsymbol{\phi}_j - \boldsymbol{\psi}_j)^T \boldsymbol{x}_i + \boldsymbol{\psi}_j^T \boldsymbol{\psi}_j - \boldsymbol{\phi}_j^T \boldsymbol{\phi}_j) \\ &= \Phi(\beta_j^T \boldsymbol{x}_i + \alpha_j) \end{aligned}$$

where $\beta_j = \dfrac{\boldsymbol{\phi}_j - \boldsymbol{\psi}_j}{\sigma_j}$ (a vector of length K), $\alpha_j = \dfrac{\boldsymbol{\psi}_j^T \boldsymbol{\psi}_j - \boldsymbol{\phi}_j^T \boldsymbol{\phi}_j}{2\sigma_j}$ (a scalar), and $\Phi(\cdot)$ denotes the CDF of a standard normal distribution.

We further assume that all vote data points are independent, given the voters' ideal points and the corresponding item parameters. Therefore, we can write out the likelihood function for the entire roll call vote matrix as follows:

$$L(\boldsymbol{Y}|\{\beta_j\}_{j=1}^J, \{\alpha_j\}_{j=1}^J, \{\boldsymbol{x}_i\}_{i=1}^N) = \prod_{i=1}^N \prod_{j=1}^J \Big(\Phi(\boldsymbol{x}_i^T \beta_j + \alpha_j)\Big)^{\mathbb{I}(y_{ij}=1)} \Big(1 - \Phi(\boldsymbol{x}_i^T \beta_j + \alpha_j)\Big)^{\mathbb{I}(y_{ij}=0)}$$

where $\mathbb{I}(\cdot)$ is an indicator function.

Building on the above modeling strategy, researchers have proposed an efficient Bayesian Gibbs sampler for estimating the parameters (Clinton, Jackman and Rivers, 2004; Jackman, 2009). Researchers place independent and conjugate normal priors on $\boldsymbol{x}_i$, α_j and β_j:

$$\begin{aligned} \boldsymbol{x}_i &\overset{\text{i.i.d}}{\sim} N_K(\boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x) \\ [\alpha_j \quad \beta_j]^T = \boldsymbol{\beta}_j &\overset{\text{i.i.d}}{\sim} N_{K+1}(\boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta) \end{aligned}$$

where $\boldsymbol{\Sigma}_x$ and $\boldsymbol{\Sigma}_\beta$ are diagonal matices.

Henceforth, we can write out the joint posterior as a product of the prior and the likelihood as

follows:

$$
\begin{aligned}
&P(\{\beta_j\}_{j=1}^J, \{\alpha_j\}_{j=1}^J, \{\boldsymbol{x}_i\}_{i=1}^N | \boldsymbol{Y}) \\
&\propto \underbrace{\prod_{i=1}^N \prod_{j=1}^J \Big(\Phi(\boldsymbol{x}_i^T \beta_j + \alpha_j)\Big)^{\mathbb{I}(y_{ij}=1)} \Big(1 - \Phi(\boldsymbol{x}_i^T \beta_j + \alpha_j)\Big)^{\mathbb{I}(y_{ij}=0)}}_{\text{likehood}} \\
&\quad \underbrace{\times \prod_{i=1}^N \phi_K(\boldsymbol{x}_i; \boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x) \prod_{j=1}^J \phi_{K+1}(\boldsymbol{\beta}_j; \boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta)}_{\text{prior}}
\end{aligned}
$$

where $\phi_K(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the PDF of a K-dimensional normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

To implement the efficient Gibbs sampler, researchers have proposed a data-augmentation step for the above joint posterior (Albert and Chib, 1993; Jackman, 2000). As in Probit regressions, we introduce a latent variable, y_{ij}^*, to represent the utility for voter i to vote "Yea" on item j. Hence, we can model a latent response as $y_{ij}^* = \alpha_j + \beta_j^T \boldsymbol{x}_i + \epsilon_{ij}$, $\epsilon_{ij} \sim N(0, 1)$, under the constraints below:

$$
\begin{cases}
y_{ij}^* > 0 & \text{if } y_{ij} = 1 \\
y_{ij}^* < 0 & \text{if } y_{ij} = 0 \\
y_{ij}^* \in \mathbb{R} & \text{if } y_{ij} \text{ is NA}
\end{cases}
$$

We store all the y_{ij}^*'s corresponding to each vote record data point in matrix $\boldsymbol{Y}^*$. A joint posterior including the latent $\boldsymbol{Y}^*$ can be expressed as follows:

$$
\begin{aligned}
&P(\boldsymbol{Y}^*, \{\beta_j\}_{j=1}^J, \{\alpha_j\}_{j=1}^J, \{\boldsymbol{x}_i\}_{i=1}^N | \boldsymbol{Y}) \\
&\propto \prod_{i=1}^N \prod_{j=1}^J \Big(\mathbb{I}(y_{ij} = 1)\mathbb{I}(y_{ij}^* > 0) + \mathbb{I}(y_{ij} = 0)\mathbb{I}(y_{ij}^* < 0)\Big) \phi_1(y_{ij}^*; \boldsymbol{x}_i^T \beta_j + \alpha_j, 1) \\
&\quad \times \prod_{i=1}^N \phi_K(\boldsymbol{x}_i; \boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x) \prod_{j=1}^J \phi_{K+1}(\boldsymbol{\beta}_j; \boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta)
\end{aligned}
$$

where $\mathbb{I}(y_{ij}^* > 0)\phi_1(y_{ij}^*; \boldsymbol{x}_i^T\beta_j + \alpha_j, 1)$ is the PDF of a univariate truncated normal distribution, which only takes positive values. Similarly, $\mathbb{I}(y_{ij}^* < 0)\phi_1(y_{ij}^*; \boldsymbol{x}_i^T\beta_j + \alpha_j, 1)$ is the PDF of a univariate truncated normal distribution, which only takes negative values.

4.3 Penalized EM algorithm for IRT Model

The above Bayesian IRT model is usually estimated by MCMC simulations, which entail large time and memory cost as the size of a binary response data matrix grows. To reduce the time and memory cost of full Bayesian simulations, researchers have proposed an EM algorithm that's based on the above joint posterior distribution for fast estimation (Imai, Lo and Olmsted, 2016). As a typical EM algorithm, this method consists of two steps: E-step and M-step. The E-step establishes the target function (often called Q-function). The Q-function is the conditional expectation of the joint log-likelihood given the parameter values saved from the last iteration. In this method, Imai, Lo and Olmsted (2016) take conditional expectation of the missing values, y_{ij}^*'s, in the latent response matrix $\boldsymbol{Y}^*$. Then in the M-step, Imai, Lo and Olmsted (2016) maximize the Q-function with respect to the parameters, $\boldsymbol{\beta}_j$'s and $\boldsymbol{x}_i$'s. They use the same normal priors for both $\boldsymbol{\beta}_j$'s and $\boldsymbol{x}_i$ as in the full Bayesian method in the optimization step.

We build on the above EM algorithm by introducing an $L1$ penalty term on each β_j. We draw on the statistics literature on sparse factor analysis and sparse principle component analysis in devising the new penalized EM algorithm (Choi, Oehlert and Zou, 2010; Lee, Huang and Hu, 2010; Hirose and Konishi, 2012; Hirose and Yamamoto, 2014, 2015). Note that no penalty is imposed on any α_j. We also replace the normal prior on $\boldsymbol{\beta}_j$'s with an improper prior, $P(\{\boldsymbol{\beta}_j\}_{j=1}^J) \propto 1$. To simplify notations, we use $\vec{x}_i$ to denote $[1 \quad \boldsymbol{x}_i]^T$.

The Q-function for the penalized EM algorithm is expressed as follows:

$$
\begin{aligned}
& Q(\{\boldsymbol{\beta}_j\}_{j=1}^J, \{\vec{x}_i\}_{i=1}^N) \\
= \; & \mathbb{E}\Big[\log P(\boldsymbol{Y}^*, \{\boldsymbol{\beta}_j\}_{j=1}^J, \{\vec{x}_i\}_{i=1}^N | \boldsymbol{Y}) \Big| \{\boldsymbol{\beta}_j^{(t-1)}\}_{j=1}^J, \{\vec{x}_i^{(t-1)}\}_{i=1}^N \Big] \\
= \; & -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^J (\boldsymbol{\beta}_j^T \vec{x}_i \vec{x}_i^T \boldsymbol{\beta}_j - 2\boldsymbol{\beta}_j^T \vec{x}_i y_{ij}^{*(t)}) - \frac{1}{2} \sum_{i=1}^N (\boldsymbol{x}_i^T \boldsymbol{\Sigma}_x^{-1} \boldsymbol{x}_i - 2\boldsymbol{x}_i^T \boldsymbol{\Sigma}_x^{-1} \boldsymbol{\mu}_x) + \text{constant}
\end{aligned}
$$

where

$$
\begin{aligned}
y_{ij}^{*(t)} &= \mathbb{E}[y_{ij}^* | \vec{x}_i^{(t-1)}, \boldsymbol{\beta}_j^{(t-1)}, y_{ij}] \\
&= \begin{cases}
\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)} + \dfrac{\phi(\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)})}{\Phi(\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)})} & \text{if } y_{ij} = 0 \\
\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)} - \dfrac{\phi(\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)})}{1 - \Phi(\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)})} & \text{if } y_{ij} = 1 \\
\vec{x}_i^{(t-1)T} \boldsymbol{\beta}_j^{(t-1)} & \text{if } y_{ij} \text{ is NA}
\end{cases}
\end{aligned}
$$

In the E-step, we minimize the above Q function with respect to $\{\vec{x}_i\}_{i=1}^N, \{\boldsymbol{\beta}_j\}_{j=1}^J$ sequentially. The optimization formulation for $\{\vec{x}_i\}_{i=1}^N$ is the same as in the previous EM algorithm:

$$
\boldsymbol{x}_i^{(t)} = \Big(\boldsymbol{\Sigma}_x^{-1} + \sum_{j=1}^J \beta_j^{(t-1)} \beta_j^{(t-1)T} \Big)^{-1} \Big(\boldsymbol{\Sigma}_x^{-1} \boldsymbol{\mu}_x + \sum_{j=1}^J \beta_j^{(t-1)} (y_{ij}^{*(t)} - \alpha_j^{(t-1)}) \Big)
$$

However, the optimization formula for $\{\boldsymbol{\beta}_j\}_{j=1}^J$ is different due to the improper prior specification.

$$
\boldsymbol{\beta}_j^{(t)} = \Big(\sum_{i=1}^N \vec{x}_i^{(t)} \vec{x}_i^{(t)T} \Big)^{-1} \Big(\sum_{i=1}^N \vec{x}_i^{(t)} y_{ij}^{*(t)} \Big)
$$

Before adding penalty, the above formula gives us the solution to $\{\boldsymbol{\beta}_j\}_{j=1}^J$ in the unconstrained M-step. As shown by the solution form, the optimal values of $\boldsymbol{\beta}_j$ is equivalent to the solution to

the optimization problem below:

$$\min_{\boldsymbol{\beta}_j} \left(y_j^{*(t)} - \vec{x}_i^{(t)T}\boldsymbol{\beta}_j\right)^T \left(y_j^{*(t)} - \vec{x}_i^{(t)T}\boldsymbol{\beta}_j\right)$$

In order to estimate the sparse item-dimension loading structure, we add an $L1$ penalty on each β_j (not on α_j) to the above optimization problem in every M-step. The regularized optimization problem is transformed into the problem below:

$$\min_{\boldsymbol{\beta}_j} \left(y_j^{*(t)} - \vec{x}_i^{(t)T}\boldsymbol{\beta}_j\right)^T \left(y_j^{*(t)} - \vec{x}_i^{(t)T}\boldsymbol{\beta}_j\right) + \lambda||\beta_j||_1$$

where λ controls the amount of penalty imposed on non-zero elements in β_j. In the statistics literature, the standard and efficient method to solve the above penalized optimization problem is the cyclic coordinate descent algorithm (Friedman et al., 2007; Friedman, Hastie and Tibshirani, 2010). The name "coordinate descent" comes from the fact that we update each dimension (coordinate) k of $\boldsymbol{\beta}_j$ given the current values of all the other β_{jk}'s. For each $\boldsymbol{\beta}_j$, we treat the optimization problem as a regularized least square problem, and we update β_{jk}'s for $k = 0, 1, \cdots, K$ for multiple iterations until convergence.

We use $\beta_{jk}^{(t)}$ to denote the current values for β_{jk} in an iteration within the coordinate descent process. Since we do not apply any penalty on α_j's (β_{j0}'s), we update α_j's as a normal least square problem in the following way:

$$\alpha_j = \beta_{j0} \leftarrow \frac{1}{N}\sum_1^N \left(y_{ij}^{*(t)} - \sum_1^K x_{ik}^{(t)}\beta_{jk}^{(t)}\right), \text{ given } \beta_{jk}^{(t)} \text{ for } k = 1, \cdots, K$$

Before we update any β_{jk} in the coordinate descent process, we define $y_{ij}^{*(t,k)}$ as the difference between $y_{ij}^{*(t)}$ and the dot product of the current ideal points and the current values of item parameters (including intercept) for all the other dimensions than k. $y_{ij}^{*(t,k)}$ is computed by the formula below:

$$y_{ij}^{*(t,k)} = y_{ij}^{*(t)} - \sum_{k \in \{0,\cdots,k-1,k+1,\cdots,K\}} x_{ik}^{(t)} \beta_{jk}^{(t)}$$

Then we update all the discrimination parameters as an $L1$ regularized least square problem as below:

$$\beta_{jk} \leftarrow S\left(\left(\sum_{1}^{N} x_{ik}^{(t)2} \right)^{-1} \sum_{1}^{N} x_{ik}^{(t)} y_{ij}^{*(t,k)}, \lambda \right)$$

where $S(z, \gamma)$ is the soft-thresholding operator below:

$$\text{sign}(z)(|z| - \gamma)_+ = \begin{cases} z - \gamma & \text{if } z > 0 \text{ and } \gamma < |z| \\ z + \gamma & \text{if } z < 0 \text{ and } \gamma < |z| \\ 0 & \text{if } \gamma \geq |z| \end{cases}$$

The penalized EM-algorithm updates $\{\vec{x}_i\}_{i=1}^N$ and $\{\boldsymbol{\beta}_j\}_{j=1}^J$ until convergence. We define the convergence criterion in a similar way as in the previous EM algorithm (Imai, Lo and Olmsted, 2016). The convergence criterion is based on large correlations of parameter estimates between two adjacent iterations under a pre-specified small threshold value, τ.[3]

$$\text{converge criterion} = \mathbb{I}[\text{Cor}(\vec{x}^{(t-1)}, \vec{x}^{(t)}) > 1 - \tau]$$

where $\vec{x}^{(t)}$ is a long vector that stores all the elements in $\{\vec{x}_i\}_{i=1}^N$ for the current iteration. We choose the value of λ by cross-validation based on Bayesian Information Criterion (BIC) (Sun et al., 2016). We start with a grid of λ values and fit the above penalized EM algorithm repeatedly under each λ value. After each trial converges, we store the estimated parameters. We compute the BIC values with the final optimal parameter values under each λ value. We choose the λ value

[3]Imai, Lo and Olmsted (2016) compute the cross-iteration correlations for both $\{\vec{x}_i\}_{i=1}^N, \{\boldsymbol{\beta}_j\}_{j=1}^J$. It's inappropriate for us to compute the cross-iteration correlations for$\{\boldsymbol{\beta}_j\}_{j=1}^J$ because there are many exact zero values in $\{\boldsymbol{\beta}_j\}_{j=1}^J$ estimates in our model. Therefore, we only compute the across-iteration correlations for $\{\vec{x}_i\}_{i=1}^N$. Based on our simulation studies, this convergence criterion works well.

that returns the least BIC value.

4.4 Results on Simulated Data

To demonstrate the effectiveness of the new penalized EM algorithm for the IRT model, we run the algorithm on simulated data and compare the estimated parameters with the ground truth. We simulate a roll call vote matrix with 50 voters and 90 items. The latent ideal points have 3 dimensions. α_j's and β_j's are first randomly drawn from a normal distribution with mean 1 and variance 0.4. Then we multiply α_j's and β_j's by -1 or 1 with equal chance. x_{ik} is sampled from a standard normal distribution. Then we impose a sparse item-dimension loading structure on the item parameter matrix. We let the first 30 items only load on the first dimension, and therefore $\beta_{jk} = 0, \forall k \in \{2, 3\}, j \in \{1, 2, \cdots, 30\}$ and $\beta_{jk} \neq 0, \forall k = 1, j \in \{1, 2, \cdots, 30\}$. Likewise, we let the second 30 items only load on the second dimension, and the last 30 items only load on the third dimension. Therefore, for the first experiment, the discrimination parameter matrix has 180 zero value elements and 90 non-zero value elements. Lastly, an element, y_{ij}, in the simulated roll call vote matrix is randomly sampled from a Bernoulli distribution. The parameter p_{ij} of the Bernoulli distribution for drawing y_{ij} is defined as $p_{ij} = \Phi(\alpha_j + \boldsymbol{x}_i^T \beta_j)$.

In sum, the penalized EM algorithm is able to accurately identify the true sparse item-dimension loading structure and consistently uncover the true multidimensional ideal points. First, we show the results on one simulated roll call vote matrix to illustrate how to run the algorithm and how to make sense of the results. We compare the results from both the previous EM algorithm and the new penalized EM algorithm to show why it's ideal to use the penalized EM algorithm in the presence a sparse item-dimension loading structure. Second, we repeat the experiments 500 times on different simulated roll call vote matrices based on different item-dimension loading structures.

First, we run the previous EM algorithm without penalty. The estimated $\boldsymbol{\beta}$ is plotted against the true $\boldsymbol{\beta}$. Figure 4.1 shows that the EM algorithm without penalty does not uncover the zero elements in the true $\boldsymbol{\beta}$. We use β_0 to represent α for an item in this section. The correlations between

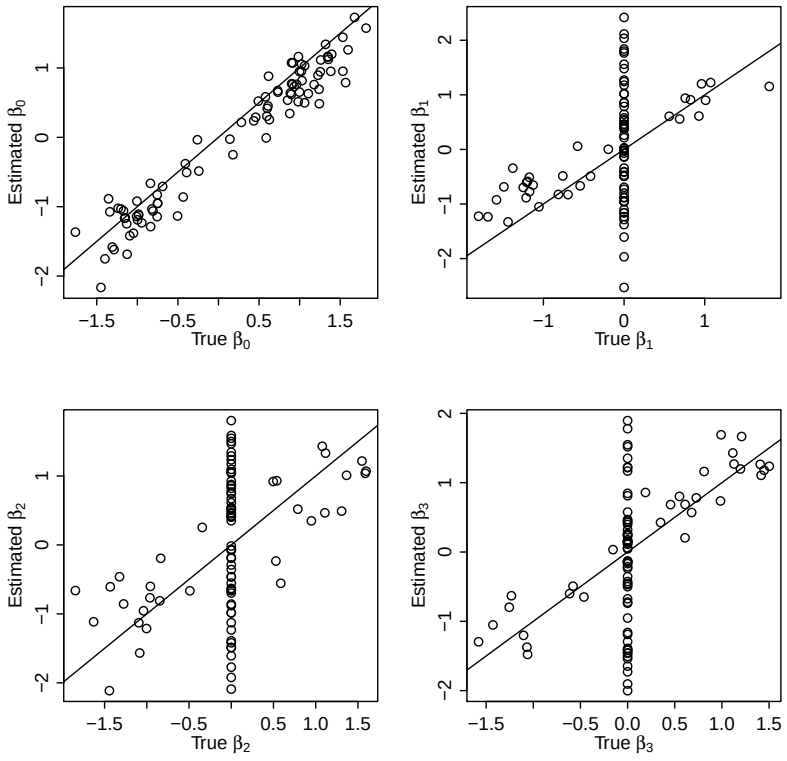

Figure 4.1: *Estimated β_k with No Penalty.* The EM algorithm without penalty does not uncover the zero elements in the true $\boldsymbol{\beta}$. The correlations between the true β_k's and the estimated β_k's, for $k = 0, 1, 2, 3$, are 0.9697469, 0.4673016, 0.4802608, 0.5763199.

the true β_k's and the estimated β_k's, for $k = 0, 1, 2, 3$, are 0.9697469, 0.4673016, 0.4802608, 0.5763199. Figure 4.2 plots the estimated x_k's from the previous EM algorithm against the true x_k's . The correlations between the true x_k's and the estimated x_k's for $k = 1, 2, 3$, are 0.4560569 ,0.4584138, 0.6555498. The above results show that the estimated item parameters and ideal points are inconsistent without penalty in the presence of a true sparse item-dimension loading structure.

Second, we run the penalized EM algorithm on the simulated roll call vote data. We start with a grid of λ values, and run the EM algorithm repeatedly under each λ value. We plot the BIC value for each trial against the λ value. We choose the λ value that returns the least BIC value. Figure 4.3 shows that the optimal value of λ is 0.46. We use the optimal λ value, 0.46, to compute the final estimated parameters from the penalized EM algorithm.

Figure 4.4 shows that the penalized EM algorithm uncovers most of the zero elements in $\boldsymbol{\beta}$.

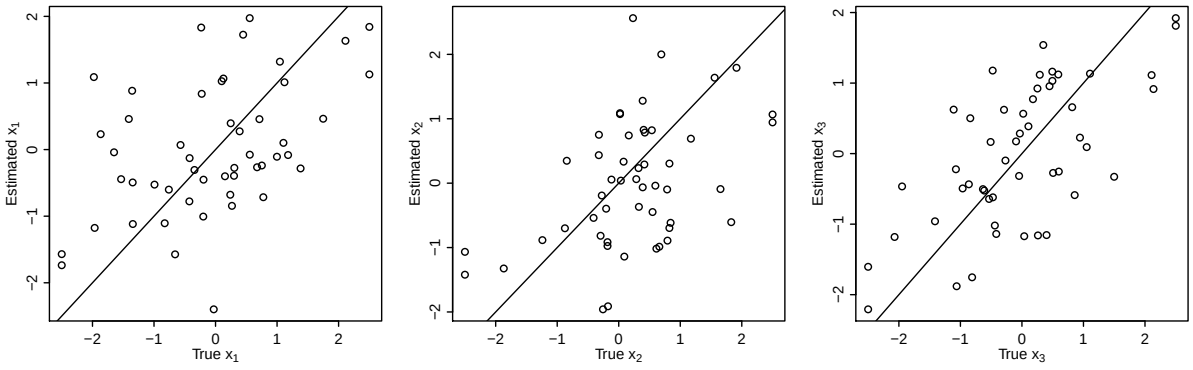

Figure 4.2: *Estimated x_k with No Penalty.* The EM algorithm without penalty does not uncover the ground truth of x_k's. The correlations between the true x_k's and the estimated x_k's for $k = 1, 2, 3$, are 0.4560569 ,0.4584138, 0.6555498.

The correlations between the true β_k's and the estimated β_k's for $k = 0, 1, 2, 3$, are 0.9615604, 0.9587055, 0.93453, 0.9320354. We can measure the performance of the penalized EM algorithm by thinking of the sparse item-dimension loading structure estimation task as a classification problem. In turn, we examine how well the penalized EM algorithm can correctly classify the true zero elements in $\boldsymbol{\beta}$ as zero. We report several classification performance metrics below. In aggregate, the penalized EM algorithm correctly classifies 82.18% of elements in $\boldsymbol{\beta}$ as either zero or non-zero. The penalized EM algorithm correctly classifies 92.22% (sensitivity metric) of the true non-zero elements in $\boldsymbol{\beta}$. The penalized EM algorithm correctly classifies 90% (specificity metric) of the true zero elements in $\boldsymbol{\beta}$. The overall F_1 score for this classification task is 0.8691. Figure 4.5 plots the estimated x_k's against the true x_k's. The correlations between the true x_k's and the estimated x_k's, for $k = 1, 2, 3$, are 0.9599019, 0.9299417, 0.9522259. The above results show that we can consistently estimate the item parameters and the ideal points with the penalized EM algorithm in the presence of a sparse item-dimension loading structure.

Third, we repeat the above roll call vote matrix simulation and penalized EM algorithm fitting for 500 times. In each data simulation, we randomly draw the item-dimension loading structure. Specifically, for each item j we first draw the number of zero discrimination parameters, t_j, from $\{0, 1, 2\}$ with equal likelihood. Then we draw t_j discrimination parameters without replacement

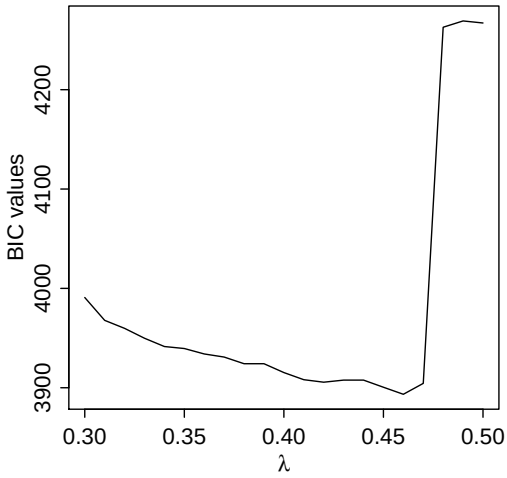

Figure 4.3: *BIC Values.* The optimal value of λ is 0.46.

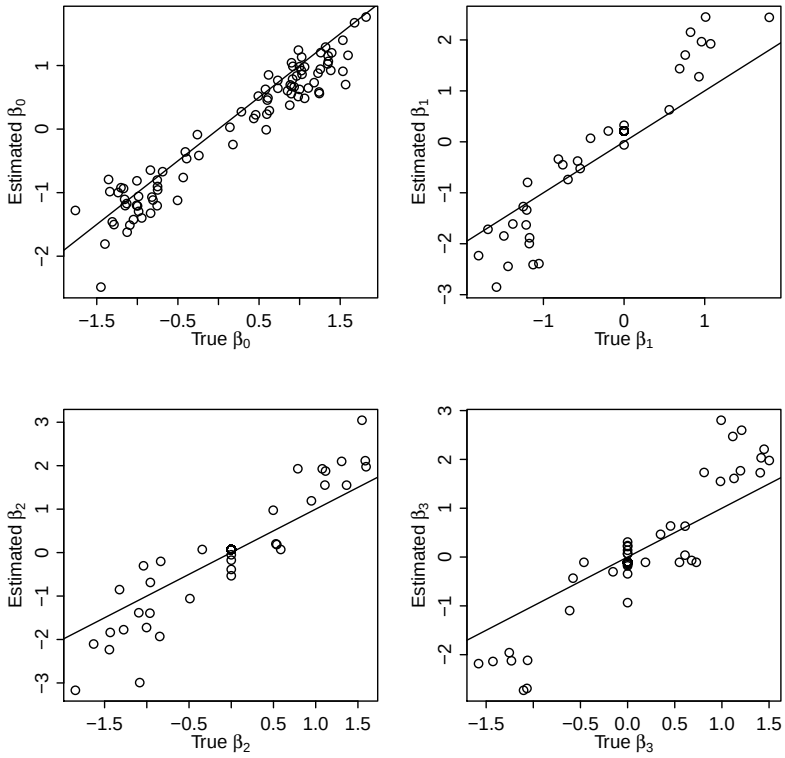

Figure 4.4: *Estimated β_k with Penalty.* The penalized EM algorithm uncovers most of the zero elements in $\boldsymbol{\beta}$. The correlations between the true β_k's and the estimated β_k's for $k = 0, 1, 2, 3$, are 0.9615604, 0.9587055, 0.93453, 0.9320354.

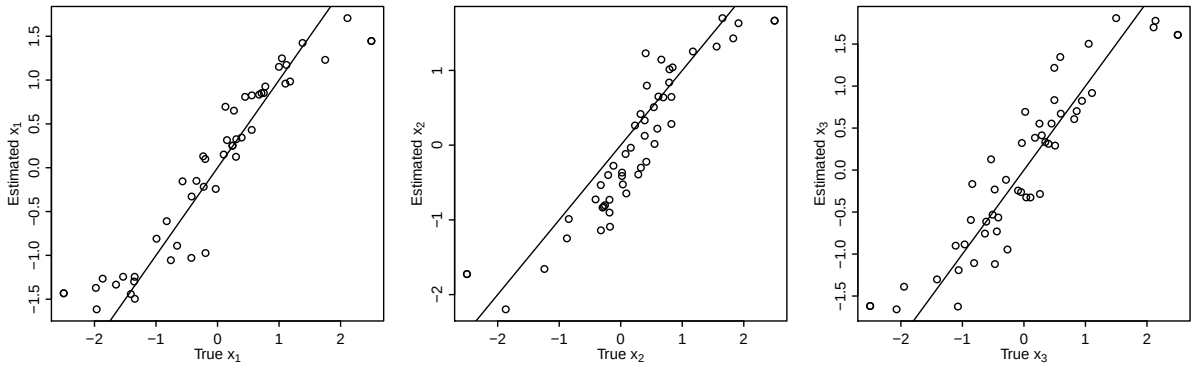

Figure 4.5: *Estimated x_k with Penalty.* The penalized EM algorithm uncovers the ground truth of $\boldsymbol{x}$. The correlations between the true x_k's and the estimated x_k's, for $k = 1, 2, 3$, are 0.9599019, 0.9299417, 0.9522259.

from $\{\beta_{j1}, \beta_{j2}, \beta_{j3}\}$ and set them to be equal to zero. Other than the distinct item-dimension loading structures, we use the same simulation steps to generate 500 roll call vote matrices.

We fit the penalized EM algorithm with the same random seed and staring values on the 500 roll call vote matrices respectively. Because of the fixed starting values, we expect that the results for some of the replications may not be ideal. Similar to all EM algorithm methods, the proposed algorithm in this paper requires researchers to experiment with multiple starting values. However, the Monte Carlo experiments still show that the penalized EM algorithm works almost all of the times in spite of the inflexible starting values. As shown by Figure 4.6, the penalized EM algorithm can accurately recover the true values of ideal points for all three latent dimensions for nearly all the trials. Similarly, Figure 4.7 shows that the penalized EM algorithm can accurately estimate the true values of item parameters for almost all of the trials. More importantly, as illustrated by Figure 4.8, the EM algorithm does a decent job at recovering the sparse item-dimension loading structure for nearly all of the trials. The vast majority of the classification metrics, including Precision, Sensitivity, Specificity and F1 Score, exceed 80%.

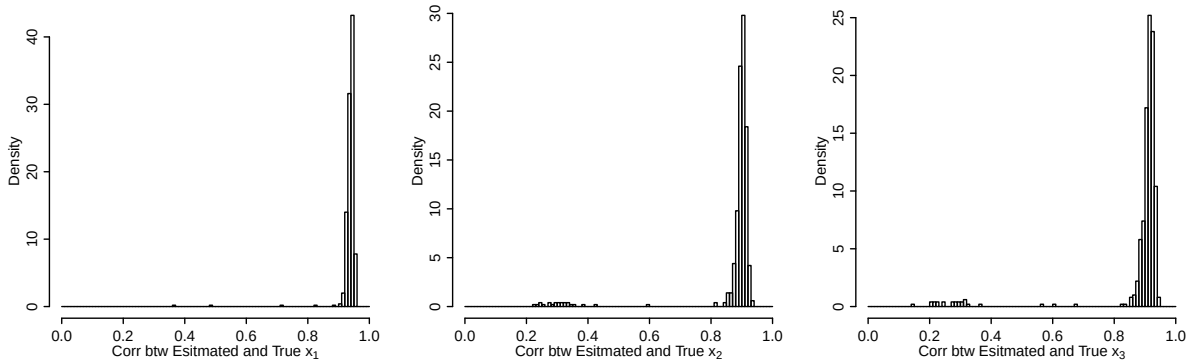

Figure 4.6: *Correlations between Estimated x_k and True x_k for 500 Trials.* The penalized EM algorithm can accurately recover the true values of ideal points for all three latent dimensions for nearly all the trials.

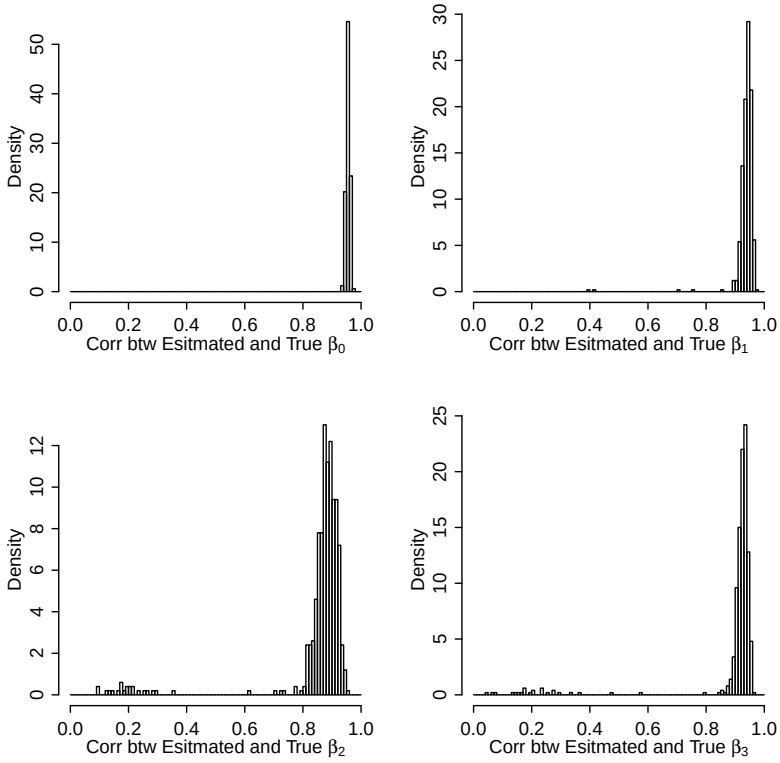

Figure 4.7: *Correlations between Estimated β_k and True β_k for 500 Trials.* The penalized EM algorithm can accurately estimate the true values of item parameters for almost all of the trials.

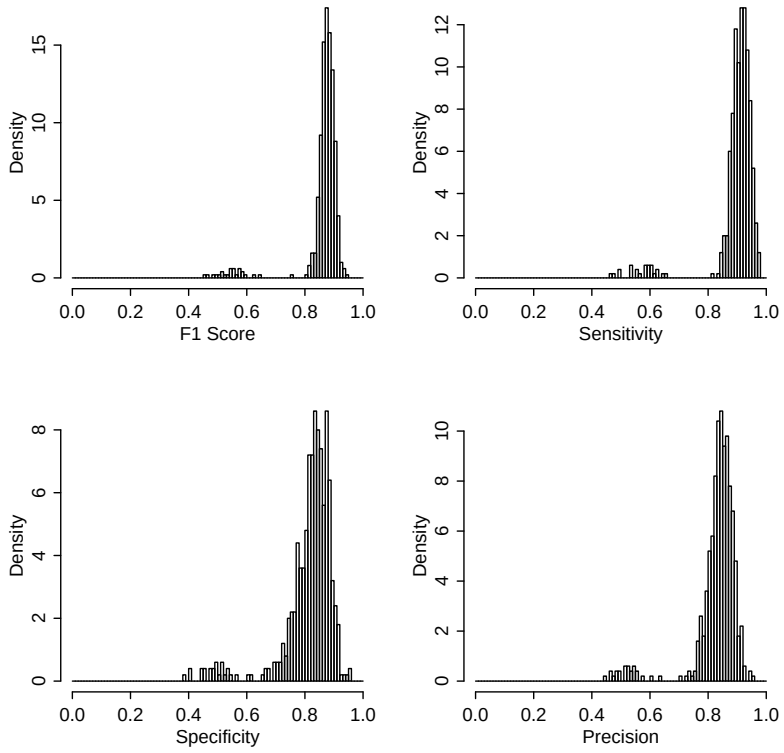

Figure 4.8: *Classification Metrics for 500 Trials.* The EM algorithm does a decent job at recovering the sparse item-dimension loading structure for nearly all of the trials. The vast majority of the classification metrics, including Precision, Sensitivity, Specificity and F1 Score, exceed 80%.

4.5 US Senator Ideal Point Estimation

We replicate a previous study on estimating 2-dimensional ideal points for US senators in Jackman (2001). In this study, Jackman emphasizes the importance of estimating the item-dimension loading structure and draws a mathematical connection between multidimensional IRT models and factor analysis models. His strategy to identify the item-dimension loading structure is to first estimate the item parameters without penalty and then examine if a discrimination parameter is significantly different than zero. He runs the full Bayesian IRT model specified above, and uses the 90% credible interval of the posterior sample of a discrimination parameter as a significance check. If the 90% credible interval of the posterior sample of a discrimination parameter includes zero, then Jackman classifies this discrimination parameter as zero. This way, he converts the original discrimination parameter matrix into a sparse matrix, and he identifies special anchoring items based on the sparse discrimination parameter matrix.

Jackman's study contributes to the literature of ideal point estimation by highlighting the importance of item-dimension loading structures. However, this method has a few potential pitfalls. First, as shown by the simulation data example, if the item-dimension loading structure is truly sparse, then an estimation method without penalty cannot consistently estimate the ideal points or the item parameters. Therefore, the full Bayesian approach may not be able to accurately estimate the item parameters in the first place, and this makes the significance check based on the posteriors of the item parameters less credible. Second, the criterion for the significance check is arbitrary, be it 90% or 95%, and we do not have a clear way to choose the optimal credible interval level.

We replicate this study with the penalized EM algorithm. Our goal is to show that the penalized EM algorithm is able to definitively identify an item-dimension loading structure and the special anchoring items. Then we closely examine the anchoring items and explain what these anchoring items mean substantively. We use the US Senate 105th Session's roll call vote data as in the original study. As in the original study, we remove all the bills where less than 3 senators are in the voting minority, because these extremely lopsided votes do not provide much information on a senator's ideal point. We end up with 100 senators and 486 bills in the data.

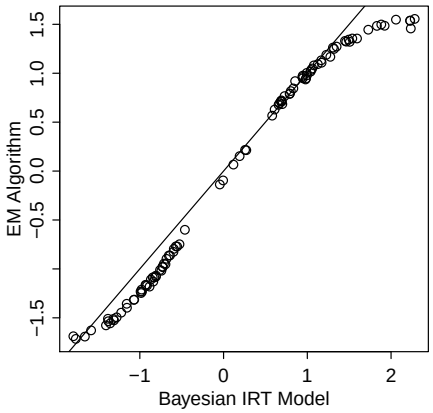

Figure 4.9: *1-D Ideal Point Estimates from Bayesian and EM IRT Models.* The estimated unidimensional ideal points from full Bayesian and EM IRT models are highly correlated.

First, we run a unidimensional IRT model with the previous EM algorithm without penalty. We use the starting values -3 for Senator Paul Wellstone (D-MN) and 3 for Senator Jim Inhofe (R-OK) due to their well-known and relatively outlying ideological leanings. We save the estimated unidimensional ideal points for all the senators. For comparison, we also run a full Bayesian unidimensional IRT model on the same data set. Figure 4.9 shows that the estimated unidimensional ideal points from full Bayesian and EM IRT models are highly correlated. We use the above unidimensional ideal point estimates as the starting values for the first dimension ideal points when we run the 2-dimensional EM IRT model. The second dimension of the ideal points starting values are randomly generated from a standard normal distribution. In preparation, we run the 2-dimensional EM IRT model without penalty with the above starting ideal point values, and save the 2-dimensional ideal point estimates.

We use the above 2-dimensional ideal point estimates as the starting values for running the penalized EM algorithm. We start with a grid of λ values, a sequence from 0.3 to 0.6 by step 0.01. We run the 2-dimensional penalized EM algorithm under each λ value. Figure 4.10 shows that the optimal λ value is 0.47. We use the optimal λ value 0.47 to run the final 2-dimensional penalized EM algorithm on the roll call vote data. In Figure 3 of the original paper (Jackman, 2001), Jackman

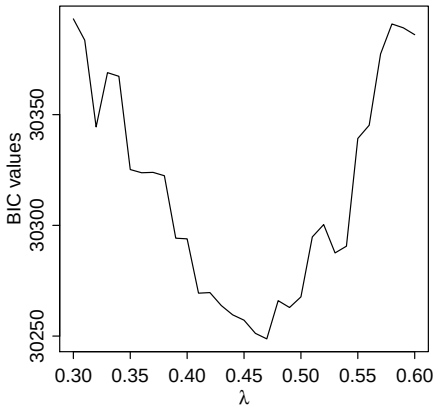

Figure 4.10: *BIC Values for Replication Study.* The optimal λ value is 0.47.

Senator	Estimated 1st D	Original 1st D	Estimated 2nd D	Original 2nd D
BREAUX	-0.5514101	-	0.2886047	+
LOTT	1.142636	+	-0.1434896	+
MCCAIN	0.640696	+	-0.3509587	-
ASHCROFT	1.076252	+	-1.040977	-
CHAFEE	0.1636766	+	0.5292594	+
DASCHLE	-1.155316	-	0.353784	+
WELLSTONE	-1.465541	-	-0.07373871	-

Table 4.1: *2-D Ideal Points from Penalized EM Algorithm and Signs of Original Ideal Points.* The penalized EM algorithm returns the same signs for the ideal points of the selected senators except for Senator Lott (R-MS)'s second dimension ideal point.

reports the 2-dimensional ideal points of a few representative senators.[4] The 2-dimensional ideal points of these senators from the penalized EM algorithm have almost the same signs as those in the original study. Table 4.1 shows that the penalized EM algorithm returns the same signs for the ideal points of the selected senators except for Senator Lott (R-MS)'s second dimension ideal point.

In the original study, Jackman identifies 282 bills that only load on the first dimension and 12 bills that only load on the second dimension. The penalized EM algorithm finds 306 pure-first-dimension bills and 40 pure-second-dimension bills. To show the meanings of the two dimen-

[4]Jackman does not report the exact values of the ideal point estimates in the paper. Judging by the plot, we can only tell the signs of the ideal points in the original study.

sions, we study the two groups of anchoring items for the first and second dimensions respectively. Substantively, the 306 pure-first-dimension bills include almost all kinds of policy issues, so it's difficult to summarize the first dimension by the pure-first-dimension bills' contents. The 40 pure-second dimension bills have more coherent substantive meanings. There are two policy topics that repeatedly show up in the pure-second-dimension issues: appropriations and foreign policy. Bill No.63, 92, 96, 181, 185, 209, 235, 251, 257, 264, 283, 298, 337, 338, 339, 342, 416, 512, 558, and 612 are all about appropriations or budget issues. Bills that are related to foreign policy include bill No. 25, 184, 185, 339, 342, 345, 410, and 546. The summaries of the above bills are reported in the tables in Appendix C.

To further differentiate the first and second dimensions of ideal points, we employ senators' party ID data to check if there is a clear difference in the degree of partisanship between the pure-first-dimension and the pure-second-dimension bills. If we can predict each senator's vote on a bill by her party ID fairly well, then we regard this bill as more partisan. On the contrary, if we cannot accurately predict the votes on a bill based on senators' party IDs, then we think of this bill as less partisan. We run 306 logistic regression models with senators' party ID as the predictor and senators' votes on a bill as the response for each pure-first-dimension bill. Then we use the learned logistic regression model to predict senators' votes on each pure-first-dimension bill. For this two-class prediction problem, we use the AUROC (Area Under the Receiver Operating Characteristics) metric to evaluate the prediction performance. We repeat the above procedure on the 40 pure-second-dimension bills.

As shown by Figure 4.11, we can use only party ID data to decently predict most of the pure-first-dimension bills, whereas the party ID variable does not give us much leverage for predicting votes for the pure-second-dimension bills. Figure 4.12 more clearly contrasts the predictability of votes based on party ID and partisanship for the pure-first-dimension and pure-second-dimension bills. This finding provides further evidence to the theory in congress studies: one dimension of ideal points should capture the partisanship of a legislator, and the other dimensions should capture a legislator's preference other than voting along the party line (Poole and Rosenthal, 1985). By

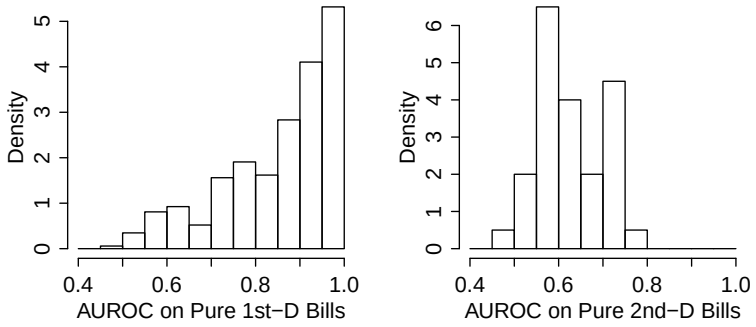

Figure 4.11: *Histogram of AUROC Values on Anchoring Items.* We can use only party ID data to decently predict most of the pure-first-dimension bills, whereas the party ID variable does not give us much leverage for predicting votes for the pure-second-dimension bills.

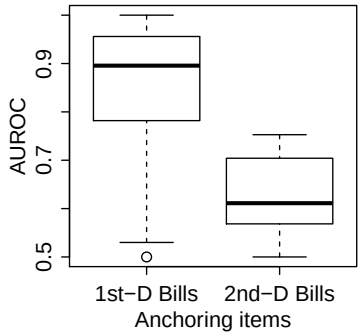

Figure 4.12: *Boxplot of AUROC Values on Anchoring Items.* This boxplot more clearly contrasts the predictability of votes based on party ID and partisanship for the pure-first-dimension and pure-second-dimension bills. The pure-first-dimension bills are much more partisan than the pure-second-dimension bills.

using the penalized EM algorithm, we are able to identify the bills that are inherently partisan or less partisan, and these anchoring items provide stronger evidence for the latent dimension meanings.

4.6 Conclusion

We have proposed an innovative penalized EM algorithm for estimating sparse item-dimension loading structures in multidimensional IRT models. The new penalized EM algorithm identifies a sparse item-dimension loading structure for multidimensional IRT models by applying an $L1$ penalty on discrimination parameters in model estimation. The sparse item-dimension loading structure aids us in identifying the anchoring items for each dimension and provides information on what each dimension means. In addition, the penalized EM algorithm has the flexibility to allow a discrimination parameter to be exactly zero. Hence, the penalized EM algorithm can consistently estimate the multidimensional ideal points when the data generation process is truly based on a sparse item-dimension loading structure. We have used simulation data to demonstrate how the penalized EM algorithm can accurately recover both the true item-dimension loading structure and ideal points. Then, we have also replicated a previous study on the 105th US Senate roll call vote data to show how the penalized EM algorithm can help us infer latent dimension meanings.

CHAPTER 5

Conclusion

In this dissertation, we have proposed three new statistical models for analyzing discrete choice data. These new models all fall into the category of Bayesian latent variable models, and respectively contribute to three political methodology literatures: pairwise comparison models, Dirichlet process mixture models, and Bayesian IRT models. The newly proposed models improve the existing models by relaxing unidimensional assumptions, modeling temporal dependence of cluster numbers, and supporting richer structures of latent variables. We have applied the three new models to survey data or roll call vote data. These models uncovered the underlying patterns from the opinion or vote data, and therefore shed important light on actors' preferences and behaviors. In these empirical studies, we have derived new insights about important political science topics, such as political biases in people's perceptions of COVID-19 related statements, politicization of human rights in the United Nations, and multiple issue dimensions of legislators' ideal points. Below we reiterate our findings and contributions in the three chapters above.

In the second chapter, we have proposed a new multidimensional pairwise comparison model to measure multidimensional latent attributes and respondent-specific perceptual parameters. We have applied this new model to original survey data where respondents are asked to judge the truthfulness of pairs of statements on COVID-19. We find a weak correlation between the actual truthfulness of a statement and respondents' perceptions of truthfulness. More importantly, we find that the political valence of statements is largely responsible for the variation in perceived truthfulness. Co-partisanship between a respondent and the speaker of a statement predicts higher

perceived truthfulness. We have also observed associations between the respondent-specific perceptual parameters and respondents' practice of mask-wearing or social distancing.

In the third chapter, we have proposed the Dynamic Dirichlet Process Mixture (DDPM) model to identify voting coalitions with roll call votes across multiple periods, along with post-processing methods to analyze the posterior samples from the DDPM model. We have applied these methods to the United Nations General Assembly (UNGA) human rights roll call vote data from 1992 to 2017. We identify human rights voting coalitions in the UNGA after the Cold War, and the polarizing resolutions that divide countries into different coalitions. Through this study, we find a clear separation between developing and developed countries in human rights voting. Moreover, we find the EU as a stable coalition among the developed countries, whereas some Latin American countries tend to form a small coalition, apart from a large developing country coalition. We also find that many polarizing resolutions repeatedly show up across years, and the lines of conflict lie in both debates on general human rights principles and human rights violation reports on specific countries.

In the fourth chapter, we have proposed an innovative penalized EM algorithm for estimating sparse item-dimension loading structures in multidimensional IRT models. The sparse item-dimension loading structure aids us in identifying the anchoring items for each dimension, which provide information on what each dimension means. In addition, the penalized EM algorithm has the flexibility to allow a discrimination parameter to be exactly zero. Hence, the penalized EM algorithm can consistently estimate the multidimensional ideal points when the data generation process is truly based on a sparse item-dimension loading structure. We have replicated a previous study on the 105th US Senate roll call vote data to show how the penalized EM algorithm can help us identify anchoring items for latent dimensions and thereby infer latent dimension meanings.

APPENDIX A

COVID-19 Statements and Respondent Descriptive Statistics

A.1 COVID-19 Statements

In the following 7 tables, we report the 42 unique statements used in the survey. We show each statement in the same format as in the survey. In the "Statement" column, we report a statement's content, the speaker, the date when the statement was made, and the channel through which the statement was made. We also report the unique IDs, the truthfulness ratings given by PolitiFact website, the grouped truthfulness ratings, and the partisanship labels for all the statements. We group "pants-on-fire", "false", and "mostly-false" statements as "low-truth" statements. We group "half-true", "mostly-true", and "true" statements as "high-truth" statements.[1]

[1]Note that in 2011 Politifact changed the label on their third lowest truthfulness rating from "barely true" to "mostly false" (https://www.politifact.com/article/2011/jul/27/-barely-true-mostly-false/). While the ratings of our statements that appear on the Politifact website reflect this change, it appears that their underlying database that we scraped still uses "barely true" for this category. We use "mostly false" in this paper, but the underlying data we rely upon uses "barely true".

Table A.1: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1001	"Joe Biden was in charge of the H1N1 Swine Flu epidemic which killed thousands of people. The response was one of the worst on record. Our response is one of the best, with fast action of border closings & a 78% Approval Rating, the highest on record. His was lowest!" This statement was made by Donald Trump on March 12, 2020 in a tweet.	barely-true	low-truth	right
1002	"We've still had more deaths to the flu this year than we've had COVID-19." This statement was made by Dan Forest on April 24, 2020 in an interview.	false	low-truth	right
1003	"There was no real scientific basis for believing that social distancing would be necessary, since it had never been studied." This statement was made by Laura Ingraham on May 4, 2020 in a TV segment.	false	low-truth	right
1004	"We've tested more than every country combined." This statement was made by Donald Trump on April 27, 2020 in comments made during a White House briefing.	pants-fire	low-truth	right
1005	"Nancy Pelosi's coronavirus bill includes tax credits for solar and wind energy, retirement plans for community newspaper employees, $300 million for PBS, climate change studies, and more." This statement was made by Donald Trump on March 24, 2020 in a campaign ad.	mostly-true	high-truth	right
1006	"Longstanding Food and Drug Administration regulations created barriers to the private industry creating a test quickly for the coronavirus." This statement was made by Dan Crenshaw on March 13, 2020 in a tweet.	true	high-truth	right

Table A.2: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1007	"The United States is actually screening fewer people (for the coronavirus than other countries) because we don't have appropriate testing." This statement was made by Lou Dobbs on March 2, 2020 in a TV interview.	mostly-true	high-truth	right
1008	"During the 2009 swine flu outbreak, Biden made reckless comments unsupported by science & the experts. The Obama Admin had to clean up his mess & apologize for his ineptitude." This statement was made by Donald Trump on March 12, 2020 in a tweet.	mostly-true	high-truth	right
1009	"Three Chinese nationals were apprehended trying to cross our Southern border illegally. Each had flu-like symptoms. Border Patrol quickly quarantined them and assessed any threat of coronavirus." This statement was made by Charlie Kirk on February 24, 2020 in a tweet.	mostly-true	high-truth	right
1010	"Regarding the risks of coronavirus transmission on an airplane, it's as safe as an environment as you're going to find." This statement was made by Gary Kelly on May 3, 2020 in an interview on CBS.	barely-true	low-truth	neutral
1011	"The April 22 jump in COVID-19 cases was related to the election." This statement was made by Chris Larson on April 23, 2020 in a Facebook post.	false	low-truth	neutral
1012	"They're furloughing nurses in hospitals in western New York state." This statement was made by Jerome Adams on March 27, 2020 in a television interview.	false	low-truth	neutral

Table A.3: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1013	"Silver Solution used on strains of coronavirus 'totally eliminate it. Kills it. Deactivates it.' " This statement was made by Sherill Sellman on February 12, 2020 in remarks on the Jim Bakker show.	pants-fire	low-truth	neutral
1014	"It's actually the safest time to fly." This statement was made by Ainsley Earhardt on March 13, 2020 in the Fox & Friends show.	false	low-truth	neutral
1015	"Small trials to test convalescent plasma therapy for coronavirus patients seem to have had some degree of success." This statement was made by Jamie Nadler on April 14, 2020 in a newspaper interview.	true	high-truth	neutral
1016	"Covid-19 is now the leading cause of death in the United States." This statement was made by Mandy Cohen on April 20, 2020 in a press conference.	half-true	high-truth	neutral
1017	"Some states are only getting 50 tests per day, and the Utah Jazz got 58." This statement was made by Michael Dougherty on March 12, 2020 in a tweet.	true	high-truth	neutral
1018	"You're more likely to die of influenza right now than the 2019 coronavirus." This statement was made by Dr. Drew Pinsky on February 3, 2020 in a Daily Blast Live segment.	half-true	high-truth	neutral

Table A.4: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1019	"(Barack Obama) set up anti-pandemic programs in 47 vulnerable countries as a way to protect against something like the coronavirus, exactly. Do you know that Trump closed 37 of them?" This statement was made by Joy Behar on March 9, 2020 in an episode of ABC's "The View".	false	low-truth	left
1020	"The health insurance industry has agreed to waive all co-payments for coronavirus treatments." This statement was made by Donald Trump on March 11, 2020 in a White House speech.	false	low-truth	right
1021	"The World Health Organization offered the testing kits that they have available and to give it to us now. We refused them. We did not want to buy them." This statement was made by Joe Biden on March 15, 2020 in a Democratic primary debate.	barely-true	low-truth	left
1022	"There was no effort to get American experts into China after it announced the coronavirus, and we had one person in-country (and Trump) pulled him out of the country." This statement was made by Joe Biden on March 27, 2020 in a virtual town hall on CNN.	barely-true	low-truth	left
1023	"If you line up all the countries that have done (Covid-19) testing on a per-capita basis, we're at the bottom of the list." This statement was made by Bobby Scott on April 2, 2020 in a radio interview.	false	low-truth	left
1024	"Republicans have shown themselves willing to cut millions off their health insurance and eliminate preexisting condition protections for millions more, even in the middle of this public health crisis." This statement was made by Barack Obama on April 14, 2020 in his online video endorsement of Joe Biden for president.	true	high-truth	left

Table A.5: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1025	"Some states, like Montana and Nebraska, are getting more than $300,000 in federal stimulus money per reported COVID-19 case. New York is the hardest-hit state and yet we are getting only about $12,000 per case." This statement was made by Andrew Cuomo on April 12, 2020 in a tweet.	mostly-true	high-truth	left
1026	"The Trump Administration promised 27 million tests by the end of March. As of now, only 4 million have been completed." This statement was made by Joe Biden on April 21, 2020 in a tweet.	half-true	high-truth	left
1027	"President Donald Trump's actions on the coronavirus: No. 1, he fired the pandemic team two years ago. No. 2, he's been defunding the Centers for Disease Control." This statement was made by Michael Bloomberg on February 26, 2020 in a CNN town hall.	half-true	high-truth	left
1028	"45 nations had already moved to enforce travel restrictions with China before the president moved." This statement was made by Joe Biden on April 5, 2020 in an ABC "This Week" interview.	half-true	high-truth	left
1029	"Dr. Anthony Fauci has known for 15 years that chloroquine and hydroxychloroquine will not only treat a current case of coronavirus but prevent future cases." This statement was made by a blogger on April 27, 2020 in an article.	false	low-truth	right
1030	"President Barack Obama signed the medical appliance tax bill that forced companies to outsource manufacturing of masks, gowns, gloves and ventilators to China, Europe and Russia to avoid the tax." This statement was made by a Facebook user on April 23, 2020 in a viral image circulating on social media attributing comments to Barack Obama.	pants-fire	low-truth	right

Table A.6: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1031	"Africans living in China now being forced to sleep outside in the cold as Chinese nationals blame them for the rising number of new coronavirus cases in the country." This statement was made by a Facebook user on April 16, 2020 in a post on Facebook.	true	high-truth	right
1032	"A California surfer was alone, in the ocean, when he was arrested for violating the state's stay-at-home order." This statement was made in a viral image on April 8, 2020 in a Facebook post.	mostly-true	high-truth	right
1033	"Herd immunity is probably why California has far fewer COVID-19 deaths than New York." This statement was made by a Facebook user on April 10, 2020 in a Facebook post.	false	low-truth	neutral
1034	"The Asian, Hong Kong, swine and bird flus each killed more people than coronavirus." This statement was made by a Facebook user on April 6, 2020 in a Facebook post.	false	low-truth	neutral
1035	"COVID-19 is here to stay and we need to accept that and be prepared to deal with COVID long term." This statement was made by a Facebook user on April 26, 2020 in a Facebook post.	mostly-true	high-truth	neutral
1036	"Hospitals get paid more to list patients as COVID-19." This statement was made by a Facebook user on April 10, 2020 in a Facebook post.	half-true	high-truth	neutral

Table A.7: Coronavirus / COVID-19 Statement Table

ID	Statement	PolitiFact Rating	Grouped Rating	Partisanship
1037	"Unemployment now pays $24/hour, even if your wages were lower. Why don't essential people forced to still work get $24, too?" This statement was made by a Facebook user on March 29, 2020 in a Facebook post.	half-true	high-truth	neutral
1038	"Mike Pence was caught on a hot mic delivering empty boxes of PPE to a nursing home and pretended they were heavy." This statement was made by a blogger on May 8, 2020 in a blog post.	false	low-truth	left
1039	"Trump said hundreds of governors are calling him & we only have 50." This statement was made by a Facebook user on April 21, 2020 in a Facebook post.	false	low-truth	left
1040	"Donald Trump would receive $17 million for three hotels closed for four days under Republican bill! How in the hell is this right?!" This statement was made by a Facebook user on March 23, 2020 in a Facebook post.	false	low-truth	left
1041	"The CDC issued its first warning on Jan 8. Trump held campaign rallies on Jan 9, Jan 14, Jan 28, Jan 30, Feb 10, Feb 19, Feb 20, Feb 21, & Feb 28. He golfed on Jan 18, Jan 19, Feb 1, Feb 15, Mar 7, Mar 8. The first time he admitted the coronavirus might be a problem was Mar 13." This statement was made by a Facebook user on March 31, 2020 in a Facebook post.	half-true	high-truth	left
1042	"On February 7, the WHO warned about the limited stock of PPE. That same day, the Trump administration announced it was sending 18 tons of masks, gowns and respirators to China." This statement was made by a Facebook user on March 31, 2020 in a Facebook post.	true	high-truth	left

A.2 Descriptive Statistics

We report various summary statistics about respondent characteristics in the tables below.

Percentile	5%	25%	50%	75%	95%
Birth Year	1953	1974	1983	1992	2000

Table A.8: *Birth Year.* The median respondent birth year is 1983. The first quartile is 1974, and the third quartile is 1992. The 5th percentile is 1953, and the 95 percentile is 2000.

Gender	Male	Female
Number	1170	1446

Table A.9: *Gender.* The respondent sample has a roughly even distribution in terms of gender. Females make a slight majority in the sample.

Education Level	High School Unfinished	High School Graduate	Some College	Bachelor's Degree	Graduate school
Number	69	523	694	886	448

Table A.10: *Education Level.* A plurality of the respondents have a bachelor's degree. The respondent sample has good representation for high school graduates, people with some college education, and people with graduate school education.

Race	White	Black	Hispanic	Asian	Native American	Middle Eastern	Mixed	Other
Number	1806	361	222	111	28	10	69	12

Table A.11: *Race.* The vast majority of the respondents are white. The respondent sample also has good representation for Black and Hispanic people.

Ideology	Very Liberal	Liberal	Somewhat Liberal	Middle of the Road	Somewhat Conservative	Conservative	Very Conservative
Number	363	393	263	791	260	285	265

Table A.12: *Ideology.* A plurality of the respondents self-identify with the middle-of-the-road ideology. The distribution of liberals and conservatives are roughly even. There are more liberals than conservatives in the respondent sample.

Partisanship	Strong Democrat	Not very strong Democrat	Lean Democrat	Independent	Lean Republican	Not very strong Republican	Strong Republican	Not Sure
Number	617	297	155	381	127	265	607	148

Table A.13: *Partisanship.* The distribution of liberals and conservatives are roughly even.

Percentile	5%	25%	50%	75%	95%
Response Time in Seconds	192	360	591	911	2023

Table A.14: *Percentiles of Response Time in Seconds.* More than half of the respondents spent more than 10 minutes on the survey. A quarter of the respondents spent between 6 and 10 minutes. The other quarter of the respondents spent less than 4 minutes.

APPENDIX B

Voting Coalition Membership Tables and Polarizing Resolution Tables

B.1 Voting Coalition Membership Tables

We report the estimated voting coalition memberships from 1992 to 2017 in Table B.1 through Table B.26 .

Table B.1: Coalition Membership in 1992

Coalition	Member
1	United States of America, Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Liechtenstein, Portugal, Germany, Austria, Hungary, Czechoslovakia, Italy, Russian Federation, Estonia, Latvia, Lithuania, Finland, Sweden, Norway, Denmark, Iceland, Israel, Japan, Australia, New Zealand, Marshall Islands
2	Cuba, Colombia, Guyana, Cyprus, Mauritania, Niger, Côte D'Ivoire, Liberia, Cameroon, Nigeria, Gabon, Angola, Zimbabwe, Namibia, Libya, Sudan, Iran (Islamic Republic of), Syrian Arab Republic, Jordan, Saudi Arabia, Afghanistan, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Thailand, Lao People's Democratic Republic, Viet Nam, Malaysia, Brunei Darussalam, Philippines, Indonesia
3	Jamaica, Barbados, Saint Lucia, Saint Vincent and the Grenadines, Antigua and Barbuda, Mexico, Belize, Guatemala, Honduras, El Salvador, Costa Rica, Venezuela, Bolivarian Republic of, Suriname, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Gambia (Islamic Republic of the), Mali, Benin, Togo, Rwanda, Botswana, Swaziland, Mauritius, Algeria, Nepal, Singapore
4	Panama, Spain, Malta, Greece, Ukraine, Belarus, Azerbaijan, Turkey, Kazakhstan, Republic of Korea, Samoa

Table B.2: Coalition Membership in 1993

Coalition	Member
1	United States of America, Canada, Dominican Republic, Argentina, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Liechtenstein, Portugal, Germany, Poland, Austria, Czech Republic, Slovakia, Italy, The former Yugoslav Republic of Macedonia, Slovenia, Bulgaria, Russian Federation, Latvia, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Israel, Japan, Australia, New Zealand, Marshall Islands, Micronesia (Federated States of)
2	Bahamas, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Saint Vincent and the Grenadines, Antigua and Barbuda, Mexico, Belize, Honduras, Nicaragua, Costa Rica, Colombia, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Chile, Uruguay, Cyprus, Benin, Ghana, Rwanda, Botswana, Swaziland, Mauritius, Algeria, Saudi Arabia, Nepal, Singapore
3	Cuba, Guinea Bissau, Mali, Mauritania, Niger, Côte D'Ivoire, Sierra Leone, Togo, Nigeria, Kenya, Mozambique, Zimbabwe, Namibia, Lesotho, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Syrian Arab Republic, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Thailand, Lao People's Democratic Republic, Viet Nam, Malaysia, Brunei Darussalam, Philippines, Indonesia
4	Panama, Spain, Malta, Greece, Ukraine, Belarus, Armenia, Azerbaijan, Kazakhstan, Republic of Korea

Table B.3: Coalition Membership in 1994

Coalition	Member
1	United States of America, Canada, Argentina, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Liechtenstein, Spain, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Italy, The former Yugoslav Republic of Macedonia, Slovenia, Greece, Bulgaria, Estonia, Finland, Sweden, Norway, Denmark, Iceland, Israel, Japan, Australia, New Zealand, Marshall Islands
2	Bahamas, Haiti, Jamaica, Trinidad and Tobago, Grenada, Antigua and Barbuda, Saint Kitts and Nevis, Mexico, Belize, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Cabo Verde, Zambia, South Africa, Botswana, Mauritius, Algeria, Fiji
3	Panama, Colombia, Cyprus, Mali, Benin, Mauritania, Niger, Côte D'Ivoire, Ghana, Togo, Cameroon, Nigeria, Gabon, Congo, Uganda, Kenya, United Republic of Tanzania, Ethiopia, Angola, Mozambique, Zimbabwe, Namibia, Swaziland, Tunisia, Libya, Iran (Islamic Republic of), Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Bahrain, Kyrgyzstan, China, India, Pakistan, Bangladesh, Sri Lanka, Maldives, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	Uruguay, Malta, Ukraine, Belarus, Kazakhstan, Republic of Korea
5	Republic of Moldova, Russian Federation, Tajikistan

Table B.4: Coalition Membership in 1995

Coalition	Member
1	United States of America, Canada, Ecuador, Argentina, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Portugal, Germany, Austria, Czech Republic, Slovakia, Italy, Slovenia, Greece, Bulgaria, Republic of Moldova, Romania, Russian Federation, Lithuania, Finland, Sweden, Denmark, Iceland, Israel, Kazakhstan, Japan, Australia, New Zealand, Solomon Islands, Samoa
2	Bahamas, Haiti, Jamaica, Trinidad and Tobago, Antigua and Barbuda, Mexico, Belize, Honduras, El Salvador, Panama, Venezuela, Bolivarian Republic of, Guyana, Suriname, Peru, Brazil, Bolivia (Plurinational State of), Uruguay, Malta, Cyprus, Ukraine, Belarus, South Africa, Botswana, Mauritius, Mongolia, Singapore
3	Colombia, Mali, Benin, Niger, Ghana, Togo, Cameroon, Nigeria, Gabon, Uganda, Kenya, Zimbabwe, Swaziland, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Lebanon, Jordan, Saudi Arabia, China, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Malaysia, Brunei Darussalam, Philippines, Indonesia, Papua New Guinea, Fiji

Table B.5: Coalition Membership in 1996

Coalition	Member
1	United States of America, Israel, Marshall Islands, Micronesia (Federated States of)
2	Canada, Argentina, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Austria, Hungary, Czech Republic, Slovakia, Italy, The former Yugoslav Republic of Macedonia, Slovenia, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Finland, Sweden, Norway, Denmark, Iceland, Japan, Australia, New Zealand, Samoa
3	Bahamas, Jamaica, Trinidad and Tobago, Saint Lucia, Antigua and Barbuda, Mexico, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Uruguay, Malta, Zambia, South Africa, Botswana, Turkey, Mongolia, Vanuatu, Solomon Islands
4	Panama, Colombia, Cabo Verde, Guinea Bissau, Senegal, Benin, Mauritania, Niger, Burkina Faso, Ghana, Togo, Cameroon, Nigeria, Uganda, Kenya, United Republic of Tanzania, Burundi, Ethiopia, Zimbabwe, Algeria, Tunisia, Sudan, Syrian Arab Republic, Lebanon, Jordan, China, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Cambodia, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
5	San Marino, Cyprus, Russian Federation, Belarus, Kazakhstan, Republic of Korea, Fiji

Table B.6: Coalition Membership in 1997

Coalition	Member
1	United States of America, Bulgaria, Israel, Marshall Islands, Micronesia (Federated States of)
2	Canada, Argentina, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Croatia, Slovenia, Greece, Republic of Moldova, Romania, Estonia, Lithuania, Armenia, Finland, Sweden, Norway, Denmark, Iceland, Japan, Australia, New Zealand
3	Bahamas, Jamaica, Trinidad and Tobago, Barbados, Antigua and Barbuda, Mexico, Belize, Guatemala, Honduras, El Salvador, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Uruguay, Malta, Russian Federation, Ethiopia, South Africa, Botswana, Mongolia, Vanuatu, Samoa
4	Cuba, Grenada, Panama, Colombia, Suriname, Cabo Verde, Guinea Bissau, Mali, Senegal, Benin, Mauritania, Niger, Côte D'Ivoire, Guinea, Burkina Faso, Liberia, Sierra Leone, Ghana, Togo, Cameroon, Nigeria, Uganda, Kenya, United Republic of Tanzania, Angola, Mozambique, Zimbabwe, Namibia, Swaziland, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Kuwait, United Arab Emirates, China, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Lao People's Democratic Republic, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia, Papua New Guinea
5	The former Yugoslav Republic of Macedonia, Cyprus, Ukraine, Belarus, Kyrgyzstan, Kazakhstan, Republic of Korea, Solomon Islands

Table B.7: Coalition Membership in 1998

Coalition	Member
1	United States of America, Israel, Marshall Islands
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Slovakia, Italy, San Marino, Albania, Croatia, Slovenia, Greece, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Finland, Sweden, Norway, Denmark, Iceland, Kyrgyzstan, Republic of Korea, Japan, Australia, New Zealand
3	Bahamas, Jamaica, Trinidad and Tobago, Barbados, Antigua and Barbuda, Guatemala, Panama, Venezuela, Bolivarian Republic of, Guyana, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Uruguay, Malta, The former Yugoslav Republic of Macedonia, Cyprus, Russian Federation, Ukraine, Senegal, Ethiopia, Botswana, Saudi Arabia, Vanuatu, Samoa
4	Saint Lucia, Mexico, Colombia, Suriname, Belarus, Cabo Verde, Guinea Bissau, Equatorial Guinea, Mali, Benin, Mauritania, Niger, Côte D'Ivoire, Sierra Leone, Ghana, Togo, Cameroon, Nigeria, Gabon, Chad, Djibouti, Eritrea, South Africa, Namibia, Algeria, Tunisia, Libya, Sudan, Egypt, Lebanon, Jordan, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Lao People's Democratic Republic, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia

Table B.8: Coalition Membership in 1999

Coalition	Member
1	United States of America, Israel, Marshall Islands
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Croatia, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand
3	Bahamas, Haiti, Jamaica, Trinidad and Tobago, Dominica, Grenada, Mexico, El Salvador, Panama, Colombia, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Uruguay, The former Yugoslav Republic of Macedonia, Russian Federation, Ukraine, Belarus, South Africa, Mauritius, Tajikistan, Mongolia, Solomon Islands, Samoa
4	Cuba, Saint Lucia, Suriname, Mali, Benin, Côte D'Ivoire, Guinea, Burkina Faso, Ghana, Togo, Cameroon, Gabon, Kenya, United Republic of Tanzania, Ethiopia, Eritrea, Angola, Namibia, Swaziland, China, Democratic People's Republic of Korea, India, Myanmar, Cambodia, Lao People's Democratic Republic, Singapore
5	Antigua and Barbuda, Cabo Verde, Senegal, Sierra Leone, Nigeria, Djibouti, Mozambique, Zambia, Zimbabwe, Botswana, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Jordan, Saudi Arabia, Kuwait, Bahrain, United Arab Emirates, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Nepal, Thailand, Malaysia, Brunei Darussalam, Philippines, Indonesia, Papua New Guinea

Table B.9: Coalition Membership in 2000

Coalition	Member
1	United States of America, Israel, Marshall Islands, Micronesia (Federated States of)
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, The former Yugoslav Republic of Macedonia, Croatia, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand
3	Cuba, Saint Lucia, Antigua and Barbuda, Belarus, Azerbaijan, Cabo Verde, Senegal, Guinea, Burkina Faso, Ghana, Togo, Nigeria, Chad, Uganda, Kenya, United Republic of Tanzania, Burundi, Djibouti, Ethiopia, Mozambique, Namibia, Botswana, Swaziland, Comoros, Algeria, Tunisia, Libya, Sudan, Egypt, Jordan, Qatar, United Arab Emirates, China, India, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Lao People's Democratic Republic, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	Dominican Republic, Barbados, Grenada, Mexico, Guatemala, Honduras, Nicaragua, Panama, Colombia, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Uruguay, Russian Federation, South Africa, Mauritius, Papua New Guinea, Solomon Islands, Fiji, Samoa

Table B.10: Coalition Membership in 2001

Coalition	Member
1	United States of America, Israel
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, The former Yugoslav Republic of Macedonia, Croatia, Slovenia, Greece, Bulgaria, Republic of Moldova, Romania, Latvia, Lithuania, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Japan, Australia, New Zealand
3	Cuba, Haiti, Saint Lucia, Belarus, Cabo Verde, Equatorial Guinea, Mali, Senegal, Benin, Mauritania, Burkina Faso, Sierra Leone, Ghana, Togo, Nigeria, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Angola, Mozambique, Zambia, South Africa, Namibia, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Kuwait, Bahrain, Qatar, United Arab Emirates, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Cambodia, Lao People's Democratic Republic, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Grenada, Mexico, Belize, Guatemala, Panama, Colombia, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Uruguay, Russian Federation, Kazakhstan, Maldives
5	Chile, Argentina, Malta, Cyprus, Ukraine, Armenia, Republic of Korea

Table B.11: Coalition Membership in 2002

Coalition	Member
1	United States of America, Israel, Nauru, Marshall Islands
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Turkey, Republic of Korea, Japan, Australia, New Zealand
3	Bahamas, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Grenada, Antigua and Barbuda, Mexico, Belize, Guatemala, Honduras, El Salvador, Costa Rica, Panama, Colombia, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Uruguay, Armenia, Senegal, Mauritius, Kuwait
4	Cuba, Haiti, Saint Lucia, Belarus, Azerbaijan, Cabo Verde, Sao Tome and Principe, Mali, Mauritania, Côte D'Ivoire, Guinea, Burkina Faso, Ghana, Togo, Nigeria, Congo, Uganda, Kenya, United Republic of Tanzania, Somalia, Djibouti, Ethiopia, Eritrea, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Botswana, Morocco, Algeria, Tunisia, Libya, Sudan, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Bahrain, Qatar, United Arab Emirates, Oman, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Thailand, Cambodia, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
5	Saint Vincent and the Grenadines, Peru, Argentina, Russian Federation, Kazakhstan, Papua New Guinea, Solomon Islands, Fiji

Table B.12: Coalition Membership in 2003

Coalition	Member
1	United States of America, Israel, Australia
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, New Zealand
3	Bahamas, Dominican Republic, Mexico, Guatemala, Honduras, El Salvador, Nicaragua, Costa Rica, Panama, Colombia, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Uruguay, Russian Federation, Rwanda, Kazakhstan, Thailand, Fiji, Nauru, Samoa
4	Cuba, Jamaica, Trinidad and Tobago, Barbados, Grenada, Saint Lucia, Antigua and Barbuda, Belize, Venezuela, Bolivarian Republic of, Guyana, Suriname, Belarus, Azerbaijan, Guinea Bissau, Mali, Senegal, Mauritania, Niger, Côte D'Ivoire, Guinea, Burkina Faso, Ghana, Togo, Nigeria, Kenya, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Botswana, Madagascar, Mauritius, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, China, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Maldives, Nepal, Cambodia, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia

Table B.13: Coalition Membership in 2004

1	United States of America, Canada, Israel, Australia, Nauru, Marshall Islands
2	Bahamas, Cuba, Jamaica, Trinidad and Tobago, Barbados, Dominica, Saint Lucia, Saint Vincent and the Grenadines, Belize, Colombia, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Belarus, Azerbaijan, Cabo Verde, Mali, Senegal, Benin, Mauritania, Côte D'Ivoire, Guinea, Burkina Faso, Sierra Leone, Ghana, Togo, Nigeria, United Republic of Tanzania, Somalia, Djibouti, Ethiopia, Eritrea, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Madagascar, Mauritius, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Turkmenistan, Kyrgyzstan, Uzbekistan, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Cambodia, Lao People's Democratic Republic, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
3	Haiti, Dominican Republic, Grenada, Guatemala, Honduras, Nicaragua, Costa Rica, Peru, Uruguay, Russian Federation, Cameroon, Uganda, Kenya, Thailand, Papua New Guinea, Samoa
4	Mexico, El Salvador, Panama, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Armenia, Guinea Bissau, Fiji
5	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, New Zealand

Table B.14: Coalition Membership in 2005

Coalition	Member
1	United States of America, Canada, Israel, Australia, Palau
2	Bahamas, Cuba, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Antigua and Barbuda, Belize, Panama, Colombia, Venezuela, Bolivarian Republic of, Guyana, Suriname, Russian Federation, Azerbaijan, Cabo Verde, Mali, Senegal, Benin, Guinea, Burkina Faso, Ghana, Togo, Cameroon, Nigeria, Uganda, Kenya, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Botswana, Mauritius, Morocco, Algeria, Tunisia, Libya, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Kyrgyzstan, Uzbekistan, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Nepal, Thailand, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
3	Haiti, Dominican Republic, Saint Vincent and the Grenadines, Mexico, Guatemala, Honduras, El Salvador, Nicaragua, Costa Rica, Ecuador, Peru, Brazil, Bolivia (Plurinational State of), Paraguay, Chile, Argentina, Uruguay, Armenia, Mauritania, Vanuatu, Fiji
4	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, The former Yugoslav Republic of Macedonia, Croatia, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, New Zealand, Micronesia (Federated States of), Samoa

Table B.15: Coalition Membership in 2006

Coalition	Member
1	United States of America, Canada, Israel, Australia, Micronesia (Federated States of)
2	Bahamas, Mexico, Guatemala, Honduras, El Salvador, Nicaragua, Panama, Ecuador, Peru, Brazil, Paraguay, Chile, Argentina, Uruguay, Armenia, Samoa
3	Cuba, Jamaica, Trinidad and Tobago, Barbados, Antigua and Barbuda, Belize, Costa Rica, Colombia, Venezuela, Bolivarian Republic of, Guyana, Suriname, Bolivia (Plurinational State of), Russian Federation, Belarus, Mali, Senegal, Benin, Mauritania, Niger, Guinea, Sierra Leone, Ghana, Nigeria, Congo, Djibouti, Ethiopia, Mozambique, Zambia, South Africa, Namibia, Comoros, Mauritius, Morocco, Algeria, Libya, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Thailand, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	Dominican Republic, Burundi, Malawi, Fiji
5	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Turkey, Republic of Korea, Japan, New Zealand

Table B.16: Coalition Membership in 2007

Coalition	Member
1	United States of America, Canada, Israel, Australia, Marshall Islands, Palau
2	Bahamas, Mexico, Guatemala, Honduras, El Salvador, Costa Rica, Panama, Colombia, Peru, Brazil, Chile, Argentina, Uruguay, Russian Federation, Armenia, Kazakhstan, Thailand
3	Cuba, Haiti, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Antigua and Barbuda, Belize, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Belarus, Mali, Senegal, Benin, Mauritania, Niger, Guinea, Burkina Faso, Ghana, Togo, Nigeria, Congo, Uganda, United Republic of Tanzania, Djibouti, Eritrea, Mozambique, Zambia, Zimbabwe, South Africa, Lesotho, Botswana, Mauritius, Morocco, Algeria, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Kyrgyzstan, Uzbekistan, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Nepal, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Czech Republic, Slovakia, Italy, San Marino, Malta, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, New Zealand

Table B.17: Coalition Membership in 2008

Coalition	Member
1	United States of America, Israel, Marshall Islands, Palau
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand
3	Cuba, Dominican Republic, Jamaica, Barbados, Grenada, Antigua and Barbuda, Nicaragua, Venezuela, Bolivarian Republic of, Ecuador, Bolivia (Plurinational State of), Belarus, Azerbaijan, Guinea Bissau, Mali, Senegal, Mauritania, Ghana, Togo, Congo, Kenya, Ethiopia, Eritrea, Angola, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Swaziland, Mauritius, Morocco, Algeria, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Kyrgyzstan, Uzbekistan, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Myanmar, Sri Lanka, Maldives, Nepal, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia
4	Mexico, Honduras, Costa Rica, Panama, Brazil, Paraguay, Chile, Argentina, Uruguay, Armenia, Botswana, Turkey, Kazakhstan, Timor-Leste
5	Guatemala, Colombia, Russian Federation, Côte D'Ivoire, Cameroon, Thailand

Table B.18: Coalition Membership in 2009

Coalition	Member
1	United States of America, Israel
2	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand
3	Bahamas, Cuba, Dominican Republic, Jamaica, Trinidad and Tobago, Dominica, Saint Lucia, Saint Vincent and the Grenadines, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Paraguay, Belarus, Azerbaijan, Mali, Senegal, Benin, Mauritania, Niger, Ghana, Nigeria, Kenya, Ethiopia, Zimbabwe, South Africa, Mauritius, Morocco, Algeria, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Kyrgyzstan, Kazakhstan, China, Democratic People's Republic of Korea, India, Pakistan, Bangladesh, Sri Lanka, Maldives, Nepal, Thailand, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Indonesia
4	Mexico, Guatemala, El Salvador, Costa Rica, Panama, Colombia, Peru, Argentina, Uruguay, Russian Federation, Côte D'Ivoire, Liberia, Burundi, Botswana
5	Chile, Turkey

Table B.19: Coalition Membership in 2010

Coalition	Member
1	United States of America, Canada, Israel, Marshall Islands, Palau
2	Bahamas, Mexico, Guatemala, Honduras, El Salvador, Costa Rica, Panama, Colombia, Peru, Paraguay, Chile, Argentina, Uruguay, Papua New Guinea, Fiji
3	Cuba, Haiti, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Saint Vincent and the Grenadines, Belize, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Greece, Russian Federation, Belarus, Azerbaijan, Guinea Bissau, Gambia (Islamic Republic of the), Mali, Senegal, Mauritania, Niger, Guinea, Liberia, Ghana, Togo, Nigeria, Uganda, Kenya, United Republic of Tanzania, Somalia, Djibouti, Eritrea, Mozambique, Zambia, Zimbabwe, Malawi, South Africa, Namibia, Lesotho, Botswana, Swaziland, Comoros, Morocco, Algeria, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Uzbekistan, Kazakhstan, China, Democratic People's Republic of Korea, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Maldives, Nepal, Thailand, Cambodia, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia, Solomon Islands
4	Grenada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand, Samoa

Table B.20: Coalition Membership in 2011

Coalition	Member
1	United States of America, Canada, Israel, Palau
2	Bahamas, Haiti, Jamaica, Barbados, Saint Lucia, Mexico, Belize, Guatemala, Honduras, El Salvador, Costa Rica, Colombia, Peru, Chile, Argentina, Uruguay, Benin, Côte D'Ivoire, Liberia, Botswana, Mauritius, Tunisia, Libya, Kazakhstan, Maldives, Thailand, Solomon Islands, Fiji
3	Cuba, Dominican Republic, Trinidad and Tobago, Saint Vincent and the Grenadines, Antigua and Barbuda, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Russian Federation, Belarus, Mali, Senegal, Guinea, Burkina Faso, Ghana, Uganda, Kenya, Djibouti, Ethiopia, Zambia, South Africa, Morocco, Algeria, Sudan, Iran (Islamic Republic of), Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Kyrgyzstan, Uzbekistan, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Nepal, Cambodia, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Philippines, Indonesia
4	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand
5	Armenia, Samoa

Table B.21: Coalition Membership in 2012

Coalition	Member
1	United States of America, Canada, Israel
2	Bahamas, Cuba, Haiti, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Grenada, Saint Lucia, Saint Vincent and the Grenadines, Antigua and Barbuda, Belize, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Uruguay, Cabo Verde, Mali, Benin, Mauritania, Niger, Côte D'Ivoire, Guinea, Nigeria, Congo, Uganda, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Angola, Zambia, Zimbabwe, South Africa, Namibia, Lesotho, Botswana, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Turkmenistan, Tajikistan, Kyrgyzstan, Uzbekistan, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Nepal, Thailand, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia, Solomon Islands, Tuvalu
3	Mexico, Guatemala, El Salvador, Costa Rica, Colombia, Peru, Paraguay, Chile, Argentina, Russian Federation, Armenia, Liberia, Togo, Burundi, Turkey, Fiji
4	Honduras, Panama, Cameroon, Papua New Guinea, Vanuatu
5	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, Australia, New Zealand, Samoa

Table B.22: Coalition Membership in 2013

Coalition	Member
1	United States of America, Canada, Israel, Australia, Palau
2	Bahamas, Mexico, Guatemala, Costa Rica, Colombia, Peru, Chile, Argentina, Armenia, Turkey, Samoa
3	Cuba, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Saint Vincent and the Grenadines, Antigua and Barbuda, Belize, El Salvador, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Brazil, Bolivia (Plurinational State of), Uruguay, Russian Federation, Cabo Verde, Sao Tome and Principe, Guinea Bissau, Gambia (Islamic Republic of the), Mali, Senegal, Benin, Mauritania, Niger, Côte D'Ivoire, Guinea, Liberia, Togo, Nigeria, Congo, Uganda, United Republic of Tanzania, Somalia, Djibouti, Ethiopia, Eritrea, Angola, Mozambique, Zimbabwe, South Africa, Namibia, Lesotho, Swaziland, Comoros, Mauritius, Morocco, Algeria, Tunisia, Libya, Sudan, Iran (Islamic Republic of), Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Turkmenistan, Kyrgyzstan, Uzbekistan, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Sri Lanka, Maldives, Nepal, Thailand, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia, Solomon Islands, Fiji
4	Honduras, Panama, Paraguay, Cameroon, Papua New Guinea
5	United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Republic of Korea, Japan, New Zealand

Table B.23: Coalition Membership in 2014

Coalition	Member
1	Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Albania, Montenegro, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Finland, Sweden, Norway, Denmark, Iceland, Israel, Japan, Australia, New Zealand, Palau
2	Bahamas, Cuba, Dominican Republic, Jamaica, Trinidad and Tobago, Barbados, Saint Lucia, Antigua and Barbuda, Belize, El Salvador, Nicaragua, Venezuela, Bolivarian Republic of, Guyana, Suriname, Ecuador, Brazil, Chile, Uruguay, Mali, Senegal, Niger, Guinea, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Angola, Mozambique, Zimbabwe, South Africa, Namibia, Seychelles, Morocco, Algeria, Tunisia, Libya, Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Kyrgyzstan, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Maldives, Nepal, Thailand, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Philippines, Indonesia, Solomon Islands, Tuvalu
3	Mexico, Guatemala, Honduras, Costa Rica, Panama, Colombia, Peru, Paraguay, Argentina, Russian Federation, Armenia, Liberia, Togo, Cameroon, Central African Republic, Papua New Guinea

Table B.24: Coalition Membership in 2015

Coalition	Member
1	United States of America, Canada, United Kingdom of Great Britain and Northern Ireland, Ireland, Netherlands, Belgium, Luxembourg, France, Monaco, Liechtenstein, Switzerland, Spain, Andorra, Portugal, Germany, Poland, Austria, Hungary, Czech Republic, Slovakia, Italy, San Marino, Malta, Albania, Montenegro, The former Yugoslav Republic of Macedonia, Croatia, Bosnia and Herzegovina, Slovenia, Greece, Cyprus, Bulgaria, Republic of Moldova, Romania, Estonia, Latvia, Lithuania, Ukraine, Georgia, Finland, Sweden, Norway, Denmark, Iceland, Turkey, Israel, Republic of Korea, Japan, Australia, New Zealand, Palau
2	Cuba, Dominican Republic, Jamaica, Trinidad and Tobago, Grenada, Saint Lucia, Saint Vincent and the Grenadines, Mexico, Belize, El Salvador, Nicaragua, Costa Rica, Guyana, Ecuador, Brazil, Bolivia (Plurinational State of), Chile, Argentina, Uruguay, Cabo Verde, Guinea Bissau, Mali, Senegal, Mauritania, Niger, Guinea, Sierra Leone, Nigeria, Gabon, Chad, Congo, Uganda, Kenya, United Republic of Tanzania, Djibouti, Ethiopia, Eritrea, Angola, Mozambique, Zambia, Zimbabwe, South Africa, Namibia, Botswana, Mauritius, Morocco, Algeria, Tunisia, Libya, Sudan, Iraq, Egypt, Syrian Arab Republic, Lebanon, Jordan, Saudi Arabia, Yemen, Kuwait, Bahrain, Qatar, United Arab Emirates, Oman, Afghanistan, Tajikistan, Kyrgyzstan, Uzbekistan, Kazakhstan, China, India, Bhutan, Pakistan, Bangladesh, Myanmar, Sri Lanka, Maldives, Nepal, Thailand, Cambodia, Lao People's Democratic Republic, Viet Nam, Malaysia, Singapore, Brunei Darussalam, Philippines, Indonesia, Papua New Guinea, Solomon Islands, Fiji
3	Guatemala, Honduras, Panama, Colombia, Peru, Paraguay, Russian Federation, Armenia, Liberia, Togo, Cameroon

Table B.25: Coalition Membership in 2016

Coalition	Member
1	Albania, Andorra, Australia, Austria, Belgium, Bosnia and Herzegovina, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Israel, Italy, Japan, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, New Zealand, Norway, Palau, Poland, Portugal, Republic of Moldova, Romania, San Marino, Slovakia, Slovenia, Spain, Sweden, Switzerland, The former Yugoslav Republic of Macedonia, Ukraine, United Kingdom of Great Britain and Northern Ireland, United States of America
2	Algeria, Antigua and Barbuda, Argentina, Bahamas, Bahrain, Bangladesh, Barbados, Benin, Bhutan, Bolivia (Plurinational State of), Brazil, Brunei Darussalam, Burundi, Chile, China, Cuba, Dominican Republic, Ecuador, Egypt, El Salvador, Ethiopia, Fiji, Guyana, India, Indonesia, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kuwait, Lao People's Democratic Republic, Lesotho, Libya, Malaysia, Maldives, Mali, Mauritania, Myanmar, Nepal, Nicaragua, Nigeria, Oman, Pakistan, Philippines, Qatar, Saint Vincent and the Grenadines, Saudi Arabia, Singapore, South Africa, Sri Lanka, Sudan, Suriname, Syrian Arab Republic, Tajikistan, Thailand, Trinidad and Tobago, United Arab Emirates, United Republic of Tanzania, Uruguay, Uzbekistan, Venezuela, Bolivarian Republic of, Viet Nam
3	Cameroon, Colombia, Costa Rica, Guatemala, Honduras, Mexico, Panama, Papua New Guinea, Paraguay, Peru, Russian Federation, Togo

Table B.26: Coalition Membership in 2017

Coalition	Member
1	Cameroon, Colombia, Guatemala, Honduras, Mexico, Panama, Paraguay, Peru, Solomon Islands, Togo
2	Australia, Canada, Israel, United States of America
3	Albania, Andorra, Austria, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Greece, Hungary, Iceland, Ireland, Italy, Japan, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, New Zealand, Norway, Poland, Portugal, Republic of Korea, Republic of Moldova, Romania, San Marino, Slovakia, Slovenia, Spain, Sweden, Switzerland, The former Yugoslav Republic of Macedonia, Turkey, Ukraine, United Kingdom of Great Britain and Northern Ireland
4	Algeria, Angola, Argentina, Bahamas, Bahrain, Bangladesh, Bhutan, Bolivia (Plurinational State of), Botswana, Brazil, Brunei Darussalam, Cabo Verde, Cambodia, Chile, China, Congo, Costa Rica, Côte D'Ivoire, Cuba, Djibouti, Dominican Republic, Ecuador, Egypt, El Salvador, Equatorial Guinea, Eritrea, Ethiopia, Gabon, Guinea, Guyana, India, Indonesia, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kuwait, Kyrgyzstan, Lao People's Democratic Republic, Lesotho, Libya, Malaysia, Maldives, Mali, Mauritania, Mauritius, Morocco, Mozambique, Namibia, Nepal, Nicaragua, Oman, Pakistan, Philippines, Qatar, Russian Federation, Saint Kitts and Nevis, Saint Lucia, Saudi Arabia, Sierra Leone, Singapore, South Africa, Sri Lanka, Sudan, Syrian Arab Republic, Tajikistan, Thailand, Trinidad and Tobago, Tunisia, Uganda, United Arab Emirates, United Republic of Tanzania, Uruguay, Uzbekistan, Venezuela, Bolivarian Republic of, Viet Nam, Yemen, Zimbabwe

B.2 Polarizing Resolution Tables

We report the polarizing resolutions from 1992 to 2017 in Table B.27 through Table B.52 .

Table B.27: Polarizing Resolutions in 1992

Count	Resolution Title
1	Programme of work, Committee on Palestinian rights
2	Importance of the universal realization of the rights of peoples to self determination and of the speedy granting of independence to colonial countries and peoples for the effective guarantee and observance of human rights
3	Use of mercernaries as a means to violate human rights and to impede the exercise of the right of peoples to self determination United Nations African Institute for the Prevention of Crime and the Treatment of Offenders
4	Situation of human rights in the Islamic Republic of Iran
5	Encourages all member states to lend assistance to the committee on Palestinian rights
6	Reaffirms the inalienable right of all displaced inhabitants to return to their homes or former places of residence in the territories occupied by Israel since 1967

Table B.28: Polarizing Resolutions in 1993

Count	Resolution Title
1	Importance of the universal realization of the rights of peoples to self determination and of the speedy granting of independence to colonial countries and peoples for the effective guarantee and observance of human rights
2	Committee on the Exercise of the Inalienable Rights of the Palestinian People
3	Division for Palestinian Rights of the Secretariat
4	Use of mercenaries as a means to violate human rights and to impede the exercise of the right of peoples to self determination
5	Situation of human rights in Cuba
6	Alternative approaches and ways and means within the UN system for improving the effective enjoyment of human rights and fundamental freedoms

Table B.29: Polarizing Resolutions in 1994

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Use of mercenaries as a means to violate human rights and to impede the exercise of the right of peoples to self determination
4	Importance of the universal realization of the rights of peoples to self determination and of the speedy granting of independence to colonial countries and peoples for the effective guarantee and observance of human rights
5	Alternative approaches and ways and means within the UN system for improving the effective enjoyment of human rights and fundamental freedoms
6	Situation of human rights in the Islamic Republic of Iran
7	Situation of human rights in Cuba

Table B.30: Polarizing Resolutions in 1995

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Situation of human rights in Cuba
4	Situation of human rights in the Sudan
5	Use of mercenaries as a means to violate human rights and to impede the exercise of the right of peoples to self determination
6	Situation of human rights in the Islamic Republic of Iran
7	Respect for the right to universal freedom of travel and the vital importance of family reunification

Table B.31: Polarizing Resolutions in 1996

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
4	The right of the Palestinian people to self-determination
5	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self determination
6	Enhancement of international cooperation in the field of human rights
7	Report of the Special Committee to Investigate Israeli Practices Affecting the Human Rights of the Palestinian People and Other Arabs of the Occupied Territories, including Jerusalem
8	Situation of human rights in the Islamic Republic of Iran

Table B.32: Polarizing Resolutions in 1997

Count	Resolution Title
1	UN promotion of human rights through international cooperation and the importance of non-selectivity, impartiality and objectivity
2	Decides to confer upon Palestine, in its capacity as observer, and as contained in the annex to the present resolution, additional rights and privileges of participation in the sessions and work of the General Assembly and the international con
3	Committee on the Exercise of the Inalienable Rights of the Palestinian People
4	Division for Palestinian Rights of the Secretariat
5	Human rights situation in the Islamic Republic of Iran
6	Human rights situation in Cuba
7	Human rights situation in Nigeria
8	Human rights and terrorism
9	Use of mercenaries to violate human rights
10	Human rights situation in the Sudan

Table B.33: Polarizing Resolutions in 1998

Count	Resolution Title
1	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
2	Human rights of Palestinians in the occupied territories
3	Right of the Palestinian people to self-determination
4	Right to development
7	Committee on the Exercise of the Inalienable Rights of the Palestinian People
8	Division for Palestinian Rights

Table B.34: Polarizing Resolutions in 1999

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
4	Reaffirms the right of all persons displaced as a result of the June 1967 and subsequent hostilities to return to their homes or former places of residence in the territories occupied by Israel since 1967
5	Israeli Practices Affecting the Human Rights of the Palestinian People in the Occupied Palestinian Territory, including Jerusalem
6	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
7	The right of the Palestinian people to self-determination
8	Human rights and terrorism
9	Human rights and unilateral coercive measures
10	The right of development

Table B.35: Polarizing Resolutions in 2000

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
4	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
5	The right of the Palestinian people to self-determination
6	Situation of human rights in the Islamic Republic of Iran
7	Reaffirms the right of all persons displaced as a result of the June 1967 and subsequent hostilities to return to their homes or former places of residence in the territories occupied by Israel since 1967
8	Israeli Practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including Jerusalem

Table B.36: Polarizing Resolutions in 2001

Count	Resolution Title
1	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
2	Reaffirms the right of all persons displaced as a result of the June 1967 and subsequent hostilities to return to their homes or former places of residence in the territories occupied by Israel since 1967
3	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including Jerusalem
4	The right of Palestinian people to self-determination
5	Equitable geographical distribution of the membership of the human rights treaty bodies
6	Human rights and unilateral coercive measures
7	The right to development
8	The right to food
9	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
10	Comprehensive implementation of and follow-up to the World Conference against Racism, Racial Discrimination, Xenophobia and Related Intolerance

Table B.37: Polarizing Resolutions in 2002

Count	Resolution Title
1	Requests the Department of Public Information of the Secretariat to assist in the worldwide dissemination of accurate and comprehensive information in support for the inalienable rights of the Palestinian people
2	Reaffirms the right of all persons displaced as a result of the June 1967 and subsequent hostilities to return to their homes or former places of residence in the territories occupied by Israel since 1967
3	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
4	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
5	Human rights and unilateral coercive measures
6	Situation of human rights in Iraq
7	Situation of human rights in the Democratic Republic of the Congo
8	The Acting President
9	Globalization and its impact on the full enjoyment of all human rights
10	Situation of human rights in the Sudan
11	Work of the Special Committee to Investigate Israeli Practices Affecting the Human Rights of the Palestinian People and Other Arabs of the Occupied Territories

Table B.38: Polarizing Resolutions in 2003

Count	Resolution Title
1	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
2	The right to development
3	Human rights and unilateral coercive measures
4	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
5	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
6	Committee on the Exercise of the Inalienable Rights of the Palestinian People
7	Division for Palestinian Rights of the Secretariat
9	Situation of human rights in the Democratic Republic of the Congo
10	Globalization and its impact on the full enjoyment of all human rights
12	Promotion of peace as a vital requirement for the full enjoyment of all human rights by all

Table B.39: Polarizing Resolutions in 2004

Count	Resolution Title
1	Situation of human rights in Turkmenistan
2	Human rights and unilateral coercive measures
3	Human rights and terrorism
4	Globalization and its impact on the full enjoyment of all human rights
5	Equitable geographical distribution in the membership of the human rights treaty bodies
6	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
7	Division for Palestinian Rights of the Secretariat
8	Committee on the Exercise of Inalienable Rights of the Palestinian People
9	Situation of human rights in the Democratic Republic of the Congo
10	Respect for the purposes and principles contained in the Charter of the United Nations to achieve international cooperation in promoting and encouraging respect for human rights and for fundamental freedoms and in solving international problems

Table B.40: Polarizing Resolutions in 2005

Count	Resolution Title
1	Situation of human rights in Uzbekistan
2	Situation of human rights in Turkmenistan
3	The right to development
4	Human rights and unilateral coercive measures
5	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
6	Situation of human rights in the Islamic Republic of Iran
7	Globalization and its impact on the full enjoyment of all human rights
8	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
9	Promotion of peace as a vital requirement for the full enjoyment of all human rights by all
10	Respect for the principles of national sovereignty and diversity of democratic systems in electoral processes as an important element for the promotion and protection of human rights
11	Division for Palestinian Rights of the Secretariat

Table B.41: Polarizing Resolutions in 2006

Count	Resolution Title
1	Israeli settlements in the Occupied Palestinian Territories, including East Jerusalem, and the occupied Syrian Golan
2	The occupied Syrian Golan
3	Human rights an unilateral coercive measures
4	Respect for the right of universal freedom of travel and the vital importance of family reunification
5	Composition for the staff of the Office of the UN High Commissioner for Human Rights
6	Globalization and its impact on the full enjoyment of all human rights
7	The right of the palestine people to self-determination
8	Use of mercenaries as a means of violating human rights and impeding the excercise of the right of peoples to self-determination
9	Inadmissibility of certain practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance
11	Situation of human rights in Belarus
12	The right of development
13	The human rights situation arising from the recent Israeli military operations in Lebanon
14	Situation of human rights in the Islamic Republic of Iran

Table B.42: Polarizing Resolutions in 2007

Count	Resolution Title
1	Commitee on the Exercice of Inalienable Rights of the Palestinian People
2	Division for Palestine Rights of the Secretariat
3	Report of the Human Rights Council on the preparations for the Durban Review Conference
4	Human rights and unilateral coercive measures
5	The right of the Palestinian People to self-determination
6	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
7	The right to development
8	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
9	Report of the Human Rights Council
10	Promotion of peace as a vital requirement for the full enhoyment of all human rights by all
11	Globalization and its impact on the full enjoyment of all human rights
12	Respect for the purposes and principles contained in the Charter of the United Nations to achieve internationlal cooperation in promoting and encouraging respect for human rights and fundamental freedoms in solving international problems
13	Work of the Special Committee to Investigate Israeli Practices Affecting the Human Rights of the Palestinian People and Other Arabs of the Occupied Territories

Table B.43: Polarizing Resolutions in 2008

Count	Resolution Title
1	Division for Palestinian Rights of the Secretariat
2	Committee on the Excercise of the Inalienable Rights of the Palestinian People
3	The occupied Syrian Golan
4	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
5	Applicability of the Geneva Convention relative to the Protection of Civilian Persons in Time of War, of 12 August 1249, to the Ocuupied Palestinian Territory, including East Jerusalem, and the other occupied Arab territories
6	Israeli settlements in the Occupied Palestinian Territory, including East Jerusalem, and the occupied Syrian Golan
7	Situation of human rights in the Democratic People's Republic of Korea
8	Promotion of a democratic and equitable international order
9	Extrajudicial, summary or arbitrary executions
10	Human rights and unilateral coercive measures
11	The right to development
12	Situation of human rights in the Islamic Republic of Iran
13	Equitable geographical distribution in the membership of the human rights treaty bodies

Table B.44: Polarizing Resolutions in 2009

Count	Resolution Title
1	Follow-up to the report of the United Nations Fact-Finding Mission on the Gaza Conflict
2	Committee on the Exercise of the Inalienable Rights of the Palestinian People
3	Division for Palestinian Rights of the Secretariat
4	Work of the Special Committee to Investigate Israel Practices Affecting the Human Rights of the Palestinian People and Other Arabs of the Occupied Territories
5	Israeli settlements in the Occupied Palestinian Territory, including East Jerusalem, and the occupied Syrian Golan
6	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
7	Applicability of the Geneva Convention relative to the Protection of Civilian Persons in Time of War, of 12 August 1949, to the Occupied Palestinian Territory, including East Jerusalem, and the other occupied Arab territories
8	The occupied Syrian Golan
9	Inadmissibility of certain practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance
10	Globalization and its impact on the full enjoyment of all human rights
11	Human Rights and unilateral coercive measures
12	The right to development

Table B.45: Polarizing Resolutions in 2010

Count	Resolution Title
1	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
2	The right of the Palestinian people to self-determination
3	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
4	Globalization and its impact on the full enjoyment of all human rights
5	Human rights and unilateral coercive measures
6	The right to development
7	Promotion of a democratic and equitable international order
8	Report of the Human Rights Council
9	Review of the Human Rights Council
10	Committee on the Exercise of the Inalienable Rights of the Palestinian People
11	Division for Palestinian Rights of the Secretariat
12	Inadmissibility of certain practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance

Table B.46: Polarizing Resolutions in 2011

Count	Resolution Title
1	Inadmissibility of certain practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance
2	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
3	The right of the Palestinian people to self-determination
4	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
5	Promotion of equitable geographical distribution in the membership of the human rights treaty bodies
6	Human rights and cultural diversity
7	The right to development
8	Human rights and unilateral coercive measures
9	Promotion of a democratic and equitable international order
10	Globalization and its impact on the full enjoyment of all human rights
11	Applicability of the Geneva Convention relative to the Protection of Civilian Persons in Time of War, of 12 August 1949, to the Occupied PalestinianTerritory, including East Jerusalem, and the other occupied Arab territories
12	Israeli settlements in the Occupied Palestinian Territory, including East Jerusalem, and the occupied Syrian Golan

Table B.47: Polarizing Resolutions in 2012

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Division for Palestinian Rights of the Secretariat
3	Applicability of the Geneva Convention relative to the Protection of Civilian Persons in Time of War, of 12 August 1949, to the Occupied Palestinian Territory, including East Jerusalem, and the other occupied Arab territories
4	Israeli settlements in the Occupied Palestinian Territory, including East Jerusalem, and the occupied Syrian Golan
5	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
6	The occupied Syrian Golan
7	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
8	The right of the Palestinian people to self-determination
9	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self determination
10	Globalization and its impact on the full enjoyment of all human rights
11	Human rights and unilateral coercive measures
12	The right to development

Table B.48: Polarizing Resolutions in 2013

Count	Resolution Title
1	Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem
2	Israeli settlements in the Occupied Palestinian Territory, including East Jerusalem, and the occupied Syrian Golan
3	Applicability of the Geneva Convention relative to the Protection of Civilian Persons in Time of War, of 12 August 1949, to the Occupied Palestinian Territory, including East Jerusalem, and the other occupied Arab territories
4	Promotion of a democratic and equitable international order
5	Globalization and its impact on the full enjoyment of all human rights
6	Promotion of equitable geographical distribution in the membership of the human rights treaty bodies
7	Human rights and unilateral coercive measures
8	Human rights and cultural diversity
9	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
10	Global efforts for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action

Table B.49: Polarizing Resolutions in 2014

Count	Resolution Title
1	A global call for concrete action for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
2	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
3	Globalization and its impact on the full enjoyment of all human rights
4	Promotion of peace as a vital requirement for the full enjoyment of all human rights by all
5	Human rights and unilateral coercive measures
6	Division for Palestinian Rights of the Secretariat
7	Combating glorification of Nazism, neo-Nazism and other practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance
8	Committee on the Exercise of the Inalienable Rights of the Palestinian People
9	Report of the Human Rights Council

Table B.50: Polarizing Resolutions in 2015

Count	Resolution Title
1	Committee on the Exercise of the Inalienable Rights of the Palestinian People
2	Human rights and unilateral coercive measures
3	Promotion of equitable geographical distribution in the membership of the human rights treaty bodies
4	Human rights and cultural diversity
5	Globalization and its impact on the full enjoyment of all human rights
6	Division for Palestinian Rights of the Secretariat
7	Combating glorification of Nazism, neo-Nazism and other practices that contribute to fuelling contemporary forms of racism, racial discrimination, xenophobia and related intolerance
8	A global call for concrete action for the total elimination of racism, racial discrimination, xenophobia and related intolerance and the comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action
9	Use of mercenaries as a means of violating human rights and impeding the exercise of the right of peoples to self-determination
10	Report of the Human Rights Council

Table B.51: Polarizing Resolutions in 2016

Count	Resolution Title
1	A/71/251 66b - Comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action. - RACIAL DISCRIMINATION–PROGRAMME IMPLEMENTATION
2	A/71/251 67 - Right of peoples to self-determination. - SELF-DETERMINATION OF PEOPLES
3	A/71/251 68b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT
4	A/71/251 68b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/71/251 68b[12] - SANCTIONS–INTERNATIONAL RELATIONS
5	A/71/251 68b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/71/251 68b[6] - GLOBALIZATION–HUMAN RIGHTS
6	A/71/251 68b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/71/251 68b[31] - DEMOCRACY
7	A/71/251 66a - Elimination of racism, racial discrimination, xenophobia and related intolerance. - RACIAL DISCRIMINATION–ELIMINATION
8	A/71/251 35 - Question of Palestine. - PALESTINE QUESTION
9	A/71/251 68c - Human rights situations and reports of special rapporteurs and representatives. - HUMAN RIGHTS–REPORTS

Table B.52: Polarizing Resolutions in 2017

Count	Resolution Title
1	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/72/251 72b[6] - GLOBALIZATION–HUMAN RIGHTS
2	A/72/251 67 - Report of the Human Rights Council. - UN. HUMAN RIGHTS COUNCIL–REPORTS
3	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/72/251 72b[12] - SANCTIONS–INTERNATIONAL RELATIONS
4	A/72/251 70b - Comprehensive implementation of and follow-up to the Durban Declaration and Programme of Action. - RACIAL DISCRIMINATION–PROGRAMME IMPLEMENTATION
5	A/72/251 70a - Elimination of racism, racial discrimination, xenophobia and related intolerance. - RACIAL DISCRIMINATION–ELIMINATION
6	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/72/251 72b[11] - RIGHT TO DEVELOPMENT
7	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/72/251 72b[30] - RIGHT TO CULTURE
8	A/72/251 71 - Right of peoples to self-determination. - SELF-DETERMINATION OF PEOPLES
9	A/72/251 54 - Report of the Special Committee to Investigate Israeli Practices Affecting the Human Rights of the Palestinian People and Other Arabs of the Occupied Territories. - TERRITORIES OCCUPIED BY ISRAEL–HUMAN RIGHTS
10	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT
11	A/72/251 38 - Question of Palestine. - PALESTINE QUESTION
12	A/72/251 72b - Human rights questions, including alternative approaches for improving the effective enjoyment of human rights and fundamental freedoms. - HUMAN RIGHTS ADVANCEMENT;A/72/251 72b[31] - DEMOCRACY

APPENDIX C

The Pure Second-Dimension Bills' Summaries

The 40 pure second-dimension bills have more coherent substantive meanings. There are two policy topics that repeatedly show up in the pure-second-dimension bills: appropriations and foreign policy. Bill No. 63, 92, 96, 181, 185, 209, 235, 251, 257, 264, 283, 298, 337, 338, 339, 342, 416, 512, 558, and 612 are all about appropriations or budget issues. Bills that are related to foreign policy include Bill No. 25, 184, 185, 339, 342, 345, 410, and 546. The summaries of the above bills are reported in the tables below.

Bill Number	Bill Summary
63	SENATE AGREED TO A MOTION TO ADVANCE S. 672, MAKING SUPPLEMENTAL APPROPRIATIONS AND RESCISSIONS FOR THE FISCAL YEAR ENDING SEPTEMBER 30,1997, TO THE THIRD READING.
92	SENATE PASSED H. CON. RES. 84, ESTABLISHING THE CONGRESSIONAL BUDGET FOR THE UNITED STATES GOVERNMENT FOR FISCAL YEAR 1998 AND SETTING FORTH APPROPRIATE BUDGETARY LEVELS FOR FISCAL YEARS 1999, 2000, 2001, AND 2002.
96	SENATE AGREED TO THE CONFERENCE REPORT ON H.CON.RES. 84, ESTABLISHING THE CONGRESSIONAL BUDGET FOR THE UNITED STATES GOVERNMENT FOR FISCAL YEAR 1998AND SETTING FORTH APPROPRIATE BUDGETARY LEVELS FOR FISCAL YEARS 1999,2000, 2001, AND 2002.
181	SENATE REJECTED THE ALLARD AMENDMENT NO. 891, TO DECREASE THE AMOUNT OF FUNDS AVAILABLE TO OPIC FOR ADMINISTRATIVE EXPENSES TO CARRY OUT THE CREDIT AND INSURANCE PROGRAMS.
185	SENATE PASSED S. 955, MAKING APPROPRIATIONS FOR FOREIGN OPERATIONS,EXPORT FINANCING, AND RELATED PROGRAMS FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.

Table C.1: Pure-2nd-D Bills on Appropriations (Part 1)

Bill Number	Bill Summary
209	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 2015, TO PROVIDE FOR RECONCILIATION PURSUANT TO SUBSECTIONS (B)(1) AND (C) OF SECTION 105 OF THE CONCURRENT RESOLUTION ON THE BUDGET FOR FISCAL YEAR 1998.
235	SENATE PASSED S. 1061, MAKING APPROPRIATIONS FOR THE DEPARTMENTS OF LABOR, HEALTH, AND HUMAN SERVICES, AND EDUCATION, AND RELATED AGENCIES FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
251	SENATE PASSED H.R. 2107, MAKING APPROPRIATIONS FOR THE DEPARTMENT OF THE INTERIOR AND RELATED AGENCIES FOR THE FISCAL YEAR ENDING SEPTEMBER 30,1998.
257	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 2209, MAKING APPROPRIATIONS FOR THE LEGISLATIVE BRANCH FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
264	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 2378, MAKING APPROPRIATIONS FOR THE TREASURY DEPARTMENT, THE UNITED STATES POSTAL SERVICE, THE EXECUTIVE OFFICE OF THE PRESIDENT, AND CERTAIN INDEPENDENT AGENCIES, FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
283	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 2107, MAKING APPROPRIATIONS FOR THE DEPARTMENT OF THE INTERIOR AND RELATED AGENCIES FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
298	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 2264, MAKING APPROPRIATIONS FOR THE DEPARTMENTS OF LABOR, HEALTH AND HUMAN SERVICES,AND EDUCATION, AND RELATED AGENCIES FOR THE FISCAL YEAR ENDING SEPTEMBER30, 1998.
337	SENATE REJECTED MCCAIN AMENDMENT NO. 2063, TO ELIMINATE CERTAIN SPENDING ITEMS FROM THE BILL.

Table C.2: Pure-2nd-D Bills on Appropriations (Part 2)

Bill Number	Bill Summary
338	SENATE REJECTED GRAMM/SANTORUM AMENDMENT NO. 2104, TO ESTABLISH THAT ONLY THAT PORTION OF BUDGET AUTHORITY PROVIDED IN THIS ACT THAT IS OBLIGATED DURING FISCAL YEAR 1998 SHALL BE DESIGNATED AS AN EMERGENCY REQUIREMENT PURSUANT TO THE BALANCED BUDGET AND EMERGENCY DEFICIT CONTROL ACT OF 1985.
339	SENATE TABLED FEINGOLD AMENDMENT NO. 2121, TO REMOVE THE EMERGENCY DESIGNATION FOR THE SUPPLEMENTAL APPROPRIATIONS TO FUND INCREMENTAL COSTS OF CONTINGENCY OPERATIONS IN BOSNIA.
342	SENATE AGREED TO THE MCCONNELL MODIFIED AMENDMENT NO. 2100, TO PROVIDE SUPPLEMENTAL APPROPRIATIONS FOR THE INTERNATIONAL MONETARY FUND FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
416	SENATE PASSED THE CONFERENCE REPORT ON H.R. 3579, MAKING EMERGENCY SUPPLEMENTAL APPROPRIATIONS FOR RECOVERY FROM NATURAL DISASTERS, AND FOR OVERSEAS PEACEKEEPING EFFORTS, FOR THE FISCAL YEAR ENDING SEPTEMBER 30,1998.
512	SENATE PASSED H.R. 4112, MAKING APPROPRIATIONS FOR THE LEGISLATIVE BRANCH FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1999, AS AMENDED.
558	SENATE PASSED H.R. 4104, MAKING APPROPRIATIONS FOR THE TREASURY DEPARTMENT, THE UNITED STATES POSTAL SERVICE, THE EXECUTIVE OFFICE OF THE PRESIDENT, AND CERTAIN INDEPENDENT AGENCIES, FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1999.
612	SENATE AGREED TO THE CONFERENCE REPORT ON H.R. 4328, THE OMNIBUS CONSOLIDATED AND EMERGENCY SUPPLEMENTAL APPROPRIATIONS ACT, 1999

Table C.3: Pure-2nd-D Bills on Appropriations (Part 3)

Bill Number	Bill Summary
25	SENATE TABLED HOLLINGS AMENDMENT NO. 19, TO REQUIRE CONGRESSIONAL APPROVAL BEFORE ANY INTERNATIONAL TRADE AGREEMENT THAT HAS THE EFFECT OF AMENDING OR REPEALING STATUTORY LAW OF THE UNITED STATES LAW CAN BE IMPLEMENTED IN THE UNITED STATES.
184	SENATE REJECTED THE HUTCHINSON AMENDMENT NO. 890, TO EXPRESS THE SENSE OF THE SENATE THAT MOST FAVORED NATIONS TRADE STATUS FOR CHINA SHOULD BE REVOKED.
185	SENATE PASSED S. 955, MAKING APPROPRIATIONS FOR FOREIGN OPERATIONS,EXPORT FINANCING, AND RELATED PROGRAMS FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
339	SENATE TABLED FEINGOLD AMENDMENT NO. 2121, TO REMOVE THE EMERGENCY DESIGNATION FOR THE SUPPLEMENTAL APPROPRIATIONS TO FUND INCREMENTAL COSTS OF CONTINGENCY OPERATIONS IN BOSNIA.
342	SENATE AGREED TO THE MCCONNELL MODIFIED AMENDMENT NO. 2100, TO PROVIDE SUPPLEMENTAL APPROPRIATIONS FOR THE INTERNATIONAL MONETARY FUND FOR THE FISCAL YEAR ENDING SEPTEMBER 30, 1998.
345	SENATE REJECTED S.J. RES. 42, TO DISAPPROVE THE CERTIFICATION OF THE PRESIDENT UNDER SECTION 490(B) OF THE FOREIGN ASSISTANCE ACT OF 1961REGARDING FOREIGN ASSISTANCE FOR MEXICO DURING THE FISCAL YEAR 1998.
410	SENATE REJECTED THE WARNER AMENDMENT NO. 2322, TO PROVIDE FOR A THREE-YEAR PAUSE IN FURTHER NATO EXPANSION AFTER ADMISSION OF POLAND,HUNGARY, AND THE CZECH REPUBLIC.
546	SENATE FAILED TO TABLE TO THE HUTCHINSON AMENDMENT NO. 3124, TO CONDEMN THOSE OFFICIALS OF THE CHINESE COMMUNIST PARTY, THE GOVERNMENT OF THE PEOPLE'S REPUBLIC OF CHINA, AND OTHER PERSONS WHO ARE INVOLVED IN THE ENFORCEMENT OF FORCED ABORTIONS BY PREVENTING SUCH PERSON FROM ENTERING OF REMAINING IN THE UNITED STATES, AND TO EXPRESS THE SENSE OF THE CONGRESS THAT THE PRESIDENT SHOULD MAKE FREEDOM OF RELIGION ONE OF THE MAJOR OBJECTS OF UNITED STATES FOREIGN POLICY WITH RESPECT TO CHINA.

Table C.4: Pure-2nd-D Bills on Foreign Policy

BIBLIOGRAPHY

Abrajano, Marisa A, Christopher S Elmendorf and Kevin M Quinn. 2018. "Labels vs. pictures: Treatment-mode effects in experiments about discrimination." *Political Analysis* 26(1):20–33.

Ahler, Douglas J and David E Broockman. 2018. "The delegate paradox: Why polarized politicians can represent citizens best." *The Journal of Politics* 80(4):1117–1133.

Albert, James H and Siddhartha Chib. 1993. "Bayesian analysis of binary and polychotomous response data." *Journal of the American Statistical Association* 88(422):669–679.

Albert, James H. and Siddhartha Chib. 1995. "Bayesian residual analysis for binary response regression models." *Biometrika* 82(4):747–759.

Aldous, David J. 1985. Exchangeability and related topics. In *École d'Été de Probabilités de Saint-Flour XIII—1983*. Springer pp. 1–198.

Antoniak, Charles E. 1974. "Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems." *The Annals of Statistics* pp. 1152–1174.

Bachman, Jerald G. and Patrick M. O'Malley. 1984. "Yea-saying, nay-saying, and going to extremes: Black-White differences in response styles." *Public Opinion Quarterly* 48(2):491–509.

Bailey, Michael A. 2013. "Is today's court the most conservative in sixty years? Challenges and opportunities in measuring judicial preferences." *The Journal of Politics* 75(3):821–834.

Bailey, Michael A, Anton Strezhnev and Erik Voeten. 2017. "Estimating dynamic state preferences from United Nations voting data." *Journal of Conflict Resolution* 61(2):430–456.

Bailey, Michael A and Erik Voeten. 2018. "A two-dimensional analysis of seventy years of United Nations voting." *Public Choice* 176(1-2):33–55.

Balakrishnan, Suhrid and Sumit Chopra. 2012. "Two of a kind or the ratings game? Adaptive pairwise preferences and latent factor models." *Frontiers of Computer Science* 6(2):197–208.

Becker, Raphael N, Arye L Hillman, Niklas Potrafke and Alexander H Schwemmer. 2015. "The preoccupation of the United Nations with Israel: Evidence and theory." *The Review of International Organizations* 10(4):413–437.

Benoit, Kenneth, Kevin Munger and Arthur Spirling. 2019. "Measuring and explaining political sophistication through textual complexity." *American Journal of Political Science* 63(2):491–508.

Besant, Alexander and Sébastien Malo. 2009. "Dim Prospects for the United Nations Human Rights Council." *Yale J. Int'l Aff.* 4:144.

Bitto, Angela and Sylvia Frühwirth-Schnatter. 2019. "Achieving shrinkage in a time-varying parameter model framework." *Journal of Econometrics* 210(1):75–97.

Blackwell, David, James B MacQueen et al. 1973. "Ferguson distributions via Pólya urn schemes." *The Annals of Statistics* 1(2):353–355.

Böckenholt, Ulf and Rung-Ching Tsai. 2001. "Individual differences in paired comparison data." *British Journal of Mathematical and Statistical Psychology* 54(2):265–277.

Boockmann, Bernhard and Axel Dreher. 2011. "Do human rights offenders oppose human rights resolutions in the United Nations?" *Public Choice* 146(3-4):443–467.

Bradley, Ralph Allan and Milton E. Terry. 1952. "Rank analysis of incomplete block designs: I. The method of paired comparisons." *Biometrika* 39(3/4).

Brady, Henry E. 1985. "The perils of survey research: Inter-personally incomparable responses." *Political Methodology* 11(June):269–290.

Broockman, David E. 2016. "Approaches to studying policy representation." *Legislative Studies Quarterly* 41(1):181–215.

Carlson, David and Jacob M. Montgomery. 2017. "A pairwise comparison framework for fast, flexible, and reliable human coding of political texts." *American Political Science Review* 111(4):835–843.

Carroll, J Douglas and Geert De Soete. 1991. "Toward a new paradigm for the study of multiattribute choice behavior: Spatial and discrete modeling of pairwise preferences." *American Psychologist* 46(4):342.

Carroll, Royce, Jeffrey B Lewis, James Lo, Keith T 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.

Carter, Chris K and Robert Kohn. 1994. "On Gibbs sampling for state space models." *Biometrika* 81(3):541–553.

Cattelan, Manuela. 2012. "Models for paired comparison data: A review with emphasis on dependent data." *Statistical Science* pp. 412–433.

Chib, Siddhartha and Edward Greenberg. 1995*a*. "Understanding the Metropolis-Hastings algorithm." *The American Statistician* 49(November):327–336.

Chib, Siddhartha and Edward Greenberg. 1995*b*. "Understanding the Metropolis-Hastings algorithm." *The American Statistician* 49(4):327–335.

Choi, Jang, Gary Oehlert and Hui Zou. 2010. "A penalized maximum likelihood approach to sparse factor analysis." *Statistics and its Interface* 3(4):429–436.

Clinton, Joshua, Simon Jackman and Douglas Rivers. 2004. "The statistical analysis of roll call data." *American Political Science Review* 98(2):355–370.

Connor, Robert J and James E Mosimann. 1969. "Concepts of independence for proportions with a generalization of the Dirichlet distribution." *Journal of the American Statistical Association* 64(325):194–206.

Dahl, David B. 2003. "An improved merge-split sampler for conjugate Dirichlet process mixture models." *Technical Report* 1:086.

David, Herbert Aron. 1963. *The Method of Paired Comparisons*. Vol. 12 London.

Dittrich, Regina, Brian Francis, Reinhold Hatzinger and Walter Katzenbeisser. 2007. "A paired comparison approach for the analysis of sets of Likert-scale responses." *Statistical Modeling* 7(1):3–28.

Eberhardt, Jennifer L, Paul G Davies, Valerie J Purdie-Vaughns and Sheri Lynn Johnson. 2006. "Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes." *Psychological Science* 17(5):383–386.

Enelow, James M and Melvin J Hinich. 1984. *The Spatial Theory of Voting: An Introduction*. CUP Archive.

Escobar, Michael D. 1994. "Estimating normal means with a Dirichlet process prior." *Journal of the American Statistical Association* 89(425):268–277.

Escobar, Michael D and Mike West. 1995. "Bayesian density estimation and inference using mixtures." *Journal of the American Statistical Association* 90(430):577–588.

Escobar, Michael D and Mike West. 1998. Computing nonparametric hierarchical models. In *Practical Nonparametric and Semiparametric Bayesian Statistics*. Springer pp. 1–22.

Eysenbach, Gunther. 2020. "How to fight an infodemic: The four pillars of infodemic management." *Journal of Medical Internet Research* 22(6):e21820.

Fariss, Christopher J. 2014. "Respect for human rights has improved over time: Modeling the changing standard of accountability." *American Political Science Review* 108(2):297–318.

Ferguson, Thomas S. 1973. "A Bayesian analysis of some nonparametric problems." *The Annals of Statistics* pp. 209–230.

Freedman, Rosa. 2013. *The United Nations Human Rights Council: A Critique and Early Assessment*. Routledge.

Friedman, Jerome, Trevor Hastie, Holger Höfling, Robert Tibshirani et al. 2007. "Pathwise coordinate optimization." *The Annals of Applied Statistics* 1(2):302–332.

Friedman, Jerome, Trevor Hastie and Rob Tibshirani. 2010. "Regularization paths for generalized linear models via coordinate descent." *Journal of Statistical Software* 33(1):1.

Frühwirth-Schnatter, Sylvia. 1994. "Data augmentation and dynamic linear models." *Journal of Time Series Analysis* 15(2):183–202.

Frühwirth-Schnatter, Sylvia. 2006. *Finite Mixture and Markov Switching Models*. Springer Science & Business Media.

Frühwirth-Schnatter, Sylvia and Gertraud Malsiner-Walli. 2019. "From here to infinity: sparse finite versus Dirichlet process mixtures in model-based clustering." *Advances in Data Analysis and Classification* 13(1):33–64.

Galvão, Jane. 2020. "COVID-19: The deadly threat of misinformation." *The Lancet Infectious Diseases* .

Gormley, Isobel Claire and Thomas Brendan Murphy. 2008. "Exploring voting blocs within the Irish electorate: A mixture modeling approach." *Journal of the American Statistical Association* 103(483):1014–1027.

Gormley, Isobel Claire, Thomas Brendan Murphy et al. 2008. "A mixture of experts model for rank data with applications in election studies." *The Annals of Applied Statistics* 2(4):1452–1477.

Habibi, Don A. 2007. "Human rights and politicized human rights: A utilitarian critique." *Journal of Human Rights* 6(1):3–35.

Hart, P Sol, Sedona Chinn and Stuart Soroka. 2020. "Politicization and polarization in COVID-19 news coverage." *Science Communication* 42(5):679–697.

Hirose, Kei and Michio Yamamoto. 2014. "Estimation of an oblique structure via penalized likelihood factor analysis." *Computational Statistics & Data Analysis* 79:120–132.

Hirose, Kei and Michio Yamamoto. 2015. "Sparse estimation via nonconcave penalized likelihood in factor analysis model." *Statistics and Computing* 25(5):863–875.

Hirose, Kei and Sadanori Konishi. 2012. "Variable selection via the weighted group lasso for factor analysis models." *Canadian Journal of Statistics* 40(2):345–361.

Holloway, Steven. 1990. "Forty years of United Nations General Assembly voting." *Canadian Journal of Political Science* 23(2):279–296.

Hovet, Thomas. 1958. *Bloc Politics in the United Nations*. Center for International Studies, Massachusetts Institute of Technology.

Hug, Simon. 2016. "Dealing with human rights in international organizations." *Journal of Human Rights* 15(1):21–39.

Hug, Simon and Richard Lukács. 2014. "Preferences or blocs? Voting in the United Nations Human Rights Council." *The Review of International Organizations* 9(1):83–106.

Imai, Kosuke, James Lo and Jonathan Olmsted. 2016. "Fast estimation of ideal points with massive data." *American Political Science Review* 110(4):631–656.

Ishwaran, Hemant and Lancelot F James. 2001. "Gibbs sampling methods for stick-breaking priors." *Journal of the American Statistical Association* 96(453):161–173.

Ishwaran, Hemant and Mahmoud Zarepour. 2000. "Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models." *Biometrika* 87(2):371–390.

Islam, Md Saiful, Tonmoy Sarkar, Sazzad Hossain Khan, Abu-Hena Mostofa Kamal, SM Murshid Hasan, Alamgir Kabir, Dalia Yeasmin, Mohammad Ariful Islam, Kamal Ibne Amin Chowdhury, Kazi Selim Anwar et al. 2020. "COVID-19–related infodemic and its impact on public health: A global social media analysis." *The American Journal of Tropical Medicine and Hygiene* 103(4):1621.

Jackman, Simon. 2000. "Estimation and inference are missing data problems: Unifying social science statistics via Bayesian simulation." *Political Analysis* 8(4):307–332.

Jackman, Simon. 2001. "Multidimensional analysis of roll call data via Bayesian simulation: Identification, estimation, inference, and model checking." *Political Analysis* 9(3):227–241.

Jackman, Simon. 2009. *Bayesian Analysis for the Social Sciences*. Vol. 846 John Wiley & Sons.

Jain, Sonia and Radford M Neal. 2004. "A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model." *Journal of Computational and Graphical Statistics* 13(1):158–182.

Jasra, Ajay, Chris C Holmes and David A Stephens. 2005. "Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling." *Statistical Science* pp. 50–67.

Jessee, Stephen A. 2009. "Spatial voting in the 2004 presidential election." *American Political Science Review* 103(1):59–81.

Johnson, Timothy R and Kristine M Kuhn. 2013. "Bayesian Thurstonian models for ranking data using JAGS." *Behavior Research Methods* 45(3):857–872.

Kim, Younghoon, Wooyeol Kim and Kyuseok Shim. 2017. "Latent ranking analysis using pairwise comparisons in crowdsourcing platforms." *Information Systems* 65:7–21.

King, Gary, Christopher J.L. Murray, Joshua A. Salomon and Ajay Tandon. 2004. "Enhancing validity and cross-cultural comparability of measurement in survey research." *American Political Science Review* 98:191–207.

Lee, Seokho, Jianhua Z Huang and Jianhua Hu. 2010. "Sparse logistic principal components analysis for binary data." *The Annals of Applied Statistics* 4(3):1579.

Lijphart, Arend. 1963. "The analysis of bloc voting in the General Assembly: A critique and a proposal." *American Political Science Review* 57(4):902–917.

Luif, Paul. 2003. *EU Cohesion in the UN General Assembly*. Vol. 49 European Union Institute for Security Studies Paris.

MacEachern, Steven N. 1994. "Estimating normal means with a conjugate style Dirichlet process prior." *Communications in Statistics-Simulation and Computation* 23(3):727–741.

MacEachern, Steven N and Peter Müller. 1998. "Estimating mixture of Dirichlet process models." *Journal of Computational and Graphical Statistics* 7(2):223–238.

Macon, Kevin T, Peter J Mucha and Mason A Porter. 2012. "Community structure in the United Nations General Assembly." *Physica A: Statistical Mechanics and its Applications* 391(1-2):343–361.

Malsiner-Walli, Gertraud, Sylvia Frühwirth-Schnatter and Bettina Grün. 2016. "Model-based clustering based on sparse finite Gaussian mixtures." *Statistics and Computing* 26(1-2):303–324.

Malsiner-Walli, Gertraud, Sylvia Frühwirth-Schnatter and Bettina Grün. 2017. "Identifying mixtures of mixtures using Bayesian estimation." *Journal of Computational and Graphical Statistics* 26(2):285–295.

Martin, Andrew D and Kevin M Quinn. 2002. "Dynamic ideal point estimation via Markov chain Monte Carlo for the US Supreme Court, 1953–1999." *Political Analysis* 10(2):134–153.

Martin, Andrew, Kevin Quinn and Jong Hee Park. 2011. "MCMCpack: Markov Chain Monte Carlo in R." *Journal of Statistical Software* 42(9):1–21.
URL: *https://www.jstatsoft.org/v042/i09*

Maydeu-Olivares, Albert and Ulf Böckenholt. 2005. "Structural equation modeling of paired-comparison and ranking data." *Psychological Methods* 10(3):285.

Maydeu-Olivares, Alberto and Anna Brown. 2010. "Item response modeling of paired comparison and ranking data." *Multivariate Behavioral Research* 45(6):935–974.

McAlister, Kevin. 2021. "Disagreement and dimensionality: A varying dimensions approach to roll call scaling in the U.S. Congress.".

Miller, Jeffrey W and Matthew T Harrison. 2013. "A simple example of Dirichlet process mixture inconsistency for the number of components." *Advances in Neural Information Processing Systems* 26:199–206.

Müller, Peter, Abel Rodriguez et al. 2013. Dirichlet Process. In *Nonparametric Bayesian Inference*. IMS and ASA pp. 23–41.

Murtagh, Fionn and Pierre Legendre. 2014. "Ward's hierarchical agglomerative clustering method: which algorithms implement Ward's criterion?" *Journal of Classification* 31(3):274–295.

Neal, Radford M. 2000. "Markov chain sampling methods for Dirichlet process mixture models." *Journal of Computational and Graphical Statistics* 9(2):249–265.

Newcombe, Hanna, Michael Ross and Alan G Newcombe. 1970. "United Nations voting patterns." *International Organization* 24(1):100–121.

Normand, Roger and Sarah Zaidi. 2008. *Human Rights at the UN: The Political History of Universal Justice*. Indiana University Press.

Oishi, Shigehiro, Jungwon Hahn, Ulrich Schimmack, Phanikiran Radhakrishan, Vivian Dzokoto and Stephen Ahadi. 2005. "The measurement of values across cultures: A pairwise comparison approach." *Journal of Research in Personality* 39:299–305.

Park, Ju Yeon. 2021. "When do politicians grandstand? Measuring message politics in committee hearings." *The Journal of Politics* 83(1):000–000.

Pauls, Scott D and Skyler J Cranmer. 2017. "Affinity communities in United Nations voting: Implications for democracy, cooperation, and conflict." *Physica A: Statistical Mechanics and its Applications* 484:428–439.

Petris, Giovanni, Sonia Petrone and Patrizia Campagnoli. 2009. Dynamic linear models. In *Dynamic Linear Models with R*. Springer pp. 31–84.

Phelps, Andrew S., David M. Naeger, Jesse L. Courtier, Jack W. Lambert, Peter A. Marcovici, Javier E. Villanueva-Meyer and John D MacKenzie. 2015. "Pairwise comparison versus Likert scale for biomedical image assessment." *American Journal of Roentgenology* 204(1):8–14.

Pitman, Jim. 1996. "Some developments of the Blackwell-MacQueen urn scheme." *Lecture Notes-Monograph Series* pp. 245–267.

Pomeroy, Caleb, Niheer Dasandi and Slava Jankin Mikhaylov. 2019. "Multiplex communities and the emergence of international conflict." *PLoS One* 14(10).

Poole, Keith T. 2000. "Nonparametric unfolding of binary choice data." *Political Analysis* 8(3):211–237.

Poole, Keith T. 2001. "The geometry of multidimensional quadratic utility in models of parliamentary roll call voting." *Political Analysis* 9(3):211–226.

Poole, Keith T and Howard Rosenthal. 1985. "A spatial model for legislative roll call analysis." *American Journal of Political Science* pp. 357–384.

Poole, Keith T and Howard Rosenthal. 1991. "Patterns of Congressional Voting." *American Journal of Political Science* pp. 228–278.

Quinn, Kevin M. 2004. "Bayesian factor analysis for mixed ordinal and continuous responses." *Political Analysis* 12(4):338–353.

Richard Hahn, P, Carlos M Carvalho and James G Scott. 2012. "A sparse factor analytic probit model for congressional voting patterns." *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 61(4):619–635.

Roberts, Gareth O, Andrew Gelman, Walter R Gilks et al. 1997. "Weak convergence and optimal scaling of random walk Metropolis algorithms." *The Annals of Applied Probability* 7(1):110–120.

Rubin, Donald B. 1976. "Inference and missing data." *Biometrika* 63(3):581–592.

Schnakenberg, Keith E and Christopher J Fariss. 2014. "Dynamic patterns of human rights practices." *Political Science Research and Methods* 2(1):1–31.

Sethuraman, Jayaram. 1994. "A constructive definition of Dirichlet priors." *Statistica Sinica* pp. 639–650.

Shor, Boris, Christopher Berry and Nolan McCarty. 2010. "A bridge to somewhere: Mapping state and congressional ideology on a cross-institutional common space." *Legislative Studies Quarterly* 35(3):417–448.

Spirling, Arthur and Kevin Quinn. 2010. "Identifying intraparty voting blocs in the UK House of Commons." *Journal of the American Statistical Association* 105(490):447–457.

Stark, Stephen, Oleksandr S Chernyshenko and Fritz Drasgow. 2005. "An IRT approach to constructing and scoring pairwise preference items involving stimuli on different dimensions: The multi-unidimensional pairwise-preference model." *Applied Psychological Measurement* 29(3):184–203.

Suchman, L. and B. Jordan. 1990. "Interactional troubles in face to face survey interviews (with comments and rejoinder)." *Journal of the American Statistical Association* 85(March):232–253.

Sun, Jianan, Yunxiao Chen, Jingchen Liu, Zhiliang Ying and Tao Xin. 2016. "Latent variable selection for multidimensional Item Response Theory models via L1 regularization." *Psychometrika* 81(4):921–939.

Tahk, Alexander. 2018. "Nonparametric ideal-point estimation and inference." *Political Analysis* 26(2):131–146.

Thurstone, Louis L. 1927. "A law of comparative judgment." *Psychological Review* 34(4):273.

Treier, Shawn and Simon Jackman. 2008. "Democracy as a latent variable." *American Journal of Political Science* 52(1):201–217.

Trump, Donald. 2018. "Remarks by President Trump to the 73rd Session of the United Nations General Assembly.".
URL: *https://www.whitehouse.gov/briefings-statements/remarks-president-trump-73rd-session-united-nations-general-assembly-new-york-ny/*

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

Voeten, Erik. 2004. "Resisting the lonely superpower: Responses of states in the United Nations to US dominance." *The Journal of Politics* 66(3):729–754.

Voeten, Erik, Anton Strezhnev and Michael Bailey. 2009. "United Nations General Assembly Voting Data.".
URL: *https://doi.org/10.7910/DVN/LEJUQZ*

Wang, Wen-Chung, Xue-Lan Qiu, Chia-Wen Chen, Sage Ro and Kuan-Yu Jin. 2017. "Item Response Theory models for ipsative tests with multidimensional pairwise comparison items." *Applied Psychological Measurement* 41(8):600–613.

West, Mike and Jeff Harrison. 2006. *Bayesian Forecasting and Dynamic Models*. Springer Science & Business Media.

Yu, Philip LH and Lucia KY Chan. 2001. "Bayesian analysis of wandering vector models for displaying ranking data." *Statistica Sinica* pp. 445–461.