doi:10.1017/S000305541600037X © American Political Science Association 2016

Fast Estimation of Ideal Points with Massive Data

KOSUKE IMAI *Princeton University*
JAMES LO *University of Southern California*
JONATHAN OLMSTED *The NPD Group*

Estimation of ideological positions among voters, legislators, and other actors is central to many subfields of political science. Recent applications include large data sets of various types including roll calls, surveys, and textual and social media data. To overcome the resulting computational challenges, we propose fast estimation methods for ideal points with massive data. We derive the expectation-maximization (EM) algorithms to estimate the standard ideal point model with binary, ordinal, and continuous outcome variables. We then extend this methodology to dynamic and hierarchical ideal point models by developing variational EM algorithms for approximate inference. We demonstrate the computational efficiency and scalability of our methodology through a variety of real and simulated data. In cases where a standard Markov chain Monte Carlo algorithm would require several days to compute ideal points, the proposed algorithm can produce essentially identical estimates within minutes. Open-source software is available for implementing the proposed methods.

INTRODUCTION

Estimation of ideological positions among voters, legislators, justices, and other actors is central to many subfields of political science. Since the pioneering work of Poole and Rosenthal (1991, 1997), a number of scholars have used spatial voting models to estimate ideological preferences from roll-call votes and other data in the fields of comparative politics and international relations as well as American politics (e.g., Bailey, Kamoie, and Maltzman 2005; Bailey, Strezhnev, and Voeten 2015; Bonica 2014; Clinton and Lewis 2008; Hix, Noury, and Roland 2006; Ho and Quinn 2010; Londregan 2007; McCarty, Poole, and Rosenthal 2006; Morgenstern 2004; Spirling and McLean 2007; Voeten 2000). These and other substantive applications are made possible by numerous methodological advancements including Bayesian estimation (Clinton, Jackman, and Rivers 2004; Jackman 2001), optimal classification (Poole 2000), dynamic modeling (Martin and Quinn 2002), and models with agenda setting or strategic voting (Clinton and Meirowitz 2003; Londregan 1999).

With the increasing availability of data and methodological sophistication, researchers have recently turned their attention to the estimation of ideological preferences that are comparable across time and institutions. For example, Bailey (2007) measures ideal points of U.S. presidents, senators, representatives, and Supreme Court justices on the same scale over time (see also Bailey 2013; Bailey and Chang 2001). Similarly, Shor and McCarty (2011) compute the ideal points of the state legislators from all U.S. states and compares them with members of Congress (see also Battista, Peress, and Richman 2013; Shor, Berry, and McCarty 2011). Finally, Bafumi and Herron (2010) estimate the ideological positions of voters and their members of Congress in order to study representation while Clinton et al. (2012) compare the ideal points of agencies with those of presidents and congressional members.

Furthermore, researchers have begun to analyze large data sets of various types. For example, Slapin and Proksch (2008) develop a statistical model that can be applied to estimate ideological positions from textual data. Proksch and Slapin (2010) and Lowe et al. (2011) apply this and other similar models to the speeches of European Parliament and the manifestos of European parties, respectively (see also Kim, Londregan, and Ratkovic 2014, who analyze the speeches of legislators in U.S. Congress). Another important new data source is social media, which often come in massive size. Bond and Messing (2015) estimate the ideological preferences of 6.2 million Facebook users while Barberá (2015) analyze more than 40 million Twitter users. These social media data are analyzed as network data, and a similar approach is taken to estimate ideal points from data on citations using court opinions (Clark and Lauderdale 2010) and campaign contributions from voters to politicians (Bonica 2014).

These new applications pose a computational challenge of dealing with data sets that are orders of magnitude larger than the canonical single-chamber roll-call matrix for a single time period. Indeed, as Table 1 shows, the past decade has witnessed a significant rise in the use of large and diverse data sets for ideal point estimation. While most of the aforementioned works are based on Bayesian models of ideal points, standard

Kosuke Imai is Professor, Department of Politics and Center for Statistics and Machine Learning, Princeton University, Princeton, NJ 08544. Phone: 609–258–6601 (kimai@princeton.edu), URL: http://imai.princeton.edu.

James Lo is Assistant Professor, Department of Political Science, University of Southern California, Los Angeles, CA 90089 (lojames@usc.edu).

Jonathan Olmsted is Solutions Manager, NPD Group, Port Washington, NY 11050 (jpolmsted@gmail.com).

We thank Simon Jackman, Will Lowe, Michael Peress, and Marc Ratkovic for their helpful discussions, Kevin Quinn for providing a replication data set and code, Yuki Shiraito for his assistance with the Japanese survey data, and the editors and anonymous reviewers for their helpful comments. The proposed methods are implemented through open-source software `emIRT`, which is available as an R package at the Comprehensive R Archive Network (CRAN; http://cran.r-project.org/package=emIRT). Replication code for this article is available at Imai, Lo, and Olmsted (2016).

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

TABLE 1. Recent Applications of Ideal Point Models to Various Large Data Sets

	Number of Subjects	Number of Items	Data Types
Parliaments			
DW-NOMINATE scores (1789–2012)	37,511	46,379	U.S. Congress
Common Space scores (1789–2012)	11,833	90,609	U.S. Congress
Hix, Noury, and Roland (2006)	2,000	12,000	European Parliament
Shor and McCarty (2011)	6,201	5,747	U.S. state legislatures
Bailey, Strezhnev, and Voeten (2015)	2,187	7,335	United Nations
Courts (and other institutions)			
Martin and Quinn scores (1937–2013)	697	5,164	U.S. Supreme Court
Bailey (2007)	27,795	2,750	US. Supreme Court, Congress, Presidents
Voters (and politicians)			
Gerber and Lewis (2004)	2.8 million	12	referendum votes
Bafumi and Herron (2010)	8,848	4,391	survey & roll calls
Tausanovitch and Warshaw (2013)	275,000	311	survey
Bonica (2014)	4.2 million	78,363	campaign contributions
Social Media			
Bond and Messing (2015)	6.2 million	1,223	Facebook
Barberá (2015)	40.2 million	1,465	Twitter
Texts and Speeches			
Clark and Lauderdale (2010)	1,000	1,000	U.S. Supreme Court citations
Proksch and Slapin (2010)	25	8995	German manifestos
Lewandowski et al. (2015)	1,000	250,000	European party manifestos

Notes: The past decade has witnessed a significant rise in the use of large data sets for ideal point estimation. Note that "# of subjects" should be interpreted as the number of ideal points to be estimated. For example, if a legislator serves for two terms and are allowed to have different ideal points in those terms, then this legislator is counted as two subjects.

Markov chain Monte Carlo (MCMC) algorithms can be prohibitively slow when applied to large data sets. As a result, researchers are often unable to estimate their models using the entire data and are forced to adopt various shortcuts and compromises. For example, Shor and McCarty (2011) fit their model in multiple steps using subsets of the data whereas Bailey (2007) resorts to a simpler parametric dynamic model in order to reduce computational costs (p. 441) (see also Bailey 2013). Since a massive data set implies a large number of parameters under these models, the convergence of MCMC algorithms also becomes difficult to assess. Bafumi and Herron (2010), for example, express a concern about the convergence of ideal points for voters (footnote 24).

In addition, estimating ideal points over a long period of time often imposes a significant computational burden. Indeed, the use of computational resources at supercomputer centers has been critical to the development of various NOMINATE scores.[1] Similarly, estimation of the Martin and Quinn (2002) ideal point estimates for U.S. Supreme Court justices over 47 years took over five days to estimate. This suggests that while these ideal point models are attractive, they are often practically unusable for many researchers who wish to analyze a large-scale data set.

In this article, we propose a fast estimation method for ideal points with massive data. Specifically, we develop the Expectation-Maximization (EM) algorithms (Dempster, Laird, and Rubin 1977) that either exactly or approximately maximize the posterior distribution under various ideal point models. The main advantage of EM algorithms is that they can dramatically reduce computational time. Through a number of empirical and simulation examples, we demonstrate that in cases where a standard MCMC algorithm would require several days to compute ideal points, the proposed algorithm can produce essentially identical estimates within minutes. The EM algorithms also scale much better than other existing ideal point estimation algorithms. They can estimate an extremely large number of ideal points on a laptop within a few hours whereas current methodologies would require the level of computational resources only available at a supercomputer center to do the same computation.

We begin by deriving the EM algorithm for the standard Bayesian ideal point model of Clinton, Jackman, and Rivers (2004). We show that the proposed algorithm produces ideal point estimates which are essentially identical to those from other existing methods. We then extend our approach to other popular ideal point models that have been developed in the literature. Specifically, we develop an EM algorithm for the model with mixed ordinal and continuous outcomes

[1] The voteview website notes that the DW-NOMINATE and Common Space DW-NOMINATE scores are computed using the Rice terascale cluster. See http://voteview.com/dwnominate.asp and http://voteview.com/dwnomjoint.asp (accessed on November 10, 2014).

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

(Quinn 2004) by applying a certain transformation to the original parametrization. We also develop an EM algorithm for the dynamic model (Martin and Quinn 2002) and the hierarchical model (Bafumi et al. 2005). Finally, we propose EM algorithms for ideal point models based on textual and network data.

For dynamic and hierarchical models as well as the models for textual and network data, an EM algorithm that directly maximizes the posterior distribution is not available in a closed form. Therefore, we rely on variational Bayesian inference, which is a popular machine learning methodology for fast and approximate Bayesian estimation (see Wainwright and Jordan (2008) for a review and Grimmer (2011) for an introductory article in political science). For each case, we demonstrate the computational efficiency and scalability of the proposed methodology by applying it to a wide range of real and simulated data sets. Our proposed algorithms complement a recent application of variational inference to combine ideal point estimation with topic models (Gerrish and Blei 2012). We implement the proposed algorithms via an open-source R package, **emIRT** (Imai, Lo, and Olmsted 2015), so that others can apply them to their own research.

In the item response theory literature, the EM algorithm is used to maximize the marginal likelihood function where ability parameters, i.e., ideal point parameters in the current context, are integrated out (Bock and Aitkin 1981). In the ideal point literature, Bailey (2007) and Bailey and Chang (2001) use variants of the EM algorithm in their model estimation. The M steps of these existing algorithms, however, do not have a closed-form solution. In this article, we derive closed-form EM algorithms for popular Bayesian ideal point models. This leads to faster and more reliable estimation algorithms.

Finally, an important and well-known drawback of these EM algorithms is that they do not produce uncertainty estimates such as standard errors. In contrast, the MCMC algorithms are designed to fully characterize the posterior, enabling the computation of uncertainty measures for virtually any quantities of interest. Moreover, the standard errors based on variational posterior are often too small, underestimating the degree of uncertainty. While many applied researchers tend to ignore estimation uncertainty associated with ideal points, such a practice can yield misleading inference. To address this problem, we apply the parametric bootstrap approach of Lewis and Poole (2004) (see also Carroll et al. 2009). Although this obviously increases the computational cost of the proposed approach, the proposed EM algorithms still scale much better than the existing alternatives. Furthermore, researchers can reduce this computational cost by a parallel implementation of bootstrap on a distributed system. We note that since our models are Bayesian, it is rather unconventional to utilize bootstrap, which is a frequentist procedure. However, one can interpret the resulting confidence intervals as a measure of uncertainty of our Bayesian estimates over repeated sampling under the assumed model.

STANDARD IDEAL POINT MODEL

We begin by deriving the EM algorithm for the standard ideal point model of Clinton, Jackman, and Rivers (2004). In this case, the proposed EM algorithm maximizes the posterior distribution without approximation. We illustrate the computational efficiency and scalability of our proposed algorithm by applying it to roll-call votes in recent U.S. Congress sessions, as well as to simulated data.

The Model

Suppose that we have N legislators and J roll calls. Let y_{ij} denote the vote of legislator i on roll call j where $y_{ij} = 1$ ($y_{ij} = 0$) implies that the vote is in the affirmative (negative) with $i = 1, \ldots, N$ and $j = 1, \ldots, J$. Abstentions, if present, are assumed to be ignorable such that these votes are missing at random and can be predicted from the model using observed data (see Rosas and Shomer 2008). Furthermore, let $\mathbf{x}_i$ represent the K-dimensional column vector of ideal point for legislator i. Then, if we use y^*_{ij} to represent a latent propensity to cast a "yea" vote where $y_{ij} = \mathbf{1}\{y^* > 0\}$, the standard K-dimensional ideal point model is given by

$$y^*_{ij} = \alpha_j + \mathbf{x}_i^\top \boldsymbol{\beta}_j + \epsilon_{ij}, \tag{1}$$

where $\boldsymbol{\beta}_j$ is the K-dimensional column vector of item discrimination parameters and α_j is the scalar item difficulty parameter. Finally, ϵ_{ij} is an independently, identically distributed random utility and is assumed to follow the standard normal distribution.

For notational simplicity, we use $\tilde{\boldsymbol{\beta}}_j^\top = (\alpha_j, \boldsymbol{\beta}_j^\top)$ and $\tilde{\mathbf{x}}_i^\top = (1, \mathbf{x}_i^\top)$ so that equation (1) can be more compactly written as

$$y^*_{ij} = \tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\beta}}_j + \epsilon_{ij}. \tag{2}$$

Following the original article, we place independent and conjugate prior distributions on $\mathbf{x}_i$ and $\tilde{\boldsymbol{\beta}}_j$, separately. Specifically, we use

$$p(\mathbf{x}_1, \ldots, \mathbf{x}_N) = \prod_{i=1}^{N} \phi_K(\mathbf{x}_i; \boldsymbol{\mu}_{\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{x}}) \quad \text{and}$$

$$p(\tilde{\boldsymbol{\beta}}_1, \ldots, \tilde{\boldsymbol{\beta}}_J) = \prod_{j=1}^{J} \phi_{K+1}\left(\tilde{\boldsymbol{\beta}}_j; \boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}\right), \tag{3}$$

where $\phi_k(\cdot;\cdot)$ is the density of a k-variate Normal random variable, $\boldsymbol{\mu}_{\mathbf{x}}$ and $\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}$ represent the prior mean vectors, and $\boldsymbol{\Sigma}_{\mathbf{x}}$ and $\boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}$ are the prior covariance matrices.

Given this model, the joint posterior distribution of $(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J)$ conditional on the roll-call matrix

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

$\mathbf{Y}$ is given by

$$
\begin{aligned}
& p\left(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J \mid \mathbf{Y}\right) \\
& \propto \prod_{i=1}^{N}\prod_{j=1}^{J}\left(\mathbf{1}\{y_{ij}^* > 0\}\mathbf{1}\{y_{ij} = 1\}\right. \\
& \quad \left. + \mathbf{1}\{y_{ij}^* \leq 0\}\mathbf{1}\{y_{ij} = 0\}\right)\ \phi_1\left(y_{ij}^*; \tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\beta}}_j, 1\right) \\
& \quad \times \prod_{i=1}^{N} \phi_K\left(\mathbf{x}_i; \boldsymbol{\mu}_{\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{x}}\right) \prod_{j=1}^{J} \phi_{K+1}\left(\tilde{\boldsymbol{\beta}}_j; \boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}\right), \quad \textbf{(4)}
\end{aligned}
$$

where $\mathbf{Y}$ and $\mathbf{Y}^*$ are matrices whose element in the ith row and jth column is y_{ij} and y_{ij}^*, respectively. Clinton, Jackman, and Rivers (2004) describe the MCMC algorithm to sample from this joint posterior distribution and implement it as the `ideal()` function in the open-source R package **pscl** (Jackman 2012).

The Proposed Algorithm

We derive the EM algorithm that maximizes the posterior distribution given in equation (4) without approximation. The proposed algorithm views $\{\mathbf{x}_i\}_{i=1}^N$ and $\{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J$ as parameters and treats $\mathbf{Y}^*$ as missing data. Specifically, at the tth iteration, denote the current parameter values as $\{\mathbf{x}_i^{(t-1)}\}_{i=1}^N$ and $\{\tilde{\boldsymbol{\beta}}_j^{(t-1)}\}_{j=1}^J$. Then, the E step is given by the following so-called "Q function," which represents the expectation of the log joint posterior distribution,

$$
\begin{aligned}
& Q(\{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J) \\
& = \mathbb{E}\Big[\log p(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J \mid \mathbf{Y}) \mid \mathbf{Y}, \\
& \qquad \{\mathbf{x}_i^{(t-1)}\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j^{(t-1)}\}_{j=1}^J\Big] \\
& = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{J}\left(\tilde{\boldsymbol{\beta}}_j^\top \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\beta}}_j - 2\tilde{\boldsymbol{\beta}}_j^\top \tilde{\mathbf{x}}_i y_{ij}^{*(t)}\right) \\
& \quad -\frac{1}{2}\sum_{i=1}^{N}\left(\mathbf{x}_i^\top \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}\mathbf{x}_i - 2\mathbf{x}_i^\top \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}\boldsymbol{\mu}_{\mathbf{x}}\right) \\
& \quad -\frac{1}{2}\sum_{j=1}^{J}\left(\tilde{\boldsymbol{\beta}}_j^\top \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\tilde{\boldsymbol{\beta}}_j - 2\tilde{\boldsymbol{\beta}}_j^\top \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}\right) + \text{const.} \quad \textbf{(5)}
\end{aligned}
$$

where

$$
\begin{aligned}
y_{ij}^{*(t)} &= \mathbb{E}\left(y_{ij}^* \mid \mathbf{x}_i^{(t-1)}, \tilde{\boldsymbol{\beta}}_j^{(t-1)}, y_{ij}\right) \\
&= \begin{cases} m_{ij}^{(t-1)} + \frac{\phi(m_{ij}^{(t-1)})}{\Phi(m_{ij}^{(t-1)})} & \text{if } y_{ij} = 1 \\ m_{ij}^{(t-1)} - \frac{\phi(m_{ij}^{(t-1)})}{1-\Phi(m_{ij}^{(t-1)})} & \text{if } y_{ij} = 0 \\ m_{ij}^{(t-1)} & \text{if } y_{ij} \text{ is missing} \end{cases} \quad \textbf{(6)}
\end{aligned}
$$

with $m_{ij}^{(t-1)} = (\tilde{\mathbf{x}}_i^{(t-1)})^\top \tilde{\boldsymbol{\beta}}_j^{(t-1)}$.

Straightforward calculation shows that the maximization of this Q function, i.e., the M step, can be achieved via the following two conditional maximization steps:

$$
\begin{aligned}
\mathbf{x}_i^{(t)} &= \left(\boldsymbol{\Sigma}_{\mathbf{x}}^{-1} + \sum_{j=1}^{J}\boldsymbol{\beta}_j^{(t-1)}\boldsymbol{\beta}_j^{(t-1)\top}\right)^{-1} \\
&\quad \times \left(\boldsymbol{\Sigma}_{\mathbf{x}}^{-1}\boldsymbol{\mu}_{\mathbf{x}} + \sum_{j=1}^{J}\boldsymbol{\beta}_j^{(t-1)}(y_{ij}^{*(t)} - \alpha_j^{(t-1)})\right), \quad \textbf{(7)}
\end{aligned}
$$

$$
\begin{aligned}
\tilde{\boldsymbol{\beta}}_j^{(t)} &= \left(\boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1} + \sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)}\left(\tilde{\mathbf{x}}_i^{(t)}\right)^\top\right)^{-1} \\
&\quad \times \left(\boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}} + \sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)} y_{ij}^{*(t)}\right). \quad \textbf{(8)}
\end{aligned}
$$

The algorithm repeats these E and M steps until convergence. Given that the model is identified up to an affine transformation, we use a correlation-based convergence criteria where the algorithm terminates when the correlation between the previous and current values of all parameters reaches a prespecified threshold.[2]

Finally, to compute uncertainty estimates, we apply the parametric bootstrap (Lewis and Poole 2004). Specifically, we first estimate ideal points and bill parameters via the proposed EM algorithm. Using these estimates, we calculate the choice probabilities associated with each outcome. Then, we randomly generate roll-call matrices given these estimated outcome probabilities. Where there are missing votes, we simply induce the same missingness patterns. This is repeated a sufficiently large number of times and the resulting bootstrap replicates for each parameter are used to characterize estimation uncertainty.

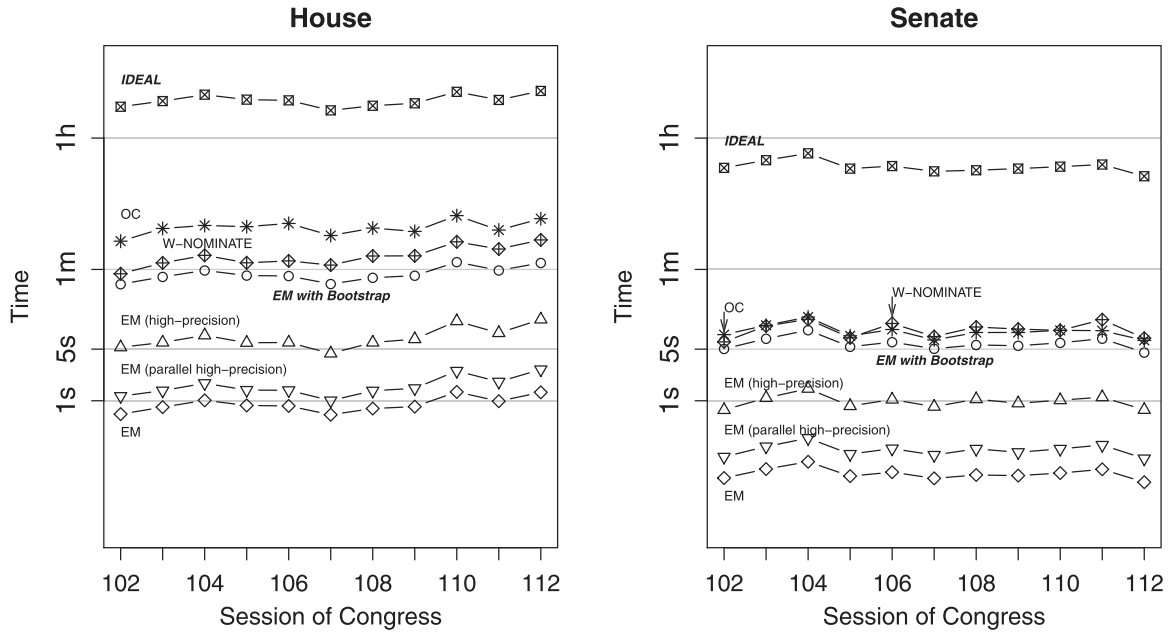
An Empirical Application

To assess its empirical performance, we apply the proposed EM algorithm to roll-call voting data for both the Senate and the House of Representatives for sessions of Congress 102 through 112. Specifically, we compare the ideal point estimates and their computation time from the proposed algorithm to those from three other methods; the MCMC algorithm implemented as `ideal()` in the R package **pscl** (Jackman 2012), the alternating maximum likelihood estimator implemented as `wnominate()` in the R package **wnominate** (Poole

[2] Most methods in the literature, including those based on Aitkin's acceleration and the gradient function, take this approach. As a reviewer correctly points out, however, a small difference can imply a lack of progress rather than convergence. Following the current literature, we recommend that researchers address this problem partially by employing a strict correlation criteria such as the correlation less than $1 - 10^{-6}$ and by using different starting values in order to avoid getting stuck in local maxima.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 1. Comparison of Computational Performance across the Methods



Notes: Each point represents the length of time required to compute estimates where the spacing of time on the vertical axis is based on the log scale. The proposed EM algorithm, indicated by "EM," "EM (high precision)," "EM (parallel high precision)," and "EM with Bootstrap" is compared with "W-NOMINATE" (Poole et al. 2011), the MCMC algorithm "IDEAL" (Jackman 2012), and the nonparametric optimal classification estimator "OC" (Poole et al. 2012). The EM algorithm is faster than the other approaches whether focused on point estimates or also estimation uncertainty. Algorithms producing uncertainty estimates are labeled in bold, italic type.

et al. 2011), and the nonparametric optimal classification estimator implemented as `oc()` in the R package **oc** (Poole et al. 2012). For all roll-call matrices, we restrict attention to just those legislators with at least 25 observed votes on nonunanimous bills. In all cases, we assume a single spatial dimension.

We caution that the comparison of computational efficiency presented here is necessarily illustrative. The performance of any numerical algorithm may depend on starting values, and no absolute convergence criteria exists for any of the algorithms we examine. For the MCMC algorithm, we run the chain for 100,000 iterations beyond a burn-in period of 20,000 iterations. Inference is based on a thinned chain where we keep every 100 draws. While the length of chain and its thinning interval are within the recommended range by Clinton, Jackman, and Rivers (2004), we emphasize that any convergence criteria used for deciding when to terminate the MCMC algorithm is somewhat arbitrary. The default diffuse priors, specified in `ideal()` from the R package **pscl**, are used and propensities for missing votes are not imputed during the data-augmentation step. In particular, we assume a single dimensional normal prior for all ideal point parameters with a mean of 0 and a variance of 1. For the bill parameters, we assume a two-dimensional normal prior with a mean vector of 0's and a covariance matrix with each variance term equal to 25 and no covariance. The standard normalization of ideal points, i.e., a mean of 0 and a standard deviation of 1, is used for local identification. The MCMC algorithm produces measures of uncertainty for the parameter estimates and so we distinguish it from those algorithms which produce only point estimates by labeling it in Figure 1 with bold, italic typeface.

For the proposed EM algorithm, we use random starting values for the ideal point and bill parameters. The same prior distributions as the MCMC algorithm are used for all parameters. We terminate the EM algorithm when each block of parameters has a correlation with the values from the previous iteration larger than $1 - p$. With one spatial dimension, we have three parameter blocks: the bill difficulty parameters, the bill discrimination parameters and the ideal point parameters. Following Poole and Rosenthal (1997) (see p. 237), we use $p = 10^{-2}$. We also consider a far more stringent criterion where $p = 10^{-6}$, requiring parameters to correlate greater than 0.999999. In the following results, we focus on the latter "high precision" variant, except in the case of computational performance where results for both criteria are presented (labeled "EM" and "EM (high precision)," respectively). The results from the EM algorithm do not include measures of uncertainty. For this, we include the "EM with Bootstrap" variant which uses 100 parametric bootstrap replicates. To distinguish this algorithm from those which produce just point estimates, it is labeled in Figure 1 with bold, italic typeface.

Finally, for W-NOMINATE, we do not include any additional bootstrap trials for characterizing

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

uncertainty about the parameters, so the results will only include point estimates. For optimal classification, we use the default setting of `oc()` in the R package **oc**. Because the Bayesian MCMC algorithm is stochastic and the EM algorithm has random starting values, we run each estimator 50 times for any given roll-call matrix and report the median of each performance measurement.

We begin by examining the computational performance of the EM algorithm.[3] Figure 1 shows the time required for the ideal point estimation for each Congressional session in the House (left panel) and the Senate (right panel). Note that the vertical axis is on the log scale. Although the results are only illustrative for the aforementioned reasons, it is clear that the EM algorithm is by far the fastest. For example, for the 102nd House of Representatives, the proposed algorithm, denoted by "EM," takes less than one second to compute estimates, using the same convergence criteria as W-NOMINATE. Even with a much more stringent convergence criteria "EM (high-precision)," the computational time is only six seconds. This contrasts with the other algorithms, which require much more time for estimation. Although the direct comparison is difficult, the MCMC algorithm is by far the slowest, taking more than 2.5 hours. Because the MCMC algorithm produces standard errors, we contrast the performance of "IDEAL" with "EM with Bootstrap" and find that obtaining 100 bootstrap replicates requires just under one minute. Even if more iterations are desired, over 10,000 bootstrap iterations could be computed before approaching the time required by the MCMC algorithm.

The W-NOMINATE and optimal classification estimators are faster than the MCMC algorithm but take approximately one and 2.5 minutes, respectively. These methods do not provide measures of uncertainty, and all of the point-estimate EM variants are over ten times faster. Last, the EM algorithm is amenable to parallelization within each of the three update steps. The open-source implementation that we provide supports this on some platforms (Imai, Lo, and Olmsted 2015). And, the parallelized implementation performs well. For any of these roll-call matrices, using eight processor cores to estimate the parameters instead of just one core reduces the required timeto completion to about one-sixth of the single core time.

We next show that the computational gain of the proposed algorithm is achieved without sacrificing the quality of estimates. To do so, we directly compare individual-level ideal point estimates across the methods. Figure 2 shows, using the 112th Congress, that except for a small number of legislators, the estimates from the proposed EM algorithm are essentially identical to those from the MCMC algorithm (left column) and W-NOMINATE (right column). The within-party correlation remains high across the plots, indicating that a strong agreement among these estimates up to affine transformation.[4] The agreement with the results based on the MCMC algorithm is hardly surprising given that both algorithms are based on the same posterior distribution. The comparability of estimates across methods holds for the other sessions of Congress considered.

For a small number of legislators, the deviation between results for the proposed EM algorithm and both the MCMC algorithm and W-NOMINATE is not negligible. However, this is not a coincidence—it results from the degree of missing votes associated with each legislator.[5] The individuals for whom the estimates differ significantly all have no position registered for more than 40% of the possible votes. Examples include President Obama, the late Congressman Donald Payne (Democrat, NJ), and Congressman Thomas Massie (Republican, KY).[6] With a small amount of data, the estimation of these legislators' ideal points is sensitive to the differences in statistical methods.

We also compare the standard errors from the proposed EM algorithm using the parametric bootstrap with those from the MCMC algorithm. Because the output from the MCMC algorithm is rescaled to have a mean of zero and standard deviation of 1, we rescale the EM estimates to have the same sample moments. This affine transformation is applied to each bootstrap replicate so that the resulting bootstrap standard errors are on the same scale as those based on the MCMC algorithm. Figure 3 does this comparison using the estimates for the 112th House of Representatives. The left panel shows that the standard errors based on the EM algorithm with the bootstrap (the vertical axis) are only slightly smaller than those from the MCMC algorithm (the horizontal axis) for most legislators. However, for a few legislators, the standard errors from the EM algorithm are substantially smaller than those from the MCMC algorithm. The right panel of the figure shows that these legislators have extreme ideological preferences.

To further examine the frequentist properties of our standard errors based on parametric bootstrap, we conduct a Monte Carlo simulation where roll-call votes are

[3] All computation in this article is completed on a cluster computer running Red Hat Linux with multiple 2.67-GHz Intel Xeon X5550 processors. Unless otherwise noted, however, the computation is done utilizing a single processor to emulate the computational performance of our algorithms on a personal computer.

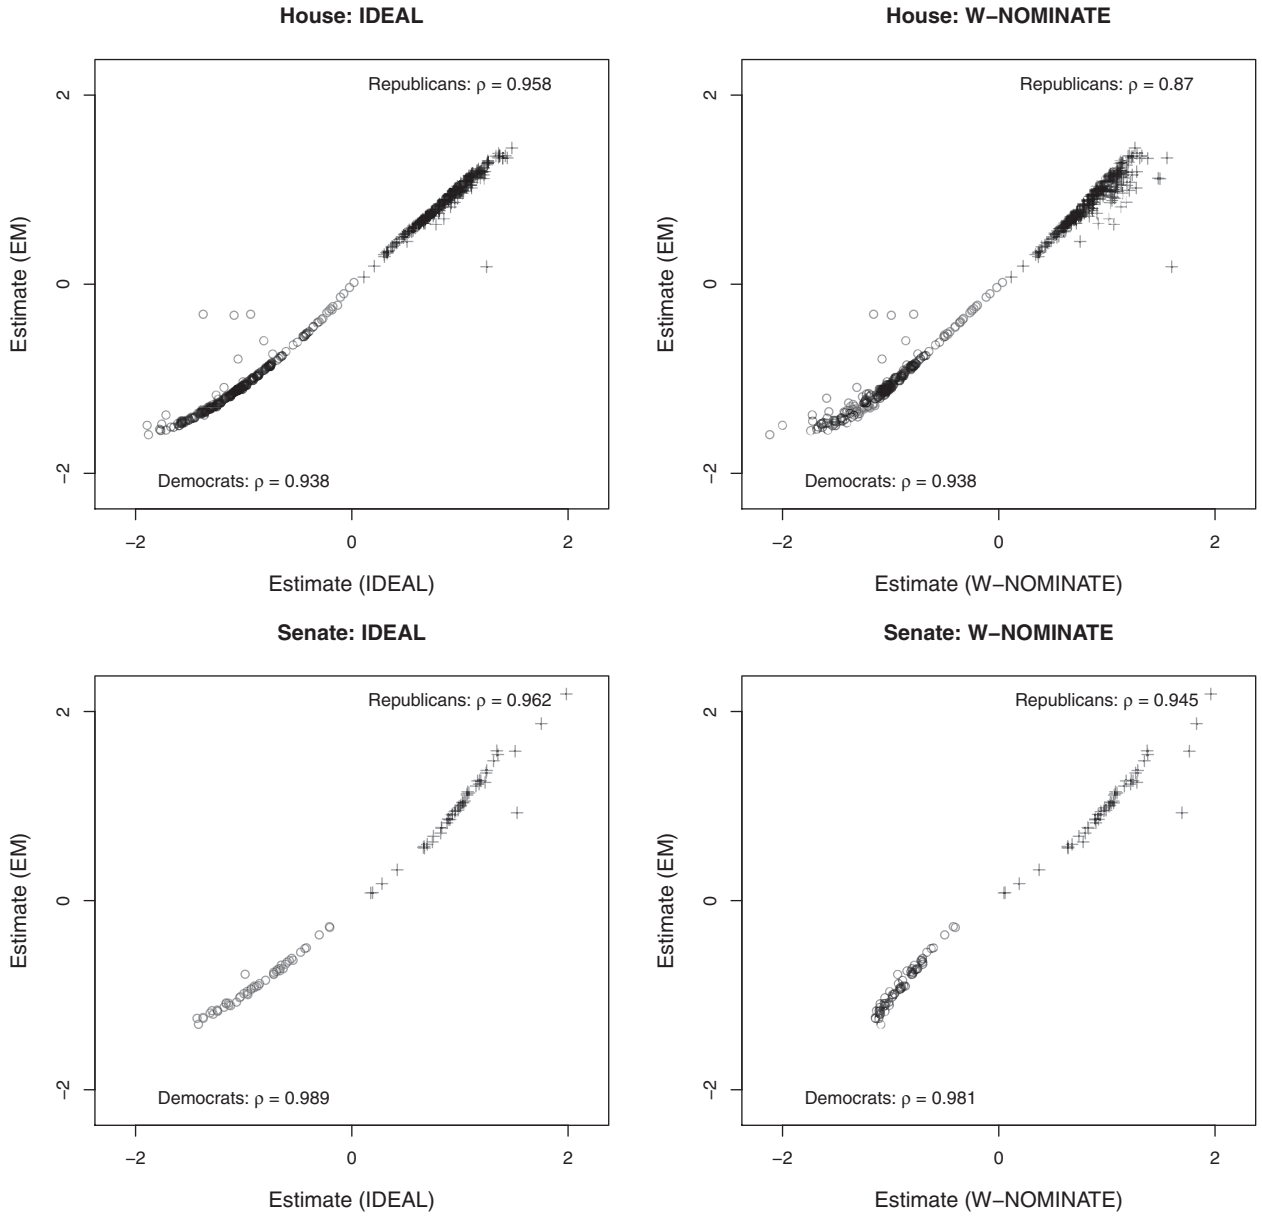
[4] The same Pearson correlations between the MCMC algorithm and W-NOMINATE for Republicans and Democrats are, respectively, 0.96 and 0.98 in the House. In the Senate, they are 0.99 and 0.98. The Spearman correlations within party between the EM algorithm and the nonparametric optimal classification algorithm are 0.77 and 0.77 in the House. They are 0.93 and 0.87 in the Senate. The same within-party correlations between OC and the W-NOMINATE and MCMC algorithms range from 0.80 to 0.93 in the House and 0.85 to 0.98 in the Senate.

[5] The discrepancies in these estimates are present even when the MCMC algorithm is executed by imputing missing votes rather than simply dropping them.

[6] Observations for the president are generated by interpreting statements of support or opposition (including vetoes) as votes. While these are not uncommon, they are issued far less frequently than Congress takes roll calls. Congressman Payne, on the other hand, passed away in the middle of the 112th session of Congress. Finally, Congressman Massie was sworn into office as the representative for Kentucky's 4th congressional district after winning a special election in November 2012 with just three months left in the 112th session.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 2. Comparison of Estimated Ideal Points across the Methods for the 112th Congress



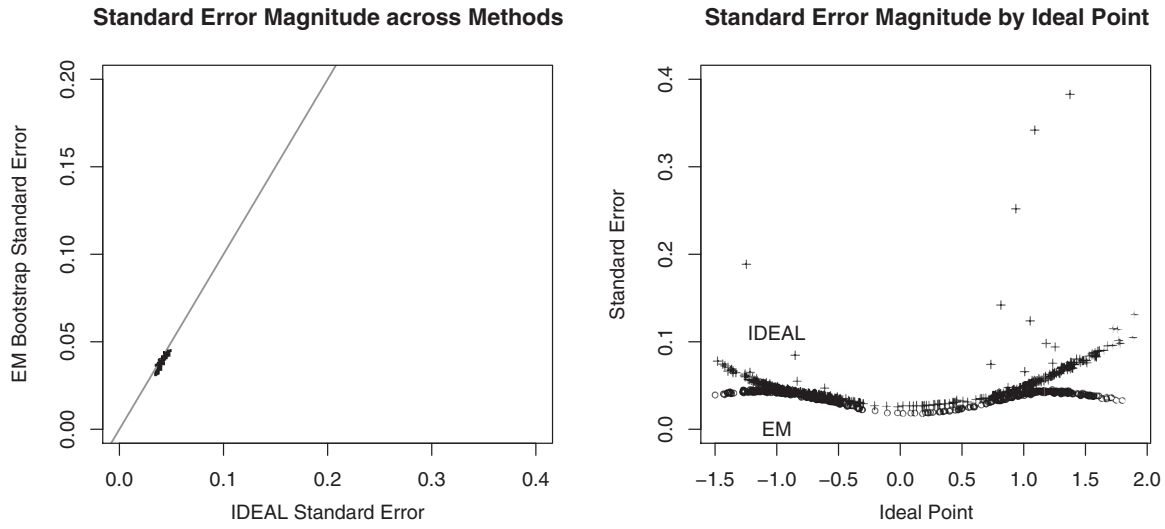
Notes: Republicans are shown with crosses while Democrats are indicated by hollow circles. The proposed EM algorithm is compared with the MCMC algorithm "IDEAL" (left column; Jackman 2012) and "W-NOMINATE" (right column; Poole et al. 2011). For each of these, the estimates are rescaled to a common scale for easy comparison across methods and chambers. Pearson correlation coefficients within parties are also reported, but are unaffected by the rescaling. The proposed algorithm yields the estimates that are essentially identical to those from the other two methods.

simulated consistent with data from the 112th House of Representatives. That is, we use the estimates obtained from these data as truths and simulate roll-call data 1,000 times according to the model. When simulating the data, the same missingness pattern as the observed data from the 112th Congress is used. For each of the 1,000 simulated roll call data sets, we estimate the ideal points and compute the standard error based on 100 parametric bootstrap replicates. We then estimate the bias of the resulting standard error for each legislator as the average difference between the bootstrap standard error and the standard deviation of estimated ideal points across 1,000 simulations.

Figure 4 shows the estimated biases. The left panel of the figure shows that the biases of the standard errors are not systematically related to the extremity of ideological positions. Thus, the divergence between the MCMC and parametric bootstrap standard errors for extreme legislators observed in Figure 3 does not necessarily suggest that the latter is underestimated.

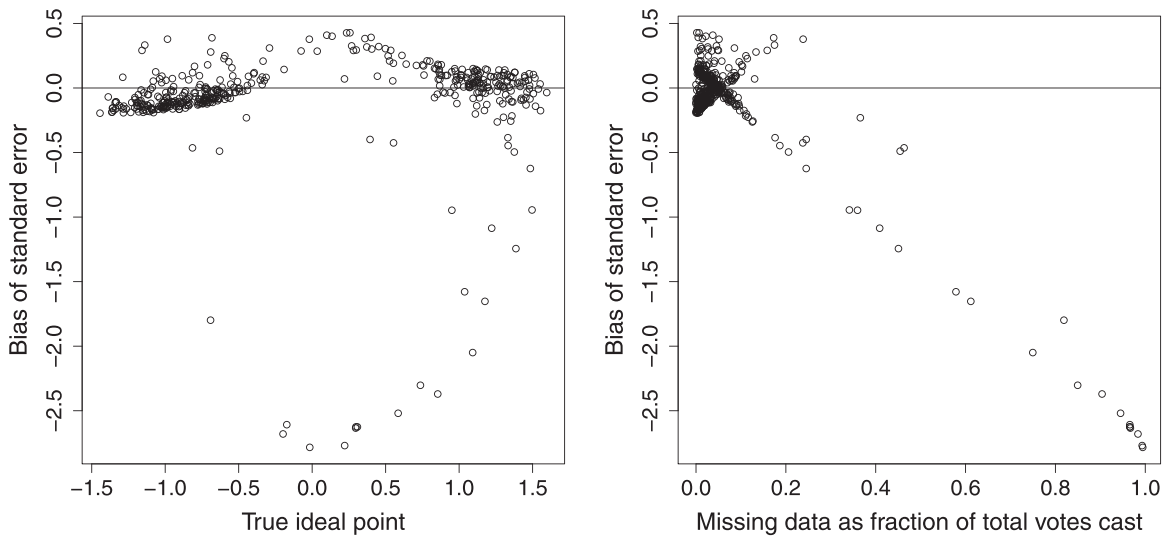
https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 3. Comparison of Standard Errors between the Proposed EM Algorithm and the Bayesian MCMC Algorithm using the 112th House of Representatives



Notes: The standard errors from the EM algorithm are based on the parametric bootstrap of 1,000 replicates. The left plot shows that the proposed standard errors (the vertical axis) are similar to those from the MCMC algorithms (the horizontal axis) for most legislators. For some legislators, the MCMC standard errors are much larger. The right panel shows that these legislators tend to have extreme ideological preferences: estimates from the Bayesian MCMC algorithm are shown with crosses and those from the proposed EM algorithm are shown with hollow circles.

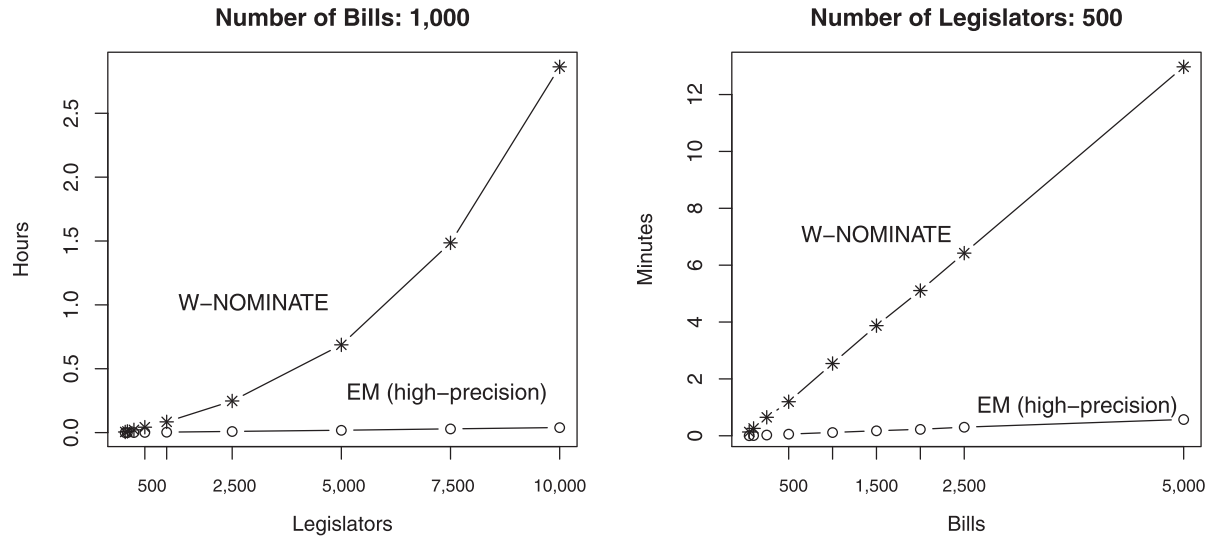
FIGURE 4. Bias of Standard Error based on the Parametric Bootstrap



Notes: The results are based on a Monte Carlo simulation where roll-call data are simulated using estimates from the 112th House of Representatives as truth. When simulating the data, the same missing data pattern as that in the data from the 112th Congress is used. A total of 1,000 roll call data sets are simulated, and each simulated data set is then bootstrapped 100 times to obtain standard errors. The estimated bias is computed as the average difference between the bootstrap standard error and the standard deviation of estimated ideal points across 1,000 simulations. The left panel shows that the estimated biases of the parametric bootstrap standard errors are not systematically related to ideological extremity. Instead, as the right panel shows, these biases are driven by the prevalence of missing data for legislators. The standard errors are significantly underestimated for those legislators with a large number of missing data.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 5. Comparison of Changing Performance across the Methods as the Dimensions of Roll-Call Matrix Increase



Notes: Estimation time is shown on the vertical axis as the number of legislators increases (left panel) and the number of bills increases (right panel). Values are the median times over 25 replications. "EM (high precision)" is more computationally efficient than W-NOMINATE (Poole et al. 2011) especially when the roll-call matrix is large.

Instead, as shown in the right panel of Figure 4, the biases of parametric bootstrap standard errors are driven by the amount of missing data. The plot shows that the standard errors are significantly underestimated for those legislators with a large amount of missing data.

Simulation Evidence

Next, we use Monte Carlo simulation to assess the computational scalability of the proposed EM algorithm as the dimensions of the roll-call matrix get larger. The empirical application in the previous subsection demonstrated the poor scalability of the MCMC algorithm. Hence, we compare the performance of our algorithm with W-NOMINATE. Figure 5 shows the median performance of both the EM algorithm with high precision and "W-NOMINATE" over 25 Monte Carlo trials for various numbers of legislators and bills. In the left panel, the number of bills is fixed at 1,000 and the number of legislators ranges from 50 to 10,000. In the right panel, the number of legislators is fixed at 500, and the number of bills ranges from 50 to 5,000. In both cases, the data-generating process follows the single dimensional ideal point model where ideal points are generated according to the standard normal distribution and the bill parameters follow from a normal distribution with mean zero and standard deviation of 10. These parameter values are chosen so that the true parameter values explain around 85% percent of the observed votes—a level of classification success comparable to the in-sample fit obtained in contemporary sessions of Congress.

The computational efficiency of the proposed algorithm can be seen immediately. Even with 10,000 legislators and 1,000 bills, convergence at high precision is achieved in less than 15 minutes. The runtime of our algorithm increases only linearly as the number of ideal points increases. This contrasts with W-NOMINATE, whose required computation time grows exponentially as the number of legislators increases. For example, both the EM algorithm and W-NOMINATE require less than 5 minutes to estimate the parameters associated with a roll-call matrix with 1,000 legislators and 1,000 bills. However, when the number of legislators increases to 10,000, W-NOMINATE takes around 2.5 hours while the EM algorithm requires less than 15 minutes. The difference is less stark when the number of bills increase (right panel). Even here, however, the EM algorithm is more computationally efficient, especially when the number of bills is large.

IDEAL POINT MODEL WITH MIXED BINARY, ORDINAL, AND CONTINUOUS OUTCOMES

We extend the EM algorithm developed above to the ideal point model with mixed binary, ordinal, and continuous outcomes. Quinn (2004) develops an MCMC algorithm for fitting this model and implements it as the `MCMCmixfactanal()` function in the open-source R package **MCMCpack** (Martin, Quinn, and Park 2013). The EM algorithm for the ordinal probit model, which is closely related to this model, poses a special challenge because its E step is not available in closed form. Perhaps for this reason, to the best of our knowledge,

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

the EM algorithm has not been developed for the ordinal probit model in the statistics literature.

In this section, we first show that the E step can be derived in a closed form so long as the outcome variable only has three ordinal categories. With a suitable transformation of parameters, we derive an EM algorithm that is analytically tractable. We then consider the cases where the number of categories in the outcome variable exceeds three and the outcome variable is a mix of binary, ordinal, and continuous variables. Finally, we apply the proposed algorithm to a survey of Japanese politicians and voters.

The Model with a Three-category Ordinal Outcome

We consider the same exact setup as the standard model introduced above, with the exception that the outcome variable now takes one of the three ordered values, i.e., $y_{ij} \in \{0, 1, 2\}$. In this model, the probability of each observed choice is given as follows:

$$\Pr(y_{ij} = 0) = \Phi(\alpha_{1j} - \mathbf{x}_i^\top \boldsymbol{\beta}_j), \tag{9}$$

$$\Pr(y_{ij} = 1) = \Phi(\alpha_{2j} - \mathbf{x}_i^\top \boldsymbol{\beta}_j) - \Phi(\alpha_{1j} - \mathbf{x}_i^\top \boldsymbol{\beta}_j), \tag{10}$$

$$\Pr(y_{ij} = 2) = 1 - \Phi(\alpha_{2j} - \mathbf{x}_i^\top \boldsymbol{\beta}_j), \tag{11}$$

where $\alpha_{2j} > \alpha_{1j}$ for all $j = 1, 2, \ldots, J$. The model can be written using the latent propensity to agree y_{ij}^* for respondent i as

$$y_{ij}^* = \mathbf{x}_i^\top \boldsymbol{\beta}_j + \epsilon_{ij}, \tag{12}$$

where $\epsilon_{ij} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ and these latent propensities are connected to the observed outcomes through the following relationship:

$$y_{ij} = \begin{cases} 0 & \text{if } y_{ij}^* < \alpha_{1j} \\ 1 & \text{if } \alpha_{1j} \le y_{ij}^* < \alpha_{2j} \\ 2 & \text{if } \alpha_{2j} \le y_{ij}^* \end{cases}. \tag{13}$$

As in the standard ideal point model, we treat abstention as missing at random.

Following the literature, we assume the same normal independent prior distribution on $(\{\boldsymbol{\beta}_j\}_{j=1}^J, \{\mathbf{x}_i\}_{i=1}^N)$ as the one used for the standard binary model. For $(\{\alpha_{1j}, \alpha_{2j}\}_{j=1}^J)$, we assume an improper uniform prior with the appropriate ordering restriction $\alpha_{1j} < \alpha_{2j}$. Then, the joint posterior distribution is given by

$$\begin{aligned} & p(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\alpha_{1j}, \alpha_{2j}, \boldsymbol{\beta}_j\}_{j=1}^J \mid \mathbf{Y}) \\ & \propto \prod_{i=1}^N \prod_{j=1}^J \big[\mathbf{1}\{y_{ij}^* < \alpha_{1j}\}\mathbf{1}\{y_{ij} = 0\} \\ & \quad + \mathbf{1}\{\alpha_{1j} \le y_{ij}^* < \alpha_{2j}\}\mathbf{1}\{y_{ij} = 1\} \\ & \quad + \mathbf{1}\{y_{ij}^* \ge \alpha_{2j}\}\mathbf{1}\{y_{ij} = 2\}\big]\, \phi_1(y_{ij}^*; \mathbf{x}_i^\top \boldsymbol{\beta}_j, 1) \\ & \quad \times \prod_{i=1}^N \phi_K(\mathbf{x}_i; \boldsymbol{\mu}_{\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{x}}) \prod_{j=1}^J \phi_K(\boldsymbol{\beta}_j; \boldsymbol{\mu}_{\boldsymbol{\beta}}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}}). \end{aligned} \tag{14}$$

The Proposed Algorithm

To develop an EM algorithm that is analytically tractable, we employ the following one-to-one transformation of parameters:

$$\tau_j = \alpha_{2j} - \alpha_{1j} > 0, \tag{15}$$

$$\alpha_j^* = -\frac{\alpha_{1j}}{\tau_j}, \tag{16}$$

$$\boldsymbol{\beta}_j^* = \frac{\boldsymbol{\beta}_j}{\tau_j}, \tag{17}$$

$$z_{ij}^* = \frac{y_{ij}^* - \alpha_{1j}}{\tau_j}, \tag{18}$$

$$\epsilon_{ij}^* = \frac{\epsilon_{ij}}{\tau_j}. \tag{19}$$

Then, the simple algebra shows that the model can be rewritten as

$$\Pr(y_{ij} = 0) = \Phi(-\tau_j \alpha_j^* - \tau_j \mathbf{x}_i^\top \boldsymbol{\beta}_j^*), \tag{20}$$

$$\begin{aligned} \Pr(y_{ij} = 1) = {} & \Phi(-\tau_j \alpha_j^* + \tau_j - \tau_j \mathbf{x}_i^\top \boldsymbol{\beta}_j^*) \\ & - \Phi(-\tau_j \alpha_j^* - \tau_j \mathbf{x}_i^\top \boldsymbol{\beta}_j^*), \end{aligned} \tag{21}$$

$$\Pr(y_{ij} = 2) = 1 - \Phi(-\tau_j \alpha_j^* + \tau_j - \tau_j \mathbf{x}_i^\top \boldsymbol{\beta}_j^*), \tag{22}$$

where the latent variable representation is given by

$$z_{ij}^* = \alpha_j^* + \mathbf{x}_i^\top \boldsymbol{\beta}_j^* + \epsilon^* \text{ with } \epsilon_{ij}^* \overset{\text{indep.}}{\sim} \mathcal{N}(0, \tau_j^{-2}). \tag{23}$$

Under this parametrization, the relationship between the observed outcome y_{ij} and the latent variable z_{ij}^* is given by

$$y_{ij} = \begin{cases} 0 & \text{if } z_{ij}^* < 0 \\ 1 & \text{if } 0 \le z_{ij}^* < 1 \\ 2 & \text{if } 1 \le z_{ij}^* \end{cases}. \tag{24}$$

Thus, the consequence of this reparametrization is that the threshold parameters $(\alpha_{1j}, \alpha_{2j})$ are replaced with the intercept term α_j^* and the heterogeneous variance parameter τ_j^{-2}.

To maintain conjugacy, we alter the prior distribution specified in equation (14). In particular, we use the

https://doi.org/10.1017/S0003055416000037X Published online by Cambridge University Press

following prior distribution:

$$p\left(\{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J, \{\tau_j^2\}_{j=1}^J, \{\mathbf{x}_i\}_{i=1}^N\right) = \prod_{j=1}^J \phi_{K+1}(\tilde{\boldsymbol{\beta}}_j; \boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}) \times \mathcal{G}\left(\tau_j^2; \frac{\nu_\tau}{2}, \frac{s_\tau}{2}\right) \prod_{i=1}^N \phi_K(\mathbf{x}_i; \boldsymbol{\mu}_{\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{x}}), \quad (25)$$

where $\tilde{\boldsymbol{\beta}}_j = (\alpha_j^*, \boldsymbol{\beta}_j^*)$ and $\mathcal{G}(\tau_j^2; \nu_\tau/2, s_\tau/2)$ is the Gamma distribution with $\nu_\tau/2 > 0$ and $s_\tau/2 > 0$ representing the prior shape and rate parameters, respectively. This change in prior distribution alters the model but so long as the prior is diffuse and the size of data is large the resulting inference should not differ much.

Given this setup, we derive the EM algorithm to maximize the posterior distribution. We take the analytical strategy similar to the one used for the standard ideal point model. The resulting algorithm is described in the final section.

Mixed Binary, Ordinal, and Continuous Outcomes

Here we consider how to apply, possibly after some modification, the EM algorithm developed above to a more generel case with mixed binary, ordinal, and continuous outcomes. If the number of ordered categories in outcome exceeds 3, we collapse them into three categories for the sake of analytical tractability and computational efficiency.[7] For example, responses to a survey question on the five-point Likert scale, i.e., `strongly disagree, disagree, neither agree nor disagree, agree, strongly agree`, may be converted into a three-point scale by combining `strongly disagree` and `disagree` into a single category and `strongly agree` and `agree` into another category. Researchers must carefully judge the extent of tradeoff between the loss of information and the computational speed for each application.

Next, suppose that y_{ij} is a binary outcome for a particular observation (i, j). Then, we consider the following relationship between this observed outcome and the latent propensity z_{ij}^*; $z_{ij} < 1 \Longleftrightarrow y_{ij} = 0$ and $z_{ij} \geq 0 \Longleftrightarrow y_{ij} = 1$. Under this assumption, the E step becomes

$$z_{ij}^{*(t)} = \begin{cases} m_{ij}^{(t-1)} - \frac{1}{\tau_j^{(t-1)}}\lambda\left(-m_{ij}^{(t-1)} + 1, \tau_j^{(t-1)}\right) & \text{if } y_{ij} = 0 \\ m_{ij}^{(t-1)} + \frac{1}{\tau_j^{(t-1)}}\lambda\left(m_{ij}^{(t-1)}, \tau_j^{(t-1)}\right) & \text{if } y_{ij} = 1 \end{cases} \quad (26)$$

and

$$\left(z_{ij}^{*2}\right)^{(t)} = \begin{cases} \left(z_{ij}^{*(t)}\right)^2 + \frac{1}{\left(\tau_j^{(t-1)}\right)^2}\left[1 - \frac{1-m_{ij}^{(t-1)}}{\tau_j^{(t-1)}}\lambda \times \left(m_{ij}^{(t-1)} - 1, \tau_j^{(t-1)}\right) - \left\{\lambda\left(m_{ij}^{(t-1)} - 1, \tau_j^{(t-1)}\right)\right\}^2\right] & \text{if } y_{ij} = 0 \\ \left(z_{ij}^{*(t)}\right)^2 + \frac{1}{\left(\tau_j^{(t-1)}\right)^2}\left[1 - \lambda\left(-m_{ij}^{(t-1)}, \tau_j^{(t-1)}\right) \times \left\{\lambda\left(-m_{ij}^{(t-1)}, \tau_j^{(t-1)}\right) + m_{ij}^{(t-1)}\tau_j^{(t-1)}\right\}\right] & \text{if } y_{ij} = 1 \end{cases} \quad (27)$$

where if y_{ij} is missing, we set $z_{ij}^{*(t)} = m_{ij}^{(t-1)}$ and $(z_{ij}^{*2})^{(t)} = (z_{ij}^{*(t)})^2 + (\tau_j^{(t-1)})^{-2}$. Other than these modifications, the rest of the EM algorithm stays identical.

Finally, it is straightforward to extend this model to also include a continuous outcome as done by Quinn (2004). In that case, set the first and second moments of the latent propensity as $z_{ij}^{*(t)} = y_{ij}$ and $(z_{ij}^{*2})^{(t)} = y_{ij}^2 + (\tau_j^{(t-1)})^{-2}$ for this observation. The rest of the EM algorithm remains unchanged.

An Empirical Application

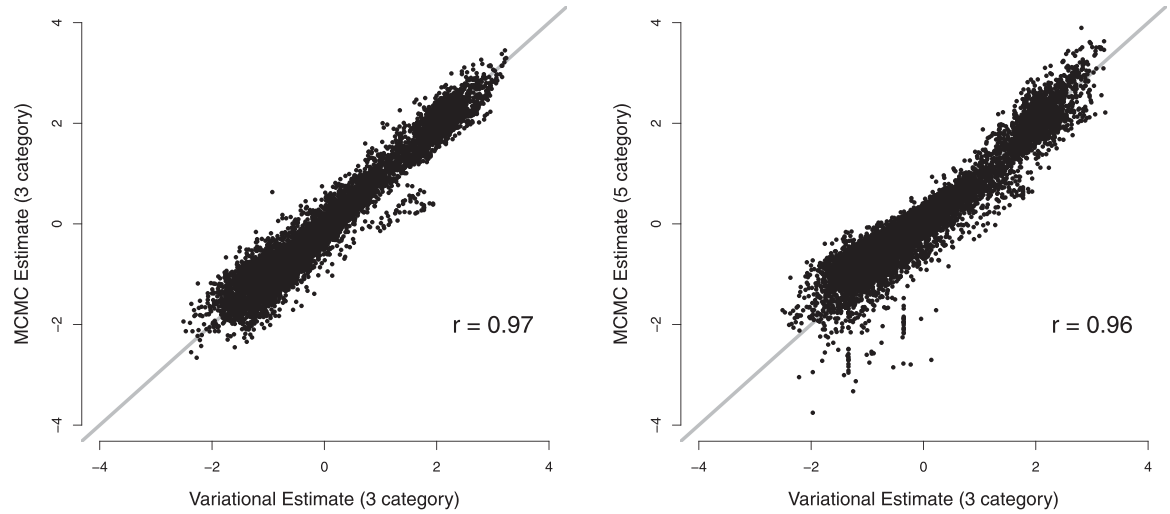
We apply the ordinal ideal point model to the survey data of the candidates and voters of Japanese Upper and Lower House elections. This Asahi-Todai Elite survey was conducted by the University of Tokyo in collaboration with a major national newspaper, the Asahi Shimbun, covering all candidates (both incumbents and challengers) for the eight elections that occurred between 2003 and 2013. In six out of eight waves, the survey was also administered to a nationally representative sample of voters with the sample size ranging from approximately 1,100 to about 2,000. The novel feature of the data is that there are a set of common policy questions, which can be used to scale both politicians and voters over time on the same dimension. Another important advantage of the data is a high response rate among politicians, which exceeded 85%. Such a high response rate is obtained in large part because the survey results are published in the Asahi Shimbun whose circulation is approximately eight million (see Hirano et al. 2011, for more details).

All together, the data set we analyze contains a total of $N = 19{,}443$ respondents, including 7,734 politicians and 11,709 voters. Here, we count multiple appearances of the same politician separately because an ideal point will be estimated separately for each wave. There are $J = 98$ unique questions in the survey, most of which consisted of questions asking for responses on a five-point Likert scale. We apply the proposed EM

[7] In theory, one could develop a variational EM algorithm to deal with more than three categories in the ordinal ideal point model. However, unlike the three category case presented in this article, it would require another layer of approximation and hence one must examine the appropriateness of that approximation.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 6. Comparison of Ideal Point Estimates from the EM and Markov Chain Monte Carlo (MCMC) Algorithms for Japanese Politicians Using the Asahi-Todai Elite Survey



Notes: The figures compare the EM estimates (horizontal axis) against MCMC estimates (vertical axis). The EM estimates use a coarsened three category response, which is compared against the MCMC estimates based on the same three category response (left panel) and the original five category response (right panel). Overall correlation between the EM and MCMC estimates are high, exceeding 0.95 in both cases.

algorithm after coarsening each response into three categories (disagree, neutral, and agree). For the purpose of comparison, we developed and used a customized C language implementation of the MCMC algorithm. The data set in this application was too large for `MCMCmixfactanal()` from the R package **MCMCpack** to handle. One model uses the full range of categories found in the data without coarsening, and the other model uses the same coarsened responses as done for our algorithm. Obtaining 10,000 draws from the posterior distribution using the MCMC algorithm takes 4 hours and 24 minutes (5 category) or 3 hours and 54 minutes (3 category). In contrast, estimation using our proposed EM algorithm takes 164 iterations and only 68 seconds to complete where the algorithm is iterated until the correlation of parameter values between two consecutive iterations reaches $1 - 10^{-6}$.

Figure 6 compares the estimated ideal points of politicians based on our EM algorithm (vertical axis) against those obtained from the standard MCMC algorithm (horizontal axis). As explained earlier, the EM estimates are based on the coarsened three category response, while we present the MCMC estimates using the original five category response (right panel) as well as the same coarsened three category response (left panel). The plots show that these two algorithms produce essentially identical estimates, achieving a correlation greater than 0.95. In addition, for this data set, coarsening the original five category response into a three category response does not appear to have a significant impact on the degree of correlation between the two sets of ideal point estimates of Japanese politicians.

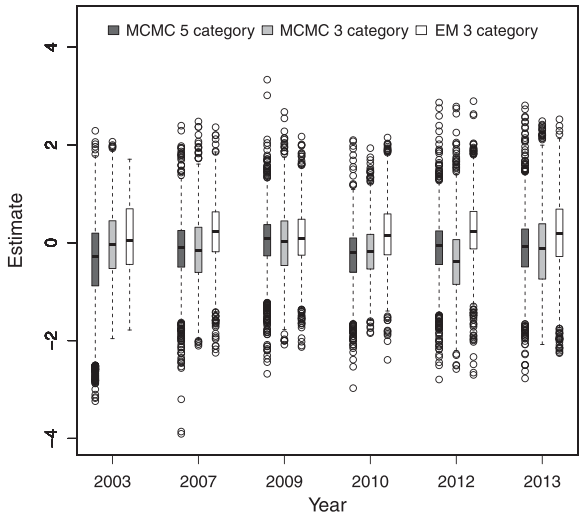
Figure 7 compares the estimated ideal points of voters for each wave of survey, obtained from our EM algorithm (white box plots) and the standard MCMC algorithm (light and dark grey box plots for the coarsened three category and original five category responses, respectively). Across all six waves of the survey, the three algorithms give similar distributions of estimated ideal points. The differences across the algorithms lie mostly in their estimated ideal points for a small subset of voters who answer too few questions. For example, the 2003 survey included only two policy questions, and 306 respondents from this survey gave the same two responses. For these respondents, our EM algorithm produces an identical ideal point estimate of -0.89 whereas the MCMC algorithm gives a set of ideal points ranging from about -3 to 0, mainly due to the imprecise nature of posterior mean point estimates when votes are not informative. Overall, the results suggest that our EM algorithm recovers virtually identical estimates to those derived via the standard MCMC algorithm but with substantial savings in time.

DYNAMIC IDEAL POINT MODEL

We next consider the dynamic ideal point model of Martin and Quinn (2002), who characterized how the ideal points of supreme court justices change over time. The authors develop an MCMC algorithm for fitting this model and make it available as the `MCMCdynamicIRT1d()` function in the open-source R package **MCMCpack** (Martin, Quinn, and Park 2013). This methodology is based on the dynamic linear

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 7. Comparing the Distributions of Estimated Ideal Points between the EM and Markov Chain Monte Carlo (MCMC) Algorithms for Japanese Voters across Six Waves of the Asahi-Todai Elite Survey



Notes: White box plots describe the distribution of the EM estimates whereas light and dark grey box plots represent the MCMC estimates for the coarsened three category and original five category responses, respectively. Across all waves, these three algorithms produce similar estimates of ideal points.

modeling approach and is more flexible than polynomial time trend models considered by other scholars (see, e.g., DW-NOMINATE, Bailey 2007).

Nevertheless, this flexibility comes at a significant computational cost. In particular, Martin and Quinn (2002) report that using a dedicated workstation it took over 6 days to estimate ideal points for U.S. Supreme Court justices over 47 years (footnote 12). Because of this computational burden, Bailey (2007) resorts to a simpler parametric dynamic model (p. 441). In addition, unlike the two models we considered above, no closed-form EM algorithm is available for maximizing the posterior in this case. Therefore, we propose use of variational inference that approximates posterior inference by deriving a variational EM algorithm. We show that the proposed algorithm is orders of magnitude faster than the standard MCMC algorithm and scales to a large data set while yielding the estimates that are similar to those obtained from the standard MCMC algorithm.

The Model

Let y_{ijt} be an indicator variable representing the observed vote of legislator i on roll call j at time t where $y_{ijt} = 1$ ($y_{ijt} = 0$) represents "yea" ("nay"). There are a total of N unique legislators, i.e., $i = 1, \ldots, N$, and for any given time period t, there are J_t roll calls, i.e., $j = 1, \ldots, J_t$. Finally, the number of time periods is T, i.e., $t = 1, \ldots, T$. Then, the single-dimensional ideal point model is given by

$$\Pr(y_{ijt} = 1) \;=\; \Phi(\alpha_{jt} + \beta_{jt} x_{it}) \;=\; \Phi(\tilde{\mathbf{x}}_{it}^{\top} \tilde{\boldsymbol{\beta}}_{jt}), \quad \textbf{(28)}$$

where x_{it} is justice i's ideal point at time t, and α_{jt} and β_{jt} represent the item difficulty and item discrimination parameters for roll call j in time t, respectively. Note that as before we use a vector notation $\tilde{\mathbf{x}}_{it} = (1, x_{it})$ and $\tilde{\boldsymbol{\beta}}_{jt} = (\alpha_{jt}, \beta_{jt})$. As before, the model can be rewritten with the latent propensity y^*_{ijt},

$$y^*_{ijt} \;=\; \tilde{\mathbf{x}}_{it}^{\top} \tilde{\boldsymbol{\beta}}_{jt} + \epsilon_{ijt}, \quad \textbf{(29)}$$

where $\epsilon_{ijt} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ and $y_{ijt} = 1$ ($y_{ijt} = 0$) if $y^*_{ijt} > 0$ ($y^*_{ijt} \leq 0$).

As done in the standard dynamic linear modeling framework, the dynamic aspect of the ideal point estimation is specified through the following random walk prior for each legislator i:

$$x_{it} \mid x_{i,t-1} \overset{\text{indep.}}{\sim} \mathcal{N}(x_{i,t-1},\ \omega_x^2) \quad \textbf{(30)}$$

for $t = \underline{T}_i, \underline{T}_i + 1, \ldots, \overline{T}_i - 1, \overline{T}_i$ where $\underline{T}_i$ is the first time period legislator i enters the data and $\overline{T}_i$ is the last time period the legislator appears in the data, i.e., $1 \leq \underline{T}_i \leq \overline{T}_i \leq T$. In addition, we assume $x_{i,\underline{T}_i - 1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_x, \Sigma_x)$ for each legislator i.

Finally, given this setup, with the independent conjugate prior on $\tilde{\boldsymbol{\beta}}_{jt}$, we have the following joint posterior distribution:

$$\begin{aligned} & p(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{t=1}^T \mid \mathbf{Y}) \\ & \propto \prod_{i=1}^{N} \prod_{t=\underline{T}_i}^{\overline{T}_i} \prod_{j=1}^{J_t} \big(\mathbf{1}\{y^*_{ijt} > 0\} \mathbf{1}\{y_{ijt} = 1\} \\ & \quad + \mathbf{1}\{y^*_{ijt} \leq 0\} \mathbf{1}\{y_{ijt} = 0\} \big)\ \phi_1\left(y^*_{ijt}; \tilde{\mathbf{x}}_{it}^{\top} \tilde{\boldsymbol{\beta}}_{jt}, 1\right) \\ & \quad \times \prod_{i=1}^{N} \left\{ \phi_1(x_{i,\underline{T}_i - 1}; \mu_x, \Sigma_x) \prod_{t=\underline{T}_i}^{\overline{T}_i} \phi_1(x_{it}; x_{i,t-1}, \omega_x^2) \right\} \\ & \quad \times \prod_{t=1}^{T} \prod_{j=1}^{J_t} \phi_2(\tilde{\boldsymbol{\beta}}_{jt}; \boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}), \end{aligned} \quad \textbf{(31)}$$

where $\mathbf{x}_i = (x_{i\underline{T}_i}, \ldots, x_{i\overline{T}_i})$ for $i = 1, \ldots, N$.

We propose a variational EM algorithm for the dynamic ideal point model summarized above. The variational inference makes factorization assumptions and approximates the posterior inference by minimizing the Kullback-Leibler divergence between the true posterior distribution and the factorized distribution (see Wainwright and Jordan (2008) for a review and Grimmer (2011) for an introductory article in political science). In the current context, we assume the following

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

factorization assumption:

$$q(\mathbf{Y}^*, \{\mathbf{x}_i\}_{i=1}^N, \{\tilde{\boldsymbol{\beta}}_j\}_{t=1}^T) = \prod_{i=1}^{N}\prod_{t=\underline{T}_i}^{\overline{T}_i} q(y_{it}^*) \prod_{i=1}^{N} q(\mathbf{x}_i) \prod_{t=1}^{T}\prod_{j=1}^{J_t} q(\tilde{\beta}_{jt}) \quad (32)$$

which basically assumes the independence across parameters. Importantly, we do not assume the independence between x_{it} and $x_{it'}$ so that we do not sacrifice our ability to model dynamics of ideal points. We also do not assume a family of approximating distributions. Rather, our results show that the optimal variational distribution belongs to a certain parametric family. The proposed variational EM algorithm is described in the final section while the detailed derivation is given in Appendix C.[8]

An Empirical Application

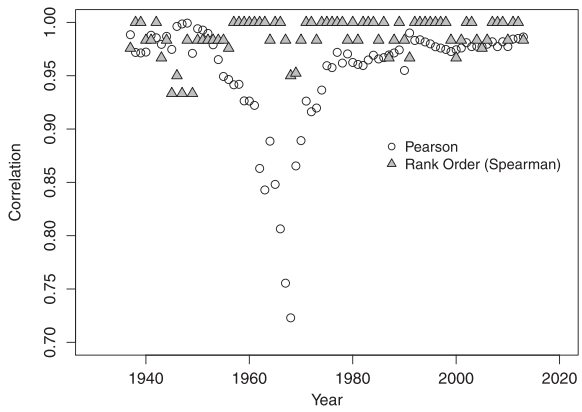
We apply the proposed variational EM algorithm for estimating the dynamic ideal point to the voting data from the U.S. Supreme court (from October 1937 to October 2013). The data set includes 5,164 votes on court cases by 45 distinct justices over 77 terms, resulting in the estimation of 697 unique ideal points for all justice-term combinations. The same data set was used to compute the ideal point estimates published as the well-known Martin-Quinn scores at http://mqscores.berkeley.edu/ (July 23, 2014 Release version).

We set the prior parameters using the replication code, which was directly obtained from the authors. In particular, the key random-walk prior variance parameter ω_x^2 is set to be equal to 0.1 for all justices. Note that this choice differs from the specification in Martin and Quinn (2002) where Douglas' prior variance parameter was set as $\omega_x^2 = 0.001$ because of his ideological extremity and the small number of cases he heard towards the end of his career. This means that Douglas's ideal point estimate is fixed at his prior mean of -3.0, but in the results we report below this constraint is not imposed.

We use the same prior specification and apply the proposed variational EM algorithm as well as the standard MCMC algorithm implemented via the `MCMCdynamicIRT1d()` function from **MCMCpack**. For the MCMC algorithm, using the replication code, 1.2 million iterations took just over 5 days of computing time. In contrast, our variational EM algorithm took only four seconds. To obtain a measure of estimation uncertainty, we use the parametric bootstrap approach of Lewis and Poole (2004) to create 100 replicates and construct bias-corrected 95% bootstrap confidence intervals. Note that even with this bootstrap procedure, the computation is done within several minutes.

We begin by examining, for each term, the correlation of the resulting estimated ideal points for nine justices between the proposed variational inference algorithm and MCMC algorithm. Figure 8 presents both Pearson's correlations and Spearman's rank-order correlations. Overall, the correlations are high, exceeding 95% in most cases. In particular, for many terms, rank-order correlations are equal to unity, indicating that the two algorithms produce justices' estimated ideal points whose rank order is identical. We note that a significant drop in Pearson correlation between 1960 and 1975 is driven almost entirely by the extreme MCMC estimates of Douglas' position in these years, which correspond to the final years of his tenure. And yet, even in these years, the rank-order correlations remain high.

FIGURE 8. Correlation of the Estimated Ideal Points for each Term between the Variational EM and Markov Chain Monte Carlo (MCMC) Algorithms



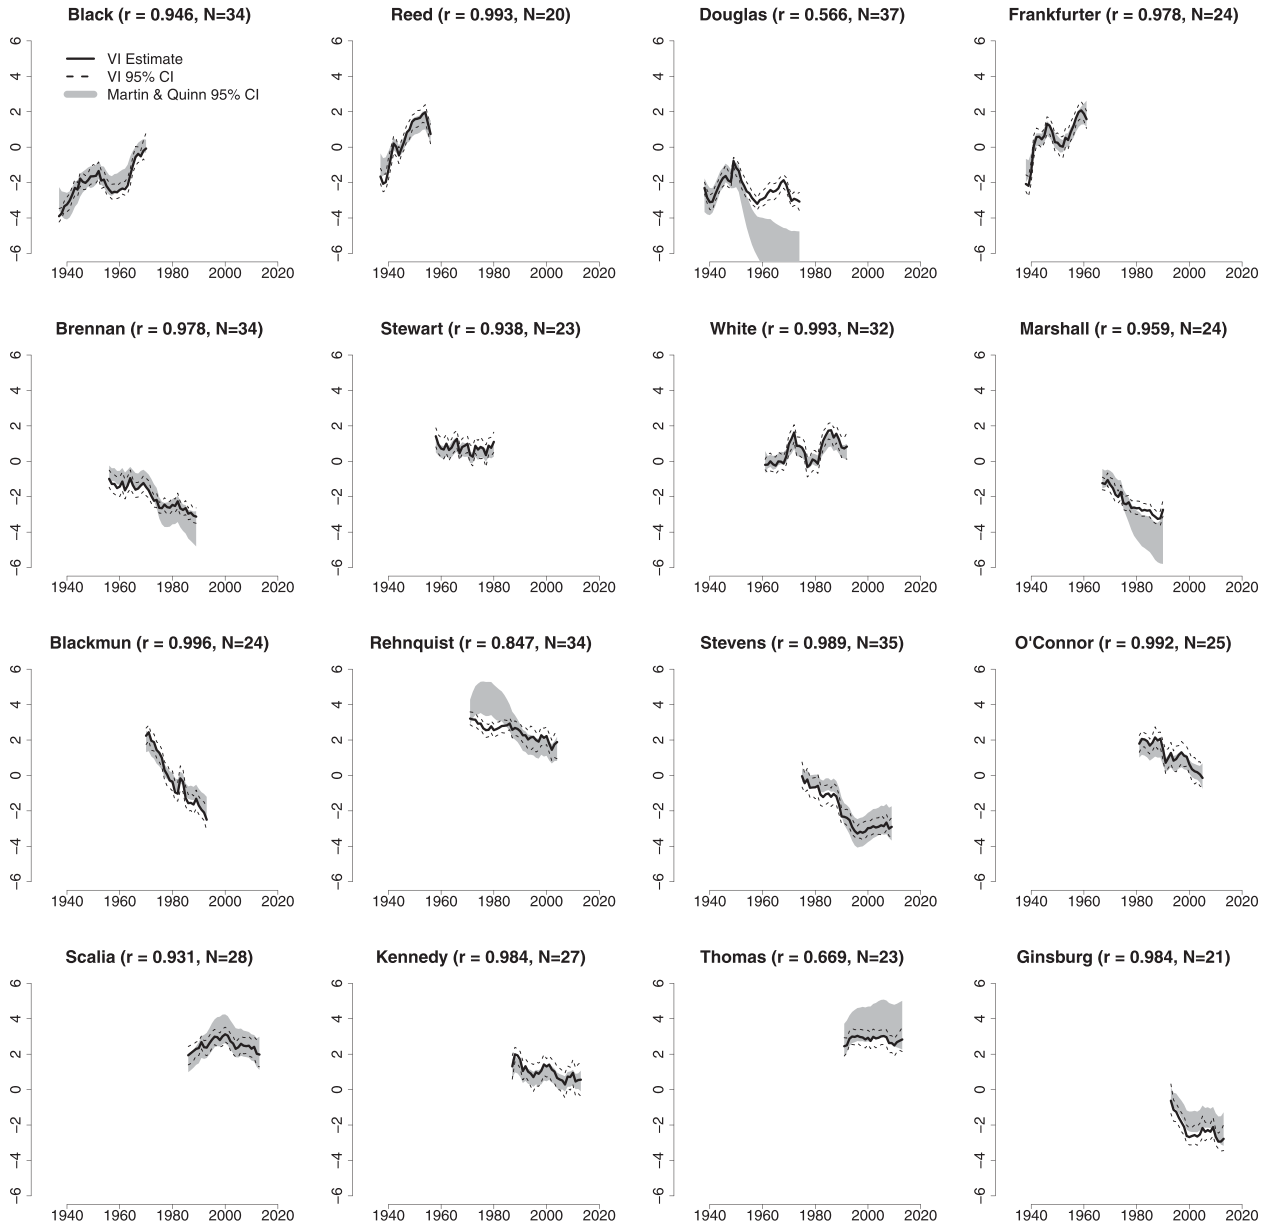
Notes: Open circles indicate Pearson correlations, while grey triangles represent Spearman's rank-order correlations. Overall, the correlations are high, exceeding 95% in most cases. The poorer Pearson correlations around 1969 are driven largely by Douglas' ideological extremity (see Figure 9).

Figures 9 present time series plots of the estimated ideal points for the 16 justices who served the longest periods of time in our study. Solid lines indicate the variational estimates, while the dashed lines indicate their 95% confidence intervals based on the parametric bootstrap. The grey polygons represent the 95% credible intervals obtained from the MCMC algorithm. For almost all justices, the movement of estimated ideal points over time is similar between the two algorithms. Indeed, for most justices, the correlation between the two sets of estimates is high, often exceeding .95. A notable exception to this is Douglas, where we see his ideal point estimates based on the MCMC algorithm becomes more extreme as time passes. The behavior observed here is consistent with earlier warnings about Douglas' ideological extremity and the fact that he cast only a small number of votes in the final years of his career (Martin and Quinn 2002). The correlation across all ideal points between the two sets of estimates increases from 0.93 to 0.96, once we exclude Douglas. Overall, our proposed variational inference algorithm

[8] For completeness, we also derive the variational EM algorithms for the standard ideal point model (Appendix A) and the ideal point model with an ordinal outcome (Appendix B).

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 9. Ideal Point Estimates for 16 Longest-serving Justices based on the Variational Inference (VI) and Markov Chain Monte Carlo (MCMC) Algorithm



Notes: The VI point estimates are indicated by solid lines while the dashed lines indicate its 95% confidence intervals based on the parametric bootstrap. We also present the 95% Bayesian confidence intervals as grey polygons. The horizontal axis indicates year and the vertical axis indicates estimated ideal points. For each justice, we also compute the Pearson's correlation between the two sets of the estimates. Overall, the correlations between the two sets of estimates are high except Douglas who is ideologically extreme and has only a small number of votes in the final years of his career.

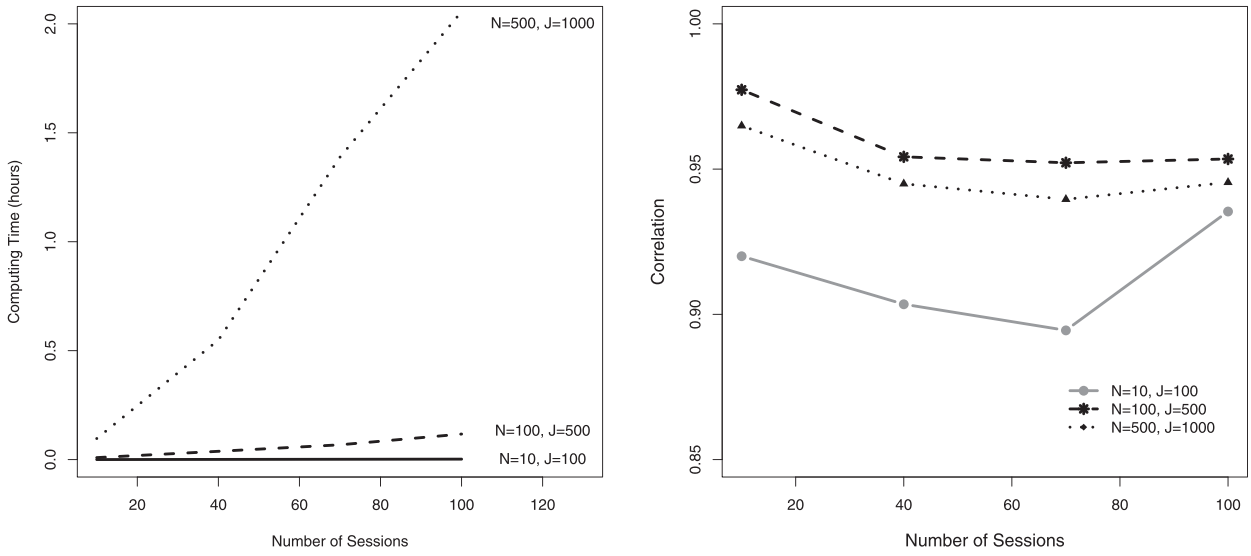
produces the estimates of ideal points that are close to the Martin-Quinn score but with significantly less computing time.

Simulation Evidence

We further demonstrate the computational scalability of the proposed variational EM algorithm through a series of simulations. We generate a number of roll call matrices that vary in size. These include roll call matrices that have $N = 10$ legislators and $J = 100$ roll calls per session (roughly corresponding to the size of the U.S. Supreme Court), roll-call matrices with $N = 100$ and $J = 500$ (roughly corresponding to the size of the U.S. Senate), and roll-call matrices with $N = 500$ and $J = 1{,}000$ (roughly corresponding in size to the U.S. House). We also vary the total number of sessions,

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 10. Scalability and Accuracy of the Proposed Variational Inference for the Dynamic Ideal Point Model



Notes: The left panel presents run times of the proposed variational EM algorithm for fitting the dynamic ideal point model. We consider three different simulation scenarios where the number of legislators N varies from 10 to 500 and the number of roll calls per session J ranges from 100 to 1,000. The number of sessions T is shown on the horizontal axis, with all N legislators assumed to vote on all J bills in every session. The vertical axis indicates the time necessary to fit the dynamic ideal point model for each data set through the proposed algorithm. Even with the largest data set we consider ($N = 500$, $J = 1{,}000$, and $T = 100$), the algorithm can estimate a half million ideal points in about two hours. The right panel shows the (Pearson) correlation between the estimated ideal points and their true values. In almost all cases, the correlation exceeds 0.95.

ranging from $T = 10$ to $T = 100$. Thus, the largest roll-call matrix represents a scenario that all members of the U.S. House vote on 1,000 bills during each of 100 consecutive sessions! As we show next, even in this extreme case, our algorithm runs in about 25 minutes, yielding the estimated ideal points that are close to the true values.

We then apply our variational EM algorithm and record the amount of time needed to estimate the model, as well as correlation between the true and recovered ideal points. In the simulation, all legislators serve throughout all periods, whose ideal points in the first period follow the standard normal distribution. Independence across legislators is assumed as done in the model, and their subsequent ideal points are generated as a random walk with $\omega_x^2 = 0.1$ for all legislators. Item difficulty and discrimination parameters in all sessions were drawn from uniform $(-1.5, 1.5)$ and $(-5.5, 5.5)$ distributions respectively. While parallelization of the algorithm is trivial and would further reduce run times, we do not implement it for this calculation. As before, convergence is assumed to be achieved when correlation across all parameters across consecutive iterations is greater than $1 - 10^{-6}$.[9]

The left panel of Figure 10 displays the amount of time needed for each simulation, with the total number of sessions T given on the horizontal axis. As a benchmark comparison, MCMC replication code provided by Martin and Quinn (2002) took over five days to estimate ideal points for U.S. Supreme Court justices over 77 years ($N = 45$, $T = 77$, and $J = 5{,}164$). For the scenario with $N = 10$ legislators and $J = 100$ roll calls per session, estimation is completed under a minute regardless of the number of sessions. Similarly, for the scenarios with 100 legislators and 500 roll calls per session, computation is completed in a matter of minutes regardless of the number of sessions. Computation only begins to significantly increase with our largest scenario of 500 legislators and 1,000 roll calls per session. But even here, for 100 sessions, the variational EM algorithm converges in under 25 minutes.

The right panel of the figure presents, for each simulation scenario, the correlation between the variational estimates of ideal points and their true values across all legislators and sessions. The plot demonstrates that the correlation exceeds .95 throughout all the simulations except the case where the size of roll call matrix is small. Even in this case, the correlation is about .90, which suggests the reasonable accuracy of the variational estimates under the dynamic ideal point model.

HIERARCHICAL IDEAL POINT MODEL

Finally, we consider the hierarchical ideal point model where the ideal points are modeled as a linear function of covariates (Bafumi et al. 2005). Like the

[9] To reduce Monte Carlo error, estimates for cases where $N = 10$ are repeated 25 times, with median run times and correlations reported in the figure.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

dynamic ideal point model, there is no closed-form EM algorithm that directly maximizes the posterior distribution. Therefore, we apply variational inference to approximate the posterior distribution. We derive the variational EM algorithm and demonstrate its computational efficiency and the accuracy of approximation through empirical and simulation studies.

The Model

Let each distinct vote be denoted by the binary random variable y_l where there exist a total of L such votes, i.e., $\ell \in \{1, \ldots, L\}$. Each vote y_ℓ represents a vote cast by legislator $i[\ell]$ on bill $j[\ell]$ ($y_\ell = 1$ and $y_\ell = 0$ representing "yea" and "nay," respectively) where $i[\ell] \in \{1, \ldots, N\}$ and $j[\ell] \in \{1, \ldots, J\}$. Thus, there are a total of N legislators and J bills. Finally, let $g[i[\ell]]$ represent the group membership of legislator $i[\ell]$ where $g[i[\ell]] \in \{1, \ldots, G\}$ and G indicates the total number of groups.

The hierarchical model we consider has the following latent variable structure with the observed vote written as $y_\ell = \mathbf{1}\{y^*_\ell > 0\}$ as before:

$$y^*_\ell = \alpha_{j[\ell]} + \beta_{j[\ell]} x_{i[\ell]} + \epsilon_\ell \quad \text{where} \quad \epsilon_\ell \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \tag{33}$$

$$x_{i[\ell]} = \boldsymbol{\gamma}^\top_{g[i[\ell]]} \mathbf{z}_{i[\ell]} + \eta_{i[\ell]} \text{ where } \eta_{i[\ell]} \overset{\text{indep.}}{\sim} \mathcal{N}\left(0, \sigma^2_{g[i[\ell]]}\right), \tag{34}$$

where $\boldsymbol{\gamma}_{g[i[\ell]]}$ is an M-dimensional vector of group-specific coefficients, $\mathbf{z}_{i[\ell]}$ is an M-dimensional vector of legislator-specific covariates, which typically includes one for an intercept, and $\sigma^2_{g[i[\ell]]}$ is the group-specific variance.

One important special case of this model is a dynamic ideal point model with a parametric time trend, an approach used to compute DW-NOMINATE scores (Poole and Rosenthal 1997) and adopted by some scholars (e.g., Bailey 2007). In this case, the $i[\ell]$ represents a legislator session, e.g., John Kerry in 2014, and $g[i[\ell]]$ indicates the legislator, John Kerry, whereas $\mathbf{z}_{i[\ell]}$ may include the number of sessions since the legislator took office as well as a constant for the intercept term. Then, the ideal points are modeled with a linear time trend. In addition, including the square term will allow one to model ideal points using a quadratic time trend. Note that in this setting the time trend is estimated separately for each legislator.

The model is completed with the following conjugate prior distribution:

$$\tilde{\boldsymbol{\beta}}_{j[\ell]} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}), \tag{35}$$

$$\boldsymbol{\gamma}_{g[i[\ell]]} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{\gamma}}, \boldsymbol{\Sigma}_{\boldsymbol{\gamma}}), \tag{36}$$

$$\sigma^2_{g[i[\ell]]} \overset{\text{i.i.d.}}{\sim} \mathcal{IG}\left(\frac{\nu_\sigma}{2}, \frac{s^2_\sigma}{2}\right), \tag{37}$$

where $\tilde{\boldsymbol{\beta}}_{j[\ell]} = (\alpha_{j[\ell]}, \boldsymbol{\beta}_{j[\ell]})$ and $\mathcal{IG}(\nu, s^2)$ represents the inverse-gamma distribution with scale and shape parameters equal to ν and s^2, respectively.

It is convenient to rewrite the model in the following reduced form:

$$y^*_\ell = \alpha_{j[\ell]} + \beta_{j[\ell]} \boldsymbol{\gamma}^\top_{g[i[\ell]]} \mathbf{z}_{i[\ell]} + \beta_{j[\ell]} \eta_{i[\ell]} + \epsilon_\ell \tag{38}$$

Then, the joint posterior distribution is given by

$$\begin{aligned} & p(\mathbf{Y}^*, \{\tilde{\boldsymbol{\beta}}_k\}_{k=1}^J, \{\boldsymbol{\gamma}_m\}_{m=1}^G, \{\eta_n\}_{n=1}^N \mid \mathbf{Y}) \\ & \propto \prod_{\ell=1}^L \prod_{k=1}^J \prod_{n=1}^N \prod_{m=1}^G \big(\mathbf{1}\{y^*_{ij} \geq 0, y_{ij} = 1\} \\ & + \mathbf{1}\{y^*_{ij} < 0, y_{ij} = 0\}\big) \\ & \times \phi_1(y^*_\ell; \alpha_k + \beta_k \boldsymbol{\gamma}_m^\top \mathbf{z}_n + \beta_k \eta_n, 1)^{\mathbf{1}\{j[\ell]=k, i[\ell]=n, g[i[\ell]]=m\}} \\ & \times \prod_{k=1}^J \phi_2(\tilde{\boldsymbol{\beta}}_k; \boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}) \prod_{n=1}^N \prod_{m=1}^G \phi_1(\eta_n; 0, \sigma^2_m)^{\mathbf{1}\{g[n]=m\}} \\ & \times \prod_{m=1}^G \mathcal{IG}\left(\sigma^2_m; \frac{\nu_\sigma}{2}, \frac{s^2_\sigma}{2}\right). \end{aligned} \tag{39}$$

For this hierarchical model, there is no closed-form EM algorithm that can directly maximize the posterior distribution given in equation (39). Therefore, as done in the case of the dynamic model, we seek the variational approximation. The factorization assumption we invoke is given by the following:

$$\begin{aligned} & q\left(\mathbf{Y}^*, \{\tilde{\boldsymbol{\beta}}_k\}_{k=1}^J, \{\boldsymbol{\gamma}_m, \sigma^2_m\}_{m=1}^G, \{\eta_n\}_{n=1}^N\right) \\ & = \prod_{\ell=1}^L q(y^*_\ell) \prod_{k=1}^J q(\tilde{\boldsymbol{\beta}}_k) \prod_{m=1}^G q(\boldsymbol{\gamma}_m) q(\sigma^2_m) \prod_{n=1}^N q(\eta_n). \end{aligned} \tag{40}$$

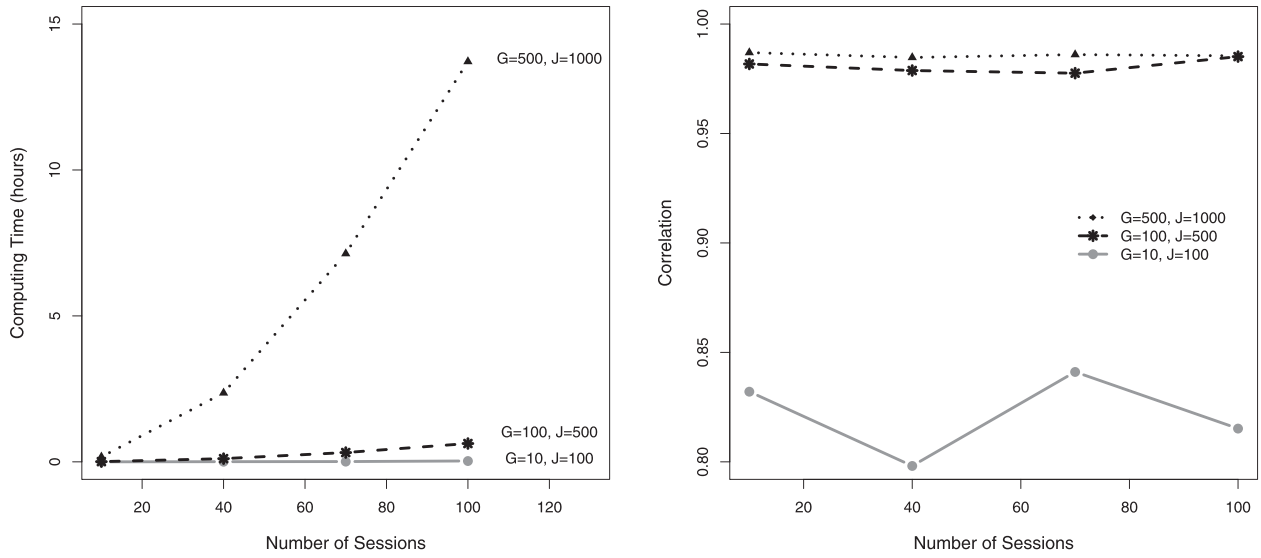
Under this factorization assumption, we can derive the variational EM algorithm that approximates the joint posterior distribution by maximizing the lower bound. Note that aside from the factorization assumption no additional assumption is made to derive the proposed algorithm. The proposed algorithm is described in the final section while the derivation is given in Appendix D.

Simulation Evidence

We conduct a simulation study to demonstrate the computational scalability and accuracy of the proposed variational EM algorithm. To do this, we generate roll-call matrices that vary in size following the simulation study for dynamic models where the number of legislators is now replaced by the number of groups G instead. Each group has N different ideal points to be estimated, and three covariates $\mathbf{z}_{i[\ell]}$ are observed for each ideal point, i.e., $M = 3$. Finally, we construct the simulation such that each group votes on the same set

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 11. Scalability and Accuracy of the Proposed Variational Inference for the Hierarchical Ideal Point Model



Notes: The left panel presents run times of the proposed variational EM algorithm for fitting the hierarchical ideal point model. We consider three different simulation scenarios where the number of groups G varies from 10 to 500 and the number of bills (per group) J ranges from 100 to 1,000. The number of ideal points to be estimated (per group) N is shown on the horizontal axis, with all G groups assumed to vote on all J bills but within each group different legislators vote on different subsets of the bills. In the largest data set we consider ($N = 100$, $J = 1{,}000$, and $G = 500$), our algorithm can estimate a hundred thousand ideal points in about 14 hours. The right panel shows the (Pearson) correlation between the estimated ideal points and their true values.

of J bills but within each group different members vote on different subsets of the bills.

The intercepts for ideal points follow another uniform distribution with $(-1, 1)$, while item discrimination parameters were both drawn uniformly from $(-0.2, 0.2)$. The group-level variance parameters $\sigma^2_{g[i[\ell]]}$ were set to 0.01 for all groups. We use diffuse priors for item difficulty and discrimination parameters as well as for group-level coefficients. Specifically, the prior distribution for these parameters is the independent normal distribution with a mean of zero and a standard deviation of five. For group-level variance parameters, we use a semi-informative prior such that they follow the inverse-gamma distribution with $\nu_0 = 2$ and $s^2 = 0.02$.

When compared to the other models considered in this article, we find the hierarchical model to be computationally more demanding. To partially address this issue, we parallelize the algorithm wherever possible and implement this parallelized code using eight cores through OpenMP in this simulation study. We also use a slightly less stringent convergence criteria than in the other cases where we check whether the correlations for bill parameters and group-level coefficients across their consecutive iterations is greater than $1 - 10^{-3}$. We find that applying a stricter convergence criteria does not significantly improve the quality of the resulting estimates.

We consider three different sets of simulation scenarios where the number of groups G varies from 10 to 500 and the number of bills (per group) J ranges from 100 to 1,000. Figure 11 shows the results. In the left plot, the vertical axis represents the runtime of our algorithm in hours, while the horizontal axis shows the size of each group N, i.e., the number of ideal points to be estimated per group. Our variational EM algorithm scales well to a large data set. In the largest data set we consider ($N = 100$, $J = 1{,}000$, and $G = 500$), for example, the proposed algorithm can estimate a hundred thousand ideal points in only about 14 hours.

In the right plot of Figure 11, we plot the correlation between the estimated ideal points and their true values for each simulation scenario.[10] The quality of estimates appear to depend on the number of groups with the simulations with a larger number of groups yielding almost perfect correlation. When the number of groups is small, however, we find that the correlations are weaker and the results are highly dependent on prior specification. This is a well-known feature of Bayesian hierarchical models (Gelman 2006) and the ideal point models appear to be no exception in this regard.

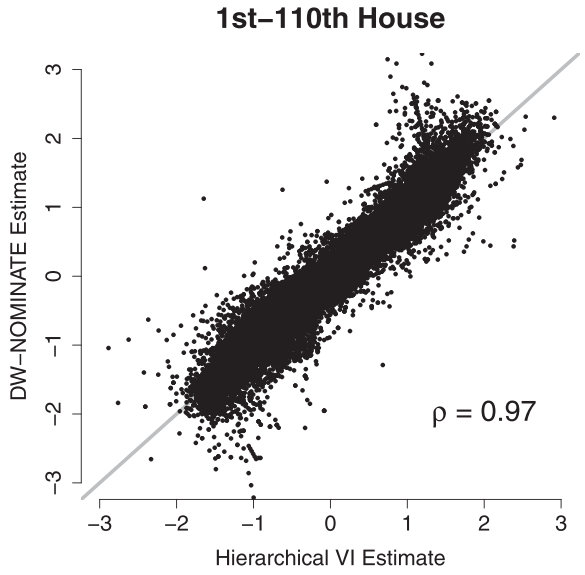
An Empirical Illustration

As noted earlier, DW-NOMINATE scores adopt a linear time trend model for legislators. A model

[10] To account for Monte Carlo error, the simulations for cases where $N = 10$ are repeated 25 times, with median run times and correlations reported in the figure. The standard error of this correlation ranges from 0.08 to 0.11.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 12. Correlation between DW-NOMINATE Estimates and the Proposed Hierarchical Ideal Point Estimates for the 1st–110th Congress



Note: These ideal point estimates are quite similar with a correlation of 0.97.

essentially equivalent to this model can be estimated as a special case of our general hierarchical model, in which the covariate $z_{i[l]}$ is the term served by a particular legislator and the ideal point noise parameter $\eta_{i[l]}$ is fixed at 0.[11] We analyze the roll-call data from the 1st–100th U.S. House and show empirically that the proposed variational EM algorithm for this model produces the ideal point estimates essentially similar to DW-NOMINATE scores. We specify the prior parameters as $\nu_\sigma = 10^8$ and $s_\sigma^2 = 10^{-8}$, which effectively fix the noise parameter as desired, and use the same starting values as those used in DW-NOMINATE. An additional constraint we impose that is consistent with DW-NOMINATE is that legislators who serve less than four terms do not shift ideal points over time.

Our model includes $G = 10{,}474$ groups (i.e., legislators) with $I = 36{,}177$ different ideal points, estimated using $J = 48{,}381$ bills. Estimation of the model using eight threads required just under 5 hours of computing time. This run time could be considerably reduced, for example, by not updating $\eta_{i[\ell]}$ and σ_m^{-2}, which are fixed at zero. Figure 12 shows the estimated ideal points from the hierarchical model, plotted against the corresponding DW-NOMINATE estimates. The two sets of ideal points correlate at 0.97, thus validating the ability of the hierarchical model to reproduce DW-NOMINATE's linear time trend ideal point model.

[11] There are other minor differences relating to different utility functions (Carroll et al. 2009; 2013).

IDEAL POINT MODELS FOR TEXTUAL AND NETWORK DATA

In recent years, political scientists began to develop and apply ideal point models to new types of data, going beyond roll call votes and survey data. They include text data (e.g., Kim, Londregan, and Ratkovic 2014; Lauderdale and Herzog 2014; Lowe et al. 2011; Slapin and Proksch 2008) and network data such as campaign contributions(Bonica 2013; 2014), court citations (Clark and Lauderdale 2010), and social media data (Barberá 2015; Bond and Messing 2015). We expect that the applications of ideal point models to these new types of "big data" will continue to increase over the next few years. In this section, we demonstrate that our approach can be extended to these models. In particular, we consider the popular "Wordfish" model of Slapin and Proksch (2008) for text analysis and an ideal point model commonly used for network data analysis.

Fast Estimation of an Ideal Point Model for Textual Data

Suppose that we have a corpus of K documents, each of which is associated with one of N actors, and there are J unique words possibly after pre-processing the corpus (e.g., stemming). Slapin and Proksch (2008) propose to analyze the $(J \times K)$ term-document matrix $\mathbf{Y}$ using a variant of ideal point model called the Wordfish model (see also Lowe et al. 2011, for a related model). Their substantive application is the analysis of manifestos to estimate ideological positions of political parties.

The Generalized Wordfish Model. The original Wordfish model only allows one document per actor. Here, we generalize this model by enabling multiple documents per actor. Let y_{jk} denote the (j, k) element of the term-document matrix $\mathbf{Y}$, representing the frequency of term j in document k. Then, our *generalized Wordfish model* is defined as

$$p(y_{jk} \mid \alpha_j, \beta_j, \psi_k, x_i) = \text{Poisson}(\lambda_{jk}), \quad \textbf{(41)}$$

$$\lambda_{jk} = \exp(\psi_k + \alpha_j + \beta_j x_{i[k]}), \quad \textbf{(42)}$$

where ψ_k represents the degree of verboseness of document k, α_j is the overall frequency of term j across all documents, β_j is the discrimination parameter for term k, and finally $x_{i[k]}$ represents the ideological position of the actor to whom document k belongs.

Although the original model is developed under the frequentist framework, we consider the Bayesian formulation by specifying a set of independent prior distributions:

$$p(\tilde{\boldsymbol{\beta}}_j) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}), \quad \textbf{(43)}$$

$$p(\psi_k) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_\psi, \sigma_\psi^2), \quad \textbf{(44)}$$

https://doi.org/10.1017/S0003055416000037X Published online by Cambridge University Press

$$p(x_i) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_x, \sigma_x^2), \tag{45}$$

where $\tilde{\boldsymbol{\beta}}_j = (\alpha_j, \beta_j)$ is a vector of term parameters. The joint posterior distribution is therefore given by

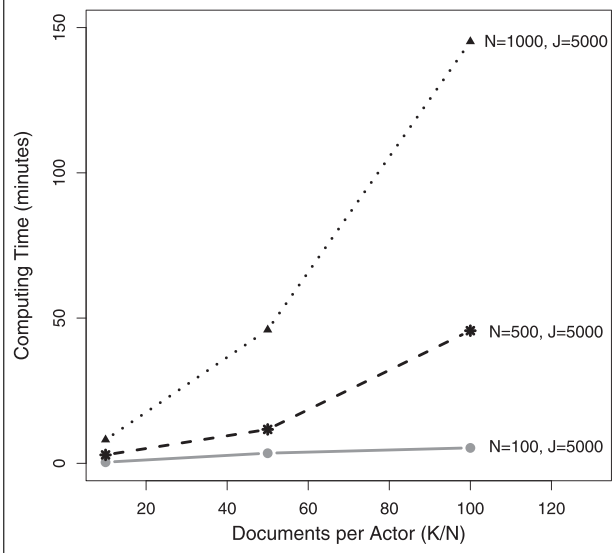
$$\begin{aligned} & p\left(\{\psi\}_{k=1}^K, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J, \{x_{i=1}^N\} \mid \mathbf{Y}\right) \\ & \propto \left[\prod_{j=1}^J \left\{\prod_{k=1}^K p(y_{jk} \mid \psi_k, \tilde{\boldsymbol{\beta}}_j, x_{i[k]}) p(\psi_k)\right\} p(\tilde{\boldsymbol{\beta}}_j)\right] \\ & \times \prod_{i=1}^N p(x_i). \end{aligned} \tag{46}$$

In Appendix E, we derive the EM algorithm for the local variational inference under the following factorization assumption:

$$\begin{aligned} & q\left(\{\psi_k\}_{k=1}^K, \{\tilde{\boldsymbol{\beta}}_j\}_{j=1}^J, \{x_{i=1}^N\}\right) \\ & = \prod_{j=1}^J q(\tilde{\boldsymbol{\beta}}_j) \prod_{k=1}^K q(\psi_k) \prod_{i=1}^N q(x_i). \end{aligned} \tag{47}$$

A Simulation Study. We conduct a simulation study to demonstrate the computational scalability and accuracy of the proposed variational EM algorithm. Here, we generate text-document matrices that vary in size. We consider three different sets of simulation scenarios where the number of actors N varies from 100 to 1,000, while the number of words are fixed at $J = 5{,}000$. We also vary the number of documents linked to each legislator, K/N, from 10 to 100. Therefore, in this simulation study, the total number of documents, K, ranges from 1,000 to 100,000. Note that the number of parameters to be estimated in this simulation $(K + 2J + N)$ varies from 11,100 to 111,000. Results are shown in Figure 13, with run times on the vertical axis in minutes. In the largest data set we consider $(N = 1{,}000,\ J = 5{,}000,\ K/N = 100)$, the proposed algorithm completes estimation in under 2.5 hours. For all simulations above, the correlation between the estimated ideal points and their true values exceeds 0.99.

FIGURE 13. Scalability of the Proposed Variational Inference for the Generalized Wordfish Model



Notes: The left panel presents run times of the proposed variational EM algorithm for fitting the generalized Wordfish model. We consider three different simulation scenarios where the number of actors N varies from 100 to 1,000, while the number of words J is fixed to 5,000. The number of documents per actor K/N is shown on the horizontal axis. The vertical axis indicates the time necessary to fit the generalized Wordfish model for each data set through the proposed algorithm. For all cases above, the correlation exceeds 0.99.

Fast Estimation of an Ideal Point Model for Network Data

We next consider the estimation of ideal points from the network data including citation, social media, and campaign contribution data. These data provide the information about a set of nodes and edges. For example, in the court citation data (Clark and Lauderdale 2010), a node represents a court opinion and the existence of a directed edge from one opinion to another implies a citation. Similarly, in social media data analyzed by Barberá (2015) and Bond and Messing (2015), nodes are Twitter and Facebook users and edges indicate whether they follow each other. For campaign contribution data (Bonica 2014), nodes are candidates and voters and edges represent the donations from voters to politicians.

The Network Ideal Point Model. Here, we consider the model developed by Barberá (2015) to analyze more than four million Twitter users. Let $y_{ij} = \mathbf{1}\{y_{ij}^* > 0\}$ represent whether Twitter user i follows politician j where y_{ij}^* is the latent propensity where $i = 1, 2, \ldots, N$ and $j = 1, 2, \ldots, J$. The *network ideal point model* is given by

$$y_{ij}^* = \alpha_j + \beta_i - \|x_i - z_j\|^2 + \epsilon_{ij}, \tag{48}$$

where $\|\cdot\|$ represents the Euclidian norm, and ϵ_{ij} is the error term.[12] We assume that ϵ_{ij} follows the standard normal distribution. The ideal points of Twitter user i and politician j are denoted by x_i and z_j, respectively. The model assumes that twitter user i are more likely to follow politician j if their ideal points are similar. The model parameters, α_j and β_i, represents the overall degree to which politician j is followed and Twitter user i follows politicians, respectively.

[12] The original model of Barberá (2015) includes a normalizing constant γ and is given by $y_{ij}^* = \alpha_j + \beta_i - \gamma\|x_i - z_j\|^2 + \epsilon_{ij}$. However, it is clear that γ is not identifiable. Hence, we do not consider this parameter.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

The model is completed by the following specification of prior distributions:

$$p(\{\alpha_j\}_{j=1}^J) = \prod_{j=1}^J \mathcal{N}\left(\mu_\alpha, \sigma_\alpha^2\right), \quad \textbf{(49)}$$

$$p(\{\beta_i\}_{i=1}^N) = \prod_{i=1}^N \mathcal{N}\left(\mu_\beta, \sigma_\beta^2\right), \quad \textbf{(50)}$$

$$p(\{x_i\}_{i=1}^N) = \prod_{i=1}^N \mathcal{N}\left(\mu_x, \sigma_x^2\right), \quad \textbf{(51)}$$

$$p(\{z_j\}_{j=1}^J) = \prod_{j=1}^J \mathcal{N}\left(\mu_z, \sigma_z^2\right). \quad \textbf{(52)}$$

Together, the joint posterior distribution conditional on the observed $(N \times J)$ matrix $\mathbf{Y}$ is given by

$$\begin{aligned} & p(\mathbf{Y}^*, \{\alpha_j\}_{j=1}^J, \{\beta_i\}_{i=1}^N, \{x_i\}_{i=1}^N, \{z_j\}_{j=1}^J \mid \mathbf{Y}) \\ & \propto \prod_{i=1}^N \prod_{j=1}^J \left(\mathbf{1}\{y_{ij}^* > 0\}\mathbf{1}\{y_{ij} = 1\}\right. \\ & \left. + \mathbf{1}\{y_{ij}^* \leq 0\}\mathbf{1}\{y_{ij} = 0\}\right) \\ & \times \phi_1\left(y_{ij}^*; \alpha_j + \beta_i - \|x_i - z_i\|^2\right) \\ & \times \prod_{i=1}^N \left\{\phi_1\left(x_i; \mu_x, \sigma_x^2\right)\phi_1\left(\beta_i; \mu_\beta, \sigma_\beta^2\right)\right\} \\ & \times \prod_{j=1}^J \left\{\phi_1\left(\alpha_j; \mu_\alpha, \sigma_\alpha^2\right)\phi_1\left(z_j; \mu_z, \sigma_z^2\right)\right\}. \end{aligned} \quad \textbf{(53)}$$

In Appendix F, we derive the variational EM algorithm for the above ideal point model for network data under the following factorization assumption:

$$\begin{aligned} & q(\mathbf{Y}^*, \{\alpha_j\}_{j=1}^J, \{\beta_i\}_{i=1}^N, \{x_i\}_{i=1}^N, \{z_j\}_{j=1}^J) \\ & = \prod_{i=1}^N \prod_{j=1}^J q(y_{ij}^*) \prod_{i=1}^N q(x_i)q(\beta_i) \prod_{j=1}^J q(\alpha_i)q(z_i). \end{aligned} \quad \textbf{(54)}$$

The variational distributions for the ideal points for users and politicians are based on the second-order Taylor approximation.

An Empirical Study. We test the performance of our estimation algorithm by applying it to a subset of the U.S. Twitter data made available by Barberá (2015). The original data set includes $N = 301{,}537$ voters following $J = 318$ political elites. Since the replication archive recommends using a subset of $N = 10{,}000$ voters following $J = 176$ political elites, we proceed with this smaller data set instead. Even with this subset of the original data, it took 6.5 days on our machine to sample just 500 posterior draws using an MCMC algorithm via **RStan** software.[13] In contrast, our variational EM algorithm was able to complete the estimation for this same data set within 35 minutes even though we did not use any parallelization. This demonstrates the scalability of our algorithm when compared to a standard MCMC algorithm.

In Figure 14, we compare our ideal point estimates with those of Barberá (2015) for both voters (left panel) and political elites (right panel). In both cases, we observe that the variational EM algorithm produces essentially the same estimates as the Markov chain Monte Carlo algorithm as implemented via **RStan** software.

CONCLUDING REMARKS

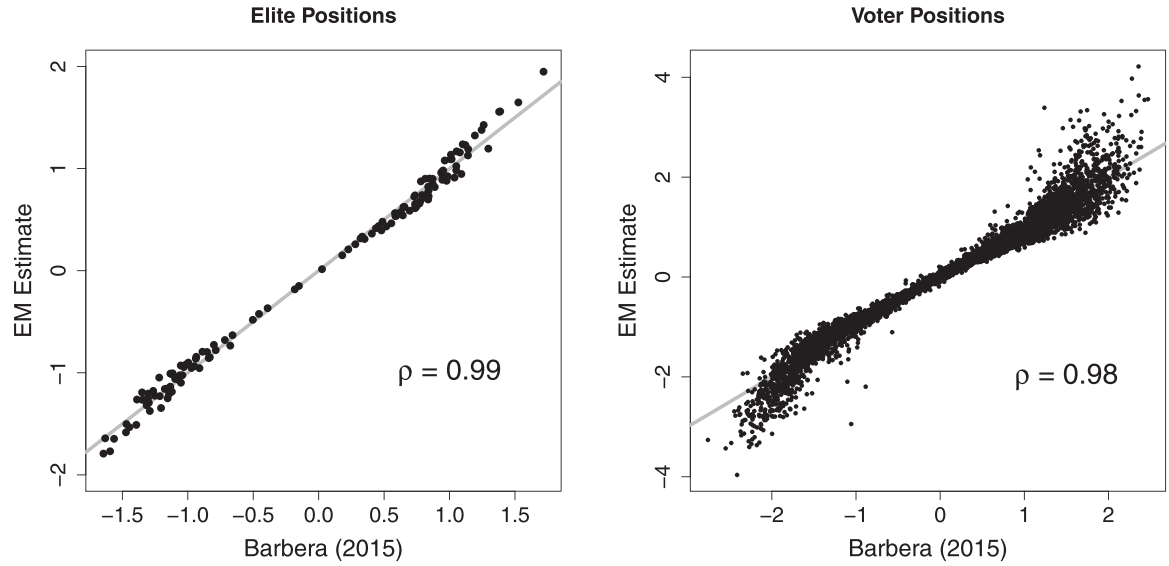
Political ideology is at the core of political science theories across all subfields. And yet, when conducting empirical analyses, it is impossible to directly observe political ideology. Instead, researchers must infer ideological positions of various actors from their behavior or expressed attitudes. Quantitative analysis of political ideology begins with the specification of a measurement model that formally connects latent ideological positions with observed behavior and attitudes. Over the last couple of decades, ideal point estimation methods based on spatial voting models and item response theory have been the main workhorse for quantitative researchers in political science to measure political ideology. These models have been used to analyze the voting in U.S. Congress (e.g., Poole and Rosenthal 1997), courts (e.g., Martin and Quinn 2002), other legislatures (e.g., Hix, Noury, and Roland 2006; Londregan 2007), and United Nations assembly (e.g., Bailey, Strezhnev, and Voeten 2015; Voeten 2000). Beyond roll call votes, the methods are also applied to survey data (e.g., Clinton and Lewis 2008), campaign contributions (e.g., Bonica 2014), party manifestos (e.g., Lowe et al. 2011), speeches (e.g., Proksch and Slapin 2010), and social media (e.g., Bond and Messing 2015).

Over the last decade, political science, like other social science disciplines, witnessed the "big data revolution" where empirical researchers are collecting increasingly large data sets of diverse types. These rich data sets allow researchers to answer the questions they were previously unable to tackle and often enable to employ more realistic but complicated modeling strategies. While the available computational power is steadily increasing, the amount of data available to social scientists and the degree of methodological sophistication are growing at an even faster rate. As a result, researchers are unable to estimate the models of their choice within a reasonable amount of time and are often forced to make a compromise by adopting a feasible and yet undesirable statistical procedure.

[13] Barberá (2015) reports that the approximate run time on a more advanced high performance computer at NYU took approximately 18 hours using parallelization.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

FIGURE 14. Comparison of Our Ideal Point Estimates with Those of Barberá (2015)



Notes: Our ideal point estimates are based on the variational EM algorithm whereas those of Barberá (2015) are based on the Markov chain Monte Carlo algorithm as implemented in RStan. The left panel compares the two sets of ideal points for $J = 176$ political elites whereas the right panel conducts the same comparison for $N = 10{,}000$ voters. In both cases, the correlation between the two sets of estimates is very high. In addition, the variational EM algorithm only took 35 minutes to complete the estimation whereas RStan took 6.5 days to obtain 500 posterior draws on the same computer.

In this article, we develop fast estimation algorithms for ideal points with massive data. These algorithms overcome the computational bottleneck created by massive data in the ideal point measurement context. Specifically, we develop the expectation-maximization (EM) algorithms that maximize the posterior distribution. When such an algorithm is not available in a closed form, we derive a variational EM algorithm that approximates posterior inference. Through empirical and simulation studies, we show that the proposed methodology improves the computational efficiency by orders of magnitude without sacrificing the accuracy of the resulting estimates.[14] With this new methodology, researchers can estimate ideal points from massive data on their laptop within minutes rather than running other estimation algorithms for days on a high-performance computing cluster.

We predict that this line of methodological research will become essential for the next generation of empirical political science research. The political science data now come in a variety of form—textual data, network data, and spatial-temporal data to name a few—and in a large quantity. To efficiently extract useful information from these data will require the development of scalable statistical estimation techniques like the ones proposed in this article.

THE DETAILS OF THE PROPOSED EM ALGORITHMS

In this section, we present the details of the proposed EM algorithms. The derivation of these algorithms is given in the Supplementary Appendix.

The Ordinal Ideal Point Model

Below, we describe the EM algorithm for the three-category ordinal ideal point model. The latent variable updates are equal to

$$z_{ij}^{*\,(t)} = \mathbb{E}\left(z_{ij}^{*} \mid \mathbf{x}_i^{(t-1)}, \tau_j^{(t-1)}, \tilde{\boldsymbol{\beta}}_j^{(t-1)}, y_{ij}\right)$$
$$= \begin{cases} m_{ij}^{(t-1)} - \frac{1}{\tau_j^{(t-1)}}\lambda\left(m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right) & \text{if } y_{ij} = 0, \\ m_{ij}^{(t-1)} + \frac{1}{\tau_j^{(t-1)}}\delta\left(1 - m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right) & \text{if } y_{ij} = 1, \\ m_{ij}^{(t-1)} + \frac{1}{\tau_j^{(t-1)}}\lambda\left(1 - m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right) & \text{if } y_{ij} = 2, \end{cases} \quad \textbf{(55)}$$

where $m_{ij}^{(t-1)} = (\tilde{\mathbf{x}}_i^{(t-1)})^\top \tilde{\boldsymbol{\beta}}_j^{(t-1)}$, $\lambda(m, \tau) = \phi(m\tau)/\{1 - \Phi(m\tau)\}$ and $\delta(m, \tau) = \{\phi(m\tau) - \phi((1-m)\tau)\}/\{\Phi((1-m)\tau) + \Phi(m\tau) - 1\}$. If y_{ij} is missing, then we set

[14] We do note mean this as a general statement for other models and applications. We find that the factorization assumption is appropriate for the ideal point models where we used a variational approach. For other problems, gains in computational efficiency may come at an unacceptable cost of estimation accuracy.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

$z_{ij}^{*(t)} = m_{ij}^{(t-1)}$. The required second moment is given by

$$\left(z_{ij}^{*2}\right)^{(t)} = \mathbb{E}(z_{ij}^{*2} \mid \mathbf{x}_i^{(t-1)}, \tau_j^{(t-1)}, \tilde{\boldsymbol{\beta}}_j^{(t-1)}, y_{ij})$$
$$= \begin{cases} \left(z_{ij}^{*(t)}\right)^2 + \frac{1}{\left(\tau_j^{(t-1)}\right)^2}\left[1 + \tau_j^{(t-1)} m_{ij}^{(t-1)} \lambda\left(m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right)\right. & \\ \quad \left. - \left\{\lambda\left(m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right)\right\}^2\right] & \text{if } y_{ij} = 0 \\ \left(z_{ij}^{*(t)}\right)^2 + \frac{1}{\left(\tau_j^{(t-1)}\right)^2} & \\ \quad \times \left[1 - \frac{\tau_j^{(t-1)}\left\{m_{ij}^{(t-1)}\phi\left(m_{ij}^{(t-1)}\tau_j^{(t-1)}\right) + \left(1 - m_{ij}^{(t-1)}\right)\phi\left(\left(1 - m_{ij}^{(t-1)}\right)\tau_j^{(t-1)}\right)\right\}}{\Phi\left(\left(1 - m_{ij}^{(t-1)}\right)\tau_j^{(t-1)}\right) + \Phi\left(m_{ij}^{(t-1)}\tau_j^{(t-1)}\right) - 1}\right. & \\ \quad \left. - \left\{\delta\left(1 - m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right)\right\}^2\right] & \text{if } y_{ij} = 1 \\ \left(z_{ij}^{*(t)}\right)^2 + \frac{1}{\left(\tau_j^{(t-1)}\right)^2}\left[1 - m_{ij}^{(t-1)} \lambda\left(1 - m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right)\right. & \\ \quad \left. \times \left\{\lambda\left(1 - m_{ij}^{(t-1)},\ \tau_j^{(t-1)}\right) - \left(1 - m_{ij}^{(t-1)}\right)\tau_j^{(t-1)}\right\}\right] & \text{if } y_{ij} = 2 \end{cases} \quad \textbf{(56)}$$

where if y_{ij} is missing, then we set $(z_{ij}^{*(t)})^2 = (m_{ij}^{(t-1)})^2 + (\tau_j^{(t-1)})^{-2}$.

Finally, the M step consists of the following conditional maximization steps:

$$\mathbf{x}_i^{(t)} = \left(\boldsymbol{\Sigma}_{\mathbf{x}}^{-1} + \sum_{j=1}^{J}\left(\tau_j^{(t-1)}\right)^2 \boldsymbol{\beta}_j^{(t-1)}\boldsymbol{\beta}_j^{(t-1)\top}\right)^{-1} \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}\boldsymbol{\mu}_{\mathbf{x}} + \sum_{j=1}^{J}\left(\tau_j^{(t-1)}\right)^2 \boldsymbol{\beta}_j^{(t-1)}(z_{ij}^{*(t)} - \alpha_j^{(t-1)}), \quad \textbf{(57)}$$

$$\tilde{\boldsymbol{\beta}}_j^{(t)} = \left(\boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1} + \left(\tau_j^{(t-1)}\right)^2 \sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)}\left(\tilde{\mathbf{x}}_i^{(t)}\right)^\top\right)^{-1} \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}} + \left(\tau_j^{(t-1)}\right)^2 \sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)} z_{ij}^{*(t)}, \quad \textbf{(58)}$$

$$\left(\tau_j^{(t)}\right)^2 = \frac{N + \nu_\tau - 2}{s_\tau^2 + \left(\tilde{\boldsymbol{\beta}}_j^{(t)}\right)^\top\left\{\sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)}\left(\tilde{\mathbf{x}}_i^{(t)}\right)^\top\right\}\tilde{\boldsymbol{\beta}}_j^{(t)} - 2\left(\tilde{\boldsymbol{\beta}}_j^{(t)}\right)^\top\sum_{i=1}^{N}\tilde{\mathbf{x}}_i^{(t)} z_{ij}^{*(t)} + \sum_{i=1}^{N}\left(z_{ij}^{*2}\right)^{(t)}}. \quad \textbf{(59)}$$

The Dynamic Ideal Point Model

The algorithm consists of three steps. First, the latent propensity update step is based on the following optimal approximating distribution:

$$q(y_{ijt}^*) = \begin{cases} \mathcal{TN}(m_{ijt}, 1, 0, \infty) & \text{if } y_{ijt} = 1 \\ \mathcal{TN}(m_{ijt}, 1, -\infty, 0) & \text{if } y_{ijt} = 0 \end{cases} \quad \textbf{(60)}$$

with $m_{ijt} = \mathbb{E}(\tilde{\mathbf{x}}_{it})^\top\mathbb{E}(\tilde{\boldsymbol{\beta}}_{jt})$. Then, the updated mean of y_{ijt}^* is given by

$$\mathbb{E}(y_{ijt}^*) = \begin{cases} m_{ijt} + \frac{\phi(m_{ijt})}{\Phi(m_{ijt})} & \text{if } y_{ijt} = 1 \\ m_{ijt} - \frac{\phi(m_{ijt})}{1 - \Phi(m_{ijt})} & \text{if } y_{ijt} = 0 \end{cases}. \quad \textbf{(61)}$$

For abstention (i.e., missing y_{ijt}), we set $q(y_{ijt}^*) = \mathcal{N}(m_{ijt}, 1)$ and $\mathbb{E}(y_{ijt}^*) = m_{ijt}$.

Second, the variational distribution for $\tilde{\boldsymbol{\beta}}$ is given by

$$q(\tilde{\boldsymbol{\beta}}_{jt}) = \mathcal{N}(\mathbf{B}_{jt}^{-1}\mathbf{b}_{jt},\ \mathbf{B}_{jt}^{-1}), \quad \textbf{(62)}$$

where $\mathbf{b}_{jt} = \Sigma_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}} + \sum_{i\in\mathcal{I}_t}\mathbb{E}(\tilde{\mathbf{x}}_{it})\mathbb{E}(y_{ijt}^*)$ and $\mathbf{B}_{jt} = \Sigma_{\tilde{\boldsymbol{\beta}}}^{-1} + \sum_{i\in\mathcal{I}_t}\mathbb{E}(\tilde{\mathbf{x}}_{it}\tilde{\mathbf{x}}_{it}^\top)$ with $\mathcal{I}_t = \{i : \underline{T}_i \le t \le \overline{T}_i\}$. Note that the summation is taken over $\mathcal{I}_t$, the set of legislators who are present at time t.

Finally, we consider the variational distribution of dynamic ideal points. Here, we rely on the forward-backward algorithm derived for the variational Kalman filtering. Specifically, we first use the forward recursion to compute

$$x_{it} \mid \ddot{y}_{i1}, \ldots, \ddot{y}_{it} \overset{\text{indep.}}{\sim} \mathcal{N}(c_{it},\ C_{it}), \quad \textbf{(63)}$$

where $\ddot{\beta}_t = \sqrt{\sum_{j=1}^{J_t}\mathbb{E}(\beta_{jt}^2)}$, $\ddot{y}_{it} = \{\sum_{j=1}^{J_t}\mathbb{E}(y_{ijt}^*)\mathbb{E}(\beta_{jt}) - \mathbb{E}(\beta_{jt}\alpha_{jt})\}/\ddot{\beta}_t$, $c_{it} = c_{i,t-1} + K_t(\ddot{y}_{it} - \ddot{\beta}_t c_{i,t-1})$ and $C_{it} = (1 - K_t\ddot{\beta}_t)\Omega_t$ with $\Omega_t = \omega_x^2 + C_{i,t-1}$, $K_t = \ddot{\beta}_t\Omega_t/S_t$ and $S_t = \ddot{\beta}_t^2\Omega_t + 1$. We recursively compute these quantities by setting $c_{i0} = \mu_x$ and $C_{i0} = \Sigma_x$. Then, combined with the backward recursion, we can derive the following variational distribution:

$$x_{it} \mid \ddot{y}_{i\underline{T}_i}, \ldots, \ddot{y}_{i\overline{T}_i} \overset{\text{indep.}}{\sim} \mathcal{N}(d_{it},\ D_{it}), \quad \textbf{(64)}$$

where $d_{it} = c_t + J_t(d_{t+1} - c_{it})$ and $D_{it} = C_{it} + J_t^2(D_{t+1} - \Omega_{t+1})$ with $J_t = C_{it}/\Omega_{t+1}$. The recursive computation is done by setting $d_{i\overline{T}_i} = c_{i\overline{T}_i}$ and $D_{i\overline{T}_i} = C_{i\overline{T}_i}$. Thus, the required first and second moments of x_{it} can be easily obtained.

The Hierarchical Ideal Point Model

The proposed EM algorithm cycles through the following updating steps until convergence. First, we update the variational distribution for the latent propensities y_ℓ^* for all $\ell = 1, \ldots, L$:

$$q(y_\ell^*) = \begin{cases} \mathcal{TN}(m_\ell, 1, -\infty, 0) & \text{if } y_\ell = 1 \\ \mathcal{TN}(m_\ell, 1, 0, \infty,) & \text{if } y_\ell = 0 \\ \mathcal{N}(m_\ell, 1) & \text{if } y_\ell \text{ is missing} \end{cases}, \quad \textbf{(65)}$$

where $m_\ell = \mathbb{E}(\alpha_{j[\ell]}) + \mathbb{E}(\beta_{j[l]})\mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]})^\top\mathbf{z}_{i[\ell]} + \mathbb{E}(\eta_{i[\ell]})\mathbb{E}(\beta_{j[\ell]})$. The required moment update step

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

is given by

$$\mathbb{E}(y_\ell^*) = \begin{cases} m_\ell + \frac{\phi(m_\ell)}{\Phi(m_\ell)} & \text{if} \quad y_\ell = 1 \\ m_\ell - \frac{\phi(m_\ell)}{1-\Phi(m_\ell)} & \text{if} \quad y_\ell = 0 \\ m_\ell & \text{if} \quad y_\ell \text{ is missing} \end{cases}. \quad (66)$$

Next, we update the first and second moments of the ideal point error term η_n using the following variational distribution:

$$q(\eta_n) = \mathcal{N}(A_n^{-1}a_n,\ A_n^{-1}), \quad (67)$$

where $A_n = \mathbb{E}(\sigma_{g[n]}^{-2}) + \sum_{\ell=1}^{L} \mathbf{1}\{i[\ell] = n\}\mathbb{E}(\beta_{j[\ell]}^2)$ and $a_n = \sum_{\ell=1}^{L} \mathbf{1}\{i[\ell] = n\}\{\mathbb{E}(y_\ell^*)\mathbb{E}(\beta_{j[\ell]}) - \mathbb{E}(\alpha_{j[\ell]}\beta_{j[\ell]}) - \mathbb{E}(\beta_{j[\ell]}^2)\mathbb{E}(\boldsymbol{\gamma}_{g[n]})^\top \mathbf{z}_n\}$. Thus, the required moments are given by $\mathbb{E}(\eta_n) = A_n^{-1}a_n$ and $\mathbb{E}(\eta_n^2) = A_n^{-1} + (A_n^{-1}a_n)^2$.

Third, we derive the variational distribution for the item parameters. This distribution is equal to

$$q(\tilde{\boldsymbol{\beta}}_k) = \mathcal{N}(\mathbf{B}_k^{-1}\mathbf{b}_k, \mathbf{B}_k^{-1}), \quad (68)$$

where $\mathbf{B}_k = \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1} + \sum_{\ell=1}^{L} \mathbf{1}\{j[\ell] = k\}\mathbb{E}(\tilde{\mathbf{x}}_{i[\ell]}\tilde{\mathbf{x}}_{i[\ell]}^\top)$ and $\mathbf{b}_k = \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}} + \sum_{\ell=1}^{L} \mathbf{1}\{j[\ell] = k\}\mathbb{E}(y_\ell^*)\mathbb{E}(\tilde{\mathbf{x}}_{i[\ell]})$ with $\tilde{\mathbf{x}}_{i[\ell]} = (1, \boldsymbol{\gamma}_{g[i[\ell]]}^\top \mathbf{z}_{i[\ell]} + \eta_{i[\ell]})^\top$. We note that $\mathbb{E}(\tilde{\mathbf{x}}_{i[\ell]}) = (1, \mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]})^\top \mathbf{z}_{i[\ell]} + \mathbb{E}(\eta_{i[\ell]}))^\top$ and

$$\mathbb{E}\left(\tilde{\mathbf{x}}_{i[\ell]}\tilde{\mathbf{x}}_{i[\ell]}^\top\right) = \begin{pmatrix} 1 & \mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]})^\top \mathbf{z}_{i[\ell]} + \mathbb{E}(\eta_{i[\ell]}) \\ \mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]})^\top \mathbf{z}_{i[\ell]} + \mathbb{E}(\eta_{i[\ell]}) & \mathbf{z}_{i[\ell]}^\top \mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]}\boldsymbol{\gamma}_{g[i[\ell]]}^\top)\mathbf{z}_{i[\ell]} + 2\mathbb{E}(\boldsymbol{\gamma}_{g[i[\ell]]})^\top \mathbf{z}_{i[\ell]}\mathbb{E}(\eta_{i[\ell]}) + \mathbb{E}(\eta_{i[\ell]}^2) \end{pmatrix}. \quad (69)$$

This gives the required moment update, $\mathbb{E}(\tilde{\boldsymbol{\beta}}_k) = \mathbf{B}_k^{-1}\mathbf{b}_k$.

Fourth, the variational distribution for the group-level coefficients is given by

$$q(\boldsymbol{\gamma}_m) = \mathcal{N}(\mathbf{C}_m^{-1}\mathbf{c}_m,\ \mathbf{C}_m^{-1}), \quad (70)$$

where $\mathbf{C}_m = \boldsymbol{\Sigma}_{\boldsymbol{\gamma}}^{-1} + \sum_{\ell=1}^{L} 1\{g[i[\ell]] = m\}\mathbb{E}(\beta_{j[\ell]}^2)\mathbf{z}_{i[\ell]}\mathbf{z}_{i[\ell]}^\top$ and $\mathbf{c}_m = \boldsymbol{\Sigma}_{\boldsymbol{\gamma}}^{-1}\boldsymbol{\mu}_{\boldsymbol{\gamma}} + \sum_{\ell=1}^{L} 1\{g[i[\ell]] = m\}\mathbf{z}_{i[\ell]}[\mathbb{E}(\beta_{j[\ell]})\{\mathbb{E}(y_\ell^*) - \mathbb{E}(\alpha_{j[\ell]})\} - \mathbb{E}(\beta_{j[\ell]}^2)\mathbb{E}(\eta_{i[\ell]})]$. Thus, the required moment updates are given by $\mathbb{E}(\boldsymbol{\gamma}_m) = \mathbf{C}_m^{-1}\mathbf{c}_m$ and $\mathbb{E}(\boldsymbol{\gamma}_m\boldsymbol{\gamma}_m^\top) = \mathbf{C}_m^{-1} + \mathbf{C}_m^{-1}\mathbf{c}_m\mathbf{c}_m^\top\mathbf{C}_m^{-1}$.

Finally, we derive the variational distribution for the group-level variance parameters. This distribution is equal to

$$q(\sigma_m^2) = \mathcal{IG}\left(\frac{\nu_\sigma + \sum_{n=1}^{N} \mathbf{1}\{g[n] = m\}}{2}, \times\frac{1}{2}\left(s_\sigma^2 + \sum_{n=1}^{N} \mathbf{1}\{g[n] = m\}\mathbb{E}(\eta_n^2)\right)\right), \quad (71)$$

where the desired moment update is given by $\mathbb{E}(\sigma_m^{-2}) = [\nu_\sigma + \sum_{n=1}^{N} \mathbf{1}\{g[n] = m\}]/[s_\sigma^2 + \sum_{n=1}^{N} \mathbf{1}\{g[n] = m\}\mathbb{E}(\eta_n^2)]$. These updating steps are repeated until convergence.

The Generalized Wordfish Model

The approximate update distribution for the document verboseness parameter ψ_k for $k = 1, \ldots, K$ is given by

$$q(\psi_k) \approx \mathcal{N}(A_k^{-1}a_k, A_k^{-1}) \quad (72)$$

where $a_k = \sum_{j=1}^{J}[y_{jk} - \exp(\xi_{jk})\{1 - \xi_{jk} + \mathbb{E}(\alpha_j) + \mathbb{E}(\beta_j)\mathbb{E}(x_{i[k]})\}] + \sigma_\psi^{-2}\mu_\psi$ and $A_k = \sum_{j=1}^{J} \exp(\xi_{jk}) + \sigma_\psi^{-2}$. Next, the variational distribution for ideal points x_n for $n = 1, \ldots, N$ is given by

$$q(x_n) \approx \mathcal{N}(B_n^{-1}b_n, B_n^{-1}), \quad (73)$$

where $b_n = \sum_{j=1}^{J}\sum_{k=1}^{K} \mathbf{1}\{i[k] = n\}\mathbb{E}(\beta_j)\{y_{jk} - \exp(\xi_{jk})(1 + \mathbb{E}(\alpha_j) - \xi_{jk} + \mathbb{E}(\psi_k))\} + \sigma_x^{-2}\mu_x$ and $B_n = \sum_{j=1}^{J}\sum_{k=1}^{K} \mathbf{1}\{i[k] = n\}\exp(-\xi_{jk})/\mathbb{E}(\beta_j^2) + \sigma_x^{-2}$. Finally, the update for the term parameters $\tilde{\boldsymbol{\beta}}_j$ for $j = 1, \ldots, J$ is given by

$$q(\tilde{\boldsymbol{\beta}}_j) \approx \mathcal{N}(\mathbf{C}_j^{-1}\mathbf{c}_j, \mathbf{C}_j^{-1}), \quad (74)$$

where $\mathbf{c}_j = \{y_{jk} - \exp(\xi_{jk})(1 - \xi_{jk} + \mathbb{E}(\psi_k))\}\mathbb{E}(\tilde{\mathbf{x}}_{i[k]}) + \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}\boldsymbol{\mu}_{\tilde{\boldsymbol{\beta}}}$ and $\mathbf{C}_j = \sum_{k=1}^{K} \exp(\xi_{jk})\mathbb{E}(\tilde{\mathbf{x}}_{i[k]}\tilde{\mathbf{x}}_{i[k]}^\top) + \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\beta}}}^{-1}$.

The Network Ideal Point Model

The variational distribution for the latent propensity is given by the following truncated normal distribution:

$$q(y_{ij}^*) = \begin{cases} \mathcal{TN}(m_{ij}, 1, 0, \infty) & \text{if } y_{ij} = 1 \\ \mathcal{TN}(m_{ij}, 1, -\infty, 0) & \text{if } y_{ij} = 0 \end{cases}, \quad (75)$$

where $m_{ij} = \mathbb{E}(\alpha_j) + \mathbb{E}(\beta_i) - \mathbb{E}(x_i^2) - \mathbb{E}(z_j^2) + 2\mathbb{E}(x_i)\mathbb{E}(z_j)$.

The variational distribution for the user-specific intercept is given by

$$q(\beta_i) = \mathcal{N}(B^{-1}b_i,\ B^{-1}), \quad (76)$$

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

where $B = J + 1/\sigma_\beta^2$ and $b_i = \mu_\beta/\sigma_\beta^2 + \sum_{j=1}^{J}\left(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\alpha_j) + \mathbb{E}(x_i^2) - 2\mathbb{E}(x_i)\mathbb{E}(z_j) + \mathbb{E}(z_j^2)\right)$. The update step for the users' ideal points is based on the following variational distribution:

$$q(\mathbf{x}) = \prod_{i=1}^{N} q(x_i) \approx \prod_{i=1}^{N} \mathcal{N}(D_i^{-1}d_i, D_i^{-1}), \quad \textbf{(77)}$$

where $D_i = \mathbb{E}(f''(\hat{x}_i))/2$ and $d_i = \{\mathbb{E}(f''(\hat{x}_i))\hat{x}_i - \mathbb{E}(f'(\hat{x}_i))\}/2$. The relevant expectations are given by

$$\begin{aligned}\mathbb{E}(f'(x)) = \frac{2}{\sigma_x^2}(x - \mu_x) \\ + 4\sum_{j=1}^{J}\big\{(x - \mathbb{E}(z_j))\big(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\alpha_j) - \mathbb{E}(\beta_i)\big) \\ + \big(x^3 - 3x^2\mathbb{E}(z_j) + 3x\mathbb{E}(z_j^2) - \mathbb{E}(z_j^3)\big)\big\},\end{aligned} \quad \textbf{(78)}$$

$$\begin{aligned}\mathbb{E}(f''(x)) = \frac{2}{\sigma_x^2} + 4\sum_{j=1}^{J}\big\{\big(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\alpha_j) - \mathbb{E}(\beta_i)\big) \\ + 3\left(x^2 - 2x\mathbb{E}(z_j) + \mathbb{E}(z_j^2)\right)\big\}.\end{aligned} \quad \textbf{(79)}$$

The updates for the politician-specific intercept and politicians' ideal points are similar to the ones for users,

$$q(\alpha_j) = \mathcal{N}(A^{-1}a_j, A^{-1}), \quad \textbf{(80)}$$

where $A = N + 1/\sigma_\alpha^2$ and $a_j = \mu_\alpha/\sigma_\alpha^2 + \sum_{i=1}^{N}(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\beta_i) + (\mathbb{E}(x_i^2) - 2\mathbb{E}(x_i)\mathbb{E}(z_j) + \mathbb{E}(z_j^2)))$,

$$q(z_j) \approx \mathcal{N}(E_j^{-1}e_j, E_j^{-1}),$$

where $E_j = \mathbb{E}(g''(\hat{z}_j))/2$ and $e_j = \{\mathbb{E}(g''(\hat{z}_j))\hat{z}_j - \mathbb{E}(g'(\hat{z}_j))\}/2$. The relevant expectations are given by

$$\begin{aligned}\mathbb{E}(g'(z)) = \frac{2}{\sigma_z^2}(z - \mu_z) \\ -4\sum_{i=1}^{N}\big\{(\mathbb{E}(x_i) - z)\left(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\alpha_j) - \mathbb{E}(\beta_i)\right) \\ + \left(\mathbb{E}(x_i^3) - 3\mathbb{E}(x_i^2)z + 3\mathbb{E}(x_i)z^2 - z^3\right)\big\},\end{aligned}$$

$$\begin{aligned}\mathbb{E}(g''(z)) = \frac{2}{\sigma_z^2} + 4\sum_{i=1}^{N}\big\{\big(\mathbb{E}(y_{ij}^*) - \mathbb{E}(\alpha_j) - \mathbb{E}(\beta_i)\big) \\ + 3\left(\mathbb{E}(x_i^2) - 2\mathbb{E}(x_i)z + z^2\right)\big\}.\end{aligned}$$

SUPPLEMENTARY MATERIAL

To view supplementary material for this article, please visit https://doi.org/10.1017/S000305541600037X

REFERENCES

Bafumi, Joseph, Andrew Gelman, David K. Park, and Noah Kaplan. 2005. "Practical Issues in Implementing and Understanding Bayesian Ideal Point Estimation." *Political Analysis* 13: 171–87.

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

Bailey, Michael. 2007. "Comparable Preferences across Time and Institutions for the Court, Congress, and Presidency." *American Journal of Political Science* 51: 433–48.

Bailey, Michael A. 2013. "Is Today's Court the Most Conservative in Sixty Years? Challenges and Opportunities in Measuring Judicial Preferences." *Journal of Politics* 75: 821–34.

Bailey, Michael, and Kelly H. Chang. 2001. "Comparing Presidents, Senators, and Justices: Interinstitutional Preference Estimation." *The Journal of Law, Economics, and Organization* 17: 477–506.

Bailey, Michael A., Brian Kamoie, and Forrest Maltzman. 2005. "Signals from the Tenth Justice: The Political Role of the Solicitor General in the Supreme Court Decision Making." *American Journal of Political Science* 49: 72–85.

Bailey, Michael A., Anton Strezhnev, and Erik Voeten. 2015. "Estimating Dynamic State Preferences from United Nations Voting Data." *Journal of Conflict Resolution*.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23: 76–91.

Battista, James Coleman, Michael Peress, and Jesse Richman. 2013. "Common-Space IdealPoints, Committee Assignments, and Financial Interests in the State Legislatures." *State Politics & Policy Quarterly* 13: 70–87.

Bock, R. Darrell, and Murray Aitkin. 1981. "Marginal Maximum Likelihood Estimation of Item Parameters: Application of an EM Algorithm." *Psychometrika* 46: 443–59.

Bond, Robert, and Solomon Messing. 2015. "Quantifying Social Medias Political Space: Estimating Ideology from Publicly Revealed Preferences on Facebook." *American Political Science Review* 109: 62–78.

Bonica, Adam. 2013. "Ideology and Interests in the Political Marketplace." *American Journal of Political Science* 57: 294–311.

Bonica, Adam. 2014. "Mapping the Ideological Marketplace." *American Journal of Political Science* 58: 367–87.

Carroll, Royce, Jeffrey B. Lewis, James Lo, and Keith T. Poole. 2009. "Measuring Bias and Uncertainty in DW-NOMINATE Ideal Point Estimates via the Parametric Bootstrap." *Political Analysis* 17: 261–75.

Carroll, Royce, Jeffrey B. Lewis, James Lo, Keith T. Poole, and Howard Rosenthal. 2009. "Comparing NOMINATE and IDEAL: Points of difference and Monte Carlo tests." *Legislative Studies Quarterly* 34: 555–91.

Carroll, Royce, Jeffrey B. Lewis, James Lo, Keith T. Poole, and Howard Rosenthal. 2013. "The Structure of Utility in Spatial Models of Voting." *American Journal of Political Science* 57: 1008–28.

Clark, Tom S., and Benjamin Lauderdale. 2010. "Locating Supreme Court Opinions in Doctrine Space." *American Journal of Political Science* 54: 871–90.

Clinton, Joshua D., Anthony Bertelli, Christian R. Grose, David E. Lewis, and David C. Nixon. 2012. "Separated Powers in the United States: The Ideology of Agencies, Presidents, and Congress." *American Journal of Political Science* 56: 341–54.

Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. "The Statistical Analysis of Roll Call Data." *American Political Science Review* 98: 355–70.

Clinton, Joshua D., and David E. Lewis. 2008. "Expert Opinion, Agency Characteristics, and Agency Preferences." *Political Analysis* 16: 3–20.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press

Clinton, Joshua D., and Adam Meirowitz. 2003. "Integrating Voting Theory andRoll Call Analysis: A Framework." *Political Analysis* 11: 381–96.
Dempster, Arthur P., Nan M. Laird, and Donald B. Rubin. 1977. "Maximum Likelihood from Incomplete Data Via the EM Algorithm (with Discussion)." *Journal of the Royal Statistical Society, Series B, Methodological* 39: 1–37.
Gelman, Andrew. 2006. "Prior Distributions for Variance Parameters in Hierarchical Models." *Bayesian Analysis* 1: 515–33.
Gerber, Elisabeth R., and Jeffrey B. Lewis. 2004. "Beyond the Median: Voter Preferences, District Heterogeneity, and Political Representation." *Journal of Political Economy* 112: 1364–83.
Gerrish, Sean M., and David M. Blei. 2012. "How They Vote: Issue-Adjusted Models of Legislative Behavior." *Advances in Neural Information Processing Systems* 25: 2762–70.
Grimmer, Justin. 2011. "An Introduction to Bayesian Inference via Variational Approximations." *Political Analysis* 19: 32–47.
Hirano, Shigeo, Kosuke Imai, Yuki Shiraito, and Masaaki Taniguchi. 2011. "Policy Positions in Mixed Member Electoral Systems:Evidence from Japan." Working Paper available at http://imai.princeton.edu/research/japan.html.
Hix, Simon, Abdul Noury, and Gérard Roland. 2006. "Dimensions of Politics in the European Parliament." *American Journal of Political Science* 50: 494–511.
Ho, Daniel E., and Kevin M. Quinn. 2010. "Did a Switch in Time Save Nine?" *Journal of Legal Analysis* 2: 1–45.
Imai, Kosuke, James Lo, and Jonathan Olmsted. 2015. "emIRT: EM Algorithms for Estimating Item Response Theory Models." available at the Comprehensive R Archive Network (CRAN). http://CRAN.R-project.org/package=list.
Imai, Kosuke, James Lo, and Jonathan Olmsted. 2016. "Replication data for: Fast Estimation of Ideal Points with Massive Data." doi:10.7910/DVN/HAU0EX. The Dataverse Network.
Jackman, Simon. 2001. "Multidimensional Analysis of Roll Call Data via Bayesian Simulation: Identification, Estimation, Inference, and Model Checking." *Political Analysis* 9: 227–41.
Jackman, Simon. 2012. *pscl: Classes and Methods for R Developed in the Political Science Computational Laboratory, Stanford University*. Department of Political Science, Stanford University, Stanford, California: Stanford University. R package version 1.04.4.
Kim, In Song, John Londregan, and Marc Ratkovic. 2014. Voting, Speechmaking, and the Dimensions of Conflict in the US Senate. Technical Report. Department of Politics, Princeton University.
Lauderdale, Benjamin E., and Alexander Herzog. 2014. Measuring Political Positions from Legislative. Technical Report. London School of Economics and Political Science.
Lewandowski, Jirka, Nicolas Merz, Sven Regel, and Pola Lehmann. 2015. *manifestoR: Access and Process Data and Documents of the Manifesto Project.* R package version 1.1-1. http://CRAN.R-project.org/package=manifestoR
Lewis, Jeffrey B., and Keith T. Poole. 2004. "Measuring Bias and Uncertainty in Ideal Point Estimates via the Parametric Boostrap." *Political Analysis* 12 (2): 105–27.
Londregan, John B. 1999. "Estimating Legislators' Preferred Points." *Political Analysis* 8: 35–56.
Londregan, John B. 2007. *Legislative Institutions and Ideology in Chile*. Cambridge, England: Cambridge University Press.
Lowe, Will, Kenneth Benoit, Slava Mikhaylov, and Michael Laver. 2011. "Scaling Policy Preferences from Coded Political Texts." *Legislative Studies Quarterly* 36: 123–55.
Martin, Andrew D., and Kevin M. Quinn. 2002. "Dynamic Ideal Point Estimation via Markov chain Monte Carlo for the U.S. Supreme Court, 1953–1999." *Political Analysis* 10: 134–53.
Martin, Andrew D., Kevin M. Quinn, and Jong Hee Park. 2013. *MCMCpack: Markov chain Monte Carlo MCMC Package*. http://cran.r-project.org/web/packages/MCMCpack
McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge, MA: MIT Press.
Morgenstern, Scott. 2004. *Patterns of Legislative Politics: Roll-Call Voting in Latin America and the United States*. Cambridge, England: Cambridge University Press.
Poole, Keith T. 2000. "Nonparametric Unfolding of Binary Choice Data." *Political Analysis* 8: 211–37.
Poole, Keith, Jeffrey Lewis, James Lo, and Royce Carroll. 2011. "Scaling Roll Call Votes with wnominate in R." *Journal of Statistical Software* 42: 1–21. http://www.jstatsoft.org/v42/i14/
Poole, Keith, Jeffrey Lewis, James Lo, and Royce Carroll. 2012. *oc: OC Roll Call Analysis Software.* R package version 0.93. http://CRAN.R-project.org/package=oc
Poole, Keith T., and Howard Rosenthal. 1997. *Congress: A Political Economic History of Roll Call Voting*. New York: Oxford University Press.
Poole, Keither T., and Howard Rosenthal. 1991. "Patterns of Congressional Voting." *American Journal of Political Science* 35: 228–78.
Proksch, Sven-Oliver, and Jonathan B. Slapin. 2010. "Position Taking in European Parliament Speeches." *British Journal of Political Science* 40: 587–611.
Quinn, Kevin M. 2004. "Bayesian Factor Analysis for Mixed Ordinal and Continuous Responses." *Political Analysis* 12: 338–53.
Rosas, Guillermo, and Yael Shomer. 2008. "Models of Nonresponse in Legislative Politics." *Legislative Studies Quarterly* 33: 573–601.
Shor, Boris, Christopher Berry, and Nolan McCarty. 2011. "A Bridge to Somewhere: Mapping State and Congressional Ideology on a Cross-institutional Common Space." *Legislative Studies Quarterly* 35: 417–48.
Shor, Boris, and Nolan McCarty. 2011. "The Ideological Mapping of American Legislatures." *American Political Science Review* 105: 530–51.
Slapin, Jonathan B., and Sven-Oliver Proksch. 2008. "A Scaling Model for Estimating Time-Series Party Positions from Texts." *American Journal of Political Science* 52: 705–22.
Spirling, Arthur, and Iain McLean. 2007. "UK OC OK? Interpreting Optimal Classification Scores for the U.K. House of Commons." *Political Analysis* 15: 85–96.
Tausanovitch, Chris, and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities." *Journal of Politics* 75: 330–42.
Voeten, Erik. 2000. "Clashes in the Assembly." *International Organization* 54: 185–215.
Wainwright, Martin J., and Michael I. Jordan. 2008. "Graphical Models, Exponential Families, and Variational Inference." *Foundations and Trends in Machine Learning* 1: 1–310.

https://doi.org/10.1017/S000305541600037X Published online by Cambridge University Press