

ORIGINAL ARTICLE

Analyzing decision records from committees*

Moritz Marbach 

Graduate School of Economic and Social Sciences (GESS), University of Mannheim, Mannheim, 68159, Germany
Corresponding author. Email: moritz.marbach@tamu.edu

(Received 11 August 2016; revised 12 April 2017; accepted 19 April 2017; first published online 5 April 2021)

Abstract

In the absence of a complete voting record, decision records are an important data source to analyze committee decision-making in various institutions. Despite the ubiquity of decision records, we know surprisingly little about how to analyze them. This paper highlights the costs in terms of bias, inefficiency, or inestimable effects when using decision instead of voting records and introduces a Bayesian structural model for the analysis of decision-record data. I construct an exact likelihood function that can be tailored to many institutional contexts, discuss identification, and present a Gibbs sampler on the data-augmented posterior density. I illustrate the application of the model using data from US state supreme court abortion decisions and UN Security Council deployment decisions.

Keywords: Bayesian; Random Utility Models; Discrete Choice Models

At every level of politics, from a city council meeting up to the United Nations Security Council, committees—groups of representatives—make rules, monitor and enforce compliance. While many of these committees adopt decisions by some form of voting, the absence of a complete voting record is an unfortunate but common feature of many of them. A large majority of domestic and international institutions such as courts, central banks, or intergovernmental organizations does not publish voting records consistently.¹

While the reasons for the lack of a voting record are plentiful, the consequences for quantitative empirical research are the same: making inferences regarding how observables are related to committee members' vote choices is challenging. In a search for a means to make inferences, some studies have turned to committees' decision records. A decision record can generally be defined as a list detailing the adoption or rejection decisions of a committee as a whole. Using these data, such studies estimate the effect of observables on the probability of the committee adopting or rejecting a decision to learn about the effect of these observables on members' vote choice.

*For helpful comments and suggestions, I thank Stephen Chaudoin, Tom Clark, Songying Fang, Carlo Horz, Simon Hug, Matt Loftis, Will Lowe, Lars Mäder, Nikolay Marinov, Christoph Mikulaschek, Simon Munzert, Adam Scharf, Curtis Signorino, Tilko Swalve, and Anna Wilke. I am also grateful to the members of my dissertation committee, Thomas König, Thomas Gschwend, Jeff Gill, and Daniel Stegmüller.

¹For example, according to an analysis of data from 93 central banks that responded to a questionnaire from the Bank of England in the 1990s, only six publish voting records (Japan, Korea, Poland, Sweden, the United Kingdom, and the United States) (Fry *et al.*, 2000, Chart 7.3). In a survey conducted by this author of 12 international organizations mentioned in Schermers and Blokker (2011) that have nonplenary organs and operate on a global scale, eight use a show of hands (ITU, UPU, ILO, UNESCO, WHO, IMO, WMO, and IAEA), two voting by assent (IMF, IBRD/IDA “World Bank”) and one a secret vote (ICAO) as ordinary vote-casting procedure. Except for FAO on some issues, none uses recorded voting as an ordinary voting procedure, but the rules of procedure typically allow committee members to request a recorded vote (exceptions are the IMF, World Bank, and WMO).

A typical example of this strategy is the literature on UN peace operations. The central puzzle in this literature is why the UN Security Council deploys UN peace operations in some conflicts but not others. One major line of inquiry is to evaluate whether UN Security Council permanent members' self-interest—captured by variables such as previous colonial relations, military alliances, and trade relationships—prevents more decisive actions by the Council.² Since most votes on UN peace operation deployments are unavailable, studies in this literature cannot estimate the effect of (measured) self-interest on permanent members' vote choice but estimate only the effect of *average* (measured) self-interest on the Council's decision to approve or reject UN peace operations.

This paper is about how to analyze decision records and the relative costs of using decision instead of voting records. Typically, as, for example, in the literature on UN peace operations, the decision record is assumed to be drawn from a convenient stochastic distribution, which allows the analyst to employ a standard model for inference (e.g., a probit model). Deviating from this reduced-form approach, I introduce a Bayesian structural model deriving the exact stochastic distribution of decision record data from the vote-choice distributions that determine a decision. To arrive at the structural likelihood function of the observed data, I model each unobserved vote choice with an ordinary probit model: the choice to vote one way or the other is a function of observable variables and a vector of coefficients. However, since choices are unobserved, I integrate out the actual vote choices to arrive at a likelihood function that is a function of observables, coefficients, and the institutional context but not of the unobserved vote choices. I highlight the intimate connection between the likelihood function and the little-known Poisson's Binomial distribution (Wang, 1993) and its relationship to the bivariate probit with partial observability (Poirier, 1980; Przeworski and Vreeland, 2002) and discuss (classical) parametric identification. I derive a suitable Gibbs sampler to simulate from the exact posterior density. This Gibbs sampler is implemented in this author's open-source R-package *consilium* that accompanies this paper.

The Bayesian structural model clarifies the main methodological challenge with decision records, incorporates additional information about the structure of the data-generating process, and has practical advantages. First, it makes the costs of (partial) aggregation transparent. As I discuss in detail, these costs include not estimable member-specific effects, an increase in posterior uncertainty, and, in some circumstances, an aggregation bias. These costs can be mitigated by including partially observed votes, which is computationally straightforward within the structural model but infeasible in a reduced-form model. Furthermore, the structural model allows the analyst to calculate vote-choice probabilities, which is also infeasible with a reduced-form model. Perhaps surprisingly, vote-choice probabilities are not linear functions of adoption probabilities. This is because adoption probabilities are conditional probabilities with respect to the institutional context, while vote-choice probabilities are unconditional probabilities. To the extent that the analyst aims to learn how observables are related to members' vote choices or intends to make comparisons across institutional contexts, the structural model is a more suitable way of analyzing decision-record data. Finally, I also show that the correct reduced-form model is not necessarily the one that is typically estimated in practice.

I conduct Monte Carlo experiments to verify that the model works as expected, and I replicate a study by Caldarone *et al.* (2009) on US state supreme courts to contrast the inference when a voting record is used with the inference when I artificially delete (a subset of) the recorded votes retaining the decision record. To highlight the advantages of the structural model relative to a reduced-form model, I return to the example of the UN Security Council and estimate whether a UN Security Council member is more likely to support the deployment of a UN Blue Helmet operation if it has strong trade relationships with the conflict location. I conclude this paper with a short discussion on the (types of) institutions for which one can successfully compile a decision record in the first place and then apply the partial m-probit.

²For example, Gilligan and Stedman (2003); Hultman (2013); Stojek and Tir (2015).

1. Modeling decision records

I consider a setting with a committee of M members ($i = 1, \dots, M$) and J decisions ($j = 1, \dots, J$). A member's vote is a binary random variable, $y_{ij} \in \{0, 1\}$ corresponding to members' binary vote choices to adopt or reject a proposal (*no* or *yes*). Crucially, the votes are not observed. The vote of each member is governed by a vector of K covariates (observables), denoted \mathbf{x}_{ij} . While the analyst does not observe the votes, he or she observes the binary outcome of the voting, which I denote with $b_j \in \{0, 1\}$, where b_j is zero if the proposal was rejected. A generic dataset that clarifies the notation appears in Table 1.

Note that, if the votes had been observed, the data could be analyzed with standard discrete-choice models. The aggregation of the voting record complicates matters here, and it is this complication that I address.

My setting is different from that of ecological studies since the covariates are not aggregated but fully observed, the dependent variable is binary instead of continuous or categorical, and the number of vote choices is much smaller. The setting also differs from aggregate studies, where the analyst usually observes only a sample of the members.³ The setting I consider is one in which the values for all covariates for all members are available to the analyst.

1.1 Model statement

Let \mathbf{X}_j be an $M \times K$ matrix that collects all covariates for all M members for each decision j , and let \mathbf{y}_j be the vector of length M collecting all votes, y_{1j}, \dots, y_{Mj} , for the corresponding proposal. I refer to this vector as the vote profile.⁴ I define \mathbf{y}_j^* as the vector of latent utilities for M members to support a decision j . An element of this vector is the latent utility of member i , denoted y_{ij}^* . Member i votes *yes* if $y_{ij}^* \geq 0$. For simplicity, I assume that the latent utility is a linear function of the covariates with the corresponding parameter vector $\boldsymbol{\beta}$.

Let the voting rule that governs the adoption or rejection of a proposal be a q -rule, with a majority threshold \mathcal{R} , such as a simple majority rule or a supermajority rule.⁵ If the number of votes, that is, $\sum_{i=1}^M y_{ij}$, is less than \mathcal{R} , the rejection decision ($b_j = 0$) is realized; otherwise, the decision to adopt is realized ($b_j = 1$). Using this notation, the model can be written as follows:

$$\begin{aligned}
 \mathbf{y}_j^* &= \mathbf{X}_j \boldsymbol{\beta} + \boldsymbol{\epsilon}_j \\
 \boldsymbol{\epsilon}_j &\overset{iid}{\sim} \boldsymbol{\phi}(0, \mathbf{1}) \\
 y_{ij} &= \begin{cases} 0 & \text{if } y_{ij}^* < 0 \\ 1 & \text{otherwise} \end{cases}, \\
 b_j &= \begin{cases} 0 & \text{if } \sum_{i=1}^M y_{ij} < \mathcal{R} \\ 1 & \text{otherwise} \end{cases},
 \end{aligned} \tag{1}$$

where $\boldsymbol{\phi}(0, \mathbf{1})$ is the standard multivariate normal density. The model rests on two assumptions:

³Aggregate data analysis is a growing body of literature in biostatistics (Wakefield and Salway, 2001; Hanseuse and Wakefield, 2008), but see Glynn *et al.* (2008) for a social science application. Aggregate studies differ from ecological studies in two key respects: they incorporate additional, partially available individual-level data, and they model the aggregate outcome based on models of individual behavior (Wakefield and Salway, 2001).

⁴Throughout the text, I follow the convention denoting vectors and matrices in bold letters.

⁵I use a simple, constant q -rule to reduce notational clutter but, as will become clear, the model and the Gibbs sampler can handle decision-specific simple rules. Simple rules can be characterized by decisive coalitions (the set of all vote profiles that lead to the adoption of a proposal) and encompass many generally used voting rules, including weighted-majority rules (weighted q -rules) and veto-majority rules (collegial rules). However, plurality rule and Borda count are not simple rules. For a formal definition of voting rules, see Austen-Smith and Banks (1999, Chap. 3.1).

Table 1. A generic dataset for a committee with M members, having made J decisions

Member	Decision	Observed		Unobserved Vote
		Covariates	Outcome	
1	1	x_{11}	} b_1	y_{11}
2	1	x_{21}		y_{21}
⋮	⋮	⋮		⋮
M	1	x_{M1}		y_{M1}
⋮	⋮	⋮	⋮	⋮
1	J	x_{1J}	} b_J	y_{1J}
2	J	x_{2J}		y_{2J}
⋮	⋮	⋮		⋮
M	J	x_{MJ}		y_{MJ}

The observed decision outcome (b_j) is realized given a voting rule and the (unobserved) votes (y_{ij}). For each member–decision combination, there is a vector of covariates (x_{ij}).

(1) coefficients are shared across all committee members, and (2) vote choices are conditionally independent. The latter assumption corresponds to the familiar sincere-voting assumption made typically in ideal-point models (e.g., Poole and Rosenthal, 1985; Clinton *et al.*, 2004). As will become clear from Section 1.3, these two assumptions are necessary for classical identification of the likelihood. However, they could be relaxed if a partially observed voting record is available to the analyst (see Section 4.3).

In most applications, vote choices will not be fully independent after conditioning on observables. However, this will not necessarily distort the inference as long as the correlation among vote choices is induced by unobservables that are independent of the covariate for which the analyst wants to estimate marginal effects. In this situation, the unobservables are said to be neglected heterogeneity that will only rescale the coefficient estimates in the same way that neglected heterogeneity affects probit models (e.g., Wooldridge 2001: 470). I provide more details in the Supplementary Information (SI-C).

I have also made two additional assumptions that could easily be relaxed. First, I assumed that the voting rule by which the committee makes decisions is known with certainty and followed strictly. Second, proposals are conditionally independent. I relax the latter assumption by modeling the unobserved heterogeneity across groups with a random intercept in the Supplementary Information (SI-D). The former assumption might be relaxed by modeling \mathcal{R} parametrically. I leave this extension to future work.

I refer to the model above as a multivariate probit model with partial observability or, for short, the partial m-probit. Multi- or k -variate probit models are usually employed to allow for correlated choices by estimating the correlation matrix from the data. Similar to the selection model for continuous outcomes popularized by Heckman (1976), bivariate probit models as selection models, for instance, allow for correlated error terms across a sample selection and a structural equation with binary outcomes (Dubin and Rivers, 1989). The problem addressed by the partial m-probit is not one of correlated (sequential) choices but of the nonobservability of the simultaneous choices.

1.2 Likelihood and prior density

The probability of observing a decision is the sum over the probabilities of the vote profiles that could have realized it. The probability of each of these vote profiles is the product over the individual choice probabilities, which are—as in a probit model—a linear function of covariates and parameters. The product over all decision probabilities yields the likelihood of the data. Next, I

define the probability of one vote profile and the sets of hypothetical vote profiles that can realize a particular decision outcome. Using these two definitions, I state the likelihood of the data.

Using the assumption of independent choice making, the probability of observing a vote profile \mathbf{y}_j is the product over the individual choice probabilities for proposal j or, equivalently, integrating over the latent utility in each dimension on the interval that corresponds to the observed vote choice. Formally, this

$$\begin{aligned}
 f(\mathbf{y}_j, \mathbf{X}_j | \boldsymbol{\beta}) &= \int_{p_{1j}} \dots \int_{p_{Mj}} \boldsymbol{\phi}(\mathbf{y}_j^* | \mathbf{X}_j \boldsymbol{\beta}) d\mathbf{y}_j^* \\
 &= \boldsymbol{\Phi}_{\mathcal{P}(\mathbf{y}_j)}(\mathbf{X}_j \boldsymbol{\beta}),
 \end{aligned}
 \tag{2}$$

where $\boldsymbol{\phi}(\cdot)$ is the M -dimensional multivariate normal density and p_{ij} is the interval that corresponds to the vote choice y_{ij} in the profile \mathbf{y}_j , that is, $p_{ij} = [0, \infty)$ if $y_{ij} = 1$ and $p_{ij} = (-\infty, 0)$ if $y_{ij} = 0$. To write this more compactly, I define $\mathcal{P}(\mathbf{y}_j)$ as the function that generates all p_{1j}, \dots, p_{Mj} given \mathbf{y}_j and let $\boldsymbol{\Phi}_{\mathcal{P}(\mathbf{y}_j)}(\cdot)$ be the implied distribution function.

Let $\tilde{\mathbf{y}}$ be a hypothetical vote profile and let $V(1)$ be the set of all hypothetical vote profiles for which $\sum_i \tilde{y}_i \geq \bar{R}$ holds. In other words, this set contains all vote profiles that realize an adoption outcome ($b_j = 1$). Let $V(0)$ be the complement set. Both sets are always finite but potentially large. For example, in the case of the UN Security Council, $V(1)$ is of size 848 and $V(0)$ of size 31,920.⁶

Using these two definitions, I can write the probability for $b_j = 1$ (and its complement) as the sum over the probabilities for all hypothetical vote profiles that can realize $b_j = 1$ ($b_j = 0$) and, after additionally relying on the conditional independence assumption across proposals, the likelihood is obtained by taking the product over all decisions. Formally, this is

$$\mathcal{L}(\boldsymbol{\beta} | \mathbf{X}, \mathbf{b}) = \prod_j \sum_{\tilde{\mathbf{y}} \in V(b_j)} \left[\boldsymbol{\Phi}_{\mathcal{P}(\tilde{\mathbf{y}})}(\mathbf{X}_j \boldsymbol{\beta}) \right].
 \tag{3}$$

Bayesian inference complements the likelihood with a prior density for the parameters (the coefficients). I follow convention and assume that they are jointly normal with a prior mean \mathbf{b}_0 and a diagonal covariance matrix \mathbf{B}_0 . The posterior density is proportional to the product of the likelihood function in Equation 3 and to the prior density.

The structure of the likelihood function is surprisingly general and can accommodate much more specific decision records than those with binary adoption/rejection information. Suppose, for example, that, in addition to knowing that the proposal passed, an analyst also knows that it passed with some vote margin. In this case, the set of permissible vote profiles V in Equation 3 can be substantially reduced. In fact, if the analyst knows how each and every member voted, that is, if there is a voting record, then V shrinks to a set with a single vote profile. In this case, Equation 3 reduces to a multivariate probit model, which is, since the covariance matrix is assumed to be the identity matrix, an ordinary probit model with $J \times M$ observations. There is also nothing in the structure of the likelihood that precludes the amount of information from varying across decisions. This implies that a partially observed voting record can be accommodated within the likelihood function without difficulty or formal extensions.

Finally, it is worth placing this model in the broader context of the (statistical) literature. The bivariate probit with partial observability by Poirier (1980) and the bilateral cooperation model by

⁶Since the five permanent members must always agree, the number of adopting coalitions is identical to the number of adopting coalitions among the ten nonpermanent members, which is given by $\sum_{m=4}^{10} \binom{10}{m}$. The number of rejecting coalitions is given by the difference between the total number of coalitions among all members minus the number of adopting coalitions, that is, $\sum_{m=0}^{15} \binom{15}{m} - \sum_{m=4}^{10} \binom{10}{m}$.

Przeworski and Vreeland (2002) emerge as special cases of Equation 3 if $M = 2$ and the voting rule is unanimity. In a recent contribution, Poirier (2014) extended his 1980 model to the case of $M > 2$ but remains focused on the case of unanimity.⁷ More importantly, each factor in the likelihood above is the (complementary) cumulative density function of Poisson's Binomial distribution⁸ parameterized with a set of probit functions (proofs for both statements appear in SI-A).

1.3 Identification

Before I continue with the computation of the posterior distribution, I discuss the (classical) parametric identification of the likelihood. A likelihood is said to be (parametrically) identified if a unique set of estimates exists for the parameters of a model. In the Supplementary Information (SI-B), I show that the conditional mean for the likelihood in Equation 3 is always identified. The system of nonlinear equations that maps the structural parameters β to the reduced-form conditional means is identified under some conditions. Using a linearization of this system with a first-order Taylor series expansion, I show that it is identified if the aggregate design matrix has full rank. The aggregate design matrix of dimension $J \times K$ results from stacking the J vectors that result from column-averaging all X_j matrices on top of each other.

The classical parametric identification condition is empirically verifiable by checking if the design matrix of the model (X), after averaging all variables for each decision, has linear independent columns. Trivially, this condition will fail whether the design matrix before averaging does not have full rank. However, it will also fail if the design matrix has full rank and a variable exhibits variation within but not across decisions. In that case, the variable will be constant after averaging as well as a linear combination of the intercept. This renders the aggregate design matrix less than full rank, and the effect of the respective variable and the intercept are not separately identifiable. In practice, this implies, for example, that, for a committee with constant membership, fixed effects for members or member-specific effects are unidentifiable and consequently cannot be estimated.

The identification condition is based on the linearization of a system of nonlinear equations. Consequently, there might be instances where the condition of a full-rank aggregate design matrix holds but a unique set of parameters still does not exist. This is problematic for frequentist inference because, for example, the properties of the maximum likelihood estimator are at least inconvenient for unidentified likelihoods. However, in a Bayesian analysis, unidentified likelihoods are of less concern since the posterior density will still be proper if proper priors are used. The only consequence is that the marginal posterior density of the intercept will be, in the worst case, perfectly negatively correlated with the marginal posterior density of the unidentified effect. From a theoretical perspective, this is not a problem, but in practice, it means that the Gibbs sampler presented in the next section will be very slow in exploring the posterior density, which is why an identified likelihood is advantageous for a Bayesian analysis.

2. Posterior computation

As in most Bayesian models, the posterior density cannot be marginalized analytically, which prompts me to construct a Gibbs sampler to simulate from the density and use the samples to characterize the density with a desirable degree of accuracy. A Gibbs sampler requires derivation

⁷Poirier (1980, 2014) allows for member-specific effects and correlated choices among agents. However, his theoretical results suggest that, even under a unanimity voting rule, member-specific effects and the correlation among agents' choices are at best only partially identifiable from the data. See also Section 1.3 in this paper.

⁸In the most comprehensive paper on the distribution, Wang (1993) follows a reviewer suggestion and refers to the density as Poisson's binomial density. While this name choice is much less confusing than earlier conventions, an even more descriptive name might be "heterogeneous binomial density" since it emerges from the convolution of heterogeneous Bernoulli densities, and the standard binomial density is a special case if all choice probabilities are identical.

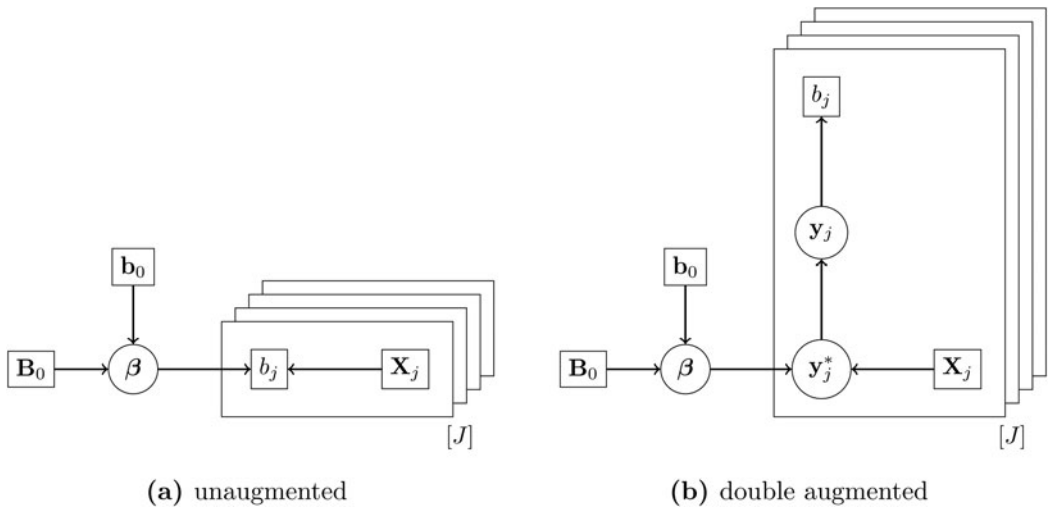


Figure 1. Two directed acyclic graphs of the partial m-probit. (a) Unaugmented. (b) Double augmented.

of the full conditional densities for all unknown quantities in the model. To derive them, I use a theorem by Lauritzen *et al.* (1990), who show that, if a joint density (such as a posterior density) can be written as a directed acyclic graph (DAG), the full conditionals are given by a simple formula (see SI-D).

A DAG representation of the posterior density appears in Figure 1(a). Each node in this graph is a random variable. Rectangular nodes indicate observed variables (the data and hyperparameters), while circle nodes represent unobserved variables (parameters). An arrow indicates the dependencies between these variables, and the plates indicate the J replications. The graph is acyclic since it has no cyclic dependency structure.

The conditional for β in Figure 1(a) is not a member of a known parametric family from which samples can be easily drawn. To arrive at full conditionals that are easy to sample from, I follow a data augmentation strategy (Tanner and Wong, 1987) and explicitly introduce two variables from the derivation of the likelihood. The augmented DAG appears in Figure 1(b). The first augmentation is identical to the Albert–Chib augmentation in a Bayesian (multivariate) probit model (Albert and Chib, 1993; Chib and Greenberg, 1998), explicitly introducing y_j^* , the latent utility, in the model. The second augmentation augments the latent utility with y_j , the unobserved votes. Because of this sequential augmentation, I refer to the Gibbs sampler as a double-augmented Gibbs sampler.

Applying the result from Lauritzen *et al.* (1990) cited above yields three full conditionals for the three unobserved variables in the DAG. The conditional for β can then be written as follows:

$$\begin{aligned}
 f(\beta | \mathbf{b}_0, \mathbf{B}_0, \mathbf{y}^*, \mathbf{y}, \mathbf{b}, \mathbf{X}) &\propto f(\beta | \mathbf{b}_0, \mathbf{B}_0) \times \prod_j f(y_j^* | X_j, \beta) \\
 &= \phi((\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{X})^{-1}(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{X}'\mathbf{y}^*), \\
 &\quad (\mathbf{B}_0^{-1} + \mathbf{X}'\mathbf{X})^{-1}).
 \end{aligned}
 \tag{4}$$

The two other conditionals and their sampling algorithms are given in the Supplementary Information (SI-D).

It is not a coincidence that the functional form of the conditional for β is exactly the same as the conditional for an ordinary probit model and a Bayesian normal regression model when the

same prior for β is chosen. The primary difference between a probit model, a partial m-probit, and a normal regression is that only in the latter case is the variable y^* fully observed. In the other two cases, y^* is observed only in a coerced fashion. However, the precise nature of the coercion is irrelevant once the data are augmented. In fact, the very purpose of the data-augmentation strategy is to render the coefficients conditionally independent of the coerced data.

The Gibbs sampler, which I refer to as the double-augmented Gibbs sampler, is an iterative sampling from the conditionals until convergence (see SI-D for the details). It has a very intuitive sequence: (1) choose some starting value for the coefficients; (2) conditional on these values, the covariates, and the decision record, draw vote profiles for all decisions; (3) conditional on the vote profiles and the covariates, draw the vector of latent utilities for all decisions; (4) conditional on the latent utilities and the covariates, draw the coefficients; and (5) repeat until convergence.

The Gibbs sampler is implemented in an open-source R-package *consilium*, which accompanies this paper. I also conducted Monte Carlo experiments to verify that the Gibbs sampler (and its implementation) obtains samples from the posterior density and to provide some insights into the computational costs of the model (see SI-F).

3. Aggregation costs

Whenever data are aggregated, the analyst pays a price in terms of (a) effects that cannot be estimated, (b) posterior uncertainty (efficiency), and (c) bias for the estimable effects. What are the costs when analyzing a decision record relative to an analysis with a voting record? The discussion on identification has highlighted that member-specific effects in committees with constant membership cannot be estimated with decision-record data. This is in sharp contrast to voting records where member-specific effects can be estimated. If such effects are the object of inquiry, decision records cannot be used. Moreover, even if the effect of interest is assumed to be shared, its inference might be hampered if the analyst suspects relevant, unobserved member-specific heterogeneity. While such heterogeneity could be modeled with varying intercepts in an analysis of voting records, it is infeasible with decision records.

For estimable effects, posterior uncertainty and aggregation bias are further potential costs. Aggregation bias, as discussed in the classical ecological inference literature (Erbring, 1989; King, 1997), is a form of confounding with the group-assignment variable. Since the number of groups equals the number of observations in the aggregated sample, adjustment strategies, that is, weighting with or conditioning by the group-assignment variable, are not feasible. However, if the group-assignment variable is chosen at random, grouping cannot lead to bias, the classical example for aggregation bias being spatially aggregated data on vote choice and race in mass elections. To the extent that electoral districts are drawn with perfect knowledge about vote choice and race in an election, the effect of race on vote choice in the same election cannot be inferred without bias.

For aggregation bias to be a threat to inference with decision records, the process of assigning members to decisions (the “groups”) must be a function of members’ vote choices on a proposal and some unmeasured covariate. If that were the case, then the proposal-assignment vector would be a confounder for which we cannot adjust and aggregation bias is unavoidable.⁹ While membership in a committee is presumably a function of (expected) vote choices and potentially some unmeasured covariates, the committee’s membership is usually constant over a certain period. Within this period of constant membership, aggregation bias cannot occur.

Beyond aggregation bias, there is also the issue of posterior uncertainty since aggregation reduces the effective sample size. While posterior uncertainty might seem secondary, it becomes

⁹In SI-E, I illustrate in a proof for the bivariate case with vague priors that the posterior mean of the coefficient is not a function of the vote choices (but only the decision record) if the proposal-assignment vector is orthogonal to the latent utility and the covariates. The proof follows directly from classical results in ecological inference (e.g., Erbring, 1989; Palmquist, 1993; King, 1997) that give extensions to the multivariate case.

paramount once the aggregation reduces information to a point where no variation is left to draw inference from. For instance, in an institution where all members have a high (low) average probability of voting one way or the other, there is a chance that the decision record will exhibit no variation and the posterior equals the prior.¹⁰

4. Advantages of the model

Unsurprisingly, the structural model tends to produce more efficient estimates since the amount of information in the estimation is larger. More importantly, the structural model allows one (a) to choose the correct reduced-form specification, (b) to estimate vote-choice probabilities instead of adoption probabilities, and (c) to combine partially observed voting records with decision records.

4.1 Choosing specifications

Decision records are used for empirical inference on a regular basis with convenient models such as a probit. However, perhaps surprisingly, the specification that is usually chosen is not the reduced-form complement to the structural model outlined in the previous section. As an example, consider this simple partial m-probit:

$$p(b_j = 1) = \sum_{\tilde{y} \in V(1)} \Phi_{P(\tilde{y})}(\beta_0 + \mathbf{x}_j \beta_1) \quad (5)$$

and one reduced-form complement with $z_j = \sum_i x_{ij}$:

$$p(b_j = 1) = \Phi(\beta'_0 + z_j \beta'_1), \quad (6)$$

where one might scale z_j by dividing by M , which then makes z_j the average of \mathbf{x}_j .¹¹

However, typically, the sum in Equation 6 is not taken over all members but only a subset. For example, in studies on the UN Security Council, measures of political or economic closeness between the conflict location and the permanent members are included (e.g., an indicator for a defense alliance), although the Council consists of the five permanent and ten nonpermanent members (e.g., Gilligan and Stedman, 2003; Mullenbach, 2005; Beardsley and Schmidt, 2012; Hultman, 2013; Stojek and Tir, 2015).

However, leaving out parts of the membership introduces measurement error in z_j .¹² As in any other setting with errors in variables, the resulting coefficient estimates will be biased. Moreover, the estimated effect cannot generally be interpreted as a member-specific effect since, as shown in Section 1.3, there is no variation in decision-record data that can identify member-specific effects.

4.2 Estimating vote-choice probabilities

Both the structural model and the reduced-form model allow one to estimate the predicted probability of observing the adoption of a proposal (the “adoption probability”). These predicted probabilities can be used to characterize how much a one-unit increase in a covariate changes

¹⁰The aforementioned Monte Carlo experiments in SI-F also provide some intuition about the increase in posterior uncertainty that comes with aggregation.

¹¹Strictly speaking, the probit model (or logit model) has no theoretical basis as a reduced-form complement but is often used in practice. A theoretical justified reduced-form complement is the linear regression model that provides the best linear approximation to the nonlinear conditional expectation function of the data.

¹²Only if the means of those members included and those left out coincide for all decisions is no measurement error introduced.

the adoption probability. In addition to the adoption probability, the structural model also allows one to calculate the predicted probability of a supportive vote choice (the “vote choice probability”). This quantity is typically calculated when one analyzes a voting record and can be used to describe how a one-unit change in a covariate changes the vote-choice probability.

While the adoption probability can be of considerable interest in some situations (e.g., if the analyst intends to predict the adoption of proposals), it must be recognized that it is not only a function of the coefficients and the covariates but also of the institutional structure (the size of the membership and the majority threshold). Consequently, it is a conditional probability whose magnitude, as it turns out, is not a linear function of the vote-choice probability.

To illustrate, consider a committee of 20 members with various majority thresholds between 11 (a simple majority) and 20 (unanimity). To simplify matters without loss of generality, suppose also that the vote-choice probability for all members is homogeneous at 0.75. The vote-choice probabilities are shown with a solid line in Figure 2. The figure also shows, corresponding to each of these vote-choice probabilities, the implied adoption probabilities conditional on the 10 majority thresholds (dashes). While the vote-choice probabilities are constant across committees of various sizes, the adoption probabilities are a monotone, but nonlinear, function of the vote-choice probabilities.

The monotonicity of the adoption probability with respect to the vote-choice probability is good news because it suggests that the direction of any effect on the vote-choice probability can always be inferred from the direction of the effect on the adoption probability. However, the nonlinearity also suggests that the adoption probability cannot be easily compared across different institutional contexts. Figure 2 illustrates that, even in the absence of differences in vote-choice probabilities in two different institutional contexts, adoption probabilities will vary if the membership or majority threshold differs.

Furthermore, the magnitude of the adoption probability can be a very poor indicator of the magnitude of the vote-choice probability. Figure 2 illustrates that the closer the majority threshold moves toward unanimity, the smaller the adoption probability becomes up to the point where it is minuscule. All the while, the vote-choice probability remains constant. This emphasizes that it is quite important to define what the quantity of interest is when analyzing decision records. If the analyst’s interest is in understanding how covariates change the vote-choice probability, the structural model is the more promising approach.

4.3 Including a partially observed voting record

The discussions on the likelihood function and the Gibbs sampler have already highlighted that including a partially observed voting record is very easy when using the structural model but infeasible when using a reduced-form model. Ordering the proposals for which only the decision record is available from $j = 1, \dots, K$ and the proposals for which a voting record is available from $j = K + 1, \dots, J$, the two-component likelihood function with parameter vector β takes the following form:

$$\hat{\mathcal{L}}(\beta | \mathbf{X}, \mathbf{b}, \mathbf{Y}) = \prod_{j=1}^K \sum_{\tilde{y} \in V(b_j)} \left(\Phi_{P(\tilde{y})}(\mathbf{X}_j \beta) \right) \cdot \prod_{j=K+1}^J \left(\Phi_{P(y_j)}(\mathbf{X}_j \beta) \right), \tag{7}$$

where \mathbf{Y} denotes the stacked matrix of all observed voting profiles. The Gibbs sampler is easy to expand by simply dropping the sampling of the vote profiles for those proposals where a voting record is available. It is fairly intuitive that the posterior inference from this likelihood will be more certain than the posterior inference from the likelihood in Equation 3.

Including a partially observed voting record can also reduce aggregation bias. In the Supplementary Information (SI-E), I show that the familiar missing-at-random (MAR) condition

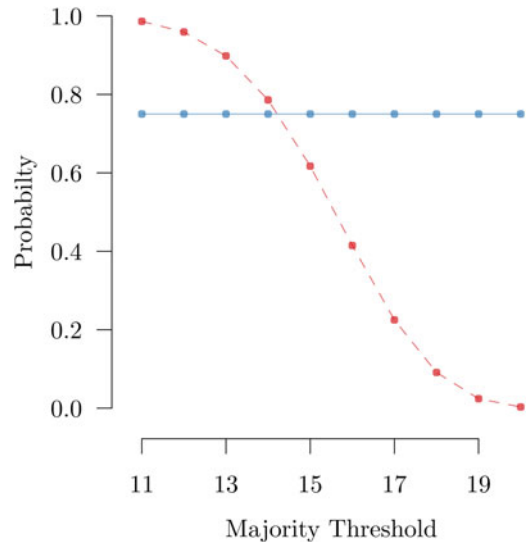


Figure 2. Illustration of the relationship between vote-choice and adoption probabilities for a committee of 20 members. The solid line indicates the different vote-choice probabilities and the dashed line the corresponding adoption probabilities.

from the literature on missing data (Little and Rubin, 2002) is a necessary assumption to reduce aggregation bias. In particular, it is necessary that, conditional on the covariates, the observability of the recorded votes is random. If this assumption is fulfilled, aggregation bias will be removed from the estimates. Complementarily, if the observed voting record is a nonrandom subset, it might cause selection bias if incorporated.

Another benefit of including a partially observed voting record is that one can relax the assumption of shared effects across all committee members. These effects are obviously identifiable from voting records and, as discussed in Section 1.3, unidentifiable with decision records. Consequently, if member-specific effects are of interest and included in the model, the identifying variation to estimate these effects will come from the variation in the partially observed voting record. The conditional-independence assumption with respect to vote choices could also be relaxed for the same reasons.

The ability to supplement a decision record with a partially observed voting record can also have advantageous consequences for data collection. Consider, for example, a situation where the analyst wishes to collect another sample of votes from a voting record to decrease the posterior uncertainty but, as it happens, collecting such a sample proves quite expensive. To avoid these costs, the analyst could instead collect a large sample from the decision record. To the extent that collecting a large sample from a decision record is much cheaper than a sample from the voting record, this reduces the costs of data collection.

5. Replication: US State Supreme Court decisions

I replicate a study by Caldarone *et al.* (2009) to contrast the coefficient estimates when a voting record is used with the coefficient estimates used when I artificially delete (some of) the recorded votes and only use the decision record in the analysis. Caldarone *et al.* (2009) test the prediction “that nonpartisan elections increase the incentives of judges to cater to voters’ ideological leanings” (p. 563). To test their prediction, the authors assemble a dataset of US state supreme court decisions on abortion for the period from 1980 to 2006. They collect these data for all state supreme courts for which judges face contested statewide elections. Their dataset contains 19 state supreme courts (which vary in size between five and nine judges) and a total of 85 abortion decisions.

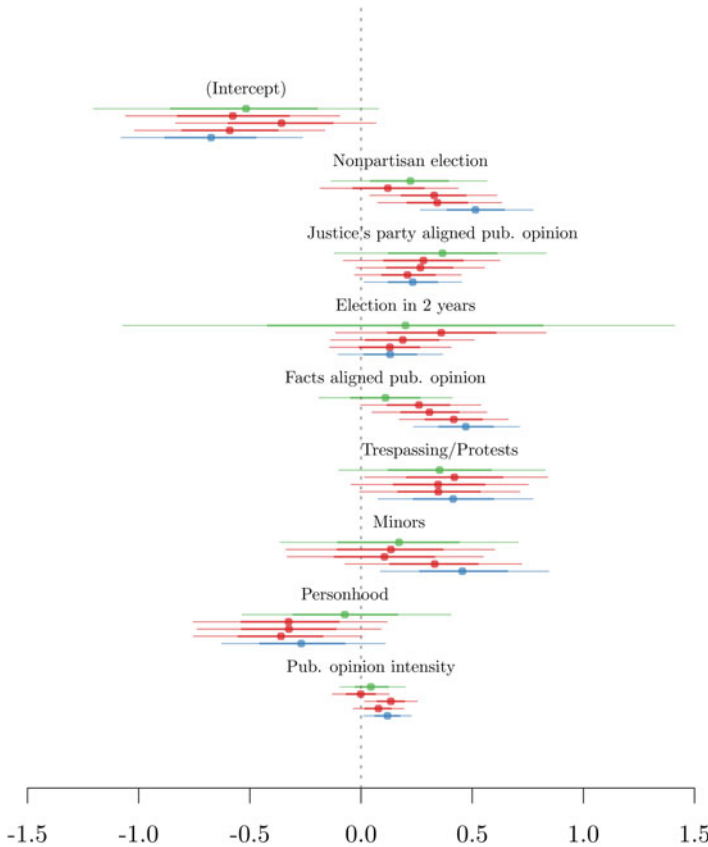


Figure 3. Regression results from a Bayesian partial m-probit model with a decision record (row 1) and a partially observed voting records (row 2: 25 percent observed, row 3: 50 percent, row 4: 75 percent) as well as a Bayesian probit model with justices’ voting record (row 5). While the dots indicate the posterior mean, the segments represent the 95 and 68 percent posterior intervals, respectively.

The dependent variable in the authors’ analysis is a regular justice’s vote. Using state-level opinion data, the authors code each justice’s vote as either popular (if it leans toward the state’s public opinion) or unpopular. Consequently, the dependent variable takes a 1 if the justice votes “pro-choice” and the state leans “pro-choice” or if he votes “pro-life” and the state leans “pro-life” (Caldarone *et al.*, 2009: 565). In the authors’ dataset, 261 votes are popular (43 percent). The authors’ independent variable of interest is a binary variable indicating whether a supreme court justice was elected in a nonpartisan election. Of the 85 abortion decisions, 39 were made in a partisan electoral environment (46 percent).

A replication of the authors’ baseline specification (model 1 in their table) using a Bayesian probit model appears as the lower row (row 5) in the coefficient plot in Figure 3. The upper row (row 1) instead shows the results produced when I retained only a binary variable indicating whether the courts passed a popular decision by majority rule and estimated the same specification using the partial m-probit.¹³ Dropping all votes leaves me with 36 popular rulings (42 percent). In essence, dropping all votes reduces the number of observations for the left-hand side of the regression equation to 85, while it leaves the observations on the right-hand side unaffected ($N = 605$).

For the main variable of interest, nonpartisan election, the posterior probability that there is a positive effect of nonpartisan elections is still 0.9 even after dropping all votes and despite the sharp decrease in available information on the left-hand side of the regression equation. The

¹³For both models, I use vague normal priors centered at 0 with a variance of 10. The regression table in SI-H contains detailed information on the Gibbs sampling parameters and convergence.

estimated effects for the two controls, which exhibit within-case variance, are notable. The effect of elections in two years is estimated with a similar posterior mean but with considerably larger posterior uncertainty. The effect of the justices' party being aligned with public opinion is estimated to be a little larger and to have more posterior uncertainty.

One benefit of the structural model is that it allows one to combine a partially observed voting record with a decision record to decrease the costs of aggregation. To demonstrate this, I re-estimate the partial m-probit with random samples of recorded votes and the same prior. The results appear again in the same coefficient plot (row 2–4). The upper bars (row 2) show the estimates when, in addition to the decision record, 25 percent of all votes are observed, followed by the estimates for 50 and 75 percent. As expected, the greater the number of recorded votes included in the analysis is, the higher the similarity will be between the estimates of the partial m-probit and the ordinary probit. For most variables, the trend toward the probit estimates and the decrease in posterior uncertainty appear to be quite linear (e.g., for nonpartisan election or the justices' party alignment). However, for some, there is a significant payoff for observing some votes compared to no votes (e.g., elections in two years). This suggests that, at least in some situations, collecting a few votes to supplement the decision record can greatly improve the quality of the estimates.

6. Application: Trade and UN operations

A major line of inquiry in the literature on the UN Security Council aims to understand Council members' motives in involving themselves in third-party conflicts within the framework of the United Nations (e.g., Gilligan and Stedman, 2003; Hultman, 2013; Stojek and Tir, 2015). Are the members more likely to support a UN Blue Helmet operation in conflicts where they expect economic or political gains from a swift end to the conflict? I reconsider this question by estimating the effect of trade relationships between the members and the territories in conflict, highlighting the advantages of using the partial m-probit.

To conduct this analysis, I use a revised version of the cross-sectional panel dataset by Hultman (2013), which combines the UCDP/PRIO Armed Conflict dataset (Gleditsch *et al.*, 2002) with the dataset on third-party interventions by Mullenbach (2005). Focusing on intrastate conflicts that occurred outside the territories of the Council's permanent members, the effective number of observations is 885, nested in 102 conflicts. There are 17 conflicts for which the UN Security Council deployed a UN operation.

I interpret each observation as an instance where each of the 15 Council members¹⁴ must decide to support or oppose the deployment of a UN operation. Consequently, the unit of analysis in my dataset is a UN Security Council member's binary support choice per conflict-year. I supplement these data with information about the size of total trade (export and imports) between a Council member and the conflict location (Barbieri *et al.*, 2009).¹⁵

There is no complete voting record from the UN Security Council. While some votes from the UN Security Council are on record and could be incorporated, these recorded votes constitute a selected sample from the set of all votes. This is because the Council conveys "in public only to adopt resolutions already agreed upon" (Cryer, 1996: 518). "By the time the resolutions come to a vote, it is usually known by all how much support there will be for each" (Luard, 1994: 19). Most conflicts are never discussed in the Council or they are discussed but the Council cannot agree on whether to deploy a UN operation. Consequently, recorded votes only occur in very particular circumstances (if the Council agrees to deploy) and incorporating these recorded votes is likely to result in a selection bias.

¹⁴The UN Security Council uses a veto-majority rule. Specifically, nine out of 15 member states must approve a proposal, but each of the five permanent members (China, France, Russia, the United Kingdom, and the United States) has a veto.

¹⁵For a detailed description of the data, see SI-G.

Table 2. Regression results from a Bayesian partial m-probit model (model 1, $N = 15 \times 885$) and seven Bayesian probit models (model 2–8, $N = 885$) each with posterior means and 95 percent posterior intervals in parentheses

Model	Variable	Posterior
Model 1	log(Trade)	–1.09 [– 2.61, 0.15]
Model 2	log(Trade) of all mbrs.	–1.02 [– 3.05, 0.66]
Misspecified		
Model 3	log(Trade) of P5 only	–1.40 [– 3.55, 0.26]
Model 4	log(Trade) of US only	–1.02 [– 3.02, 0.71]
Model 5	log(Trade) of UK only	–2.97 [– 6.36, – 0.50]
Model 6	log(Trade) of FR only	–0.15 [– 1.60, 1.63]
Model 7	log(Trade) of RU only	–2.09 [– 4.75, – 0.24]
Model 8	log(Trade) of CN only	–1.31 [– 3.34, 0.41]

All models include covariates, varying intercepts and B-splines ($df = 3$).

I condition on a set of common causes to decrease the threat of confounding and also include a varying intercept for the conflict location. In order to account for annual and conflict-period trends, I include two B-splines (with the deployment year and the period of the conflict). Except for the binary independent variables, I center and scale all variables by twice their standard deviation before estimating each model, which aids in the construction of weakly informative, normal priors centered at 0 and a variance of 5.

The estimates appear in [Table 2](#) in the row labeled model 1 (see also SI-H, for the full table and details on the Gibbs sampling parameters and convergence). The estimates suggest that an increase in trade between a Council member and the conflict location decreases a member's probability of supporting a UN operation. The posterior probability for this effect to be negative is 0.95.

To illustrate the difference between the inference from the partial m-probit and a reduced-form model, I aggregate the data to a dataset of conflict-years. In the aggregated dataset, the trade variable measures the total trade of all Council members with the conflict location. The estimates from a probit model appear in [Table 2](#) in the row labeled model 2. As expected, the sign of the association is identical to model 1. Interestingly, the posterior probability for this association to be negative is only 0.89—reflecting that the partial m-probit delivers more efficient estimates. Notice that the magnitude of the coefficient from model 2 provides no information about how trade between a Council member and the conflict location decreases a member's probability of supporting a UN operation. This information is only available from the partial m-probit estimate.

Typical studies on the UN Security Council¹⁶ do not include covariates that measure the variation of a concept across *all* members but, rather, usually focus on the permanent five (the P5). To illustrate that this can lead to a misleading inference in the trade case, I estimate the effect of total trade of the P5 leaving out the contribution from the ten nonpermanent members (see row labeled model 3 in the [Table 2](#)) and include each P5 trade share separately (rows labeled models 4–8). As explained in [Section 4.1](#), none of these estimates can be interpreted as estimates of the

¹⁶See, for example, Gilligan and Stedman (2003); Mullenbach (2005); Fortna (2008); Beardsley and Schmidt (2012); Hultman (2013) and Stojek and Tir (2015).

effect of trade on the respective members' vote choices (or the heterogeneous effect of trade on members in general). Instead, the estimates from the models 3–8 can be interpreted as a version of the estimates in model 2 but contaminated by measurement error.

The results here are at odds with the recent analysis by Stojek and Tir (2015). Using data from Fortna (2008) and a logit model on UN peacekeeping deployment, they estimate a positive effect of the P5 total trade volume on the probability of deployment. While their unit of analysis is a ceasefire analysis, the positive effect they estimate is largely driven by conflicts in which permanent members are directly involved (e.g., the Northern Ireland conflict), while the data I use exclude all conflicts that occur in the territory of the permanent member states.

7. Discussion

Analyzing a decision record instead of a voting record is not something for which one would hope. The aggregation of vote choices by a voting rule increases the uncertainty of estimable effects and may even bias them. It also prohibits the estimation of member-specific effects. However, confronted with the choice between abstaining from an analysis or relying on decision records, an analyst might still prefer the latter. In this paper, I argue that, if the analyst decides to examine the decision record, his or her analysis can be improved by turning to a structural model instead of opting for a convenient reduced-form model.

In this paper, I highlight several advantages of the structural model; however, the most important might be that it allows one to bring partially observed voting records into the analysis. *Inter alia*, the replication of the study by Caldarone *et al.* (2009) highlights that there are large benefits in terms of efficiency when analyzing a decision record jointly with a voting record sample even if the later is small. Beyond efficiency, such a joined analysis opens the route to estimate member-specific effects as well as reduce potential aggregation bias. This suggests that effort should be made to collect a sample of votes from archival documents or committee members' personal notes. Potentially, even if no explicit voting record is provided in existing documents, it might be still feasible to reconstruct a small set of votes with high confidence based on in-depth qualitative research.

Beyond the question of which model to use to analyze an available decision record, one might wonder for which (types of) institutions one can successfully compile a decision record in the first place and then apply the partial m-probit. While a systematic listing is beyond the scope of this paper, a few examples might highlight that decision records are either directly available from particular institutions or can be compiled based on available knowledge about these institutions.

A decision record is typically available from institutions where members vote on a regular basis on issues but decide not to publish these votes. While I have artificially created a decision record in the case of the US state supreme courts in Section 5, international courts in particular (e.g., the European Court of Justice or the European Court of Human Rights) typically publish only the decisions on each case but not judges' votes.¹⁷ Another example in this category is central banks other than those of the US and UK where voting records are published.

However, even committees that do not explicitly vote on each decision may adopt proposals by acclamation on a regular basis, which gives rise to a decision record that can be analyzed. The UN Security Council analyzed in Section 6 is a case in point: the Council explicitly only votes on the deployment of UN peace operations that are known to pass but implicitly rejects all UN peace operations in ongoing conflicts by never advancing them to the voting stage in the first place. Another example is the IMF Executive Board, which approves loans by acclamation instead of voting and whose decision record has been analyzed previously using reduced-form models (Broz and Hawes, 2006; Copelovitch, 2010; Breen, 2013).

¹⁷See, for example, Carrubba *et al.* (2008) and Helfer and Voeten (2014) for studies using decision records from international courts.

However, not every institution's decision record will be suitable for analysis nor will it always be possible to compile a decision record in the first place. The ability to compile a decision record when it is not directly published by an institution depends on the availability of a natural agenda that defines the issues under consideration at the respective institution. In the case of the UN Security Council, for example, studies assume that the agenda is defined by the set of ongoing conflicts. Suitable decision records are those where the conflict between committee members across decisions evolves around a binary decision "to do something or not". However, if the conflict across decisions is determined by a conflict over how much to do (and consequently something is always done), the analysis of decisions records will provide little further insights into the institution.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2021.11>.

References

- Albert JH and Chib S** (1993) Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* **422**: 669–679.
- Austen-Smith D and Banks JS** (1999) *Positive Political Theory I: Collective Preference*. Ann Arbor: University of Michigan Press.
- Barbieri K, Keshk OMG and Pollins BM** (2009) Trading data evaluating our assumptions and coding rules. *Conflict Management and Peace Science* **26**: 471–491.
- Beardsley K and Schmidt H** (2012) Following the flag or following the charter? Examining the determinants of UN involvement in International crises, 1945–2002. *International Studies Quarterly* **56**: 33–49.
- Breen M** (2013) *The Politics of IMF Lending*. Basingstoke: Palgrave Macmillan.
- Broz LJ and Hawes MB** (2006) U.S. domestic politics and International Monetary Fund policy. In Hawkins DG, Lake DA, Nielson DL and Tierney MJ (eds). *Delegation and Agency in International Organizations*. Cambridge: Cambridge University Press, pp. 77–106.
- Caldarone RP, Canes-Wrone B and Clark TS** (2009) Partisan labels and democratic accountability: an analysis of State Supreme Court abortion decisions. *Journal of Politics* **71**: 560–573.
- Carrubba CJ, Gabel M and Hankla C** (2008) Judicial behavior under political constraints: evidence from the European Court of Justice. *American Political Science Review* **102**: 435–452.
- Chib S and Greenberg E** (1998) Analysis of multivariate probit models. *Biometrika* **85**: 347–361.
- Clinton JD, Jackman S and Rivers D** (2004) The statistical analysis of roll call data. *American Political Science Review* **98**: 355–370.
- Copelovitch MS** (2010) *The International Monetary Fund in the Global Economy: Banks, Bonds, and Bailouts*. Cambridge: Cambridge University Press.
- Cryer R** (1996) The Security Council and Article 39: a threat to coherence. *Journal of Conflict and Security Law* **1**: 161–195.
- Dubin JA and Rivers D** (1989) Selection bias in linear regression, logit and probit models. *Sociological Methods and Research* **18**: 360–390.
- Erbring L** (1989) Individuals Writ large: an epilogue on the "Ecological Fallacy". *Political Analysis* **1**: 235–269.
- Fortna VP** (2008) *Does Peacekeeping Work?* Princeton: Princeton University Press.
- Fry M, Julius D, Mahadeva L, Roger S and Sterne G** (2000) Key issues in the choice of monetary policy framework. In Mahadeva L and Sterne G (eds). *Monetary Policy Frameworks in a Global Context*. London: Routledge, pp. 1–216.
- Gilligan M and Stedman SJ** (2003) Where do the peacekeepers go?. *International Studies Review* **5**: 37–54.
- Gleditsch NP, Wallensteen P, Eriksson M, Sollenberg M and Strand H** (2002) Armed conflict 1946–2001: a new dataset. *Journal of Peace Research* **39**: 615–637.
- Glynn A, Wakefield J, Handcock MS and Richardson TS** (2008) Alleviating linear ecological bias and optimal design with subsample data. *Journal of the Royal Statistical Society. Series A* **171**: 179–202.
- Hanseuse S J-P A and Wakefield J** (2008) The combination of ecological and case-control data. *Journal of the Royal Statistical Society. Series B* **70**: 73–93.
- Heckman JJ** (1976) The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement* **5**: 475–492.
- Helfer LR and Voeten E** (2014) International courts as agents of legal change: evidence from LGBT rights in Europe. *International Organization* **68**: 77–110.
- Hultman L** (2013) UN peace operations and protection of civilians: cheap talk or norm implementation. *Journal of Peace Research* **50**: 59–73.
- King G** (1997) *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton: Princeton University Press.

- Lauritzen SL, Dawid AP, Larsen BN and Leimer H-G** (1990) Independence properties of directed markov fields. *Networks* **20**: 491–505.
- Little RJA and Rubin DB** (2002) *Statistical Analysis with Missing Data*. New York: Wiley.
- Luard E** (1994) *The United Nations: How It Works and What It Does*, 2nd Edn. London: Maxmillan Press.
- Mullenbach MJ** (2005) Deciding to keep peace. An analysis of international influences on the establishment of third-party peacekeeping missions. *International Studies Quarterly* **49**: 529–555.
- Palmquist BL** (1993) *Ecological Inference, Aggregate Data Analysis of U.S. Elections, and the Socialist Party of America* (Ph.D. thesis). University of California, Berkeley.
- Poirier DJ** (1980) Partial observability in bivariate probit models. *Journal of Econometrics* **12**: 209–217.
- Poirier DJ** (2014) Identification in multivariate partial observability probit. *International Journal of Mathematical Modelling and Numerical Optimisation* **5**: 45–63.
- Poole KT and Rosenthal H** (1985) A spatial model for legislative roll call analysis. *American Journal of Political Science* **29**: 357–384.
- Przeworski A and Vreeland JR** (2002) A statistical model of bilateral cooperation. *Political Analysis* **10**: 101–112.
- Schermers HG and Blokker NM** (2011) *International Institutional Law*, 5th Edn., Leiden: Nijhoff.
- Stojek SM and Tir J** (2015) The supply side of united nations peacekeeping operations: trade ties and united nations-led deployments to Civil War States. *European Journal of International Relations* **21**: 352–376.
- Tanner MA and Wong WH** (1987) The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* **82**: 528–540.
- Wakefield J and Salway R** (2001) A statistical framework for ecological and aggregate studies. *Journal of the Royal Statistical Society. Series A* **164**: 119–137.
- Wang YH** (1993) On the number of successes in independent trials. *Statistica Sinica* **3**: 295–312.
- Wooldridge JM** (2001) *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press.