Appendix for
Using Item Response Theory to Improve Measurement in Strategic Management Research: An Application to Corporate Social Responsibility

August 15, 2015

In this appendix, we give a more complete account of our statistical model's assumptions and of the MCMC routine for posterior simulation. Our approach is similar to that of Martin and Quinn (2002). Indeed, their model is a special case of ours; with certain assumptions, the two approaches are identical. For the purposes of introducing the model to a new audience, and to demonstrate how existing software can be used for estimation, we maintain as much similarity to Martin and Quinn (2002) as possible and highlight points of potential departure.

In the main body of the paper, the latent utility received by firm i for choosing to adopt policy j in time period t is represented by

$$z_{i,j,t} = \alpha_{j,t} + \beta_{j,t}\rho_{i,t} + \varepsilon_{i,j,t}.$$

All of the terms except for $\varepsilon_{i,j,t}$ are discussed in the main body of the paper. Following common practice, we assume that $\varepsilon_{i,j,t}$ is identically and independently distributed with mean zero and variance one. Our aim is to estimate the α, β, and ρ terms. To discuss how we do so, we simplify notation. Collect the respective $\alpha_{j,t}$ terms for all the observables in all the time periods into a vector $\boldsymbol{\alpha}$, and do the same for the respective $\beta_{j,t}$ terms into a vector $\boldsymbol{\beta}$. Collect all the latent trait scores of all firms in all time periods into a vector $\boldsymbol{\rho}$. Finally, collect all the decisions on all observables by all observed firms in all time periods into a vector $\boldsymbol{d}$. We adopt a Bayesian inference protocol, so our object of interest is the posterior density

$$p\left(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho} \mid \boldsymbol{d}\right) \;\propto\; p\left(\boldsymbol{d} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}\right) p\left(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}\right).$$

Here the symbol $\propto$ denotes proportionality: the term on the right-hand side of the expression is the numerator familiar from Bayes' rule, where $p\left(\boldsymbol{d} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}\right)$ is the *likelihood function* and $p\left(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}\right)$ is the *prior distribution*. Thus, the expression above states that the posterior distribution is proportional to the product of the likelihood of the decisions conditional on the underlying parameters and the prior beliefs about those parameters. For the likelihood function, the assumptions made about the error terms imply that the decisions have Bernoulli likelihood

$$p\left(\boldsymbol{d} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\rho}\right) \;\propto \prod_{t=1}^{T} \prod_{j \in J_t} \prod_{i \in I_j} \Phi\left(\alpha_{j,t} + \beta_{j,t}\rho_{i,t}\right)^{\delta_{i,j,t}} \left(1 - \Phi\left(\alpha_{j,t} + \beta_{j,t}\rho_{i,t}\right)\right)^{1-\delta_{i,j,t}},$$

where $\delta_{i,j,t} = 1$ if $D = A$, $\delta_{i,j,t} = 0$ if $D = R$, and Φ is the standard normal cumulative distribution function. Here J_t is the set of all observables in time period t and I_j is the set of all firms observed for observable j.

For identification purposes, the assumed prior distributions must be semi-informative. The literature has adopted a standard set of semi-informative prior for such purposes, which we adopt here. First, we are concerned with the change in D-SOCIAL-KLD scores over time and, as such, must model a mapping from one time period's D-SOCIAL-KLD score into the next. Firms' initial D-SOCIAL-KLD score in (unobserved) time period 0, $\rho_{i,0}$, is distributed

$$\rho_{i,0} \sim \mathcal{N}\left(m_{i,0}, C_{i,0}\right),$$

where we specify $m_{i,0}$ and $C_{i,0}$ numerically below. With the anchor in place, a *random walk* process dictates how one period's D-SOCIAL-KLD score maps into the next period's:

$$\rho_{i,t} \sim \mathcal{N}\left(\rho_{i,t-1}, \Delta_{\rho_{i,t}}\right).$$

$\Delta_{\rho_{i,t}}$, the variance of the normal distribution from which a D-SOCIAL-KLD score is drawn, dictates how closely information from the previous period relates to information in the current period. If it is very small, then the time series of D-SOCIAL-KLD scores for a firm over time approaches a constant value. If it is very large, then the time series is essentially unrelated to itself across time. Martin and Quinn (2002) observe that this is a "happy median" between one extreme (not modeling changes over time at all) and the other (not allowing one time period's responsibility to be related to the next).

We assume that $\alpha_{j,t}$ and $\beta_{j,t}$ are drawn from a multivariate normal distribution,

$$\begin{bmatrix} \alpha_{j,t} \\ \beta_{j,t} \end{bmatrix} \sim \mathcal{N}_2\left(\boldsymbol{b}_0, \boldsymbol{B}_0\right),$$

for all observables across all time periods. Importantly, we do *not* model dynamic effects in policy-specific attributes. Instead, we treat each observable "as a new case" in each year. To be sure, the analyst could utilize a random walk model like we utilize for the D-SOCIAL-KLD score. But, the marginal costs and benefits of adopting a given policy, which presumably fluctuate over time, are just the sort of thing we aim to learn more about. Put differently, we know that firms have sticky policies over time, but we *don't* know what CSR is. That is the very aim of the enterprise, and so allowing fluctuations seems the best approach for the learning process. Put differently, and in the terms stated above, we assume that the variance of a random-walk process underlying changes in α and β over time is *infinite*. This is an assumption that we make to maximize similarity to the extant IRT literature, but it can be relaxed—and time dependency in policy-specific parameters can be modeled—in future work.

The joint posterior distribution is simulated via Markov chain Monte Carlo (MCMC) methods. The algorithm, developed by Albert and Chib (1993), proceeds in three parts:

1. First, for all time periods, all firms, and all observables, we simulate $z_{i,j,t}$. Given the assumptions

of the model, the distribution of $z_{i,j,t}$ is given by

$$p\left(z_{i,j,t} \mid d_{i,j,t}, \rho_{i,t}, \alpha_{j,t}, \beta_{j,t}\right) = \begin{cases} \mathcal{N}_{[0,\infty]}\left(\alpha_{j,t} + \beta_{j,t}\rho_{i,t}, 1\right), & d_{i,j,t} = 1 \\ \mathcal{N}_{[-\infty,0]}\left(\alpha_{j,t} + \beta_{j,t}\rho_{i,t}, 1\right), & d_{i,j,t} = 0 \\ \mathcal{N}\left(\alpha_{j,t} + \beta_{j,t}\rho_{i,t}, 1\right), & d_{i,j,t} = NA. \end{cases}$$

That is, if $d_{i,j,t}$ is 1, the latent utility is distributed truncated normal at zero on the left; alternatively, if $d_{i,j,t}$ is 0, the latent utility is distributed truncated normal at zero on the right. For missing observations (which constitute a large proportion of our data), we make no truncation assumptions. Note that this step discriminates between observed and unobserved data: observed data are drawn from the appropriate truncated normal distributions, whereas unobserved data are drawn from the full normal distribution. Collect all of the latent utilities into a vector $\boldsymbol{z}$.

2. For all time periods and all firms, we simulate the $\alpha_{j,t}$ and $\beta_{j,t}$ terms. Let $\boldsymbol{\rho}^*_{i,t} = \begin{bmatrix} 1 \\ \rho_{i,t} \end{bmatrix}$, and let $\boldsymbol{\rho}^*_t$ be the $|I_j| \times 2$ matrix formed by stacking all these vectors for all firms. Then the observable-level parameters are distributed bivariate normal:

$$p\left(\begin{bmatrix} \alpha_{j,t} \\ \beta_{j,t} \end{bmatrix} \mid \boldsymbol{d}, \boldsymbol{z}, \boldsymbol{\rho}\right) = \mathcal{N}_2\left(\mathbb{E}\left[\left(\boldsymbol{\rho}^*_t\right)' \boldsymbol{z}_{i,t} + \boldsymbol{B}_0^{-1}\boldsymbol{b}_0\right], \left(\boldsymbol{\rho}^*_t\right)' \boldsymbol{\rho}^*_t + \boldsymbol{B}_0^{-1}\right).$$

3. For all firms in all time periods, simulate the D-SOCIAL-KLD score. This is done via a forward-filtering, backward sampling algorithm described at length by Martin and Quinn (2002). They place the algorithm in the context of a general multivariate dynamic linear model—that is, one where each firm's ideal points constitute a time series—that allows for more complicated specifications of the general process of interest. Begin by rewriting the latent utility equation into a form we will call the *observation equation*:

$$z_{i,\cdot,t} - \boldsymbol{\alpha}_t = \boldsymbol{\beta}_t \boldsymbol{\rho}_{i,t} + \boldsymbol{\varepsilon}_{i,\cdot,t},$$

where $\cdot$ makes it explicit that we are considering a time period across all observables. For the dynamic ideal point estimation, we formulate the *evolution equation*

$$\rho_{i,t} = \rho_{i,t-1} + \delta_t,$$

where $\delta_t \sim \mathcal{N}\left(0, \Delta_{\rho_{i,t}}\right)$. Thus the evolution equation is a simple re-expression of the random walk process described above. With these in place, we can simulate the posterior distributions of the D-SOCIAL-KLD scores. This occurs by first sampling $\boldsymbol{\rho}_T$, where T is the total number of time periods, from $p\left(\boldsymbol{\rho}_T \mid D_T\right)$, where D_T is shorthand for *all* information available up to time T, and then by exploiting the fact that

$$p\left(\boldsymbol{\rho} \mid D_T\right) = p\left(\boldsymbol{\rho_T} \mid D_T\right) p\left(\boldsymbol{\rho}_{T-1} \mid \boldsymbol{\rho}_T, D_{T-1}\right) \cdots p\left(\boldsymbol{\rho}_0 \mid \boldsymbol{\rho}_1, D_0\right),$$

as specified by the evolution equation. The "backward sampling" technique thus begins from the final time point and works backward to the first time point. The algorithm proceeds by computing:[1]

(a) The prior mean of $\boldsymbol{\rho}_t$ given D_{t-1},

$$\boldsymbol{a}_t = \mathbf{1}_t \boldsymbol{m}_{t-1},$$

where in the first iteration $\boldsymbol{m}_{t-1}$ comes directly from the prior on $\boldsymbol{\rho}_0$ and in subsequent iterations it has already been computed;

(b) The prior variance of $\boldsymbol{\rho}_t$ given D_{t-1},

$$\boldsymbol{R}_t = \mathbf{1}_t C_{t-1} \mathbf{1}_t' + \Delta_t,$$

where C again comes from the prior in the first iteration;

[1]For simplicity but explicitness, let $\mathbf{1}_t$ be a vector of ones with as many elements as there are observables in time period t. This is the vector to which any further covariates could be added in future applications of the algorithm.

(c) The mean of the forecast:

$$\boldsymbol{f}_t = \boldsymbol{\beta}_t \boldsymbol{a}_t,$$

where we have already calculated $\boldsymbol{a}_t$ and the form comes from the evolution equation;

(d) The variance of the forecast:

$$\boldsymbol{Q}_t = \boldsymbol{\beta}_t \boldsymbol{R}_t \boldsymbol{\beta}_t' + \boldsymbol{I}_{J_t},$$

where $|J_t|$ is the number of policies in time period t;

(e) The posterior mean of $\boldsymbol{\rho}_t$ given D,

$$\boldsymbol{m}_t = \boldsymbol{a}_t + \boldsymbol{R}_t \boldsymbol{\beta}_t \boldsymbol{Q}_t^{-1} \left(\boldsymbol{z}_{t,\cdot,x} - \boldsymbol{\alpha}_t - \boldsymbol{f}_t \right),$$

where we maintain notation $(\boldsymbol{z}_{i,\cdot,t} - \boldsymbol{\alpha}_t - \boldsymbol{f}_t)$ to highlight that it is the error of the forecast; and

(f) The posterior variance of $\boldsymbol{\rho}_t$ given D_t,

$$\boldsymbol{C}_t = \boldsymbol{R}_t - \boldsymbol{R}_t \boldsymbol{\beta}_t \boldsymbol{Q}_t^{-1} \boldsymbol{Q}_t \left(\boldsymbol{R}_t \boldsymbol{\beta}_t \boldsymbol{Q}_t^{-1} \right)'.$$

This process is repeated up to T, thus providing the "forward filtering" component to the algorithm. The quantities are saved and then are utilized for the backward sampling:

(a) Sample $\boldsymbol{\rho}_T$ from a multivariate normal distribution with mean $\boldsymbol{m}_T$ and variance $\boldsymbol{C}_T$ as computed above.

(b) Calculate the mean of the conditional distribution of $\boldsymbol{\rho}_t$ given ρ_{t+1} and D_t,

$$\boldsymbol{h}_t = \boldsymbol{m}_t + \boldsymbol{C}_t \boldsymbol{1}_{t+1}' \boldsymbol{R}_{t+1}^{-1} \left(\boldsymbol{\rho}_{t+1} + \boldsymbol{a}_{t+1} \right),$$

where all quantities have again been computed in the forward filtering step.

(c) Calculate the variance of the conditional distribution of $\boldsymbol{\rho}_t$ given ρ_{t+1} and D_t,

$$\boldsymbol{H}_t = \boldsymbol{C}_t - \boldsymbol{C}_t \mathbf{1}'_{t+1} \boldsymbol{R}_{t+1}^{-1} \boldsymbol{R}_{t+1} \left(\boldsymbol{C}_t \mathbf{1}'_{t+1} \boldsymbol{R}_{t+1}^{-1} \right)',$$

which again are all known.

With these quantities in place, we draw from $p\left(\boldsymbol{\rho}_t \mid \boldsymbol{\rho}_{t+1}, D_t\right) \sim \mathcal{N}\left(\boldsymbol{h}_t, \boldsymbol{H}_t\right)$.

The entire process—forward filter to compute the quantities for the conditional distributions, then backward sample from them—is repeated many times. The algorithm is automated in `MCMCpack` for users in the R statistical computing environment. Importantly, given our massive data and the computational expense of the model, the automated algorithm makes use of parallel processing in C++ to maximize efficiency.

Given that we are working with such a large dataset, and given that so many firms have adopted so few KLD items, the routine as written in `MCMCpack` only works if appropriate care is given to setting starting values. Throughout our estimation, we use the KLD Index value for each firm divided by the number of metrics in year t as the starting values for $\boldsymbol{\rho}$; taking a similar approach, we use the mean number of firms that adopt a given policy as the starting values for $\boldsymbol{\alpha}$. We set the starting value of all the $\boldsymbol{\beta}$ terms at one. The model requires two firms to be "identifiers" of the dimension. We chose firms that consistently ranked high and low on routines estimated with different identifiers: IBM on the positive side and Bed, Bath, and Beyond on the negative side. Much like estimates from logit models where the variance is assumed, the parameters are estimated given assumed default settings: prior means $\boldsymbol{a}_0 = \boldsymbol{b}_0 = m_0 = 0$ and prior variance diagonal elements of $\boldsymbol{A}_0$, $\boldsymbol{B}_0$, and $\boldsymbol{C}_0$ of 0.1. We run the algorithm for 5,000 iterations after a burn-in period of 1,000 iterations; because draws may be autocorrelated, we save only every other iteration, leaving 2,500 draws from each posterior distribution of interest.

References

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

Martin AD, Quinn KM. 2002. Dynamic ideal point estimation via markov chain monte carlo for the U.S. Supreme Court. *Political Analysis*, **10**(2):134–153.