Family Planning Policy in the United States: The Converging Politics of Abortion and Contraception

Supplementary Materials

1 Data

Overview. Our data consist of all roll-call record votes in the 2011 and 2003 Texas legislative sessions. We also collected information on legislator and constituency-specific factors that have been suggested by previous studies to influence voting behavior and which are candidate factors for Texas. These include the legislator's party affiliation, religion, gender, and the district-level percentage of constituents who: live in a rural area, are white, are foreign-born, live in a single-parent family, live in poverty, or hold a bachelor's degree or higher. These data were obtained from Telicon, a private legislative research company based in Austin, Texas. The latter four variables are estimated by Telicon based on data from the 2007-2011 American Community Survey (ACS), and a detailed description of each can be found in the ACS data documentation codebooks. Percent white is based on U.S. Census data aggregated at the district level, and refers to the percentage of constituents who are white non-Hispanic. Percent rural is also based on U.S. Census data, and refers to the percentage of constituents living in a geographical area that is not classified as an urbanized area of 50,000 or more people, or an urban cluster of between 2,500 and 50,000 people. We do not have detailed demographic data for each Texas house district in 2003 and 2011 individually. Therefore, we use the 2007-2011 ACS data provided by Telicon as a proxy for district-level characteristics both 2003 and 2011.

We assembled the data for each year into an $N \times D$ member-vote matrix Y whose rows correspond to individual legislators and whose columns correspond to votes. (Thus N refers to the number of votes and D to the number of legislators in a given year.) Each entry y_{ij} of this matrix is an indicator for the vote of legislator i on bill j: $y_{ij} = 1$ for yes, and $y_{ij} = 0$ for any other vote (including a no or an absention). There are two such matrices, one for 2003 and one for 2011.

Tagging of family-planning bills. All potential bills concerning family planning in each session were identified using an online search of the Texas legislative record[1] using the search terms: “family planning”; “contraception”; “abortion”; “women’s health”; and “Title X”. Each author separately and independently examined the description of each bill to verify that it concerned family planning. Inter-rater agreement was high: there were three cases of disagreement out of 101 potential bills flagged by the keyword search. These bills were not analyzed as part of the family-planning subset of bills, but were retained for analysis as non-family-planning bills. This resulted in 63 family-planning bills in the 2011 session and 35 family-planning bills in the 2003 session.

2 Further details of the model

2.1 Model structure

In this section, we provide further details of the Bayesian ideal-point models that we fit to voting data from 2003 and 2011 Texas Legislatures. We fit separate models for each year. In what follows, we describe the model for a single year.

Our ideal-point model is of the same basic form recommended by [2]. It assumes that vote y_{ij} is a Bernoulli random variable such that

$$\Pr(y_{ij} = 1 \mid \alpha, \beta, f) = \Phi(\alpha_j + \beta_{j1} f_{i1} + \beta_{j2} f_{i2}) \, . \tag{1}$$

Here Φ is the cumulative distribution function of the standard normal distribution (i.e. the link function in a standard probit regression), β_{jk} are the factor loadings for bill j, and f_{ik} are the factor scores (ideal points) for legislator i. Conditional on the factor scores, the votes on bill j thus follow a generalized linear model with a probit link and regression coefficients $(\alpha_j, \beta_{j1}, \beta_{j2})$. The factor scores (f_{i1}, f_{i2}) locate legislator i in the political space, while the factor loadings (β_{j1}, β_{j2}) locate bill j in the political space.

This model can be motivated in at least two different ways. First, there is the behavioral interpretation. As [2] observes, Equation (1) corresponds to a stylized behavior model in which an individual legislator’s utility for a bill decreases quadratically with the Eucidean distance of that bill from his or her ideal point in the political space.

Second, there is the direct statistical interpretation of Equation (1) as a factor model. A factor model presumes that all of the correlation in a high-dimensional outcome vector—in this case, the set of votes cast by a single legislator—can be described in terms of shared dependence on a set of latent factors. In Equation (1), this dependence takes the form of a probit regression for the binary outcome y_{ij}. However, both the shared factors and the regression coefficients must be inferred from the data. This requires additional restrictions to ensure that the model is formally identifiable, and distinguishes factor analysis from ordinary regression.

[1] `http://www.capitol.state.tx.us`

2.2 Prior distributions and identification constraints

We adopt a Bayesian approach to fitting the model defined by Equation (1). This entails: (1) placing prior distributions on the model parameters $(\alpha_j, \beta_{j1}, \beta_{j2})$ and (f_{i1}, f_{i2}); (2) applying Bayes' rule to combine the record-vote data with these prior distributions to form the posterior distribution; and (3) using Monte Carlo simulation to take draws from the posterior distribution.

There are several reasons for taking a Bayesian approach to ideal-point modeling, with [2] highlighting two in particular. First, standard non-Bayesian ideal-point models do not allow for simple quantification of uncertainty about model parameters. A Bayesian approach, on the other hand, leads to a full posterior distribution over all parameters, and thus a complete description of our uncertainty in light of the data. Second, the use of Monte Carlo simulation elegantly avoids the difficult computional problem of fitting Equation (1) by maximum likelihood.

A third reason for taking a Bayesian approach is that the prior distributions for model parameters can be chosen to ensure that the axes of the political space parameters are easy to interpret. Our model follows in this tradition by employing highly structured prior distributions that encode the interpretation that f_{i1} is a partisanship factor and f_{i2} is a family-planning factor. Specifically, for the first (partisanship) factor scores f_{i1} and loadings β_{j1}, we use the prior distributions

$$\beta_{j1} \sim \text{Normal}(0,1) \tag{2}$$

$$f_{i1} \sim \begin{cases} \text{Normal}(1,1) & \text{if legislator } i \text{ is a Republican} \\ \text{Normal}(-1,1) & \text{if legislator } i \text{ is a Democrat}\,. \end{cases} \tag{3}$$

This is sufficient to fix the sign and interpretation of the first factor as encoding partisan voting tendencies, as any changes to the underlying space that failed to respect the broad tendency for Republicans and Democrats to have opposite signs would incur a substantial penalty from the prior. Yet it is still flexible enough to allow for factor scores consistent with a liberal Republican or a conservative Democrat, if the data warrant it.

For the second factor, we use the following priors:

$$\beta_{j2} \sim \begin{cases} \text{Normal}(0,1) & \text{if bill } j \text{ relates to family planning} \\ \delta_0 & \text{otherwise,} \end{cases} \tag{4}$$

$$f_{i2} \sim \text{Normal}(0,1)\,. \tag{5}$$

Here δ_0 represents a Dirac distribution at zero, so that $\beta_{j2} = 0$ if bill j is unrelated to family planning. This has the effect of allowing f_{i2} to influence only those small fraction of bills specifically hand-tagged as being relevant to family planning (35 in 2003, 62 in 2011, out of over 1000 total bills in each session).

For the intercepts or item-difficult parameters α_j, we use priors that are very noninformative on the probit scale: $\alpha_j \sim \text{Normal}(0, \sigma^2 = 25)$.

Further constraints on the model. We impose three additional constraints on the model, which are exceptions to the prior distributions just described. Specifically, we designate certain bills as fixed points in the political space and constrain the prior distribution so that the factors loadings of these bills are assumed known *a priori*.

Within each session, we chose one bill from each of the following three categories and fixed the factor loadings of this bill in the manner indicated.

Republican votes: a Republican-favored roll call unrelated to family planning. This bill's first loading is fixed at $\beta_{j1} = 2$. Because this is a non-family-planning bill, the constraint that $\beta_{j2} = 0$ is also assumed to hold, as described in the main paper.

Democratic votes: a Democrat-favored roll call unrelated to family planning. This bill's first loading is fixed at $\beta_{j1} = -2$. Similarly, because this is a non-family-planning bill, the constraint that $\beta_{j2} = 0$ is also assumed to hold, as described in the main paper.

Family-planning access votes: a vote restricting access to abortion services. This bill's second loading is fixed at $\beta_{j2} = -2$, with the first factor loading given a $N(0, 1)$ prior.

The intuition behind this step is to ensure that there are a small handful of bills that serve as "pole stars"—fixed locations in political space to which other bills and legislators can be compared. This has the effect of orienting the sign of both the partisanship and family-planning factors and to aid in making comparisons of factor scores and loadings across years. In particular, without the final restriction, we could flip the signs of all scores and loadings for the family-planning factor without changing the model. (For a technical discussion of model identifiability in Bayesian factor models, see [6] and [7].)

However, there are several choices of votes in each category, meaning that the choice of a single one as a pole star—even if it is an archetypal member of its category—is somewhat arbitrary. Therefore, to check that our results are robust with respect to this choice, we constructed three different batches of votes in each of 2003 and 2011 and re-ran our models under each batch. (By a "batch," we mean a set of three votes in a given year, one from each category.)

We first describe the batches, and then the way these batches were used for robustness checks. The votes chosen for each batch are described in general terms below, and listed explicitly in Tables 1 and 2.

For the Republican and Democratic categories, we picked votes on major bills or amendments across a range of controversial issues that arose in each session: Congressional redistricting, tort reform, fiscal policy, job benefits for state employees, state funding for higher education, immigration enforcement, and voter ID. These issues tend to be heavily partisan, to the extent that they are intimately tied to the electoral identities of the two parties in Texas. They were thus natural candidates for votes that could serve to orient the partisanship axis. A Republican (Democratic) vote means that a typical Republican (Democrat) in Texas could be expected to vote "yes" on the roll call.

For the family-planning category, we fixed the loadings of votes on major abortion legislation. This was an obvious choice because there was such bill in each session, with similar political debates surrounding each one. One of our paper's aims is to understand the changing politics of family-planning funding with respect to abortion politics. Fixing votes on abortion bills was the most direct way to allow such an assessment. Indeed, the fact that 2003 also had a major piece of abortion legislation was what led us to use it as a comparison for 2011 in the first place.

In 2003, we selected motions to table three proposed amendments to House Bill 15 ("The Women's Right to Know Act"). This piece of legislation imposed several new requirements on facilities that provide abortion services. The Texas Department of State Health Services summarizes the bill as follows [8]:

> During the 2003 session, the Texas Legislature passed the Woman's Right to Know Act (House Bill 15). Under this law:
>
> - A doctor who is to perform an abortion (or the doctor's agent) must tell the woman that benefits may be available to help with medical care before, during, and after childbirth.
> - The father is required to help support the child whether or not he has offered to pay for an abortion.
> - Government and private agencies can counsel the woman in preventing pregnancy, or refer her to a doctor for medications or devices to prevent pregnancy, including emergency contraception for victims of rape and incest.
> - The woman has the right to look at printed information. If she chooses to see the material the law describes, the doctor (or the doctor's agent) shall give her a copy at least 24 hours before the abortion is scheduled. The doctor (or agent) may instead mail her the materials, with delivery restricted to her, at least 72 hours before the abortion is scheduled.

Most of the amendments attempted to weaken various provisions of the bill, thereby removing or mitigating impediments to abortion-care access. We used the three amendments listed under the "Family planning" category in Table 1. Both in 2003 and 2011, we used votes on major amendments to the abortion bill, rather than the vote for final passage of the bill. We did so for the simple reason that we needed three such votes, and we wanted them all to be of a similar procedural character. In all cases, although the amendment itself was designed to promote access to abortion care, a motion to table (or kill) the amendment was actually voted upon. Therefore a "yes" vote on these roll calls can be interpreted as a vote in favor of further restrictions on access to abortion care.

In 2011, we used votes on three amendments to the controversional sonogram bill passed by the Texas House in 2011, coincidentally also numbered House Bill 15. The Department of State Health Services describes the content of the law that took effect upon passage of this bill [8]:

> Texas law says your doctor must talk to you about certain things before you can have an abortion. After you get this information, your doctor must wait 24 hours before your abortion can be performed. The law also requires that you receive a sonogram from the doctor (or agent) who will be performing your abortion at least

24 hours before the abortion is to occur. If you live more than 100 miles away from the nearest abortion provider, you can waive this requirement, but you will still be required to have a sonogram by the doctor performing the abortion. It will need to occur at least 2 hours before the abortion, rather than 24 hours before.

Once you arrive at your doctors office, and before your sonogram, you will be asked to sign a sonogram/abortion election form. This form certifies that you are aware of the law and its requirements. The doctor will also give you materials to read. One of those materials should be the Resource Directory listed on this webpage.

During your sonogram, the doctor is required to display the sonogram images and make the heart beat audible. You may decline to view the images and listen to the heartbeat. The doctor must also provide a verbal explanation of the sonogram results. Only women who certify on the sonogram/abortion election form one of the following criteria can decline to hear the verbal explanation: the pregnancy is a result of a sexual assault, incest, or other violation of the Penal Code that has been reported to law enforcement authorities or that has not been reported because she reasonably believes that to do so would put her at risk of retaliation resulting in serious bodily injury; the woman is a minor and obtaining an abortion in accordance with judicial bypass procedures under Chapter 33, Family Code; or the fetus has an irreversible medical condition or abnormality, as medically documented in the woman's medical file.

As with the Women's Right to Know Act in 2003, most of the proposed amendments to the sonogram bill attempted to weaken its provisions, thereby removing impediments to abortion-care access. We used the three motions to table amendments that are listed under the "Family planning" category in Table 2.

Robustness checks. Of the three batches listed for each year, Batch 1 was used to fit the models whose results are summarized in the main paper. Batches 2 and 3 were used to check the robustness of our results with respect to the particular batch of bills chosen as fixed. To assess this, we fit the model three times, each time fixing one of the three batches of factor loadings in the manner explained above. We then compared the legislators' estimated factor scores under the three different batches.

Figure 1 shows the results of this comparison. Each panel shows a scatter plot of the estimated factor scores under one batch, versus those in another batch. For example, the bottom left panel shows the 2011 legislators' estimated family-planning factor scores with batch 2 held fixed, versus those with batch 1 held fixed. The remaining panels show similar plots for both years, both factors, and all three pairwise combinations of batches.

The minimum pairwise correlation across all panels was 99.3%. The typical deviation of a factor score between batches is much smaller than that score's posterior standard deviation within a batch, and only slightly larger than the Monte Carlo variability of the posterior means for a fixed batch. To give further context to these figures, the batch-to-batch variability of the posterior means is smaller in relative magnitude than the bootstrapped standard errors of the ideal points ($\approx$ 1-4%) reported by [9], whose use data on the entire voting history of members of the United States Congress.

We conclude that our results are robust with respect to the particular set of bills whose factor loadings are held fixed in a given year, and that the differences due to reasonable changes in modeling assumptions are very small compared to other sources of uncertainty.

2.3 Comparison of scales across years.

As described in the main manuscript, our factor model recovers the relative locations of legislators and votes in a latent "political space" that holds for a given year. Comparing these locations across years is a much more difficult problem, and has been addressed by political scientists in a variety of ways [10]. A common approach is fix the factor scores of legislators who serve in both time periods. But for our purposes, this approach does not seem appropriate. We expect that the view of Texans and their representatives on abortion and family planning may have evolved, or even changed in major ways, over time. One legislator (Aaron Peña) switched parties between 2003 and 2011; we do not wish to pre-suppose whether the parties remained fixed while this legislator's views changed, or vice versa. For the same reason, we do not wish to impose a dynamic model requiring the assumption that the political space itself is time-invariant, as in [4].

As detailed above, we have instead fixed certain factor loadings in each year in order to facilitate comparisons of factor scores and loadings across years. For example, in batch 1, we fixed the Congressional redistricting bill in 2003 and the voter ID bill in 2011 both to have a loading of $+2$ on the partisanship factor. These were among the most controversial (non-family-planning) issues in each session. Effectively, this encodes an assumption that these bills are equally partisan in a Republican direction. Similary, we fixed major votes on the Women's Right to Know Act in 2003 and the sonogram bill in 2011 to have loadings of -2 on the family-planning factor. This encodes an assumption that these bills are equally "anti-access" to family-planning services—if not in substantive terms, then at least in terms of what they reveal about legislators' beliefs.

These assumptions allow us to reason in a more informed way about legislator scores and bill loadings across the two sessions. The technical rationale for these assumptions is that they allow meaningful comparisons of magnitudes, directions, and angles between bills in factor space. Although the scale of the factors is arbitrary and set only by assumption, any nontrivial rescaling that would change magnitudes and angles would also (contrary to assumption) change the loadings of at least one of the three bills we have fixed.

But these restrictions do not imply that a 2 on the 2003 scale means exactly the same thing as a 2 on the 2011 scale. For this equivalence to hold, not only would the bills have to be identical, but the status quos modified by the bills would also have to be identical. This plainly cannot be: for one thing, the 2003 abortion bill greatly modified the status quo upon which the 2011

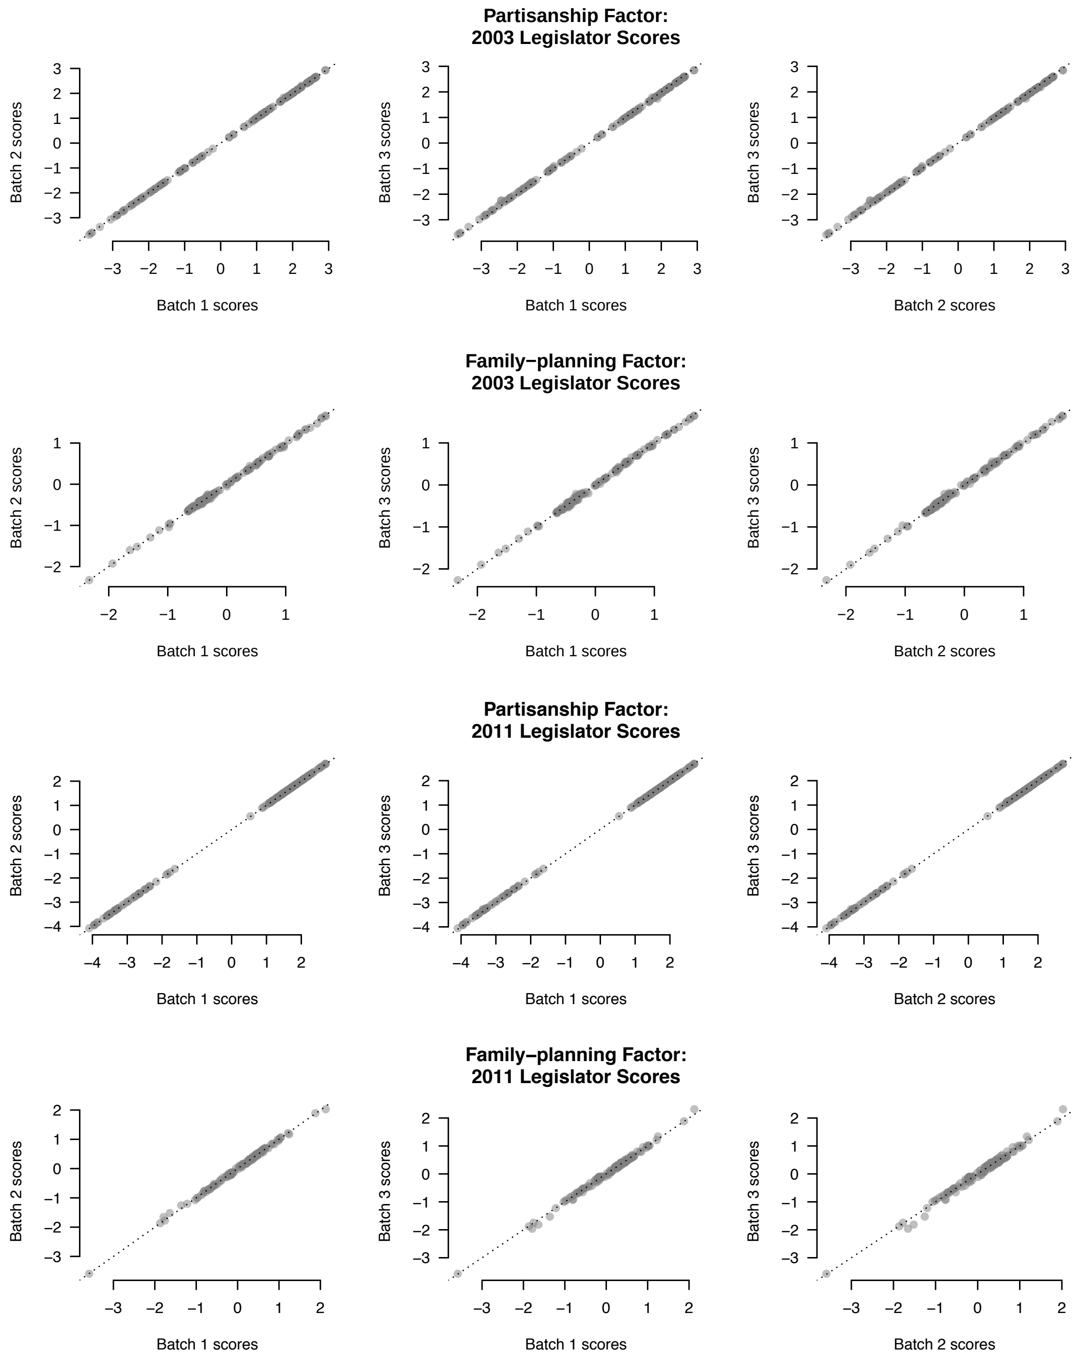

Figure 1: Pairwise scatter plots of legislator's estimated factor scores under three differents sets of constraints on the factor loadings (Batch 1, 2 and 3). The strong correlation in all panels shows that our results are very robust with respect to different possible versions of the constrained prior distributions described in Sections 2.2.

abortion bill acted. Moreover, one of our paper's main arguments is that legislators voted on family-planning legislation in a systematically different way in 2011 than in 2003. Specifically, we find that votes on family-planning funding bills looked like votes on abortion in 2011, but not in 2003. We conclude that support for family planning funding means something different in each year.

Therefore, we believe that it is neither possible nor advisable to attempt to recover a fixed "family-planning" scale whose meaning is substantively identical across years. Recovering such a scale is not an aim of our analysis, nor is it necessary for any of our arguments to hold. For example, we have asked: where do the votes on state funding for family planning in 2003 (2011) sit relative to the votes on abortion in 2003 (2011)? How much of the variation in voting on family-planning bills in 2003 (2011) can be explained by the partisanship factor in 2003 (2011)? Having fixed locations for archetypal Republican/Democratic/family-planning votes in each session is sufficient to answer these questions. We need not place our faith in the notion that the 2003 and 2011 family-planning scales are identical in all respects.

To be sure, many other hypothetical comparisons would require greater attention to scale equivalence. One such example would be an attempt to ascertain from the voting record alone whether a particular legislator has become more or less "pro-family-planning" over time, relative to the earlier version of himself or herself. But our analysis does not address any such questions.

3 Model fitting

3.1 Overview

In our Bayesian factor probit model, all quantities are estimated using the posterior distribution—specifically, the posterior mean provides a point estimate, and the posterior standard deviation provides a measure of uncertainty for each parameter.

The chief computational difficulty is that the posterior distribution itself is not available in closed form. In other words, one cannot compute the posterior distribution by directly applying Bayes' rule to yield an analytical expression. Therefore, we use a technique called Markov Chain Monte Carlo (MCMC) to simulate random draws from the posterior distribution. This allows us to visualize the posterior distribution using standard plots, such as histograms (for single parameters) or scatter plots (for pairs of parameters). This is a very common approach in Bayesian modeling; for an overview of MCMC, we refer the reader to [11].

Our MCMC sampler takes 300,000 draws from the posterior distribution. Of these 300,000, we discard the first 50,000 as a "burn-in" period. This is important because an MCMC sampler must begin from an initial guess of the parameter values. If this guess is poor, the sampler will require a burn-in period to reach equilibrium. During this period, the samples drawn are liable to be misleading about the true posterior distribution, and are typically discarded.

Of the remaining 250,000 draws, we retain every 25th draw, discarding the other 24. This "thinning" procedure produces an overall sample of 10,000 draws from the joint posterior distribution. The rationale for discarding 24 out of every 25 draws is that, by its nature, an MCMC sampler produces autocorrelated (rather than independent) draws from the posterior. This autocorrelation decays over time: the draws at steps t and $t+1$ are often highly correlated, but the draws at steps t and $t+25$ are much less correlated. Because the goal of MCMC sampling is to draw samples from the posterior that are as close to uncorrelated as possible, it is standard practice is to take many more draws than one needs and to produce a smaller, less autocorrelated sample by thinning them.[2]

We use standard diagnostics to ensure that our MCMC has reached equilibrium and has acceptably small autocorrelation, as described in [11]. Specifically, we examined trace plots and computed autocorrelation coefficients for those trace plots. While it is impossible to show a traceplot for every single parameter, we give a sampling in Figures 1 and 2. Figure 1 shows trace plots and autocorrelation plots for the partisanship factor loadings β_{j1} for four selected bills, across a range of values (high, middling, and low). Figure 2 shows the same plots for the family-planning factor scores f_{i2} for four selected legislators across a range of values (high, midding, and low). These are representative of all the trace plots we examined (with the largest autocorrelations observed for the parameters with largest absolute values). They provide evidence of acceptable convergence and Monte-Carlo sample sizes.

3.2 Technical details

The standard Bayesian approach to fitting probit models is to exploit the data-augmentation strategy of [1]. This entails introducing latent variables z_{ij} such that

$$\begin{aligned} z_{ij} &= \alpha_j + \beta_{j1} f_{i1} + \beta_{j2} f_{i2} + \epsilon_{ij} \\ \epsilon_{ij} &\sim \text{Normal}(0, 1)\,. \end{aligned}$$

The z_{ij}'s are often referred to as latent utilities, in that they connect the factor-analytic interpretation of the model with the behavioral model of legislators are utility maximizers. We then identify $y_{ij} = 1$ with the event $\{z_{ij} \geq 0\}$, and $y_{ij} = 0$ with the event $\{z_{ij} < 0\}$. Because y_{ij} follows a probit regression, we have $\Pr(z_{ij} \geq 0) = \Pr(y_{ij} = 1)$, meaning that the augmented model specifies the same likelihood function as the original model once the z_{ij}'s are integrated out.

Let D be the total number of votes, and N the total number of legislators. We may write the augmented model in matrix notation as

$$Z^T = A + BF^T + E^T\,,$$

where A is a D-vector whose jth entry is α_j, Z^T is the transpose of the $N \times D$ matrix of latent utilities having entries z_{ij}; B is the $D \times 2$ factor-loadings matrix with entries β_{jk}; F^T is the

[2]Technically the thinning step is unnecessary and one could do better by retaining all the draws. However, available computer memory often precludes this, especially for complex models. Thus thinning is a practical strategy for reducing the correlation in one's sampler while retaining a tractable memory footprint.

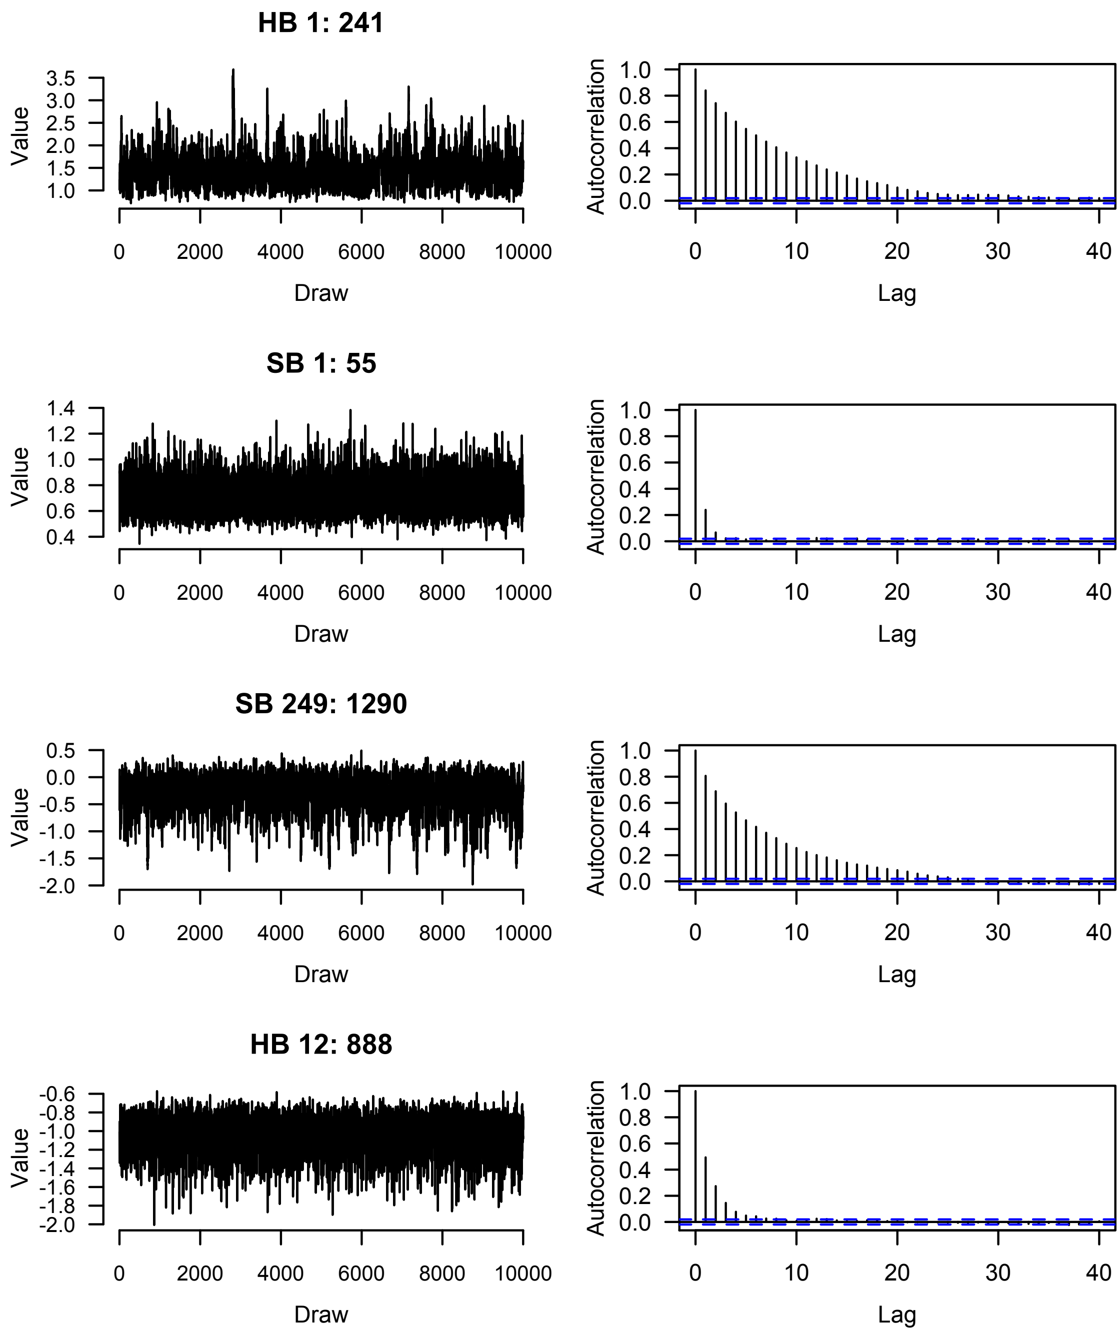

Figure 2: Trace plots and autocorrelation plots for the partisanship factor loadings β_{j1} for four selected bills. These figures show that our MCMC sampler has reached equilibrium and has collected draws from the posterior distribution with acceptably low autocorrelation. For details of interpreting an autocorrelation plot, we refer the reader to Section 5.4 of [11].

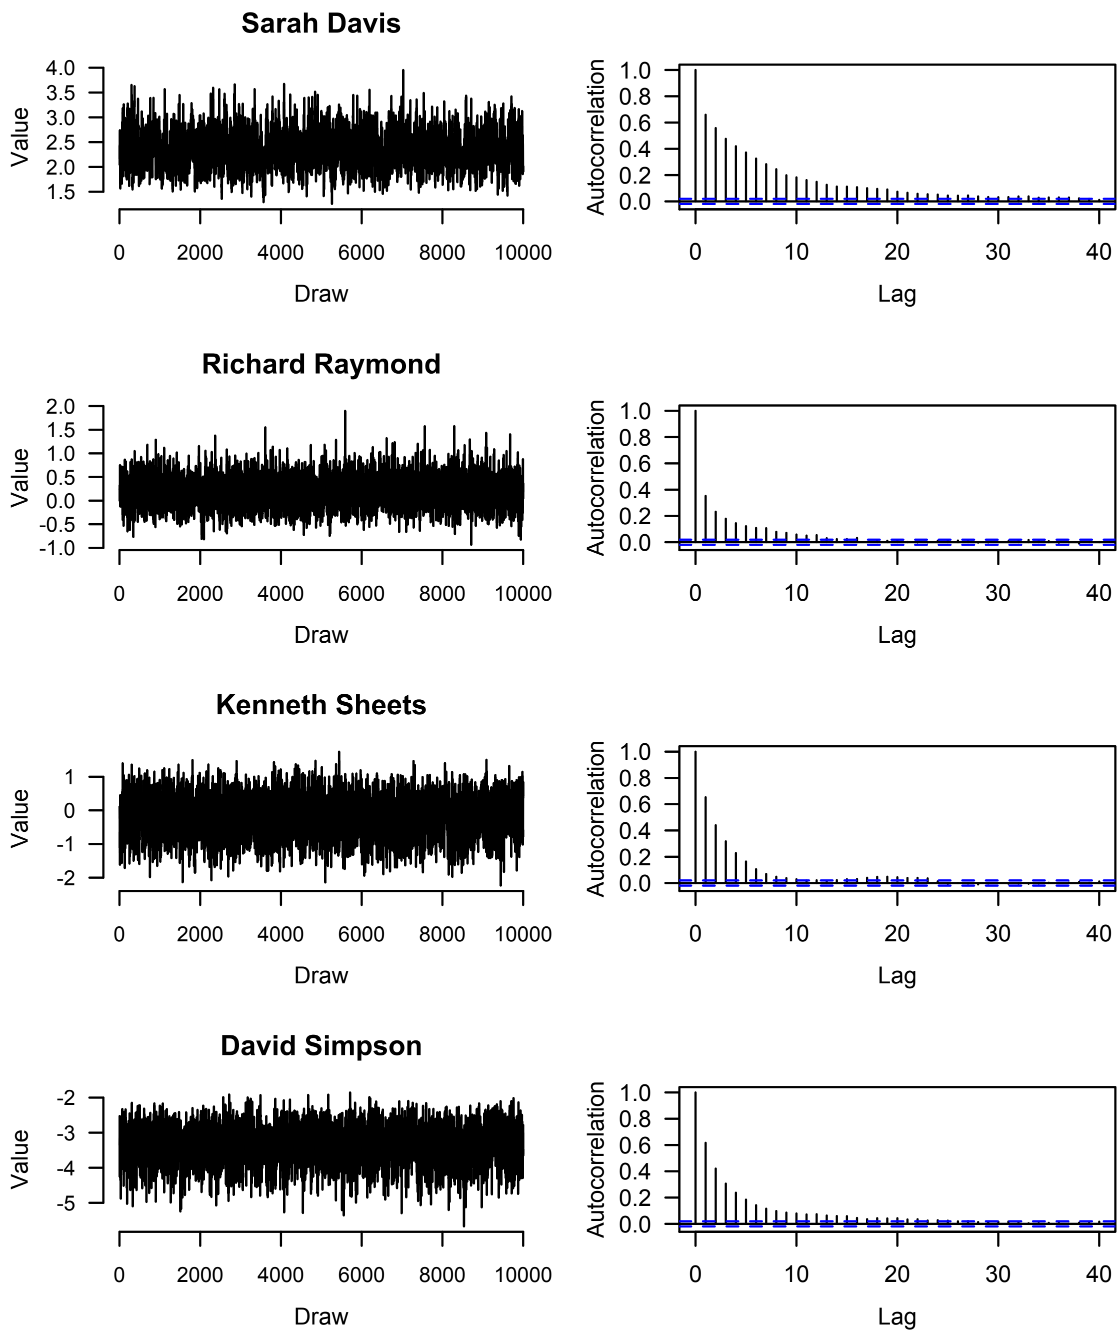

Figure 3: Trace plots and autocorrelation plots for the family-planning factor scores f_{i2} for four selected legislators.

transpose of the $N \times 2$ matrix comprising the factor scores f_{ik}; and E is the matrix of latent residuals ϵ_{ij}.

This latent-variable representation plays a key role in our calculation of Bayesian R^2, as discussed below. More importantly, it also allows the model to be fit straightforwardly using MCMC. Full details of the Gibbs sampling steps can be found in, for example, [2] and [3].

3.3 Calculation of Bayesian R^2 for each factor

By construction, the residuals ϵ_{ij} on the probit scale have variance 1, and the factor scores f_{i1} and f_{i2} are *a priori* independent with variance 1. This allows us to define a Bayesian version of R^2 for each factor on each bill, denoted ρ^2_{jk}, in terms of a simple Pythagorean decomposition:

$$\rho^2_{jk} = \frac{\text{var}(\beta_{jk} f_k)}{\text{var}(\beta_{j1} f_1) + \text{var}(\beta_{j2} f_2) + \text{var}(\epsilon)} = \frac{\beta^2_{jk}}{\beta^2_{j1} + \beta^2_{j2} + 1} \, .$$

Intuitively, if β_{j2} has a large magnitude relative to β_{j1} (and to the residual variance of 1 on the latent utility scale), then the family-planning factor accounts for much of the variation in the latent utilities z_{ij} for bill j, and therefore for much of the variation in voting outcomes.

We calculate each bill's ρ^2_{jk} for the partisanship factor ($k = 1$) and family-planning factor ($k = 2$) for each MCMC draw of the factor loadings, and average the draws. These ergodic averages are plotted in Figure 6 of the main manuscript.

3.4 Calculation of cut line for each bill

To calculate the cut lines show in Figure 2 of the main manuscript, and in Section 4 of this document, we form the following optimization problem. To lighten the notation, we drop the subscript j indexing votes. Let $\beta = (\beta_1, \beta_2)$ be the estimated vector factor loadings for a given vote (the "vote vector") and $f_i = (f_{i1}, f_{i2})$ the estimated factor loadings for legislator i. Let $s_i = f_i \cdot \beta$ be the inner product of f_i with the vote vector, and recall that $y_i = 1$ if legislator voted "yes", and 0 if "no." Now let

$$\hat{c} = \arg\min_{c \in \mathbb{R}} \left\{ \sum_{i: y_i = 1} \max\{0, (s_i - c)^2\} + \sum_{i: y_i = 0} \max\{0, (c - s_i)^2\} \right\} \, .$$

The solution to this one-dimensional optimization problem is easily found by a black-box numerical solver. Intuitively, $\hat{c}$ divides cuts the s_i into two groups $\{s_i \geq c\}$ and $\{s_i < c\}$. If the s_i values corresponding to yes votes are perfectly separable, then $\hat{c}$ will maximize the margin between the groups. If the s_i's are not separable, then $\hat{c}$ will minimize the squared error of those s_i's that are on the wrong side of the boundary. (Here error is to be interpreted as "distance from the boundary point.")

The cut line is then the line with intercept $\hat{c}/\beta_2$ and slope β_1/β_2, which is orthogonal to the vote vector, and passes through the point that is distance $\hat{c}$ from the origin along the vote vector.

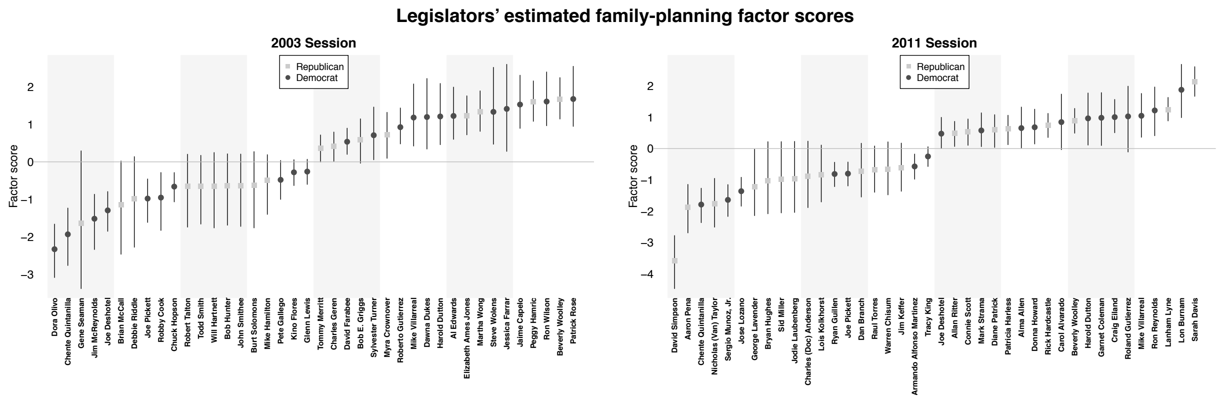

Figure 4: Posterior means and 95% credible intervals family-planning factor scores for the 2003 (left) and 2011 (right) sessions. Only the legislators with the 20 largest and 20 smallest posterior mean factor scores within each session are shown.

This is analogous to a max-margin separating hyperplane between the yes and no votes in factor space, but constrained to be perpendicular to the estimated vote vector (β_1, β_2).

4 Additional figures and tables

Figure 4 shows the estimates and error bars of the family-planning factor scores (f_{i2}) for 40 legislators (the 20 with the largest positive values, and the 20 with the largest negative values) in the 2003 and 2011 sessions. Recall that legislators with a family-planning factor score near zero tend to vote on family-planning issues in a way that is indistinguishable from their voting on other issues. A negative (positive) factor score implies that the legislator tends to vote for bills that restrict (promote) access to family planning services, adjusting for the legislator's own demonstrated partisan voting tendencies.

Figures 5 and 6 show further examples of cut lines for specific bills in 2003 and 2011, respectively. Within each figure, the top three panels show votes on abortion, while the bottom three panels show votes on state funding for contraception. These further examples demonstrate that the broad trend depicted in Figure 2 (particularly Panels A2–D2) of the main paper hold for other choices of abortion and contraception bills.

Finally, Tables 3 through 10 give the regression coefficients for all 8 models of family-planning factor scores versus constituency characteristics (2003 versus 2011, Republicans versus Democrats, and full model versus AIC-selected model). We provide these coefficients, t-statistics, and p-values for the sake of completeness. But because of the high degree of multicollinearity among the predictor variables, we are wary of over-interpreting the individual coefficients. We use these regression results only to report R^2 in the main paper, which quantifies the overall degree of predictability of the factor scores.

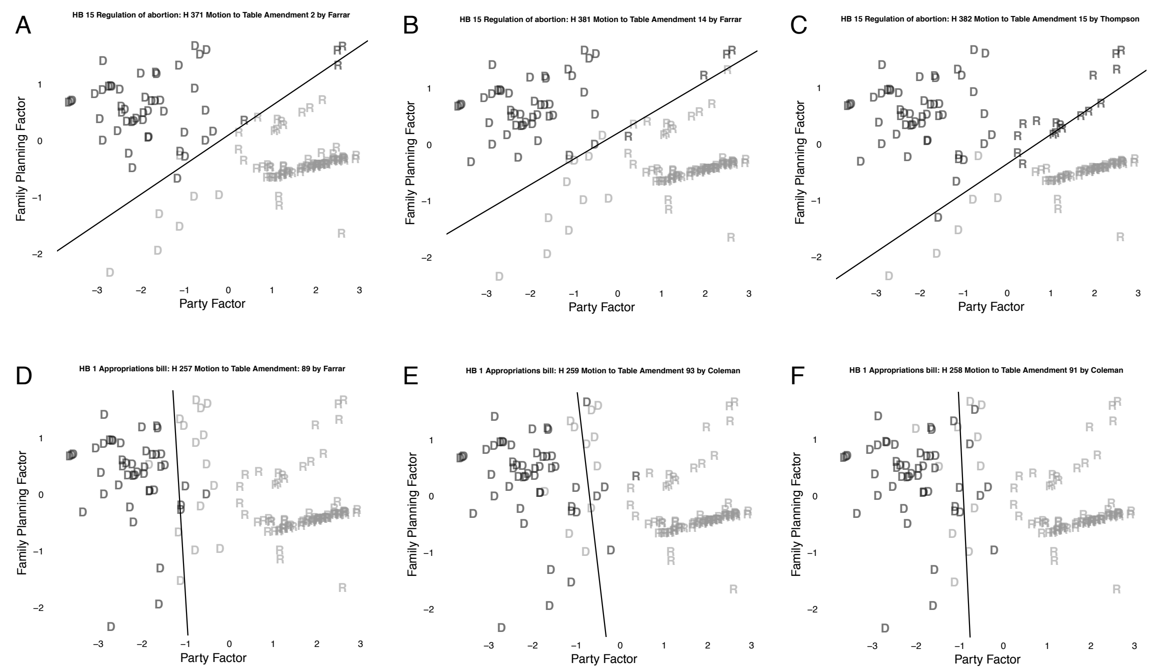

Figure 5: Six examples of family planning votes in 2003. Within each panel, the points show the legislators' estimated locations (f_{i1}, f_{i2}) in factor space. Each point is labeled R or D to indicate the legislator's party, and shaded to indicate the legislator's vote on the roll call in question (light grey for yes, dark grey for no). The cut line (which is perpendicular to the factor-loadings vector for that vote) is shown as a solid line. Panels A-C show three votes on the Women's Right to Know Act. Panels D-F show three votes on funding for contraception.

We note that our regression models are fit in a second, separate stage after the ideal-point analyses are complete. That is, they treat the posterior-mean factor scores as the fixed dependent variables in the regression. We recognize that a full Bayesian analysis might combine the regression model for the factor scores with the model for the votes in the legislature, such that uncertainty in the factor score is reflected in uncertainty over the regression estimates of the effects of covariates on the factor scores. However, we have pursued a two-stage strategy, which conditions on the posterior means (Bayes estimates) of the factor scores from the measurement model, treating the ideal points as known, fixed quantities in the analysis from that point on. We do so because we do not want the second-stage linear model to influence the estimates of the factor scores directly. This would open the door to the possibility of serious model mis-specification errors that are hard to diagnose, and we would feel much less confident interpreting the results.

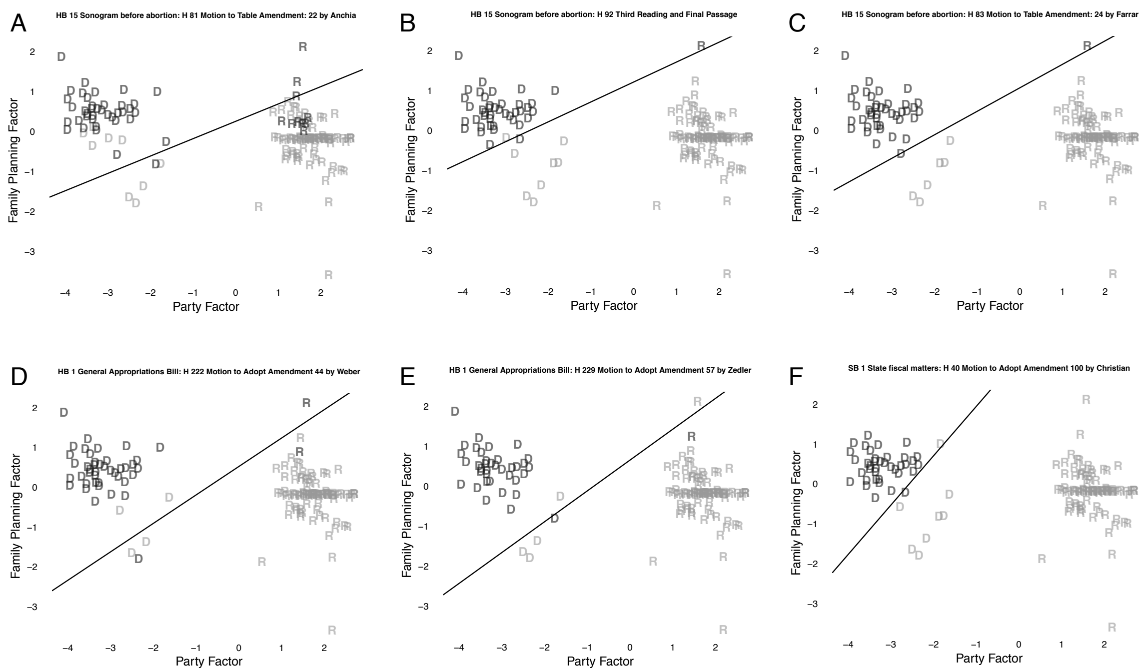

Figure 6: Six examples of family-planning votes in 2011. Within each panel, the points show the legislators' estimated locations (f_{i1}, f_{i2}) in factor space. Each point is labeled R or D to indicate the legislator's party, and shaded to indicate the legislator's vote on the roll call in question (light grey for yes, dark grey for no). The cut line (which is perpendicular to the factor-loadings vector for that vote) is shown as a solid line. Panels A-C show three votes on the sonogram bill. Panels D-F show three votes on funding for contraception.

References

[1] J. H. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–79, 1993.

[2] S. Jackman. Multidimensional analysis of roll call data via Bayesian simulation: Identification, estimation, inference, and model checking. *Political Analysis*, 9(3), 2001.

[3] P. Hahn, C. M. Carvalho, and J. G. Scott. A sparse factor-analytic probit model for Congressional voting patterns. *Journal of the Royal Statistical Society, Series C*, 61(4):619–35, 2012.

[4] A.D. Martin and K.M. Quinn. Dynamic ideal point estimation via Markov chain Monte Carlo for the U.S. Supreme Court, 1953–1999. *Political Analysis*, 10 (2):134–53, 2002.

[5] J. Geweke and G. Zhou. Measuring the pricing error of the arbitrage pricing theory. *The Review of Financial Studies*, 9:557–587, 1996.

[6] H. Lopes and M. West. Bayesian model assessment in factor analysis. *Statistica Sinica*, 14:41–67, 2004.

[7] M. West. Bayesian factor regression models in the "large p, small n" paradigm. In J. M. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith, and M. West, editors, *Bayesian Statistics 7*, pages 723–732. Oxford University Press, 2003.

[8] Texas Department of State Health Services, "The Women's Right to Know Act" website: http://www.dshs.state.tx.us/wrtk/

[9] R. Carroll, J. Lewis, J. Lo, K. Poole, and H. Rosenthal. Measuring bias and uncertainty in DW-NOMINATE ideal point estimates via the parametric bootstrap. *Political Analysis*, 17:261–75, 2009.

[10] N. McCarty, K. T. Poole, and H. Rosenthal. *Income Redistribution and the Realignment of American Politics*. American Enterprise Institute, 1997.

[11] D. Gamerman and H. Lopes. *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Chapman and Hall/CRC, 2nd edition, 2006.

Table 1: 2003 session. Three batches of votes whose factor loadings were held fixed. Batch 1 was used to fit the model whose results are quoted in the main manuscript. Batches 2 and 3 were used to check the robustness of our results to the particular bills chosen as fixed.

Batch	Category	Bill	Vote #	Summary
1	Republican	HB 9999	H 568	Gerrymandering: redistricting bill whose passage allowed the Republicans to gain a majority in the Texas delegation to the U.S. House of Representatives in the 2004 elections.
	Democrat	HB 4	H 92	Tort reform: failed Democratic motion to kill an amendment to a tort-reform bill that tightened the standards for admissibility of evidence used by plaintiffs in civil actions.
	Family planning	HB 15	H 372	Tabled an amendment that would have specifically excluded birth-control devices and oral contraceptives from the definition of "abortion" in the Women's Right to Know Act.
2	Republican	HB 2	H 11	Fiscal issue: killed a Democratic proposal to redirect state highway funds to the Children's Health Insurance Program.
	Democrat	HB 1	H 285	Public employees: failed Democratic amendment to the budget bill that would have increased funding for state-employee health plans.
	Family planning	HB 15	H 373	Tabled an amendment to the Women's Right to Know Act that would have increased to 23 weeks (from 16 weeks) the fetal age after which abortions can take place only in an ambulatory surgery center.
3	Republican	HB 2292	H 357	Medicaid/fiscal issue: killed a Democratic proposal that would have provided additional state funding for prescription drug benefits under the Public Assistant Health Benefit.
	Democrat	SB 1952	H 757	Higher education: failed Democratic amendment to a bill requiring the State Comptroller to conduct performance audits of the state's student loan program.
	Family planning	HB 15	H 374	Tabled an amendment that would have removed the 24-hour "reflection period" from the Women's Right to Know Act.

Table 2: 2011 session. Three batches of votes whose factor loadings were held fixed. Batch 1 was used to fit the model whose results are quoted in the main manuscript. Batches 2 and 3 were used to check the robustness of our results to the particular bills chosen as fixed.

Batch	Category	Bill	Vote #	Summary
1	Republican	SB 14	H 1128	Voter ID: final passage of bill that required all voters to show identification at the polls.
	Democrat	HB 12	H 884	"Sanctuary cities": failed adoption of a Democratic amendment regarding state enforcement of immigration policy.
	Family planning	HB 15	H 79	Tabled an amendment obligating the state to pay for the health care of the child or children born to a woman who, after seeking an abortion and being required to listen to the fetal heartbeat, changed her mind and refused an abortion.
2	Republican	SB 14	H 114	Voter ID: killed a Democratic amendment that would have allowed a paycheck or proof of employment to satisfy the voter ID requirements.
	Democrat	HB 275	H 1182	State fiscal issues: Democratic amendment that would have used money from the state "Rainy Day Fund" to fill a large gap in the 2011 budget.
	Family planning	HB 15	H 68	Tabled an amendment to that would have allowed a woman seeking an abortion to refuse to listen to the fetal heartbeat.
3	Republican	HCR 18	H 366	Federal fiscal issues: killed a Democratic amendment regarding Social Security and Medicare to a House resolution urging the federal government to adopt a balanced budget.
	Democrat	HB 2229	H 535	Health care: a bill that would have established and funded a Texas HIV Medication Advisory Committee as part of the Department of State Health Services.
	Family planning	HB 15	H 71	Tabled an amendment that would have allowed a woman to avoid certain requirements of HB 15 by signing a waiver in the presence of her physician.

Table 3: Family-planning factor scores versus constituency characteristics for Democrats in 2003 (full model fit by ordinary least squares).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-1.01	0.938	-1.077	0.286
sexMale	-0.177	0.275	-0.644	0.523
Catholic	0.224	0.215	1.042	0.302
pctWhite	0.01	0.01	0.972	0.336
pctNoncitizen	0.025	0.016	1.605	0.114
pctSingleParentFamilies	0.024	0.017	1.424	0.16
pctRural	-0.01	0.01	-1.007	0.318
pctBachelorsDegree	0.008	0.018	0.419	0.677
pctLivingInPoverty	0.005	0.02	0.232	0.818

Table 4: Family-planning factor scores versus constituency characteristics for Democrats in 2003 (reduced model fit by minimizing AIC).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-0.919	0.473	-1.942	0.057
pctNoncitizen	0.022	0.01	2.119	0.038
pctSingleParentFamilies	0.031	0.014	2.296	0.025

Table 5: Family-planning factor scores versus constituency characteristics for Republicans in 2003 (full model fit by OLS).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-0.952	0.549	-1.733	0.087
sexMale	-0.382	0.129	-2.969	0.004
Catholic	0.165	0.153	1.076	0.285
pctWhite	0.004	0.006	0.618	0.538
pctNoncitizen	0	0.017	-0.016	0.987
pctSingleParentFamilies	0.019	0.018	1.071	0.287
pctRural	-0.001	0.005	-0.208	0.836
pctBachelorsDegree	0.012	0.008	1.529	0.13
pctLivingInPoverty	-0.003	0.018	-0.19	0.85

Table 6: Family-planning factor scores versus constituency characteristics for Republicans in 2003 (reduced model fit by minimizing AIC).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-0.274	0.159	-1.723	0.088
sexMale	-0.375	0.123	-3.054	0.003
pctBachelorsDegree	0.01	0.004	2.746	0.007

Table 7: Family-planning factor scores versus constituency characteristics for Democrats in 2011 (full model fit by ordinary least squares).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	0.639	0.741	0.862	0.394
sexMale	0.054	0.182	0.295	0.77
Catholic	-0.324	0.188	-1.718	0.093
pctWhite	0.016	0.01	1.586	0.121
pctNoncitizen	0.024	0.012	1.948	0.058
pctSingleParentFamilies	0.02	0.014	1.49	0.144
pctRural	-0.009	0.013	-0.687	0.496
pctBachelorsDegree	-0.026	0.013	-1.962	0.057
pctLivingInPoverty	-0.048	0.015	-3.224	0.003

Table 8: Family-planning factor scores versus constituency characteristics for Democrats in 2011 (reduced model fit by minimizing AIC).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	0.458	0.658	0.696	0.491
Catholic	-0.305	0.183	-1.667	0.103
pctWhite	0.013	0.009	1.479	0.147
pctNoncitizen	0.027	0.01	2.629	0.012
pctSingleParentFamilies	0.025	0.011	2.201	0.033
pctBachelorsDegree	-0.023	0.012	-1.869	0.069
pctLivingInPoverty	-0.05	0.014	-3.596	0.001

Table 9: Family-planning factor scores versus constituency characteristics for Republicans in 2011 (full model fit by OLS).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-0.793	0.651	-1.218	0.226
sexMale	-0.284	0.166	-1.712	0.09
Catholic	-0.024	0.198	-0.122	0.903
pctWhite	0	0.007	-0.03	0.976
pctNoncitizen	-0.02	0.019	-1.061	0.291
pctSingleParentFamilies	0.023	0.017	1.349	0.181
pctRural	0.003	0.005	0.533	0.596
pctBachelorsDegree	0.018	0.009	2.03	0.045
pctLivingInPoverty	-0.008	0.016	-0.54	0.59

Table 10: Family-planning factor scores versus constituency characteristics for Republicans in 2011 (reduced model fit by minimizing AIC).

Variable	Estimate	Std. Error	t statistic	p value
(Intercept)	-0.124	0.21	-0.591	0.556
sexMale	-0.321	0.16	-2.009	0.047
pctNoncitizen	-0.019	0.013	-1.431	0.156
pctBachelorsDegree	0.013	0.005	2.53	0.013