MODGIRT: Multidimensional Dynamic Scaling of Aggregate Survey Data

Supplemental Material

Elissa Berwick Devin Caughey

Table of contents

```r
dir <- "/Users/elissa/Library/CloudStorage/Dropbox/Projects/mdgirt/"
update_files <- function(replication_path, paper_path) {
  file_list <-
    list.files(
      path = replication_path,
      full.names = TRUE
    )
  file.copy(from = file_list, to = paper_path, overwrite = TRUE)
}
update_files(
  paste0(dir, "replication/figures"),
  paste0(dir, "paper/final/figures")
)
update_files(
  paste0(dir, "replication/tables"),
  paste0(dir, "paper/final/tables")
)
```

1 Derivation of MODGIRT model

This section provides a more detailed derivation of the MODGIRT model.

Let subject i's binary response to question q be

$$y_{iq} = \begin{cases} 1 \text{ if } \beta_q' \theta_i + \epsilon_{iq} > \alpha_q \\ 0 \text{ otherwise.} \end{cases} \tag{1}$$

Under the conventional identification restriction $\epsilon_{iq} \sim \mathcal{N}(0, 1)$, the probability of a positive response is

$$\Pr(y_{iq} = 1 \mid \alpha_q, \beta_q, \theta_i) = \Pr(\beta_q' \theta_i - \alpha_q + \epsilon_{iq} > 0) \tag{2}$$
$$= \Phi(\beta_q' \theta_i - \alpha_q). \tag{3}$$

This is the individual-level probit IRT model.

We assume that ideal points are distributed multivariate normal around a group-specific mean vector $\bar{\theta}_g$:

$$\theta_{g[i]} \sim \mathcal{N}_D(\bar{\theta}_g, \Sigma_\theta),$$

where unlike a standard multinomial distribution Σ_θ is a D-by-D variance-covariance matrix.

This assumption allows us to derive $p_{gq} = \Pr(y_{iq} = 1 \mid \alpha_q, \beta_q, \bar{\theta}_{g[i]}, \Sigma_\theta)$. Given multivariate normality within groups,[1]

$$\beta_q' \theta_{g[i]} - \alpha_q \sim \mathcal{N}(\beta_q' \bar{\theta}_g - \alpha_q, \beta_q' \Sigma_\theta \beta_q),$$

and because ϵ_{iq} is an independent standard normal variable,

$$\beta_q' \theta_{g[i]} - \alpha_q + \epsilon_{iq} \sim \mathcal{N}(\beta_q' \bar{\theta}_g - \alpha_q, 1 + \beta_q' \Sigma_\theta \beta_q).$$

Thus, the probability that subject i randomly sampled from group g gives a positive answer to question q is

$$p_{gq} = \Phi\left(\frac{\beta_q' \bar{\theta}_g - \alpha_q}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}}\right). \tag{4}$$

If responses are conditionally independent,[2] the number of positive answers in group g is distributed

$$s_{gq} \sim \text{Binomial}(n_{gq}, p_{gq}).$$

The binary group-level IRT model just described can be extended to multiple ordered response categories using an ordinal cumulative model (Samejima 1997), which defines the probability of selecting response option $k \in 1 \ldots K_q$ as

$$\Pr(y_{iq} = k) = \Pr(y_{iq} > k - 1) - \Pr(y_{iq} \geq k).$$

Under the ordinal variant of the model, the probability of subjects in group g selecting category k on item q is

$$\Pr(y_{gq} = k) = \Phi\left(\frac{\beta_q' \bar{\theta}_g - \alpha_{q,k-1}}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}}\right) - \Phi\left(\frac{\beta_q' \bar{\theta}_g - \alpha_{q,k}}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}}\right) = p_{gqk},$$

where the $K_q + 1$ thresholds are ordered

$$-\infty = \alpha_{q,0} < \alpha_{q,1} < \cdots < \alpha_{q,K_q-1} < \alpha_{q,K_q} = \infty.$$

[1] Affine transformations of the multivariate normal distribution operate such that if $X \sim \mathcal{N}(\mu, \Sigma)$ and $\mathbf{Y} = \mathbf{c} + \mathbf{B}\mathbf{X}$, then $\mathbf{Y} \sim \mathcal{N}(\mathbf{c} + \mathbf{B}\mu, \mathbf{B}\Sigma\mathbf{B}')$.

[2] That is, independent conditional on the item parameters and the group means and covariances. This independence is violated if respondents answer more than one question each.

Letting $\mathbf{p}_{gq} = (p_{gq1}, \dots, p_{gqK_q})$, the response counts $\mathbf{s}_{gq} = (s_{gq1}, \dots, s_{gqK_q})$ can be modeled as

$$\mathbf{s}_{gq} \sim \text{Multinomial}(\mathbf{p}_{gq}).$$

Dropping constants, the unnormalized probability mass function for this model is

$$\text{Multinomial}(s_{gqk} \mid p_{gqk}) = \prod_{k=1}^{K} p_{gqk}^{s_{gqk}},$$

where s_{gqk} may take non-integer values if observations are weighted (e.g., to adjust for nonresponse).

2 Stan code

Note that for computational efficiency, we calculate p_{gqk} using Stan's `ordered_probit_lupmf(k | eta, c)` function, where

$$\begin{aligned}
\texttt{k} &= k, \\
\texttt{eta} &= \frac{\beta_q' \bar{\theta}_g}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}}, \text{ and} \\
\texttt{c} &= \left(\frac{\alpha_{q,1}}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}}, \dots, \frac{\alpha_{q,K-1}}{\sqrt{1 + \beta_q' \Sigma_\theta \beta_q}} \right).
\end{aligned}$$

The full Stan code code for the MODGIRT model is printed below.

```
functions {
  /* De-mean and 'whiten' (cov = I) XX */
  matrix whiten(matrix XX) {
    matrix[rows(XX), cols(XX)] DM;
    matrix[cols(XX), cols(XX)] SS;
    matrix[cols(XX), cols(XX)] PP;
    matrix[cols(XX), cols(XX)] WW;
    for (d in 1:cols(XX)) {
      DM[ : , d] = XX[ : , d] - mean(XX[ : , d]); /* de-mean each column */
    }
    SS = crossprod(DM) ./ (rows(XX) - 1.0); /* covariance of XX */
    PP = inverse_spd(SS); /* precision of XX */
    WW = cholesky_decompose(PP); /* Cholesky decomposition of precision */
    return DM * WW; /* de-meaned and whitened XX */
```

```
  }
}
data {
  int<lower=1> T; // number of periods
  int<lower=1> G; // number of groups
  int<lower=1> Q; // number of items
  int<lower=1> K; // max number of answer options
  int<lower=1> D; // number of latent dimensions
  array[T, G, Q, K] real<lower=0> SSSS; // (weighted) # of responses
  array[Q, D] int beta_nonzero; // loading point restrictions
  array[Q, D] int beta_sign; // loading sign restrictions
}
parameters {
  array[Q] ordered[K - 1] z_alpha; // thresholds (difficulties)
  array[Q, D] real beta_free; // unconstrained discriminations
  array[Q, D] real<lower=0> beta_pos; // sign-constrained discriminations
  array[T, G, D] real z_bar_theta;
  vector<lower=0>[D] sd_theta; // within-group SD of theta
  corr_matrix[D] corr_theta; // within-group corr of theta across dims
  vector<lower=0>[D] sd_bar_theta_evol; // evolution SD of bar_theta
  corr_matrix[D] corr_bar_theta_evol;   // cross-dim transition corr
}
transformed parameters {
  array[T, Q] vector[K - 1] alpha; // thresholds (difficulty)
  array[T, Q] real alpha_drift;     // question-specific time trends
  matrix[Q, D] beta;
  array[T] matrix[G, D] bar_theta; // group ideal point means
  cov_matrix[D] Sigma_theta; // within-group variance-covariance
  cov_matrix[D] Omega; // transition variance-covariance
  Sigma_theta = quad_form_diag(corr_theta, sd_theta);
  Omega = quad_form_diag(corr_bar_theta_evol, sd_bar_theta_evol);
  matrix[D, D] chol_Omega = cholesky_decompose(Omega);
  for (q in 1:Q) {
    for (d in 1:D) {
      if (beta_sign[q, d] == 0) {
        beta[q, d] = beta_nonzero[q, d] * beta_free[q, d];
      } else if (beta_sign[q, d] > 0) {
        beta[q, d] = beta_nonzero[q, d] * beta_pos[q, d];
      } else if (beta_sign[q, d] < 0) {
        beta[q, d] = -1.0 * beta_nonzero[q, d] * beta_pos[q, d];
      }
    }
```

```
  }
  for (t in 1:T) {
    if (t == 1) {
      /* Make period 1 ideal points orthogonal and mean zero */
      bar_theta[t][1:G, 1:D] =
        whiten(to_matrix(z_bar_theta[t, 1:G, 1:D]));
    }
    if (t > 1) {
      for (g in 1:G) {
        vector[D] bt_vec_tm1 = to_vector(bar_theta[t-1][g, 1:D]);
        vector[D] zbt_t = to_vector(z_bar_theta[t, g, 1:D]);
        vector[D] bt_vec_t = bt_vec_tm1 + chol_Omega * zbt_t;
        bar_theta[t][g, 1:D] = to_row_vector(bt_vec_t);
      }
    }
    for (q in 1:Q) {
      alpha_drift[t, q] = 0; // could estimate but here set to 0.
      alpha[t, q][1:(K - 1)] = z_alpha[q][1:(K - 1)] + alpha_drift[t, q];
    }
  }
}
model {
  /* Priors */
  to_array_1d(z_bar_theta[1:T, 1:G, 1:D]) ~ std_normal();
  to_array_1d(beta_free[1:Q, 1:D]) ~ std_normal();
  to_array_1d(beta_pos[1:Q, 1:D]) ~ std_normal();
  for (q in 1:Q) {
    z_alpha[q][1:(K - 1)] ~ std_normal();
  }
  sd_theta ~ cauchy(0, 1);
  corr_theta ~ lkj_corr(2);
  sd_bar_theta_evol ~ cauchy(0, .1);
  corr_bar_theta_evol ~ lkj_corr(2);
  /* Likelihood */
  if (K > 1) {
    /* ordinal outcomes */
    for (t in 1:T) {
      for (q in 1:Q) {
        real denom; // denominator of linear predictor
        vector[K - 1] cuts; // ordered probit cutpoints
          real sbs;
          sbs = quad_form(Sigma_theta[1:D, 1:D], to_vector(beta[q][1:D]));
```

```
        denom = sqrt(1 + sbs);
        cuts = alpha[t, q][1:(K - 1)] / denom;
        for (g in 1:G) {
          real eta; // linear predictor
          eta = to_row_vector(beta[q][1:D])
                * to_vector(bar_theta[t, g, 1:D]) / denom;
          for (k in 1:K) {
            if (SSSS[t, g, q, k] > 0) {
              /* Add SSSS[t, g, q, k] log ordinal probit densities */
              target += SSSS[t, g, q, k] * ordered_probit_lupmf(k | eta, cuts);
            }
          }
        }
      }
    }
  }
}
```

3 Simulation results

3.1 Local independence

The MODGIRT model assumes that responses are independent conditional on the item and group parameters. This assumption of local independence will be violated if subjects respond to multiple items. To assess the effect of this violation, we vary the number of items each individual responds to, examining effects for $A = 10$, $A = 20$, and $A = 40$ responses (out of 50 total items) while holding constant the total number of observed responses $R = 100{,}000$. We held R constant so that we could observe how the model responds in situations where the total amount of data remains the same but number of unique individuals providing that information decreases. Mean squared errors, correlation, and average coverage of 90% intervals for these simulations appear in Table 1 alongside results for the baseline scenario, which are highlighted in blue.

The main consequence of violating the local independence assumption is on coverage of group ideal point estimates. When only one question is answered, a 90% interval is too conservative, covering the true ideal point 95% of the time. This rate declines to 92% when ten questions are answered and 90% when twenty questions are answered. When forty questions are answered the estimated 90% intervals are more seriously over-confident, covering the true value of the ideal point only 85% of the time. Mean squared error in estimates of ideal points also increases as more questions are answered, rising to 0.006 with $A = 40$. However, if 80% of responses were available in an applied setting, it would be feasible to estimate an individual-level model instead. In settings where the MODGIRT model is most useful, the effect of violating local independence on coverage is minor. Meanwhile, coverage of discrimination parameters remains conservative.

Table 1: Effect of violating local independence assumption on MODGIRT model fit

Parameter type	A	N	G	R	Mean MSE	Correlation	Mean 90% CI coverage
Difficulty	1	2,000	50	100,000	0.007	0.997	0.902
Difficulty	10	200	50	100,000	0.007	0.997	0.898
Difficulty	20	100	50	100,000	0.007	0.997	0.899
Difficulty	40	50	50	100,000	0.007	0.997	0.902
Ideal point	1	2,000	50	100,000	0.003	0.999	0.950
Ideal point	10	200	50	100,000	0.003	0.998	0.920
Ideal point	20	100	50	100,000	0.004	0.998	0.897
Ideal point	40	50	50	100,000	0.006	0.997	0.853
Discrimination	1	2,000	50	100,000	0.013	0.994	0.940
Discrimination	10	200	50	100,000	0.013	0.994	0.941
Discrimination	20	100	50	100,000	0.013	0.994	0.937
Discrimination	40	50	50	100,000	0.012	0.994	0.948

3.2 Heteroskedasticity

MODGIRT also assumes that within-group variance of $\theta_{g[i]}$ is constant. To examine the effect of heteroskedasticity across groups, we randomly draw a vector of scaling coefficients $\tau_g \sim 1 + \rho \cdot \mathcal{U}(-1, 1)$ for each group and use it to vary the variance-covariance matrix such that $\Sigma_{\theta g} = Diag(\tau_g)\Sigma_\theta Diag(\tau_g)$. This setup enables us to control the amount of heteroskedasticity across groups by setting the factor ρ. The smaller ρ is, the more similar the within-group variance structure will be across groups, with $\rho = 0$ implying homoskedasticity.

We apply the MODGIRT model to datasets generated under these new conditions to see how well the simulated parameters can be recovered, first with $\rho = 0.5$ and then with $\rho = 0.9$. Mean-squared error, correlation between simulated and estimated parameters, and mean coverage appear in Table 2 alongside results for the baseline scenario (highlighted in blue). When $\rho = 0.5$, mean-squared error for group ideal points is slightly higher than in the baseline scenario, but the coverage of 90% credible intervals is 89%. However, when we increase heteroskedasticity by setting $\rho = 0.9$, coverage for group ideal points drops below 90%.[3] Mean-squared errors for difficulty and item discrimination parameters also increase with ρ, though at a slower rate than errors for ideal points. Coverage for both parameter types decreases with ρ,

[3]The larger ρ corresponds to a situation where within-group standard deviations can be up to 90% smaller or larger than in the average group. In other words, the scaling coefficients τ_g are distributed $\mathcal{U}(0.1, 1.9)$.

though for discrimination parameters the coverage rate remains above 90% even when $\rho = 0.9$. Average correlations between simulated and estimated values are consistently around 0.99 for all parameter types.

Table 2: Effect of group-level heterogeneity on MODGIRT model fit

Parameter type	Rho	G	N	R	A	Mean MSE	Correlation	Mean 90% CI coverage
Difficulty	0	50	2,000	100,000	1	0.007	0.997	0.902
Difficulty	0.5	50	2,000	100,000	1	0.008	0.996	0.888
Difficulty	0.9	50	2,000	100,000	1	0.014	0.993	0.846
Ideal point	0	50	2,000	100,000	1	0.003	0.999	0.950
Ideal point	0.5	50	2,000	100,000	1	0.008	0.996	0.891
Ideal point	0.9	50	2,000	100,000	1	0.016	0.992	0.801
Discrimination	0	50	2,000	100,000	1	0.013	0.994	0.940
Discrimination	0.5	50	2,000	100,000	1	0.016	0.992	0.936
Discrimination	0.9	50	2,000	100,000	1	0.022	0.990	0.915

3.3 Groups and individuals

We conduct a series of additional simulations to examine how well the MODGIRT model performs when we vary the number of groups, the number of total responses, and the number of individuals in each group. Table 3 shows results for these further simulations, with results for the baseline scenario highlighted in blue. In the first set, we reduce the number of groups to $G = 5$ while all other conditions remain the same as in the assumed data-generating process. This actually reduces mean-squared error and increases coverage on estimates of group ideal points relative to the baseline scenario, because if we maintain a constant total amount of data $R = 100{,}000$, having fewer groups implies that there are more individuals in each group. As a result, group ideal points are estimated with greater precision. If we instead maintain $N_g = 2{,}000$ with $G = 5$, reducing the total number of responses to $R = 10{,}000$, mean-squared errors do increase. However, in both cases coverage rates remain at or above 90%.[4]

In the second set of simulations, we return the number of groups to $G = 50$ but vary the number of individuals in each group N_g. We estimate the model with $1{,}000 < N_g < 3{,}000$, $N_g = 200$, and $100 < N_g < 300$. Results are reported in Table 3. While both varying and reducing N_g increase mean-squared error on all parameter types, reducing group size has a more serious effect on errors. Reducing group size also slightly decreases correlations. But neither reducing nor varying N_g appears to affect the coverage rates of 90% intervals, which are

[4]One additional repercussion of reducing G is that simulations tend to have more divergent transitions. Over $M = 200$ simulations with $G = 5$ and $R = 100{,}000$, all had at least one divergent transition. In simulations where $G = 50$, at least half of the simulations have no divergent transitions. Models fit with a small number of groups should probably also use a smaller adapt delta to compensate.

still above 90%. These results show that a having a larger number of groups, more individuals per group, and less variation across groups all result in more accurate estimates.

Table 3: Effect of varying the number of groups, number of responses, and the number of individuals in groups on MODGIRT model fit

Parameter type	G	N Min	N Max	R	Mean MSE	Correlation	Mean 90% CI coverage
Difficulty	5	20,000	20,000	100,000	0.009	0.996	0.907
Difficulty	5	2,000	2,000	10,000	0.039	0.980	0.899
Difficulty	**50**	**2,000**	**2,000**	**100,000**	**0.007**	**0.997**	**0.902**
Difficulty	50	1,000	3,000	100,000	0.007	0.997	0.908
Difficulty	50	200	200	10,000	0.036	0.982	0.896
Difficulty	50	100	300	10,000	0.036	0.982	0.903
Ideal point	5	20,000	20,000	100,000	0.003	0.998	0.966
Ideal point	5	2,000	2,000	10,000	0.008	0.997	0.972
Ideal point	**50**	**2,000**	**2,000**	**100,000**	**0.003**	**0.999**	**0.950**
Ideal point	50	1,000	3,000	100,000	0.005	0.998	0.950
Ideal point	50	200	200	10,000	0.032	0.985	0.950
Ideal point	50	100	300	10,000	0.035	0.983	0.950
Discrimination	5	20,000	20,000	100,000	0.020	0.991	0.947
Discrimination	5	2,000	2,000	10,000	0.065	0.969	0.952
Discrimination	**50**	**2,000**	**2,000**	**100,000**	**0.013**	**0.994**	**0.940**
Discrimination	50	1,000	3,000	100,000	0.014	0.993	0.945
Discrimination	50	200	200	10,000	0.056	0.973	0.951
Discrimination	50	100	300	10,000	0.058	0.972	0.952

4 Supplemental information for reanalysis of Cavaillé and Trump (2015)

The majority of the data for this application come from the replication archive for Cavaillé and Trump (2015), which can be accessed at https://doi.org/10.7910/DVN/L4MGG5. Further details on the data including the full text for all questions are provided in the article and its supplemental materials, available at https://doi.org/10.1086/678312. We added responses to items beyond those available in the replication archive by downloading the complete versions of the British Social Attitudes datasets used by Cavaillé and Trump from the UK Data Service.

4.1 Descriptive statistics

Table 4 reports basic descriptive statistics on the panel data used in this application.

Table 4: Descriptive statistics for the British Social Attitudes data used in this analysis

Total years	23
Total groups	30
Total items	32
Total responses	768,417
Ave items per year	16
Ave items per group	32
Ave responses per year	33,409
Ave responses per group	25,614
Ave responses per item	24,013
Pct item-year missing	51
Pct item-groups missing	0
Pct item-group-year missing	54

4.2 Dimensionality

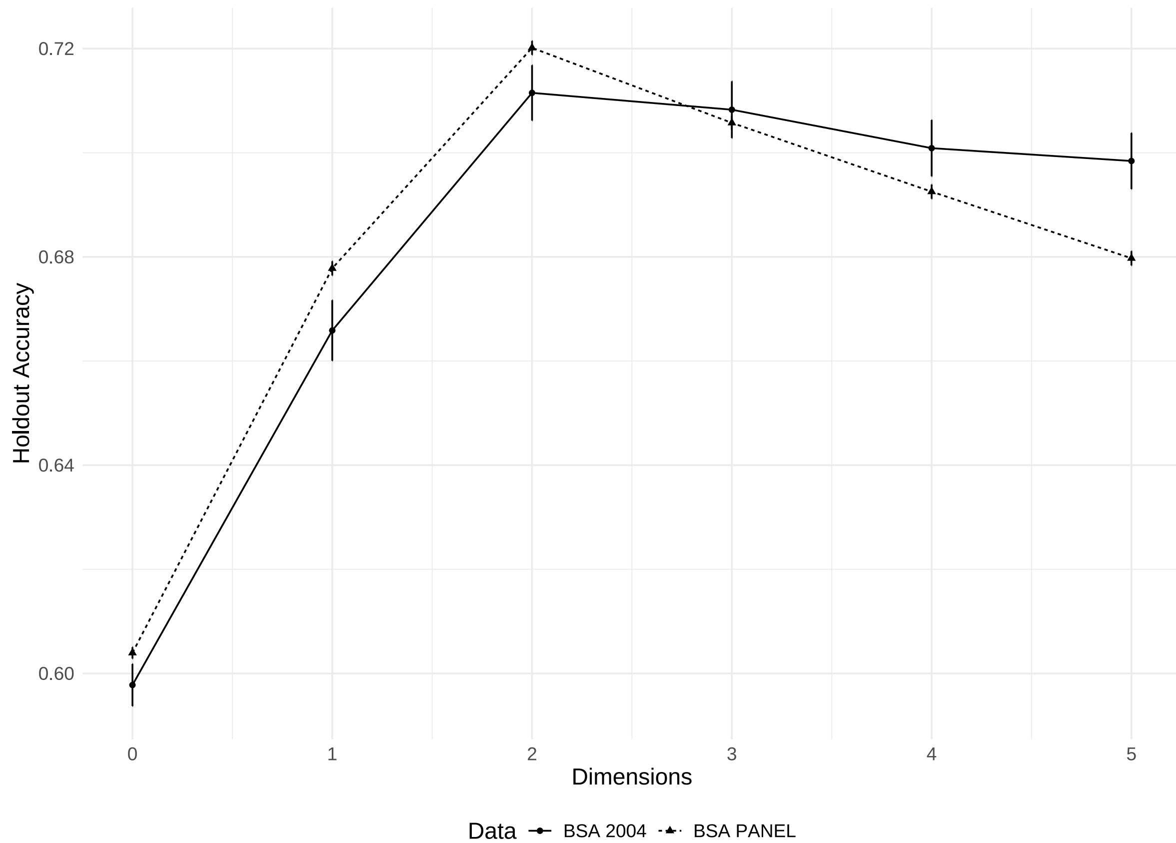

Figure 1: Holdout accuracy by dimension for models for BSA data (both 2004 and panel) shows that two dimensions improve over one dimension, and adding a third dimension does not improve over two dimensions. Accuracy estimated using methods and replication code from Marble and Tyler (2022).

5 Supplemental information for reanalysis of Caughey, O'Grady, and Warshaw (2019a)

The data for this application were obtained from the replication archive for Caughey, O'Grady, and Warshaw (2019a), which can be accessed at https://doi.org/10.7910/DVN/H9XGEB. For our analysis, we subset the data to begin in 1999. Details on the data are provided in the article's supplemental materials (SM), which are available at https://doi.org/10.1017/S0003055419000157. Pages 19–20 of the SM describe the data sources, and pages 21–32 provide details on the individual items. We identify items using the same labels as the COW SM.

Table 5: Descriptive statistics for the Caughey, O'Grady, and Warshaw (2019a) data used in this analysis

Total biennia	9
Total countries	27
Total groups	162
Total items	100
Total responses	1,817,270
Ave items per biennium	32
Ave items per country	83
Ave items per group	83
Ave responses per biennium	201,919
Ave responses per country	67,306
Ave responses per group	11,218
Ave responses per item	18,173
Pct item-biennia missing	71
Pct item-countries missing	24
Pct item-groups missing	24
Pct item-group-biennia missing	80

5.1 Descriptive statistics

Table 5 reports basic descriptive statistics on the data subset we analyze.

5.2 Item discriminations

Figure 2 reports the varimax-rotated discrimination for each item on each factor (i.e., dimension). The item responses are coded so that higher values indicate greater conservatism, and each factor is oriented so that the average discrimination on that factor has a positive sign.

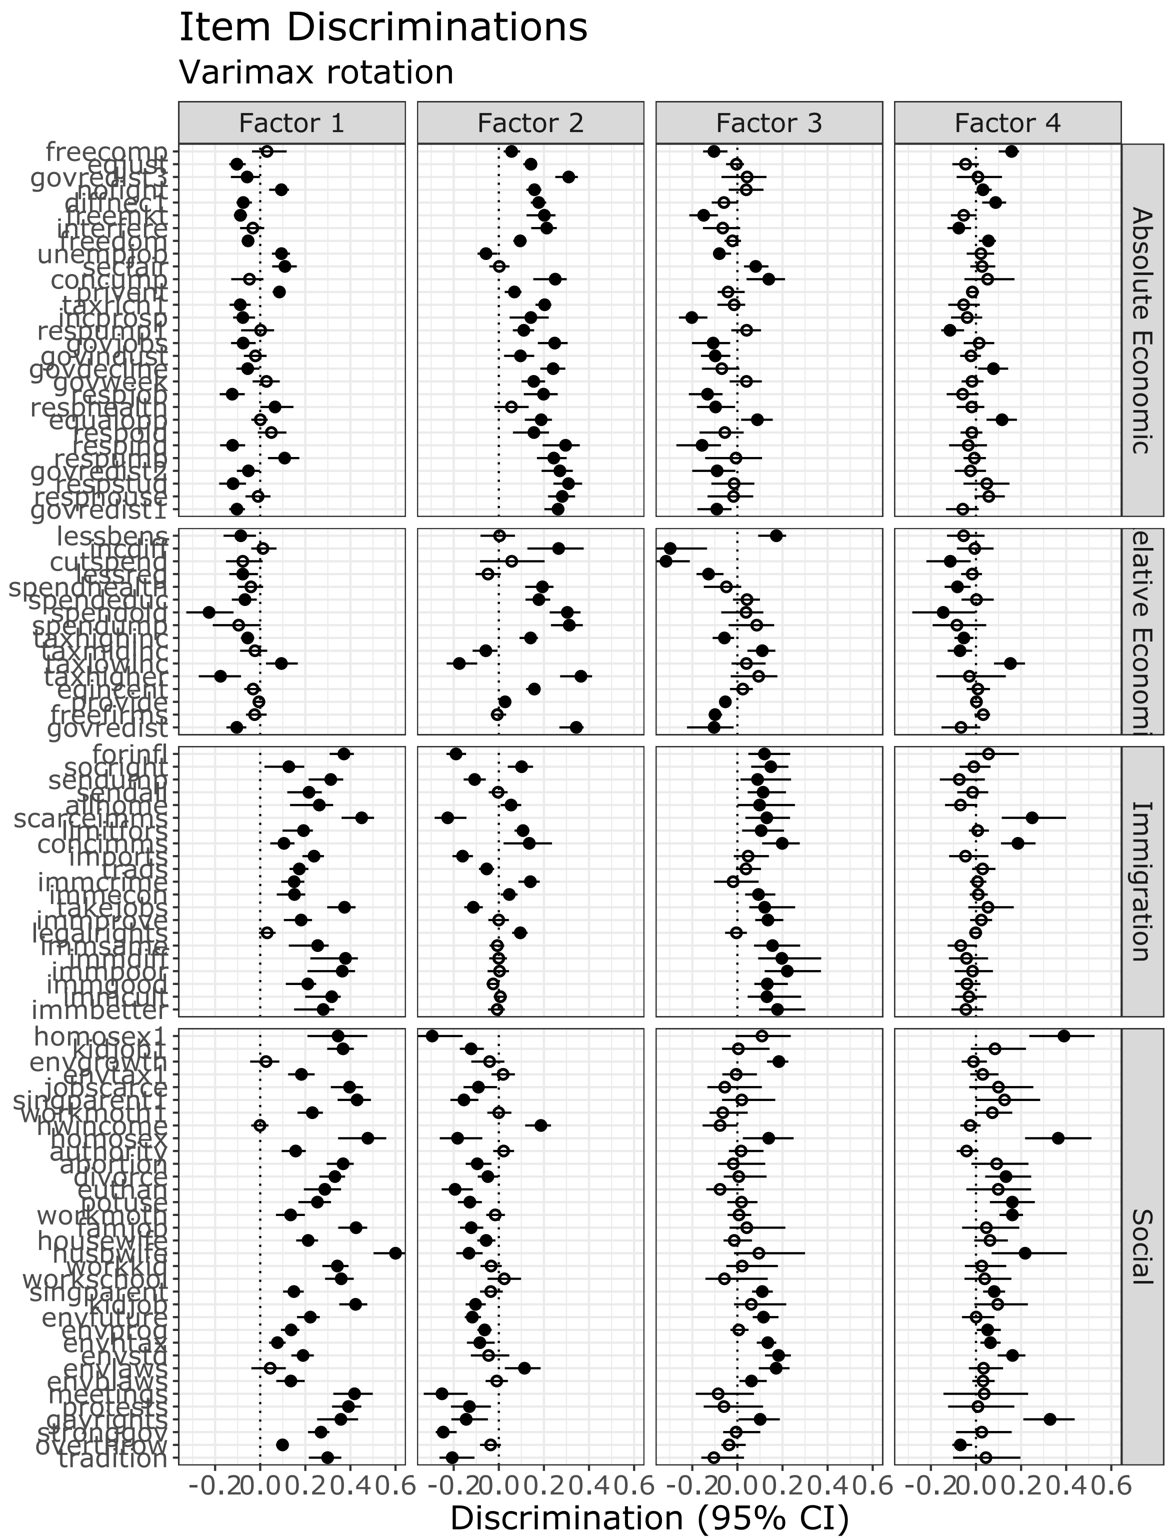

Figure 2: Item discrimination estimates by domain and factor.

5.3 Country plots

Figure 3 through Figure 6 show the estimated ideal points for each country on each of the four dimension over time.

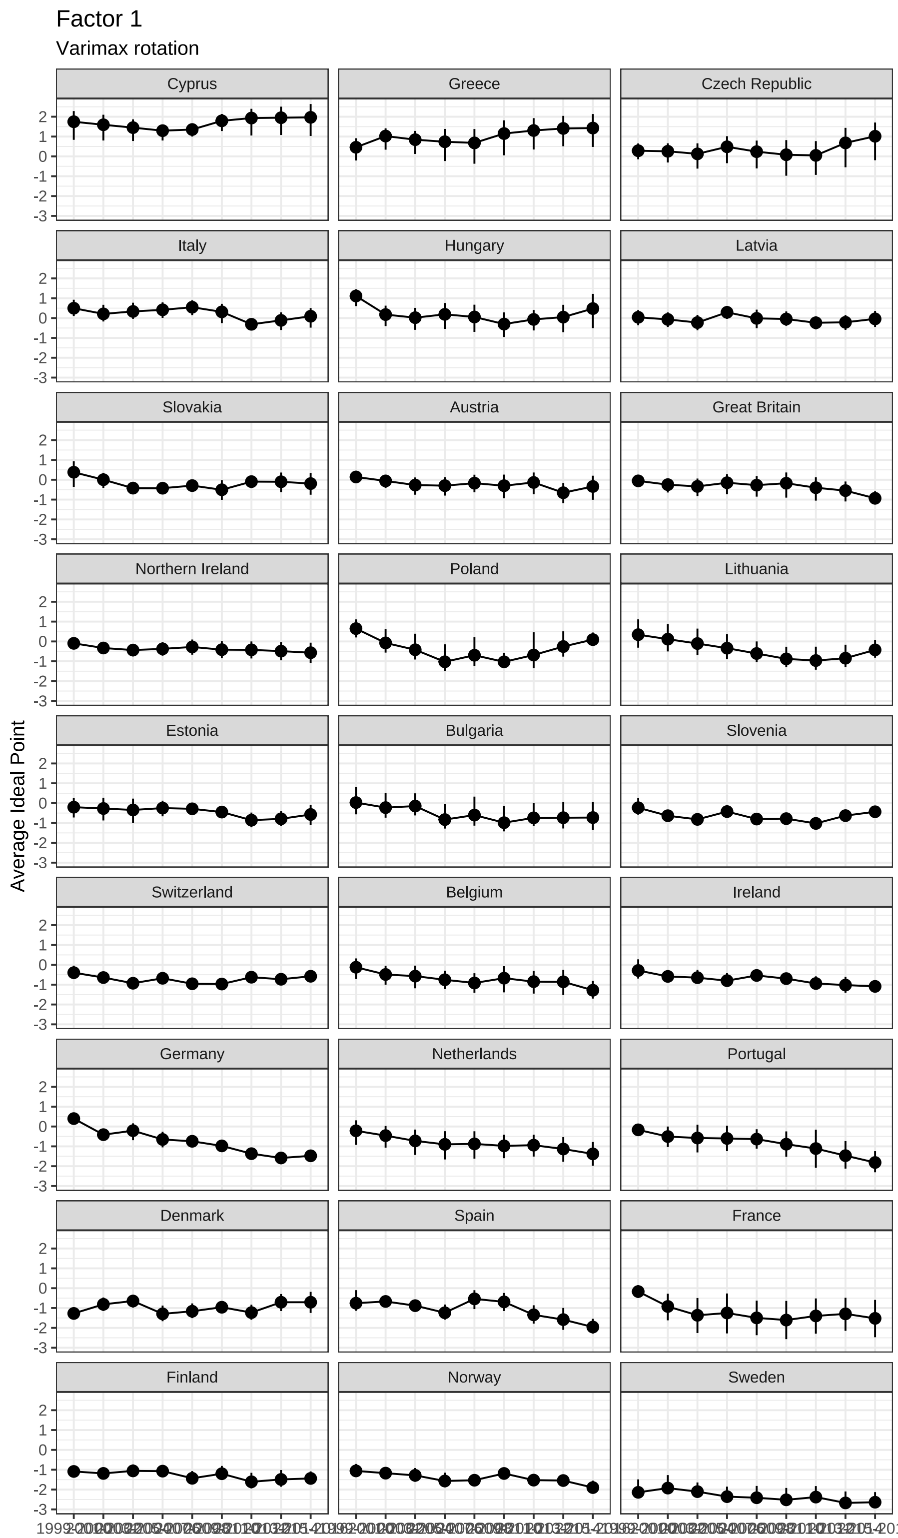

Figure 3: Country plots: Dimension 1

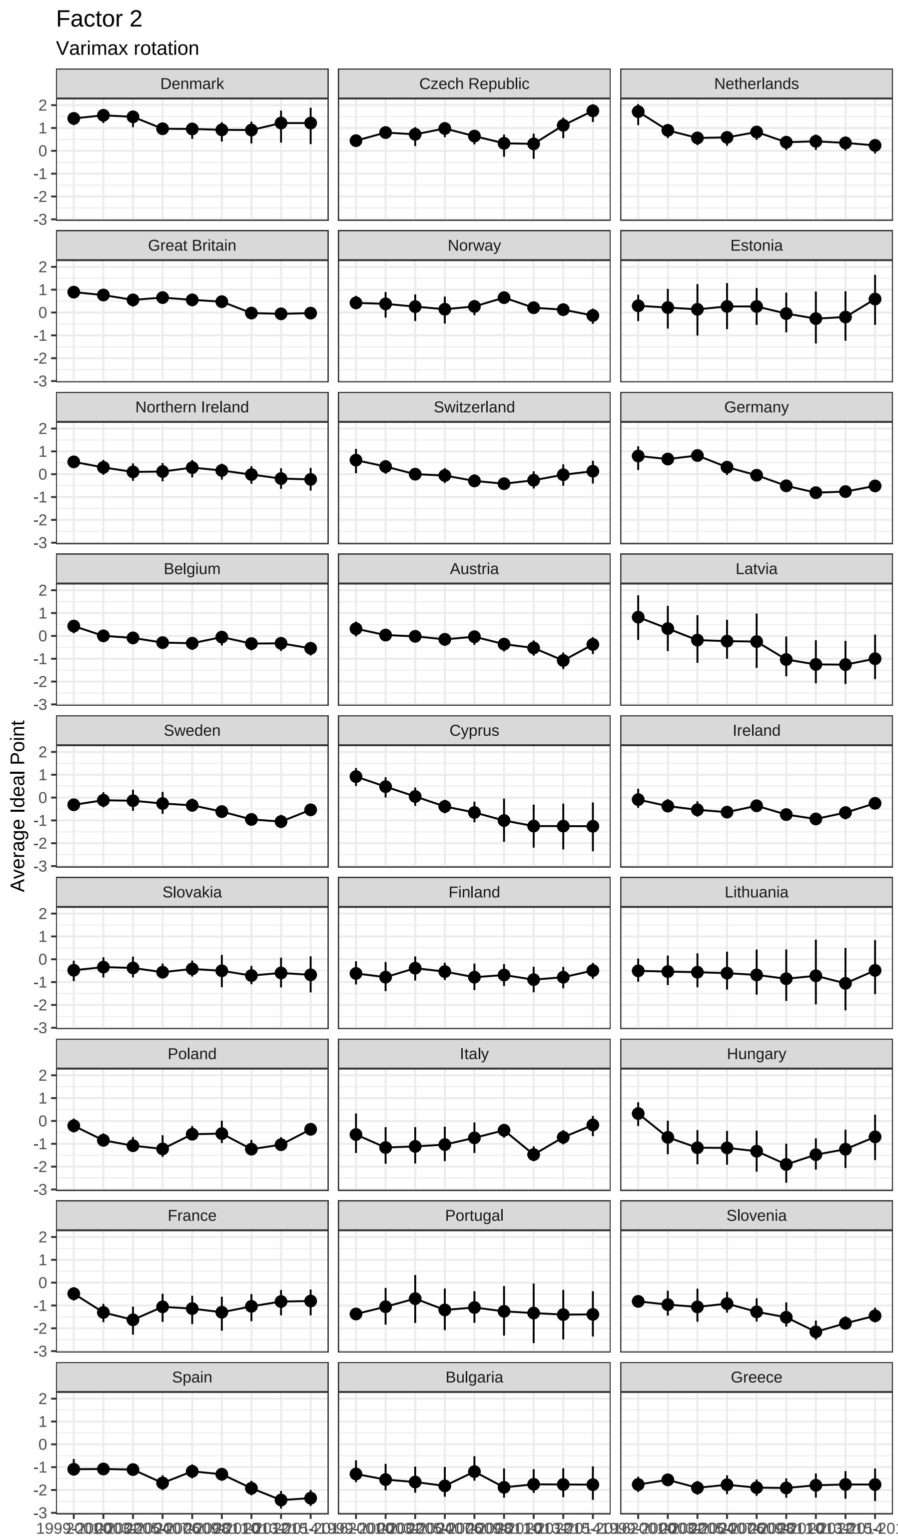

Figure 4: Country plots: Dimension 2

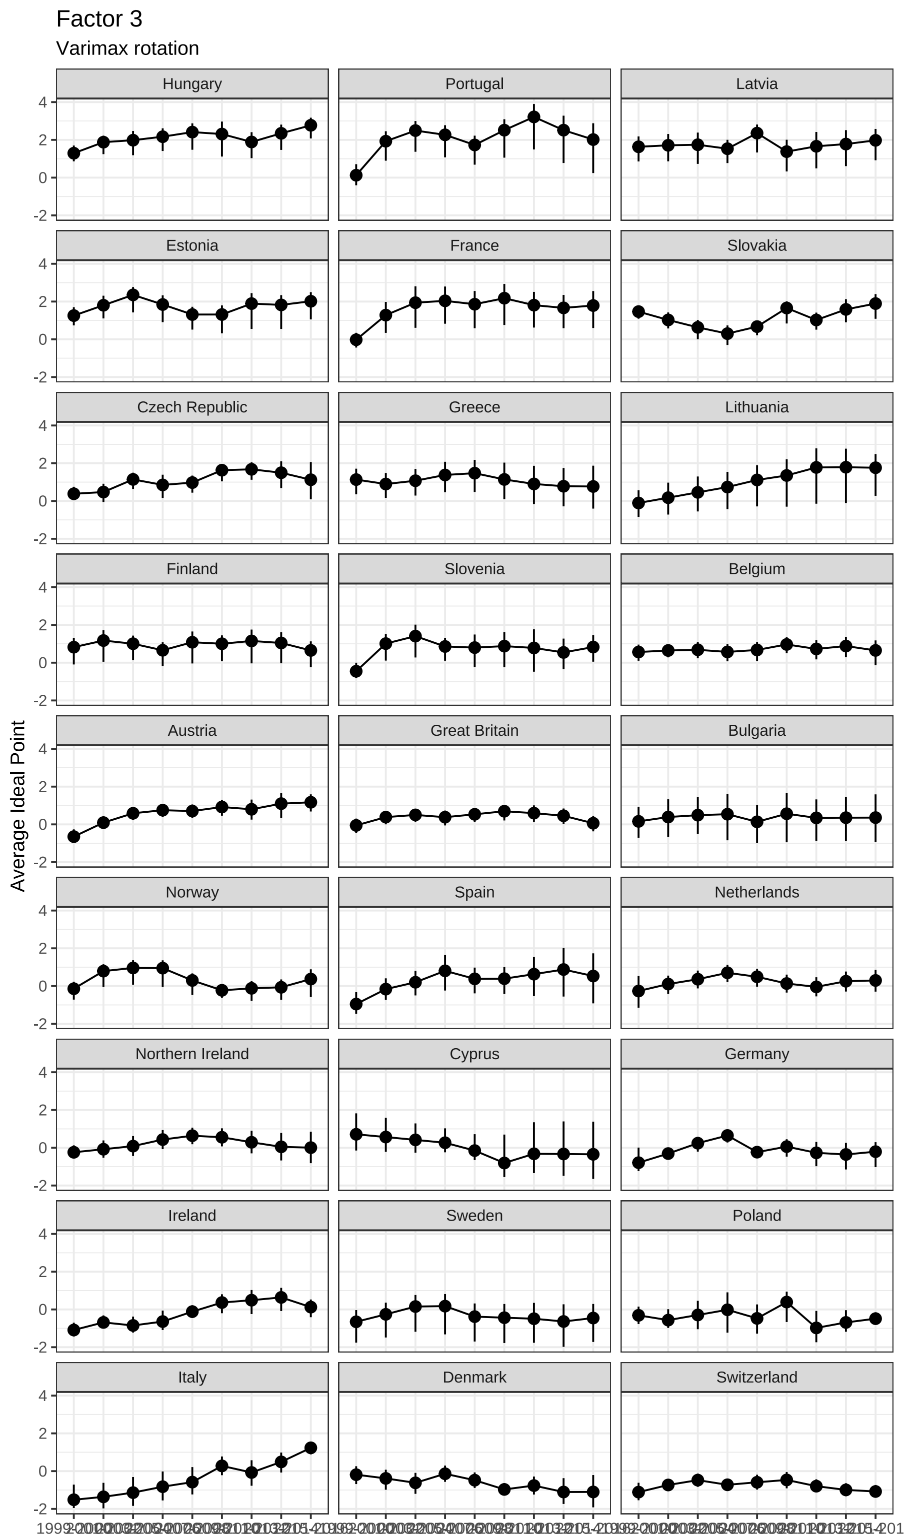

Figure 5: Country plots: Dimension 3

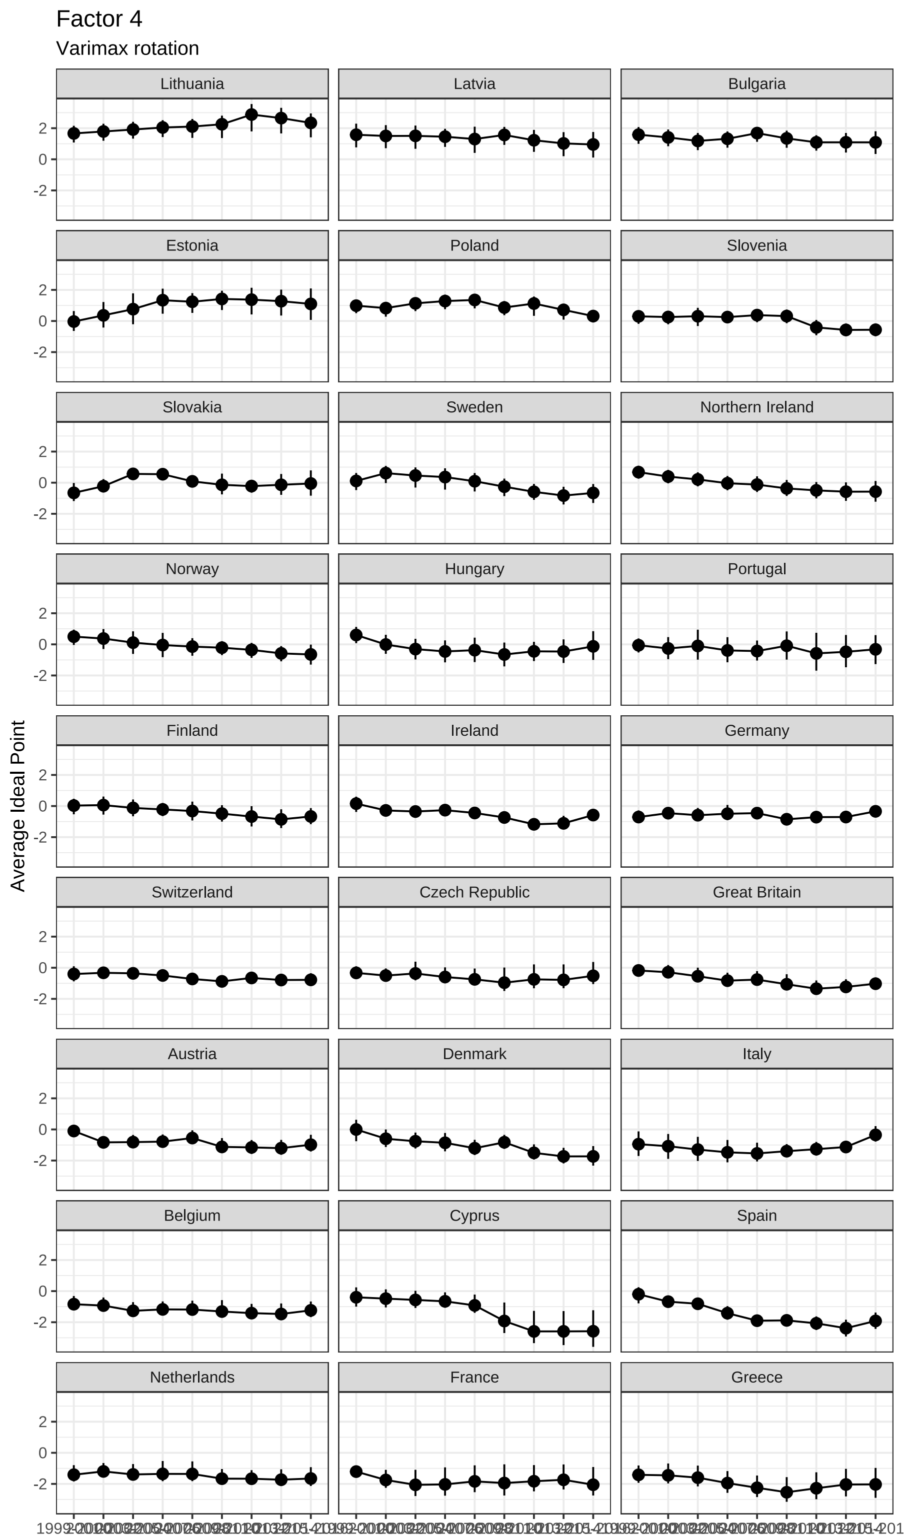

Figure 6: Country plots: Dimension 4

5.4 Covariance plots

Figure 7 shows estimated within-group covariance across the four dimensions.

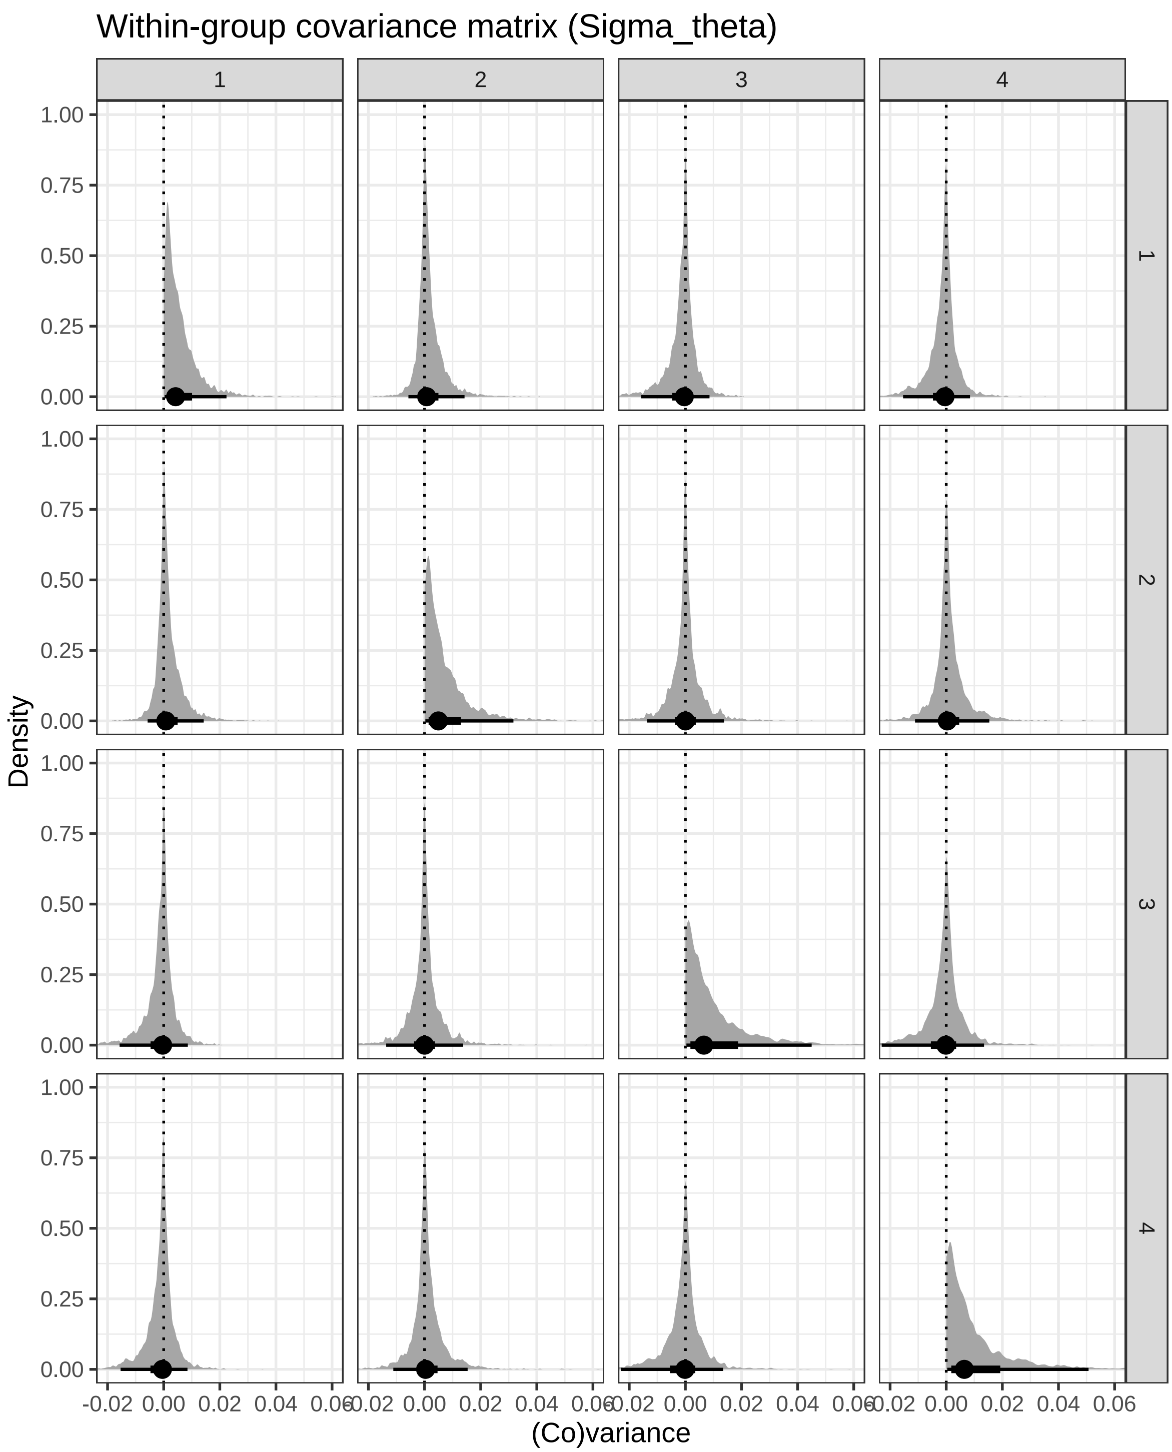

Figure 7: Covariance Plots

6 Multidimensional ideology in Spain

In this application, we use the MODGIRT model to distinguish two dimensions of political preferences in Spain: the classic cleavage over socioeconomic issues and an orthogonal cleavage related to the power and structure of the state. MODGIRT is especially useful in this setting because national surveys tend not to include enough items for subject-level multidimensional scaling, and international surveys do not have enough respondents to make accurate subnational inferences. MODGIRT permits the inclusion of many more items included on different surveys fielded at different points in time. It therefore allows us to track ideological trends in both dimensions over two decades.

Scholars consistently describe Spanish politics as at least a two dimensional space, where one dimension consists of centralization preferences rather than traditional economic and cultural issues (Elias, Szöcsik, and Zuber 2015). The centralization dimension is especially key to politics in regions of Spain characterized by more significant significant center-periphery cleavages, such as the Basque Country and Catalonia (Fernández-Albertos 2002; Balcells i Ventura 2007). Centralization preferences in Spain complicate not only party politics, but also individuals' own self-descriptions. Regional attachments lead individuals to report more progressive left-right self-placements than their responses to more traditional ideological items would anticipate (Dinas 2012; Galais and Serrano 2020).[5]

How then can researchers parsimoniously assess claims regarding ideological differences between regions or ideological change over time? Given the difficulty of multidimensional scaling, past research has turned to regional manifestos and spatial models that describe how voters position parties relative to themselves (Alonso, Cabeza, and Gómez 2015; Rivero 2015). Other scholarship looks at how signals of ideological and national positions shape how voters evaluate candidates (Liñeira, Muñoz, and Rico 2021). Yet all of these solutions focus on the relationship between parties and voters. They rely on parties to define the political space, rather than voters themselves. MODGIRT takes a more direct approach based on Spaniards' actual expressed preferences.

6.1 Estimation and identification

Based on our interest in distinguishing centralization preferences from orthogonal dimensions, we estimated a two-dimensional model, with item responses dichotomized at their midpoints for efficiency.[6] We defined groups as the cross-classification of region, gender, age (16–37, 38–60, 61+) and education (post-secondary or not). The survey data used to estimate the model

[5]Less has been written about the role of centralization preferences in the other historic nationalities of Navarre and Galicia. As Section 7.1 shows, Navarrese respondents look similar to Basques and Catalans on left-right self-placement, but Galicians report positions closer to the Spanish average. We therefore expect Navarre but not Galicia to behave similarly to the Basque Country and Catalonia.

[6]See Section 7.3 for evidence that two dimensions provide a reasonable fit.

include over 100,000 responses spread over surveys from 2000–2020.[7] We allowed group ideal points to vary by biennia, starting with 2000-2001 and ending with 2018-2020. As before, we used the varimax-RSP algorithm to initially identify the model. Because scholars generally agree that centralization defines one dimension of Spanish politics, we rotated the parameter draws such that a item regarding preferences for regional autonomy, *reg auto*, loaded only on the first dimension.

[7]See Section 7.2 for descriptions of the component surveys and a breakdown of responses by region.

6.2 Interpretation of the latent dimensions

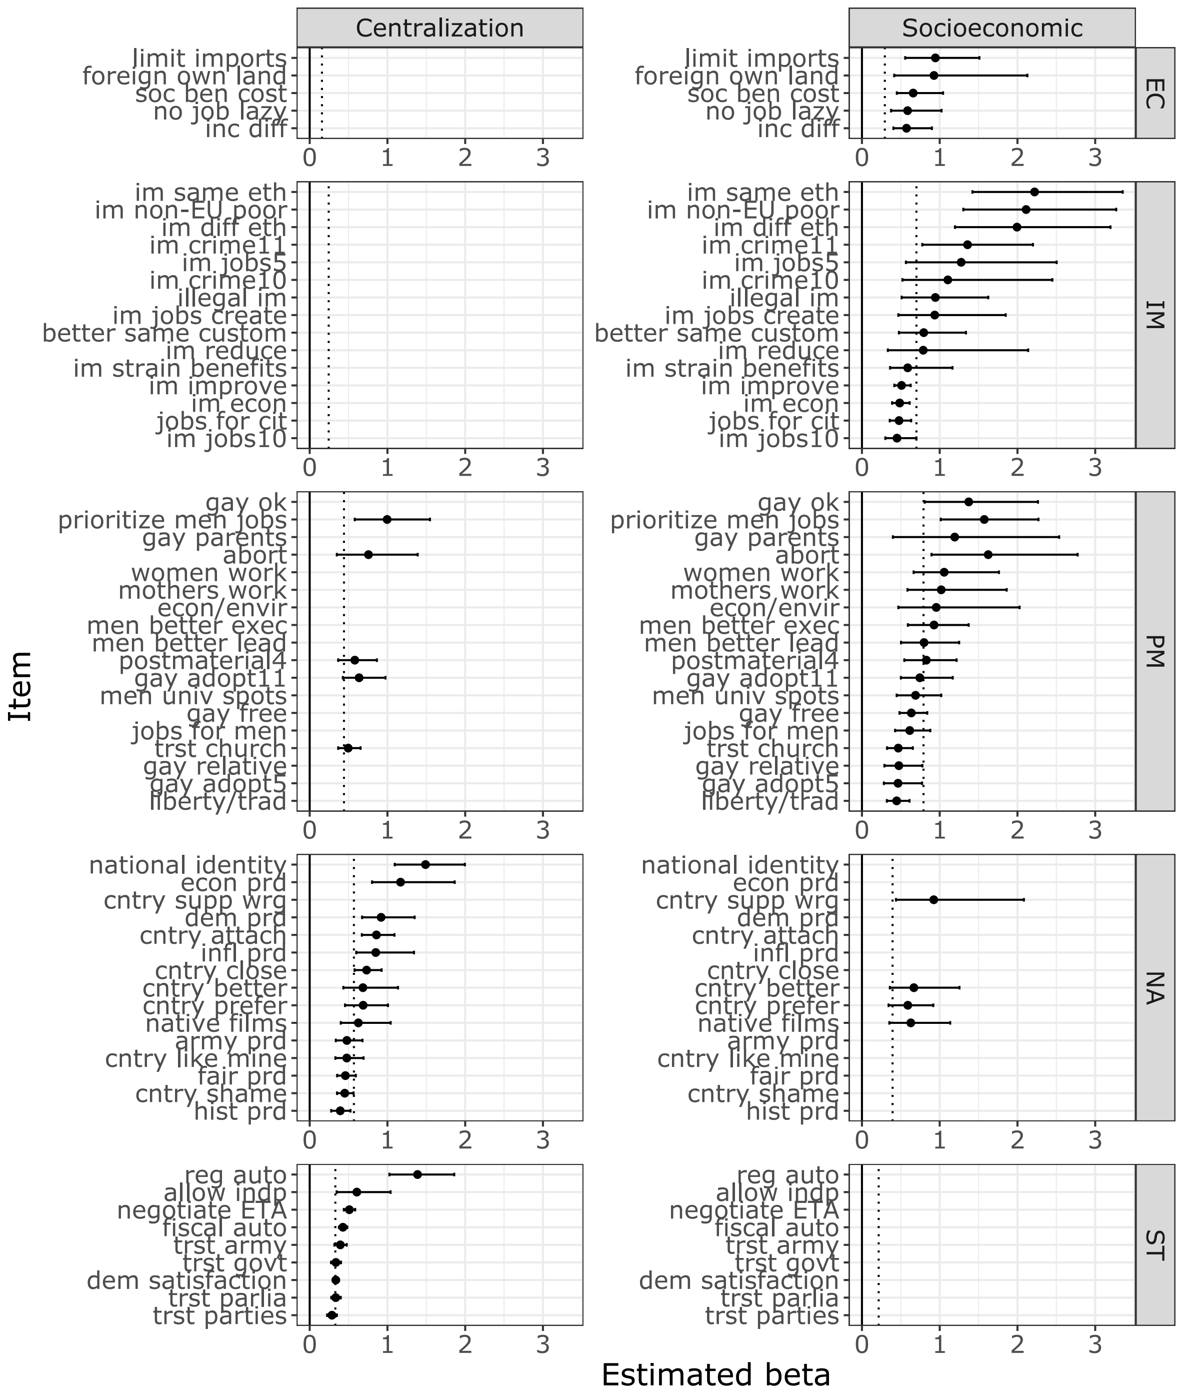

Figure 8: Loadings on each dimension of consistently highly informative items by issue domain.

To better understand what types of items load onto each dimension, we first classified every item into one of five issue domains: Economics (EC), Immigration (IM), Nationalism (NA), Postmaterialism (PM), and Statism (ST), as well as a separate category for left-right self-placement (LR). Because the full dataset includes 133 items, Figure 8 shows the estimated loadings on each dimension for a subset of highly informative items (points), as well as the mean loading across all items on each issue domain (vertical lines). We define items as highly informative if their estimates are above the dimensional median on at least 99% of draws.

Highly informative items for the first dimension come primarily from the "nationalism", "statism", and "post-material" domains. In addition to the rotation target *reg auto*, which is in the statism domain, the most informative item on the first dimension is a nationalism item that asks whether individuals identify more with Spain or their Autonomous Community. The post-material items that are highly informative for the first dimension include trust in the church and support for abortion, though these items load on the second dimension as well. We label this dimension "centralization" ideology, since it is based primarily on preferences for the structure of the state. But position on the first dimension also reflects trust in state institutions and other traditional power-holders.

What distinguishes the second dimension from the first is the greater influence of items from the economic and immigration domains, and the reduced role of statism items. No items from the economic or immigration domains have consistently high loadings on the first dimension, but many of these items are highly informative for the second dimension. These include the idea that social benefits cost too much and beliefs about how immigration affects crime. Because this dimension combines concerns related to economics, immigration, and cultural issues, we interpret it as reflecting latent socioeconomic ideology.[8]

6.3 Convergent validation of latent measures

To demonstrate the validity of our measures of centralization and socioeconomic ideology (Adcock and Collier 2001), we examine how closely they are associated with individual questions in the same issue area. We selected one question from each of the substantive domains described above (Economic, Immigration, Postmaterial, Nationalism, Statism, Self-identification). In each domain, we picked the question that had the highest number of responses in order to be able to produce the most accurate possible estimates of sub-group opinions.[9] These items were then removed from the dataset and were not used in fitting the model. We expected the questions from the nationalism, statism, and self-placement groups to be closely associated with the latent centralization dimension, and the questions from the economic, immigration, and postmaterial domains to be associated with the socioeconomic domain.

[8]See Section 7.4 for the predictors of "conservativism" on both dimensions.

[9]Sub-groups here are defined by the intersection of the age, gender, and education categories used in the model, as well as two groupings of regions: one for the Basque Country, Catalonia and Navarre, and another for all other Autonomous Communities. We did not use region or period to ensure a sufficient number of respondents in each sub-group.

Figure 9 shows how each of the six questions correlates with sub-groups' latent scores on both dimensions. As expected, "pride in nationality" (Nationalism), "trust in parliament" (Statism) and "right-wing self-placement" (Self-placement) correlate very strongly with centralization ideology, but have little correlation with socioeconomic ideology. Of the remaining three questions, "tradition unimportant" (Postmaterial) and "services/taxes tradeoff" (Economic) are weakly correlated with centralization, while "immigrants enrich culture" (Immigration) is not correlated with centralization at all. However, all of these items do correlate with socioeconomic ideology, though the relationship is not quite as strong in the case of "tradition unimportant." This analysis not only shows that scores on our latent dimensions explain variation between groups on these held-out questions in a manner consistent with our interpretation of the two dimensions, it also reinforces how self-reported ideology in Spain reflects centralization preferences more than socioeconomic ones.

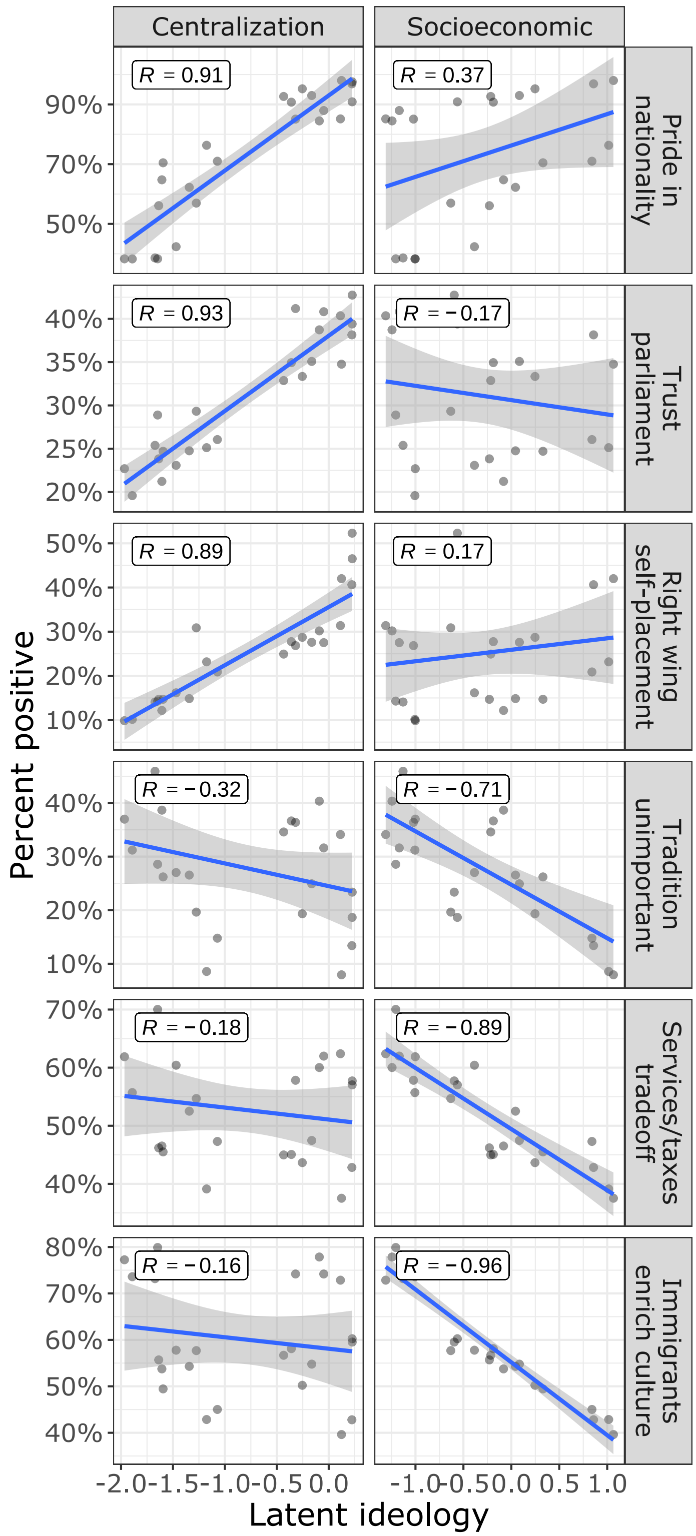

Figure 9: Convergent validation examines correlations between latent estimates and held-out questions from each issue domain.

6.4 Dynamic latent ideology

Once we interpret the first dimension as latent centralization ideology and the second dimension as latent socioeconomic ideology, we can use the estimated averages for cross-sectional comparisons. In Figure 10 we use the MODGIRT estimates to examine average post-stratified regional latent ideology in each time period.

A few important time trends stand out in Figure 10, which reflect salient events in Spanish politics. The upper panel reveals marked decline in support for the central state between 2010–11 and 2012–2013, which is likely associated with the 15-M movement in Spain (also known as the *Indignados* movement). This was a series of mass anti-austerity protests beginning in May 2011, which contributed to the founding of the left-wing populist party *Podemos* in January 2014. While opposition to the government's economic reforms were the immediate cause of the 15-M protests, they were defined above all by their populist rejection of the Spanish party system and traditional political elites. For this reason, it is unsurprising that the rise of 15-M is associated with a rejection of the power of the central state. Interestingly, this downturn does not seem as pronounced in the Basque Country, Catalonia and Navarre, which start out relatively unsupportive of centralized power and become gradually less supportive over time. In other regions, however, support for centralization subsequently recovers to its earlier levels.

The lower panel of Figure 10 shows trends in socioeconomic ideology. Here a downturn in conservatism is clearly associated with the onset of the global financial crisis in 2008. On this dimension, conservatism begins trending downward at the time of the crisis and continues downward in the years after. This is consistent with arguments that recent rise of the far-right party Vox is more associated with trends on the centralization dimension than on the socioeconomic dimension. The ability of our latent measures to capture known dynamics of Spanish politics offers construct validation for these measures.

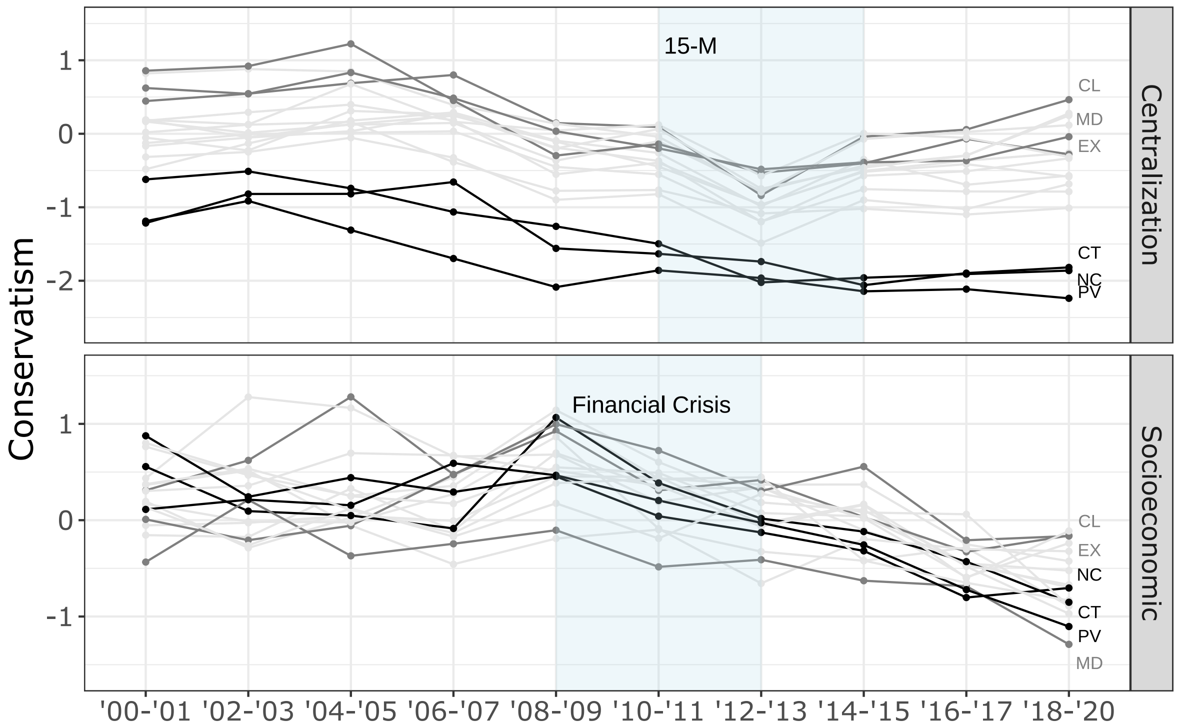

Figure 10: Time trends in average post-stratified latent regional ideology 2000-2020. A few regions are highlighted for comparison purposes: the historical nationalities of the Basque Country (BC), Catalonia (CT), and Navarre (NC), as well as Madrid (MD), Castile-Leon (CL), and Extremadura (EX).

The results presented here illustrate how the MODGIRT model enables comprehensive comparison across groups in multidimensional contexts. Taking a group-level approach to sparse Spanish survey data allows us to untangle the two dimensions of Spanish politics and examine how each has changed over time. Using our method, future analyses of the multidimensional structure of Spanish preferences can delve deeper into how cross-regional differences have evolved over time, and how they relate to changes in the Spanish party system.

7 Spanish public opinion supplemental information

7.1 Left-right self-placement in Spain

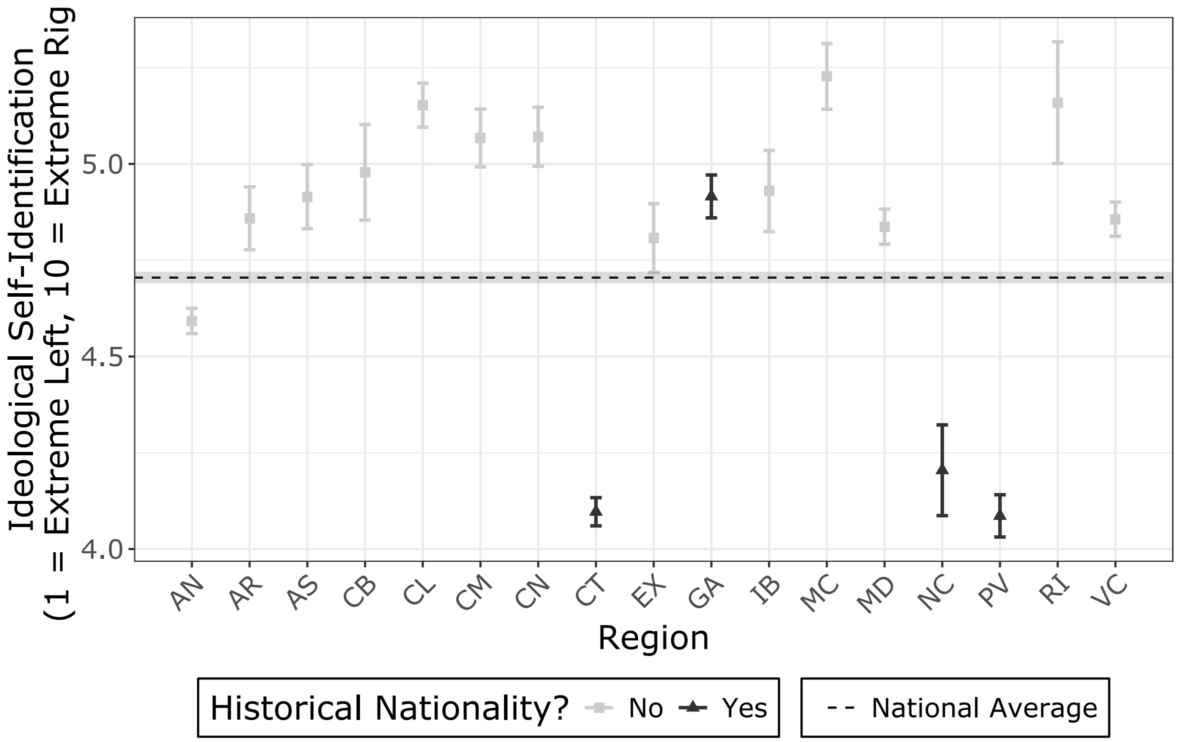

Figure 11: Average left-right self-placement in the Basque Country, Catalonia, and Navarre is significantly left of the Spanish average. Data from Spanish fiscal and postelection surveys 2000-2016

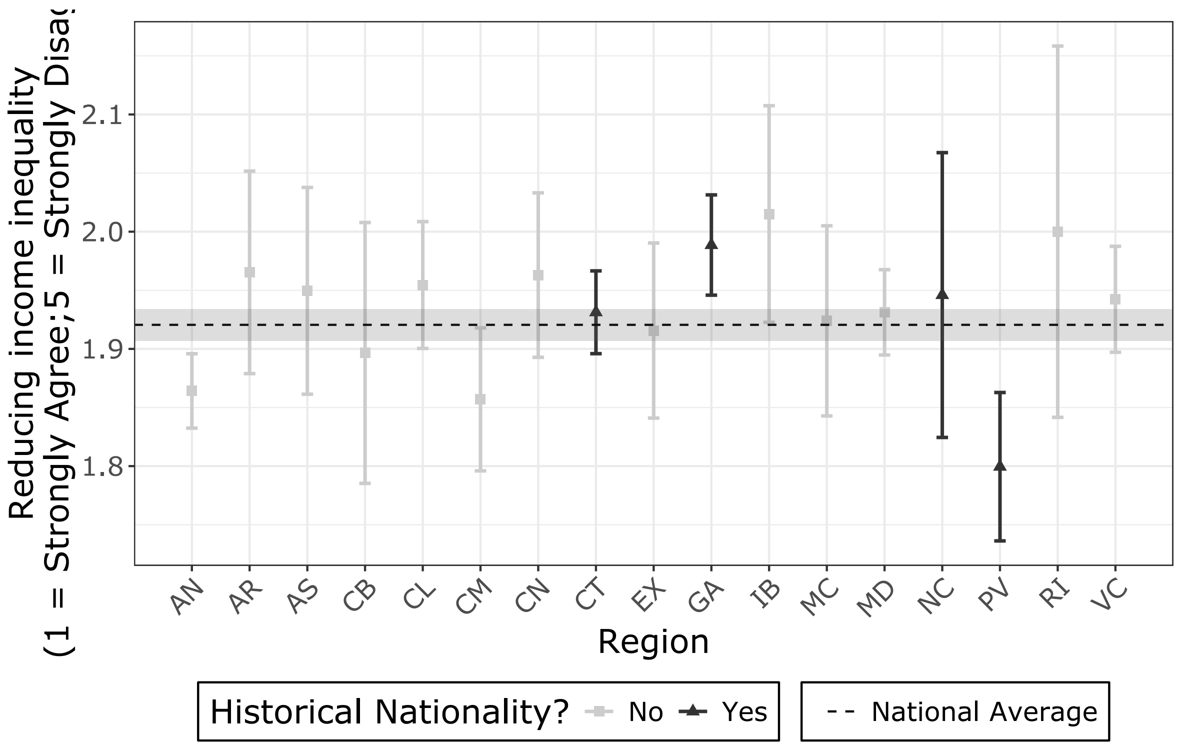

Figure 12: Average responses to questions about economic issues in the Basque Country, Catalonia, and Navarre are not significantly left of the Spanish average. Data from ESS 2002-2020.

7.2 Component surveys for Spanish model

We combine responses from Spanish post-election and fiscal surveys administered by the Centro de Investigaciones Sociológicas (Fiscal surveys annually 2000–2017; Post-election surveys 2008–2016), as well as Spanish subsets of the European Social Survey (2002–2019), International Social Survey Program National Identity Modules (2003–2013), European Values Survey (2008–2017), and World Values Survey (2000 – 2011). The full dataset used for estimation covers 93,676 responses spread over surveys from 2000–2020. A breakdown of responses by survey and Spanish region appears in Table 6.

Table 6: Spanish survey data by region and data source

	Survey						
CCAA	CIS Fiscal	CIS Postelection	ESS	EVS	ISSP	WVS	Total
Andalusia	7,696	4,885	3,168	1,165	639	1,105	18,658
Aragon	1,337	890	507	196	112	188	3,230

Asturias	1,154	642	445	171	96	187	2,695
Balearic Islands	868	609	344	137	73	126	2,157
Basque Country	2,259	1,415	800	336	184	351	5,345
Canary Islands	1,761	1,254	633	275	159	263	4,345
Cantabria	585	427	222	87	51	85	1,457
Castile-La Mancha	1,926	1,308	773	289	147	248	4,691
Castile-Leon	2,663	1,738	1,015	404	242	396	6,458
Catalonia	7,035	3,054	2,540	1,024	575	983	15,211
Extremadura	1,126	744	435	175	88	158	2,726
Galicia	2,958	1,955	1,445	445	242	438	7,483
Madrid	5,785	2,106	2,270	849	489	829	12,328
Murcia	1,234	732	469	190	95	174	2,894
Navarre	591	407	231	87	48	84	1,448
Rioja	297	204	103	43	34	43	724
Valencia	4,656	2,163	1,607	669	383	659	10,137
Total	43,931	24,533	17,007	6,542	3,657	6,317	101,987

7.3 Dimensionality

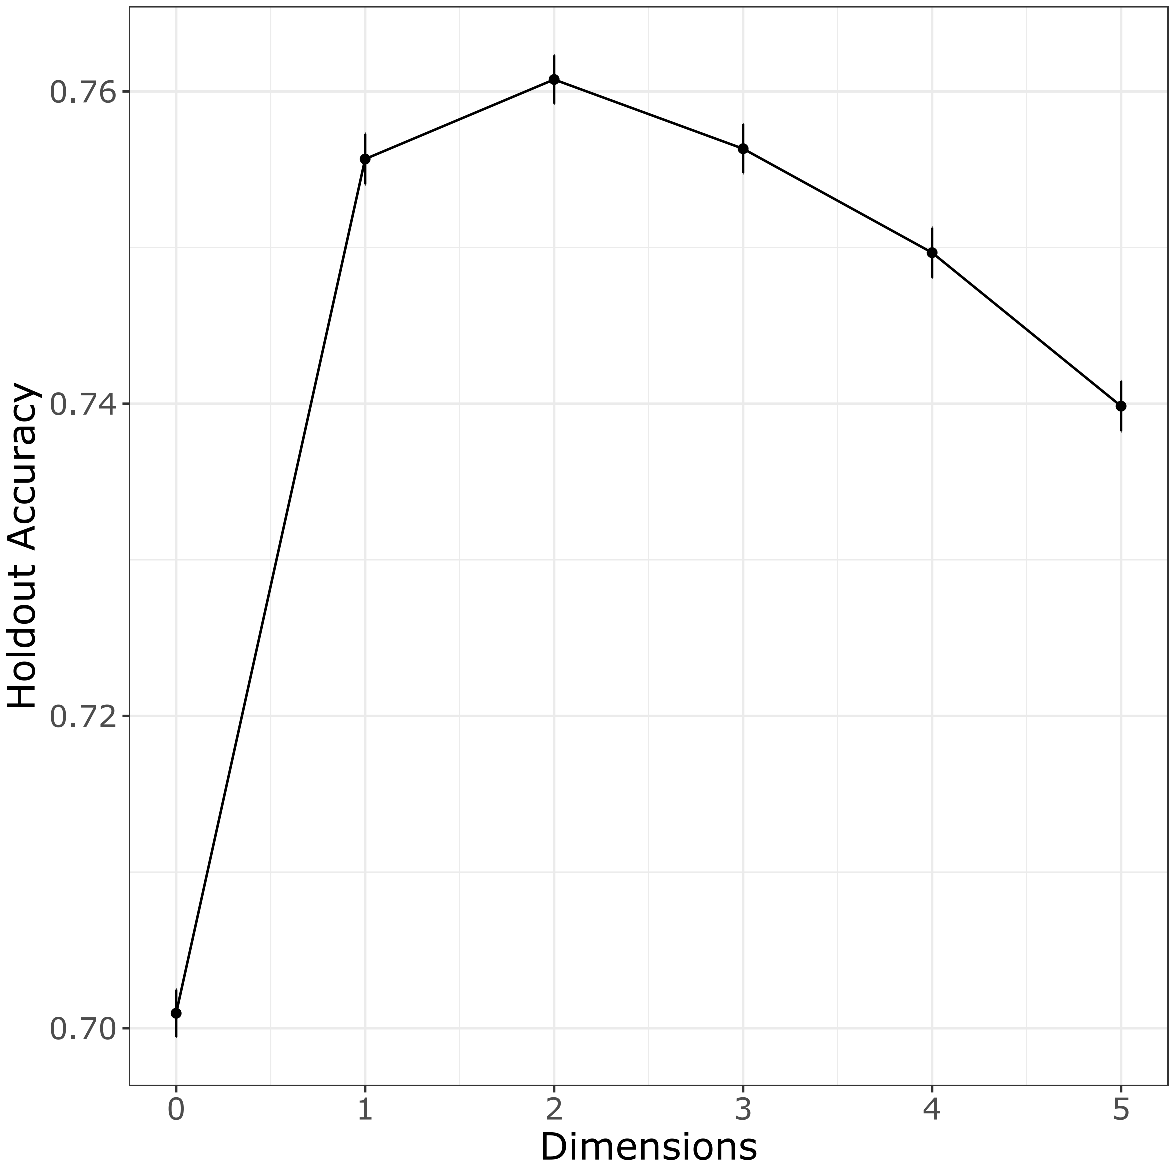

Figure 13: Holdout accuracy by dimension for models for Spanish opinion data shows that two dimensions improve over one dimension, and adding a third dimension does not much improve over two. Accuracy estimated using methods and replication code from Marble and Tyler (2022).

7.4 Predictors of Latent Conservatism

Examining conditional associations with "conservativism" on both dimensions confirms this interpretation of the dimensions. We determine conditional association by regressing conservativism in each dimension on the groups used in the model (region, gender, education, age group and period) within each iteration, and then averaging over iterations. Consistent with known trends, Figure 14 indicates that being older has a positive conditional association with conservatism on the socioeconomic second dimension, while being college educated has a negative conditional association with conservativism on this dimension. Age again has a positive conditional association with conservatism on the first centralization dimension that reflects concerns with the role of the state, while hailing from a historic nationality (the Basque Country, Navarre, Catalonia and Galicia) has the expected strong negative association with centralization conservativism.

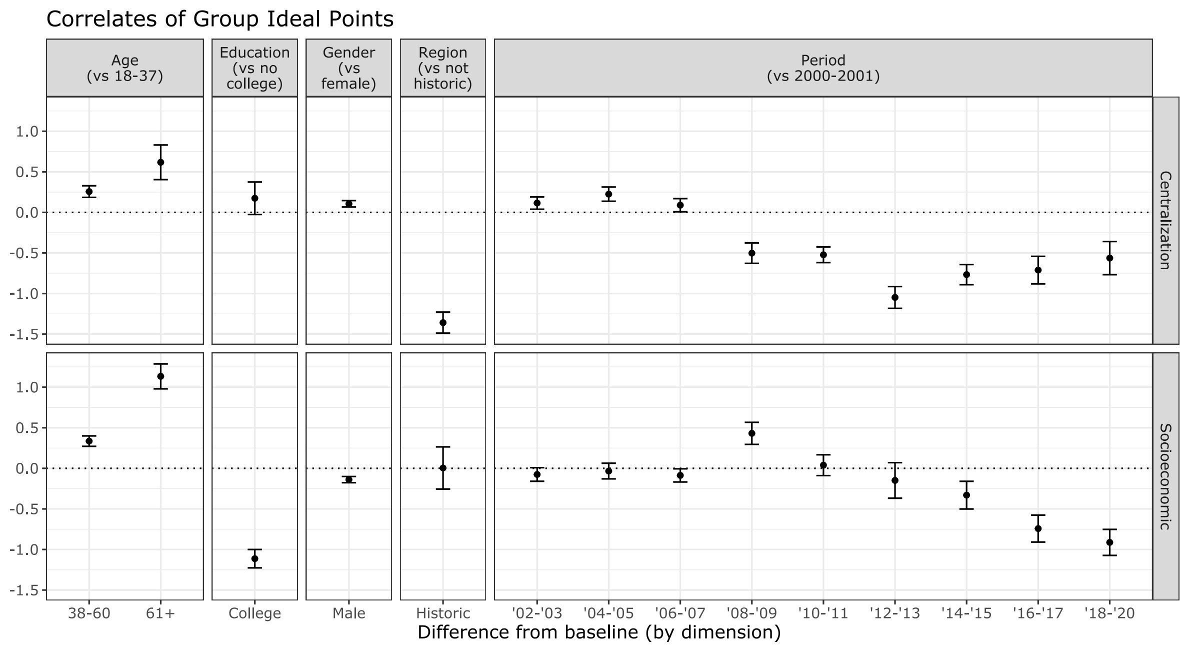

Figure 14: Predictors of conservativism

8 Data availability

Replication code for this article is available at https://doi.org/10.7910/DVN/UUPSCM (Berwick and Caughey 2024).

An R package for implementing the MODGIRT model, including the pre-processing and postprocessing steps needed to prepare the data and analyze the posterior samples, is available at https://github.com/devincaughey/dbmm (Berwick, Caughey, and Sasaki 2024).

Replication datasets analyzed in this article from Cavaillé and Trump (2015) are available at https://doi.org/10.7910/DVN/L4MGG5 (Cavaille 2018). The replication datasets were supplemented by re-analysis of the original data from the British Social Attitudes Survey (NatCen Social Research 2024).

Replication datasets from Caughey, O'Grady, and Warshaw (2019a) are available at https://doi.org/10.7910/DVN/H9XGEB (Caughey, O'Grady, and Warshaw 2019b).

References

Adcock, R., and D. Collier. 2001. "Measurement Validity: A Shared Standard for Qualitative and Quantitative Research." *American Political Science Review* 95 (3): 529–46. doi:10.1017/S0003055401003100.

Alonso, S., L. Cabeza, and B. Gómez. 2015. "Parties' Electoral Strategies in a Two-Dimensional Political Space: Evidence from Spain and Great Britain." *Party Politics* 21 (6): 851–65. doi:10.1177/1354068815597576.

Balcells i Ventura, L. 2007. "¿Es El Voto Nacionalista Un Voto de Proximidad o Un Voto de Compensación? Una Nueva Aproximación 'Espacial' Al Voto En Dos Dimensiones." *Revista Española de Ciencia Política*, no. 16: 6188. https://recyt.fecyt.es/index.php/recp/article/view/37440.

Berwick, E., and D. Caughey. 2024. "Replication Data for: Berwick and Caughey, 'MODGIRT: Multidimensional Dynamic Scaling of Aggregate Survey Data'." Harvard Dataverse. doi:10.7910/DVN/UUPSCM.

Berwick, E., D. Caughey, and T. Sasaki. 2024. "Dbmm: Dynamic Bayesian Measurement Models." *GitHub Repository*. https://github.com/devincaughey/dbmm/tree/modgirt; GitHub.

Caughey, D., T. O'Grady, and C. Warshaw. 2019a. "Policy Ideology in European Mass Publics, 1981–2016." *American Political Science Review* 113 (3): 674–93.

———. 2019b. "Replication Data for: Policy Ideology in European Mass Publics, 1981–2016." Harvard Dataverse. doi:10.7910/DVN/H9XGEB.

Cavaille, C. 2018. "Replication Data for The Two Facets of Social Policy Preferences." Harvard Dataverse. doi:10.7910/DVN/L4MGG5.

Cavaillé, C., and K.-S. Trump. 2015. "The Two Facets of Social Policy Preferences." *Journal of Politics* 77 (1): 146–60.

Dinas, E. 2012. "Left and Right in the Basque Country and Catalonia: The Meaning of Ideology in a Nationalist Context." *South European Society and Politics* 17 (3): 467485.

Elias, A., E. Szöcsik, and C.I. Zuber. 2015. "Position, Selective Emphasis and Framing: How Parties Deal with a Second Dimension in Competition." *Party Politics* 21 (6): 839–50. doi:10.1177/1354068815597572.

Fernández-Albertos, J. 2002. “Votar En Dos Dimensiones: El Precio Del Nacionalismo y La Ideología En El Comportamiento Electoral Vasco, 1993-2001.” *Revista Española de Ciencia Política* 6: 153181.

Galais, C., and I. Serrano. 2020. “The Effects of Regional Attachment on Ideological Self-Placement: A Comparative Approach.” *Comparative European Politics* 18 (4): 487–509. doi:10.1057/s41295-019-00196-z.

Liñeira, R., J. Muñoz, and G. Rico. 2021. “Inferring Party Positions Across Issue Dimensions.” *Party Politics* 27 (5): 1031–43. doi:10.1177/1354068820912653.

Marble, W., and M. Tyler. 2022. “The Structure of Political Choices: Distinguishing Between Constraint and Multidimensionality.” *Political Analysis* 30 (3): 328–45.

NatCen Social Research. 2024. “British Social Attitudes Survey.” UK Data Service. doi:10.5255/UKDA-Series-200006.

Rivero, G. 2015. “Heterogeneous Preferences in Multidimensional Spatial Voting Models: Ideology and Nationalism in Spain.” *Electoral Studies* 40 (December): 136–45. doi:10.1016/j.electstud.2015.06.002.

Samejima, F. 1997. “Graded Response Model.” In *Handbook of Modern Item Response Theory*, edited by W.J. van der Linden and R.K. Hambleton, 85–100. New York: Springer.