Exposition of the Dynamic Group-level Item Response Theory (DGIRT) Model

Devin Caughey

June 28, 2016

The dynamic group-level item response theory (DGIRT) model is a Bayesian method for estimating subpopulation groups' average conservatism (or other trait) from individuals' responses to dichotomous questions. It is "dynamic" both in the sense that groups are allowed to evolve over time and in the sense that the model "borrows strength" from other time periods, to a degree specified by the user. The DGIRT model is a modified version the hierarchical group-level IRT model described by Caughey and Warshaw (2015). It differs from the latter mainly in that DGIRT models groups' evolution directly though a dynamic linear model rather than indirectly through a hierarchical model. The model is implemented in the R package **dgirt** (Dunham, Caughey, and Warshaw 2016).

Let θ_i denote individual i's score on some latent trait, which for the sake of exposition we will call "conservatism." Let $y_{iq} \in \{0, 1\}$ indicate i's dichotomous response to question q, where $y_{iq} = 1$ indicates the conservative response option.[1] Under the assumptions of the standard one-dimensional probit IRT model (e.g., Fox 2010), individual i's response to question q is a probabilistic function of i's conservatism θ_i as well as the question's "difficulty" α_q, which captures the base level of support for the question, and its "difficulty" β_q, which represents the strength of its relationship with conservatism. Specifically, $y_{iq} = 1$ if

$$\beta_q \theta_i - \alpha_q + \epsilon_{iq} > 0 \tag{1}$$

1. An ordinal question with $L > 2$ levels can be handled by transforming it into a set of $L - 1$ dichotomous variables indicating whether i's response is above the $L - 1$ lowest levels.

and 0 otherwise, where $\epsilon_{iq} \overset{iid}{\sim} \mathrm{N}(0,1)$. Individual i's probability $\pi_{iq} = \Pr(y_{iq} = 1)$ of responding conservatively to question q is thus

$$\pi_{iq} = \Phi(\beta_q \theta_i - \alpha_q), \tag{2}$$

where the normal distribution function Φ maps $\beta_q \theta_i - \alpha_q$ to the unit interval.[2] If every individual answers many questions, then we can use a standard IRT model to infer their conservatism from their question responses via the sampling model

$$y_{iq} \overset{iid}{\sim} \mathrm{Bernoulli}(\pi_{iq}). \tag{3}$$

Even if this condition is not satisfied—that is, if respondents answer as few as one question each—it is often possible to draw inferences about *average* conservatism in different subpopulation groups by aggregating the responses of different individuals to different questions. To do so, it is helpful to reparamaterize the individual-level IRT model as

$$\pi_{iq} = \Phi\left(\frac{\theta_i - \kappa_q}{\sigma_q}\right), \tag{4}$$

where $\kappa_q = \alpha_q/\beta_q$ and $\sigma_q = 1/\beta_q$. Let g index groups and let $\bar{\theta}_g$ denote average conservatism in group g. Under the assumption that conservatism has an iid normal and homoskedastic distribution within groups—i.e., $\theta_{i[g]} \overset{iid}{\sim} \mathrm{N}(\bar{\theta}_g, \sigma_\theta^2)$—the probability that a randomly sampled member of group g gives a conservative response to question q is

$$\pi_{gq} = \Phi\left(\frac{\bar{\theta}_g - \kappa_q}{\sqrt{\sigma_q^2 + \sigma_\theta^2}}\right), \tag{5}$$

where σ_θ is the standard deviation of θ_i within groups. We can then connect (5) to the data through the sampling model

$$s_{gq} \overset{iid}{\sim} \mathrm{Binomial}(n_{gq}, \pi_{gq}), \tag{6}$$

2. Note that the assumption $\epsilon_{iq} \overset{iid}{\sim} \mathrm{N}(0,1)$ is violated if y_{iq} is affected by other individual-level traits correlated with (but distinct from) conservatism (i.e., if the true model is multidimensional).

where n_{gq} is group g's total number of non-missing responses to question q and s_{gq} is the number of those responses that are conservative.[3] Together, (5) and (6) constitute a static group-level IRT model, which can be used to infer groups' average conservatism $\bar{\theta}_g$ (see Mislevy 1983).

We add dynamics to this model by allowing groups' conservatism (and optionally other parameters) to change between time periods and modeling their temporal evolution with a dynamic linear model (DLM). Specifically, we model $\bar{\theta}_{gt}$ as a function of its value in the previous period ($\bar{\theta}_{g,t-1}$), a year-specific intercept common to all groups (ξ_t), and a vector of observed group attributes ($\boldsymbol{x}_{g\cdot}$):

$$\bar{\theta}_{gt} \overset{iid}{\sim} \mathrm{N}(\bar{\theta}_{g,t-1}\delta_{\bar{\theta}t} + \xi_t + \boldsymbol{x}'_{g\cdot}\boldsymbol{\gamma}_t, \sigma^2_{\bar{\theta}t}). \tag{7}$$

The standard deviation $\sigma_{\bar{\theta}t}$ is estimated from the data and allowed to evolve across years, as is the within-group SD $\sigma_{\theta t}$. The posterior estimates of $\bar{\theta}_{gt}$ are a compromise between this prior and the likelihood implied by Equations (5) and (6). When a lot of survey data are available for a given year, the likelihood will be given most weight by the model. If no survey data are available at all, the prior acts as a predictive model that imputes $\bar{\theta}_{gt}$.

In addition to allowing groups to evolve over time, it is possible to allow the relationship between conservatism and question responses to evolve as well.[4] This is accomplished by allowing κ_q to vary by period and modeling its evolution indirectly using the local-level DLM

$$\alpha_{qt} \overset{iid}{\sim} \mathrm{N}(\alpha_{q,t-1}, \sigma^2_{\gamma}), \tag{8}$$

where transition variance σ^2_{γ} is estimated from the data. This "evolving item" version of the model holds constant each question's discrimination but allows the "difficulty" of a conservative response to vary between time periods. The other time-indexed parameters in (5) and (7) are allowed to evolve in similar

3. Following Ghitza and Gelman (2013) and Caughey and Warshaw (2015, 202–3), we adjust the raw values of s_{gq} and n_{gq} to account for survey weights and for respondents who answer multiple questions.

4. This may be desirable if individual question series exhibit idiosyncratic trends.

fashion:

$$\log(\sigma_{\theta t}) \overset{iid}{\sim} \mathrm{N}(\log(\sigma_{\theta t-1}), \sigma_\sigma^2) \tag{9}$$

$$\log(\sigma_{\bar{\theta} t}) \overset{iid}{\sim} \mathrm{N}(\log(\sigma_{\bar{\theta} t-1}), \sigma_\sigma^2) \tag{10}$$

$$\delta_{\bar{\theta} t} \overset{iid}{\sim} \mathrm{N}(\delta_{\bar{\theta} t-1}, \sigma_\delta^2) \tag{11}$$

$$\xi_{qt} \overset{iid}{\sim} \mathrm{N}(\xi_{q,t-1}, \sigma_\gamma^2) \tag{12}$$

$$\gamma_{pt} \overset{iid}{\sim} \mathrm{N}(\gamma_{p,t-1}\delta_{\gamma t} + \boldsymbol{z}'_{p\cdot t}\boldsymbol{\nu}_t, \sigma_\gamma^2). \tag{13}$$

All of the above are simple local-level DLMs except the last, which models the evolution of the hierarchical parameters γ_{pt} as a function of a vector of period-parameter-specific attributes $\boldsymbol{z}_{p\cdot t}$. This may be useful if γ_{pt} is an intercept for a geographic unit (e.g., a state) and one wishes borrow strength from observably similar units (e.g., states with a similar per-capita income). The parameters in the transition model for γ_{pt} themselves evolve over time as follows:

$$\delta_{\gamma t} \overset{iid}{\sim} \mathrm{N}(\delta_{\gamma t-1}, \sigma_\delta^2) \tag{14}$$

$$\nu_{ht} \overset{iid}{\sim} \mathrm{N}(\nu_{h,t-1}, \sigma_\delta^2). \tag{15}$$

To identify the location and scale of the model, in each iteration we transform the difficulty parameters so that their mean is 0 in the first period, and do likewise for the discrimination so that their product is 1:

$$\tilde{\alpha}_{qt} = \alpha_{qt} - Q^{-1}\sum_{q=1}^{Q}\alpha_{q1} \tag{16}$$

$$\tilde{\beta}_q = \beta_q\left(\prod_{q=1}^{Q}\beta_q\right)^{-1/Q}. \tag{17}$$

The transformed parameters $\tilde{\alpha}_{qt}$ and $\tilde{\beta}_q$ are then re-parameterized as κ_{qt} and σ_q, as explained above.

The priors for the first period and for temporally constant parameters are as follows:

$$\bar{\theta}_{g1} \overset{iid}{\sim} \mathrm{N}(\xi_1 + \boldsymbol{x}'_{g\cdot}\boldsymbol{\gamma}_1, \sigma^2_{\bar{\theta}1}) \tag{18}$$

$$\alpha_{q1} \overset{iid}{\sim} \mathrm{N}(0, 1) \tag{19}$$

$$\log(\beta_q) \overset{iid}{\sim} \mathrm{N}(0, 1) \tag{20}$$

$$|\sigma_{\theta 1}| \overset{iid}{\sim} \mathrm{Cauchy}(0, 2.5) \tag{21}$$

$$|\sigma_{\bar{\theta}1}| \overset{iid}{\sim} \mathrm{Cauchy}(0, 2.5) \tag{22}$$

$$\xi_{q1} \overset{iid}{\sim} \mathrm{N}(0, 10) \tag{23}$$

$$\gamma_{p1} \overset{iid}{\sim} \mathrm{N}(\boldsymbol{z}'_{p\cdot 1}\boldsymbol{\nu}_{h0}, \sigma^2_{\gamma}) \tag{24}$$

$$\delta_{\bar{\theta}1} \overset{iid}{\sim} \mathrm{N}(m_{\delta_{\bar{\theta}}}, s_{\delta_{\bar{\theta}}}) \tag{25}$$

$$\delta_{\gamma t} \overset{iid}{\sim} \mathrm{N}(0.5, 0.5) \tag{26}$$

$$\nu_{h1} \overset{iid}{\sim} \mathrm{N}(0, 10) \tag{27}$$

$$\nu_{h0} \overset{iid}{\sim} \mathrm{N}(0, 10) \tag{28}$$

$$|\sigma_{\delta}| \overset{iid}{\sim} \mathrm{Cauchy}(0, s_{\sigma_\delta}) \tag{29}$$

$$|\sigma_{\sigma}| \overset{iid}{\sim} \mathrm{Cauchy}(0, s_{\sigma_\sigma}) \tag{30}$$

$$|\sigma_{\gamma}| \overset{iid}{\sim} \mathrm{Cauchy}(0, 2.5), \tag{31}$$

where $m_{\delta_{\bar{\theta}}}$, $s_{\delta_{\bar{\theta}}}$, s_{σ_δ}, and s_{σ_σ} are specified by the analyst.

References

Caughey, Devin, and Christopher Warshaw. 2015. "Dynamic Estimation of Latent Opinion Using a Hierarchical Group-Level IRT Model." *Political Analysis* 23 (2): 197–211.

Dunham, James, Devin Caughey, and Christopher Warshaw. 2016. *dgirt: Dynamic Group-level IRT Models.* `R` package. `https://jamesdunham.github.io/dgirt/`.

Fox, Jean-Paul. 2010. *Bayesian Item Response Modeling: Theory and Applications.* Springer.

Ghitza, Yair, and Andrew Gelman. 2013. "Deep Interactions with MRP: Election Turnout and Voting Patterns Among Small Electoral Subgroups." *American Journal of Political Science* 57 (3): 762–776.

Mislevy, Robert J. 1983. "Item Response Models for Grouped Data." *Journal of Educational Statistics* 8 (4): 271–288.