Advance Access publication May 7, 2013

Political Analysis (2013) 21:314–333
doi:10.1093/pan/mpt001

Modeling Dynamic Preferences: A Bayesian Robust Dynamic Latent Ordered Probit Model

Daniel Stegmueller
Department of Government, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK
e-mail: mail@daniel-stegmueller.com

Edited by R. Michael Alvarez

Much politico-economic research on individuals' preferences is cross-sectional and does not model dynamic aspects of preference or attitude formation. I present a Bayesian dynamic panel model, which facilitates the analysis of repeated preferences using individual-level panel data. My model deals with three problems. First, I explicitly include feedback from previous preferences taking into account that available survey measures of preferences are categorical. Second, I model individuals' initial conditions when entering the panel as resulting from observed and unobserved individual attributes. Third, I capture unobserved individual preference heterogeneity both via standard parametric random effects and a robust alternative based on Bayesian nonparametric density estimation. I use this model to analyze the impact of income and wealth on preferences for government intervention using the British Household Panel Study from 1991 to 2007.

1 Introduction

Individuals' political and economic preferences typically exhibit patterns of both stability and change (e.g., Wlezien 1995). On the one hand, preferences are often very highly correlated over time. But, on the other hand, preferences can change in response to external events, such as income shocks, becoming unemployed, or experiencing an economic crisis. To capture the dynamics of preferences—their stability and their change—an appropriate modeling strategy involves the use of individual-level panel data and dynamic panel models, in which past preferences influence current preferences via a first-order Markov process. Panel data are increasingly being used in political science, in the form of both long-term household panels, such as the British Household Panel Survey, and election panels, such as the Cooperative Campaign Analysis Project. *Linear* dynamic panel models are also well known in political science (for an introduction, see Wawro 2002). However, the application of these models to modeling dynamic preferences is not straightforward.

Three central issues arise when modeling preference dynamics: categorical preference measures, endogenous initial observations, and individual heterogeneity. First, although political scientists conceive of preferences as continuous, available survey data on preferences are usually ordered categorically, often using rather coarse categories. The nonlinear nature of preference measures prohibits direct application of established linear dynamic panel models (e.g., Arellano and Bond 1991; Blundell and Bond 1998) and instead requires a dynamic model for categorical data for both the dependent variable and the feedback process. Second, because initial observations—an individual's preference stated when entering the panel—are endogenous to the preference formation process under study, one should explicitly model initial observations in nonlinear panel models

Author's note: I am indebted to Jeff Gill, Michael Malecki, Michael Becher, Tom Snijders, Simon Jackman, Thomas Gschwend, Vera Troeger, Adam Ziegfeld, Charlotte Cavaille, Jeroen Vermunt, and Martyn Plummer, as well as conference and seminar participants in Chicago, Berlin, Cologne, Oxford, and Tilburg, for helpful comments and criticisms. Equal thanks are due to my reviewers and the editors. Furthermore, I thank the Oxford Supercomputing Centre for resources and support. This paper was written while I was a Prize Research Fellow at Nuffield College, University of Oxford. Replication materials are available at http://hdl.handle.net/1902.1/20433. Supplementary materials for this article are available on the *Political Analysis* Web site.

© The Author 2013. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

(Heckman 1981b; Nerlove, Sevestre, and Balestra 2008). Third, unobserved individual heterogeneity must also be modeled explicitly in order to capture unobserved or unmeasured effects of individual characteristics such as motivation or ability. When modeling heterogeneity via Gaussian random effects—as is standard in virtually all hierarchical models in political science—inferences can be sensitive to this specific distributional assumption and should be checked using a more flexible model specification.[1] Standard fixed effects estimation strategies are unavailable due to the presence of a lagged dependent (endogenous) variable in the nonlinear model (see, e.g., Nickell 1981; Heckman 1981b; Arellano and Carrasco 2003).

I present a Bayesian robust latent dynamic ordered probit model specification, which tackles these three problems. First, it captures the categorical nature of survey-based preference measures by using an ordered probit specification, in which a continuous latent preference variable generates observed survey responses. Most existing categorical dynamic panel models specify the lagged dependent variable as categorical, which implies the unrealistic assumption that current continuous *preferences* are influenced by past categorical survey *responses*. In contrast, I specify feedback from previous preferences to current ones as also arising from latent preferences, thus appropriately distinguishing between continuous concept and categorical survey items. Second, I model initial conditions using a simultaneous equation specification, in which individuals' initial observations depend on observed covariates, background information such as parents' education, and unobserved individual-specific effects. Third, I present robust specifications for the distribution of unobserved heterogeneity. I specify hierarchical or multilevel models with both Gaussian and t-distributed random effects. To relax these parametric assumptions, I employ Bayesian nonparametric density estimation for flexible estimation of the random effects distribution using Dirichlet process priors (for recent applications of Bayesian nonparametrics in political science, see Imai, Lu, and Strauss 2008; Gill and Casella 2009; Grimmer 2010; Spirling and Quinn 2010).

The article proceeds as follows. In the next section, I set up the hierarchical latent dynamic panel model and discuss my treatment of initial conditions, the specification of priors, and possible model extensions. Next, I present robust random effects specifications using Dirichlet process priors. I illustrate the model by an example from the political economy of redistribution preferences—where studies are usually cross-sectional and ignore both unobserved heterogeneity and dynamics. I analyze the impact of income and wealth on preferences for government intervention using the British Household Panel Study from 1991 to 2007, which repeatedly measures individual preferences for nearly two thousand individuals. I discuss results arising from the model specification using standard Gaussian random effects and illustrate how to conduct robustness tests using the flexible Dirichlet process random effects model. The last section concludes the article.

2 Latent Dynamic Model

A dynamic analysis of individual behavior or preferences has three features not present in cross-sectional studies. First, individual preferences show a certain degree of persistence. Although cross-sectional studies provide a snapshot of individuals in time, modeling the dynamics of preferences using panel data provides an explicit model of how preferences change over time. A straightforward theoretical specification posits that preferences are persistent, which creates correlated observations within the same individual. In other words, "[. . .] preferences remain unchanged unless something happens to change them [. . .]" (Wlezien 1995, 989). Thus, a dynamic model of preferences should include a *persistence parameter* capturing this correlation.

[1] Dynamic panel models for ordinal data are not widely developed in political science. Theoretical work and applications exist in biostatistics, medicine, and finance (e.g., Lunn, Wakefield, and Racine-Poon 2001; Müller and Czado 2005; Hasegawa 2009; Varin and Czado 2010; Czado, Heyn, and Müller 2011), but are developed with long time series in mind, and are not concerned with initial conditions in short panels of individuals (note that the start of medical studies often does coincide with the start of the data-generating process). Pang (2010) presents a model for repeated categorical data using correlated residuals. However, extending the model to include dynamic feedback is not straightforward due to the special status of initial conditions (cf. supplementary A). Pudney (2006, 2008) presents a model for dynamic ordinal data using Gaussian random effects in a maximum likelihood framework.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Second, some individual characteristics, such as intelligence or motivation, can have a strong influence on preferences or attitudes, but are unobserved or unobservable to the researcher. This *individual heterogeneity* is captured via individual constants, which I specify as random effects (I discuss robustness of distributional assumptions in Section 3). It is well known that if heterogeneity is present in the true data-generating process but ignored in the estimated model, the degree of preference persistence will be overestimated (see Heckman 1981a). Conversely, ignoring persistence leads researchers to overstate the extent of heterogeneity. Thus, a completely specified model of dynamic preferences has to include both components.[2]

Third, a sample of individuals, be it cross-sectional or a panel, provides only a time-limited observation window. Individuals start forming their beliefs and preferences a long time before one starts observing them. The fact that individuals do not enter a study with an "empty mind," that is, the problem of initial conditions, has to be included in the model. Those three features are important when interpreting the effect of shocks, such as becoming unemployed, on preferences. Estimating the effect of such shocks from cross-sectional data, ignoring preference persistence as well as individual heterogeneity, might lead to erroneous conclusions.

2.1 *Modeling Dynamics*

Concepts such as preferences and attitudes are not inherently discrete. The fact that one works with categorical variables is usually simply due to methodological limitations in data collection and measurement (McKelvey and Zavoina 1975). Consequently, preferences should be specified as a *latent variable* z_t, which represents the underlying continuous concept that generates observed categorical scores y_t (e.g., Greene and Hensher 2010). Since from the conceptual perspective of preferences there is no reason to expect that current *continuous* preferences depend on past preference *categories*, we also need the latent variable to appear on the right-hand side of our dynamic panel model (Heckman 1978; Müller and Czado 2005; Pudney 2008). In other words, feedback from past preferences to current ones should be specified as arising from z_{t-1}, not y_{t-1}.[3]

Thus, following Albert and Chib (1993), I model observed responses in category c ($c = 1, \ldots, C$) of observed variable y_{it} ($i = 1, \ldots, N$; $t = 0, \ldots, T$) as being generated by an underlying continuous latent variable z_{it} and a vector of threshold parameters τ such that

$$y_{it} = c \text{ if } z_{it} \in (\tau_{c-1}, \tau_c]. \tag{1}$$

To capture the ordinal nature of observed preference scores, threshold parameters are constrained to be monotonically increasing,

$$-\infty = \tau_0 < \tau_1 = 0 < \tau_2 < \ldots < \tau_{C-1} < \tau_C = \infty; \tag{2}$$

and $\tau_1 = 0$ to identify the model (assuming that an overall constant will be included in the model; see Albert and Chib 1993; Johnson and Albert 1999).

Now, the dynamic model for latent preferences z_{it} can be written as

$$z_{it} = \phi z_{it-1} + \beta' \boldsymbol{x}_{it} + \xi_i + \epsilon_{it},\ t = 1, \ldots, T, \tag{3}$$

where ϕ captures the degree of preference persistence, that is, the extent to which current preferences depend on previous ones. β is a vector of regression parameters for matrix $\boldsymbol{x}_{it}$ of possibly time-varying covariates and an overall constant. Errors are decomposed into an individual-specific time-constant random effect ξ_i and stochastic disturbances ϵ_{it}, which vary over individuals and

[2] The importance of distinguishing persistence (or state dependence) and heterogeneity has been well established in economics (e.g., Heckman 1981a; Keane 1997; Vella and Verbeek 1998; Arulampalam 2000). For recent discussions of its relevance to political science, see Wawro (2002) and Bartels et al. (2011).

[3] One of my reviewers rightly pointed out that other mechanisms could introduce dependence on past preferences, for example, when individuals "adapt" to repeatedly presented categories. If the objective of an analysis is to study these survey method effects, the model can be extended, for example, by including dummy response categories in addition to the latent variable (see Heckman 1978 for a detailed discussion of continuous and categorical lagged dependent variables).

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

survey waves. For identification, the variance of the stochastic errors, distributed $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, has to be fixed. I set $\sigma_\epsilon^2 = 1$, yielding an ordered probit specification.[4]

Unobserved individual heterogeneity is modeled via random effects, which are drawn from a normal distribution centered at zero with estimated variance σ_ξ^2:

$$\xi_i \sim N(0, \sigma_\xi^2). \tag{4}$$

This specification can be seen as a multilevel or hierarchical model, with responses nested within individuals. The presence of random effects induces correlations between responses of the same individual over time (Rabe-Hesketh and Skrondal 2008).[5] The proportion of total variance that is due to individual random effects, after accounting for preference persistence, can be estimated by

$$\rho = \frac{\sigma_\xi^2}{(1 + \sigma_\xi^2)}. \tag{5}$$

This provides a useful indicator of the relevance of unobserved individual differences, ignored in cross-sectional analyses.

2.2 *Modeling Initial Observations*

The previous discussion indicates that one generally assumes preference or attitude formation to be a continuous ongoing process. However, panel data provide only a limited window into this process. Clearly, the first panel observation of an individual does not coincide with the first time he or she has ever formed a preference. To the contrary, most researchers would argue that individuals start forming preferences at a very young age and are influenced by parental characteristics, such as education, and by both observed and unobserved individual characteristics.[6] Thus, modeling initial observations has special relevance in a (short) dynamic panel model, as one's "assumption about the initial observations plays a crucial role in interpreting the model" (Anderson and Hsiao 1981, 598).[7]

Nerlove, Sevestre, and Balestra (2008, 11–12) argue that initial observations should be modeled by a specification similar to the one affecting the remaining observations—that is, as depending on observed individual characteristics in $\boldsymbol{x}_i$, while possibly including additional background variables $\boldsymbol{v}_i$, such as parental education or the region of upbringing. Furthermore, to capture the dependence of the initial observation on unobserved individual characteristics, one should specify an arbitrary correlation with the individual-specific effect ξ_i (Harris, Matyas, and Sevestre 2008; Nerlove, Sevestre, and Balestra 2008). In specifying an explicit model for endogenous initial observations, I follow Heckman (1981a, 1981b), who specifies an approximation for the fist (latent) observation $z_{i0}|\boldsymbol{x}_i, \xi_i$ as

$$z_{i0} = \boldsymbol{\delta}' \boldsymbol{w}_i + \lambda \xi_i + \epsilon_{i0}, \tag{6}$$

where $\boldsymbol{w}_i = (\boldsymbol{x}_{i0}, \boldsymbol{v}_i)$ is a vector of initial observation covariates composed of an individual's covariate values at sample entry $\boldsymbol{x}_{i0}$ and additional background information $\boldsymbol{v}_i$. As noted above, initial observations are already shaped by unobserved individual characteristics, which

[4] As usual, errors are assumed independent, $\mathrm{Cov}(\epsilon_{is}, \epsilon_{it}) = 0 \forall s \neq t$, and uncorrelated with covariates, $\mathrm{Cov}(\epsilon_{it}, x_{it}) = 0$.

[5] I employ standard assumptions of normal random effects, that is, they are assumed to be independent of stochastic errors: $\mathrm{Cov}(\xi_i, \epsilon_{it}) = 0$, and independent of $\boldsymbol{x}_{it}$: $\mathrm{Cov}(\xi_i, \boldsymbol{x}_{it}) = 0$. The latter assumption is principally unverifiable. Thus, Pudney (2008) suggests to regard this as a normalization and interpret effects of covariates $\boldsymbol{x}_i^*$ (those covariates in $\boldsymbol{x}_{it}$, which are time constant) as a combination of the true effect of $\boldsymbol{x}_i^*$ and the part of the random effect ξ_i that can be proxied by a linear function of $\boldsymbol{x}_i^*$. The estimated random effects variance σ_ξ^2 is then interpreted as variation not predicted by $\boldsymbol{x}_i^*$. Alternatively, the model might be extended to allow for correlated random effects (Mundlak 1978; Wooldridge 2002).

[6] Models that ignore this problem and specify initial conditions as exogenous can lead to severely biased estimates of the most central parameters of a dynamic panel model, namely, individual random effects and preference persistence (e.g., Heckman 1981b; Fotouhi 2005; Arulampalam and Stewart 2009).

[7] As Anderson and Hsiao (1981, 598) note, this is a problem specific to the short dynamic panels (such as household or election panels), since one cannot credibly assume that $T \to \infty$.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Heckman's specification captures by including the individual-specific effect ξ_i with a scale factor λ that allows for a different effect magnitude of unobserved characteristics on initial preferences.[8] Finally, ϵ_{i0} is a random disturbance term at the initial condition assumed uncorrelated with other errors, that is, $\text{Cov}(\epsilon_{i0}, \epsilon_{it}) = 0$, $\forall t > 0$. Monte Carlo evidence indicates that this approximation works well in short panels (Heckman 1981a; Akay 2012).[9] A somewhat more detailed discussion can be found in supplementary appendix A.

Jointly estimating (1)–(4) and (6) yields a model that deals with four of the five central problems outlined in the introduction. The dynamic model is supposed to capture serial correlation of responses given at different points in time by the same individual (e.g., Beck and Katz 1996). An estimate of this correlation is given by ρ defined in equation (5). To test for the remaining autocorrelation, latent residuals (Albert and Chib 1995) can be used. I calculate the remaining residual correlation as

$$\hat{r} = \frac{\sum_{i=1}^{N}\sum_{t=2}^{T}\mu_{it}\mu_{it-1}}{\sum_{i=1}^{N}\sum_{t=2}^{T}\mu_{it}^2}, \tag{7}$$

where μ_{it} stands for the linear predictor used in equation (3). If the specification succeeds in modeling individuals' correlated responses over time, then $\hat{r}$ should be close to zero.

2.3 *Prior Specifications*

Model specification is completed by assigning (hyper-) priors to all parameters.[10] Priors for intercept and parameters of individual characteristics, in both dynamics and initial condition equations, are diffuse with mean zero and large variance to yield regression-type estimates:

$$\beta, \delta \sim N(0, 100). \tag{8}$$

I use a normally distributed prior for ϕ, the parameter capturing persistence of preferences. I set a prior mean of 0.5 indicating an *a priori* expectation that persistence is not zero, but use a very large variance to yield a diffuse prior:

$$\phi \sim N(0.5, 100). \tag{9}$$

More informative priors might be preferable in some applications, for example, by restricting ϕ using an uniform prior on $U(-1, 1,)$.

My hyperprior for the variance of individual random effects is uniform on the standard deviation, bounded between zero and ten. Gelman (2006) recommends this prior over the more commonly used inverse Gamma specification (Spiegelhalter et al. 1997):

$$\sqrt{\sigma_\xi^2} \sim U(0, 10). \tag{10}$$

[8] It facilitates a simple specification test of the appropriateness of assuming independence of initial conditions and unobserved individual effects: this assumption is rejected if $\lambda \neq 0$. This parametrization is sometimes called a factor-analytic formulation of random effects (e.g., Skrondal and Rabe-Hesketh 2004). Alternatively, one could introduce a second set of random effects with fixed variance in equation (6), and estimate the covariance between them and those in equation (3). The present formulation is somewhat more intuitive and allows for a more straightforward test of exogeneity by testing the parameter λ instead of a covariance.

[9] Alternative approximations, such as Wooldridge (2005), would specify the distribution of $\xi_i | y_{i0}, \boldsymbol{x}_{it}$, that is, simply include the first panel observation among the regressors. This approximation is computationally easier to implement than Heckman's solution, which explains its predominance in applied research. However, if one specifies the preferences as latent constructs, the variable one would need for conditioning on (z_{i0}) is not observable (Pudney 2006, 8). As another disadvantage, this approximation usually works less well in short panels (Akay 2012).

[10] Note that, as in every Bayesian analysis, sensitivity analyses for values of the hyperparameters should be carried out. For an overview of robustness check strategies, see Gill (2008a, 199f.). Basic regression-type priors can be checked by using different variances. For more specific or complex priors, I describe sensitivity check strategies in the text.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

However, this prior has the disadvantage of assigning equal probability to unrealistically large random effect variances. Although this can be seen as representing very little *a priori* knowledge, some researchers might prefer a more informed specification using inverse gamma priors

$$\sigma_{\xi}^{-2} \sim \Gamma(a_0, b_0), \tag{11}$$

with values for a_0 and b_0 chosen using knowledge or expectations of the variation of the individual-specific effects. I provide examples of such an analysis in supplementary appendix D.

An uninformative prior for the random effect scale factor in the initial condition equation (6) is a normal distribution centered at zero and with large variance:

$$\lambda \sim N(0, 100). \tag{12}$$

To ensure that thresholds follow the monotonicity constraint given in equation (2), I specify thresholds recursively ensuring that each subsequent threshold is larger than the previous one by adding a positive value υ_τ. This is achieved by drawing υ_τ from a distribution with positive support such as an exponential distribution (cf. Jackman 2009).[11] The first threshold is normalized to zero for identification; in a model without overall intercept, it can be drawn from a normal distribution centered at zero with large variance:

$$\tau_1 = 0 \tag{13}$$

$$\tau_c = \tau_{c-1} + \upsilon_\tau, \quad c = 2, \ldots, C-1 \tag{14}$$

$$\upsilon_\tau \sim \text{Exp}(1). \tag{15}$$

2.4 *Model Extensions*

Given its hierarchical nature, the model can be extended straightforwardly to capture higher-order nesting by adding random effects for the relevant grouping factor. For example, individuals nested within families (e.g., Winkelmann 2005) or regions j ($j = 1, \ldots, J$) can be modeled by extending equation (3) to

$$z_{ijt} = \phi z_{ijt-1} + \beta' \boldsymbol{x}_{ijt} + \xi_i + \psi_j + \epsilon_{ijt},$$

where ξ_i is the individual-specific effect, and ψ_j represents the regional random effect. Initial conditions are still modeled via equation (6). This is now a three-level model with responses nested in individuals nested in regions. Region random effects are distributed $\psi_j \sim N(0, \sigma_\psi^2)$ with an appropriate hyperprior such as $\sigma_\psi \sim U(0, c)$.

3 Robust Random Effects

The discussion in the previous section assumed normally distributed random effects. This assumption goes almost unnoticed as it is standard in the majority of random effects or "multilevel" models in the social sciences. However, assumptions about the distribution of individual random effects ξ_i are not innocuous and can have important *substantive* implications for panel data analysis.[12] When using a normal distribution as random effects prior, the well-known shrinkage property of hierarchical models (Robert 2007, chap. 10; Gill 2008a, 183) pulls individuals with extreme ξ_i values toward one common mean. Multimodality or interesting patterns of random effects might be obscured. Checks of the normality assumption cannot be carried out using the already shrunken residuals (Kyung, Gill, and Casella 2010).

[11] Here, I use an exponential distribution with rate one, but other parametrizations are possible depending on one's *a priori* expected distance between thresholds. My specification expects a distance of one, which is close to the difference observed in a simple ordered probit regression. An alternative strategy for an ordering constraint is to order thresholds at each step of the MCMC sampler.

[12] For a similar argument in the context of marketing models, see Rossi, Allenby, and Mcculloch (2005, chap. 5); see Navarro et al. (2006) for experimental psychology.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

In this section, I describe two strategies for a more robust estimation of individual heterogeneity: (1) accommodating more extreme individual random effects by specifying a distribution with heavier tails, such as a *t*-distribution with small degrees of freedom (Lange, Little, and Taylor 1989) and (2) estimating the random effects distribution nonparametrically using the Dirichlet process priors (e.g., Gill and Casella 2009).

3.1 t-*Distributed Random Effects*

As an alternative to the normal distribution, a *t*-distribution can be used as a robust prior for random effects. A *t*-distribution with small degrees of freedom has heavier tails and accommodates more extreme random effect values (cf. Lange, Little, and Taylor 1989; Gelman et al. 2004, chap. 17). Thus, changing the distributional specification in equation (4) to

$$\xi_i \sim t(0, \sigma_\xi^2, \mathrm{df}) \tag{16}$$

yields a model with *t*-distributed random effects. However, estimating the degrees of freedom from the data—for example, by assigning a uniform prior—is often rather difficult. For my goal of checking the robustness of the normal random effects assumption, choosing a small value, such as four degrees of freedom, is more appropriate (cf. Gelman et al. 2004, 446).

3.2 *Dirichlet Process Random Effects*

A more flexible alternative to assuming normally distributed random effects consists in estimating the random effects distribution non- or semiparametrically. In the simpler linear dynamic panel case, a fixed effects approach can be employed without distributional assumptions; however, this is unavailable for the current model (e.g., Heckman 1981b; Nickell 1981). Thus, when random effects have to be used, Arellano and Carrasco (2003) argue that (p. 126) "a semiparametric random effects specification may represent a useful compromise" between fixed and random effects.

In a frequentist framework, nonparametric estimation can be accomplished by using finite mixtures of normals or by approximating the random effects distribution by a finite number of mass points (e.g., Heckman and Singer 1984; Lindsay 1995; Aitkin 1999; Eckstein and Wolpin 1999; Vermunt 2004). When applied to substantive research questions, a central problem consists in how to choose the number of mixtures or mass points (Laird 1978; Follmann and Lambert 1989; Skrondal and Rabe-Hesketh 2004, 181f.; Vermunt, Tran, and Magidson 2008).

In a fully Bayesian analysis, instead of assuming a distribution G for the random effects, one can place a Dirichlet process (DP) prior (Ferguson 1973, 1974) on G itself to indicate uncertainty about its shape (e.g., Kleinman and Ibrahim 1998; Gill and Casella 2009):

$$\xi_i \sim G \tag{17}$$

$$G \sim \mathrm{DP}(\alpha, G_0). \tag{18}$$

A Dirichlet process is characterized by two components. The base distribution G_0 is the expectation of G—the distribution one would have used in a non-DP model (Escobar 1995, 98). In my current application, this is the zero-centered normal distribution with estimated variance. The precision or dispersion parameter α determines the dispersion of the prior for G over its mean G_0 (Müller and Quintana 2004). Thus, using a Dirichlet process prior, each set of individual random effects $\{\xi_1, \ldots, \xi_N\}$ drawn from G lies in a set of K distinct values or "subclusters" (with $K \leq N$) sampled from G_0: $\{\zeta_1, \ldots, \zeta_K\}$.[13] For each number of realized subclusters at any particular step

[13] The term "subcluster" is used to indicate that clustering is done nonparametrically and not based on the substantive criteria (cf. Kyung, Gill, and Casella 2010).

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

of an MCMC sampler, random effects ξ_i are drawn from the set $\{\zeta_1, \ldots, \zeta_K\}$ via multinomial sampling. Define subcluster membership indicators $S = \{s_1, \ldots, s_N\}$, which are $s_i = k$ if $\xi_i = \zeta_k$; and $m_k = \#\{s_i = k\}$ as the number of random effects that share the same value ζ_k (i.e., they belong to the same subcluster k).[14]

To illustrate the working of the Dirichlet process, I describe the assignment of random effect ξ_i of a particular individual to a subcluster k, conditional on all remaining random effects $\xi_{[i]} = \{\xi_1, \ldots, \xi_{i-1}, \xi_{i+1}, \ldots, \xi_N\}$ being already assigned. Denote by $S_{[i]}$ the specific configuration of $N-1$ random effects into $K_{[i]}$ subclusters existing at this point, with $m_{[i],k} = \#\{s_i = k, k \neq i\}$ giving the number of individuals sharing a common value $\zeta_{[i],k}$. The conditional prior for ξ_i is (see Hanson, Branscum, and Johnson 2005 or Dunson, Pillai, and Park 2007 for details)

$$[\xi_i|\xi_{[i]}, K_{[i]}, S_{[i]}, \alpha] \sim \frac{\alpha}{\alpha + N - 1} G_0 + \frac{1}{\alpha + N - 1} \sum_{k \neq i} \delta(\xi_k) \tag{19}$$

$$\sim \frac{\alpha}{\alpha + N - 1} G_0 + \frac{1}{\alpha + N - 1} \sum_{k=1}^{K_{[i]}} m_{[i],k} \delta(\zeta_{[i],k}), \tag{20}$$

where $\delta(\cdot)$ now represents the Dirac delta function yielding a single value at its argument. In other words, ξ_i forms a new subcluster with probability $\alpha/\alpha + N - 1$, in which case it is drawn from G_0. Else, it gets value $\zeta_{[i],k}$ of an existing subcluster with multinomial probability according to $N_{[i],k}/\alpha + N - 1$. If one imagines a stream of individual random effects to be assigned, this leads to a *preferential attachment* clustering structure: as the number of individuals grows, the probability that a new individual is assigned to an already existing subcluster is proportional to the subcluster's size. The probability that a new individual forms a new subcluster of the Dirichlet process is proportional to α, and if that happens, values for ξ_i are generated according to the base distribution G_0 (Müller, Quintana, and Rosner 2007).

The realized numbers of subclusters K are stochastic and are governed by α, which can itself be estimated from the data (see below). The role of α can be visualized by inspecting its relationship with the expected number of subclusters (Hanson, Branscum, and Johnson 2005), which can be approximated as (Antoniak 1974; Escobar 1995)

$$E(k|\alpha, n) \approx \alpha \log[(\alpha + N)/\alpha]. \tag{21}$$

Figure 1 plots the expected number of subclusters as a function of the number of individuals for different values of α. This nicely illustrates the logarithmic nature of the preferential attachment property of the Dirichlet process and conforms to intuitions about the relationship between the number of different subclusters and the number of individuals: As more and more individuals are observed, the chance of observing new and unexpected random effect values increases, but at a decreasing rate.

In the dynamic panel model with random effects, considered here, the set of parameters in the base distribution is simply $G_0 = \{p(\sigma_\xi^2)\}$ with a uniform hyperprior $\sigma_\xi \sim U(0, 10)$ as before. Thus, the marginal distribution—averaging over all possible G—yields a mixture of normal distributions with the number of subclusters K randomly varying between 1 and N (see Kleinman and Ibrahim 1998 for a similar setup).[15] The individual-specific random effect variance parameters are either selected from the $K_{[i]}$ existing values $\zeta_k = \sigma_{\xi,k}^2$ drawn from G_0 or created via a fresh values draw from G_0. A more detailed technical discussion of the Dirichlet process and its implementation is available in supplementary appendix B.

[14] Thus, using a Dirichlet process prior provides discrete realizations from the infinite space of prior distributions with probability one (Ghosh and Ramamoorthi 2003; Müller and Quintana 2004). A more detailed discussion can be found in supplementary appendix B.

[15] In practical implementations using a truncated Dirichlet process, the number of subclusters is restricted to some truncation value $T \ll N$. See supplementary appendix B for details.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

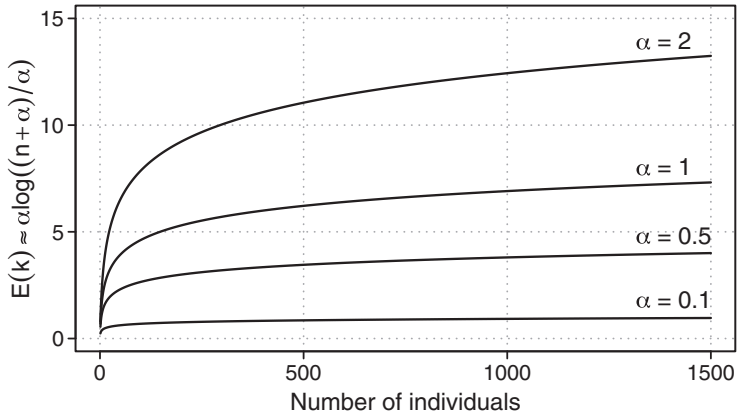


Fig. 1 Expected number of subclusters as function of sample size and Dirichlet process precision parameter α.

3.3 *Estimating Dispersion Parameter* α *from the Data*

The dispersion parameter, α, is a central parameter of the model. Higher values of α increase not only the number of expected subclusters but also the rate with which new ones are created by the Dirichlet process. Given the absence of clear prior expectations about values of α, its value can be determined by the data yielding a mixture of Dirichlet processes (Antoniak 1974). In a fully Bayesian context, this is achieved by assigning it a hyperprior:

$$\alpha \sim \Gamma(a_0, b_0). \tag{22}$$

The gamma distribution is a common choice for this problem (Escobar and West 1998; Jara, Garcia-Zattera, and Lesaffre 2007); however, its parameters do not allow for an intuitive prediction of its effect on the model.[16] Kottas, Müller, and Quintana (2005) provide an approximation to the relationship between Γ-prior parameters and expectation and variance of the number of subclusters K, which can be used to choose semi-informed prior values (for more details, see supplementary appendix C). I select parameters for the gamma hyperprior so that they yield eight *a priori* expected clusters with a SD of 4, which yields parameters $a_0 = 5.16$ and $b_0 = 4.54$ for the gamma prior. To check the sensitivity of this specification, I also used values that lead to a prior expectation of half the number of clusters ($a_0 = 0.921$ and $b_0 = 1.435$). In an alternative strategy (and robustness test), one can forgo estimation of α and instead fix it to a set of prespecified values, for example, $\alpha = \{0.5, 1, 2, 10\}$, in order to determine the robustness of one's estimates to increasingly larger numbers of random effects subclusters. The approximations given in equation (21) and Fig. 1 can serve as guidelines relating values of α to expected subclustering and one's sample size.

4 Application: Dynamic Preferences for Redistribution

A recent wave of research in (comparative) political economy has augmented macro-level studies of redistribution by concentrating on individual-level factors influencing redistribution preferences (see, among many, Iversen and Soskice 2001; Moene and Wallerstein 2001; Alesina and Angeletos 2005; Alesina and La Ferrara 2005; Cusack, Iversen, and Rehm 2005; Scheve and Stasavage 2006; Shayo 2009; Rehm 2011; Rehm, Hacker, and Schlesinger 2012). Studies examining

[16] Specifying an essentially flat prior for computational reasons is common in political science applications (Jackman 2000; but see Jackman and Western 1994), but is of somewhat questionable value here. Even medium-sized values of α lead to a large number of clusters, which in the limiting case creates one cluster per individual—essentially defying the purpose of the hierarchical setup. Therefore, I argue to use a semi-informed prior specification (Gill and Casella 2009, 3) for the DP precision parameter. Kyung, Gill, and Casella (2010) provide alternative strategies of sampling the concentration parameter.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

preferences for redistribution and government intervention in the economy are usually cross-sectional and ignore dynamic aspects of preference formation.[17] As a consequence, estimates of key variables, such as the effect of job loss (as in Cusack, Iversen, and Rehm 2008), might be influenced by unobserved factors, such as ability and motivation, as well as by persistent preferences.[18]

In this section, I present a short study of the dynamics of individual redistribution preferences, by applying the model outlined before to repeated measurements of individuals' preferred level of government intervention. More specifically, I examine individual responses to the question of whether government has the obligation to provide jobs. This survey item correlates strongly with other widely used measures of general redistribution preferences.[19] I examine the effects of income and wealth and "socioeconomic shocks" such as becoming unemployed or getting divorced. For a recent summary of the theoretical relevance of these factors, see Alesina and Giuliano (2011).

4.1 *Data and Variables*

I use data from the British Household Panel Survey, conducted between 1991 and 2008, which provides measurements of my dependent variable on seven occasions. I use the original (Essex) sample and create a balanced panel using individuals who provide responses to all seven waves.[20] This provides me with data on 1958 individuals observed over a span of 17 years. Responses to the item "It is the government's responsibility to provide a job for everyone who wants one" are captured using the usual five-point strongly agree—strongly disagree scale.[21] Since both extreme ends of the response categories are rather sparsely populated, I combine categories to yield a clear three-category response vector, which indicates if preferred levels of government activity should stay the same (0) or should be increased (1) or decreased (−1).[22] Thus, the relationship between observed responses and the latent preference variable is given by

$$y_{it} = \begin{cases} -1 & \text{if} \quad z_{it} < \tau_1 = 0 \\ 0 & \text{if} \quad \tau_1 = 0 < z_{it} < \tau_2 \\ 1 & \text{if} \quad \tau_2 < z_{it} \, . \end{cases}$$

Income is captured by both household income and the share of a respondent's income of the total household income. I measure income as real equivalent household income, that is, it is deflated using the consumer price index with base year 2005 and adjusted for household size using the modified OECD equivalence scale (Hagenaars, de Vos, and Zaidi 1994). I decompose income into a time-varying and a time-constant part. Thus, I estimate both a level and a shock effect, which mirrors the theoretical idea of permanent and transient income components (Friedman 1957). More precisely, observed income w_{it} is decomposed as $w_{it} = \bar{w}_i + (w_{it} - \bar{w}_i)$ with appropriately specified regression weights for both terms. Household wealth is captured by the estimated value of a respondent's house. Definitions and descriptive statistics of all other independent variables used in the analysis can be found in Table 1. Following Gelman (2008), in all models estimated below I centered and scaled all continuous variables by dividing by 2 SDs (which makes them roughly comparable to binary covariates).

[17] But see recent research based on experimental evidence, for example, Margalit (2013) and Neustadt (2010).
[18] This should not be read as a critique of this particular paper, given that the authors' interest lies in a comparative analysis (where panel data are unavailable).
[19] Its correlation with a latent preference measure of several redistribution items, following the methodology of Stegmueller (2011) using data for the UK from the International Social Survey Programme, is 0.64.
[20] Items are available in waves A, C, E, G, J, N, and Q. Estimating the model using multiple imputation for missing values provides results that are substantively similar to the ones presented here, as does an analysis, which uses an unbalanced panel of respondents who participated in at least three waves.
[21] Categories are labeled as strongly agree; agree; neither agree nor disagree; disagree; strongly disagree.
[22] Note that in single index models, such as this one, consistency of the estimates is not hampered by combining categories. See Franses and Cramer (2010) for a further discussion on combining categories in ordered response models. Furthermore, this dependent variable clearly represents a situation where linear models are not appropriate.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Table 1 Descriptive statistics of independent variables

Name	*Description*	*Mean*	*SD*
Income	Equivalent household income (in 10,000 £)		
Permanent	Permanent income component	4.590	2.390
Transitory	Transitory income component	0.000	2.558
Income share	R's share of total HH income	0.532	0.300
House value	Estimated house value (in 100,000 £)	1.181	1.397
Owner	House-owned outright or with mortgage	0.819	0.385
Unemployed	Unemployed	0.033	0.178
Union member	Union member	0.231	0.421
Divorced	Divorced	0.059	0.235
HH size	Size of household	3.202	1.272
N kids	Number of kids in HH	0.906	1.084
Female	Gender: female	0.511	0.500
Age	Age in years	3.963	0.960
Non-white[a]	Ethnic group non-white	0.032	0.175
Education[a,b]			
Degree	University degree	0.199	0.399
A levels	A level or higher national diploma	0.193	0.394
O levels	O level or GCSE	0.411	0.492
London[a]	R grew up in greater London area	0.100	0.300
Parents' jobs[a,c]	Parents' job status		
Unskilled	Blue collar, unskilled jobs	0.163	0.369
Skilled	Blue collar, skilled jobs	0.223	0.416
White collar	White collar	0.145	0.352
Self-employed	Self-employed	0.144	0.351
N rows		13,706	
N individuals		1,958	

[a]Variables are time constant.
[b]Reference category: no or primary education.
[c]Reference category: Managers, Salariat.
[d]Equivalized using OECD scale; deflated using consumer price index, 2005 prices.

4.2 *Results*

First, I describe results obtained from estimating the model described in Section 2 assuming normally distributed random effects. I use a 66% subsample of individuals from the full sample. Results are obtained by MCMC sampling using two chains run for 500,000 iterations thinned by a factor of 25. Two hundred thousand previous iterations are discarded as burn-in. The model is implemented using JAGS (version 3.1.0) with a truncation threshold of 20 (see the discussion of the truncated Dirichlet process in supplementary appendix B).[23] Diagnostics suggested by Brooks and Roberts (1998) and Gelman and Rubin (1992) do not show signs of absence of "convergence."[24]

Resulting estimates are shown in Table 2, where I provide posterior means and standard deviations as well as 95% highest posterior density regions (HPD). Concentrating on central dynamic parameters, I find a significant amount of preference persistence: ϕ is estimated as 0.23 with a small posterior standard deviation. An estimated random effect variance, σ_{ξ}^2, of 0.83 ± 0.09 underscores the importance of controlling for unobserved individual heterogeneity. The proportion of the total

[23] A second run with a truncation value of 40 yields a maximum posterior sampled value for K of 17, which indicates that a truncation level of $T = 20$ was appropriate (see supplementary appendix B).
[24] The posterior samples converge early, but I ran the sampler for longer, providing more draws for the thresholds in order to avoid nonconvergence in this part of the model (cf. Gill 2008b). I conducted an "insurance run" (Gill 2008b, 173): running the sampler for twice as many iterations. Estimates for all key model parameters are virtually identical with the largest difference being 0.0019. All code and diagnostics are available in the author's dataverse.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Table 2 Posterior summary for hierarchical dynamic latent ordered probit model

Initial conditions		*Mean*	*SD*	*95% HPD*		*Dynamics*		*Mean*	*SD*	*95% HPD*	
δ_1	[Permanent income]	−0.422	0.132	−0.679	−0.162	β_1	[Permanent income]	−0.404	0.071	−0.546	−0.267
δ_2	[Transitory income]	−0.033	0.122	−0.273	0.202	β_2	[Transitory income]	−0.059	0.038	−0.133	0.018
δ_3	[R's Income share]	−0.015	0.107	−0.225	0.194	β_3	[R's Income share]	−0.113	0.051	−0.214	−0.016
δ_4	[House value]	−0.219	0.228	−0.656	0.234	β_4	[House value]	−0.115	0.052	−0.220	−0.016
δ_5	[House owner]	−0.073	0.102	−0.267	0.133	β_5	[House owner]	−0.026	0.047	−0.117	0.067
δ_6	[HH size]	0.229	0.185	−0.143	0.579	β_6	[HH size]	0.085	0.079	−0.072	0.237
δ_7	[N kids in HH]	−0.047	0.168	−0.381	0.278	β_7	[N kids in HH]	−0.021	0.065	−0.151	0.105
δ_8	[Union member]	0.155	0.100	−0.045	0.349	β_8	[Union member]	0.148	0.049	0.054	0.244
δ_9	[Age]	−0.551	0.115	−0.779	−0.326	β_9	[Age]	−0.314	0.056	−0.424	−0.202
δ_{10}	[Female]	0.271	0.090	0.094	0.446	β_{10}	[Female]	0.235	0.064	0.112	0.360
δ_{11}	[Divorced]	0.124	0.237	−0.340	0.589	β_{11}	[Divorced]	0.058	0.087	−0.111	0.225
δ_{12}	[Unemployed]	−0.398	0.216	−0.815	0.030	β_{12}	[Unemployed]	0.140	0.100	−0.055	0.334
δ_{13}	[Non-white]	0.317	0.282	−0.231	0.871	β_{13}	[Non-white]	0.606	0.171	0.272	0.941
δ_{14}	[Degree]	−0.524	0.175	−0.866	−0.181	β_{14}	[Degree]	−0.669	0.105	−0.879	−0.469
δ_{15}	[A levels]	−0.521	0.169	−0.849	−0.185	β_{15}	[A-levels]	−0.465	0.102	−0.671	−0.269
δ_{16}	[O levels]	−0.431	0.143	−0.712	−0.155	β_{16}	[O-levels]	−0.321	0.085	−0.493	−0.158
δ_{17}	[Parents: unskilled]	0.464	0.138	0.189	0.732	β_0	[Intercept]	0.213	0.043	0.131	0.299
δ_{18}	[Parents: skilled]	0.233	0.117	0.004	0.462	τ_2		0.818	0.019	0.781	0.855
δ_{19}	[Parents: white-collar]	0.299	0.125	0.055	0.544	ϕ		0.226	0.026	0.176	0.277
δ_{20}	[Parents: self-empl.]	0.115	0.130	−0.137	0.373	σ_ξ^2		0.829	0.092	0.654	1.013
δ_{21}	[London]	0.047	0.139	−0.220	0.326	$\hat{r}$		−0.009	0.012	−0.031	0.015
λ		1.185	0.098	0.995	1.379	ρ		0.452	0.027	0.398	0.506
Deviance							25,705				
DIC							26,259				
Posterior predictive p-value							0.528				

Note. Based on 20,000 MCMC draws. Threshold τ_1 fixed at 0. Balanced panel, 9044 rows, 1292 individuals. Posterior predictive check calculated for mean number of preference changes.

variance that is due to unobserved individual factors, ρ, is estimated as $45 \pm 3\%$. Thus, almost half of the difference in preferences between individuals is due to unobserved factors such as ability or motivation (which remains hidden in cross-sectional studies). Clearly, more research is needed to capture such unobserved individual characteristics.

As described in Section 2, a specification test for the independence of initial conditions and unobserved individual effects is obtained by testing whether λ is equal to zero. This is clearly rejected by an estimate of 1.19 and an HPD region far away from zero. In other words, initial conditions should be modeled as endogenous to individual (observed and unobserved) characteristics. Relevant covariates in the initial equation are age, income, and notably education, as well as presample information on parental background. For example, individuals who grew up in a working-class household already have substantively higher preferences for government intervention at the start of the panel.

In a dynamic panel model, a central quantity of interest are long-run or steady-state relationships between z and $\boldsymbol{x}$ taking preference persistence into account. Since I fixed the scale of the error variance to one, steady-state effects are calculated as $\beta/(1-\phi)$. Using 5,000 draws from the relevant parameters' posterior distributions, I calculate posterior means and standard deviations of steady-state effects, displayed in Table 3. For easier interpretation, I provide them both in the metric of the latent-dependent variable z and calculated as first differences in predicted probabilities of preferring more government intervention resulting from a unit change in a covariate. For discrete variables, this reflects a change from zero to one, and for continuous variables, this represents a change of 2 SDs (cf. Gelman 2008).

Long-run estimates of wealth captured by permanent income and house value show a strong and negative relationship with preferences for government intervention. All else equal, a unit change of permanent income reduces an individual's probability to opt for more government intervention by 17 ± 3 percentage points. It is noteworthy that income shocks have little effect on preferences and are statistically indistinguishable from zero. I find the same for the estimated long-run effect of becoming unemployed, which is large but has a posterior density that includes zero. This also holds for its parameter estimates displayed in Table 2. Excluding all income effects from the model does not change this finding. This points to the relevance of including preference persistence

Table 3 Steady-state effects (calculated on the scale of the latent variable z and as predicted probability of responding in the highest category; posterior means and standard deviations)

	z-metric		$P(y_i = 1)$	
	Mean	*SD*	*Mean*	*SD*
Permanent income	−0.521	0.093	−0.172	0.031
Transitory income	−0.076	0.050	−0.025	0.016
R's Income share	−0.146	0.066	−0.049	0.022
House value	−0.149	0.068	−0.049	0.022
House owner	−0.034	0.060	−0.014	0.026
Household size	0.110	0.103	0.036	0.034
N kids in HH	−0.027	0.085	−0.009	0.028
Union member	0.192	0.064	0.065	0.022
Age	−0.406	0.074	−0.134	0.024
Female	0.304	0.083	0.100	0.027
Divorced	0.074	0.111	0.027	0.039
Unemployed	0.178	0.131	0.064	0.046
Non-white	0.785	0.221	0.294	0.085
Degree	−0.866	0.138	−0.233	0.030
A levels	−0.601	0.134	−0.174	0.034
O levels	−0.416	0.111	−0.134	0.035

Note. Calculated using 5,000 simulated values. Predicted probabilities represent unit change in variable holding all else constant.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

and (especially) unobserved individual heterogeneity in studies of individual preferences. It is this specification of unobserved heterogeneity which I turn to next.

4.3 *Robust Random Effects Results*

To check the robustness of my random effects specification, I reestimated my model using the strategies outlined in Section 3. A model with *t*-distributed random effects with 4 df produces a lower estimate of the random effects variance, σ_ξ^2, of 0.77 with a 95% HPD region ranging from 0.68 to 0.86. However, all other model parameters, including the preference persistence parameter ϕ, are estimated at virtually the same values (at 2 significant figures). When using a more flexible density estimate of the random effects distribution using a Dirichlet Process prior, more differences emerge.[25]

Figure 2 plots a kernel density estimate of the distribution of random effect estimates (more precisely their posterior expectation) from the Dirichlet process hierarchical model. Clearly, the distribution of random effects differs from the traditionally made normal assumption, being slightly skewed and more peaked. However, there is no clear evidence of multimodality or the existence of extreme random effects in the tails of the distribution. This suggests that central model parameters might not be too strongly affected by differences in random effect estimates.

To illustrate differences in parameter estimates that emerge when using different random effect prior specifications, I plot posterior distributions of some selected parameters in Fig. 3. It shows posterior distributions of the initial condition random effects scale factor λ, preference persistence ϕ, and estimates of age and being female obtained using normal, *t*, and DP random effects. Estimates of the scale factor and preference persistence are indistinguishable between normal and *t*-distributed random effects. However, they are larger under the DP prior specification, especially for preference persistence. Results for substantive covariates also differ when using a flexible DP prior specification. My estimate for the influence of age on preferences for government intervention becomes smaller, while the posterior distribution of being female is clearly shifted to the right, indicating an even stronger effect. Nonetheless, the magnitude of these differences is limited and other covariate estimates are somewhat less affected than the ones shown here.

To assess if these differences change one's substantive results, it is advisable to focus again on steady-state estimates calculated from the model. In Table 4, panel (A), I provide steady-state effects from the DP random effects model. As before, I calculate them in the metric of the latent variable and as predicted probabilities of preferring more government involvement. In panel (B),

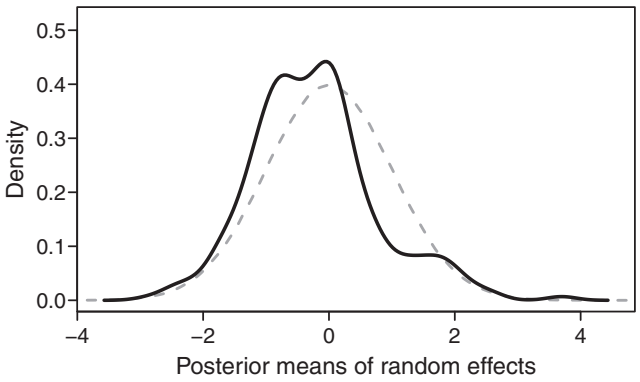


Fig. 2 Distribution of random effects using a Dirichlet process prior. Density estimate of posterior means of random effects ξ_i evaluated over grid of two hundred points. Normal distribution (dashed lines) shown for comparison.

[25] A full table of parameter estimates for the DP prior model is given in appendix Table A1 of this article.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

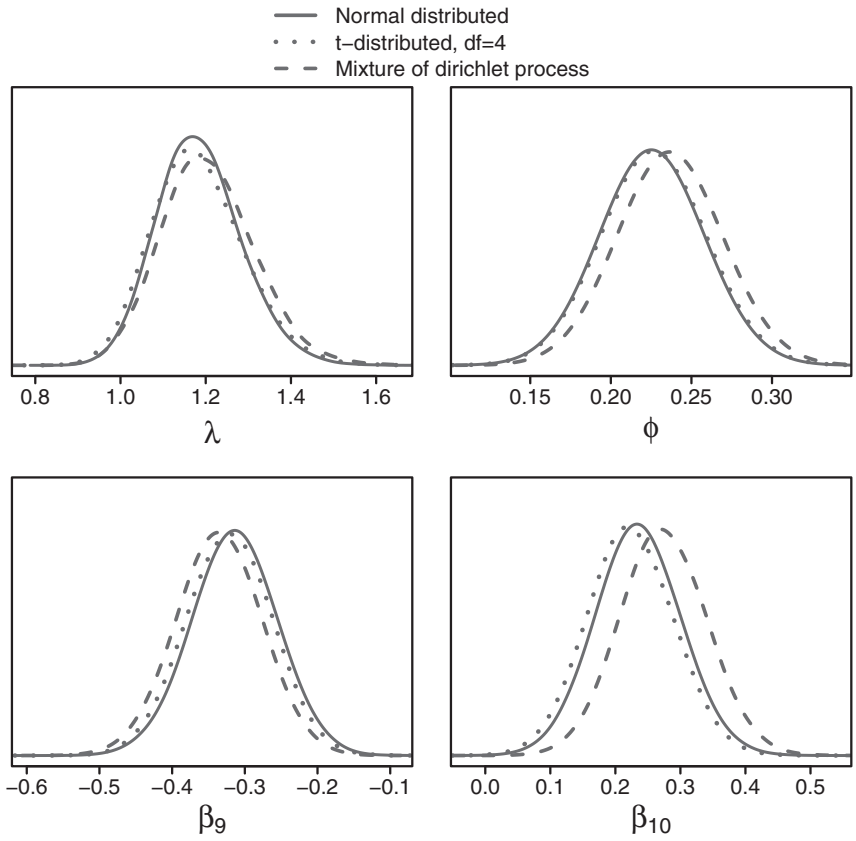


Fig. 3 Consequences of different random effect prior specifications. Posterior distributions of selected parameters obtained using a normal distribution; a *t*-distribution with four degrees of freedom; and mixture of Dirichlet processes as prior.

Table 4 Steady-state estimates from Dirichlet Process random effects model (panel (A): estimated steady-state effects; panel (B): the difference to normal random effects model; posterior means and standard deviations)

	(A) Estimates				*(B) Difference to normal RE*			
	z-metric		$P(y_i = 1)$		*z-metric*		$P(y_i = 1)$	
	Mean	*SD*	*Mean*	*SD*	*Mean*	*SD*	*Mean*	*SD*
Permanent inc.	–0.516	0.088	–0.187	0.031	0.005	0.130†	–0.015	0.043†
Transitory inc.	–0.081	0.051	–0.030	0.019	–0.006	0.071†	–0.004	0.025†
Income share	–0.137	0.065	–0.050	0.024	0.007	0.091†	–0.002	0.033†
House value	–0.145	0.068	–0.053	0.025	0.004	0.097†	–0.003	0.033†
House owner	–0.025	0.062	–0.013	0.029	0.007	0.086†	0.002	0.039†
Household size	0.142	0.105	0.052	0.038	0.034	0.147†	0.016	0.051†
N kids in HH	–0.041	0.087	–0.015	0.031	–0.014	0.120†	–0.006	0.042†
Union member	0.193	0.063	0.072	0.024	0.002	0.089†	0.007	0.033†
Age	–0.440	0.075	–0.160	0.028	–0.033	0.105†	–0.026	0.037†
Female	0.359	0.086	0.131	0.031	0.054	0.119†	0.031	0.039†
Divorced	0.077	0.112	0.028	0.042	0.002	0.157†	0.002	0.057†
Unemployed	0.181	0.129	0.069	0.050	0.000	0.180†	0.004	0.068†
Non-white	0.635	0.224	0.245	0.084	–0.145	0.315†	–0.051	0.119†
Degree	–0.844	0.134	–0.264	0.037	0.017	0.196†	–0.031	0.045†
A levels	–0.614	0.134	–0.202	0.040	–0.016	0.189†	–0.029	0.051†
O levels	–0.422	0.109	–0.152	0.038	–0.006	0.152†	–0.018	0.051†

Note. Calculated using 5,000 simulated values. Predicted probabilities represent unit change in variable holding all else constant. Differences in panel (B) calculated as DP random effects estimates—normal random effects estimates. Difference estimates whose 95% HPD interval includes zero are marked by a †.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

I calculate the difference to the steady-state estimates based on normal random effects, shown in Table 3, and mark the difference in estimates whose 95% HPD interval contains zero indicated by dagger. I find that differences are especially marked for estimates of time-constant covariates. The difference between the estimated effects of holding an advanced degree is almost three percentage points, whereas the effect of being non-white differs by five percentage points. However, when taking uncertainties of my estimates into account, these differences appear to not be statistically relevant: in each and every case, the 95% highest posterior density interval of the difference contains zero. Thus, in this particular application, one can conclude that substantive results obtained with a "simple" Gaussian random effects specification are robust to violations of the distributional assumption of unobserved individual heterogeneity.

5 Conclusion

The central aim of this article is to present a modeling strategy for analyzing the dynamics of individual preferences or attitudes using panel data. I employ the idea of an underlying latent continuous variable, which generates observed categorical preference measures. The dynamics of the model are also specified on the level of the latent variable since it should be one's latent past preference—not observed survey scores—providing feedback to current preferences. Furthermore, I explicitly model initial observations, following the approach suggested by Heckman (1981a, 1981b). I capture unobserved individual heterogeneity using random effects and discuss possible shortcomings of the usual distributional assumptions. I employ a distinctively "Bayesian" solution to this problem, which is to specify a prior over possible random effect distributions, in order to capture uncertainty about its true form. This yields flexible nonparametric density estimation of random effects, which I use to assess the robustness of my findings.

Applying the model to data on individuals' preferences for government intervention over a span of 17 years clearly shows the necessity of employing a hierarchical dynamic panel modeling approach. First, I find a significant level of preference persistence. In other words, individuals' preferences are "sticky," and covariate estimates will be biased when ignoring this fact. Second, initial conditions matter. Individuals enter the panel study with preferences already shaped by presample variables and observed and unobserved characteristics. Third, nearly half of the total variation in preferences is due to unobserved individual factors such as motivation or ability. Using both parametric and semi-parametric random effects specifications, I show that these findings are robust to distributional assumptions.

Existing political science research on individual preferences and attitudes using cross-sectional data should be augmented into the time domain to explicitly study dynamic implications of theories. Using panel data and an appropriate dynamic model provides the tools to generate new insights into how individual preferences evolve over time, how they are shaped by observed and unobserved individual characteristics, and how individuals adjust their preferences in reaction to socioeconomic shocks.

Appendix: DP Random Effects Estimates

Table A1 shows estimated parameters of the model with Dirichlet process random effects. α is the estimated dispersion parameter of the Dirichlet process; K represents the sampled value of the number of clusters at each MCMC step.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Table A1 Posterior summary for hierarchical dynamic latent ordered probit model with Dirichlet process random effects

Initial conditions		*Mean*	*SD*	*95% HPD*		*Dynamics*		*Mean*	*SD*	*95% HPD*	
δ_1	[Permanent income]	−0.403	0.130	−0.665	−0.155	β_1	[Permanent inc.]	−0.393	0.067	−0.526	−0.265
δ_2	[Transitory income]	−0.072	0.124	−0.312	0.172	β_2	[Transitory inc.]	−0.061	0.038	−0.135	0.016
δ_3	[Income share]	0.099	0.113	−0.118	0.326	β_3	[Income share]	−0.104	0.050	−0.203	−0.007
δ_4	[House value]	−0.402	0.241	−0.873	0.071	β_4	[House equity]	−0.110	0.052	−0.214	−0.011
δ_5	[House owner]	0.004	0.105	−0.204	0.209	β_5	[House owner]	−0.020	0.047	−0.111	0.074
δ_6	[HH size]	0.385	0.192	0.014	0.770	β_6	[HH size]	0.108	0.079	−0.050	0.260
δ_7	[N kids in HH]	−0.129	0.172	−0.471	0.204	β_7	[N kids in HH]	−0.032	0.066	−0.160	0.097
δ_8	[Union member]	0.171	0.101	−0.027	0.370	β_8	[Union member]	0.148	0.048	0.053	0.241
δ_9	[Age]	−0.720	0.126	−0.971	−0.479	β_9	[Age]	−0.336	0.056	−0.446	−0.225
δ_{10}	[Female]	0.485	0.113	0.265	0.710	β_{10}	[Female]	0.274	0.064	0.147	0.400
δ_{11}	[Divorced]	0.047	0.241	−0.427	0.509	β_{11}	[Divorced]	0.057	0.085	−0.100	0.231
δ_{12}	[Unemployed]	−0.408	0.221	−0.834	0.035	β_{12}	[Unemployed]	0.137	0.099	−0.058	0.331
δ_{13}	[Non-white]	0.134	0.293	−0.443	0.703	β_{13}	[Non-white]	0.482	0.171	0.152	0.820
δ_{14}	[Degree]	−0.499	0.173	−0.833	−0.156	β_{14}	[Degree]	−0.643	0.101	−0.845	−0.449
δ_{15}	[A levels]	−0.546	0.171	−0.873	−0.203	β_{15}	[A-levels]	−0.469	0.102	−0.666	−0.267
δ_{16}	[O levels]	−0.485	0.142	−0.760	−0.202	β_{16}	[O-levels]	−0.324	0.082	−0.480	−0.160
δ_{17}	[Parents: unskilled]	0.481	0.139	0.207	0.751	β_0	[Intercept]	0.405	0.073	0.264	0.551
δ_{18}	[Parents: skilled]	0.264	0.118	0.034	0.496	κ_2		0.811	0.018	0.777	0.844
δ_{19}	[Parents: white-collar]	0.324	0.127	0.077	0.572	ϕ		0.237	0.026	0.185	0.288
δ_{20}	[Parents: self-empl.]	0.160	0.131	−0.088	0.425	$\hat{r}$		−0.010	0.012	−0.034	0.012
δ_{21}	[London]	0.054	0.139	−0.224	0.320	α		1.476	0.508	0.583	2.502
λ		1.206	0.109	0.994	1.418	K		11.944	3.199	6.000	18.000
Deviance					25,704						
DIC					26,192						
Posterior predictive *p*-value					0.551						

Note: Based on 20,000 MCMC draws. Threshold τ_1 fixed at zero. Balanced panel: N1 = 9044, N2 = 1292.

References

Aitkin, Murray. 1999. A general maximum likelihood analysis of variance components in generalized linear models. *Biometrics* 55:117–28.

Akay, Alpaslan. 2012. Finite-sample comparison of alternative methods for estimating dynamic panel data models. *Journal of Applied Econometrics* 27:1189–204.

Albert, James H., and Siddhartha Chib. 1993. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88:669–79.

Albert, Jim, and Siddhartha Chib. 1995. Bayesian residual analysis for binary response regression models. *Biometrika* 82:747–59.

Alesina, Alberto, and Eliana La Ferrara. 2005. Preferences for redistribution in the land of opportunities. *Journal of Public Economics* 89:897–931.

Alesina, Alberto, and George-Marios Angeletos. 2005. Fairness and redistribution. *American Economic Review* 95:960–80.

Alesina, Alberto, and Paola Giuliano. 2011. Preferences for redistribution. In *Handbook of social economics*, eds. Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 93–131. San Diego, CA: North-Holland.

Anderson, T. W., and C. Hsiao. 1981. Estimation of dynamic models with error components. *Journal of the American Statistical Association* 76:598–606.

Antoniak, Charles E. 1974. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Annals of Statistics* 2:1152–74.

Arellano, Manuel, and Stephen Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies* 58:277.

Arellano, Manuel, and Raquel Carrasco. 2003. Binary choice panel data models with predetermined variables. *Journal of Econometrics* 115:125–57.

Arulampalam, Wiji. 2000. Unemployment persistence. *Oxford Economic Papers* 52:24–50.

Arulampalam, Wiji, and Mark B. Stewart. 2009. Simplified implementation of the Heckman estimator of the dynamic probit model and a comparison with alternative estimators. *Oxford Bulletin of Economics and Statistics* 71:659–81.

Bartels, Brandon L, Janet M. Box-Steffensmeier, Corwin D. Smidt, and Rene M. Smith. 2011. The dynamic properties of individual-level party identification in the United States. *Electoral Studies* 30:210–22.

Beck, Nathaniel, and Jonathan N. Katz. 1996. Nuisance vs. substance: Specifying and estimating time-series-cross-section models. *Political Analysis* 6:1–36.

Blundell, Richard, and Stephen Bond. 1998. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics* 87:115–43.

Brooks, Stephen P., and Gareth O. Roberts. 1998. Convergence assessment techniques for Markov chain Monte Carlo. *Statistics and Computing* 8:319–35.

Cusack, Thomas, Torbern Iversen, and Phillip Rehm. 2005. Risks at work: The demand and supply sides of government redistribution. *Oxford Review of Economic Policy* 22:365–89.

———. 2008. Economic shocks, inequality, and popular support for redistribution. In *Democracy, inequality, and representation: A comparative perspective*, eds. Pablo Beramendi and Christopher J. Anderson, 203–31. New York: Russell Sage Foundation.

Czado, Claudia, Anette Heyn, and Gernot Müller. 2011. Modeling individual migraine severity with autoregressive ordered probit models. *Statistical Methods and Application* 20:101–21.

Dunson, David B, Natesh Pillai, and Ju-Hyun Park. 2007. Bayesian density regression. *Journal of the Royal Statistical Society B* 69:163–83.

Eckstein, Zvi, and Kenneth Wolpin. 1999. Why youths drop out of high school: The impact of preferences, opportunities, and abilities. *Econometrica* 67:1295–339.

Escobar, Michael D. 1995. Nonparametric Bayesian methods in hierarchical models. *Journal of Statistical Planning and Inference* 43:97–106.

Escobar, Michael D., and Mike West. 1998. Computing Bayesian nonparametric hierarchical models. In *Practical nonparametric and semiparametric Bayesian statistics*, eds. Dipak K. Dey, Peter Müller, and Debajyoti Sinha, 1–22. Springer.

Ferguson, Thomas S. 1973. A Bayesian analysis of some nonparametric problems. *Annals of Statistics* 1:209–30.

———. 1974. Prior distributions on spaces of probability measures. *Annals of Statistics* 2:615–29.

Follmann, Dean A., and Diane Lambert. 1989. Generalizing logistic regression by nonparametric mixing. *Journal of the American Statistical Association* 84:295–300.

Fotouhi, Ali Reza. 2005. The initial conditions problem in longitudinal binary process: A simulation study. *Simulation Modelling Practice and Theory* 13:566–83.

Franses, Philip Hans, and J. S. Cramer. 2010. On the number of categories in an ordered regression model. *Statistica Neerlandica* 64:125–8.

Friedman, M. 1957. *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press.

Gelman, Andrew. 2006. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* 1:515–34.

———. 2008. Scaling regression inputs by dividing by two standard deviations. *Statistics in Medicine* 27:2865–73.

Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin. 2004. *Bayesian Data Analysis*. Boca Raton, FL: Chapman & Hall.

Gelman, Andrew, and Donald Rubin. 1992. Inference from iterative simulation using multiple sequences. *Statistical Science* 7:457–511.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Ghosh, J. K., and R. V. Ramamoorthi. 2003. *Bayesian nonparametrics*. New York: Springer.
Gill, Jeff. 2008a. *Bayesian methods: A social and behavioral sciences approach*. Boca Raton, FL: Chapman & Hall.
———. 2008b. Is partial-dimension convergence a problem for inferences from MCMC algorithms? *Political Analysis* 16:153–78.
Gill, Jeff, and George Casella. 2009. Nonparametric priors for ordinal Bayesian social science models: Specification and estimation. *Journal of the American Statistical Association* 104:1–12.
Greene, William, and David Hensher. 2010. *Modeling ordered choices: A primer*. Cambridge, UK: Cambridge University Press.
Grimmer, Justin. 2010. An introduction to Bayesian inference via variational approximations. *Political Analysis* 19:32–47.
Hagenaars, A., K. de Vos, and M. A. Zaidi. 1994. *Poverty statistics in the late 1980s: Research based on micro-data*. Luxembourg: Office for Official Publications of the European Communities.
Hanson, Timothy E., Adam J. Branscum, and Wesley O. Johnson. 2005. Bayesian nonparametric modeling and data analysis: An introduction. In *Handbook of statistics*, Vol. 25, 245–78. Amsterdam: Elsevier.
Harris, Mark N., Laszlo Matyas, and Patrick Sevestre. 2008. Dynamic models for short panels. In *The econometrics of panel data: Fundamentals and recent developments in theory and practice*, eds. Laszlo Matyas and Patrick Sevestre, 249–78. Berlin, Germany: Springer.
Hasegawa, Hikaru. 2009. Bayesian dynamic panel-ordered probit model and its application to subjective well-being. *Communications in Statistics: Simulation and Computation* 38:1321–47.
Heckman, James J. 1978. Dummy endogeneous variables in a simultaneous equation system. *Econometrica* 46:931–59.
———. 1981a. Heterogeneity and state dependence. In *Studies in labor markets*, ed. Sherwin Rosen, 91–140. Chicago: University of Chicago Press.
———. 1981b. The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process. In *Structural analysis of discrete data with econometric applications*, eds. C. F. Manski and Daniel McFadden, 179–95. Cambridge, MA: MIT Press.
Heckman, James J., and B. Singer. 1984. A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica* 52:271–320.
Imai, Kosuke, Ying Lu, and Aaron Strauss. 2008. Bayesian and likelihood inference for 2 x 2 ecological tables: An incomplete-data approach. *Political Analysis* 16:41–69.
Iversen, Torben, and David Soskice. 2001. An asset theory of social policy preferences. *American Political Science Review* 95:875–93.
Jackman, Simon. 2000. Estimation and inference are missing data problems: Unifying social science statistics via Bayesian simulation. *Political Analysis* 8:307–32.
———. 2009. *Bayesian analysis for the social sciences*. New York: Wiley.
Jackman, Simon, and Bruce Western. 1994. Bayesian inference for comparative research. *American Political Science Review* 88:412–23.
Jara, Alejandro, Maria Jose Garcia-Zattera, and Emmanuel Lesaffre. 2007. A Dirichlet process mixture model for the analysis of correlated binary responses. *Computational Statistics and Data Analysis* 51:5402–15.
Johnson, Valen E., and Jim H. Albert. 1999. *Ordinal data modeling*. New York: Springer.
Keane, M. P. 1997. Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics* 15:310–27.
Kleinman, Ken P., and Joseph G. Ibrahim. 1998. A semiparametric Bayesian approach to the random effects model. *Biometrics* 54:921–38.
Kottas, Athanasios, Peter Müller, and Fernando Quintana. 2005. Nonparametric Bayesian modeling for multivariate ordinal data. *Journal of Computational and Graphical Statistics* 14:610–25.
Kyung, Minyung, Jeff Gill, and George Casella. 2010. Estimation in Dirichlet random effects models. *Annals of Statistics* 38:979–1009.
Laird, N. 1978. Nonparametric maximum likelihood estimation of a mixture distribution. *Journal of the American Statistical Association* 73:805–11.
Lange, Kenneth L, Roderick J. A. Little, and Jeremy M. G. Taylor. 1989. Robust statistical modeling using the *t*-distribution. *Journal of the American Statistical Association* 84:881–96.
Lindsay, Bruce. 1995. *Mixture models: Theory, geometry and applications*. Hayward, CA: Institute of Mathematical Statistics.
Lunn, David J., Jon Wakefield, and Amy Racine-Poon. 2001. Cumulative logit models for ordinal data: A case study involving allergic rhinitis severity scores. *Statistics in Medicine* 20:2261–85.
Margalit, Yotam. 2013. Explaining social policy preferences: Evidence from the great recession. *American Political Science Review* 107:80–103.
McKelvey, Richard D., and William Zavoina. 1975. A statistical model for the analysis of ordinal level dependent variables. *Journal of Mathematical Sociology* 4:103–20.
Moene, Kark Ove, and Michael Wallerstein. 2001. Inequality, social insurance, and redistribution. *American Political Science Review* 95:859–74.
Müller, Gernot, and Claudia Czado. 2005. An autoregressive ordered probit model with application to high-frequency financial data. *Journal of Computational and Graphical Statistics* 14:320–38.
Müller, Peter, and Fernando Quintana. 2004. Nonparametric Bayesian data analysis. *Statistical Science* 19:95–110.
Müller, Peter, Fernando Quintana, and Gary L. Rosner. 2007. Semiparametric Bayesian inference for multilevel repeated measurement data. *Biometrics* 63:280–89.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press

Mundlak, Yair. 1978. On the pooling of time-series and cross-section data. *Econometrica* 46:69–85.
Navarro, Daniel J., Thomas L. Griffiths, Mark Steyvers, and Michael D. Lee. 2006. Modeling individual differences using Dirichlet processes. *Journal of Mathematical Psychology* 50:101–22.
Nerlove, Marc, Patrick Sevestre, and Pietro Balestra. 2008. Introduction. In *The econometrics of panel data: Fundamentals and recent developments in theory and practice*, eds. Laszlo Matyas and Patrick Sevestre, 3–22. Berlin: Springer.
Neustadt, Ilja. 2010. Do religious beliefs explain preferences for income redistribution? Experimental evidence. University of Zurich, Socioeconomic Institute working paper 1009.
Nickell, Stephen. 1981. Biases in dynamic models with fixed effects. *Econometrica* 49:1417–26.
Pang, Xun. 2010. Modeling heterogeneity and serial correlation in binary time-series cross-sectional data: A Bayesian multilevel model with AR(p) errors. *Political Analysis* 18:470–98.
Pudney, Stephen. 2006. The dynamics of perception: Modelling subjective well-being in a short panel. ISER working paper 2006–27.
———. 2008. The dynamics of perception: Modelling subjective well-being in a short panel. *Journal of the Royal Statistical Society A* 171:21–40.
Rabe-Hesketh, Sophia, and A. Skrondal. 2008. Generalized linear mixed effects models. In *Longitudinal data analysis: A handbook of modern statistical methods*, eds. Garret Fitzmaurice, Marie Davidian, Geert Verbeke, and Geert Molenberghs, 79–106. Boca Raton, FL: Chapman & Hall.
Rehm, Philipp. 2011. Risk inequality and the polarized American electorate. *British Journal of Political Science* 41:363–87.
Rehm, Philipp, Jacob S. Hacker, and Mark Schlesinger. 2012. Insecure alliances: Risk, inequality, and support for the welfare state. *American Political Science Review* 106:386–406.
Robert, Christan P. 2007. *The Bayesian choice: From decision-theoretic foundations to computational implementation*. New York: Springer.
Rossi, Peter E., Greg M. Allenby, and Robert Mcculloch. 2005. *Bayesian statistics and marketing*. Chichester, UK: Wiley.
Scheve, Kenneth, and David Stasavage. 2006. Religion and preferences for social insurance. *Quarterly Journal of Political Science* 1:255–86.
Shayo, Moses. 2009. A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political Science Review* 103:147–74.
Skrondal, Anders, and Sophia Rabe-Hesketh. 2004. *Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models*. Boca Raton, FL: Chapman & Hall.
Spiegelhalter, D. J., A. Thomas, N. Best, and W. R. Gilks. 1997. *BUGS: Bayesian inference using Gibbs sampling manual*. Cambridge, UK: Medical Research Council Biostatistics Unit.
Spirling, Arthur, and Kevin M. Quinn. 2010. Identifying intraparty voting blocs in the U.K. House of Commons. *Journal of the American Statistical Association* 105:447–57.
Stegmueller, Daniel. 2011. Apples and oranges? The problem of equivalence in comparative research. *Political Analysis* 19:471–87.
Varin, Cristiano, and Claudia Czado. 2010. A mixed autoregressive probit model for ordinal longitudinal data. *Biostatistics* 11:127–38.
Vella, F., and Marno Verbeek. 1998. Whose wages do unions raise? A dynamic model of unionism and wage rate determination for young men. *Journal of Applied Econometrics* 13:163–83.
Vermunt, Jeroen. 2004. An EM algorithm for the estimation of parametric and nonparametric hierarchical nonlinear models. *Statistica Neerlandica* 58:220–33.
Vermunt, Jeroen, Bac Tran, and Jay Magidson. 2008. Latent class models in longitudinal research. In *Handbook of longitudinal research, design, measurement, and analysis*, ed. Scott Menard, 373–85. Waltham, MA: Academic Press.
Wawro, G. 2002. Estimating dynamic panel data models in political science. *Political Analysis* 10:25–48.
Winkelmann, Rainer. 2005. Subjective well-being and the family: Results from an ordered probit model with multiple random effects. *Empirical Economics* 30:749–61.
Wlezien, Christopher. 1995. The public as thermostat: Dynamics of preferences for spending. *American Journal of Political Science* 39:981–1000.
Wooldridge, Jeffrey M. 2002. *Econometric analysis of cross section and panel data*. Cambridge, MA: MIT Press.
———. 2005. Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics* 20:39–54.

https://doi.org/10.1093/pan/mpt001 Published online by Cambridge University Press