Electoral Studies 33 (2014) 52–62

Contents lists available at SciVerse ScienceDirect

Electoral Studies

journal homepage: www.elsevier.com/locate/electstud

Bayesian hierarchical age-period-cohort models with time-structured effects: An application to religious voting in the US, 1972–2008☆

Daniel Stegmueller

Department of Government, University of Essex, Colchester, CO4 3SQ United Kingdom

ARTICLE INFO

Article history:
Received 15 May 2013
Accepted 24 June 2013

Keywords:
Age-period-cohort
Hierarchical models
State-space models
Religion
Voting

ABSTRACT

To examine dynamics of political processes using repeated cross-section data, effects of age, cohort, and time period have to be disentangled. I propose a Bayesian dynamic hierarchical model with cohort and period effects modeled as random walk through time. It includes smoothly time-varying effects of covariates, allowing researchers to study changing effects of individual characteristics on political behavior. It provides a flexible functional form estimate of age by integrating a semi-parametric approach in the hierarchical model. I employ this approach to examine religious voting in the United States using repeated cross-sectional surveys from 1972 to 2008. I find starkly differing nonlinear trends of de- and re-alignment among different religious denominations.

1. Examining socio-political processes

Students of political behavior are often not only interested in relationships between individual characteristics and political outcomes, but want to describe evolution of these relationships over time. In absence of panel data, repeated cross-sectional surveys are used to examine the dynamics of social and political behavior (e.g., among many, De Graaf et al., 2001; Brooks and Manza, 2004; Elff, 2007, 2009). When analyzing change, a well-known problem is the tightly knotted relationship between age, time period and cohort membership, as outlined by Neundorf and Niemi (2014). Recent proposals (Yang and Land, 2006, 2008), which use hierarchical models to disentangle these effects by simultaneously nesting individuals in cohorts and time periods, have become popular (for a recent application see Caren et al. (2011) and Smets and Neundorf (2014)).

In this paper, I propose to extend these various models to ameliorate some of their potential shortcomings. First, previous contributions have simply assumed that both time and cohort random effects are exchangeable. The assumption that the 'order' of periods or cohorts does not matter seems hardly justifiable: both cohort and period effects are (by definition) changes occurring in sequential time. In contrast, I propose a hierarchical model where random effects are time-structured by modeling them as a random walk through time. Second, hitherto proposed hierarchical age-period-cohort (APC) specifications model time period and cohort differences in the dependent variable, but assume time-homogeneity of effects of relevant individual attributes, such as religion or class. However, changes in effects of these

☆ I thank my editors, Anja Neundorf and Richard Niemi, as well as my anonymous reviewer for helpful suggestions and criticisms. Furthermore, I thank NanDirk de Graaf, Thomas Gschwend, Simon Jackman, Larry Bartels, Bent Nielsen, and participants at the ESRA 2011 session "Analyzing social change with repeated cross-sections I" for helpful discussions. As always, they are not to be implicated by my shortcomings. Additional material is available at www.daniel-stegmueller.com/BHAPC.

E-mail address: mail@daniel-stegmueller.com.

0261-3794/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.electstud.2013.06.005

variables are often of central interest to researchers (e.g., Manza and Brooks, 1997; De Graaf et al., 2001; Brooks and Manza, 2004; Elff, 2009; Jansen et al., 2011). Borrowing from state-space modeling, I extend my model to estimate smoothly time-varying effects of central covariates.[1] Finally, in most applications researchers specify flexible time and cohort effects, but treat age as linear or quadratic. I detail how to estimate the functional form of age more flexibly by employing a nonparametric strategy based on penalized splines. I specify my model in a Bayesian framework, which allows me to implement this various extensions under a common hierarchical model framework. Furthermore, using the Bayesian inferential paradigm, researchers do not have to rely on the interpretation of random effects as random samples from a (super-)population, which is questionable when working with APC models. Instead, inferences are made based only on actually observed periods and cohorts.[2]

I demonstrate the practical application of my proposed model specification by analyzing trends in religious voting in the United States. Researchers and pundits alike have put renewed focus on the role of religion in shaping the political landscape of the US (Wald et al., 2005; Wald and Wilcox, 2006). Following the perception of increased polarization, specific interest lies in examining changes in traditional party alignments of different denominations (e.g., Brooks and Manza, 2004). I use repeated cross-sectional surveys from 1972 to 2008 with detailed measures of religious denominations. I find no effect of cohort membership, but starkly differing nonlinear trends of de- and re-alignment among different religious denominations.

The paper is structured in two central parts. The next section contains a detailed discussion of hierarchical age-period-cohort models and my proposed extensions. I also discuss a strategy for comparing models of differing complexity. Section 3 contains the application of my model to denominational vote trends in the US. I discuss data and imputation strategy and present detailed results. Finally, Section 4 concludes the paper with a short discussion.

2. Dynamic hierarchical model

The key to a combined analysis of effects of age, cohort, and time period is to account for the fact that individuals experience the same time period but are also simultaneously members of a specific cohort. To capture this dependence structure, recent contributions (e.g., Yang and Land, 2006, 2008) propose the use of cross-classified random effect models (Snijders and Bosker, 2012; Browne et al., 2001; Rasbash and Browne, 2008).[3] One models response y^*_{itk} of individual i in the cell defined by the cross-classification of time period t ($t = 1,\ldots,T$) and cohort k ($k = 1,\ldots,K$) as a function of age a_i, covariates x_i, and time (p_t) and cohort (c_k) effects:

$$y_{itk} = a_i + p_t + c_k + x_i$$

More specifically, I propose to model responses as resulting from an individual's age a_i, some theoretically relevant variables w_i, a set of controls for socio-economic background characteristics x_i including an overall constant (i.e., for all i, $x_{i1} = 1$), and random effects for time periods ζ_t and cohort membership ξ_k:

$$y^*_{itk} = f(a_i) + \sum_{r=1}^{R} \beta'_r w_{ir} + \sum_{s=1}^{S} \gamma'_s x_{is} + \zeta_t + \xi_k + \varepsilon_{itk} \qquad (1)$$

Here, $f(a)$ is some smooth functional form estimate of age effects (discussed in detail later), β_r and γ_s are effect coefficients for theoretically central and control variables, respectively, and ε_{itk} is a white noise error term. In this general formulation, the model is applicable to outcomes that are either continuous, dichotomous or ordered categorical, by suitably specifying the link between observed y and latent $y*$ (Greene and Hensher, 2010). Since my application uses a binary dependent variable, the following discussion refers to this situation. I use a probit model, obtained by assuming that a latent variable – normally distributed with unit variance – generates observed outcomes via some threshold or utility mechanism (Albert and Chib, 1993; Jackman, 2000). Thus an individual responds $y_{itk} = 0$ if $y^*_{itk} < 0$ and $y_{itk} = 1$ if $y^*_{itk} \geq 0$.

Model specification is completed by specifying a distribution for cohort and time period random effects. Existing proposals (e.g., Yang, 2006) rely on the standard assumption of normally distributed random effects. For example, cohort random effects are assumed to be drawn from a normal distribution with zero mean and estimated variance σ^2:

$$\xi \sim N(0, \sigma^2).$$

While convenient, I argue that this specification is somewhat ignorant about our knowledge of social processes. It seems implausible *a priori*, that each new cohort of individuals or each new time period, is formed in a social and political vacuum – independent of the past. Rather, most theoretical accounts of socio-political change emphasize continuity between cohorts or time points, as well as elements of change. Following these arguments, I propose to add structure to cohort and time period random effects by including systematic dependencies between neighboring periods and cohorts, while still allowing for the occurrence of sudden change.

2.1. Time-structured random effects

To yield a more realistic treatment of period and cohort succession, I model random effects as following a nonstationary second order random walk transition process (Besag et al., 1995; Knorr-Held and Rainer, 2001):

[1] These model specifications are increasingly common in public health and disease mapping research, where researchers face the age-period-cohort problem with grouped data (such as cancer incidence rates). See, e.g., Knorr-Held (2000); Knorr-Held and Rainer (2001); Schmid and Held (2004) and the references therein.

[2] Add to this the general intellectual attractiveness of Bayesian inference (see, e.g., Jackman, 2009, ch. 1 or Jaynes, 2003). Yang and Land (2008) provide further specific arguments for using a Bayesian approach for APC models.

[3] See Smets and Neundorf (2014) in this issue for a more detailed discussion of this model. See Nielsen and Nielsen (2010) for a detailed discussion of identification problems in APC models.