Social Science & Medicine 101 (2014) 176–180

Contents lists available at ScienceDirect

Social Science & Medicine

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

Commentary

Don't birth cohorts matter? A commentary and simulation exercise on Reither, Hauser, and Yang's (2009) age–period–cohort study of obesity

Andrew Bell [a], [b], [*], Kelvyn Jones [a], [b]

[a] School of Geographical Sciences, University of Bristol, University Road, Bristol, BS8 1SS, UK
[b] Centre for Multilevel Modelling, University of Bristol, 2 Priory Road, Bristol, BS8 1TX, UK

ARTICLE INFO

Article history:
Available online 19 September 2013

Keywords:
Age–period–cohort models
Obesity
Collinearity
Model identification

ABSTRACT

Reither, Hauser, and Yang (2009) use a Hierarchical Age–Period–Cohort model (HAPC – Yang & Land, 2006) to assess changes in obesity in the USA population. Their results suggest that there is only a minimal effect of cohorts, and that it is periods which have driven the increase in obesity over time. We use simulations to show that this result may be incorrect. Using simulated data in which it is cohorts, rather than periods, that are responsible for the rise in obesity, we are able to replicate the period-trending results of Reither et al. In this instance, the HAPC model misses the true cohort trend entirely, erroneously finds a period trend, and underestimates the age trend. Reither et al.'s results may be correct, but because age, period and cohort are confounded there is no way to tell. This is typical of age–period–cohort models, and shows the importance of caution when any APC model is used. We finish with a discussion of ways forward for researchers wishing to model age, period and cohort in a robust and non-arbitrary manner.

Introduction

The desire to separate age, period and cohort (APC) effects has been a key feature of both the medical and social sciences for a number of decades (Ryder, 1965). For at least the same period, levels of obesity have been rising at a continuous rate, to the point that in 1997 it was classified by the World Health Organisation as a global epidemic (Caballero, 2007). In 2009, Reither, Hauser, and Yang (2009) used the recently developed Hierarchical Age–Period–Cohort (HAPC) model (Yang & Land, 2006) to assess the relative importance of periods and cohorts in the development of the obesity epidemic. Whilst they found some significant cohort effects, the implication of their results was "that period effects were principally responsible for the obesity epidemic" (Reither et al., 2009: 1445), and this result was repeated by Yang and Land (2013:215–222).

However, the possibility of separating APC effects is beset by an 'identification problem' due to the fact that age, period and cohort when taken together are perfectly collinear. In this paper we show that the HAPC model does not solve this identification problem, and therefore that the results found by Reither et al. should be treated with some scepticism.

The purpose of this paper is twofold. The first substantive contribution is to add to the growing debate in epidemiology regarding the causes of, and therefore possible solutions to, the obesity epidemic. Whether periods or cohorts are responsible for changes in obesity is of profound importance because it should affect how policy interventions are targeted. The second, methodological, contribution is to assess the capabilities of age–period–cohort models, and the dangers of using these models without critical forethought regarding their limits. In this we are building on previous work (Bell & Jones, 2013a, 2013b, 2013c; Glenn, 2005; Luo, 2013; Luo & Hodges, submitted for publication) questioning the capabilities of the HAPC model and other methodological innovations to disentangle APC effects.

We first outline the identification problem and Yang and Land's proffered solution to it. Second we briefly review the literature on the development of the obesity epidemic. Third we outline our simulation design which we use to show that the results found by Reither et al. could have been created by a different data generating process (DGP). Finally, we discuss the implications of this both within obesity research and beyond, considering ways forward for researchers wishing to use techniques like the HAPC model to make robust conclusions regarding APC effects.

DOI of original article: http://dx.doi.org/10.1016/j.socscimed.2009.08.040.

[*] Corresponding author.
E-mail address: andrew.bell@bristol.ac.uk (A. Bell).

0277-9536/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.socscimed.2013.09.004

The APC identification problem and Yang and Land's HAPC model

The conceptual distinction between age, period and cohort is well known (Bell & Jones, submitted for publication; Suzuki, 2012). However despite this, there remains the problem of statistically modelling the three effects because of the mathematical dependency between them:

$$\text{Age} = \text{Period} - \text{Cohort} \tag{1}$$

As such, if we know the value of two of the terms, we will always know the value of the third. From an 'experimental' standpoint, therefore, it is impossible to hold two of APC constant whilst varying the third. Because of this, each of the following DGPs (and an infinite number more) would produce identical values for a dependent variable Y:

$$Y = (1^*\text{Age}) + (1^*\text{Period}) + (1*\text{Cohort}) \tag{2a}$$

$$Y = 2^*\text{Age} + 2^*\text{Cohort} \tag{2b}$$

$$Y = 2^*\text{Period} \tag{2c}$$

Given such data, therefore, it would not be possible to tell which DGP actually produced the data. These three instances presented here have very different substantive meanings, yet it would not be possible to tell which of the three actually produced the data at hand.[1] It is for this reason that many see a solution to the identification problem to be a logical impossibility:

> "The continued search for a statistical technique that can be mechanically applied always to correctly estimate the effects is one of the most bizarre instances in the history of science of repeated attempts to do the logically impossible."
>
> (Glenn, 2005:6)

Despite this, numerous supposed solutions to the identification problem have been proposed, each of which imposes some kind of constraint on the model (Mason, Mason, Winsborough, & Poole, 1973; Sasaki & Suzuki, 1987; Tu, Smith, & Gilthorpe, 2011; Yang, Schulhofer-Wohl, Fu, & Land, 2008). The problem arises when these constraints are not clearly stated, are applied arbitrarily on the basis of statistical necessity, and are not grounded in any kind of substantive theory. The models are generally very sensitive to such constraints and as such can provide extremely misleading results when those constraints are not precisely justified and appropriate.

Yang and Land's proposed solution is to use a cross-classified multilevel model, which treats age as a fixed effect and periods and cohort groups as random effects – contexts in which individuals reside. The model can thus be specified (in the continuous Y case) as:

$$y_{i(j_1j_2)} = \beta_{0j_1j_2} + \beta_1\text{Age}_{i(j_1j_2)} + \beta_2\text{Age}^2_{i(j_1j_2)} + e_{i(j_1j_2)}$$

$$\beta_{0j_1j_2} = \beta_0 + u_{1j_1} + u_{2j_2}$$

$$e_{i(j_1j_2)} \sim N\left(0, \sigma_e^2\right), u_{1j_1} \sim N\left(0, \sigma_{u1}^2\right), u_{2j_2} \sim N\left(0, \sigma_{u2}^2\right) \tag{3}$$

The dependent variable, $y_{i(j_1j_2)}$ is measured for individuals i in period j_1 and cohort j_2. The 'micro' model has linear and quadratic age terms, with coefficients β_1 and β_2 respectively; a constant ($\beta_{0j_1j_2}$) that varies across both periods and cohorts; and a level 1 residual error term ($e_{i(j_1j_2)}$). The macro model defines the intercept in the micro model by a non-varying overall intercept β_0, and a residual term for each of period and cohort. The period, cohort and level-1 residuals are all assumed to follow Normal distributions, each with variances that are estimated.

Putting age in the fixed part and period and cohort in the random part is conceptually attractive; but also, it is argued by Yang and Land that this distinction solves the identification problem:

> "An HAPC framework does not incur the identification problem because the three effects are not assumed to be linear and additive at the same level of analysis" (Yang & Land, 2013:191)

In addition to this, Yang and Land suggest that the inclusion of the quadratic term for age helps to further resolve the identification problem:

> the underidentification problem of the classical APC accounting model has been resolved by the specification of the quadratic function for the age effects.
>
> (Yang & Land, 2006:84)

However, it has been shown elsewhere that this methodological advance in fact amounts to another constraint (Luo & Hodges, submitted for publication), and simulation studies have shown that the use of this model, without critical forethought, can lead to misleading results (Bell & Jones, submitted for publication).

The obesity epidemic

Historically, obesity was a rare affliction, predominantly affecting those of high socio-economic status (Caballero, 2007). However, levels of obesity increased throughout the twentieth century, particularly amongst those of lower socio-economic status and education levels (Visscher, Snijder, & Seidell, 2010). A number of reasons for this have been proposed, including the more sedentary lifestyle associated with the technological advances of the modern world (Rokholm, Baker, & Sorensen, 2010), and the greater availability, portion size and fat content of food (Hill & Peters, 1998). However, the question remains as to whether it is via periods or cohorts that these changes occur. If the former, it would suggest that changes in lifestyle have affected all age groups equally, resulting in bad diets and low levels of exercise for all individuals. In contrast, the latter would suggest that these cultural changes particularly affect people in their formative years, and these changes have affected their behaviour and possibly their physiological resistance to obesity throughout their subsequent life-course. In the same vein, interventions to the obesity epidemic should be similarly targeted to the groups most affected. If cohorts are responsible for changes in obesity, then policy interventions should be focused on children in their formative years because interventions targeted at adults are likely to be ineffectual.

Reither et al. argue that the obesity epidemic is the result predominantly of periods, and their results are shown graphically in the third column of Fig. 1. They argue that "the pattern of predicted probabilities for U.S. adults shows a monotonic increase over time, with no sign of abatement in recent periods of observation" (Reither et al., 2009:1443). Similarly, Allman-Farinelli, Chey, Bauman, Gill, and James (2008) find that period effects are the driving force of changes in their APC analysis of obesity in Australia, whilst Rokholm et al. (2010:843) argue that the slight levelling off of the obesity epidemic observed in recent years "occurred at approximately the same time for different age groups". However, other studies find evidence that cohorts have the greater influence on obesity: for example Olsen, Baker, Holst, and Sorensen (2006) find that non-linearities in cohort trends match for different age

---

[1] The technical consequence of this is that a regression with age, period and cohort as linear independent variables will not be estimable (at least with OLS) because the design matrix $\mathbf{X}^\mathbf{T}\mathbf{X}$ cannot be inverted.