



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/87660/>

Version: Accepted Version

Article:

Bell, A.J. and Jones, K. (2013) The impossibility of separating age, period and cohort effects. *Social science & medicine*, 93. 163 - 165. ISSN: 0037-7856

<https://doi.org/10.1016/j.socscimed.2013.04.029>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Commentary: The impossibility of separating age, period and cohort effects

Andrew Bell^{a b} and Kelvyn Jones^{a b}

^aSchool of Geographical Sciences
University of Bristol
University Road
Bristol
BS8 1SS
UK

^bCentre for Multilevel Modelling
University of Bristol
2 Priory Road
Bristol
BS8 1TX
UK

Draft – please do not cite without permission

Last updated: 19th April 2013

Corresponding Author:

Andrew Bell
School of Geographical Sciences
University of Bristol
University Road
Bristol
BS8 1SS
UK
Email: andrew.bell@bristol.ac.uk

Keywords

Age-period-cohort models, model identification, collinearity.

Research Highlights

- Age, period and cohort (APC) trends cannot be disentangled mechanically
- Explicit assumptions must be made for APC models to be identified
- Imposing arbitrary assumptions leads to arbitrary model results
- Assumptions should be based on strong theory and be stated explicitly

Abstract

This commentary discusses the age-period-cohort identification problem. It shows that, despite a plethora of proposed solutions in the literature, no model is able to solve the identification problem because the identification problem is inherent to the real-world processes being modelled. As such, we cast doubt on the conclusions of a number of papers, including one presented here (Page et al., this issue). We conclude with some recommendations for those wanting to model age, period and cohort in a compelling way.

1 Introduction

This commentary addresses this issue of statistically modelling separate age, period and cohort (APC) effects. The issue has been hotly debated for decades (Ryder, 1965), in sociology (Glenn, 1976; Mason et al., 1973), medical science (Osmond & Gardner, 1989; Robertson & Boyle, 1986) and elsewhere. However, the publication of an article in this issue (Page et al., this issue) alongside more recent methodological developments in APC modelling (Tu et al., 2011; Yang & Land, 2006; Yang et al., 2008), shows that there is still profound interest in modelling and discerning APC effects.

This commentary does not directly critique Page et al.'s paper in terms of its substantive conclusions; rather, we address the key methodological issues in modelling APC effects. However the implications of our argument call into question the results found by Page et al, and should act as a warning for others researchers wishing to disentangle APC effects in a meaningful way .

We first outline what APC effects are substantively, and describe the identification problem which makes them so difficult to model. We then outline some proposed solutions to the identification problem, including that used by Page et al, and explain why they will only work in very specific and arguably usually unrealistic circumstances. The commentary finishes with some recommendations for researchers wishing to model APC effects.

2 Age, period and cohort effects

The difference between age effects, period effects and cohort effects is well explicated by this fictional dialogue by Suzuki (2012, p. 452):

A: I can't seem to shake off this tired feeling. Guess I'm just getting old. [Age effect]

B: Do you think it's stress? Business is down this year, and you've let your fatigue build up.
[Period effect]

A: Maybe. What about you?

B: Actually, I'm exhausted too! My body feels really heavy.

A: You're kidding. You're still young. I could work all day long when I was your age.

B: Oh, really?

A: Yeah, young people these days are quick to whine. We were not like that. [Cohort effect]

It is worth considering these in terms of suicide, the substantive subject area of Page et al.'s paper (this issue). An age effect could entail that the risk of suicide increases as an individual gets older. A cohort effect could suggest that people born or socialized at a certain time (say, during a war) are more likely to commit suicide throughout their life (due, for example, to psychological problems caused in childhood across that cohort). A cohort effect could also be a source of more continuous change, whereby new cohorts are less likely to commit suicide than those before them (perhaps because of continuously improving living conditions over time). Finally, a period effect could suggest that suicide is more likely during a particular event (such as an economic downturn) that affects everybody irrespective of age.

The problem lies in attempting to disentangle the three effects. Age, period and cohort are exactly mathematically confounded, such that:

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

As such, it is impossible conceptually to hold two of the terms constant without holding the third term constant as well. In other words, the three terms cannot simply be put into a regression analysis as there would not be a single solution from any given dataset. For example, (as shown by Bell & Jones, under review-a) consider a data generating process (DGP) where equal age, period and cohort effects, of magnitude 1, contribute to the generation of the Y-variable (which could, taking the Page et al example, be the log of the probability of suicide):

$$Y = Age + Period + Cohort \quad (2)$$

Because of the dependency shown in equation (1), we could substitute *Period* with *Age + Cohort* to give

$$Y = 2 * Age + 2 * Cohort \quad (3)$$

We could then substitute *Age + Cohort* with *Period* to give

$$Y = 2 * Period \quad (4)$$

So, all three of the above DGPs (equations 2-4) would produce the same data. Equally, a regression model using this data to attempt to estimate coefficients associated with these three terms could produce any of these three sets of parameter estimates (and, in fact, an infinite number more) with equally good model fit. Any data that include age, period or cohort effects will be exactly collinear in this way. This is why many have dismissed APC modelling as a 'futile quest' (Glenn, 1976; see also Goldstein, 1979):

"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, p. 6)

3 Solutions to the identification problem, and why they do not work

Despite this scepticism, there have been numerous attempts to find a way around the identification problem. The most common, and that suggested by Mason et al. (1973), is to constrain certain parameters in a model to be equal. So, each age group, period group and cohort group is included in a regression model as a dummy variable, but 2 age groups (or period groups, or cohort groups) are combined together as if they were a single group. This breaks the exact collinearity and allows the model to be identified. However, as Glenn points out (1976, 2005), "the linear dependence is

broken *in the statistical model only* and not in the real world” (Glenn, 2005 p.14, emphasis in original). In the real world, the dependence of equation 1, present in equations 2-4, remains. So whilst this model will produce *an* answer (say, one of equations 2-4), there is no way of knowing if it has found the correct answer unless there is a good reason to think that the identifying constraint imposed is correct. If we cannot make this assumption, the solution found will be as arbitrary as the constraint being imposed.

The ‘solution’ used by Page et al. (this issue) is to model coarse groupings of age and cohort compared to single-year groupings of period. This is effectively a version of the method described above. Multiple constraints are imposed on the parameters, such that the effects of age within each age grouping are assumed to be equal (and each of the cohort effects within a given cohort grouping are also assumed equal). No theory is used to justify this grouping and so we can assume that the results found are as arbitrary as the constraints imposed. For example, a researcher may find the results shown in equation 4 (that the Y-variable is caused solely by a period effect), but it is just as likely that, in fact, equations 2 (a mix of all three APC effect) or 3 (a mix of age and cohort effects) created the data.

Other general solutions to the identification problem have been proposed in recent years. For example, Yang and Land (2006) claim that, by treating period and cohort as cross-classified contexts within a multilevel model, “it is clear that the underidentification problem of the classical APC accounting model has been resolved” (p. 84). Yet, simulation studies have shown that this too is not the case, and the model is subject to the same biases as those described above (Bell & Jones, under review-a; Luo & Hodges, under review). The logical or mathematical confounding (Goldstein, 1979) of age, period and cohort cannot be solved by a trick of data management or model structuring because the confounding is inherent to the APC processes that are being modelled.

4 Recommendations

So far, there is little good news for the researcher hoping to find separate APC effects. However, we believe that theorising and finding age, period and cohort is often very important in social science, and that they can be modelled, so long as the assumptions that are being made by the model are justified by theory and stated explicitly.

It is often the case that we can assume that continuous period effects are non-existent. It seems to us that theory often indicates that progress over time is the result of cohort succession (where new cohorts are different to those that came before) rather than period effects. For example, social attitudes are usually the result of new cohorts with a different attitude replacing older cohorts. There are rarely continuous changes in the world that lead to a shift in attitudes across all ages and cohorts. If we can assume that there are no period trends, and we make clear that we are making that assumption, we are able to model age and cohort in a model and make relatively robust findings (Bell & Jones, under review-b; for example see McCulloch, 2012). It is worth reiterating however that this assumption cannot be made on the basis of the data – it can only be made on the basis of the researcher's understanding of the theoretical causal process being studied.

Where it is assumed that continuous period trends are non-existent, discrete period effects, caused by specific events rather than continuous trends, may still be prevalent. For example, a war in a given period is likely to push up mortality rates for the duration of that war, which would then decline when the war finishes. Yang and Land's (2006) model mentioned above, with the addition of a cohort trend specified in the fixed part of the model (Bell & Jones, under review-a), is able to model these discrete period effects (along with discrete cohort effects) as random effects in a multilevel cross-classified model. These are not affected by the identification problem in the same way as continuous trends that function across all periods. The researcher can thus assess how the variable of interest varies by discrete periods and cohorts, net of any age and cohort trends.

As such, if it is judged that certain assumptions can be made, for example that there is no period trend in the DGP, then we are able to make robust inference about APC effects, if those assumptions are indeed justified. However, where those assumptions are not explicitly declared, or are made arbitrarily on the basis of statistical necessity rather than theory, then readers should be sceptical of the presented results.

5 References

- Bell, A., & Jones, K. (under review-a). Another 'futile quest'? A simulation study of Yang and Land's Hierarchical Age-Period-Cohort model. *Working paper*. Available at <http://www.mendeley.com/download/public/19138881/5392730501/d2fb15102177bd6fd8a32cab9e4193776576e41c/dl.pdf> [Accessed 19th April 2013].
- Bell, A., & Jones, K. (under review-b). Current practice in the modelling of Age, Period and Cohort effects with panel data: a commentary on Tawfik et al (2012), Clarke et al (2009), and McCulloch (2012). *Working paper*. Available at <http://www.mendeley.com/download/public/19138881/5392730511/2f76331fbd618b02e52ff33515ff809b01fdab7c/dl.pdf> [Accessed 19th April 2013].
- Glenn, N.D. (1976). Cohort Analysts Futile Quest - Statistical Attempts to Separate Age, Period and Cohort Effects. *American Sociological Review*, 41, 900-904.
- Glenn, N.D. (2005). *Cohort Analysis*. London: Sage.
- Goldstein, H. (1979). Age, period and cohort effects - a confounded confusion. *Journal of Applied Statistics*, 6, 19-24.
- Luo, L., & Hodges, J. (under review). The cross-classified age-period-cohort model as a constrained estimator. *Working paper*. Available at <http://paa2013.princeton.edu/papers/132093> [Accessed 16th April 2013].
- Mason, K.O., Mason, W.M., Winsborough, H.H., & Poole, K. (1973). Some Methodological Issues in Cohort Analysis of Archival Data. *American Sociological Review*, 38, 242-258.
- McCulloch, A. (2012). Cohort variation in the membership of voluntary organisations in Great Britain, 1991-2007. *Working paper*. Available at <http://www.significancemagazine.org/SpringboardWebApp/userfiles/sig/file/Paper.pdf> [Accessed 3rd Spetember 2012].
- Osmond, C., & Gardner, M.J. (1989). Age, Period, and Cohort Models - Non-Overlapping Cohorts Dont Resolve the Identification Problem. *American Journal of Epidemiology*, 129, 31-35.
- Page, A., Milner, A., Morrell, S., & Taylor, R. (this issue). The role of under-employment and unemployment in recent birth cohort effects in australian suicide. *Social Science & Medicine*, this issue.
- Robertson, C., & Boyle, P. (1986). Age, Period and Cohort Models - the Use of Individual Records. *Statistics in medicine*, 5, 527-538.
- Ryder, N.B. (1965). The Cohort as a Concept in the Study of Social-Change. *American Sociological Review*, 30, 843-861.
- Suzuki, E. (2012). Time changes, so do people. *Social Science & Medicine*, 75, 452-456.
- Tu, Y.K., Smith, G.D., & Gilthorpe, M.S. (2011). A New Approach to Age-Period-Cohort Analysis Using Partial Least Squares Regression: The Trend in Blood Pressure in the Glasgow Alumni Cohort. *Plos One*, 6.

- Yang, Y., & Land, K.C. (2006). A mixed models approach to the age-period-cohort analysis of repeated cross-section surveys, with an application to data on trends in verbal test scores. *Sociological Methodology*, 36, 75-97.
- Yang, Y., Schulhofer-Wohl, S., Fu, W.J.J., & Land, K.C. (2008). The intrinsic estimator for age-period-cohort analysis: What it is and how to use it. *American Journal of Sociology*, 113, 1697-1736.