Methods for disentangling period and cohort changes in mortality risk over the twentieth century: comparing graphical and modelling approaches

July 20, 2022

Authors

1. Jones, P.M. - University of Derby (ORCID 0000-0001-5173-3245)
2. Minton, J. - NHS Health Scotland (ORCID 0000-0003-1207-6259)
3. Bell, A. - University of Sheffield (ORCID 0000-0002-8268-5853)

Corresponding author

Phil Mike Jones

- p.jones2@derby.ac.uk

Declarations

- This paper is a pre–print of a manuscript accepted for publication with the journal *Quality and Quantity*

Funding

The authors wish to acknowledge funding received from the British Academy Skills Innovator Award, award reference SK150011

Availability of data and material

- Data used are available from Human Mortality Database (https://mortality.org).
- Code used for data analysis are available from: https://doi.org/10.5281/zenodo.6866402

Abstract

This paper explores changes in age–specific mortality risk across periods and cohorts during the twentieth century in the developed world. We use and compare two approaches—one

graphical (Lexis plots) and one statistical (an adapted Hierarchical age–period–cohort model)—that control out overall trends in mortality, to focus on discrete changes associated with specific events.

Our analyses point to a number of key global and local events in the Twentieth Century associated with period and/or cohort effects, including the World Wars and the influenza pandemic of 1918–19. We focus particularly on the UK but look at other countries where results are particularly noteworthy, either substantively or methodologically. We also find a decline in mortality in many western countries, specifically in the 1948 birth cohort, which may be associated with the development of post–war social welfare policies, the economic investment in Europe by the United States, the accessibility of antibiotics such as penicillin, and, in the UK, the founding of the NHS.

We finish by considering the advantages and disadvantages of using the two methods with different sorts of data and research questions.

Keywords

- Age-period-cohort
- Lexis surface
- mortality
- social epidemiology
- welfare state

Introduction

Age, period, and cohort methods attempt to disentangle three ways that societies can change over time: as individuals age, as time passes, and as birth cohorts replace one another. This paper compares two such approaches, one statistical and one graphical, using a worked example of mortality in the Twentieth Century.

Age–standardised mortality risk decreased during the twentieth century for most groups of people, in most places in the world (The World Bank, 2017). Separate from this overall downward trend there are annual deviations in mortality. Some of this deviation will occur naturally without any particular cause: everything varies. However, some of this deviation will be due to important influences or events whose effects are worthy of study.

In this paper we explore *deviations* in age–specific mortality risk during the twentieth century in a number of developed countries, considering both global and country–specific patterns. Our interest is not in the long–run downward trend in mortality but in the discrete changes in mortality seen as a result of global and national events. The twentieth century was notable for remarkable social and technological progress, as well as catastrophic global conflict, and we explore how these affected mortality. Many of these events were global, whilst others were geographically specific.

Deviations over time in mortality can occur in two ways. First, period effects affect everyone

at the time who is exposed to the event. Second, cohort effects occur when events affect specific generations of people born at a particular time in history. Throughout their lives individuals in affected birth cohorts benefit from or are hindered by these events that occurred in their formative years. Whilst age effects also exist (the risk of death increases broadly exponentially as age increases, an important consideration in ageing societies), there are fewer deviations caused by specific ages, with the notable exception of the 'accident hump' in males around age 20 (Heligman and Pollard, 1980).

We use this example of mortality to compare two approaches to APC analysis, one visual and one statistically modelled. First, we use Lexis plots to show the patterns in annual changes in age–specific mortality in all developed countries with data available, to see fine–grained differences between different combinations of period and cohort effects. Second, we use a modified version of the hierarchical age–period–cohort model (HAPC) (Yang and Land, 2006) in part to find the statistical significance of such patterns, and we compare different approaches to setting up these models. In both cases we do not consider long–running linear age–period–cohort (APC) trends; instead we focus only on deviations from those trends, avoiding the issue of the APC identification problem (Glenn, 2005). Our focus is predominantly UK–based, but we consider other countries where the results are particularly interesting, either substantively or methodologically.

This paper thus makes both substantive and methodological contributions. Substantively, we point to a number of key occasions in the twentieth century that had period and/or cohort effects, both global and geographically specific, including the effects of the World Wars, the flu pandemic of 1918, and the post–World War II social welfare policies, such as the establishment of the NHS in the UK. Methodologically, we present novel graphical and statistical techniques for finding discrete APC effects whilst removing long–run effects. We compare the advantages and disadvantages of each approach, and consider how the approaches can potentially be combined into a broader methdological framework.

Literature

Age, period, and cohort effects on mortality

Suzuki (2012, p. 452) outlines the following fictional dialogue to illustrate the difference between age, period, and cohort effects:

> 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]

Age is the measurement of time passed since birth. Period is 'historical time' when the measurement was taken, so represents a snapshot of all people, of all ages, in the study at that instance (Goldstein, 1979, p. 19; Suzuki, 2012, p. 452). A cohort refers to:

> . . . those individuals (human or otherwise) who experienced a particular event during a specified period of time. The kind of cohort most often studied by social scientists is the human *birth cohort*, that is, those persons born during a given year, decade, or other period of time (Glenn, 2005, p. 2, original emphasis).

Ryder argues that "[e]ach cohort has a distinctive composition and character reflecting the circumstances of its unique origination and history" (1965, p. 845).

Each of age, period, and cohort can have effects on individuals. Considering mortality as the outcome of interest, an age effect might mean that the risk of death increases or decreases as a person gets older. A period effect could be caused by an event that affected people at a particular snapshot in time, for example a war, disease, or economic recession causing increased likelihood of death across individuals of all ages at that point in time. A cohort effect might manifest as subsequent cohorts having incrementally lower mortality risk than earlier cohorts, perhaps because of improvements in living standards in their formative years. However, it could also occur as a result of events which have an impact on people in their formative years — an effect that stays with those people throughout their lives. For this paper we are primarily interested in period and cohort effects, since the (increasing) effect of age on mortality is relatively well established, and there are fewer reasons to expect discrete effects that apply to most specific age groups (as opposed to long–run gradual changes over the life course).

We anticipate being able to detect period effects for significant events such as war, famine, or epidemic because more deaths are observed at the time of the event. Literature on developmental plasticity (Gluckman, Hanson and Buklijas, 2010) suggests cohort effects on mortality over the life course are also plausible. Developmental plasticity as a theory is primarily adopted and advanced through the Developmental Origin of Health and Disease (DOHaD) hypothesis and life course epidemiology (Hanson and Gluckman, 2016). These hypothesise that an individuals' developmental environment affects the structure, physiology, and function of organs and systems throughout the individual's life (Fall *et al.*, 1995; Wadsworth and Kuh, 1997; Ben–Shlomo and Kuh, 2002; Ben–Shlomo, Cooper and Kuh, 2016; Hardy and Tilling, 2016; Newman, 2016). 'Better' *in utero* and early–life environment leads to longer, healthier lives, while lower quality early–life environments lead to shorter, less healthy lives (Hertzman, 1999, p. 85). For instance, links between prenatal malnutrition and low birth weight, neonatal mortality, cardio–vascular disease, coronary heart disease, ischaemic heart disease, and hypertension have been demonstrated (Hales and Barker, 1992).

Under this paradigm a stimulus—such as economic circumstances, sudden improvement in healthcare, and so on—can have biological and physiological effects on the individual that last

throughout their life course, which has been shown to affect their morbidity and mortality. If the same stimulus affects a large number of individuals from the same or similar cohorts in the same way, patterns of mortality will be seen throughout the lives of the cohort members as they age.

There may also be cohort effects which do not become apparent at birth, but later in life. This could be because formative years occur long after birth; for instance, with smoking uptake the age of exposure is much older than birth (Schöley and Willekens, 2017, p. 633). It could also be because cohort effects are delayed and only appear long after exposure. As such there may be a higher risk of psychological and physiological trauma among older cohorts which may manifest as differences in mortality later in their life course, with earlier life events being the cause.

Events that affected mortality in the twentieth century

A number of significant events occurred in the twentieth century, both globally and nationally, that are likely to have affected population mortality in the developed world, both as period effects and as cohort effects. Here we briefly discuss four that we see as particularly important: World War I; the 1918–19 influenza pandemic; World War II; and the enormous social welfare progression that occurred in many countries following the end of the second world war, including the formation of the National Health Service (NHS) in 1948 in the UK.

We would expect a period effect increase in mortality associated with the First World War of 1914–1918. For the most part we would expect this to be limited to military personnel in countries participating in the war, but we might expect to see a period effect in the civilian population in countries with high civilian casualties, such as those in continental Europe. In other countries such as the UK, civilians were not directly affected by the conflict but effects of deteriorating environmental conditions may be detectable. A cohort effect among those born during the conflict is also plausible, for example because of poor maternal nutrition, exposure to disease, maternal stress, or otherwise inadequate early–life health care as a result of the conflict.

It is also possible that the reverse could be true. There is evidence that war, or rather the threat of war, led to improvements in public health in the early twentieth century, especially for expectant mothers and young children, as the state sought to ensure sufficient numbers of healthy combatants should war break out (Dwork, 1987). Similarly, Winter and Prost argue that the Great War resulted in *lower* mortality among British males aged over 40 (2005, p. 160). In sum, World War I likely had multifaceted effects on mortality, both instantaneous (period) and long–run for those in their formative years at the time (cohort).

The 1918–19 influenza pandemic is likely to result in detectable period effects as recent estimates have put the number of deaths from this disease at 50 million worldwide, or approximately five per cent of the global population (Patterson and Pyle, 1991; Johnson and Mueller, 2002). Approximately 250,000 died in the UK. Cohort effects for those born during the outbreak (1918 to early 1919) are also well established in the literature. Increased incidence of cardiovascular disease (Mazumder *et al.*, 2010), decreases in life expectancy at birth (Noymer and Garenne, 2000), and increases in socio–economic deprivation (Almond,

2006) have been demonstrated in cohorts in the United States born with prenatal exposure to the disease. Of course, it is difficult to tell apart cohort effects of the war and the influenza pandemic given their temporal proximity. In the case of period effects the different age and gender of those theorised to be affected by each give a clue as to what caused each (with young men most likely to be affected by the war, whilst the effects of the influenza pandemic affected both men and women, and a broader age range).

Even populations that diverged following the influenza pandemic, such as those of East and West Germany, show remarkably similar mortality 'scars' (Minton, Vanderbloemen and Dorling, 2013) in cohorts born in 1918–1919:

> ...those born in early 1919 who were exposed prenatally to the most virulent phase in the Fall of 1918, had lifetime defecits in economic productivity and in education, as well as excess work disability, which suggests developmental impairments or lifetime health issues (Mazumder *et al.*, 2010, p. 26).

Following the First World War, both female and male children born in the group of cohorts between approximately 1926 and 1945 have been found to experience a rapid improvement in mortality, which slowed for subsequent generations born after 1945 (Willets, 2004). The cause of this 'golden' cohort effect is not known, but it is hypothesised that a combination of factors led to their improved mortality compared to preceeding and subsequent generations. Most in this birth cohort were not old enough to have been involved in World War II, and post–war rationing led to an improved diet for this cohort. They also likely benefited from the development of the welfare state, declining smoking prevalence, and being born during a period of relatively low fertility (Willets, 2004).

We anticipate a detectable period–related increase in mortality during World War II for both military and civilian populations. Civilian populations are likely to be more affected than in World War I, due to the changing nature of warfare, specifically the increase in bombings of civilians made possible by advances in technology. However, as with World War I, we would expect the larger effect to be found among young men.

As well as the period effects there could also be cohort effects among individuals born during World War II in some contexts. Specific events such as the Siege of Leningrad and the Dutch *Hongerwinter*, where significant numbers of individuals perished, have been shown to be associated with period and cohort mortality increases in the affected populations. Survivors of the Siege of Leningrad had a significantly higher risk of dying from breast cancer (Koupil *et al.*, 2009), ischaemic heart disease, or stroke (Sparén *et al.*, 2004) compared to those born during the same period who were not exposed to the seige. Similarly survivors of the Dutch *Hongerwinter* who were part of the Dutch Famine Birth Cohort Study were more likely to have blunted cardiovascular and cortisol stress responses, which are in turn associated with a range of adverse health outcomes (Carroll *et al.*, 2017). Other studies have shown a lack of effect on other morbidities, however: participants in the Leningrad Siege study did not appear to be at greater risk of diabetes (Stanner *et al.*, 1997), whilst the risk of coronary heart disease may be mediated by obesity in adulthood (Stanner *et al.*, 1997, para. 17).

Following the Second World War, many Western nations implemented a number of progressive policies aimed at improving population health and wellbeing. In the UK, these covered a

range of social issues, such as National Insurance, housing, education, and child welfare, as well as the nationalisation of a number of key industries. Perhaps the most prominent example was the formation of the NHS in 1948 (Rivett, 1998). This involved a comprehensive reorganisation and rationalisation of medical provision, and treatment became free at the point of access for all. This included previously marginalised groups, such as working–class women, for whom treatment had previously been limited due to the prohibitive cost (Webster, 2002). A detectable period effect of reduced mortality is plausible at this time; although no new treatments were immediately developed with the founding of the NHS, existing treatments were suddenly accessible to everyone regardless of ability to pay.

A cohort effect is also plausible for cohorts born around this time in the UK in particular. Limited availability of antenatal and perinatal care—critical periods for the child—prior to the introduction of the NHS is likely to have adversely affected the developmental trajectory of many children. With the NHS, pregnant women could now access antenatal care, for the first time often provided by general practitioners, and give birth in hospital. Increasing the opportunities for intervention at critical periods *in utero* could result in improved health and reduced mortality over the whole life course for the infant. Similar effects could be found in other countries, associated with other social welfare policies introduced at a similar time.

Moreover, exposing pregnant mothers to the health care system through antenatal care and a hospital birth may have the cultural effect of 'normalising' the use of medical care. If this contributed to earlier detection of disease or illness this cultural effect could have benefits to the mortality of children born under the NHS throughout their lives, for whom seeing a doctor became part of their early socialisation. The NHS, along with other public health improvements in the UK and elsewhere, are likely to have resulted in lower mortality for people born in those post–war years onwards.

Methods

Mortality data for 40 countries[1] with data available for the twentieth century were obtained from *The Human Mortality Database* (University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany), 2017). This provides full demographic data on mortality rates, deaths, and populations, for all ages and for all years since at least 1900 for many developed countries (although the data goes further back it is less reliable, so we have not used this older data). Our aim is to use this data to analyse discrete, non–continuous changes in mortality rates, net of any long–run improvements in mortality.

Here we present two methods: one visual and one statistical. First, we use Lexis plots of mortality change (the change in the mortality rate for a given age from one year to the next). We use mortality change for a given age, rather than mortality, in order to remove long–run changes in mortality over time. Lexis surface diagrams have long been used in demography to depict cohort information as well as period and the event of interest (Derrick, 1927; Kermack, McKendrick and Mckinlay, 1934; Carstensen, 2006; Healy, 2018). Lexis

[1]Countries included were those with consistent mortality data available for the whole of the twentieth century. The full list of countries is available in the online appendix.

diagrams were produced using the `Lattice` package for `R`, version 0.20-45 (Sarkar, 2008). These plots were made for all countries in the Human Mortality Database; although only some are shown in this paper, the rest can be found in the online appendix.

Interpreting Lexis diagrams, especially using them to disentangle age, period, and cohort effects, in the presence of a 'linear drift' is problematic and therefore controversial as the linear drift tends to account for the majority of variation in mortality (Murphy, 2010, p. 371). However, this is not a problem here, as we focus on non–continuous, discrete effects, and long–run changes in mortality are removed by modelling *change* in mortality rates, rather than the mortality rate itself.

An additional advantage of this approach is that it allows us to see period and cohort effects that only affect specific age groups. However, as a descriptive approach it cannot quantify the level of uncertainty around those effects given the data that we have, and often patterns are difficult to see when there is a lot of random variation. What it does do, is allow researchers to identify possible patterns and then choose a modelling approach that suits the quantification of those patterns.

One approach that could be taken is to adapt a Lee–Carter style model to allow it to model similar APC trends. In general, Lee–Carter models have been used for the purpose of forecasting evolving mortality rates, and so are often used by actuaries and demographers where that is the focus of interest. Where these models have been extended to allow the modelling of, for instance, cohort–type features (see (Renshaw and Haberman, 2006)) this has generally been for the purpose of evaluating and validating forecasting models, rather than those features being the primary purpose of fitting those models. An effective strategy for comparing out–of–sample fit between models is demonstrated by Hyndman and Koehler (2006) and Pascariu, Lenart and Canudas–Romo (2019), and we consider such approaches important for comparing demographic forecasting approaches. However, in practice a model with an *a priori* specification of structure and variables which correspond directly to readily interpretable sociological or epidemiological quantities of interest can be immensely valuable for researchers whose aims are to understand the processes which gave rise to the observations, even if the in– or out–of–sample fit of the model is poorer than for models with less directly interpretable parameters[2]. As such we do not take this approach, aiming instead for a model which explicitly parameterizes and identifies APC features. These approaches are complementary but distinct, most notably in that our aims and framing are more sociological and epidemiological than actuarial.

Instead we use modified hierarchical age–period–cohort (HAPC) models constructed for countries or sub–regions of interest that control for the linear trends in APC, allowing us to focus on discrete, non–continuous change.

[2]As an example of this a difference–in–differences (DiD) model for time series data comparing intervention and control populations is often more valuable for users than a model based on smoothed splines or polynomial terms, even if the latter leads to improved fit, because for a DiD model the intervention effect is explicitly modelled and interpretable to the user. Similarly we argue that our model specification, which explicitly includes age, period, and cohort terms to estimate, may be especially valuable for understanding the substantive processes which may have given rise to the data observed, even if a Lee–Carter based model has superior fit.

The original version of the HAPC model (Yang and Land, 2006) treats the age effect as a fixed effect polynomial, with the period and cohort effects as cross–classified random effects. The model can be specified as (for a continuous outcome variable):

$$y_{i(j_1,j_2)} = \beta_{0j_1,j_2} + \beta_1 Age_{i(j_1,j_2)} + \beta_2 Age^2_{i(j_1,j_2)} + \epsilon_{i(j_1,j_2)} \tag{1}$$

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

$$\epsilon_{i(j_1,j_2)} \sim N(0, \sigma^2_e), u_{1j_1} \sim N(0, \sigma^2_{u1}), u_{2j_2} \sim N(0, \sigma^2_{u2}) \tag{3}$$

where $y_{i(j_1,j_2)}$ is the dependent variable (in our case age–cohort specific mortality from the previous year) for individual (or in our case age–period measurement) i in cohort group j_1 and year of measurement j_2. u_{1j_1} represents the cohort random effects and u_{2j_2} the period random effects, both of which are assumed to be normally distributed, as is the level one residual term ($\epsilon_{i(j_1,j_2)}$).

When considering age, period, and cohort there is a problem that by knowing two variables we can perfectly predict the other: age equals period minus cohort, so the three variables have only two degrees of freedom. This is referred to as the 'identification problem' (Glenn, 2005; Bell and Jones, 2013). The HAPC model (Reither, Hauser and Yang, 2009), as well as the 'intrinsic estimator' (Yang and Land, 2006; Yang *et al.*, 2008), are attempts to statistically separate the three compondents. Unfortunately both of these models have been shown to apportion linear trends in ways that often do not fit with the true data generating processes (DGPs) (Luo, 2013; Bell and Jones, 2014a, 2014b, 2018; Luo and Hodges, 2015).

In our case, however, we are interested only in non–linear period and cohort stochastic fluctuations, once the age, period and cohort long–run trends are controlled. As such we can control for these trends in the fixed part of the HAPC model, leaving only discrete deviations in the random part of the model (see Chauvel, Leist and Ponomarenko (2016)). Whilst we cannot control for all three of APC in the fixed part of the model because of the identification problem, controlling for two of APC will control out the linear component of the third by default. Our first version of this model can therefore be specified as follows:

$$MortalityChange_{i(j_1,j_2)} = \beta_{0j_1,j_2} + \beta_1 Age_{i(j_1,j_2)} + \beta_2 Age^2_{i(j_1,j_2)} + \epsilon_{i(j_1,j_2)} \tag{4}$$

$$\beta_{0j_1,j_2} = \beta_0 + Period_{j_1} + u_{1j_1} + u_{2j_2} \tag{5}$$

Here $MortalityChange_{i(j_1j_2)}$ is the change in mortality rate for a specific age group, in comparison to the previous year, for age–year cell i in year j_1 and birth year j_2. This is the same as Equations (1) to (3), but with the addition of a Period term in the fixed part of the model, which means all APC linear trends will be absorbed from the period and cohort residuals into the fixed part of the model. However, because we are using a measure of

mortality change (as opposed to the number of deaths) we would not expect to see much in the way of linear trends in any case.

A downside of this approach is that because we are using age–period cells as our units of analysis, we cannot account for the differences in size of the different groups, and so our measures of uncertainty will be somewhat inaccurate (a cell of 10 people is treated the same as a cell of 10,000 people). An alternative approach would be to model the number of deaths, controlling for the size of the population. To do this we use a Poisson model for the number of deaths in a given age–year cell. We additionally use an offset of the expected number of deaths given the population size of that cell, if deaths were distributed evenly across the population. The inclusion of the offset means that we are effectively modelling the mortality rate by taking account of the population size in our estimation of uncertainty (Jones *et al.*, 2015). Thus, our model is specified as follows:

$$Deaths_{i(j_1,j_2)} \sim Poisson(\pi_{i(j_1,j_2)}) \tag{6}$$

$$Log_e(\pi_{i(j_1,j_2)}) = Log_e(E_{i(j_1,j_2)}) + \beta_{0j_1,j_2} + \beta_1 Age_{i(j_1,j_2)} + \beta_2 Age^2_{i(j_1,j_2)} \tag{7}$$

$$\beta_{0j_1,j_2} = \beta_0 + \beta_3 Period_{j_2} + u_{j_1} + u_{j_2} \tag{8}$$

$$u_{j_1} \sim N(0, \sigma^2_{u_1}); u_{j_2} \sim N(0, \sigma^2_{u_2}) \tag{9}$$

$$Var(Deaths_{i_(j_1,j_2)}|\pi_{i_(j_1,j_2)}) = \pi_{i_(j_1,j_2)} \tag{10}$$

There are a number of key differences between this model and that specified in Equations (1)–(3). First, as stated above, we use a Poisson model with log link function, meaning we assume that the level 1 variance is equal to the estimated mean deaths ($\pi_{i(j_1,j_2)}$), and we model deaths with an offset, Expected Deaths $E_{i_(j_1,j_2)}$ so we are effectively modelling death rates (see Jones *et al.*, 2015). We also include $Period_{j2}$ in the fixed part of the model, as in Equation (5). Between this and the $Age_{i(j1j2)}$ variable, we are controlling for all linear effects of age, period, and cohort because of the exact dependency between the three terms[3].

In both of the models above we cannot trust the estimates of β_1 or β_3 (because they will incorporate any cohort linear effects if they exist in the DGP), but we are not particularly interested in their estimates. We can say that u_{1j_1} and u_{2j_2} will be accurate estimates of deviations from the long–run trends in periods and cohorts (whatever they are), and we can be confident (linear) APC trends will not be included in those estimates. However, there may be some long–run, but not linear, trends remaining in these residual estimates which

[3]These models assume the Poisson (level 1) residuals are not overdispersed - we would encourage researchers to check this when using the Poisson link function

should not be interpreted as their meaning will depend on the trends controlled out in the fixed part of the model.

We removed data for individuals aged 91 years and over from our analysis, and removed data for birth years before 1900. In both cases there were significant problems with the data prior to this date and at older ages, as well as artefacts from imputation. See Section 5.4 of the HMD methods protocol for methods used consistently in the database for older populations aged 90+ (Wilmoth *et al.*, 2021). The models were fitted in MLwiN (Charlton *et al.*, 2017) using `R` and the `R2MLwiN` package (Zhang *et al.*, 2016) using MCMC (Browne, 2017), with a 500,000 burn–in and 1,000,000 iterations.

Full HAPC results tables can be found in an online appendix. It should be noted, however, that the APC fixed terms should not be interpreted because of the APC identification problem. Full replication code can also be found in the online appendix.

Results

In this section we present findings predominantly from England and Wales, with comparisons with other countries where useful, as a case study. Figures for all countries are available in the online appendix. The performance of the models for other countries is comparable to those for England and Wales. We have also written a short comparison of three countries, which can be found in the paper's online appendix.

Figure 1 shows a Lexis surface for mortality change in England and Wales, with blue and green representing a decline in mortality, and red and orange representing an increase in mortality on the previous year for a given age of person. Cohort effects appear as diagonal 'scars' (Minton, Vanderbloemen and Dorling, 2013), emanating through age–time upwards and rightwards from the affected birth cohorts, whilst period effects appear as vertical scars.

A red line followed by a blue line might represent temporary excess mortality caused by an event such as the influenza pandemic. A blue line followed by a red line would represent a temporary decrease in mortality, that later returned to its previous level. A mild winter might exhibit such an effect if excess winter deaths are lower than neighbouring years. Figure 1 shows evidence of both period and cohort effects in England and Wales. Whilst there are some notable differences between this figure and the equivalents for other countries, this presents a good starting point given the fullness of data and some key features that are present in other countries as well.

In addition to these effects, it is possible to see longer–lasting changes in age–specific mortality change, where a decrease in mortality change is not followed by an increase, and vice–versa. A lone red line represents a long–term increase in mortality rate, for example caused by an enduring economic crash and recession. A blue line without a corresponding red line would represent a long–term decrease in mortality rate, for example due to a medical advancement.

There are some cohort effects visible on the Lexis surface plot (Figure 1) for females and males born approximately every ten years between approximately 1840 and 1900 (upper left

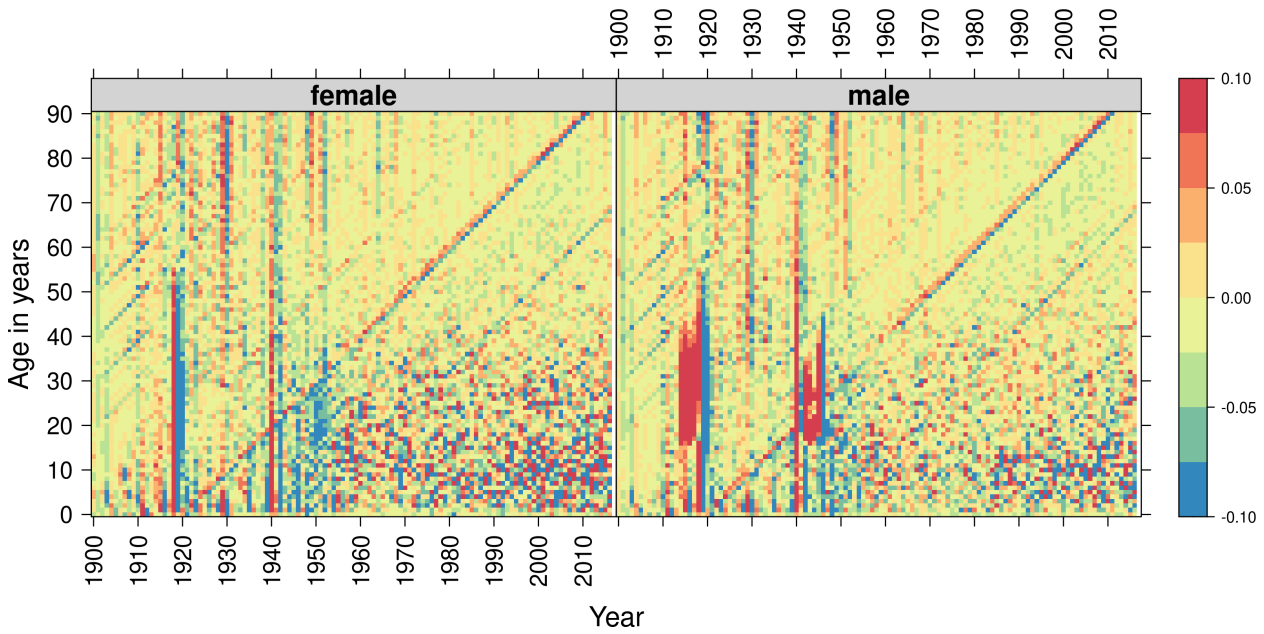

Figure 1: England and Wales total population Lexis surface plot for annual change in age-specific mortality. Red signifies worsening mortality compared to previous year; blue signifies improved mortality

quadrant). We believe these are spurious and a result of data imputation from the decennial census, partly because they are not detected in the HAPC models, which we discuss below.

The equivalent HAPC model for England and Wales produces year (period) and cohort residuals. These are shown in Figures 2 and 3 respectively for continuous-Y models with change in mortality as the outcome variable. Figures 4 and 5 respectively show the residuals of the Poisson model with deaths as the outcome variable, both for the change in mortality rate models and the death count Poisson models. The period residuals can be interpreted as the deviation in a given year from the overall linear period trend, which is controlled out in the fixed part of the model. The cohort residuals can be interpreted as the deviation for a given birth cohort, again from the overall linear cohort trend.

It should be noted that, for the Poisson models, there are continuous trends visible in both Figure 4 and 5 which have not been completely controlled–out in the fixed part of the model, including a rather dramatic increase in mortality seen in the later cohorts in Figure 5. These are continuous effects that are non–linear and so were not controlled (for example, quadratic and cubic effects). Given these are not interpretable without knowing what the linear portions of these effects are, these should not be interpreted, and only discrete, sudden changes around these continuous curves should be analysed. Their presence is perhaps a disadvantage of the approach when the outcome includes non–linear continuous trends, unless an appropriate functional form can be used to absorb those trends. Because the outcome has been detrended by modelling year–by–year change in the other models, this is not a problem. However, in both models, a number of features can be identified which we discuss now.

In Figures 1, 3, and 5 there is a noticeable cohort effect with increased mortality in the cohort

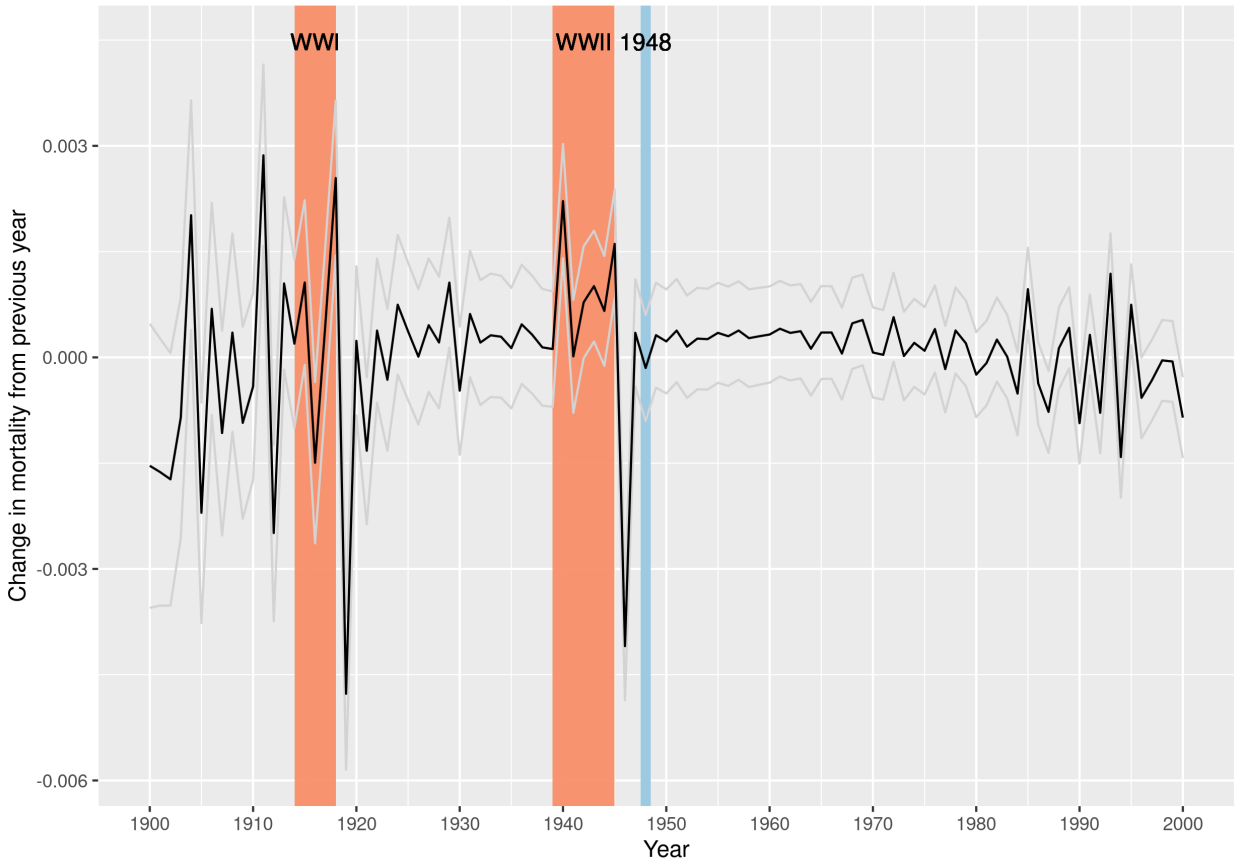

Figure 2: Plot of year (period) residuals in England and Wales for males, from the adapted continuous-Y HAPC model of change in mortality rate. The residuals can be interpreted as the deviation from the overall (and unknown) linear period trend.

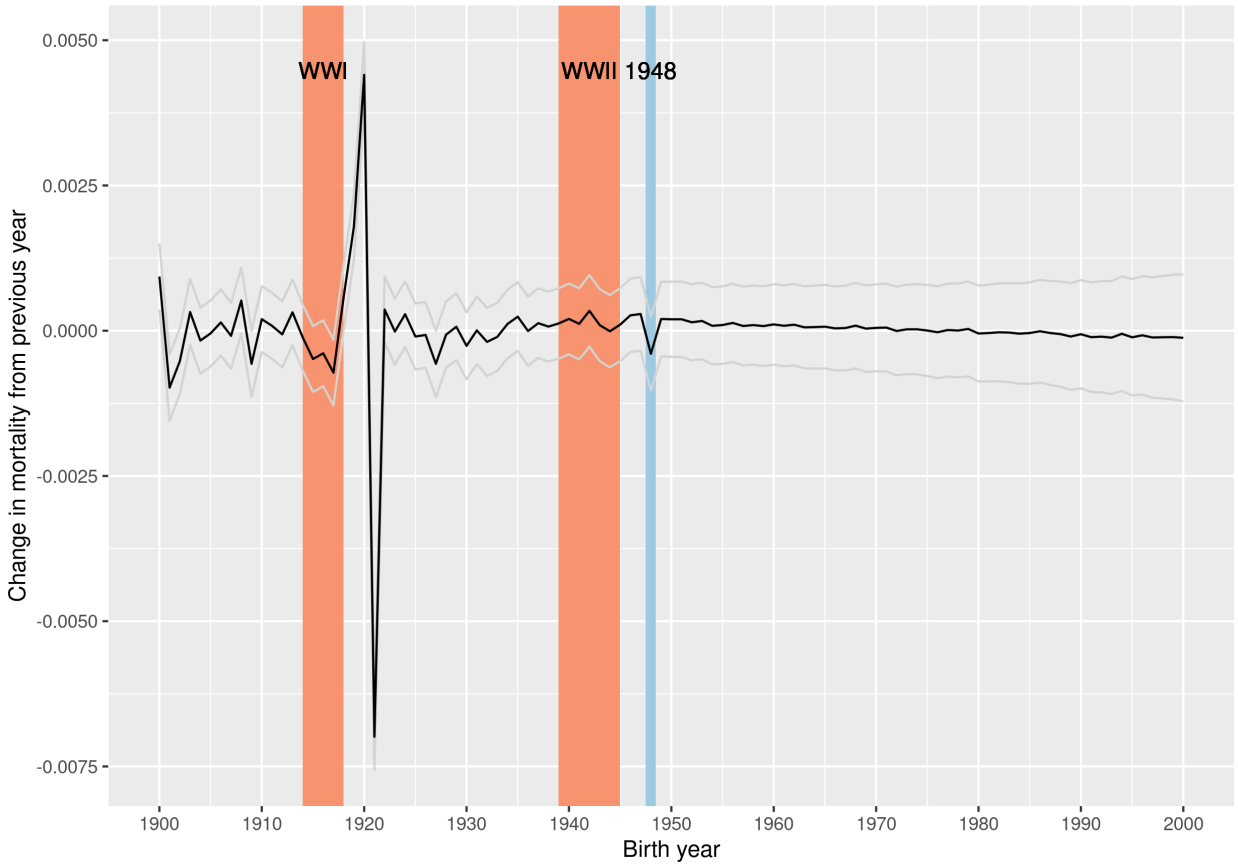

Figure 3: Plot of birth year (cohort) residuals in England and Wales for males, from the adapted continuous-Y HAPC model of change in mortality rate. The residuals can be interpreted as the deviation from the overall (and unknown) linear trend in cohorts.

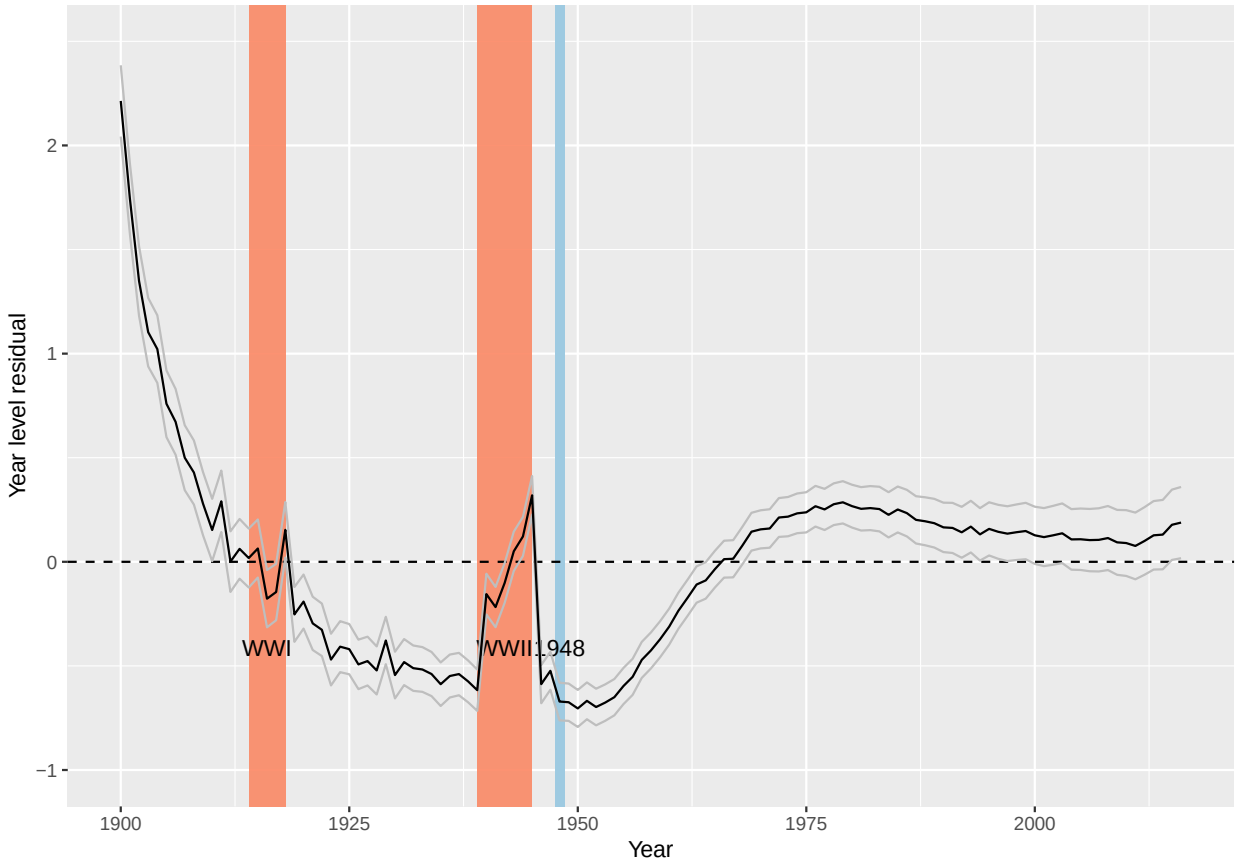

Figure 4: Plot of year (period) residuals in England and Wales for males, from the adapted HAPC Poisson model of mortality rate. The residuals can be interpreted as the deviation from the overall (and unknown) linear period trend.

Figure 5: Plot of birth year (cohort) residuals in England and Wales for males, from the adapted HAPC Poisson model of mortality rate. The residuals can be interpreted as the deviation from the overall (and unknown) linear trend in cohorts. The strong rise in later cohorts is likely an artefact of the non-linear continuous effects; it should not be interpreted.

born in 1918. Whilst this could be in part due to World War I, given the lack of effect for those born earlier in the war, and the similar effects found for both males and females, it seems likely that this is primarily the result of the 1918–19 influenza pandemic. This effect is noticeable in that it appears almost universal across all countries with sufficient data quality to identify such an effect, including countries that were less affected by the influenza outbreak, for example Australia where the pandemic affected the country later and to a lesser extent than European countries (Curson and McCracken, 2006).

A period effect is also clearly visible around the year 1918 in females and males under the age of about 55, in all countries with data going back that far. A sharp increase in mortality is followed by a commensurately sharp decrease, indicating a sudden increase in deaths caused by the pandemic which then returned to the previous level. For males there is an additional effect on mortality in the preceding years for those between ages 15 and 35 in Great Britain and Italy. This high increase in mortality is concentrated in young men during the entirety of the First World War, reflecting the increasing deadliness of this conflict for military personnel. A number of other countries that we might expect similar effects for (for example France or Germany) have missing data at around the time of World War I.

Literature on the cohorts born around 1931 (1926–1945) suggests it may have been possible to find a positive effect of being born around these times (Willets, 2004). However, we do not see clear evidence of such a cohort effect.

Another period effect appears around the Second World War. In Great Britain the population from birth to old age exhibits higher period mortality in the year around 1940, contemporaneous with *The Blitz*. This suggests either civilians suffered greater exposure to the conflict or environmental conditions worsened during this time, or both. Although that specific pattern does not appear in other countries, some countries involved in World War II do show increases in mortality for young men. This seems more extensive than the equivalent effect of World War I, affecting in particular Finland, Great Britain, Italy, and the Netherlands (again there was limited data for France and Germany).

The Netherlands also appears to show an increase in mortality associated with World War II for the whole population. Based on the plot for The Netherlands a decline in period mortality around World War II is detectable in the Dutch population (Figure 6). The Lexis surface plot shows increased mortality for all ages and both sexes during the Second World War, but for a greater time period beginning in 1940, and in particular in 1945. This suggests a greater exposure to the conflict for the Dutch civilian population than the UK population or other countries with mortality data. As The Netherlands was occupied from May 1940 until 1944–1945 this is to be expected.

Of particular interest is the mortality rate in the year 1945, where the increase in mortality spans a much greater age range. By the end of 1944 much of The Netherlands south of the Waal was liberated, but areas north of the Waal, included the densely populated coastal provinces, remained occupied until 1945. It is these areas that suffered the *Hongerwinter* (Warmbrunn, 1963, pp. 14–17). Therefore, it is possible that much of the increase in mortality rate observed in 1945 could be because of the famine in occupied areas of The Netherlands, before the mortality rate recovered following the end of the Second World War.

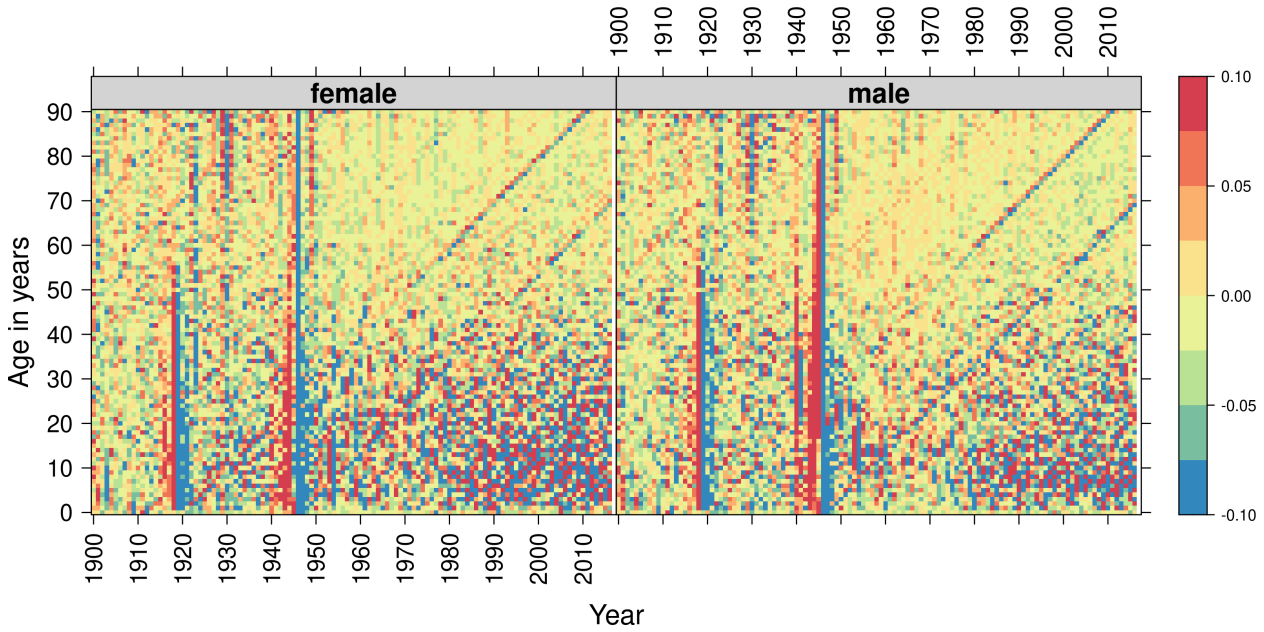

Figure 6: Lexis surface plot of mortality for The Netherlands

Both males and females show improved mortality rates in the years immediately following the end of World War II. There does not appear to be evidence of a cohort effect for those born during World War II in any countries (either positive or negative).

A less obvious, but nonetheless present, change in cohort mortaltiy rate is observed among those born in the year 1948 in a number of countries. In England and Wales this is visible in Figures 1 (a diagonal line originating from 1948), 3 and 5. Similar effects are visible in Canada (Figure 7) and the USA (Figure 8). In each case a small reduction in mortality is evident and this is not followed by a comparable increase in mortality in the following cohorts. The effect is small, but does suggest people born in those countries in 1948 and later experienced a lower mortality rate throughout their lifecourse than individuals born even just one year previously.

The obvious change that occurred in England and Wales, as with the rest of the UK, at this time was the formation of the NHS. If the NHS is indeed the cause of this improvement in cohort mortality, the implication is that being born under the NHS institution gave an advantage in terms of mortality. Whilst those born just prior to 1948 lived the majority of their lives under the NHS, they did not appear to receive this benefit.

This could be because pre–natal and early life care are particularly important in improving mortality for individuals throughout their lives. Alternatively (or additionally) the NHS may have had a cultural effect on those born under it—and their parents—making them more likely to seek treatment through it throughout their lives.

Whilst the localisation of this effect to 1948 implies the NHS is important it is not the only possible explanation. The winter of 1946–1947 in Europe was especially harsh with fuel and food shortages reported from late January 1947. If the severity of this winter affected the

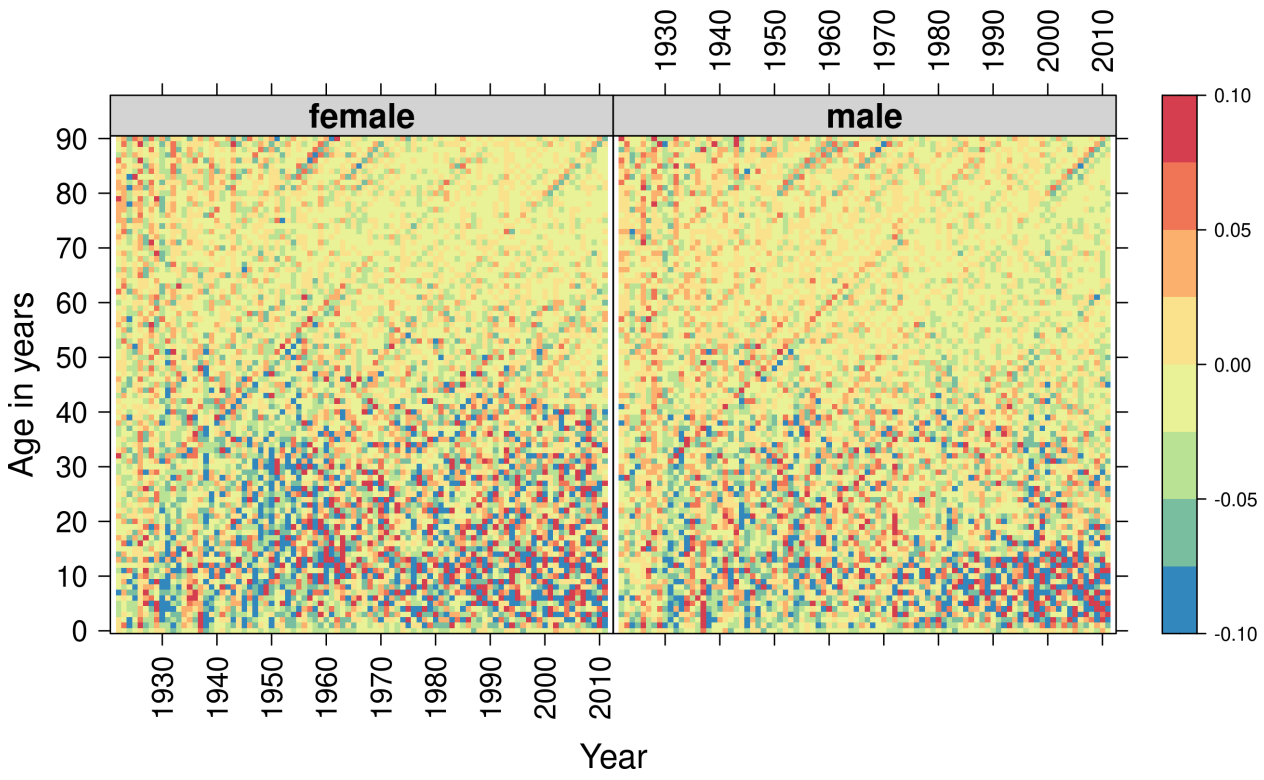

Figure 7: Canada Lexis surface plot

nutrition available to pregnant mothers it may have also affected the later morbidity and mortality of their children born up to early 1948. It is possible the lower mortality in 1948 is partially explained by the returning to background levels of mortality after an increase in late 1947. However, if this were the case, we would expect to see a paired banding of constrasting colours (red, then blue) as seen in the case of the 1919 birth cohort, rather than the single blue line seen for the 1948 cohort.

The presence of the 1948 effect in countries other than England and Wales perhaps suggests a more global explanation. First, all of these countries implemented health and welfare policy after the war, and the finding could be a result of a more general improvement in health and welfare provision as a result of these. For instance, in the UK the formation of the NHS was situated within a context of high employment, the implementation of welfare policies such as the National Assistance Act (1948) – which was itself an addition to the National Insurance Act 1946 which introduced social protections, nationalisation of energy and rail transport, and substantial financial aid from the United States in 1946 and 1947 (Medlicott, 1967; Hill, 1970, p. 291). Similar social welfare improvements in other countries may have led to similar improvements in mortality. However, this does not provide a clear reason why this would happen specifically in 1948, and not the years immediately before or after.

Second, penicillin was first produced in bulk during the early 1940s, but became more accessible to patients as costs were driven down during the mid– to late–1940s. It is possible that penicillin became more accessible in 1948 in the UK, as well as in the US and other countries, leading to reductions in cohort and period mortality. Penicillin could have been

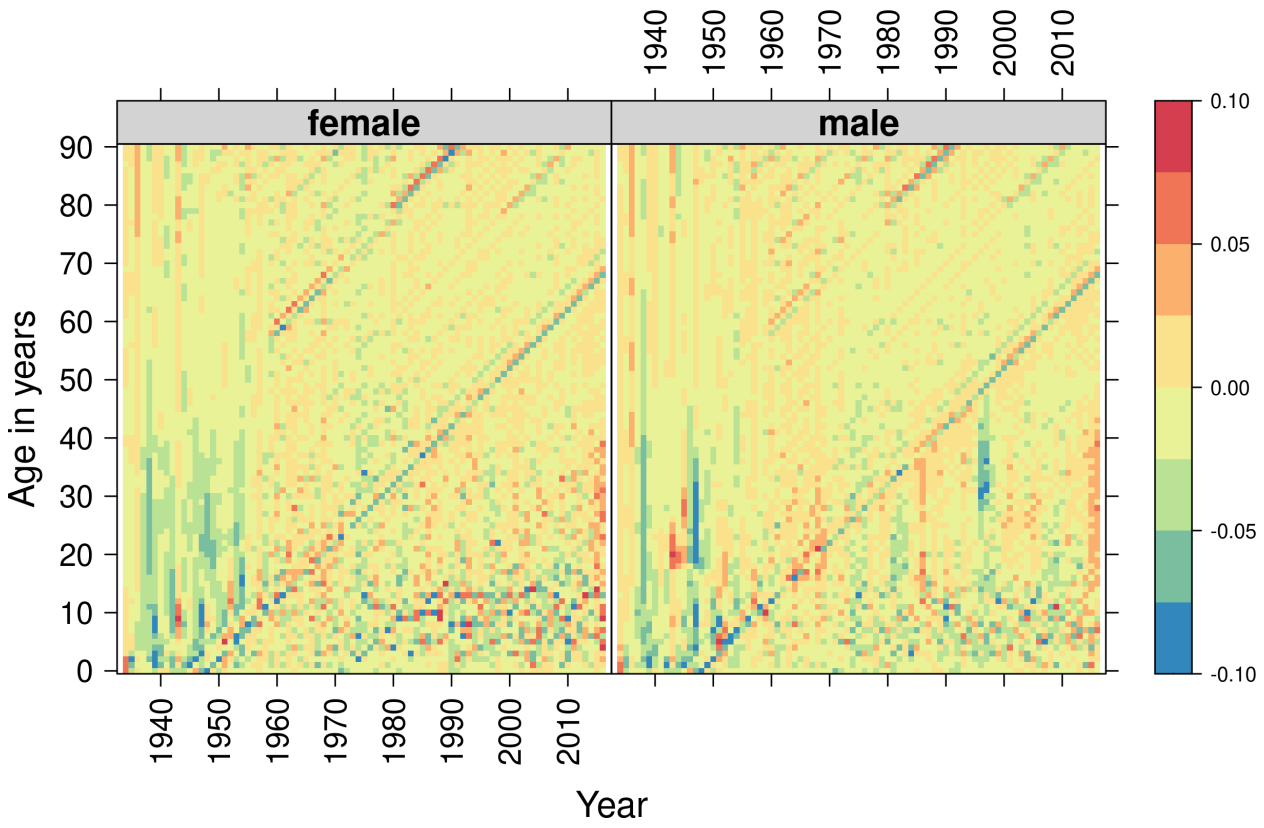

Figure 8: USA Lexis surface plot

used both to treat young children, and also to treat new mothers (particularly for postpartum infections), improving survival rates of mothers in labour and thus, plausibly, the life outcomes of their children. However, penicillin also became more accessible at the same time in countries such as Portugal (Bell, Rui Pita and Pereira, 2017) which showed an increase in cohort mortality in 1948 (see appendix) suggesting, if penicillin availability were partly responsible for decline in cohort and period mortality, the picture is complicated by other factors. This work is very much exploratory and more work would be needed to confirm these hypotheses.

Comparing the approaches: which works best?

For the most part the two methods produce results in agreement with each other: they both find specific period and cohort effects relating to events, such as the 1919 flu pandemic and the World Wars. It should also be noted that both approaches are intrinsically exploratory—so neither should be used to test specific hypotheses about the presence of particular cohort or period effects. Rather they provide opportunities to explore the temporal patterns in the data. In that sense both methods 'work'. However, it is clear that there are advantages and disadvantages to both that are worthy of discussion.

The Lexis plots have the advantage of being unconstrained by the model parameters that are set. They allow for unanticipated interactions between APC, as seen for instance with the period effects of the World Wars which affected only a particular age group and gender. The Lexis plots also do not rely on some of the assumptions that the models are constrained by, for example normality of residuals or linearity of main effects. The main limitation of the Lexis plots in comparison to the modelled approach is the lack of information about uncertainty in the results that are produced. Where we are using population–level data, as here, this is less of a problem since there will likely be little uncertainty in the results found. With other data, for example survey data, this is likely to be more of a problem with results found that are actually caused by chance alone, and patterns missed in the 'long grass' of natural variability. There is also scope to combine the effects found in different countries on to single Lexis 'curvature' plots, allowing for interesting cross–national comparison (Acosta and van Raalte, 2019).

Conversely the modelled approach does produce measures of uncertainty: confidence intervals relating to the period and cohort residuals, although these are potentially less accurate when the assumptions of the model are problematic. This is particularly evident in the Poisson models where continuous trends remain even after the inclusion of the linear APC terms in the fixed part of the model. It would seem sensible, therefore, to only use these models where the dependent variable is lacking in such trends, or can be de–trended by calculating change as we have done in our Normal model. A further disadvantage is a lack of flexibility in comparison to the Lexis approach: any interactions for example would need to be explicitly modelled, whereas these can be explored more readily with a Lexis plot.

In general, certainly for this data, we find the Lexis plots are more effective than the HAPC model for the exploration of the data that we are using them for. However, with other data and outcomes which are, for instance, noisier—making it difficult to find trends in the Lexis plot—the HAPC model might be more appropriate if there are no trends in the residuals.

An approach that potentially combines the two approaches is outlined by Minton (2021). There a Lexis plot could be used to identify key features in the data which then could be explicitly modelled. The model residuals can then be plotted in a Lexis plot to see the extent to which the model 'explains' those features. Such a model could, in fact, incorporate features of Lee–Carter style models where the data deems them appropriate. Of course the model is then only as good as the researchers' reading of the data, and the features of the Lexis plot would still need to be understood substantively. However it provides a potentially useful way to formalise features in the data seen visually, in model form.

Conclusions

This paper has explored period and cohort effects on mortality in developed countries during the twentieth century, using Lexis surface plots and hierarchical age–period–cohort models. The paper makes both a substantive and methodological contribution. Substantively, we have shown where key events appear to have affected national mortality rates, both as period effects and cohort effects. In particular, World Wars I and II both appear to have had period effects on male mortality, whilst the influenza pandemic of 1918–1919 appears to have had both a period and cohort effect on mortality across a number of countries. There also appears to be a cohort effect associated with 1947 in the Netherlands and a cohort effect, this time reducing mortality, associated with 1948 in a number of countries including Great Britain although the cause of this remains uncertain.

Methodologically this paper has shown the value of APC analysis of non–linear stochastic variation, both using statistical methods (such as the adapted HAPC model) and graphical techniques (such as the use of Lexis diagrams). These techniques can be used to assess a range of outcomes across the health and social sciences, wherever age, period, and cohort stochastic effects are of interest. There is the potential for further work to assess different ways our modelling approach could be adapted, reducing the misspecification seen where non–linear APC trends remain in the residuals. A comparison between these sorts of models, and Lee–Carter models, would also be worthwhile, revealing the ways in which they complement each other and could potentially be combined to produce more robust inference.

Of course, our results are only as accurate as the data we have used, and so some of our results could be driven by inaccuracies or inconsistencies in the data. Our results could be in part related to artefacts in the way some of the HMD data is imputed for some countries. Alternatively it could be a result of 'phantoms' in the data (Cairns *et al.*, 2016) relating to different distributions of birth registrations throughout each year, which could in turn affect the accuracy of our mortality predictions for particular cohorts.

Acknowledgements

The authors wish to acknowledge funding received from the British Academy Skills Innovator Award, award reference SK150011. Thank you to Kelvyn Jones, Danny Dorling, Mel Bartley, and Charles Pattie for constructive comments in the preparation of this manuscript.

References

Acosta, E. and van Raalte, A. (2019) 'APC curvature plots: Displaying nonlinear age–period–cohort patterns on Lexis plots', *Demographic Research*, 41, pp. 1205–1234. doi:10.4054/DemRes.2019.41.42.

Almond, D. (2006) 'Is the 1918 influenza pandemic over? Long–term effects of *in utero* influenza exposure in the post–1940 U.S. population', *Journal of Political Economy*, 114(4), pp. 672–712.

Bell, A. and Jones, K. (2013) 'The impossibility of separating age, period and cohort effects', *Social Science & Medicine*, 93, pp. 163–165. doi:10.1016/j.socscimed.2013.04.029.

Bell, A. and Jones, K. (2014a) 'Another 'futile quest'? A simulation study of Yang and Land's hierarchical age–period–cohort model', *Demographic Research*, 30, pp. 333–360.

Bell, A. and Jones, K. (2014b) 'Don't birth cohorts matter? A commentary and simulation exercise on Reither, Hauser and Yang's (2009) age—period—cohort study of obesity', *Social Science & Medicine*, 101, pp. 176–180. doi:10.1016/j.socscimed.2013.09.004.

Bell, A. and Jones, K. (2018) 'The hierarchical age–period–cohort model: Why does it find the results that it finds', *Quality & Quantity*, (2), pp. 783–799.

Bell, V., Rui Pita, J. and Pereira, A.L. (2017) 'Regulação, circulação e distribuição da penicilina em Portugal (1944-1954)', *Dynamis*, 37(1), pp. 159–186.

Ben–Shlomo, Y., Cooper, R. and Kuh, D. (2016) 'The last two decades of life course epidemiology, and its relevance for research on ageing', *International Journal of Epidemiology*, 45(4), pp. 973–988. doi:10.1093/ije/dyw096.

Ben–Shlomo, Y. and Kuh, D. (2002) 'A life course approach to chronic disease epidemiology: Conceptual models, empirical challenges and interdisciplinary perspectives', *International Journal of Epidemiology*, 31, pp. 285–293.

Browne, W.J. (2017) *MCMC estimation in MLwiN version 3.0.* Centre for Multilevel Modelling, University of Bristol.

Cairns, A.J.G. *et al.* (2016) 'Phantoms never die: Living with unreliable population data', *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 179(4), pp. 975–1005.

Carroll, D. *et al.* (2017) 'The behavioural, cognitive, and neural corollaries of blunted cardiovascular and cortisol reactions to acute psychological stress', *Neuroscience & Biobehavioural Reviews*, 77, pp. 74–86.

Carstensen, B. (2006) 'Age–period–cohort models for the Lexis diagram', *Statistics in Medicine*, 26, pp. 3018–3045.

Charlton, C. *et al.* (2017) *MLwiN version 3.0.* Centre for Multilevel Modelling, University of Bristol. Available at: http://www.bristol.ac.uk/cmm/software/mlwin/.

Chauvel, L., Leist, A. and Ponomarenko, V. (2016) 'Testing persistence of cohort effects in the epidemiology of suicide: An age–period–cohort hysteresis model', *PLoS ONE* [Preprint]. Available at: https://doi.org/10.1371/journal.pone.0158538.

Curson, P. and McCracken, K. (2006) 'An Australian perspective of the 1918–1919 influenza pandemic', *New South Wales Public Health Bulletin*, 17. doi:https://doi.org/10.1071/NB06025.

Derrick, V.P.A. (1927) 'Observations on (1) errors of age in the population statistics of England Wales, and (2) the changes in mortality indicated by the national records', *Journal of the Institute of Actuaries (1886-1994)*, 58(2), pp. 117–159. Available at: http://www.jstor.

org/stable/41137191.
Dwork, D. (1987) *War is good for babies and other young children: A history of the infant and child welfare movement in England 1898–1918*. London: Tavistock Publications.
Fall, C.H.D. *et al.* (1995) 'Weight in infancy and prevalence of coronary heart disease in adult life', *BMJ*, 310.
Glenn, N.D. (2005) *Cohort analysis*. 2nd edn. Thousand Oaks, CA; London: Sage.
Gluckman, P.D., Hanson, M.A. and Buklijas, T. (2010) 'A conceptual framework for the developmental origins or health and disease', *Journal of Developmental Origins of Health and Disease*, 1(1), pp. 6–18.
Goldstein, H. (1979) 'Age, period and cohort effects: A confounded confusion', *Journal of Applied Statistics*, 6(1), pp. 19–24.
Hales, C.N. and Barker, D.J.P. (1992) 'Type 2 (non–insulin–dependent) diabetes mellitus: The thrifty phenotype hypothesis', *Diabetologia*, 35, pp. 595–601. doi:10.1007/BF00400248.
Hanson, M. and Gluckman, P. (2016) 'Commentary: Developing the future: Life course epidemiology, DOHaD and evolutionary medicine', *International Journal of Epidemiology*, 45(4), pp. 993–996.
Hardy, R. and Tilling, K. (2016) 'Commentary: The use and misuse of life course models', *International Journal of Epidemiology*, 45(4), pp. 1003–1005. doi:10.1093/ije/dyw101.
Healy, K. (2018) 'Heatmap of mortality rates'. Available at: https://kieranhealy.org/blog/archives/2018/12/04/heatmaps-of-mortality-rates/ (Accessed: 24 November 2020).
Heligman, L. and Pollard, J.H. (1980) 'The age pattern of mortality', *Journal of the Institute of Actuaries*, 107(1), pp. 49–80. doi:10.1017/S0020268100040257.
Hertzman, C. (1999) 'The biological embedding of early experience and its effects on health in adulthood', *Annals New York Academy of Sciences*, 896, pp. 85–95.
Hill, C.P. (1970) *British economic and social history: 1700–1964*. 3rd edn. London: Edward Arnold.
Hyndman, R. and Koehler, A. (2006) 'Another look at measures of forecast accuracy', *International Journal of Forecasting*, 22(4), pp. 679–688. doi:https://doi.org/10.1016/j.ijforecast.2006.03.001.
Johnson, N.P. and Mueller, J. (2002) 'Updating the accounts: Global mortality of the 1918–1920 "Spanish" influenza pandemic', *Bull Hist Med*, 76(1), pp. 105–15.
Jones, K. *et al.* (2015) 'Modelling the occupational assimilation of immigrants by ancestry, age group and generational differences in Australia: A random effects approach to a large table of counts', *Quality and Quantity*, 49, pp. 2595–2615. doi:https://doi.org/10.1007/s11135-014-0130-8.
Kermack, W.O., McKendrick, A.G. and Mckinlay, P.L. (1934) 'Death–rates in Great Britain and Sweden: Some general regularities and their significance', *The Lancet*, 223(5770), pp. 698–703. doi:https://doi.org/10.1016/S0140-6736(00)92530-3.
Koupil, I. *et al.* (2009) 'Cancer mortality in women and men who survived the siege of Leningrad (1941-–1944)', *International Journal of Cancer*, 124(6), pp. 1416–1421.
Luo, L. (2013) 'Assessing validity and application scope of the intrinsic estimator approach to the age–period–cohort problem', *Demography*, 50(6), pp. 1945–1967.
Luo, L. and Hodges, J.S. (2015) 'Block constraints in age–period–cohort models with unequal–width intervals', *Sociological Methods & Research*, 45(4), pp. 700–726.
Mazumder, B. *et al.* (2010) 'Lingering prenatal effects of the 1918 influenza pandemic on

cardiovascular disease', *Journal of Developmental Origins of Health and Disease*, 1(1), pp. 26–34.
Medlicott, W.N. (1967) *Contemporary England: 1914–1964*. London: Longmans.
Minton, J. (2021) 'The Lexis surface: A tool and workflow for better reasoning about population data', in Bell, A. (ed.) *Age, period and cohort effects: Statistical analysis and the identification problem*. New York: Routledge.
Minton, J., Vanderbloemen, L. and Dorling, D. (2013) 'Visualizing Europe's demographic scars with coplots and contour plots', *International Journal of Epidemiology*, 42(4), pp. 1164–1176. doi:10.1093/ije/dyt115.
Murphy, M. (2010) 'Reexamining the dominance of birth cohort effects on mortality', *Population and Development Review*, 36(2), pp. 365–390. Available at: http://www.jstor.org/stable/25699065.
Newman, A.B. (2016) 'Commentary: Life course epidemiology: The past two decades and future directions', *International Journal of Epidemiology*, 45(4), pp. 1013–1015. doi:10.1093/ije/dyw102.
Noymer, A. and Garenne, M. (2000) 'The 1918 influence epidemic's effects on sex differentials in mortality in the United States', *Popul Dev Rev*, 26(3), pp. 565–581.
Pascariu, M., Lenart, A. and Canudas–Romo, V. (2019) 'The maximum entropy mortality model: Forecasting mortality using statistical moments', *Scandinavian Actuarial Journal*, 2019(8), pp. 661–685. doi:10.1080/03461238.2019.1596974.
Patterson, K.D. and Pyle, G.F. (1991) 'The geography and mortality of the 1918 influenza pandemic', *Bull Hist Med*, 65(1), pp. 4–21.
Reither, E.N., Hauser, R.M. and Yang, Y. (2009) 'Do birth cohorts matter? Age–period–cohort analyses of the obesity epidemic in the United States', *Social Science & Medicine*, 69, pp. 1439–1448.
Renshaw, A.E. and Haberman, S. (2006) 'A cohort–based extension to the Lee–Carter model for mortality reduction factors', *Insurance: Mathematics and Economics*, 38, pp. 556–70.
Rivett, G. (1998) *From cradle to grave: Fifty years of the NHS*. London: King's Fund Publishing.
Ryder, N.B. (1965) 'The cohort as a concept in the study of social change', *Americal Sociological Review*, 30(6), pp. 843–861.
Sarkar, D. (2008) *Lattice: Multivariate data visualization with R*. New York: Springer. Available at: http://lmdvr.r-forge.r-project.org.
Schöley, J. and Willekens, F. (2017) 'Visualizing compositional data on the Lexis surface', *Demographic Research*, 36(21), pp. 627–658. doi:10.4054/DemRes.2017.36.21.
Sparén, P. *et al.* (2004) 'Long term mortality after severe starvation during the seige of Leningrad: Prospective cohort study', *BMJ*, 328.
Stanner, S.A. *et al.* (1997) 'Does malnutrition in utero determine diabetes and coronary heart disease in adulthood? Results from the leningrad siege study, a cross sectional study', *BMJ*, 315(7119), pp. 1342–1348. doi:10.1136/bmj.315.7119.1342.
Suzuki, E. (2012) 'Time changes, so do people', *Social Science & Medicine*, 75, pp. 452–456.
The World Bank (2017) 'Death rate, crude (per 1,000 people)'. Available at: http://databank.worldbank.org/data/reports.aspx?source=2&type=metadata&series=SP.DYN.CDRT.IN (Accessed: 16 March 2017).
University of California, Berkeley (USA) and Max Planck Institute for Demographic Research

(Germany) (2017) 'Human Mortality Database'. Available at: http://www.mortality.org (Accessed: 22 March 2017).
Wadsworth, M.E.J. and Kuh, D.J.H. (1997) 'Childhood influences on adult health: A review of recent work from the British 1946 national birth cohort study, the MRC National Survey of Health and Development', *Paediatric and Perinatal Epidemiology*, 11, pp. 2–20.
Warmbrunn, W. (1963) *The Dutch under German occupation, 1940–1945*. Stanford, California; London: Stanford University Press; Oxford University Press.
Webster, C. (2002) *The National Health Service: A political history*. 2nd edn. Oxford: Oxford University Press.
Willets, R.C. (2004) 'The cohort effect: Insights and explanations', *British Actuarial Journal*, 10(4), pp. 833–877. doi:10.1017/S1357321700002762.
Wilmoth, J.R. *et al.* (2021) *Methods protocol for the human mortality database: HMD methods protocol v6*. Available at: https://www.mortality.org/Public/Docs/MethodsProtocol.pdf (Accessed: 21 December 2021).
Winter, J. and Prost, A. (2005) *The Great War in history: Debates and controversies, 1914 to the present*. New York; Cambridge: Cambridge University Press.
Yang, Y. *et al.* (2008) 'The intrinsic estimator for age–period–cohort analysis: What is it and how to use it', *Americal Journal of Sociology*, 113(6), pp. 1697–1736.
Yang, Y. and Land, K.C. (2006) 'A mixed methods 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, pp. 75–97. doi:https://doi.org/10.1111/j.1467-9531.2006.00175.x.
Zhang, Z. *et al.* (2016) 'R2MLwiN: A package to run MLwiN from within R', *Journal of Statistical Software*, 72(10). doi:http://dx.doi.org/10.18637/jss.v072.i10.