ANNUAL REVIEWS

Annual Review of Sociology

Analyzing Age-Period-Cohort Data: A Review and Critique

Ethan Fosse[1] and Christopher Winship[2]

[1]Department of Sociology, University of Toronto, Toronto, Ontario M5S 2J4, Canada; email: ethan.fosse@utoronto.ca

[2]Department of Sociology, Harvard University, Cambridge, Massachusetts 02138, USA

Annu. Rev. Sociol. 2019. 45:467–92

First published as a Review in Advance on May 15, 2019

The *Annual Review of Sociology* is online at soc.annualreviews.org

https://doi.org/10.1146/annurev-soc-073018-022616

Keywords

age-period-cohort analysis, APC, social change, mechanism-based models, partial identification, bounds, sensitivity analysis, causal inference, life course

Abstract

Age-period-cohort (APC) analysis has a long, controversial history in sociology and related fields. Despite the existence of hundreds, if not thousands, of articles and dozens of books, there is little agreement on how to adequately analyze APC data. This article begins with a brief overview of APC analysis, discussing how one can interpret APC effects in a causal way. Next, we review methods that obtain point identification of APC effects, such as the equality constraints model, Moore-Penrose estimators, and multilevel models. We then outline techniques that entail point identification using measured causes, such as the proxy variables approach and mechanism-based models. Next, we discuss a general framework for APC analysis grounded in partial identification using bounds and sensitivity analyses. We conclude by outlining a general step-by-step procedure for conducting APC analyses, presenting an empirical example examining temporal shifts in verbal ability.

1. INTRODUCTION

Age-period-cohort (APC) analysis has a long, controversial history in the social sciences. Formally, age refers to the time since birth, period refers to the calendar date at which an outcome is observed, and cohort refers to the time when an individual was born. The core idea of an APC analysis is that changes in any particular outcome can be attributed to three distinct types of causal processes (Glenn 2005, p. 11): (*a*) changes over the life course of individuals, or age effects; (*b*) changes due to the events in particular years, or period effects; and (*c*) changes due to the replacement of older cohorts of individuals with younger ones with different characteristics, or cohort effects. Typically the goal of an APC analysis is to understand social and cultural change in a given outcome by identifying the separate contributions of the causal processes associated with each of the three APC variables.

Despite the existence of hundreds, if not thousands, of articles and dozens of books, there is little agreement on how to adequately analyze APC data. The central issue has been, and continues to be, the linear relationship among age, period, and cohort, that is, that age = period − cohort. This linear dependence results in the so-called APC identification problem, or the fact that it is not possible to directly estimate the linear effect of any one of the APC variables holding the other two constant. Numerous solutions have been proposed to deal with the identification problem, but none have gained wide acceptance.

In this article, we review both older and more recent methods for analyzing APC data. We also discuss which approaches provide viable means for separating out the independent effects of age, period, and cohort and which do not. More specifically, we present a simplification of the identification problem and clarify the properties of multiple traditional APC methods. In doing so, we arrive at a number of key insights that have been either overlooked or underappreciated in the literature:

1. Age, period, and cohort variables can be understood as observed indicators for distinct sets of underlying, typically unobserved, causal processes (Clogg 1982, pp. 460–62; Heckman & Robb 1985, pp. 137–40; Mason & Fienberg 1985, pp. 45–48). Typically the goal of an APC analysis is to determine the unique contribution of these causal processes for some particular outcome.
2. A vast number of studies in sociology and related fields are affected by the APC identification problem. Too often these studies simply omit one or two of the temporal variables entirely without reference to the identification problem. The assumptions underlying these studies are often testable.
3. The fact that age, period, and cohort are exact functions of each other presents an identification problem only for the linear effects (or slopes). Nonlinear effects (or deviations from the slopes) are identified (Fienberg 2013, Fosse & Winship 2018). As such, the APC identification problem can be radically simplified as a problem of determining the value of the three linear effects.
4. Any technique that identifies point estimates for age, period, and cohort amounts to choosing a point on (or near) a multidimensional solution line in the parameter space (Luo et al. 2016; O'Brien 2014, pp. 59–91). Because only the linear effects are unidentified, this simplifies to selecting a point on the canonical solution line, which lies in just three dimensions (Fosse & Winship 2018). The canonical solution line is easily visualized using a 2D-APC graph (Fosse & Winship 2019).
5. As with any problem of causal inference, techniques for analyzing APC data as causal structures necessarily entail assumptions (Converse 1976, pp. 16–26; Glenn 2005, p. 7). The credibility of a set of APC estimates is strengthened greatly by the degree to which these

assumptions are based on social, biological, or cultural theory and/or tested against the data (Fosse & Winship 2018, 2019). Complicated and highly technical APC methods can rely on surprising, even unjustifiable theoretical assumptions.

6. With some and possibly most APC data, much can be learned about the relative importance of different causal processes without point identification. Bounds based on constraining the size, sign, or shape of one or more of the temporal effects can often provide informative results, in some cases approximating point estimates under relatively weak theoretical assumptions (Fosse & Winship 2019; cf. Manski 1990, 2003). Although bounds can be expressed mathematically, they have an equivalent visual representation in 2D-APC graphs, simplifying the task of partial identification (Fosse & Winship 2019).
7. Observed measures of the causal processes indexed by the temporal variables can help greatly in estimating the overall relative effects of age, period, and cohort on a given outcome (Winship & Harding 2008). APC models with observed measures of causal processes can be combined easily with bounding (or sensitivity) analyses, allowing one to make conclusions with even weaker assumptions.

We develop these insights in three main parts. First, we briefly discuss the history of APC analysis, the classical APC model, the APC identification problem, and the interpretation of APC effects as causal structures. Second, we review methods that obtain point identification of APC effects. In doing so, we investigate the properties of various traditional APC estimators commonly used by demographers and sociologists, including the equality constraints model, intrinsic estimator (IE), and hierarchical age-period-cohort (HAPC) model. These techniques can be understood as selecting particular points on the canonical solution line, often relying on untested or untestable assumptions that have little or no theoretical basis. We then review techniques that incorporate explicit measures of intermediate causal processes, in particular the proxy variables approach and mechanism-based models. Next, we discuss an alternative framework for APC analysis grounded in partial identification using bounding or sensitivity analyses. We conclude by outlining a general step-by-step procedure for conducting APC analyses and presenting an empirical example on changes in verbal ability.

2. OVERVIEW OF AGE-PERIOD-COHORT ANALYSIS

2.1. Foundational Works and Key Applications

The earliest known APC analysis dates back to at least the 1860s, predating the eponymous diagrams of Wilhelm Lexis. In a remarkable confluence of thought, a number of German demographers devised inventive ways to visualize mortality data indexed by temporal variables (see Keiding 2011). The modern era of APC analysis began with seminal works in epidemiology and medicine in the 1930s and later in sociology and demography in the 1950s and 1960s [Mannheim 1952 (1927), Ryder 1965]. In a now classic work, Frost (1939) noted that changes in observed tuberculosis mortality rates could be attributed to not only aging effects but also period and cohort effects (Mason & Smith 1985).[1]

Similarly, in a posthumous theoretical essay, the sociologist and philosopher Karl Mannheim [1952 (1927)] emphasized the importance of generations in understanding social change (see also Demartini 1985, Pilcher 1994, Simirenko 1966). Among Mannheim's insights was that social change occurs not only because individuals alter their views, but also because each cohort

[1]See Case (1956) for an overview of other early studies in epidemiology and medicine on cohort effects and mortality rates (e.g., Andvord 1930, Derrick 1927, Kermack et al. 1934).

is imprinted with different values and beliefs that they carry until their death [Mannheim 1952 (1927), pp. 292–302].

Similarly, in an influential article linking population studies with cultural analysis, the demographer Norman Ryder (1965) noted that while social change is partly caused by individuals changing their values and beliefs, a potentially more dramatic driver is that, through the biological processes of birth and death, groups are continually entering and leaving society. His argument appealed to a broad swath of social scientists in large part because he recognized that massive social change can occur even though no particular individual has changed (Firebaugh 1989, Firebaugh & Harley 1991, Hobcraft et al. 1982, Inglehart 1971).

Since the seminal works by Frost, Mannheim, and Ryder, social and behavioral scientists have used APC data to investigate a wide range of outcomes. There have been two parallel tracks of research. On the one hand, demographers and epidemiologists have focused on examining disease rates and health-related behaviors, including drug use (Chen et al. 2003, Kerr et al. 2004, O'Malley et al. 1984, Vedøy 2014), obesity (Diouf et al. 2010, Fu & Land 2015), cancer (Clayton & Schifflers 1987, Liu et al. 2001), and mental health (Lavori et al. 1987, Yang 2008). A main goal of these researchers has been to understand the future burden of disease as well as the risk factors related to various health conditions. On the other hand, sociologists and other social scientists have focused on a range of outcomes related to understanding social change, including verbal ability (Alwin 1991, Hauser & Huang 1997, Wilson & Gove 1999), social trust (Clark & Eisenstein 2013, Putnam 1995, Robinson & Jackson 2001, Schwadel & Stout 2012), party identification (Ghitza & Gelman 2014, Hout & Knoke 1975, Tilley & Evans 2014), and religious affiliation (Chaves 1989, Firebaugh & Harley 1991). The goal of these applications has typically been to identify the separate contributions of age, period, and cohort effects so as to understand the ways in which societies are changing.

2.2. The Classical Age-Period-Cohort Model

The temporal effects of age, period, and cohort can be represented using the classical APC (C-APC) model, also known as the multiple classification model (Mason et al. 1973, p. 243) or accounting model (Mason & Fienberg 1985, pp. 46–47, 67). To understand the C-APC, suppose we have collected data on a set of individuals and have measured each person's birth year, age, year of measurement, and value on some outcome. Typically such data are collected in an age-period table, or what is called a Lexis table (see **Table 1**). It is common in the APC literature to use index notation to keep track of the dimensions of an age-period table. Let $i = 1, \ldots, I$ represent the age groups; $j = 1, \ldots, J$ the period groups; and $k = 1, \ldots, K$ the cohort groups, with $k = j - i + I$ and $K = I + J - 1$.[2]

Table 1 Structure of a Lexis table with indices for age, period, and cohort

	Period			
Age	$j = 1$	$j = 2$	$j = 3$	$= J$
$i = 1$	$k = 5$	$k = 6$	$k = 7$	$= K$
$i = 2$	$k = 4$	$k = 5$	$k = 6$	
$i = 3$	$k = 3$	$k = 4$	$k = 5$	
$i = 4$	$k = 2$	$k = 3$	$k = 4$	
$i = 5 \quad = I$	$k = 1$	$k = 2$	$k = 3$	

Age and period are aggregated into equally spaced intervals indexed by i (rows) and j (columns), respectively. Cohorts are indexed by k (cells). Shaded cells trace the path through time for the kth cohort, where $k = 3$.

[2]Note that I is added to $j - i$ so that the cohort index begins at $k = 1$. This ensures that, for example, $i = j = k = 1$ refers to the 1st group for all three temporal measures. One could just as easily index the cohorts using $k = j - i$, but this identity would be lost.

The C-APC model is, accordingly,

$$Y_{ijk} = \mu + \alpha_i + \pi_j + \gamma_k + \epsilon_{ijk}, \quad 1.$$

where Y_{ijk} is the outcome variable to be explained, μ is the intercept, α_i represents the ith age effect, π_j represents the jth period effect, γ_k represents the kth cohort effect, and ϵ_{ijk} is the error term. We can specify the above in matrix notation as follows:

$$\mathbf{y} = \mathbf{Xb} + \epsilon, \quad 2.$$

where $\mathbf{y}$ is an $(I \times J) \times 1$ outcome vector; $\mathbf{X}$ is an $(I \times J) \times 2(I + J) - 3$ design matrix; $\mathbf{b}$ is a vector of coefficients with $2(I + J) - 3$ rows; and ϵ is an $(I \times J) \times 1$ vector representing random error. Typically in an APC model, in order to allow for nonlinear effects, age, period, and cohort are treated as categorical variables.[3] As in traditional regression analysis, there are multiple possible coding schemes, many of which necessitate applying a constraint to deal with the overparameterization due to the inclusion of the intercept. For example, one could constrain the parameters to sum to zero, such that $\sum_{i=1}^{I} \alpha_i = \sum_{j=1}^{J} \pi_j = \sum_{k=1}^{K} \gamma_k = 0$.[4]

2.3. The Age-Period-Cohort Identification Problem

It is recognized throughout the APC literature that all design matrices with APC variables are rank deficient because of the linear relationship age = period − cohort. It is this rank deficiency that results in the APC identification problem.[5] Because APC models are not fully identified, we cannot obtain point estimates for each of the age, period, and cohort linear effects (or effects that are partially a function of the linear effects). A convenient way to express the nonidentifiability problem is to note that, for any particular APC model, we can specify the linear effects as (Rodgers 1982a, p. 782)

$$\begin{aligned} \alpha^* &= \alpha + \nu, \\ \pi^* &= \pi - \nu, \\ \gamma^* &= \gamma + \nu, \end{aligned} \quad 3.$$

where the asterisk indicates an arbitrary set of estimated slopes from an APC model under a particular constraint or, equivalently, when the scalar ν is fixed to some value. As Equation 3 indicates, the estimated parameters are simple additive transformations of the true unobserved slopes α, π, and γ shifted by a single arbitrary scalar, ν.[6]

Although the linear effects of an APC model are not identifiable, several other quantities are. First and most importantly, the nonlinear effects, that is, the deviations from the linear effects,

[3] One could also use splines, treating age, period, and cohort as continuous variables (see Heuer 1997).

[4] Alternatively, one could fix the parameters at one of the levels to zero. By convention, researchers typically fix the first set of levels (e.g., $\alpha_{i=1} = \pi_{j=1} = \gamma_{k=1} = 0$) or the last set (e.g., $\alpha_{i=I} = \pi_{j=J} = \gamma_{k=K} = 0$), although another set could be chosen to satisfy the constraints.

[5] With a few exceptions (Mason & Fienberg 1985, pp. 71–72), APC analysts have only considered models with main effects for age, period, and cohort, ruling out interactions among these variables. Each interaction added to a model increases the rank deficiency of the design matrix by one, further exacerbating the identification problem. Procedures for dealing with such models are an important topic for future research.

[6] Researchers may believe that certain kinds of data, for example, cross-sectional or panel data, are immune to the identification problem. This is incorrect. In a cross-sectional study, researchers are comparing age groups in a snapshot of calendar time. The estimated linear age effect is not α, but $\widehat{\alpha} = \alpha - \gamma$. The data are not themselves informative as to whether people differ because of age or cohort effects. In a panel study, one might examine, say, a single cohort of people and observe how they change as they age. Here, the estimated linear age effect is not α, but $\widehat{\alpha} = \alpha + \pi$.

are fully identifiable [Fosse & Winship (2018) provide formal proofs]. Second, as we outline later, particular combinations of the linear effects are also identifiable. Third, the predicted values of the outcome are identical regardless of the constraint (i.e., the value of v) used to identify the model. Finally, forecasts from an APC model are not subject to the identification problem. Specifically, the predicted values of the outcome from a linear extrapolation are identifiable (Holford 2005).

2.4. Interpreting Age-Period-Cohort Effects

In order to understand the relationship between the formal structure of an APC model and typical applications, consider two examples. First, consider the possible factors involved in temporal changes in verbal ability as measured by vocabulary knowledge (see, e.g., Alwin 1991, Hauser & Huang 1997, Wilson & Gove 1999). At first glance, one may not expect that age would be a contributing factor. Some scholars, however, have noted the potential for the age composition of a population to result in observed changes in levels of verbal ability (Wilson & Gove 1999). The situation, however, is complicated, as the age effect is likely to be curvilinear: At younger ages, individuals' vocabulary may increase as they are exposed to more new words, while at older ages it might well decrease due to declines in basic cognitive functioning. Given this curvilinear relationship, the age distribution would only have a substantial effect on verbal ability if there were a relative increase or decrease in the number of middle-aged individuals. It is also not obvious that there would be period effects until one considers changes in the way people obtain information. With the advent of radio, television, the Internet, and cell phones, individuals now read less (Glenn 1994). Verbal ability, however, as measured by vocabulary has increased over time, particularly during the first half of the twentieth century (Schooler 1998). Scholars have recognized that the expansion of higher education across birth cohorts is probably the critical factor in changing levels of verbal ability (Alwin 1991). This important cohort-based change will appear as increasing levels of verbal ability.

As a second example, consider changes in the prevalence of tuberculosis, the classic example of an APC analysis being used to explain changes in disease rates (Frost 1939). Here too, age, period, and cohort are all potential factors. Again changes in the age distribution could be a factor; however, as with verbal ability, this is a complicated relationship because the young and the old would be the most susceptible to tuberculosis. As well, there are likely period and cohort effects (Mason & Smith 1985). Cohorts might well vary in their resistance and immunity, especially as more recent cohorts are healthier in general. As with infectious diseases in general, however, period effects are likely important and will shift due to the implementation of new medical technologies and diagnostic tools, as well as changes in how diseases are classified.

In the APC literature, researchers have typically viewed age, period, and cohort as indicators of distinct sets of unmeasured causes (see, e.g., Mason & Fienberg 1985, Rodgers 1990). APC variables themselves are causally ill-defined because there is no direct intervention that would change, say, cohort, to some other value while holding constant age and period. This point of view is expressed by Clogg (1982, p. 460): "age, period, and cohort are merely indicators of other variables which actually 'cause' the observed variation in the dependent variable under study. The APC framework is properly interpreted as an accounting scheme, not a 'causal model.'" Similarly, Mason & Fienberg (1985, pp. 46–47) write, "these models do not explain so much as they provide categories with which to seek explanation. For accounting models to have value, the parameterizations of the general framework must be linked to phenomena presumed to underlie the accounting categories." **Figure 1** illustrates one way to represent APC variables as causal structures in a graphical model. The simplified model in **Figure 1*a*** can be understood as shorthand for an underlying mechanism-based model, shown in **Figure 1*b***. Unobserved sets of causal processes

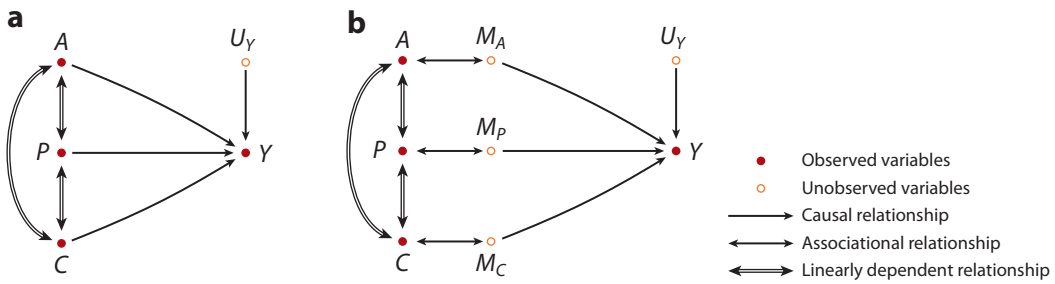

Figure 1

Graphical models of temporal variables. Panel *a* shows the simplified age-period-cohort graphical model for age (*A*), period (*P*), and cohort (*C*). The double-line, double-headed arrows denote the linearly dependent relationships among the three temporal variables. *Y* is the outcome, while U_Y refers to idiosyncratic unobserved causes. Filled circles denote observed variables, while hollow circles denote unobserved variables. Panel *b* shows the graphical model with a full set of mechanism variables (M_A, M_P, and M_C), which are unobserved. The single-line, double-headed arrows denote associational linkages between the temporal variables and unobserved causal mechanisms.

(represented by M_A, M_P, and M_C in **Figure 1*b***) are implied whenever researchers refer to "APC effects."

To illustrate how this might be conceptualized, consider a study examining temporal effects of political alienation. For example, age might be associated with the marital status of the respondent and the presence of children in the home, period with the unemployment rate and who is president, and cohort with cohort size and the lifetime amount of education obtained. In this sense one can think that in stating that an APC variable has a causal effect, one is actually making an elliptical assertion in which the full statement is that the causal processes associated with one of the APC variables has an effect on some outcome. For example, the statement "period has a causal effect on political alienation" is shorthand for stating that "period membership is associated with the unemployment rate and who is president, which change political attitudes."[7] Under this conceptualization there is no problem when we state that an APC variable has a causal effect; we are simply using shorthand for a more complicated statement.

There is still the problem of what the coefficient on an APC variable means in a standard linear equation model. An approach that clarifies the distinction between indicators and associated causes interprets the APC coefficients as estimates of differences in group means, which are potentially explained by a set of associated causes that may or may not be observed. As we show below, this approach amounts to conceptualizing an APC model as a problem in regression decomposition.

To illustrate, consider a simple model where a particular outcome is only a function of cohort and period (that is, the age linear effect is zero in the population). In addition, for simplicity, assume that both variables only have two categories and there are only linear effects, giving us the following equation:

$$Y_{jk} = \mu + \pi(\text{period}_j) + \gamma(\text{cohort}_k) + \epsilon_{jk}, \qquad 4.$$

where γ is the cohort linear effect and π is the period linear effect. Equation 4 can be interpreted as expressing the conditional mean of Y_{jk} for different period-cohort categories. As such, the coefficients π and γ equal the differences in group means for the period and cohort categories, respectively.

[7]The conceptualization here is close to that of Sen & Wasow (2016), who consider the related problem of whether race can have a causal effect. They depict race, as one might also depict gender, as a bundle of causal effects.

Now consider extending Equation 4 by estimating a regression equation that includes measures of the associated causes of the observed group differences. For example, one might attempt to explain the differences in political alienation between age groups by controlling for income and years of work experience. The goal here is to include a sufficiently detailed set of causal factors representing the process associated with age and period that the estimated coefficients for age and period are zero.

To understand how to interpret the coefficients in Equation 4, we can use regression decomposition, an extensively employed method used to assess the relative importance of different causes in explaining differences in group means on an outcome (Blinder 1973, Oaxaca 1973, Vaupel & Romo 2003). Without loss of generality, assume that there are two dichotomous variables X and Z that, when controlled for, reduce the linear coefficient of cohort, γ, to zero. One could then decompose the coefficients of these two variables on the difference in group means, γ, as follows:

$$\gamma = \underbrace{(\overline{X}_2 - \overline{X}_1)\beta_1}_{\text{I}} + \underbrace{(\overline{Z}_2 - \overline{Z}_1)\beta_2}_{\text{II}}. \qquad 5.$$

Component I consists of the estimated causal effect of X on the outcome Y, given by the coefficient β_1, and the difference in the means of X across the two cohort groups $(\overline{X}_2 - \overline{X}_1)$. Likewise, component II gives the estimated casual effect of Z on Y, given by β_2, and the difference in the means of Z across the two cohort groups $(\overline{Z}_2 - \overline{Z}_1)$. As such, components I and II in Equation 5 can be interpreted as representing the contribution of the two age-indexed causal processes, represented by X and Z, to the overall difference in the outcome Y between the two cohort groups, given by γ.[8] A similar decomposition could be conducted for π, assuming that one had the causal processes underlying observed differences across the period groups.

Now consider a full model in which all three of the APC variables are entered as inputs. Again, for simplicity, we assume that there are only linear effects. The full model is

$$Y_{ijk} = \mu + \alpha(\text{age}_i) + \pi(\text{period}_j) + \gamma(\text{cohort}_k) + \epsilon_{ijk}, \qquad 6.$$

where γ is the cohort linear effect. This equation cannot be directly estimated because of the linear dependence among the three timescales. However, Winship & Harding (2008) show that it is possible to estimate Equation 6 indirectly if one has a complete set of causes for the group differences for either age, period, or cohort—that is, controlling for a sufficient set of variables, one can assume that the coefficient on at least one of age, period, and cohort variables is zero.

There are two closely related ways to interpret the coefficients in Equation 6. One interpretation, as in Equation 4, is that the coefficients represent the differences in the group means for each of the three APC variables conditional on the other two. The strength of this interpretation is that it simply extends the interpretation of the coefficients for the two APC variables in Equation 4. The weakness is that because of linear dependence, the three conditional mean differences are not directly observed.

A second, less demanding way to interpret these coefficients is to state that they represent the total contribution or total effects of the causal processes that produce observed mean differences across the cross-classified age, period, and cohort categories. This is consistent with the view that, by referring to APC effects, researchers are referring to the total effects of causal processes believed to be associated with the temporal variables. For example, if we were comparing a single cohort at

[8]Note that the differences in group means in Equation 5 could be modeled in a way similar to that in Equation 4 by two conditional mean equations with each variable, X and Z, as the outcome. The coefficients in each equation would then equal the mean differences in components I and II in Equation 5.

two different ages, and thus at two different periods, the difference in their observed mean on the outcome would be equal to the effect of age plus the effect of period because they would differ on both age and period.

The above presentation makes many simplifying assumptions. Extensions to situations in which variables have multiple categories and where variables interact with each other is conceptually straightforward but requires quite extensive and cumbersome notation.

3. POINT IDENTIFICATION OF AGE-PERIOD-COHORT EFFECTS

As noted in Section 1, traditional approaches to APC analysis have attempted to obtain point estimates of the total effects of age, period, and cohort. Doing so always involves making a strong assumption involving a mathematical constraint that generally either is not tested or is untestable (O'Brien 2014). In this section we review the most commonly used methods for point identification. These methods fall into two broad types: approaches where measures of the causes associated with the three APC variables either are not available or are not used, and methods that rely on such measures. As our presentation makes clear, we are skeptical about the usefulness of these methods in the absence of tests of their implied constraints.

3.1. Point Identification Without Measured Causes

Point identification methods without measured causes can be divided into two groups: those that use explicit constraints, such as setting the effects associated with two categories equal to each other, and those that use mechanical constraints, such as the recently developed IE.

3.1.1. Explicit constraints. Here we consider two types of explicit constraints. The two most common techniques using explicit constraints are drop-one variable models and the equality constraints approach.

3.1.1.1. Drop-one variable. By far the most common approach among researchers, sometimes stated and other times unstated, is to drop one or more of the three APC variables (see, e.g., Brooks & Manza 1997; Firebaugh 1989; Firebaugh & Harley 1991; Putnam 1995, 2001; Voas & Chaves 2016). Dropping a variable is equivalent to constraining its linear and nonlinear effects to zero, which is typically an unnecessarily strong assumption.[9] What has not generally been recognized is that because the nonlinear effects are identifiable, constraining them to be zero is a testable assumption. As such, if one is considering dropping one of the three APC variables, one should always test whether this is consistent with the data.

3.1.1.2. Equality constraints. The most frequent approach used during the initial wave of research on APC effects in the 1970s and 1980s was an equality constraints model (Fienberg & Mason 1979, Mason et al. 1973, Mason & Fienberg 1985). The basic idea is to set the effects for two (typically adjacent) APC groups equal to each other (e.g., $\pi_j = \pi_{j+1}$).[10] For instance, in analyzing tuberculosis, Mason & Smith (1985) assume that the coefficients for ages 5–9 and ages 10–14 are equal to each other.

The main problem with the equality constraints approach is that it relies on a much stronger assumption than it might appear, because setting two groups equal to each other is tantamount

[9]If these constraints are not valid, then the estimates of the other two APC variables' linear and nonlinear effects will be inconsistent.

[10]A slight variation is to fix the first and last group to have equal effects, which sets the linear effects to zero.

to setting each of the unknown linear effects to some particular value.[11] In addition, despite the assertion by Fienberg and colleagues that equality constraints should be based on overt theoretical assumptions (Mason & Fienberg 1985, Smith et al. 1982), in practice researchers have used such constraints in arbitrary and often atheoretical ways (though for a notable exception, see Mason & Smith 1985, pp. 175–78). Furthermore, researchers may be misled by failing to recognize that models with differing equality constraints have identical fit statistics (Yang & Land 2013, pp. 65–66).

An additional problem is that researchers using the equality constraints approach have frequently relied on just-identifying constraints by, for example, setting two groups equal to each other on just one of the temporal variables. The focus on just-identifying constraints has resulted in neglect of the usefulness of applying multiple constraints. For example, one could set multiple groups equal to each other, resulting in an overidentified model. In general a model with multiple constraints will be overidentified and thus testable against the data. Further research is needed to determine the extent to which this is a fruitful direction for new analyses.

3.1.2. Mechanical constraints. Methods utilizing mechanical constraints have been popular over the past several decades (Fu 2000, Fu et al. 2011, Yang & Land 2006). By "mechanical constraint," we are referring to general-purpose methods (see Glenn 2005). For example, Yang & Land (2013, p. 76) write that "the IE is a general-purpose method of APC analysis." These techniques have led to a wave of prominent studies on various topics, including verbal ability (Yang & Land 2006), infant mortality (Powers 2013), heart disease (Lee & Park 2012), obesity (Reither et al. 2009), and perceived happiness (Yang 2008). By far the two most widely used techniques are the IE and the HAPC, but similar results can be obtained from ridge regression, partial least squares regression, and principal components regression (O'Brien 2011).

3.1.2.1. Moore-Penrose estimators. Moore-Penrose (MP) estimators, and in particular the IE, have recently gained popularity in sociology, public health, and a variety of other fields. However, these estimators are controversial because they are sensitive to the design matrix used (Fienberg 2013, Luo et al. 2016, O'Brien 2011). As shown in the mathematical appendix to Luo et al. (2016), there will always be an MP estimator based on a specific design matrix that will produce any set of estimates that are equally consistent with the data, that is, any set of points that fall on the so-called solution line (discussed in detail in Section 4.1). As such, the choice of design matrix is critical.

MP estimators are based on what is known as the Moore-Penrose inverse, a particular kind of generalized inverse. Conceptually, an MP estimator can be thought of as a two-stage least-squares, minimum-norm estimator. In the first stage, the estimator applies the least-squares criterion. Because the APC design matrix is rank deficient, there are multiple sets of least-squares estimates. In the second stage, the estimator applies the minimum norm constraint, selecting that particular set of least-squares estimates for which the square root of the sum of the squared estimates is as small as possible. This is equivalent to choosing that point on the solution line that is closest in Euclidean distance to the origin. Alternatively, an MP estimator can be thought of as a type of shrinkage estimator, where estimates are being shrunk toward zero. They are also closely related to ridge regression.

[11] As Rodgers (1982a, p. 785) has cautioned, "Although a constraint of the type described by Mason et al. (1973) and Fienberg & Mason (1979) seems trivial, in fact it is exquisitely precise and has effects that are multiplied so that even a slight inconsistency between the constraint and reality, or small measurement errors, can have very large effects on estimates." However, see also the reply by Smith et al. (1982) as well as the rejoinder by Rodgers (1982b).

MP estimators have several desirable statistical properties (Fu 2000, 2016; Fu & Hall 2006; Fu et al. 2011; Yang et al. 2004, 2008): First, an MP estimator has minimum sampling variance among all possible estimators based on its specific design matrix; second, it is an estimable function, meaning that it produces a unique set of estimates for the effects of age, period, and cohort; and finally, it is unbiased, meaning that the average of any estimates produced by an MP estimator over an infinite number of simple random samples will equal that estimator's values when it is applied to the full population data.

However, MP estimators differ from one another in a vitally important way: Because they are based on different design matrices, and because different design matrices change the location of the origin relative to the solution line, they will produce quite different estimates. In general the estimates of the linear effects produced by an MP estimator, including the IE, will vary depending on at least three aspects of the data: first, the number of APC groups in the data set, e.g., the number of periods; second, the size and sign of the nonlinearities; and third, the choice of the reference category. This lack of robustness is a serious weakness.

In a series of papers and books, Fu, Land, and their various coauthors have argued for the IE (e.g., see Yang & Land 2013), an MP estimator based on a design matrix of categorical variables parameterized as sum-to-zero effects. Fosse & Winship (2018) suggest that if one is going to use an MP estimator, one should use one based on a linearized design matrix, that is, a matrix in which the linear effects of age, period, and cohort are orthogonal to their nonlinear terms. They call this the orthogonal estimator (OE). Unlike other MP estimators, the OE estimates of the linear effects are not affected by the number of APC groups, the size or sign of the nonlinear effects, or the reference category used. However, there is no reason to think that the OE or any other MP estimator will produce reliable estimates of the actual APC effects.

In general, we are not enthusiastic about MP estimators. We are in agreement with Duncan & Stenbeck (1988), who "consider it a methodological absurdity to propose all-purpose models for 'dependent variables' as diverse as tuberculosis mortality, human fertility, educational attainment, and electoral participation." For arguments in their favor, particularly for the IE, the reader should consult the extensive literature by Fu, Land, and coauthors (for recent books on the topic, see Fu 2018, Yang & Land 2013). If one is going to use an MP estimator, we recommend using the OE because it is not affected by the number of periods or other aspects of the data structure. If one is going to estimate temporal effects with the OE (or other MP estimator), then we recommend using the OE's estimates as a starting point for APC analysis, not as an end point. In particular, as we discuss below, a researcher should use a 2D-APC plot to examine whether or not other points result in qualitatively different effects that are equally theoretically reasonable. It is possible that a set of estimates from the OE (or IE) may seem plausible. This, however, does not rule out the possibility that other estimates, which are just as consistent with the data, may be equally as credible and, as such, should be discussed.

3.1.2.2. Hierarchical age-period-cohort model. The HAPC model has been used in a variety of studies, including analyses of obesity, mental health, and verbal ability. This approach entails specifying a multilevel (or hierarchical) model with age and age-squared terms and varying intercepts for the period and cohort groups. In matrix notation the HAPC model can be represented as

$$\mathbf{y} = \mathbf{W}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}, \qquad 7.$$

where $\mathbf{y}$ is an $N \times 1$ column vector of outcome values; $\mathbf{W}$ is an $N \times 3$ design matrix for the overall intercept, age, and age-squared fixed effects (or additional higher level polynomials); $\boldsymbol{\beta}$ is a 3×1 column vector for the overall intercept, age, and age-squared fixed effects; $\mathbf{Z}$ is an $N \times (J + K)$

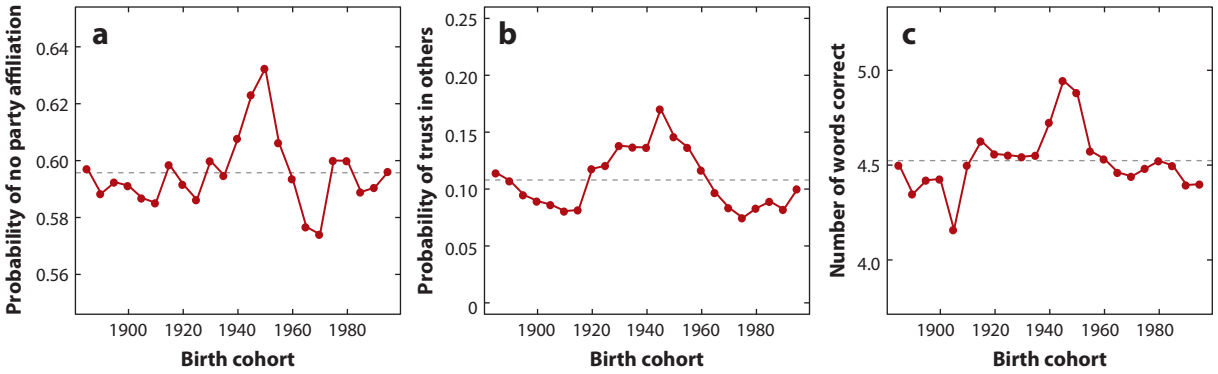

Figure 2

The hierarchical age-period-cohort (HAPC) and estimates of cohort effects. Panel *a* examines the probability of having no political party identification. Panel *b* shows the probability that one "can trust others." Panel *c* shows results for number of words correct out of a ten-item vocabulary test. All data are from the US General Social Survey (1972–2016). The horizontal dashed line indicates the overall (or grand) mean as given by the intercept in the HAPC.

design matrix for the period and cohort random effects; $\mathbf{u}$ is a $(J + K)$ column vector containing the period and cohort random effects; and $\boldsymbol{\epsilon}$ is an $N \times 1$ column vector for the individual-level error term. The vector of random effects is assumed to be normally distributed around a zero mean [i.e., each $u \sim N(0, \sigma_u^2)$, as are the individual-level error terms, i.e., $\epsilon \sim N(0, \sigma_\epsilon^2)$]. Written as a single equation, the HAPC is

$$Y_{ijk} = [\mu + \alpha_1(\text{age}) + \alpha_2(\text{age}^2)] + [\phi_j + \psi_k] + \epsilon_{ijk}, \qquad 8.$$

where μ is the mean, α_1 and α_2 are the parameters for the age and age-squared fixed effects, ϕ_j are the period random effects, ψ_k are the cohort random effects, and ϵ_{ijk} is random error. The model can be estimated using Bayesian (e.g., Markov chain Monte Carlo) or frequentist (e.g., maximum likelihood) approaches. There is a two-step procedure using residual maximum likelihood, a procedure typically used in empirical applications (see O'Brien 2017).

As discussed by several authors (Bell & Jones 2018, Luo & Hodges 2019, O'Brien 2017), the HAPC imposes a set of nonobvious constraints on the estimates. Typically, such constraints end up setting the linear cohort effect to zero or close to zero. Intuitively, the reason for this appears to be because the cohort variable has the largest number of categories, so for the same slope the cohort random effects will have greater total variance compared with the period random effects (Bell & Jones 2018, pp. 786–88; Luo & Hodges 2019, p. 24).[12] The assumption that the cohort linear effect is near zero can be confirmed visually by noting that in empirical applications, the HAPC consistently generates results with a zero linear cohort effect. For instance, in **Figure 2**, we show the cohort random effects for the probability of being a political independent, the probability of trusting other people, and the number of words correct on a ten-item vocabulary test. All results are based on the HAPC. These findings clearly show, consistent with the body of empirical research on the HAPC, that the cohort effects have a zero overall slope. In fact, the results are nearly identical to simply estimating a model with a zero linear cohort effect except without the shrinkage property of the random effects (i.e., the fact that, for cells with few observations and high internal variance, estimates shrink toward the overall mean).

[12] More or less coarse groupings of APC variables will also influence results by affecting the variation in the data (Luo & Hodges 2016).

To reiterate, the consequence of the strong assumptions imposed by the HAPC model can be seen by inspecting the published results using the HAPC; studies on verbal ability, social trust, confidence in institutions, subjective happiness, and obesity have all assumed that the cohort linear effect is zero (for example, see figure 5 in Reither et al. 2009). In contrast, Ryder, Mannheim, and other thinkers have long theorized that intercohort differences account for massive social and cultural change. Yet, the HAPC will, in general, arbitrarily fix the cohort slope near to zero, giving results that assume trendless fluctuation.

3.2. Point Identification with Measured Causes

As discussed in Section 2.4, at the heart of any APC analysis is a theoretical model that specifies how various causal processes, what we call here mechanisms,[13] are associated with each of the APC variables. What is peculiar to many APC analyses is the attempt to assess the relative importance of the three APC processes without relying on either theory or observed measures of their associated causal mechanisms.

The first instance of a mechanism-based approach in an APC analysis was by Duncan (1985), who proposed using education as a mechanism between cohort and an outcome. The existence of observed measures of associated processes has two advantages. First, with assumptions, the total effects of the APC variables may be identified. Second, in some cases, models—or more precisely some components of a model—will be overidentified, allowing one to jointly test the consistency of one's assumptions. In general the more measures one has of the associated processes and related assumptions, the more likely that the model will be at least just identified, if not overidentified. As is the case with the methods using equality constraints discussed in Section 3.1, overidentification can be of great assistance in supporting the credibility of a model's assumptions.

3.2.1. Proxy variables. Starting in the 1980s, researchers began using what is now called the proxy variables approach, which involves replacing one of the temporal variables with one or more mechanisms (Glenn 2005, Heckman & Robb 1985).[14] For example, Farkas (1977) replaced the period variable with unemployment rates to examine the female employment rate across age and cohort groups. Other examples include cohort size (Kahn & Mason 1987), cohort exposure to lead (McCall & Land 2004), and cohort smoking behavior (Preston & Wang 2006).

The proxy variables approach is typically viewed as simply substituting one of the temporal variables for another nontemporal variable. There is considerable advantage, however, in understanding this approach as assuming that there is a single mechanism, here M_P, representing all of the direct association of P with Y. As specified, the model can be estimated in two stages. Using the example of female employment rate in the previous paragraph, in the first stage, one would regress the unemployment rate on the period variable. In the second stage, one would regress the female employment rate, Y, on age, cohort, and the unemployment rate.

The model as specified makes two strong assumptions. The first assumption is that neither age nor cohort influence the unemployment rate, an assumption that is testable. The second assumption is that M_P is the only mechanism linking P with Y. This can be assessed using conventional methods by testing the hypothesis that period's nonlinear effects are zero.[15]

[13] O'Brien (2014) uses the term "characteristic" in discussing the type of models in this section.

[14] When the proxy refers to the cohort variable, O'Brien (2014) terms these "age-period-cohort characteristic" models.

[15] A problem can occur if the proxy variable only depends linearly on its APC variable. In this case, we are simply back to the original identification problem in that age, cohort, and the unemployment rate are linearly dependent.

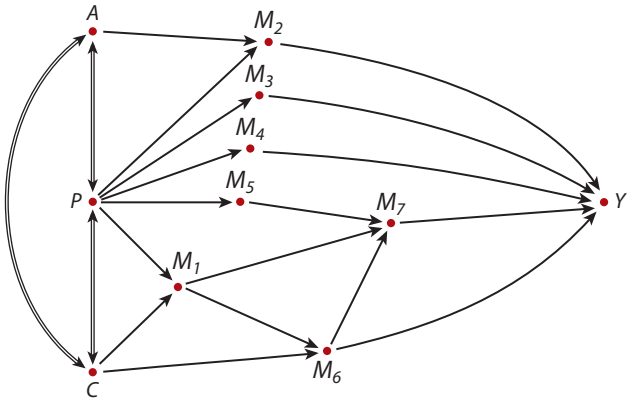

Figure 3

This graph shows a full-fledged mechanism-based age-period-cohort model of political alienation. Double-line, double-headed arrows denote the linearly dependent relationships among the three temporal variables. All variables are observed in the graph. *A*, *P*, and *C* refer to age, period, and cohort, while *Y* denotes the outcome, political alienation. The causal mechanisms are defined as follows: M_1, cohort size; M_2, church attendance; M_3, Watergate; M_4, Republican president; M_5, unemployment rate; M_6, educational level; M_7, employment status. Figure adapted with permission from Winship & Harding (2008).

3.2.2. Mechanism-based approaches. As noted above, the proxy variable approach makes two strong assumptions: that any mechanism is only affected by a single APC variable and that the association between any APC variable and the outcome runs through a single mechanism. Winship & Harding (2008) show that neither of these two assumptions are desirable or necessary. **Figure 3** shows their model for political alienation based on General Social Survey (GSS) data.

As can be seen in this figure, both period and cohort work through multiple mechanisms and cohort size is associated with both period and cohort. Winship and Harding's approach brings a full-fledged mechanism-based strategy to APC analysis. They show in particular that Pearl's (2009) front door criterion allows for specific APC models to be just-identified or even overidentified. A key result of Winship & Harding (2008) is that if one has measures of all the causal processes indexed by at least one of the APC variables, then the total effects for all three APC variables will be identified. This can be assessed by testing the hypothesis that the nonlinear effects of the variable on the outcome are zero after conditioning on its associated mechanisms. Winship and Harding's model is vastly overidentified, containing 62 degrees of freedom (if it were just identified, it would have zero degrees of freedom). They present chi-square, Bayesian information criterion (BIC), and Akaike information criterion (AIC) statistics that all show that the model fits the data very well.

In general, the more overidentified the model, the more opportunity there is for the model to be inconsistent with the data. Without additional information, all APC models with just-identifying constraints fit the data equally well. With respect to the data, these models are empirically equivalent (see Lee & Hershberger 1990). Distinctions among empirically equivalent models must be made on the basis of theoretical grounds. It is far better if one can also use data to justify one's model and its underlying assumptions.

4. PARTIAL IDENTIFICATION USING BOUNDING AND SENSITIVITY ANALYSES

As we hope the previous section has convinced the reader, point identification of the total effects of the three APC variables involves strong assumptions. Unless these assumptions are theoretically

motivated and/or tested against the data, it is unclear how one should interpret one's estimates, much less have any confidence in them. In this section we discuss an approach based on partial identification of the total effects using bounds (see Fosse & Winship 2019), a less demanding goal and, as such, one requiring weaker assumptions. Using bounds, we can achieve estimates of the total effects of the APC variables that are typically far more defensible because the strong assumptions required for point identification are not needed. That noted, as the empirical example reveals, and our analysis of other empirical examples has found (Fosse & Winship 2019), a partial identification analysis can often give bounds that are sufficiently tight that one has, in essence, obtained point identification.

4.1. The Canonical Solution Line

To understand the bounding approach, it is crucial to recognize the geometric interpretation of the linear dependence problem (O'Brien 2011, 2014). In the discussion that follows, we let $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$. These are estimable quantities that determine the location of the canonical solution line in the parameter space (that is, each data set will produce differing values of these parameters and thus different trade-offs for setting various constraints). In the absence of data, the age, period, and cohort slopes may take on any combination of values in a 3D space. For convenience, in **Figure 4*a*** we show the age-period plane defined by the identified quantity $\theta_1 = 3$, while in **Figure 4*b*** we show the period-cohort plane defined by the identified quantity $\theta_2 = -2$. The intersection of these two planes then defines a line, as demonstrated in **Figure 4*d***. This line is known as the solution line as all points on the line represent parameter estimates for α, π, and γ that are equally consistent with the data (O'Brien 2011, 2014). This is a visual representation of the APC identification problem. If there were no linear dependency, then we would have three planes intersecting at a single point in the parameter space. Any constrained set of estimates α^*, π^*, and γ^* will lie on this line, which is mathematically equivalent to assuming a particular value of the unknown overall linear component given by the scalar ν in Equation 3.

It is crucial to understand what is accomplished when data are applied to an APC model. In the simple case where there are only linear effects, the data can take us from a 3D space where all parameter values are possible to a 1D space where only combinations of estimates lying on a line are consistent with the data. This same reduction also holds if our model has nonlinear effects, because they are fully identified. This has not been widely recognized in the current APC literature in sociology and demography. As Fosse & Winship (2018) show, the solution line can always be reduced to just three dimensions. They call this the canonical solution line. Accordingly, the data go a long way toward constraining the possible estimates for the age, period, and cohort slopes, restricting an initial set of parameters that could be anywhere in a high-dimensional space representing a full set of possible linear and nonlinear effects to values lying on a single line. As such, the data have been quite informative about parameter values, just not quite as informative as we might like in the sense of providing unique estimates for the linear effects.

4.2. Bounding Regions and 2D-APC Graphs

There is a convenient way to further simplify the representation of the solution line in the previous section. Because of the linear relationships $\alpha + \pi = \theta_1$ and $\pi + \gamma = \theta_2$, we can reduce our 3D representation to just two dimensions. One way of doing this is by having the horizontal axis represent the period slope, the left vertical axis represent the age slope, and the right vertical axis represent the cohort slope. Fosse & Winship (2019) call this a 2D-APC graph.

Continuing with our example, in **Figure 5** we show a 2D-APC graph based on values of $\theta_1 = 3$ and $\theta_2 = -2$. The solution line shown in **Figure 5** is identical to that shown in **Figure 4**. We

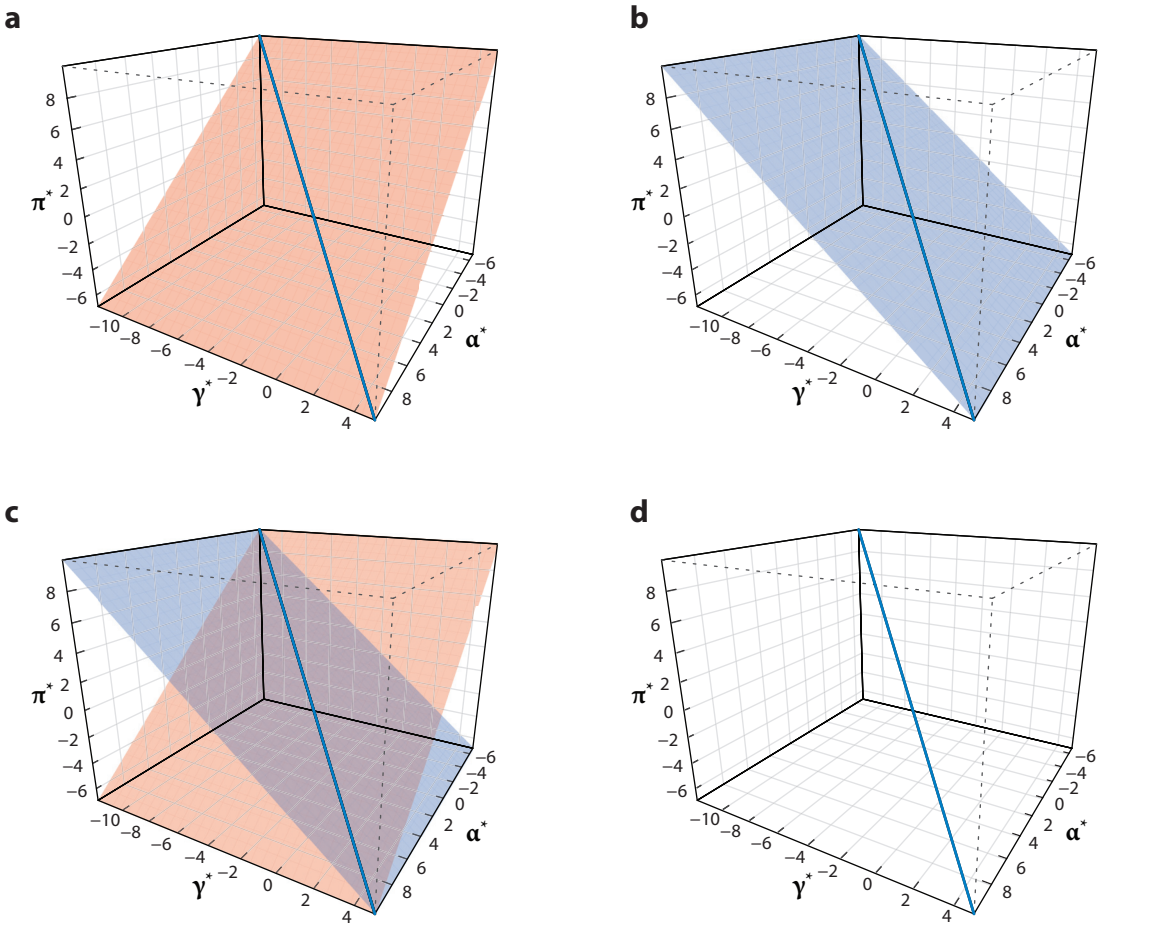

Figure 4

Geometric derivation of the canonical solution line (*blue diagonal lines*). Panel *a* shows the age-period plane (*red area*) defined by $\pi = \theta_1 - \alpha$, where $\theta_1 = 3$. Panel *b* shows a period-cohort plane (*blue area*) defined by $\pi = \theta_2 - \gamma$, where $\theta_2 = -2$. Panel *c* visualizes the intersection of the age-period and period-cohort planes, while panel *d* visualizes the solution line based on $\theta_1 = 3$ and $\theta_2 = -2$.

specify each point in the coordinate space in terms of age, period, and cohort so that, for example, the point (1, 3, −5) refers to $\alpha = 1$, $\pi = 3$, and $\gamma - 5$. The solution line runs from the upper left to the bottom right and the three dotted lines indicate, respectively, when the age, period, and cohort coefficients are equal to zero.

As shown in **Figure 5**, and as is the case for all APC graphs, the slope of the solution line relating period to age and cohort will always be −1. The values of θ_2 and θ_1 determine two things. First, the difference between θ_2 and θ_1, −5 in **Figure 5**, determines the offset between the age and cohort scales. Thus, in **Figure 5**, when $\alpha = 0$, $\gamma = -5$. Second, at the point where $\pi = 0$, θ_1 and θ_2 determine, respectively, the location of the solution line in terms of α and γ. Various traditional APC estimators can be located on the 2D-APC graph and can be understood as making particular assumptions about the linear age, period, and cohort effects (see **Supplemental Appendix A**).

Supplemental Material >

The 2D-APC graph also clarifies an important, unrecognized fact in the APC literature that by fixing the location of the solution line, the data also determine which of the eight

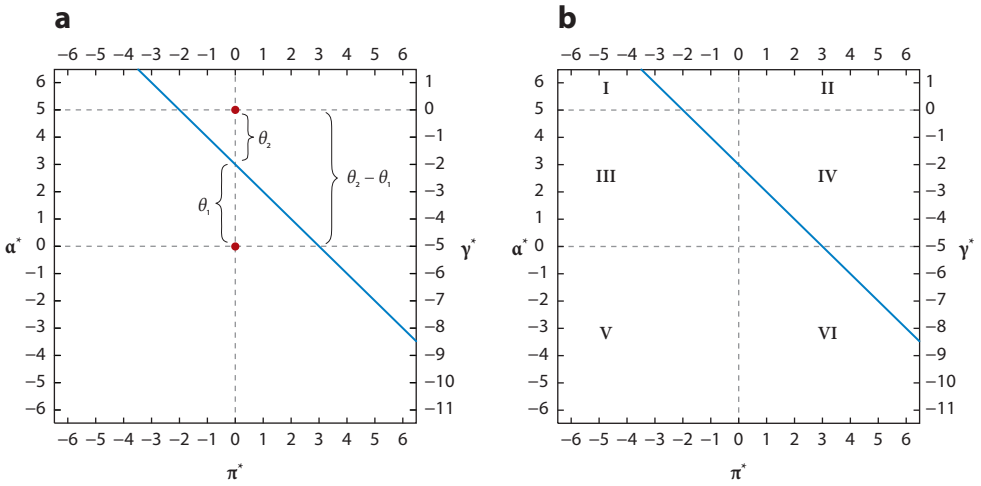

Figure 5

2D-APC graph. Panel *a* reveals the 2D-APC graph is a function of the values of θ_1 and θ_2, where $\theta_1 = \alpha + \pi$ and $\theta_2 = \pi + \gamma$, and the variables α, π, and γ signify the age, period, and cohort linear effects, respectively. As in **Figure 4**, $\theta_1 = 3$ and $\theta_2 = -2$, and the solution line is indicated in blue. The three dotted lines indicate when the age, period, and cohort linear effects are equal to zero. Panel *b* shows the various signed plotting regions: I (+, −, +), II (+, +, +), III (+, −, −), IV (+, +, −), V (−, −, −), and VI (−, +, −). Note that there is no region representing either the (−, +, +) or (−, −, +) combinations of linear effects.

combinations of positive and negative age, period, and cohort effects are empirically possible. Because the offset between age and cohort must be either positive or negative as determined by the difference between the θs, only six combinations can exist. Then, depending on the location of the solution line, as few as two and as many as four remaining combinations might be possible.

Figure 5 shows six regions of the parameter space defined by the signed combinations of the slopes (α, π, γ). These are regions I (+, −, +), II (+, +, +), III (+, −, −), IV (+, +, −), V (−, −, −), and VI (−, +, −). Note first that there is no region representing either the (−, +, +) or (−, −, +) pattern of effects. Thus, theories that posit a linear negative age but positive period and cohort effects, as well as those that posit a linear negative age and period but positive cohort effects, have been ruled out by the data alone. In addition, the solution line runs through only four out of the six regions, so we can also rule out any social or biological theory that posits that the linear age, period, and cohort effects are all positive (region II) or all negative (region V). Thus, despite the linear dependence problem, the data can eliminate a number of theoretical possibilities.

In order to make further progress in estimating the effects of age, period, and cohort, theory is needed. For the moment, assume that age, period, and cohort have only linear effects. If we are willing to make an assumption about the sign of one or more of the APC variables, we can often draw conclusions about the sign and even the size of the remaining APC variables. For example, suppose the outcome of interest is political ideology on a simple left-right scale, where higher values indicate a more left-leaning ideology. One might think it reasonable to assume that the age linear effect is negative (i.e., that people become less liberal as they age). Under this assumption, the only pattern of effects that is consistent with the solution line are those associated with region VI (−, +, −), that is, if the age effect is negative, then the period effect must be positive and the cohort effect negative. Furthermore, we would know that the period effect had to be greater than 3 and the cohort effect less than −5.

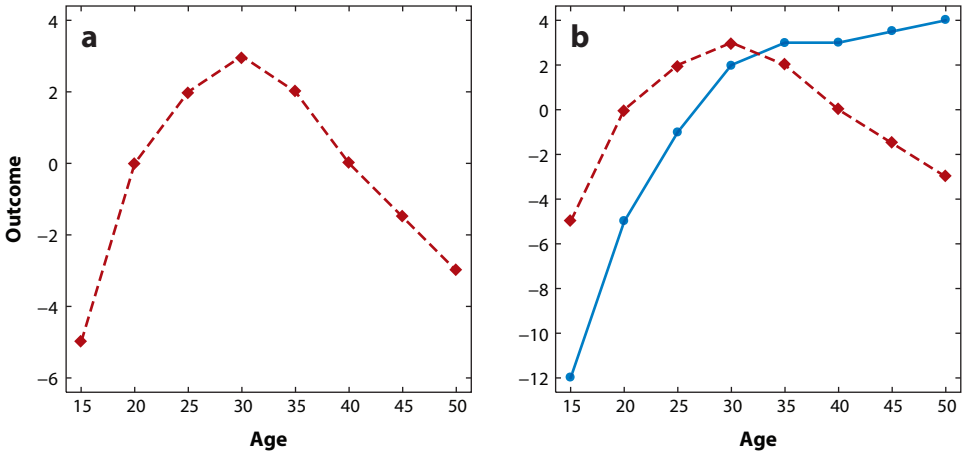

Figure 6

Specifying monotonicity constraints. Suppose we have reason to believe the overall age effect is monotonically increasing, and we want to specify a value for the linear age effect that ensures this is the case. The red dashed lines show simulated nonlinearities for age. The solid blue line in panel *b* reflects the assumption that $\alpha = 2$, which is the minimum age slope required for a monotonic increasing overall age effect.

Of course, most APC data contain linear and nonlinear effects. In addition, in most situations, theory does not imply that linear effects should be positive or negative, but rather that the effect of a variable should be monotonically increasing or decreasing over some range of the variable. As we now explain, assumptions about monotonicity bound both the sign and size of linear effects.

Assume that we have strong theoretical reasons to believe that the overall age effect is monotonically increasing. Now consider the dotted line in **Figure 6**, which graphs the nonlinear effects of age. As can be seen, the effect is first increasing and then decreasing. We want to specify a value for the linear age effect that ensures that the effect in age is monotonically increasing. This means that between any two adjacent age categories, the effect is at least flat, if not upward sloping. To do so, we need only find that pair of adjacent age categories in which the downward effect is most negative. For example, suppose that the forward differences for the age nonlinearities are $\Delta\widetilde{\alpha}_{I-1} = \{5, 2, 1, -1, -2, -1.5, -1.5\}$. The minimum of these differences is -2, which is between ages 35 and 40. To counter this downward deviation, the parameter value for the linear age term must be greater than or equal to $+2$. In this figure the solid line shows what happens to the overall age effect when the age slope is set to $+2$. As can now be seen, the overall effect of age is now monotonically increasing. This will be true as long as the linear age effect is greater than or equal to $+2$.[16]

Bounds analysis is not necessarily restricted to using a single bound. Fosse & Winship (2019) conduct an APC analysis of homicide rates in the United States. They assume that the effect of age is monotonically increasing during adolescence but monotonically decreasing after young adulthood. This leads to very narrow bounds for the effects of period and cohort. They further restrict the period effect to be nondecreasing during the crack epidemic of the second half of the 1980s. This essentially gives point estimates. An additional advantage of using multiple bounds for more than one APC variable is that there may be no portion of the solution line that is consistent

[16]With some data sets, it may be useful to smooth the observed effects in age, period, or cohort in the Lexis table to rule out the possibility of unreasonable sharp nonlinearities in the data affecting one's monotonicity analysis.

with one's assumed bounds. In this sense one is testing the theory as specified by a set of bounds against the data.

Bounding analyses can be easily incorporated with mechanism-based models to produce sensitivity analyses. A bounds analysis begins with the widest possible range of estimates, which can be narrowed down to a point estimate using progressively stronger assumptions. In contrast, a sensitivity analysis begins with a point estimate and then weakens the assumptions to consider a wider range of estimates. In other words, they come from different starting points, but bounding analyses and sensitivity analyses are essentially equivalent. To conduct a sensitivity analysis, one can first use a mechanism-based model to obtain a point estimate and then specify upper and lower bounds representing the sensitivity due to omitted pathways.[17] For example, in their study of political alienation, Winship & Harding (2008) construct bounds by theorizing about the sign and size of the effect of cohort on church attendance, which is omitted in their model due to linear dependence. The resulting bounds are small, and their main conclusion of the dominance of period effects in understanding changes in political alienation remains unchanged. As Winship & Harding (2008, p. 392) point out, it is often the case that "a partially identified model may be quite informative if the bounds on the effects of the APC variables are relatively narrow."

4.3. A General Framework for Age-Period-Cohort Analysis

The insights outlined above provide general guidelines for doing an APC analysis. A basic principle is that a researcher should attempt to learn as much as possible from the data while making the fewest assumptions possible. We recommend the following steps when conducting an APC analysis:

1. Linearized design matrix: Separate the linear from the nonlinear components using a linearized design matrix. Using the linearized design matrix, fit a model with the period linear effect fixed to zero.[18]
2. Identified quantities: Report the identifiable combinations of linear effects (e.g., $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$) that determine the location of the canonical solution line in the parameter space. Report the full set of nonlinear effects (e.g., $\widetilde{\alpha}$, $\widetilde{\pi}$, and $\widetilde{\gamma}$). Conventional significance tests and fit statistics can be applied.
3. Canonical solution line and the 2D-APC graph: Visualize the canonical solution line using a 2D-APC graph. Because the nonlinear effects are point-identified, they can be visualized using traditional graphical techniques.
4. Partial identification using bounding analyses: Specify a series of bounds using explicit theoretical assumptions about the size, sign, and/or shape of the temporal effects. If one is willing to make strong assumptions, then the bounds will reduce to point estimates on the canonical solution line. In specifying bounds, one should ask whether there is value in making stronger assumptions to obtain more precise estimates. If not, proceed to the next step.
5. If observed measures of causes are available:
   (*i*) Mechanism-based models: Using the observed measures of causes, fit one or more mechanism-based models to obtain point estimates of the temporal effects. Overidentification tests may be applied.

[17]For an excellent overview of sensitivity analyses with underidentified structural equation models, see Land & Felson (1978).

[18]Methods for constructing a linearized design matrix can be found in most textbooks on matrix algebra; also see Fosse & Winship (2018).

(*ii*) Sensitivity analyses: After fitting mechanism-based models, consider conducting a sensitivity analysis with a 2D-APC graph to assess the robustness of findings in the presence of unobserved causal pathways.

6. Development of theoretical models: Make tentative conclusions oriented toward building more detailed, richer theoretical models of temporal effects.

4.4. Empirical Example with Verbal Ability

For the purposes of illustration, we examine APC effects on verbal ability. Using the GSS, we obtain data on $N = 23,824$ respondents. We restrict the sample to white males born in the United States who are 30 years or older (i.e., who are likely no longer enrolled in college).[19] The outcome is the number of words correct on a basic vocabulary quiz of ten items. For simplicity of exposition, we assume the outcome is continuous.[20] Age and period are grouped into five-year intervals.

We followed the general procedure outlined in the previous section to analyze these data. First, we constructed a linearized design matrix that separates the linear from the nonlinear components. In our design matrix, we fixed the period linear effect to zero in order to obtain estimates of θ_1 and θ_2, which were $\widehat{\theta_1} = -0.122$ ($t = -1.466$, $p = 0.143$) and $\widehat{\theta_2} = 0.384$ ($t = 4.109$, $p < 0.001$).

Second, we inspected the nonlinear effects. F-tests indicated that the nonlinear effects are statistically significant for all three variables. Third, we graphed the canonical solution line using a 2D-APC graph (see **Figure 7*a***) as well as the nonlinear effects (see **Figure 10** in **Supplemental Appendix B**). The 2D-APC graph in **Figure 7*a*** reveals that we can rule out all parameters lying in regions II and V (note that the canonical solution line runs through a small part of region III). In other words, the three temporal linear effects cannot be all positive or all negative. We also further divided the 2D-APC graph into additional regions based on monotonicity constraints. These regions are shown in **Figure 11** in **Supplemental Appendix B**.

Next, we conducted a series of basic bounding analyses, which are shown in **Figure 7*b*,*c***. In **Figure 7*b***, we assume that the linear effect for age is nonnegative, reflecting a basic assumption that the verbal ability variable in the GSS captures crystallized rather than fluid intelligence. With this assumption, we can conclude that the period linear effect is negative and the cohort linear effect must be greater than 0.506. In **Figure 7*c***, we impose monotonicity constraints on the age effect. Specifically, we assume that the age effect increases monotonically to age 45 and is not monotonically increasing after age 45. In other words, the age effect after 45 can monotonically decrease, remain the same, or increase (but not necessarily monotonically so). This is equivalent to the assumption that $-0.031 \leq \alpha \leq 2.755$. When combined with the nonlinear effects, we obtain the bounded effects shown in **Figure 8*a–c***.

At this point we would quit our analysis if we had no measured causes. Because the GSS is a rich source of other variables, we proceed to the fifth step and conduct a mechanism-based analysis. As a basic mechanism-based model, we use years of education (ranging from zero to 20) as a mechanism between period and cohort and the outcome. This reflects the widely held belief that the rise in verbal ability scores throughout the twentieth century is at least partly attributable to the expansion of higher education (Schooler 1998). The mechanism-based model includes a full set of nonlinearities that have a direct effect on the outcome and indirect effects via the education variable. Using the product rule, the mechanism-based model produces a total linear effect of $\widehat{\gamma} = 0.888$.

[19]Note that this sample restriction increases the plausibility of no direct path between age and years of education.

[20]We obtain similar results treating the outcome as a count variable.

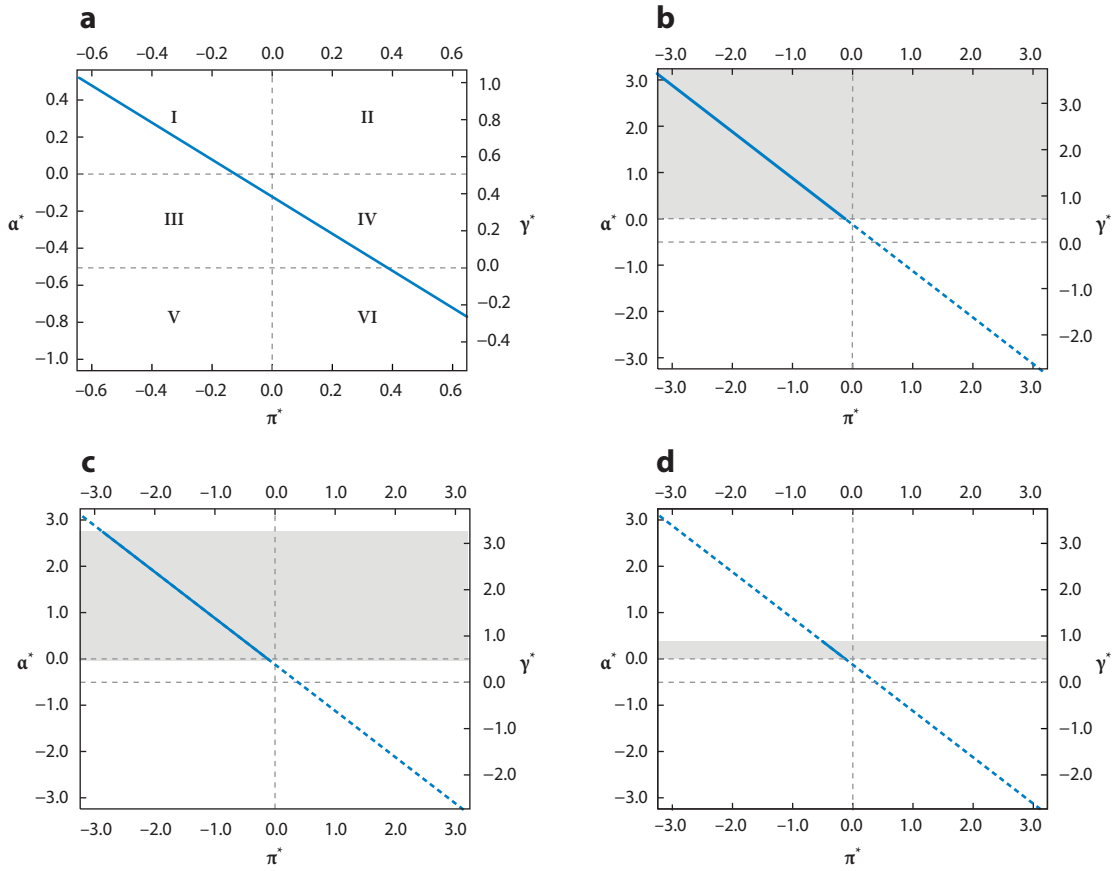

Figure 7

Verbal ability: 2D-APC graph with bounds. Panel *a* indicates the canonical solution line (*solid blue line*) on the 2D-APC graph for verbal ability, as measured by the US General Social Survey, with the boundaries for the signed plotting regions as shown in **Figure 5**. Panel *b* displays bounds on the 2D-APC under the assumption that the linear age effect is non-negative. The shaded area indicates the regions with possible solutions for each parameter, and the dashed blue line indicates portions of the canonical solution line that have been ruled out. Panel *c* displays bounds under monotonicity constraints for the age effect, such that $-0.031 \leq \alpha \leq 2.755$. Panel *d* shows bounds under the assumption that the age linear effect is positive and that $\widehat{\gamma} = 0.888$ is an upper bound on the cohort linear effect.

However, it is likely that several mechanisms are missing, so we proceed to the sixth step and conduct a bounding analysis. As argued by Glenn (1994), increases in television use and declines in newspaper consumption likely drive cohort effects in verbal ability. Glenn argues that these negatively impact verbal ability, suggesting that the estimated linear effect for cohort using just education as a mechanism is an upper bound on the true cohort linear effect.[21] We also assume that the age effect is monotonically increasing from ages 30 to 45. These bounds and the point estimate for the mechanism-based model are shown in **Figure 7*d***. When combined with the nonlinear effects, we obtain the overall temporal effects in **Figures 8*d–f***. These results reveal that cohort effects dominate in explaining changes in verbal ability. Finally, we conclude by specifying that

[21]As Glenn (1999, p. 270) contends, "the main 'something' that has offset the effects of increased education is a decline in reading."

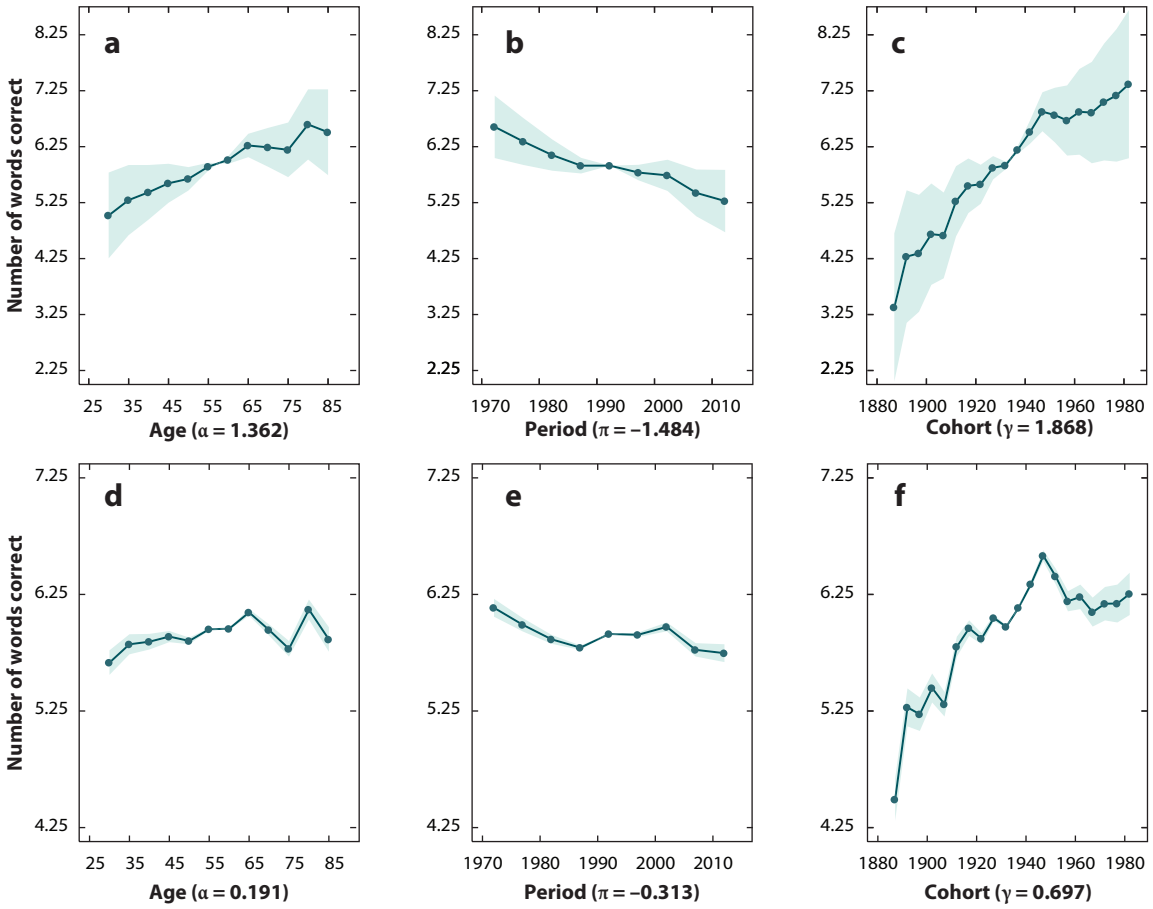

Figure 8

Verbal ability: upper and lower bounds on APC effects. Panels *a–c* show the overall bounded APC effects under monotonicity constraints for the age effect, such that $-0.031 \leq \alpha \leq 2.755$. Panels *d–f* show the overall bounded APC effects under the assumption that the age linear effect is positive and that $\widehat{\gamma} = 0.888$ is an upper bound on the cohort linear effect. Solid lines indicate upper and lower bounds, while the dotted line denotes the midpoint values between the bounds. Abbreviation: APC, age-period-cohort.

theoretical models of temporal effects should include not only education but also measures of media consumption, which are thought to be linked to the cohort timescale (Glenn 1999).

5. CONCLUSION

Social and cultural change is a fundamental topic in the social sciences, especially within sociology. The great classical thinkers in sociology—Weber, Marx, and Durkheim—all placed understanding social and cultural change at the core of their works. In sociology and demography, APC models are a popular, general set of techniques for understanding temporal effects. Over the decades, a variety of methods have been proposed to identify the point estimates of the effects of each of the three APC variables. However, there has been little agreement on the validity of any proposed approaches, and findings from APC models remain highly controversial.

In the first part of this article, we reviewed the basics of APC analysis, discussing how APC effects can be interpreted as causal structures. In the second part of the article, we examined a

range of APC techniques used by researchers. Traditional methods are only potentially useful if researchers are willing to recognize the strong assumptions behind them and justify these assumptions theoretically and, in the case of overidentified models, test them against the data. Far too seldom is this done. Researchers have been too ambitious in immediately desiring point estimates of the total effects of age, period, and cohort and then remiss in failing to recognize the costs in pursuing this goal. Much can be learned from APC data using methods that are more transparent and make much weaker assumptions. In the third part of the article, we reviewed an approach based on partial identification using bounds. The goal has been to clarify what can be learned from the data itself and what can be concluded using theoretical assumptions about the size, sign, or shape of the temporal effects. As is generally the case with methodological advances, there is the question of whether past substantive findings will be sustained or overturned. This is a critical task for future APC analysis given that many of the methods used have often relied on untested or untestable assumptions.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

Alwin DF. 1991. Family of origin and cohort differences in verbal ability. *Am. Sociol. Rev.* 56:625–38

Andvord KF. 1930. Hvad kan vi lære ved å følge tuberkulosens gang fra generasjon til generasjon? [What can we learn by following the development of tuberculosis from one generation to another?] *Norsk Mag. Lægevidenskapen* 91:642–60

Bell A, Jones K. 2018. The hierarchical age-period-cohort model: Why does it find the results that it finds? *Qual. Quantity* 52.2:783–99

Blinder A. 1973. Wage discrimination: reduced form and structural estimates. *J. Hum. Resour.* 8:436–455

Brooks C, Manza J. 1997. Social cleavages and political alignments: U.S. presidential elections, 1960 to 1992. *Am. Sociol. Rev.* 62:937–46

Case RAM. 1956. Cohort analysis of mortality rates as an historical or narrative technique. *Br. J. Prev. Soc. Med.* 10:159–71

Chaves M. 1989. Secularization and religious revival: evidence from U.S. church attendance rates, 1972–1986. *J. Sci. Study Religion* 28:464–77

Chen X, Li G, Unger JB, Liu X, Johnson CA. 2003. Secular trends in adolescent never smoking from 1990 to 1999 in California: an age-period-cohort analysis. *Am. J. Public Health* 93:2099–104

Clark AK, Eisenstein MA. 2013. Interpersonal trust: an age–period–cohort analysis revisited. *Soc. Sci. Res.* 42:361–75

Clayton D, Schifflers E. 1987. Models for temporal variation in cancer rates. II: Age-period-cohort models. *Stat. Med.* 6:469–81

Clogg CC. 1982. Cohort analysis of recent trends in labor force participation. *Demography* 19:459–79

Converse PE. 1976. *The Dynamics of Party Support: Cohort-Analyzing Party Identification*. Thousand Oaks, CA: SAGE

Demartini JR. 1985. Change agents and generational relationships: a reevaluation of Mannheim's problem of generations. *Soc. Forces* 64:1–16

Derrick VPA. 1927. Observations on (1) errors of age in the population statistics of England and Wales, and (2) the changes in mortality indicated by the national records. *J. Inst. Actuaries* 58:117–59

Diouf I, Charles MA, Ducimetière P, Basdevant A, Eschwege E, Heude B. 2010. Evolution of obesity prevalence in France: an age-period-cohort analysis. *Epidemiology* 21:360

Duncan OD. 1985. Generations, cohorts, and conformity. In *Cohort Analysis in Social Research: Beyond the Identification Problem*, ed. WM Mason, S Fienberg, pp. 289–321. New York: Springer

Duncan OD, Stenbeck M. 1988. Panels and cohorts: design and model in the study of voting turnout. *Sociol. Methodol.* 18:1–35

Fienberg SE. 2013. Cohort analysis' unholy quest: a discussion. *Demography* 50:1981–84

Fienberg SE, Mason WM. 1979. Identification and estimation of age-period-cohort models in the analysis of discrete archival data. *Sociol. Methodol.* 10:1–67

Firebaugh G. 1989. Methods for estimating cohort replacement effects. *Sociol. Methodol.* 19:243–62

Firebaugh G, Harley B. 1991. Trends in U.S. church attendance: secularization and revival or merely lifecycle effects? *J. Sci. Study Religion* 30:487–500

Fosse E, Winship C. 2018. Moore-Penrose estimators of age-period-cohort effects: their interrelationship and properties. *Sociol. Sci.* 5:304–34

Fosse E, Winship C. 2019. Bounding analyses of age-period-cohort effects. *Demography*. In press

Frost WH. 1939. The age selection of mortality from tuberculosis in successive decades. *Am. J. Hyg.* 30:90–106

Fu Q, Land KC. 2015. The increasing prevalence of overweight and obesity of children and youth in China, 1989–2009: an age-period-cohort analysis. *Popul. Res. Policy Rev.* 34:901–21

Fu W. 2018. *A Practical Guide to Age-Period-Cohort Analysis: The Identification Problem and Beyond*. Boca Raton, FL: CRC

Fu WJ. 2000. Ridge estimator in singular design with application to age-period-cohort analysis of disease rates. *Commun. Stat. Theory Methods* 29:263–78

Fu WJ. 2016. Constrained estimators and consistency of a regression model on a Lexis diagram. *J. Am. Stat. Assoc.* 111:180–99

Fu WJ, Hall P. 2006. Asymptotic properties of estimators in age-period-cohort analysis. *Stat. Probab. Lett.* 76:1925–29

Fu WJ, Land KC, Yang Y. 2011. On the intrinsic estimator and constrained estimators in age-period-cohort models. *Sociol. Methods Res.* 40:453–66

Ghitza Y, Gelman A. 2014. *The great society, Reagan's revolution, and generations of presidential voting*. Work. Pap., Dep. Political Sci., Columbia Univ., New York

Glenn ND. 1994. Television watching, newspaper reading, and cohort differences in verbal ability. *Sociol. Educ.* 67:216–30

Glenn ND. 1999. Further discussion of the evidence for an intercohort decline in education-adjusted vocabulary. *Am. Sociol. Rev.* 64:267–71

Glenn ND. 2005. *Cohort Analysis*. Thousand Oaks, CA: SAGE

Hauser RM, Huang MH. 1997. Verbal ability and socioeconomic success: a trend analysis. *Soc. Sci. Res.* 26:331–76

Heckman J, Robb R. 1985. Using longitudinal data to estimate age, period and cohort effects in earnings equations. In *Cohort Analysis in Social Research*, ed. WM Mason, S Fienberg, pp. 137–50. New York: Springer

Heuer C. 1997. Modeling of time trends and interactions in vital rates using restricted regression splines. *Biometrics* 53:161

Hobcraft J, Menken J, Preston S. 1982. Age, period, and cohort in demography: a review. *Popul. Index* 48:4–43

Holford TR. 2005. Age-period-cohort analysis. In *Encyclopedia of Biostatistics*, ed. P Armitage, T Colton. Hoboken, NJ: Wiley

Hout M, Knoke D. 1975. Change in voting turnout, 1952–1972. *Public Opin. Q.* 39:52–68

Inglehart R. 1971. The silent revolution in Europe: intergenerational change in post-industrial societies. *Am. Political Sci. Rev.* 65:991–1017

Kahn JR, Mason WM. 1987. Political alienation, cohort size, and the Easterlin hypothesis. *Am. Sociol. Rev.* 52(2): 155–69

Keiding N. 2011. Age-period-cohort analysis in the 1870s: diagrams, stereograms, and the basic differential equation. *Can. J. Stat.* 39:405–20

Kermack WO, McKendrick AG, McKinlay PL. 1934. Death-rates in Great Britain and Sweden: expression of specific mortality rates as products of two factors, and some consequences thereof. *J. Hygiene* 34:433–57

Kerr WC, Greenfield TK, Bond J, Ye Y, Rehm J. 2004. Age, period and cohort influences on beer, wine and spirits consumption trends in the U.S. National Alcohol Surveys. *Addiction* 99:1111–20
Land KC, Felson M. 1978. Sensitivity analysis of arbitrarily identified simultaneous-equation models. *Sociol. Methods Res.* 6:283–307
Lavori PW, Klerman GL, Keller MB, Reich T, Rice J, Endicott J. 1987. Age-period-cohort analysis of secular trends in onset of major depression: findings in siblings of patients with major affective disorder. *J. Psychiatr. Res.* 21:23–35
Lee HA, Park H. 2012. Trends in ischemic heart disease mortality in Korea, 1985–2009: an age-period-cohort analysis. *J. Prev. Med. Public Health* 45:323–28
Lee S, Hershberger S. 1990. A simple rule for generating equivalent models in covariance structure modeling. *Multivar. Behav. Res.* 25.3:313–34
Liu S, Semenciw R, Ugnat AM, Mao Y. 2001. Increasing thyroid cancer incidence in Canada, 1970–1996: time trends and age-period-cohort effects. *Br. J. Cancer* 85:1335
Luo L, Hodges JS. 2016. Block constraints in age-period-cohort models with unequal-width intervals. *Sociol. Methods Res.* 45.4:700–26
Luo L, Hodges JS. 2019. Constraints in random effects age-period-cohort models. arXiv:1904.07672 [stat.ME]
Luo L, Hodges J, Winship C, Powers D. 2016. The sensitivity of the intrinsic estimator to coding schemes: comment on Yang, Schulhofer-Wohl, Fu, and Land. *Am. J. Sociol.* 122:930–61
Mannheim K. 1952 (1927). The problem of generations. In *Essays on the Sociology of Knowledge*, ed. P Kecskemeti, pp. 276–322. London: Routledge & Kegan Paul
Manski CF. 1990. Nonparametric bounds on treatment effects. *Am. Econ. Rev.* 80:319–23
Manski CF. 2003. Identification problems in the social sciences and everyday life. *South. Econ. J.* 70:11–21
Mason KO, Mason WM, Winsborough HH, Poole WK. 1973. Some methodological issues in cohort analysis of archival data. *Am. Sociol. Rev.* 38:242–58
Mason WM, Fienberg SE, eds. 1985. *Cohort Analysis in Social Research*. New York: Springer
Mason WM, Smith HL. 1985. Age-period-cohort analysis and the study of deaths from pulmonary tuberculosis. In *Cohort Analysis in Social Research*, ed. WM Mason, HL Smith, pp. 151–227. New York: Springer
McCall PL, Land KC. 2004. Trends in environmental lead exposure and troubled youth, 1960–1995: an age-period-cohort-characteristic analysis. *Soc. Sci. Res.* 33(2):339–59
O'Brien RM. 2011. Constrained estimators and age-period-cohort models. *Sociol. Methods Res.* 40:419–52
O'Brien RM. 2014. *Age-Period-Cohort Models: Approaches and Analyses with Aggregate Data*. Boca Raton, FL: CRC
O'Brien RM. 2017. Mixed models, linear dependency, and identification in age-period-cohort models. *Stat. Med.* 36:2590–600
O'Malley PM, Bachman JG, Johnston LD. 1984. Period, age, and cohort effects on substance use among American youth, 1976–82. *Am. J. Public Health* 74:682–88
Oaxaca R. 1973. Male-Female Wage Differentials in Urban Labor Markets. *Int. Econ. Rev.* 14:693–709
Pearl J. 2009. *Causality*. Cambridge, UK: Cambridge Univ. Press
Pilcher J. 1994. Mannheim's sociology of generations: an undervalued legacy. *Br. J. Sociol.* 45:481–95
Powers DA. 2013. Black–white differences in maternal age, maternal birth cohort, and period effects on infant mortality in the U.S. (1983–2002). *Soc. Sci. Res.* 42:1033–45
Preston SH, Wang H. 2006. Sex mortality differences in the United States: the role of cohort smoking patterns. *Demography* 43(4):631–46
Putnam RD. 1995. Tuning in, tuning out: the strange disappearance of social capital in America. *PS* 28:664–83
Putnam RD. 2001. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster
Reither EN, Hauser RM, Yang Y. 2009. Do birth cohorts matter? Age-period-cohort analyses of the obesity epidemic in the United States. *Soc. Sci. Med.* 69:1439–48
Robinson RV, Jackson EF. 2001. Is trust in others declining in America? An age–period–cohort analysis. *Soc. Sci. Res.* 30:117–45
Rodgers WL. 1982a. Estimable functions of age, period, and cohort effects. *Am. Sociol. Rev.* 47:774–87
Rodgers WL. 1982b. Reply to comment by Smith, Mason, and Fienberg. *Am. Sociol. Rev.* 47:793–96

Rodgers WL. 1990. Interpreting the components of time trends. *Sociol. Methodol.* 20:421–38
Ryder NB. 1965. The cohort as a concept in the study of social change. *Am. Sociol. Rev.* 30:843–61
Schooler C. 1998. Environmental complexity and the Flynn effect. In *The Rising Curve: Long-Term Gains in IQ and Related Measures*, ed. U Neisser, pp. 67–79. Washington, DC: Am. Psychol. Assoc.
Schwadel P, Stout M. 2012. Age, period and cohort effects on social capital. *Soc. Forces* 91:233–52
Sen M, Wasow O. 2016. Race as a bundle of sticks: designs that estimate effects of seemingly immutable characteristics. *Annu. Rev. Political Sci.* 19:499–522
Simirenko A. 1966. Mannheim's generational analysis and acculturation. *Br. J. Sociol.* 17:292–99
Smith HL, Mason WM, Fienberg SE. 1982. Estimable functions of age, period, and cohort effects: more chimeras of the age-period-cohort accounting framework: comment on Rodgers. *Am. Sociol. Rev.* 47:787–93
Tilley J, Evans G. 2014. Ageing and generational effects on vote choice: combining cross-sectional and panel data to estimate APC effects. *Electoral Stud.* 33:19–27
Vaupel JW, Romo VC. 2003. Decomposing change in life expectancy: a bouquet of formulas in honor of Nathan Keyfitzs 90th birthday. *Demography* 40(2):201–16
Vedøy TF. 2014. Tracing the cigarette epidemic: an age-period-cohort study of education, gender and smoking using a pseudo-panel approach. *Soc. Sci. Res.* 48:35–47
Voas D, Chaves M. 2016. Is the United States a counterexample to the secularization thesis? *Am. J. Sociol.* 121:1517–56
Wilson JA, Gove WR. 1999. The age-period-cohort conundrum and verbal ability: empirical relationships and their interpretation: reply to Glenn and to Alwin and McCammon. *Am. Sociol. Rev.* 64:287–302
Winship C, Harding DJ. 2008. A mechanism-based approach to the identification of age–period–cohort models. *Sociol. Methods Res.* 36:362–401
Yang Y. 2008. Social inequalities in happiness in the United States, 1972 to 2004: an age-period-cohort analysis. *Am. Sociol. Rev.* 73:204–26
Yang Y, Fu WJ, Land KC. 2004. A methodological comparison of age-period-cohort models: the intrinsic estimator and conventional generalized linear models. *Sociol. Methodol.* 34:75–110
Yang Y, Land KC. 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. *Sociol. Methodol.* 36:75–97
Yang Y, Land KC. 2013. *Age-Period-Cohort Analysis: New Models, Methods, and Empirical Applications*. Boca Raton, FL: CRC
Yang Y, Schulhofer-Wohl S, Fu WJ, Land KC. 2008. The intrinsic estimator for age-period-cohort analysis: what it is and how to use it. *Am. J. Sociol.* 113:1697–736

Annual Review
of Sociology

Volume 45, 2019

Contents