Standard Length Article

Sociological Methodology
1–52

DOI: 10.1177/00811750231151949
http://sm.sagepub.com

The Anatomy of Cohort Analysis: Decomposing Comparative Cohort Careers

Ethan Fosse[1] and Christopher Winship[2]

Abstract

In a widely influential essay, Ryder argued that to understand social change, researchers should compare cohort careers, contrasting how different cohorts change over the life cycle with respect to some outcome. Ryder, however, provided few technical details on how to actually conduct a cohort analysis. In this article, the authors develop a framework for analyzing temporally structured data grounded in the construction, comparison, and decomposition of cohort careers. The authors begin by illustrating how one can analyze age-period-cohort (APC) data by constructing graphs of cohort careers. Although a useful starting point, the major problem with this approach is that the graphs are typically of sufficient complexity that it can be difficult, if not impossible, to discern the underlying trends and patterns in the data. To provide a more useful foundation for cohort analysis, the authors therefore introduce three distinct improvements over the purely graphical approach. First, they provide a mathematical definition of a cohort career, demonstrating how the underlying parameters of interest can be estimated using a reparameterized version of the conventional APC model. The authors call this the life cycle and social change (LC-SC) model. Second, they contrast the proposed model with two alternative three-factor APC models and all logically possible two-factor models, showing that none of these other models are adequate for fully representing Ryder's ideas. Third, the authors present the article's major accomplishment: using the LC-SC model, they show how a collection of cohort careers can be decomposed into just four basic components: a curve representing an overall intracohort trend (or life cycle change); a curve representing an overall intercohort trend (or social change); a set of common cross-period temporal fluctuations that permit variability across cohort careers; and, finally, a set of terms representing cell-specific heterogeneity (or, equivalently, interactions among age, period, and/or cohort). As the authors demonstrate, these parts can be reassembled into simpler versions of cohort careers, revealing underlying trends and patterns that may not be evident otherwise. The authors illustrate this approach by analyzing trends in political party strength in the General Social Survey.

Keywords

cohort analysis, comparative cohort careers, Norman Ryder, social change, life cycle change

Most of the methodological literature on analyzing age-period-cohort (APC) data has focused on developing techniques to identify the unique effects for age, period, and

[1]University of Toronto, Toronto, ON, USA
[2]Harvard University, Cambridge, MA, USA

Corresponding Author:
Ethan Fosse, University of Toronto, Department of Sociology, 725 Spadina Avenue, Room 348, Toronto, ON M5S 2J4, Canada
Email: ethan.fosse@utoronto.ca

cohort (for a review, see Fosse and Winship 2019a). However, the demographer and sociologist Norman B. Ryder, whose 1965 *American Sociological Review* article "The Cohort as a Concept in the Study of Social Change" is the most cited and arguably most influential work on the topic, advocated for a different approach.[1] Rather than attempting to disentangle period effects from effects attributable to age or cohort, Ryder (1965) believed that when conducting cohort analysis, one should describe and contrast how cohorts age through time (i.e., across periods), constructing what he termed "comparative cohort careers" (p. 549).[2] Importantly, Ryder's (1965, 1968, 1992) approach is based on descriptively distinguishing intracohort trends (or life cycle change) from intercohort trends (or social change), which, together, constitute a collection of comparative cohort careers. Unfortunately, however, as noted by a number of scholars (e.g., Glenn 1977; Hardy and Wilson 2002; Smith 2021), Ryder never developed a formal methodology for carrying out his approach.

Following Ryder's insights, our goal in this article is to develop an explicit framework for carrying out a descriptive analysis of temporally structured data grounded in the construction, comparison, and decomposition of cohort careers. It is crucial to understand that our approach, as a kind of cohort analysis, is distinct from APC analysis as it is conventionally understood. Whereas in an APC analysis the goal is nearly always to identify distinct "effects" for age, period, and cohort (see, e.g., Fosse and Winship 2019a), our goal is to describe and deconstruct intra- and intercohort trends.[3] As we will discuss, an intracohort trend (or life cycle change) can be estimated using a combination of age and period parameters, whereas an intercohort trend (or social change) can be estimated using a combination of cohort and period parameters. Our approach is purely descriptive in the sense that no additional information external to the data is required for identification and, accordingly, estimation. This is in contrast to a traditional APC analysis in which, for example, researchers use constraints, ideally informed by theoretical considerations, to extract unique effects for age, period, and cohort (e.g., Fosse and Winship 2019a; Mason and Fienberg 1985).

In demonstrating how to conduct an analysis of cohort careers, however, we do *not* suggest that our approach is always preferred over a conventional APC analysis (e.g., Fosse and Winship 2019a, 2019b). In general, we view our approach as being deployed in two different ways. First, there may be cases in which a purely descriptive analysis of cohort trajectories and their underlying elements is all that is needed to test a particular social theory. Second, inasmuch it does not involve any identification assumptions, our approach can be used as a first step in a traditional APC analysis that does, in fact, attempt to extract distinct effects for age, period, and cohort. The principle here is that one should start by analyzing the data making the fewest assumptions possible; then, to gain further insights, one can progressively incorporate stronger assumptions. Of course, the utility of subsequent results will depend critically on the validity of the assumptions invoked.

Regardless of how our approach is used, a key point that emerges from our analysis is that it is absolutely critical for researchers to be clear about whether they are conducting a cohort analysis or, alternatively, a more traditional APC analysis focused on identifying effects for age, period, and cohort. We believe there has been considerable ambiguity and confusion on this point in the literature. To distinguish between these two distinct approaches, we will refer to an *APC analysis* as any analysis that attempts, using information external to the data, to extract unique effects for age, period, and cohort (e.g., Fosse and Winship 2019a; Mason et al. 1973; Mason and Fienberg 1985), whereas *cohort analysis* will refer to any analysis that attempts to describe intra- and intercohort trends, or life cycle and social change (e.g., Ryder 1965, 1968).

We begin the exposition of our approach by demonstrating how one can construct cohort careers using an APC data set by stratifying on cohort and arranging a set of cell-specific summaries (e.g., means) across age groups. These summaries can then be compared within cohorts to reveal intracohort trends and across cohorts to uncover intercohort trends. This approach has the benefit of simplicity, but it is limited as it can be difficult, if not impossible, to visually disentangle linear trends, fluctuations, and heterogeneity because of interactions among age, period, and/or cohort. As an alternative, we present a mathematical framework for defining and constructing cohort careers that allows one to decompose a set of cohort careers into discrete components, not only algebraically but also graphically in terms of a set of simpler figures.

We develop our approach using the classical three-factor APC model, which is by far the most widely used way to represent time-series data organized by age, period, and cohort (Mason and Fienberg 1985). We extend the classical model by allowing for cell-specific heterogeneity or, equivalently, interactions among age, period, and/or cohort. Following previous work (e.g., Fosse and Winship 2018, 2019b; Holford 1983; Smith 2021), we also simplify this model by separating the linear from the nonlinear components for age, period, and cohort.[4] Because of these differences from the classical APC (C-APC) model, we call this the linearized APC heterogeneity (L-APCH) model. As we will show, the parameters of the L-APCH model can be grouped so that one set represents intracohort trends, or Ryder's concept of life cycle change, and another (partially overlapping) set represents intercohort trends, or Ryder's concept of social change.[5] We use this grouping as the basis of what we call the life cycle and social change (LC-SC) model.

Importantly, the LC-SC model enables one to decompose a collection of cohort careers into just four components: a life cycle (LC) curve indexed by age representing an overall intracohort trend; a social change (SC) curve indexed by birth year representing an overall intercohort trend; a set of common cross-period temporal fluctuations that permit variability across cohort careers; and, last, a set of terms representing cell-specific heterogeneity or, equivalently, interactions among age, period, and/or cohort. As we demonstrate, these various pieces can be reassembled into simpler versions of cohort careers, as well as intracohort and intercohort trends, thereby revealing underlying patterns in the data that may not be evident otherwise.

The LC-SC model is a powerful way to estimate the underlying parameters of a set of cohort careers, but a number of alternative purely descriptive models could be used. Accordingly, we evaluate two related three-factor models and all three logically possible two-factor models (age-cohort, age-period, and period-cohort). As we demonstrate, in contrast to the LC-SC model, none of these alternative models can adequately represent Ryder's concepts of intra- and intercohort trends. Furthermore, by revealing the relationship between the parameters of the LC-SC model and those of various other models, our analysis clarifies what is actually estimated by applied researchers in a wide range of studies. This, in turn, implies a set of general principles that enables one to evaluate the extent to which a given model accurately represents Ryder's overall approach. We discuss these general principles in the conclusion.

The remainder of the article is organized into five main parts. First, we discuss Ryder's thinking about cohorts in greater detail, showing how comparative cohort careers can be constructed and graphed using the means (or other summaries) of an APC data set. In doing so, we introduce our empirical example on political party strength using the General Social Survey (GSS). Second, to set the foundations for the sections that follow, we introduce models for APC analysis, showing how the conventional APC model can be reparameterized to allow for separate linear and nonlinear terms, and extended to allow for cell-level heterogeneity representing interactions among age, period, and/or cohort (Fosse and Winship 2019b; Holford 1983; Smith 2021). This results in what we call the L-APCH model. To clarify our later contributions, we briefly discuss models for APC analysis in which identification is achieved by assuming that one of the three linear (i.e., slope) parameters is zero. Third, we present models for cohort analysis, showing how intra- and intercohort trends and, accordingly, comparative cohort careers, can be represented by regrouping the parameters of the L-APCH model. We then introduce the LC-SC model, contrasting it to a closely related, yet conceptually distinct, APC model in which the period slope is assumed to equal zero. Fourth, we consider alternative models for cohort analysis, comparing the LC-SC model with two alternative three-factor models and all three logically possible two-factor models (age-cohort, age-period, and period-cohort). As we show, only the two-factor age-cohort model succeeds, albeit partially, in estimating the key parameters of interest. Fifth, and arguably most important, we demonstrate how the LC-SC model can be used to decompose cohort careers into simpler components, revealing trends, patterns, and fluctuations that are otherwise obscured. Namely, with respect to political party strength, despite considerable heterogeneity, we find evidence of a steep downward intercohort trend along with a strong upward intracohort trend. Neither of these trends are easily discerned using a purely graphical approach. We conclude with a discussion of limitations, outlining suggestions for further research as well as general guidelines for conducting a cohort analysis using other models.

CONSTRUCTING AND COMPARING COHORT CAREERS

For nearly a century, generations of social scientists have sought to understand the nature and extent of social change using APC data.[6] As noted earlier, easily the most influential work on cohort analysis is Ryder's 1965 article "The Cohort as a Concept in the Study of Social Change" (see also Ryder 1968, 1992). Ryder's crucial insight was that because of the succession of cohorts through time, the composition of a population, and thus a society, could change without any individual (or collection of individuals) changing over the life-course. He reasoned that over time, older cohorts are gradually but ineluctably replaced by new cohorts that typically have different distributional characteristics than earlier ones, sometimes dramatically so.[7] Cohorts might differ, for example, in terms of their fertility and mortality rates, health statuses, education levels, religious beliefs, or attitudes toward out-groups, to name just a few characteristics. Because of these distributional differences, the succession of cohorts through time can alter the composition of the population in absence of any individual or set of individuals changing over the life-course.

The Cohort Approach

The core aspect of cohort analysis, or what Ryder (1965:549) termed "the cohort approach," is the construction and comparison of cohort careers, or cohort-specific age-time trajectories. Specifically, building off earlier work by historians and demographers, Ryder (1965:861) argued that cohort analysis entails the study of intracohort trends, or "intracohort temporal development," as well as the examination of intercohort trends, or "intercohort temporal differentiation." As mentioned earlier, because they involve comparisons within cohorts as they age through time, intracohort trends represent life cycle change; by contrast, because they entail comparisons of successive cohorts through time, intercohort trends represent social change. We can summarize Ryder's ideas using the following general-purpose, heuristic equation:

$$\begin{matrix}\text{Comparative} \\ \text{Cohort Career}\end{matrix} := \begin{matrix}\text{Intracohort Trend} \\ \text{(Life Cycle Change)}\end{matrix} \;\&\; \begin{matrix}\text{Intercohort Trend} \\ \text{(Social Change)}\end{matrix} \qquad (1)$$

where $:=$ means "is defined by."Equation (1) simply states that a comparative cohort career is defined as some combination of terms for an intracohort trend, representing life cycle change, and those for an intercohort trend, representing social change. Note that each cohort has a "career" inasmuch as there is change within the cohort (i.e., life cycle change), which is "comparative" in the sense there is some degree of change relative to other cohorts (i.e., social change).

In practice, the main question that arises concerns how to use APC data to quantitatively describe the cohort careers outlined heuristically in equation (1). As Ryder himself noted, his 1965 essay presents few technical details on how to actually analyze APC data in keeping with his framework. However, in a rather infrequently cited article,[8] Ryder (1968) provided a brief mathematical exposition of his approach that offers some additional insight (see also Ryder 1979).[9] Let Y denote an outcome, A age, P period, and C cohort. Ryder (1968) argued that "cohort analysis is concerned with the characteristics of functions $Y = f(A, C)$ for variations in A (intracohort) and C (intercohort)," where "event Y occurs at time $P = A + C$" (p. 546). This implies a simple, straightforward procedure for constructing comparative cohort careers: simply stratify on cohort (C) and arrange a set of cell-specific summaries of the outcome (e.g., means of Y) across levels of age (A). Conditional on C, variation in the summaries of Y across levels of A (and thus also P) reveals an intracohort trend, or life cycle change. Conversely, conditional on A, variations in the summaries of Y across levels of C (and thus also P) reveals an intercohort trend, or social change.

Several points are worth noting about Ryder's general framework. First, his cohort approach is based on distinguishing intra- from intercohort trends, not on disentangling the effects of age, period, and cohort from each other. In fact, Ryder (1992) viewed his approach as at odds with that of Mason and colleagues (e.g., Mason and Fienberg 1985) and generally disavowed attempts to identify unique APC effects. As Ryder (1979) stated in no uncertain terms, "My impression of the problem of determining the separate contributions of cohort, age and period to some surface of data is that it is a question without an answer, and one that we should stop wasting time on" (p. 1).

Second, Ryder's approach entails diachronic (i.e., over-time), not synchronic (i.e., cross-sectional), comparisons. From Ryder's (1979:547) perspective, the key question concerns whether one should compare age summaries using "cohort sections" or "period sections."[10] According to Ryder (1979:547), a cohort section consists of summaries of Y across levels of A (and thus P) within a particular cohort C, whereas a period section consists of summaries of Y across levels of A (and thus C) within a particular period P. In contrast to an analysis within a cohort section, which is dynamic, a comparison of variations in an outcome within a period section is based on a set of static intraperiod differences (Ryder 1964, 1965). This, in Ryder's view, is typically not of interest to social scientists, notwithstanding the relative ease by which one can collect data from a period rather than a cohort section.[11] Thus, although an analysis using period sections is "conventional practice," the study of cohort sections is "the essence of the proposal of cohort analysis" (Ryder 1968:547).

Third, although Ryder's (1968) notation would seem to suggest age and cohort are the only truly relevant dimensions, Ryder (1972) viewed his cohort approach as incorporating information from all three temporal variables: "I have never proposed, implicitly or explicitly, that two variables be picked and the third ignored, as a practical procedure" (p. 1). This would appear to rule out techniques that focus on just two of the three variables, a topic we discuss below. Moreover, Ryder (1979) was deeply aware that P is determined given A and C, and, as such, one dimension cannot be ignored simply because it is not directly specified in some function: "The combination of any cohort identification, and a particular age for that cohort, determines the period of occurrence" (p. 5). Thus, it is necessarily the case, for example, that "members of a cohort age *pari passu* with time [i.e., period]" (Ryder 1979:547; see also Ryder 1965, 1968).

Finally, Ryder's justification for a cohort-centric approach is based, in part, on the claim that differences across cohorts are based primarily on "metabolic" rather than "mutative" processes, reflecting changes *of* individuals rather than changes *in* individuals. Specifically, Ryder viewed intracohort trends as arising mainly from changes in individuals, with intercohort trends largely due to the replacement of old individuals by new individuals with different characteristics, or to differential fertility and mortality. Even within cohorts, Ryder argued that individual-level changes tend to be most dramatic in younger ages, with changes stabilizing after young adulthood (see also Mannheim [1927/1928] 1952). This further justifies a cohort-based perspective, as events nearer the time of a cohort's birth are thought to be more formative than those occurring later in a cohort's life cycle.[12]

Empirical Example

To develop the ideas outlined above, we present a running example analyzing trends in political party strength in the United States using the GSS.[13] The specific question asked of respondents is as follows: "Generally speaking, do you usually think of yourself as a Republican, Democrat, Independent, or what?" The response consists of seven possible categories ranging from "strong Democrat" to "strong Republican," with "Independent" as the middle category.[14] Following Converse (1976), who distinguished between party strength and party affiliation, we transform these responses into a measure of strength of party identification using the following coding: 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Consistent with Converse's analysis, and for ease of exposition, we treat party strength as a continuous variable.[15] Following the convention in the literature, we group age and period into equal (five-year) intervals, with their difference used to create the cohort groups. This leads to $I = 15$ age groups ranging from 18–22 to 88 years and older, $J = 10$ period groups ranging from 1970–1974 to 2015–2019, and $K = 24$ cohort groups ranging from 1882–1886 to 1997–2001. The cell means of the resultant APC data set range from a high of 2.57 (relatively strong partisanship) to a low of 1.22 (relatively weak partisanship).

To construct a set of cohort careers, we stratify on cohort and arrange the summaries (i.e., means) for party strength across the levels of age. This is depicted in Figure 1 as a collection of cohort-specific lines, with age as the horizontal axis.[16] Following Ryder's (1965, 1968) framework, vertical differences across cohorts represent intercohort trends (or social change), and horizontal differences within cohorts represent intracohort trends (or life cycle change). Although masked by nonlinearities and additional heterogeneity, Figure 1 shows some kind of underlying patterning. Specifically, there seems to be a general increase in party strength within cohorts, but the magnitude of this trend is unclear because of the substantial variability of the careers. Likewise, there appears to be some vertical separation among the careers, suggesting a possible decline in party strength across cohorts.

There are two main limitations in the approach to constructing and comparing cohort careers outlined above (and displayed in Figure 1). First, the results are easily misinterpreted as purely attributable to age and cohort (rather than period), simply because the data happen to be organized by age and cohort. As noted earlier, Ryder viewed his approach as incorporating information from all three temporal variables and did not think one dimension should be ignored. Second, and more important, the intra- and intercohort trends are complicated by the age, period, and cohort nonlinearities as well as additional heterogeneity arising from interactions among age, period, and/or cohort. As shown in Figure 1, there is so much variability that it is difficult, for example, to discern the underlying trends and patterns in the data from short-term fluctuations. Both of these limitations suggest that the simple, informal approach to constructing and comparing cohort careers can be improved by using a parametric model, a topic to which we turn next.

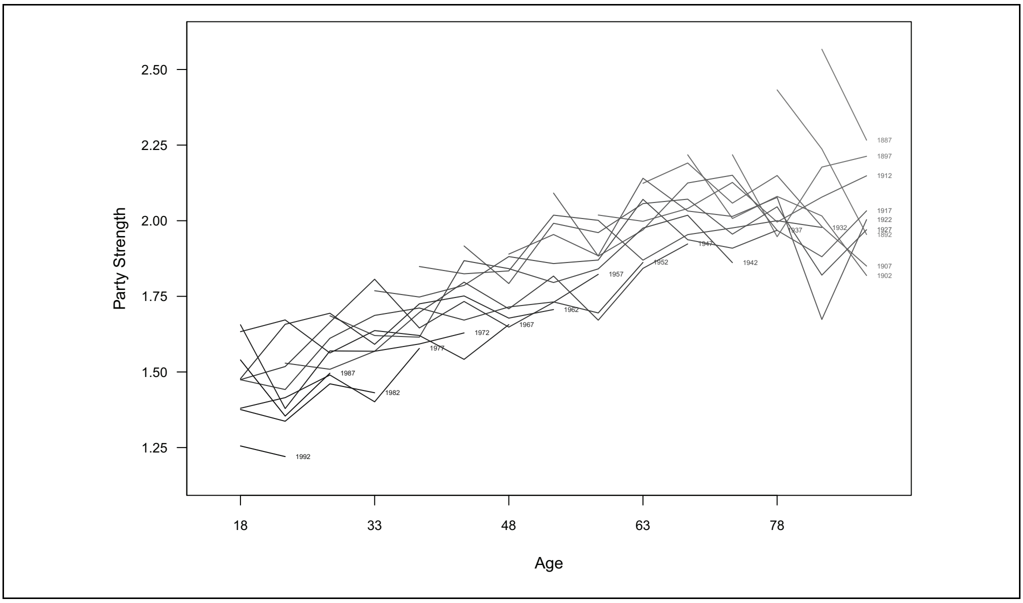

Figure 1. Comparative cohort careers.
Note: This graph shows the mean party strength for each cohort arranged across age levels. The two most extreme cohorts, which are each based on only one age group, are not displayed. Vertical differences across cohorts correspond to intercohort trends; horizontal differences within cohorts correspond to intracohort trends. Outcome is strength of party identification, treated as a continuous variable, with responses coded as 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights. Sample size is $R = 63,128$ respondents.

MODELS FOR APC ANALYSIS

Before presenting how one can use APC models to formally construct, compare, and decompose cohort careers, we first outline how such models are generally used in the literature. We begin this section on APC analysis by outlining the classical model with age, period, and cohort as inputs (see, e.g., Mason et al. 1973). We next expand this model to include additional terms distinguishing between cell- and individual-level heterogeneity, and then discuss a linearized version of this model that simplifies the identification problem (see Smith 2021). Finally, we consider three possible versions of the linearized model in which identification is achieved by assuming that one of the APC linear effects is zero.

The C-APC Model

Suppose we have categorically coded age, period, and cohort data with a continuous outcome Y and cohort calculated from age and period.[17] Let $i = 1, \ldots, I$ denote the age groups; $j = 1, \ldots, J$ the period groups; and $k = 1, \ldots, K$ the cohort groups, with $k = j - i + I$ and $K = I + J - 1$.[18] Let $r = 1, \ldots, R$ index the respondents in a time-series cross-sectional data set organized by age, period, and cohort, where R is the sample size. The *classical age-period-cohort (C-APC) model* is as follows (Fosse and Winship 2019a)[19]:

$$Y_{rijk} = \mu + \alpha_i + \pi_j + \gamma_k + \epsilon_{rijk}, \tag{2}$$

where μ is the overall mean; α_i, π_j, and γ_k are age, period, and cohort parameters; and ϵ_{rijk} is an error term centered on zero. To identify the levels of the parameters given the inclusion of the intercept, sum-to-zero constraints are applied to the parameters: $\sum_{i=1}^{I} \alpha_i = \sum_{j=1}^{J} \pi_j = \sum_{k=1}^{K} \gamma_k = 0$. However, because of the linear dependency among the time scales and the fact that the age, period, and cohort parameters combine slopes with deviations, in general none of the parameters in the C-APC model are identified, other than the overall mean (see, e.g., Fosse and Winship 2018; Mason et al. 1973; Mason and Fienberg 1985).[20]

It will prove useful in our discussion to extend the C-APC model to include terms representing cell- and individual-level heterogeneity. Typically, analysts have not decomposed the error terms (i.e., ϵ_{ijk}'s) into cell- and individual-level components. However, on an age-period array, there are $(I - 2) \times (J - 2)$ additional degrees of freedom beyond those taken up by the conventional APC model.[21] To incorporate this additional cell-specific variability, we introduce what we call the *C-APC heterogeneity* (C-APCH) model:

$$Y_{rijk} = \mu + \alpha_i + \pi_j + \gamma_k + \underbrace{\eta_{ijk} + \xi_{rijk}}_{\epsilon_{rijk}}. \tag{3}$$

The η_{ijk} terms denote cell-specific heterogeneity, or interactions among age, period, and/or cohort, and the ξ_{rijk} terms represent individual-level (within-cell) heterogeneity.

The Linearized APC Model

As noted earlier, the separate parameters for age, period, and cohort in the C-APC are generally all unidentified. This identification problem can be greatly simplified by considering a linearized form of the model where, in general, only the linear parameters for age, period, and cohort are not identified. Following previous work (Fosse and Winship 2018, 2019b; Holford 1983; Smith 2021), we use a representation of the C-APCH model that specifies separate linear and nonlinear terms, resulting in the *linearized APC heterogeneity (L-APCH) model*:

$$Y_{rijk} = \mu + \alpha(i - i^*) + \pi(j - j^*) + \gamma(k - k^*) + \tilde{\alpha}_i + \tilde{\pi}_j + \tilde{\gamma}_k + \underbrace{\eta_{ijk} + \xi_{rijk}}_{\epsilon_{rijk}}, \tag{4}$$

where the asterisks denote midpoint or referent indices $i^* = (I+1)/2$, $j^* = (J+1)/2$, and $k^* = (K+1)/2$; α, π, and γ are the slopes of age, period, and cohort; and $\tilde{\alpha}_i$, $\tilde{\pi}_j$, and $\tilde{\gamma}_k$ represent the i th age, j th period, and k th cohort nonlinearities, respectively.[22]

APC Models with Zero-Slope Constraints

It has long been understood that to identify unique effects for age, period, and cohort, the conventional APC model can be just-identified by applying a constraint on the parameters. One approach to identification is to assume that the period linear effect is equal to zero (see, e.g., Bell 2014; Bell and Jones 2015, 2021; Holford 1983:316; O'Brien 2015:50–54; Smith 2021). We refer to any model that invokes this assumption as the zero period slope (ZPS) model, which can be specified generically as follows (cf. equation 11 in Smith 2021):

$$Y_{rijk} = \mu + \alpha(i - i^*) + \gamma(k - k^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k + \eta_{ijk} + \xi_{rijk}, \tag{5}$$

where, by assumption, $\pi = 0$. This model assumes the "true" linear effects, α and γ, are recovered after constraining the period slope, π, to equal zero, as in equation (5). This is an assumption external to the data and cannot be directly verified by the data.

For the sake of completeness, and to set the discussion that follows, two alternative models are also worth noting. The first is the zero age slope (ZAS) model, which is based on the assumption that the age linear effect is equal to zero. This model can be specified generically as follows:

$$Y_{rijk} = \mu + \pi(j - j^*) + \gamma(k - k^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k + \eta_{ijk} + \xi_{rijk}, \tag{6}$$

where, by assumption, $\alpha = 0$. Second is the zero cohort slope (ZCS) model, which has the following form:

$$Y_{rijk} = \mu + \alpha(i - i^*) + \pi(j - j^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k + \eta_{ijk} + \xi_{rijk}, \tag{7}$$

where, again by assumption, $\gamma = 0$. As with the ZPS model, the ZAS and ZCS models require information external to the data for identification.

To understand the assumptions invoked by equations (5), (6), and (7), note that the APC identification problem can be defined in terms of a ''canonical solution line'' lying in three-dimensional space on which all possible slope estimates lie (Fosse and Winship 2019a, 2019b; cf. O'Brien 2011, 2015). More formally, let α^*, π^*, and γ^* denote any particular set of age, period, and cohort slopes that specify a point on the canonical solution line. The relationship between these slopes and the "true," unknown linear effects is given by $\alpha^* = \alpha + \nu$, $\pi^* = \pi - \nu$, and $\gamma^* = \gamma + \nu$, where ν is an unknown scalar. The ZPS model assumes $\pi^* = \pi = 0$, namely, that the point on the canonical solution line in which a period linear term is zero ($\pi^* = 0$) equals the "true," unknown period linear term (π). This is a single point on the canonical solution line, thereby determining the values of the age and cohort linear terms.[23] Similar reasoning applies to the ZAS and ZCS models.

MODELS FOR COHORT ANALYSIS

The foregoing is essentially an outline of the conventional approach in which the goal is to extract unique effects for age, period, and cohort by incorporating additional information external to the data, often in the form of theoretically informed constraints

(Fosse and Winship 2019b). However, as noted previously, from Ryder's perspective the goal is *not* to disentangle the separate effects of age, period, and cohorts, but rather to distinguish intra- from intercohort trends (or life cycle from social change). Identifying intracohort trends requires tracking age groups through time (i.e., across periods) conditional on cohort, whereas identifying intercohort trends requires tracking cohorts through time (i.e., across periods) conditional on age. As we will show, no identifying restrictions are needed for this approach.

In this section, we first show how the separate parameters for age, period, and cohort in the L-APCH model can, consistent with Ryder's (1968) formulation, be reindexed solely in terms of age and cohort. Next, we demonstrate how the resulting parameter sets can be estimated using what we call the LC-SC model, which is the basis for the decomposition analysis we present later. Finally, we conclude this section with a comparison of the LC-SC model with the ZPS model described previously. Although the two models share the same design matrix, crucially they differ in their estimands and, accordingly, their interpretation.

Parameterizing Comparative Cohort Careers

To parameterize comparative cohort careers using the L-APCH model, note that $j = i + k - I$ and $J = K - I + 1$.[24] Substituting for j and J in equation (4) allows us to specify an equation that is only a function of age (i and I) and cohort (k and K). This, as we reveal below, is the critical mathematical step in creating a model that is both identified and whose parameters represent life cycle and social change.[25] Replacing these values into the right-hand side of equation (4) and rearranging the terms, we obtain the following reexpression of the L-APCH model[26]:

$$\begin{aligned} Y_{rijk} &= \mu + \alpha\left(i - \frac{I+1}{2}\right) + \pi\left([i+k-I] - \frac{K-I+2}{2}\right) + \gamma\left(k - \frac{K+1}{2}\right) \\ &\qquad + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \widetilde{\gamma}_k + \eta_{i[i+k-I]k} + \xi_{ri[i+k-I]k} \\ &= \mu + (\alpha+\pi)(i - i^*) + (\gamma+\pi)(k - k^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \widetilde{\gamma}_k + \eta_{i[i+k-I]k} + \xi_{ri[i+k-I]k}, \end{aligned} \tag{8}$$

which is simply the L-APCH model with the parameters indexed by age (i and I) and cohort (k and K) but *not* period (j and J). After rearranging the terms again in equation (8), we can specify comparative cohort careers as follows:

$$\widehat{Y}_{ri[\,i-k+I]k} = \mu + \overbrace{(\alpha+\pi)(i-i^*) + \widetilde{\alpha}_i + \underbrace{\widetilde{\pi}_{i+k-I} + \eta_{ij[\,i+k-I]k}}_{}}^{\substack{\text{Intracohort Trend} \\ \text{(Life Cycle Change)}}} \underbrace{+ (\gamma+\pi)(k-k^*) + \widetilde{\gamma}_k}_{\substack{\text{Intercohort Trend} \\ \text{(Social Change)}}} \tag{9}$$

for $i = 1, \ldots, I$ in each cohort k,

which expresses the cell-specific means from an age-period table (cf. Figure 1) in terms of an age-indexed parameter set representing intracohort trends (i.e., life cycle change) and a cohort-indexed parameter set representing intercohort trends (i.e., social change).

Consistent with Ryder's formulation, the intracohort trends in equation (9) representing life cycle change involve all parameters indexed by the age groups (i), and the intercohort trends representing cohort social change involve all parameters indexed by the cohort groups (k). Specifically, an intracohort trend consists of parameters for age, α and $\widetilde{\alpha}_i$, parameters for period, π and $\widetilde{\pi}_{i+k-I}$, and heterogeneity (or interactive) terms $\eta_{i[i+k-I]k}$. This is because, within any given cohort, its members age in lockstep with calendar time, thereby reflecting both developmental (age-graded) and transient (period-specific) processes. Similarly, an intercohort trend consists of parameters for cohort, γ and $\widetilde{\gamma}_k$, parameters for period, π and $\widetilde{\pi}_{i+k-I}$, as well as heterogeneity (or interactive) terms $\eta_{i[i+k-I]k}$. This is because, within any given age stratum, successive cohorts not only enter the population at different times but are also observed at different times, thus reflecting both persistent (cohort-based) and transient (period-specific) processes.

Several additional points deserve emphasis regarding the specification of cohort careers in equation (9). First, because they reflect change over time (i.e., across periods) rather than differences within cross-sections (i.e., within periods), the two sets of parameters in equation (9) represent diachronic (i.e., over-time) *trends* rather than synchronic (i.e., cross-sectional) differences. From this perspective, the aliasing of the age and cohort slopes with the period slope is a natural, expected consequence of the fact that the analysis is diachronic, rather than synchronic. Second, in contrast to the age and cohort nonlinearities, the period nonlinearities and the heterogeneity terms in equation (9) contribute to both intra- and intercohort trends. Finally, and perhaps most important, a primary benefit of the mathematical definition above is the relative ease by which one can decompose cell means (given on the left-hand side of equation 9) into various distinct components (given on the right-hand side of equation 9). To emphasize, as we will show, these pieces can be assembled in numerous ways to construct simpler versions of the cohort careers depicted in Figure 1, thereby uncovering underlying trends and patterns in the data.

The LC-SC Model

We now describe the specific model used to estimate the key parameters of the comparative cohort careers defined in equation (9). The main parameters are estimated using the *LC-SC model*:

$$Y_{rijk} = \mu + \theta_1(i - i^*) + \theta_2(k - k^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{[i+k-I]} + \widetilde{\gamma}_k + \eta_{i[i+k-I]k} + \xi_{ri[i+k-I]k}, \tag{10}$$

where $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$. Note again that, as a result of the substitution of the sum of age and cohort indices for the period indices (i.e., $j = i + k - I$ and $J = K - I + 1$), the outcome is only a function of age, indexed by i, with corresponding parameters representing life cycle change, and cohort, indexed by k, with corresponding parameters representing social change. This model is identified (i.e., the design matrix is of full rank) as it does not contain a separate linear term for period, which is instead combined with the age and cohort linear terms.

The two linear parameters of interest in equation (10) are θ_1 and θ_2. The linear term θ_1 is the *life cycle change slope*, or the *LC slope*, which represents an overall (linear) intracohort trend (or life cycle change). Note that θ_1 is the sum of two terms from the L-APCH model: α, the linear age term, and π, the linear period term. To reiterate, this is because, as cohort members age, they both become older and enter a new time period. In contrast, θ_2 is the *social change slope*, or the *SC slope*, which represents an overall (linear) intercohort trend (or social change). Similar to the LC slope, θ_2 is the sum of two terms from the L-APCH model: γ, the linear cohort term, and π, the linear period term. Again, this is because we are comparing successive cohorts through time, so any two cohorts will differ in terms of when they entered the population and the time period in which they are observed.

Comparison with the ZPS Model

It is crucial to understand the conceptual differences between the LC-SC and ZPS models.[27] The LC-SC and ZPS models are mathematically isomorphic in that their design matrices are identical. Accordingly, the numerical estimates of the two models will equal each other. However, they are quite distinct in two critical respects.

First, the estimands of interest in the LC-SC model, the intra- and intercohort trends specified in equation (9), differ from the estimands of the ZPS model, namely, the separate effects of age, period, and cohort.[28] That is, with the ZPS model, the linear estimands of interest are the age linear effect, α, and the cohort linear effect, γ, under the assumption of no linear period effect ($\pi = 0$). Likewise, regarding the nonlinearities, with the ZPS model, the estimands of interest are the separate nonlinear effects for age, period, and cohort. By contrast, with the LC-SC model, the linear estimands of interest are the life cycle slope ($\theta_1 = \alpha + \pi$) and the SC slope ($\theta_2 = \gamma + \pi$). Regarding the nonlinearities, with the LC-SC model the focus is on how the age and period nonlinearities jointly contribute to intracohort trends (life cycle change) and how the cohort and period nonlinearities together constitute intercohort trends (social change). These differences between the two models underscore the importance of explicitly specifying the estimands of interest in a particular analysis (Lundberg, Johnson, and Stewart 2021).

Second, the two models are based on distinctly different derivations and assumptions. The ZPS model is derived by simply assuming that the period linear effect in the L-APCH model is zero, and its use has been advocated for in situations in which this assumption seems reasonable (see, e.g., Bell 2014; Bell and Jones 2015, 2021; Holford 1983:316; O'Brien 2015:50–54). The LC-SC model makes no such assumption. As shown in the discussion of equations (8) and (9), the LC-SC model is derived by reparameterizing the traditional APC model so that it is only a function of age and cohort. Rather than assumed to equal zero, in the LC-SC model the period slope π is combined with the age slope, α, as well as the slope parameter, γ, resulting in two new linear parameters, $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$.

The implications of the above differences are critical. Note that $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$ define the set of all possible age, period, and cohort linear terms that are consistent with a set of APC data. As mentioned previously, visually this results in the canonical solution line (Fosse and Winship 2019a; cf. O'Brien 2011, 2015).[29] The ZPS model assumes that $\pi^* = \pi = 0$; that is, it specifies the point on the canonical solution where the period linear term is zero ($\pi^* = 0$). As this is a single point on the canonical solution line, it determines the values of the age and cohort linear terms. Importantly, the ZPS model is only consistent with a single point on the solution line.[30] By contrast, the LC-SC model imposes no such constraint. The LC-SC model's two slopes, θ_1 and θ_2, are compatible with all points on the canonical solution line.[31]

ALTERNATIVE MODELS

The previous section outlined how, using the L-APCH model, one can reindex the parameters by age and cohort for the purpose of a cohort analysis. The resultant parameter sets, estimated using the LC-SC model, represent intra- and intercohort trends (or, equivalently, life cycle and social change). Importantly, the model is purely descriptive; no identifying assumptions are required, unlike the ZPS model. However, there are two other logically possible ways to reindex the parameters of the L-APCH model: besides age and cohort, one could group the parameters in terms of age and period or period and cohort.

In this section, we consider two alternative models for cohort analysis based on these two other ways of reindexing the L-APCH model. We also examine all three logically possible two-factor models (age-cohort, age-period, and cohort-period), comparing their estimates to those of the three different types of reindexed L-APCH models. Importantly, we show that, among the reindexed L-APCH models, only the estimates of the LC-SC model represent Ryder's ideas of life cycle and social change adequately. We also demonstrate that, among the two-factor models, the age-cohort model comes closest to representing Ryder's thinking, although its estimates will in general be biased. The two-factor age-period and cohort-period models, as well as their reindexed L-APCH counterparts, are only partially successful.

Alternative Indexing of Three-Factor Models: Age-Period and Period-Cohort

Instead of tracking some set of parameters across levels of age (conditional on cohort) and cohort (conditional on age), one could also track parameter sets across levels of age (conditional on period) and period (conditional on age), as well as across levels of period (conditional on cohort) and cohort (conditional on period). Borrowing Ryder's (1968) formulation, this means that, rather than specifying functions of the form $Y = f(A, C)$, one might specify functions of the form $Y = f(A, P)$ or $Y = f(P, C)$. However, crucially, these functions are inadequate for defining comparative cohort careers because, as Ryder recognized, they will encode, in part, intraperiod differences (or synchronic comparisons) rather than intra- or intercohort trends (or diachronic comparisons).

To illustrate this fact, consider first comparing parameter sets across ages within periods. Assuming again that we are using data collected on the basis of age and period (with cohort as the diagonal), it is by definition the case that $k = j - i + I$ and $K = I + J - 1$. Substituting these values into the right-hand side of equation (4) and using simple algebra, we obtain the following collection of intraperiod differences:

$$\widehat{Y}_{rij[j-i+I]} = \mu + \overbrace{((\alpha+\pi)-(\gamma+\pi))(i-i^*)+\widetilde{\alpha}_i+\underbrace{\widetilde{\gamma}_{j-i+I}+\eta_{ij[j-i+I]}}_{}}^{\text{Intraperiod Differences}}\underbrace{{}+(\gamma+\pi)(j-j^*)+\widetilde{\pi}_j}_{\text{Intercohort Trend (Social Change)}} \tag{11}$$

for $i = 1, \ldots, I$ in each period j,

which is essentially the L-APCH model with the parameters indexed by age (i and I) and period (j and J), but *not* cohort (k and K). Importantly, equation (11) expresses a collection of cell-specific means in terms of an age-indexed parameter set representing intraperiod differences and a period-indexed parameter set representing intercohort trends (i.e., social change).

In a similar fashion, one can compare parameter sets across cohorts within periods. Assuming again that we are using data collected on the basis of age and period (with cohort as the diagonal), it is the case that $i = j - k + I$ and $I = K - J + 1$. Substituting these values into the right-hand side of equation (4) and using simple algebra, we obtain the following collection of intraperiod differences:

$$\widehat{Y}_{r[j-k+(K-J+1)]jk} = \mu + \overbrace{((\gamma+\pi)-(\alpha+\pi))(k-k^*)+\widetilde{\gamma}_k+\underbrace{\widetilde{\alpha}_{j-k+(K-J+1)}+\eta_{[j-k+(K-J+1)]jk}}_{}}^{\text{Intraperiod Differences}}\underbrace{{}+(\alpha+\pi)(j-j^*)+\widetilde{\pi}_j}_{\text{Intracohort Trend (Life Cycle Change)}}$$

for $k = 1, \ldots, K$ in each period j, (12)

which is, similar to the previous equation, essentially the L-APCH model with the parameters indexed by cohort (k and K) and period (j and J), but *not* age (i and I). In contrast to equation (11), equation (12) expresses a collection of cell-specific means in terms of a cohort-indexed parameter set representing intraperiod differences, and a period-indexed parameter set representing intracohort trends (i.e., life cycle change).

The main limitation of using equations (11) and (12) or, more generally, functions of the form $Y = f(A, P)$ and $Y = f(P, C)$, for cohort analysis is that one parameter set will represent static, intraperiod differences rather than dynamic trends. Specifically, equation (11) is a combination of social change (indexed by period) and intraperiod differences (indexed by age), and equation (12) is a combination of life cycle change (indexed by period) and intraperiod differences (indexed by cohort). Intuitively, this is because, in each equation, one of the parameter sets is compared *within* rather than *across* periods, and thus necessarily entails a synchronic rather than diachronic analysis. We elaborate on this point in the following section, which outlines models for estimating the main parameters in equations (11) and (12).

Alternatives to the LC-SC Model: The Diff-SC and LC-Diff Models

Two other variants of the L-APCH model can be fit by expressing the parameters in terms of age and period or cohort and period, respectively. Expressing the parameters of the L-APCH model in terms of age and period results in the *intraperiod differences and social change (diff-SC) model*:

$$Y_{rijk} = \mu + (\theta_1 - \theta_2)(i - i^*) + \theta_2(j - j^*) + \tilde{\alpha}_i + \tilde{\pi}_j + \tilde{\gamma}_{j-i+I} + \eta_{ij[j-i+I]} + \xi_{rij[j-i+I]}, \tag{13}$$

where $\theta_1 - \theta_2 = (\alpha + \pi) - (\gamma + \pi) = \alpha - \gamma$ and $\theta_2 = \gamma + \pi$. The difference $\theta_1 - \theta_2$ in equation (13) is a slope of intraperiod differences, and θ_2 is simply the SC slope from the LC-SC model but indexed by period ($j = 1, \ldots, j = J$) instead of cohort ($k = 1, \ldots, k = K$). Similar to the LC-SC model, the diff-SC model is identified because it does not contain a unique linear term for cohort, which is instead absorbed into the period and cohort linear terms.

Likewise, expressing the parameters of the L-APCH model in terms of period and cohort results in the *life cycle and intraperiod differences (LC-diff) model*:

$$Y_{rijk} = \mu + \theta_1(j - j^*) + (\theta_2 - \theta_1)(k - k^*) + \tilde{\alpha}_{j-k+(K-J+1)} + \tilde{\pi}_j + \tilde{\gamma}_k + \eta_{[j-k+(K-J+1)]jk} + \xi_{r[j-k+(K-J+1)]jk}, \tag{14}$$

where $\theta_2 - \theta_1 = (\gamma + \pi) - (\alpha + \pi) = \gamma - \alpha$ and $\theta_1 = \alpha + \pi$. The difference $\theta_2 - \theta_1$ in equation (14) is a slope of intraperiod differences, and θ_1 is simply the LC slope from the LC-SC model but indexed by period ($j = 1, \ldots, j = J$) instead of age ($i = 1, \ldots, i = I$). Similar to the LC-SC model, the LC-diff model is identified because it does not include a separate linear term for age, which is instead incorporated into the period and cohort linear terms.

The LC-diff and diff-SC models provide the same estimates of the intercept and the age, period, and cohort nonlinearities as the LC-SC model. However, unlike the LC-SC model, the slopes indexed by age and cohort in equations (13) and (14), respectively, capture static differences within periods, not dynamic trends.[32] To reiterate the discussion in the previous section, this is because these slopes are estimated while conditioning on the period linear component and, as such, represent comparisons within cross-sections of time (i.e., within periods). In fact, it is only under very specific circumstances that the two synchronic models will provide unbiased estimates of the diachronic slopes. Specifically, the diff-SC model will give the correct estimate of θ_1 only if θ_2 happens to be zero, and the LC-diff model will give the correct estimate of θ_2 only if θ_1 happens to be zero.[33] Thus, for the purpose of directly estimating the parameters of equation (9), the LC-SC model is strongly preferred over the LC-diff and diff-SC models.

Two-Factor Models: Age-Cohort, Age-Period, and Period-Cohort

The LS-SC, diff-SC, and LC-diff models are APC models reindexed, respectively, in terms of age and cohort, age and period, and period and cohort. In other words, these models include parameters for age, period, and cohort, but are reindexed in terms two dimensions (e.g., age and cohort for the LS-SC model). Because these models include main parameters for age, period, and cohort, we refer to them as three-factor models.

As an alternative to three-factor APC models, researchers have also considered various two-factor models that drop the linear and nonlinear terms for one of the age, period, or cohort dimensions. For example, the age-period model, which is arguably the most widely used two-factor model, has been proposed in some version by Ohtaki, Kim, and Munaka (1990), Keyes and Li (2010), Tonda, Satoh, and Kamo (2015), and, most recently, Luo and Hodges (2020). Although less popular, a number of researchers have also adopted period-cohort models, including Firebaugh (1989),[34] Davis (1992:270–75), Voas and Chaves (2016:1551–53), and Hout (2021), in addition to age-cohort models (e.g., Yang and Land 2013:285–312). For the purpose of directly estimating the parameters of a set of cohort careers, two-factor models have several issues, which we discuss briefly in this section. Interested readers should also consult the technical discussion in Appendix A.

The age-cohort model, which, although not as widely used as the age-period or period-cohort models, has the closest relationship to the LC-SC model. In fact, the age-cohort model's design matrix is equivalent to that of the LC-SC model except the period nonlinear components have been dropped. Accordingly, like the LC-SC model, the age-cohort model will generate an LC slope for age and SC slope for cohort, thereby summarizing dynamic (i.e., over-time) linear intra- and intercohort trends, respectively. In this regard, the age-cohort model would appear relatively unproblematic for directly estimating the key parameters of a collection of cohort careers.

However, dropping the period nonlinear components results in two issues with the age-cohort model. First, although the variables for the period nonlinearities are orthogonal to the period linear component, they are correlated with the age and cohort variables. As a result, dropping the period nonlinear components gives rise to an omitted variable bias problem, producing biased estimates for the age and cohort parameters.[35] In addition, because the variables for the period nonlinearities are an inexact function of age and cohort, the period nonlinear components are only partially absorbed into age and cohort. As a consequence, part of the contribution of the period nonlinearities will be reflected in the estimated interactions between age and cohort, rather than in the main parameters of the model.[36]

Similar to the age-cohort model, the age-period and period-cohort models will also produce biased estimates of the slopes, as they leave out the nonlinear terms for cohort

Table 1. Comparative Summary of Three- and Two-Factor Models for Cohort Analysis

Function	Model	Slopes		
		Age	Period	Cohort
$f(A, C)$	LC-SC	θ_1	—	θ_2
	Age-Cohort	$\theta_1 + \xi_{\theta_1}$	—	$\theta_2 + \xi_{\theta_2}$
$f(A, P)$	Diff-SC	$\theta_1 - \theta_2$	θ_2	—
	Age-Period	$(\theta_1 - \theta_2) + \phi_{(\theta_1 - \theta_2)}$	$\theta_2 + \phi_{\theta_2}$	—
$f(P, C)$	LC-diff	—	θ_1	$\theta_2 - \theta_1$
	Period-Cohort	—	$\theta_1 + \psi_{\theta_1}$	$(\theta_2 - \theta_1) + \psi_{(\theta_2 - \theta_1)}$

Note: $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$. Third, fourth, and fifth columns indicate the parameters estimated, where the ξ, ϕ, and ψ terms indicate biases due to dropping the period, cohort, or age nonlinear components, respectively. For details, see Appendix A.

and age, respectively; as such, the contributions of the excluded nonlinear components will be reflected in the estimated interaction terms instead of the main parameters. An additional problem, which is arguably even more serious, is that the estimated slopes from these models will, in part, capture intraperiod (cross-sectional) differences instead of dynamic trends. Specifically, the age-period model, similar to the diff-SC model, will generate a set of intraperiod differences indexed by age, and the period-cohort model, like the LC-diff model, will produce a set of intraperiod differences indexed by cohort. Intuitively, this is because both models condition on the period linear component, so the remaining linear component (either age or cohort) is estimated as a set of cross-sectional differences rather than over-time trends. Thus, the age-period model fails to adequately capture life cycle change, and the period-cohort model fails to capture social change across cohorts. We thus advise against using either age-period or period-cohort models as part of a cohort analysis.

Table 1 summarizes the main properties of the three- and two-factor models discussed in the preceding sections. The first column provides information on what variables are used to index the model (either *A*, *P*, or *C*). For example, the LC-SC model (a three-factor model) is indexed by age and cohort (*A* and *C*), as is the age-cohort model (a two-factor model). The second column indicates the name of the model. The next three columns summarize the slopes indexed by age, period, and cohort, respectively, that are generated by the various models. In addition to various combinations of θ_1 and θ_2, the expressions for each of the two-factor models include bias terms that result from omitting some set of nonlinear components. This is best understood by comparing each two-factor model with its three-factor analog. For example, the slope estimates for the diff-SC and age-period models both contain terms involving $(\theta_1 - \theta_2)$ and θ_2. In addition, in the case of the two-factor age-period model, there are two bias terms, $\phi_{(\theta_1 - \theta_2)}$ and ϕ_{θ_2}. These bias terms result from the fact that the two-factor age-period model does not adjust for the cohort nonlinearities, leading to an omitted variable problem. In the age-cohort model, the omitted variables are the nonlinear period components; in the period-cohort model, the omitted variables are the nonlinear age components.

THE ANATOMY OF COHORT ANALYSIS: DECOMPOSING COHORT CAREERS

At this point, we have the set of conceptual and methodological tools not only to construct and compare a set of cohort careers, but also to decompose them into various underlying, theoretically distinct components. Specifically, we will now show how any set of cohort careers can be fruitfully decomposed into a life cycle curve, a social change curve, period fluctuations, and a set of terms representing cell-specific heterogeneity or, equivalently, interactions among age, period, and/or cohort. Furthermore, the combination of an age curve, period fluctuations, and cell-specific heterogeneity terms results in a collection of intracohort trends representing life cycle change, and the combination of a cohort curve, period fluctuations, and cell-specific heterogeneity terms results in a set of intercohort trends representing social change.[37]

To understand the need for such a decomposition, consider again Figure 1, which displays the cohort careers using the GSS party strength data. Recall that these careers are constructed using the informal approach, namely, by stratifying on cohort and plotting cell-specific summaries (i.e., means) across levels of age. Unfortunately, the results in Figure 1 are hampered by the fact that because there is so much heterogeneity, it is difficult, if not impossible, to discern the underlying trends and patterns in the data. It is thus useful to decompose the cell means in Figure 1 into distinct parts using the LC-SC model. These various pieces can be reassembled into simpler variants of the comparative cohort careers under study, revealing underlying intra- and intercohort trends, as well as the contributions of different types of variability.

In the following sections we present cohort careers that are built from the LC and SC slopes and then incorporate, in sequence, age and cohort nonlinearities, period fluctuations, and cell-specific heterogeneity (or interactions among age, period, and/or cohort). We then conclude this section with an overview of the four basic components that can be used to summarize any set of cohort careers.

Initial Building Blocks of Cohort Careers: θ_1 and θ_2

The simplest way to represent intra- and intercohort trends is in terms of the LC and SC slopes indexed by age and year of birth, respectively. Accordingly, we define *slopes-only comparative cohort careers* using the following parametric expression[38]:

$$\overbrace{\theta_1(i - i^*)}^{\substack{\text{Intracohort Trend} \\ \text{(Life Cycle Change)}}} + \underbrace{\theta_2(k - k^*)}_{\substack{\text{Intercohort Trend} \\ \text{(Social Change)}}} \quad \text{for } i = 1, \ldots, I \text{ in each cohort } k, \tag{15}$$

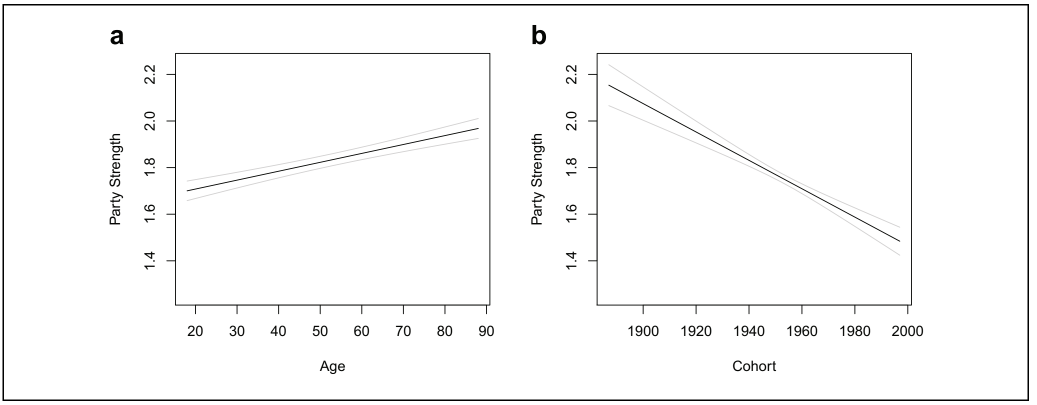

Figure 2. LC and SC slopes (intra- and intercohort trends).
Note: The life cycle (LC) slope (a), indexed by age, is an overall (linear) intracohort trend (i.e., life cycle change). The social change (SC) slope (b), indexed by year of birth, is an overall (linear) intercohort trend (i.e., social change). Upper and lower lines denote 95 percent confidence intervals. Outcome is strength of party identification, treated as a continuous variable, with responses coded as 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

where the LC slope (θ_1) denotes an overall (linear) intracohort trend for a given cohort, k and the SC slope (θ_2) denotes an overall (linear) intercohort trend for a given age group i.[39] To aid in the interpretation of equation (15), we first provide visualizations of the slopes separately and then combined, using the GSS data on political party strength.

Figure 2 provides line plots of the LC and SC slopes estimated using the LC-SC model, with upper and lower lines denoting 95 percent confidence intervals. Both slopes are statistically significant well beyond conventional levels. The estimated LC slope displayed in Figure 2a reveals a strong increase in party identification within cohorts ($\widehat{\theta}_1 = 0.019$, $S.E. = 0.025$), and the SC slope shown in Figure 2b indicates a sharp drop in party strength as successive cohorts are compared through time ($\theta_2 = -0.030$, $S.E. = 0.033$). In short, results shows that party affiliation has tended to strengthen throughout the life-course, but it has also weakened considerably because of social change.

Figure 3 combines the two graphs in Figure 2, displaying the slopes-only careers for a selected subset of cohorts. Similar to the results in Figure 1, horizontal comparisons within cohorts reflect an overall (linear) intracohort trend (which is represented by the LC slope), and vertical comparisons across cohorts in Figure 2 reflect an overall (linear) intercohort trend (which is represented by the SC slope). Specifically, the steeper the estimated slope of θ_1 in Figure 2a, the greater the overall (linear) intracohort trend (i.e., life cycle change). If θ_1 were zero, then Figure 2a would simply be a flat line, and Figure 3 would consist of a series of horizontal lines. In this case, there would be no horizontal differences within cohorts, reflecting a situation in which there is no overall (linear) intracohort trend or, equivalently, no life cycle change. By contrast, the steeper the slope of θ_2 in Figure 2b, the greater the overall (linear) intercohort trend (i.e., social

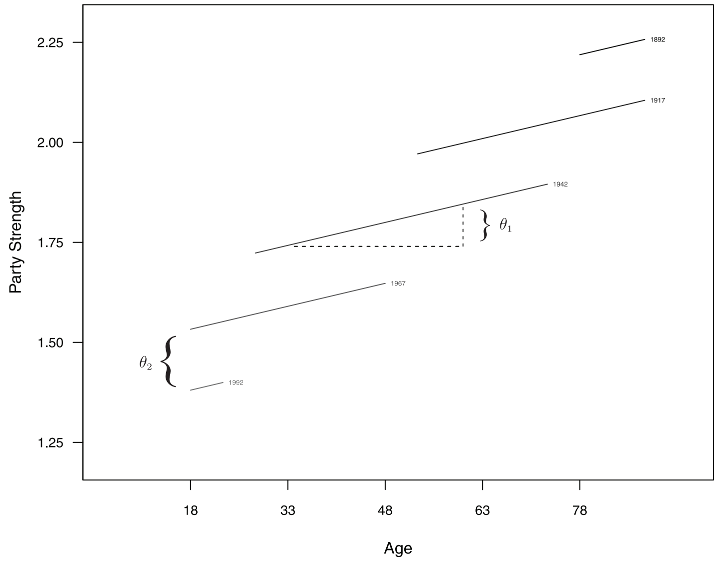

Figure 3. Slopes-only comparative cohort careers.
Note: This graph displays slopes-only comparative cohort careers for selected cohorts. Careers are estimates of $\theta_1(i - i^*) + \theta_2(k - k^*)$ for each selected cohort k. Annotations indicate the spacing of the cohorts vertically is governed by θ_2, and the slope for each cohort is a function of θ_1. Outcome is strength of party identification, treated as a continuous variable, with responses coded as 0 (independent or apolitical), 0 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

change). If θ_2 were zero, then Figure 2b would simply be a flat line and the slopes in Figure 3 would collapse on top of each other. In this case, there would be no vertical differences between the cohort careers, reflecting a scenario in which there is no overall (linear) intercohort trend or, equivalently, an absence of social change.

Incorporating Age and Cohort Nonlinearities

So far we have only considered cohort careers in which the intra- and intercohort trends are represented by slopes. However, we can also incorporate age and cohort nonlinearities, providing a richer, more detailed account of life cycle and social change. Including these nonlinearities in equation (15) defines *curves-only comparative cohort careers*:

$$\overbrace{\theta_1(i - i^*) + \widetilde{\alpha}_i}^{\substack{\text{Intracohort Trend} \\ \text{(Life Cycle Change)}}} + \underbrace{\theta_2(k - k^*) + \widetilde{\gamma}_k}_{\substack{\text{Intercohort Trend} \\ \text{(Social Change)}}} \text{ for } i = 1, \ldots, I \text{ in each cohort } k, \tag{16}$$

which represents the intracohort trend for a given cohort k in terms of the *LC curve* (indexed by age), or $\theta_1(i - i^*) + \widetilde{\alpha}_i$, and the intercohort trend for a given age group i in

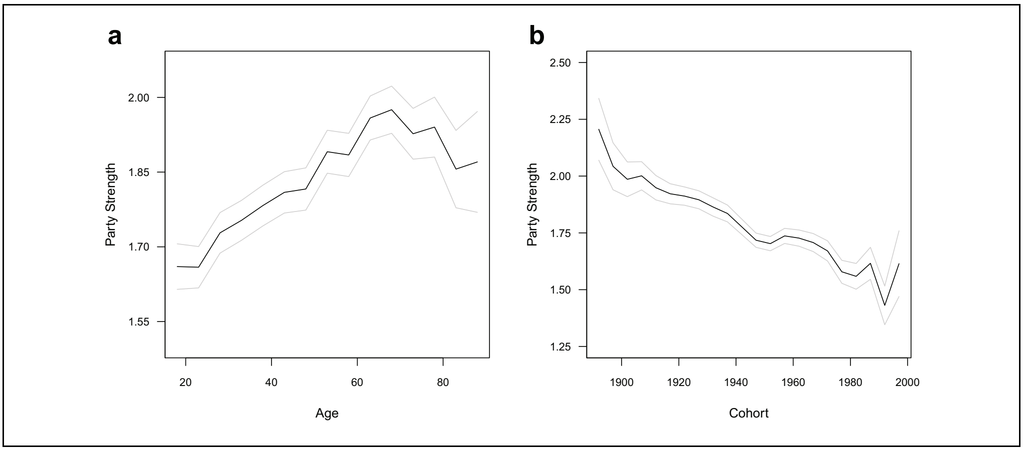

Figure 4. LC and SC curves (intra- and intercohort trends).
Note: The life cycle (LC) curve (a), or $\theta_1(i - i^*) + \widetilde{\alpha}_i$, for age groups $i = 1, \ldots, I$, represents an overall intracohort trend. The social change (SC) curve (b), or $\theta_2(k - k^*) + \widetilde{\gamma}_k$, for cohort groups $k = 1, \ldots, K$, represents an overall intercohort trend. Outcome is strength of party identification, treated as a continuous variable, with responses coded as 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

terms of the *SC curve* (indexed by birth year), or $\theta_2(k - k^*) + \widetilde{\gamma}_k$. Despite its relative simplicity, the expression in equation (16) is a highly informative, compact summary of life cycle and social change. As with the slopes-only cohort careers discussed in the previous section, to assist in the interpretation of equation (16), we present visualizations of the curves separately and then together.[40]

Figure 4 displays line plots of the LC and SC curves estimated using the LC-SC model, with upper and lower lines denoting 95 percent confidence intervals. Specifically, Figure 4a displays, for age groups $i = 1, \ldots, I$, estimates of the LC curve, which represents an overall intracohort trend. Figure 4b shows estimates of the SC curve, which represents an overall intercohort trend, for cohort groups $k = 1, \ldots, K$. These are two of the four principal components of our decomposition.

Broadly supporting Converse's (1976) arguments, we find that party strength tends to increase over the life-course, with a decline around age 65.[41] Similar to the results for the SC slope, Figure 4b reveals that, as we compare successive cohorts through time, there is a nearly monotonic decline in party strength. This has important implications for theories of voting behavior, indicating a dramatic generational disaffiliation from the major two-party system in the United States (cf. Abramson 1976).

Figure 5 combines Figures 4a and 4b, displaying the curves-only comparative cohort careers defined in equation (16). The horizontal differences within cohorts are a function of the LC curve, which represents an overall intracohort trend. Because each cohort's career is a section of the LC curve, the careers in Figure 5 are depicted as a collection of parallel curves that have the same shape, for each cohort-specific section, as the curve in Figure 4a. Findings show a clear increase in strength of party

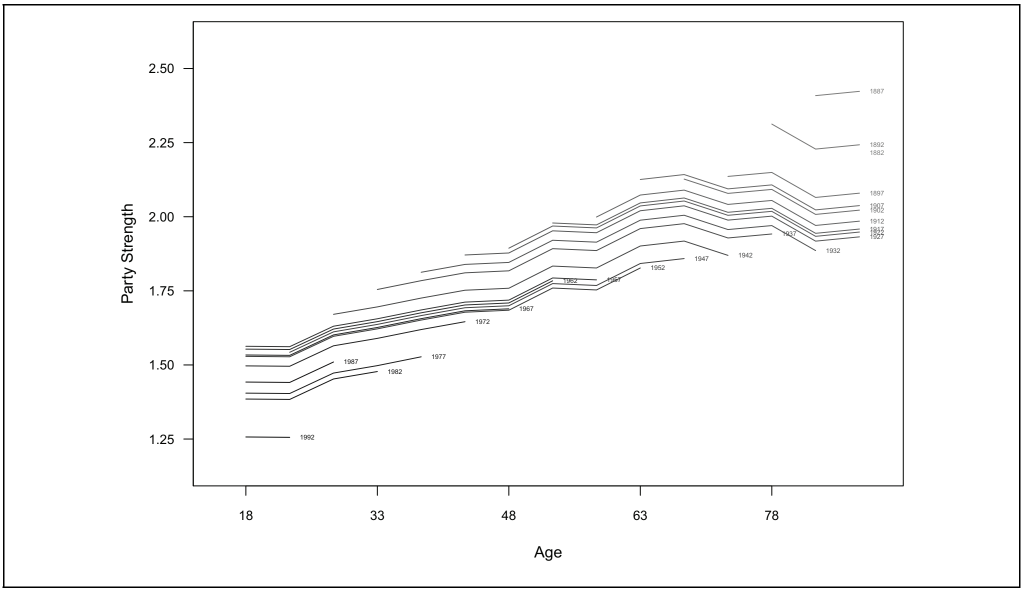

Figure 5. Curves-only comparative cohort careers.
Note: This graph displays a collection of curves-only comparative cohort careers. Each curve is a function of $\theta_1(i - i^*) + \widetilde{\alpha}_i + \theta_2(k - k^*) + \widetilde{\gamma}_k$ for $i = 1, \ldots, I$ in each cohort k. Vertical differences between cohorts are a function of $\theta_2(k - k^*) + \widetilde{\gamma}_k$, and horizontal differences within cohorts are a function of $\theta_1(i - i^*) + \widetilde{\alpha}_i$. The two most extreme cohorts, which are each based on only one age group, are not displayed. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

identification and then a decline starting around age 65. The vertical differences between the curves are given by the SC curve in Figure 4b, which represents an overall intercohort trend.[42] As in Figure 4, Figure 5 shows that more recent cohorts exhibit substantially weaker party identification than do earlier cohorts.

Incorporating Period Fluctuations

Besides the age and cohort nonlinearities, we can also include the period nonlinearities. As we will discuss, the period nonlinearities have a dual role, appearing in expressions for both intra- and intercohort trends, reflecting an inherently "interactive" quality with respect to age and cohort. Including the period nonlinearities along with the LC and SC curves in equation (16) defines what we term *adjusted comparative cohort careers*:

$$\overbrace{\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}}^{\substack{\text{Intracohort Trend} \\ \text{(Life Cycle Change)}}} \hspace{-4.5em} \underbrace{\phantom{\widetilde{\pi}_{i+k-I}} + \theta_2(k - k^*) + \widetilde{\gamma}_i}_{\substack{\text{Intercohort Trend} \\ \text{(Social Change)}}} \quad \text{for} \quad i = 1, \ldots, I \text{ in each cohort } k. \qquad (17)$$

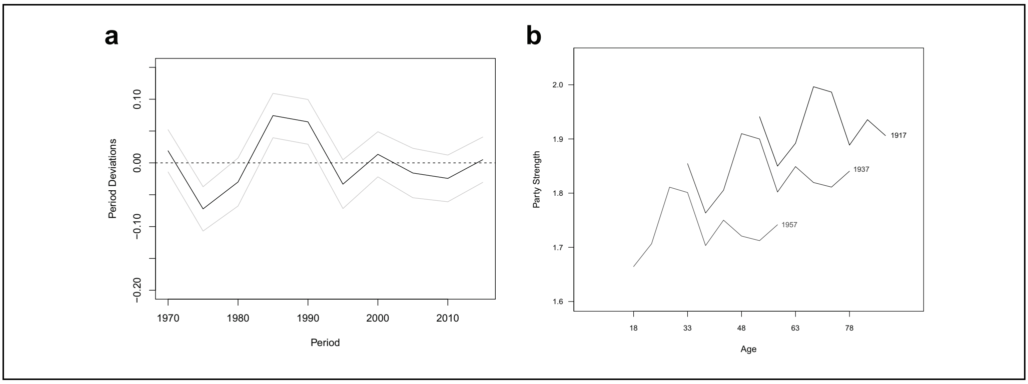

Figure 6. Period nonlinearities for selected cohorts.
Note: (a) Overall estimated period nonlinearities for groups $i = 1, \ldots, I$, with the intercept set to zero. (b) Estimated period nonlinearities for three selected cohorts (1917–1921, 1937–1941, and 1957–1961), with calculations based on the adjusted local social change curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{i+k-I}$, for $i = 1, \ldots, I$ in each selected cohort k. The peak corresponds to the nonlinearity in the 1985–1989 period group, which is experienced at different ages for each cohort, resulting in the observed staggered pattern. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

These careers are "adjusted" in the specific sense that they are purged of cell-specific heterogeneity (i.e., joint interaction terms). The intracohort trend for a given cohort k in equation (17) is now represented by the *adjusted local LC curve*, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$, and the intercohort trend for a specific age group i is denoted by the *adjusted local SC curve*, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I}$.[43] Again, these intra- and intercohort trends are "adjusted" inasmuch as they are purged of cell-specific heterogeneity. Additionally, these curves are "local" in the sense that, by incorporating the period nonlinearities, they are particular (or "local") to each age or cohort group.[44]

To understand how the cohort careers defined in equation (17) differ from those outlined in equations (15) and (16), we focus first on analyzing the period nonlinearities, which have a unique relationship with respect to age and cohort. Figure 6 displays the estimated period nonlinearities using the LC-SC model. Figure 6a shows the party nonlinearities for all observed periods, with the intercept set to zero. This is the third key component of our decomposition analysis. Results indicate statistically significant period nonlinearities, with a "bump" in party strength in the late 1980s and early 1990s. This is how period nonlinearities are typically presented in the APC literature: as a single line plot showing deviations across levels of period (see, e.g., Fosse and Winship 2019a, 2019b).

However, Figure 6a obscures a fundamental, albeit not widely appreciated, feature of APC data: cohorts experience the period nonlinearities at different ages. To put it another way, as reflected in the period nonlinearities, the summary value of the outcome (i.e., mean party strength) for each age group differs depending on the level of the cohort variable. Inherent in any APC data, therefore, is a special kind of "interaction" between age and period that is a function of cohort.

This feature of APC data can be illustrated using the party strength data. Figure 6b displays the estimated period nonlinearities for three selected cohorts (1917–1921, 1937–1941, and 1957–1961). The peak corresponds to the nonlinearity in the 1985–1989 period group, which is experienced at different ages for each cohort, resulting in the observed staggered pattern.[45] Each set of period nonlinearities has the exact same shape as that shown in Figure 6a. However, each cohort's particular set of period nonlinearities must, by virtue of the deterministic relationship among age, period, and cohort, be aligned with a specified age group. This is indicated by the indexing of the period nonlinearities in equation (17), which are a function of both age (i) and cohort (k). For example, for cohort k, the jth period nonlinearity occurs in the ith age group. However, for cohort $k+1$, a more recent cohort, the jth period nonlinearity occurs in the $(i-1)$ age group, an earlier age category. This is simply another way of stating that individuals born at later times must necessarily experience the period nonlinearities at earlier ages. Visually, this is a matter of "moving" the period nonlinearities in Figure 6b to the left for more recent cohorts and to the right for older cohorts.

The period nonlinearities, which appear at different age groups for each cohort, allow the intracohort trends to vary across cohorts and the intercohort trends to differ across age groups. Figure 7 displays the adjusted comparative cohort careers using the GSS party strength data. As noted earlier, the adjusted local LC and SC curves, which represent intra- and intercohort trends particular (or local) to each cohort and age group, respectively, constitute adjusted comparative cohort careers (equation 17). Importantly, the *only* difference between Figures 5 and 7 is that the latter includes the period nonlinearities, and the former does not. Specifically, the cohort careers in Figure 7 are a function of three distinct components: first, the LC slope, θ_1, and the nonlinear age terms, $\widetilde{\alpha}_i$, which together constitute an overall intracohort trend; second, the SC slope, θ_2, and the nonlinear cohort terms, $\widetilde{\gamma}_i$, which combined constitute an overall intercohort trend; and, third, the nonlinear period terms, $\tilde{\pi}_i$, which allow the intra- and intercohort trends to vary across cohort and age groups, respectively. It is these separate pieces that, in combination, result in the set of cohort careers displayed in Figure 7.

Incorporating Interactions

So far we have only considered the main parameters for age, period, and cohort. However, sociologists have also been interested in understanding possible interactions among age, period, and/or cohort. For example, Ryder (1965:848–49) pointed out how age and period interacted in the case where war deferentially affected the cohorts of men who were likely to be in the military. The classic case is perhaps Elder's (1974) life-course analysis of cohorts who came of age during the Great Depression. More recently, Neil and Sampson (2021) showed how cohorts who were young during the crack epidemic were more likely to have extended criminal careers; and Shen et al. (2020) revealed the resulting incarceration rates for those cohorts, and corresponding

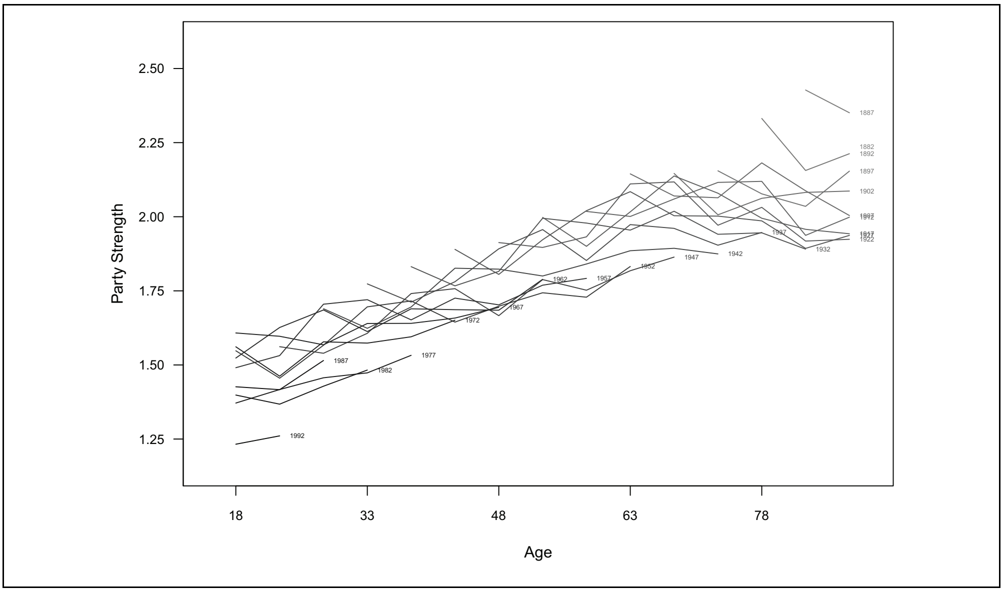

Figure 7. Adjusted comparative cohort careers.
Note: This graph displays a collection of adjusted comparative cohort careers. Each curve is a function of $\theta_1(i - i^*) + \widetilde{\alpha}_i + + \widetilde{\pi}_{i+k-I} + \theta_2(k - k^*) + \widetilde{\gamma}_k$ for $i = 1, \ldots, I$ in each cohort k. Vertical differences between cohorts are a function of $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{i+k-I}$, and horizontal differences within cohorts are a function of $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$. The two most extreme cohorts, which are each based on only one age group, are not displayed. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

long-term prison sentences, were an important factor in the rise of mass incarceration. To cite another example, Jager et al. (2022) demonstrated that changes in the legal drinking age affected cohorts differently depending on their age at the time the law was enacted.

Our discussion of the LC-SC model to this point can be thought of as involving a model with a set of main parameters for age, period, and cohort. The last component of our decomposition are the terms for cell-specific heterogeneity, or the η_{ijk}'s, which capture interactions among age, period, and/or cohort. We are only able to provide a brief overview here of how to analyze such interactions. A full treatment requires a complete, separate article, which is an important task for future research.[46] As defined in equation (10), the η_{ijk}'s are orthogonal to all other terms in the LC-SC model. The situation here is analogous to a traditional analysis-of-variance model, where one first specifies a set of main parameters and then considers interaction terms that are orthogonal to the main parameters (see, e.g., Rosnow and Rosenthal 1995).

Including interaction terms along with the parameters in equation (17) defines what we call *unadjusted comparative cohort careers*:

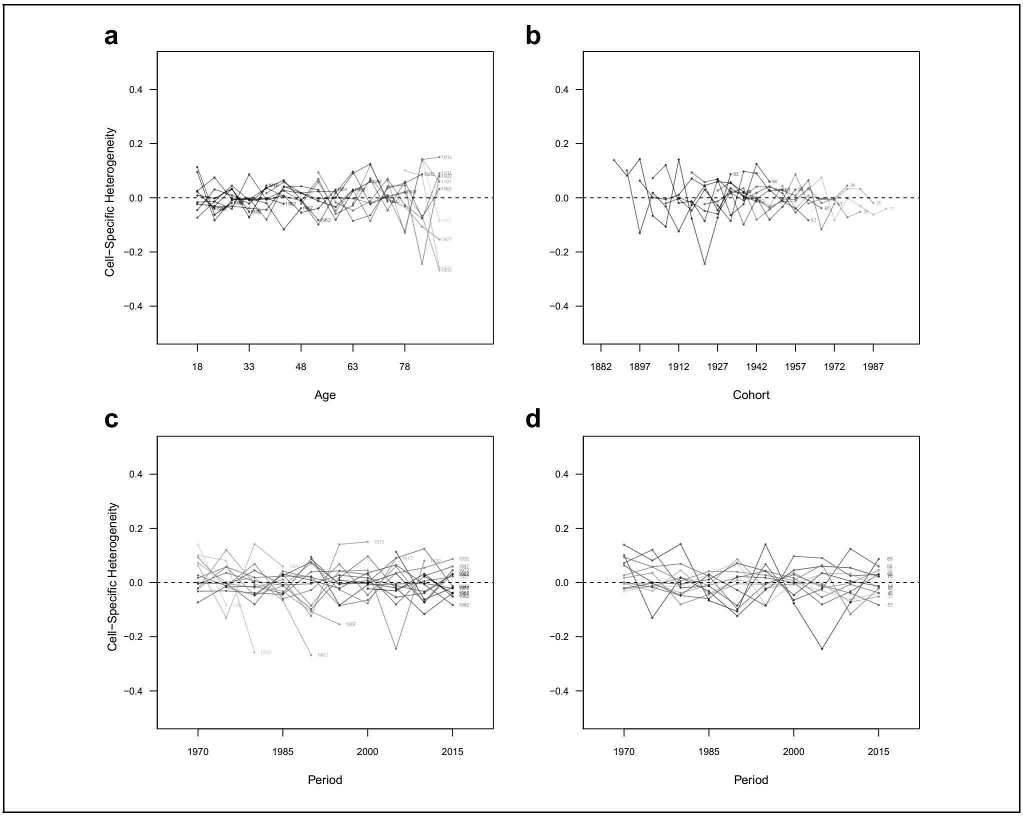

Figure 8. Cell-specific heterogeneity by age, period, and cohort.

Note: Each panel displays cell-specific heterogeneity within age or cohort groups across levels of age, period, or cohort: (a) heterogeneity within cohort groups across age levels, (b) heterogeneity within age groups across cohort levels, (c) heterogeneity within cohort groups across period levels, and (d) heterogeneity within age groups across period levels. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

$$\overbrace{\theta_1(i-i^*)+\widetilde{\alpha}_i+\rlap{$\underbrace{\phantom{\widetilde{\pi}_{i+k-I}+\eta_{i[i+k-I]k}+\theta_2(k-k^*)+\widetilde{\gamma}_i}}_{\substack{\text{Intercohort Trend}\\ \text{(Social Change)}}}$}\widetilde{\pi}_{i+k-I}+\eta_{i[i+k-I]k}}^{\substack{\text{Intracohort Trend}\\ \text{(Life Cycle Change)}}}+\theta_2(k-k^*)+\widetilde{\gamma}_i \quad \text{for} \quad i=1,\ldots,I \text{ in each cohort } k, \tag{18}$$

which is equivalent to the definition of cohort careers given in equation (9).[47] These cohort careers are “unadjusted” in the sense that they are estimated by calculating and then arranging, for each cohort, cell-specific means across levels of age, as shown in Figure 1. The only difference between Figure 1 and Figure 7 is that the former includes the cell-specific heterogeneity terms and the latter does not.

The cell-specific heterogeneity, or the η_{ijk} terms, are displayed in Figures 8a to 8d. Specifically, Figure 8a displays heterogeneity within cohort groups across age levels,

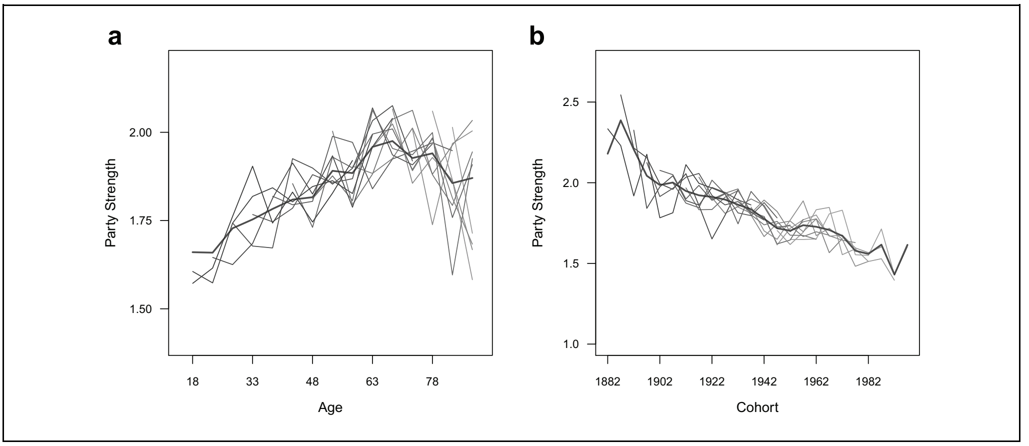

Figure 9. Spaghetti plots of unadjusted local LC and SC curves (intra- and intercohort trends).

Note: (a) Estimates of unadjusted local life cycle (LC) curves, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \eta_{i[i+k-I]k}$, for $i = 1, \ldots, I$ in each cohort k. Each curve represents an intracohort trend for a particular cohort. Gray lines denote the local LC curves; the black line indicates the LC curve. (b) Estimates of unadjusted local social change (SC) curves, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I} + \eta_{i[i+k-I]k}$, for cohort groups $k = 1, \ldots, K$ in each age group i. Each curve represents an intercohort trend for a particular age group. Gray lines denote the local SC curves; the black line indicates the SC curve. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

Figure 8b shows heterogeneity within age groups across cohort levels, Figure 8c displays heterogeneity within cohort groups across period levels, and Figure 8d shows heterogeneity within age groups across period levels. Each panel of Figure 8 consists of the same set of η_{ijk} terms; the only difference is in the choice of the x-axis. These heterogeneity terms are the final component of our decomposition.

Overall, at least in the case of the party strength data, there is no clear structure to the cell-specific heterogeneity, or equivalently, interactions between age, period, and/or cohort. In fact, the distribution of the heterogeneity is approximated by the normal curve, with only a few outlying deviations (see Figure B5 in Appendix B). However, as shown in Figure 8a, there is some evidence that older ages in earlier cohorts exhibit larger fluctuations than do younger ages in later cohorts. This is not particularly surprising given the small cell sizes among the oldest age groups in the earliest cohorts.

An additional way to examine the cell-specific heterogeneity is to separately plot, in a way similar to that in the previous sections, the intra- and intercohort trends in equation (18). Figure 9 displays spaghetti plots of intra- and intercohort trends after incorporating the cell-specific heterogeneity. Specifically, Figure 9a shows estimates of each *unadjusted local LC curve*, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \eta_{i[i+k-I]k}$, for $i = 1, \ldots, I$ in a given cohort k, which represents an intracohort trend for a particular cohort group. Gray lines in Figure 9a denote the unadjusted local LC curves, which are also displayed as a trellis plot in Figure B6 in Appendix B. The black line indicates the overall LC curve, which is also displayed in Figure 4a. In contrast, Figure 9b displays estimates of each *unadjusted local SC curve*, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I} + \eta_{i[i+k-I]k}$ for cohort groups $k = 1, \ldots, K$ in a given age group i, which represents an intercohort trend for a particular age group. Gray lines in Figure 9b denote the unadjusted local SC curves, which are also displayed as a trellis plot in Figure B7 in Appendix B. The black line indicates the overall SC curve, which is also shown in Figure 4b. In comparing the unadjusted local curves in Figure 9 with their adjusted counterparts (see Figures B2, B3, and B4 in Appendix B), it is clear additional variability is introduced by the heterogeneity terms, although the results do not suggest any overall patterning or particularly informative "bumps" beyond those captured by main parameters for age, period, and cohort.

Measures of fit are consistent with the findings in Figures 8 and 9, suggesting that, in the case of the party strength data, one should drop the η_{ijk}'s from the LC-SC model. Both the Akaike information criterion (AIC) and Schwarz's Bayesian information criterion (BIC) indicate that an LC-SC model without the extra heterogeneity terms provides a better fit to the data than a model that includes these terms. Specifically, the AIC and BIC values for the LC-SC model are 180,201.6 and 178,834.6 (on the basis of 150 degrees of freedom), and the corresponding values for the LC-SC model without the heterogeneity terms are, respectively, 179,179.1 and 178,753.6 (on the basis of 46 degrees of freedom).[48] Given that a simpler model without the heterogeneity terms better fits the data, we can conclude that the cohort careers in Figure 7, which are purged of cell-specific heterogeneity, provide a parsimonious, informative representation of the main trends and patterns in the data. However, this may not always be the case, and in some populations (or subsets of populations) this additional heterogeneity may be highly informative, reflecting singular events or complex patterns that would otherwise be overlooked.

SUMMARY: FOUR MAIN COMPONENTS

The previous sections illustrated how a set of cohort careers depicted in Figure 1 can be decomposed into six different components, namely, the LC and SC slopes, age, period, and cohort nonlinearities, and a set of terms representing cell-specific heterogeneity or, equivalently, interactions among age, period, and/or cohort. As shown in equations (15) to (18) and Figures 3, 5, and 7, these pieces can then be assembled in various ways to construct simpler types of comparative cohort careers, showing underlying trends and patterns in the data that may be concealed using the informal, purely graphical approach.

However, a careful inspection of equation (9) reveals a further simplification is possible. Instead of six different parts, the cell means in Figure 1 can be fruitfully decomposed into just four components, as shown in Figure 10: an LC curve representing an overall intracohort trend or change over the life cycle, an SC curve representing an overall intercohort trend or social change, a set of common cross-period temporal fluctuations that permit variability across cohort careers, and a set of terms denoting cell-specific heterogeneity. Somewhat remarkably, these four components, depicted in Figure 10, are all that are needed to provide a concise summary of the main findings from any collection of comparative cohort careers. Specifically, as indicated by

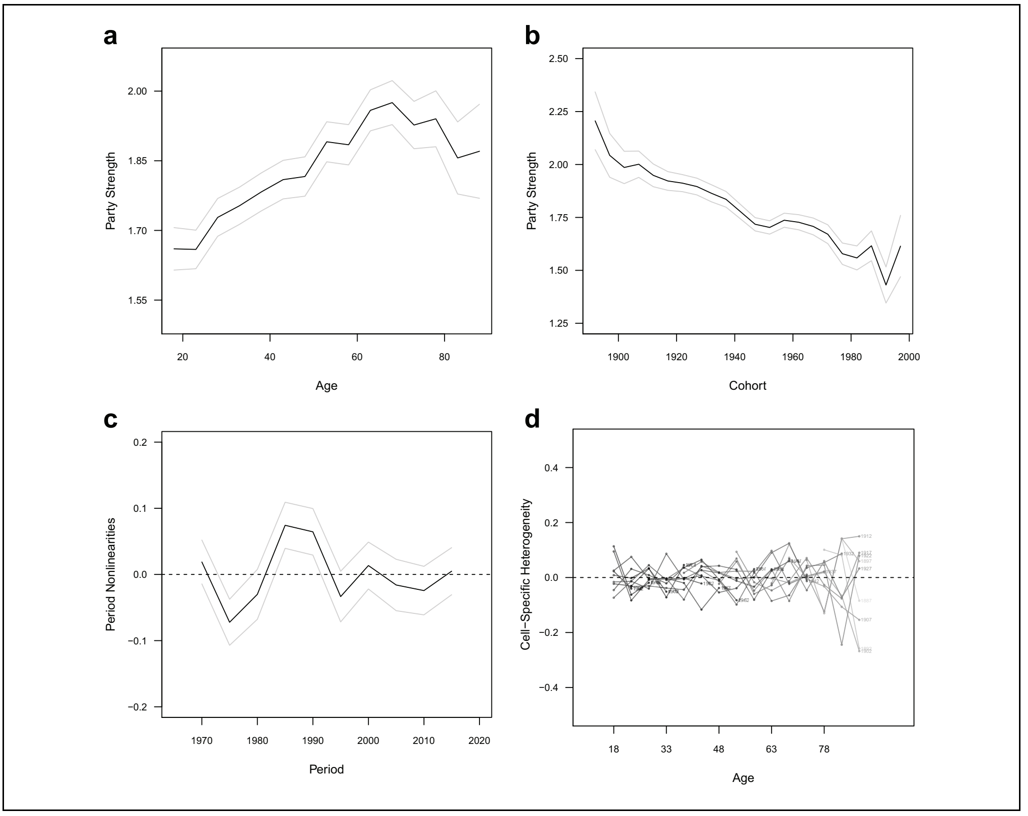

Figure 10. Four main components of comparative cohort careers.
Note: (a) Estimates of the life cycle curve, or $\theta_1(i - i^*) + \widetilde{\alpha}_i$, for age groups $i = 1, \ldots, I$, which represents an overall intracohort trend. (b) Estimates of the social change curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k$, for cohort groups $k = 1, \ldots, K$, which represents an overall intercohort trend. (c) Estimates of the period nonlinearities, or $\widetilde{\pi}_j$, for period groups $j = 1, \ldots, J$. (d) Estimates of the terms representing cell-specific heterogeneity, or η_{ijk}, for all observed age-period combinations $i = 1, \ldots, I$ and $j = 1, \ldots, I$ (and thus $k = 1, \ldots, K$). Upper and lower lines in (a), (b), and (c) denote 95 percent confidence intervals. Random jitter added to cell-specific heterogeneity in (d) to reveal overlapping estimates. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

equation (9) (see also Figure 9), the combination of the LC curve, period deviations, and terms for cell-specific heterogeneity results in a set of intracohort trends, and the combination of the SC curve, period deviations, and cell-specific heterogeneity terms results in a set of intercohort trends. We now review each of these four components.

The LC and SC curves, as shown in Figures 10a and 10b, capture overall intra- and intercohort trends (see also Figure 4). When combined, they constitute the curves-only cohort careers, as defined in equation (16) and displayed in Figure 5. In the case of party strength, within cohorts there is a general strengthening in party affiliation as cohort members age through time; in contrast, as successive cohorts are compared through time, results reveal a steep drop in party strength. In short, we find strong evidence of considerable life cycle and social change.

The period fluctuations, as displayed in Figure 10c, introduce additional structured heterogeneity, allowing the intra- and intercohort trends to differ across cohort and age groups, respectively (see also Figure 6). As discussed earlier, this reflects the inherently “interactive” quality of the period nonlinearities, which are experienced at different ages for different cohorts. The adjusted cohort careers, which include the period nonlinearities along with the overall LC and SC curves, are defined in equation (17) and displayed in Figure 7 (see also Figure B3 and Figure B4 in Appendix B). In contrast to the curves-only cohort careers, each cohort now has its own particular shape, with some cohorts experiencing bigger “bumps” than others. In particular, the period nonlinearities reveal a noticeable increase in party strength during the Reagan and Bush presidencies.

Finally, as shown in Figure 10d, the cell-specific heterogeneity introduces additional cross-cohort variability, albeit without a fixed patterning, unlike the period nonlinearities (see also Figure 8). The unadjusted cohort careers, which are defined in equation (18), incorporate this extra variability, resulting in curves that are equivalent to those obtained by calculating and arranging the cell-specific summaries (means) for each cohort by age, as displayed in Figure 1. Consistent with a visual inspection of the plots in Figures 8 and 9, measures of fit suggest this additional heterogeneity is not particularly informative in the case of party strength.

CONCLUSIONS

In keeping with Ryder’s (1965) goals for cohort analysis, in this article we outlined a formal mathematical framework for constructing, comparing, and decomposing a set of cohort careers. Using data from the GSS on strength of political party affiliation, we first illustrated how one can informally construct cohort careers by stratifying on cohort and ordering cell-specific summaries (e.g., mean party strength) across levels of age. Each trajectory is a cohort career because it reflects the development of a cohort as it ages through time (i.e., across periods), but also comparative because it exhibits differentiation as successive cohorts are compared through time (i.e., across periods). Unfortunately, cohort careers constructed using this approach, although useful for an initial analysis, are problematic in two respects: first, by conflating a number of theoretically distinct components, they can render invisible, or at least indeterminate, underlying trends and patterns in the data; second, by casting the careers in terms of age and cohort, results are easily misinterpreted as exclusively attributable to just age and cohort rather than all three dimensions.

To provide a more illuminating approach to cohort analysis, we introduced three main improvements. First, we presented a mathematical definition of comparative cohort careers that incorporates parameters explicitly representing all three temporal dimensions and enables a simple but powerful decomposition of cohort careers into multiple, theoretically distinct components. To estimate the parameters of a set of cohort careers, we introduced the LC-SC model. Although closely related to a model that assumes a zero period linear effect, the LC-SC model is distinct in that it is entirely identified, reflecting joint sets of parameters that can be straightforwardly estimated from a set of APC data.

Second, we contrasted the properties of the LC-SC model with various other models for cohort analysis, including two alternative three-factor models and all three logically possible two-factor models (age-cohort, age-period, and period-cohort). As we demonstrated, the models that condition on period (or, more precisely, the period linear component) will generate slopes that represent static (i.e., cross-sectional) intraperiod differences rather than dynamic (i.e., over-time) trends. This is inconsistent with the goal of decomposing a set of cohort careers into intra- and intercohort trends (i.e., life cycle and social change), which are inherently dynamic, not static, quantities. Moreover, in general, two-factor models will produce biased results because one of the temporal dimensions has been dropped. For these reasons, for a cohort analysis, we strongly prefer the LC-SC model over the alternative three- and two-factor models considered in this article.

Third, we presented the article's major accomplishment: using the LC-SC model, we illustrated how comparative cohort careers can be decomposed into just four components, namely, an LC curve representing an overall intracohort trend or change over the life cycle, an SC curve representing an overall intercohort trend or social change, a set of common cross-period temporal fluctuations that enable variability across cohort careers, and a set of terms representing cell-specific heterogeneity. As we illustrated, these distinct parts can be reassembled into simpler cohort careers, revealing underlying trends and patterns that are otherwise easily obscured. In considering various cohort careers, we also illustrated a fundamental, but oft overlooked, property of APC data: because the period nonlinearities are experienced by different cohorts at different ages, such temporally organized data are inherently "interactive." Accordingly, as demonstrated with the party strength data, the period nonlinearities can inject considerable variability into the cohort careers, in effect relaxing the restriction that each cohort's career is some section of an overall LC curve (i.e., an overall intracohort trend).

Our analyses of the GSS party strength data uncovered a number of important trends and patterns. Overall, results indicated a considerable increase in party strength within cohorts, with a slight decline in late adulthood, lending support to Converse's (1976) argument that partisanship tends to increase over the life-course. However, our analyses also uncovered a remarkably steep drop in party strength as we compared cohorts through time, indicating the presence of substantial social change. Additionally, we found some variability due to the period nonlinearities, with a noticeable upward "bump" in party strength in the late 1980s and early 1990s. Because of the inherently "interactive" nature of APC data, this marked deviation altered the shape of the careers of different cohorts at different ages. Finally, although there are a few potentially suggestive deviations (e.g., the relative increase in party strength for the cohort born between 1912 and 1916), overall we found relatively little patterning attributable to cell-specific heterogeneity, or equivalently, the interactions among age, period, and/or cohort.

Notwithstanding the contributions outlined above, there are several limitations of our approach that point to potentially beneficial directions for further research. First, although we compared the LC-SC model to two alternative three-factor models and all three possible two-factor models, there are certainly other models that can be used to analyze APC data (see, e.g., Mayer and Huinink 1990; Schulhofer-Wohl and Yang 2016). The properties of these models and their significance for realizing Ryder's cohort approach should be considered and compared with the models discussed in this article. Second, Ryder's cohort approach is based on distinguishing intra- from intercohort trends, not on identifying unique effects for age, period, and cohort. However, in some cases, uncovering separate effects may be of primary concern. Thus, additional research should consider how attempts at identifying distinct APC effects and Ryder's cohort approach can complement each other and be integrated into a single analysis. Third, Ryder's cohort approach is based on comparing the life cycles of cohorts, which are distinct from the life cycles of individuals (for related points, see Bernardi, Huinink, and Settersten 2019). For researchers interested in individual-level life cycle change, a cohort analysis may not be appropriate because, unlike individuals, cohorts are populations, and thus can undergo compositional changes (e.g., because of migration in and out of the cohort distribution). Fourth, although we demonstrated how the cell-specific heterogeneity from the LC-SC model can be incorporated into a cohort analysis, we did so at a high level of generality. Future research should explore the various ways of parameterizing and visualizing this extra variability, considering the extent to which different, possibly contrasting, techniques might provide fresh insight into underlying patterns in the data. Finally, although we briefly touched on the issue in our discussion of cell-specific heterogeneity, deliberately did not discuss in detail statistical tests for determining the presence or absence of various trends and patterns in the data. Instead, because our goal was to show how APC data can be analyzed by breaking it down into its component parts, we chose to use the full set of parameters in the LC-SC model.[49] Nonetheless, the assessment of model fit with APC data is an important topic warranting further discussion.[50]

Although we could not hope to assess the suitability of all APC models for a Ryderian cohort analysis, our study suggests several guiding principles that can be applied regardless of the model used. First, rather than trying to disentangling the unique effects of age, period, and cohort on some outcome, Ryder's approach is based on distinguishing intra- from intercohort trends. Consistent with Ryder's (1968) formulation, this implies that one should reexpress the parameters in terms of the age and cohort groups (cf. equation 9). Conditional on cohort, tracking how the parameters change across age groups (and thus also periods) reveals an intracohort (or life cycle change); in contrast, given a particular age level, tracking how the parameters change across cohorts (and thus also periods) reveals an intercohort (of social change). Second, such an analysis, as a general rule, does not involve conditioning on the period linear component, or a variable containing the period linear component. As we demonstrated (see Table 1 and Appendix A), by conditioning on such a variable, one generates a synchronic rather than a diachronic estimate. Third, a Ryderian cohort approach entails incorporating terms for all three dimensions (age, period, and cohort) into a model, not just one or two. In practice, this precludes two-factor models, unless extreme care is taken in the interpretation of the parameters.[51] Finally, Ryder's approach is centered on the construction and comparison of cohort careers, which are most easily displayed as a series of line graphs, with summaries (e.g., means) plotted for each cohort as a function of age (e.g., Figure 1). Similar plots can, in principle, be created for any number of models with age, period, and/or cohort as inputs.

With the above principles in mind, it is our hope that the construction, comparison, and decomposition of cohort careers will, in line with Ryder's (1965) vision, play a central role in any analysis of temporally structured data. Ryder's cohort approach, as developed here, has the potential to uncover fundamental insights on the nature of life cycle and social change on any number of outcomes, including party strength. In particular, as we demonstrated, the decomposition of cohort careers into theoretically distinct components using the LC-SC model uncovered underlying trends and patterns in the data that, with purely graphical approaches or other parametric models, may have been misinterpreted, if not concealed outright.

APPENDIX A: INTERPRETING THE PARAMETERS OF TWO-FACTOR MODELS

In this appendix, we clarify the interpretation of the parameters from all three logically possible two-factor models (age-period, period-cohort, age-cohort). We first present each two-factor model, outlining the relationship between each model's parameters and those from a corresponding model that includes all three factors (see equations 10, 13, and 14). Next, using the age-period model as an example, we show how these relationships can be derived using matrix algebra and the logic of omitted variable bias. To avoid confusion with corresponding terms in the three-factor models introduced in the main text, in this appendix we use asterisks to denote the parameters from two-factor models.

Age-Period Model

Assume we have APC data that is indexed by $i = 1, \ldots, I$ age rows and $j = 1, \ldots, J$ period columns, with $k = 1, \ldots, K$ cohort diagonals, where $k = j - i + I$ and $K = I + J - 1$. The age-period model fit to such data has the following general form:

$$Y_{rijk} = \mu^* + \alpha_i^* + \pi_j^* + \eta_{ij}^* + \epsilon_{rij}^*, \tag{A1}$$

where μ^* is the intercept; α_i^* and π_j^* are parameters for age and period using sum-to-zero deviation (or "effect") coding; η_{ij}^* denotes group-level heterogeneity terms; and ϵ_{rij}^* is individual-level error. The simplest way to obtain the η_{ij}^* terms representing group-level heterogeneity is to calculate the residuals relative to a fully saturated model (see, e.g., Keyes and Li 2010:780; Ohtaki et al. 1990:119). Alternatively, one can specify all pairs of interactions between age and period (see, e.g., Keyes and Li 2010:780; Luo and Hodges 2020).[52]

The age-period model outlined above is equivalent to the following:

$$Y_{rijk} = \underbrace{(\mu + \phi_{\mu})}_{\mu^*} + \underbrace{((\theta_1 - \theta_2) + \phi_{(\theta_1 - \theta_2)})(i - i^*) + (\widetilde{\alpha}_i + \phi_{\widetilde{\alpha}_i})}_{\alpha_i^*}$$
$$+ \underbrace{(\theta_2 + \phi_{\theta_2})(j - j^*) + (\widetilde{\pi}_j + \phi_{\widetilde{\pi}_j})}_{\pi_j^*} + \underbrace{(\eta_{ijk} + \nu_{ijk})}_{\eta_{ij}^*} + \underbrace{\epsilon_{rijk}}_{\epsilon_{rij}^*}, \quad \text{(A2)}$$

where μ is the intercept; $\theta_1 - \theta_2 = \alpha - \gamma$ is the intraperiod slope (indexed by age); θ_2 is the SC slope; $\widetilde{\alpha}_i$ is the *i*th age nonlinearity; $\widetilde{\pi}_j$ is the *j*th period nonlinearity; ϕ_{μ}, $\phi_{(\theta_1 - \theta_2)}$, ϕ_{θ_2}, $\phi_{\widetilde{\alpha}_i}$, and $\phi_{\widetilde{\pi}_j}$ are bias terms for the intercept, intraperiod slope, SC slope, *i*th age nonlinearity, and *j*th period nonlinearity; η_{ijk} denotes terms for unique cell-specific heterogeneity; ν_{ijk} denotes terms for unique cohort-attributed heterogeneity; and ϵ_{rijk} is individual-level error. The terms in brackets below equation (A2) denote the corresponding parameters from the age-period model presented in equation (A1).

Three main points stand out from equation (A2). First, the intercept, age, and period parameters will all have some degree of bias because of the exclusion of the cohort nonlinearities from the age-period model.[53] Second, assuming there is no bias due to the exclusion of the cohort nonlinearities, either because the cohort nonlinearities are zero or the cohort variables are unrelated to the included variables, then the underlying age and period slopes of the age-period model will equal those from the diff-SC model. Finally, the group-level heterogeneity terms η_{ij}^* from the age-period model equal the sum of the unique cell-specific heterogeneity terms from the diff-SC model, η_{ijk}, and the unique cohort-attributed heterogeneity terms ν_{ijk}. The cohort-attributed heterogeneity terms are simply the predicted values from the parameters for the cohort nonlinearities (i.e., the excluded variables) using that part of the cohort variables that is unassociated with the variables for the intercept, age, and period terms (i.e., the included variables).

Period-Cohort Model

Suppose we have APC data that are indexed by $j = 1, \ldots, J$ period rows and $k = 1, \ldots, K$ cohort columns, with $i = 1, \ldots, I$ age diagonals, where $i = j - k + K$ and $I = J + K - 1$. The period-cohort model fit to such data has the following general form:

$$Y_{rijk} = \mu^* + \pi_j^* + \gamma_k^* + \eta_{jk}^* + \epsilon_{rjk}^*, \quad \text{(A3)}$$

where μ^* is the intercept; π_j^* and γ_k^* are parameters for period and cohort using sum-to-zero deviation coding; η_{jk}^* denotes group-level heterogeneity terms; and ϵ_{rjk}^* is individual-level error. The η_{jk}^* terms can be obtained in a way similar to that for the age-period model.

The period-cohort model outlined in equation (A3) is equivalent to the following:

$$Y_{rijk} = \underbrace{(\mu + \psi_{\mu})}_{\mu^*} + \underbrace{(\theta_1 + \psi_{\theta_1})(j - j^*) + (\widetilde{\pi}_j + \psi_{\widetilde{\pi}_j})}_{\pi_j^*}$$
$$+ \underbrace{((\theta_2 - \theta_1) + \psi_{(\theta_2 - \theta_1)})(k - k^*) + (\widetilde{\gamma}_k + \psi_{\widetilde{\gamma}_k})}_{\gamma_k^*} + \underbrace{(\eta_{ijk} + \nu_{ijk})}_{\eta_{jk}^*} + \underbrace{\epsilon_{rijk}}_{\epsilon_{rjk}^*}, \quad \text{(A4)}$$

where μ is the intercept; $\theta_1 = \alpha + \pi$ is the LC slope; $\theta_2 - \theta_1 = \gamma - \alpha$ is the intraperiod slope (indexed by cohort); $\widetilde{\pi}_j$ is the *j*th period nonlinearity; $\widetilde{\gamma}_k$ is the *k*th cohort nonlinearity; ψ_{μ}, ψ_{θ_1},

$\psi_{(\theta_2-\theta_1)}$, $\psi_{\tilde{\pi}_j}$, and $\psi_{\tilde{\gamma}_k}$ are bias terms for the intercept, LC slope, intraperiod slope, jth period nonlinearity, and kth cohort nonlinearity; η_{ijk} denotes terms for unique cell-specific heterogeneity; ν_{ijk} denotes terms for unique age-attributed heterogeneity; and ϵ_{rijk} is individual-level error. The terms in brackets below equation (A4) denote the corresponding parameters from the period-cohort model displayed in equation (A3).

As with the age-period model, three main conclusions follow from equation (A4). First, as indicated by the presence of the ψ parameters, the intercept, period, and cohort parameters will be biased because of the exclusion of the age nonlinearities from the period-cohort model. Second, assuming that excluding the age nonlinearities results in no bias, then the underlying period and cohort slopes will equal those from the LC-diff model. Last, the group-level heterogeneity terms η^*_{jk} equal the sum of the unique cell-specific heterogeneity terms from the LC-diff model, η_{ijk}, and the unique age-attributed heterogeneity terms ν_{ijk}. Similar to the age-period model, the age-attributed heterogeneity terms are just the predicted values from the parameters for the age nonlinearities (i.e., the excluded variables) using that part of the age variables that is unrelated to the variables for the intercept, period, and cohort terms (i.e., the included variables).

Age-Cohort Model

The remaining two-factor model is the age-cohort model. Assume we have APC data that are indexed by $i = 1, \ldots, I$ age rows and $k = 1, \ldots, K$ cohort columns, with $j = 1, \ldots, J$ period diagonals, where $j = i + k - 1$ and $J = I + K - 1$. The age-cohort model fit to this data has the following form:

$$Y_{rijk} = \mu^* + \alpha^*_i + \gamma^*_k + \eta^*_{ik} + \epsilon^*_{rik}, \tag{A5}$$

where μ^* is the intercept; α^*_i and γ^*_k are parameters for age and cohort using sum-to-zero deviation coding; η^*_{ik} denotes group-level heterogeneity terms; and ϵ^*_{rik} is individual-level error. Again, the η^*_{ik} terms can be obtained in a way similar to that for the age-period or period-cohort models.

The age-cohort model presented in equation (A5) is equivalent to the following:

$$\begin{aligned} Y_{rijk} = (\underbrace{\mu + \xi_\mu}_{\mu^*}) + \underbrace{(\theta_1 + \xi_{\theta_1})(i - i^*) + (\widetilde{\alpha}_i + \xi_{\widetilde{\alpha}_i})}_{\alpha^*_i} \\ + \underbrace{(\theta_2 + \xi_{\theta_2})(k - k^*) + (\widetilde{\gamma}_k + \xi_{\widetilde{\gamma}_k})}_{\gamma^*_k} + (\underbrace{\eta_{ijk} + \nu_{ijk}}_{\eta^*_{jk}}) + \underbrace{\epsilon_{rijk}}_{\epsilon^*_{rik}}, \end{aligned} \tag{A6}$$

where μ is the intercept; $\theta_1 = \alpha + \pi$ is the LC slope; $\theta_2 = \gamma + \pi$ is the SC slope; $\widetilde{\alpha}_i$ is the ith age nonlinearity; $\widetilde{\gamma}_k$ is the kth cohort nonlinearity; ξ_μ, ξ_{θ_1}, ξ_{θ_2}, $\xi_{\widetilde{\alpha}_i}$, and $\xi_{\widetilde{\gamma}_k}$ are bias terms for the intercept, LC slope, SC slope, ith age nonlinearity, and k th cohort nonlinearity; η_{ijk} denotes terms for unique cell-specific heterogeneity; ν_{ijk} denotes terms for unique period-attributed heterogeneity; and ϵ_{rijk} is individual-level error. The terms in brackets below equation (A6) refer to the corresponding parameters from the age-cohort model shown in equation (A5).

As with the other two-factor models, there are three main takeaways from equation (A6). First, the intercept, age, and cohort parameters will be biased because of the exclusion of the age nonlinearities from the age-cohort model. Second, assuming that excluding the period nonlinearities produces no bias, then the underlying age and cohort slopes will equal those from the LC-SC model. Finally, the group-level heterogeneity terms η^*_{ik} equal the sum of the unique cell-specific heterogeneity terms from the LC-SC model, η_{ijk}, and the unique period-attributed heterogeneity terms ν_{ijk}. The period-attributed heterogeneity terms are, like those for the other two-factor models, simply the predicted values from the parameters for the period nonlinearities (i.e., the excluded variables) using that part of the period variables that is unrelated to the variables for the intercept, age, and cohort terms (i.e., the included variables).

Derivation of Relationships Using Matrix Algebra: Age-Period Model

In this section, we show how the relationships outlined above can be derived using matrix algebra and the logic of omitted variable bias. We illustrate the derivation using the age-period model, but similar calculations can be applied to the period-cohort and age-cohort models. Suppose we fit the age-period model (equation A1) on a data set indexed by $i = 1, \ldots, I$ age rows and $j = 1, \ldots, J$ period columns, with $k = 1, \ldots, K$ cohort diagonals, where $k = j - i + I$ and $K = I + J - 1$. To reveal the underlying structure of the model, it is useful to express equation (A1) as a linearized age-period model with group-level heterogeneity, which decomposes each deviation from the overall mean into its constitutive linear and nonlinear components:

$$Y_{rij} = \mu^* + \alpha^*(i - i^*) + \widetilde{\alpha}^*_i + \pi^*(j - j^*) + \widetilde{\pi}^*_j + \eta^*_{ij} + \epsilon^*_{rij}, \tag{A7}$$

where the parameters are the same as in equation (A1) except that α^* denotes the age slope, $\widetilde{\alpha}^*_i$ the ith age deviation from the overall mean, π^* the period slope, and $\widetilde{\pi}^*_j$ the jth period deviation from the overall mean. Because equation (A7) is the same as equation (A1) but with a different coding scheme, we will refer to them interchangeably as an age-period model in the discussion that follows.

Let $\mathbf{y}$ denote an $R \times 1$ column vector of outcome values, $\mathbf{1}$ an $R \times 1$ column vector of 1's, $\mathbf{A}$ an $R \times (I - 1)$ matrix of age orthogonal polynomials, $\mathbf{P}$ an $R \times (J - 1)$ matrix of period orthogonal polynomials, and $\widetilde{\mathbf{C}}$ an $R \times (K - 2)$ matrix of cohort orthogonal polynomials with no linear component. Using matrix notation, the age-period model can be expressed as follows:

$$\mathbf{y} = \mathbf{1}\mu^* + \mathbf{A}\boldsymbol{\alpha}^* + \mathbf{P}\boldsymbol{\pi}^* + \boldsymbol{\eta}^* + \boldsymbol{\epsilon}^*, \tag{A8}$$

where μ^* is again the intercept, $\boldsymbol{\alpha}^*$ is an $(I - 1) \times 1$ column vector of linear and nonlinear age parameters, $\boldsymbol{\pi}^*$ is a $(J - 1) \times 1$ column vector of linear and nonlinear period parameters, $\boldsymbol{\eta}^*$ is an $R \times 1$ column vector of group-level heterogeneity parameters, and $\boldsymbol{\epsilon}^*$ is an $R \times 1$ column vector of individual-level error terms.

For purposes of comparison, the diff-SC model can be specified in matrix form as follows:

$$\mathbf{y} = \mathbf{1}\mu_y + \mathbf{A}\boldsymbol{\alpha} + \mathbf{P}\boldsymbol{\pi} + \widetilde{\mathbf{C}}\widetilde{\boldsymbol{\gamma}} + \boldsymbol{\eta}_y + \boldsymbol{\epsilon}, \tag{A9}$$

where μ_y is the intercept, $\boldsymbol{\alpha}$ is an $(I - 1) \times 1$ column vector of linear and nonlinear age parameters, $\boldsymbol{\pi}$ is a $(J - 1) \times 1$ column vector of linear and nonlinear period parameters, $\widetilde{\boldsymbol{\gamma}}$ is a $(K - 2) \times 1$ column vector of nonlinear cohort parameters, $\boldsymbol{\eta}_y$ is an $R \times 1$ column vector of cell-specific heterogeneity terms, and $\boldsymbol{\epsilon}$ is an $R \times 1$ column vector of individual-level error terms.

To understand how the age-period model is related to the diff-SC model, we need to specify an auxiliary equation that expresses the association between the variables included in the age-period model and those excluded from the age-period model but included in the diff-SC model. To begin, note that we can define a matrix **S** of dimension $(I+J-1) \times (K-2)$ as follows:

$$\mathbf{S} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\widetilde{\mathbf{C}}, \tag{A10}$$

where $\mathbf{X} = [\mathbf{1}\ \mathbf{A}\ \mathbf{P}]$ is an $R \times (I+J-1)$ matrix of 1's, orthogonal age polynomials, and orthogonal period polynomials; and $\widetilde{\mathbf{C}}$ is again an $R \times (K-2)$ matrix of orthogonal cohort polynomials with no linear component.

The matrix **S** can be interpreted as a collection of parameters representing the relationships between the variables included in the age-period model (**X**) and those variables excluded from the age-period model but included in the diff-SC model ($\widetilde{\mathbf{C}}$). We can define an auxiliary equation compactly as $\widetilde{\mathbf{C}} = \mathbf{XS} + \mathbf{U}_{\tilde{C}}$ or, equivalently:

$$\widetilde{\mathbf{C}} = \mathbf{1s}_{\mu} + \mathbf{AS}_{A} + \mathbf{PS}_{P} + \mathbf{U}_{\tilde{C}}, \tag{A11}$$

where $\mathbf{s}_{\mu}$ is a $1 \times (K-2)$ row vector of parameters, $\mathbf{S}_A$ is an $(I-1) \times (K-2)$ matrix of parameters, $\mathbf{S}_P$ is a $(J-1) \times (K-2)$ matrix of parameters, and $\mathbf{U}_{\tilde{C}}$ is an $R \times (K-2)$ matrix of error terms representing that part of $\widetilde{\mathbf{C}}$ unrelated to the variables included in the age-period model. Note that $\mathbf{s}_{\mu}$ is simply the first row of **S**, $\mathbf{S}_A$ is rows 2 to I of **S**, and $\mathbf{S}_P$ is rows $I+1$ to $I+J-1$ of **S**.

We are now in a position to substitute equation (A11) into equation (A9), which will clarify the meaning of the parameters of equation (A8). After substituting and rearranging terms, we obtain the following equation:

$$\mathbf{y} = \mathbf{1}\underbrace{(\mu_y + \mathbf{s}_{\mu}\tilde{\boldsymbol{\gamma}})}_{\mu^*} + \mathbf{A}\underbrace{(\boldsymbol{\alpha} + \mathbf{S}_A\tilde{\boldsymbol{\gamma}})}_{\boldsymbol{\alpha}^*} + \mathbf{P}\underbrace{(\boldsymbol{\pi} + \mathbf{S}_P\tilde{\boldsymbol{\gamma}})}_{\boldsymbol{\pi}^*} + \underbrace{(\boldsymbol{\eta}_y + \mathbf{U}_{\tilde{C}}\tilde{\boldsymbol{\gamma}})}_{\boldsymbol{\eta}^*} + \underbrace{\boldsymbol{\epsilon}_y}_{\boldsymbol{\epsilon}^*}, \tag{A12}$$

which shows how the diff-SC model is related to the age-period model.[54] As indicated above, the intercept, age, and period parameters from the age-period model will all have some degree of bias because of the exclusion of the cohort nonlinearities. Depending on the structure of the data, the orthogonal cohort polynomials in $\widetilde{\mathbf{C}}$ will be more or less related to the vector **1**, orthogonal age polynomials **A**, and orthogonal period polynomials **P**. If these relationships are strong, then the bias will be large, and the parameter estimates from the age-period and the diff-SC models will differ, possibly quite substantially. If these relationships are weak, then the bias will be relatively small, such that the intercept, age, and period parameters of the age-period model will be approximately equal to those from the diff-SC model. Similarly, the vector of group-level heterogeneity terms $\boldsymbol{\eta}^*$ in the age-period model, which can be interpreted as age-period interactions, is equal to a weighted sum of the cell-specific heterogeneity terms $\boldsymbol{\eta}_y$ and the cohort nonlinearities $\tilde{\boldsymbol{\gamma}}$, with the cohort nonlinearities weighted by $\mathbf{U}_{\tilde{C}}$, or that part of the excluded variables (i.e., the orthogonal cohort polynomials) unrelated to the variables included in the age-period model.

APPENDIX B: SUPPLEMENTAL GRAPHS

Figure B1. Age and cohort nonlinearities.
Note: (a) Age nonlinearities. (b) Cohort nonlinearities. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Intercept is set to zero. Results adjusted using sampling weights.

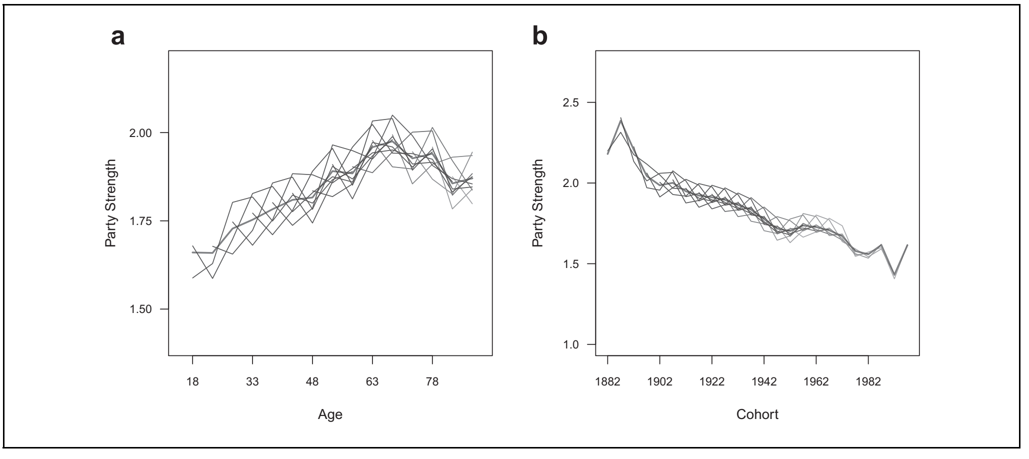

Figure B2. Spaghetti plots of adjusted local LC and SC curves (intra- and intercohort trends).
Note: (a) Estimates of adjusted local life cycle (LC) curves, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$, for $i = 1, \ldots, I$ in each cohort k. Each curve represents an intracohort trend for a particular cohort group. Gray lines denote the local LC curves; the black line indicates the overall LC curve. (b) Estimates of adjusted local social change (SC) curves, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I}$, for cohort groups $k = 1, \ldots, K$ in each age group i. Each curve represents an intercohort trend for a particular age group. Gray lines denote the local SC curves; the black line indicates the overall SC curve. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

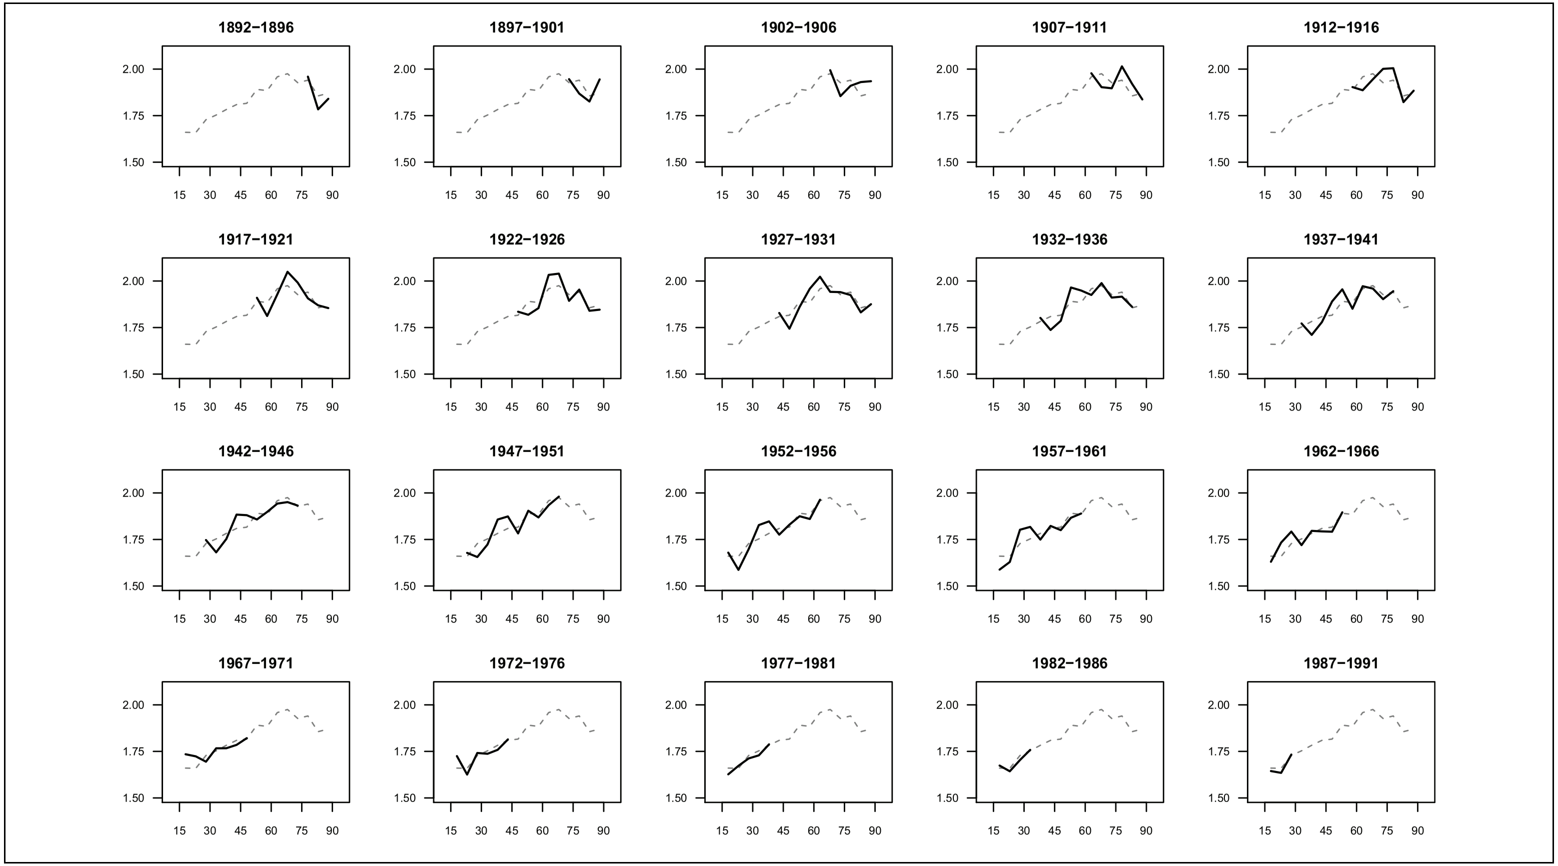

Figure B3. Trellis plot of adjusted local LC curves (intracohort trends).

Note: Each plot displays estimates of the adjusted local life cycle (LC) curve, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$, for $i = 1, \ldots, I$ in a given cohort k, which represents an intracohort trend for a particular cohort group. Solid lines denote adjusted local LC curves; the dashed lines denote the LC curve, which represents an overall intracohort trend. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

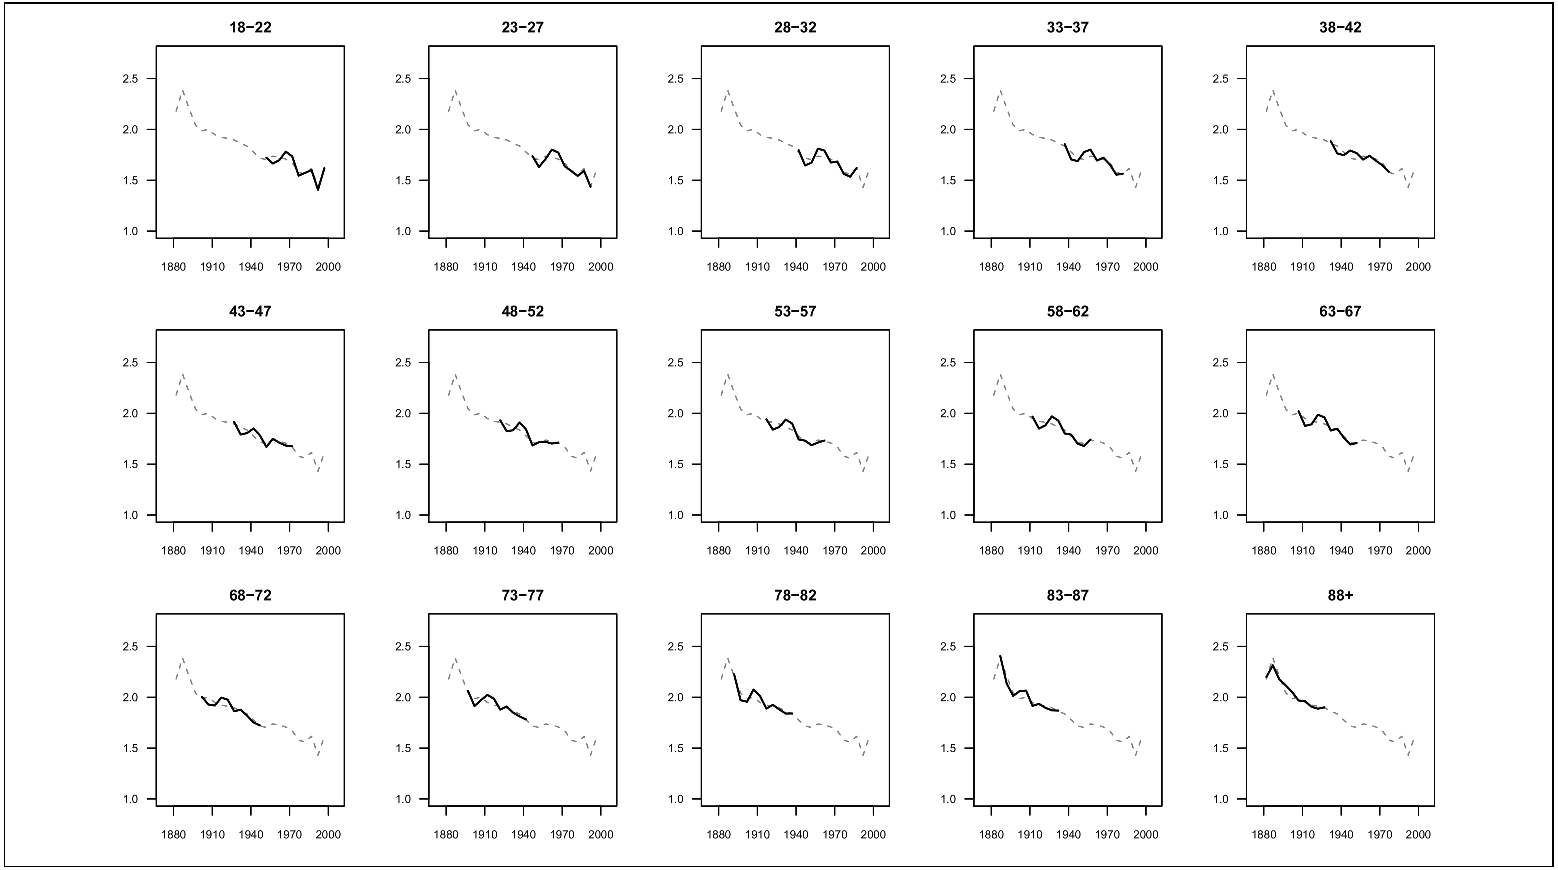

Figure B4. Trellis plot of adjusted local SC curves (intercohort trends).

Note: Each plot displays estimates of the adjusted local social change (SC) curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I}$, for cohort groups $k = 1, \ldots, K$ in a given age group i, which represents an intercohort trend for a particular age group. Solid lines denote unadjusted local SC curves; the dashed lines denote the SC curve, which represents an overall intercohort trend. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

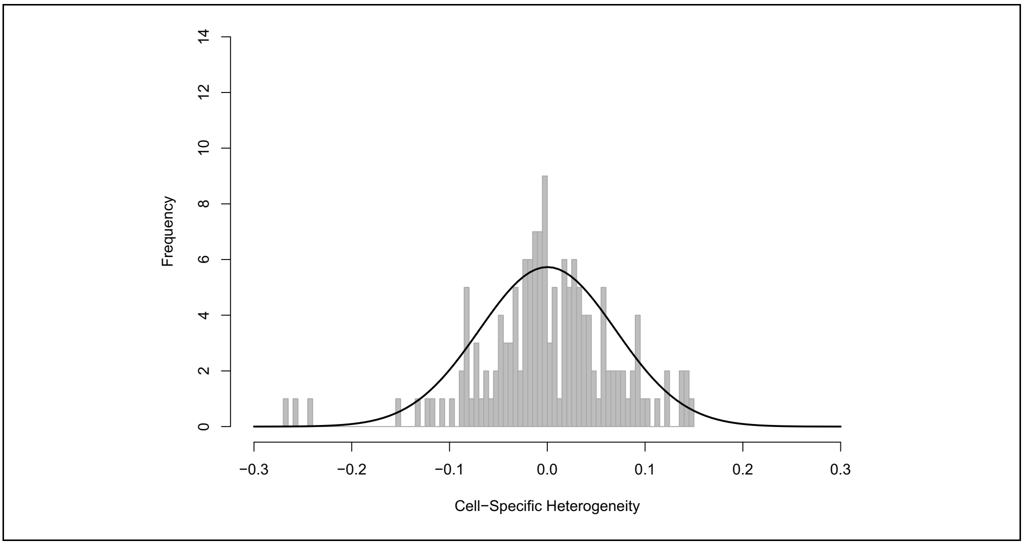

Figure B5. Histogram of cell-specific heterogeneity.

Note: Histogram of cell-specific heterogeneity overlaid with a normal curve. Normal curve is based on the observed mean ($\approx$0) and standard deviation (0.070) of the cell-specific heterogeneity. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights. Three outlying values on the left-hand side of the distribution reflect sparseness of observations among older age groups in earlier cohorts (see also Figure 8 in the main text).

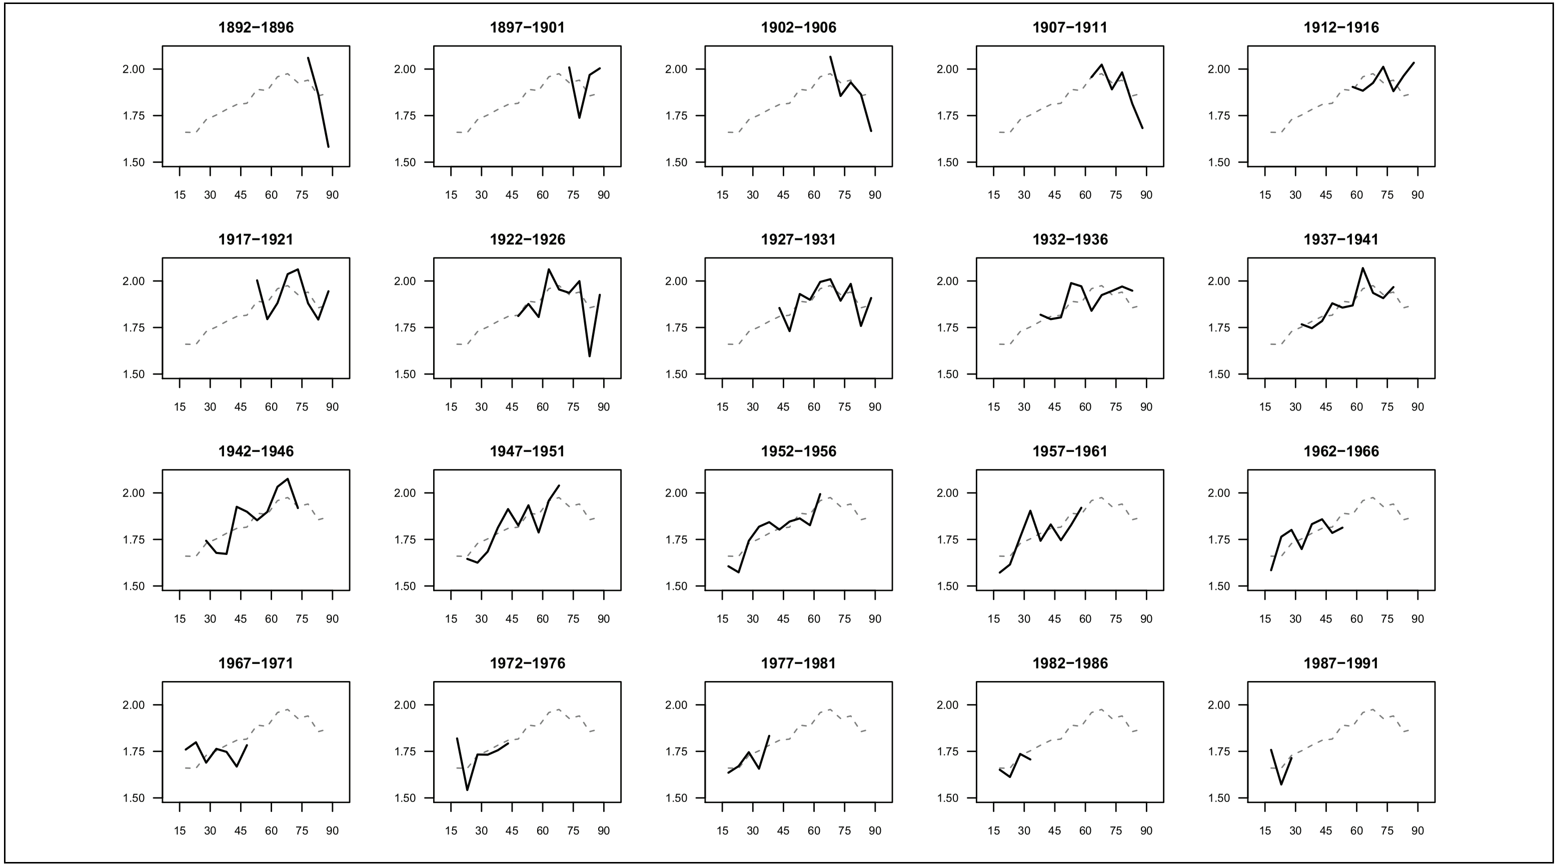

Figure B6. Trellis plot of unadjusted local LC curves (intracohort trends).
Note: Each plot displays estimates of the unadjusted local life cycle (LC) curve, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \eta_{i[i+k-I]k}$, for $i = 1, \ldots, I$ in a given cohort k, which represents an intracohort trend for a particular cohort group. Solid lines denote unadjusted local LC curves; the dashed lines denote the LC curve, which represents an overall intracohort trend. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

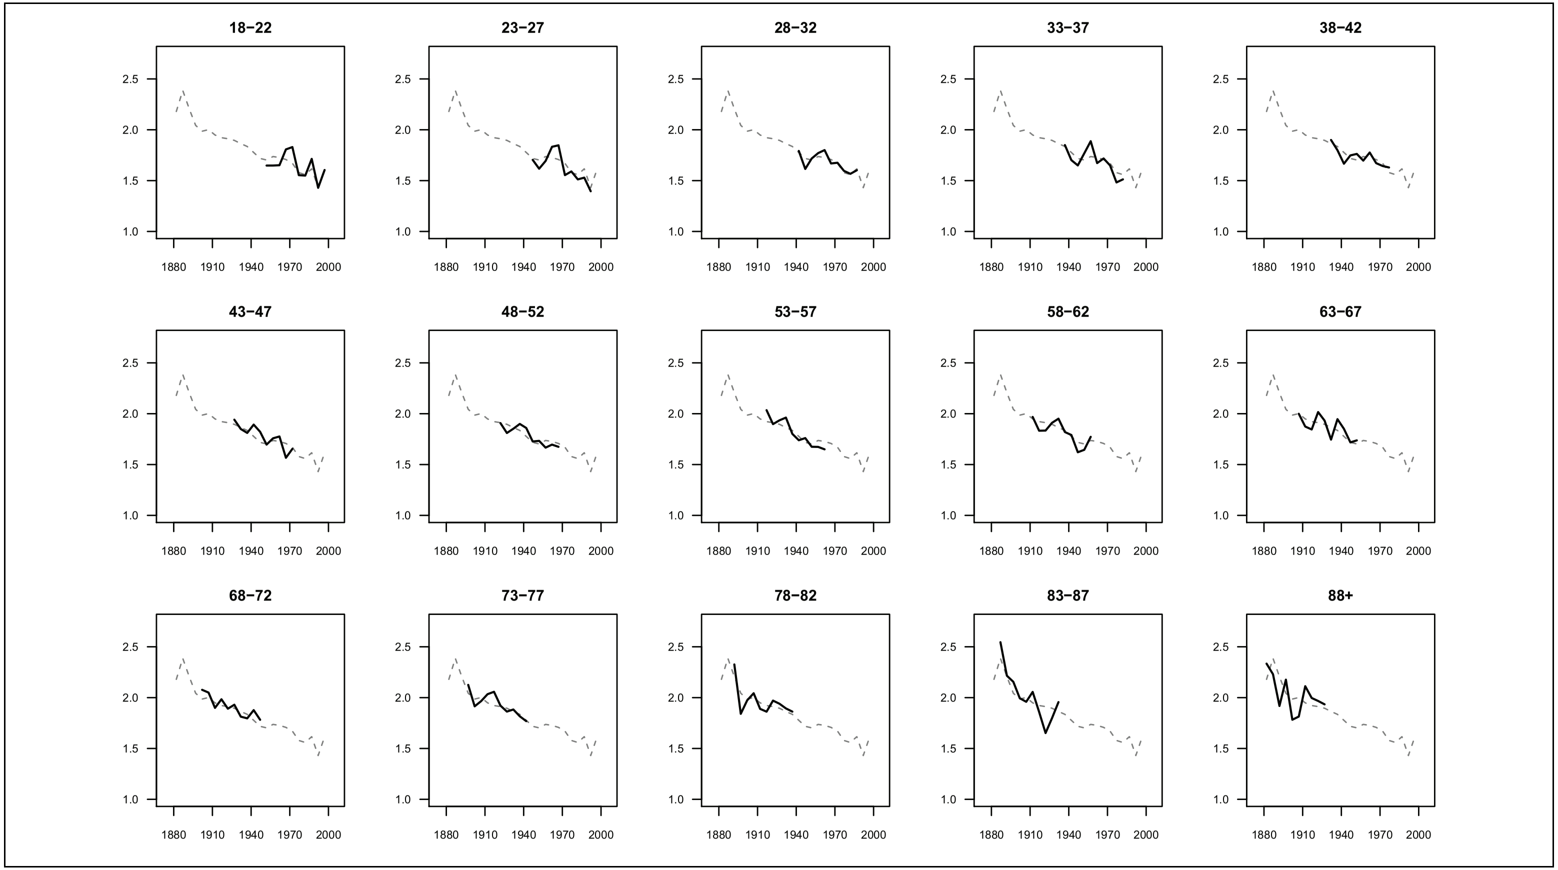

Figure B7. Trellis plot of unadjusted local SC curves (intercohort trends).

Note: Each plot displays estimates of the unadjusted local social change (SC) curve, or $\theta_2(k-k^*)+\widetilde{\gamma}_k+\widetilde{\pi}_{k+i-I}+\eta_{i[i+k-I]k}$, for cohort groups $k=1,\ldots,K$ in a given age group i, which represents an intercohort trend for a particular age group. Solid lines denote unadjusted local SC curves; the dashed lines denote the SC curve, which represents an overall intercohort trend. Outcome is strength of party identification, treated as a continuous variable, with responses coded 0 (independent or apolitical), 1 (leaning independent), 2 (weak partisan), and 3 (strong partisan). Results adjusted using sampling weights.

ORCID iDs

Ethan Fosse https://orcid.org/0000-0002-7901-5234
Christopher Winship https://orcid.org/0000-0002-0632-241X

Notes

1. According to Google Scholar, Ryder's 1965 article had been cited at least 4,286 times as of August 2022. Other articles on cohort analysis are at most cited fewer than 100 times, a moderate number are cited a few hundred times, and only several are cited more than a thousand times; none of these articles approaches 4,000 citations.
2. Unless otherwise stated or implied, we will generally use *calendar time* or *time* synonymously with *period*.
3. In using the term *effects*, the APC literature has conventionally not intended that age, period, and cohort be understood as having direct causal effects. Instead, these variables are typically viewed as surrogate indices for underlying causal factors that actually generate the outcome of interest (Clogg 1982). We follow this convention here.
4. With some exceptions, we generally use *components* and *variables* to refer to columns of the design matrix, and *terms* to refer to model parameters.
5. A similar grouping can be used with the C-APC model, but not only does this model not incorporate additional heterogeneity terms, but the linear and nonlinear parameters for each dimension (age, period, and cohort) will be combined.
6. For example, demographers and epidemiologists have examined disease rates and health-related behaviors, including drug use (O'Malley, Bachman, and Johnston 1984; Vedøy 2014), obesity (Diouf et al. 2010; Fu and Land 2015), cancer (Clayton and Schifflers 1987; Liu et al. 2001), and mental health (Lavori et al. 1987). Sociologists, political scientists, and other social scientists have similarly been concerned with a range of outcomes related to understanding social change, including verbal ability (Alwin 1991; Hauser and Huang 1997; Wilson and Gove 1999), social trust (Clark and Eisenstein 2013; Putnam 1995; Robinson and Jackson 2001; Schwadel and Stout 2012), party identification (Knoke and Hout 1974; Tilley and Evans 2014), and religious affiliation (Chaves 1989; Firebaugh and Harley 1991).
7. For a similar statement regarding how scientific knowledge changes, see Planck (1949:33, 97). We are indebted to Herb Smith for this point.
8. Google Scholar indicates 51 citations as of August 2022.
9. For clarity, we altered his notation to be consistent with the current APC literature.
10. Given the importance of age in examining fertility, Ryder's substantive area of interest, he took it for granted that one would examine variation across age.
11. Suppose, for example, that yearly data are organized into an age-period array. A single cohort section can take upward of 80 years to collect, whereas a period section can be collected in a single year.
12. Recent research generally confirms Ryder's views on this point, supporting a "settled dispositions" model rather than an "automatic updating" model of individual-level change (Vaisey and Kiley 2021; see also Corning and Schuman 2015).
13. Reflecting the convention in the APC literature, we focus on analyzing the repeated cross-sections of the GSS (Davern et al. 2021). However, our approach extends straightforwardly to panel data in which the same individuals are tracked through time (i.e., across periods). Intuitively, the reason for this is that our methods are based on the cell means of a Lexis table created by the cross-classification of age by period, and such a table can be easily constructed using panel data. However, when using panel data, one would want to take account of the intraindividual correlation across observations in calculating standard errors and carrying out statistical tests.
14. Some years have a small handful of "don't know" responses. These observations are dropped from our analysis.
15. Substantively similar results are obtained treating party strength as an ordered categorical variable.

16. Examples of such age-cohort displays in the APC literature date at least as far back as Andvord (1930) and Frost (1939). For more recent examples, see Figure 1.2 in O'Brien (2015) and Figure 3 in Robinson and Jackson (2001).
17. Reflecting the fact that cohort is calculated from age and period, we will assume the data are organized into an age-period array. For simplicity, we also assume the age and period categories are of equal width.
18. Note that I is added to $j-i$ so the cohort index begins at $k=1$. This ensures that, for example, $i=j=k=1$ refers to the first group for all three temporal scales. One could just as easily index the cohorts using $k=j-i$, but this identity would be lost.
19. Throughout, we will also refer to this widely used three-factor APC model as the "conventional," "classical," or "classic" APC model.
20. However, as Smith (2021) noted, if I, J, and/or K are odd, then under conventional "normalization" assumptions, the corresponding middle parameters $\alpha_{(I+1)/2}$, $\pi_{(J+1)/2}$, and/or $\gamma_{(K+1)/2}$ will be identified.
21. To understand how this is calculated, note that the C-APC model (with, say, a zero period slope constraint) takes up $1+(I-1)+(J-2)+(K-1)=1+(I-1)+(J-2)+((I+J-1)-1)$ or $2(I+J)-4$ degrees of freedom. However, there are $I \times J$ degrees of freedom in total. Thus, there are $IJ-(2I+2J-4)=IJ-2I-2J+4=(I-2)(J-2)$ additional degrees of freedom beyond the C-APC model.
22. As with the C-APCH model, with the L-APCH model each cell in an age-period array is a function of a unique combination of parameters. For example, the ith age parameter in the C-APCH model is represented in the L-APCH model by the overall age slope along with a unique parameter for the ith age nonlinearity: $\alpha_i=(i-i^*)\alpha+\widetilde{\alpha}_i$.
23. That is, by assuming $\pi^*=\pi=0$, using simple algebra one can show that $\nu=\pi=0$ and accordingly that $\alpha^*=\alpha+\pi=\alpha$ and $\gamma^*=\gamma+\pi=\gamma$. This is just stating that assuming the value of any one slope determines the values of the other two.
24. This is based on data collected directly on age and period, with cohort calculated as a function of age and period.
25. One could apply a similar substitution using the C-APC or C-APCH model. These models can likewise be used to represent intra- and intercohort trends (i.e., life cycle and social change). However, the advantage of the L-APCH model over the C-APC and C-APCH models is that, by separating the linear from the nonlinear parameters, we can present a more fine-grained decomposition of cohort careers.
26. This substitution is analogous to substituting (age + cohort) for period in a more basic slopes-only linear regression model: $Y=\mu+\alpha\text{Age}+\pi\text{Period}+\gamma\text{Cohort}+\epsilon=\mu+(\alpha+\pi)\text{Age}+(\gamma+\pi)\text{Cohort}+\epsilon$. As a result, after substitution, the period slope, π, is combined with *both* the age and cohort slopes, α and γ, respectively.
27. We would like to thank the reviewers of an earlier version of this article for pointing out the similarities between the LC-SC and ZPS models and challenging us to explain the ways in which they differ. This was very helpful in advancing our own thinking about the relationship between cohort analysis and APC analysis.
28. This means the columns of the design matrices have different labels and the parameters, specifically the linear terms, have different meanings.
29. We have two simple linear equations, $\theta_1=\alpha+\pi$ and $\theta_2=\gamma+\pi$. There are, however, three unknown APC parameters (α, π, and γ). As such, any assumption or knowledge about one of the three APC slope terms determines the values of the other two. Fosse and Winship (2018, 2019a, 2019b) showed how the canonical solution line can be graphed into what they called a two-dimensional APC graph, where the two θ's determine the relative position of the vertical age and cohort axes with respect to each other, as well as the intercept for the canonical solution line.
30. Of course, the assumption that $\pi=0$ invoked by the ZPS model could very well be wrong. Holford (1983) described the consequence of this (notation altered slightly for clarity): "One possibility is to assume that period is linear, $\pi=0$. In this case the linear age is actually an estimate of $\alpha+\pi$, while

the linear cohort is an estimate of $\gamma + \pi$. Hence, the linear age and cohort parameters are biased by whatever the true value of π happens to be" (p. 316).

31. That is, for any particular assumption about the "true" slopes under a just-identifying constraint, it is always the case that $\theta_1 = \alpha^* + \pi^* = \alpha + \pi$ and $\theta_2 = \gamma^* + \pi^* = \gamma + \pi$. This is because the ν's cancel out. For example, $\alpha^* + \pi^* = (\alpha + \nu) + (\pi - \nu) = \alpha + \pi = \theta_1$.
32. To reinforce the distinction between synchronic and diachronic measures, we will not use the terms "trends" or "change" to refer to any synchronic quantity. The reason should be clear: any trend or change must occur *through* calendar time, not within a cross-section of calendar time.
33. That is, $\theta_1 - \theta_2 = \theta_1$ only if $\theta_2 = 0$, and $\theta_2 - \theta_1 = \theta_2$ only if $\theta_1 = 0$.
34. Relevant to this discussion is the debate between Firebaugh (1989) and Rodgers (1990). Firebaugh (1989:244) argued that a slopes-only period-cohort model estimates a "cohort replacement effect," whereas Rodgers stated that it combines a cohort linear effect and an age linear effect—in our notation, it is estimating $\theta_2 - \theta_1$, that is, a synchronic (cohort) slope representing a set of intraperiod differences. Firebaugh (1997) acknowledged Rodgers's point in later work, noting it is "important not to equate the cross-cohort slope" in a period-cohort model "with a cohort *effect*," because the slope "could reflect age effects as well as cohort effects" (p. 66). See also Firebaugh (1990, 2008).
35. More precisely, estimates of the age and cohort parameters will be inconsistent because the influence of the period parameters on the outcome have not been partialed out. In other words, one's estimates of the age and cohort parameters will differ from their true values even in a sample of infinite size. See Appendix A for a detailed discussion of this issue.
36. One may be tempted to claim that the heterogeneity estimated from the specified age-cohort interactions is identical to that arising from the period nonlinearities, but this requires assuming that there is no additional heterogeneity beyond that attributable to the main parameters for age, period, and cohort. This is an unnecessarily strong assumption that can be tested against the data.
37. For a novel decomposition of an APC model, see Smith (2021).
38. For simplicity of presentation, we exclude the intercept from the definitions of the various parametric expressions in the following sections. However, unless otherwise stated, all visualizations in this article include the intercept term.
39. Note, however, that θ_1 is the same for all cohort levels $k = 1, \ldots, K$, and θ_2 is the same for all age groups $i = 1, \ldots, I$.
40. Because of space constraints, the age and cohort nonlinearities used to construct the LC and SC curves are shown in Figure B1 in Appendix B, with upper and lower lines indicating 95 percent confidence intervals. Notwithstanding some imprecision, we find substantial nonlinear variation for both age and cohort. As discussed by O'Brien (2015), Fosse and Winship (2019a), and Smith (2021), the presence and statistical significance of substantial nonlinearities is testable against the data. Caution is needed, however, in assessing nonlinearities in the tails due to sparseness. This is particularly the case for cohort, which is unbalanced. As shown in Figure B1b, an examination of the 95 percent confidence intervals reveals considerable imprecision in the estimation of the cohort nonlinearities at the two end points.
41. Converse did not posit a decline in party strength over the life cycle but rather a flattening of the curve in older adulthood. However, the decline in Figure 4a could be due to differential mortality by age.
42. In contrast to Figure 3, the vertical distance between any two cohorts is allowed to differ because of the inclusion of the cohort nonlinearities.
43. In Appendix B (see Figures B2, B3, and B4), we display spaghetti and trellis plots of the adjusted local LC and SC curves.
44. By contrast, the curves-only comparative cohort careers in equation (16) are composed of the LC and SC curves. These curves are the same for each cohort and age group, except that most cohorts will experience only a section of the overall LC curve. This simply reflects the unbalanced nature of the cohorts when data are collected based on age and period, as is the case with most time-series cross-sectional data, such as the GSS.

45. To avoid visual clutter due to overlapping period nonlinearities, in Figure 6b we opted to display the adjusted local SC curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{i+k-I}$, for $i = 1, \ldots, I$ in each selected cohort k. This effectively spreads out the period nonlinearities vertically so the line plots are not on top of each other. The variation within cohorts is still entirely given by the period nonlinearities, and the (vertical) variation across cohorts is governed by all the terms in the adjusted local SC curve.
46. Interactions beyond the main parameters of an APC model are potentially quite important. Yet, their specification and interpretation require considerable care. This is in part because such "pure" interactions are easily conflated with the main parameters of an APC model (cf. Rosnow and Rosenthal 1995). For example, as we illustrated in the previous section, the main parameters for period appear to "interact" with respect to age and cohort. (A similar pattern occurs for the cohort main parameters with respect to age and period, as well as the age main parameters with respect to cohort and period.) An additional complication is that there are limited degrees of freedom beyond the main parameters for age, period, and cohort, so only a restricted set of interactions can be parameterized (for a discussion, see Mason and Fienberg 1985). Finally, there are important statistical issues as to how to test for interactions, particularly in large data sets in which conventional tests almost always indicate statistically significant effects at conventional levels. Fortunately, fit measures such as the Akaike information criterion and Schwarz's Bayesian information criterion, as well the use of test data sets and cross-validation, provide appropriate ways to deal with this (see, e.g., Weakliem 2016). For additional issues with interpreting interactions more generally, see Hainmueller, Mummolo, and Xu (2019).
47. For simplicity of presentation, as previously mentioned, we excluded the intercept from the parametric expressions discussed in this section.
48. For the purposes of comparing fit, to represent the LC-SC model with terms for cell-specific heterogeneity we specify a fully saturated model, namely, a set of dummy variables for all cells of an age-period array, which takes up $I \times J = 150$ degrees of freedom. We then compare this to a reduced model with just the intercept and the main parameters for age, period, and cohort, which takes up $2(I - J) - 4 = 2(15 + 10) - 4 = 46$ degrees of freedom.
49. If our goal were to provide a more parsimonious representation of the data, then we would have focused on a simpler model. The AIC suggests that one should fit a model that includes all terms except those representing the additional cell-specific heterogeneity, whereas the adjusted BIC and BIC suggest that one should fit a model that includes just the linear components and the period nonlinearities.
50. A general issue in any APC analysis will be what constitutes an adequate model of the data. One can certainly use traditional statistical tests based on classic inference theory, such as the t test or F test. However, with large samples, typical of nonaggregated APC data, such tests will almost always indicate that the only acceptable model is that which includes the full set of parameters. In recent years, researchers have often used other model criteria, such as AIC, BIC, or the adjusted BIC (see Burnham and Anderson 2004; Weakliem 2016). These criteria provide heavier penalties for model complexity, that is, the number of parameters that are included in a model. The AIC, with a penalty factor of $2q$, where q is the number of parameters, provides the weakest penalty. The BIC, with a penalty function of $\log(R)q$, provides the strongest penalty, where R is the sample size. The adjusted BIC's penalty falls between that of the AIC and BIC. Unfortunately, there is no one best model fit criterion. Rather, researchers need to decide what is the right balance in their analysis between capturing the detailed richness of their data versus achieving descriptive parsimony.
51. The parameters of a two-factor model can be interpreted with reference to terms for all three factors, thereby clarifying the meaning of the parameters in the two-factor model (see Appendix A).
52. If the data are balanced such that there are an equal number of observations in each age-period cell, then, using sum-to-zero deviation coding or orthogonal polynomial coding, the residuals will be equivalent to specifying a full set of age-period interactions. The reason for this is that in such a setting, the columns for the age-period interactions will be orthogonal to the main age and period columns.

53. The individual-level ϵ_{rij} error term is unbiased. This is because the age-period model is saturated, so the individual-level error will be the same as that from the diff-SC model, which is also saturated.
54. Equation (A12) looks complicated, but it is just a multivariate extension of the standard formula for omitted variables bias. That is, in a regression equation with variables X_1 and X_2, where X_2 is omitted, $\beta_1 = \beta_1 + S_{21}\beta_2$ where $S_{21} = (X'_1X_1)^{-1}X'_1X_2$ or, equivalently, the regression of X_2 on X_1.

References

Abramson, Paul R. 1976. "Generational Change and the Decline of Party Identification in America: 1952–1974." *American Political Science Review* 70:469–78.

Alwin, Duane F. 1991. "Family of Origin and Cohort Differences in Verbal Ability." *American Sociological Review* 56:625–38.

Andvord, Kr. F. 1930. "Der Verlauf der Tuberkulose durch Generationen." *Lung* 75:552–63.

Bell, Andrew. 2014. "Life-Course and Cohort Trajectories of Mental Health in the UK, 1991– 2008: A Multilevel Age–Period–Cohort Analysis." *Social Science & Medicine* 120(1):21–30.

Bell, Andrew, and Kelvyn Jones. 2015. "Age, Period and Cohort Processes in Longitudinal and Life Course Analysis: A Multilevel Perspective." Pp. 197–213 in *A Life Course Perspective on Health Trajectories and Transitions*, edited by C. Burton-Jeangros, S. Cullati, A. Sacker, and D. Blane. Cham, Switzerland: Springer.

Bell, Andrew, and Kelvyn Jones. 2021. "Multilevel Models for Age–Period–Cohort Analysis." Pp. 23–40 in *Age, Period and Cohort Effects: Statistical Analysis and the Identification Problem*, edited by A. Bell. New York: Routledge.

Bernardi, Laura, Johannes Huinink, and Richard A. Settersten. 2019. "The Life Course Cube: A Tool for Studying Lives." *Advances in Life Course Research* 41:100258.

Burnham, Kenneth P., and David R. Anderson. 2004. "Multimodel Inference: Understanding AIC and BIC in Model Selection." *Sociological Methods & Research* 33(2):261–304.

Chaves, Mark. 1989. "Secularization and Religious Revival: Evidence from U.S. Church Attendance Rates, 1972–1986." *Journal for the Scientific Study of Religion* 28(4):464–77.

Clark, April K., and Marie A. Eisenstein. 2013. "Interpersonal Trust: An Age–Period–Cohort Analysis Revisited." *Social Science Research* 42(2):361–75.

Clayton, D., and E. Schifflers. 1987. "Models for Temporal Variation in Cancer Rates. II: Age–Period–Cohort Models." *Statistics in Medicine* 6(4):469–81.

Clogg, Clifford C. 1982. "Cohort Analysis of Recent Trends in Labor Force Participation." *Demography* 19(4):459–79.

Converse, Philip E. 1976. *The Dynamics of Party Support: Cohort-Analyzing Party Identification*. Beverly Hills, CA: Sage.

Corning, Amy, and Howard Schuman. 2015. *Generations and Collective Memory*. Chicago: University of Chicago Press.

Davern, Michael, Rene Bautista, Jeremy Freese, Stephen L. Morgan, and Tom W. Smith. 2021. "General Social Survey Cross-Section." Technical report (machine-readable data file). Chicago: NORC.

Davis, James A. 1992. "Changeable Weather in a Cooling Climate atop the Liberal Plateau: Conversion and Replacement in Forty-Two General Social Survey Items, 1972–1989." *Public Opinion Quarterly* 56:261–306.

Diouf, Ibrahima, Marie Aline Charles, Pierre Ducimetière, Arnaud Basdevant, Evelyne Eschwege, and Barbara Heude. 2010. "Evolution of Obesity Prevalence in France: An Age-Period-Cohort Analysis." *Epidemiology* 21(3):360–65.

Elder, Glen H. 1974. *Children of the Great Depression: Social Change in Life Experience*. Chicago: University of Chicago Press.

Firebaugh, Glenn. 1989. "Methods for Estimating Cohort Replacement Effects." *Sociological Methodology* 19:243–62.

Firebaugh, Glenn. 1990. "Replacement Effects, Cohort and Otherwise: Response to Rodgers." *Sociological Methodology* 20:439–46.

Firebaugh, Glenn. 1997. *Analyzing Repeated Surveys*. Newbury Park, CA: Sage.

Firebaugh, Glenn. 2008. *Seven Rules for Social Research*. Princeton, NJ: Princeton University Press.

Firebaugh, Glenn, and Brian Harley. 1991. "Trends in U.S. Church Attendance: Secularization and Revival, or Merely Lifecycle Effects?" *Journal for the Scientific Study of Religion* 30(4):487–500.

Fosse, Ethan, and Christopher Winship. 2018. "Moore–Penrose Estimators of Age–Period–Cohort Effects: Their Interrelationship and Properties." *Sociological Science* 5(14):304–34.

Fosse, Ethan, and Christopher Winship. 2019a. "Analyzing Age-Period-Cohort Data: A Review and Critique." *Annual Review of Sociology* 45:467–92.

Fosse, Ethan, and Christopher Winship. 2019b. "Bounding Analyses of Age-Period-Cohort Effects." *Demography* 56(5):1975–2004.

Frost, Wade Hampton. 1939. "The Age Selection of Mortality from Tuberculosis in Successive Decades." *American Journal of Hygiene* 30(Section A):91–96.

Fu, Qiang, and Kenneth C. Land. 2015. "The Increasing Prevalence of Overweight and Obesity of Children and Youth in China, 1989–2009: An Age–Period–Cohort Analysis." *Population Research and Policy Review* 34(6):901–21.

Glenn, Norval D. 1977. *Cohort Analysis*. Beverly Hills, CA: Sage.

Hainmueller, Jens, Jonathan Mummolo, and Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis* 27(2):163–92.

Hardy, Melissa, and Andrea E. Wilson. 2002. "Cohort Change." Pp. 241–45 in *Encyclopedia of Aging*, Vol. 1, edited by D. J. Ekerdt. New York: Macmillan Reference.

Hauser, Robert M., and Min-Hsiung Huang. 1997. "Verbal Ability and Socioeconomic Success: A Trend Analysis." *Social Science Research* 26(3):331–76.

Holford, Theodore R. 1983. "The Estimation of Age, Period and Cohort Effects for Vital Rates." *Biometrics* 39(2):311–24.

Hout, Michael. 2021. "America's Liberal Social Climate and Trends: Change in 283 General Social Survey Variables between and within US Birth Cohorts, 1972–2018." *Public Opinion Quarterly* 85(3): 1009–49.

Jager, Justin, Katherine M. Keyes, Daye Son, Megan E. Patrick, Jonathan Platt, and John E. Schulenberg. 2022. "Age 18–30 Trajectories of Binge Drinking Frequency and Prevalence across the Past 30 Years for Men and Women: Delineating When and Why Historical Trends Reversed across Age." *Development and Psychopathology*. doi:10.1017/S0954579421001218.

Keyes, Katherine M., and Guohua Li. 2010. "A Multiphase Method for Estimating Cohort Effects in Age-Period Contingency Table Data." *Annals of Epidemiology* 20(10):779–85.

Knoke, David, and Michael Hout. 1974. "Social and Demographic Factors in American Political Party Affiliations, 1952–72." *American Sociological Review* 39(5):700–13.

Lavori, Philip W., Gerald L. Klerman, Martin B. Keller, Theodore Reich, John Rice, and Jean Endicott. 1987. "Age-Period-Cohort Analysis of Secular Trends in Onset of Major Depression: Findings in Siblings of Patients with Major Affective Disorder." *Journal of Psychiatric Research* 21(1):23–35.

Liu, S., R. Semenciw, A. M. Ugnat, and Y. Mao. 2001. "Increasing Thyroid Cancer Incidence in Canada, 1970–1996: Time Trends and Age-Period-Cohort Effects." *British Journal of Cancer* 85(9):1335–39.

Lundberg, Ian, Rebecca Johnson, and Brandon M. Stewart. 2021. "What Is Your Estimand? Defining the Target Quantity Connects Statistical Evidence to Theory." *American Sociological Review* 86(3): 532–65.

Luo, Liying, and James S. Hodges. 2020. "The Age-Period-Cohort-Interaction Model for Describing and Investigating Inter-cohort Deviations and Intra-cohort Life-Course Dynamics." *Sociological Methods & Research* 51(3):1164–210.

Mannheim, Karl. [1927/1928] 1952. "The Problem of Generations." Pp. 276–322 in *Karl Mannheim: Essays*, edited by P. Kecskemeti. London: Routledge.

Mason, Karen Oppenheim, William M. Mason, H. H. Winsborough, and W. Kenneth Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38(2): 242–58.

Mason, William M., and Fienberg, Stephen E., eds. 1985. *Cohort Analysis in Social Research: Beyond the Identification Problem*. New York: Springer.

Mayer, Karl Ulrich, and Johannes Huinink. 1990. "Age, Period, and Cohort in the Life Course: A Comparison of Classical A-P-C Analysis with Event History Analysis or Farewell to Lexis?" Pp. 211–32 in *Data Quality in Longitudinal Research*, edited by D. Magnusson and L. R. Bergman. New York: Cambridge University Press.

Neil, Roland, and Robert J. Sampson. 2021. "The Birth Lottery of History: Arrest over the Life Course of Multiple Cohorts Coming of Age, 1995–2018." *American Journal of Sociology* 126(5):1127–78.

O'Brien, Robert M. 2011. "Constrained Estimators and Age-Period-Cohort Models." *Sociological Methods & Research* 40(3):419–52.

O'Brien, Robert. 2015. *Age-Period-Cohort Models: Approaches and Analyses with Aggregate Data*. Boca Raton, FL: CRC.

Ohtaki, Megu, Dong-Kyu Kim, and Masaki Munaka. 1990. "A Nonparametric Method for Estimating Interaction Effect of Age and Period on Mortality." *Environmental Health Perspectives* 87:115–21.

O'Malley, Patrick M., Jerald G. Bachman, and Lloyd D. Johnston. 1984. "Period, Age, and Cohort Effects on Substance Use among American Youth, 1976–82." *American Journal of Public Health* 74(7): 682–88.

Planck, Max. 1949. *Scientific Autobiography and Other Papers*. New York: Philosophical Library.

Putnam, Robert D. 1995. "Tuning In, Tuning Out: The Strange Disappearance of Social Capital in America." *PS: Political Science and Politics* 28(4):664–83.

Robinson, Robert V., and Elton F. Jackson. 2001. "Is Trust in Others Declining in America? An Age–Period–Cohort Analysis." *Social Science Research* 30(1):117–45.

Rodgers, Willard L. 1990. "Interpreting the Components of Time Trends." *Sociological Methodology* 20: 421–38.

Rosnow, Ralph L., and Robert Rosenthal. 1995. "'Some Things You Learn Aren't So': Cohen's Paradox, Asch's Paradigm, and the Interpretation of Interaction." *Psychological Science* 6(1):3–9.

Ryder, Norman B. 1964. "The Process of Demographic Translation." *Demography* 1(1):74–82.

Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30(6):843–61.

Ryder, Norman B. 1968. "Cohort Analysis." Pp. 546–50 in *International Encyclopedia of the Social Sciences*, Vol. 2, edited by D. L. Sills. New York: Macmillan.

Ryder, Norman B. 1972. October 18. "Letter from Norman B. Ryder to Karen Oppenheim Mason."

Ryder, Norman B. 1979. "Commentary on Cohorts, Periods, and Ages." Unpublished manuscript, Office of Population Research, Princeton University.

Ryder, Norman B. 1992. "Cohort Analysis." Pp. 227–31 in *Encyclopedia of Sociology*, edited by E. F. Borgatta and M. L. Borgatta. New York: Macmillan.

Schulhofer-Wohl, Sam, and Yang Yang. 2016. "Modeling the Evolution of Age and Cohort Effects." Pp. 313–35 in *Dynamic Demographic Analysis*, edited by R. Schoen. New York: Springer International.

Schwadel, P., and M. Stout. 2012. "Age, Period and Cohort Effects on Social Capital." *Social Forces* 91(1):233–52.

Shen, Yinzhi, Shawn D. Bushway, Lucy C. Sorensen, and Herbert L. Smith. 2020. "Locking Up My Generation: Cohort Differences in Prison Spells over the Life Course." *Criminology* 58(4):645–77.

Smith, Herbert L. 2021. "Age-Period-Cohort Analysis: What Is It Good For?" Pp. 176–205 in *Age, Period and Cohort Effects: Statistical Analysis and the Identification Problem*, edited by A. Bell. New York: Routledge.

Tilley, James, and Geoffrey Evans. 2014. "Ageing and Generational Effects on Vote Choice: Combining Cross-Sectional and Panel Data to Estimate APC Effects." *Electoral Studies* 33(1):19–27.

Tonda, Tetsuji, Kenichi Satoh, and Ken-ichi Kamo. 2015. "Detecting a Local Cohort Effect for Cancer Mortality Data Using a Varying Coefficient Model." *Journal of Epidemiology* 25(10):639–46.

Vaisey, Stephen, and Kevin Kiley. 2021. "A Model-Based Method for Detecting Persistent Cultural Change Using Panel Data." *Sociological Science* 8(5):83–95.

Vedøy, Tord F. 2014. “Tracing the Cigarette Epidemic: An Age-Period-Cohort Study of Education, Gender and Smoking Using a Pseudo-panel Approach.” *Social Science Research* 48(1):35–47.

Voas, David, and Mark Chaves. 2016. “Is the United States a Counterexample to the Secularization Thesis?” *American Journal of Sociology* 121(5):1517–56.

Weakliem, David L. 2016. *Hypothesis Testing and Model Selection in the Social Sciences*. New York: Guilford.

Wilson, James A., and Walter R. Gove. 1999. “The Intercohort Decline in Verbal Ability: Does It Exist?” *American Sociological Review* 64(2):253–66.

Yang, Yang, and Kenneth C. Land. 2013. *Age-Period-Cohort Analysis: New Models, Methods, and Empirical Applications*. Boca Raton, FL: CRC.

Author Biographies

Ethan Fosse is an Assistant Professor of Sociology and Associate Director of the Data Sciences Institute at the University of Toronto. He received his PhD from Harvard University and previously worked as a Postdoctoral Research Associate at Princeton University. He has published research on social change, computational and quantitative methods, and cohort analysis in a number of journals, including *Demography*, *Sociological Science*, and the *Annual Review of Sociology*. He is currently working on a book with Cambridge University Press (coauthored with Christopher Winship) on methods for analyzing APC data.

Christopher Winship is the Diker-Tishman Professor of Sociology at Harvard University. Prior to coming to Harvard in 1992, he was professor of sociology, statistics, and economics at Northwestern University. He received a PhD in sociology from Harvard in 1977 and a BA in sociology and mathematics from Dartmouth in 1972. Along with Steve Morgan, he is the author of *Counterfactuals and Causal Inference* (2007, 2015), which has sold more than 30,000 copies. With Neil Gross and Isaac Reed, he is coeditor of *The New Pragmatist Sociology* (2022). He has also written extensively about community policing in Boston.