sociological science

Dissecting the Lexis Table: Summarizing Population-Level Temporal Variability with Age–Period–Cohort Data

Ethan Fosse

University of Toronto

Abstract: Since Norman Ryder's (1965) classic essay on cohort analysis was published more than a half century ago, scores of researchers have attempted to uncover the separate effects of age, period, and cohort (APC) on a wide range of outcomes. However, rather than disentangling period effects from those attributable to age or cohort, Ryder's approach is based on distinguishing intra-cohort trends (or life-cycle change) from inter-cohort trends (or social change), which, together, constitute comparative cohort careers. Following Ryder's insights, in this article I show how to formally summarize population-level temporal variability on the Lexis table. In doing so, I present a number of parametric expressions representing intra- and inter-cohort trends, intra-period differences, and Ryderian comparative cohort careers. To aid the interpretation of results, I additionally introduce a suite of novel visualizations of these model-based summaries, including 2D and 3D Lexis heat maps. Crucially, the Ryderian approach developed in this article is fully identified, complementing (but not replacing) conventional approaches that rely on theoretical assumptions to parse out unique APC effects from unidentified models. This has the potential to provide a common base of knowledge in a literature often fraught with controversy. To illustrate, I analyze trends in social trust in the U.S. General Social Survey from 1972 to 2018.

Keywords: Lexis table; cohort analysis; Norman Ryder; social change; life-cycle change

For at least a century, social scientists in a wide range of fields have sought to understand social change using age–period–cohort (APC) analysis. As it is conventionally understood, the objective of an APC analysis is to use time-series cross-sectional data to identify the distinct contributions of age, period (or survey year), and cohort (or birth year) on some outcome of interest.[1] Although the cohort concept has ancient roots,[2] Ryder's (1965) classic essay on the relationship between the cohort concept and social change inaugurated the modern era of APC analysis in sociology and demography.

Drawing from earlier work by historians and demographers, Ryder (1965) outlined two main aspects of what he labeled "the cohort approach" (P. 549): first, the study of intra-cohort trends, or "intra-cohort temporal development"; second, the examination of inter-cohort trends, or "intercohort temporal differentiation" (P. 861). Because they involve comparisons within cohorts as they age through time (i.e., across periods), intra-cohort trends represents life-cycle change; by contrast, because they entail comparisons of successive cohorts through time (i.e., across periods), inter-cohort trends represent social change.[3] Together, intra- and inter-cohort trends, or, equivalently, life-cycle and social change, constitute what Ryder (1965) famously called "comparative cohort careers," or cohort-specific age–time trajectories (P. 861).

Citation: Fosse, Ethan. 2023. "Dissecting the Lexis Table: Summarizing Population-Level Temporal Variability with Age–Period–Cohort Data." Sociological Science 10: 150-196.

Received: April 7, 2022

Accepted: May 9, 2022

Published: March 13, 2023

Editor(s): Arnout van de Rijt, Richard Breen

DOI: 10.15195/v10.a5

In short, then, Ryder's cohort approach is essentially the depiction of life-cycle and social change, with cohort as a general unifying, analytic concept (see also Ryder 1968, 1992).

Ryder's essay is now considered a citation classic in sociology and demography, and his insights have inspired generations of researchers to take up APC analysis, including many of the leading figures of quantitative sociology (Alwin 1991; Alwin and McCammon 2003; Clogg 1982; Clogg and Shockey 1985; Duncan 1985; Duncan and Stenbeck 1988; Fienberg and Mason 1979; Firebaugh 1989, 1990, 2008; Glenn 1977, 2005; Knoke and Hout 1974; Mason et al. 1973; Mason and Fienberg 1985; Reither et al. 2009; Rodgers 1982, 1990; Yang and Land 2013). Yet, despite significant methodological innovations, considerable controversy remains on how to obtain meaningful results from APC data. This is in part because Ryder, quite uncharacteristically given his foundational contributions to mathematical demography (1964; 1980; 1983), provided virtually no technical details on how to actually conduct a cohort analysis. As Hardy and Wilson (2002) rightly observed, "while Ryder's classic essay developing the connection between cohort succession and social change was extremely influential, it contained few specifics on the techniques of research" (P. 243).[4]

Progress has furthermore been hampered by a number of conceptual and methodological issues with APC data, most notably the deterministic relationship among the temporal variables. Simply put, period is the sum of age and cohort. As a consequence, in what is known as the APC identification problem, there is simply not enough information to determine the unique contributions of each variable on any given outcome. Conventional regression analysis cannot be used to estimate the parameters; instead, one must incorporate additional information into the model by, for example, fixing one of the coefficients to zero, applying parametric constraints, or specifying informative prior distributions over one or more of the parameters. There are now, decades after Ryder's classic essay was first published, a dizzying array of methods available to researchers wishing to extract unique effects from APC data, including Moore–Penrose estimators, structural equation models, time-series approaches, ridge and lasso regressions, multilevel models, and Bayesian regressions, among many other techniques.[5]

Although much has been written on various methods for identifying unique temporal effects in light of the identification problem, Ryder, in fact, advocated for a different strategy. Instead of disentangling period effects from those attributable to age or cohort,[6] Ryder's cohort approach is based on distinguishing intra-cohort trends (or life-cycle change) from inter-cohort trends (or social change), which, together, compose Ryderian comparative cohort careers. In other words, as Glenn (1976) correctly recognized, "a complete disentangling" of age, period, and cohort "is not necessary" to understand social change, and "it is this use of cohort analysis which Ryder (1965) recommends in his classic essay" (P. 903).[7] To underscore this critical point, I will refer to an *APC analysis* as any approach that attempts to derive unique "effects" for age, period, and cohort (e.g., Fosse and Winship 2019b; Mason et al. 1973; Mason and Fienberg 1985), whereas *cohort analysis* will refer to any approach that attempts to describe intra- and inter-cohort trends, or life-cycle and social change (e.g., Ryder 1965, 1968).

In this article, I show formally how the classic APC model can be used in line with Ryder's vision for cohort analysis, enabling researchers to parsimoniously summarize population-level temporal variability on the Lexis table in terms of intra- and inter-cohort trends (or life-cycle and social change), as well as Ryderian comparative cohort careers.[8] Additionally, although such an analysis is not a core part of Ryder's cohort approach, I reveal how one can adapt the classic APC model to formally summarize intra-period differences. Crucially, the overall framework developed in this article is based on fully identified models, complementing (but not replacing) conventional approaches that rely on assumptions, ideally informed by theory, to parse out unique APC "effects" (e.g., Fosse and Winship 2019a,b). As such, this has the potential to provide a common base of knowledge in a literature often fraught with controversy. Due to space limitations, I focus on interpreting the classic APC model, which is by far the most commonly used in the literature. However, as discussed below, the general principles outlined in this article can be applied to any number of other APC models.

The rest of this article is organized as follows. First, I outline the basic contours of Ryder's vision for analyzing cohorts and clarify the differences among intra-cohort, inter-cohort, and intra-period comparisons. Second, I demonstrate how the conventional APC model can be deployed for the purpose of summarizing population-level temporal variability, outlining three different models: diachronic, age synchronic, and cohort synchronic. I also discuss how these models are related to each other, revealing under what conditions the two synchronic models provide estimates identical to the diachronic model. Next, to illustrate how the classic APC model can be used for a Ryderian analysis, I examine trends in social trust in the United States using data from the U.S. General Social Survey (GSS). Then, drawing from Ryder's cohort approach, I discuss four main types of summaries from the diachronic and synchronic APC models: intra-cohort trends, inter-cohort trends, Ryderian comparative cohort careers, and intra-period differences. I also introduce a number of novel visualizations of these model-based summaries, including two-dimensional heat maps and three-dimensional Lexis surfaces. Additionally I discuss related model-based approaches in the sociological literature, including analyses of nonlinearities (Acosta and van Raalte 2019), an approach outlined by Duncan (1981), and a technique put forth by Firebaugh (1989, 1990, 1997, 2008). These methods can, in some sense, be interpreted as special cases of the Ryderian approach outlined in this article. Finally, I sketch possible extensions and future directions for research, as well as the limitations of the methods outlined in this article.

Diachronic Trends and Synchronic Differences

Consider first, for ease of exposition, a simple slopes-only APC model (cf. Mason et al. 1973:243). Let Y denote a continuous outcome, A age in years, P period (or survey year), and C cohort (or birth year). The basic slopes-only APC model can be represented by[9]

$$Y = \mu + \alpha A + \pi P + \gamma C + \epsilon, \quad (1)$$

where μ is the intercept; α, π, and γ are the age, period, and cohort slopes; and ϵ is an error term. Typically the goal in APC analysis is to understand the separate contributions of age, period, and cohort for a particular outcome (e.g., Mason et al. 1973). As is well known, however, the coefficients in Equation (1) cannot be estimated due to a linear dependency among the time scales, such that $p = a + c$. Accordingly, researchers must apply some information external to the data to identify unique temporal effects (e.g., Fosse and Winship 2019b).

Yet, as noted in the introduction, Ryder's cohort approach is based not on disentangling period effects from those effects attributable to age or cohort, but rather on distinguishing intra-cohort trends (or life-cycle change) from inter-cohort trends (or social change). The key insight is that, from a Ryderian perspective, the relationship of age and cohort to calendar time (i.e., period) is either *diachronic* (from Greek "dia-" meaning "through" and "khronos" meaning "time") or *synchronic* (from Greek "syn-" meaning "together" and "khronos" meaning "time").[10] In other words, observationally, we can compare age or cohort *through* calendar time (i.e., diachronically), resulting in dynamic trends, or *within* a particular cross-section of calendar time (i.e., synchronically), thereby generating static differences. For a similar point, see Riley (1973).

Perhaps unsurprisingly, synchronic comparisons are not a core feature of Ryder's cohort approach, which is instead based on diachronic comparisons. As Ryder (1965) declared in no uncertain terms, an analysis based on synchronic measures "destroys individual sequences," "inhibits dynamic inquiry," "diverts attention from process," and "fosters the illusion of immutable structure" (P. 859). Fortunately, as Ryder correctly recognized, inasmuch it is based on diachronic rather than synchronic comparisons, these problems "can be avoided by using the cohort approach" (P. 859).

The basic slopes-only APC model of Equation (1) can be easily altered for estimating diachronic trends in line with Ryder's vision for cohort analysis. The key idea is to re-index the parameters with respect to A and C, as suggested by Ryder (1968). This is accomplished by substituting P with $A + C$ into Equation (1) and rearranging terms, resulting in the *diachronic slopes-only model*:

$$Y = \mu + \theta_1 A + \theta_2 C + \epsilon, \tag{2}$$

where μ is the intercept, $\theta_1 = \alpha + \pi$, $\theta_2 = \gamma + \pi$, and ϵ is an error term. Note that Equation (2) is fully identified and hence estimable. Despite its simplicity, Equation (2) is crucial for understanding the relationship between APC models and Ryder's cohort approach.

To clarify the interpretation of Equation (2) as well as its relevance for cohort analysis, Table 1 summarizes θ_1 and θ_2 as well as the two logically possible differences, $\theta_1 - \theta_2$ and $\theta_2 - \theta_1$. As shown in the first row of Table 1, the first parameter is $\theta_1 = \alpha + \pi$, the diachronic inter-age slope or, for short, the *diachronic age slope*. The diachronic age slope reflects an *intra-cohort trend* because we are observing how, within cohorts, populations of individuals are aging through calendar time (i.e., across periods). Using Ryder's (1965) terminology, θ_1 is a measure of "intra-cohort temporal development throughout the life-cycle" (P. 861), or, simply, life-cycle change.

Table 1: Overview of diachronic and synchronic slope estimands

Slope estimand	Relation to calendar time	Within-group comparisons	Between-group comparisons	Ryderian interpretation	
$\theta_1 = \alpha + \pi$	Diachronic	Intra-cohort trend	Inter-age trend	Life-cycle change	Comparative cohort career
$\theta_2 = \gamma + \pi$	Diachronic	Intra-age trend	Inter-cohort trend	Social change	
$\theta_1 - \theta_2 = \alpha - \gamma$	Synchronic	Intra-period differences	Inter-age differences		
$\theta_2 - \theta_1 = \gamma - \alpha$	Synchronic	Intra-period differences	Inter-cohort differences		

Notes: The diachronic age slope, which represents an intra-cohort trend, is $\theta_1 = \alpha + \pi$, where α is the age slope and π is the period slope. The diachronic cohort slope, which represents an inter-cohort trend, is $\theta_2 = \gamma + \pi$, where α is the age slope and π is the period slope. The synchronic age and cohort slopes, which represent intra-period differences, are given by $\theta_1 = \theta_2 = \alpha - \gamma$ and $\theta_2 - \theta_1 = \gamma - \alpha$, respectively. Highlighted cells indicate "Ryderian" terminology for characterizing comparisons across levels of APC data.

The second parameter, as shown in the second row of Table 1, is $\theta_2 = \gamma + \pi$, the diachronic inter-cohort slope or, for short, the *diachronic cohort slope*. The diachronic cohort slope reflects an *inter-cohort trend* because we are comparing successive cohorts through calendar time (i.e., across periods). In Ryder's (1965) phraseology, θ_2 is a measure of "inter-cohort temporal differentiation" (P. 861), or, equivalently, social change. Combined, as shown in the last column of Table 1, intra-cohort and inter-cohort trends (or life-cycle and social change) constitute Ryderian comparative cohort careers.

The bottom rows of Table 1 outline the interpretation of the two logically possible differences between θ_1 and θ_2. As displayed in the third row of Table 1, $\theta_2 - \theta_1 = \gamma - \alpha$ defines the synchronic inter-age slope or, for short, the *synchronic age slope*. Similarly, the difference $\theta_1 - \theta_2 = \alpha - \gamma$ defines the synchronic inter-cohort slope or, for short, the *synchronic cohort slope*. Rather than dynamic trends, the synchronic age and cohort slopes reflect static *intra-period differences*.[11] Intuitively this is because these slopes are based on comparisons within cross-sections of time (i.e., within periods).

The synchronic age and cohort slopes can be calculated as the difference between the diachronic age and cohort slopes estimated in Equation (2). Alternatively, they can be estimated directly by re-indexing Equation (1) with respect either to A and P or to C and P, respectively. Specifically, re-indexing Equation (1) with respect to A and P (by substituting C with $P - A$ and rearranging terms) gives the *age synchronic slopes-only model*:

$$Y = \mu + A(\theta_1 - \theta_2) + P\theta_2 + \epsilon. \tag{3}$$

Similarly, re-indexing Equation (1) with respect to P and C (by substituting A with $P - C$ and rearranging terms) results in the *cohort synchronic slopes-only model*:

$$Y = \mu + P\theta_1 + C(\theta_2 - \theta_1) + \epsilon. \quad (4)$$

The age synchronic slopes-only model (Eq. [3]) generates estimates of the age synchronic slope and the diachronic cohort slope, whereas the cohort synchronic slopes-only model (Eq. [4]) provides estimates of the diachronic age slope and the cohort synchronic slope. Intuitively, both models provide cross-sectional estimates for age and cohort, respectively, because they condition on calendar time (or period).

To provide additional insight on the interpretation of the slope parameters in Table 1, Figure 1 shows an age–period Lexis heat map based on simulated data. Each cell of Figure 1 displays the mean of the outcome based on values of $\theta_1 = 0.08$ and $\theta_2 = 0.10$, with selected age (row), period (column), and cohort (diagonal) groups highlighted for the purposes of exposition. The three highlighted sections illustrate the various ways in which the underlying linear trends of a Lexis table can be dissected. First, the diagonal section in Figure 1 identifies those individuals born in 1920.[12] Because cohorts have been constructed based on age and period, we only observe this cohort for a section of its entire life-cycle, from age 50 in 1970 to age 60 in 1980. As this cohort ages through time, the outcome shifts across both age (rows) and period (columns) groups. In other words, comparisons within this section reflect an intra-cohort trend (i.e., the diachronic age slope). Accordingly, tracing from the upper left to lower right, the difference between adjacent cells, which are five years apart, is equal to $0.08 \times 5 = 0.40$, where $\theta_1 = \alpha + \pi = 0.08$. For example, at age 50 in period 1970 the expected mean outcome is -1.50, whereas at age 55 in period 1975 it is -1.10, an increase of 0.40.

Second, the horizontal section in Figure 1 subsets to those individuals aged 30 years. As we compare successive cohorts through time, the outcome shifts across both cohort (diagonals) and period (columns) groups. That is, comparisons within this section reflect an inter-cohort trend (i.e., diachronic cohort slope). Following from left to right, the difference between adjacent cells, which are again five years apart, equals $0.10 \times 5 = 0.40$, where $\theta_2 = \gamma + \pi = 0.10$. For instance, for the cohort born in 1940 and observed in 1970, the expected mean outcome is -1.10, whereas for the cohort born in 1945 and observed in 1975 the expected mean outcome is -0.60, an increase of 0.50.

Lastly, the vertical section in Figure 1 identifies those individuals observed in the period 1990. As we compare adjacent cells within this vertical section, the outcome shifts across both age (rows) and cohort (diagonals) groups. Specifically, comparing age groups from young to old (i.e., low to high age levels) also necessarily entails comparing cohort groups from recent to old (i.e., high to low cohort levels). Because these comparisons occur within a cross-section, they reflect static intra-period differences (i.e., synchronic age or cohort slopes) rather than dynamic trends. Following from top to bottom (i.e., low to high age levels), the difference between adjacent cells, which again are five years apart, is $-0.02 \times 5 = -0.10$, where $\theta_1 - \theta_2 = -0.02$. By contrast, following from bottom to top (i.e., low to high cohort levels), the difference between adjacent cells, which are again five years apart, is $0.02 \times 5 = 0.10$, where $\theta_2 - \theta_1 = 0.02$. For instance, for those aged 15 years and

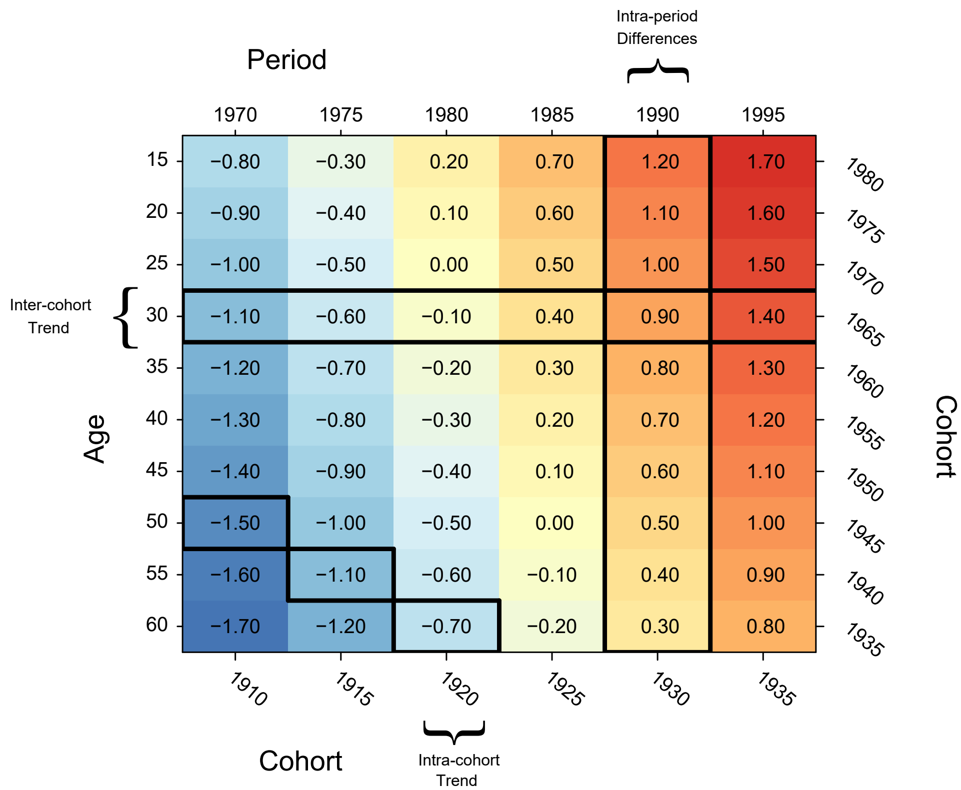

Figure 1: Two-dimensional age–period–cohort Lexis heat map of diachronic trends and synchronic differences. *Notes:* Each cell gives the mean outcome based on simulated data using values of $\theta_1 = 0.08$ and $\theta_2 = 0.10$. Three sections are highlighted: the diachronic intra-cohort trend for the birth cohort born in 1920 (diagonal section), the diachronic inter-cohort trend for those aged 30 years (horizontal section), and intra-period differences for those in the 1990 period (vertical section).

born in 1990 the expected mean outcome is 1.20, whereas for those aged 20 years and born in 1990 the expected mean outcome is 1.10. Comparing the 15-year-old age group with the 20-year-old age group in 1990 (and thus the 1975 cohort with the 1970 cohort), this difference is a decrease of 0.10. Conversely, comparing the 1970 cohort with the 1975 cohort in 1990 (and thus the 20-year-old age group with the 15-year-old age group), this difference is an increase of 0.10.

In short, instead of the separate linear effects for age, period, and cohort, the building blocks of a Ryderian analysis are the diachronic age and cohort slopes, denoted by θ_1 and θ_2. These two parameters represent intra- and inter-cohort trends, respectively. By contrast, their cross-sectional counterparts, the synchronic age slope $\theta_1 - \theta_2$ and the synchronic cohort slope $\theta_2 - \theta_1$, represent intra-period differences, which, as noted above, are not a core part of Ryder's cohort approach. The diachronic age slope and cohort slopes correspond to comparisons within diagonal and horizontal sections, respectively, of an age–period Lexis table, whereas synchronic age and cohort slopes correspond to comparisons within vertical sections.

Diachronic and Synchronic Temporal Models

So far the discussion has focused on interpreting the linear trends underlying a set of APC data. In practice, APC researchers typically fit a more flexible model than one with just age, period, and cohort linear components. In this section I first introduce three main diachronic and synchronic APC models, showing how they are related to each other, and then apply these models to the empirical example, an examination of temporal trends in social trust in the U.S. GSS. I focus on interpreting the conventional APC model because it is by far the model most commonly used in the literature. However, the general principles of a Ryderian analysis outlined in the previous section could be applied to any number of parametric models used by APC analysts.

Classical and Linearized APC Models

Suppose we have categorically coded age, period, and cohort data for a set of n respondents.[13] We 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 $K = I + J - 1$.[14] The *classical APC (C-APC) model* is (Fosse and Winship 2019a)

$$Y_{ijk} = \mu + \alpha_i + \pi_j + \gamma_k + \epsilon_{ijk}, \tag{5}$$

where μ is the overall mean; α_i, π_j, and γ_k are age, period, and cohort parameters; and ϵ is an error term. 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, the C-APC model is still not identified even after assuming sum-to-zero constraints (e.g., see Mason and Fienberg 1985; Yang and Land 2013).

An alternative representation of the C-APC model orthogonally decomposes the linear from the nonlinear components. We can accordingly specify the *linearized APC (L-APC) model* with the form (Fosse and Winship 2018)

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

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 for age, period, and cohort; $\tilde{\alpha}$, $\tilde{\pi}$, and $\tilde{\gamma}$ represent age, period, and cohort nonlinearities; and again ϵ_{ijk} is an error term.

Note that the C-APC and L-APC models are equivalent representations of temporal data grouped by age, period, and cohort. As with the C-APC model, each cell in an age–period array is modeled by a unique combination of parameters. For example, the ith age parameter in the C-APC model is represented in the L-APC model by the overall age slope along with a unique parameter for the ith age nonlinearity: $\alpha_i = (i - i^*)\alpha + \widetilde{\alpha}_i$. In other words, each age parameter α_i is decomposed into the sum of a common parameter α representing the age slope for the entire array, with a value shifting across rows (or age categories) as a function of the age index i, and a unique parameter $\widetilde{\alpha}_i$, which is a nonlinearity specific to each row (or age category). We can similarly decompose each of the period and cohort parameters in the C-APC model into linear and nonlinear terms. Like the C-APC model, the L-APC model also cannot be identified due to the linear dependency among the time scales. Because they are simply alternative representations of the same underlying model, I will refer to the C-APC and L-APC models interchangeably as the "classical," "classic," or "conventional" APC model.[15]

The Diachronic L-APC Model and Synchronic Alternatives

As noted previously, the great majority of work in APC analysis has focused on developing techniques for identifying unique APC effects. However, as with the basic slopes-only APC model, one can use the conventional APC model to estimate diachronic trends in line with Ryder's cohort approach. The idea, as before, is to re-index the C-APC model with respect to age and cohort by substituting period with age and cohort.[16] After rearranging terms, this results in the *diachronic L-APC model*:

$$Y_{ijk} = \mu + \theta_1(i - i^*) + \theta_2(k - k^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} + \widetilde{\gamma}_k + \epsilon_{i[i+k-I]k}, \tag{7}$$

where $\theta_1 = \alpha + \pi$ is the diachronic age slope and $\theta_2 = \gamma + \pi$ is the diachronic cohort slope. The diachronic L-APC model is fully identified (i.e., the design matrix is of full rank) and thus can be estimated using APC data.

As with the basic slopes-only APC model discussed in the previous section, one can also estimate the two synchronic variants of the L-APC model. First, re-indexing the L-APC model with respect to age and period results in the *age synchronic L-APC model*:

$$Y_{ijk} = \mu + (\theta_1 - \theta_2)(i - i^*) + \theta_2(j - j^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_{j-i+I} + \epsilon_{ij[j-i+I]}, \tag{8}$$

where $\theta_1 - \theta_2 = \alpha - \gamma$ is the synchronic age slope and $\theta_2 = \gamma + \pi$ is again the diachronic cohort slope, but indexed by period ($j = 1, \ldots, j = J$) rather than cohort ($k = 1, \ldots, k = K$). Second, re-indexing the L-APC model with respect to period and cohort results in the *cohort synchronic L-APC model*:

$$\begin{aligned} Y_{ijk} = \mu + \theta_1(j - j^*) + (\theta_2 - \theta_1)(k - k^*) + \widetilde{\alpha}_{j-k+(K-J+1)} \\ + \widetilde{\pi}_j + \widetilde{\gamma}_k + \epsilon_{[j-k+(K-J+1)]jk}, \end{aligned} \tag{9}$$

where $\theta_2 - \theta_1 = \gamma - \alpha$ is the synchronic cohort slope and $\theta_1 = \alpha + \pi$ is again the diachronic age slope, but indexed by period ($j = 1, \ldots, j = J$) rather than age ($i = 1, \ldots, i = I$). As with their slopes-only counterparts discussed in the previous section (Eqs. [3] and [4]), the age synchronic L-APC model (Eq. [8]) and the cohort synchronic L-APC model (Eq. [9]) provide synchronic estimates for age and cohort, respectively, because both models involve conditioning on the period linear component.

For the purposes of summarizing APC data within a Ryderian framework, the diachronic L-APC model is generally preferred over the age and cohort synchronic

Table 2: Relationships among diachronic and synchronic slopes

Diachronic age slope:	If $\theta_1 = 0$,	then	$\theta_2 - \theta_1$	$=$	θ_2
	If $\theta_1 > 0$,	then	$\theta_2 - \theta_1$	$<$	θ_2
	If $\theta_1 < 0$,	then	$\theta_2 - \theta_1$	$>$	θ_2
Diachronic cohort slope:	If $\theta_2 = 0$,	then	$\theta_1 - \theta_2$	$=$	θ_1
	If $\theta_2 > 0$,	then	$\theta_1 - \theta_2$	$<$	θ_1
	If $\theta_2 < 0$,	then	$\theta_1 - \theta_2$	$>$	θ_1
Synchronic age slope:	If $\theta_1 - \theta_2 = 0$,	then	θ_1	$=$	θ_2
	If $\theta_1 - \theta_2 > 0$,	then	θ_1	$>$	θ_2
	If $\theta_1 - \theta_2 < 0$,	then	θ_1	$<$	θ_2
Synchronic cohort slope:	If $\theta_2 - \theta_1 = 0$,	then	θ_2	$=$	θ_1
	If $\theta_2 - \theta_1 > 0$,	then	θ_2	$>$	θ_1
	If $\theta_2 - \theta_1 < 0$,	then	θ_2	$<$	θ_1

Notes: $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$.

L-APC models. The reason is that, although all three models provide identical estimates of the age, period, and cohort nonlinearities as well as the intercept, the age and cohort slopes of the the two synchronic models will not reflect intra- and inter-cohort trends (i.e., life-cycle and social change) except under very strong assumptions on the absence of a linear trend across or within cohorts.[17] To understand why this is the case, consider Table 2, which outlines the relationships among the diachronic and synchronic slopes and, accordingly, the three L-APC models defined in Equations (7), (8), and (9).

The top panel in Table 2 shows how, for different values of the diachronic age slope, the synchronic and diachronic cohort slopes are systematically related. As shown in the first line of the top panel, the synchronic cohort slope will equal the diachronic cohort slope only if the diachronic age slope is zero. Under this condition, the cohort slope in Equation (9) will be an unbiased estimate for the cohort slope in Equation (7). Otherwise, as shown in the second and third lines of the top panel, if the diachronic age slope is greater (or less) than zero, then the synchronic cohort slope will be less (or greater) than the diachronic cohort slope, and accordingly the cohort slope in Equation (9) will provide an underestimate (or overestimate) of the cohort slope in Equation (7).

The middle panel in Table 2 reveals the relations the relationships between the synchronic and diachronic age slopes for different values of the diachronic cohort slope. As indicated in the first line of the middle panel, if the diachronic cohort slope is zero, then the synchronic age slope will equal the diachronic age slope. Under this scenario, the age slope in Equation (8) will be an unbiased estimate of the age slope in Equation (7). By contrast, as shown in the second and third lines of the middle panel, if the diachronic cohort slope is greater (or less) than zero, then the synchronic age slope will be less (or greater) than the diachronic age slope, and accordingly the age slope in Equation (8) will provide an underestimate (or overestimate) of the age slope in Equation (7).

The bottom panel in Table 2 shows the relationships between the diachronic slopes for different values of the synchronic slopes. As shown in the first and fourth lines of the bottom panel, if the synchronic age (or cohort) slope is zero, then the diachronic age and cohort slopes are the same. Otherwise, if the synchronic age slope is negative (or, equivalently, the synchronic cohort slope is positive), then the diachronic age slope is less than the diachronic cohort slope in Equation (7). By contrast, if the synchronic age slope is positive (or, equivalently, the synchronic cohort slope is negative), then the diachronic age slope is greater than the diachronic cohort slope in Equation (7).

Two main conclusions follow from the relationships in Table 2. First, for the purposes of a Ryderian cohort analysis, in general the diachronic L-APC model is preferred over the age or cohort synchronic L-APC models. The reason is that, except under very strong (and often testable) assumptions, the age and cohort slopes of the synchronic L-APC models represent static (cross-sectional) comparisons, not dynamic trends.[18] The age synchronic slope will be an unbiased estimate of the diachronic age slope only if there is no linear inter-cohort trend or social change; likewise, the cohort synchronic slope will be an unbiased estimate of the diachronic cohort slope if there is no linear intra-cohort trend or life-cycle change. Although it will naturally depend on the particular application, these conditions seem unlikely to hold in most cases.

Second, the age and cohort synchronic L-APC models, although not providing direct estimates of intra- and inter-cohort trends, respectively, are, in fact, informative about the relative magnitude of intra- versus inter-cohort trends (i.e., life-cycle versus social change). If the synchronic age slope is negative (or, equivalently, the synchronic cohort slope is positive), then the linear intra-cohort trend is less than that for the inter-cohort trend; that is, the overall (linear) life-cycle change is less than that for social change. By contrast, if the synchronic age slope is positive (or, equivalently, the synchronic cohort slope is negative), then the linear intra-cohort trend is greater than that for the inter-cohort trend; in other words, the overall (linear) life-cycle change is greater than that for social change. This topic is explored later in this article in the discussion of parametric expressions for intra-period differences (see Eqs. [20] and [21]).

Example: Social Trust in the United States

To illustrate how APC data can be used to summarize population-level temporal variability in line with Ryder's vision, I examine trends in social trust using data from the U.S. GSS. Respondents were asked the following question: "Generally speaking, would you say that most people can be trusted or that you can't be too careful in dealing with people?" The outcome is coded so that $1 =$ "can trust" while $0 =$ "can't be too careful" or "depends." The sample size is $R = 41,126$ respondents, and all results are adjusted using appropriate sampling weights. There are $I = 14$ age groups, $J = 10$ period groups, and $K = 23$ cohort groups. Age groups were coded into five-year intervals beginning with 18 to 22 and ending with 83 and older,

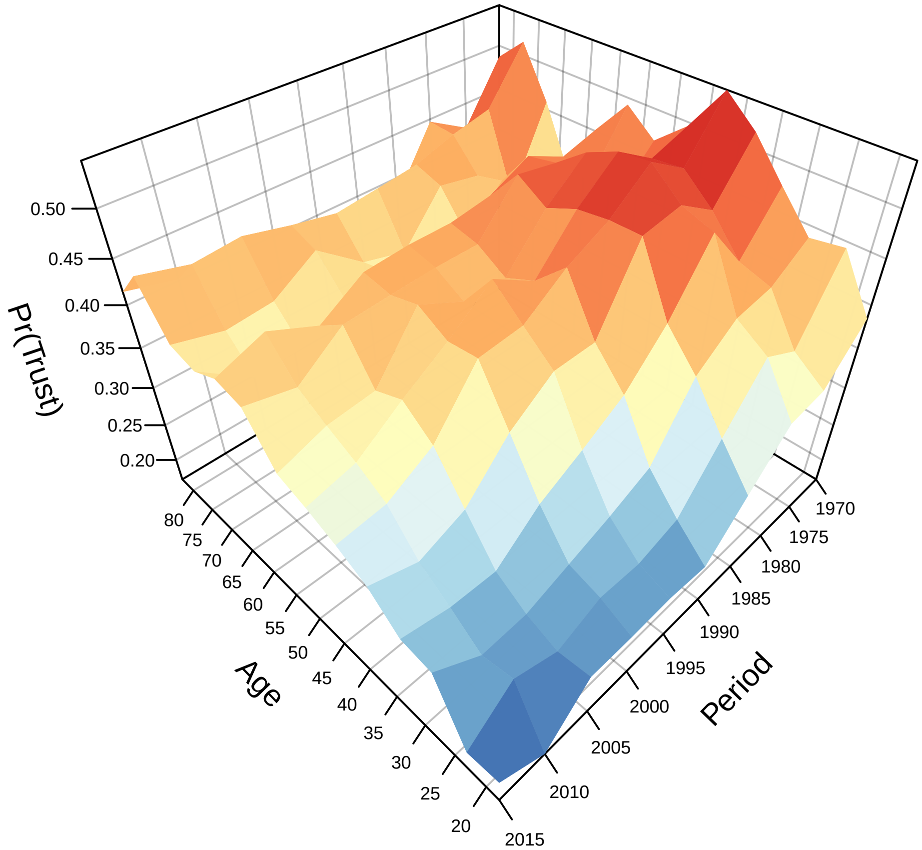

Figure 2: Three-dimensional age–period–cohort Lexis surface. *Notes:* Each cell gives the predicted probability of the outcome from a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Sample size $R = 41,126$ respondents. Results adjusted using sampling weights.

whereas period groups were coded into five-year intervals beginning with 1970 to 74 and ending with 2015 to 2019. Cohort groups, calculated as the difference between the period and age groups, began with 1887 to 1891 and ended with 1997 to 2001.

Figure 2 shows a three-dimensional age–period–cohort Lexis surface of the predicted probability of social trust using the diachronic L-APC model.[19] The predicted probability in each cell is based on a unique combination of age, period, and cohort values. Results show considerable temporal variability, with the lowest levels of social trust among young people in more recent cohorts who are also observed in more recent periods, and higher levels of social trust among middle-aged people in earlier cohorts who are accordingly observed in earlier periods.

An alternative representation of the surface of predicted probabilities is shown in Figure 3, which shows a two-dimensional Lexis heat map. This figure highlights the three main types of descriptive comparisons summarized in the third and fourth columns of Table 1. First, there is the intra-cohort trend, which is equivalent to

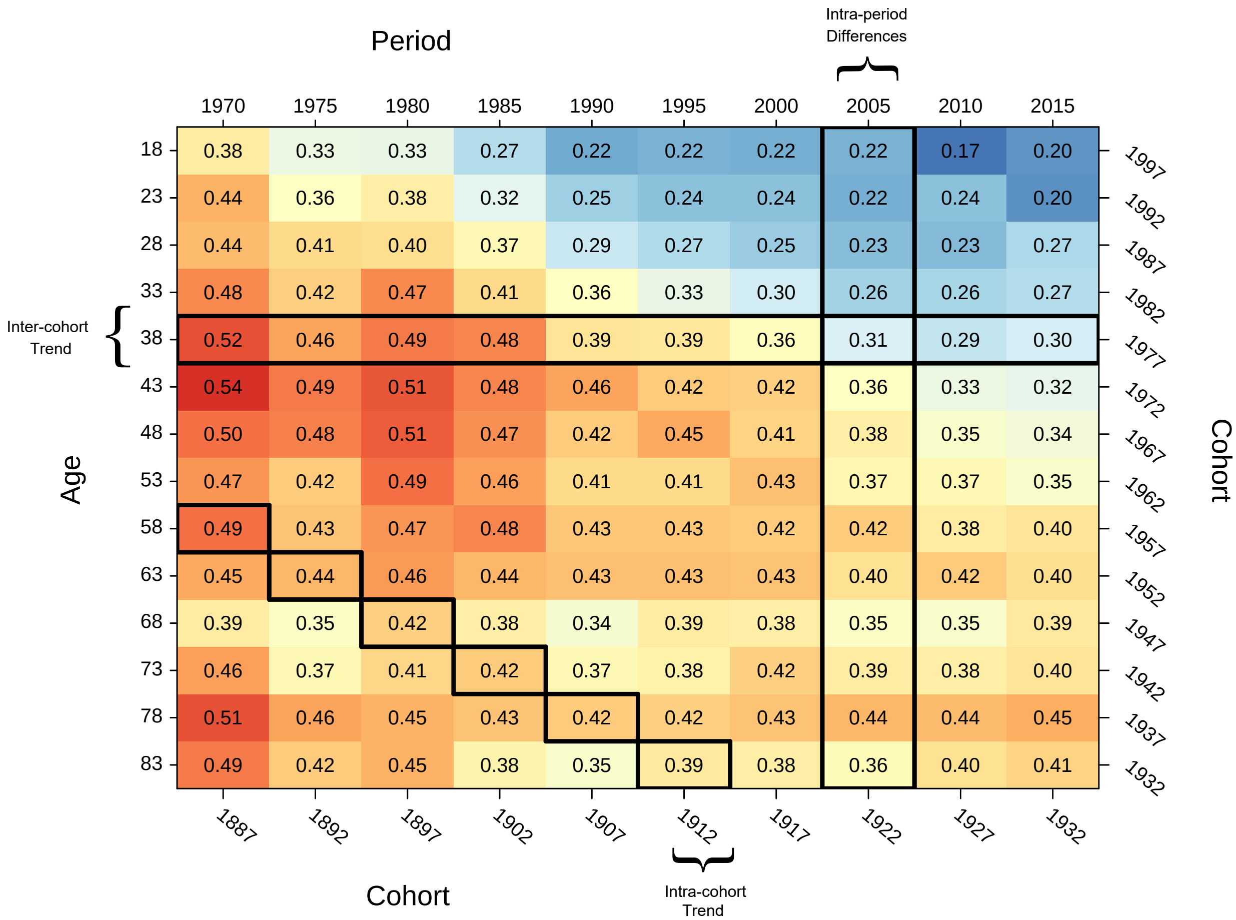

Figure 3: Two-dimensional Lexis heat map of diachronic trends and synchronic differences, plotting predicted probabilities from a diachronic L-APC logistic regression model. *Notes:* Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Sample size $R = 41,126$ respondents. Results adjusted using sampling weights. Three sections are highlighted: the diachronic intra-cohort trend for the 1912-to-1916 birth cohort (diagonal section), the diachronic inter-cohort trend for the 38-to-42 age group (horizontal section), and intra-period differences for the 2005-to-2009 period group (vertical section).

stratifying by cohort and comparing age groups through time. For example, as shown by the highlighted diagonal in Figure 3, one can subset to the cohort with a midpoint birth year of 1912 and compare the probability of social trust among those individuals aged 58 to 62 years and observed from 1970 to 1974 with that of those aged 83 and older and observed from 1995 to 1999. Second, there is the inter-cohort trend, which is equivalent to stratifying by age and comparing cohorts through time. For instance, as indicated by the highlighted row in Figure 3, one can subset to those individuals aged 38 to 42 years and compare the probability of social trust among those individuals with a midpoint birth year of 1932 and observed between 1970 and 1974 with that of those with a midpoint birth year of 1977 and observed between 2015 and 2019. Finally, there are intra-period differences. This is equivalent to stratifying by period and comparing age-cum-cohort groups. For example, as

Table 3: Key parameters from diachronic and synchronic L-APC models

	Log-odds ratio			Odds ratio	
Parameter	Coef.	Std. error	p value	Coef.	95% CI
Intercept (μ)	-0.558	0.027	< 0.001	0.573	$(0.542, 0.604)$
Diachronic age slope (θ_1)	-0.192	0.055	0.001	0.826	$(0.741, 0.920)$
Diachronic cohort slope (θ_2)	-0.677	0.071	< 0.001	0.508	$(0.442, 0.584)$
Synchronic cohort slope ($\theta_2 - \theta_1$)	-0.485	0.066	< 0.001	0.615	$(0.541, 0.700)$
Synchronic age slope ($\theta_1 - \theta_2$)	0.485	0.066	< 0.001	1.625	$(1.423, 1.848)$

Notes: $\theta_1 = \alpha + \pi$, $\theta_2 = \gamma + \pi$. CI = confidence interval. Estimates based on logistic diachronic and synchronic L-APC models. Outcome is social trust, coded so that $1 =$ "can trust" while $0 =$ "can't be too careful" or "depends." Sample size $R = 41,126$ respondents. Results adjusted using sampling weights.

shown by the highlighted column in Figure 3, one can subset to the 2005-to-2010 period and compare the probability of social trust among those individuals aged 18 to 22 years (with a midpoint birth year of 1987) with that of those aged 83 or older (with a midpoint birth year of 1922). Again, because we are comparing individuals within a single cross-section, this analysis is static rather than dynamic.

Underlying the patterns in Figure 3 are the diachronic and synchronic slopes discussed previously, which can be estimated using the diachronic and synchronic L-APC models shown in Equations (7), (8), and (9). Table 3 shows the estimates for the intercept as well as various diachronic and synchronic slopes. The diachronic age slope, which represents an intra-cohort trend, indicates that there is a decline in social trust as we compare age groups through time (log-odds ratio: -0.192; $p = 0.001$). Likewise, the diachronic cohort slope, which represents an inter-cohort trend, reveals that there is a decline in social trust as we compare cohorts through time (log-odds ratio: -0.677; $p < 0.001$). The synchronic cohort slope, representing intra-period differences, indicates that the difference between the diachronic cohort and age slopes is large and statistically significant (log-odds ratio: -0.485; $p < 0.001$). This is also reflected in the synchronic age slope, which has the same size but opposite sign as the synchronic cohort slope (log-odds ratio: 0.485; $p < 0.001$). In short, the estimated slopes indicate considerable life-cycle and social change in the data, with a particularly steep decline as we compare cohorts through time.

Table 4 compares results for the diachronic L-APC model and various submodels. The first row presents fit statistics for the diachronic L-APC model, which is the baseline model for the Wald tests. The second and third rows are models that drop the diachronic age and cohort slopes, respectively. Rows four to six present models that drop the age, period, and cohort nonlinearities, respectively. The last three rows of Table 4 present fit statistics for models that drop two of the nonlinearities. Wald tests indicate that all of the nonlinearities are statistically significant at conventional levels. Visual inspection of the nonlinearities also reveal notable patterns, especially for age and cohort (see Figure 3 in the online supplement). The Akaike information criterion (AIC) and Bayesian information criterion (BIC) both suggest that one

Table 4: Fit statistics of diachronic L-APC model versus selected submodels

APC model	LLV	AIC	BIC	Wald test	
				χ^2	p value
$\theta_1 + \theta_2 + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k$	53,117.38	53,224.11	53,584.85	Baseline model	
$\theta_2 + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k$	53,131.07	53,235.51	53,586.44	12.09	< 0.001
$\theta_1 + \widetilde{\alpha}_i + \widetilde{\pi}_j + \widetilde{\gamma}_k$	53,184.20	53,291.28	53,664.97	90.74	< 0.001
$\theta_1 + \theta_2 + \widetilde{\pi}_j + \widetilde{\gamma}_k$	53,235.62	53,313.29	53,553.67	94.17	< 0.001
$\theta_1 + \theta_2 + \widetilde{\alpha}_i + \widetilde{\gamma}_k$	53,167.29	53,253.88	53,540.97	39.29	< 0.001
$\theta_1 + \theta_2 + \widetilde{\alpha}_i + \widetilde{\pi}_j$	53,286.66	53,341.18	53,508.23	142.26	< 0.001
$\theta_1 + \theta_2 + \widetilde{\alpha}_i$	53,329.00	53,363.26	53,457.75	174.95	< 0.001
$\theta_1 + \theta_2 + \widetilde{\pi}_j$	53,561.43	53,586.66	53,589.56	348.94	< 0.001
$\theta_1 + \theta_2 + \widetilde{\gamma}_k$	53,289.79	53,347.32	53,512.47	136.15	< 0.001

Notes: $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$. LLV = log-likelihood value. AIC = Akaike information criterion. BIC = Bayesian information criterion. Calculations based on a diachronic L-APC logistic regression model and selected submodels. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Sample size $R = 41,126$ respondents. Results adjusted using sampling weights.

should include a model with nonlinearities. AIC suggests that the best-fitting model includes both diachronic slopes and all of the nonlinearities. As expected, BIC, which penalizes the inclusion of additional parameters, prefers a more parsimonious model that includes both slopes but just the age nonlinearities.

Summarizing Population-Level Temporal Variability

As shown in previous section, the conventional APC model can be adapted for summarizing intra- and inter-cohort trends, or life-cycle and social change. The next task is to determine what parametric expressions are useful for parsimoniously summarizing, in line with Ryder's (1965) goals for cohort analysis, population-level temporal variability on the Lexis table. Table 5 summarizes the main parametric expressions that can be estimated using Equations (7), (8), and (9). Following the general Ryderian framework for cohort analysis outlined in Table 1, these expressions are categorized into four main types. The top panel of Table 5 lists expressions for intra-cohort trends (or, equivalently, life-cycle change), the second panel lists expressions for inter-cohort trends (or, equivalently, social change), the third panel lists expression for Ryderian comparative cohort careers (which unite life-cycle with social change), and the last panel displays expressions for intra-period differences. In general, the expressions within each category are listed in

Table 5: Summarizing population-level temporal variability with an APC model

Ryderian terminology	Specific summary	Parametric expression	
Intra-cohort trends (life-cycle change)	Diachronic age slope	$\theta_1(i - i^*)$	for all i
	Diachronic age curve	$\theta_1(i - i^*) + \widetilde{\alpha}_i$	for all i
	Age–period Lexis surface	$\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j$	for combinations of i, j
	Local diachronic age curves	$\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$	for all i in each cohort k
Inter-cohort trends (social change)	Diachronic cohort slope	$\theta_2(k - k^*)$	for all k
	Diachronic cohort curve	$\theta_2(k - k^*) + \widetilde{\gamma}_k$	for all k
	Cohort–period Lexis surface	$\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_j$	for combinations of k, j
	Local diachronic cohort curves	$\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I}$	for all k in each age i
Ryderian comparative cohort careers	Curves-only comparative cohort careers	$\phi_k + \theta_1(i - i^*) + \widetilde{\alpha}_i$	for all i in each cohort k
	Adjusted comparative cohort careers	$\phi_k + \widetilde{\pi}_{i+k-I} + \theta_1(i - i^*) + \widetilde{\alpha}_i$	for all i in each cohort k
Intra-period differences	Synchronic age slope	$(\theta_1 - \theta_2)(i - i^*)$	for all i
	Synchronic age curve	$(\theta_1 - \theta_2)(i - i^*) + \widetilde{\alpha}_i$	for all i
	Local synchronic age differences	$(\theta_1 - \theta_2)(i - i^*) + (\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I})$	for all i in each period j
	Synchronic cohort slope	$(\theta_2 - \theta_1)(k - k^*)$	for all k
	Synchronic cohort curve	$(\theta_2 - \theta_1)(k - k^*) + \widetilde{\gamma}_k$	for all k
	Local synchronic cohort differences	$(\theta_2 - \theta_1)(k - k^*) + (\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I})$	for all k in each cohort k

Notes: $\theta_1 = \alpha + \pi$ and $\theta_2 = \gamma + \pi$. The quantity ϕ_k is equal to $\theta_2(k - k^*) + \widetilde{\gamma}_k$, which is a single value for a given cohort k.

increasing order of complexity in the sense a larger number of terms from the conventional APC model are included.

The second and third columns of Table 5 list the specific summaries and parametric expressions that fall under each Ryderian category, whereas the fourth column lists the dimensions over which the parameters are successively calculated or compared. There are several key points about the terminology used in the second column, which is more precise than Ryder's language or that commonly deployed by researchers. This precision is crucial so as to provide a tighter coupling between theory and method as well as to avoid fundamental misinterpretations that appear in the APC literature.[20] First, throughout "slopes" refers to estimands based on only a linear component, whereas "curves" refers to estimands based on linear as well as nonlinear components. Slopes and curves can be calculated and compared across a single temporal variable. By contrast, "surfaces" are expressions that entail calculating and comparing two temporal variables simultaneously and are thus most naturally visualized as two-dimensional heat maps or three-dimensional surfaces. For instance, the diachronic age slope is based only on θ_1, whereas the diachronic age curve is based on θ_1 as well as the set of $\widetilde{\alpha}_i$'s, or age nonlinearities. The age–period Lexis surface also includes the set of $\widetilde{\pi}_j$'s, introducing another temporal dimension (i.e., period) to the expression. As discussed previously, the terms "diachronic" and "synchronic" refer to whether or not we are comparing successive parameters using θ_1 or θ_2, which are diachronic in that they capture age or cohort differences across periods, or $\theta_2 - \theta_1$ or $\theta_1 - \theta_2$, which are synchronic in that they reflect static differences within periods.

Second, the expressions in the second column of Table 5 refer to "inter-" group comparisons, with the particular groups that are successively compared listed in the last column of Table 5. However, for simplicity the "inter-" prefix has been omitted from the names of the expressions. The reason for this is that the diachronic/synchronic distinction is only necessary for differentiating between inter-group comparisons, not intra-group comparisons.[21] Accordingly, by virtue of employing the adjective "diachronic," we are, in fact, signaling that an expression entails "inter-" group comparisons, with parameters successively calculated and compared across some specified set of groups. Thus, for example, rather than labeling the expression in the second row of Table 5 the "diachronic inter-age curve," we call it the "diachronic age curve" for short. This is an inter-age comparison, because parameters are calculated and compared across successive age groups, that is, for all $i = 1, \ldots, I$, as indicated by the phrase "for all i" in the last column of Table 5.

As well, the second column of Table 5 avoids referring to "intra-" group comparisons. The reason for this is that, once nonlinearities are incorporated, referring to intra-group comparisons is inherently more vague than referring to inter-group comparisons, whether diachronic or synchronic. To see this, take, for example, the expression for the diachronic age curve in the second row of Table 5. Within cohorts there are both age and period parameters that can be calculated and successively compared. If we were to call the expression for the diachronic age curve an intra-cohort curve, it would be unclear whether or not we were calculating and comparing parameters across age or period (or both), or if age or period nonlinearities (or both) were included. A similar ambiguity would arise if we were to refer to various other parametric expressions in terms of intra- rather than inter-group comparisons.

Finally, in Table 5 we have included several additional terms that differentiate the various expressions, including "local," "comparative," "careers," and "adjusted." Local expressions are diachronic age or cohort curves that, by including the period nonlinearities, are unique to each cohort or age group. The reason is that the period nonlinearities, which are experienced by different ages for different cohorts, allow the curves to vary across cohorts. By contrast, comparative expressions are those that directly include the parameters θ_1 and θ_2 (not their difference). Note the two Ryderian comparative cohort career expressions in the third panel of Table 5 include the diachronic age curve, which is the "career" aspect of these expressions. By contrast, the "comparative" aspect is given by the quantity ϕ_k, or $\theta_2(k - k^*) + \widetilde{\gamma}_k$, which is equal to a single value for a given cohort k. Lastly, "adjusted" refers to cohort careers incorporating all of the age, period, and cohort nonlinearities in the diachronic L-APC model, which will generally differ from those calculated using the raw (or "unadjusted") summaries in the data.[22]

Having noted the above, one can view Ryder's terminology in the first column of Table 5 as a general cohort-based, descriptive framework for interpreting the estimates from the classical APC model. Distinguishing between intra- and inter-cohort trends, and comparing the careers of successive cohorts, is crucial for relating the conventional APC model to general theories of life-cycle and social change popular in sociology and demography (e.g., see Elder 1974; Elder and George 2016; Goldstone 1991; Inglehart 1990; Mannheim [1927/1928] 1952; Tilly 1984; see also Ryder 1965, 1968, 1992). The more specific terminology in the second column of Table 5, however, is useful for clarifying the differences among the various kinds of intra- and inter-cohort trends, comparative cohort careers, and intra-period differences, as well as to link specific language to the core elements of each parametric expression, namely, what is calculated, what is compared, and what is conditioned upon. Following this framework, in the subsequent sections I use the data on social trust to examine the four main kinds of summaries that can be derived from a Lexis table (cf. the four panels of Table 5): first, intra-cohort trends (i.e., life-cycle change); second, inter-cohort trends (i.e., social change); third, Ryderian comparative cohort careers, which combine life-cycle with social change; and, lastly, intra-period differences. The first three types of summaries constitute the core of Ryder's cohort approach. (In Appendix C of the online supplement I also discuss in detail adjusted age, period, and cohort marginal curves, which condition on all of the nonlinearities included in the diachronic L-APC model.) I review each of these in turn, presenting both mathematical formulas and novel visualizations using the data on social trust.

Intra-cohort Trends (Life-Cycle Change)

The top panel of Table 5 lists the set of expressions for describing intra-cohort trends or, equivalently, life-cycle change. The simplest representation of an intra-cohort trend is the diachronic age slope, which is given by $\theta_1 = \alpha + \pi$ in Equation (7). However, a more realistic depiction of an intra-cohort trend would incorporate the age nonlinearities. The *diachronic age curve*, which adds the age nonlinearities to the diachronic age slope, is defined as

$$\theta_1(i - i^*) + \widetilde{\alpha}_i \text{ for } i = 1, \ldots, I, \tag{10}$$

which again is an overall measure of how cohorts age through time. Figure 4 shows the predicted probability of social trust using the diachronic age and cohort curves. As shown in panel (a), the diachronic age slope indicates that there is a general decline in the predicted probability of social trust as we compare age groups through time. This decline in predicted probability reflects the large, negative diachronic age slope in Table 3 (log-odds ratio: -0.192; $p = 0.001$). Panel (b), which displays the diachronic age curve, indicates that the predicted probability of trust increases in early adulthood and declines in later adulthood, with a peak in middle age. Together these graphs reveal considerable life-cycle change within cohorts.

So far we have considered only adding the age nonlinearities to the diachronic age slope, which results in the diachronic age curve. For an even richer depiction of the trends, one can also incorporate the period nonlinearities. The *age–period Lexis surface*, which adds the age nonlinearities to the diachronic age curve, is defined as[23]

$$\big(\alpha(i - i^*) + \widetilde{\alpha}_i\big) + \big(\pi(j - j^*) + \widetilde{\pi}_j\big) = \theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j \tag{11}$$

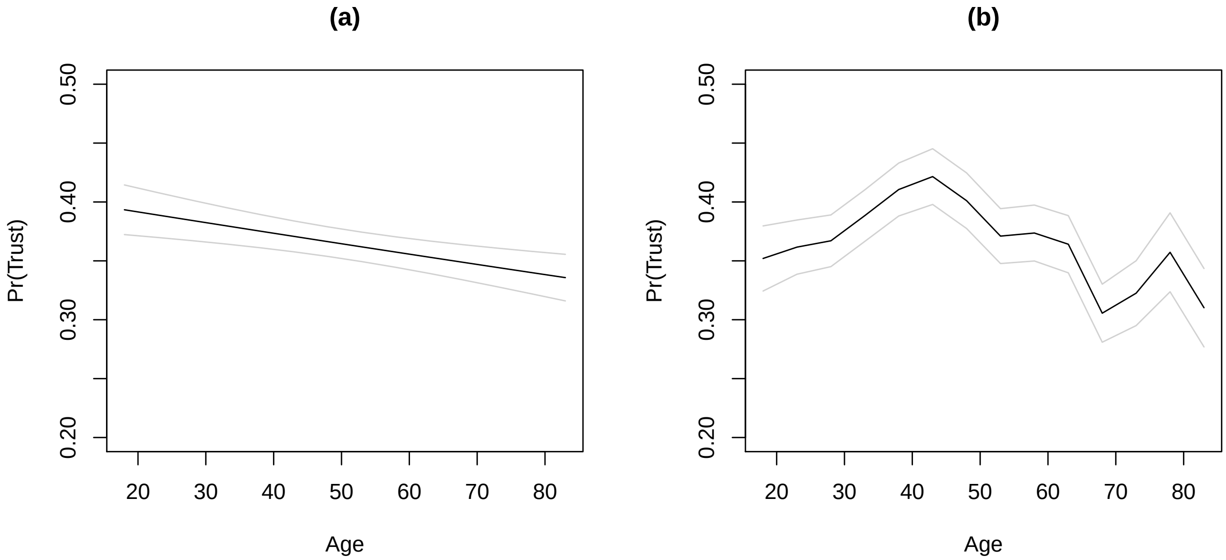

Figure 4: Diachronic age slope and curve. *Notes:* (a) displays the diachronic age slope, whereas (b) displays the diachronic age curve. Calculations are based on $\theta_1(i - i^*)$ and $\theta_1(i - i^*) + \widetilde{\alpha}_i$, respectively, for $i = 1, \ldots, I$. Parameter estimates are derived from a diachronic L-APC logistic regression model. Upper and lower bounds denote 95 percent confidence intervals. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

for all observed combinations of $i = 1, \ldots, I$ and $j = 1, \ldots, J$. Equation (11) summarizes a surface that varies across levels of age and period, but not cohort. Accordingly, this surface captures life-cycle change, or, equivalently, a set of intra-cohort trends. Figure 5 displays two- and three-dimensional representations of the age–period Lexis surface. The surface in both panels reflects the pattern of the diachronic age curve in Figure 11(b), with trust peaking in middle adulthood and declining thereafter. However, there is now additional variability across periods, with lower levels of trust in particular periods, such as 1975 to 1979 and 1990 to 1994. Despite this variability, however, it is clear that the diachronic age curve dominates the patterns in the data.

Some care, however, is required when interpreting the results in Figure 5. The age–period Lexis surface provides quite different summaries depending on whether or not one compares predicted probabilities across age groups within periods (i.e., columns) or cohorts (i.e., diagonals). Consider first the comparison of probabilities across successive age groups within each period. Typically comparing age groups within a period would be highly problematic, because this would imply a set of static rather than dynamic differences. However, the age–period Lexis surface shown in Figure 5 is calculated such that the diachronic cohort slope is zero. As discussed in Appendix A of the online supplement, when the diachronic cohort slope is zero, the synchronic age slope is identical to the diachronic age slope. Thus, the probabilities of successive ages within a particular period of the age–period Lexis surface in Figure 5 are equivalent to those from a diachronic age curve with a period-specific nonlinearity.[24]

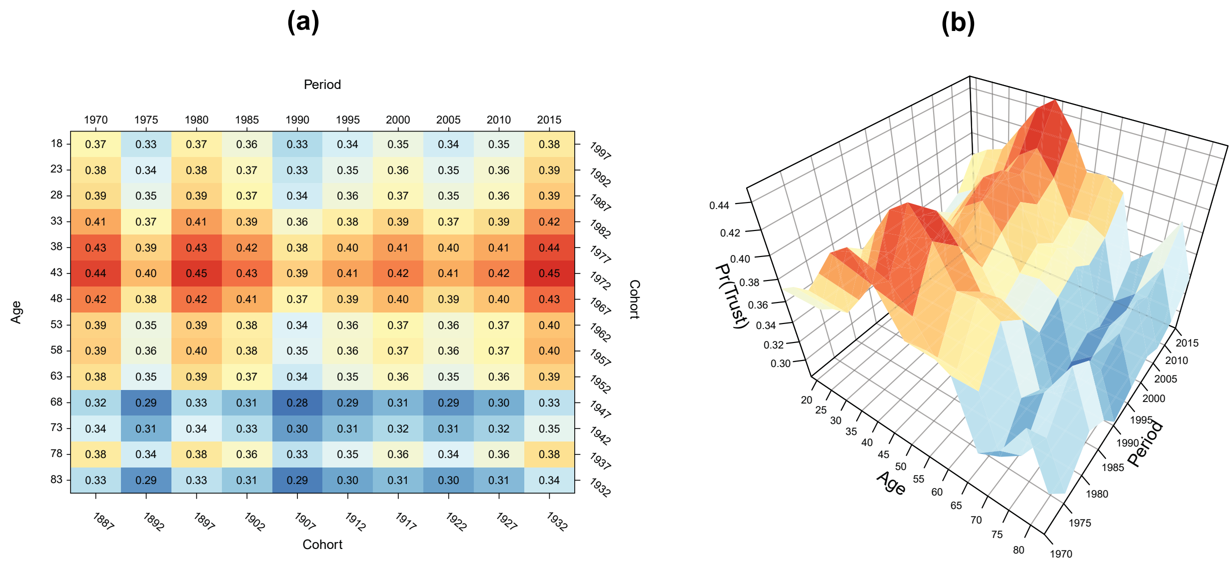

Figure 5: Two- and three-dimensional age–period Lexis surfaces. *Notes:* (a) and (b) display two-dimensional and three-dimensional Lexis surfaces, respectively, of predicted probabilities of the outcome for all observed combinations of age and period after adjusting for cohort. Calculations are based on $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j$ for all observed combinations of $i = 1, \ldots, I$ and $j = 1, \ldots, J$. All estimates are derived from a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

Besides comparing probabilities across successive age groups within each period, one can also do so within each cohort. These probabilities are equivalent to those from *local diachronic age curves*, which are defined as

$$\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I} \quad \text{for } i = 1, \ldots, I \text{ in each cohort } k, \tag{12}$$

where the period nonlinearities are indexed by both age and cohort. Equation (12) summarizes life-cycle change, or the intra-cohort trend, for each cohort; it is in this sense that the curves are "local." The local diachronic age curves will usually differ from the overall diachronic age curve, given by Equation (10), in two respects. First, the period nonlinearities $\widetilde{\pi}_{i+k-I}$, which are indexed by cohort (as well as age), will usually alter, sometimes quite considerably, the shape of each cohort-specific diachronic age curve. Second, even in the absence of period nonlinearities, most cohorts will typically be represented by only a section of the overall diachronic age curve. The reason is that, assuming the data take the form of an age–period Lexis table, the number of age groups will vary systematically across cohorts, with the first and last cohorts consisting of just a single age group.

Figure 6 shows the local diachronic age curves for social trust. Panel (a) displays a spaghetti plot of local diachronic age curves, whereas panel (b) displays the local diachronic age curves overlaid with the (overall) diachronic age curve. Only cohorts with three or more age groups are displayed. To reveal more clearly the variability of the intra-cohort trends, Figure 7 displays a trellis plot of the local

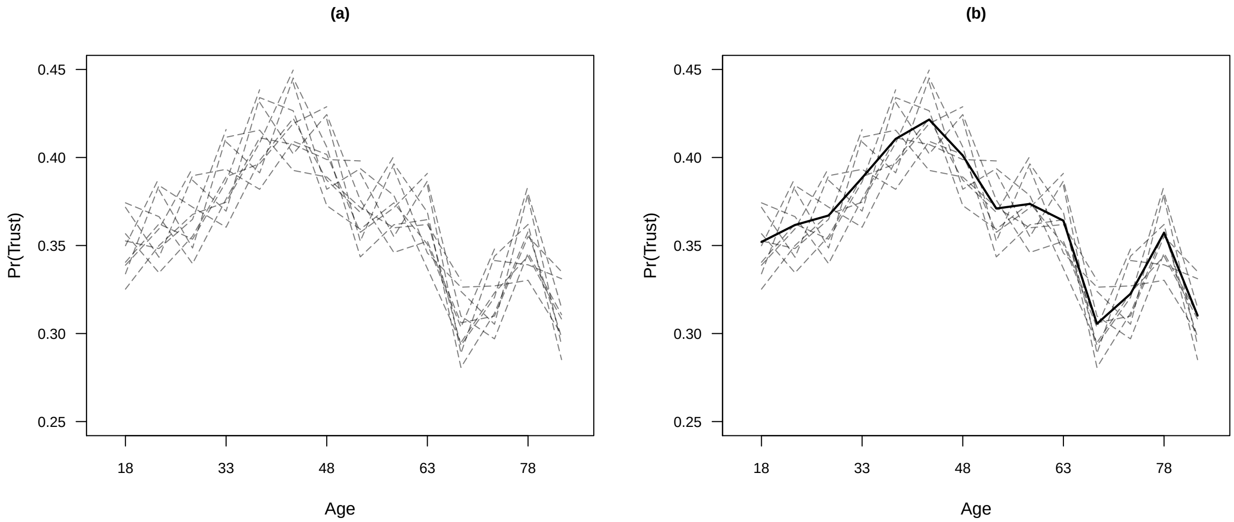

Figure 6: Spaghetti plot of local diachronic age curves. *Notes:* (a) displays a spaghetti plot of local diachronic age curves. (b) displays local diachronic age curves overlaid with the (overall) diachronic age curve. Calculations for the local curves are based on $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$ for $i = 1, \ldots, I$ in each cohort k. Local curves are shown with dashed lines, whereas the overall curve is shown with a bold solid line. Only cohorts with three or more age groups are displayed. All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

diachronic age curves for social trust. Each panel displays the (overall) diachronic age curve overlaid with local diachronic age curves. Note that the local curves are shown with solid lines, whereas the overall curve is shown with a dashed line. Again, only cohorts with three or more age groups are displayed. Results show that cohorts exhibit considerable variation in intra-cohort trends, with some cohorts experiencing steep declines and others relatively large increases.

As shown in the top panel of Table 5, there are, in short, a number of parametric expressions for representing intra-cohort trends, or life-cycle change. The diachronic age slope and curve are parsimonious representations that reflect overall changes as we compare age groups through time. These entail straightforward visualizations as line graphs. Incorporating the period nonlinearities, which reflect specific temporal contexts, requires calculating and visualizing either an age–period Lexis surface or a set of local diachronic age curves. Both approaches are helpful in revealing the unique life-cycle changes of particular cohorts.

Inter-cohort Trends (Social Change)

The second panel of Table 5 lists the set of expressions for describing inter-cohort trends or, equivalently, social change. The simplest representation of an inter-cohort trend is the diachronic cohort slope, which is denoted by $\theta_2 = \gamma + \pi$ in Equation (7). Although this expression is appealing in its simplicity, a more informative characterization of an inter-cohort trend would include the cohort nonlinearities. The *diachronic cohort curve*, which overlays the diachronic cohort slope with cohort

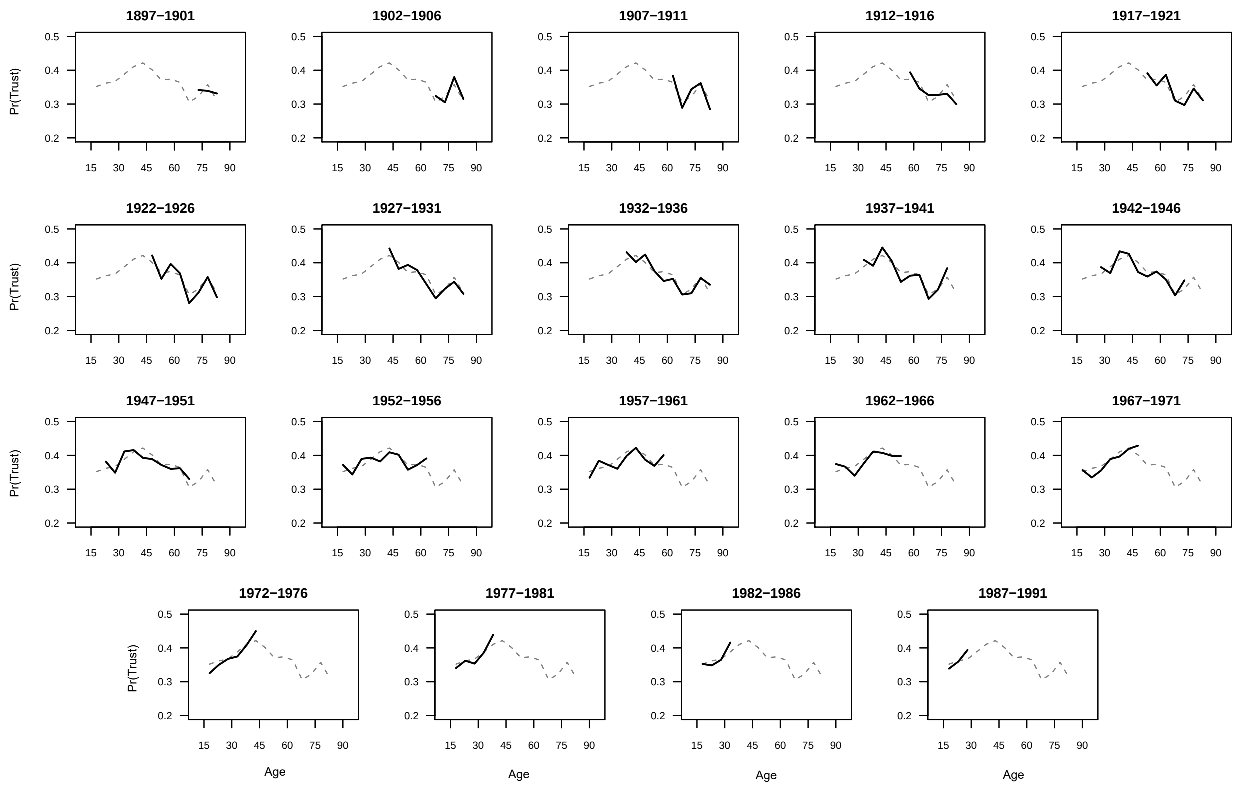

Figure 7: Trellis plot of local diachronic age curves. *Notes:* Each panel displays the (overall) diachronic age curve overlaid with local diachronic age curves. Horizontal axes are age in years, whereas vertical axes are predicted probability of social trust. Calculations for the local curves are based on $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$ for $i = 1, \ldots, I$ in each cohort k. Local curves are shown with solid lines, whereas the overall curve is shown with a dashed line. Only cohorts with three or more age groups are displayed. All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

nonlinearities, is defined as

$$\theta_2(k - k^*) + \widetilde{\gamma}_k \text{ for } k = 1, \ldots, K, \tag{13}$$

which is an overall measure of how cohorts differ through time. That is, Equation (13) represents social change as distinct from life-cycle change. Figure 8 displays the diachronic slope and cohort curves for social trust. Panel (a) shows a steep decline in the predicted probability of social trust, reflecting the large, negative diachronic cohort slope in Table 3 (log-odds ratio: -0.677; $p < 0.001$). Panel (b) shows that, as we compare cohorts through time, the predicted probability of social trust exhibits trendless fluctuation but then declines precipitously, especially among successive cohorts born after World War II. In summary, then, Figure 8 provides clear evidence of substantial social change, supporting the claim that there has been a collapse in social trust (e.g., see Robinson and Jackson 2001).

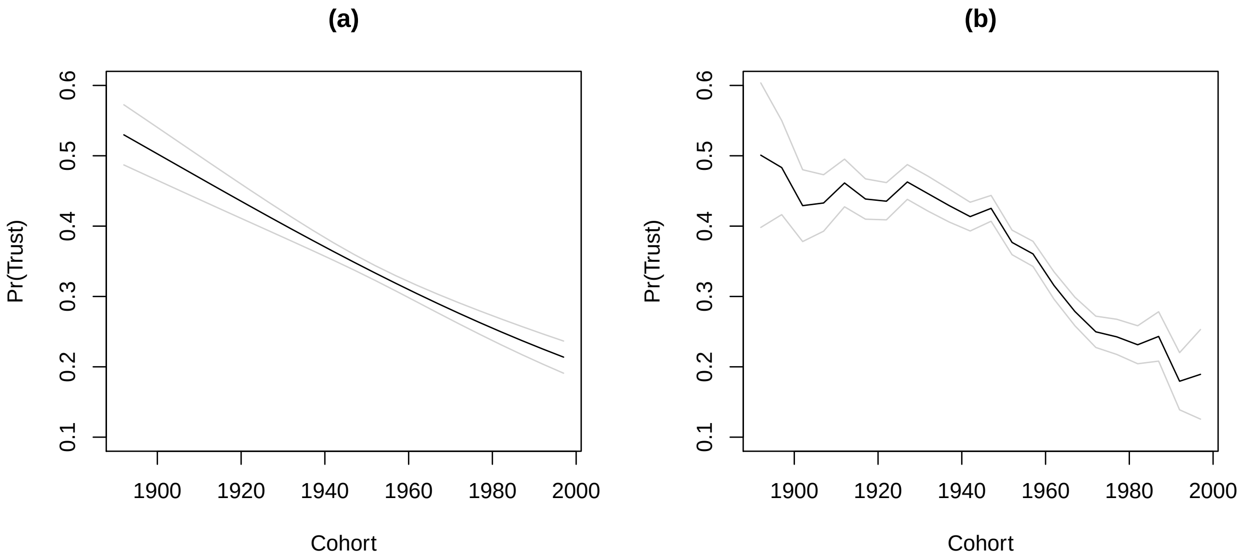

Figure 8: Diachronic cohort slope and curve. *Notes:* (a) displays the diachronic cohort slope, whereas (b) displays the diachronic cohort curve. Calculations are based on $\theta_2(k-k^*)$ and $\theta_2(k-k^*)+\widetilde{\gamma}_k$, respectively, for $k=1,\ldots,K$. Parameter estimates are derived from a diachronic L-APC logistic regression model. Upper and lower bounds denote 95 percent confidence intervals. Outcome is social trust, coded so that $1=$ "can trust" while $0=$ "can't be too careful" or "depends." Results adjusted using sampling weights.

By including the cohort nonlinearities, the diachronic cohort curve provides a more informative summary of social change than the diachronic cohort slope. However, for a more detailed representation of inter-cohort trends, one can include the period nonlinearities as well as those for cohort. The *cohort–period Lexis surface*, which summarizes social change in terms of joint cohort–period parameters, is defined as[25]

$$\left(\gamma(k-k^*)+\widetilde{\gamma}_k\right)+\left(\pi(j-j^*)+\widetilde{\pi}_j\right)=\theta_2(k-k^*)+\widetilde{\gamma}_k+\widetilde{\pi}_j \tag{14}$$

for all observed combinations of $k=1,\ldots,K$ and $j=1,\ldots,J$. Equation (14) defines a surface that varies across levels of cohort and period, but not age. Thus, this surface reflects social change, or, equivalently, a set of inter-cohort trends. Figure 5 displays two- and three-dimensional cohort–period Lexis surfaces. The surface in both panels reflects the pattern of the diachronic cohort curve in Figure 8(b), with trust dropping steeply across cohorts born after World War II. Additional variability is introduced by the period nonlinearities, but this heterogeneity is negligible compared with the overwhelmingly steep decline observed as we compare cohorts across periods.

Similar to the age–period Lexis surface, the cohort–period Lexis surface provides somewhat different summaries depending on whether or not one compares predicted probabilities across cohorts within periods (i.e., columns) or ages (i.e., rows). Consider first the successive comparison of cohorts within a given period. Usually comparing cohorts within a period would result in a set of static differences, thereby giving a distorted view of social change. However, the cohort–period Lexis

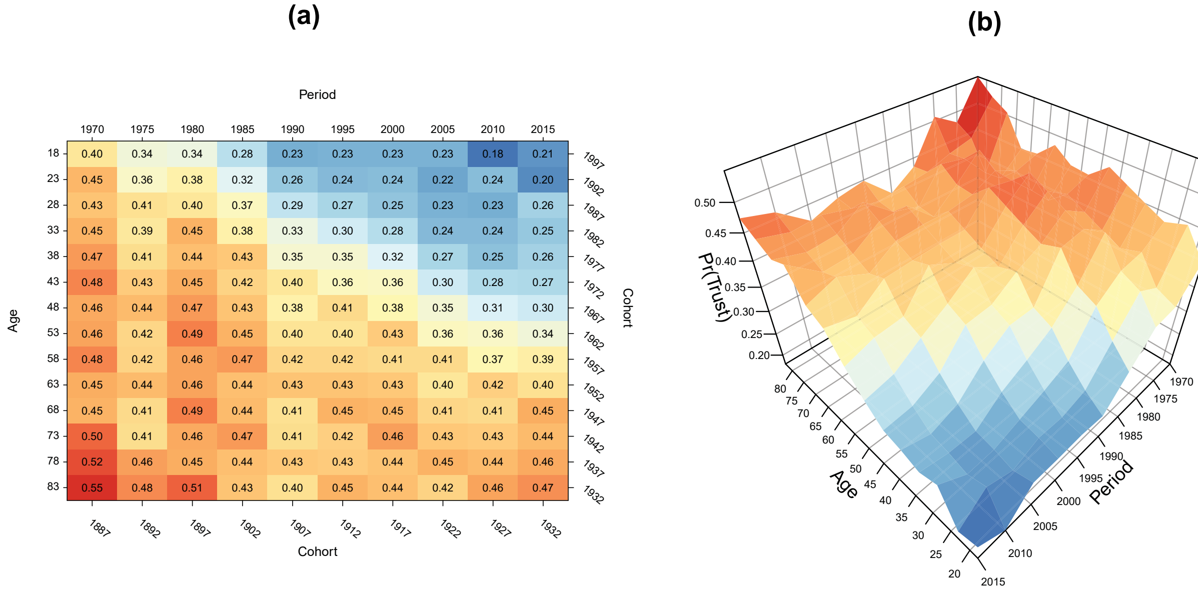

Figure 9: Two- and three-dimensional cohort–period Lexis surfaces. *Notes:* (a) and (b) display two-dimensional and three-dimensional Lexis surfaces, respectively, of predicted probabilities of the outcome for all observed combinations of cohort and period after adjusting for age. Calculations are based on $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_j$ for all observed combinations of $k = 1, \ldots, K$ and $j = 1, \ldots, J$. All estimates are derived from a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

surface in Figure 9 is calculated such that the diachronic age slope is zero. However, as detailed in Appendix A of the online supplement, when the diachronic age slope is zero, the synchronic cohort slope equals the diachronic cohort slope. Thus, the probabilities of successive cohorts within a given period of the cohort–period Lexis surface in Figure 9 are equivalent to those from a diachronic cohort curve, albeit with a single period nonlinearity.[26]

Besides comparing the probabilities in Figure 9 across cohorts within periods, one can also compare the probabilities across cohorts within ages. These probabilities are equivalent to those from *local diachronic cohort curves*, which are defined as

$$\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I} \quad \text{for } k = 1, \ldots, K \text{ in each age } i, \tag{15}$$

where again the period nonlinearities are indexed by both cohort and age. These curves are "local" in that they summarize the pattern of social change for each age group. Similar to the local diachronic age curves, the local diachronic cohort curves will usually differ from the overall diachronic cohort curve shown in Equation (13). This is because the the period nonlinearities $\widetilde{\pi}_{k+i-I}$, which are indexed by age (as well as cohort), will typically alter the shape of each age-specific diachronic cohort curve. Moreover, because we are calculating the curve for each age group, the length of each local diachronic curve will equal the number of period groups, not the number of cohort groups. Thus, for younger age groups, the local diachronic cohort curve will overlap a more recent section of the overall diachronic cohort

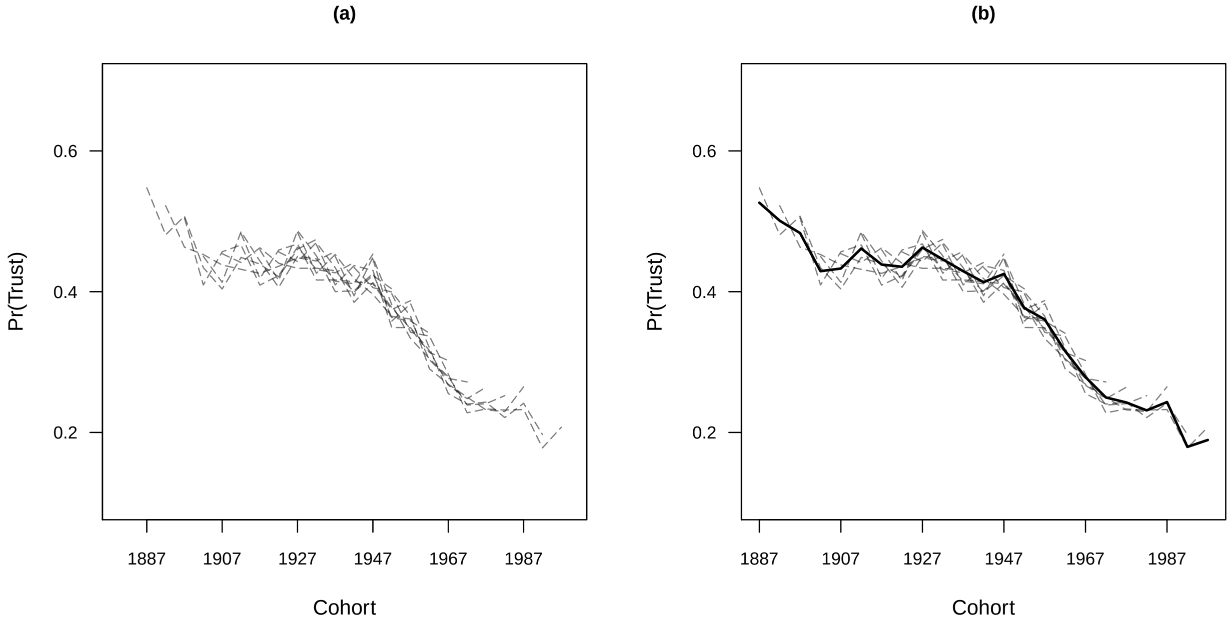

Figure 10: Local diachronic cohort curves. *Notes:* (a) displays a spaghetti plot of local diachronic cohort curves. (b) displays local diachronic cohort curves with the (overall) diachronic cohort curve overlaid. Calculations are based on $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i-I}$ for $k = 1, \ldots, K$ in each age group i. Local curves are shown with dashed lines, whereas the overall curve is shown with a bold solid line. Each local curve has a length equal to the number of periods in the data ($J = 10$). All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

curve. For older age groups, the local diachronic cohort curve will overlap an older section of the overall diachronic cohort curve.

Figure 10 shows the local diachronic cohort curves for social trust. Panel (a) displays a spaghetti plot of local diachronic cohort curves, whereas panel (b) displays the local diachronic cohort curves overlaid with the (overall) diachronic cohort curve. Calculations are based on $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{k+i}$ for $k = 1, \ldots, K$ in each age group i. Local curves are shown with dashed lines, whereas the overall curve is shown with a bold solid line. Each local curve has a length equal to the number of periods in the data ($J = 10$). To reveal more detail, Figure 11 displays a trellis plot of the local diachronic cohort curves. Each panel displays the (overall) diachronic cohort curve overlaid with local diachronic cohort curves. Local curves are shown with solid lines, whereas the overall curve is shown with a dashed line. Findings show that there is only relatively modest variation in age-specific social change.

In summary, as shown in the middle panel of Table 5, there are several parametric expressions, of varying complexity, for representing inter-cohort trends (or social change). The diachronic cohort slope and curve parsimoniously summarize the social change that is observed as we compare successive cohorts through time. These summaries are easily visualized as line graphs. Including the period nonlinearities, however, requires calculating a cohort–period Lexis surface, which can be viewed as a set of local diachronic cohort curves. Both the cohort–period Lexis surface and

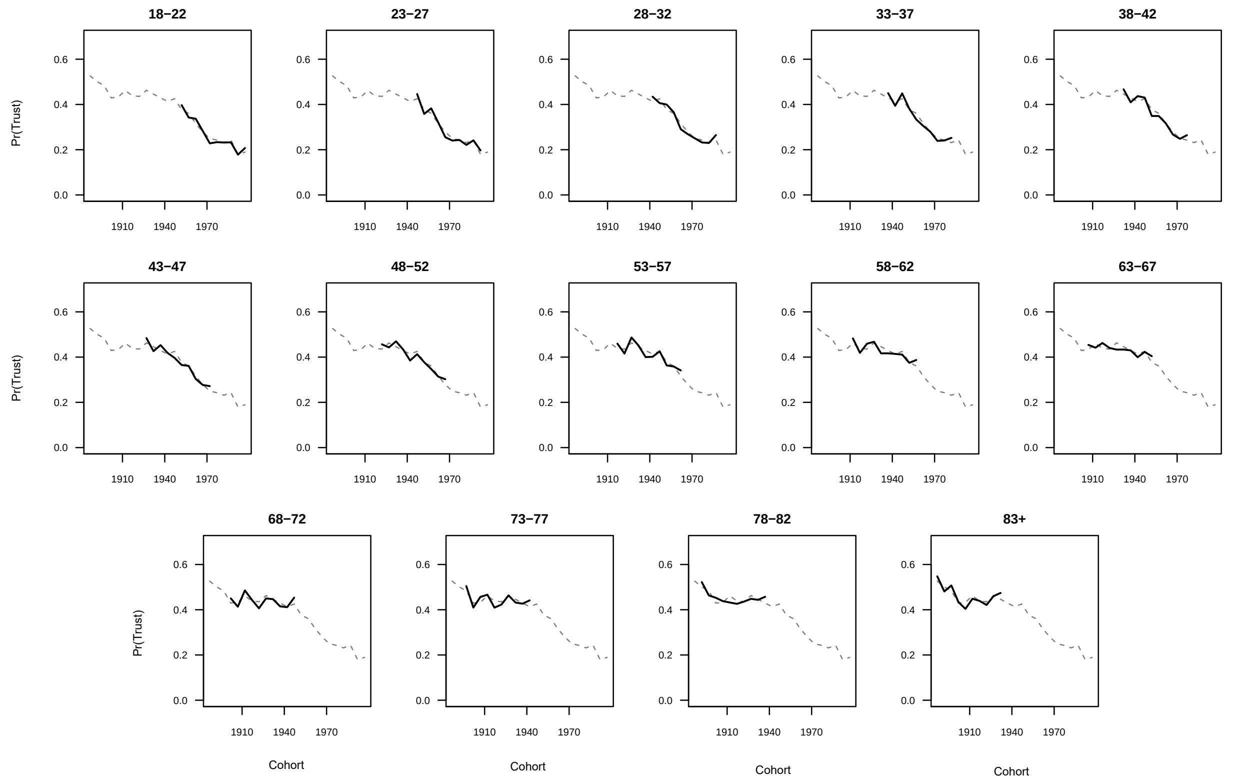

Figure 11: Trellis plot of local diachronic cohort curves. *Notes:* Each panel displays the (overall) diachronic cohort curve overlaid with local diachronic cohort curves. Horizontal axes are age in years, whereas vertical axes are predicted probability of social trust. Calculations are based on $\theta_2(k-k^*)+\widetilde{\gamma}_k+\widetilde{\pi}_{k+i-I}$ for $k=1,\ldots,K$ in each age group i. Local curves are shown with solid lines, whereas the overall curve is shown with a dashed line. Each local curve has a length equal to the number of periods in the data ($J=10$). All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that $1=$ "can trust" while $0=$ "can't be too careful" or "depends." Results adjusted using sampling weights.

local diachronic curves are crucial in revealing the pattern of social change unique to specific age groups.

Ryderian Comparative Cohort Careers

The third panel of Table 5 lists two expressions for Ryderian comparative cohort careers, which are composed of intra- and inter-cohort trends, or, equivalently, life-cycle and social change. Using the diachronic L-APC model, *adjusted comparative cohort careers* are defined as

$$\overbrace{\phi_k + \widetilde{\pi}_{i+k-I}}^{\substack{\text{Inter-cohort trend} \\ \text{(social change)}}} + \theta_1(i-i^*) + \widetilde{\alpha}_i \quad \text{for } i=1,\ldots,I \text{ in each cohort } k, \tag{16}$$

where the intra-cohort trend (life-cycle change) is $\underbrace{\widetilde{\pi}_{i+k-I} + \theta_1(i-i^*) + \widetilde{\alpha}_i}_{\text{Intra-cohort trend (life-cycle change)}}$.

where $\phi_k = \theta_2(k - k^*) + \widetilde{\gamma}_k$. The collection of parameters $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$, when calculated over age levels $i = 1, \ldots, I$ for each cohort k, represents a set of intra-cohort trends (i.e., life-cycle change), whereas the collection of parameters $\phi_k + \widetilde{\pi}_{i+k-I}$, when compared across K cohorts, represents a set of inter-cohort trends (i.e., social change). Together, these parameters define cohort-specific age trajectories through time, or comparative cohort careers. Each trajectory is a "career" in that it reflects the development of a cohort as it ages through time, but also "comparative" in that it is uniquely differentiated due to its particular location in time, which distinguishes it from all other cohorts.

To clarify how Equation (16) operates, it is helpful to break it down into three separate parts with three different functions. First, the parameters $\theta_1(i - i^*) + \widetilde{\alpha}_i$ define an overall curve that is invariant to the value of k. Second, ϕ_k, or $\theta_2(k - k^*) + \widetilde{\gamma}_k$, determines the level of the kth cohort's career. That is, for any given cohort k, the parameter ϕ_k is just a single value, so it effectively acts as an intercept that varies across cohorts. Finally, the period fluctuations, denoted by $\widetilde{\pi}_{i+k-I}$, have a dual role, contributing both to variability within a cohort as well as variability across cohorts. This is reflected in the fact that the $\widetilde{\pi}_{i+k-I}$ parameters are indexed by both i and k, or age and cohort.

The $\widetilde{\pi}_{i+k-I}$ parameters in Equation (16) allow the shape of a career to vary in complex ways across cohorts. In some instances admitting such heterogeneity may be desirable, but in other applications it may obscure the underlying pattern of life-cycle change across cohorts that is, from a Ryderian perspective, of fundamental analytic interest.[27] Researchers desiring a more parsimonious representation can focus on interpreting the curves-only comparative careers that are purged of period fluctuations. As shown in the third panel of Table 5, *curves-only comparative cohort careers* are defined as

$$\overbrace{\phi_k}^{\substack{\text{Inter-cohort trend} \\ \text{(social change)}}} + \underbrace{\theta_1(i - i^*) + \widetilde{\alpha}_i}_{\substack{\text{Intra-cohort trend} \\ \text{(life-cycle change)}}} \quad \text{for} \quad i = 1, \ldots, I \text{ in each cohort } k, \tag{17}$$

where $\phi_k = \theta_2(k - k^*) + \widetilde{\gamma}_k$. Again, the set of parameters $\theta_1(i - i^*) + \widetilde{\alpha}_i$, when calculated over age levels $i = 1, \ldots, I$ for each cohort k, represents an intra-cohort trend (i.e., life-cycle change), whereas ϕ_k, when compared across K cohorts, represents an inter-cohort trend (i.e., social change). Because the cohort careers are purged of period fluctuations, the shape of the curve is the same for all cohorts. However, if cohorts are unbalanced, as is typically the case with APC data in sociology and demography, most cohorts will be observed to experience only a section of the overall curve.

Careful inspection of Equations (16) and (17) reveal that they are generalizations of various diachronic curves. Regarding Equation (16), the comparative cohort career is a generalization of the local diachronic age and cohort curves. If $\theta_2(k - k^*) + \widetilde{\gamma}_k$ is zero, then Equation (16) is equal to the equation for the local diachronic age curve, or $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_{i+k-I}$, which is calculated over age levels $i = 1, \ldots, I$ for each cohort k. Similarly, if $\theta_1(i - i^*) + \widetilde{\alpha}_i$ is zero, then Equation (16) is equal to the equation for the local diachronic cohort curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_{i+k-I}$.

The main difference is that the local diachronic cohort curves are calculated over cohorts $k = 1, \ldots, K$ for each age group i, whereas comparative cohort careers are calculated over age levels $i = 1, \ldots, I$ for each cohort k. In a similar way Equation (17) can be viewed as a generalization of the diachronic age and cohort curves. If $\theta_2(k - k^*) + \widetilde{\gamma}_k$ is zero, then Equation (17) to equal to the diachronic age curve, or $\theta_1(i - i^*) + \widetilde{\alpha}_i$. Likewise, if $\theta_1(i - i^*) + \widetilde{\alpha}_i$ is zero, then Equation (17) is equal to that for the diachronic cohort curve, or $\theta_2(k - k^*) + \widetilde{\gamma}_k$.

Figure 12(a) displays the comparative cohort careers for social trust, whereas Figure 12(b) displays the curves-only comparative cohort careers. In each graph, the careers of the earliest cohorts are shown in the upper right, and the latest cohorts are shown in the lower left. The only difference between panels (a) and (b) is that the latter is purged of period fluctuations that vary across age and cohort groups. The vertical differences between careers reflect inter-cohort trends (i.e., social change), whereas the shapes of the careers reflect intra-cohort trends (i.e., life-cycle change). If there were no social change—that is, no differentiation as we compared cohorts through time—then all careers would collapse on top of one other and both graphs would resemble the diachronic age curve displayed in Figure 4(a). The only difference is that most cohorts would be some particular section of the diachronic age curve, because the duration of each cohort career is equal to the number of observed age categories, which is typically fewer than the actual number of age categories. If there were no life-cycle change—that is, no intra-cohort development as cohorts age through time—then all careers would be horizontal lines with a length equal to the kth cohort's number of observed age categories. This hypothetical scenario is shown in Figure 13 using data on social trust.

The vertical distances of the horizontal lines from each other are governed by the values of ϕ_k. For example, as shown in Figure 13, the $k = 14$ cohort, with a midpoint birth year of 1952, is represented by the parameter $\phi_{k=14}$ with a corresponding predicted probability of 0.38, whereas the $k = 16$ cohort, with a midpoint birth year of 1962, is represented by the parameter $\phi_{k=16}$ with a corresponding predicted probability of 0.32. Thus, the vertical distance between the careers of these two cohorts is given by $0.32 - 0.38 = -0.06$. Note that if there were neither life-cycle nor social change then the graph would simplify further, appearing as a single horizontal line equal to the predicted probability of the intercept. Most cohorts, however, would occupy just a particular section of the horizontal line, with a length again equal to the number of observed (rather than actual) age categories.

In short, the two expressions for Ryderian comparative cohort careers in Table 5 unite information on both intra- and inter-cohort trends, or life-cycle and social change. Cohort careers can be easily—yet powerfully—visualized using a series of line graphs, as shown in Figure 12. The vertical spread indicates the degree of social change, whereas the steepness of the curves reveals the extent of life-cycle change. In general, curves-only cohort careers can help highlight underlying trends and patterns that would otherwise be obscured by the inclusion of the period nonlinearities.

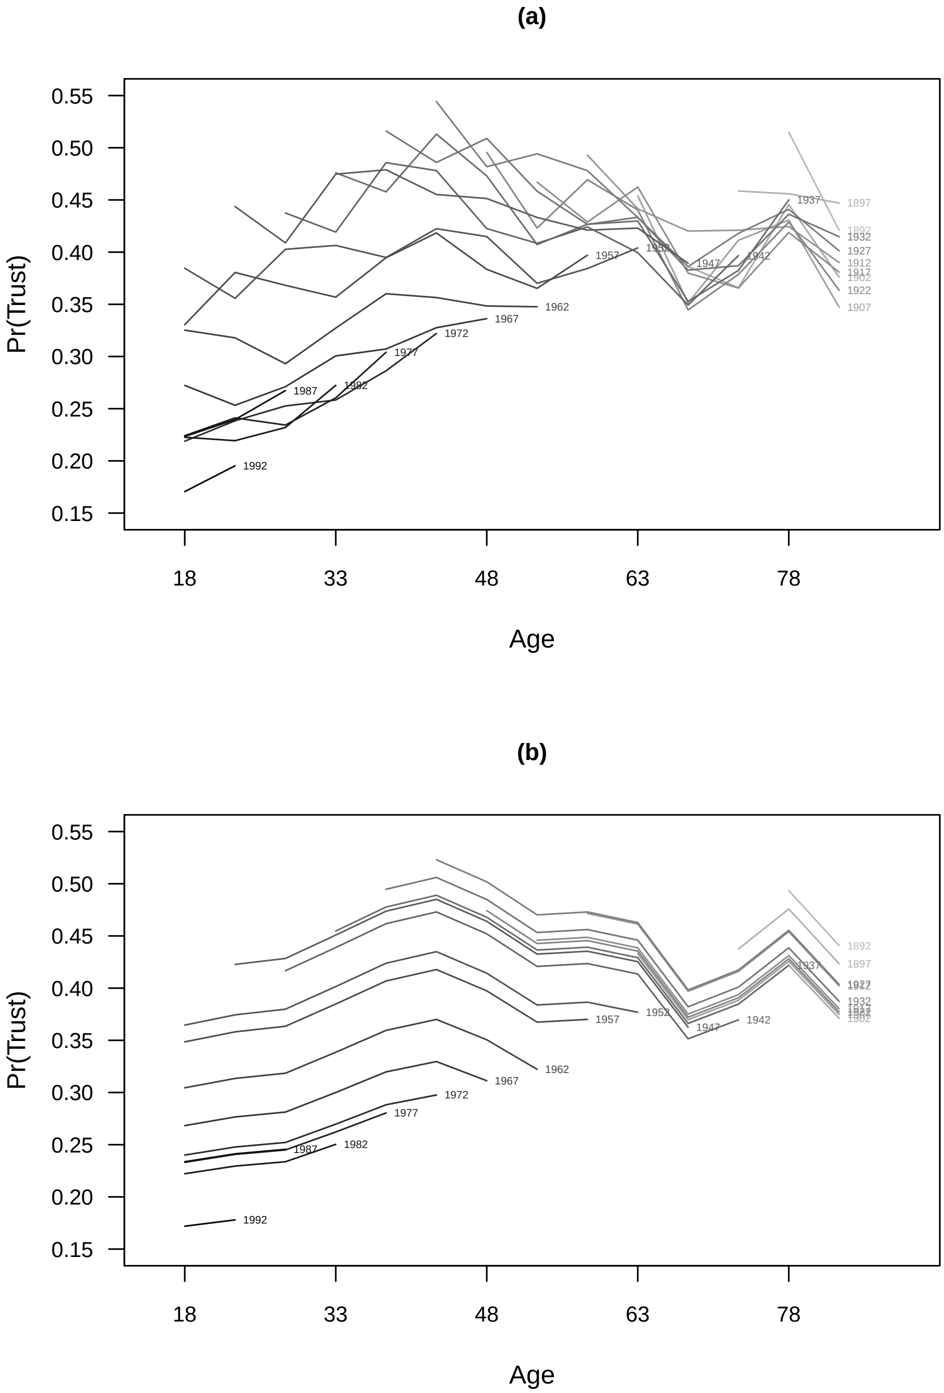

Figure 12: Ryderian comparative cohort careers. *Notes:* (a) displays comparative cohort careers, or $\phi_k + \widetilde{\pi}_{i+k-I} + \theta_1(i - i^*) + \widetilde{\alpha}_i$ for $i = 1, \ldots, I$ in each cohort k. (b) displays curves-only comparative cohort careers, or $\phi_k + \theta_1(i - i^*) + \widetilde{\alpha}_i$ for $i = 1, \ldots, I$ in each cohort k. Note that $\phi_k = \theta_2(k - k^*) + \widetilde{\gamma}_k$. Only cohorts with two or more age categories are displayed. All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

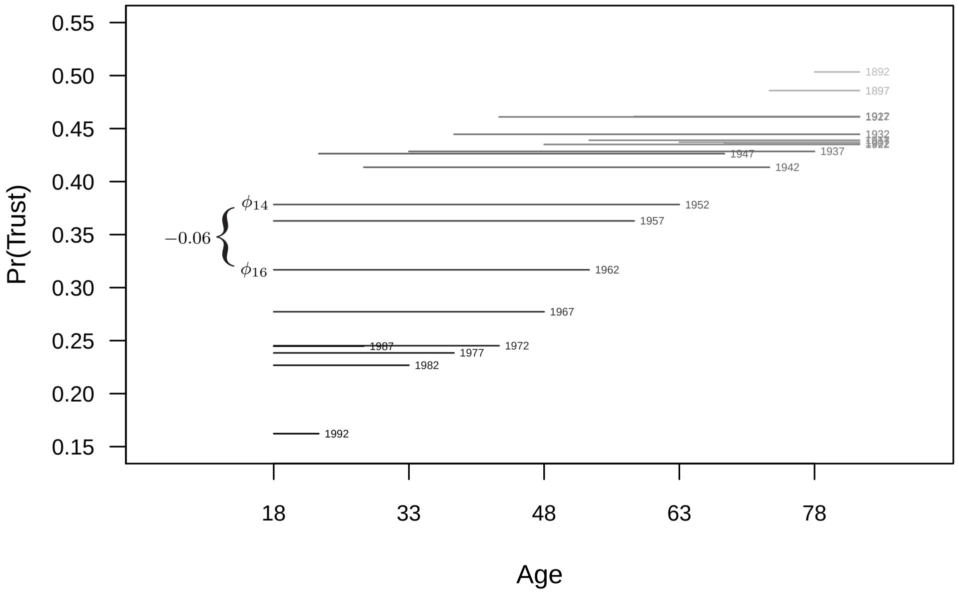

Figure 13: Comparative cohort careers in absence of life-cycle change. *Notes:* Graph displays comparative cohort careers under the hypothetical scenario of no life-cycle change (i.e., no intra-cohort trends). The vertical distance from the intercept for each horizontal line is given by $\phi_k = \theta_2(k - k^*) + \widetilde{\gamma}_k$. The length of each horizontal line is equal to the number of observed age categories. For example, the $k = 14$ cohort, with a midpoint birth year of 1952, has a parameter $\phi_{k=14}$ with a corresponding predicted probability of 0.38. By contrast, the $k = 16$ cohort, with a midpoint birth year of 1962, has a parameter $\phi_{k=16}$ with a corresponding predicted probability of 0.32. Thus, the vertical distance between the careers of these two cohorts is given by the difference in predicted probabilities, or $0.32 - 0.38 = -0.06$. Only cohorts with two or more age categories are displayed. Cohort groups are labeled using midpoint values. All estimates are based on a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

Intra-period Differences

Finally, the bottom panel of Table 5 displays various expressions for intra-period differences. These are not a core part of Ryder's approach but can provide insight, albeit indirectly, into intra- and inter-cohort trends (and thus life-cycle and social change). The simplest expressions are the age and cohort synchronic slopes, or $\theta_1 - \theta_2 = \alpha - \gamma$ and $\theta_2 - \theta_1 = \gamma - \alpha$, respectively. However, one can also estimate synchronic age and cohort curves by adding their respective nonlinearities. The *synchronic age curve* is given by

$$(\theta_1 - \theta_2)(i - i^*) + \widetilde{\alpha}_i \quad \text{for } i = 1, \ldots, I. \tag{18}$$

Similarly, the *synchronic cohort curve* is given by

$$(\theta_2 - \theta_1)(k - k^*) + \widetilde{\gamma}_k \quad \text{for } k = 1, \ldots, K. \tag{19}$$

Both curves, because they are estimated from models that adjust for the period linear component, are based on slopes that capture static differences rather than dynamic

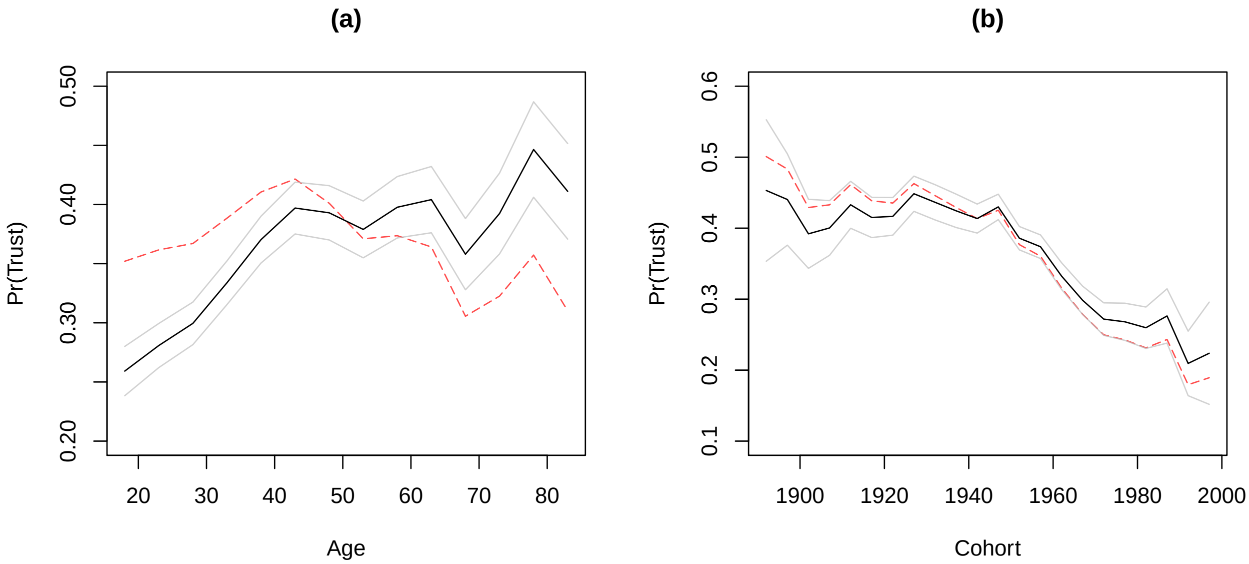

Figure 14: Synchronic age and cohort curves. *Notes:* (a) displays the synchronic age curve with an estimated slope of $\theta_1 - \theta_2 = 0.485$. Overlaid red dashed line denotes the diachronic age curve with an estimated slope of $\theta_1 = -0.192$. (b) displays the synchronic cohort curve with an estimated slope of $\theta_2 - \theta_1 = -0.485$. Overlaid red dashed line denotes the diachronic cohort curve with an estimated slope of $\theta_2 = -0.677$. Estimates are based on age and cohort synchronic L-APC logistic regression models, respectively. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

changes. Note again that the slopes of the curves are identical in magnitude but with opposing signs. That is, one just needs to multiply one of the slopes by negative one to obtain the other slope (e.g., $(\theta_1 - \theta_2)(-1) = (\theta_2 - \theta_1)$).

Panels (a) and (b) of Figure 14 display the synchronic age and cohort curves for the social trust data. Both curves are overlaid with their diachronic counterparts. The main problem with synchronic curves is that they are easily interpreted as dynamic comparisons rather than static differences. As discussed previously, using synchronic curves to estimate diachronic curves requires very strong assumptions on the absence of overall (linear) life-cycle or social change (see Table 2 and the accompanying discussion). In most applications these assumptions will not hold and synchronic curves will be biased estimates of diachronic curves. For example, as shown panel (a) of Figure 14, the slope of the synchronic age curve greatly overestimates the diachronic age curve, implying that cohorts of people are generally more trusting as they age through time. This is not true, however, as shown by the dashed line, which denotes the diachronic age curve. Likewise, as shown in panel (b) of Figure 14, the slope of the synchronic cohort curve underestimates the decline in social trust across cohorts.[28]

A more principled way of examining synchronic age and cohort measures is to acknowledge that they reflect differences between designated levels of age and cohort within a given period. Age–cohort comparisons within a particular period, or *local synchronic age differences*, are defined as

$$\alpha_i - \gamma_{j-i+I} = (\theta_1 - \theta_2)(i - i^*) + (\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I}) \quad \text{for all } i \text{ in each period } j, \qquad (20)$$

where the index for cohort is a function of varying age and period levels such that $k = j - i + I$. Equation (20) can be interpreted as the difference, for designated pairs of age and cohort levels in a given period, between the diachronic age and cohort curves. Likewise, cohort–age comparisons within a particular period, or *local synchronic cohort differences*, are defined as

$$\gamma_k - \alpha_{j-k+I} = (\theta_2 - \theta_1)(k - k^*) + (\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I}) \quad \text{for all } k \text{ in each period } j, \qquad (21)$$

where the index for age is a function of varying period and cohort levels such that $i = j - k + I$. Similar to Equation (20), Equation (21) can be interpreted as the difference, for designated pairs of cohort and age levels in a given period, between the diachronic cohort and age curves.

The local synchronic age and cohort differences reveal the relative magnitude of life-cycle versus social change in each period.[29] If local synchronic age or cohort differences are zero, then this indicates that, for designated age–cohort comparisons, parameters representing life-cycle change are the same as those for social change in a given period. That is, the slopes of the local synchronic differences in Equations (20) and (21) will be zero if the diachronic age and cohort slopes are the same, such that $\theta_1 = \theta_2$ and thus $\theta_1 - \theta_2 = \theta_2 - \theta_1 = 0$. Likewise, the nonlinearities will be zero if designated pairs of age and cohort nonlinearities are the same, such that $\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I} = 0$ for a given age i and period j or $\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I} = 0$ for a given cohort k and period j.

However, if local synchronic age differences are positive (or negative), then this reveals that, for designated age–cohort comparisons, parameters representing life-cycle change are greater (or less) than those for social change in a given period. For example, the slopes of the local age synchronic differences in Equations (20) will be positive if the diachronic age slope is greater than the diachronic cohort slope, such that $\theta_1 > \theta_2$ and thus $\theta_1 - \theta_2 > 0$. Likewise, the nonlinearities will be positive if the designated age nonlinearities are greater than those for cohort, such that $\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I} > 0$ for a given age i and period j. By contrast, if local synchronic cohort differences are positive (or negative), then this indicates that, for designated cohort–age comparisons in a given period, parameters representing social change are greater (or less) than those for life-cycle change. For instance, the slopes of the local cohort synchronic differences in Equation (21) will be positive only if the diachronic cohort slope is greater than the diachronic age slope, such that $\theta_2 > \theta_1$ and thus $\theta_2 - \theta_1 > 0$. Similarly, the nonlinearities will be positive only if the designated cohort nonlinearities are greater than those for age, such that $\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I} > 0$ for a given cohort k and period j.

Figures 15 and 16 display trellis plots for local synchronic age and cohort differences, respectively. The top horizontal axes denote age in years, bottom horizontal axes denote cohort in years, and vertical axes indicate the difference in the predicted probability of social trust. The pairs of age and cohort levels that are compared vary based on age, which is a sliding window across periods. Verti-

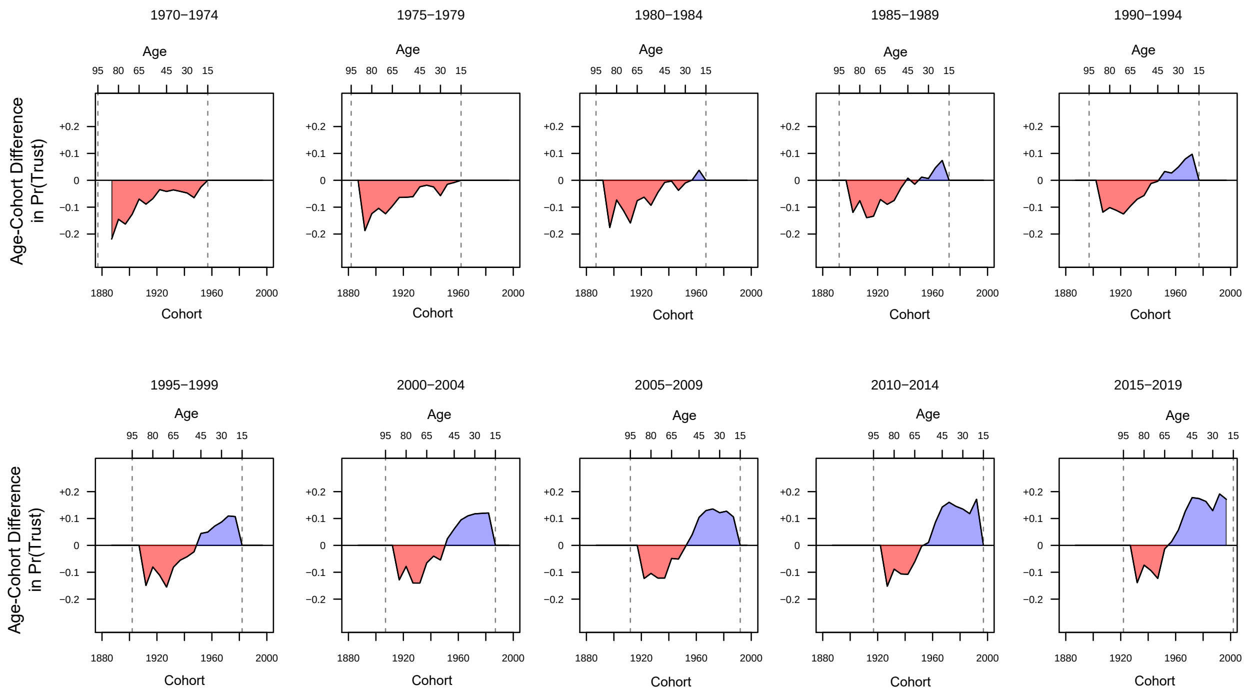

Figure 15: Trellis plot of local synchronic age differences. *Notes:* Each panel displays the difference in predicted probability of social trust for designated pairs of age and cohort levels. Calculations are based on $(\theta_1 - \theta_2)(i - i^*) + (\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I})$ for $i = 1, \ldots, I$ in each period j. Top horizontal axes are age in years, bottom horizontal axes are cohort in years, and vertical axes are the difference in the predicted probability of social trust. The pairs of age and cohort levels that are compared vary based on age, which is a sliding window across periods. Vertical dotted lines denote the beginning and end points of the age window in each period. Positive (or negative) values indicate that, for particular age–cohort comparisons in a given period, parameters representing life-cycle change are greater (or less) than those for social change. All estimates are based on the age synchronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

cal dotted lines denote the beginning and end points of the age window in each period. Each panel of Figure 15 displays the difference in predicted probability of social trust for designated pairs of age and cohort levels, with the calculations based on $(\theta_1 - \theta_2)(i - i^*) + (\widetilde{\alpha}_i - \widetilde{\gamma}_{j-i+I})$ for $i = 1, \ldots, I$ in each period j. Likewise, each panel of Figure 16 displays the difference in predicted probability of social trust for designated pairs of cohort and age levels, with the calculations based on $(\theta_2 - \theta_1)(k - k^*) + (\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I})$ for $k = 1, \ldots, K$ in each period j.

Figure 15 reveals that, for the defined age–cohort combinations, the predicted probability of social trust is lower for life-cycle change than social change for earlier periods, especially from 1970 to 1984. However, since the 1990s, the probability of social trust is much larger for life-cycle change than social change. This is particularly the case for younger age–cohort combinations, as indicated by the growing positive "bulge" across the panels of Figure 15. Figure 16, which is based on comparing cohort–age combinations, shows the inverse pattern of Figure 15. Specifically, in earlier periods the probability of social trust is greater for social than

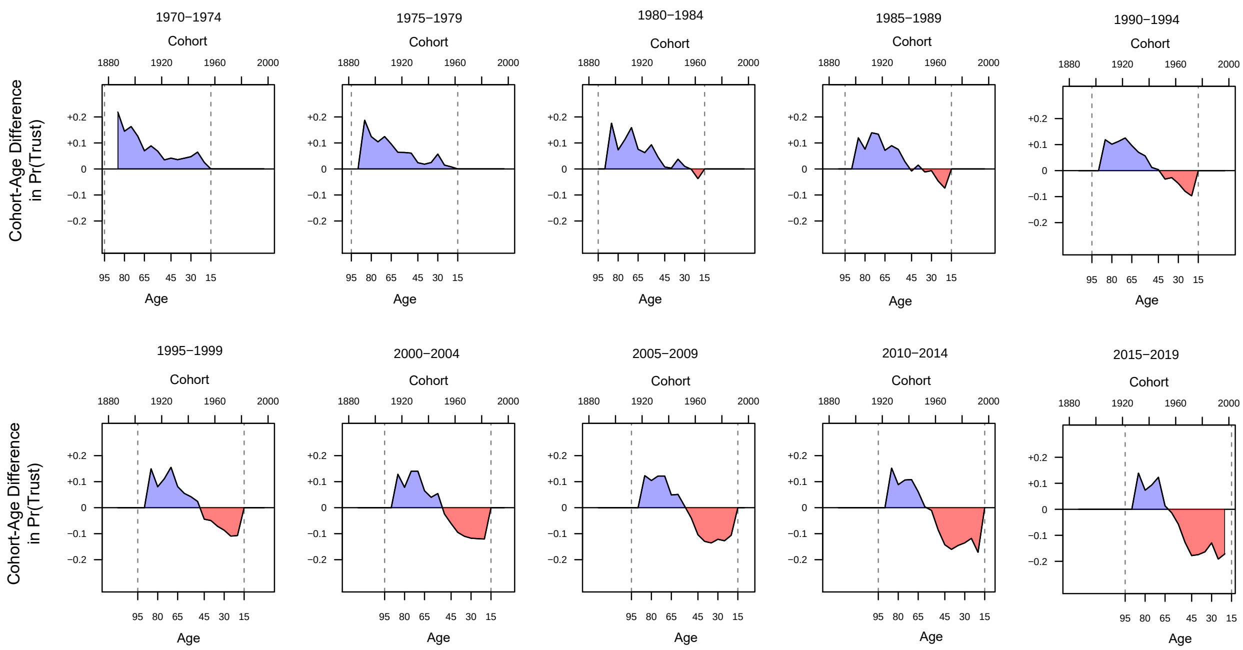

Figure 16: Trellis plot of local synchronic cohort differences. *Notes:* Each panel displays the difference in predicted probability of social trust for designated pairs of cohort and age levels. Calculations are based on $(\theta_2 - \theta_1)(k - k^*) + (\widetilde{\gamma}_k - \widetilde{\alpha}_{j-k+I})$ for $k = 1, \ldots, K$ in each period j. Top horizontal axes are cohort in years, bottom horizontal axes are age in years, and vertical axes are the difference in the predicted probability of social trust. The pairs of age and cohort levels that are compared vary based on age, which is a sliding window across periods. Vertical dotted lines denote the beginning and end points of the age window in each period. Positive (or negative) values indicate that, for particular cohort–age comparisons in a given period, parameters representing social change are greater (or less) than those for life-cycle change. All estimates are based on the cohort synchronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

life-cycle change, whereas in later periods the probability is lower for social than life-cycle change, as indicated by the growing negative "bulge." The results from Figures 15 and 16 are consistent with those in Figures 4(b) and 8(b), which reveal a relatively modest decline as cohorts age through time (i.e., life-cycle change) but a steep collapse as we compare cohorts through time (i.e., social change).

In summary, extreme care should be exercised when calculating and displaying synchronic age and cohort measures. Importantly, synchronic slopes and curves represent static intra-period differences, not dynamic trends. As shown in Figure 14, attempting to use synchronic measures in place of their diachronic counterparts will give a misleading account of life-cycle and social change. A more principled approach is to calculate local synchronic age and cohort differences. Although these represent static intra-period differences, by comparing these differences across periods, one can extract, albeit indirectly, some information about the relative magnitude of life-cycle and social change. These local differences can be visualized as a series of "bulges" that vary across periods, as shown in Figures 15 and 16.

Generalization of Related Approaches

The previous sections outlined how the classic APC model can be adapted for summarizing population-level temporal variability in keeping with Ryder's vision for cohort analysis. In this section I discuss three related approaches: first, the analysis and graphical presentation of APC nonlinearities (e.g., O'Brien 2015:110–12); second, a reformulation of the APC model by Duncan (1981); and, finally, a linear decomposition proposed by Firebaugh (1989). Each of these techniques can be viewed, in some sense, as special cases of the model-based approach discussed in this article.

Due to space constraints, I restrict my discussion to techniques derived from the conventional APC model, which is a "three-dimensional" model in that it incorporates all three temporal variables. Thus, I omit from consideration "two-dimensional" models, such as age–period models or, more generally, row–column models, including those with row–column interactions (e.g., Luo and Hodges 2022). Similarly, because they are not explicitly based on the classic APC model, I do not discuss the wide range of informal graphical techniques for summarizing APC data (for excellent reviews, see Robertson and Boyle [1998:1325–31] and Yang and Land [2013:56–61]). Formal analyses of the relations between three- and two-factor APC models as well as the relations between the parameters of the classic APC model and informal graphical techniques are likely fruitful topics for further research. The latter is especially important inasmuch as informal graphical displays, as opposed to graphical displays derived from well-defined parametric expressions, may be particularly misleading (for discussions on the limitations of informal graphical approaches, see Yang and Land [2013:59–60] and Holford [1991:426–28]).

Models of APC Nonlinearities

Various researchers have emphasized the importance of modeling APC nonlinearities (or deviations) apart from the linear components. As noted by some analysts, under the classic APC model the nonlinearities are identified and can, at least in some instances, reveal meaningful patterns in the data (Holford 1991:445–46; O'Brien 2015:110–12). Typically the nonlinearities for age, period, and cohort are presented separately along their respective dimensions, as shown in Appendix B of the online supplement. However, as demonstrated by Acosta and van Raalte (2019), additional insight can be gained by visualizing APC nonlinearities jointly on a Lexis surface.

Figure 17(a), for example, displays two-dimensional Lexis heat maps of the age and period nonlinearities for social trust, with calculations based on $\widetilde{\alpha}_i + \widetilde{\pi}_j$ for all observed combinations of $i = 1, \ldots, I$ and $j = 1, \ldots, J$. Likewise, Figure 17(b) displays a two-dimensional Lexis heat map of cohort and period nonlinearities, with calculations based on $\widetilde{\gamma}_k + \widetilde{\pi}_j$ for all observed combinations of $k = 1, \ldots, K$ and $j = 1, \ldots, J$. Results show considerable temporal patterning, with pronounced "bumps" in social trust for those individuals in middle age and for cohorts born in the middle of the twentieth century.

Three points are worth noting about the results in Figure 17. First, although informative in some examples (e.g., see Acosta and van Raalte 2019:1218–24), by

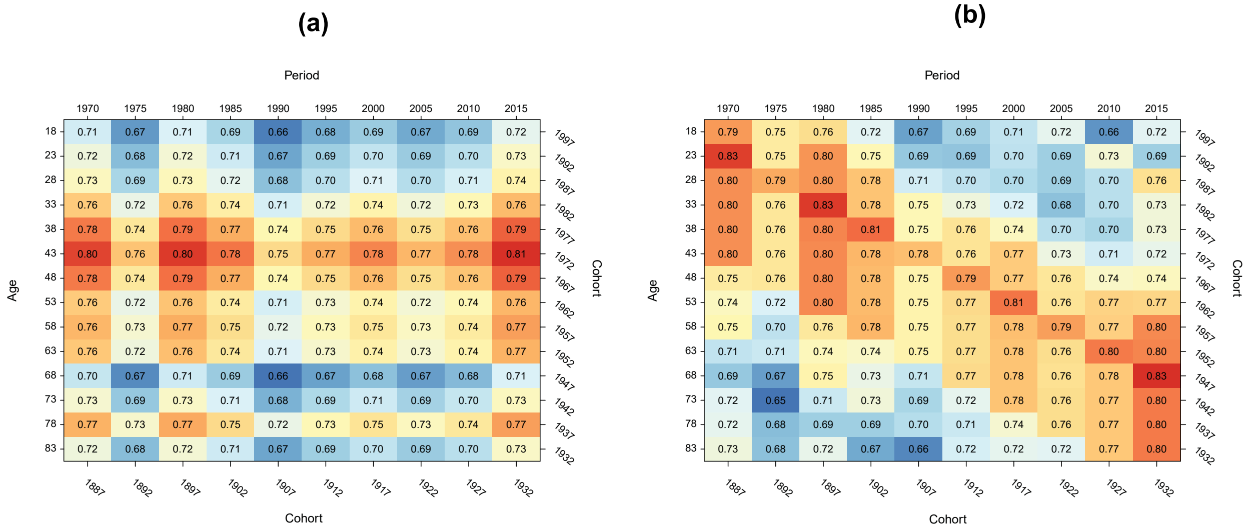

Figure 17: Two-dimensional Lexis heat maps of age–period and cohort–period nonlinearities. *Notes:* (a) displays a two-dimensional Lexis heat map of age and period nonlinearities, with calculations based on $\widetilde{\alpha}_i + \widetilde{\pi}_j$ for all observed combinations of $i = 1, \ldots, I$ and $j = 1, \ldots, J$. (b) displays a two-dimensional Lexis heat map of cohort and period nonlinearities, with calculations based on $\widetilde{\gamma}_k + \widetilde{\pi}_j$ for all observed combinations of $k = 1, \ldots, K$ and $j = 1, \ldots, J$. All estimates are derived from a diachronic L-APC logistic regression model. Outcome is social trust, coded so that 1 = "can trust" while 0 = "can't be too careful" or "depends." Results adjusted using sampling weights.

definition the nonlinearities in Figure 17 can reveal only local shifts, not overall changes (or trends). Second, the curvature plots in Figure 17 can be interpreted as special cases of the age–period and cohort–period Lexis surfaces discussed previously. Specifically, Figure 17(a) is identical to the age–period Lexis surface in Figure 5 except the diachronic age slope is zero; similarly, Figure 17(b) is identical to the cohort–period Lexis surface in Figure 9 except the diachronic cohort slope is zero. Finally, in principle there is additional variability beyond the conventional model that could be modeled and displayed in Figure 17.[30] This may render the model less parsimonious (because more parameters would be included) but possibly more informative depending on the particular substantive example. Extending the conventional APC model to explicitly incorporate such additional heterogeneity is a potentially promising area of methodological development.

Duncan's Reformulation of the APC Model

In an unpublished manuscript ("A Reformulation in the APC Model"), Otis Dudley Duncan (1981) suggested that the age–period parameters of the C-APC model could represent Ryder's concept of aggregate, cohort-specific life-cycle change and thereby enable the analysis of comparative cohort careers. Specifically, Duncan (1981:1) wrote that, "for the same cohort k at one time period (and age interval) earlier," the C-APC model is

$$y_{i-1,j-1,k} = \mu + \alpha_{i-1} + \pi_{j-1} + \gamma_k. \tag{22}$$

Duncan further observed that taking the difference between a cohort k at age i and period j versus age $i-1$ and $j-1$ yields

$$y_{ijk} - y_{i-1,j-1,k} = (\alpha_i - \alpha_{i-1}) + (\pi_j - \pi_{j-1}) \tag{23}$$

or

$$\Delta y_{ij} = \Delta\alpha_i + \Delta\pi_j \text{ for a given } k, \tag{24}$$

where Δ refers to first-differences between age and period levels. Duncan (1981) argued that this "serves to describe the 'trajectory' of the cohort entering the population at time $j = i + k$" and that the "trajectory has a remarkably simple form reflecting solely the separate row effects (of age) and column effects (of period)" (P. 1). Quoting from an unpublished paper by Ryder (1979), Duncan concluded that Equation (24) corresponds "to Ryder's 'assignment,' which is 'the characterization of a cohort's performance over the life cycle,' so that the 'dependent variable' becomes 'the entire curve of the variable across the age span, for a particular cohort"' (Pp. 1–2). In fact, Equation (24) is the difference in predicted outcomes between age i and period j versus age $i-1$ and period $j-1$ in cohort k, thereby providing information about a section of a local diachronic age curve. To my knowledge, Duncan was the first researcher to recognize, albeit partially, that some of the parameters of the conventional APC model could be used to achieve Ryder's stated goals for cohort analysis.

Unfortunately, however, Duncan's approach is of limited utility for a Ryderian analysis. It appears that Duncan never realized how the conventional APC model can be used to represent diachronic cohort trends or synchronic differences, nor did he recognize that the parameters of the model can be separated into linear and nonlinear terms, thus yielding an expanded set of population-level summaries. Perhaps most problematic, Duncan (1981) erroneously believed, because the cohort parameters have dropped out, that Equation (24) proved that cohort is a "vacuous" variable that one should "abolish" (P. 2).[31] As he concluded: "I suggest that the APC model be recognized for what it is, an additive model of age and period effects on cohort trajectories, and not a model of 'cohort effects' at all" (Duncan 1981:2). Ryder's (1981) response to Duncan's conjecture was to the point: "Forgive me, but I doubt that this is helpful." Instead, Ryder (1981) interpreted Duncan as simply stratifying on (and thus adjusting for) cohort: "My reading of what you have done is to single out a lifetime record of a particular cohort, and express the change in the value of a dependent variable as a consequence of the passage of the cohort from period to period and from age to age."

Firebaugh's Linear Decomposition

Firebaugh (1989, 1990, 1997, 2008) has proposed what he calls "linear decomposition" to examine social change in population-level data.[32] Let Y denote an outcome, P a continuous period variable measured in years, and C a continuous cohort variable, also measured in years. Firebaugh recommended estimating the following model:

$$Y_{ijk} = \mu + \beta_1 P + \beta_2 C + \epsilon_{ijk}, \tag{25}$$

where μ is the intercept, β_1 is the slope for the continuous period variable P, β_2 is the slope for the continuous cohort variable C, and ϵ_{ijk} is an error term. According to Firebaugh (1997; see also 2008:197), β_1 indicates "intracohort change," whereas β_2 indicates "cross-cohort or intercohort change—the average difference between adjacent cohorts" (Pp. 24–25). In fact, Equation (25) is equivalent to the cohort synchronic slopes-only model displayed in Equation (4), with $\beta_1 = \theta_1 = \alpha + \pi$ and $\beta_2 = \theta_2 - \theta_1 = \gamma - \alpha$.

The main limitation of Equation (25) for a Ryderian analysis is that the cohort slope is synchronic, not diachronic, describing static comparisons within periods rather dynamic change across periods. This is in part why Norval Glenn (2005), in a volume on methods for analyzing APC data, avoided discussing Firebaugh's approach. As Glenn (2005) wrote, "I do not describe Firebaugh's techniques here because I now believe that neither they nor similar means of decomposition is [*sic*] very helpful for understanding change" (P. 36).[33] An additional limitation is that Equation (25) ignores nonlinearities for age, period, and cohort, discarding information that would otherwise provide a richer, and potentially more revealing, depiction of population-level temporal variability.

In short, modeling and graphing APC nonlinearities, Duncan's reformulation of the APC model, and Firebaugh's "linear decomposition" are related approaches that may provide insights on various patterns in APC data. In particular, as shown by Acosta and van Raalte (2019), graphical displays of APC nonlinearities can reveal substantively meaningful "bumps" in the data. By contrast, the techniques developed by Duncan and Firebaugh are of more limited utility. Duncan's reformulation is simply equivalent to a comparison within a section of a local diachronic age curve, whereas Firebaugh's linear model is just the synchronic cohort model with the restriction that the nonlinearities for age, period, and cohort are all zero. Accordingly, the proposed approach is both more general and flexible than either of these techniques.

Discussion and Conclusion

Drawing on the insights from Ryder's (1965) classic essay, in this article I have shown formally how the conventional APC model can be used in accordance with Ryder's goals for cohort analysis, enabling researchers to parsimoniously summarize population-level temporal variability on the Lexis table. The diachronic age slope, surface, and curves involve comparisons across age groups through time within cohorts (i.e., intra-cohort trends), thereby representing life-cycle change. By contrast, the diachronic cohort slope, surface, and curves entail comparisons of cohorts through time (i.e., inter-cohort trends), thus representing social change. Ryderian comparative cohort careers can be viewed as combining information on both intra- and inter-cohort trends and, accordingly, life-cycle and social change. Finally, synchronic age and cohort measures, representing intra-period differences, provide information on the relative magnitude of life-cycle and social change within periods. Together these summaries constitute a formal, unified framework for describing the main trends and patterns on a Lexis table in a way that honors Ryder's vision for analyzing APC data.

Although researchers would benefit from calculating and visualizing all of the expressions outlined in this article, the most illuminating are what I term the "big three" summaries of population-level temporal variability. These are the diachronic age curve (Eq. [10]), diachronic cohort curve (Eq. [13]), and both of the Ryderian comparative cohort careers (Eqs. [16] and [17]). The diachronic age curve summarizes how age groups differ through time, or life-cycle change. It can also be interpreted as an overall measure of intra-cohort change. By contrast, the diachronic cohort curve summarizes how cohorts differ through time, or social change. Finally, comparative cohort careers, whether the curves-only variant or not, depict, borrowing Ryder's terminology, diachronic intra-cohort development (i.e., life-cycle change) as well as diachronic inter-cohort differentiation (i.e., social change). These "big three" summaries, which can and should be visualized as in Figures 4, 8, and 12, provide, in a parsimonious way, the most essential information on population-level temporal variability in any given time-series cross-sectional data set. It is absolutely crucial to understand that, although expressed using the classical APC model, these model-based summaries are purely descriptive in that they do not rely on the introduction of information external to the data for estimation.[34] This has the great virtue that, consistent with Ryder's goal (1965, 1979), such summaries (and accompanying visualizations) can provide a basis for consensus across multiple sociological subfields.

If one must choose among these summaries, then arguably the most important is the calculation and display of comparative cohort careers (see Figure 12), which compactly depict the interplay of life-cycle and social change. This reflects Ryder's (1968) insight that the study of comparative cohort careers is "the most important application of cohort analysis" (P. 546). Alternatively, if one were to include one additional summary beyond the "big three," then I recommend presenting the adjusted marginal period curve (see Eq. [6] in the online supplement). Briefly, the adjusted marginal period curve is a composite diachronic summary that captures change in toto as a function of underlying life-cycle and social change.[35] I also strongly advise against estimating and displaying synchronic age and cohort curves (Eqs. [18] and [19]), which are easily misinterpreted as representing intra- and inter-cohort trends, respectively. As demonstrated previously, synchronic curves will give biased estimates of diachronic curves except under very strong (but typically testable) assumptions on the absence of overall (linear) life-cycle or social change (see Table 2). As a more principled way to represent intra-period differences, I recommend presenting and visualizing local synchronic differences (Eqs. [20] and [21]), which reflect the relative magnitude of life-cycle versus social change in each period.

There are several limitations of the current approach that point to potentially fruitful areas of further research. First, the Ryderian approach outlined in this article is based on distinguishing intra- from inter-cohort trends, not on deriving unique APC effects, which may nonetheless be of primary interest in some applications. Accordingly, future work should examine the extent to which a Ryderian approach can complement analyses that attempt to identify separate APC effects, and vice versa. Second, this article focused on analyzing APC data, defined as time-series cross-sectional data organized by age, period, and cohort. Further research should consider applying and comparing the approach outlined in this article using other kinds of data structures (e.g., panel data with multiple cohorts). As well, additional work should contrast the life-cycle changes observed in time-series cross-sectional data with those observed in panel data. In general, they should be similar, but they are likely to differ due to changes in the compositions of cohorts. Although the composition of cohorts will change in both panel data and time-series cross-sectional data due to mortality, one would expect the latter to also be subject to patterns of migration. Finally, this article concentrated on developing a Ryderian cohort approach using the conventional APC model, which is by far the most commonly used in the literature, and outlined how three related techniques can be viewed as generalizations of this approach. However, as noted previously, there are undoubtedly other models that can be used that are consistent, at least in part, with Ryder's theoretical framework (e.g., see Harding and Jencks 2003; Luo and Hodges 2022; Schulhofer-Wohl and Yang 2016). Further studies should compare the models used in this article with these alternative models.

There are also a number of possible methodological extensions to the approach outlined in this article. First, one can easily incorporate additional variables and examine heterogeneity across these variables. For example, ethnicity (or race), social class, and gender could be included, and variability in life-cycle and social change could be examined and graphed. Second, the approach outlined in this article relies on categorical age, period, and cohort variables with intervals of equal width. However, despite widespread practice, there are noteworthy limitations to using categorical APC data.[36] Future work should treat APC data as continuous, expressing the diachronic and synchronic L-APC models in terms of parametric regression splines (see Heuer 1997). Finally, the various model-based summaries discussed in this article can be incorporated into a framework using stochastic counterfactuals, as suggested by Fosse and Winship (2019a). For example, conditional on age, cohort and period jointly define an observed difference (i.e., disparity or gap). Various mechanisms could be proposed and examined to reduce this difference, thereby explaining social change (cf. Jackson and VanderWeele 2018).

Many of the findings in the APC literature remain quite controversial, and there are ongoing debates about the utility of various techniques. There is little doubt that disputes on the appropriate way to analyze APC data are unlikely to abate anytime soon. However, it is the viewpoint of this author that much would be gained by analyzing APC data using the models, expressions, and visualizations outlined here. Building on the seminal insights of Ryder (1965), as well as other writings in his corpus, the Ryderian cohort-based approach developed in this article offers a rigorous, transparent, and parsimonious way of formally summarizing times-series cross-sectional data in the APC framework. With judicious application, the proposed method has the potential to guide the development of theory, aid in the accumulation of evidence, and build a common base of knowledge in a literature often fraught with controversy.

Notes

1 Time-series cross-sectional data entail repeated observations on the same population. I will refer to time-series cross-sectional data organized by age, period, and cohort generically as "APC data." The interpretation of repeated cohorts of individuals tracked over time, which has received less attention in the APC literature, is beyond the scope of this article, but the general principles will be applicable.

2 Nash (1978:1–2) traces the concept of generation to ancient Greece, as revealed in Homer's *Iliad*, whereas Burnett (2010:11–14) locates the concept further back in time to ancient Egypt.

3 Throughout I will refer to "calendar time" or "time" as synonymous with "periods."

4 For a similar point, see Glenn (1977:22–23).

5 Fosse and Winship (2019a) provide a review and generalization of these methods, showing that a large class of estimators can be viewed as special cases of a bounding analysis under various theoretical assumptions.

6 There are frequent references to the "effects" of temporal variables in the APC literature (e.g., Mason and Fienberg 1985). Note, however, that age, period, and cohort resemble variables such as race or gender, which may be considered "non-manipulable" and thus of dubious causal status by at least some who adopt, broadly speaking, the counterfactual framework of causality (Rubin 1986; cf. Pearl 2009). Clarifying the interpretation of APC parameters using the language and notation of contemporary causal inference is an important topic for further research.

7 However, Glenn (2003) admitted that he never developed a formal means to analyze APC data in the way that Ryder had envisioned.

8 For the purposes of this article, I will focus on a Lexis table with age as the rows and period as the columns, with cohorts on the diagonals. This is natural because data are typically collected on age and period, with cohort as the derived dimension.

9 Throughout I assume that the data consist of individuals indexed from $r = 1, \ldots, R$, where R is the total sample size. However, for simplicity of exposition in all equations I omit the subscripts indexing individuals in the data.

10 By contrast, in analyses that attempt to derive unique temporal effects, the relationship of age and cohort to calendar time (i.e., period) is *apochronic* (from Greek "apo-" meaning "separate" and "khronos" meaning "time"). The reason is that, after fixing calendar time to some particular value, we are comparing values of age or cohort independent of calendar time itself (i.e., apochronically).

11 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.

12 This is a simplification. If birth cohorts are calculated using yearly data for age and period, then individuals may have been born within a range of up to two years. For example, the 1920 birth cohort is a midpoint value referring to individuals born between 1919 and 1921.

13 For simplicity we also assume that the age and period categories are of equal width. This approach could likewise be easily extended to continuous data using regression splines (see Heuer 1997).

14 Note that I is added to $j - i$ so that the cohort index begins at $k = 1$. This ensures that, for example, $i = j = k = 1$ refers to the 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.

15 This model could also be referred to as the *life-cycle and social change (LC-SC) model* inasmuch as it can be used to represent intra- and inter-cohort trends or, equivalently, life-cycle and social change.

16 With the L-APC model applied to data collected based on age and period, this is accomplished by replacing j with $i + k - I$ and J with $K - I + 1$.

17 For a similar reason, the diachronic slopes-only model (Eq. [2]) is preferred over the age and cohort synchronic slopes-only models (Eqs. [3] and [4]).

18 In APC data, these assumptions are typically testable simply because one can fit the diachronic L-APC model instead of the age or cohort synchronic L-APC models.

19 Note that all three models (Eqs. [7], 8, and [9]) give the same predicted probabilities. This is because each model fits the data equally well.

20 For example, Firebaugh (1989, 1990) refers to an "inter-cohort" slope estimand but does not distinguish clearly between diachronic and synchronic variants of the slope. His proposed model, as demonstrated later, in fact generates a synchronic slope estimate, which is precisely what Ryder cautioned against in his writings.

21 Distinguishing between diachronic and synchronic "intra-" group comparisons is unnecessary, because intra-age and intra-cohort comparisons are *always* diachronic, whereas an intra-period comparison is *always* synchronic. Thus, to label, for instance, an expression a "diachronic intra-cohort slope" burdens it with unneeded verbiage.

22 The reason for this, as discussed later, is that there is additional cell-specific heterogeneity on an age–period array that is not typically captured by the diachronic L-APC model. Incorporating such extra variability into the diachronic L-APC model and thus the parametric expressions outlined in this article is a subject for further research.

23 After substituting $j = i + k - I$ and $j^* = i^* + k - I$, where k is some reference cohort, note that $\alpha(i - i^*) + \pi(j - j^*) = \theta_1(i - i^*)$.

24 Specifically, $\theta_1(i - i^*) + \widetilde{\alpha}_i + \widetilde{\pi}_j$ for ages $i = 1, \ldots, I$ in a given period j. Note that this curve can be viewed as "synthetic" in that no actual cohort experiences just a single period nonlinearity over its entire life course.

25 After substituting $j = i + k - I$ and $j^* = i + k^* - I$, note that $\gamma(k - k^*) + \pi(j - j^*) = \theta_2(k - k^*)$, where i is some reference age group.

26 Specifically, the probabilities across cohorts within a particular period are equivalent to those from $\theta_2(k - k^*) + \widetilde{\gamma}_k + \widetilde{\pi}_j$ in a specified period j.

27 Although he wrote relatively little on how to conduct cohort analysis, in his unpublished writings Ryder appeared to advocate the separation of cohort careers from period fluctuations (see Ryder 1979).

28 Although the bias for the synchronic cohort curve is not particularly large in this example, this is not necessarily always the case.

29 Note that the local synchronic differences are relative and are in general not informative about the magnitude of the parameters.

30 In an age–period table there are $(I - 2) \times (J - 2)$ additional parameters that can be included. To understand how this is calculated, note that the diachronic L-APC model, which is based on the conventional APC model, takes up $1 + (I - 1) + (J - 2) + (K - 1) = 1 + (I - 1) + (J - 2) + ((I + J - 1) - 1)$ or $2(I + J) - 4$ parameters. However, there are $I \times J$ possible parameters in total. Thus, there are $IJ - (2I + 2J - 4) = IJ - 2I - 2J + 4 =$

$(I-2)(J-2)$ additional parameters that can be included beyond the diachronic L-APC model. Using similar calculations, we can conclude that if we had a period–cohort table or an age–cohort table, then we could include $(J-2)\times(K-2)$ or $(I-2)\times(K-2)$ additional parameters, respectively.

31 For this reason, Duncan also believed that Equation (24) somehow allows one to identify the unique effects of age and period. This is false, as Ryder (1981) correctly observed in a letter to Duncan: "If the APC problem can be characterized as that of two equations in three unknowns, the consequence of looking at the history of a particular cohort is to collapse the problem into that of one equation in two unknowns. Your [Equation (24)] is intractable in the sense that there is no way of making an observation about the effect of a change in period which is not likewise an observation about the effect of a change in age."

32 His approach is "meant to be descriptive, not causal" (Firebaugh 2008:199).

33 Glenn (2005) argued that Firebaugh's "decomposition is meaningful only in the absence of age effects on the dependent variable" (P. 36). This is putatively because the cohort slope in Equation (25) is synchronic and is thus equal to $\theta_2 - \theta_1$ or $\gamma - \alpha$.

34 More formally, the summaries are based on a fully identified model with a design matrix that is of full rank.

35 Specifically, the curve is a weighted sum of the diachronic age and cohort slopes, along with period-by-period fluctuations.

36 For example, if the temporal variables are coded using different widths (e.g., age is coded as five-year groups and period is coded as two-year groups), then this can generate an artifactual zigzag pattern in the nonlinearities (see Holford 2006). Moreover, the nonlinearities can fluctuate wildly due to sparseness. This is especially the case for cohort, because data are typically collected based on age and period and thus there tend to be very few individuals in the extreme cohorts.

References

Acosta, Enrique, and Alyson van Raalte. 2019. "APC Curvature Plots: Displaying Nonlinear Age-Period-Cohort Patterns on Lexis Plots." *Demographic Research* 41:1205–34. https://doi.org/10.4054/DemRes.2019.41.42.

Alwin, Duane F. 1991. "Family of Origin and Cohort Differences in Verbal Ability." *American Sociological Review* 56(5):625–38. https://doi.org/10.2307/2096084.

Alwin, Duane F., and Ryan J. McCammon. 2003. "Generations, Cohorts, and Social Change." In *Handbook of the Life Course*, edited by Jeylan T. Mortimer and Michael J. Shanahan, Handbooks of Sociology and Social Research, pp. 23–49. New York: Kluwer Academic Publishers. https://doi.org/10.1007/978-0-306-48247-2_2.

Burnett, Judith. 2010. *Generations: The Time Machine in Theory and Practice*. Burlington, VT: Ashgate.

Clogg, Clifford C. 1982. "Cohort Analysis of Recent Trends in Labor Force Participation." *Demography* 19:459–79. https://doi.org/10.2307/2061013.

Clogg, Clifford C., and James W. Shockey. 1985. "The Effect of Changing Demographic Composition on Recent Trends in Underemployment." *Demography* 22:395–414. https://doi.org/10.2307/2061068.

Duncan, Otis Dudley. 1981. "A Reformulation of the APC Model." Department of Sociology, University of Arizona.

Duncan, Otis Dudley. 1985. "Generations, Cohorts, and Conformity." In *Cohort Analysis in Social Research*, edited by William M. Mason and Stephen E. Fienberg, pp. 289–321. New York: Springer. https://doi.org/10.1007/978-1-4613-8536-3_9.

Duncan, Otis Dudley, and Magnus Stenbeck. 1988. "Panels and Cohorts: Design and Model in the Study of Voting Turnout." *Sociological Methodology* 18:1–35. https://doi.org/10.2307/271043.

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

Elder, Glen H., Jr., and Linda K. George. 2016. "Age, Cohorts, and the Life Course." In *Handbook of the Life Course, Volume II*, edited by Michael J. Shanahan, Jeylan T. Mortimer, and Monica Kirkpatrick Johnson, Handbooks of Sociology and Social Research, pp. 59–86. New York: Springer. https://doi.org/10.1007/978-3-319-20880-0_3.

Fienberg, Stephen E., and William M. Mason. 1979. "Identification and Estimation of Age-Period-Cohort Models in the Analysis of Discrete Archival Data." *Sociological Methodology* 10:1–67. https://doi.org/10.2307/270764.

Firebaugh, Glenn. 1989. "Methods for Estimating Cohort Replacement Effects." *Sociological Methodology* 19:243–62. https://doi.org/10.2307/270954.

Firebaugh, Glenn. 1990. "Replacement Effects, Cohort and Otherwise: Response to Rodgers." *Sociological Methodology* 20:439–46. https://doi.org/10.2307/271093.

Firebaugh, Glenn. 1997. *Analyzing Repeated Surveys*. Quantitative Applications in the Social Sciences. Thousand Oaks, CA: Sage Publications.

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

Fosse, Ethan, and Christopher Winship. 2018. "Moore–Penrose Estimators of Age–Period–Cohort Effects: Their Interrelationship and Properties." *Sociological Science* 5:304–34. https://doi.org/10.15195/v5.a14.

Fosse, Ethan, and Christopher Winship. 2019a. "Analyzing Age-Period-Cohort Data: A Review and Critique." *Annual Review of Sociology* 45:467–92. https://doi.org/10.1146/annurev-soc-073018-022616.

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

Glenn, Norval D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period and Cohort Effects." *American Sociological Review* 41(5):900–904. https://doi.org/10.2307/2094738.

Glenn, Norval D. 1977. *Cohort Analysis*. Thousand Oaks, CA: Sage Publications.

Glenn, Norval D. 2003. "Distinguishing Age, Period, and Cohort Effects." In *Handbook of the Life Course*, edited by Jeylan T. Mortimer and Michael J. Shanahan, Handbooks of Sociology and Social Research, pp. 465–76. New York: Springer. https://doi.org/10.1007/978-0-306-48247-2_21.

Glenn, Norval D. 2005. *Cohort Analysis*. Quantitative Applications in the Social Sciences. Thousand Oaks, CA: Sage Publications, 2nd edition.

Goldstone, Jack. 1991. *Revolution and Rebellion in the Early Modern World*. University of California Press.

Harding, David, and Christopher Jencks. 2003. "Changing Attitudes toward Premarital Sex: Cohort, Period, and Aging Effects." *Public Opinion Quarterly* 67(2):211–26. https://doi.org/10.1086/374399.

Hardy, Melissa, and Andrea E. Wilson. 2002. "Cohort Change." In *Encyclopedia of Aging*, edited by David J. Ekerdt, volume 1, pp. 241–45. New York: Macmillan Reference.

Heuer, Carsten. 1997. "Modeling of Time Trends and Interactions in Vital Rates Using Restricted Regression Splines." *Biometrics* 53(1):161. https://doi.org/10.2307/2533105.

Holford, Theodore R. 1991. "Understanding the Effects of Age, Period, and Cohort on Incidence and Mortality Rates." *Annual Review of Public Health* 12:425–57. https://doi.org/10.1146/annurev.pu.12.050191.002233.

Holford, Theodore R. 2006. "Approaches to Fitting Age-Period-Cohort Models with Unequal Intervals." *Statistics in Medicine* 25(6):977–93. https://doi.org/10.1002/sim.2253.

Inglehart, Ronald. 1990. *Culture Shift in Advanced Industrial Society*. Princeton University Press.

Jackson, John W., and Tyler J. VanderWeele. 2018. "Decomposition Analysis to Identify Intervention Targets for Reducing Disparities." *Epidemiology* 29(6):825–35. https://doi.org/10.1097/ede.0000000000000901.

Knoke, David, and Michael Hout. 1974. "Social and Demographic Factors in American Political Party Affiliations, 1952–72." *American Sociological Review* 39(5):700–713. https://doi.org/10.2307/2094315.

Luo, Liying, and James S. Hodges. 2022. "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. https://doi.org/10.1177/0049124119882451.

Mannheim, Karl. [1927/1928] 1952. "The Problem of Generations." In *Karl Mannheim: Essays*, edited by Paul Kecskemeti, pp. 276–322. 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. https://doi.org/10.2307/2094398.

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

Nash, Laura L. 1978. "Concepts of Existence: Greek Origins of Generational Thought." *Daedalus* 107(4):1–21.

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

Pearl, Judea. 2009. *Causality*. Cambridge University Press.

Reither, Eric N., Robert M. Hauser, and Yang Yang. 2009. "Do Birth Cohorts Matter? Age-Period-Cohort Analyses of the Obesity Epidemic in the United States." *Social Science & Medicine* 69(10):1439–48. https://doi.org/10.1016/j.socscimed.2009.08.040.

Riley, Matilda White. 1973. "Aging and Cohort Succession: Interpretations and Misinterpretations." *Public Opinion Quarterly* 37(1):35–49. https://doi.org/10.1086/268058.

Robertson, Chris, and Peter Boyle. 1998. "Age–Period–Cohort Models of Chronic Disease Rates. II: Graphical Approaches." *Statistics in Medicine* 17(12):1325–40. https://doi.org/10.1002/(SICI)1097-0258(19980630)17:12%3C1325::AID-SIM854%3E3.0.CO;2-R.

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. https://doi.org/10.1006/ssre.2000.0692.

Rodgers, Willard. 1982. "Trends in Reported Happiness within Demographically Defined Subgroups, 1957–78." *Social Forces* 60(3):826–42. https://doi.org/10.1093/sf/60.3.826.

Rodgers, Willard L. 1990. "Interpreting the Components of Time Trends." *Sociological Methodology* 20:421–38. https://doi.org/10.2307/271092.

Rubin, Donald B. 1986. "Comment: Which Ifs Have Causal Answers." *Journal of the American Statistical Association* 81(396):961–2. https://doi.org/10.1080/01621459.1986.10478355.

Ryder, Norman B. 1964. "The Process of Demographic Translation." *Demography* 1(1):74–82. https://doi.org/10.1007/BF03208446.

Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30(6):843–61. https://doi.org/10.2307/2090964.

Ryder, Norman B. 1968. "Cohort Analysis." In *International Encyclopedia of the Social Sciences*, edited by David L. Sills, volume 2, pp. 546–50. New York: Macmillan.

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

Ryder, Norman B. 1980. "Components of Temporal Variations in American Fertility." In *Demographic Patterns in Developed Societies*, edited by R. W. Hirons, volume 19, pp. 15–54. London: Taylor & Francis.

Ryder, Norman B. 1981. Letter from Norman B. Ryder to Otis Dudley Duncan.

Ryder, Norman B. 1983. "Cohort and Period Measures of Changing Fertility." In *Determinants of Fertility in Developing Countries*, edited by Rodolfo A. Bulatao and Ronald D. Lee, volume 2, pp. 737–56. New York: Academic Press.

Ryder, Norman B. 1992. "Cohort Analysis." In *Encyclopedia of Sociology*, edited by Edgar F. Borgatta and Marie L. Borgatta, pp. 227–31. New York: Macmillan.

Schulhofer-Wohl, Sam, and Yang Yang. 2016. "Modeling the Evolution of Age and Cohort Effects." In *Dynamic Demographic Analysis*, edited by Robert Schoen, pp. 313–35. New York: Springer. https://doi.org/10.1007/978-3-319-26603-9_16.

Tilly, Charles. 1984. *Big Structures, Large Processes, Huge Comparisons*. New York: Russell Sage Foundation.

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

Acknowledgments: The author thanks Jessica Harris, Sebastien Parker, and Christopher Winship.

Ethan Fosse: Department of Sociology, University of Toronto. E-mail: ethan.fosse@utoronto.ca.