

Cross-Group Differences in Age, Period, and Cohort Effects:  
A Bounding Approach to the Gender Wage Gap

Ohjae Gowen  
*Harvard University*

Ethan Fosse  
*University of Toronto*

Christopher Winship  
*Harvard University*

## Abstract

For decades, researchers have sought to understand the separate contributions of age, period, and cohort (APC) on a wide range of outcomes. However, a major challenge in these efforts is the linear dependence among the three time scales. Previous methods have been plagued by either arbitrary assumptions or extreme sensitivity to small variations in model specification. In this article, we present an alternative method that achieves partial identification by leveraging additional information about subpopulations (or strata) such as race, gender, and social class. Our first goal is to introduce the Cross-Strata Linearized APC (CSL-APC) model, a re-parameterization of the traditional APC model that focuses on cross-group variations in effects instead of main effects. Similar to the traditional model, the linear cross-strata APC effects are not identified. The second goal is to show how Fosse and Winship's (2019) bounding approach can be used to address the identification problem of the CSL-APC model, allowing one to partially identify cross-group differences in effects. This approach often involves weaker assumptions than previously used techniques, and in some cases can lead to highly informative bounds. To illustrate our method, we examine differences in temporal effects on wages between men and women in the United States.

## Introduction

The social sciences have long focused on studying changes in disparities among different groups—including but not limited to those based on race, gender, education, and social class—in various outcomes. At the core of much of this research is the objective of identifying the relative importance of aging, period-related, and cohort-based factors in creating trends in disparities observed among groups. For instance, in the context of wage disparities between men and women, the causes could range from changes in the age distribution, to period-related factors like changes in economic and employment opportunities, to cohort-based factors that reflect enduring generational differences in educational attainment or underlying values and attitudes.

However, as is well known, it is extremely difficult to identify the independent effects of age, period, and cohort (APC) variables in a given data set.<sup>1</sup> This is due to the APC identification problem, or the fact that each of the temporal scales is a linear function of the other, such that  $\text{Period} = \text{Age} + \text{Cohort}$  (for a discussion, see O’Brien 2015; Fosse and Winship 2019a). For example, if a researcher follows a single birth cohort of men and women, one cannot tell whether observed changes in the gender gap in earnings are entirely due to age or period effects because the age and period variables have advanced in parallel. Similarly, when comparing two cohorts of the same age, one cannot tell whether changes in a gender earnings gap are due to cohort or period effects because the two cohorts differ not only in their year of birth but also in the year in which their earnings are observed, reflecting a possible period effect.

Despite the difficulties posed by the identification problem, Fosse and Winship (2019b) have shown that much can be learned from the data using what they call a bounds approach. Specifically, they show how theoretically driven assumptions about the size, sign, and/or shape of one or more of the three underlying APC effects over a range of the data can be used to derive bounds on the parameters of interest. In some cases, depending on the nature of the data and the particular assumptions invoked, the bounds obtained can be very narrow (i.e., highly informative).

We have two primary objectives in this article. The first is to introduce what we call the Cross-Strata Linearized APC (CSL-APC) model, a reparameterization of the traditional APC model that is uniquely suited for the analysis of group disparities. Like the traditional APC model, the CSL-APC model is used to specify the possible separate effects of age, period, and cohort. However, instead of explaining an overall outcome in a population, the CSL-APC model is used to explain the difference in an outcome across strata (or subpopulations) such as gender, social class, and geographic region. We do this by defining the between-group differences in age, period, and cohort effects as the estimands of interest. However, the CSL-APC model, similar to the traditional APC model, is not fully identified. Accordingly, our second objective is to show how the bounding approach of Fosse and Winship (2019b) can be extended to examine between-group differences in APC effects. We build on their core idea that theoretically driven assumptions can help set bounds on temporal

---

<sup>1</sup>By adopting the language of “effects,” we refer to the putative bundles of underlying causal factors that are proxied by age, period, and cohort (Clogg 1982). This is distinct from “trends,” which are observed patterns in the data that vary over calendar time (or period). For a discussion, see Fosse and Winship (2023).

effects, thereby achieving partial identification. A potential advantage of focusing on partial identification of between-group differences in effects is that it changes the nature of the theoretical assumptions required. However, instead of assumptions about the effects of age, period, and cohort on the overall levels of an outcome, assumptions are made about the effects of the three APC variables on *differences between groups*. In some cases, these theoretical assumptions may be more plausible than separate assumptions for each group as in a typical APC analysis.

To illustrate our approach, we analyze the wage gap between U.S. men and women using annual supplemental data from the Current Population Survey from 1976 to 2019. Under a limited set of assumptions, our analysis shows that cohort replacement effects have driven continued progress in women’s relative pay, but that this progress has been partially offset by stagnating period effects since the 1990s. These results provide valuable insights into the dramatic change in gender wage inequality over the past four decades, as well as the slowing of the “gender revolution” since the 1990s (England et al. 2020).

In the following sections, we first present our empirical example and then briefly review the most commonly used APC techniques for identifying the sources of change in cross-group disparities. We then discuss the related literature that attempts to explain temporal shifts in the gender wage gap. Next, we introduce the CSL-APC model and show how it can be used to construct a 2D APC graph, which is a crucial component of our bounds approach. We then analyze the effects of age, period, and cohort on the gender wage gap. For conceptual clarity, we consider a case where point identification is achieved by assuming that one of the three linear APC effects, specifically the effect of age, is the same between women and men. However, this assumption is unrealistic in our particular example. Therefore, to overcome this limitation, we build on the CSL-APC model developing a cross-strata bounding approach that relies on much weaker assumptions. While explaining the analytical process step by step, we demonstrate how theoretical assumptions can be used to derive bounding constraints and ultimately partially identify the cross-group differences in age, period, and cohort effects. In addition, we present two novel types of sensitivity analyses that can be used to assess the robustness of our results. To our knowledge, these sensitivity analyses are a novel methodological contribution to the APC literature.


## **Empirical Example: The Gender Wage Gap**

As one of the most widely used measures of gender inequality, the wage gap between men and women has been of great interest to scholars across the social sciences (O’Neill and Polachek 1993; Cha and Weeden 2014; Blau and Kahn 2017; Horowitz and Igielnik 2020). Our analysis focuses on identifying age, period, and cohort effects on over-time changes in the gender difference in median annual earnings. Our data consists of pooled cross-sections of the Annual Social and Economic Supplement of the Current Population Survey (CPS ASEC) from 1976 to 2019 (Flood et al. 2021). Our analytic sample includes full-time, year-round wage and salary workers between the ages of 25 and 64 (830,856 women and 1,121,562 men). Full-time year-round workers are defined as those

who worked 50 or more weeks and at least 35 hours per week in the last calendar year.<sup>2</sup> We exclude the self-employed as their income is likely to be conflated with capital income. Since the outcome is measured by respondents' annual earnings in the previous calendar year, our earnings data cover the period from 1975 to 2018. Throughout the article, we refer to the years to which the earnings information applies instead of the survey years for convenience. Earnings are adjusted for inflation in 2018 dollars. Because the top-coding scheme for earnings in the CPS has changed over time, we use the median rather than the mean to measure the gender wage gap.<sup>3</sup> Following the convention in the APC literature, we group age and period into equal five-year intervals and compute cohort from these intervals. This results in eight age categories ranging from 25-29 to 60-64, nine period categories ranging from 1975-79 to 2015-2019, and sixteen cohort categories ranging from 1915-19 to 1990-1994. We use the CPS ASEC survey sampling weights throughout.

We first present descriptive plots to illustrate trends in the U.S. gender wage gap over the 1975-2018 period (see Table 1 in Appendix C for additional descriptive statistics). The left panel of Figure 1 plots the observed log median earnings across periods. It shows that the gender gap in log median earnings has narrowed significantly over the last forty years, particularly in the 1980s. The right panel also shows that, in the 1975-1979 period, women's median earnings were 59 percent of their male counterparts on the original (unlogged) scale. This share rose to 71 percent in the 1990-1994 period. Since then, women's relative gains have slowed, reaching 80 percent in 2015-2018.<sup>4</sup>

Figure 1: Marginal Period Trends of the Gender Difference in Earnings



Notes: Log median annual earnings (left) and the female-male percentage of median earnings (right) are estimated among full-time (35+ usual work hours a week), year-round (50+ weeks a year), wage/salary workers aged 25-64. Data from CPS ASEC 1976-2019.

<sup>2</sup>As an additional robustness check, we analyzed gender differences in median hourly earnings instead of annual earnings. We also repeated the analysis of hourly earnings with the extended sample of full-time employees who worked at least half a year (26 weeks) in the last calendar year. The results are robust to these different strategies for measuring the gender wage gap. Figures 1 and 2 in Appendix D present the results.


<sup>3</sup>The substantive conclusions remain the same when we analyze mean wage differences instead of those based on the median. See Figure 3 in Appendix D for the results.

<sup>4</sup>Previous research has shown stagnation since the 1990s in other outcomes related to women's labor market status, such as occupational desegregation (England et al. 2020).

The above patterns raise the important question of whether these observed trends are driven by underlying age, period, or cohort effects. Age effects may have driven the trends if age composition of full-time workers has changed in favor of women's relative pay.<sup>5</sup> If period effects are dominant, this suggests that contemporaneous society-wide shifts, such as changing norms and workplace policies regarding women's employment, were the key factor in the wage convergence. Alternatively, if cohort effects are the main explanation, it implies that the entrance of new cohorts of men and women into employment where the gender wage gap is smaller, and the exit of older individuals where it is larger, explains the observed decrease in the overall wage gap. Additionally, it is of interest to determine whether the slowdown of gender wage convergence is due to age, period, or cohort effects.

Figure 2 documents the age-graded patterns of the gender wage gap for selected birth cohorts. As the figure shows, the gender gap has been decreasing steadily across cohorts. At the same time, when comparing cohorts within the same age range (e.g. at age 40 along the vertical dashed line), we can see that the distance between cohorts is generally narrowing. These changes cannot necessarily be attributed solely to underlying cohort effects, as cohorts compared at a given age differ both in the year they were born and in the period during which their outcomes are observed. Thus, observed differences may be the result of cohort effects or period effects.

Figure 2: Age-Graded Patterns of the Observed Gender Difference in Earnings Among Selected Cohorts



Notes: Age-graded patterns of the female-male percentage of median earnings are estimated among full-time (35+ usual work hours/week) year-round (50+ weeks/year) wage/salary workers aged 25-64. Selective birth cohorts are presented for illustration. The vertical dashed line is drawn to illustrate the degree of the gender wage gap for each cohort when the cohort members were all at the same age of 40.

<sup>5</sup>However, the age composition of our sample has not changed substantially over time, such that the correlation between age and period is less than 0.1. Therefore, changing age composition is not likely the main source of the trends.

## Previous Literature

So far, our analysis has focused only on the observed trends in the gender wage gap. However, as we have emphasized, a key goal of APC analysis is to understand the underlying distinct effects of age, period, and cohort on the outcome of interest. Before turning to how we will use the CSL-APC model to analyze the distinct contributions of age, period, and cohort on the gender wage gap, we highlight in this section two distinct bodies of literature relevant to our proposed method and empirical analyses. We first discuss previous scholarship that has used various APC methods to uncover differences in distinct temporal effects across subpopulations (or strata) such as race, gender, and social class. We then outline research that has attempted to identify potential factors causing over-time variability in the gender wage gap, focusing on research in the United States.

### Incorporating Strata into APC Analysis

In general, researchers have relied on three approaches to extract unique effects for age, period, and cohort while accounting for variability across strata. One approach is to fit separate APC models for each stratum and then compare estimates across strata (e.g., Yang and Land 2013:125-169; Masters et al. 2014). A second, closely related approach is to fit a single APC model with a set of interaction terms that allow for variability in effects across strata (e.g., Yang 2008; Pampel and Hunter 2012). An advantage of this approach when using sample data is that one can easily conduct statistical tests of cross-strata differences in the effects by examining the statistical significance of the interaction terms. Finally, a third method is to construct an outcome that is a difference between two contrasting strata and then fit a conventional APC model with age, period, and cohort as inputs (e.g., O'Brien 2015:106-112). These parameters will capture cross-strata differences in the effects rather than the main effects.<sup>6</sup> Although limited in that only two groups can be compared at a time, this approach has the advantage of simplicity and, because the focus is only on identifying cross-strata variability in the effects, it may entail weaker assumptions than either of the first two approaches. We adopt this last approach in presenting our CSL-APC model.

Regardless of which approach is adopted, a key decision entails how to obtain identification in light of the linear dependence among the three temporal scales. We discuss the four most common approaches to obtain identification among sociological and demographic studies that explicitly examine variability in effects across strata. First, as suggested by Mason and colleagues (1973), one can constrain two or more effects to be equivalent (e.g., Riebler and Held 2010, 2012). For example, Mason and Smith (1985) analyze the effects of age, period, and cohort on tuberculosis. To identify their model, they assume that the coefficients for ages 5 – 9 and ages 10 – 14 are equal. The main problem with this approach is that any particular equality constraint, while seemingly trivial, is actually a very strong assumption, since it is tantamount to assuming a particular value for the unknown linear effect. An additional limitation is that in practice it is usually difficult to

---

<sup>6</sup>This approach is equivalent to modeling the interaction terms in the second approach above for two contrasting strata (see Appendix A).

theoretically justify any particular equality constraint over any other equality constraint (Kupper et al. 1985).

A second approach is to categorize age, period, and cohort variables in a dataset so that they do not retain exact linear dependence. This is often done by using categories of different lengths for age, period, and cohort (Underwood et al. 2022). For example, Campbell and Pearlman (2013) examine the same substantive question as in this article, namely the gender wage gap. Using CPS data, they find that cohort replacement has driven the gender convergence of earnings. Although their article provides useful insights for our study, Holford (2006) and Luo and Hodges (2016) show that this approach amounts to imposing a set of implicit equality constraints that typically lack theoretical justification. Moreover, as with the equality constraints approach, seemingly trivial reparameterizations can produce dramatically different results depending on how category intervals are defined (Luo and Hodges 2016).

A third approach is to use a hierarchical APC (HAPC) model (Yang and Land 2006, 2013). For example, Pampel and Hunter (2012) use the HAPC model to examine how the education gap in support for environmental spending varies across successive cohorts, controlling for age and period. The main limitation of the HAPC model, as demonstrated by Luo and Hodges (2020), is that it generally leads to a zero linear cohort effect (see also O’Brien 2017; Bell and Jones 2018).<sup>7</sup> In most applications, the assumption of a zero linear effect for cohort is inconsistent with prevailing theories on the importance of cohort replacement (e.g., Ryder 1965).

Finally, a fourth approach is to use the Intrinsic Estimator or IE (Fu 2000, 2016; Yang et al. 2008), a type of Moore-Penrose estimator (Fosse and Winship 2018). For example, Masters et al. (2014) use the IE to examine period and cohort effects on adult mortality, focusing on the Black-white mortality gap. The main limitation of the IE is that the estimates depend on how the data are coded (O’Brien 2015; Luo 2013; Fosse and Winship 2018). As shown in Luo et al. (2016), there are always multiple coding schemes where the IE produces different sets of parameter estimates equally consistent with the data.

In summary, current APC methods for incorporating cross-strata variability are limited, typically relying on strong implicit assumptions that are, in many cases, divorced from theoretical considerations. To address this, we propose an alternative approach that extends the bounding approach of Fosse and Winship (2019b) to explore cross-strata differences in APC effects (hereafter, “cross-strata effects”). This has several advantages. First, because it is based on partial rather than point identification, the bounding approach generally involves much weaker assumptions than methods based on point identification (Manski 2007). In fact, all of the above methods can be seen as special cases of the bounding approach where very strict assumptions are used to obtain narrow bounds (which are equivalent to point identification in the limit). Second, the bounding approach is highly flexible, allowing for a variety of assumptions about the size, sign, and/or shape of one or more of the temporal effects. Third, unlike many other techniques for APC analysis, the

---

<sup>7</sup>More precisely, the HAPC model will typically impose a zero linear effect on whichever of the temporal scales has the most categories in the data. When a conventional age-by-period Lexis table is the data input, this will be cohort.



bounding approach is based on a parameterization that clearly separates the identified from the unidentified parts of the model (see Fosse and Winship 2019a). It can thus take full advantage of estimates of nonlinear effects, which are identifiable because they are not linearly dependent. Assumptions based on monotonic temporal effects over some range of the data are particularly useful in this sense because in many cases they are easier to justify than constraints that assume exactly equal effects. For example, one might assume that criminality increases in early adolescence and then declines from the mid-twenties onward, or that the prevalence of prostate cancer increases monotonically with age (Fosse and Winship 2019b). Finally and relatedly, the separation of the unidentified from the identified parameters provides for a transparent link between theory and the estimates, allowing researchers to trace the consequences of particular theoretical assumptions on the conclusions obtained in any given application.

### **Potential Causes of the Gender Wage Gap**

As shown above, the wage gap between men and women in the United States has narrowed substantially since the 1980s, although the pace of narrowing slowed in the 1990s (England et al. 2020). A dominant explanation for this convergence emphasizes the increasingly similar levels of human capital that women and men bring to the labor force. Research shows that women's human capital characteristics, such as educational attainment, job tenure, and work experience, have grown at a faster rate relative to men's (O'Neill and Polachek 1993). This uneven pace of human capital accumulation explains a significant portion of the decline in the gender wage gap over the past four decades (Blau and Kahn 2017).

However, gender convergence in wage-raising human capital characteristics is not the only source of the narrowing wage gap. Declining fertility rates have also contributed to the rise in women's relative pay, as fewer children mean that women are less likely to experience a potential wage loss from having children (Budig and England 2001; Killewald and Cricco 2020). In addition, declining unionization and the rise of service economy have disproportionately suppressed wages in traditionally male-dominated occupations and manufacturing sectors with high unionization rates (Blau and Kahn 1997, 2017; Borghans et al. 2014). There is also evidence that gender discrimination in earnings has declined to some extent, in part due to government anti-discrimination policies in the 1960s and 1970s (Kurtulus 2012).

There is less consensus on why the progress in women's relative earnings has slowed since the 1990s. Given that the degree of occupational sex segregation has remained more or less the same since the 1990s, various factors leading to gender differences in occupational choice may have slowed the rate of growth in women's relative earnings (Blau et al. 2006). Relatedly, the emerging ideologies of intensive mothering since the mid-1990s may have aggravated a career penalty for working mothers who face more pressure from home (Hays 1996). Some studies suggest that the Family and Medical Leave Act of 1993 may have suppressed women's earnings growth by impeding their accumulation of work experience, but the empirical evidence is unclear (Blau and Kahn 2017).

Seemingly gender-neutral changes in the wage structure—for example, increasing wage returns to long work hours or high-skilled tasks—have been shown to benefit the relative earnings of men, who disproportionately occupy jobs with such conditions in the labor market (Blau and Kahn 2006; Cha and Weeden 2014).

While we do not aim to adjudicate between these competing explanations for observed changes in the gender wage gap, they provide useful insights into why age, period, and cohort may have distinct effects on widening or narrowing the gap. On the one hand, changes in the earnings gap driven by human capital and family demographics are likely to manifest as cohort effects. This would be particularly the case for the effects driven by education and occupational choice, for which intra-cohort variation among workers is likely to be limited. On the other hand, effects driven by changes in the wage structure or policy enforcement would presumably manifest as period effects on the wage gap, since their impact on the labor force is not likely to be limited to workers of a particular birth cohort. Age effects will reflect underlying family demographics and career patterns, but we would generally not expect age effects to drive the observed changes in the gender wage gap across periods (as in Figure 1) unless there has been a significant shift in the age structure of the workforce (see note 5).

## Modeling Cross-Strata Temporal Effects

In this section, we present an extension of the conventional APC model by including additional interaction terms between the strata variable and the variables representing age, period, and cohort. This allows us to examine variations in temporal effects across different levels of the strata variable, such as gender in our example. We then present a simplified version of this model, which we refer to as the Cross-Strata Linearized APC (CSL-APC) model. As we discuss below, this model allows for the focus on cross-strata differences in unidentified parameters (i.e., linear effects) as well as the visualization of these parameters in a 2D APC plot, which is the basis for our bounding approach.

### Classical and Linearized APC Models

Suppose we have an age-period Lexis table with cohorts on the diagonals. Each cell of the table represents a value of a continuous outcome  $Y$ . Following the convention in the literature, we will treat age, period, and cohort as categorical variables.<sup>8</sup> Let  $i = 1, \dots, I$  denote the age groups,  $j = 1, \dots, J$  the period groups, and  $k = 1, \dots, K$  the cohort groups, with  $k = j - i + I$  and  $K = I + J - 1$ . The *Classical APC model* is represented by the following equation (Fosse and Winship 2019a):

$$Y_{ijk} = \mu + \alpha_i + \pi_j + \gamma_k + \eta_{ijk}, \quad (1)$$

where  $Y_{ijk}$  is the cell value;  $\mu$  is the intercept;  $\alpha_i$  denotes the  $i$ th age effect ( $1, \dots, I$ );  $\pi_j$  denotes the  $j$ th period effect ( $1, \dots, J$ );  $\gamma_k$  denotes the  $k$ th cohort effect ( $1, \dots, K$ ); and  $\eta_{ijk}$  denotes a cell-specific

---

<sup>8</sup>For simplicity, we will also assume that the age and period categories are of equal width and that we have only aggregated data (i.e., there is no individual-level variability within the cells).

error term on the Lexis table. To identify the intercept, the age, period, and cohort parameters are constrained to sum to zero such that  $\sum_i^I \alpha_i = \sum_j^J \pi_j = \sum_k^K \gamma_k = 0$ .

An alternative specification is the *Linearized APC (L-APC) model*, which divides the overall temporal effects into linear and nonlinear effects (Holford 1983; Fosse and Winship 2019a). The L-APC model is given by:

$$Y_{ijk} = \mu + \alpha(i - i^*) + \pi(j - j^*) + \gamma(k - k^*) + \tilde{\alpha}_i + \tilde{\pi}_j + \tilde{\gamma}_k + \eta_{ijk}, \quad (2)$$

where  $\alpha$ ,  $\pi$ , and  $\gamma$  denote the age, period, and cohort linear effects, respectively; and  $\tilde{\alpha}_i$ ,  $\tilde{\pi}_j$ , and  $\tilde{\gamma}_k$  refer to nonlinear effects for the  $i$ th age,  $j$ th period, and  $k$ th cohort categories, respectively. The age, period, and cohort categories are centered around the midpoint indices marked by the asterisks,  $i^* = (I + 1)/2$ ,  $j^* = (J + 1)/2$  and  $k^* = (K + 1)/2$ , such that the age, period, and cohort parameters, as in Equation 1, satisfy the zero-sum constraint. The parameterization in Equation 2 has the main advantage of clarifying the nature of the identification problem: While the intercept and the nonlinear effects are identified, the linear effects  $\alpha$ ,  $\pi$ , and  $\gamma$  are not (Fienberg and Mason 1979; Holford 1983).

### Cross-Strata Linearized APC Model

Suppose  $s = 1, \dots, S$  indices some set of  $S$  strata, or subpopulations of substantive interest (e.g., race, gender, class, geographic region, and so on). The L-APC model can be generalized to incorporate cross-strata differences, leading to what we call the *Stratified Linearized APC (SL-APC) model*:

$$Y_{ijk s} = \mu_s + \alpha_s(i - i^*) + \pi_s(j - j^*) + \gamma_s(k - k^*) + \tilde{\alpha}_{is} + \tilde{\pi}_{js} + \tilde{\gamma}_{ks} + \eta_{ijk s}, \quad (3)$$

which is identical to Equation 2 except now the intercept, linear effects, nonlinear effects, and error terms are allowed to vary across levels of the strata variable. As noted previously, our interest lies in identifying the *differences* in the APC effects across strata, rather than the stratum-specific effects (see also Appendix A for another interpretation).

Focusing on the differences requires a reformulation of Equation 3. The simplest approach is to reformulate Equation 3 in terms of differences between any two selected strata.<sup>9</sup> For our example, in which gender is the strata variable, we will let  $s = 2$  refer to women and  $s = 1$  to men. This gives us the following differenced model equation, which we refer to as the *Cross-Strata Linearized APC (CSL-APC) model*:

$$\begin{aligned} Y_{ijk[s=2]} - Y_{ijk[s=1]} \\ = \left( \mu_2 + \alpha_2(i - i^*) + \pi_2(j - j^*) + \gamma_2(k - k^*) + \tilde{\alpha}_{i2} + \tilde{\pi}_{j2} + \tilde{\gamma}_{k2} + \eta_{ijk2} \right) - \left( \mu_1 + \alpha_1(i - i^*) + \pi_1(j - j^*) + \gamma_1(k - k^*) + \tilde{\alpha}_{i1} + \tilde{\pi}_{j1} + \tilde{\gamma}_{k1} + \eta_{ijk1} \right) \end{aligned}$$

or, equivalently,

$$\Delta Y_{ijk} = \Delta \mu + \Delta \alpha(i - i^*) + \Delta \pi(j - j^*) + \Delta \gamma(k - k^*) + \Delta \tilde{\alpha}_i + \Delta \tilde{\pi}_j + \Delta \tilde{\gamma}_k + \Delta \eta_{ijk}, \quad (4)$$

---

<sup>9</sup>Alternatively, as we show in Appendix A, one can specify a model with a set of interactions between the linear and nonlinear components and the levels of the strata variable. With this formulation, the interactions will capture the cross-strata differences in the temporal effects.

where  $\Delta Y_{ijk}$  denotes the cross-strata difference in outcomes;  $\Delta\mu$  is the difference in intercepts;  $\Delta\alpha$ ,  $\Delta\pi$ , and  $\Delta\gamma$  are the cross-strata differences in the linear effects;  $\Delta\tilde{\alpha}_i$ ,  $\Delta\tilde{\pi}_j$ , and  $\tilde{\gamma}_k$  are the differences in the nonlinear effects (or deviations from the linear effects);  $\Delta\eta_{ijk}$  is the difference in cell-specific error terms between the two strata.

The CSL-APC model, like the traditional APC model, is not identified without further assumptions because it is still the case that  $\text{Period} = \text{Age} + \text{Cohort}$  ( $j - j^* = i - i^* + k - k^*$ ). However, focusing only on identifying the differences in APC effects across strata, as opposed to both the strata-specific main effects and the differences in effects in Equation 3, requires assumptions only about the differences in effects.<sup>10</sup> This often entails weaker theoretical assumptions, intuitively because fewer are unidentified. We discuss this issue further below in our analysis of the gender gap in earnings.

The CSL-APC model is easily applied to aggregated APC data. Assuming that the data have been collected based on age and period, researchers can first construct two age-period Lexis tables, one for each group to be compared. For example, in our application one table would be a set of log median earnings for men while the other would be a set of log median earnings for women. Then a new Lexis table is created that is the difference between the two tables (in our case, women's log median earnings minus those of men). This differenced Lexis table, in which the cells denote cross-strata differences in the outcome, is then the data object used to fit the CSL-APC model. The data preparation procedure is conceptually identical when researchers have individual-level data as in our case. Appendix B provides the Lexis table used in our analysis.

### The Cross-Strata Canonical Solution Line

As noted above, the parameters for the linear components in Equation 4 are not identified. However, we will show that theoretically driven assumptions can place bounds on these parameters, allowing partial identification of the cross-strata APC effects. These bounds can be represented algebraically or graphically; both approaches are equivalent (Fosse and Winship 2019b). For simplicity, we present a graphical representation throughout the rest of the article (see Table 2 in Appendix C for an algebraic representation of the bounds).

An important graphical tool for understanding how the bounds work is the so-called *canonical solution line*, a line showing all possible estimates of the cross-strata linear APC effects (Fosse and Winship 2019a, 2019b). To understand the canonical solution line, it is important to be first aware of the two underlying parameters on which it is based:  $\theta_1$  and  $\theta_2$ . Although the parameters for the linear components are not identified, certain combinations of the parameters are (Holford 1983). Because of the identity  $\text{Period} = \text{Age} + \text{Cohort}$ , any APC model can be written as a function of just two of the three APC variables. Replacing the period index ( $j - j^*$ ) with the age ( $i - i^*$ ) and cohort ( $k - k^*$ ) indices in Equation 4 gives:

<sup>10</sup>For a detailed discussion of identification issues in a two group model using an alternative parameterization, see Nielsen (Unpublished).

$$\Delta Y_{i[i+k-l]k} = \Delta\mu + \theta_1(i - i^*) + \theta_2(k - k^*) + \Delta\tilde{\alpha}_i + \Delta\tilde{\pi}_{i+k} + \Delta\tilde{\gamma}_i + \Delta\eta_{i[i+k]k} \quad (5)$$

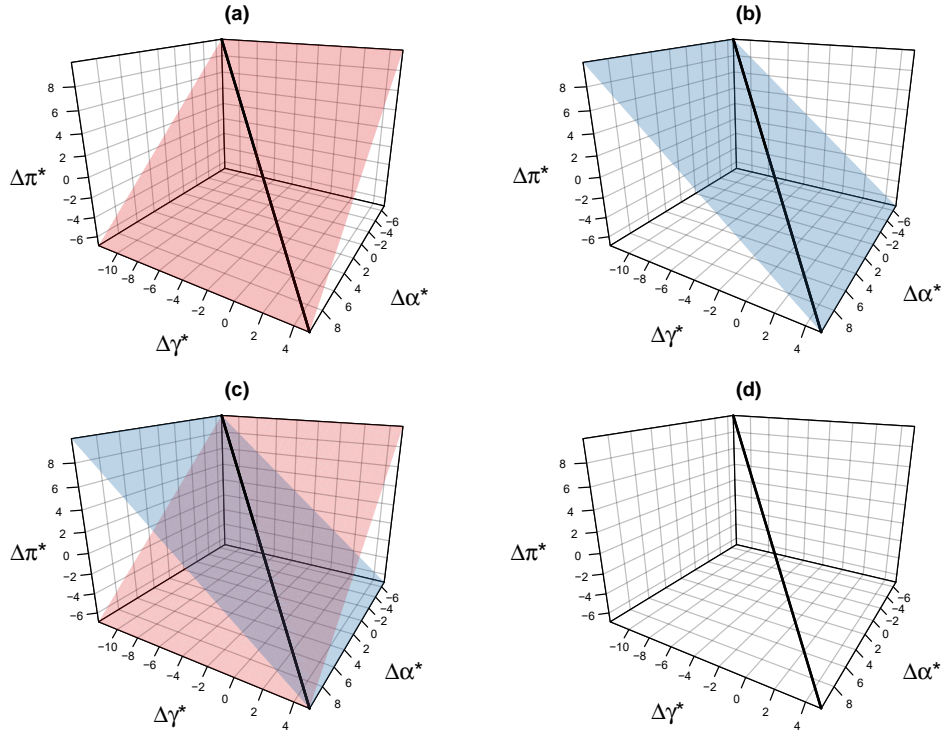
where, in this application,  $\theta_1 = \Delta\alpha + \Delta\pi$  and  $\theta_2 = \Delta\pi + \Delta\gamma$ . Following Fosse and Winship (2023), we interpret  $\theta_1$  as the *life cycle slope* because it describes the overall linear trend in group disparities across age levels (within any given cohort); likewise,  $\theta_2$  is the *social change slope* because it describes the overall linear trend in group disparities across cohorts (within any given age group). Note that in Equation 5 there are effectively only two linear parameters, not three. As such, both  $\theta_1$  and  $\theta_2$  are identified and can be estimated from the data.

Using our data on the gender wage gap in the United States, the least squares estimate of  $\theta_1$  is  $-0.002$  ( $p > 0.05$ ), meaning that a 10-year change in age is associated with a 0.002 (about 0.2 percent) decrease in women's earnings relative to men. The estimated  $\theta_2$  is  $0.084$  ( $p < 0.001$ ), meaning that a 10-year cohort change is associated with a 0.084 (about 8.7 percent) increase in women's relative earnings (see Table 3 in Appendix C for the full regression results). The number of observations is 72, which is simply the total number of age-period cells in the Lexis table defined by eight age categories and nine period categories ( $8 \times 9 = 72$ ). Appendix A and Appendix B explain the estimation process in more detail.

Next, it is essential to understand that the estimated  $\theta_1$  and  $\theta_2$  can severely restrict the possible estimates of the cross-strata age, period, and cohort linear effects. To see this, note that two equations  $\theta_1 = \Delta\alpha + \Delta\pi$  and  $\theta_2 = \Delta\pi + \Delta\gamma$  involve three unknowns of cross-strata linear effects (or slopes). Among the initial sets of differenced parameters, which could be anywhere in three-dimensional parameter space, these two equations, together with estimates of  $\theta_1$  and  $\theta_2$ , can constrain the possible estimates to only certain combinations of values that lie on a single line. This insight has not been widely recognized in the current APC literature in sociology and demography. As shown by Fosse and Winship (2018), the APC solution space can always be reduced to a one-dimensional space, or the *cross-strata canonical solution line*—the simplest geometric representation of the APC identification problem.

To illustrate this fact, consider Figure 3, which is based on simulated values of  $\theta_1 = 3$  and  $\theta_2 = -2$ . Panels (a) and (b) of Figure 3 display the age-period plane defined by the identified quantity  $\theta_1 = 3$  and the period-cohort plane defined by the identified quantity  $\theta_2 = -2$ , respectively. These two planes intersect to form a line, as shown in Figures 3 (c) and (d). Each point on this line represents a set of estimates for the parameters  $\Delta\alpha$ ,  $\Delta\pi$ , and  $\Delta\gamma$  that are consistent with the data. This visualization also represents the APC identification problem, since the absence of a linear dependence would cause three respective planes to intersect at a single point in parameter space.


Figure 3: Geometric Derivation of the Cross-Strata Canonical Solution Line



*Notes:* Representation of the canonical solution line using values of  $\theta_1 = 3$  and  $\theta_2 = -2$ . The vertical axis represents a range of cross-strata period linear effects, while the horizontal axes represent ranges for cross-strata cohort and age linear effects, respectively. Age-period plane is defined by  $\Delta\pi = \theta_1 - \Delta\alpha$ , where  $\theta_1 = 3$ . Period-cohort plane is defined by  $\Delta\pi = \theta_2 - \Delta\gamma$ , where  $\theta_2 = -2$ . The intersection of the two planes in the parameter space of cross-strata linear effects defines the canonical solution line.

Fosse and Winship (2019b) further demonstrate how the canonical solution line can be represented in two-dimensional space without loss of information. They call this a *2D-APC plot*, as shown in Figure 4. This two-dimensional representation is possible because of the linear dependence among the three temporal scales that occurs even in the analysis of cross-strata differences in effects.

Figure 4: 2D APC Plot of the Cross-Strata Canonical Solution Line



*Notes:* This figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left vertical axis represents a range of cross-strata age linear effects, the horizontal axis a range of cross-strata period linear effects, and the right vertical axis a range of cross-strata cohort linear effects. The axis labels are denoted with asterisks (\*) to indicate that the axes move along all possible linear effects consistent with data (the solid line) instead of the “true” linear effects. The dashed lines refer to the points where each respective axis is equal to zero. The solid line indicates the canonical solution line denoting all possible cross-strata linear effects consistent with the data. The empty circles represent the age-period (upper) and period-cohort (lower) origins where the respective set of axes are equal to zero.

Figure 4 illustrates three key elements of the data: the slope of the canonical solution line, which is always negative one, the direction and scale of the axes, and the values of  $\theta_1$  and  $\theta_2$ . While the slope and direction of the axes are equivalent for all temporally structured data, the values of  $\theta_1$  and  $\theta_2$  can vary depending on the  $\Delta\alpha$ ,  $\Delta\pi$  and  $\Delta\gamma$  parameters. These values determine the location of the canonical solution line in the 2D APC plot relative to the age-period and period-cohort origins.

### Point Identification: Equality-of-Effects Assumption

In the following sections, we outline how to use this 2D APC plot, along with theoretical assumptions, to bound the cross-strata differences in age, period, and cohort effects. To provide conceptual clarity on how to bound the parameters, we first consider the simplest case in which one imposes the assumption of equal effects across strata, thereby yielding point estimates. While we present this approach for illustrative purposes, we view this as a very strong assumption that may be theoretically justified only in particular applications.

The equality-of-effects (EOE) assumption involves presuming that one of the overall age, period, or cohort effects is the same across two different groups, thereby allowing for the identifica-

tion of differences in the other two temporal effects.<sup>11</sup> In general, we contend that two conditions must be met to justify the EOE assumption: first, there needs to be a strong theoretical reason to support the invariance of a total effect for age, period, or cohort (i.e., both the linear and nonlinear effects) across subpopulations; second, the nonlinear effects between the two groups must be observationally equivalent within some degree of uncertainty. We elaborate on each of these points below.

With respect to the first condition, it is critical to have a strong theoretical rationale to support the assumption of invariance of a total temporal effect across subpopulations. For example, epidemiological theories may arguably predict that certain biological mechanisms determine the age-related patterns of some health-related outcomes. If the mechanisms do not differ across subpopulations, it is plausible to assume that they are affected by the same age effects. For example, Riebler and Held (2010) examine APC effects on chronic obstructive pulmonary disease-specific mortality rates among men in England and Wales from 1950 to 1999. Their analysis assumes that men in England and men in Wales are exposed to the same general age effects, leading to the identification of period- and cohort-specific differences in mortality rates between the two groups of men.<sup>12</sup>

Figure 5 illustrates the impact of the EOE assumption for age on identification in the case of the gender wage gap. By assuming that the age effects are the same, we also assume that the age linear effects are the same for men and women, thereby identifying the between-gender differences in period and cohort effects on earnings. Graphically, point identification is achieved at the point where the dashed solution line intersects the red solid line specified by the assumption (i.e.,  $\Delta\alpha^* = 0$ ). As a result, among the innumerable sets of parameter estimates along the canonical solution line, a particular set of the estimates for cross-strata linear effects (or slopes) is identified.

However, we are not aware of any theory that justifies the equality of age effects on earnings between men and women, nor for period or cohort effects. It is well known that the age-related shifts in earnings are drastically different between men and women. This is because life-cycle patterns of family and demographic behaviors are closely intertwined with women's earnings, even conditional on full-time employment. Moreover, there is reason to believe that the period and cohort effects operate differently for male and female earnings. A large body of research evidences that the social environment has changed in favor of women's relative pay and that the composition of the female working population has shifted substantially across cohorts (see Blau and Kahn 2017 for a review).


---

<sup>11</sup>The EOE assumption is equivalent to a strategy adopted by Riebler and Held (2010, 2012). They show that if two groups are assumed to have the same total effect for one of either age, period, or cohort (including *both* linear and nonlinear effects), then group differences are identified for the other two temporal effects. Although their claim is correct, as we elaborate below, it is stronger than it needs to be and can be partially tested against the data. Regarding the first point, because only the cross-group linear effects are unidentified, an assumption about the equality of linear effects across the two groups is sufficient for identification. Moreover, the assumption that the total effects are equal ignores the fact that the nonlinear effects are identified and can be compared empirically.

<sup>12</sup>It is important to note that the identified cross-group effects of period and cohort are only as valid as the theory underlying this EOE assumption.



Figure 5: Identifying Constraint Based on the Age Equality-of-Effects Assumption with a 2D-APC Plot




*Notes:* The figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left vertical axis represents a range of cross-strata age linear effects, the horizontal axis a range of cross-strata period linear effects, and the right vertical axis a range of cross-strata cohort linear effects. The dashed line indicates the canonical solution line denoting all possible cross-strata linear effects consistent with the data. The dot refers to the point where the solution line intersects the red solid line (i.e., the cross-strata age linear effect equals zero) as stated by the EOE assumption for age.

The second condition for justifying the EOE assumption is that the nonlinear effects of the temporal scale of interest (either age, period, or cohort) need to be similar between two groups within some degree of uncertainty. In practice this means that the overall shape of the nonlinear effects should be similar between the groups being compared.<sup>13</sup> We propose three specific ways to detect the possible (in)equality of the nonlinear effects, using the gender wage gap as an example.

First, as an informal “test,” researchers can graphically examine whether each cross-strata nonlinearity (e.g.,  $\Delta\tilde{\alpha}_i$  for each  $i = 1, \dots, I$ ) is close to zero. The estimated cross-strata nonlinear effects of age, period, and cohort are shown in Figure 6. As can be observed, it does *not* seem to be the case that the nonlinear effects on the gender wage gap are close to zero, either for age, period, or cohort. If they were, we would see straight flat lines in the three panels.

<sup>13</sup>Certainly some part of the nonlinear effects will reflect noise. Although this is not the focus of our discussion here, one solution is to smooth the nonlinear effects by, for example, dropping higher-order polynomials or imposing strong zero-centered prior distributions over the nonlinearities (e.g., see Fosse 2021).

Figure 6: Estimated Cross-Strata Nonlinear Effects of Age, Period, and Cohort on the Gender Wage Gap



Notes: The panels (a), (b), and (c) show the estimated cross-strata nonlinear age, period, and cohort effects on the gender wage gap, respectively. The dotted lines indicate the mean value (i.e., intercept) of the gender wage gap.

Second, in addition to an informal visual inspection, one could perform a formal statistical test of nonlinear effects. Specifically, an F-test can be used to assess whether nonlinear effects of age, period, or cohort are jointly different between two groups (e.g.,  $\tilde{\alpha}_i = 0$  for each  $i$ ). Although the inability to reject the null hypothesis of zero (i.e., equality between two groups) does not necessarily confirm that the nonlinear effects are equal, the test at least provides face validity to the claim that the nonlinear effects are not substantially different between the two groups, especially when the sample size is large. In our example, the F-tests shown in Table 1 indicate that we can reject the null of no differential nonlinear effects with respect to age, period, and cohort.

Table 1: Joint F-Test Results of Cross-Strata Nonlinear Effects and Model Fit Statistics

Test Type	Test of Nonlinear Effects			Model Fit	
	$H_0$ in F-test	$\chi^2$ Statistic	$p$ -value	AIC	BIC
Full CSL-APC Model	-	-	-	-363.93	-293.35
No Between-Gender Age Nonlinearities	$\Delta \tilde{\alpha}_i = 0$ for all $i$	516	$p < 0.001$	-189.68	-132.77
No Between-Gender Period Nonlinearities	$\Delta \tilde{\pi}_j = 0$ for all $j$	132.38	$p < 0.001$	-275.43	-220.79
No Between-Gender Cohort Nonlinearities	$\Delta \tilde{\gamma}_k = 0$ for all $k$	197.07	$p < 0.001$	-266.71	-228.01

Lastly, researchers can rely on model fit statistics to examine whether omitting the cross-strata nonlinear effects significantly reduces model fit. The Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC) values are presented in Table 1. We can see that the full CSL-APC model with parameters for all of the between-gender nonlinear effects shows the best model fit

(i.e. the smallest AIC/BIC statistics) in the case of the wage gap. This means that it is difficult to claim that the cross-strata nonlinear effects do not play a role in shaping the gender wage gap.

In sum, when researchers have solid evidence that the linear effect is the same and the nonlinear effects between two groups are observationally indistinguishable for one of the three APC variables, they can make use of the EOE assumption to identify cross-group differences in the remaining two temporal effects. However, this is not applicable in the case of the gender wage gap, as there is no strong evidence to suggest that the age, period, or cohort effects are equal between men and women. The results of the three tests conducted also do not support that the nonlinear effects are the same for both groups. In the following, we present our more general bounding approach to the identification problem, which can flexibly incorporate more realistic assumptions supported by theories of the gender wage gap. The EOE assumption-based solution outlined above can be considered as one special, restricted case of this bounding approach.

## **Partial Identification: Bounding Analysis of Cross-Strata APC Effects**

Our cross-strata bounding approach builds on the framework proposed by Fosse and Winship (2019b), which is based on the insight that constraints implied by theoretical claims, along with information from the data, can be used to bound one or more of the APC effects. Despite its flexibility in incorporating theoretical assumptions, Fosse and Winship's bounding approach has not been readily applied to the temporal analysis of group disparities. We adapt their framework to develop a bounding approach that partially identifies cross-strata differences in APC effects (or "cross-strata effects"). It is critical to understand that although the analytical procedures might appear similar, our estimands differ substantially from those of a group-specific APC analysis; rather than overall main effects, our approach focuses on identifying cross-group differences in effects.<sup>14</sup> Below we elaborate on the process of cross-strata bounding analysis step by step.

### *Step 1: List Set of Credible Theoretical Assumptions*

The bounding approach requires assumptions about the size, sign, or shape of the cross-strata APC effects. These assumptions serve to constrain the possible region of the cross-strata canonical solution line shown in Figure 4. Because the assumptions must be informed by theories about the differences in temporal effects between groups, researchers should carefully review their assumptions in light of various (and possibly conflicting) theories to ensure that the assumptions are as credible as possible.

With the flexibility of a bounding approach, researchers can invoke a wide range of assumptions about cross-strata age, period, or cohort effects. We illustrate three types of assumptions that we

---

<sup>14</sup>An alternative strategy to our cross-strata bounding approach would be to perform a bounding analysis for each respective group and then compare the bounded APC effects between the two groups. We indeed encourage applied researchers to rely on this approach involving multiple stratum-specific bounding analyses if there are sufficient theoretical foundations to yield precisely bounded temporal effects within each comparison group. In many substantive cases, however, researchers may not be able to find such rich theoretical support separately for each group.

think are particularly useful for applied researchers:

1. *Monotonic Effects Assumption:* This assumption states that differences in age, period, or cohort effects across strata are monotonically increasing (or decreasing) over a specified range. The wider the specified range, the stronger the assumption involving monotonicity of cross-strata effects. An assumption of monotonic cross-strata effects involves imposing a constraint on the total cross-strata effects (i.e., the combination of the linear and nonlinear effects) of age, period, or cohort over the specified range. Because we can identify and estimate the nonlinear effects, the monotonic effects assumption implies a constraint on the unknown linear cross-strata effect. If one of the three cross-strata linear effects (or slopes) is bounded, this implies bounds on the other two slopes. This is simply an extension of the idea that for any given value of the cross-strata linear effect, the canonical solution line determines the value of the remaining two linear effects. As a result, bounds on all cross-strata effects are identified.
2. *Non-Monotonic Effects Assumption:* This assumption is that the cross-strata effects do not increase (or decrease) monotonically over a given range of ages, periods, or cohorts. The cross-strata effects may decrease, remain constant, or increase over the specified range, but the increase (or decrease) is not monotonic. In this sense, the non-monotonic effects assumption is mutually exclusive with, and the opposite of, the monotonic effects assumption. Unlike the monotonic effects assumption, the non-monotonic effects assumption becomes weaker as the specified range increases. However, like the monotonic effects assumption, the non-monotonic effects assumption imposes a constraint on the total cross-strata effects of age, period, or cohort over the range specified by the researcher. Because the nonlinear effects are identified, this results in a bound on one of the cross-strata APC linear effects. This in turn yields bounds on the other two temporal effects.
3. *Linear Effects Assumption:* This assumption states that over the full range observed in the data, one of the linear age, period, or cohort effects is assumed to diverge, converge, or remain the same across groups. In practice, this means that one is specifying the direction of one of the underlying cross-strata linear effects (or slopes). Importantly, this assumption is limited to the linear effects and does not imply that the total age, period, or cohort effects always diverge, converge, or remain the same across groups. The EOE assumption explained in the previous section is a more restrictive case of this assumption, where not only a slope but also the corresponding cross-strata nonlinear effects are assumed to be zero.

The above three types of assumptions are quite general and can be applied to any number of contexts. In practice, theoretical assumptions will vary depending on the outcome of interest and the subpopulations being compared. Based on the literature reviewed earlier, we propose the following assumptions in the context of the gender gap in earnings.

1. *Monotonic Age Effects (25-34)*: We assume that the differential age effects between men and women will expand the wage gap from age 25 to 34. This is because this age range overlaps substantially with women's prime childbearing ages, and fertility is shown to have differential effects on the wages of men and women. The extensive literature on the motherhood wage penalty suggests that mothers' earnings relative to fathers, net of full-time employment, will decrease through pathways involving reduced work hours, employer discrimination, seeking family-friendly jobs with lower earnings, and so on (Budig and England 2001; Correll et al. 2007; Killewald and García-Manglano 2016; Yu and Kuo 2017). Conversely, a related literature suggests that men will receive a fatherhood wage premium through this age range (Hodges and Budig 2010; Killewald 2013).
2. *Non-Monotonic Age Effects (35-49)*: Another assumption we rely on is that the gender differences in age effects will not necessarily increase the wage gap in a monotonic fashion from 35 to 49. The gap may increase over a shorter age range (from 35-39 to 40-44 or from 40-44 to 45-49), but it will not increase monotonically from age 35 to 49 (not *both* from age 35-39 to 40-44 *and* from age 40-44 to 45-49). This assumption about the difference in age effects between women and men, like the previous one, is based on life-cycle patterns of fertility. As children reach school age, the childcare burden on parents tends to decrease, especially for mothers between 35 to 49. Even conditional on full-time employment, mothers are likely to increase their working hours or move to higher paid positions in the workforce (Musick et al. 2020).
3. *Period Linear Effect*: We also impose a sign constraint on the gender difference in period linear effect (or period slope), assuming that it is positive (i.e., leads to a smaller wage gap). This assumption is based on a large body of research claiming that the social environment has, at least to some extent, become more favorable to women's pay relative to men's since the 1970s (see Blau and Kahn 2017). For example, anti-discrimination practices and changes in laws and policies likely increased women's relative pay on average over this period (Kurtulus 2012). Declining unionization rates and a declining share of manufacturing jobs, along with an increase in well-paying service sector jobs, may have reduced the relative earnings of men, who were disproportionately employed in manufacturing and in sectors with high unionization rates (Blau and Kahn 1997; Borghans et al. 2014). Thus, we expect these differential impacts on men and women to manifest as a positive period linear effect on women's relative pay.
4. *Cohort Linear Effect*: Lastly, we also consider an assumption that the cohort slope is positive (i.e., leads to a smaller wage gap). This assumption is based on human capital theory, which predicts that shifts in the female labor force composition across cohorts have led to an overall increase in women's relative pay (Goldin 2021). This is evidenced by the fact that the educational attainment of recent female cohorts is higher than that of their earlier coun-

terparts, and that the rate of educational expansion has been much greater for women than for men. The duration of work interruptions, an important determinant of earnings loss, is also shorter on average for recent cohorts of women, partly due to their lower fertility (Blau and Kahn 2017; Killewald and Cricco 2020). These factors are considered to be primarily cohort-specific characteristics and are therefore expected to result in a positive linear effect on the gender wage gap.

### *Step 2. Estimate and Display Set of Linear and Nonlinear Effects*

In the second step, we will begin by estimating  $\theta_1$ ,  $\theta_2$ , and the cross-strata nonlinear effects. These quantities are all identified and can therefore be estimated from the data. Specifically, to estimate these parameters, we fit the model in Equation 5 to an age-period Lexis table in which the cell values are outcome differences between the two groups. The estimation results for our example are reported in Table 3 in Appendix C. We then use the values of  $\hat{\theta}_1$  and  $\hat{\theta}_2$  to construct the canonical solution line. The cross-strata canonical solution line and nonlinear effects can be presented visually, as shown in the 2D APC plot (Figure 4) and cross-strata nonlinear effect plots (Figure 6).

### *Step 3. Compute Bounds on the Age, Period, and Cohort Effects*

The third step in the analysis is to determine the bounds on the age, period, and cohort effects based on the specified assumptions. By calculating the minimum and maximum slope values for the temporal scale of interest that are consistent with a given assumption, researchers can assess the limits and variability of the possible parameter estimates. Although it is possible to evaluate each assumption separately, in our example we choose to apply the first two age-related assumptions simultaneously because of their common theoretical basis.

The general principle of the bounding approach is that researchers start with the assumptions that are most credibly supported by theories. In our case, we believe that the life-cycle patterns of fertility reasonably support our assumptions about age effects, particularly given the large literature demonstrating the existence of a motherhood wage penalty. Additionally, since we are primarily interested in distinguishing between period and cohort effects, it seems natural to apply the age-related assumptions first.

Our first assumption is that women's earnings have declined relative to men's earnings between ages 25 and 34. For this assumption to hold, the age slope must be less than about 0.116, which is twice the difference between the nonlinear effects of ages 25-29 and ages 30-34 (see Figure 6 and Appendix C Table 3).<sup>15</sup> If the age slope is equal to or greater than 0.116, then women's relative earnings would increase or at least remain constant over this age range, violating this assumption.


Our second assumption is that women's relative earnings do not necessarily decline monotonically from age 35 to 49. For this assumption to be satisfied, the age slope must not take the value that results in a monotonic decline in women's relative earnings over this age range. In this age

<sup>15</sup>We multiply by two just to make the coefficient scales consistent (i.e., a 10-year change).

span, the most positive shift in the nonlinear effects is 0.031 (between ages 40-44 and ages 45-49; see Figure 6). Therefore, the age slope must be equal to or greater (or less negative) than  $-0.031$ ; otherwise, the negative age slope will offset the positive shift in the nonlinear effects and result in monotonically decreasing relative earnings for women.

From these two assumptions we can derive bounds on the minimum ( $-0.031$ ) and maximum ( $0.116$ ) values of the age slope. These bounds on the linear effects can be easily visualized using a 2D APC plot, as shown in Figure 7. This graph shows that the possible region of the canonical solution line is now restricted to the range between  $-0.031$  and  $0.116$  in terms of the age slope.

Figure 7: Upper and Lower Bounds of Cross-Strata APC Linear Effects on the Gender Wage Gap



*Notes:* The figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left vertical axis represents a range of cross-strata age linear effects, the horizontal axis a range of cross-strata period linear effects, and the right vertical axis a range of cross-strata cohort linear effects. The dashed line indicates all possible linear effects consistent with data. The solid line in the colored region refers to the feasible region of cross-strata linear effects given the first two assumptions about the shape of the cross-strata age effects.

Given the constraints on the age slope, we can also compute the minimum and/or maximum values of the period and cohort slopes. Since the estimated sum of the between-gender age and period linear effects ( $\hat{\theta}_1$ ) is  $-0.002$ , the resulting minimum and maximum values of the between-gender period linear effect are  $-0.118$  and  $0.029$ , respectively. Similarly, the bounds on the period slope translate into the bounds on the cohort slope (from  $0.054$  to  $0.202$ ), since we estimated the sum of the period and cohort slopes to be  $0.084$ . These calculations are easily visualized in Figure 7. Since the plot is based on the values of  $\hat{\theta}_1$  and  $\hat{\theta}_2$ , restricting the range of the age slope automatically translates into restricting the ranges of the period and cohort slopes. As a result, we have now

partially identified the between-gender linear effects of age, period, and cohort based on the two theoretical assumptions about the differential age effects between women and men.

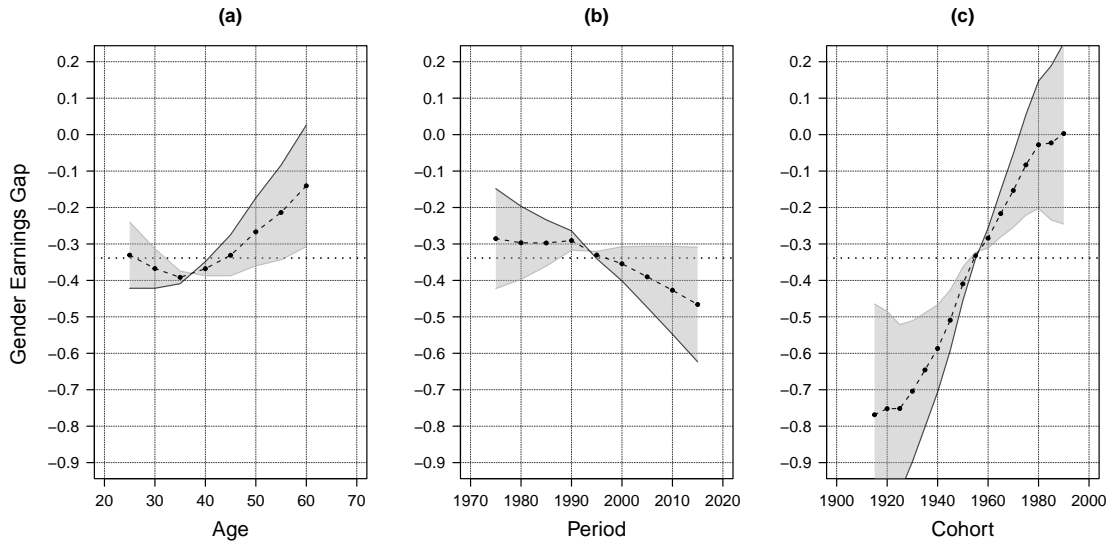
#### Step 4. Partial Identification Based on Bounding Constraints

Based on the constraints on the APC slopes, we can construct bounds on the overall age, period, and cohort effects producing cross-group disparities. The following Equation 6 formulates the bounded APC effects. While its format is similar to Equation 4, we add a bounding scalar  $\nu$  to each cross-strata linear effect as follows:

$$\Delta Y_{ijk} = \Delta\mu + (\Delta\alpha + \nu)(i - i^*) + (\Delta\pi - \nu)(j - j^*) + (\Delta\gamma + \nu)(k - k^*) + \Delta\tilde{\alpha}_i + \Delta\tilde{\pi}_j + \Delta\tilde{\gamma}_k + \Delta\eta_{ijk}, \quad (6)$$

Setting different values of  $\nu$  yields different possible values on the solution line. If  $\nu$  is set to the maximum value, the age slope will be the maximum, the period slope will be the minimum, and the cohort slope will be the maximum. If  $\nu$  is set to the minimum value, the opposite is true. The range of a total age effect in a given age category  $((\Delta\alpha + \nu)(i - i^*) + \Delta\tilde{\alpha}_i)$  can then be constructed and plotted, as can the period and cohort effects. In the case of the gender wage gap, the bounded cross-strata APC effects based on the two age-related assumptions are shown in Figure 8. The results suggest that the cross-strata cohort effects may have played a crucial role in increasing women's relative earnings compared to the cross-strata period effects. However, the bounds of the overall between-gender APC effects remain wide, precluding a more precise understanding of each temporal effect.

Figure 8: Bounded Cross-Strata Age, Period, Cohort Effects on the Gender Wage Gap



*Notes:* The shaded areas represent the bounded effects of age (a), period (b), and cohort (c) on the gender wage gap based on the two age-related assumptions. The dotted lines follow the mid-points in each shaded area. The dark bold lines along one end of the shaded areas depict one possibility of the APC effects where the between-gender age linear effect is most positive, the between-gender period linear effect is most negative, and the between-gender linear cohort effect is most positive within the bounded ranges.




*Step 5. Repeat Steps 3-4 Invoking an Additional Theoretical Assumption*

Although the results so far provide a meaningful conclusion about the relative importance of period and cohort effects in driving gender wage convergence, the assumptions about age effects alone do not provide a precise understanding of the temporal effects on the gender wage gap, especially the period effects. Researchers can repeat Steps 3 and 4, relying on additional theoretical assumptions that may provide a narrower range of bounds on the cross-strata temporal effects.

In our example, we proceed with the assumption that the period slope is positive ( $\Delta\pi > 0$ ). We now have three bounding constraints at hand: 1)  $\Delta\alpha > -0.031$ ; 2)  $\Delta\alpha < 0.116$ ; and 3)  $\Delta\pi > 0$ . The resultant bounds in all slopes can be computed manually. Alternatively, these constraints can be depicted in the 2D-APC plot, serving to illustrate further restrictions in the possible region of the solution line. Figure 9 shows the 2D-APC plot where all three bounding constraints are introduced. The possible minimum and maximum values are  $-0.031$  and  $-0.002$  for the age slope, zero and  $0.029$  for the period slope, and  $0.054$  and  $0.084$  for the cohort slope. Note that the resultant constraints on the slopes are different from the ones originally specified by the assumptions because the bounds for one slope are mutually defined by constraints on the other slopes. The figure shows that the bounds on each slope have significantly narrowed down compared to Figure 7 due to the newly introduced constraint from the positive period slope.

Figure 9: Upper and Lower Bounds of Cross-Strata  
APC Linear Effects on the Gender Wage Gap

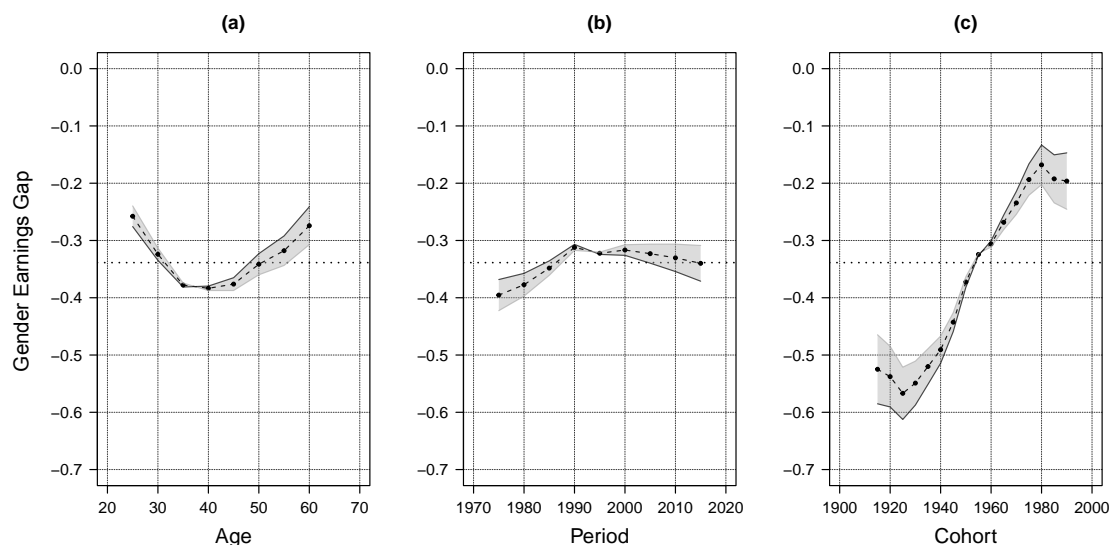


*Notes:* The figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. The solid line in the overlapping colored regions (red indicates age-related assumptions; blue indicates a period-related assumption) refers to the possible area of the linear effects if the three assumptions about age and period effects were satisfied.

Figure 10 shows the bounded APC effects based on the three theoretical assumptions about age and period effects. We assumed that 1) women's relative earnings decline from ages 25-29 to 30-34; 2) women's relative earnings do not necessarily decline monotonically from ages 35-39 to 45-49; and 3) the between-gender period linear effect is positive such that it is favorable to gender wage convergence. The estimated APC effects shown in the figure prove to be almost as informative as the point estimates. The results suggest that progress in gender convergence in earnings has been largely driven by cohort effects, and that cohort effects have continued to the present. This may be due either to progress in human capital accumulation among the female labor force entering the labor market, to the exit from the labor market of earlier female cohorts who tend to share the characteristics associated with lower average earnings (e.g. high fertility rates), or to declining economic prospects among the newly entering male labor force. It is reassuring that this finding on cohort effects is not based on theoretical assumptions about cohort effects, but is derived from the age and period-related assumptions alone. This strategy for making confident inferences about temporal effects is related to the principles we described in Step 1 about the order in which different theoretical assumptions need to be introduced: Researchers may want to start with the most credible assumption that does not involve the temporal scale in which researchers are most

interested.

Figure 10: Bounded Cross-Strata Age, Period, Cohort Effects on the Gender Wage Gap



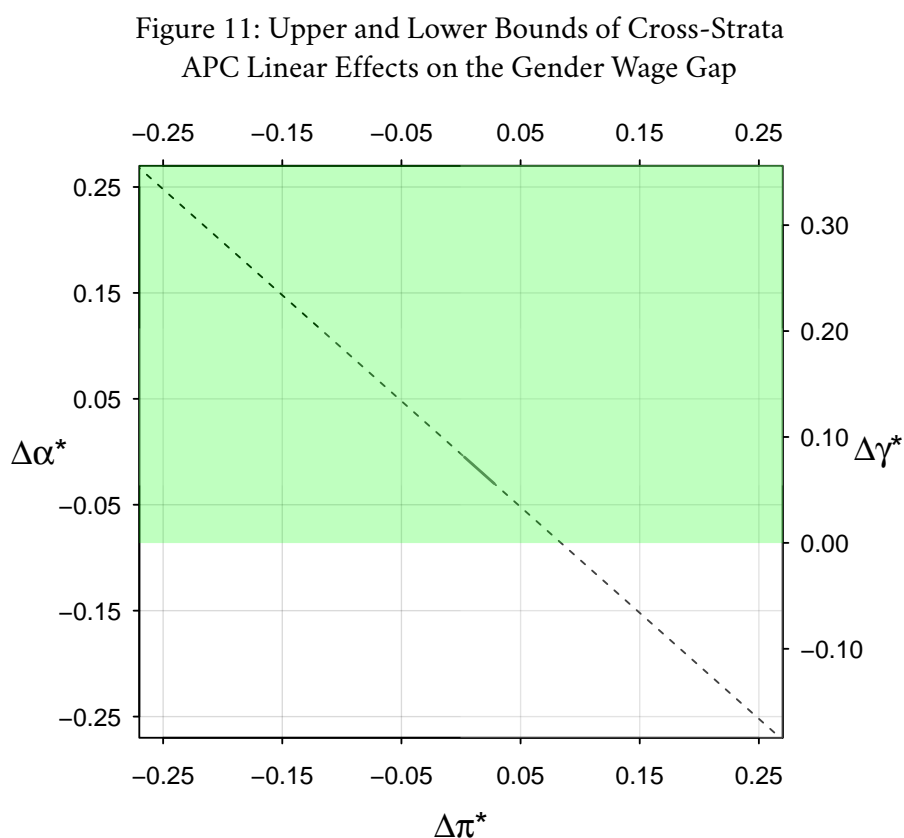
*Notes:* The shaded areas represent the bounded effects of age (a), period (b), and cohort (c) on the gender wage gap based on the three assumptions about age and period effects. The dotted lines follow the mid-points in each shaded area. The dark bold lines along one end of the shaded areas depict one possibility of the APC effects where the linear age effect is most positive, the linear period effect is most negative, and the linear cohort effect is most positive within the bounded ranges.

A new finding in Figure 10 as compared to Figure 8 is that the slowing gender convergence in earnings since the 1990s is likely due to period effects that have been stagnant or may have even become slightly negative since the beginning of the 1990s. The period effects since the 1990s are consistent with the argument that labor demand shifts in terms of industries and occupations favoring women decreased in the 1990s than in the 1980s (Blau and Kahn 2006). Prior scholarship also posits that the ideology and practices of intensive mothering appeared since the mid-1990s (Hays 1996), which may have prevented mothers from making more wage gains by increasing labor demands from home. An increase in wage returns to overwork could have also benefited men's wage gains who disproportionately work long hours (Cha and Weeden 2014).

Finally, age effects on women's relative pay are estimated to follow the U-shaped pattern. Age effects have the most severe effects on the gender wage gap during age 35 to 49. After about age 50, women begin to show a recovery in their earnings relative to men's.

Researchers may want to further introduce the remaining set of theoretical assumptions if their estimation results are not precise enough to answer the research question at hand. They might want to do so also when they are almost equally confident about two assumptions and wonder if an additional introduction of the remaining assumption may affect the estimation results. Since the estimation results about each unique temporal effect are very precise in our example, it is no longer necessary to adopt further identifying constraints. Still, we can also test out the positive cohort slope assumption, as we view this assumption as equally compelling as the positive period slope

assumption. Figure 11 describes the constraint newly introduced by the positive cohort slope assumption. The solid line indicates the solution line bounded by the previous age and period-related assumptions. The figure shows that the positive cohort slope assumption does not further narrow down the possible region of the solution line, as the constraint imposed by the assumption is located outside the existing possible region of solution line (the solid line). Therefore, the resultant bounded effects of age, period, and cohort are the same as in Figure 10. The results confirm that the order of the assumption introduced (the period slope assumption vs. the cohort slope assumption) does not affect the estimation results. These results also increase our confidence in the conclusion with regards to the key importance of cohort replacement in determining changes in women's relative pay since the conclusion does not rely on any assumption about the cohort effect itself.



*Notes:* The figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. The solid line in the colored region refers to the possible area of the linear effects if the previous three assumptions about age and period effects were satisfied. The colored region indicates the bounds of linear APC effects imposed by the positive cohort-slope assumption.

In sum, the bounded effects shown in Figure 10 support the conclusion that cohort replacement has been fundamental to the observed convergence in earnings between women and men. The slower rate of change since the 1990s appears to be driven by stagnating or maybe even slightly declining period effects. These findings shed light on the temporal process by which a categorical inequality in the labor market, in this case gender inequality, can change and face challenges in

making further progress.<sup>16</sup>

## Sensitivity Analyses

As has become more common in empirical social science research, researchers are expected to carry out sensitivity analyses in order to assess the robustness of their findings. This is not something Fosse and Winship (2019b) considered. We discuss two types of sensitivity analyses. First, theoretical assumptions made by the researchers may be too strong to be justified, and the bounds in the temporal effects might be biased as a result. Second, researchers might be concerned if the constraint resulting from a monotonic or non-monotonic effect assumption is driven by an extreme component in one of the nonlinear effects.

First, test results implying the inconsistency of the assumption set with the data can serve as a sufficient condition for the invalidity of the assumption set.<sup>17</sup> Researchers can test the validity of an assumption set based on the constraints imposed on the canonical solution line. The logic is that if the assumption set were not feasible, bounding constraints implied by the assumption set would not be consistent with any combination of cross-strata linear age, period, and cohort effects. Graphically, this means that there would be no remaining region of the solution line that is consistent with the assumption set. For example, suppose that we assumed the *monotonically increasing* effects of period on women's relative pay, instead of the positive period slope. That is, the period effects contribute to increasing women's relative pay across every successive period, and there is no interval where the more recent period's effect is the same (stagnant) or less positive (decrease) than the preceding period. This monotonic increase assumption in terms of period, in addition to the previous two age assumptions, would give the following bounding constraints: 1)  $\Delta\alpha > -0.031$ ; 2)  $\Delta\alpha < 0.116$ ; and 3)  $\Delta\pi > 0.035$ .


These constraints are depicted in the 2D-APC plot in Figure 12. The figure demonstrates that there is no region of the solution line consistent with these constraints. This is because the period slope is bounded between  $-0.118$  and  $0.029$  by the age constraints, while the constraint that  $\Delta\pi > 0.035$  places the minimum value of the period slope outside this bound. The results show that the monotonic-increase assumption on period effects, in combination with the two age assumptions, are not consistent with data.

---

<sup>16</sup>There has been some debate about whether selection into full-time work may have driven the observed time trends in the gender wage gap. However, a recent study by Blau et al. (2021) reviews previous evidence on this debate and casts doubt on this interpretation.

<sup>17</sup>Still, it is important to note that the test results *not* indicating the inconsistency do not necessarily ensure the validity of the assumption set, nor is it viable to assess the validity of *each* theoretical assumption among the assumption set against data separately.

Figure 12: Upper and Lower Bounds of Cross-Strata  
APC Linear Effects on the Gender Wage Gap




*Notes:* The figure is based on the  $\theta_1$  estimate of  $-0.002$  and the  $\theta_2$  estimate of  $0.084$ . The left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. There is no region where all three colored regions (each representing the age (red), period (blue), or cohort (green) bound) overlap, meaning that there is no possible area of the linear effects if all the assumptions introduced here were satisfied.

Second, researchers might be concerned if the resultant constraint is driven by a single extreme component of a nonlinear effect. The concern can arise because bounding constraints concerning monotonic/non-monotonic effects are determined by the segment of nonlinear effects that has the maximum or minimum slope. This segment might differ to an unusual degree from other neighboring segments of the nonlinear effects. Researchers can inspect what nonlinear effect is directly responsible for the constraint made and evaluate the extent to which they are uncertain about the estimate of this nonlinear effect.

In our empirical example, we have two assumptions that state (non-)monotonicity of the age effects. Figure 13 shows which of the nonlinear effects is associated with the resultant constraint on the bounded age effects. For the first assumption involving the monotonic decrease in women's relative earnings over the age range 25 to 34, the corresponding interval linking the age category 25-29 and the category 30-34 is colored red in the upper panel. For the second assumption about the non-monotonic decrease in women's relative earnings, the corresponding intervals from age 35-39 to 45-49 are colored red in the bottom panel. Among the intervals colored red, those directly responsible for the bounding constraint are expressed with solid lines, and the other red lines are expressed with dotted lines. By visualizing these intervals, researchers can evaluate their confi-

dence in the intervals that are directly associated with the bounding constraint (i.e., the red solid lines). For example, if the slope of the red solid line is quite different from neighboring intervals in a way that is not supported by theories, they may wonder if any bias has been induced into any of the estimated nonlinear effects linking the red solid line. Likewise, if they are uncertain about the estimated nonlinear effect mostly due to sparse observations in the corresponding temporal category, especially at the tails of the cohort categories, they can smooth the estimated nonlinear effects by reducing the degree of polynomials and repeat the analysis. Since our data set comprises a large number of respondents, we are confident in the estimated nonlinear effects. Also, the red solid line in the bottom panel of Figure 13 is not substantially different in its steepness from the preceding interval (i.e., the red dotted line), which relieves the concern about the credibility of the constraint imposed by the second assumption.

Figure 13: Estimated Cross-Strata Nonlinear Effects Governing the Bounds Imposed by (Non-)Monotonicity Assumptions



*Notes:* In the upper panel, the interval governing the bounds imposed by the age monotonicity assumption is denoted with a red solid line. In the bottom panel, the interval involving the age non-monotonicity assumption is colored red (either with a dotted or solid line), and the interval directly relevant to the bounds imposed by the assumption is denoted with the red solid line.

## Discussion and Conclusion

The analysis of group disparities has been a central empirical agenda in sociology and other social science disciplines. Disparities in socially valued resources across race, ethnicity, gender, education, geographic region, and other dimensions of social categories reflect social inequality produced by underlying stratification processes (Grusky 1994). The extent of group differences in public opinion is also a primary focus of scholars interested in social polarization and cohesion (Hout et al. 2022). Sociologists have also sought to understand the broader process of social change through the study of group differences, which can illuminate, for example, how attitudes toward emerging social issues diffuse across socioeconomic groups (Esping-Andersen and Billari 2015; Pampel and Hunter 2012).

In this article, we first developed the Cross-Strata Linearized APC (CSL-APC) model, which focuses on understanding the separate cross-group effects of age, period, and cohort on an observed disparity. Similar to the traditional APC model, the cross-strata linear age, period, and cohort effects are not identified in the CSL-APC model. To address the identification problem, we show how the bounding approach of Fosse and Winship (2019b) can be extended to examine differences in age, period, and cohort effects across strata. We define cross-strata effects as the estimands of interest, rather than the separate effects for each group, and use theoretically driven assumptions to achieve partial identification of these temporal effects. Our approach involves setting bounds on the cross-strata effects, starting with identifying what can be known from the data alone with as few restrictions as possible. Using data on the wage gap between U.S. men and women and under a limited set of assumptions, our analysis shows that cohort replacement effects have driven continued progress in women's relative pay. Yet, this progress has been partially offset by stagnating period effects since the 1990s. These results are generally consistent with those of Campbell and Pearlman (2013), although we support this conclusion based on a weaker set of theoretically driven assumptions.<sup>18</sup>

The approach outlined in this article has several advantages. First, our method is more general and flexible because it allows for various constraints on the size, shape, or sign of one or more of the cross-strata parameters, rather than specifying only one type of constraint. Second, our constraints typically involve weaker theoretical assumptions than those commonly used in previous methods and, at least in some applications, can provide quite narrow bounds solely from general theoretical assumptions about life course effects. In particular, by focusing on cross-strata differences rather than the main effects, we can obtain more credible results in many applications where the assumptions made are more plausible for cross-strata differences in effects than for strata-specific effects. Third, we provide two types of sensitivity analyses that researchers can use to assess the credibility of their results, a novel contribution to the APC literature. Finally, we have outlined a general five-step procedure for bounding APC effects to guide future research, and some of the steps can

---

<sup>18</sup>The assumption for their main analysis is that the age, period, and cohort effects are the same within each five-year category of the temporal scales, which is effectively a kind of equality constraints approach.



be useful to any practitioner analyzing temporally structured data.

Further research on cross-strata differences in APC effects could focus on a few key areas. First, we have focused on model identification but have not addressed the quantification of uncertainty due to sampling variability. Although our CPS data have a large sample size and we are less concerned about the uncertainty of our results, sampling variability can make bounds too broad when small samples are used. Therefore, we suggest that researchers consider developing techniques to quantify this uncertainty, such as those based on bootstrapping, especially when using small samples. Second, a fruitful direction for future research is to conduct a sensitivity analysis of a cross-strata APC model that includes mechanisms or proxy variables (Winship and Harding 2008; see also Fosse and Winship 2019a). This approach would involve starting with strong assumptions about mechanisms—first aiming for a point estimate—and then weakening those assumptions, yielding a set of upper and lower bounds on the cross-strata effects. This may be useful for assessing the robustness of the estimated results when important mechanisms or proxies are thought to be missing from the data. Lastly, future research can explore the Bayesian interpretation of our framework by imposing prior distributions on one or more of the parameters of the CSL-APC model. A Bayesian approach would require just as strong assumptions as the approach used here, but some analysts may be attracted to a Bayesian approach in part because it offers a wider range of ways to constrain the set of possible linear effects (e.g., see Fosse 2021).

In conclusion, the methods presented in this article provide a coherent, step-by-step approach to partially identifying cross-strata effects using the least stringent assumptions possible. It must be recognized, however, that the results of any APC analysis, including the cross-strata approach presented in this article, are not fully verifiable or falsifiable from the data alone. In contrast to purely descriptive analyses that focus, for example, on estimating observed intra- and intercohort trends or marginal period trends (see, e.g., Fosse 2023; Fosse and Winship 2023), APC analyses necessarily require assumptions that are external to the data (Fienberg 2013). The validity of these assumptions ultimately depends on the soundness of the social, biological, or cultural theory on which they are based, which may be flawed and lead to erroneous conclusions. Therefore, it is essential that researchers conducting a cross-strata APC analysis not only emphasize the tentative nature of their findings, but also carefully triangulate their results using different sets of credible, theoretically-based assumptions. Our bounding approach provides just such a conceptual framework and methodology for researchers interested in leveraging information about subpopulations to uncover cross-group age, period, and cohort effects.

## References

- Bell, Andrew and Kelvyn Jones. 2018. "The Hierarchical Age-Period-Cohort Model: Why Does It Find the Results That It Finds?" *Quality & Quantity* 52:783–799. doi:10.1007/s11135-017-0488-5.
- Blau, Francine D., Mary C. Brinton, and David B. Grusky, eds. 2006. *The Declining Significance of Gender?* New York, NY: Russell Sage Foundation.
- Blau, Francine D. and Lawrence M. Kahn. 1997. "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s." *Journal of Labor Economics* 15:1–42. doi:10.1086/209845.
- . 2006. "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence." *ILR Review* 60:45–66. doi:10.1177/001979390606000103.
- . 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55:789–865. doi:10.1257/jel.20160995.
- Blau, Francine D., Lawrence M. Kahn, Nikolai Boboshko, and Matthew L. Comey. 2021. "The Impact of Selection into the Labor Force on the Gender Wage Gap." NBER Working Paper Series.
- Borghans, Lex, Bas Ter Weel, and Bruce A. Weinberg. 2014. "People Skills and the Labor-Market Outcomes of Underrepresented Groups: New Evidence on Gender and the Labor Market." *ILR Review* 67:287–334. doi:10.1177/001979391406700202.
- Budig, Michelle J. and Paula England. 2001. "The Wage Penalty for Motherhood." *American Sociological Review* 66:204–225. doi:10.2307/2657415.
- Campbell, Colin and Jessica Pearlman. 2013. "Period Effects, Cohort Effects, and the Narrowing Gender Wage Gap." *Social Science Research* 42:1693–1711. doi:10.1016/j.ssresearch.2013.07.014.
- Cha, Youngjoo and Kim A. Weeden. 2014. "Overwork and the Slow Convergence in the Gender Gap in Wages." *American Sociological Review* 79:457–484. doi:10.1177/0003122414528936.
- Clogg, Clifford C. 1982. "Cohort Analysis of Recent Trends in Labor Force Participation." *Demography* 19:459–479. doi:10.2307/2061013.
- Correll, Shelley J., Stephen Benard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112:1297–1339. doi:10.1086/511799.
- Elbers, Benjamin. 2020. "Orthogonal Polynomial Contrasts and Applications to Age-Period-Cohort Models." doi:10.31235/osf.io/xrbgv.
- England, Paula, Andrew Levine, and Emma Mishel. 2020. "Progress toward Gender Equality in the United States Has Slowed or Stalled." *Proceedings of the National Academy of Sciences* 117:6990–6997. doi:10.1073/pnas.1918891117.
- Esping-Andersen, Gøsta and Francesco C. Billari. 2015. "Re-Theorizing Family Demographics." *Population and Development Review* 41:1–31. doi:10.1111/j.1728-4457.2015.00024.x.
- Fienberg, Stephen E. 2013. "Cohort Analysis' Unholy Quest: A Discussion." *Demography* 50:1981–1984. doi:10.1007/s13524-013-0251-z.

- 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. doi:10.2307/270764.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven J. Ruggles, Robert Warren, and Michael Westberry. 2021. *Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [Dataset]*. Minneapolis, MN: IPUMS. doi:10.18128/D030.V9.0.
- Fosse, Ethan. 2021. "Bayesian Models of Age-Period-Cohort Effects." In *Age, Period, and Cohort Effects: The Identification Problem and Beyond*, edited by Andrew Bell. London: Routledge Press.
- . 2023. "Dissecting the Lexis Table: Summarizing Population-Level Temporal Variability with Age-Period-Cohort Data." *Sociological Science* .
- Fosse, Ethan and Christopher Winship. 2018. "Moore–Penrose Estimators of Age–Period–Cohort Effects: Their Interrelationship and Properties." *Sociological Science* 5:304–334. doi:10.15195/v5.a14.
- . 2019a. "Analyzing Age-Period-Cohort Data: A Review and Critique." *Annual Review of Sociology* 45:467–492. doi:10.1146/annurev-soc-073018-022616.
- . 2019b. "Bounding Analyses of Age-Period-Cohort Effects." *Demography* 56:1975–2004. doi:10.1007/s13524-019-00801-6.
- . 2023. "The Anatomy of Cohort Analysis: Decomposing Ryderian Comparative Cohort Careers." *Sociological Methodology* .
- Fox, John. 2016. *Applied Regression Analysis and Generalized Linear Models*. Los Angeles, CA: SAGE Publications, 3rd edition.
- Fu, Wenjiang J. 2000. "Ridge Estimator in Singular Design with Application to Age-Period-Cohort Analysis of Disease Rates." *Communications in Statistics - Theory and Methods* 29:263–278. doi:10.1080/03610920008832483.
- . 2016. "Constrained Estimators and Consistency of a Regression Model on a Lexis Diagram." *Journal of the American Statistical Association* 111:180–199. doi:10.1080/01621459.2014.998761.
- Goldin, Claudia D. 2021. *Career and Family: Women's Century-Long Journey Toward Equity*. Princeton, NJ: Princeton University Press.
- Grusky, David B. 1994. "The Contours of Social Stratification." In *Social Stratification: Class, Race, and Gender in Sociological Perspective*, edited by David B. Grusky, Social Inequality Series, pp. 3–35. Boulder, CO: Westview Press.
- Hays, Sharon. 1996. *The Cultural Contradictions of Motherhood*. New Haven, CT: Yale University Press.
- Hodges, Melissa J. and Michelle J. Budig. 2010. "Who Gets the Daddy Bonus?: Organizational Hegemonic Masculinity and the Impact of Fatherhood on Earnings." *Gender & Society* 24:717–745. doi:10.1177/0891243210386729.
- Holford, Theodore R. 1983. "The Estimation of Age, Period and Cohort Effects for Vital Rates." *Biometrics* 39:311–324. doi:10.2307/2531004.

- . 2006. “Approaches to Fitting Age-Period-Cohort Models with Unequal Intervals.” *Statistics in Medicine* 25:977–993. doi:10.1002/sim.2253.
- Horowitz, Juliana and Ruth Igielnik. 2020. “A Century After Women Gained the Right To Vote, Majority of Americans See Work To Do on Gender Equality.” Technical report, Pew Research Center.
- Hout, Michael, Stuart Perrett, and Sarah K. Cowan. 2022. “Stasis and Sorting of Americans’ Abortion Opinions: Political Polarization Added to Religious and Other Differences.” *Socius* 8:237802312211176. doi:10.1177/23780231221117648.
- Killewald, Alexandra. 2013. “A Reconsideration of the Fatherhood Premium: Marriage, Coresidence, Biology, and Fathers’ Wages.” *American Sociological Review* 78:96–116. doi:10.1177/0003122412469204.
- Killewald, Alexandra and Nino Cricco. 2020. “Have Changing Family Demographics Narrowed the Gender Wage Gap?” Paper presented at the annual meeting of the Population Association of American, Virtual, April 23.
- Killewald, Alexandra and Javier García-Manglano. 2016. “Tethered Lives: A Couple-Based Perspective on the Consequences of Parenthood for Time Use, Occupation, and Wages.” *Social Science Research* 60:266–282. doi:10.1016/j.ssresearch.2016.03.007.
- Kupper, Lawrence L., Joseph M. Janis, Azza Karmous, and Bernard G. Greenberg. 1985. “Statistical Age-Period-Cohort Analysis: A Review and Critique.” *Journal of Chronic Diseases* 38:811–830. doi:10.1016/0021-9681(85)90105-5.
- Kurtulus, Fidan Ana. 2012. “Affirmative Action and the Occupational Advancement of Minorities and Women During 1973–2003.” *Industrial Relations* 51:213–246. doi:10.1111/j.1468-232X.2012.00675.x.
- Luo, Liying. 2013. “Assessing Validity and Application Scope of the Intrinsic Estimator Approach to the Age-Period-Cohort Problem.” *Demography* 50:1945–1967. doi:10.1007/s13524-013-0243-z.
- Luo, Liying, James Hodges, Christopher Winship, and Daniel Powers. 2016. “The Sensitivity of the Intrinsic Estimator to Coding Schemes: Comment on Yang, Schulhofer-Wohl, Fu, and Land.” *American Journal of Sociology* 122:930–961. doi:10.1086/689830.
- Luo, Liying and James S. Hodges. 2016. “Block Constraints in Age–Period–Cohort Models with Unequal-width Intervals.” *Sociological Methods & Research* 45:700–726. doi:10.1177/0049124115585359.
- . 2020. “Constraints in Random Effects Age-Period-Cohort Models.” *Sociological Methodology* 50:276–317. doi:10.1177/0081175020903348.
- Manski, Charles F. 2007. *Identification for Prediction and Decision*. Cambridge, MA: Harvard University Press.
- 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* 38:242. doi:10.2307/2094398.

- Mason, William M. and Herbert L. Smith. 1985. "Age-Period-Cohort Analysis and the Study of Deaths from Pulmonary Tuberculosis." In *Cohort Analysis in Social Research*, edited by William M. Mason and Stephen E. Fienberg, pp. 151–227. New York, NY: Springer-Verlag.
- Masters, Ryan K., Robert A. Hummer, Daniel A. Powers, Audrey Beck, Shih-Fan Lin, and Brian Karl Finch. 2014. "Long-Term Trends in Adult Mortality for U.S. Blacks and Whites: An Examination of Period- and Cohort-Based Changes." *Demography* 51:2047–2073. doi:10.1007/s13524-014-0343-4.
- Musick, Kelly, Megan D. Bea, and Pilar Gonalons-Pons. 2020. "His and Her Earnings Following Parenthood in the United States, Germany, and the United Kingdom." *American Sociological Review* 85:639–674. doi:10.1177/0003122420934430.
- Nielsen, Bent. 2022. "Two-Sample Age-Period-Cohort Models with an Application to Swiss Suicide Rates." Department of Economics, Nuffield College, University of Oxford, Oxford, UK. Unpublished manuscript.
- O'Brien, Robert M. 2015. *Age-Period-Cohort Models: Approaches and Analyses with Aggregate Data*. Boca Raton, FL: CRC Press.
- . 2017. "Mixed Models, Linear Dependency, and Identification in Age-Period-Cohort Models." *Statistics in Medicine* 36:2590–2600. doi:10.1002/sim.7305.
- O'Neill, June and Solomon Polachek. 1993. "Why the Gender Gap in Wages Narrowed in the 1980s." *Journal of Labor Economics* 11:205–228. doi:10.1086/298323.
- Pampel, Fred C. and Lori M. Hunter. 2012. "Cohort Change, Diffusion, and Support for Environmental Spending in the United States." *American Journal of Sociology* 118:420–448. doi:10.1086/666506.
- Riebler, Andrea and Leonhard Held. 2010. "The Analysis of Heterogeneous Time Trends in Multivariate Age-Period-Cohort Models." *Biostatistics* 11:57–69. doi:10.1093/biostatistics/kxp037.
- Riebler, Andrea, Leonhard Held, Håvard Rue, and Matthias Bopp. 2012. "Gender-Specific Differences and the Impact of Family Integration on Time Trends in Age-Stratified Swiss Suicide Rates." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 175:473–490. doi:10.1111/j.1467-985X.2011.01013.x.
- Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30:843–861.
- Underwood, Ted, Kevin Kiley, Wenyi Shang, and Stephen Vaisey. 2022. "Cohort Succession Explains Most Change in Literary Culture." *Sociological Science* 9:184–205. doi:10.15195/v9.a8.
- Winship, Christopher and David J. Harding. 2008. "A Mechanism-Based Approach to the Identification of Age-Period-Cohort Models." *Sociological Methods & Research* 36:362–401. doi:10.1177/0049124107310635.
- Yang, Yang. 2008. "Social Inequalities in Happiness in the United States, 1972 to 2004: An Age-Period-Cohort Analysis." *American Sociological Review* 73:204–226. doi:10.1177/000312240807300202.

- Yang, Yang and Kenneth C. Land. 2006. "A Mixed Models Approach to the Age-Period-Cohort Analysis of Repeated Cross-Section Surveys, with an Application to Data on Trends in Verbal Test Scores." *Sociological Methodology* 36:75–97. doi:10.1111/j.1467-9531.2006.00175.x.
- . 2013. *Age-Period-Cohort Analysis: New Models, Methods, and Empirical Applications*. Boca Raton, FL: CRC Press.
- Yang, Yang, Sam Schulhofer-Wohl, Wenjiang J. Fu, and Kenneth C. Land. 2008. "The Intrinsic Estimator for Age-Period-Cohort Analysis: What It Is and How to Use It." *American Journal of Sociology* 113:1697–1736. doi:10.1086/587154.
- Yu, Wei-hsin and Janet Chen-Lan Kuo. 2017. "The Motherhood Wage Penalty by Work Conditions: How Do Occupational Characteristics Hinder or Empower Mothers?" *American Sociological Review* 82:744–769. doi:10.1177/0003122417712729.

## Appendix A: Details on the CSL-APC Model

The presentation of the models in the main text are quite general in that we do not specify exactly how the linear and nonlinear components have been constructed. Following previous research (e.g., Holford 1983; Fosse and Winship 2019a), we will use orthogonal polynomials such that, for example,  $a_L$  denotes the linear component,  $a_2$  denotes the quadratic component,  $a_3$  the cubic component, and so forth.<sup>19</sup> This implies the following re-expression of the L-APC model:

$$\begin{aligned} Y_{ijk} &= \mu + \alpha(i - i^*) + \pi(j - j^*) + \gamma(k - k^*) + \tilde{\alpha}_i + \tilde{\pi}_j + \tilde{\gamma}_k \\ &= \mu + \alpha a_L + \pi p_L + \gamma c_L + \sum_{i=1}^{I-1} \alpha_i a_i + \sum_{j=1}^{J-1} \pi_j p_j + \sum_{k=1}^{K-1} \gamma_k c_k + \eta_{ijk}, \end{aligned} \quad (7)$$

where  $a_L$ ,  $p_L$ , and  $c_L$  are the age, period, and cohort linear components with corresponding linear effects  $\alpha$ ,  $\pi$ , and  $\gamma$ ;  $a_2, \dots, a_{I-1}$  are the age nonlinear components with corresponding nonlinear effects  $\alpha_2, \dots, \alpha_{I-1}$ ;  $p_2, \dots, p_{J-1}$  are the period nonlinear components with corresponding nonlinear effects  $\pi_2, \dots, \pi_{J-1}$ ;  $c_2, \dots, c_{K-1}$  are the cohort nonlinear components with corresponding nonlinear effects  $\gamma_2, \dots, \gamma_{K-1}$ ; and  $\eta_{ijk}$  denote the cell-specific error terms.

Likewise, the CSL-APC model is also quite general and there are various ways of parameterizing cross-strata differences in APC effects. While one can allow parameters to vary across any number of levels of a strata variable in principle, we have two levels for gender in our case. Let an indicator (dummy) gender variable  $G$  coded as  $g = 1$  for women and  $g = 0$  for men. Including and interacting this strata variable with the linear and nonlinear components for age, period, and cohort results in the following model, which is analogous to the SL-APC model in Equation 3 in the main text:<sup>20</sup>

$$\begin{aligned} Y_{ijk} &= \mu + \alpha a_L + \pi p_L + \gamma c_L + \sum_{i=1}^{I-1} \alpha_i a_i + \sum_{j=1}^{J-1} \pi_j p_j + \sum_{k=1}^{K-1} \gamma_k c_k + \mu_G G + \\ &\quad \alpha_G(a_L G) + \pi_G(p_L G) + \gamma_G(c_L G) + \sum_{i=1}^{I-1} \alpha_{Gi}(a_i G) + \sum_{j=1}^{J-1} \pi_{Gj}(p_j G) + \sum_{k=1}^{K-1} \gamma_{Gk}(c_k G) + \eta_{ijk}, \end{aligned} \quad (8)$$

where  $\alpha_G$ ,  $\pi_G$ , and  $\gamma_G$  are interaction effects between the gender variable and the age, period, and cohort linear components, respectively;  $\alpha_{Gi}$ ,  $\pi_{Gj}$ , and  $\gamma_{Gk}$  are interaction effects between the gender variable and the nonlinear components for age, period, and cohort, respectively.

There are two closely related ways of interpreting the interaction terms in Equation 8. On the one hand, the interaction effects can be interpreted as representing differences in age, period, and cohort effects between the strata. For example,  $\alpha_G$  can be interpreted as the difference between the age linear effect for women and the age linear effect for men. On the other hand, the interaction effects can be interpreted as representing the cross-strata outcome disparity for varying values of age, period, and cohort. For example, the parameters for  $\alpha_G$  can be interpreted as the gender “effect” on the outcome (i.e., between-gender wage disparity) for varying levels of age. In general, we focus on interpreting the parameters as cross-strata differences in age, period, and cohort effects,

<sup>19</sup>Note that the identity in the below equation may be only approximate depending on how the design matrix is constructed. In our analyses we use QR decomposition to construct the nonlinear components. Because the elements of the design matrix can still be quite large, for the purposes of numerical stability we include the additional step of norming each of the columns of the design matrix representing the nonlinear components. We use a weighted version of orthogonal polynomials to make sure polynomial terms are perpendicular to lower-order terms in our empirical data (Elbers 2020).

<sup>20</sup>If we had multiple levels of strata, then we would specify an expanded set of interaction terms on the right-hand side of the equation with indicator variables  $X_1$ ,  $X_2$ , and so on.

but researchers may opt for the second interpretation depending on the substantive application.<sup>21</sup>

We are now ready to show how the CSL-APC model is derived from a variant of the model in Equation 8, or the SL-APC model in Equation 3 in the main text. After substituting for  $G = 1$  and  $G = 0$ , we can express the CSL-APC model as follows:

$$\begin{aligned}
Y_{ijk[G=1]} - Y_{ijk[G=0]} &= \\
&\left( \mu_G \times 1 + \alpha_G(a_L \times 1) + \pi_G(p_L \times 1) + \gamma_G(c_L \times 1) + \sum_{i=1}^{I-1} \alpha_{Gi}(a_i \times 1) + \sum_{j=1}^{J-1} \pi_{Gj}(p_j \times 1) + \sum_{k=1}^{K-1} \gamma_{Gk}(c_k \times 1) + \epsilon_{ijk[G=1]} \right) \\
&- \left( \mu_G \times 0 + \alpha_G(a_L \times 0) + \pi_G(p_L \times 0) + \gamma_G(c_L \times 0) + \sum_{i=1}^{I-1} \alpha_{Gi}(a_i \times 0) + \sum_{j=1}^{J-1} \pi_{Gj}(p_j \times 0) + \sum_{k=1}^{K-1} \gamma_{Gk}(c_k \times 0) + \epsilon_{ijk[G=0]} \right) \\
&= \mu_G + \alpha_G a_L + \pi_G p_L + \gamma_G c_L + \sum_{i=1}^{I-1} \alpha_{Gi} a_i + \sum_{j=1}^{J-1} \pi_{Gj} p_j + \sum_{k=1}^{K-1} \gamma_{Gk} c_k + (\eta_{ijk[G=1]} - \eta_{ijk[G=0]}),
\end{aligned}$$

or, in a more compact general form:

$$\Delta Y_{ijk} = \Delta\mu + \Delta\alpha(i - i^*) + \Delta\pi(j - j^*) + \Delta\gamma(k - k^*) + \Delta\tilde{\alpha}_i + \Delta\tilde{\pi}_j + \Delta\tilde{\gamma}_k + \Delta\eta_{ijk}, \quad (9)$$

which is equivalent to Equation 4 in the main text.<sup>22</sup>

---

<sup>21</sup>Our case is a specific example of the general conceptual issue that appears when interpreting interaction effects. Suppose there is an interactive effect of a continuous variable  $X$  and a binary group indicator  $G$  on an outcome  $Y$ , namely,  $Y = \mu + \beta_1 X + \beta_2 G + \beta_3 (X \times G) + \epsilon$ . One interpretation of  $\beta_3$ , which aligns with the first interpretation above, is that the effect of  $X$  is different between the two groups indicated by  $G$ . Another interpretation, more consistent with our second interpretation, is that the between-group difference in the outcome indicated by  $G$  varies depending on the level of  $X$  (see Fox 2016:140-150)

<sup>22</sup>The estimated nonlinear effects for each age, period, cohort category as presented in Figure 6 are predicted values based on the parameter estimates for an intercept and orthogonal polynomials.



## Appendix B: Lexis Table

Table 1: Lexis Table of the Gender Wage Gap

Age	Period								
	1975-79	1975-79	1975-79	1975-79	1975-79	1975-79	1975-79	1975-79	1975-79
25-29	-0.35 (0.004)	-0.30 (0.009)	-0.23 (0.006)	-0.15 (0.008)	-0.14 (0.009)	-0.09 (0.009)	-0.10 (0.009)	-0.09 (0.006)	-0.12 (0.011)
30-34	-0.49 (0.005)	-0.39 (0.007)	-0.31 (0.008)	-0.26 (0.009)	-0.24 (0.005)	-0.20 (0.005)	-0.16 (0.006)	-0.15 (0.010)	-0.20 (0.007)
35-39	-0.59 (0.009)	-0.50 (0.009)	-0.44 (0.008)	-0.34 (0.010)	-0.34 (0.009)	-0.29 (0.005)	-0.27 (0.011)	-0.22 (0.007)	-0.19 (0.007)
40-44	-0.61 (0.013)	-0.58 (0.006)	-0.49 (0.010)	-0.39 (0.002)	-0.36 (0.011)	-0.32 (0.009)	-0.29 (0.007)	-0.26 (0.007)	-0.25 (0.010)
45-49	-0.63 (0.013)	-0.60 (0.010)	-0.53 (0.009)	-0.45 (0.009)	-0.37 (0.013)	-0.35 (0.010)	-0.33 (0.010)	-0.32 (0.010)	-0.28 (0.010)
50-54	-0.60 (0.011)	-0.58 (0.014)	-0.53 (0.015)	-0.48 (0.012)	-0.44 (0.011)	-0.36 (0.012)	-0.28 (0.013)	-0.33 (0.010)	-0.27 (0.009)
55-59	-0.57 (0.015)	-0.61 (0.013)	-0.56 (0.012)	-0.49 (0.016)	-0.44 (0.012)	-0.40 (0.010)	-0.31 (0.011)	-0.30 (0.014)	-0.27 (0.012)
60-64	-0.52 (0.009)	-0.51 (0.023)	-0.52 (0.012)	-0.44 (0.023)	-0.42 (0.019)	-0.36 (0.019)	-0.39 (0.012)	-0.30 (0.014)	-0.26 (0.012)

Notes: The rows indicate age categories, and the columns indicate period categories. The input in each cell denotes the gender difference (=female–male) in log median annual earnings in the respective age-period category. Standard errors for each difference are presented in parentheses.

In presenting our CSL-APC model, we assumed that we have only aggregated data (see note 8). But researchers may have individual-level sample data to construct a Lexis table of cross-strata differences, as we do in our empirical example using the CPS data. In these cases, the Lexis table needs to be estimated from the individual-level sample data. To do so, we conducted a regression of log earnings on age, period, sex, and all the two-way and three-way interactions between these variables. Since we are modeling a *median* difference between men’s and women’s earnings conditional on age and period, we relied on a (conditional) quantile regression at the median. The CPS ASEC survey sampling weights are applied to the regression so that the estimated median differences represent the population characteristic. Predicted marginal “effects” of gender, which are allowed to vary by age and period, can be computed from the fitted median regression and are used as the cell values in the above Lexis table. If researchers are interested in modeling a cross-strata difference in *mean* values, they can instead conduct an OLS regression to estimate the corresponding Lexis table.

A potential advantage of using individual-level data as opposed to aggregated data is that researchers can have a better knowledge of the precision with which each cell value is estimated. For example, the gender-specific variance of log earnings, which is likely not available in aggregated data, affects the precision of the estimated cell values. Taking differential precision into account can yield an efficiency gain as compared to OLS where cell variance is assumed to be constant, especially in small samples like the Lexis table ( $N=72$ ). We therefore incorporate the estimated standard errors of cell values (presented in parentheses in Table 1 above) when fitting a weighted least squares regression of the CSL-APC model. The weight of each cell is computed as  $1/se^2$ . With that said, the point estimates are very similar to OLS estimates of the CSL-APC model, and the substantive conclusions from the bounding APC analysis remain the same.

## Appendix C: Additional Tables

Table 1: Descriptive Statistics of the CPS Sample

Variable	Men				Women			
	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
Age	41.7	10.3	25	64	41.88	10.4	25	64
Period	1999.3	12.5	1976	2019	2000.8	12.0	1976	2019
Cohort	1957.7	15.5	1912	1994	1958.9	15.0	1912	1994
Earnings	69,437.2	60,473.6	1.6	1,958,398.9	48,314.5	40,898.2	1.3	1,401,398.9
Obsv.	1,121,562				830,856			

*Notes:* Mean values and their standard deviations are computed for the weighted sample using the CPS ASEC survey sampling weights. Earnings refer to respondents' annual earnings in the last calendar year.

Table 2: Bounding Formulas for Cross-Strata Differences in APC Slopes

Age Bounds:	$\alpha_{\min} \leq \Delta\alpha \leq \alpha_{\max}$ $\theta_1 - \alpha_{\max} \leq \Delta\pi \leq \theta_1 - \alpha_{\min}$ $(\theta_2 - \theta_1) + \alpha_{\min} \leq \Delta\gamma \leq (\theta_2 - \theta_1) + \alpha_{\max}$
Period Bounds:	$\theta_1 - \pi_{\max} \leq \Delta\alpha \leq \theta_1 - \pi_{\min}$ $\pi_{\min} \leq \Delta\pi \leq \pi_{\max}$ $\theta_2 - \pi_{\max} \leq \Delta\gamma \leq \theta_2 - \pi_{\min}$
Cohort Bounds:	$(\theta_1 - \theta_2) + \gamma_{\min} \leq \Delta\alpha \leq (\theta_1 - \theta_2) + \gamma_{\max}$ $\theta_2 - \gamma_{\max} \leq \Delta\pi \leq \theta_2 - \gamma_{\min}$ $\gamma_{\min} \leq \Delta\gamma \leq \gamma_{\max}$

*Notes:* Age, period, and cohort slopes are  $\alpha$ ,  $\pi$ , and  $\gamma$ , respectively, with  $(\cdot)_{\min}$  and  $(\cdot)_{\max}$  denoting minimum and maximum values of the bounds. We denote  $\theta_1 = \Delta\alpha + \Delta\pi$ ,  $\theta_2 = \Delta\gamma + \Delta\pi$ ,  $\theta_1 - \theta_2 = \Delta\alpha - \Delta\gamma$ , and  $\theta_2 - \theta_1 = \Delta\gamma - \Delta\alpha$ .

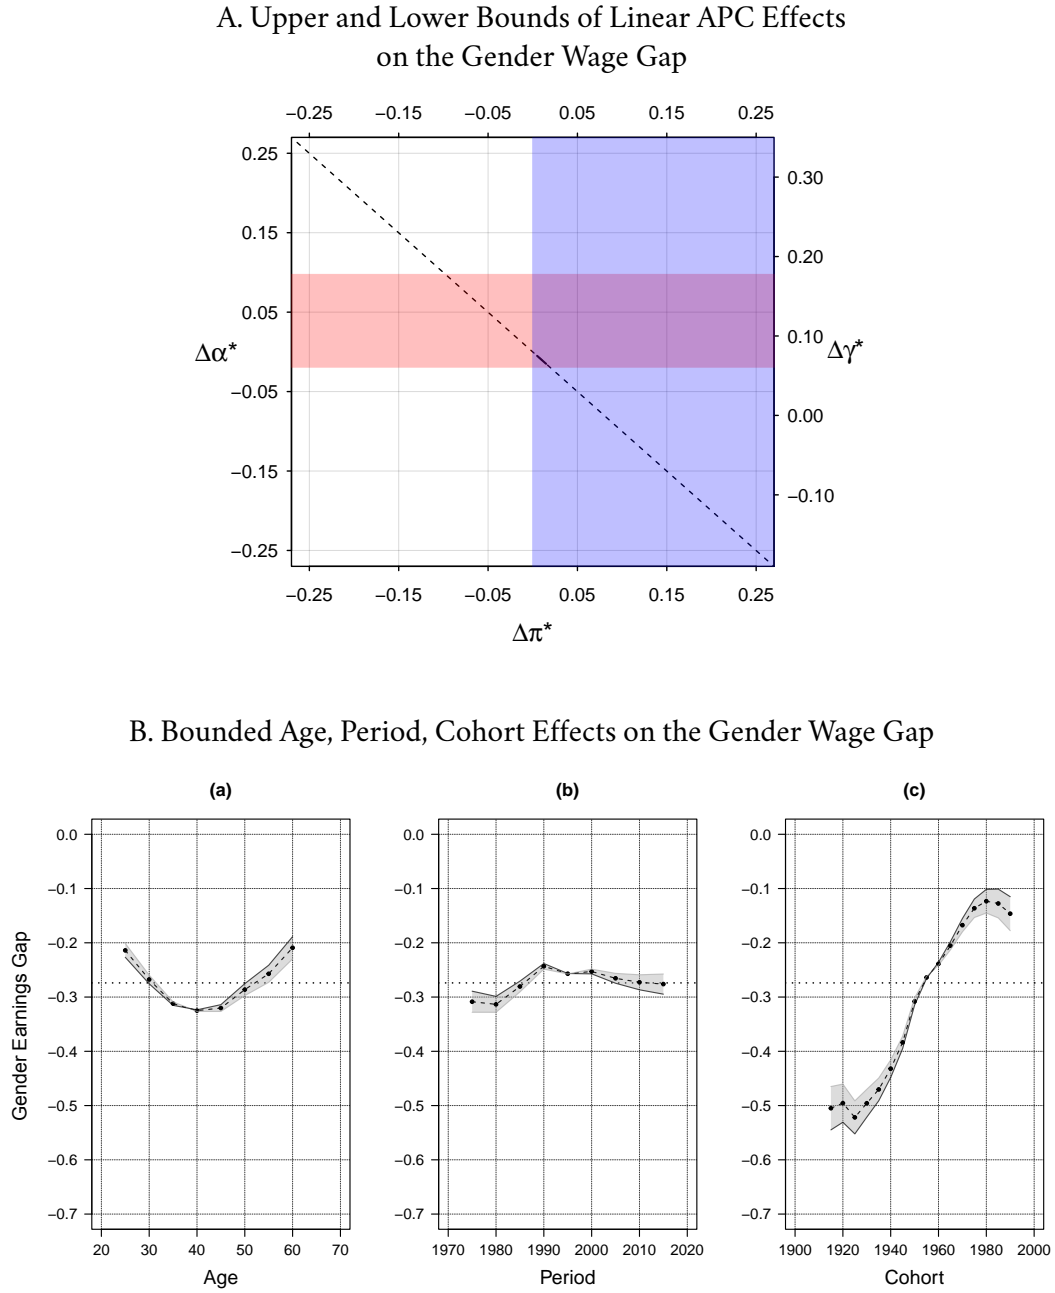
Table 3: Estimated Thetas and Nonlinear Age, Period, and Cohort Effects on the Gender Wage Gap

Parameter	Coefficient	Std. Error	95% CI:	
			Lower Bound	Upper Bound
$\Delta\mu$	-0.339	0.002	-0.342	-0.335
$\Delta\theta_1$	-0.002	0.002	-0.007	0.002
$\Delta\theta_2$	0.084	0.001	0.081	0.086
$\Delta\tilde{\alpha}_2$	0.152	0.007	0.138	0.166
$\Delta\tilde{\alpha}_3$	-0.026	0.006	-0.039	-0.013
$\Delta\tilde{\alpha}_4$	-0.006	0.006	-0.019	0.007
$\Delta\tilde{\alpha}_5$	0.011	0.006	-0.001	0.023
$\Delta\tilde{\alpha}_6$	-0.003	0.006	-0.015	0.009
$\Delta\tilde{\alpha}_7$	0.009	0.006	-0.003	0.021
$\Delta\tilde{\pi}_2$	-0.063	0.006	-0.075	-0.052
$\Delta\tilde{\pi}_3$	0.009	0.006	-0.003	0.020
$\Delta\tilde{\pi}_4$	0.014	0.006	0.002	0.025
$\Delta\tilde{\pi}_5$	-0.014	0.006	-0.025	-0.002
$\Delta\tilde{\pi}_6$	0.003	0.006	-0.008	0.014
$\Delta\tilde{\pi}_7$	0.008	0.006	-0.004	0.019
$\Delta\tilde{\pi}_8$	-0.011	0.006	-0.023	0.000
$\Delta\tilde{\gamma}_2$	-0.038	0.015	-0.069	-0.007
$\Delta\tilde{\gamma}_3$	-0.150	0.014	-0.178	-0.123
$\Delta\tilde{\gamma}_4$	0.046	0.014	0.018	0.074
$\Delta\tilde{\gamma}_5$	-0.031	0.014	-0.059	-0.004
$\Delta\tilde{\gamma}_6$	-0.046	0.013	-0.073	-0.018
$\Delta\tilde{\gamma}_7$	0.031	0.013	0.005	0.058
$\Delta\tilde{\gamma}_8$	0.049	0.014	0.022	0.077
$\Delta\tilde{\gamma}_9$	0.012	0.013	-0.014	0.038
$\Delta\tilde{\gamma}_{10}$	-0.020	0.012	-0.044	0.003
$\Delta\tilde{\gamma}_{11}$	0.013	0.010	-0.007	0.032
$\Delta\tilde{\gamma}_{12}$	0.012	0.009	-0.006	0.030
$\Delta\tilde{\gamma}_{13}$	-0.006	0.008	-0.023	0.010
$\Delta\tilde{\gamma}_{14}$	-0.009	0.008	-0.024	0.007
$\Delta\tilde{\gamma}_{15}$	-0.003	0.006	-0.016	0.010
Adjusted R-squared		0.98		
Number of Cells		72		

*Notes:* Nonlinearity parameters are estimated for varying degrees of orthogonal polynomials, respectively, as denoted by their subscripts (see Appendix A and Appendix B for more details of the estimation process). The estimated nonlinear effects presented in Figure 6 are predicted values based on the above nonlinearity and intercept estimates.

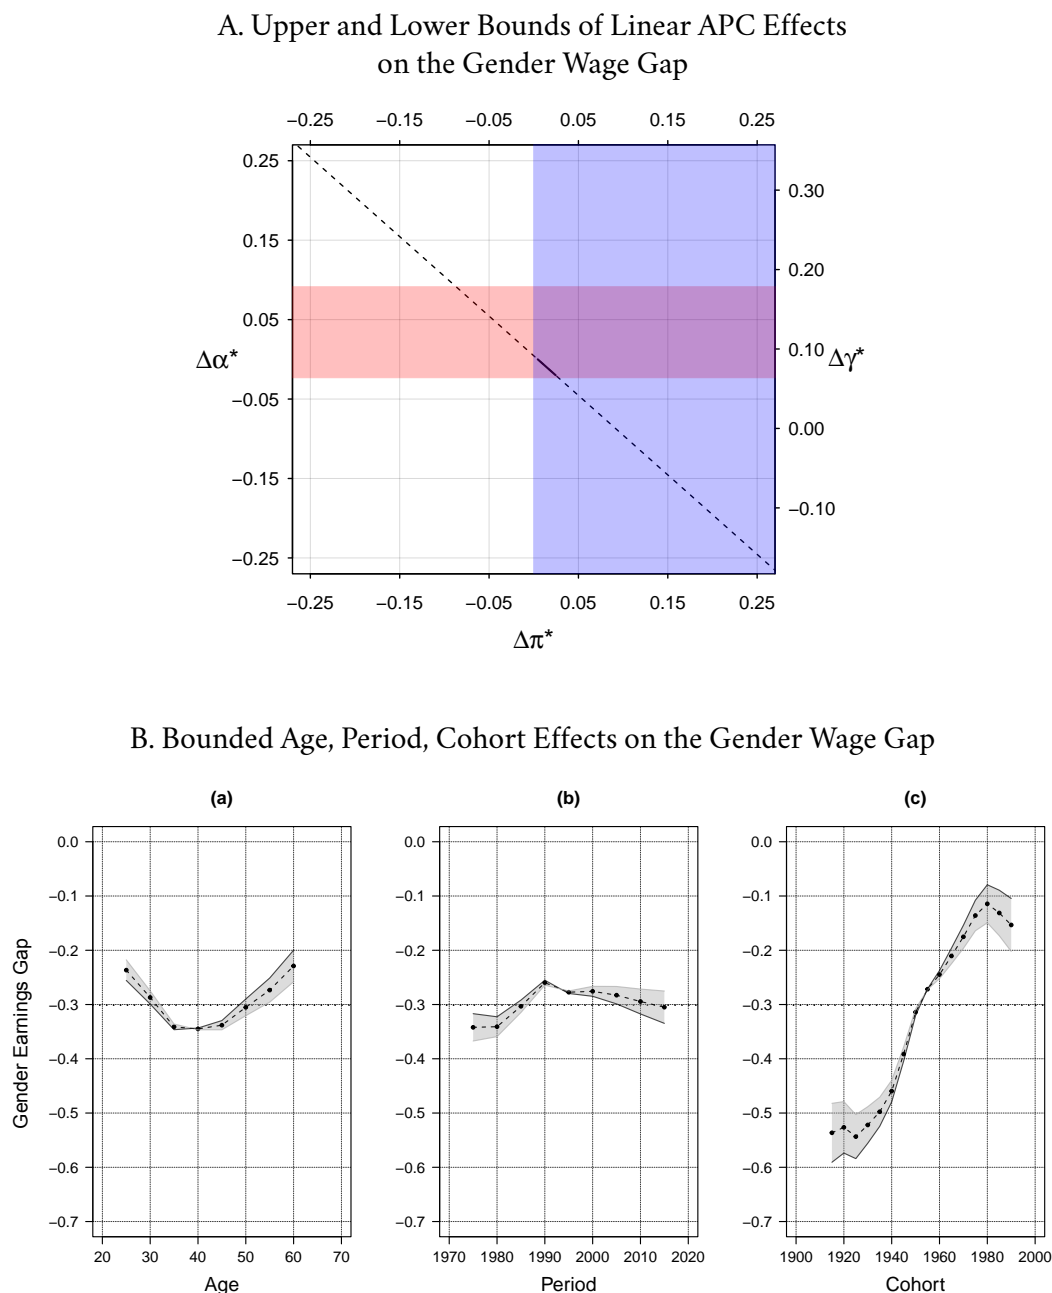
## Appendix D: Additional Figures

Figure 1. Bounding Analysis Results of Gender Differences in Median Hourly Wages



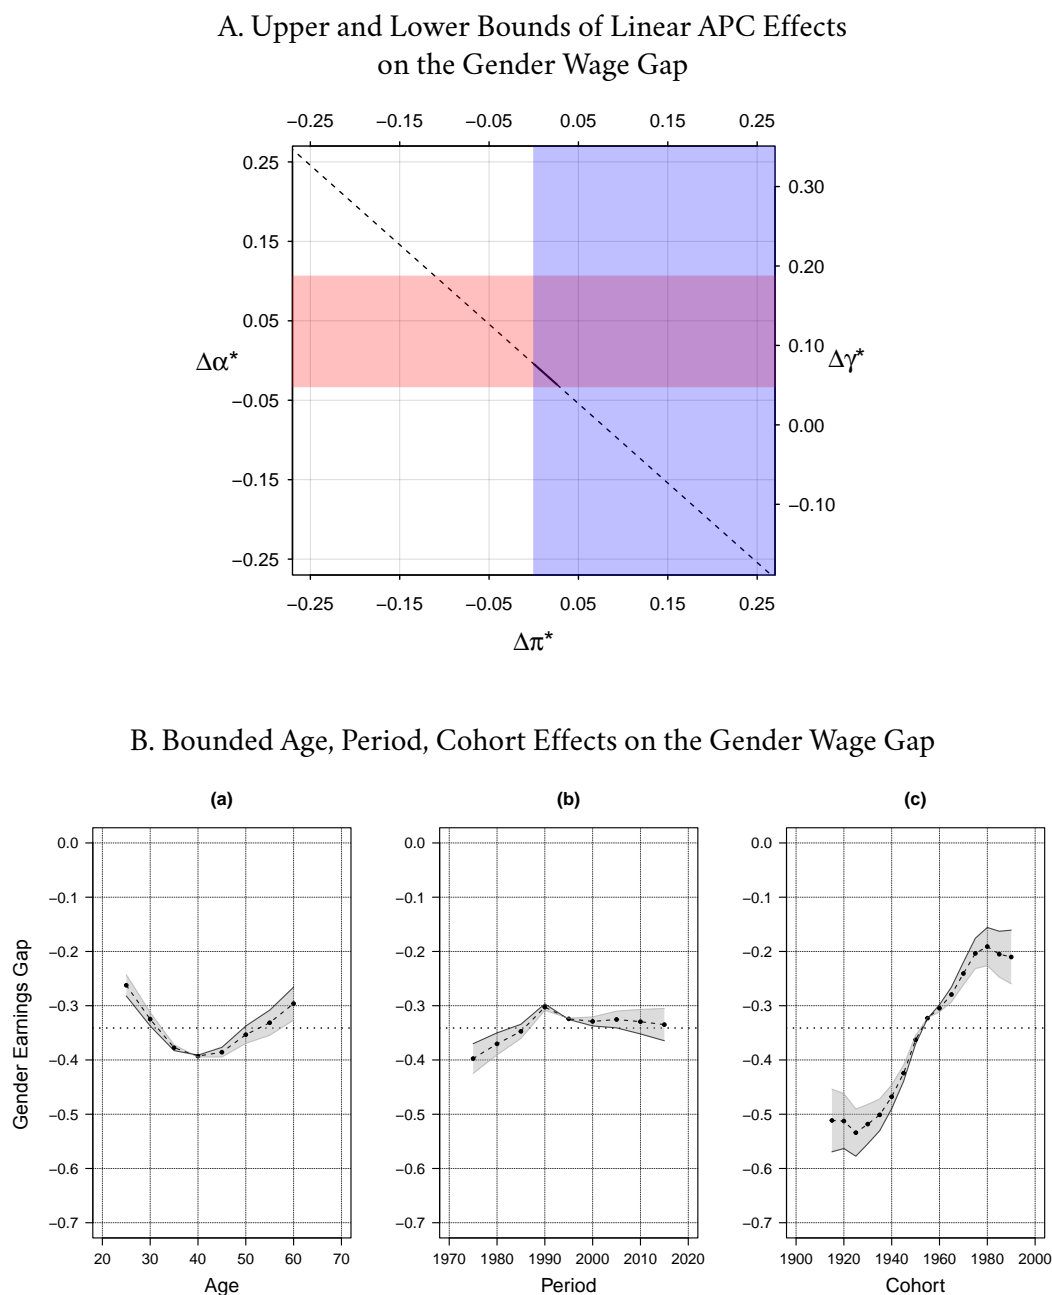
*Notes:* Hourly wages instead of annual earnings are analyzed. In the upper panel (A), the left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. The solid line in the overlapping colored regions (red indicates age-related assumptions; blue indicates a period-related assumption) refers to the possible area of the linear effects if the three assumptions about age and period effects were satisfied. In the bottom panel (B), the shaded areas represent the bounded effects of age (a), period (b), and cohort (c) on the gender wage gap based on the three assumptions about age and period effects. The dotted lines follow the mid-points in each shaded area. The dark bold lines along one end of the shaded areas depict one possibility of the APC effects where the linear age effect is most positive, the linear period effect is most negative, and the linear cohort effect is most positive within the bounded ranges.

Figure 2. Bounding Analysis Results of Gender Differences in Median Hourly Wages, Including Half-Year-Round Workers



*Notes:* Hourly wages instead of annual earnings are analyzed, and full-time wage/salary workers who worked at least 26 weeks in the last calendar year are included. In the upper panel (A), the left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. The solid line in the overlapping colored regions (red indicates age-related assumptions; blue indicates a period-related assumption) refers to the possible area of the linear effects if the three assumptions about age and period effects were satisfied. In the bottom panel (B), the shaded areas represent the bounded effects of age (a), period (b), and cohort (c) on the gender wage gap based on the three assumptions about age and period effects. The dotted lines follow the mid-points in each shaded area. The dark bold lines along one end of the shaded areas depict one possibility of the APC effects where the linear age effect is most positive, the linear period effect is most negative, and the linear cohort effect is most positive within the bounded ranges.

Figure 3. Bounding Analysis Results of Gender Differences in Mean Annual Earnings



*Notes:* Differences in log mean annual earnings between men and women are analyzed instead of median differences. In the upper panel (A), the left y-axis represents a linear age effect, the x-axis indicates a linear period effect, and the right y-axis represents a linear cohort effect. The dashed line indicates all possible linear effects consistent with data. The solid line in the overlapping colored regions (red indicates age-related assumptions; blue indicates a period-related assumption) refers to the possible area of the linear effects if the three assumptions about age and period effects were satisfied. In the bottom panel (B), the shaded areas represent the bounded effects of age (a), period (b), and cohort (c) on the gender wage gap based on the three assumptions about age and period effects. The dotted lines follow the mid-points in each shaded area. The dark bold lines along one end of the shaded areas depict one possibility of the APC effects where the linear age effect is most positive, the linear period effect is most negative, and the linear cohort effect is most positive within the bounded ranges.