

Age-Period-Cohort Analysis: Critiques and Innovations

A Dissertation
SUBMITTED TO THE FACULTY OF
UNIVERSITY OF MINNESOTA
BY

Liyang Luo

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

John Robert Warren, Advisor

July 2015

© Liying Luo 2015
All rights reserved

ACKNOWLEDGEMENTS

At the time of the preliminary examination, which concerned a completely different topic with a very different committee, I had no clues about what the dissertation is now. Every piece of the dissertation and every little progress that I have made would not be possible without the help from many people. Exhausting the list of the wonderful men and women that I am deeply grateful to would be a daunting task that I am not capable to accomplish, but I will at least acknowledge some of them for their support over the years.

I could not image Rob being an advisor and mentor that I respect (and like) so much when he ignored me in an elevator of the Social Sciences tower during my first year here. Amazingly, he turned out to be *the* most important person shaping my life in the United States. He taught me everything that I need to know about doing good science, surviving the publication process, handling stressful occasions, and balancing work and family. He serves as a template for what I hope to become as a scholar and mentor. Words, either in English or in Chinese, cannot express how grateful I am for everything that he has done to help and support me, but perhaps I can at least do one thing: to be even 51% the mentor that Rob was to me when it comes to mentoring.

The most valuable thing that I got from the graduate training in biostatistics is not the degree, but that I developed good working relationship and friendship with Jim (and his wife, Chi-ping). Jim is a fine scientist and explains things extremely well. He has read drafts of every paper that I first-authored and listened patiently my ideas and complaints. Jim, thank you very much.

I am equally grateful to other committee members of my committee. Chris Winship, Ann Meier, and Ross Macmillan provided very helpful comments on the dissertation and my other research. Chris deserves special thanks for his advice and guidance about how to respectfully disagree. It is a true privilege working with each of them.

Other people who have left unforgettable marks on my work include Bob O'Brien (a remarkably generous and patient friend), Adrian Raftery, and Steven Fienberg. I am lucky to have a group of wonderful interdisciplinary scholars including Carolyn Liebler, Caren Arbeit, Julia Drew, Cathy Fitch, and Susan Mason providing feedback and comments on nearly every paper I have written.

I want to thank my friends Wenjie Liao (Our friendship must be destiny.), Erika Busse, Vicente, Mickey, Meg Krausch, Cyrus, Yu-ju Chien, and Wen Fan, who deeply understand my joys and pains. You are the windows through which I see a whole new world. I would also like to thank Andy Halpern-Manners and Elaine Hernandez for raising the bar for doing good research and being a beautiful couple.

Lastly but (of course) not least, I thank my family, Xiao (an engineer and neuroscientist who has little prejudice against social science research), Yinuo (my beautiful 19-month old son who grows me up), my dad, my mom, and parents-in law (I have been abusing them for providing unpaid child care). I love you.

*To my father,
whom I love and struggle with so much*

ABSTRACT

Researchers have frequently attempted to decompose temporal trends in social, demographic, economic, and health outcomes into three aspects of time processes: age, period, and cohort. The analytical problem that has faced analysts for decades is that these three distinct processes are linearly related to each other ($\text{cohort} = \text{period} - \text{age}$), so disaggregation of temporal trends has to rely on statistical assumptions that are difficult to verify. In this dissertation, I critically evaluate the validity and application scope of two commonly used age-period-cohort (APC) methods: the Intrinsic Estimator and the Cross-Classified Fixed/Random Effects Model. I identify the methodological and theoretical limitations of these methods and conclude that these methods should not be used for estimating the underlying age, period, and cohort patterns without explicit theoretical justification.

What should researchers do? Drawing on the literature of demography, sociology, and statistics, I develop a new method, called the age-period-cohort-interaction (APC-I) model, for analyzing age, period, and cohort variations. Unlike other APC methods, the APC-I model is fully identified and does not rely on problematic statistical assumptions. It also relaxes this assumption in other methods about within-cohort dynamics.

I use the new APC-I model to analyze the 1962 to 2014 data from the Current Population Survey March Supplement to investigate age and period patterns and deviations between cohorts and dynamics within a cohort's life course in labor force participation (LFP) for white and black men and women. I found that while men's LFP was sensitive to social and economic events such as economic recessions and wars, the

effects of these events may not carry on to their later ages. However, there are substantial variations in women's LFP associated with cohort membership that cannot be explained by pure age and period main effects. I also found that while white women's LFP rates caught up with and exceed those of black women by 1980, after adjusting for educational attainment, the racial differences in participation rates among women were substantially reduced after the late 1980s. In addition, the results suggested that a great deal of the period trend and cohort deviations in black women's LFP can be explained by changes in their educational attainment. This is less true for white women; the cohort deviations in participation rates remained after adjusting for education. Surprisingly, there was little evidence supporting an association between changes in marital status and the temporal trends in LFP; the shape of the age, period, and cohort patterns in LFP did not seem to change qualitatively after controlling for current marital status. This finding suggests that the temporal variation in LFP may stem from changes in the behaviors of subgroups of the population other than changes in the marriage composition of the population.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS.....	i
ABSTRACT.....	iv
TABLE OF CONTENTS.....	vi
LIST OF TABLES.....	viii
LIST OF FIGURES	x
Chapter 1 : INTRODUCTION.....	1
Chapter 2 : CRITIQUE OF THE INTRINSIC ESTIMATOR.....	4
The Identification Problem.....	7
The Constrained Approach: IE and CGLM	11
The Linear Constraint Implied by IE	15
The Implications of IE's Constraint: Is IE an unbiased and consistent estimator?.....	25
Biasedness	25
Consistency.....	28
Application Scope: IE vs. CGLM	31
Conclusion and Discussion	36
Appendix	39
Chapter 3 : CRITIQUE OF CROSS-CLASSIFIED MODELS.....	40
The Identification Problem and Cross-Classified APC Models.....	42
The Multiple Block Constraints (MBCs) that CCFEM and CCREM Implicitly Assumes	45
Simulations: Effect of Multiple Block Constraints (MBCs) and Application Scope of CCFEM and CCREM.....	53

Conclusion and Discussion	62
Chapter 4 : A NEW APC MODEL.....	66
Cohort Theories and the APC Accounting Model: Disparities between conceptualization and operationalization	67
Cohort Analysis and Age-Period-Cohort Framework	67
Methodological Critique: What the APC accounting model estimates	69
Theoretical Critique: How cohort effects are defined	71
Intra-Cohort Dynamics: Constant, cumulative, or compensatory?	73
Towards Paradigm Shift: A new model	75
Model Specification.....	76
Estimation and Testing	81
Discussion and Conclusion	87
Appendix	89
Chapter 5 : APPLICATION OF THE APC-I MODEL TO TEMPORAL TRENDS IN LABOR FORCE PARTICIPATION	92
Background	94
Data	97
Results	98
Conclusion and Discussion	119
Chapter 6 : CONCLUSION AND FUTURE RESEARCH	123
REFERENCES	128

LIST OF TABLES

Table 2.1. Different Values of s and the Corresponding Parameters.	10
Table 2.2 Tabular Data: Unobserved parameters and fitted values from Table 2.1's three different parameter vectors.	10
Table 2.3. Simulation Results: IE estimates for three datasets.	21
Table 2.4. Simulation Results: Inconsistent IE estimates as the number of periods increases.	29
Table 3.1. Design Matrix of CCFEM for Equal Interval-Width Age, Period, and Cohort Groups.	48
Table 3.2. Design Matrix of CCFEM for Equal Interval-Width Age, Period, and Cohort Groups by Constraining Two Adjacent Cohorts to Be Equal.	49
Table 3.3. Design Matrix of CCREM for Equal Interval-Width Age, Period, and Cohort Groups.	50
Table 3.4. Design Matrix of CCREM for Equal Interval-Width Age, Period, and Cohort Groups by Constraining Two Adjacent Cohorts to Be Equal.	51
Table 3.5. Examples of True Age, Period, and Cohort Effects.	54
Table 3.6. Simulation Results: CCFEM and CCREM estimates for Table 3.5's three selections (sets of true age, period, and cohort effects).	56
Table 4.1. Unobserved Parameters in Models (1) and (6).	80
Table 4.2. Testing Theories about Intra-Cohort Dynamics Using the Three-Step Procedure.	85
Table 5.1. Descriptive Statistics for All Analytic Variables, Current Population Survey March Supplement, 1962-2014.	99
Table 5.2. Estimated Age and Period Main Effects on Labor Force Participation, with and without Adjustment for Education and Marital Status, in March CPS, 1962-2014.	101

Table 5.3. Estimates of Age-by-Period Interaction in Model 1a for Black Women’s Labor Force Participation in March CPS, 1962-2014	106
Table 5.4. Estimates of Age-by-Period Interaction, Deviation Magnitude F Test, Consistency t Test, and Intra-Cohort Slope t Test Results in Model 1a for Black Women’s Labor Force Participation in March CPS, 1962-2014.....	107
Table 5.5. Estimated Average Deviation between Cohorts and Intra-Cohort Linear Slope in Labor Force Participation, with and without Adjustment for Education, in March CPS, 1962-2014.	110

LIST OF FIGURES

Figure 2.1. Simulation Results: IE vs. CGLM estimates for age, period, and cohort effects.	35
Figure 2.2. Simulation Results: Inconsistent IE estimates.....	39
Figure 3.1. Simulation Results: CCFEM estimates using two- and five-year cohort groups for data from four scenarios.....	60
Figure 3.2. Simulation Results: CCREM estimates using two- and five-year cohort groups for data from four scenarios.....	61
Figure 5.1. Age and Period Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014.....	100
Figure 5.2. Age and Period Main Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014	117
Figure 5.3. Age and Period Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014.....	118

CHAPTER 1 : INTRODUCTION

Social scientists are often concerned with how individual attitudes, behaviors, and health outcomes vary across time. For example, have death rates in the U.S. declined across birth cohorts? Do Americans become politically more liberal or conservative as they get older and has this changed in recent years? Have American men's and women's labor force participation rates decreased or increased over the past decades? Answering questions like these requires analysts to consider simultaneously the roles of three distinct dimensions of time: *age* (how old people are at the time of interview), *period* (the year in which they are interviewed), and *cohort* (in these examples, the year in which they were born). Moreover, in a society in which individual biographies are shaped by social characteristics such as race, gender, and socioeconomic status, individuals who differ in these respects are likely to have divergent age, period, and cohort patterns in social, demographic, and economic outcomes. Therefore, investigating age, period, and cohort patterns can provide new insights about how aging, social changes, and population processes interact with social institutions such as schools and families to produce inequality.

To separate the independent effects of age, period, and cohort, Mason et al. (1973) proposed an age-period-cohort (APC) accounting model. Unfortunately, this APC accounting model suffers from an identification problem: the value for one of the three variables is completely determined by the other two: $\text{cohort} = \text{period} - \text{age}$. That is, researchers have sought to understand three dimensions of time, yet one dimension is an exact function of the other two dimensions. As a result, valid estimates for the age,

period, and cohort effects in the APC accounting model are not possible without additional constraints.

I have conducted four original studies to address methodological, theoretical, and substantive issues in APC research. The first two studies, consisting of the first two chapters of the dissertation, identifies and discusses the methodological limitations of statistical approaches that are commonly used in APC analysis. Many methods have been developed to circumvent the APC identification problem, some of which require rigorous theoretical thinking (see, e.g., Winship and Harding 2008) whereas others are purely technical solutions. Two technical solutions, the Intrinsic Estimator (IE) and the Hierarchical APC (HAPC) models, have gained much popularity because it is believed that they solve the identification problem without external information. However, I show in Chapters 2 and 3 that the IE and the HAPC models both rely on statistical assumptions that are difficult or impossible to verify, and researchers should not use these methods unless theoretical or external information is available for assessing the validity of these assumptions.

What should we do if we want to investigate age, period, and cohort patterns? In the third study, I propose an alternative approach, called the APC-I model, that does not rely on problematic statistical assumptions and is tied more closely to concepts about what cohort represents. Moreover, the new method allows researchers to examine within-cohort dynamics, an important type of cohort-related variation that has been ignored in research using the APC accounting model. I describe the theoretical foundation and technical detail of the APC-I model in Chapter 4. The fourth study reported in Chapter 5

presents an empirical application of the method to investigate age, period, and cohort changes in American men's and women's labor force participation rates using the 1962 to 2014 Current Population Survey data.

CHAPTER 2 : CRITIQUE OF THE INTRINSIC ESTIMATOR

For over a century, social scientists have attempted to separate cohort effects from age and period effects on various social phenomena including mortality, disease rates, and inequality (e.g., Mason et al. 1973; Holford 1983; Fu 2000; O'Brien 2000; Winship and Harding 2008). Whereas *age effects* represent the variation associated with growing older, *period effects* refer to effects due to social and historical shifts such as economic recessions and prevalent unemployment that affect all age groups simultaneously. Cohort refers to a group of people who experience an event such as birth at the same age. *Cohort effects* are defined as the formative effects of social events on individuals at a specific period during their life course (Ryder 1965). Age-period-cohort (APC) models, where the three variables are simultaneously considered in a statistical equation, have been the conventional framework for quantifying age, period, and cohort effects. Unfortunately, such APC models suffer from a logical *identification problem*: once any two of the three variables (age, period, and cohort) are known, the value of the third is determined; this is because $\text{Cohort} = \text{Period} - \text{Age}$. Because of this exact linear dependency, there exist no valid estimates of the distinct effects of the three variables.

Various methods have been developed to address this identification problem. For example, Mason et al. (1973) introduced the APC multiple classification model and suggested the Constrained Generalized Linear Model (CGLM) as a means of estimating the independent effects of age, period, and cohort. More recently, Fu (2000) and Yang and colleagues (2004) proposed a new APC method, called the *Intrinsic Estimator* (IE). They recommended IE as “a general-purpose method of APC analysis with potentially

wide applicability in the social sciences” (Yang et al. 2008:1699) on the grounds that IE has desirable statistical properties such as unbiasedness and consistency.

However, in this chapter I show that IE cannot be used to recover the true age, period, and cohort effects because IE, like CGLM, imposes a constraint on parameter estimation that is difficult to verify using theories or empirical evidence; that is, the validity of IE relies on assumptions that are very difficult to verify in applied practice. In this sense, IE is no better than CGLM. In fact, IE is equivalent to the Principal Component Estimator, an estimator with a potential for bias that was noted by its developer (Kupper et al. 1985). Unfortunately, this has not been understood by the community of demographers, sociologists, and epidemiologists who have used IE in a wide variety of research applications. As I demonstrate below, many researchers have misunderstood what IE actually estimates and how IE estimates should be interpreted, resulting in inappropriate applications of IE in empirical research and potentially misleading substantive conclusions.

This chapter contributes to the literature in two ways: First, although O’Brien (2011a) clarified that IE assumes a special constraint – the null-vector constraint – on parameters, it is challenging for researchers to fully appreciate and evaluate the appropriateness of this constraint when applying IE in substantive studies. In this chapter, I derive an easily-understood form of IE’s constraint on the *linear* components of age,

period, and cohort effects so the implications of using IE to estimate the true age, period, and cohort effects can be better understood.¹

Second, while scholars agree that IE is a constrained estimator, they debate whether IE can provide reliable estimates of the true age, period, and cohort trends (see Fu et al. 2011; O'Brien 2011b). I address this debate using several types of simulated data generated based on social theories. By comparing IE estimates to the true effects in various circumstances, I show that IE does not work better than CGLM for recovering the true age, period, and cohort trends in empirical research.

This paper is organized as follows. I begin with an introduction of the APC multiple classification model and the identification problem. While reviewing the methodological challenge that has hampered APC research for decades, this section establishes a framework for discussing the nature and limitations of different constrained APC estimators including IE and CGLM. I then review how IE's developers have described IE and how applied researchers have understood and used it in substantive studies; the two are often not the same. As a result, many scholars have misunderstood IE, so that this technique has been misused in empirical research. To clarify this common misunderstanding and avoid further misuse, in the section "The Linear Constraint Implied by IE," I derive the constraint that IE imposes on the *linear* components of age, period, and cohort effects. In the "Application Scope" section that follows the technical

¹ One way to characterize the effects of an interval variable like time is to break the effect into two components: linear and non-linear (curvature or deviations from linearity) trends. It has been known at least since Holford (1983) that the linear components of age, period, and cohort effects cannot be estimated without constraints because they are not identified. In contrast, non-linear age, period, and cohort trends can be estimated without bias.

discussion of IE's linear constraint, I use simulations to demonstrate how this constraint affects estimation of age, period, and cohort effects. Based on these mathematical derivations and simulation evidence, I conclude that IE cannot and should not be used to estimate true age, period, and cohort effects.

The Identification Problem

I first review the identification problem that IE and other constrained estimators are intended to address to develop a framework for understanding the nature of these methods. In APC analysis, researchers have conventionally used the Analysis of Variance (ANOVA) model to separate the independent age, period, and cohort effects:

$$g(E(Y_{ij})) = \mu + \alpha_i + \beta_j + \gamma_k, \quad (1)$$

for age groups $i = 1, 2, \dots, a$, periods $j = 1, 2, \dots, p$, and cohorts $k = 1, 2, \dots, (a + p - 1)$, where $\sum_{i=1}^a \alpha_i = \sum_{j=1}^p \beta_j = \sum_{k=1}^{a+p-1} \gamma_k = 0$. $E(Y_{ij})$ denotes the expected value of the outcome of interest Y for the i th age group in the j th period of time; g is the "link function"; α_i denotes the mean difference from the global mean μ associated with the i th age category; β_j denotes the mean difference from μ associated with the j th period; γ_k denotes the mean difference from μ due to the membership in the k th cohort. The usual ANOVA constraint applies where the sum of coefficients for each effect is set to zero.

For a normally distributed outcome Y_{ij} , the ANOVA model above can also be written in a generic regression fashion:

$$Y = Xb + \varepsilon, \quad (2)$$

where Y is a vector of outcomes; X is the design matrix; b denotes a parameter vector with elements corresponding to the effects of age, period, and cohort groups; and ε

denotes random errors with distribution centered on zero. Then the estimated age, period, and cohort effects can be obtained using the ordinary least squares (OLS) method:

$$\hat{b} = (X^T X)^{-1} X^T Y. \quad (3)$$

Unfortunately, the inverse of the matrix $(X^T X)^{-1}$ does not exist because of the age-period-cohort linear dependency, so the parameter vector b is inestimable. This is the identification problem in APC analysis: no unique set of coefficients can be obtained because an infinite number of solutions give *identical* fits to the data.

This identification problem can be shown more explicitly. For simplicity, suppose the data we have are perfect, without random or measurement errors, so that $\varepsilon = 0$; then the problem is mathematical rather than statistical, and the regression model is:

$$Y = Xb. \quad (4)$$

Due to the linear dependency between age, period, and cohort, there exists a nonzero vector b_0 , a linear function of the design matrix X , such that the product of the design matrix and the vector equals zero:

$$Xb_0 = 0. \quad (5)$$

In other words, b_0 represents the null space of the design matrix X , which has dimension equal to one. (The null space has dimension one by the specification of model (1), and the value of b_0 is given below.) It follows that the parameter vector b can be decomposed into components:

$$b = b_1 + s \cdot b_0, \quad (6)$$

where s is an arbitrary real number corresponding to a specific solution to equation (4), and b_1 is a linear function of the parameter vector b , corresponding to the projection of b

on the non-null space of the design matrix X , orthogonal to the null space. b_1 and b_0 are thus orthogonal to each other. That is, b_1 is the part of b that is in the non-null space of the design matrix X , orthogonal (perpendicular) to the null space, so that b_0 is orthogonal to b_1 , i.e., $b_1 \cdot b_0 = 0$.

Given equations (4) and (6), the following equation must hold:

$$Y = Xb = X(b_1 + s \cdot b_0) = Xb_1 + s \cdot Xb_0. \quad (7)$$

But $Xb_0 = 0$ and thus $s \cdot Xb_0 = 0$, so equation (7) is true for all values of s . That is, s can be *any* real number, and each distinct value of s gives a distinct solution to equation (4). Therefore, an infinite number of possible solutions for b exist, and no solution can be deemed the uniquely preferred or “correct” solution without additional constraints on b .

To illustrate, suppose the data have three age groups, three periods, and five cohorts and that error is zero for ease of presentation (and without loss of generality). Table 2.1 presents three different parameter vectors $b^T = (u, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$ arising from three different values of s , namely 0, 2, and 10. In Table 2.2’s top panel, the observed value in each cell is represented in terms of the unknown parameters α_i, β_j , and γ_k . Table 2.2’s bottom panel shows the fitted values $u + \alpha_i + \beta_j + \gamma_k$ based on Table 2.1’s three different s ’s in the same tabular form as above. Note that these three sets of fitted values are *identical* although the parameter vectors in Table 2.1 differ. In fact, these parameter vectors are not just different; their age and period effects change directions depending on s , and the data cannot distinguish between different s ’s.

Table 2.1. Different Values of s and the Corresponding Parameters

s	Age			Period			Cohort				
	α_1	α_2	α_3	β_1	β_2	β_3	γ_1	γ_2	γ_3	γ_4	γ_5
0	2	0	-2	-1	0	1	-1	-0.5	0	0.5	1
2	0	0	0	1	0	-1	-5	-2.5	0	2.5	5
10	-8	0	8	9	0	-9	-21	-10.5	0	10.5	21

Note. s is an arbitrary real number corresponding to a specific solution to equation (4). Numbers in each row are a set of age, period, and cohort coefficients corresponding to a specific value of s .

Table 2.2 Tabular Data: Unobserved parameters and fitted values from Table 2.1's three different parameter vectors

		Period		
		1	2	3
Unobserved Parameters	Age 1	$\mu + \alpha_1 + \beta_1 + \gamma_3$	$\mu + \alpha_1 + \beta_2 + \gamma_4$	$\mu + \alpha_1 + \beta_3 + \gamma_5$
	Age 2	$\mu + \alpha_2 + \beta_1 + \gamma_2$	$\mu + \alpha_2 + \beta_2 + \gamma_3$	$\mu + \alpha_2 + \beta_3 + \gamma_4$
	Age 3	$\mu + \alpha_3 + \beta_1 + \gamma_1$	$\mu + \alpha_3 + \beta_2 + \gamma_2$	$\mu + \alpha_3 + \beta_3 + \gamma_3$
Observed Values	Age 1	11	12.5	14
	Age 2	8.5	10	11.5
	Age 3	6	7.5	9

Note. The bottom panel presents *identical* observed values produced by the three different parameter vectors in Table 2.1.

Taken together, Tables 2.1 and 2.2 show that for a single dataset, an infinite number of possible solutions for age, period, and cohort effects exist, and each solution corresponds to a specific value of s . Therefore, any solution, or alternatively, none of these solutions, can be viewed as reflecting the “true” effects even though different values of s give radically different age, period, and cohort effects. In social science research, data inevitably contain random and/or measurement errors so researchers will not have the perfect fit of the idealized data above; however, the fundamental identification problem remains. Various methods have been developed to address the identification problem and find a set of uniquely preferred estimates. In the section below, I will consider IE and other solutions to the identification problem that impose a constraint on b .

The Constrained Approach: IE and CGLM

A large body of literature dating back to the 1970s has addressed the identification problem. Mason et al. (1973) explicated the “identification problem” in APC analysis and proposed the Constrained Generalized Linear Model (CGLM), a coefficient-constrained approach that has been used as a conventional method for APC analysis. This method places at least one identifying restriction on the parameter vector b in equation (2). Usually the effects of the first two age groups, periods, or cohorts are constrained to be equal based on theoretical or external information. With this additional constraint, the APC model becomes just-identified and unique OLS and maximum likelihood (ML) estimators exist. However, such theoretical information often does not exist or cannot easily be verified. Also different choices of identifying constraint can produce widely

different estimates for age, period, and cohort effects. That is, CGLM estimates are quite sensitive to the choice of constraints (Rodgers 1982a, 1982b; Glenn 2005).

More recently, a group of scholars has developed a new APC estimator, called the *Intrinsic Estimator* (IE). They argued that IE has clear advantages over CGLM (called “CGLIM” in Yang et al. 2008) and can produce valid estimates of the true age, period, and cohort effects (see Fu 2000, 2006; Yang et al. 2004, 2008). The most compelling evidence they provided to support this claim is simulation results where IE and CGLM estimates were compared to the true effects of age, period, and cohort (see Yang et al. 2008:1718-1719). They concluded that IE outperforms CGLM because IE estimates are closer to the true parameters that generate the data than CGLM (*ibid.*:1719-1722).

This evidence could easily be interpreted as confirmation that IE produces unbiased estimates of the true age, period, and cohort effects. Unfortunately, few clarifications are provided and the developers of IE are sometimes unclear about what IE actually estimates themselves. For example,

“for a finite number of time periods of data, the IE produces an unbiased estimate of the coefficient vector.” (Yang 2008:400)

“Because of its estimability and unbiasedness properties, the IE may provide a means of accumulating reliable estimates of the trends of coefficients across the categories of the APC accounting model.” (Yang et al. 2008:1711)

“[T]he IE, by its very definition and construction, satisfies the estimability condition. ... If other estimators do indeed satisfy the estimability

condition, then they also produce unbiased estimates of the A, P, and C effect coefficients. If not, then the estimates they produce are biased.”

(*ibid.*:1710)

“[P]erhaps most importantly for empirical applications of APC analysis, the IE produces estimated age, period, and cohort coefficients and their standard errors in a direct way, without the necessity of choosing among a large array of possible constraints on coefficients that may or may not be appropriate for a particular analysis.” (Yang et al. 2004:105)

Many researchers doing substantive APC analyses have interpreted these and other statements to mean that IE produces unbiased estimates of true age, period, and cohort effects. Consequently, they have used IE in empirical research to address substantive issues including mortality, disease, and religious activity (e.g., Keyes and Miech 2013; Winkler and Warnke 2012; Schwadel 2011; Langley et al. 2011; Miech et al. 2011). These authors seem convinced that IE produces unbiased estimates of age, period, and cohort effects. For example,

“[r]ecent advances in modeling APC effects with repeated cross-sectional data allow age, period, and cohort effects to be simultaneously estimated without making subjective choices requiring constraining data or dropping age, period, or cohort indicators from the model. In particular, APC intrinsic estimator models provide unbiased estimates of regression coefficients for age groups, time periods, and birth cohorts (Fu, 2000).”

(Schwadel 2011:183)

“[T]he intrinsic estimator provides unbiased estimates of age, period, and cohort effects.” (*ibid.*:184)

“The IE model has been recommended as a better alternative to the widely discussed constrained generalized linear model (CGLM) (Yang et al. 2004). We used the IE model to estimate individual effects of age, period, and cohort for males and females separately.” (Langley et al. 2011:106)

“The IE is an approach that places a constraint on the model, but not a constraint that affects the estimation of regression parameters for age, period, and cohort in any way. That is, the regression parameter estimates are unbiased by the constraint placed, and a unique set of regression estimates can be estimated.” (Keyes and Miech 2013:2)

Unfortunately, claims of this sort are incorrect; as I demonstrate below, IE does impose constraints that are as consequential as those imposed by CGLM. To help researchers better understand the constraint imposed by IE and make informed decisions in choosing an APC estimator, I will first derive an easily-understood form of IE’s constraint. Because an unbiased and consistent estimator is desirable and necessary to produce reliable and valid results, I will then address how IE’s constraint affects these key properties: unbiasedness (Is the expectation of IE the “true” age, period, and cohort effects?) and consistency (As the sample size increases, does IE converge to the “true” effects?).

The Linear Constraint Implied by IE

To understand IE's constraint and its implications for estimation, it is helpful to review IE's conceptual foundation and computational algorithm. IE can be viewed as an extension of Principal Component (PC) Analysis, a multi-purpose technique that can be used to deal with identification problems when explanatory variables are highly correlated. By transforming correlated explanatory variables to a set of orthogonal linear combinations of these variables, called principal components, PC analysis can be a useful tool for reducing data redundancy and developing predictive models.

In contrast, the goal of IE is neither data reduction nor prediction, but estimation of the effects of, and capturing the general trends of, age, period, and cohort.² IE's computational algorithm includes five steps: (a) transform the design matrix X to the PC space using its eigenvector matrix; (b) in the PC space, identify the "null eigenvector" – the special eigenvector that corresponds to an eigenvalue of zero – and the corresponding null subspace (with one dimension) and non-null subspace (with $m - 1$ dimensions, where m denotes the number of coefficients to be estimated); (c) in the non-null subspace of $m - 1$ dimensions, regress the outcome of interest using OLS or ML on the $m - 1$ PCs to obtain $m - 1$ coefficient estimates; (d) extend the $m - 1$ coefficient estimates to the whole PC space of dimension m by adding an element corresponding to the null

² It is important to distinguish data reduction or prediction from coefficient estimation. Because the identification problem does not prevent us from obtaining a set of solutions with good fit to the data, we can still make good predictions. The PC technique treats such problems as data redundancy and allows us to obtain one solution. However, as noted above, none of these solutions is the uniquely preferred solution, the solution that APC techniques including IE aim to discover. Therefore, providing a solution for the purpose of prediction is not the same as finding a uniquely preferred solution for *estimation of separate age, period, and cohort effects*.

eigenvector direction and arbitrarily setting it to zero; and (e) use the eigenvector matrix to transform the extended coefficient vector estimated in the PC space, including the added zero element, back to the original age-period-cohort space to obtain estimates for age, period, and cohort effects (see Yang 2004; Yang et al. 2008).³

The fourth step, “extend the $m - 1$ coefficient estimates to the whole PC space of dimension m by adding an element corresponding to the null eigenvector direction and arbitrarily setting it to zero,” carries the key assumption of the IE approach to APC analysis. This assumption is implicit yet has major implications for the validity and application of the IE approach. Specifically, setting the “coefficient of the null eigenvector”, s , to zero is equivalent to assuming

$$b \cdot b_0 = 0, \quad (8)$$

i.e., the projection of b on b_0 is zero, where b and b_0 were defined in equation (6). Kupper and colleagues (1985) provided a closed-form representation for the eigenvector b_0 . Using vector notation,⁴

$$b_0 = (0, A, P, C)^T, \quad (9)$$

where

$$A = \left(1 - \frac{1+a}{2}, \dots, (a-1) - \frac{1+a}{2}\right)$$

³ Alternatively, Yang (2008) described the computational algorithm of IE as follows: after obtaining $r-1$ coefficients in the PC space (w_2, \dots, w_r), “[s]et coefficient w_1 equal to 0 and transform the coefficients vector $w = (w_1, \dots, w_r)^T$ ” (Appendix, p413), where w_1 corresponds to the null eigenvector direction.

⁴ Yang et al. (2004, 2008) use $b_0^* = \frac{b_0}{\|b_0\|}$, where $\|b_0\|$ is the length of b_0 , so b_0^* has a length of 1. b_0 is used in this paper because it is simply a multiple of b_0^* and is simpler for exposition and computation.

$$P = -1 \cdot \left(1 - \frac{1+p}{2}, \dots, (p-1) - \frac{1+p}{2}\right)$$

$$C = \left(1 - \frac{a+p}{2}, \dots, (a+p-2) - \frac{a+p}{2}\right).$$

For example, when $a = 3$ and $p = 3$, that is, for three age groups and three time periods, b_0 is

$$b_0 = (0, -1, 0, 1, 0, -2, -1, 0, 1)^T, \quad (10)$$

where $A = (-1, 0)$, $P = (1, 0)$, and $C = (-2, -1, 0, 1)$.

What does equation (8) mean? What is the specific form of this constraint for datasets with varying number of age, period, and cohort groups? To illustrate, suppose that age, period, and cohort each have effects on the outcome variable that show a linear trend. Denote these trends as k_a, k_p , and k_c , respectively, the intercepts for the three variables as i_a, i_p , and i_c , and the overall mean as μ . Thus the effects associated with the three age categories are $i_a, i_a + k_a$, and $i_a + 2 \cdot k_a$, respectively. Similarly, the effects related to the three periods are $i_p, i_p + k_p$, and $i_p + 2 \cdot k_p$, respectively. For the five cohorts, the effects are $i_c, i_c + k_c, i_c + 2 \cdot k_c, i_c + 3 \cdot k_c$, and $i_c + 4 \cdot k_c$, respectively. Then the parameter vector, b , can be written as:

$$b = (\mu, i_a, i_a + k_a, i_p, i_p + k_p, i_c, i_c + k_c, i_c + 2 \cdot k_c, i_c + 3 \cdot k_c)^T, \quad (11)$$

where the last category of each variable is omitted as the reference group. According to the constraint for age effects in model (1), we know that

$$\sum_{i=1}^a \alpha_i = i_a + (i_a + k_a) + (i_a + 2 \cdot k_a) = 3 \cdot i_a + 3 \cdot k_a = 0, \quad (12)$$

which implies that

$$i_a = -k_a. \quad (13)$$

Similarly, it can be shown using the constraint for period and cohort effects in model (1) that

$$i_p = -k_p, \quad (14)$$

and

$$i_c = -2 \cdot k_c. \quad (15)$$

Using equations (13), (14), and (15), equation (11) can be simplified as:

$$b = (\mu, -k_a, 0, -k_p, 0, -2 \cdot k_c, -k_c, 0, k_c)^T. \quad (16)$$

Since the constraint that IE implicitly imposes is $b \cdot b_0 = 0$, by equations (8), (10) and (16), the specific form of IE's linear constraint (LC) for APC data with three age categories, three periods, and five cohorts are

$$\begin{aligned} b \cdot b_0 &= \mu \cdot 0 + (-k_a) \cdot (-1) + 0 \cdot 0 + (-k_p) \cdot 1 + 0 \cdot 0 + (-2 \cdot k_c) \cdot (-2) + (-k_c) \cdot \\ &(-1) + 0 \cdot 0 + k_c \cdot 1 = k_a - k_p + 6 \cdot k_c = 0. \end{aligned} \quad (17)$$

In other words, when age, period, and cohort show linear trends, IE's implicit constraint is that these linear trends must satisfy equation (17). If, in fact, the true age, period, and cohort trends do not satisfy this equation, then the implicit LC imposed by IE is incorrect.

To illustrate the implications of IE's LC, I simulate normally distributed data as follows. For those at age i in period j , the mean response is $10 + k_a \cdot age_i + k_p \cdot period_j + k_c \cdot cohort_{ij}$ and the standard deviation of error ε equals 0.1. The number of age and period groups is fixed at three each. I consider three sets of true $k_a, k_p,$ and k_c : (a) $k_a = 1, k_p = 7, k_c = 1$; (b) $k_a = 1, k_p = 7, k_c = 10$; and (c) $k_a = 3, k_p = 1, k_c = 4$. For each selection of true $k_a, k_p,$ and k_c , I simulate 1,000 such data sets by drawing

random errors. As shown in Table 2.3, for dataset 1 the true effects for the three age categories are $-1, 0$, and 1 , respectively, so k_a , the linear trend in age effects, equals 1 . The period effects are $-7, 0$, and 7 , respectively, so k_p is 7 . Similarly, since the cohort effects are $-2, -1, 0, 1$, and 2 , k_c is 1 . Note that for this dataset,

$$k_a - k_p + 6 \cdot k_c = 1 - 7 + 6 \cdot 1 = 0, \quad (18)$$

i.e., the relationship between the linear trends in the true age, period, and cohort effects satisfies equation (17), the LC implicit in IE. However, for datasets 2 and 3 generated by the other sets of true k_a, k_p , and k_c in Table 2.3, equation (17) does not hold. Specifically, for the second set, $k_a = 1, k_p = 7$, and $k_c = 10$, so

$$k_a - k_p + 6 \cdot k_c = 1 - 7 + 6 \cdot 10 = 54 \neq 0; \quad (19)$$

And for the third set, $k_a = 3, k_p = 1$, and $k_c = 4$, so

$$k_a - k_p + 6 \cdot k_c = 3 - 1 + 6 \cdot 4 = 26 \neq 0. \quad (20)$$

Table 2.3 presents IE estimates, averaged over the 1,000 simulated datasets, for the three sets of age, period, and cohort effects. The bias of IE is estimated by the difference between the truth and the averaged IE estimates. Table 2.3 shows that for dataset 1, IE yields good estimates because the true k_a, k_p , and k_c in the data satisfy equation (17), the implicit LC that IE imposes. Specifically, the estimated slopes for age, period, and cohort are $\hat{k}_a = 0.999$, $\hat{k}_p = 7.001$, and $\hat{k}_c = 1.000$, respectively. In contrast, IE returns highly biased estimates, very different from the true effects, for the second and third datasets because the true k_a, k_p , and k_c do not satisfy IE's LC. For example, for datasets 2 and 3, the estimated age effects, averaged over the 1,000

simulations, show a downward trend ($\hat{k}_a = -5.750$ for dataset 2 and $\hat{k}_a = -2.582$ for dataset 3) when the true trend is upward (the true age slopes are $k_a = 1$ for dataset 2 and $k_a = 3$ for dataset 3).

Table 2.3. Simulation Results: IE estimates for three datasets

		Dataset 1			Dataset 2			Dataset 3		
		Truth	IE	Bias	Truth	IE	Bias	Truth	IE	Bias
Age	1	-1	-0.997	0.003	-1	5.747	6.747	-3	0.249	3.249
	2	0	-0.002	-0.002	0	0.002	0.002	0	0.000	0.000
	3	1	0.999	-0.001	1	-5.749	-6.749	3	-0.249	-3.249
Period	1	-7	-6.999	0.001	-7	-13.75	-6.750	-1	-4.250	-3.250
	2	0	-0.002	-0.002	0	0.002	0.002	0	-0.002	-0.002
	3	7	7.002	0.002	7	13.748	6.748	1	4.252	3.252
Cohort	1	-2	-2.001	-0.001	-20	-6.497	13.503	-8	-1.500	6.500
	2	-1	-0.998	0.002	-10	-3.253	6.747	-4	-0.750	3.250
	3	0	-0.001	-0.001	0	0.002	0.002	0	0.000	0.000
	4	1	1.004	0.004	10	3.250	-6.750	4	0.750	-3.250
	5	2	1.996	-0.004	20	6.498	-13.502	8	1.500	-6.500

Note. For each dataset, the IE estimates are averaged over 1,000 simulations. The bias of IE is evaluated by the difference between the true effects and the IE estimates, averaged over 1,000 simulations. Equation (17) holds for dataset 1, but does not hold for datasets 2 and 3.

Note that equation (17) is derived for the simplest scenario where the age, period, and cohort trends are purely linear. For more complex scenarios where these trends are not purely linear, IE's constraint depends on the non-linear components of the age, period, and cohort effects.⁵ For example, suppose that age, period, and cohort each have effects on the outcome of interest that include a linear and a quadratic trend. Denote the quadratic trends as k'_a, k'_p , and k'_c , respectively. Using the same derivation above, the specific form of IE's constraint for APC data with three age categories, three periods, and five cohorts is

$$(k_a - k_p + 6 \cdot k_c) + \frac{5}{3}(k'_a - k'_p + 12 \cdot k'_c) = 0. \quad (21)$$

That is, when age, period, and cohort effects include quadratic components, these effects must satisfy equation (21) in order for IE to yield good estimates. Equation (17) can be viewed as a special case of equation (21) when there are no quadratic or higher-order non-linear components in the age, period, and cohort effects. Alternatively, because the linear dependency between age, period, and cohort does not affect the identification of nonlinear effects, IE's constraint can be said to *bind* only on the linear age, period, and cohort trends, and the specific value of the constraint on the linear effects is determined by the non-linear effects, which are estimable.

⁵ The constraint imposed by IE depends on how model (2) is parameterized. If the model is parameterized in terms of orthogonal polynomial contrasts for each of the age, period, and cohort effects, as in Holford (1983), then IE imposes a constraint solely on the linear contrasts of age, period, and cohort effects irrespective of any non-linear trends that are present. The parameterization used here is more common, e.g., Kupper et al. (1985), and in this parameterization, the constraint on the linear components of the age, period, and cohort effects depends on the non-linear components when both components are present.

For any coefficient-constraint approach such as CGLM and IE, “the choice of constraint is *the* crucial determinant of the accuracy in the estimated age, period, and cohort effects” (Kupper et al. 1985:822). Since the constraint assumption strongly affects estimation results, no matter what constraint a statistical method assumes, that method produces good estimates only when its assumption approximates the true structure of the data under investigation. It follows that when there are three age groups, three periods, and five cohorts and their effects are purely linear, IE can only yield accurate estimates when these linear effects of age, period, and cohort satisfy equation (17). Unfortunately, researchers usually have no *a priori* knowledge about true age, period, and cohort effects that would allow them to evaluate whether the constraint implied in equation (17) holds. Therefore, researchers cannot assess whether IE produces unbiased estimates of age, period, and cohort effects for their data. Thus IE is no better than CGLM in this respect.

More importantly, the exposition above indicates that the LC assumed by IE also depends on the design matrix X , i.e., on the number of age, period, and cohort groups. For example, if we add one age group to our example, such that we now have four age groups, three periods, and six cohorts, then following the same derivation used above, the LC implied by IE is

$$b \cdot b_0 = 2.75 \cdot k_a - k_p + 11.25 \cdot k_c = 0, \quad (22)$$

or

$$b \cdot b_0 = (k_a - k_p + 6 \cdot k_c) + (1.75 \cdot k_a + 5.25 \cdot k_c) = 0. \quad (23)$$

Compared to equation (17) for the case of three age groups, three periods, and five cohorts, equations (22) and (23) show that adding an age group dramatically changes the

constraint so that the true effects satisfying IE's LC with three age categories no longer satisfy this LC when an age category is added. Readers can verify that increasing or reducing the number of periods or cohorts also greatly alters IE's LC.

These examples demonstrate that not only does IE rely on a constraint like CGLM does, but unlike CGLM — where the constraint (e.g., equal effects for the first two age groups) is explicit and rationalized by theoretical account or side information — the LC of IE is implicit and varies depending on the number of age, period, and cohort groups. Although this constraint has been described as minimal (e.g., Schwadel 2011; Yang et al. 2008), in fact, as shown, it can have major implications for the quality of substantive results.

Theoretically speaking, the limitation of IE results from a misinterpretation of the constraint that IE imposes on parameter estimation. It is true that b_0 , the null eigenvector, is determined by the design matrix, but it is incorrect to conclude that therefore b_0 “should not play any role in the estimation of effect coefficients” (Yang et al. 2008:1705). Rather, both the null eigenvector and non-null eigenvectors (with nonzero eigenvalues) are determined by the design matrix, that is, by the number of age, periods, and cohort groups. To this extent, it is no less likely that the data contain a significant component in the b_0 direction than in the directions of the non-null eigenvectors. The fact that s , the coefficient for b_0 , can be *any* real number without changing the fitted values Xb simply means that variation in Y in the direction of b_0 is not estimable. If the data have variation in this direction, IE will mistakenly attribute that variation to other columns in the design matrix, causing significant errors in estimation.

The Implications of IE’s Constraint: Is IE an unbiased and consistent estimator?

Because IE imposes a constraint on the linear age, period, and cohort trends, IE yields reliable estimates only when the true trends satisfy its constraint. However, Yang and colleagues argue that “[b]ecause of its estimability and unbiasedness properties, the IE may provide a means of accumulating reliable estimates of the trends of coefficients across the categories of the APC accounting model” (*ibid.*:1711). In the discussion below, I clarify that IE is not an unbiased estimator of the “true” age, period, and cohort effects. I also use concrete examples to illustrate that IE is not consistent and explain why IE appears to be converging to the truth in Yang et al. (2008)’s article. This section may be particularly helpful for non-technical researchers.

Biasedness

By definition, an estimator δ is an unbiased estimator of a parameter θ if the expectation of δ over the distribution that depends on θ is equal to θ , or $E_{\theta}(\delta) = \theta$. It follows that, for an unbiased APC estimator, its expectation must be the true effects of age, period, and cohort.⁶ Per this definition, if IE is an unbiased estimator, the expected value of IE must be the true age, period, and cohort effects. The following mathematical computation shows, however, that the expectation of the IE estimator is not the true effects unless those true effects happen to satisfy IE’s implicit constraint.

⁶ Yang and colleagues have used “unbiasedness” in a different sense to mean that the expectation of IE is equal to b_I , the projection of parameter vector b onto the non-null space of design matrix X (e.g., see *ibid.*:1709). This is an important distinction because the true parameter vector b can be very different from its projection b_I onto the non-null space, the vector that IE actually estimates. Because APC analysts are usually interested in estimating the true age, period, and cohort effects, the classic concept of unbiasedness is more relevant to APC research than that used by IE’s proponents. Thus I use “unbiasedness” in its classic sense in the following discussion.

As noted in the section above, the key computation of IE is to extend the coefficient estimates in the PC space, b'

$$(b')^T = (b'_0, b'_1, b'_2, \dots, b'_{m-1}) \quad (24)$$

by adding a zero element such that

$$(b'_{new})^T = (b'_0, b'_1, b'_2, \dots, b'_{m-1}, 0), \quad (25)$$

where b'_{new} corresponds to the projection of the coefficient vector b in the non-null space, i.e., b_1 in equation (6). IE then transforms the extended coefficient vector b'_{new} including the added zero element, back to the original age-period-cohort space to obtain coefficient estimates for age, period, and cohort.

Given that OLS and ML estimators have been proven unbiased in simpler — identifiable — problems with normally distributed errors as in equation (2), and since IE uses these methods to obtain estimates for b_1 , whose projection in the PC space corresponds to the extended coefficient vector b'_{new} , IE yields unbiased estimates for b_1 . In other words,

$$E(b_{IE}) = b_1. \quad (26)$$

Based on the preceding discussion of the identification problem, the true parameter space b can be decomposed into two orthogonal subspaces corresponding to b_1 and b_0 in equation (6), which is equivalent to

$$b_1 = b - s \cdot b_0. \quad (27)$$

Substituting equation (27) in (26) results in

$$E(b_{IE}) = b_1 = b - s \cdot b_0. \quad (28)$$

Equation (28) means that the expectation of the IE estimator will be different from the true effects b unless $s \cdot b_0 = 0$, i.e., unless $s = 0$. IE assumes $s = 0$; thus, IE is a biased estimator when the true value of s is anything but 0. The larger the absolute value of s , the more biased the IE estimates become.

For researchers who wish to investigate age, period, and cohort effects for the purposes of substantive demographic, social, or other applied research, there exists little theoretical or empirical knowledge about the value of s and what b_0 , the “null eigenvector,” may imply about the outcome variable. In specific applications, then, IE must be assumed to be biased, resulting in misleading conclusions about the true age, period, and cohort effects unless proven otherwise.

Note that IE’s developers argue that IE satisfies the “estimability criterion” proposed by Kupper et al. (1985), so IE is in that sense an unbiased estimator. However, estimability of a function of b implies unbiased estimation only of the estimable function of b , not necessarily of the true parameter b itself. b_1 , the projection of the parameter vector onto the non-null space, is indeed an estimable function of b , the true parameter vector, and thus IE is an unbiased estimator of b_1 . But IE is a biased estimator for the true APC effects when b_1 is different from b . Therefore, it is not accurate to say that “Kupper et al. (1985) ... suggested that an estimable function satisfying this condition resolves the identification problem” as claimed in Yang and associates (2008:1703). To emphasize, estimability in the non-null space does not imply unbiasedness in estimating the *true* age, period, and cohort effects. Discovering a set of estimable functions is not the same as solving the identification problem.

Consistency

In statistics, for an estimator δ to be a consistent estimator of an unknown parameter space θ , δ must converge in probability to θ as the sample size grows. If δ is unbiased, consistency usually follows immediately. A biased estimator can be consistent if its bias decreases as the sample size increases. However, the bias of IE, $s \cdot b_0$, does not necessarily shrink as the sample size grows. Thus, IE is not a consistent estimator of the coefficient vector b .

This theoretical argument can be illustrated with simulations. I simulate normally-distributed data using the same function as that for Dataset 1 in Table 2.3: For those at age i in period j , the mean response is $10 + 1 \cdot age_i + 7 \cdot period_j + 1 \cdot cohort_{ij}$ and standard deviation of error = 0.1. I begin with three age groups and three periods, and then increase the number of periods to six and 12, respectively. For each scenario, I simulate 1,000 such datasets by drawing random errors. If IE is a consistent estimator, as the number of periods increases, the resulting estimates should get closer and closer to the true effects that we know based on the simulation function.

Table 2.4 presents the IE estimates, averaged over 1,000 datasets, for the three scenarios in which the number of periods is set at three, six, and 12, respectively. It shows that the IE estimates are not converging to the truth and the bias appears to increase as the number of periods increases from three to 12. Specifically, when p , the number of periods, equals six and 12, although IE correctly captures the direction of the age, period, and cohort trends, there is no evidence that these estimates are converging to the truth; the estimated age, period, and cohort slopes are $\hat{k}_a = 2.144$, $\hat{k}_p = 5.857$, and

$\hat{k}_c = 2.144$, respectively, when $p = 6$; $\hat{k}_a = 3.017$, $\hat{k}_p = 4.983$, and $\hat{k}_c = 3.017$ when p increases to 12. In fact, even with an unrealistically large number of periods (e.g., 100 periods), as I show in Appendix Figure A2.1, the IE estimates do not appear to converge to the truth.

The developers of IE correctly note that the estimation of period and cohort effects will not improve with more time periods because “adding a period to the data set does not add information about the previous periods or about cohorts not present in the period just added” (*ibid.*: 1718). However, when they simulated data, the IE estimates for age effects did appear to become closer and closer to the true values as the number of periods increased. They simulated data using the following function:

$$y_{ij} \sim \text{Poisson}\{\exp[0.3 + 0.1(\text{age}_i - 5)^2 + 0.1 \sin(\text{period}_j) + 0.1 \cos(\text{cohort}_{ij}) + 0.1 \sin(10 \cdot \text{cohort}_{ij})]\}. \quad (29)$$

Table 2.4. Simulation Results: Inconsistent IE estimates as the number of periods increases

		Periods=3			Periods=9			Periods=12		
		Truth	IE	Bias	Truth	IE	Bias	Truth	IE	Bias
Age	1	-1	-0.997	0.003	-1	-2.144	-1.144	-1	-3.016	-2.016
	2	0	-0.002	-0.002	0	0.000	0.000	0	-0.000	-0.000
	3	1	0.999	-0.001	1	2.144	1.144	1	3.017	2.017
Period	1	-7	-6.999	0.001	-17.5	-14.642	2.858	-38.5	-27.407	11.093
	2	0	-0.002	-0.002	-10.5	-8.783	1.717	-31.5	-22.427	9.073
	3	7	7.002	0.002	-3.5	-2.931	0.569	-24.5	-17.441	7.059
	4				3.5	2.928	-0.572	-17.5	-12.462	5.038
	5				10.5	8.785	-1.715	-10.5	-7.470	3.030
	6				17.5	14.643	-2.857	-3.5	-2.493	1.007
	7							3.5	2.494	-1.006
	8							10.5	7.473	-3.027
	9							17.5	12.455	-5.045
	10							24.5	17.441	-7.059
	11							31.5	22.425	-9.075
	12							38.5	27.412	-11.088
Cohort	1	-2	-2.001	-0.001	-3.5	-7.497	-3.997	-6.5	-19.612	-13.112
	2	-1	-0.998	0.002	-2.5	-5.362	-2.862	-5.5	-16.590	-11.090
	3	0	-0.001	-0.001	-1.5	-3.214	-1.714	-4.5	-13.575	-9.075
	4	1	1.004	0.004	-0.5	-1.071	-0.571	-3.5	-10.558	-7.058
	5	2	1.996	-0.004	0.5	1.074	0.574	-2.5	-7.542	-5.042
	6				1.5	3.215	1.715	-1.5	-4.525	-3.025
	7				2.5	5.353	2.853	-0.5	-1.506	-1.006
	8				3.5	7.502	4.002	0.5	1.509	1.009
	9							1.5	4.529	3.029
	10							2.5	7.542	5.042
	11							3.5	10.559	7.059
	12							4.5	13.575	9.075
	13							5.5	16.589	11.089
	14							6.5	19.604	13.104

It appears that IE estimates of the age effects converge to the true effects in this simulation as the number of periods increases because IE's implicit LC is not satisfied by the "true" age, period, and cohort effects in the simulation mechanism (29) with five periods ($b \cdot b_0 = -0.339$), but the true effects do approximately satisfy the LC ($b \cdot b_0 = -0.036$) when the number of periods increases to 50. In other words, IE appears to perform better as the number of periods increases *not* because IE is a consistent procedure but because the true effects used in the data-generating function (29) conform better to IE's implicit LC as the number of periods increases.

For demographic or social data where the linear trends in the three variables are unknown, adding more periods or cohorts promises nothing about the accuracy of the coefficient estimation for either age or period or cohort effects. That is, even with a sufficiently large sample, researchers using IE to estimate the true age, period, and cohort effects are not guaranteed to have desirable results that are close to the true values.

Application Scope: IE vs. CGLM

The preceding discussions of IE's linear constraint (LC) and statistical properties are fairly technical. In this section, I will use several types of simulated data to illustrate how the implicit LC of IE affects its ability to recover the underlying age, period, and cohort effects in social science research⁷. This exercise is important because scholars have

⁷Yang and colleagues have used empirical data, where the true effects are unknown, to assess the properties and performance of IE (see *ibid.*:1712-1716). However, it is logically impossible to assess the performance of an estimator when the true effects are unknown. If such a cross-model validation of IE for a specific empirical dataset were to show that IE yields reasonable estimates, this can only depend on having selected examples that are consistent with the IE's constraint. Therefore, cross-model comparisons using empirical data are not an appropriate method to validate IE.

debated the application scope of IE in empirical research. As Fu and associates (2011) suggested, “the important statistical issue about APC modeling is how to identify the trend that helps to resolve the real-world problem for a given APC data set” (p455). So I examine whether, compared to CGLM, IE yields better (if not unbiased) estimates of the true age, period, and cohort patterns that may be observed in empirical research.

IE’s developers provided simulations in which IE estimates are closer to the true age, period, and cohort effects than CGLM results. This, they argued, supports their conclusion that IE has clear advantages over CGLM. However, as noted above, the true age, period, and cohort effects in Yang et al.’s (2008) simulation in fact approximately satisfy the LC that IE imposes ($b \cdot b_0 = -0.036$)⁸. For age, period, and cohort effects that do not satisfy IE’s implicit constraint, IE will not necessarily perform better than CGLM and may perform much worse. Thus, IE is no better than CGLM because the restriction that IE imposes is essentially no different from the constraints assumed in CGLM.

To illustrate, I show simulations, as Yang and colleagues did, to compare the CGLM and IE estimates. However, here the data-generating mechanisms satisfy the constraint assumed by CGLM but not the constraint assumed by IE. Moreover, I simulate data from four models that embody specific social theories and thus conform to empirical

⁸ While Yang and colleagues correctly pointed out that IE estimates the *projection* of the “true” effects onto the non-null space, they compared IE estimates to the “true” *parameters*, not to the projection (see *ibid.*:1718-1722). This is key, because the true parameter vector can be very different from its projection onto the non-null space (the vector that IE actually estimates). That is, what IE actually estimates can be very different from the true APC effects if the true effects do not at least approximately satisfy the LC implicit in IE.

reality. The first dataset is simulated to represent the observation that overall health for adults deteriorates as they grow older, and that while recent development in health knowledge and technology have improved health conditions for the entire population, people born in more recent years are likely to be healthier than older cohorts. On the other hand, the demographic literature has also suggested that age, period, or cohort effects may not all exist (Alwin 1991; Winship and Harding 2008; Fabio et al. 2006; Preston and Wang 2006). Accordingly, the other three simulations approximate likely empirical situations where one of the three variables has little impact on the outcome variable.

Specifically, I fix the number of age groups at nine and periods at 50 in all of these simulations with little loss of generality. I then generate 1,000 datasets from each of the following four models:

$$y_{ij} \sim \text{Normal}\{10 + 2 \cdot \text{age}_i - 0.5 \cdot \text{age}_i^2 + 1 \cdot \text{period}_j - 0.015 \cdot \text{period}_j^2 + 0.15 \cdot \text{cohort}_{ij} + 0.03 \cdot \text{cohort}_{ij}^2, \sigma = 0.1\} \quad (30)$$

$$y_{ij} \sim \text{Normal}\{10 + 1 \cdot \text{period}_j - 0.015 \cdot \text{period}_j^2 + 0.15 \cdot \text{cohort}_{ij} + 0.03 \cdot \text{cohort}_{ij}^2, \sigma = 0.1\} \quad (31)$$

$$y_{ij} \sim \text{Normal}\{10 + 2 \cdot \text{age}_i - 0.5 \cdot \text{age}_i^2 + 0.15 \cdot \text{cohort}_{ij} + 0.03 \cdot \text{cohort}_{ij}^2, \sigma = 0.1\} \quad (32)$$

$$y_{ij} \sim \text{Normal}\{10 + 2 \cdot \text{age}_i - 0.5 \cdot \text{age}_i^2 + 1 \cdot \text{period}_j - 0.015 \cdot \text{period}_j^2, \sigma = 0.1\} \quad (33)$$

For instance, in equation (30), the outcomes for people with age i in period j are normally distributed with mean $(10 + 2 \cdot age_i - 0.5 \cdot age_i^2 + 1 \cdot period_j - 0.015 \cdot period_j^2 + 0.15 \cdot cohort_{ij} + 0.03 \cdot cohort_{ij}^2)$ and standard deviation $\sigma = 0.1$. In equations (31), (32), and (33), one of the age, period, and cohort effects is not present while the effects for the other two variables are the same as in equation (30). Note that none of these models satisfies IE's constraint; specifically, for the first model, $b \cdot b_0 = 115.01$; for the second, third, and last model, $b \cdot b_0 = 115.72, 130.41$ and 16.12 , respectively.

Figure 2.1 compares, for the simulated data from the four models, IE estimates and CGLM estimates using two different constraints. The IE estimates, averaged over 1,000 datasets, are largely away from the true effects for all models because for all four models, the constraint that IE assumes is not satisfied. For example, in Scenario 3 in Figure 2.1, when there is no period effect in the data-generating mechanism (32), the IE estimates suggest a substantially positive period trend on top of inaccurate estimates for age and cohort effects. In contrast, the CGLM assuming equal age effects for the first and third age groups produces close estimates for all four models. It is equally important to note that the performance of the CGLM estimator also depends on whether its assumption approximates the truth. For instance, in Scenario 4, whereas the CGLM that assumes equal age effects for the first and third group yields good estimates, the same method with a different constraint, i.e., the age effects are the same for the first and second groups, results in biased estimates.

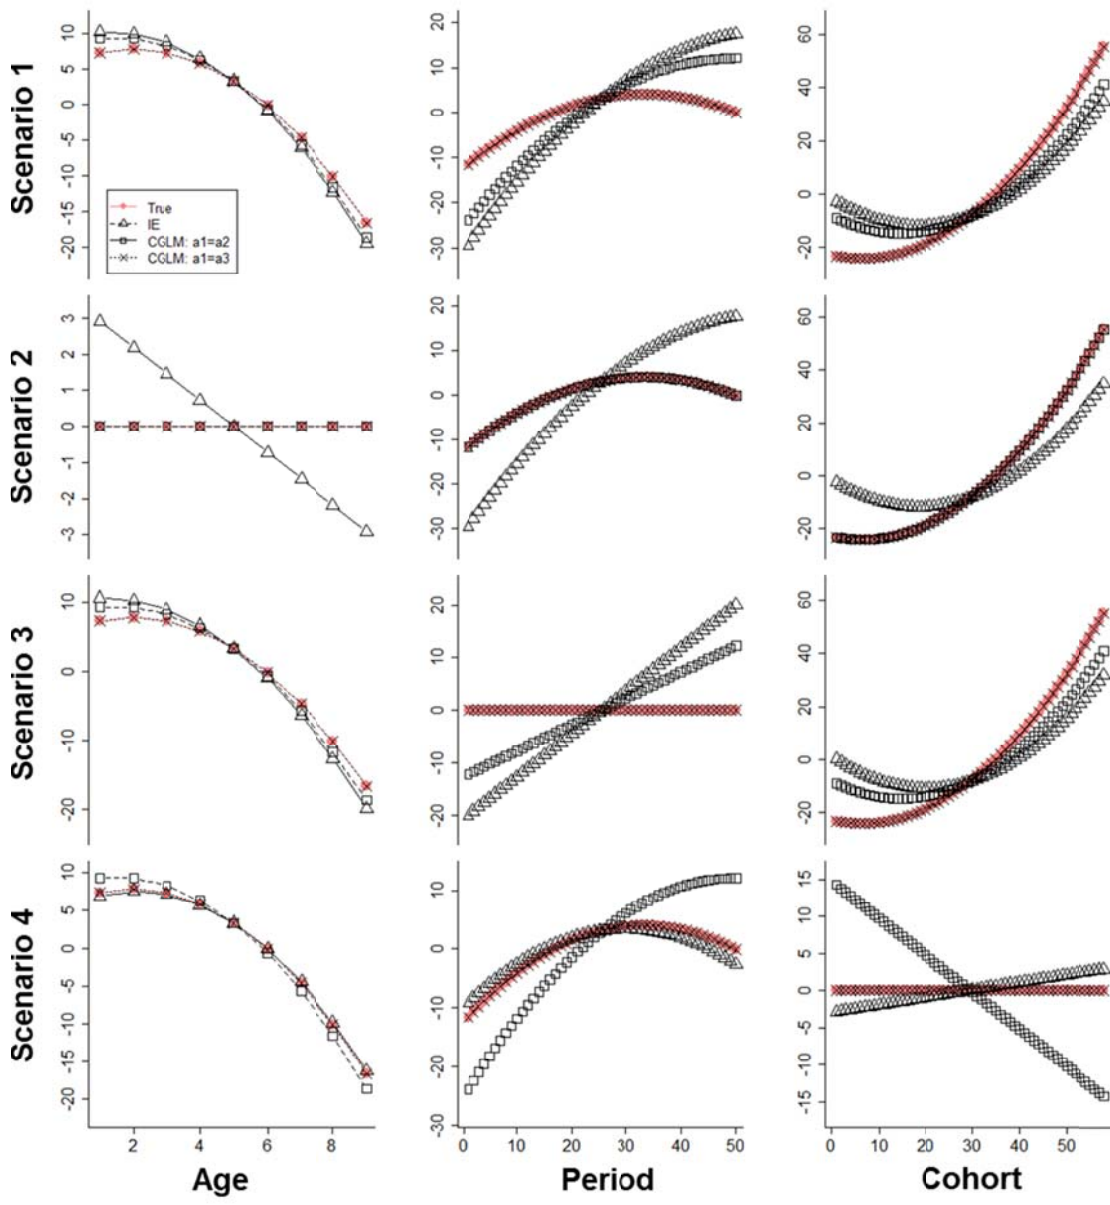


Figure 2.1. Simulation Results: IE vs. CGLM estimates for age, period, and cohort effects.

In sum, it must be concluded that a) if there is *a priori* information or theoretical justification, the constrained solution that corresponds to such information (e.g., CGLM estimates assuming equal effects for the first and third age groups in data-generating functions (30) to (33)) will yield better estimates than IE, and b) without such *a priori* knowledge, IE is not necessarily better than other constrained estimators including CGLM. Without such knowledge, neither IE nor CGLM results are valid.

Conclusion and Discussion

In this chapter, I focus on the Intrinsic Estimator (IE), a statistical method intended to separate the independent effects of age, period, and cohort on various outcomes. I have discussed the nature and application scope of IE theoretically and illustrated it with simulated data. This chapter has shown that IE assumes a specific constraint on the linear age, period, and cohort effects. This assumption not only depends on the number of age, period, and cohort groups, but also is extremely difficult, if not impossible, to verify in empirical research. This feature of IE is no different from the constraint assumed in CGLM except that the CGLM constraint does not change automatically as the numbers of age, period, and cohort groups change. The conclusion is that IE is not an unbiased or consistent estimator of the “true” age, period, and cohort effects. Therefore, for demographers and social scientists whose goal is to understand the “true, simultaneously independent effects” of age, period, and cohort, IE’s strategy of circumventing the identification problem can yield biased and potentially misleading estimates.

There is no doubt that Yang and associates have revitalized APC research and inspired many scholars. However, IE is nothing new in APC analysis. Kupper and his

colleagues introduced the IE solution to APC analysts, calling this solution as the Principal Component Estimator (PCE) (Kupper et al. 1983:2795-2797). As O'Brien (2011a:420) noted, such an estimator "produces coefficients identical to those of the recently introduced intrinsic estimator." However, instead of concluding that IE is preferable to CGLM, Kupper et al. (1983) clearly stated that PCE (that is, IE) "could lead to more bias than the use of some other constraints" (p2797). As a result, Kupper and associates did not advocate PCE/IE as a general solution, then or subsequently.

Generally speaking, PCE/IE or any other constrained estimator provides just *one* possible solution from the infinite number of solutions for an under-determined problem (i.e., the rank deficiency problem in APC analysis). That said, the PCE/IE solution should not be regarded as *the true solution* or *the uniquely preferred solution* without theoretical justification. In fact, the statistical literature has recognized a variety of constrained estimators including other types of generalized inverse solutions. It is important for demographers and sociologists to understand that the PCE/IE estimates are not necessarily better (i.e., closer to the true parameters) than other constrained estimators.

What should well-intentioned researchers, who wish to investigate the age, period, and cohort patterns, do? On the one hand, several alternative methods have been developed, some of which are more theoretically driven, taking external information into account⁹, while others are statistical approaches¹⁰. Although each of these methods has advantages and limitations and a thorough examination is a topic for future research, I

⁹ Examples include "Age-Period-Cohort Characteristic Models" developed by O'Brien (2000) and the "mechanism-based approach" proposed by Winship and Harding (2008).

¹⁰ E.g., "Cross-Classified Random Effects Models" created by Yang and Land (2006, 2008).

caution that purely statistical techniques are unlikely to yield accurate estimates. The methodological problem of IE and its non-trivial implications for empirical research identified in this paper are not unique to IE. The biostatistics literature shows that use of the APC model (1), regardless of estimation technique, precludes valid estimation as well as meaningful interpretations of the linear components of age, period, and cohort effects (see, e.g., Holford 1983; Kupper et al. 1985). Therefore, my position is to encourage development of APC models that are informed by social theories and thus different from model (1) in basic structure.

On the other hand, although the statistical difficulty in quantifying independent effects of age, period, and cohort was recognized long ago, decades of effort has only resulted in unsatisfactory solutions. Thus it is not unreasonable to ask: Is this unusual challenge suggesting a problem that is not statistical but theoretical in nature? In other words, is the identification problem pointing to a more fundamental problem in the theoretical framework of APC analysis? Should the answers to these questions be positive, the identification problem inherent in model (1) “is a blessing for social science” (Heckman and Robb, 1985:144) because it warns scientists that they want something — a general statistical decomposition of data — for nothing.

Appendix

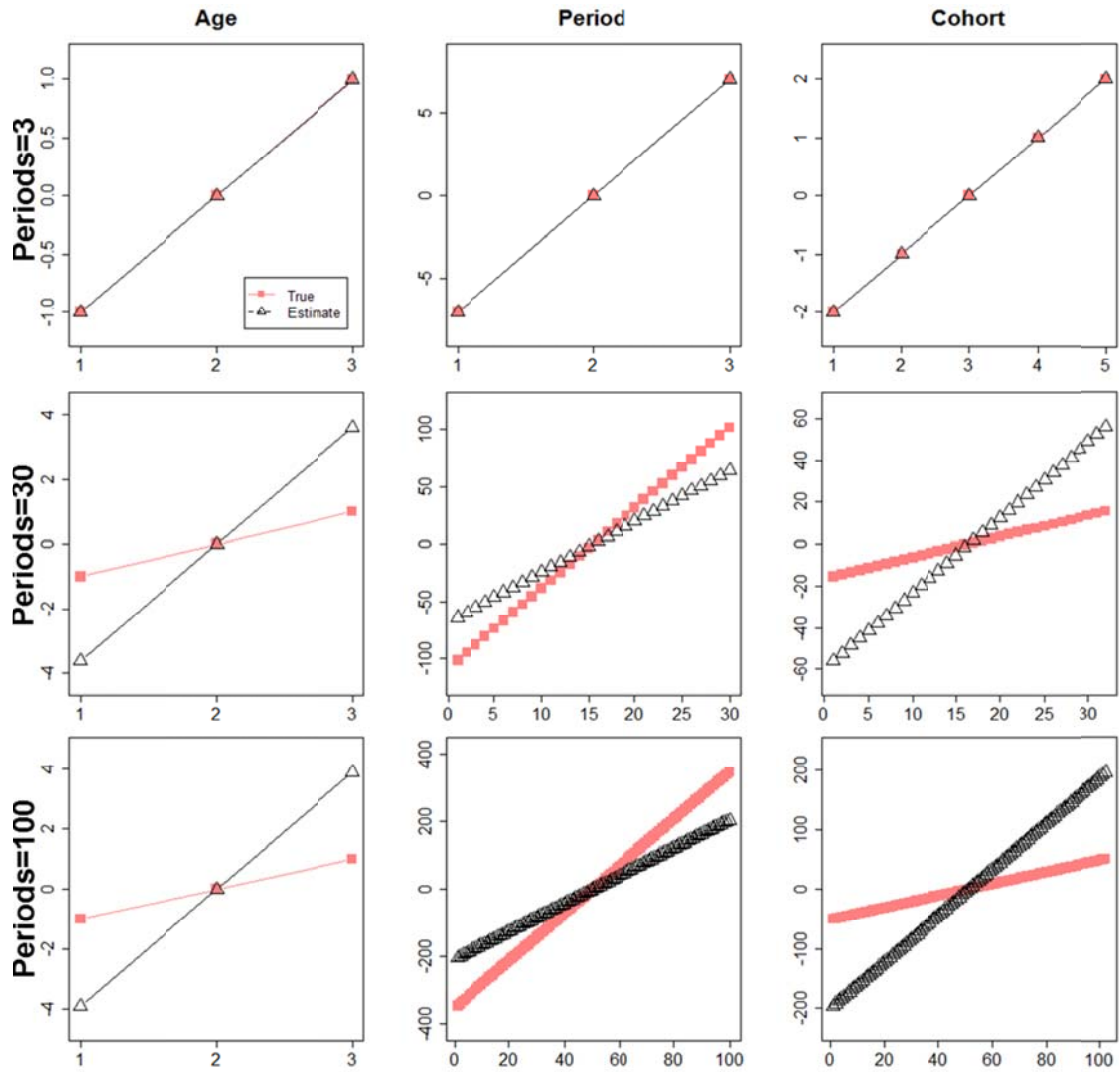


Figure 2.2. Simulation Results: Inconsistent IE estimates.

CHAPTER 3 : CRITIQUE OF CROSS-CLASSIFIED MODELS

Among the various methods proposed to address the identification problem, cross-classified APC models¹¹, including Cross-Classified Fixed Effects Models (CCFEM) and Cross-Classified Random Effects Models (CCREM), developed by Yang and Land (2006, 2008) are probably the most popular and most widely used APC method. These scholars argued that the cross-classified method, using unequal interval widths for age, period, and cohort groups, breaks the exact linear dependency and thus solves the identification problem and provides valid estimates for age, period, and cohort effects (Yang and Land 2008:302-303). Since its introduction, this method has been adopted by researchers addressing important substantive issues including mortality, happiness, religious activities, and obesity (see, e.g., Pampel and Hunter 2012; Masters 2012; Masters et al. 2012; Schwadel 2010; Reither et al. 2009; Yang 2008).

However, this chapter shows that the age-period-cohort linear dependency gives rise to an inherent identification problem that cannot be solved by changing interval widths or the model setup. In other words, although the unequal interval-width age, period, and cohort groups used in the cross-classified APC models are not exactly linearly dependent in the same simple way as in equal interval-width APC models, it is still true that the independent effects of age, period, and cohort cannot be estimated without constraints. The cross-classified strategy implicitly uses *multiple* constraints to choose one estimate from the infinite number of possible estimates. The resulting estimates can be highly biased and substantive conclusions may be misleading when any

¹¹ sometimes called Hierarchical APC models.

of the multiple constraints is not satisfied by the “true” age, period, or cohort effects. Because external information verifying these constraints is scant or non-existent in empirical research, neither CCFEM nor CCREM results should be deemed valid without an explicit rationale justifying these implicit constraints.

This chapter contributes to the literature in two ways. First, it enriches the theoretical discussion about the identification problem in APC analysis. APC literature focuses on APC models with equal-width intervals, and methodological problems in APC models with unequal-width intervals are not yet fully understood. Drawing on two articles, I show that the identification problem remains for cross-classified APC models. Second, despite the popularity of the cross-classified models, the constraints in these models and their implications have not been thoroughly examined or understood. This chapter fills this gap by identifying the multiple constraints implicit in CCFEM and CCREM, illustrating the implications of these constraints using simulations, and examining the effect on estimates in situations pertinent to sociologists. By making the multiple constraints implicit in CCFEM and CCREM explicit to readers, this chapter should help researchers better understand the nature of these models and make informed decisions in choosing an APC estimator.

This chapter is organized as follows. To contextualize the methods under investigation, I briefly review the identification problem in APC analysis, cross-classified APC models, and other constrained estimators. Yang and Land (henceforth Y&L) have claimed that cross-classified methods break the identification problem; I show that the identification problem remains intact and the cross-classified approach is simply a

constrained estimator that places multiple constraints on the parameter vector. These sections are fairly technical; readers may skip them on a first reading and go directly to “Simulation Results”.

The “Simulation Results” section uses simulation experiments, where the true age, period, and cohort effects are known, to assess the performance of CCFEM and CREM under various circumstances. If cross-classified methods solve the identification problem as Y&L claim, CCFEM and CREM estimates should be close to the true age, period, and cohort effects, which are known because I am simulating data from known mechanisms. Yet our simulations show that CCFEM and CREM estimates can be highly biased, i.e., systematically very different from the true age, period, and cohort effects. The conclusion is that CCFEM and CREM do not in fact solve the identification problem and empirical researchers should be aware of the non-trivial implications of the constraints implicit in CCFEM and CREM.

The Identification Problem and Cross-Classified APC Models

To separate independent age (A), period (P), and cohort (C) effects, researchers have conventionally used an analysis of variance (ANOVA) model:

$$g(E(Y_{ij})) = \mu + \alpha_i + \beta_j + \gamma_k, \quad (1)$$

for age groups $i = 1, 2, \dots, a$, periods $j = 1, 2, \dots, p$, and cohorts $k = j - i + a = 1, 2, \dots, a + p - 1$, where $\sum_{i=1}^a \alpha_i = \sum_{j=1}^p \beta_j = \sum_{k=1}^{a+p-1} \gamma_k = 0$. $E(Y_{ij})$ denotes the expected value of the outcome Y for the i th age group in the j th period of time; g is the “link function” linking the expected value of Y_{ij} to the effects; α_i denotes the mean difference from the global mean μ associated with, or the effect of, the i th age category;

β_j denotes the effect of the j th period; γ_k denotes the effect of the k th cohort. The conventional ANOVA constraint applies, under which the sum of coefficients for each effect is set to zero. In using model (1), researchers usually assume the interval widths for A, P, and C groups are equal; if the interval width is five years, then the age groups could be 0-4, 5-9, ..., the periods 1980-1984, 1985-1989, ..., and the cohorts, 1890-1894, 1895-1899,

When the outcome Y_{ij} is treated as a normal (Gaussian) random variable, $g(E(Y_{ij}))$ is usually $E(Y_{ij})$, and model (1) can also be written in a generic regression fashion:

$$Y = Xb + \varepsilon, \quad (2)$$

where Y is a vector of outcomes; X is the design matrix; b denotes a parameter vector whose elements are the intercept and the effects of the A, P, and C groups; and ε denotes random errors. Then estimated A, P, and C effects can be obtained using the ordinary least squares method:

$$\hat{b} = (X^T X)^{-1} X^T Y. \quad (3)$$

However, it is well-known that the inverse $(X^T X)^{-1}$ does not exist because of the age-period-cohort linear dependency; the design matrix X has rank one less than full rank and thus the parameter vector b is inestimable without external information. This is the identification problem in APC analysis: infinitely many \hat{b} have *identical* fits to the data. Note that this identification problem is over and above the identifiability problem addressed by the usual sum-to-zero constraint or by omitting one category for each effect.

Various methods have been proposed to address the identification problem, many of which impose a constraint on the parameter vector b . Constrained estimators include the Constrained Generalized Linear Model (CGLM) proposed by Mason et al. (1973) and the Intrinsic Estimator (IE) introduced by Fu (2000) and Yang and associates (2008). Constrained estimators have been criticized for relying on external information to specify constraints when such information generally does not exist or cannot be verified (Rodgers 1982a, 1982b, Glenn 2005, Luo 2013a, Luo et al. 2014, O'Brien 2011a, 2011b).

Y&L (2006, 2008) proposed a new method that apparently involves no such constraints, called the cross-classified APC method. Specifically, for individual-level data where each person's exact age and cohort membership and the time of data collection are known, Y&L proposed creating A, P, and C groups with different interval widths so the exact linear dependency between these variables disappears. For example, when the individual-level data are represented using one-year age groups, two-year periods, and five-year cohorts, one cannot determine a person's age from the period and that person's cohort.

These authors then proposed analyzing these cross-classified data treating period and cohort effects as either fixed effects or random effects, termed CCFEM and CCREM respectively. The matrix form of CCFEM and CCREM is:

$$Y = W\beta + Zu + \varepsilon, \quad (4)$$

where W and Z are the design matrices for the age effects and for the period and cohort effects, respectively. In CCFEM, β and u are just regression coefficients, so fundamentally CCFEM is no different from the classic APC accounting model. In

CCREM, each element of the random effects is modeled as normally distributed around a zero mean, i.e., $u \sim N(0, G)$ and $\varepsilon \sim N(0, \sigma^2)$, which has the effect of “shrinking” the estimates of these coefficients towards zero (Hodges 2013). I explain this point in the “Conclusion and Discussion” section.

CCFEM and CCREM have been adopted by many researchers because it appears that unlike conventional APC estimators, they do not require difficult-to-verify assumptions to produce estimates of the A, P, and C effects. However, in the following sections, I show that in fact CCFEM and CCREM impose *multiple* constraints on the parameter vector (β, u) . This feature of CCFEM and CCREM is no different from the constraints in methods such as CGLM and IE. I will identify the constraint implicit in CCFEM and CCREM and then use simulations to illustrate the effect of applying this method in empirical research.

The Multiple Block Constraints (MBCs) that CCFEM and CCREM Implicitly Assumes

Although most APC literature centers on the estimation problem for models with equal-width intervals, models with unequal-width intervals have been discussed since the 1970s. This literature seems to have made little impact, because many researchers apparently believe that the identification problem vanishes with unequal-width intervals. To clarify this misunderstanding, I first review two key articles addressing identification with unequal-width intervals and then show how the cross-classified approach chooses an estimate by imposing constraints.

Fienberg and Mason (1979) were the first to characterize identification in APC models with unequal-width intervals, showing that as with equal-width intervals, only nonlinear components of the A, P, and C effects are estimable. Specifically, they demonstrated that “[n]ot only does this identification problem remain, but there is an additional age-cohort dependency” using an example where the period intervals were as wide as age intervals (1979:38). They found that, as in the equal-width case, a restriction on the linear effects is necessary to identify model (1). But, unlike the simpler case with equal-width intervals, “[n]ot just one but two restrictions are necessary to identify all the effects” and “the two restrictions cannot be placed arbitrarily” (1979:39), that is, the constraints have to be chosen carefully because many pairs of constraints do not identify the model.

Holford (1983, 2006) went beyond the foregoing, noting that using unequally spaced intervals leads to problems specific to the age and cohort effects in addition to the classic identification problem in APC model (1). In other words, using unequal-width intervals retains the old difficulty in identifying model (1) and creates a new one.

This literature implies that CCFEM and CCREM must be imposing two or more constraints to estimate A, P, and C effects. What are those constraints? Suppose that each person’s exact age and birth cohort are known. To apply CCFEM and CCREM, reorganize the data so there are a age groups, p periods, and c cohorts, and the A, P, and C intervals have widths l , m , and n years, respectively. To estimate independent A, P, and C effects, CCFEM and CCREM assume the l year-specific age effects within each of the a aggregated age groups are equal, the m year-specific period effects within each of the p

period groups are equal, and the n year-specific cohort effects in each of the c cohort groups are also equal. That is,

$$\alpha_{i1} = \alpha_{i2} = \dots = \alpha_{in}, \quad (5)$$

$$\beta_{j1} = \beta_{j2} = \dots = \beta_{jm}, \quad (6)$$

and

$$\gamma_{k1} = \gamma_{k2} = \dots = \gamma_{kn}, \quad (7)$$

for age groups $i = 1, 2, \dots, a$, periods $j = 1, 2, \dots, p$, and cohorts $k = 1, 2, \dots, c$. Thus the assumption implicit in CCFEM and CCREM is that the true A, P, and C effects satisfy equations (5) through (7). I call these constraints the cross-classified APC models' "Multiple Block Constraints" (MBCs). If, in fact, the true effects do not satisfy *any* of these equations, then the MBCs implicit in CCFEM or CCREM are incorrect. As we will see, failure to satisfy these assumptions can have large counterintuitive effects on estimates.

For illustration, suppose I have three age categories, three periods, and thus five cohorts so $a = 3, p = 3$ and $c = 5$. For each age-by-period combination, there is one observation. Table 3.1 gives the design matrix X for CCFEM using the sum-to-zero parameterization. With the last group omitted for each effect, the design matrix X consists of $\{1 + a - 1 + p - 1 + a + p - 2\}$ columns. If more A, P, and C categories or more than one observation in each age-by-period combination are included, Table 3.1 can be expanded, with each new category or observation corresponding to an additional column or row in X . Age or period can be treated as a continuous predictor, which does not affect the discussion below.

Table 3.1. Design Matrix of CCFEM for Equal Interval-Width Age, Period, and Cohort Groups.

Intercept	Age		Period		Cohort			
	α_1	α_2	β_1	β_2	γ_1	γ_2	γ_3	γ_4
1	1	0	1	0	0	0	1	0
1	1	0	0	1	0	0	0	1
1	1	0	-1	-1	-1	-1	-1	-1
1	0	1	1	0	0	1	0	0
1	0	1	0	1	0	0	1	0
1	0	1	-1	-1	0	0	0	1
1	-1	-1	1	0	1	0	0	0
1	-1	-1	0	1	0	1	0	0
1	-1	-1	-1	-1	0	0	1	0

The design matrix in Table 3.1 has rank one less than full rank because Cohort = Period – Age. That is, the number of linearly independent columns is one less than the number of columns. One way to make model (1) estimable, i.e., to modify the design matrix in Table 3.1 so that it has full rank, is to constrain pairs of adjacent cohorts to have equal effects, that is, to assume $\gamma_1 = \gamma_2$ and $\gamma_3 = \gamma_4$. Let $\gamma'_1 = \gamma_1 = \gamma_2$ and $\gamma'_2 = \gamma_3 = \gamma_4$. With this assumption, the design matrix of model (1) has full rank and can be written as in Table 3.2.

The design matrix in Table 3.2 is identical to that for CCFEM with three one-year age groups, three one-year periods, and three two-year cohorts, that is, $l = 1, m = 1$ and $n = 2$. In other words, applying CCFEM to this example using two-year cohorts is equivalent to assuming the two one-year cohorts within each of the two-year cohort

groups have identical effects. If age or period is also grouped into multi-year categories, CCFEM places additional constraints on the age or period effects, forcing the single-year age or period groups within each multi-year group to have equal effects.

Table 3.2. Design Matrix of CCFEM for Equal Interval-Width Age, Period, and Cohort Groups by Constraining Two Adjacent Cohorts to Be Equal.

Intercept	Age		Period		Cohort	
	α_1	α_2	β_1	β_2	γ'_1	γ'_2
1	1	0	1	0	0	1
1	1	0	0	1	0	1
1	1	0	-1	-1	-1	-1
1	0	1	1	0	1	0
1	0	1	0	1	0	1
1	0	1	-1	-1	0	1
1	-1	-1	1	0	1	0
1	-1	-1	0	1	1	0
1	-1	-1	-1	-1	0	1

Note. The design matrix shown in this table is identical to the design matrix for the fixed effects model with 1-year age groups, 1-year periods, and 2-year cohorts.

A similar exposition applies to CCREM with age treated as a fixed effect and period and cohort as random effects. Table 3.3 gives the design matrices for the fixed and random effects, equation (4)'s W and Z , using equal-width intervals for the data above. Like the CCFEM design matrix in Table 3.1, the CCREM design matrix in Table 3.3 does not have full rank. To estimate A, P, and C effects, one can constrain two adjacent cohorts to have equal effects, i.e., to constrain $\gamma_1 = \gamma_2$ and $\gamma_3 = \gamma_4$. Let $\gamma'_1 = \gamma_1 = \gamma_2$

and $\gamma'_2 = \gamma_3 = \gamma_4$. Then the CCREM design matrix has full rank and can be rewritten as in Table 3.4. The design matrix in Table 3.4 is identical to that for CCREM using single-year age groups, single-year periods, and two-year cohorts. In other words, applying CCREM to the data set above using two-year cohorts is equivalent to assuming that the two one-year cohorts within each of the two-year cohort groups have identical effects.

Table 3.3. Design Matrix of CCREM for Equal Interval-Width Age, Period, and Cohort Groups.

Intercept	Fixed Effects		Random Effects							
	Age		Period			Cohort				
	α_1	α_2	β_1	β_2	β_3	γ_1	γ_2	γ_3	γ_4	γ_5
1	1	0	1	0	0	0	0	1	0	0
1	1	0	0	1	0	0	0	0	1	0
1	1	0	0	0	1	0	0	0	0	1
1	0	1	1	0	0	0	1	0	0	0
1	0	1	0	1	0	0	0	1	0	0
1	0	1	0	0	1	0	0	0	1	0
1	-1	-1	1	0	0	1	0	0	0	0
1	-1	-1	0	1	0	0	1	0	0	0
1	-1	-1	0	0	1	0	0	1	0	0

Note. The design matrix shown in this table is no different from that for the fixed effects model in Table 3.1 except that the random effects part of this design matrix includes columns for β_3 and γ_5 that are omitted from Table 3.1.

Table 3.4. Design Matrix of CCREM for Equal Interval-Width Age, Period, and Cohort Groups by Constraining Two Adjacent Cohorts to Be Equal.

Intercept	Fixed Effects		Random Effects					
	Age		Period			Cohort		
	α_1	α_2	β_1	β_2	β_3	γ'_1	γ'_2	γ'_3
1	1	0	1	0	0	0	1	0
1	1	0	0	1	0	0	1	0
1	1	0	0	0	1	0	0	1
1	0	1	1	0	0	1	0	0
1	0	1	0	1	0	0	1	0
1	0	1	0	0	1	0	1	0
1	-1	-1	1	0	0	1	0	0
1	-1	-1	0	1	0	1	0	0
1	-1	-1	0	0	1	0	1	0

Note. The design matrix shown in this table is no different from that for the fixed effects model in Table 3.2 except that the random effects part of this design matrix includes columns for β_3 and γ'_3 that were omitted from Table 3.2. It is also identical to the design matrix for random effects models with 1-year age groups, 1-year periods, and 2-year cohorts.

This exposition shows that the MBCs implicit in CCFEM and CCREM depend on the design matrix, i.e., on the number of A, P, and C groups, when their interval widths change. For example, if I use a three-year interval for cohort groups, so I now have three one-year age groups, three one-year periods, and two three-year cohorts, then following the derivation above, the MBCs implied by CCFEM and CCREM are $\gamma_1 = \gamma_2 = \gamma_3$ and $\gamma_4 = \gamma_5$. Compare these with the constraints for the model with three two-year cohorts: increasing the interval width for cohort changes the constraints on cohort effects so that

true effects satisfying the MBCs with two-year cohorts do not satisfy the MBCs with three-year cohorts. Readers can verify that increasing or decreasing the width of age or period intervals also alters the MBCs that CCFEM and CCREM assume.

These examples demonstrate that CCFEM and CCREM rely on a constraint like CGLM does but unlike CGLM, where the constraint (e.g., equal effects for the first two age groups) is explicit, the MBCs used by CCFEM and CCREM are implicit and change depending on the width of the A, P, and C groups. Surprisingly, these MBCs can have large and non-intuitive consequences for estimating A, P, and C effects, as I show in next section.

In sum, CCFEM or CCREM are constrained estimators that place *multiple* equality constraints on the A, P, and C effects. Because a constraint determines estimates in the APC problem, a constrained method produces valid estimates only when its assumption approximates the true structure of the data under investigation. For any coefficient-constraint approach, “the choice of constraint is the crucial determinant of the accuracy in the estimated age, period, and cohort effects” (Kupper et al. 1985:822). In particular, only when the true effects are equal within each of the multi-year A, P, and C categories can CCFEM and CCREM yield accurate estimates. If the true A, P, and C effects do not satisfy this assumption, then estimates from CCFEM and CCREM may be highly distorted. However, researchers usually have no *a priori* knowledge about the relationship among A, P and C effects. Therefore, unfortunately, they cannot know when CCFEM and CCREM can be applied; in this respect, these methods are no better than any other constrained method. To the extent that CCFEM and CCREM impose more than

one constraint on the A, P, and C effects, it means that they require *more*, not less, side information than CGLM.

Simulations: Effect of Multiple Block Constraints (MBCs) and Application Scope of CCFEM and CCREM

This section's simulation results are intended to be a straightforward illustration of the preceding discussion about the Multiple Block Constraints (MBCs) implicit in CCFEM and CCREM. I first use simple simulations to illustrate the implications of the MBCs and then use simulated data embodying specific social theories to examine how the MBCs affect estimates of the true A, P, and C trends in social research¹². I show how the CCFEM and CCREM estimates differ from the “true” effects of the data-generating mechanism and depend on the MBCs, i.e., different intervals widths for A, P, and C groups.

To illustrate the implications of the MBCs, I first simulate normally distributed individual-level data as follows. For persons at age i in period j , the mean response is $0 + \alpha_i + \beta_j + \gamma_{ij}$ and the standard deviation of error ε equals 1¹³. The number of age and period groups is fixed at three each in the micro data, so there are three age groups, three

¹² Y&L claimed to use empirical data, where the true effects are unknown, to validate cross-classified methods (see Yang et al. 2008: 1712-1716). However, it is logically impossible to assess an estimator's ability to estimate true effects when the true effects are unknown.

¹³ I do not simulate each new hypothetical dataset by making a new draw from the random effects in CCREM for each data set, because A, P, and C effects cannot be random effects as understood by Scheffé (1959). Rather, these effects are of interest because they are presumed to have some true values, which are unknown to us and which empirical studies are intended to estimate, and a model with random effects is just another way to estimate these unknowns.

periods, and five cohorts at the individual level. I consider three sets of true α_i , β_j , and γ_{ij} , shown in Table 3.5. For each selection of true α_i , β_j , and γ_{ij} , I simulated 10,000 individuals to give one simulated data set. I repeated the process 100 times to give 100 data sets, each with 10,000 observations. To avoid meaningless artifacts, each combination of age and period has the same number of observations.

Table 3.5. Examples of True Age, Period, and Cohort Effects.

Selection	Age			Period			Cohort				
	α_1	α_2	α_3	β_1	β_2	β_3	γ_1	γ_2	γ_3	γ_4	γ_5
1	-1	0	1	-1	0	1	-1.5	-1.5	0	0	1.5
2	-1	0	1	-1	0	1	-3	-1.5	0	1.5	3
3	-1	0	1	-1	0	1	2	1	0	-1	-2

Note. The different rows (“selections”) are different true age, period, and cohort effects.

To apply cross-classified APC models to the simulated data, I created one-year age groups, one-year periods and two-year cohorts so the resulting design matrix has full rank. I then include the three aggregated age, period, and cohort variables as fixed effects in CCFEM. CCREM is implemented using fixed age effects and random period and cohort effects.

Note that for Table 3.5’s Selection #1 (i.e., first true A, P, and C effects), the A, P, and C effects satisfy equations (5) through (7), the MBCs implicit in CCFEM and CCREM. In contrast, for Selections #2 and #3, the cohort effects do not satisfy the MBCs implicit in the cross-classified APC model.

Table 3.6 shows the CCFEM and CCREM estimates, averaged over the 100 simulated data sets, for the three selections, i.e., sets of true α_i , β_j , and γ_{ij} , along with the true effects. The biases of CCFEM and CCREM are estimated by the difference {estimates averaged over 100 simulated data sets} minus {true values}. For the first selection, CCFEM and CCREM give good estimates because the true effects satisfy the MBCs implicit in the two methods. In contrast, for the second and third selections, CCFEM and CCREM give highly biased estimates because the true effects do not satisfy the MBCs. For the second selection, the estimated period effects show a downward trend; the true trend is upward. Similarly, for the third selection, the true age effects trend is upward, while the CCFEM and CCREM estimates suggest the outcome does not vary with age. For the second and third selections, the cohort effect estimates are effectively zero, while the true cohort effects have increasing and decreasing trends, respectively.

Table 3.6. Simulation Results: CCFEM and CCREM estimates for Table 3.5's three selections (sets of true age, period, and cohort effects).

		Selection #1			Selection #2			Selection #3		
		Truth	CCFEM	CCREM	Truth	CCFEM	CCREM	Truth	CCFEM	CCREM
Age	1	-1	-1.002	-0.999	-1	-2.500	-2.499	-1	-0.002	0.001
	2	0	0.001	-0.001	0	0.001	-0.002	0	0.000	-0.002
	3	1	1.001	1.000	1	2.499	2.501	1	0.002	0.000
Period	1	-1	-1.002	-0.999	-1	0.498	0.498	-1	-1.998	-2.002
	2	0	0.000	-0.001	0	0.000	0.000	0	0.000	0.002
	3	1	1.003	1.000	1	-0.498	-0.499	1	1.998	2.000
Cohort	1	-1.5	-1.503	-1.500	-3	-0.001	-0.001	2	0.004	0.000
	2	-1.5	-1.503	-1.500	-1.5	-0.001	-0.001	1	0.004	0.000
	3	0	-0.001	-0.001	0	0.000	-0.001	0	-0.001	0.002
	4	0	-0.001	-0.001	1.5	0.000	-0.001	-1	-0.001	0.002
	5	1.5	1.504	1.501	3	0.002	0.002	-2	-0.004	-0.002

The A, P, and C effects in the simulations above were purely linear; they were intended as simple examples that obviously obey or violate the constraints (the MBCs) implicit in CCFEM and CCREM. The simulated data that follow show nonlinear effects and are designed to mimic data patterns that would be expected in four scenarios implied by specific social theories. To test whether these theories reflect social realities, it is essential to estimate the A, P, and C effects without distortion. The first scenario is simulated to represent the observation that while body weight tends to increase with age and peak around age 60, obesity rates have increased in recent decades in the US (Ogden 2006, Flegal et al. 2002, Mokdad et al. 2001). At the same time, the fetal over-nutrition theory posits that increasing *in utero* exposure to maternal obesity may lead to an inter-generational increase in obesity (Cole et al. 2008, Gillman 2004). This argument would imply that the obesity epidemic should manifest as cohort effects, as each successive cohort is at higher risk for obesity. On the other hand, the epidemiological, sociological, and demographic literatures also suggest that A, P, or C effects may not all exist (Keyes et al. 2009, Fabio et al. 2006, Preston and Wang 2006, Raftery et al. 1995, Ryder 1965). Accordingly, the other three scenarios use simulated data representing situations in which A, P, and C respectively has little impact on the outcome.

Specifically, to simulate individual-level data, I fixed the number of age groups at 20 and periods at 20 in all of these simulations. For each age-and-period combination, I simulated 25 individuals, so the sample size for each data set is 10,000. I then generated 100 such individual-level data sets, each with 10,000 observations, from each of the

following four scenarios, in which the mean $E(Y_{ij})$ is a function of age, period, and cohort:

$$0.3 \cdot age_i - 0.01 \cdot age_i^2 - 0.04 \cdot period_j + 0.02 \cdot period_j^2 + 0.35 \cdot cohort_{ij} - 0.0015 \cdot cohort_{ij}^2 \quad (8)$$

$$0.04 \cdot period_j + 0.02 \cdot period_j^2 + 0.35 \cdot cohort_{ij} - 0.0015 \cdot cohort_{ij}^2 \quad (9)$$

$$0.3 \cdot age_i - 0.01 \cdot age_i^2 + 0.35 \cdot cohort_{ij} - 0.0015 \cdot cohort_{ij}^2 \quad (10)$$

$$0.3 \cdot age_i - 0.01 \cdot age_i^2 - 0.04 \cdot period_j + 0.02 \cdot period_j^2 \quad (11)$$

In each scenario, the data are normally distributed with these means and error standard deviation $\sigma = 1$. For instance, in Scenario 1, body weight outcomes for people in age i in period j are normally distributed with a mean as in equation (8) and standard deviation $\sigma = 1$. In Scenarios 2 through 4, with mean functions (9), (10), and (11), the A, P, and C effect respectively is not present while the other two effects are the same as in function (8).

To apply cross-classified APC models, I created one-year age groups, two-year periods and five-year cohorts as in Y&L (2008); the resulting design matrix has full rank. To examine whether the estimates depend on the width of the A, P, and C groups, I also analyzed the simulated data using two-year intervals for cohorts.

Figures 3.1 and 3.2 show the CCFEM and CCREM estimates for the simulated data from the four scenarios, using two- and five-year interval widths for cohorts, along with the true A, P, and C effects specified in functions (8) to (11). The CCFEM and CCREM estimates shown in these figures are averaged over 100 simulated data sets, so the bias of these methods is estimated as {estimates averaged over 100 simulated data

sets} minus {true values}. In Figures 3.1 and 3.2, the CCFEM and CCREM estimates suffer from two types of errors: errors in estimating the overall trends in A, P, and C effects, and errors in estimating the trends within each aggregate A, P, and C groups. For the overall A, P, and C trends, the CCFEM and CCREM estimates are largely different from the true effects because in all four scenarios, the MBCs that CCFEM and CCREM assume are not satisfied. For example, in Scenario 1 with substantial A, P, and C effects, CCFEM and CCREM estimate an age trend contrary to the true trend for the 15th and older age groups, and they estimate a U-shaped period trend, contrary to the increasing trend in the true effects. Similarly, although the true cohort effects show a strong positive trend, CCFEM and CCREM estimate hardly any cohort effect. In Scenarios 2, 3, and 4 which lack the A, P, and C effect respectively, CCFEM and CCREM estimate a large trend for the effect that is truly absent, while also estimating the other two effects inaccurately.

Figures 3.1 and 3.2 also show that CCFEM and CCREM estimates depend on the width chosen for the cohort intervals. This effect is most striking in Scenarios 3 and 4, in which either the period or cohort effect is, in truth, absent. Readers can verify that changing the interval width for the age or period groups can also alter the estimated A, P and/or C trends. Clearly, the cross-classified approach not only retains the identification problem inherent in APC model (1) but creates new difficulties. Therefore, in specific applications, the cross-classified method must be assumed to be biased, resulting in potentially misleading conclusions about the true A, P, and C effects unless proven otherwise.

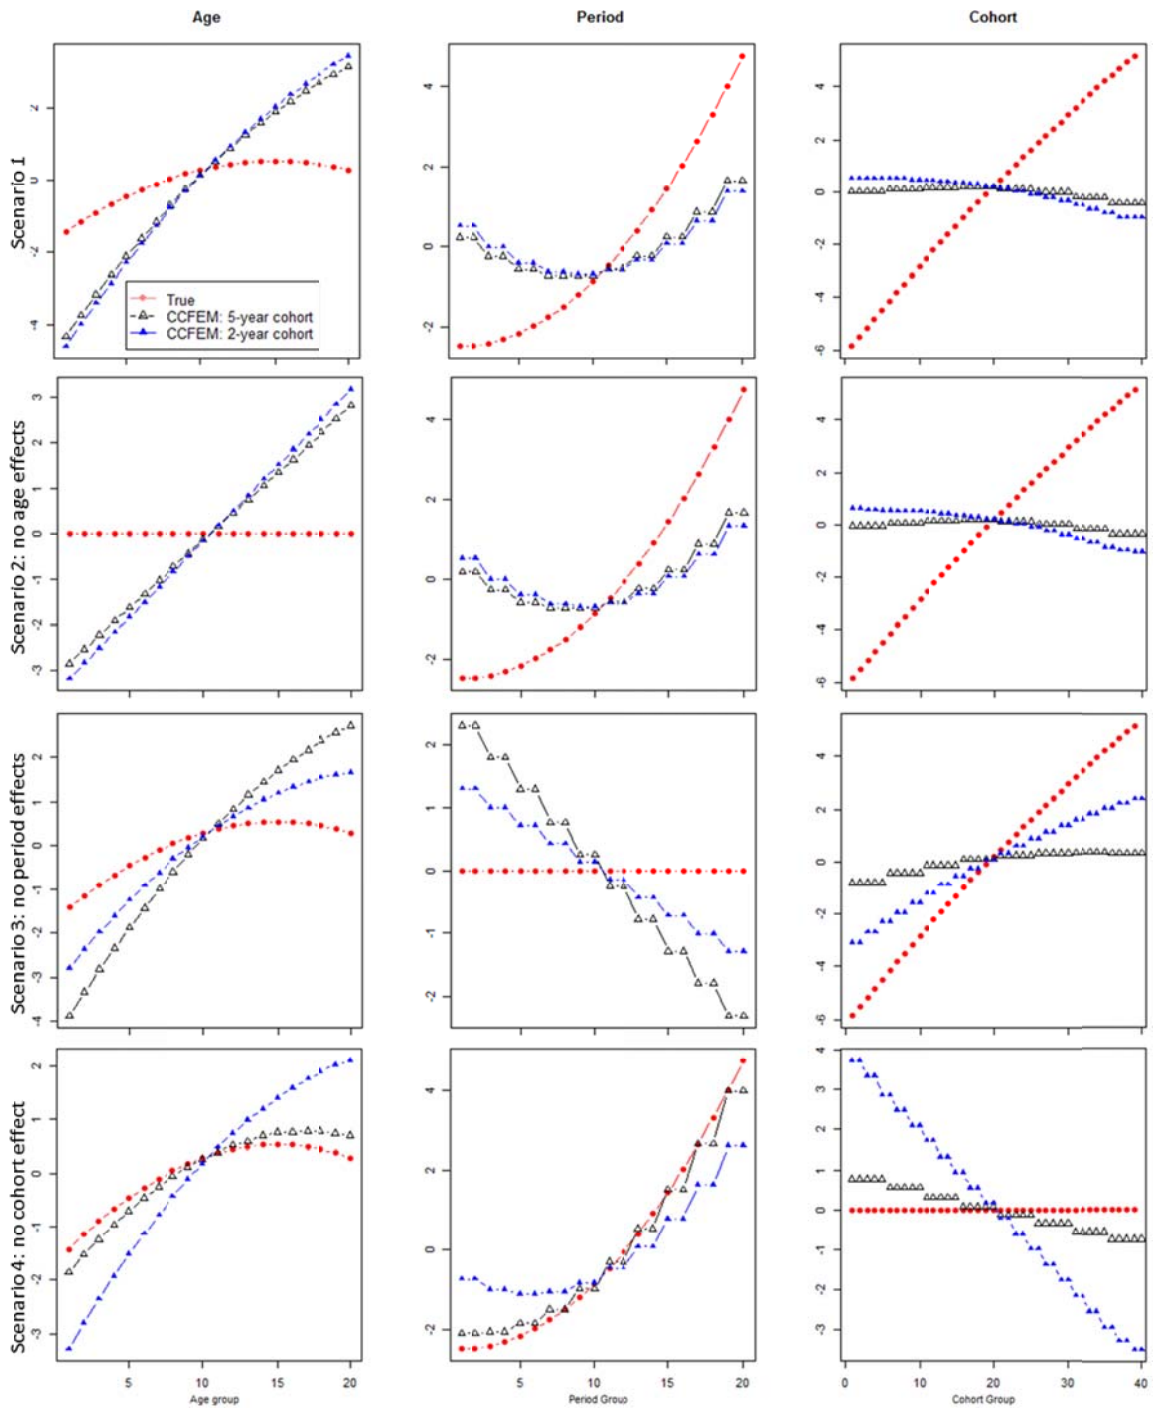


Figure 3.1. Simulation Results: CCFEM estimates using two- and five-year cohort groups for data from four scenarios.

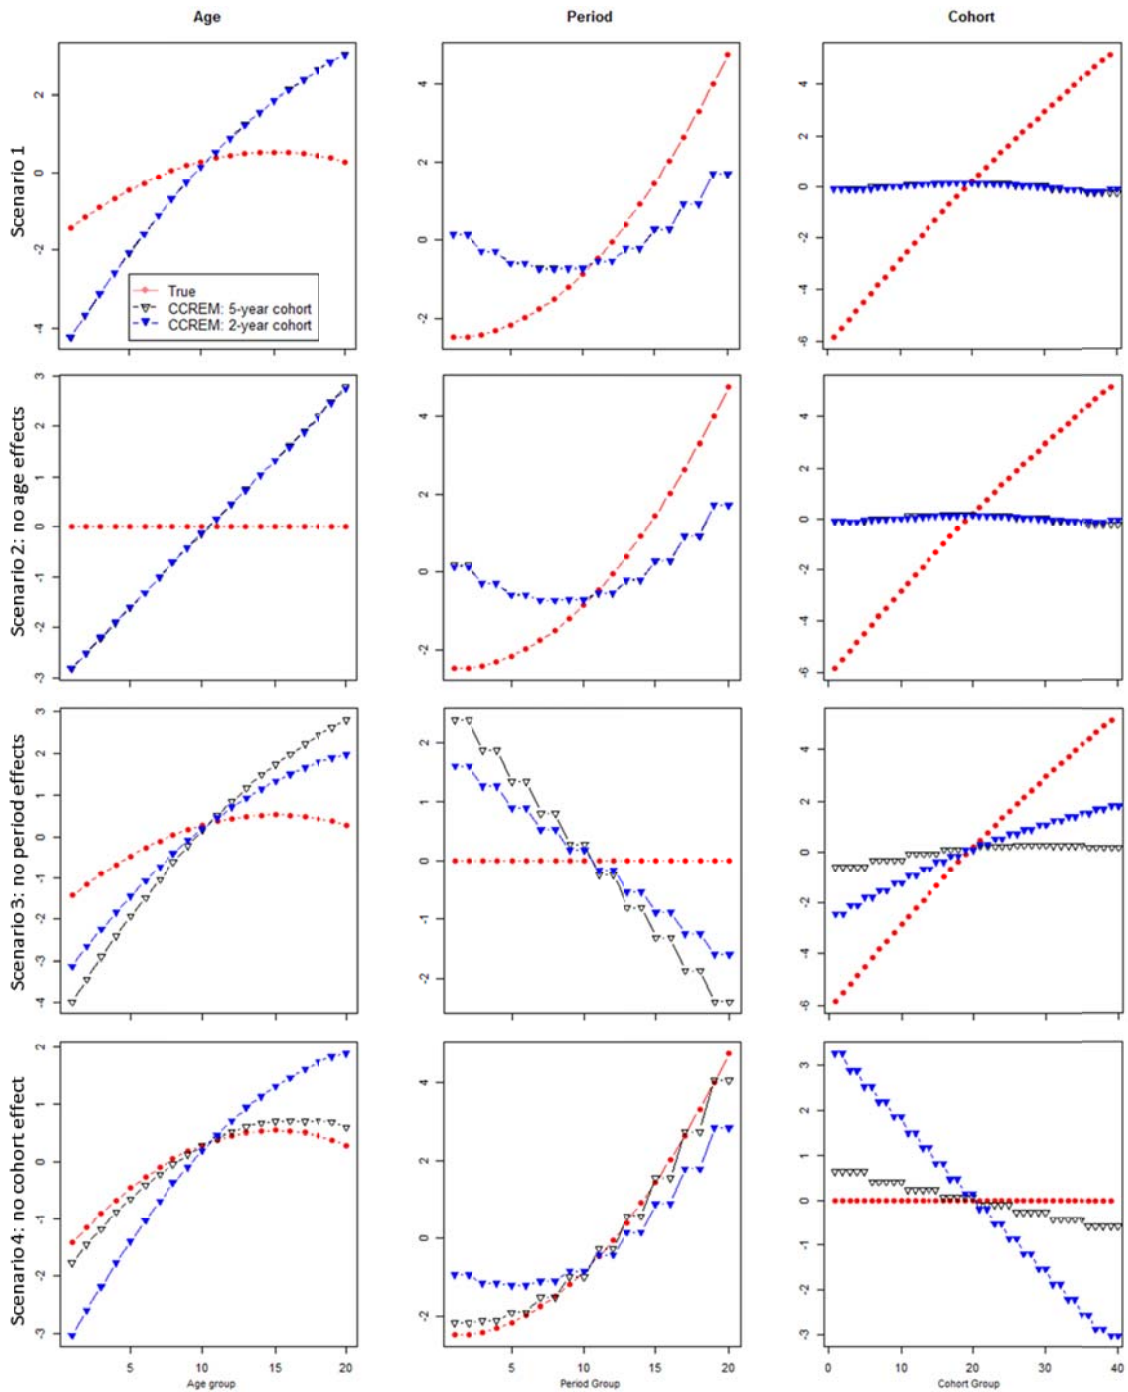


Figure 3.2. Simulation Results: CREM estimates using two- and five-year cohort groups for data from four scenarios.

Conclusion and Discussion

This chapter focuses on the Cross-Classified Fixed Effects Model (CCFEM) and Cross-Classified Random Effects Model (CCREM), two statistical methods intended to separate age, period, and cohort effects on various outcomes. I have shown that, like the APC accounting model, CCFEM and CCREM are not identifiable without constraints, and using different interval widths for age, period, and cohort groups amounts to assuming Multiple Block Constraints (MBCs). These MBCs are extremely difficult, if not impossible, to verify in empirical research. This aspect of CCFEM and CCREM is qualitatively identical to the constraint assumed in CGLM for equal-width age groups, periods, and cohorts except that CGLM usually imposes just one constraint, which is explicit. Moreover, the simulations imply that CCFEM and CCREM do not give unbiased estimators of true age, period, and cohort effects, so that these methods are potentially misleading.

I emphasize that the identification problem is intrinsic in any APC model that attempts to estimate $\{1 + (a - 1) + (p - 1) + (a + p - 2)\}$ parameters because there are only $\{1 + (a - 1) + (p - 1) + (a + p - 2) - 1\}$ unique parameters. This mathematical and logical argument implies that the identification of an APC model cannot be achieved with variable manipulation, i.e., using categorical or continuous variables, creating finer or coarser grouping, etc. As I have shown, although Y&L claim that unequal-width grouping has solved the problem, in fact the true source of the “solution” is the assumption that the effects are identical within the aggregated age, period, and cohort groups.

Y&L are inconsistent about why CCFEM and CCREM appear to be identified. In some expositions they include both linear and quadratic age terms in the model and claim that “it is clear that the underidentification problem of the classical APC accounting model has been resolved by the specification of the quadratic function for the age effects” (2008:84). However, adding a quadratic age term alone does not break the *linear* dependency noted above and thus does not solve the identification problem; the linear age, period, and cohort effects are exactly related in APC models with unequal interval widths, as they are in models with equal interval widths. In other articles they seem aware of the problem of MBCs, noting that “results may be sensitive to the choice of interval widths” (Zheng, Yang, and Land 2011:960). Unfortunately, despite this caution, they and other researchers continue to use unequal-width age, period, and cohort intervals without recognizing the substantial consequences of the MBCs for estimation results.

Another argument put forth by Y&L about the seeming immunity of CCFEM and CCREM to the identification problem is that these models are hierarchical models and fundamentally different from the APC accounting model (Yang and Land 2013:191). In fact, CCFEM is no different from the APC accounting model except the data are not aggregated in CCFEM, which has no bearing on identification. CCREM includes fixed age effects and random period and cohort effects, which makes it look different from the accounting model. However, as I have explained and the APC literature (Bell and Jones 2013a, 2013b, 2014a, 2014b, Fienberg 2013, O’Brien 2014, O’Brien et al. 2008) has documented, the identification problem arises from the linear relationship between age, period, and cohort and cannot be solved by using random effects model.

It is worth noting that although unequal-width intervals have always been used in CCREM, CCREM may be identified without the MBCs because of another type of constraint. O'Brien et al. (2008) and Bell and Jones (2013b, 2014a, 2014b) note that treating the effects of a variable as random is analogous to assuming no linear trends in these effects. Also, declaring the period or cohort effects to be a random effect is simply adding a constraint (Hodges 1998, 2013), albeit with a form that is less familiar than setting the first two age effects to be equal as in CGLM. A full treatment of this point is beyond the present chapter's scope, but I briefly describe the idea. In conventional analyses of mixed effect models in the form of model (4) (e.g., in R, SAS, or STATA), estimates of σ^2 and G are obtained by maximizing the restricted likelihood, then β and u are estimated by minimizing this equation:

$$(Y - W\beta - Zu)^T(Y - W\beta - Zu)/\hat{\sigma}^2 + u^T \hat{G}^{-1}u. \quad (12)$$

The first term in eq. (12) is the residual sum of squares in a standard regression, divided by $\hat{\sigma}^2$; the second term, sometimes called a penalty (as in "penalized likelihood" or "penalized regression"), constrains the estimates of the random effects by shrinking them toward zero. Thus the random effect distribution is itself a constraint on effects modeled as random effects. This constraint is distinct from and additional to the MBCs. I emphasize that because the identification problem remains in CCREM, even with the constraint implied by CCREM's random effects, changing the MBCs can change CCREM's estimates, as shown in Figures 3.1 and 3.2.

Y&L have suggested that CCREM should be preferred to CCFEM if CCREM has smaller standard error estimates for coefficient estimates (these are available as part of

the output from statistical packages). However, using standard errors to compare the performance of different APC models is inappropriate because this ignores biases in coefficient estimates, which are “the primary reason why patterns in estimated age, period, and cohort effects vary so much as a function of the additional linear constraint” (Kupper et al. 1985:822).

The idea of using individual data to achieve identifiability is not new in the APC literature. For example, Boyle and Robertson (1987: 733) argued that “[a] solution to this problem lies in using non-aggregated data, i.e., data available in the form of individual records” and accordingly Boyle and colleagues (1983, 1987) and Robertson and Boyle (1986) proposed an APC model which, they claim, assumed no arbitrary constraints. However, this individual-records approach does not in fact solve the problem (see, e.g., Clayton and Schiffers 1987, Tango 1988, Osmond and Gardner 1989), and Robertson and Boyle (1998:1311) later noted that their method is problematic.

I argue that the statistical problem in the APC accounting model is in fact theoretical in nature; it warns researchers that the parameters in the APC accounting model may well be different from the concepts of age, period, and cohort effects theorized in the demographic and sociological literature. In the next chapter, I introduce a new APC model that is closely tied to the theoretical idea of cohort and does not incur the identification problem.

CHAPTER 4 : A NEW APC MODEL

Drawing on the literature of sociology, demography, and biostatistics, I develop and introduce a new age-period-cohort model, called the age-period-cohort-interaction (APC-I) model, that can be used to investigate inter- and intra-cohort changes for both aggregated and individual-level data. The specification of this new model is informed by demographic and sociological theories. The model is fully identified and thus does not suffer from the identification problem that has hampered traditional APC models for decades. The APC-I model gives valid estimates that reflect the theoretical ideas of age, period, and cohort effects on various social, demographic, economic, and health outcomes. The traditional APC accounting model assumes that cohort effects are established at birth or during early childhood and do not change over the life course, so even if it were not problematic on technical grounds, it could not be used to examine competing life course hypotheses including the cumulative advantage/disadvantage hypothesis and the compensation hypothesis, in which the outcome of interest, such as political outlook, death rates, and happiness, is dynamic across the life course *within* cohorts (Dannefer 2003, 1987; Hobcraft, Menken, and Preston 1982). The APC-I model, in contrast, is flexible enough to allow tests of these important life course theories about within-cohort variation.

This chapter proceeds as follows. I begin by contrasting two types of cohort effects: the concept of cohort effect that interests demographers and sociologists and the type of cohort effects that are estimated in the classic APC accounting model. The discrepancies between the two imply the theoretical nature of the identification problem

in APC analysis. Next, I introduce the new APC-I model, provide theoretical and methodological rationales for it, describe how it is specified, and explain how inter- and intra-cohort effects can be estimated and tested. I conclude by discussing the application of the APC-I model in particular situations, its connections with other APC models, and its limitations.

Cohort Theories and the APC Accounting Model: Disparities between conceptualization and operationalization

Cohort Analysis and Age-Period-Cohort Framework

Researchers in various disciplines are interested in how outcomes of interest vary across time in a society in which individual biographies are shaped by social events and historic shifts. For example, demographers are interested in social conditions that affect temporal trends in divorce rates in the United States (Kennedy and Ruggles 2014). Sociologists of religion have attempted to test the theory of secularization, which refers to the decline of religion in modern societies (Chaves 1989; Firebaugh and Harley 1991). Until the 1970s, research on temporal processes was dominated by an age-period paradigm, a paradigm that only considers shifts across age groups and time periods. *Age* is arguably one of the most important factors in social science research: a wide range of research has documented that many social, demographic, economic, and health outcomes change as one gets older (Borella et al. 2011; Cole 1971; Elder 1975; Lynch 2004; Riley 1987). At the same time, social and historical changes, captured as a package by *period effects*, can affect individual outcomes including political views, vocabulary knowledge, and health conditions (Peng 1987; Smith 1990; Wilson and Gove 1999; Winship and Harding 2008).

Demographers and sociologists have challenged this age-period paradigm, arguing that this type of research ignores an important dimension of temporal processes: *cohort*. A cohort refers to a group of individuals who experience a significant event like birth, marriage, or graduation at the same time. Cohort is a key concept and useful analytical tool because cohort patterns reflect the formative effects of exposure to social events during critical ages that act persistently over time (Ryder 1965). Social science literature has demonstrated the importance of cohort; omitting cohort in analyzing temporal trends may lead to spurious conclusions about age and period patterns. Therefore, answering questions about temporal processes of demographic, social, economic, and health outcomes requires analysts to simultaneously consider the distinct patterns of age, period, and cohort.

To separate the independent age, period, and cohort effects, Mason et al. (1973) specified an analysis of variance (ANOVA) model, which they titled the age-period-cohort (APC) accounting model¹⁴:

$$g(E(Y_{ij})) = \mu + \alpha_i + \beta_j + \gamma_k, \quad (1)$$

for age groups $i = 1, 2, \dots, a$, periods $j = 1, 2, \dots, p$, and cohorts $k = 1, 2, \dots, (a + p - 1)$, where $\sum_{i=1}^a \alpha_i = \sum_{j=1}^p \beta_j = \sum_{k=1}^{a+p-1} \gamma_k = 0$. $E(Y_{ij})$ denotes the expected value of the outcome Y for the i th age group in the j th period of time; g is the “link function”; α_i denotes the mean difference from the global mean μ associated with the i th age category; β_j denotes the mean difference from μ associated with the j th period; γ_k denotes the mean difference from μ due to the membership in the k th cohort. The usual ANOVA

¹⁴ This is called an “accounting” model because it is not intended for causal analysis.

constraint applies where the sum of coefficients for each effect is set to zero. Unfortunately, the APC accounting model has methodological and theoretical limitations, which the next two sections discuss.

Methodological Critique: What the APC accounting model estimates

The APC accounting model's methodological problem can be illustrated more explicitly using a generic form for the statistical model. Suppose that the outcome of interest is normally distributed, then model (1) can be written as follows:

$$Y = Xb + \varepsilon, \quad (2)$$

where Y is a vector of outcomes; X is the design matrix implied by model (1); b denotes a parameter vector whose elements correspond to the effects of age, period, and cohort groups; and ε denotes random errors with distribution centered on zero. Because of the linear dependency between age, period, and cohort, the design matrix X has rank one less than full, so an infinite number of solutions (estimates) for b fit any data Y equally well. That is, the data cannot distinguish different estimation results, so a constraint must be imposed in order to choose one set of estimates. This problem is called the APC identification problem.

This methodological challenge is inherent in any APC model that attempts to separate the independent effects of age, period, and cohort and thus cannot be solved by changing the model set up (e.g., using random effects for period and cohort as in Yang and Land 2008) or variable manipulation (e.g., using unequal interval width for age, period, and cohort groups as in Robertson and Boyle 1986; Sarma et al. 2011; Sarma et al. 2012). The identification problem is well recognized and its consequences have been

discussed extensively (Bell and Jones 2013a, 2013b, 2014a, 2014b; Fienberg and Mason 1985; Holford 1983, 2006; Kupper et al. 1983, 1985; Luo 2013a; Luo and Hodges Forthcoming). Apparently, empirical information that derived from the data cannot help either because the problem is circular: researchers do the analysis to learn precisely the kind of information needed to justify any such constraint. Therefore, scholars have emphasized that the choice of the constraint must be based on theoretical grounds or external information (Fienberg 2013; Glenn 2005; Luo 2013b; O'Brien 2013), but such theoretical information often does not exist. More importantly, even when a constraint can be justified on theoretical grounds, the estimated “cohort effects” obtained from the APC accounting model (1) can be difficult to interpret. To illustrate, suppose that each of the age, period, and cohort effect has linear and quadratic trends, then model (1) can be written as

$$Y = \beta_0 + \beta_1 a + \beta_2 a^2 + \beta_3 p + \beta_4 p^2 + \beta_5 c + \beta_6 c^2 + \varepsilon,^{15} \quad (3)$$

where Y is the outcome, β_0 denotes the grand mean, and $\beta_1, \beta_2, \dots, \beta_6$ denotes the coefficients for linear and quadratic age, period, and cohort terms. Because *cohort* = *period* – *age*, replacing cohort terms with age and period results in

$$Y = \beta_0 + \beta_1 a + \beta_2 a^2 + \beta_3 p + \beta_4 p^2 + \beta_5 (p - a) + \beta_6 (p - a)^2 + \varepsilon. \quad (4)$$

Simple algebra then gives

$$Y_{ij} = \beta_0 + \beta_1 a + \beta_2 a^2 + \beta_3 p + \beta_4 p^2 + \beta_5 (p - a) + \beta_6 (a^2 + p^2 - 2 \cdot a \cdot p)^2 + \varepsilon_{ij}. \quad (5)$$

¹⁵ The use of a continuous term to index cohort membership may seem odd to APC analysts; I use this strategy only to demonstrate the implication of the age-period-cohort linear dependency for estimating and interpreting cohort effects in the classic APC accounting model.

Eq. (5) shows that the “cohort effects” that APC model (1) attempts to estimate in fact involve linear age and period effects, quadratic age and period effects, and most crucially an age-by-period interaction. Eq. (5) is revealing because it shows that even when researchers *can* choose a set of estimates (i.e., a constraint on β) based on theoretical grounds, the resulting estimates for cohort effects are a combination of linear and nonlinear age and period effects and their interaction. This is unfortunate because the APC accounting model is designed to simultaneously isolate the “independent” effects of age, period, and cohort, but apparently it has not achieved this goal.

The identification problem discussed above has been characterized as being methodological in nature. With that focus, inadequate attention has been given to the theoretical problem that creates the methodological problem in the APC accounting model. In the following subsection, I argue that the APC accounting model fails not so much because of the identification problem but because it makes a conceptual error by assuming that there *are* independent, additive age, period, and cohort effects in the phenomena of interest.

Theoretical Critique: How cohort effects are defined

In his seminal work on cohort analysis, Norman Ryder (1965) offered a theoretical vision about how cohort effects manifest:

“The aggregate by which the society counterbalances attrition is the birth cohort, those persons born in the same time interval and aging together. Each new cohort makes fresh contact with the contemporary social heritage and carries the impress of the encounter through life. ... The new

cohorts provide the opportunity for social change to occur. They do not cause change; they permit it. If change does occur, it differentiates cohorts from one another, and the comparison of their careers becomes a way to study change. The minimal basis for expecting interdependency between intercohort differentiation and social change is that change has variant import for persons of unlike age [emphasis added], and that the consequences of change persist in the subsequent behavior of these individuals and thus of their cohorts.”(1965: 844)

He further elaborated three basic notions on which cohort analysis rest:

“persons of age a in time t are those who were age $a-1$ in time $t-1$; transformations of the social world modify people of different ages in different ways; the effects of these transformations are persistent. In this way a cohort meaning is implanted in the age-time specification.” (1965: 861)

According to this conceptualization, a cohort effect is *defined* as the interaction between age and period effects, where “interaction” is understood in the sense used by statisticians. A social or historical transformation that has uniform consequences for people of all ages can have no cohort effect; likewise, an age-related process that works the same way in all time periods also cannot have a cohort effect. Conceptually, this differs from thinking about cohort as having independent effects net of period and age effects. While researchers have sought (at least implicitly) to isolate the independent effect of cohort among people who are equivalent with respect to age and period, in the

new APC model that I introduce below, *I conceptualize cohort as the degree to which age and period effects are moderated by one another.*

What does this alternate notion of cohort mean for describing and explaining temporal trends in demographic, social, economic, and health outcomes? Instead of assuming that period effects do not exist or that cohort has independent effects net of age and period effects, I argue that a researcher should begin by explicitly describing the degree to which age effects vary across time periods or equivalently, the extent to which period effects vary across age groups. Then, if the effects of period are the same across age groups and or equivalently, if the effects of age are the same across periods, the researcher must look for explanations for trends in the outcome of interest that do not rely on cohort processes, i.e., that are consistent with this empirical pattern. On the other hand, if such moderating effects are present, then the researcher must seek explanations that are consistent with this empirical pattern. It seems very likely, for example, that temporal changes in church attendance have occurred differently in different age groups; older people's church-going activity is probably less amenable to change, and the church attendance of younger people may be declining. If so, church attendance is a cohort characteristic and might explain cohort trends in Americans' political views—but this is the case only if the effects of period vary by age and or vice versa.

Intra-Cohort Dynamics: Constant, cumulative, or compensatory?

Another theoretical limitation of the APC accounting model and its variants (e.g., the cross-classified random effects models) is that they assume cohort effects are constant across the life course (Hobcraft et al. 1982). That is, previous research using these

models not only assumes cohort has an independent effect net of age and period, but also that this cohort effect does not change for individuals across their life course. However, such constant cohort effects may not be plausible. For example, being a young adult when the civil rights movement swept through America may have a long-term effect on individuals' political views, but it is not necessary to assume that those effects persist into later life for that birth cohort. That is, intra-cohort dynamics, an important type of cohort-related variation, are ignored using previous APC models.

Fortunately, under Ryder's (1965) conceptualization of cohort and in the new model introduced below, the assumption of time-constant cohort effects can be relaxed so that competing theoretical ideas about intra-cohort variation can be examined in empirical studies. Specifically, the "cumulative advantage/disadvantage" theory (Dannefer 1987, 2003; DiPrete and Eirich 2006) resonates with the "Matthew Effect" (Merton 1968) and the saying "the rich get richer; and the poor get poorer" (Entwisle, Alexander, and Olsen 2001). It posits that the initial advantages/disadvantages of people with different capacities, resources, and structural locations are incremental or cumulative over the life course so that the gaps between the advantaged and disadvantaged tend to widen over the life course. For example, the protective effects of higher education on overall health may be multiplicative over the life course as highly educated people can utilize more resources and have better opportunities than the less educated to maintain a healthier life style and behaviors. In contrast, the compensatory hypothesis, represented by "the survivor effect" in mortality research (Hobcraft et al. 1982), argues that a harsh environment in early life may eliminate vulnerable individuals so that cohort would show

higher death rates and worse general health in young ages but lower mortality rates and better health when they are old. I show how to use the APC-I model to test this important but neglected type of cohort-related variation.

Towards Paradigm Shift: A new model

The preceding discussion of the methodological limitations of the APC accounting model and associated estimation techniques is not to deny the theoretical importance or explanatory power of the concept of a cohort. The point, rather, is that any search for an ultimate statistical solution under the APC accounting framework, attempting to estimate cohort effects independent of age and period effects, is a “futile” and “unholy” quest (Glenn 1976:900 and Fienberg 2013: 1981, respectively). All forms of the APC accounting models, including the cross-classified APC method proposed by Yang and Land (2008), have serious limitations because they conceive of cohorts in a way that departs from the concept as described by Ryder (1965) and because they assume that cohort effects are constant across the life course. To solve these problems, researchers must move beyond the accounting framework and precipitate a paradigm shift (Kuhn 1996).

I propose a new model that is conceptually and methodologically different from other APC methods: an APC model that explicitly estimates and tests cohort effects as the age-by-period interactions. Each of the hypotheses about intra-cohort dynamics, “constant effects”, “cumulative advantage/disadvantage,” or “compensation”, corresponds to a specific structure of the age-by-period interaction and is thus an alternative to focusing on nonlinear cohort effects as suggested by Holford (1992) and

Fienberg (2013). This new APC model is closely tied to theoretical ideas about cohort effects, it is fully identified, and it is flexible enough to test various hypotheses about changes within cohorts. Thus I suggest that it is a step towards a paradigm shift in APC research. In this chapter, I describe how the new model is specified and suggest appropriate estimation and testing techniques. In Chapter 5, I show how the new model can be used to test theoretical ideas about inter- and intra-cohort changes using the example of labor force participation.

Model Specification

The APC accounting model (1) implies that an independent cohort effect can be present even when the effects of period apply equivalently to all age groups. However, as discussed above, sociological and demographic theories suggest that cohort effects cannot be observed unless period effects differ between age groups. Informed by this theoretical insight, I propose a new model, called the age-period-cohort-interaction (APC-I) model, that treats cohort effects as a specific form of the age-by-period interactions. The general form of this model can be written as

$$g(E(Y_{ij})) = \mu + \alpha_i + \beta_j + \alpha\beta_{ij(k)}, \quad (6)$$

where g, Y_{ij}, μ, α_i and β_j are defined as in model (1) and $\alpha\beta_{ij(k)}$ denotes the interaction of the i th age group and j th period group, corresponding to the effect of the k th cohort. Note that the effect of one cohort includes multiple age-by-period interaction terms $\alpha\beta_{ij(k)}$ that lie on the same diagonal in a table with age groups in rows and periods in columns. Model (6) differs from model (1) in the way cohort effects are modeled; here, cohort effects are considered as a specific form of the age-by-period interaction (I return

to this point in the next paragraph). In statistics, the interaction between two variables describes the differential effects of one variable depending on the level of the other variable (Scheffé 1959). In APC research, this means that if the temporal patterns of the outcome can be attributed to cohorts, significant age-by-period interactions are present. When cohort membership does not affect the outcome—that is, when the effects of historical or social shifts (period effects) are uniform across age categories—then age-by-period interactions are not present.

The view that cohort effects can be quantified as age-by-period interactions has not gone unnoticed in the APC literature (Clogg 1982; Holford 1983; Fienberg and Mason 1985). For example, Clogg noted that cohort effects “are special kinds of A-P interaction” (p. 467). Also, Holford (1983) argued that “[a] model which assumes that ... there is an additive effect due to age, period and cohort is in itself arbitrary. I might instead have considered interactions, but in fact if I look at interactions among any two factors, the third factor spans a subspace of that interaction space.” (p. 322) Technically it is indisputable that the effects of any choice of third variable from among age, period, and cohort can be expressed as the interaction between the other two variables. In this sense, the APC-I model appears to privilege age and period effects and “discriminate against” cohort effects by including age and period main effects and reducing cohort effects to the interaction of the other two. I make this choice based on the following theoretical and methodological considerations. I choose to include age main effects because researchers are usually interested in a general age pattern that many individuals follow. Period main effects are included to represent the kind of impacts of social

changes that affect everyone in the society. The decision to explicitly quantify cohort effects a specific form of age-by-period interaction is informed by the literature on how cohort effects are conceptualized in relation to age and period effects.

In fact, the conceptual motivation of the APC accounting model is similar to the APC-I model: “the inclusion of a set of cohort effects in this kind of model [the APC accounting model] is a way to get a simple and parsimonious description of age by period interactions” (Fienberg and Mason 1985: 71), although this simple and parsimonious description comes at the price of the identification problem and cohort effects that are constant over the life course. To illustrate, suppose I have a normally-distributed outcome Y with five age categories and five periods. Table 4.1’s top panel represents the expected value of the outcome $E(Y_{ij})$ in each cell in terms of the unknown effects α_i and β_j in the APC accounting model (1) and includes $5 - 1 = 4$ independently-varying estimates for the 5 cohort effects. Table 4.1’s bottom panel represents the expected value in each cell in terms of the parameters α_i, β_j and $\alpha\beta_{ij(k)}$ in the APC-I model (6). The latter includes $(5 - 1) \cdot (5 - 1) = 16$ independently-varying estimates for the $5 \cdot 5 = 25$ age-by-period categories, where the remaining $25 - 16 = 9$ quantities are computed using the usual ANOVA constraints.

Consider, for example, the 5th cohort in the diagonal that runs from the upper-left to the lower-right. The effect of belonging to that cohort in the top panel of Table 4.1, γ_5 , corresponds to five elements in the age-by-period interaction, $\alpha\beta_{11(5)}, \alpha\beta_{22(5)}, \alpha\beta_{33(5)}, \alpha\beta_{44(5)}$, and $\alpha\beta_{55(5)}$, in the bottom panel. These five age-by-period interaction terms are unrestricted in model (6), meaning that they can take on any

values (subject to summing to zero down columns and across rows). In model (1), these five age-by-period interaction terms are replaced by a single parameter γ_5 , conforming to a particular theory about changes over the life course within a cohort. Therefore, the APC accounting model, at least conceptually, may be viewed as a special case of the APC-I model effect that attempts to recover a special type of cohort effect by replacing the $(a \cdot p)$ age-by-period interactions with $(a + p - 1)$ cohort categories. However, this parsimony is costly: its price is the model's identifiability and ability to investigate within-cohort dynamics. The developers of the APC accounting model recognized this limitation (Fienberg and Mason 1985: 70, 84), but unfortunately many APC researchers have taken the accounting model as the final word and thus focused on solving the identification problem.

In the "Estimation and Testing" subsection below I describe statistical procedures that I developed to estimate and test cohort effects characterized by age-by-period interactions. This section is fairly technical, so one can skip it on a first reading. In outline, I describe a three-step procedure for first testing the age-by-period interaction overall, second testing the interactions that correspond to each cohort, and third testing inter- and intra-cohort differences.

Table 4.1. Unobserved Parameters in Models (1) and (6).

		Period					
		1	2	3	4	5	
Parameters in Model (1)	Age	1	$\mu + \alpha_1 + \beta_1 + \alpha\beta_5$	$\mu + \alpha_1 + \beta_2 + \alpha\beta_6$	$\mu + \alpha_1 + \beta_3 + \alpha\beta_7$	$\mu + \alpha_1 + \beta_4 + \alpha\beta_8$	$\mu + \alpha_1 + \beta_5 + \alpha\beta_9$
		2	$\mu + \alpha_2 + \beta_1 + \alpha\beta_4$	$\mu + \alpha_2 + \beta_2 + \alpha\beta_5$	$\mu + \alpha_2 + \beta_3 + \alpha\beta_6$	$\mu + \alpha_2 + \beta_4 + \alpha\beta_7$	$\mu + \alpha_2 + \beta_5 + \alpha\beta_8$
		3	$\mu + \alpha_3 + \beta_1 + \alpha\beta_3$	$\mu + \alpha_3 + \beta_2 + \alpha\beta_4$	$\mu + \alpha_3 + \beta_3 + \alpha\beta_5$	$\mu + \alpha_3 + \beta_4 + \alpha\beta_6$	$\mu + \alpha_3 + \beta_5 + \alpha\beta_7$
		4	$\mu + \alpha_4 + \beta_1 + \alpha\beta_2$	$\mu + \alpha_4 + \beta_2 + \alpha\beta_3$	$\mu + \alpha_4 + \beta_3 + \alpha\beta_4$	$\mu + \alpha_4 + \beta_4 + \alpha\beta_5$	$\mu + \alpha_4 + \beta_5 + \alpha\beta_6$
		5	$\mu + \alpha_5 + \beta_1 + \alpha\beta_1$	$\mu + \alpha_5 + \beta_2 + \alpha\beta_2$	$\mu + \alpha_5 + \beta_3 + \alpha\beta_3$	$\mu + \alpha_5 + \beta_4 + \alpha\beta_4$	$\mu + \alpha_5 + \beta_5 + \alpha\beta_5$
Parameters in Model (6)	Age	1	$\mu + \alpha_1 + \beta_1 + \alpha\beta_{11(5)}$	$\mu + \alpha_1 + \beta_2 + \alpha\beta_{12(6)}$	$\mu + \alpha_1 + \beta_3 + \alpha\beta_{13(7)}$	$\mu + \alpha_1 + \beta_4 + \alpha\beta_{14(8)}$	$\mu + \alpha_1 + \beta_5 + \alpha\beta_{15(9)}$
		2	$\mu + \alpha_2 + \beta_1 + \alpha\beta_{21(4)}$	$\mu + \alpha_2 + \beta_2 + \alpha\beta_{22(5)}$	$\mu + \alpha_2 + \beta_3 + \alpha\beta_{23(6)}$	$\mu + \alpha_2 + \beta_4 + \alpha\beta_{24(7)}$	$\mu + \alpha_2 + \beta_5 + \alpha\beta_{25(8)}$
		3	$\mu + \alpha_3 + \beta_1 + \alpha\beta_{31(3)}$	$\mu + \alpha_3 + \beta_2 + \alpha\beta_{32(4)}$	$\mu + \alpha_3 + \beta_3 + \alpha\beta_{33(5)}$	$\mu + \alpha_3 + \beta_4 + \alpha\beta_{34(6)}$	$\mu + \alpha_3 + \beta_5 + \alpha\beta_{35(7)}$
		4	$\mu + \alpha_4 + \beta_1 + \alpha\beta_{41(2)}$	$\mu + \alpha_4 + \beta_2 + \alpha\beta_{42(3)}$	$\mu + \alpha_4 + \beta_3 + \alpha\beta_{43(4)}$	$\mu + \alpha_4 + \beta_4 + \alpha\beta_{44(5)}$	$\mu + \alpha_4 + \beta_5 + \alpha\beta_{45(6)}$
		5	$\mu + \alpha_5 + \beta_1 + \alpha\beta_{51(1)}$	$\mu + \alpha_5 + \beta_2 + \alpha\beta_{52(2)}$	$\mu + \alpha_5 + \beta_3 + \alpha\beta_{53(3)}$	$\mu + \alpha_5 + \beta_4 + \alpha\beta_{54(4)}$	$\mu + \alpha_5 + \beta_5 + \alpha\beta_{55(5)}$

Estimation and Testing

The conceptual idea of characterizing cohort effects as age-by-period interactions described in the preceding section could be implemented in more than one way. In this section, I describe the way I estimate and test these effects and the interpretations that it allows. However, other ways are possible, and I briefly consider them in Chapter 6.

With cohort effects represented as the age-by-period interaction, testing hypotheses about inter- and intra-cohort variation is equivalent to examining a specific form of—that is, specific patterns and structures in—the diagonal cells of an age-by-period cross-classification. Specifically, in the APC-I model, variation between cohorts can be examined by testing the average difference between the groups of age-by-period interactions that lie along the $(a + p - 1)$ diagonals of the age-by-period cross-classification. Variation within cohorts can be investigated by imposing a restriction on the group of age-by-period interactions that corresponds to a cohort of interest, so, for example, testing the APC accounting model’s hypothesis of a constant γ_5 is equivalent to testing a specific pattern in $\alpha\beta_{11}$, $\alpha\beta_{22}$, $\alpha\beta_{33}$, $\alpha\beta_{44}$, and $\alpha\beta_{55}$. I describe below a three-step procedure for investigating age and period effects and inter- and intra-cohort dynamics. The next section demonstrates this procedure with an empirical example. The Appendix provides exemplary R code for the tests in Steps 1 through 3.

Step 1. A global F test¹⁶: Are there variations in the outcome of interest associated with cohort membership that cannot be explained by age and period effects? To answer

¹⁶ For normally-distributed outcomes, the F test is the likelihood-ratio test. For outcomes with non-Gaussian distributions, e.g., Poisson and Bernoulli distribution, the likelihood-

this question, fit model (6), which includes main age effects, main period effects, and their interactions. Then test the variation attributable to the age-by-period interactions, with $(a - 1) \times (p - 1)$ degrees of freedom. A significant F statistic indicates that cohort effects *may* be present. Note that a significant global F test does not characterize cohort effects, nor is it a sufficient condition for the existence of cohort effects. For example, the interaction might appear significant because the data deviate from pure age and period main effects but do so in a haphazard manner, i.e., one that has no reasonable interpretation as a cohort pattern.

However, a significant test is a necessary condition for cohort effects: A non-significant result suggests that the age-by-period interaction does not explain much variation in the outcome, so the reduced model with only age and period main effects fits the data as well as the model with the full interaction. In other words, a non-significant test suggests that the effects of social events are not different for individuals of different ages so that there is no evidence that cohort membership matters for the outcome of interest. In this case, there is no need to do the tests in Steps 2 or 3, which concern cohort patterns.

Step 2. Deviation magnitude F tests: Does membership in a specific cohort matter?

I address this question using an F test to test a hypothesis about each subset of the age-by-period interactions that corresponds to a cohort. This F test examines the magnitude of cohort-specific deviation from age and period main effects; that is, whether that group of age-by-period interactions, taken together, explains a significant proportion of

ratio test can be used for Steps 1 and 2. The Appendix gives R code for implementing both types of tests.

variation in the outcome. If the F test rejects the null hypothesis, one may conclude that membership in that cohort has effects on the outcome of interest. However, these F tests do not allow researchers to distinguish to what extent or in what ways cohorts differ from each other in the outcome of interest. Steps 3.1 and 3.2 include two t tests for characterizing between-cohort differences and within-cohort dynamics.

Step 3.1. Deviation consistency t tests¹⁷. For each cohort that significantly deviates from age and main effects, based on the deviation magnitude F test, compute the average of the age-by-period interaction terms contained in that cohort and use a t test¹⁸ to examine the consistency of cohort-specific deviation. These averages and associated t test results can be used to assess patterns across cohorts in the outcome.

Step 3.2. t tests for intra-cohort variation. For each cohort whose life course deviates from that defined by the age and period main effects, based on the deviation magnitude F test, conduct a t test¹⁹ of the linear (and quadratic if desirable) orthogonal polynomial contrast of that cohort's age-by-period interaction terms to investigate whether the advantages or disadvantages of members of that cohort cumulate, remain stable, or disappear in their life course.

Table 4.2 provides a guideline about how to use the results of Steps 3.1 and 3.2 to evaluate three theoretical ideas about within-cohort dynamics, “constant effects”, “cumulative advantage/disadvantage”, and “compensation hypothesis”. Specifically, the

¹⁷ Because there are multiple tests involved, the researcher needs to adjust these tests for multiple comparisons.

¹⁸ See the Appendix for the formula for the standard error for with the average of these interaction effects.

¹⁹ See the Appendix for formula for the contrasts and associated standard errors.

data can be considered to support the cumulative advantage/disadvantage hypothesis when a given cohort's average (Step 3.1) and linear trend (Step 3.2) have the same sign, as shown in Table 4.2's upper-left and lower-right cells. When a cohort's average and linear trend have opposite signs, as in Table 4.2's upper-right and lower-left cells, this supports the "convergence" hypothesis: a cohort's initial advantage/disadvantage is disappearing as that cohort ages. When a cohort's average effect is not statistically significant but its linear trend is significant, this favors compensation theory. If the linear trend is not significant but the average effect is significant, then the constant effects hypothesis—the hypothesis implicit in the APC accounting model—seems plausible. If neither the average nor the linear trend is significant, it could mean that there is no clear pattern in cohort variation, and the significant deviation magnitude F test is a result of some kind of deviation that does not conform to any theoretical idea of cohort effects.

Table 4.2. Testing Theories about Intra-Cohort Dynamics Using the Three-Step Procedure.

		Sign of Intra-Cohort Linear Trend (Step 3.2)		
		+	0	-
Sign of Average Cohort Effect (Step 3.1)	+	cumulative advantage	constant effects	converging
	0	compensatory	no clear pattern	compensatory
	-	converging	constant effects	cumulative disadvantage

Three remarks about the three-step procedure: First, the idea of using these test statistics in APC analysis is not new. For example, Clayton and Schiffers (1987) recommended using deviance or the likelihood-ratio criterion to choose among an age-only model, an age-period model, an age-cohort model, or a full age-period-cohort model. Also, Yang (2008) suggested comparing these models, though concluding that certain test outcomes justified use of a constrained approach like the intrinsic estimator. However, the purpose of the global F test proposed here is neither model selection nor verifying technical constraints on the unknown age, period, and cohort parameters. Because I consider cohort effects as the interaction between age and period, the global F test on the age-by-period interaction in model (6) serves as an explicit measure of and necessary condition for cohort effects to be considered present.

Second, in the presence of significant age-by-period interactions, researchers should use caution in interpreting estimated age and period main effects. In general,

there are two types of interactions: quantitative interactions, in which the trend of the outcome in age, say, has the same direction for all periods but periods differ in the strength of the trend; and cross-over or qualitative interactions, in which the trend of the outcome in age has a different direction depending on period. It is difficult to interpret main effects in a meaningful way when cross-over or qualitative interactions are present. However, one can still interpret main effects in the presence of quantitative interactions, as the average trend. See Aiken and West (1991) and Jaccard and Turrisi (2003) for detailed discussions on this topic. In the APC-I method, age and period main effects are still interpretable with quantitative age-by-period interactions. With qualitative age-by-period interactions, it is advisable to refrain from reifying estimated age and period main effects because the effect of period depends qualitatively on age, to such an extent that a general age pattern or period pattern, applying to the whole population, is not meaningful. In my view, the most sensible analytical strategy in the presence of qualitative age-by-period interactions is a comparison of cohort-specific age-graded trajectories because each cohort has a distinct aging process and is subject to the impacts of social change in a way differing from other cohorts.

Third, researchers should be careful about interpreting cohort effects when less than three age-by-period cells make up the cohort's diagonal in the age-period classification table because it may be misleading to treat a trend determined by so few data points (e.g., two age-by-period cells), usually for the youngest or the oldest cohorts, as a good indicator of the general trend for that cohort across its life course. The more

age-by-period cells observed for a cohort, the more informative the estimates are for understanding changes within that cohort.

Discussion and Conclusion

Despite the conceptual merits and explanatory power of age, period, and cohort, traditional age-period-cohort (APC) models that are designed to separate the independent effects of the three variables suffer from an identification problem. As a result, no valid estimates of age, period, or cohort effects can be ascertained. While this identification problem has been considered a methodological challenge, I argue that the identification problem is theoretical in nature: The cohort effects conceptualized in the demographic and sociological literatures and those estimated in traditional APC models are not the same, which gives rise to the technical problem.

In this chapter, I developed a new APC model, the APC-I model, that is more closely tied to the concept of cohort effects by explicitly modeling cohort effects as specific structures of the age-by-period interaction. The APC-I model has two advantages. First, like any two-way ANOVA model with an interaction, this model is identified, so it avoids the identification problem. Second, unlike traditional APC models that implicitly assume cohort effects are static through the life course, the APC-I model relaxes this assumption, allowing researchers to investigate within-cohort dynamics. In Chapter 5, I will demonstrate how the new APC-I model and the testing strategies described in this chapter can be used to examine age, period, and in particular, inter- and intra-cohort changes in white and black men's and women's labor force participation.

Considering cohort effects as the differential effects of period (or age) depending on age groups (or time periods) has important implications for study designs for cohort research. Technically, the presence of qualitative or cross-over interactions—in which the trend in the effects of one variable has a different direction depending on the level of the other variable—poses a challenge for interpretation of main effects. In some cases and disciplines, interactions are difficult to interpret because they often have no substantive meanings. However, in the APC-I model, the age-by-period interaction terms may represent substantively meaningful cohort effects. This implies that when qualitative age-by-period interactions are operative, the most sensible study design would be comparing life course trajectories of cohorts because each cohort has a distinct age or period pattern so that there is no general age or period trend.

Appendix

Appendix 4.1: Exemplary R code for implementing Steps 1 through 3.

```
options(contrasts=c("contr.sum","contr.poly"), na.action = na.omit)
attach(data)

A # number of age groups
P # number of period groups
C = A + P - 1 # number of cohort groups
covn = 0 # number of covariates

temp = glm(data = data, inlfc ~ acc * pcc, family)
r6 = temp$coefficients
r6se = summary(temp)$coef[,"Std. Error"]
r6p = summary(temp)$coef[,"Pr(>|z|)"]

##### computing "full" interactions index #####
T = array(rep(0, A*P*(A-1)*(P-1)), dim=c(A*P, (A-1)*(P-1)))
ind1 = A*1:(P-1)
ind2 = (A*(P-1)+1):(A*P-1)
ind3 = A*P
ind = c(ind1,ind2,ind3)
newind = 1:(A*P)
newind = newind[-ind]

T[newind,] = diag((A-1)*(P-1))
T[ind1,] = -diag(P-1)[,rep(1:(P-1),each=A-1)]
T[ind2,] = -diag(A-1)[,rep(1:(A-1),P-1)]
T[ind3,] = -rep(1,(A-1)*(P-1))

##### computing "full" interactions contrast #####
iatemp = vcov(temp6)[(covn+A+P): length(r6), (covn+A+P): length(r6)]
iavcov = T%%iatemp%%t(T)
df = temp6$df.residual

iaesti = as.vector(T%%r6[(covn+A+P): length(r6)])
iase = sqrt(diag(iavcov))
iap = pt(-abs(iaesti/iase), df)*2

cindex = array(rep(0, A*P), dim = c(A, P))
for (j in 1:P){
  cindex [,j] = seq((A+j-1),j, -1)
}
```

```

cohortindex = as.vector(cindex)
ia  = as.data.frame(cbind(iaesti,iase,iap,cohortindex))

##### inter-cohort changes #####
cint  = rep(NA, C)
cintse = rep(NA, C)
cintt  = rep(NA, C)
cintp  = rep(NA, C)

for (k in 1:C){
  O = sum(cindex == k)
  k1 = rep(1/O, O)
  k2 = rep(0, A*P)
  k2[cindex == k] = k1

  contresti = k2%%iaesti
  contrse = sqrt(t(k2)%%iavcov%%k2)
  t = contresti/contrse
  if (t > 0){
    p = 2*pt(t, df, lower.tail=F)
  } else {
    p = 2*pt(t, df, lower.tail=T)
  }

  cint[k]  = contresti
  cintse[k] = contrse
  cintt[k] = t
  cintp[k] = p
}

cgroup = seq(1, C)
cohortint = cbind(cgroup, cint, cintse, cintt, cintp)

##### intra-cohort changes #####
cslope  = rep(NA, C)
cslopese = rep(NA, C)
cslopet  = rep(NA, C)
cslopep  = rep(NA, C)

poly = 1
for (k in (poly+1):(C-poly)){
  o = sum(cindex == k)
  k1 = contr.poly(o)
  k2 = rep(0, A*P)

```

```

k2[cindex == k] = k1[,poly]

contresti = k2%%iaesti
contrse = sqrt(t(k2)%%iavcov%%k2)
t = contresti/contrse
if (t > 0){
  p = 2*pt(t, df, lower.tail=F)
} else {
  p = 2*pt(t, df, lower.tail=T)
}
cslope[k] = contresti
cslopese[k] = contrse
cslopet[k] = t
cslopep[k] = p
}

cgroup = seq(1, C)
cohortslope = cbind(cgroup, cslope, cslopese, cslopet, cslopep)
cohortslope

```

CHAPTER 5 : APPLICATION OF THE APC-I MODEL TO TEMPORAL TRENDS IN LABOR FORCE PARTICIPATION

Temporal trends in white and black women's labor force participation (LFP) have been the subject of much research (see, e.g., Connelly 1992; Farkas 1977; Hollister and Smith 2014; Macunovich 2012; Treas 1987). Decomposing these trends into different dimensions of time, namely age, period, and cohort, can provide important clues about demographic, social, and economic factors that give rise to these temporal trends in women's LFP. However, while LFP is obviously a function of age, there is no consensus among scholars about the extent to which the temporal patterns in women's LFP are driven by social and historical shifts (period effects) or by population metabolism (cohort effects) (Clogg 1982; Percheski 2008).

Meanwhile, in contrast to the huge volume of scholarly work on women's LFP, limited attention has been given to changes in men's LFP by race over the past decades. This underdevelopment of the literature on men's LFP implies that scholars have missed the opportunity to obtain a more comprehensive view of temporal patterns in LFP in general and to advance knowledge about how gender and race/ethnicity differences in this important aspect of life manifest across different dimensions of time and the implications for gender inequality in earnings, housework, and child care.

In this chapter, using the APC-I model developed and described in Chapter 4, I will address the substantive limitations in the LFP literature by examining temporal patterns in white and black men's and women's LFP. This research not only describes

mean differences in LFP between gender and race groups but also highlights the different shapes of their age trajectory, period pattern, and cohort deviations in LFP.

Moreover, for a number of reasons, the few studies on LFP that adopted a cohort perspective focused on differences between cohorts. This ignores *intra*-cohort dynamics, i.e., dynamics within cohorts across their life course. In this study, in addition to examining differences between cohorts, I consider intra-cohort variation by investigating whether advantages or disadvantages of a cohort in LFP are constant, cumulative, or compensatory over the life course *over* and *above* the general age and period pattern. Specifically, I attempt to answer three sets of research questions:

1. For each gender (male and female) and race/ethnic group (white and black), (a) to what extent does their LFP vary as a function of the three dimensions of time, i.e., age, period, and cohort? (b) Which cohorts have especially high and especially low LFP rates relative to their age and period? (c) Over and above age and period patterns, are cohort effects on LFP constant, cumulative, or compensatory over the life course?
2. For each gender and race/ethnic group, in what ways have the two most significant changes, in educational attainment and marital status, affected those trends? Would LFP rates be stagnant or lower had education or marriage not changed in the United States? That is, are some cohorts more likely to consistently engage in the labor force than others, regardless of their age? Or do different cohorts demonstrate distinct trajectories or patterns of LFP as they age?

3. Have gender and racial inequalities in LFP remained the same, increased, or decreased in each of the three dimensions of time, before and after adjusting for the influence of the factors mentioned above?

Background

Few trends in post-World War II American society are as striking as the rise in women's LFP, although it has leveled off or even declined since 2000 (Bureau of Labor Statistics 2014). These changes in LFP among women have enormous implications for child rearing, gender inequality in earnings, couples' work-family arrangements, and mental health. Scholars have searched for social and demographic factors that have (sometimes differentially) contributed to the temporal change (see, e.g., Bianchi 2011; Fernández 2013; Goldin 1990, 2006; Hollister and Smith 2013). For example, in her series of papers, Goldin (1989, 1990, 2006) documented and analyzed women's LFP for the past two centuries and explored how these changes were intertwined with changes in education, in wages, in the household and workplace, and in attitudes towards women. Boushey (2008) and Bradbury and Katz (2005) found varying effects of having children on LFP across time periods and between marital statuses. Percheski (2008) identified an inter-cohort increase in employment rates among college-educated professional women.

Women's decisions to participate in the labor force, however, are not made solely based on their own will, needs, or characteristics; rather, their decisions are likely to depend at least partially on the attitude of their partners—to a large extent, men—and important others including family and neighbors, towards women working outside the home and to their socioeconomic status, including LFP status. Likewise, among

heterosexual couples, men's LFP is likely to be affected by their female partners' education, financial status, and attitudes about work-family arrangements. Therefore, an explicit consideration of both men's and women's LFP in conjunction is useful for better understanding temporal variation in LFP in the United States and for suggesting demographic and social forces that underlie this change.

Unfortunately, men's LFP has received limited scholarly attention; although men's LFP in the US declined after World War II, the literature on LFP among men is less developed than that on women's LFP. Interest may be smaller because the magnitude of secular or period trends in LFP among men is not as profound as for women's LFP. However, period variation, or lack thereof, represents only one dimension of time processes and thus is insufficient evidence on which to make any definitive conclusion about temporal variation.

In this chapter, I begin with a comparative description of those temporal trends in LFP for white and black men and women using the APC-I model. I then explore how changes in social and demographic factors are associated with temporal variation in LFP. An exhaustive investigation of all possible factors is beyond the scope of this research. I focus on educational attainment and marital status; the LFP literature has well-documented the ways in which changes in these two social and demographic factors may affect LFP. Specifically, using the APC-I model, I analyzed the 1962 to 2014 Current Population Survey data (described in detail in the next section) to examine the following hypotheses.

Hypothesis 1: While LFP is positively associated with educational attainment for all gender and race groups, the implication of this positive association for the age, period, and cohort patterns in LFP varies between white and black men and women. If educational attainment had not increased during the past decades, the decrease in LFP rates among men would have been more pronounced (Hypothesis 1.1). For women's LFP, Goldin (1989) argued that there are long-standing race/ethnicity differences, which can be traced back to a "double legacy" of slavery. She hypothesized that while black people tend to have lower LFP rates than white people due to lower levels of education, black women may have had higher LFP rates than white women because the norms and expectations developed under slavery about women's work were different from those of whites and those norms were carried into the post-Emancipation era. However, as educational attainment and the percentage of the population that is college-educated has increased more among whites than among blacks, I hypothesize that the black-white gap in LFP among women should have been reversed in recent years with white women having higher LFP rates than black women (Hypothesis 1-2).

Hypothesis 2: Because marriage implies greater family responsibilities for women than men, the decline in marriage rates in the United States, which implies that a smaller percentage of women would assume the burden of housekeeping chores and child rearing, was associated with the increase in women's but not in men's LFP (Hypothesis 2-1). This marriage penalty on LFP, however, is greater for white women than for black women because fewer black women are married and fewer children were born to married black mothers (Hypothesis 2-2).

Data

I use data from the 1962 through 2014 Current Population Survey (CPS) March Supplement (as disseminated by IPUMS-CPS). The CPS is a monthly survey conducted by the Census Bureau and the Bureau of Labor Statistics. A battery of questions on demographic information and labor force participation is fielded very month. The focal outcome is labor force participation (LFP). Every year since 1962, CPS has asked respondents whether they participated in the labor force during the week prior to the interview. Being in the labor force (coded 1) means the respondents “were at work; held a job but were temporarily absent from work due to factors like vacation or illness; were seeking work; or were temporarily laid off from a job during the reference period” (Mariam et al. 2010). The respondents were otherwise out of the labor force (coded 0). Age and year of interview are ascertained in every survey. I exclude respondents with missing data on LFP, age, survey year, gender, or race²⁰, giving a sample of 2,470,428 records²¹ for white males, 264,385 for black males, 2,727,462 for white females, and 351,734 for black females.

To examine how changes in educational attainment and current marital status affect age, period, and cohort patterns in LFP, I further exclude respondents with missing data on years of schooling and marital status, resulting in a slightly smaller sample size for each subgroup. I constructed 13 age groups (18-19, 20-24, 25-29, ... 70-74, and 75

²⁰ While race categories in the CPS changed in 1988, 2003, and 2013, the black and white categories are comparable across the full 1962-2014 time periods.

²¹ A person who appears in one March CPS will also appear in an adjacent March CPS. Therefore, the sample size in this research refers to the number of individual records, not respondents.

and older), 11 periods (1962-1964, 1965-1969, 1970-1974, ..., 2005-2009, and 2010-2014), and thus 23 birth cohorts (1885, 1890, ..., 1990, 1994)²². Table 5.1 presents descriptive statistics for the outcome variable and for gender, race, educational attainment, marital status, and the three time-related predictors (age, period, and cohort).

Results

I analyzed the CPS data using the APC-I model described in Chapter 4²³. For each gender and race group, I began with a model, labelled Model *a*, without including any covariates; these models describe “raw” age and period effects on LFP and deviations associated with cohort membership for white and black men and women. I then present models labelled *b* and *c* that add educational attainment and marital status, respectively. These analyses were intended to depict age, period, and cohort patterns in LFP after considering the effects of each of the two factors; that is, in what ways have changes in educational attainment and marital status influenced age, period, and cohort patterns in LFP for each gender and race group?

²² In a table of five-year age groups and five-year periods, a birth cohort is defined by diagonals of the age-period cross-classification table and extends over a nine-year interval. For example, the observations in 1975 through 1979 for people in the 30 to 34 age group describe the birth cohort of 1941 to 1949. Conventionally, each cohort is identified by its mid or central birth year (e.g., Mason and Winsborough 1973; O’Brien 2011). I follow this practice so, for example, the 1945 cohort refers to the group of people born between 1941 and 1949. When so defined, birth cohorts overlap with adjacent cohorts. This overlap is usually ignored in statistical modeling (Kupper et al. 1985).

²³ Because the CPS uses a complex sampling strategy, all analyses used the weighting variable “WTSUPP” provided by IPUMS-CPS.

Table 5.1. Descriptive Statistics for All Analytic Variables, Current Population Survey March Supplement, 1962-2014

Description	N	Mean	S.D.
Labor force participation (LABFORCE; 0=not in labor force, 1=in labor force)	6,106,893	0.64	(0.48)
Age at time of survey (AGE)	6,106,893	42.99	(18.02)
Period (YEAR)	6,106,893	-	-
Birth cohort (YEAR - AGE)	6,106,893	-	-
Gender (SEX; 0=male, 1=female)	6,106,893	0.53	(0.50)
Race (RACE; 0=white, 1=black)	5,814,187	0.11	(0.31)
Educational attainment (EDUC; 5=graduate school; 4=college; 3=some college; 2=high school; 1=less than high school)	6,054,568	-	-
Marrital status (MARST; 0=never married, separated, divorced, or widowed; 1=currently married)	6,106,350	-	-

Note. Analysis includes respondents who participated in the 1962 through 2014 CPS surveys March Supplement for whom labor force participation status, year of birth, gender, and race, educational attainment, and marital status are available. Words in all caps are CPS variable names.

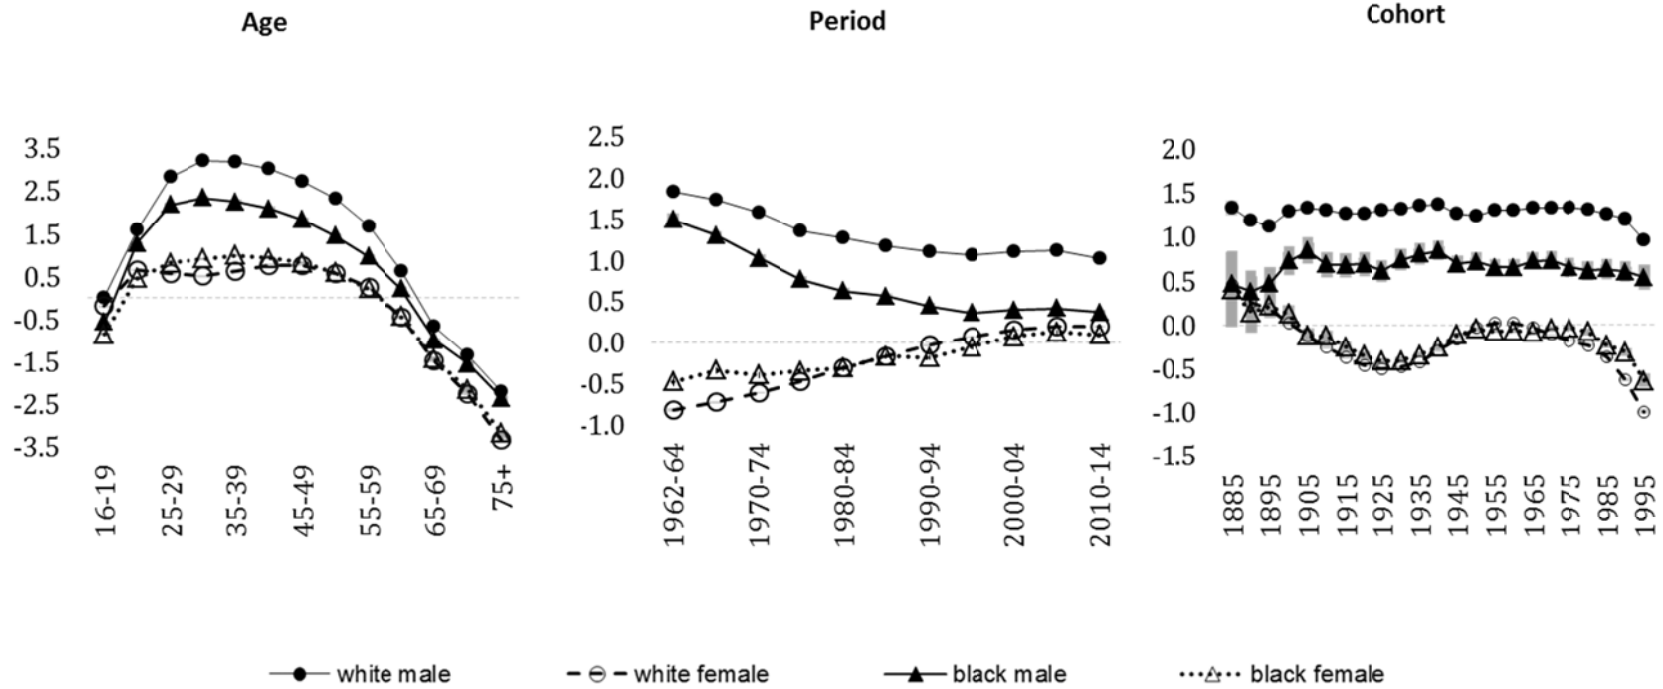


Figure 5.1. Age and Period Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014

Note: Age and period main effects estimates derived from Models 1a through 4a plus the intercept for each gender and race group in Tables 5.2. Cohort deviation estimates derived from the “Average deviation” columns in Table 5.5 plus the intercept estimates in Model *a* in Table 5.2. Values represent REML coefficients coded to sum to zero. Grey bars depict 95% confidence intervals. Analysis includes CPS respondents who participated in the 1962 through 2014 CPS surveys for whom labor force participation status and age are available.

Table 5.2. Estimated Age and Period Main Effects on Labor Force Participation, with and without Adjustment for Education and Marital Status, in March CPS, 1962-2014

		Men					
		White			Black		
		Model 1a	Model 1b	Model 1c	Model 2a	Model 2b	Model 2c
Intercept		1.260 ***	1.601 ***	1.137 ***	0.645 ***	1.221 ***	0.669 ***
	<H.S.	—	-0.574 ***	—	—	-0.835 ***	—
	H.S.	—	-0.214 ***	—	—	-0.287 ***	—
Education	Some Col.	—	-0.270 ***	—	—	-0.188 ***	—
	B.A.	—	0.384 ***	—	—	0.551 ***	—
	>B.A.	—	0.674 ***	—	—	0.759 ***	—
Marital Status	currently married	—	—	-0.473 ***	—	—	-0.515 ***
	never or SDW	—	—	0.473 ***	—	—	0.515 ***
	16-19	-1.261 ***	-1.100 ***	-0.684 ***	-1.198 ***	-1.026 ***	-0.719 ***
	20-24	0.280 ***	0.292 ***	0.685 ***	0.580 ***	0.535 ***	0.933 ***
	25-29	1.517 ***	1.503 ***	1.670 ***	1.408 ***	1.361 ***	1.570 ***
	30-34	1.868 ***	1.837 ***	1.878 ***	1.564 ***	1.493 ***	1.615 ***
	35-39	1.856 ***	1.824 ***	1.807 ***	1.477 ***	1.404 ***	1.479 ***
	40-44	1.660 ***	1.635 ***	1.585 ***	1.321 ***	1.266 ***	1.283 ***
Age	45-49	1.423 ***	1.405 ***	1.333 ***	1.099 ***	1.074 ***	1.044 ***
	50-54	1.013 ***	0.993 ***	0.897 ***	0.749 ***	0.733 ***	0.665 ***
	55-59	0.387 ***	0.381 ***	0.253 ***	0.284 ***	0.287 ***	0.159 **
	60-64	-0.668 ***	-0.685 ***	-0.836 ***	-0.444 ***	-0.428 ***	-0.598 ***
	65-69	-1.961 ***	-1.973 ***	-2.159 ***	-1.620 ***	-1.591 ***	-1.849 ***
	70-74	-2.628 ***	-2.635 ***	-2.821 ***	-2.194 ***	-2.155 ***	-2.421 ***
	75+	-3.485 ***	-3.477 ***	-3.608 ***	-3.026 ***	-2.952 ***	-3.162 ***
	1962-64	0.549 ***	0.709 ***	0.482 ***	0.817 ***	1.069 ***	0.717 ***
	1965-69	0.457 ***	0.619 ***	0.380 ***	0.636 ***	0.844 ***	0.537 ***
	1970-74	0.294 ***	0.443 ***	0.220 ***	0.347 ***	0.537 ***	0.271 ***
	1975-79	0.097 ***	0.232 ***	0.045 ***	0.059 **	0.225 ***	0.026
	1980-84	-0.010	0.114 ***	-0.036 ***	-0.079 ***	0.067 **	-0.090 ***
Period	1985-89	-0.117 ***	0.000	-0.115 ***	-0.175 ***	-0.047 *	-0.157 ***
	1990-94	-0.183 ***	-0.231 ***	-0.158 ***	-0.296 ***	-0.335 ***	-0.244 ***
	1995-99	-0.226 ***	-0.400 ***	-0.183 ***	-0.350 ***	-0.540 ***	-0.297 ***
	2000-04	-0.244 ***	-0.436 ***	-0.187 ***	-0.293 ***	-0.541 ***	-0.247 ***
	2005-09	-0.265 ***	-0.471 ***	-0.191 ***	-0.317 ***	-0.602 ***	-0.251 ***
	2010-14	-0.353 ***	-0.579 ***	-0.257 ***	-0.349 ***	-0.675 ***	-0.265 ***
Cohort		(See Table 5.5)					
N		2,470,428	2,448,163	2,448,163	264,385	262,131	264,385

Note. Analysis includes CPS respondents who participated in the 1962 through 2014 CPS surveys in years for which labor force participation status and year of birth are available. Then, samples are restricted to respondents for whom all data are available. Figures represent REML regression coefficients coded to sum to zero. ***=p<0.001 ; ** = p < 0.01 ; * = p < 0.05

Table 5.2 (continued). Estimated Age and Period Main Effects on Labor Force Participation, with and without Adjustment for Education and Marital Status, in March CPS, 1962-2014

		Women					
		White			Black		
		Model 3a	Model 3b	Model 3c	Model 4a	Model 4b	Model 4c
Intercept		-0.229 ***	0.146 ***	-0.169 ***	-0.180 ***	0.413 ***	-0.171 ***
	<H.S.	—	-0.720 ***	—	—	-0.963 ***	—
	H.S.	—	-0.212 ***	—	—	-0.297 ***	—
Education	Some Col.	—	-0.049 ***	—	—	-0.003	—
	B.A.	—	0.300 ***	—	—	0.556 ***	—
	>B.A.	—	0.681 ***	—	—	0.707 ***	—
Marital Status	currently married	—	—	0.348 ***	—	—	-0.031 ***
	never or SDW	—	—	-0.348 ***	—	—	0.031 ***
	16-19	0.042 ***	0.247 ***	-0.313 ***	-0.642 ***	-0.441 ***	-0.622 ***
	20-24	0.903 ***	0.853 ***	0.795 ***	0.631 ***	0.544 ***	0.640 ***
	25-29	0.837 ***	0.767 ***	0.908 ***	0.994 ***	0.896 ***	0.992 ***
	30-34	0.759 ***	0.691 ***	0.892 ***	1.076 ***	0.986 ***	1.070 ***
	35-39	0.855 ***	0.801 ***	1.004 ***	1.176 ***	1.101 ***	1.169 ***
	40-44	0.994 ***	0.954 ***	1.141 ***	1.111 ***	1.055 ***	1.102 ***
Age	45-49	0.984 ***	0.960 ***	1.124 ***	1.002 ***	0.966 ***	0.994 ***
	50-54	0.808 ***	0.797 ***	0.935 ***	0.761 ***	0.759 ***	0.754 ***
	55-59	0.456 ***	0.453 ***	0.551 ***	0.397 ***	0.420 ***	0.391 ***
	60-64	-0.235 ***	-0.228 ***	-0.196 ***	-0.255 ***	-0.230 ***	-0.258 ***
	65-69	-1.258 ***	-1.239 ***	-0.196 ***	-1.264 ***	-1.224 ***	-1.264 ***
	70-74	-2.027 ***	-1.993 ***	-0.196 ***	-1.965 ***	-1.924 ***	-1.959 ***
	75+	-3.120 ***	-3.063 ***	-0.196 ***	-3.022 ***	-2.909 ***	-3.010 ***
	1962-64	-0.593 ***	-0.362 ***	-0.561 ***	-0.281 ***	0.000	-0.288 ***
	1965-69	-0.511 ***	-0.299 ***	-0.480 ***	-0.201 ***	0.080 ***	-0.208 ***
	1970-74	-0.373 ***	-0.178 ***	-0.347 ***	-0.223 ***	0.036	-0.227 ***
	1975-79	-0.231 ***	-0.061 ***	-0.215 ***	-0.175 ***	0.049 *	-0.176 ***
	1980-84	-0.063 ***	0.085 ***	-0.060 ***	-0.139 ***	0.051 *	-0.139 ***
Period	1985-89	0.073 ***	0.198 ***	0.068 ***	-0.008	0.149 ***	-0.007
	1990-94	0.208 ***	0.139 ***	0.197 ***	0.021	-0.043 *	0.024
	1995-99	0.319 ***	0.111 ***	0.304 ***	0.168 ***	-0.081 ***	0.171 ***
	2000-04	0.382 ***	0.144 ***	0.363 ***	0.270 ***	-0.034 *	0.273 ***
	2005-09	0.395 ***	0.127 ***	0.371 ***	0.294 ***	-0.064 ***	0.298 ***
	2010-14	0.395 ***	0.095 ***	0.360 ***	0.274 ***	-0.144 ***	0.279 ***
Cohort		(See Table 5.5)					
N		2,727,462	2,702,798	2,726,969	351,734	348,980	351,698

Recall that the APC-I model includes age and period main effects and their interaction terms. Table 5.2 reports and Figure 5.1 illustrates estimated age and period main effects on LFP without adjusting for educational attainment or marriage (Model *a*). An inspection of the estimated age and period effects shows that age trends in LFP rates do not differ in a *cross-over* or *qualitative* manner²⁴ depending on period²⁵, so describing a general age trend and a general period trend, i.e., the main effects, is meaningful. Not surprisingly, LFP increases with age through midlife and then declines thereafter for all gender and race groups. White men had the highest LFP rates across all ages, while black men had lower participation but exceeded white and black women. The estimated period effects for the models in Table 5.2 and Figure 5.1 suggest that there was a decline in men's LFP, especially from the 1960s to late 1990s. In contrast, women's LFP was gradually increasing so that it almost reached the same level as black men's by the 2010s. In general, the magnitude of the period effects is smaller than that of the age effects.

The magnitude of decrease or increase in LFP over time differed between white and black people within both gender groups. The decline in LFP was greater among black men than white men. Although black women had higher participation rates before

²⁴In general, there are two types of interactions: quantitative interactions, in which the trend of the outcome in one variable has the same direction across levels of another variable but differs in the strength of the trend; and cross-over or qualitative interactions, in which the trend of the outcome in one variable has a different direction depending on the level of another variable. It is difficult to interpret main effects in a meaningful way when cross-over or qualitative interactions are present. However, one can still interpret main effects in the presence of quantitative interactions, as the average trend. See Aiken and West (1991) and Jaccard and Turrisi (2003) for detailed discussions on this topic.

²⁵ See Fig A5.1 in Appendix 5.1.

1980, the increase for white women was faster so that they had a somewhat higher rates after 1980.

To answer the question about whether and in what ways cohort membership affects LFP, I followed the three-step procedures described in Chapter 4 to examine the age-by-period interaction terms in the APC-I model. Specifically, the global F statistics for age-by-period interactions (Step 1) on LFP for white men, white women, black men, and black women are 966.58, 1627.2, respectively, and they are all statistically significant ($p < 0.0001$). These F statistics suggest that for each gender and race subgroup, the model that includes the age-by-period interactions fits better than the model with age and period main effects only. I thus concluded that there may be cohort effects on LFP rates and proceeded to Step 2 to identify which cohort deviates significantly from that the LFP trajectory defined by age and period main effects.

The results of the F and t tests in Steps 2 and 3 are presented in Tables 5.5 and 5.6 for each gender and race group, and I will discuss them in detail later. Consider the example of black women for the purpose of illustrating computation, testing, and interpretation of cohort deviations in the APC-I model. Table 5.3 presents estimated age-by-period interaction terms in the APC-I Model 1a for black women, with rows defined by age groups and columns by time periods. As shown in Table 4.1 in Chapter 4, the interactions terms on each diagonal corresponds to the effects, or deviations, associated with a cohort relative to the age and period main effects. For example, the age-by-period interaction terms that lie on the diagonal running from the cell for people aged 25-29 in 1962-64 to that for those aged 65-69 in 2010-14 (i.e., -0.482, -0.399, ... , -0.133)

correspond to the deviation effects for the cohort born around 1935. To better illustrate the implementation of Steps 2 and 3 for investigating cohort deviations, Table 5.4 rearranges Table 5.3's estimates for age-by-period interaction terms, with rows defined by age groups and columns defined by cohorts.

The bottom three rows show the F and t statistics in Steps 2 and 3 for the group of interaction terms for each cohort. As described in Chapter 4, Step 2, the deviation magnitude F test for a given cohort compares the fit of the saturated model with all main effects and $(a - 1) \cdot (p - 1)$ age-by-period interaction terms versus a reduced model with main effects and $(a - 1) \cdot (p - 1) - o$ interaction terms, where o is the number of age-by-period interaction terms for that cohort. These F tests indicate, generally speaking, whether the LFP of a cohort deviates significantly from the pattern defined by age and period main effects. The F test results for black women's LFP are shown in the third row from the bottom in Table 5.4, labelled "Cohort Deviation Magnitude F Test" at the far left. For example, for the 1935 cohort, the F statistic of 212.603 ($p < 0.0001$) indicates that the group of age-by-period interactions for this cohort—i.e., that lie on the diagonal corresponding to the 1935 cohort in Table 5.3—taken all together explains a significant portion of variation in LFP for black women.

Table 5.3. Estimates of Age-by-Period Interaction in Model 1a for Black Women’s Labor Force Participation in March CPS, 1962-2014

		Period										
		1962-64	1965-69	1970-74	1975-79	1980-84	1985-89	1990-94	1995-99	2000-04	2005-09	2010-14
Age	16-19	0.114	0.094 *	0.112 **	0.121 **	0.031	0.230 ***	0.020	0.109 **	-0.092 **	-0.280 ***	-0.460 ***
	20-24	-0.199 ***	-0.060	0.044	0.041	0.047	0.061	0.003	0.149 ***	0.038	-0.106 ***	-0.019
	25-29	-0.482 ***	-0.414 ***	-0.112 **	0.173 ***	0.158 ***	0.075 *	0.005	0.273 ***	0.199 ***	0.126 ***	-0.001
	30-34	-0.424 ***	-0.399 ***	-0.263 ***	-0.029	0.246 ***	0.185 ***	0.109 **	0.164 ***	0.241 ***	0.170 ***	0.001
	35-39	-0.438 ***	-0.363 ***	-0.216 ***	-0.108 *	0.120 **	0.232 ***	0.186 ***	0.190 ***	0.244 ***	0.096 **	0.057
	40-44	-0.312 ***	-0.273 ***	-0.265 ***	-0.216 ***	0.114 *	0.246 ***	0.294 ***	0.108 **	0.160 ***	0.089 *	0.056
	45-49	-0.120	-0.210 ***	-0.254 ***	-0.127 **	0.101 *	0.148 **	0.126 **	0.077 *	0.110 **	0.127 ***	0.022
	50-54	0.163 *	-0.039	-0.124 **	-0.174 ***	-0.158 ***	-0.001	0.080	0.143 ***	0.129 ***	-0.040	0.024
	55-59	0.133	0.166 **	0.121 *	-0.027	-0.058	-0.121 **	-0.191 ***	-0.037	-0.020	0.038	-0.004
	60-64	0.348 ***	0.171 **	0.183 **	0.000	0.086	-0.128 **	-0.136 **	-0.258 ***	-0.203 ***	-0.041	-0.021
	65-69	0.368 ***	0.384 ***	0.161 *	0.205 **	-0.065	-0.386 ***	-0.152 *	-0.379 ***	-0.183 **	-0.115 *	0.163 ***
	70-74	0.204	0.530 ***	0.468 ***	0.136	-0.137	-0.314 **	-0.065	-0.321 ***	-0.414 ***	-0.138	0.050
	75+	0.645 ***	0.413 **	0.147	0.004	-0.486 **	-0.227	-0.278 *	-0.218 *	-0.209 *	0.075	-0.133

Note: Estimates represent REML regression coefficients for age-by-period interaction terms coded to sum to zero across age and period groups. ***=p<0.001 ; ** = p < 0.01 ; * = p < 0.05 .

Table 5.4. Estimates of Age-by-Period Interaction, Deviation Magnitude F Test, Consistency t Test, and Intra-Cohort Slope t Test Results in Model 1a for Black Women’s Labor Force Participation in March CPS, 1962-2014

		Birth Cohort											
		1885	1890	1895	1900	1905	1910	1915	1920	1925	1930	1935	1940
Age	16-19												
	20-24												-0.199 ***
	25-29											-0.482 ***	-0.414 ***
	30-34										-0.424 ***	-0.399 ***	-0.263 ***
	35-39									-0.438 ***	-0.363 ***	-0.216 ***	-0.108 *
	40-44								-0.312 ***	-0.273 ***	-0.265 ***	-0.216 ***	0.114 *
	45-49							-0.120	-0.210 ***	-0.254 ***	-0.127 **	0.101 *	0.148 **
	50-54						0.163 *	-0.039	-0.124 **	-0.174 ***	-0.158 ***	-0.001	0.080
	55-59					0.133	0.166 **	0.121 *	-0.027	-0.058	-0.121 **	-0.191 ***	-0.037
	60-64				0.348 ***	0.171 **	0.183 **	0.000	0.086	-0.128 **	-0.136 **	-0.258 ***	-0.203 ***
	65-69			0.368 ***	0.384 ***	0.161 *	0.205 **	-0.065	-0.386 ***	-0.152 *	-0.379 ***	-0.183 **	-0.115 *
	70-74		0.204	0.530 ***	0.468 ***	0.136	-0.137	-0.314 **	-0.065	-0.321 ***	-0.414 ***	-0.138	0.050
	75+	0.645 ***	0.413 **	0.147	0.004	-0.486 **	-0.227	-0.278 *	-0.218 *	-0.209 *	0.075	-0.133	
	Step 2:												
	Cohort Deviation Magnitude F Test	5.674 *	10.794 **	38.268 ***	74.137 ***	34.060 ***	30.544 ***	15.350 *	101.307 ***	181.026 ***	262.535 ***	212.603 ***	190.769 ***
Step 3.1:													
Step 1:	Cohort Deviation Consistency t Test	0.645 **	0.309 *	0.348 ***	0.301 ***	0.023	0.059	-0.099 **	-0.157 ***	-0.223 ***	-0.231 ***	-0.192 ***	-0.086 ***
Step 3.2:													
	Intra-Cohort Linear Slope t Test	NA	0.148	-0.156	-0.212	-0.402 ***	-0.338 ***	-0.228 *	0.055	0.132	0.197 **	0.240 ***	0.260 ***

Note: Point estimates represent REML regression coefficients for age-by-period interaction terms coded to sum to zero across age and period groups. F statistics and t derived using the method described in Chapter 4. ***= $p < 0.001$; ** = $p < 0.01$; * = $p < 0.05$.

Table 5.4 (continued). Estimates of Age-by-Period Interaction, Deviation Magnitude F Test (Step 2), Consistency t Test, and Intra-Cohort Slope t Test Results in Model 1a for Black Women's Labor Force Participation in March CPS, 1962-2014

		Birth Cohort										
		1945	1950	1955	1960	1965	1970	1975	1980	1985	1990	1995
Age	16-19	0.114	0.094 *	0.112 **	0.121 **	0.031	0.230 ***	0.020	0.109 **	-0.092 **	-0.280 ***	-0.460 ***
	20-24	-0.060	0.044	0.041	0.047	0.061	0.003	0.149 ***	0.038	-0.106 ***	-0.019	
	25-29	-0.112 **	0.173 ***	0.158 ***	0.075 *	0.005	0.273 ***	0.199 ***	0.126 ***	-0.001		
	30-34	-0.029	0.246 ***	0.185 ***	0.109 **	0.164 ***	0.241 ***	0.170 ***	0.001			
	35-39	0.120 **	0.232 ***	0.186 ***	0.190 ***	0.244 ***	0.096 **	0.057				
	40-44	0.246 ***	0.294 ***	0.108 **	0.160 ***	0.089 *	0.056					
	45-49	0.126 **	0.077 *	0.110 **	0.127 ***	0.022						
	50-54	0.143 ***	0.129 ***	-0.040	0.024							
	55-59	-0.020	0.038	-0.004								
	60-64	-0.041	-0.021									
	65-69	0.163 ***										
	70-74											
	75+											
	Step 2:											
	Cohort Deviation Magnitude F Test	83.257 ***	166.318 ***	92.513 ***	90.901 ***	74.362 ***	101.158 ***	80.511 ***	22.635 ***	19.053 ***	64.130 ***	159.644 ***
	Step 3.1:											
Step 1:	Cohort Deviation Consistency t Test	0.059 ***	0.130 ***	0.095 ***	0.107 ***	0.088 ***	0.150 ***	0.119 ***	0.069 ***	-0.066 **	-0.149 ***	-0.460 ***
Global F Test		13.555 ***										
	Step 3.2:											
	Intra-Cohort Linear Slope t Test	0.090	-0.096 *	-0.114 **	0.004	0.050	-0.075	0.030	-0.053	0.064	0.185 ***	NA

The deviation consistency t test (Step 3.1) results reported in the second row from the bottom examines whether the average of each cohort's group of age-by-period interaction terms is significantly different from 0. Again, take the 1935 cohort of black women for example: the average cohort deviation is -0.192 and the t statistic is statistically significant, meaning that on average, the 1935 cohort had lower LFP than what we would expect relative to their ages and periods.

The t test results for intra-cohort slope (Step 3.2) for black women's LFP are shown in the last row in Table 5.4. Using the previous example of the 1935 cohort, this group of black women had an intra-cohort slope of 0.24, which was estimated as the linear contrast of the age-by-period interaction terms contained in that cohort. This negative slope was statistically significant using the method for computing standard errors provided in the Appendix in Chapter 4. It implies that although on average, this cohort had lower-than-expected LFP rates relative to their ages and periods, they seemed catching up by being more likely to participate at older than at younger ages.

Table 5.5. Estimated Average Deviation between Cohorts and Intra-Cohort Linear Slope in Labor Force Participation, with and without Adjustment for Education, in March CPS, 1962-2014

		White Men					
		Model 1a		Model 1b		Model 1c	
		Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope
Cohort	1885	0.016	NA	0.064	NA	0.151 **	NA
	1890	-0.150 ***	-0.061	-0.103 ***	-0.089 *	-0.038	-0.033
	1895	-0.201 ***	-0.002	-0.152 ***	-0.033	-0.124 ***	0.015
	1900	-0.015	-0.287 ***	0.019	-0.278 ***	0.050 **	-0.311 ***
	1905	0.018	-0.344 ***	0.067 ***	-0.376 ***	0.064 ***	-0.386 ***
	1910	-0.025	-0.419 ***	0.005	-0.391 ***	-0.004	-0.447 ***
	1915	-0.034 *	-0.337 ***	0.002	-0.353 ***	-0.038 **	-0.377 ***
	1920	-0.052 ***	-0.342 ***	-0.037 *	-0.344 ***	-0.082 ***	-0.333 ***
	1925	0.026 *	-0.231 ***	0.021	-0.237 ***	-0.009	-0.243 ***
	1930	0.055 ***	0.035	0.039 **	-0.009	0.002	0.047
	1935	0.073 ***	-0.310 ***	0.077 ***	-0.332 ***	0.035 **	-0.197 ***
	1940	0.095 ***	0.505 ***	0.064 ***	0.455 ***	0.040 ***	0.558 ***
	1945	-0.025 **	0.808 ***	-0.082 ***	0.803 ***	-0.053 ***	0.808 ***
	1950	-0.073 ***	0.433 ***	-0.136 ***	0.469 ***	-0.069 ***	0.435 ***
	1955	-0.023 **	-0.045 *	-0.059 ***	0.023	0.008	-0.048 *
	1960	-0.053 ***	-0.452 ***	-0.069 ***	-0.368 ***	-0.009	-0.450 ***
	1965	-0.007	-0.545 ***	-0.019	-0.457 ***	0.034 ***	-0.546 ***
	1970	0.002	-0.598 ***	-0.003	-0.523 ***	0.039 ***	-0.585 ***
	1975	0.010	-0.542 ***	0.052 ***	-0.558 ***	0.040 ***	-0.511 ***
	1980	0.010	-0.551 ***	0.092 ***	-0.646 ***	0.045 ***	-0.479 ***
1985	-0.016	-0.380 ***	0.081 ***	-0.488 ***	0.016	-0.272 ***	
1990	-0.043 ***	0.007	0.098 ***	-0.058 ***	-0.060 ***	0.071 ***	
1995	-0.307 ***	NA	-0.107 ***	NA	-0.395 ***	NA	

Note. Analysis includes CPS respondents who participated in the 1962 through 2014 CPS surveys in years for whom labor force participation status and year of birth are available. Then, samples are restricted to respondents for whom all data are available. Figures in the “Average deviation” columns represent the averages of the group of age-by-period interaction estimates for each cohort. Figures in the “Intra-cohort linear slope” columns represent the linear slope of the group of age-by-period interaction estimates for each cohort. ***= $p < 0.001$; ** = $p < 0.01$; * = $p < 0.05$

Table 5.5 (continued). Estimated Average Deviation between Cohorts and Intra-Cohort Linear Slope in Labor Force Participation, with and without Adjustment for Education, in March CPS, 1962-2014

		Black Men					
		Model 2a		Model 2b		Model 2c	
		Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope
Cohort	1885	-0.093	NA	-0.172	NA	0.073	NA
	1890	-0.398 ***	-0.040	-0.454 ***	-0.039	-0.246 *	-0.021
	1895	-0.323 ***	0.059	-0.322 ***	-0.010	-0.212 **	0.105
	1900	-0.001	-0.123	-0.010	-0.137	0.098	-0.148
	1905	0.137 **	0.023	0.138 **	0.037	0.181 ***	-0.026
	1910	-0.065	-0.400 **	-0.053	-0.378 **	-0.054	-0.413 ***
	1915	-0.001	-0.112	0.041	-0.133	-0.012	-0.187
	1920	-0.010	-0.074	0.022	-0.002	-0.027	-0.107
	1925	-0.054	0.049	-0.022	0.103	-0.098 **	0.013
	1930	0.037	0.079	0.066	0.019	-0.001	0.088
	1935	0.097 **	-0.562 ***	0.146 ***	-0.706 ***	0.039	-0.333 *
	1940	0.165 ***	-0.056	0.164 ***	-0.113	0.072 **	0.064
	1945	-0.020	0.404 ***	-0.058 *	0.397 ***	-0.071 **	0.354 ***
	1950	0.017	0.206 ***	-0.050 *	0.218 ***	0.004	0.157 **
	1955	-0.083 ***	0.056	-0.137 ***	0.114 *	-0.033	0.027
	1960	-0.049 *	-0.134 **	-0.093 ***	-0.080	0.017	-0.169 ***
	1965	0.011	-0.193 ***	-0.029	-0.123 *	0.077 ***	-0.183 ***
	1970	0.025	-0.228 ***	-0.024	-0.162 **	0.080 ***	-0.181 ***
	1975	-0.057 *	-0.336 ***	-0.065 *	-0.363 ***	-0.012	-0.248 ***
	1980	-0.112 ***	-0.515 ***	-0.049	-0.607 ***	-0.050	-0.353 ***
1985	-0.083 **	-0.350 ***	0.013	-0.481 ***	-0.051	-0.215 ***	
1990	-0.075 **	-0.195 ***	0.083 **	-0.293 ***	-0.122 ***	-0.131 **	
1995	-0.212 ***	NA	0.040	NA	-0.327 ***	NA	

Table 5.5 (continued). Estimated Average Deviation between Cohorts and Intra-Cohort Slope in Labor Force Participation, with and without Adjustment for Education, in March CPS, 1962-2014

		White Women					
		Model 3a		Model 3b		Model 3c	
		Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope
Cohort	1885	0.653 ***	NA	0.662 ***	NA	0.590 ***	NA
	1890	0.487 ***	-0.026	0.511 ***	-0.085	0.423 ***	-0.023
	1895	0.428 ***	-0.179 ***	0.429 ***	-0.208 ***	0.362 ***	-0.161 ***
	1900	0.249 ***	-0.210 ***	0.240 ***	-0.216 ***	0.187 ***	-0.197 ***
	1905	0.108 ***	-0.382 ***	0.098 ***	-0.371 ***	0.059 ***	-0.365 ***
	1910	-0.030 **	-0.379 ***	-0.028 *	-0.387 ***	-0.064 ***	-0.365 ***
	1915	-0.141 ***	-0.278 ***	-0.116 ***	-0.280 ***	-0.153 ***	-0.263 ***
	1920	-0.242 ***	-0.169 ***	-0.214 ***	-0.183 ***	-0.235 ***	-0.150 ***
	1925	-0.280 ***	0.062 *	-0.250 ***	0.032	-0.259 ***	0.097 ***
	1930	-0.257 ***	0.351 ***	-0.228 ***	0.304 ***	-0.227 ***	0.386 ***
	1935	-0.214 ***	0.446 ***	-0.188 ***	0.389 ***	-0.187 ***	0.387 ***
	1940	-0.086 ***	0.430 ***	-0.087 ***	0.378 ***	-0.039 ***	0.394 ***
	1945	0.045 ***	0.181 ***	0.003	0.195 ***	0.079 ***	0.157 ***
	1950	0.185 ***	0.067 ***	0.121 ***	0.096 ***	0.206 ***	0.038 *
	1955	0.241 ***	-0.137 ***	0.187 ***	-0.098 ***	0.251 ***	-0.163 ***
	1960	0.240 ***	-0.272 ***	0.207 ***	-0.220 ***	0.239 ***	-0.290 ***
	1965	0.177 ***	-0.256 ***	0.146 ***	-0.197 ***	0.166 ***	-0.258 ***
	1970	0.123 ***	-0.172 ***	0.091 ***	-0.127 ***	0.102 ***	-0.159 ***
	1975	0.063 ***	0.055 ***	0.071 ***	0.013	0.027 ***	0.061 ***
	1980	0.025 **	0.136 ***	0.080 ***	0.011	-0.032 ***	0.121 ***
1985	-0.128 ***	0.220 ***	-0.040 ***	0.089 ***	-0.213 ***	0.148 ***	
1990	-0.403 ***	0.149 ***	-0.258 ***	0.044 **	-0.488 ***	0.050 **	
1995	-0.801 ***	NA	-0.563 ***	NA	-0.806 ***	NA	

Table 5.5 (continued). Estimated Inter- and Intra-Cohort Effects on Labor Force Participation, with and without Adjustment for Education, in March CPS, 1962-2014

		Black Women					
		Model 4a		Model 4b		Model 4c	
		Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope	Average deviation	Intra-cohort linear slope
Cohort	1885	0.645 **	NA	0.706 **	NA	0.658 **	NA
	1890	0.309 *	0.148	0.163	0.253	0.323 **	0.146
	1895	0.348 ***	-0.156	0.328 ***	-0.198	0.360 ***	-0.159
	1900	0.301 ***	-0.212	0.271 ***	-0.203	0.311 ***	-0.217
	1905	0.023	-0.402 ***	0.035	-0.463 ***	0.029	-0.405 ***
	1910	0.059	-0.338 ***	0.083 *	-0.409 ***	0.063	-0.343 ***
	1915	-0.099 **	-0.228 *	-0.068 *	-0.170	-0.098 **	-0.233 *
	1920	-0.157 ***	0.055	-0.118 ***	0.066	-0.158 ***	0.050
	1925	-0.223 ***	0.132	-0.169 ***	0.149	-0.227 ***	0.126
	1930	-0.231 ***	0.197 **	-0.195 ***	0.134	-0.237 ***	0.194 **
	1935	-0.192 ***	0.240 ***	-0.162 ***	0.188 **	-0.198 ***	0.256 ***
	1940	-0.086 ***	0.260 ***	-0.073 ***	0.156 *	-0.094 ***	0.269 ***
	1945	0.059 ***	0.090	0.026	0.073	0.052 ***	0.096
	1950	0.130 ***	-0.096 *	0.062 ***	-0.082	0.127 ***	-0.093 *
	1955	0.095 ***	-0.114 **	0.024	-0.069	0.096 ***	-0.116 **
	1960	0.107 ***	0.004	0.045 **	0.084 *	0.112 ***	0.002
	1965	0.088 ***	0.050	0.039 *	0.121 **	0.094 ***	0.046
	1970	0.150 ***	-0.075	0.091 ***	-0.027	0.156 ***	-0.076
	1975	0.119 ***	0.030	0.118 ***	-0.005	0.127 ***	0.034
	1980	0.069 ***	-0.053	0.129 ***	-0.183 ***	0.077 ***	-0.041
1985	-0.066 **	0.064	0.041	-0.097 *	-0.060 **	0.077 *	
1990	-0.149 ***	0.185 ***	0.031	0.053	-0.149 ***	0.196 ***	
1995	-0.460 ***	NA	-0.147 ***	NA	-0.469 ***	NA	

The deviation magnitude F statistics are all statistically significant ($p < 0.05$; numbers not shown here). The results of the deviation consistency and intra-cohort slope t tests for each gender and race/ethnicity group are summarized in Table 5.5. The cohort pattern in Figure 5.1 was created by plotting the estimates in the “Average deviation” columns in Table 5.5 plus the intercept estimates in Model a Table 5.2 with 95% confidence intervals for each gender and race group. It shows that differences between cohorts in average deviation among white and black men were not substantial, although many are statistically significant because of the large sample size of the March CPS; exceptions include the second and third oldest cohorts and the youngest and second youngest cohorts of white and black men who had especially low LFP rates. Differences in average cohort deviation were more pronounced among women: Relative to the pattern determined by age and period main effects, deviation associated with cohort in LFP decreased between the 1885 and 1930 cohorts, increased afterwards until the 1980 cohort for black women and the 1955 cohort for white women, and then decreased for more recent cohorts.

While the average deviation from age and period main effects associated each cohort provides one way to assess cohort-related variation, the life-course dynamics within each cohort are still hidden behind these averages. In some cases, depending on average deviation may miss the opportunity to show theoretically important and empirically interesting information about within-cohort dynamics. Did these average deviations in LFP among black and white women remain stable, decrease or increase over the life course of those cohorts? According to the “accumulative advantage”

hypothesis (Dannefer 1987; Dannefer 2003; Hobcraft, Menken and Preston 1982), cohorts with high LFP rates should have progressively higher rates across the life course as they accumulate more experiences, skills, and resources. The *t* test for intra-cohort slope (Step 3.2) examines this hypothesis by testing the linear contrast of the corresponding age-by-period interactions. The results of this *t* test reported in Table 5.5²⁶ show limited support for the “cumulative (dis)advantage” hypothesis.

Specifically, while white women in the 1955 through 1975 birth cohorts had higher LFP rates (relative to age and period main effects) than other cohorts at young ages, the intra-cohort slopes for those cohorts are significantly negative and substantial in magnitude—suggesting that these cohorts lost their relative advantage as they aged. For the 1950 and 1975 birth cohorts, the substantively trivial slopes indicate that members of these cohort were able to maintain (but not to increase) their relative advantage in LFP. Similarly, the non-significant intra-cohort slopes for the 1925 cohort suggest that this cohort remained (neither increased nor decreased) at a low level of LFP.

In contrast, the significant negative intra-cohort slopes for the 1895 through 1920 birth cohorts indicate that members of those cohorts had increasingly lower LFP (relative to their age and period) as they grew older. For the 1930 through 1940 cohorts, although on average, their LFP rates were lower than other cohorts relative to their age and period,

²⁶ I caution about the estimates for intra-cohort trends for the youngest and oldest cohorts with fewer than three age-by-period interaction terms on the cohort diagonals in the age-by-period cross-classifications; the effects estimates of these cohorts are determined by only two age-by-period interaction terms, so the linear trend in these effects may be different from the trend that would be observed if more age-by-period interactions were available for these cohorts. The oldest and youngest cohort have only one corresponding age-by-period interaction term, so no information about intra-cohort change is available.

they were catching up as they aged. Most interestingly, for the 1945 and 1980 cohorts, their LFP rates seemed “compensatory”; that is, these cohorts’ lower-than-average LFP rates at younger ages are compensated by higher-than-average rates at older ages so that the average deviation associated with the cohort membership in LFP for them were not statistically or substantively significant.

To what extent can we attribute the age, period, and cohort patterns described in Figure 5.1 to changes in demographic and socioeconomic factors? I investigate the extent to which age, period, and cohort effects in LFP can be attributed to changes in educational attainment and in marriage. Given that education is positively associated with LFP and the amount of formal schooling has increased considerably in the last century (Fischer and Hout 2006), it would be interesting to see how the change in educational attainment (differentially) affects LFP for each gender and race group. Similarly, since marital status is related to LFP and the proportion of currently married Americans has been declining, I attempt to understand whether and how this change has influenced temporal trends in white and black men’s and women’s LFP.

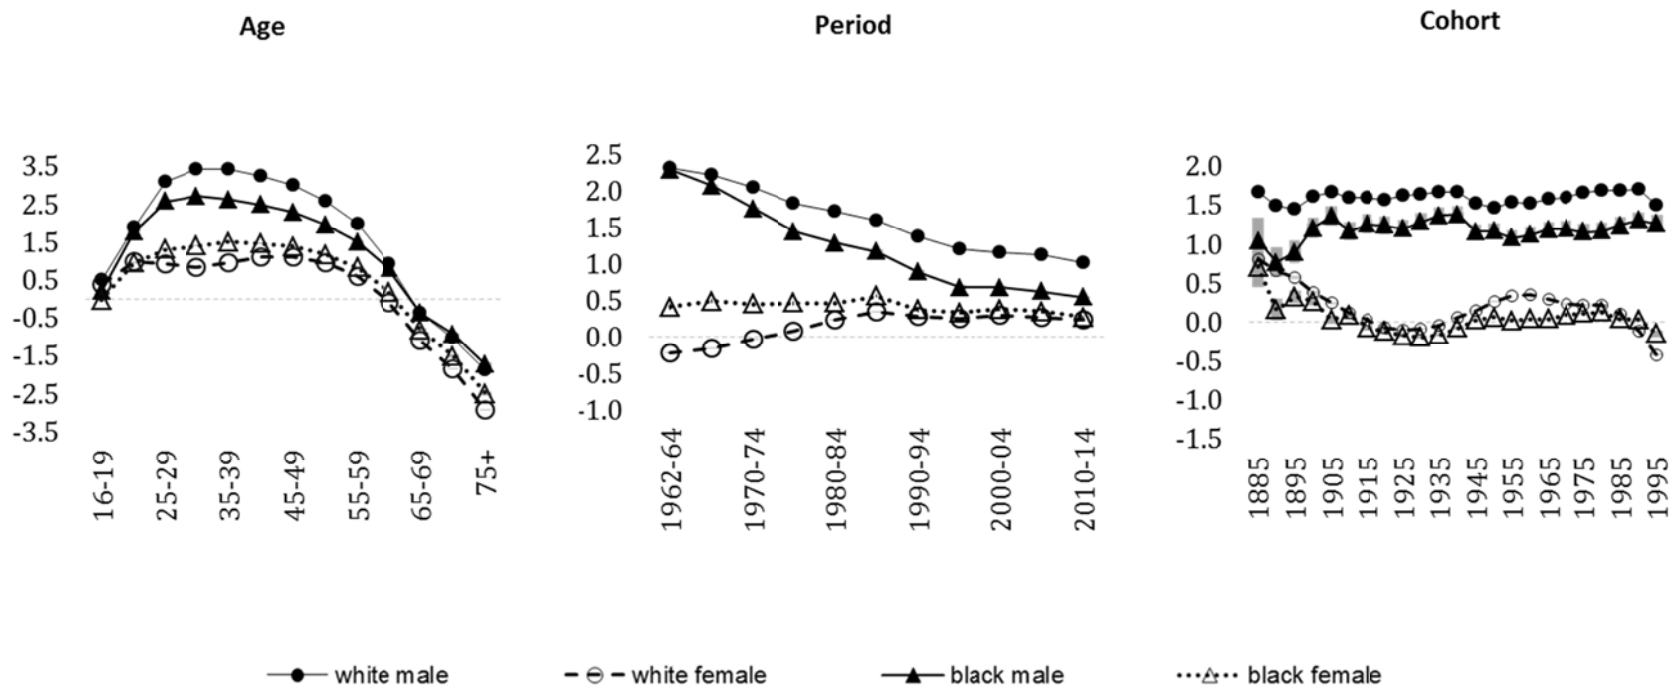


Figure 5.2. Age and Period Main Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014

Note: Age and period main effects estimates derived from Models 1b through 4b plus the intercept for each gender and race group in Tables 5.2. Cohort deviation estimates derived from the “Average deviation” columns in Table 5.5 plus the intercept estimates in Model *b* in Table 5.2. Values represent REML coefficients coded to sum to zero. Grey bars depict 95% confidence intervals. Analysis includes CPS respondents who participated in the 1962 through 2014 CPS surveys for whom labor force participation status and age are available.

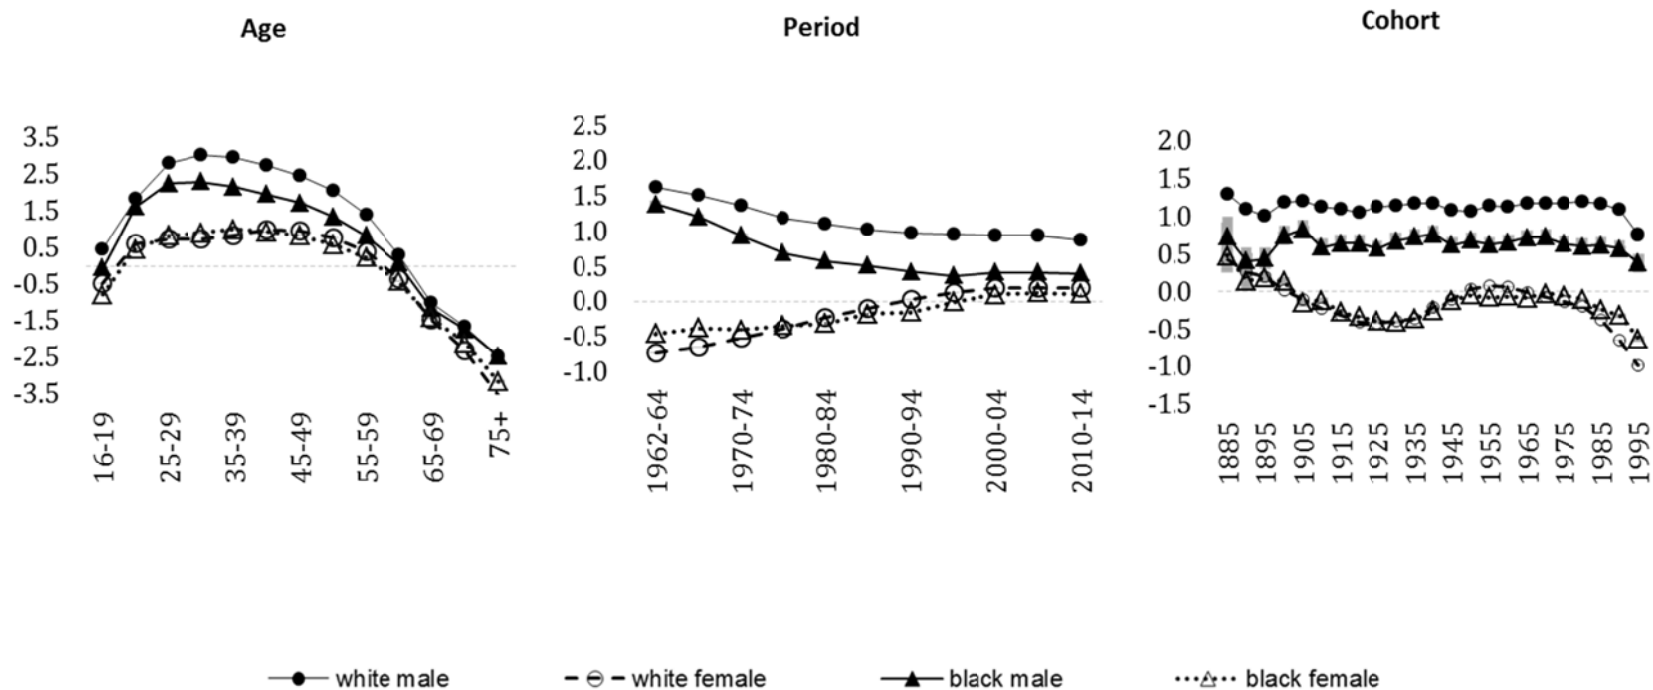


Figure 5.3. Age and Period Effects and Cohort Deviations in Labor Force Participation, Current Population Survey, March Supplement, 1962-2014

Note: Age and period main effects estimates derived from Models 1c through 4c plus the intercept for each gender and race group in Tables 5.2. Cohort deviation estimates derived from the “Average deviation” column in Table 5.5 plus the intercept estimates in Model c in Table 5.2. Values represent REML coefficients coded to sum to zero. Grey bars depict 95% confidence intervals. Analysis includes CPS respondents who participated in the 1962 through 2014 CPS surveys for whom labor force participation status and age are available.

To examine these questions, my strategy is to begin with each Model *a* in Table 5.2 and then add—in separate analyses—measures for educational attainment and marital status. In each case, I ask how the age, period, and cohort patterns noted in the models in Tables 5.2 and 5.5 are changed by holding constant the value of each factor. If I find, for example, that age-by-period interactions are no longer present after adjusting for educational attainment, then I will conclude that the cohort patterns noted above are due to changes over time in educational attainment.

Models *b* and *c* in Tables 5.2 and 5.5 report—Figures 5.2 and 5.3 illustrate—how the age and period effects and cohort deviations for white and black men’s and women’s LFP were associated with changes in educational attainment and marital status, respectively. Except for white women, people who were not currently married were more likely to participate in the labor force. However, the shift in marriage did not appear to drive the age, period, and cohort trends for any of the gender and race group, although it seems to attenuate the magnitude of age, period, and cohort effects on LFP for men and, to a lesser degree, for women.

On the one hand, as in previous research, the results show that education is positively related to LFP for all four groups. On the other hand, the change in educational attainment had differential effects on age, period, and cohort patterns for different groups. Specifically, after adjusting for the amount of formal schooling, the decline in LFP among white and black men became more pronounced, suggesting that had educational attainment not increased in the past years, men’s LFP would have declined more severely. Most interestingly, controlling for educational attainment for

black women resulted in a basically flat and even slightly negative period trend. The cohort deviations were also much less pronounced for all cohorts. It suggests that the change in black women's LFP may be largely explained by the increase in the amount of education that they have received. However, although educational attainment also "explained away" some of the period and cohort related differences in white women's participation, the upward trend across periods and fluctuation across cohorts remained, suggesting that other factors may also contribute to white women's LFP. In future research, I will explore other social and demographic factors that may give rise to the temporal variation in white women's participation.

Conclusion and Discussion

In this chapter, I examined age and period patterns and cohort deviations in labor force participation (LFP) using the 1962-2014 Current Population Survey March Supplement data. The descriptive results of the age, period, and cohort pattern in LFP suggest that men's LFP can be largely described by age and period main effects. This finding indicates that while men's LFP was sensitive to social and economic events such as economic recessions and wars, the effects of these events may not carry on to their later ages. However, there are substantial variations in women's LFP associated with cohort that cannot be explained by pure age and period main effects. In particular, the 1905-1945 and 1980 and younger cohorts of women, relative to their ages and periods, were less likely to participate in the labor force than other cohorts. That is, actual LFP rates of these cohorts were lower than that determined by the age groups and time periods that they have experienced. I speculate that the lower-than-expected participation rates of

older cohorts may be related to the lasting influences of the Great Recession and World War II that they experienced during their youth. For the younger cohorts, because there is no information about their participation beyond age 35, the lower-than-expected LFP rates may simply reflect their longer time in school and consequently delayed entrance to the labor force.

I also explored the ways in which changes in educational attainment and marital status have affected LFP for white and black men and women. The results lent support for Hypothesis 1-1 about the protective effects of education on men's LFP rates. The magnitude of the education-adjusted decline was larger among black men than white men; that is, if educational attainment had not increased among black men, they would have been even less likely to participate in the labor force than we have observed.

There was also supporting evidence for Hypothesis 1-2 about education effects on women's participation in the labor force: while white women's LFP rates did catch up with and exceed those of black women, after adjusting for educational attainment, the race/ethnic differences in participation rates between these two groups seemed to gradually narrow until the late 1980s when their rates became very close.

At the same time, a great deal of the period trend and cohort deviations in black women's LFP can be explained by changes in their educational attainment. This is less true for white women; the cohort deviations in participation rates remained after adjusting for education. These results are consistent with prior research about the complexity of the temporal variation in white women's LFP.

Surprisingly, the results supported neither Hypothesis 2-1 or 2-2 about the relationship between marital status and LFP; the shape of the age, period, and cohort patterns in LFP did not appear to change significantly after controlling for current marital status. This finding suggests that the temporal variation in LFP may stem from changes in the behaviors of subgroups of the population other than changes in the marriage composition of the population.

CHAPTER 6 : CONCLUSION AND FUTURE RESEARCH

In Chapters 2 and 3, I discussed the statistical problems, i.e., biasedness and inconsistency, of the intrinsic estimator (IE) and Hierarchical APC (HAPC) methods for the age-period-cohort (APC) accounting model. However, the more fundamental issue is that even if they could produce unbiased or consistent estimates of the parameters in the APC accounting model, it is unclear what these estimates mean and how researchers should interpret them. Specifically, the APC accounting model assumes that cohort membership can have an additive effect, independent of age and period effects, so the interpretation of the estimated cohort effects should be the “net” and “pure” effects associated with cohort membership after controlling for age and period. However, this interpretation does not make sense, because there is no variation in cohort when age and period are fixed. In this sense, the identification problem is a “blessing” because it warns analysts that “a purely statistical approach to the problem is bound to fail.” (Heckman and Robb 1985: 144-5).

In Chapters 4 and 5, I painted a less-bleak picture for APC research. Specifically, I proposed a new APC method, the APC-I model, and illustrated how it can be used in an example of white and black men’s and women’s labor force participation. The APC-I model is fundamentally different from the accounting model in that it does not estimate “independent, additive” cohort effects net of age and period. Whether one considers this method a better alternative to others depending on the specific research question and/or data structure, the idea of characterizing cohort effects as the age-by-period interaction is

founded on the idea that specification and development of statistical models must be informed by theory.

The APC-I model include age and period main effects and consider cohort effects as the interaction between age and period. One may argue that this specification of cohort effects does not estimate the “linear cohort effects” in the APC accounting model (1) and is therefore deficient. However, the APC-I model is not intended to recover the age, period, and cohort effects in the APC accounting model or to solve its identification problem. I stress that the cohort effects estimated in the APC-I model naturally differ from those in the accounting model because the type of cohort effects operationalized in the APC accounting model departs from the sociological idea of what cohort effects are and when such effects can be observed. As I discussed earlier, Fienberg and Mason (1985) recognized the relationship between cohort effects and the age-by-period interactions, and they designed the APC accounting model with the intention to describe cohort effects as a particular form for the age-by-period interaction. In this sense, the APC-I model can be viewed a renewed effort to describe cohort effects as age-by-period interactions.

In Chapter 5, I used the individual-level Current Population Survey data to demonstrate how to apply the APC-I model in empirical research. The APC-I model can also be used for aggregated data such as mortality, crime, and disease rates. However, for such aggregated data, the $(a - 1) \cdot (p - 1)$ interactions between age and period are completely confounded with the error term because when arranged in the form of age-by-period cross-classification there is only one observation, i.e., no replication, per cell (e.g.,

mortality rate for age 80 in 2000). For binary and count data, one can test the interaction even though the model including it is saturated.²⁷ For continuous data modeled as normally distributed, no replication does indeed create a problem. One possible route is to perform the global and deviation magnitude F tests in Steps 1 and 2 using techniques such as Tukey's test of additivity (Tukey 1949). Another possibility is to conduct deviation magnitude F tests by adding each cohort to the main-effects-only model; that is, testing whether the overall variation associated with that cohort is significantly large compared to other cohorts in aggregate. Of course, such tests are only a partial solution to the problem of one degree of freedom, but it enables researchers to investigate whether cohort effects exist by detecting departures from a model in which the effects of age and periods are additive, i.e., not dependent on each other.

Many APC studies, including my analysis of labor force participation, have used "convenient" birth cohorts, cohorts whose memberships are not constructed based on theoretical account or prior knowledge about the unique experience of a group of people but determined by age and period intervals. Such cohorts might not experience distinctive social changes during their critical ages, so they may not be considered a cohort in a substantive sense. With data that have finer, e.g., one-year or two two-year age and period intervals, the researcher may construct more meaningful cohorts by drawing cohort boundaries based on prior knowledge about a cohort's distinctive experience. The effects of such cohorts are thus represented by age-by-period interaction terms that lie on more than one diagonal in the age-by-period table. As such, the cohort

²⁷ this is only possible because the binomial and Poisson models—the default for count data—make strong assumptions about how the variance is related to the mean

effects can be investigated by applying the three-step procedure to age-by-period interaction terms that lie on these multiple diagonals of age-by-period interaction terms.

As I have shown in Chapter 5, the APC-I model can be used to examine temporal trends in the outcome by modeling the ways in which possible explanatory factors affect these trends. To this extent, this method echoes the ideas behind the proxy variable approach (Heckman and Robb 1985), the APC-Characteristics model (O'Brien 2000), and the Mechanism-Based model (Winship and Harding 2008), which specify the theoretical mechanisms through which age, period, and/or cohort affects the outcome. The importance of theoretical thinking in informing model specification and interpretation cannot be overstated. In their insightful article, Fienberg and Mason (1985) encouraged researchers to “begin with conceptualization and attempt to move toward explicit measurement, in order to test understanding of the interaction.” (p. 83) To the extent that the APC-I model is explicitly tied to the conceptualization of cohort effects in sociological and demographic literatures, I believe that the APC-I model is promising in advancing APC research.

Lastly, although the APC-I model is designed for the APC analysis, the conceptual critiques and methodological ideas can be extended to many other fields in which focal explanatory variables are exactly related. For example, scholars of status inconsistency study the likelihood of a person attaining higher or lower socioeconomic status than their parents and the consequences of changes in status for various outcomes including happiness, marriage, and health conditions. Also, researchers of assortative mating are interested in how marriage forms between persons of the same or different

levels of educational attainment, and the implications of such educational homogeneity or heterogeneity for marriage duration, life satisfaction, and other economic and health well-beings. Despite long-standing interest in these areas among sociologists and demographers, these lines of scholarship suffer from effectively the same methodological problem as APC analysis: the third variable is completely determined by the other two. Specifically, in status inconsistency studies, status inconsistency equals adult socioeconomic status minus status of their parents; in educational homogamy research, educational difference equals husband's education minus wife's education. Several methods have been developed to address this estimation problem, but none is satisfactory from a statistical point of view (Hope 1975; Houle 2011; Sobel 1981). The APC-I model developed in this chapter can potentially be modified to address these important sociological issues.

It is puzzling that while the literature from the 1970s to the 1990s emphasized the importance of theoretical information, this tradition has given way to pure technical solutions such as the IE and HAPC. I hope my research can raise awareness about the ultimately important conversation between method and theory. Bearing the important role of theory in statistical modeling, I will explore interesting methodological and substantive issues in other lines of inquiry including mortality selection, gender inequality in health and labor market participation, and assortative mating.

REFERENCES

- Aiken, Leona S. and Stephen G. West. 1991. *Multiple Regression: Testing and Interpreting Interactions*. Newbury Park, London: SAGE Publications.
- Alwin, Duane F. 1991. "Family of Origin and Cohort Differences in Verbal Ability." *American Sociological Review* 56 (5): 625-638.
- Bell, Andrew and Kelvyn Jones. 2013a. "The Impossibility of Separating Age, Period and Cohort Effects." *Social Science & Medicine* 93: 163-165.
- . 2013b. "Bayesian Informative Priors with Yang and Land's Hierarchical Age-Period-Cohort Model." *Quality & Quantity* DOI 10.1007/s11135-013-9985-3
- . 2014a. "Don't Birth Cohorts Matter? A commentary and simulation exercise on Reither, Hauser, and Yang's (2009) age-period-cohort study of obesity." *Social Science & Medicine* 101: 176-180.
- . 2014b. "Another 'futile quest'? A simulation study of Yang and Land's Hierarchical Age-Period-Cohort model." *Demographic Research* 30:333-360.
- Bianchi, Suzanne M. 2011. "Changing Families, Changing Workplaces." *The Future of Children* 21(2):15-36.
- Borella, Erika, Paolo Ghisletta, and Anik de Ribaupierre. 2011. "Age Differences in Text Processing: The Role of Working Memory, Inhibition, and Processing Speed." *The Journals of Gerontology, Series B: Psychological Sciences and Social Sciences* 66(3):311-320.
- Boushey, Heather. 2005. "Are Women Opting Out? Debunking the Myth." Center for Economic and Policy Research, Briefing Paper.
- Bradbury, Katharine and Jane Katz. 2005. "Women's Rise—A Work in Progress: Are Professional Women Opting Out? Recent Evidence on College-Educated Women's Labor Force Participation." *Regional Review* 14(3):58-67.
- Boyle, Peter, Nicholas E. Day, and Knut Magnus. 1983. "Mathematical Modelling of Malignant Melanoma Trends in Norway, 1953-1978." *American Journal of Epidemiology* 118 (6):887-896.
- Boyle, Peter, and Chris Robertson. 1987. "Statistical Modelling of Lung Cancer and Laryngeal Cancer Incidence in Scotland, 1960-1979." *American Journal of Epidemiology* 125 (4):731-744.

Bureau of Labor Statistics, U.S. Department of Labor. 2014. *Women in the Labor Force: A Databook*.

Chaves, Mark. 1989. "Secularization and Religious Revival: Evidence from U.S. Church Attendance Rates, 1972-1986." *Journal for the Scientific Study of Religion* 28 (4): 464-477.

Clayton, D., and E. Schiffers. 1987. "Models for Temporal Variation in Cancer Rates. II: Age-Period-Cohort Models." *Statistics in Medicine* 6 (4):469-481.

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

Connell, Rachel. 1992. "The Effect of Child Care Costs on Married Women's Labor Force Participation." *The Review of Economics and Statistics* 74(1): 83-90.

Cole, Tim J., Chris Power, and Gudrun E. Moore. 2008. "Intergenerational Obesity Involves both the Father and the Mother." *The American Journal of Clinical Nutrition* 87 (5):1535-1636.

Dannefer, Dale. 1987. "Aging as Intracohort Differentiation: Accentuation, the Matthew Effect, and the Life Course." *Sociological Forum* 2 (2):211-236.

———. 2003. "Cumulative Advantage/Disadvantage and the Life Course: Cross-Fertilizing Age and Social Science Theory." *The Journals of Gerontology. Series B, Psychological Sciences and Social Sciences* 58 (6):S327-S337.

DiPrete, Thomas, and Gregory Eirich. 2006. "Cumulative Advantage as a Mechanism for Inequality: A Review of Theoretical and Empirical Developments." *Annual Review of Sociology* 32:271-297.

Elder, Glen. H. 1975. "Age Differentiation and the Life Course". *Annual Review of Sociology* 1:165-190.

Fabio, Anthony, Rolf Leober, G. K. Balasubramani, Jeffrey Roth, Wenjiang Fu, and David P. Farrington. 2006. "Why some Generations are More Violent than Others: Assessment of Age, Period, and Cohort Effects." *American Journal of Epidemiology* 164 (2):151-160.

Farkas, George. 1977. "Cohort, Age, and Period Effects upon the Employment of White Females: Evidence for 1957-1968." *Demography* 14(1):22-42.

- Fernández, Raquel. 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review* 2013, 103(1): 472–500.
- Fienberg, Stephen E. 2013. "Cohort Analysis' Unholy Quest: A Discussion." *Demography* 50 (6): 1981-1984.
- Fienberg, Stephen E., James Hodges, and Liying Luo. 2015. "Letter to the Editor: Comment on Review of Age-Period-Cohort Analysis: New Models, Methods, and Empirical Applications by Y. Yang and K.C. Land." *Journal of the American Statistical Association* 110 (509):457.
- 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.
- . 1985. "Specification and implementation of age, period, and cohort models." In: William M. Mason and Stephen E. Fienberg (eds.) *Cohort Analysis in Social Research: Beyond the Identification Problem*. New York: Springer-Verlag, pp. 44–88.
- Firebaugh, Glenn and Brian Harley. 1991. "Trends in U.S. Church Attendance: Secularization and revival, or merely lifecycle effects?" *Journal for the Scientific Study of Religion* 30:487–500.
- Fischer, Claude S., and Michael Hout. 2006. *Century of Difference: How America Changed in the Last One Hundred Years*. New York: Russell Sage Foundation.
- Flegal, Katherine M., Margaret Carroll, Cynthia Ogden, Clifford L. Johnson. 2002. "Prevalence and Trends in Obesity among US Adults, 1999-2000." *JAMA (Chicago, Ill.)* 288 (14):1723.
- Fu, Wenjiang. 2000. "Ridge Estimator in Singular Design with Applications to Age-Period-Cohort Analysis of Disease Rates." *Communications in Statistics Theory and Method* 29:263-78.
- Fu, Wenjiang J., and Peter Hall. 2006. "Asymptotic Properties of Estimators in Age-Period-Cohort Analysis." *Statistics & Probability Letters* 76 (17):1925-1929.
- Fu, Wenjiang J. 2008. "A Smoothing Cohort Model in Age-Period-Cohort Analysis with Applications to Homicide Arrest Rates and Lung Cancer Mortality Rates." *Sociological Methods & Research* 36 (3):327-361.
- Fu, Wenjiang J., Kenneth C. Land, and Yang Yang. 2011. "On the Intrinsic Estimator and Constrained Estimators in Age-Period-Cohort Models." *Sociological Methods & Research* 40 (3):453-466.

Gillman, Matthew W. 2004. "A Life Course Approach to Obesity." Pp. 189-217 in *A Life Course Approach to Chronic Disease Epidemiology* edited by D. Kuh and Y. Ben-Shlomo. Oxford: Oxford University Press.

Glenn, Norval D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period and Cohort Effects." *American Sociological Review* 41 (5):900-904.

———. 2005. *Cohort Analysis*. Thousand Oaks, Calif.: Sage Publications.

Goldin, Claudia D. 1989. "Life-Cycle Labor Force Participation by Married Women: Historical Evidence and Implications." *Journal of Labor Economics* 7:20-47.

———. 1990. *Understanding the Gender Gap : An Economic History of American Women*. Oxford University Press.

———. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review* 96 (2):1-21.

Heckman, James and Richard Robb. 1985. "Using Longitudinal Data to Estimate Age, Period, and Cohort Effects in Earnings Equations." Pp. 137-50 in *Cohort Analysis in Social Research*, edited by W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.

Hobcraft, John, Jane Menken, and Samuel Preston. 1982. "Age, Period, and Cohort Effects in Demography: A Review." *Population Index* (48)1:4-43.

Hodges, James S. 1998, "Some Algebra and Geometry for Hierarchical Models, Applied to Diagnostics (with discussion)", *Journal of the Royal Statistical Society, Series B*, 60:497-536.

Hodges, James S. 2013. *Richly Parameterized Linear Models: Additive, Time Series, and Spatial Models Using Random Effects*. Boca Raton, Florida: Chapman and Hall.

Holford, Theodore R. 1983. "The Estimation of Age, Period and Cohort Effects for Vital Rates." *Biometrics* 39 (2):311-324.

———. 1992. "Analysing the Temporal Effects of Age, Period and Cohort." *Statistical Methods in Medical Research* 1(3):317-337.

———. 2006. "Approaches to Fitting Age-Period-Cohort Models with Unequal Intervals." *Statistics in Medicine* 25 (6):977-993.

- Hollister, Matissa N., and Kristin E. Smith. 2014. "Unmasking the Conflicting Trends in Job Tenure by Gender in the United States, 1983-2008." *American Sociological Review* 71(1):159-181.
- Hope, Keith. 1975. "Models of Status Inconsistency and Social Mobility Effects." *American Sociological Review* 40(3):322-343.
- Houle, Jason N. 2011. "The Psychological Impact of Intragenerational Social Class Mobility." *Social Science Research* 40(3): 757-772.
- Jaccard, James J. and Robert Turrisi. 2003. *Interaction Effects in Multiple Regression (Quantitative Applications in the Social Sciences)*. Thousand Oaks, Calif.: Sage Publications.
- Kennedy, Sheela, and Steven Ruggles. 2014. "Breaking Up Is Hard to Count: The Rise of Divorce in the United States, 1980-2010." *Demography* 51(2):87-598.
- Keyes, Katherine, and Richard Miech. 2013. "Age, Period, and Cohort Effects in Heavy Episodic Drinking in the US from 1985 to 2009." *Drug and Alcohol Dependence*. doi: 10.1016/j.drugalcdep.2013.01.019.
- Kuhn, Thomas S. 1996. *The Structure of Scientific Revolutions*. 3rd ed. Chicago: University of Chicago Press.
- Kupper, Lawrence L., Joseph Janis, Ibrahim A. Salama, Carl N. Yoshizawa, Bernard G. Greenberg, and H. H. Winsborough. 1983. "Age-period-cohort analysis: an illustration of the problems in assessing interaction in one observation per cell data." *Communications in Statistics.Theory and Methods* 12 (23): 2779-2807.
- Kupper, Lawrence L., Joseph Janis, Azza Karmous, and Bernard G. Greenberg. 1985. "Statistical Age-Period-Cohort Analysis: A Review and Critique." *Journal of Chronic Diseases* 38 (10):811-830.
- Langley, John., Ari Samaranayaka, J. Davie, and A. J. Campbell. 2011. "Age, Cohort and Period Effects on Hip Fracture Incidence: Analysis and Predictions from New Zealand Data 1974-2007." *Osteoporosis International* 22 (1):105-111.
- Luo, Liying. 2013a. "Assessing Validity and Application Scope of the Intrinsic Estimator Approach to the Age-Period-Cohort Problem." *Demography* 50 (6): 1945-1967.
- . 2013b. "Paradigm Shift in Age-Period-Cohort Analysis: A Response to O'Brien, Fienberg, Held and Riebler, and Yang and Land." *Demography* 50 (6): 1985-1988.

- Luo, Liying, and James Hodges. Forthcoming. "Block Constraints in Age-Period-Cohort Models with Unequal-Width Intervals." *Sociological Methods & Research*.
- Luo, Liying, James S. Hodges, Christopher Winship, and Daniel Powers. 2015. "The Sensitivity of the Intrinsic Estimator to Coding Schemes: A Comment on Yang, Schulhofer-Wohl, Fu, and Land." Conditionally accepted. *American Journal of Sociology*.
- Lynch, John, George Davey Smith, Sam Harper, Marianne Hillemeier, Nancy Ross, George A. Kaplan, and Michael Wolfson. 2004. "Is Income Inequality a Determinant of Population Health? Part 2. U.S. National and Regional Trends in Income Inequality and Age- and Cause-Specific Mortality." *The Milbank Quarterly* 82 (2):355-400.
- Macunovich, Diane J. 2012. "Relative Cohort Size, Relative Income, and Married Women's Labor Force Participation: United States, 1968–2010." *Population and Development Review* 38(4):631-648.
- Mason, Karen Oppenheim, William M. Mason, H.H. Winsborough, W. Kenneth Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38 (2):242-258.
- Masters, Ryan K. 2012. "Uncrossing the U.S Black-White Mortality Crossover: The Role of Cohort Forces in Life Course Mortality Risk." *Demography* 49 (3):773-796.
- Masters, Ryan K., Robert A. Hummer, and Daniel A. Powers. 2012. "Educational Differences in U.S. Adult Mortality." *American Sociological Review* 77 (4):548-572.
- Miech, Richard, Steve Koester, and Brook Dorsey-Holliman. 2011. "Increasing US Mortality due to Accidental Poisoning: The Role of the Baby Boom Cohort." *Addiction* 106 (4):806-815.
- Miriam King, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick. 2010. Integrated Public Use Microdata Series, Current Population Survey: Version 3.0. [Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor].
- Mokdad, Ali H., Earl S. Ford, Barbara A. Bowman, William H. Dietz, Frank Vinicor, Virginia S. Bales, and James S. Marks. 2003. "Prevalence of Obesity, Diabetes, and Obesity-Related Health Risk Factors, 2001." *JAMA (Chicago, Ill.)* 289 (1):76-79.
- O'Brien, Robert. M. 2000. "Age Period Cohort Characteristic Models." *Social Science Research* 29 (1):123.

- . 2011a. "Constrained Estimators and Age-Period-Cohort Models." *Sociological Methods & Research* 40 (3):419-452.
- . 2011b. "Intrinsic Estimators as Constrained Estimators in Age-Period-Cohort Accounting Models." *Sociological Methods & Research* 40 (3):467-470.
- . 2014. "Age-Period-Cohort Models: Approaches and Analyses with Aggregate-Level Data." New York: Chapman & Hall.
- O'Brien, Robert M, Kenneth Hudson, and Jean Stockard. 2008. "A Mixed Model Estimation of Age, Period, and Cohort Effects" *Sociological Methods & Research* 36 (3):402-428.
- Ogden, Cynthia L. 2006. "Prevalence of overweight and obesity in the United States, 1999-2004." *JAMA (Chicago, Ill.)* 295 (13):1549-55.
- Osmond, Clive, and M. J. Gardner. 1989. "Age, Period, and Cohort Models. Non-Overlapping Cohorts Don't Resolve the Identification Problem." *American Journal of Epidemiology* 129 (1):31-35.
- Pampel, Fred C. and Lori M. Hunter. 2012. "Cohort Change, Diffusion, and Support for Environmental Spending in the United States." *American Sociological Review* 118(2): 420-448.
- Peng, Xizhe. "Demographic Consequences of the Great Leap Forward in China's Provinces." *Population and Development Review* 13(4): 639-670.
- Pelzer, Ben, Manfred te Grotenhuis, Rob Eisinga, and Alexander W. Schmidt-Catran. Forthcoming. "The Non-uniqueness Property of the Intrinsic Estimator in APC Models." *Demography* , 52(1): 315-327.
- Percheski, Christine. 2008. "Opting Out? Cohort Differences in Professional Women's Employment Rates from 1960 to 2005." *American Sociological Review* 73:497-517.
- Preston, Samuel H., and Haidong Wang. 2006. "Sex Mortality Differences in the United States: The Role of Cohort Smoking Patterns." *Demography* 43 (4):631-646.
- Raftery, Adrian E., Steven Lewis, and Akbar Aghajanian. 1995. "Demand or ideation? Evidence from the Iranian marital fertility decline." *Demography* 32 (2):159-182.
- Reither, Eric N., Robert M. Hauser, and Yang Yang. 2009. "Do Birth Cohorts Matter? Age-Period-Cohort Analyses of the Obesity Epidemic in the United States." *Social Science & Medicine* 69 (10):1439-1448.

- Riley, Matilda White, 1987. "On the Significance of Age in Sociology." *American Sociological Review* (52)1:1-14.
- Robertson, Chris, and Peter Boyle. 1986. "Age, Period and Cohort Models: The use of Individual Records." *Statistics in Medicine* 5 (5):527-538.
- . 1998. "Age–period–cohort Analysis of Chronic Disease Rates. I: Modelling Approach." *Statistics in Medicine* 17 (12):1305.
- Rodgers, Willard L. 1982a. "Estimable Functions of Age, Period, and Cohort Effects." *American Sociological Review* 47 (6):pp. 774-787.
- . 1982b. "Reply to Comment by Smith, Mason, and Fienberg." *American Sociological Review* 47 (6):793-796.
- Ryder, Norman B. 1965. "The Cohort as a Concept in the Study of Social Change." *American Sociological Review* 30 (6): 843-861.
- Sarma, Sisira, Rosa A. Devlin, Amardeep Thind, and Man-Kee Chu. 2012. "Canadian Family Physicians' Decision to Collaborate: Age, Period and Cohort Effects." *Social Science & Medicine* (1982) 75 (10):1811-1819.
- Sarma, Sisira, Amardeep Thind, and Man-Kee Chu. 2011. "Do New Cohorts of Family Physicians Work Less Compared to their Older Predecessors? The Evidence from Canada." *Social Science & Medicine* 72 (12):2049-2058.
- Sassi, Franco, Marion Devaux, Michele Cecchini, Elena Rusticelli. 2009. "The Obesity Epidemic: Analysis of Past and Projected Future Trends in Selected OECD Countries", *OECD Health Working Papers* 45: OECD Publishing doi: 10.1787/225215402672.
- Scheffé, Henry. 1959. *The Analysis of Variance*. New York: John Wiley and Sons.
- Schwadel, Philip. 2010. "Age, Period, and Cohort Effects on U.S. Religious Service Attendance: The Declining Impact of Sex, Southern Residence, and Catholic Affiliation." *Sociology of Religion* 71 (1):2-24.
- . 2011. "Age, Period, and Cohort Effects on Religious Activities and Beliefs." *Social Science Research* 40 (1):181-192.
- Smith, Tim. 1990. "Liberal and Conservative Trends in the United States Since World War II." *The Public Opinion Quarterly* 54(4):479-507

- Sobel, Michael E. 1981. "Diagonal Mobility Models: A Substantively Motivated Class of Designs for the Analysis of Mobility Effects." *American Sociological Review* 46 (6): 893-906.
- Tango, Toshiro. 1988. "RE: "Statistical Modelling of Lung Cancer and Laryngeal Cancer Incidence in Scotland, 1960–1979."” *American Journal of Epidemiology* 128 (3):677-678.
- Treas, Judith. 1987. "The Effect of Women’s Labor Force participation on the Distribution of Income in the United States." *Annual Review of Sociology* 13:259-288.
- Tukey, John W. 1949. "Comparing Individual Means in the Analysis of Variance." *Biometrics* 5(2):99-114.
- Wilson, James A., and Walter R. Gove. 1999. "The Intercohort Decline in Verbal Ability: Does it Exist?" *American Sociological Review* 64(2):253-266.
- Winkler, Richelle, and Keith Warnke. 2012. "The future of hunting: an age-period-cohort analysis of deer hunter decline." *Population Environment*: Springer doi:10.1007/s11111-012-0172-6
- Winship, Christopher, and David J. Harding. 2008. "A Mechanism-Based Approach to the Identification of Age--Period--Cohort Models." *Sociological Methods & Research* 36 (3): 362-401.
- Yang, Yang. 2008. "Trends in U.S. Adult Chronic Disease Mortality, 1960-1999: Age, Period, and Cohort Variations." *Demography* 45 (2):387-416.
- 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.
- . 2008. "Age–Period–Cohort Analysis of Repeated Cross-Section Surveys: Fixed Or Random Effects?" *Sociological Methods & Research* 36 (3):297-326.
- Yang, Yang, Wenjiang J. Fu, and Kenneth C. Land. 2004. "A Methodological Comparison of Age-Period-Cohort Models: The Intrinsic Estimator and Conventional Generalized Linear Models." *Sociological Methodology* 34 (1):75-110.
- Yang, Yang, Sam Schulhofer-Wohl, Wenjiang J. Fu, 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 (6):1697-1736.

Zheng, Hui, Yang Yang, and Kenneth C. Land. 2011. "Variance Function Regression in Hierarchical Age-Period-Cohort Models: Applications to the Study of Self-Reported Health." *American Sociological Review* 76(6): 955-983.