A Simulation Study of the Intrinsic Estimator for Age-Period-Cohort Analysis

Yang Yang,[*] Sam Schulhofer-Wohl, [**] and Kenneth C. Land[***]

[*]Department of Sociology, Population Research Center, and Center on Demography and Economics of Aging at NORC, The University of Chicago

[**]Department of Economics and Woodrow Wilson School of Public and International Affairs, Princeton University

[***]Department of Sociology and Center for Population Health and Aging, Duke University

Nov 5, 2007

Paper presented at the Methodology Paper Session at the Annual Meetings of the American Sociological Association at New York, NY, August, 2007. Address all correspondence to: Yang Yang, Department of Sociology, University of Chicago, 1126 E. 59th St., Chicago, IL 60637; Phone: (773) 834-1113, Fax: (773) 702-4849, e-mail: yangy@uchicago.edu

A Simulation Study of the Intrinsic Estimator

for Age-Period-Cohort Analysis

Abstract

A new approach to the statistical estimation of Age-Period-Cohort (APC) accounting models for tabular data, called the Intrinsic Estimator (IE), recently has been developed. Some finite sampling properties of the IE have been proven for inferences based on a fixed number of time periods of data. Asymptotic properties as the number of time periods of data grows without bound also have been studied. To provide further exposition, especially using straightforward and replicable numerical illustrations, this paper presents results of simulation studies of properties of the IE. Generally, the results show that the IE performs well as a statistical estimator under most conditions. We also find that some long-standing critiques of the utility of APC accounting models rest on problematic logic and misuse of models. The results of the simulation studies can inform empirical studies that use the IE for statistical estimation of APC models.

Word Count: 144

Key Words: Age-period-cohort analysis; identification problem; Intrinsic Estimator; generalized inverse; estimability; Monte Carlo simulation

Introduction

Age-period-cohort (APC) analysis has played a critical role in studying time-specific phenomena in the social sciences, including economics, demography, sociology, and political science, for the past 80 years. Broadly defined, APC analysis distinguishes three types of time-related variation: age, period, and cohort effects (Hobcraft, Menken, and Preston 1982). Such distinctions have important implications for measurement and analysis in various phenomena of interest such as consumption expenditures, fertility and mortality rates, test scores, and voting for Presidential candidates. One common goal of APC analysis is to assess the effects of one of the three factors on some outcomes of interest net of the influences of the other two time-related dimensions. The fundamental question of determining whether the process under study is some combination of age, period, and cohort phenomena points to the necessity of statistically estimating and delineating the age, period, and cohort effects (Kupper et al. 1985).

The conventional statistical method for modeling APC data in the form of rectangular age-by-time period tables is the *Age-Period-Cohort accounting/multiple classification model* (Mason et al. 1973). The major challenge of estimating separate age, period, and cohort effects using the model is the "identification problem" induced by the exact linear dependency between age, period, and cohort: period - age = cohort. This results in an infinite number of solutions that fit the data equally well. It follows that one must assign certain additional identifying constraints to obtain unique estimates of these effects. The most widely used approach to solving this problem is to place at least one equality constraint on two or more of the age, period, or cohort coefficients (e.g., Mason et al. 1973; Fienberg and Mason 1985). This has been referred to as the constrained generalized linear models (CGLIM) approach (Yang, Fu, and Land 2004). For example, one can constrain the effect coefficients of two adjacent age groups, periods, or cohorts to be equal to identify the model (see, e.g., Mason and Smith 1985; Yang et al. 2004). The main criticisms of this approach and its variants are that: 1) different equality constraints yield different effect coefficient estimates but identical model fit; and 2) estimates of the effect coefficients and thus of the patterns of change across the age, period, and cohort dimensions are sensitive to the choice of the identifying constraints, which depends on strong prior or external information that rarely exists or can be well verified (Mason and Wolfinger 2002; Robertson et al. 1999).

Recently, a new approach to estimating the APC multiple classification models was described and evaluated by Yang and colleagues (2004). Incorporating methodological developments using estimable functions in biostatistics, the new method of estimation termed the *intrinsic estimator* (IE) yields a unique solution to the model that is determined by the Moore-Penrose generalized inverse. It achieves model identification with minimal assumptions. Given the long history of debates over the existence of any solution to the APC model identification problem (Glenn 2005), critiques can be raised as to whether the IE method is based on a constraint of purely algebraic convenience and just as arbitrary as previous methods of constraints. Is the IE just another way to go wrong? In response, it can be noted that Yang et al. (2004) showed that the IE has several key statistical properties, including estimability, finite time period unbiasedness, relative efficiency, and asymptotic consistency, all of which distinguish the IE from any other estimator. But the conceptual foundations of the IE have been established statistically using mathematical proofs and thus remain abstract and potentially difficult to understand. Further exposition and analysis, especially using straightforward and replicable numerical illustrations, is needed to directly address this question.

This study fills this gap through means of model validation, or numerical assessment of the performance of the IE under various model specifications. The remainder of the paper is organized as follows. First, we briefly review the structure and properties of the IE. Second, we report and compare the numerical results from Monte Carlo simulations that are conducted systematically for the IE and CGLIM estimators of APC multiple classification models. Third, we revisit a recent critique of the utility of APC models in social research through numerical examples (Glenn 2005), show the consequence of misuse and misinterpretation of such models, and make suggestions to future research in light of findings from the simulation analysis.

The IE As A Statistical Estimator

First, the APC multiple classification model can be written in the linear regression form as:

$$Y_{ij} = \mu + \alpha_i + \beta_j + \gamma_k + \varepsilon_{ij} \qquad (1)$$

where Y_{ij} denotes the observed outcome for the *i-th* age group for $i = 1,\ldots,a$ age groups at the *j-th* time period for $j = 1,\ldots, p$ time periods of observed data, μ denotes the intercept or adjusted mean, α_i denotes the *i-th* row age effect or the coefficient for the *i-th* age group, β_j denotes the *j-th* column period effect or the coefficient for the *j-th* time period, γ_k denotes the *k-th* cohort effect or the coefficient for the *k-th* cohort for $k = 1,\ldots,(a+p-1)$ cohorts, with $k=a-i+j$, and ε_{ij} denotes the random errors with expectation $E(\varepsilon_{ij}) = 0$. Note that Eq. (1) generalizes straightforwardly to a Generalized Linear Models (GLM) framework, where it can take various alternative forms such as normal, log-linear, or logistic regression models (McCulloch and Searle 2001). It can be treated as a fixed-effects linear model after a reparametrization to center the parameters so that they sum to zero, i.e.,

$\sum_i \alpha_i = \sum_j \beta_j = \sum_k \gamma_k = 0$.

The key problem in APC analysis using model (1) is the "identification problem". Rewriting model (1) in matrix form, we have:

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

where Y is a vector of observed outcomes, X is the regression design matrix that contains dummy variable column vectors for the model parameter b of dimension $m = 1 + (a - 1) + (p - 1) + (a + p - 2)$, $b = (\mu, \alpha_1, ...\alpha_{a-1}, \beta_1, ..., \beta_{p-1}, \gamma_1, ..., \gamma_{a+p-2})^T$, and ε is a vector of random errors with mean 0 and constant diagonal variance matrix $\sigma^2 I$, with I denoting an identity matrix.[1] The ordinary least squares estimator of model (2) is the solution b of the normal equations: $\hat{b} = (X^T X)^{-1} X^T Y$. The linear relationship between the age, period and cohort variables (period = age + birth year) translates to a design matrix, X, that is one less than full column rank. This implies that $X^T X$ is singular, i.e., the inverse of $X^T X$ does not exist. It follows that solution to normal equations is not unique. Therefore, the model identification problem exists without assigning certain additional identifying constraints.

Since the work of Fienberg and Mason (1978, 1985), the conventional approach to multiple classification APC models in demography has been a *coefficients constraints approach*, which takes the form of placing (at least) one additional identifying constraint on the parameter vector b, e.g., constraining the effect coefficients of the first two periods to be equal, $\beta_1 = \beta_2$. With this one additional constraint, the model (2) is just-identified, the matrix $(X^T X)$ becomes nonsingular, and the least squares estimator exists (as do related maximum likelihood estimators for log-linear or logistic models). The main problem with this CGLIM approach is that the methodological usefulness of the method depends on strong prior information for identifying these restrictions. As has been known at least since the work of Mason and Smith (1985), and as was demonstrated with U. S. female mortality rates by Yang et al. (2004), estimates of model effect coefficients are sensitive to the choice of the identifying constraint. Thus, different choices of just-identifying linear constraint can lead to widely divergent patterns of estimated effect coefficient across the age, period, and cohort categories. This has led to a large methodological literature in demography, epidemiology, and statistics (see, e.g., the review of cohort analysis by Mason and Wolfinger, 2002) advising APC analysts of this sensitivity to model specification and that the choice of the model identifying constraint must be based on prior theoretical or empirical information that, unfortunately, rarely exists. Subsequent developments of APC modeling can largely be considered variants of the APC accounting models using different constraints. Generally speaking, strong assumptions and sensitivities of results to choice of assumptions have largely prohibited reliable and consistent findings to be revealed. And statisticians acknowledge the limitations of existing approaches and conclude that APC analysis is still in its infancy (Kupper et al. 1985; Mason and Wolfinger 2002).

So what is new about the Intrinsic Estimator (IE)? Although it is a consensus that the key problem for APC analysis is to identify an estimable function independent of constraint and uniquely determines the parameter estimates, the controversy continues whether there exists an estimable function that solves the identification problem. The conventional wisdom is that only the nonlinear, but not the linear components of the APC models are estimable (Holford 1983; Rodger 1982). As noted in Fu (2007), however, this proposition has not been supported by mathematical proofs. And it should also be observed that Kupper et al. (1985: Appendix B) mathematically derived a condition that estimable functions must satisfy and stated that an estimable function satisfying this condition resolves the identification problem. *Estimable functions* are invariant with respect to whatever solution is obtained to the normal equations. It has been shown in works that followed that the IE satisfies the Kupper et al. condition and is the unique estimable function of both the linear and nonlinear components for the multiple classification model (Fu 2000, 2007; Yang et al. 2004).

Using Yang et al.'s (2004) notation, the structure and the estimability of the IE can be shown in the following:

1) The exact linear dependency between age, period, and cohort variables in model (2) is mathematically equivalent to

$$XB_0 = 0, \tag{3}$$

which states the property that X is singular (i.e., the column space of X is less than full rank) and the product of X and some nonzero vector B_0 is 0 (see, e.g., Christensen, 2000 for this proposition). Kupper et al. (1985) showed that B_0 has the specific form that is a function of the design matrix. Specifically, $B_0 = \frac{\tilde{B}_0}{\|\tilde{B}_0\|}$, or a normalized vector of $\tilde{B}_0$, where $\tilde{B}_0 = (0, A, P, C)^T$, and $A = \left(1 - \frac{a+1}{2}, \cdots, (a-1) - \frac{a+1}{2}\right)$,

$P = \left(\frac{p+1}{2} - 1, \cdots, \frac{p+1}{2} - (p-1)\right)$, and $C = \left(1 - \frac{a+p}{2}, \cdots, (a+p-2) - \frac{a+p}{2}\right)$. It is important to note from above that the vector B_0 is fixed because it is a function solely of the number of age groups (a) and periods (p). The fact that the fixed vector B_0 is independent of the response variable Y suggests that it should not play any role in the estimation of effect coefficients. But this principle may be violated in the conventional CGLIM approach, as illustrated below.

2) The parameter space of the unconstrained vector b of the linear model (2) can be decomposed into two parts that are orthogonal or independent to each other:

$$b = b_0 + tB_0, \qquad (4)$$

where $b_0 = P_{proj} b$ is a parameter vector that is a linear function of b corresponding to the projection of b to the non-null space of X perpendicular to the null space of X defined by B_0 (the unique normalized eigenvector with norm 1 of the singular design matrix corresponding to the unique eigen value 0 as shown in Eq. (3)) and a multiple of B_0 with t being any real number (please see Figure 1 in Yang et al. [2004] for the geometric projection). The special parameter vector b_0 corresponding to $t = 0$ satisfies the geometric projection:

$$b_0 = (I - B_0 B_0^T)b \qquad (5)$$

3) It is important to note that the IE does not estimate the unconstrained parameter vector b. Rather, the IE estimates the constrained vector b_0. The above decomposition of parameter vector b means that each of the infinite number of possible estimators of parameter vector of model (2), denoted as $\hat{b}$, can be written as a linear combination:

$$\hat{b} = B + tB_0 \qquad (6)$$

where B is the IE that estimates b_0. Different linear constraints on coefficients of the b vector assign different values to t and lead to different estimates. The IE is free of such variation by setting $t = 0$ and can be obtained using the projection: $B = (I - B_0 B_0^T)\hat{b}$ based on (5) or the principal component regression algorithm (see the Appendix A).

4) Combining Eqs. (3) and (6) yields the following result:

$$X\hat{b} = X(B + tB_0) = XB + tXB_0 = XB + 0 = XB \qquad (7)$$

This shows that the IE B is the special estimator that uniquely determines the age, period, and cohort effects in the parameter subspace that is orthogonal to the null space of the singular design matrix. The IE is an estimable function in the sense that it is invariable to the choice of linear constraints on b. And no other constrained functions of b, with t not equal 0, are estimable because different non-zero values of t introduce different sizes of influences of the design matrix that are irrelevant to variations in Y to the estimates.

The estimability of the IE also follows from the fact that it satisfies the Kupper et al. (1985: 830) condition for estimability of linear functions of the parameter vector b, namely

$$l^T B_0 = 0, \qquad (8)$$

where l^T is a constraint vector (of appropriate dimension) that defines a linear function $l^T b$ of b. Note that, since the IE imposes the constraint that $t = 0$, i.e., that the arbitrary vector B_0 has zero influence, $l^T = (I - B_0 B_0^T)$ by (5). Since $B_0^T B_0 = 1$, condition (8) holds for the IE.

Yang and colleagues (2004) further showed that the IE possesses other desirable statistical properties. First, Kupper et al. (1985) noted that estimable functions are linear functions of the unidentified parameter vector that can be estimated without bias, i.e., estimable functions have unbiased estimators. Because the IE B satisfies this condition for APC models, it follows that, for a fixed number of time periods of data, the IE is an unbiased estimator. Second, the IE also has a variance smaller than that of any other estimators, unless, again, the other constraints result in $t = 0$. Therefore, it also has relative statistical efficiency. Third, the IE has statistical consistency – under suitable regularity conditions on the error term process and a fixed set of age categories, the IE will converge asymptotically to the "true" parameters that generate the sequence of APC data.

Model Validation Using Simulation Analysis

In brief, the IE possesses some valuable properties as a statistical estimator. Empirical analyses have also shown that the IE yields sensible estimated coefficients of age, period, and cohort variations in human mortality (Yang et al. 2004; Yang forthcoming 2008). However, given the long history of problems and pitfalls in proposed methods of APC analysis, it is reasonable to question whether this estimator gives numerical estimates of age, period, and cohort effect coefficients that are valid, i.e., reveal the true effects. This is a question of model validation – that is, does the identifying constraint imposed by the IE, the projection of the unconstrained APC accounting model vector onto the non-null space of X, produce estimated coefficients that capture the true age, time period, and cohort effects? Previous mathematical proofs are difficult to relate to real world situations and must rely on assumptions that may be violated. And empirical data analyses are not informative of the form of the true models, because models using different just-identifying constraints fit the data equally well. Therefore, this study conducts Monte Carlo simulation analyses that compare results from application of an estimator such as the IE or CGLIM to artificial data wherein we know the true form of the underlying model that generated the data. The simulation analyses can help to determine whether the IE indeed recovers the true parameters while CGLIM estimators do not.

We first investigate whether the IE is unbiased and is relatively efficient in samples with a fixed number of age groups and time periods. The asymptotic results in previous studies on the properties of the IE apply as the number of periods in the dataset goes to infinity. But any dataset used in practice has only a finite number of periods. Thus, we explore whether the asymptotic results give good approximations to the behavior of the IE in finite samples by simulating datasets with five, 10, and 50 periods. The basic asymptotic result we investigate is that, as the number of periods increases, estimated age effects should converge to the true age effects when using the IE but not necessarily when using other estimators. If the underlying processes generating the period-to-period changes in the observed outcomes are constant throughout the periods of the simulation, then we also expect the estimated period and cohort effects converge to the true period and cohort effects when using the IE.[2]

We conduct the simulation analyses systematically to examine the performance of the IE and alternative methods, CGLIM in particular, in reproducing the true models with all possible combinations of age, period, and cohort effects. We begin with simulations where the true model is a full APC model in which all three of the age, period, and cohort effects are present. This serves as a numerical illustration of the previous theoretical discussions of the properties of the IE provided by Yang et al. (2004). We then extend the analyses to other specifications of the true models wherein one or two effects are null.

The APC models using the CGLIM and IE approaches are estimated using Stata 9.2 through **`apc_cglim.ado`** and **`apc_ie.ado`**, respectively. Both ado-files may be obtained by typing "`ssc install apc`" on the Stata command line on any computer connected to the Internet, or downloaded from the Statistical Software Components archive at http://ideas.repec.org/s/boc/bocode/s456754.html. The programs are documented more fully in Stata help files.

Results for APC Models: True Effects of A, P, and C All Present

We fix the number of age categories in all simulations given the fact that humans have a relatively fixed life span. We let the number of age categories to be 10 without loss of generality. For a given number of periods P, we generate 1,000 datasets by Monte Carlo simulation in which the entries in the 10 × P outcome matrix are distributed according to:

$$y_{ij} \sim N(\mu, \sigma^2),$$

where[3]

$$\mu = 0.3 + 0.1(age_{ij} - 5.5)^2 + 0.1\sin(period_{ij}) + 0.1\cos(cohort_{ij}) + 0.1\sin(10 \cdot cohort_{ij})$$
$$\sigma^2 = 25$$

This equation for the data-generating process tells us what the true age, period and cohort effects are:

age effect at age a	$0.1(a - 5.5)^2$
period effect in period p	$0.1\sin(p)$

cohort effect in cohort *c*	$0.1\cos(c) + 0.1\sin(10c)$

So that the true effects have mean zero in each category in accord with the constraints on the effect coefficient specified earlier, we subtract constants from the effects listed above, where the constants are calculated as the mean effects for each category. To explore the finite time period properties of various estimators, we then estimate age, period and cohort effects in each simulated dataset for a given P using the IE and using three different CGLIM estimators: one with the first two age effects constrained to be equal (CGLIM_a), one with the first two period effects constrained to be equal (CGLIM_p), and one with the first two cohort effects constrained to be equal (CGLIM_c). To explore the large sample properties of these estimators, we let P increase from five to 10, and to 50 and repeat the simulations for each number of P.

Table 1 reports the results on age, period, and cohort effects estimated from data simulated with five time periods (P = 5). For each age, period, and cohort effect in the model, we show the true value and, for each estimator, the mean, standard deviation and mean squared error (MSE) of the estimated effect across 1,000 simulations. By comparing the mean of the simulated estimates to the true values, we can assess the degree of *unbiasedness* for each estimator. The standard deviation of the simulated estimates shows how much the estimated parameters vary from sample to sample. Smaller variance relates to *relative efficiency*. Mean squared error (MSE) is the average squared difference between the estimated parameter and the truth; this measure of accuracy takes into account both bias and variance.

[Table 1 and Figure 1 about here]

Figure 1 compares the means of IE and CGLIM estimates shown in Table 1. Two of the four estimators recover the profile of the age, period, and cohort effects qualitatively: the IE and the CGLIM_p. The other two sets of CGLIM estimates that constrain the first two age and first two cohort effects to be equal clearly fail to recover true forms of these effects because the constraints are incorrect: the first two true age effects and the first two true cohort effects are not equal; and the differences between the two true effects are large. The CGLIM_p estimator recovers the qualitative shapes of true effects more closely because constraining the coefficients of the first two period effects to be equal more closely approximates the fact that the difference between the two true effects is much smaller. Scrutiny of the numerical results in Table 1, however, suggests that CGLIM_p estimates are far off the mark in quantitative terms; only for the IE is the mean of each estimated effect close to the true value and hence unbiased. This is a direct result of the nonestimability of the CGLIM estimator. That is, any substantial departure from the IE constraint, which incorporates a large nonzero t, will not yield an estimable function and thus will induce bias in the estimates.

Nonetheless, this constrained model illustrates numerically a property established algebraically by Kupper et al. (1985: 830), namely, that, if the constraint used to just-identify the APC accounting model is, indeed, satisfied by the underlying true model, then the orthogonality condition stated above in Eq. (8) will hold and the corresponding constrained coefficient vector is estimable. With respect to the specification of the true effect coefficients in the present simulation, for example, it can be noted that, due to the periodicity of the age coefficients, several are equal. Thus, for example, if the analyst were to impose the identifying constraint a1 = a10, then the corresponding constrained coefficient vector will be estimable, and, in fact, the resulting CGLIM estimated coefficients will be within sampling and rounding error of those estimated by the IE.

Table 1 also shows that the IE exhibits substantially less sampling variation than the CGLIM estimators. The IE estimates of A, P, and C effects have standard deviations that range between 0 and 1. CGLIM_p estimates have the smallest standard deviations among all CGLIM estimates, but their standard deviations are still at least 10 times larger than those of the IE. The IE also has much smaller mean squared errors. The MSEs of the IE estimates are close to 0, whereas those of the CGLIM_p estimators can be as large as 6. All estimators have larger MSEs for the youngest and oldest cohorts because these cohorts are located at the upper and lower corners of the age by period table and have the smallest sample sizes. The CGLIM_a and CGLIM_c estimates have MSEs too large to provide reliable findings. It is noteworthy that the cohort effects are particularly poorly estimated by the CGLIM models.

To see what happens to the above estimators in cases when analysts have access to more data, we next increase the number of time periods to 10 and 50. Because the CGLIM_p continues to yield the best estimates among all CGLIM estimators, the following analysis focuses on the comparison of the CGLIM_p and the IE.

The means and standard errors of age effects estimates are shown in Figure 2 for the IE and the CGLIM_p by number of P. The means of the IE are extremely close to the true age effects for all P and rapidly approach the true age effects as P increases from five to 50. The means of the CGLIM_p also recover the true age effects well and do better with increasing P, but to a less extent than the IE. Although the means of the IE and CGLIM_p are close, the CGLIM_p shows much larger standard errors and thus statistically significant difference between the mean and the truth. Comparison of the standard errors across P (not shown) suggests decreasing sampling variations for both estimators with increasing P but much smaller variability for the IE for all P. Figure 3 further shows the advantage of the IE in terms of MSE. The IE has MSE much closer to 0 than the CGLIM_p for all P. Whereas the MSE of the IE approaches 0 as P increases, that of the CGLIM_p, although decreasing, is far above 0.

[Figures 2 and 3 about here]

Figures 4 and 5 present the mean and MSE of the IE and CGLIM_p estimates of period effects and cohort effects, respectively. Similar to the results shown earlier, the IE recovers the true effects much better for all P, increases in precision, and decreases in MSE with increasing P. In contrast, the CGLIM_p shows much larger departures from the true effects that do not decrease with increasing P. The first two time periods coefficients were constrained to be equal by the CGLIM_p, whereas in fact they increase slightly from time one to time two. As a result, the period effects estimated by the CGLIM_p rotate the true period effects (horizontal oscillations) upward. And the cohort effects estimates are rotated downward. While the IE has MSE close to 0, the CGLIM_p also produces much larger and in some cases increasing MSE with increasing P. This illustrates that linear constraints with even small deviations from the truth can result in coefficient estimates with large bias in unknown directions that will not lessen with more periods of data.

[Figures 4 and 5 about here]

Several insights follow from the above analysis:

First, the IE produces estimates of the A, P, and C effects that are more invariant to changes in the design matrix, such as additional time periods of data, than estimates produced by estimators that incorporate such influences. This precisely is because of its estimability/unbiasedness property. In this sense, the IE reduces the part of the subjectivity in the estimator that is due to the influence of fixed component determined by the shape of the data by removing it.

Second, both the IE B and any other estimator $\hat{b} = B + tB_0$ with $t \neq 0$ obtained from an equality constraint produce asymptotically consistent age effects as the number of time periods of data increase without bound. Therefore, with a large number (e.g., 50) of periods of data, differences among estimators decline and it makes little difference which identifying constraint is employed. In most empirical APC analyses, however, there usually are a small number (e.g., 5) time periods of observations available for analysis. In these cases, the differences can be substantial and an unbiased estimator should be preferred to a biased estimator, as the latter can be misleading with respect to the estimated trends.

Third, as a result of the above properties, the IE may provide a means for the accumulation of reliable estimates of the A, P, and C trends when more data become available over time, whereas the other estimators may not.

Results for APC Models: True Effects of A, P, and C Not All Present

We next investigate how well the estimators perform when one of the three sets of true age, period, and cohort effects is absent. For instance, age and period (AP) effects models with null cohort effects can have useful applications to such phenomena as fertility when the age-specific fertility rates are dependent on time period, but not birth cohort. The question naturally arises then as to how well the estimators perform when one of the three sets of effect coefficients has true effects that are equal to zero.

Using the Monte Carlo simulation from the same normal distribution shown earlier, we thus specify true effects to be the following:

age effect at age *a*	$0.1(a - 5.5)^2$
period effect in period *p*	$0.1\sin(p)$
cohort effect in cohort *c*	0

Similarly, we conducted the simulations for models with true AC, PC, A, P, and C effects by specifying the true effects of P, A, PC, AC, and AP to be zero, respectively.

Figure 6 shows the results of estimates of APC model with true A and P effects, but no C effects given five periods of data. The IE reproduces the true A and P effects remarkably well just as before and is much more superior in terms of MSE to the alternative estimator. It does show some small deviations from the true cohort effect which is zero. Figure 7 shows that such deviations decreased with increasing P. On average, the IE correctly estimates that there are no cohort effects in the data. The CGLIM_p estimator (as well as other CGLIM estimators) incorrectly finds cohort effects that are different from zero and change substantially across birth cohorts. In addition, bias and MSE of the CGLIM estimators do not decrease with more periods of data.

[Figures 6 and 7 about here]

Results are similar when we set all of the true P effects to zero in the data: only the IE reproduced true A and C effects and detected that no P effects were truly present. And the results hold for the other cases where one or two true effects are zero. It should be noted, however, that this statement applies to the performance of the IE in large samples of periods better than small samples.

Misuse of APC Models: Revisiting A Numerical Example

The above exposition and simulation analyses suggests that the IE indeed yields unbiased estimates of age, period, and cohort effects, have relative efficiency compared to alternative estimators, and converge to the true effects with increasing number of time periods. This seems to be in conflict with the age-old notion that there is no solution to the APC model identification problem because there can be any number of estimates that fit the data equally well. This notion is best represented in a recent critique of the utility of APC models in social research raised by Norval Glenn (2005). We next revisit the numerical example given by Glenn to evaluate this critique.

[Table 3 about here]

Glenn based his analysis on some hypothetical data sets that are cited here as Tables 3.1 – 3.3 for purpose of illustration. These data potentially show very different age, period, and cohort effects. The dependent variable values in Table 3.1 show obvious age variations with an increment of 5 for each successive age, but seemingly no period or cohort effects. Glenn correctly pointed out that there could be some combination of age and offsetting period and cohort effects and an infinite number of combinations of such effects can produce the pattern of variation. Data in Tables 3.2 and 3.3 show stronger period and cohort variations, respectively, and similarly can arise from many different combinations of true effects. In subsequent analyses, Glenn used the CGLIM approach to estimate APC models of these data. We present these results of the CGLIM analyses for data of Table 3.1 in Table 4. Corresponding analyses were conducted for the other two data sets but are not reported because the results are similar. Models 1 to 4 using Glenn's results confirm the point made earlier that different equality constraints result in drastically different estimates of A, P, and C effects. For purpose of comparison, we add Model 5 estimated by the IE. There are two fallacies in Glenn's interpretation of results of Models 1 to 4.

[Table 4 about here]

First, patterns in hypothetical data do not suggest true effects generating them. Glenn acknowledged that the data in Table 3.1 can be generated by any forms of the true A, P, and C effects, but then contradicted himself in the discussion of the modeling results claiming that "For this simulation experiment, I know what the effects are and can apply the Mason et al. method to the data to see how well it performs" (p. 12). Particularly, he assigned the true effects to be age effects shown in Models 1 and 2. Based on this incorrect starting point of what the true effects are, he went on to conclude that the method gives "grossly incorrect results" using certain constraints (like those in Models 3 and 4) and hence it is impossible to estimate APC effects with this method in practice when one cannot know what the right constraints are. It is clear from our simulation analysis shown above that one can only examine the performance of certain model estimators by specifying the true effects that generate the data rather than using certain data to speculate what the true effects are.

Second, the assumption that the age, period, and cohort trends in any given set of data can be best estimated by full APC models rather than by reduced models of one or two of the three effects needs to be tested. If the true effects of one or two of the three factors are null, the full model will overfit the data and produce inaccurate estimates. In addition, the full model has the model identification problem. The results shown in Models 1 to 4 given by Glenn reflect precisely this problem. And different linear constraints used to estimate the full model are bound to produce different estimates that are inaccurate in different degrees. Unlike in the simulation exercise, analysts cannot know which true effects are present and which are not, given only observed data.

One way to select among alternative models is to conduct model fit tests of whether all three of the age, period, and cohort effects are present and should be simultaneously estimated (see, e.g., Mason and Smith 1985). That is, analysts should successively estimate model with the A, P, C, AP, AC, PC, and APC sets of effect coefficients and examine the corresponding model fit statistics for improvement as additional sets and combinations of coefficients are added. This gives a sense of the relative importance of A, P, and C effects and the best model that summarizes the trends in the observed data. Accordingly, for the data of Table 3, we estimated nested models and computed model fit statistics for the three sets of data. The results shown in Table 5 suggest that the best fitting models for data sets one, two, and three are age effects only, period effects only, and cohort effects only or age and period effects only models, respectively. Because the full APC models are not the preferred models, the discussion of which identifying constraint gives the correct estimates is not productive and can be avoided given the findings from the model selection analysis.

[Table 5 about here]

There is a caveat for use of model fit tests for full APC models identified by using different constraints. As has been noted many times over the years in discussions of the

APC accounting model (see, e.g., Pullum 1977, 1980; Rogers 1982), different just-identified models will generate the same data and yield exactly the same model fit, and therefore goodness-of-fit can not be used as a criterion by which to select a just-identified model. But estimability can be so used. And the IE stands out in this respect.

If, indeed, the preferred model is a simpler model with true age effects or any other effects, will the IE be useful in revealing such true effects? We have shown in the foregoing Monte Carlo simulations that the answer is yes with some qualifications. For example, in cases when there are only true age effects, but no period and cohort effects, Figure 8 suggests that the IE does not recover the true zero effects well if we have only five periods of data and may lead to false conclusions about these effects. But as the number of periods increases, the IE largely reproduces the true zero period and cohort effects.

Revisiting the numerical example given in previous sociological literature, we find that the long-standing critique of the utility of APC accounting models rests on problematic logic and misuse of APC models. There is no adequate proof that there exists no estimable function of the linear components of the effects, only misunderstanding of what is or is not estimable.

Conclusion

The problem of obtaining reliable estimates of the patterns of simultaneous changes across age groups, time periods, and cohorts has long provided an intriguing challenge in many contexts in the social sciences. This paper has studied the performance of a new method of estimation of the APC accounting model, namely, the intrinsic estimator. The IE has passed simulation tests of validity under various circumstances and can provide a useful tool for the accumulation of scientific knowledge about the distinct effects of age, period, and cohort effects in social research.

The APC multiple classification model has been mostly widely used in demographic and epidemiologic research wherein the outcomes of interest are usually vital events that are Poisson rather than normal variates. We therefore also conducted simulation analysis sampling from Poisson distributions. The findings were identical to those reported here. In sum, the simulations show that, regardless of the sampling distribution and the sample size of the observed data, the IE is more accurate and efficient than the CGLIM. The IE performs well even in datasets with just five periods, perhaps the smallest sample size that might be used in practice. By contrast, the CGLIM estimators give incorrect results even when there are as many as fifty periods, which would be an unusually large sample in many social and demographic applications.

Is the intrinsic estimator then a "complete solution" to the identification problem in APC and similar models in social science? No. Structural identification problems are just that — points of underidentification of parameters due to the very nature of the underlying models. There is not now, and can never be, a complete resolution of such problems. But there can be variations among approaches to structural identification problems with respect to desirable statistical properties. Because of its desirable properties as a statistical estimator – including its ability, as demonstrated above, to produce good estimates of the underlying patterns of change in age, period, and cohort effects with a small number of time periods of data – the IE adds a potentially useful method to the toolkit available for these analyses.

Does our results mean that researchers should naively apply the IE to APC data and expect to obtain meaningful results? No. APC analysis is well known to be treacherous for reasons articulated by Glenn (2005) and should, in all cases, be approached with great caution and awareness of its many pitfalls. Every statistical model has its limits and will break down under some conditions. We have shown one such condition, in which the IE produces larger bias in small samples when the true effects are zero than when the true effects are not zero. We also have shown, by analyses of Glenn's (2005) numerical example, that researchers should conduct careful model selection tests before using full APC models. We have shown that simulation analysis is one avenue for model validation. Imposition of a full APC model on data when a reduced model fits the data equally well or better constitutes a model misspecification and should be avoided. On the other hand, when all three of the age, period, and cohort dimensions appear to be operative in producing a given set of tabulated rates, application of the Intrinsic Estimator may be quite useful in producing meaningful and stable estimates of the trends across the age, period, and cohort categories.

Endnotes:

---

[1] Note that the parameters α_a, β_p, and γ_{a+p-1} are not included in the parameter vector b because they can be uniquely determined by use of the sum-to-zero constraint. The use of reference categories is equivalent to the translation by a constant of the parameter estimates produced by the constraint and thus of no substantive importance.

[2] If such processes are not constant and change from period to period, regardless of the estimator used, estimated period and cohort effects cannot be expected to converge to their true values as the number of periods increases because adding a period to the dataset does not add information about the previous periods or about cohorts not present in the period just added.

[3] We chose the variance of 25 so that the sampling variability of the estimator would be visible in our graphs. Experiments with smaller and larger variances produced qualitatively similar results.

Appendix A: The Principal Components Computational Algorithm for the Intrinsic Estimator

i) Compute the eigen-vectors $u_1, \ldots, u_r$ of matrix $X^T X$, where X denotes the design matrix of model (2). Normalize them with $\|u_m\| = 1, \ldots, r$ and denote the orthonormal matrix as $U = (u_1, \ldots u_r)^T$;

ii) Identify the special eigenvector B_0 corresponding to eigenvalue 0. Denote $u_1 = B_0$ without loss of generality;

iii) Select the principal components to be the remaining eigen-vectors $u_2, \ldots, u_r$ with non-zero eigen-values;

iv) Fit a principal components regression (PCR) model with the outcome variable of interest (e.g., logged death rates) as the response using a design matrix V whose column vectors are the principal components $u_2, \ldots, u_r$, i.e., $V = (u_2, \ldots, u_r)$, to obtain the coefficients ($w_2, \ldots, w_r$);

v) Set coefficient $w_1 = 0$ and transform the coefficients vector $w = (w_1, \ldots, w_r)^T$ by the orthonormal matrix $U = (u_1, \ldots u_r)^T$ to obtain the intrinsic estimator $B = Uw$.

Note: Instead of using reference categories, the IE uses the usual ANOVA-type constraints: $\sum_i \alpha_i = \sum_j \beta_j = \sum_k \gamma_k = 0$. The computational algorithm used by the IE estimates effect coefficients for each of the $a - 1$, $p - 1$, and $a + p - 2$ age, period, and cohort categories, respectively, which is consistent with the definition of the parameter vector b. Then the IE uses the zero-sum constraints to obtain the numerical values of the omitted age, period, and cohort categories.

References

Christensen, R. 2000. *Plane Answers to Complex Questions: The Theory of Linear Models*. New York: Springer.

Fu, W.J. 2007. "A Smoothing Cohort Model in Age-Period-Cohort Analysis with Applications to Homicide Arrest Rates and Lung Cancer Mortality Rates." *Sociological Methods & Research*: in press. http://www.msu.edu/~fuw/apc/apcsmthFinal.pdf

Fu, W.J. 2000. "Ridge Estimator in Singular Design with Application to Age-Period-Cohort Analysis of Disease Rates." *Communications in Statistics--Theory and Method* 29:263-278.

Glenn, N. 2005. *Cohort Analysis*. Thousand Oaks, CA: Sage.

Hobcraft, J., J. Menken, and S. Preston. 1982. "Age, Period, And Cohort Effects in Demography: A Review." *Population Index* 48:4-43.

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

Kupper, L.L., J.M. Janis, A. Karmous, and B.G. Greenberg. 1985. "Statistical Age-Period-Cohort Analysis: A Review and Critique." *Journal of Chronic Disease* 38:811-830.

Mason, K.O., W.H. Mason, H.H. Winsborough, and K. Poole. 1973. "Some Methodological Issues in Cohort Analysis of Archival Data." *American Sociological Review* 38: 242-58.

Mason, W.H. and H.L. Smith. 1985. "Age-Period-Cohort Analysis and the Study of Deaths from Pulmonary Tuberculosis." Pp. 151-228 in *Cohort Analysis in Social Research*, edited by W.M. Mason and S.E. Fienberg. New York: Springer-Verlag.

Mason, W.H. and N.H. Wolfinger. 2002. "Cohort Analysis." Pp. 2189-2194 in *International Encyclopedia of the Social & Behavioral Sciences*. New York: Elsevier.

Fienberg, S.E. and W.M. Mason. 1978. "Identification and Estimation of Age-period-cohort Models in the Analysis of Discrete Archival Data." in *Sociological Methodology 1979* edited by K.F. Schuessler. San Francisco: Jossey-Bass, pp. 1-67.

Fienberg, S.E. and W.M. Mason. 1985. "Specification and Implementation of Age, Period, and Cohort Models." Pp. 45-88 in *Cohort Analysis in Social Research*, edited by W.M. Mason and S.E. Fienberg. New York: Springer-Verlag.

McCulloch, C. E. and S. R. Searle. 2001. *Generalized, Linear, and Mixed Models*. New York: Wiley-Interscience.

Robertson, C., S. Gandini, and P. Boyle. 1999. "Age-Period-Cohort Models: A Comparative Study of Available Methdologies." *Journal of Clinical Epidemiology* 52:569-583.

Rodger, W. L. 1982. “Estimable Functions of Age, Period, and Cohort Effects.” *American Sociological Review* 47:774-787.

Yang, Y. Forthcoming 2008. “Trends in U.S. Adult Chronic Disease Mortality: Age, Period, and Cohort Variations.” *Demography* 45. http://home.uchicago.edu/%7Eyangy/Yang_Demography08.pdf

Yang, Y., W. J. Fu, and K.C. Land. 2004. “A Methodological Comparison of Age-Period-Cohort Models: Intrinsic Estimator and Conventional Generalized Linear Models.” *Sociological Methodology* 34: 75-110.

Table 1. Simulation Results (n = 1000) of the IE and CGLIM Estimators of APC Models: P = 5

			IE			CGLIM_a (a1=a2)			CGLIM_p (p11=p12)			CGLIM_c (c1=c2)	
Variable	**True Effect**	**Mean**	**sd**	**MSE**	**Mean**	**sd**	**MSE**	**Mean**	**sd**	**MSE**	**Mean**	**sd**	**MSE**
Age													
a1	1.200	1.172	2.408	5.795	-2.198	14.035	208.3	1.357	10.877	118.2	1.858	30.111	906.2
a2	0.400	0.423	2.218	4.915	-2.198	14.035	203.5	0.567	8.572	73.4	0.956	23.455	549.9
a3	-0.200	-0.252	2.283	5.208	-2.124	9.221	88.6	-0.149	6.268	39.3	0.129	16.700	278.7
a4	-0.600	-0.603	2.362	5.573	-1.726	5.954	36.7	-0.541	4.285	18.3	-0.374	10.249	105.0
a5	-0.800	-0.797	2.331	5.426	-1.172	3.052	9.4	-0.777	2.606	6.8	-0.721	3.977	15.8
a6	-0.800	-0.780	2.312	5.340	-0.406	2.837	8.2	-0.801	2.636	6.9	-0.856	4.092	16.7
a7	-0.600	-0.643	2.259	5.101	0.481	5.801	34.8	-0.705	4.135	17.1	-0.871	10.271	105.5
a8	-0.200	-0.145	2.331	5.432	1.728	9.204	88.4	-0.248	6.440	41.4	-0.526	16.770	281.1
a9	0.400	0.358	2.235	4.990	2.979	12.591	165.0	0.213	8.599	73.9	-0.176	23.807	566.5
a10	1.200	1.267	2.302	5.296	4.637	16.174	273.2	1.081	10.675	113.9	0.581	29.482	868.7
Period[a]													
p11	-0.110	-0.023	1.429	2.046	1.475	7.157	53.7	-0.106	3.610	13.0	-0.328	13.114	171.9
p12	-0.064	-0.064	1.475	2.173	0.684	3.784	14.9	-0.106	3.610	13.0	-0.217	7.019	49.2
p13	0.032	0.000	1.462	2.136	0.000	1.462	2.1	0.000	1.462	2.1	0.000	1.462	2.1
p14	0.089	0.109	1.390	1.932	-0.639	3.790	14.9	0.151	2.743	7.5	0.262	6.788	46.1
p15	0.055	-0.022	1.486	2.211	-1.520	7.098	52.8	0.061	4.915	24.1	0.283	13.419	179.9
Cohort													
c1	-0.006	-0.245	4.714	22.257	-14.315	65.589	4502.3	0.065	41.572	1726.5	2.088	108.599	11786.3
c2	0.044	-0.092	3.467	12.024	-13.413	61.125	3913.6	0.176	39.210	1535.9	2.088	108.599	11786.1
c3	-0.203	-0.242	2.995	8.962	-12.814	57.620	3475.8	-0.014	36.896	1360.0	1.787	99.857	9965.5
c4	0.004	-0.032	2.526	6.377	-11.855	54.197	3074.9	0.154	34.487	1188.2	1.844	93.307	8700.9
c5	-0.003	0.015	2.321	5.380	-11.059	50.752	2695.4	0.160	32.539	1057.8	1.739	86.485	7475.2
c6	0.060	-0.046	2.540	6.454	-10.372	47.120	2326.9	0.058	30.080	903.9	1.525	79.842	6370.5
c7	0.147	0.172	2.647	7.001	-9.405	43.721	2000.9	0.234	27.755	769.6	1.591	73.440	5390.1
c8	-0.120	-0.105	2.445	5.973	-8.933	40.140	1687.3	-0.084	25.567	653.0	1.161	66.656	4440.2
c9	-0.007	0.060	2.456	6.031	-8.019	36.361	1385.0	0.040	23.613	557.0	1.174	60.066	3605.7
c10	-0.140	-0.111	2.268	5.139	-7.441	33.486	1173.5	-0.172	21.563	464.5	0.850	53.464	2856.6
c11	-0.010	0.052	2.537	6.436	-6.530	30.168	951.7	-0.051	19.621	384.6	0.861	46.768	2185.8
c12	0.137	0.261	2.772	7.692	-5.571	26.900	755.5	0.117	17.396	302.3	0.918	40.302	1623.2
c13	-0.008	0.208	3.268	10.718	-4.876	24.050	601.5	0.023	15.559	241.8	0.712	33.969	1153.3
c14	0.106	0.106	5.332	28.401	-4.229	22.501	524.6	-0.121	14.134	199.6	0.458	27.651	763.9

[a]Periods are labeled such that cohort = period – age.

Note: The intercept is a normalizing constant, so its estimates do not matter for evaluation of performance of certain estimators and are not presented here.

Table 2. Simulation Results (n = 1000) of the IE and CGLIM Estimators: Age Effects by Number of Time Periods

Age		True Effect	IE			CGLIM_a			CGLIM_p			CGLIM_c		
			P = 5	P = 10	P = 50	P = 5	P = 10	P = 50	P = 5	P = 10	P = 50	P = 5	P = 10	P = 50
a1	mean	1.200	1.172	1.215	1.233	-2.198	-2.389	-2.599	1.357	1.407	1.521	1.858	1.758	3.155
	sd		2.408	1.592	0.694	14.035	9.532	4.181	10.877	10.659	10.361	30.111	28.481	27.687
	MSE		5.795	2.531	0.482	208.337	103.655	31.900	118.224	113.547	107.351	906.214	810.661	769.646
a2	mean	0.400	0.423	0.414	0.381	-2.198	-2.389	-2.599	0.567	0.563	0.606	0.956	0.837	1.876
	sd		2.218	1.559	0.702	14.035	9.532	4.181	8.572	8.343	8.094	23.455	22.245	21.531
	MSE		4.915	2.429	0.492	203.540	98.553	26.461	73.435	69.571	65.484	549.880	494.557	465.310
a3	mean	-0.200	-0.252	-0.191	-0.185	-2.124	-2.193	-2.314	-0.149	-0.085	-0.025	0.129	0.111	0.883
	sd		2.283	1.553	0.671	9.221	6.094	2.682	6.268	6.088	5.775	16.700	15.821	15.377
	MSE		5.208	2.410	0.450	88.646	41.075	11.658	39.256	37.041	33.344	278.721	250.145	237.400
a4	mean	-0.600	-0.603	-0.613	-0.586	-1.726	-1.814	-1.864	-0.541	-0.549	-0.490	-0.374	-0.432	0.055
	sd		2.362	1.569	0.688	5.954	3.926	1.694	4.285	3.842	3.490	10.249	9.537	9.237
	MSE		5.573	2.461	0.473	36.680	16.870	4.464	18.344	14.747	12.182	104.997	90.894	85.667
a5	mean	-0.800	-0.797	-0.846	-0.794	-1.172	-1.247	-1.220	-0.777	-0.825	-0.762	-0.721	-0.786	-0.580
	sd		2.331	1.608	0.715	3.052	2.012	0.901	2.606	1.981	1.360	3.977	3.514	3.146
	MSE		5.426	2.585	0.510	9.446	4.245	0.988	6.786	3.921	1.849	15.810	12.333	9.934
a6	mean	-0.800	-0.780	-0.832	-0.800	-0.406	-0.431	-0.374	-0.801	-0.853	-0.832	-0.856	-0.892	-1.014
	sd		2.312	1.592	0.669	2.837	2.009	0.880	2.636	1.957	1.301	4.092	3.569	3.096
	MSE		5.340	2.534	0.447	8.194	4.168	0.955	6.943	3.829	1.693	16.734	12.735	9.619
a7	mean	-0.600	-0.643	-0.683	-0.619	0.481	0.518	0.658	-0.705	-0.747	-0.715	-0.871	-0.864	-1.260
	sd		2.259	1.552	0.673	5.801	3.789	1.718	4.135	3.950	3.516	10.271	9.638	9.253
	MSE		5.101	2.414	0.453	34.784	15.595	4.533	17.095	15.609	12.366	105.463	92.858	85.969
a8	mean	-0.200	-0.145	-0.161	-0.199	1.728	1.841	1.930	-0.248	-0.268	-0.359	-0.526	-0.463	-1.266
	sd		2.331	1.610	0.672	9.204	6.175	2.650	6.440	6.001	5.766	16.770	15.799	15.389
	MSE		5.432	2.592	0.451	88.351	42.253	11.552	41.431	35.982	33.243	281.062	249.427	237.735
a9	mean	0.400	0.358	0.423	0.389	2.979	3.226	3.370	0.213	0.274	0.165	-0.176	0.001	-1.106
	sd		2.235	1.568	0.683	12.591	8.486	3.711	8.599	8.472	8.115	23.807	22.343	21.589
	MSE		4.990	2.457	0.466	165.016	79.926	22.574	73.896	71.717	65.845	566.534	498.855	467.902
a10	mean	1.200	1.267	1.274	1.180	4.637	4.878	5.012	1.081	1.082	0.891	0.581	0.731	-0.742
	sd		2.302	1.549	0.684	16.174	10.781	4.689	10.675	10.520	10.356	29.482	28.234	27.640
	MSE		5.296	2.402	0.468	273.157	129.645	36.499	113.861	110.577	107.239	868.679	796.560	766.961

Table 3. Patterns of Data Showing Age, Period, and Cohort Effects (Glenn 2005)

Table 3.1 "Pure Age Effects" (Glenn 2005: Table 1.2)

	Year					
Age	1950	1960	1970	1980	1990	2000
20-29	50	50	50	50	50	50
30-39	55	55	55	55	55	55
40-49	60	60	60	60	60	60
50-59	65	65	65	65	65	65
60-69	70	70	70	70	70	70
70-79	75	75	75	75	75	75

Table 3.2 "Pure Period Effects" (Glenn 2005: Table 1.3)

	Year					
Age	1950	1960	1970	1980	1990	2000
20-29	30	35	40	45	50	55
30-39	30	35	40	45	50	55
40-49	30	35	40	45	50	55
50-59	30	35	40	45	50	55
60-69	30	35	40	45	50	55
70-79	30	35	40	45	50	55

Table 3.3 "Pure Cohort Effects" (Glenn 2005: Table 1.4)

	Year					
Age	1950	1960	1970	1980	1990	2000
20-29	50	55	60	65	70	75
30-39	45	50	55	60	65	70
40-49	40	45	50	55	60	65
50-59	35	40	45	50	55	60
60-69	30	35	40	45	50	55
70-79	25	30	35	40	45	50

Note: As noted by Glenn, data in Table 3.1 could show alternative true effects such as pure age effects, offsetting period and cohort effects, or a combination of age effects and offsetting period and cohort effects. The same applies to data in Tables 3.2 and 3.3.

Table 4. Regression Coefficients of APC Models of Data in Table 3.1 Estimated by the CGLIM and IE

	CGLIM*				IE
Variable	**1**	**2**	**3**	**4**	**5**
Intercept	50.0	50.0	25.0	28.7	62.5
Age					
20-29	-12.5	-12.5	0.0	**-5.0**	-11.2
30-39	-7.5	-7.5	0.0	**-5.0**	-6.7
40-49	-2.5	-2.5	**0.0**	-1.4	-2.2
50-59	2.5	2.5	**0.0**	1.1	2.2
60-69	7.5	7.5	0.0	3.6	6.7
70-79	12.5	12.5	0.0	6.4	11.2
Period					
1950	0.0	0.0	-12.5	**-5.0**	-1.3
1960	0.0	0.0	-7.5	**-5.0**	-0.8
1970	0.0	0.0	-2.5	-1.4	-0.3
1980	0.0	**0.0**	2.5	1.1	0.3
1990	0.0	**0.0**	7.5	3.6	0.8
2000	0.0	0.0	12.5	6.4	1.3
Birth Cohort					
1880	0.0	0.0	25.0	11.2	2.6
1890	0.0	0.0	20.0	10.1	2.1
1900	0.0	0.0	15.0	7.7	1.6
1910	0.0	0.0	10.0	5.3	1.1
1920	0.0	0.0	5.0	3.0	0.5
1930	0.0	0.0	0.0	0.2	0.0
1940	0.0	0.0	-5.0	-2.0	-0.5
1950	0.0	0.0	-10.0	-4.7	-1.0
1960	**0.0**	0.0	-15.0	-7.3	-1.6
1970	**0.0**	0.0	-20.0	-9.9	-2.1
1980	0.0	0.0	-25.0	-13.8	-2.6

* Adopted from Glenn (2005: Table 2.1); the regression coefficients are centered to sum to zero within age, period, and cohort categories based on the usual constraint, $\sum_i \alpha_i = \sum_j \beta_j = \sum_k \gamma_k = 0$ (see Footnote 1); coefficients highlighted in Bold are constrained to be equal.

Table 5. Model Fit Statistics for Data in Table 3

For Data in Table 3.1			
Models	**Log-Likelihood**	**DF**	**BIC**
A	**121.8**	**6**	**-107.5**
P	-128.3	6	2516.4
C	-115.8	11	1222.5
AP	123.1	11	-89.6
AC	132.6	16	-71.7
PC	132.6	16	-71.7
APC	122.6	20	-57.3

For Data in Table 3.2			
Models	**Log-Likelihood**	**DF**	**BIC**
A	-128.3	6	2517.6
P	**119.9**	**6**	**-107.5**
C	-115.8	11	1223.5
AP	121.4	11	-89.6
AC	125.4	16	-71.7
PC	129.8	16	-71.7
APC	131.5	20	-57.3

For Data in Table 3.3			
Models	**Log-Likelihood**	**DF**	**BIC**
A	-128.3	6	2518.3
P	-128.3	6	2516.9
C	**126.1**	**11**	**-89.6**
AP	**125.4**	**11**	**-89.6**
AC	133.5	16	-71.7
PC	129.4	16	-71.7
APC	124.4	20	-57.3

Note: Model fit statistics, BIC (Bayesian Information Criterion), are calculated by Stata GLM; the smaller the AIC and BIC, the better the model fit. The best fitting models for each dataset are highlighted in Bold.

Figure 1. Means of Estimates from 1000 Simulations of APC Models with P = 5: IE vs. CGLIM

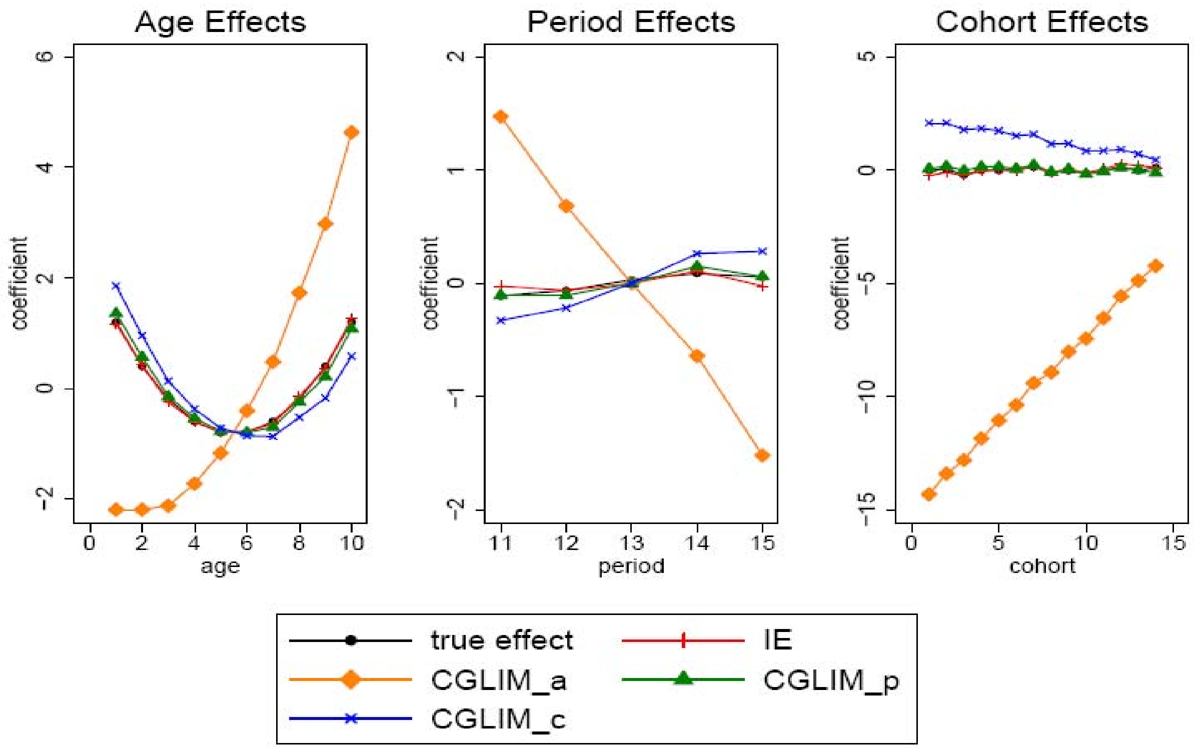

Figure 2. Means and Standard Errors of Age Effects Estimates from 1000 Simulations of APC Models by P: IE vs. CGLI M

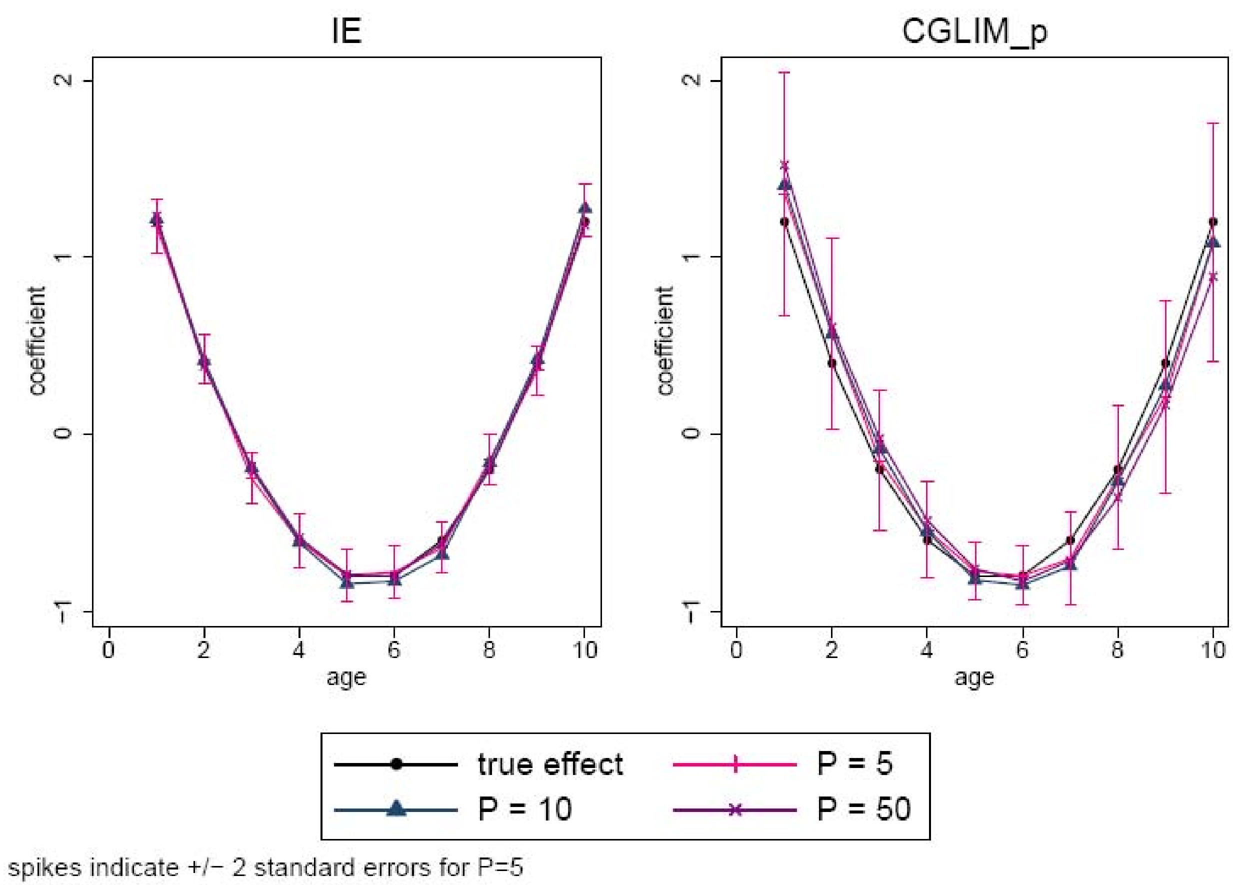

Figure 3. Mean Squared Errors of the Age Effects Estimates from 1000 Simulations of APC Models by P: IE vs. CGLIM

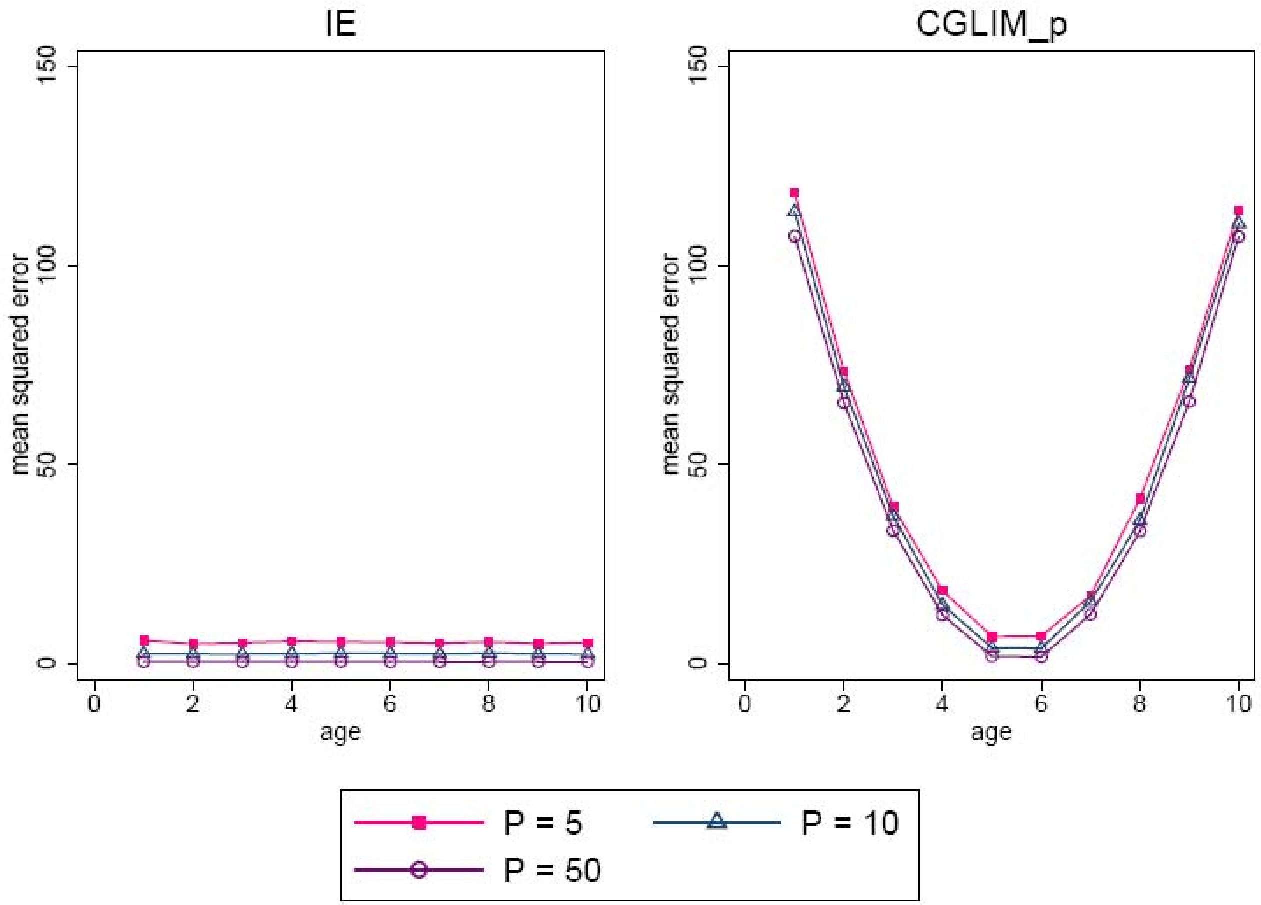

Figure 4. Period Effects Estimates from 1000 Simulations of APC Models by P: IE and CGLIM

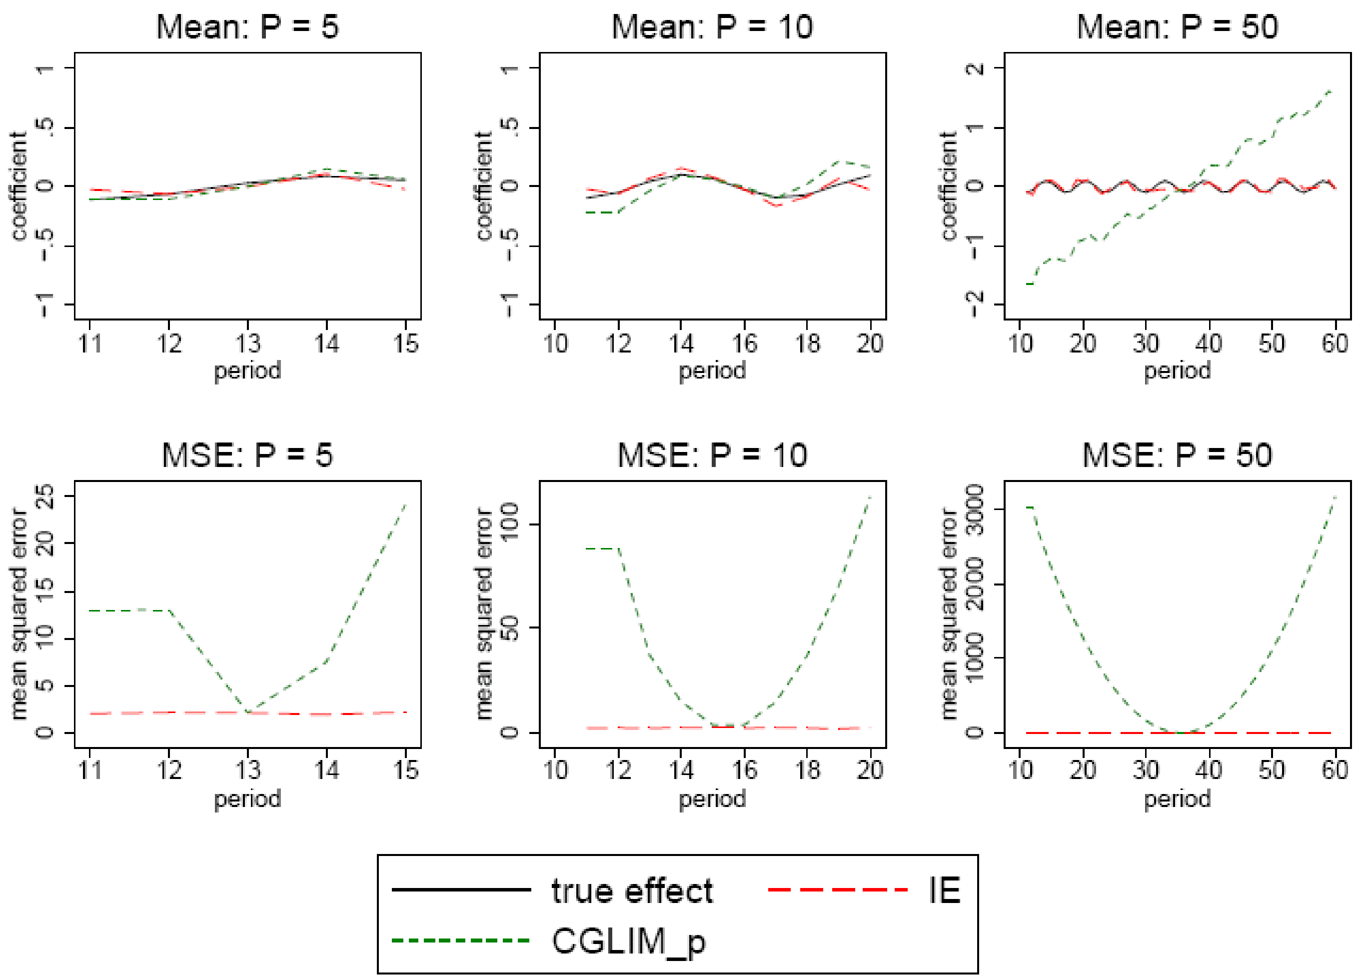

Figure 5. Cohort Effects Estimates from 1000 Simulations of APC Models by P: IE and CGLIM

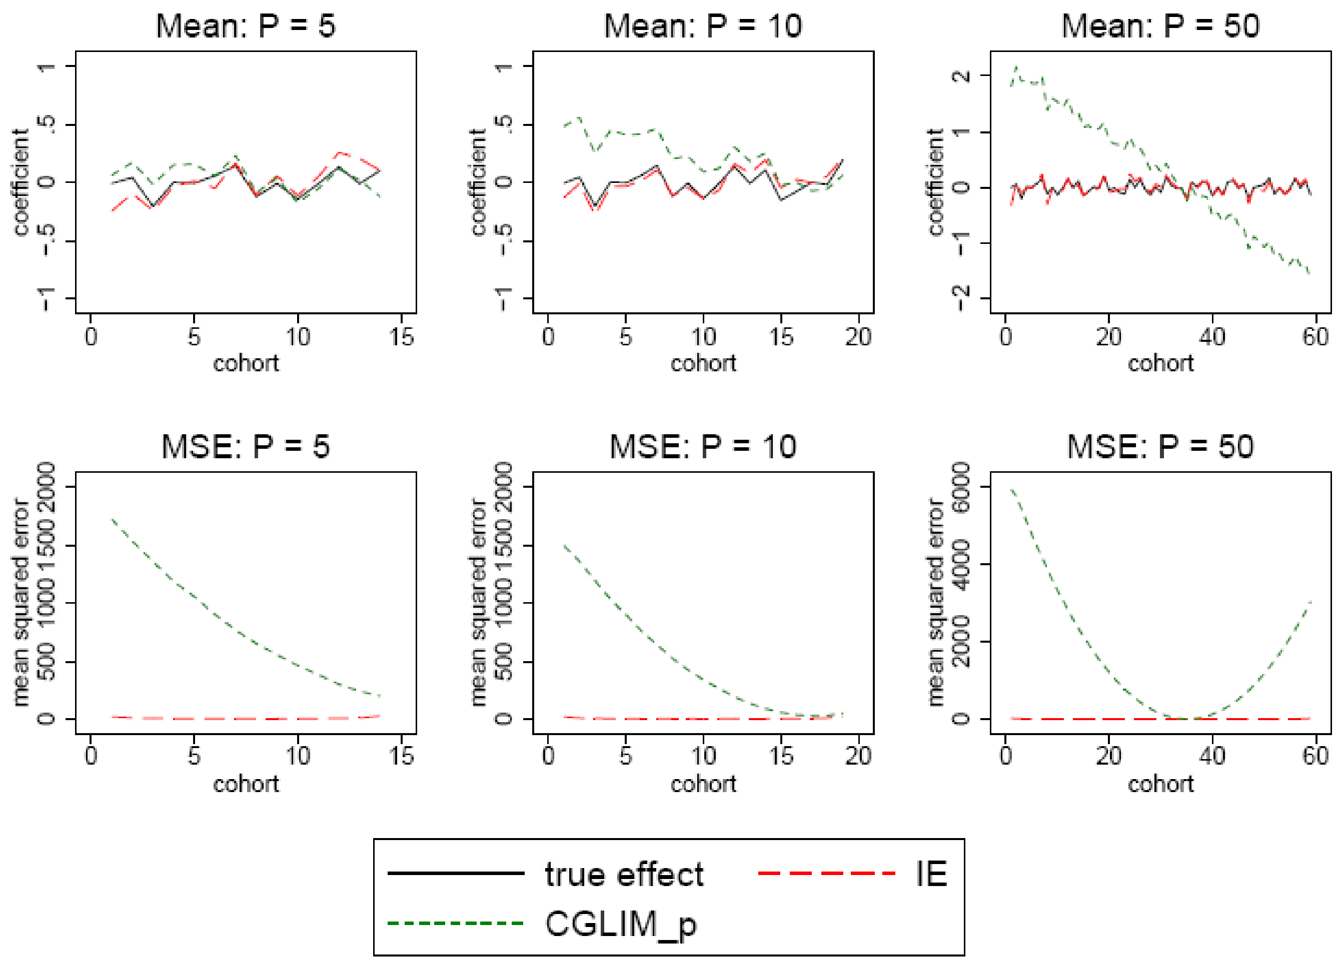

Figure 6. Results from 1000 Simulations of APC Models with P = 5: True AP Effects but No True C Effects

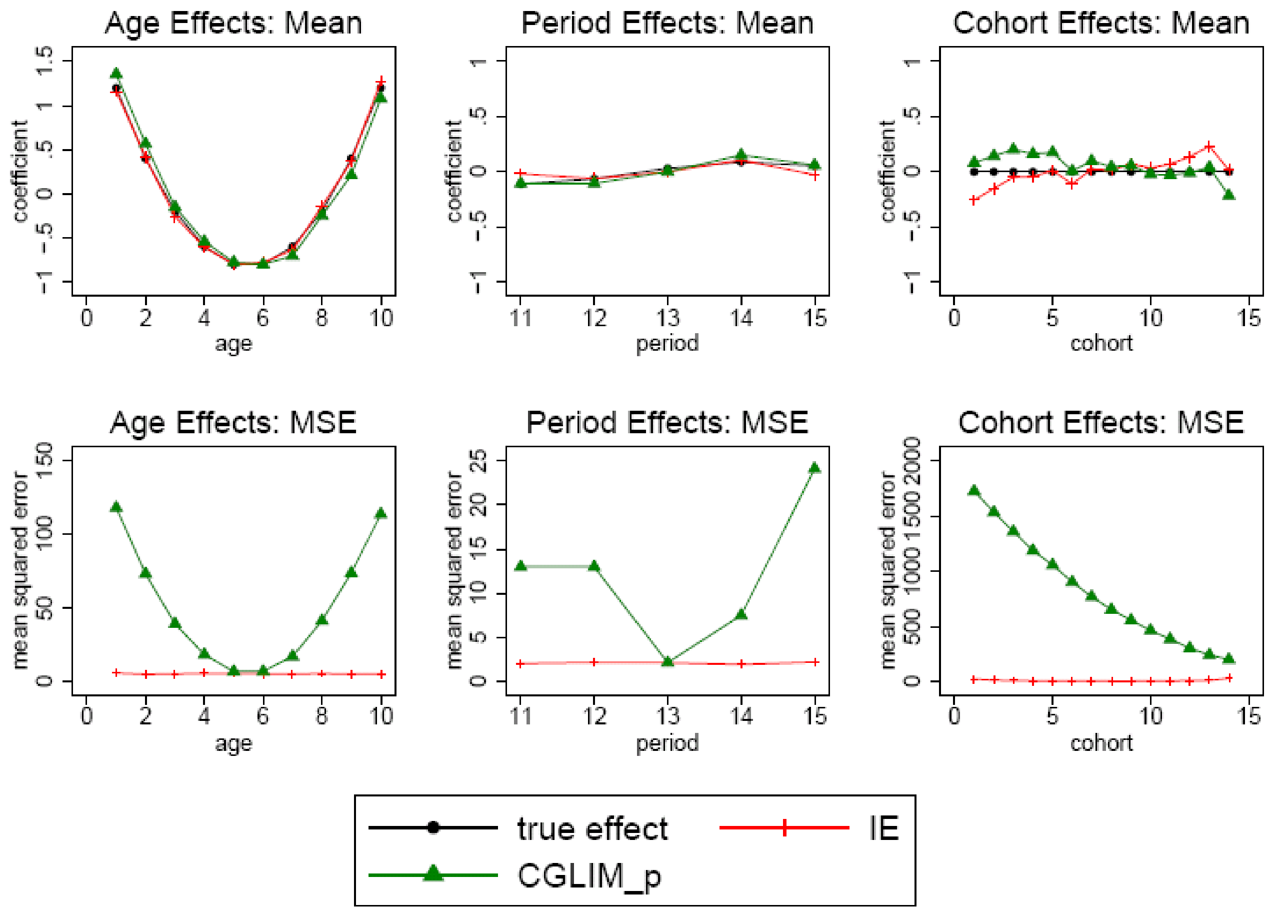

Figure 7. Cohort Effects Estimates from 1000 Simulations of APC Models by P: True AP Effects but No True C Effects

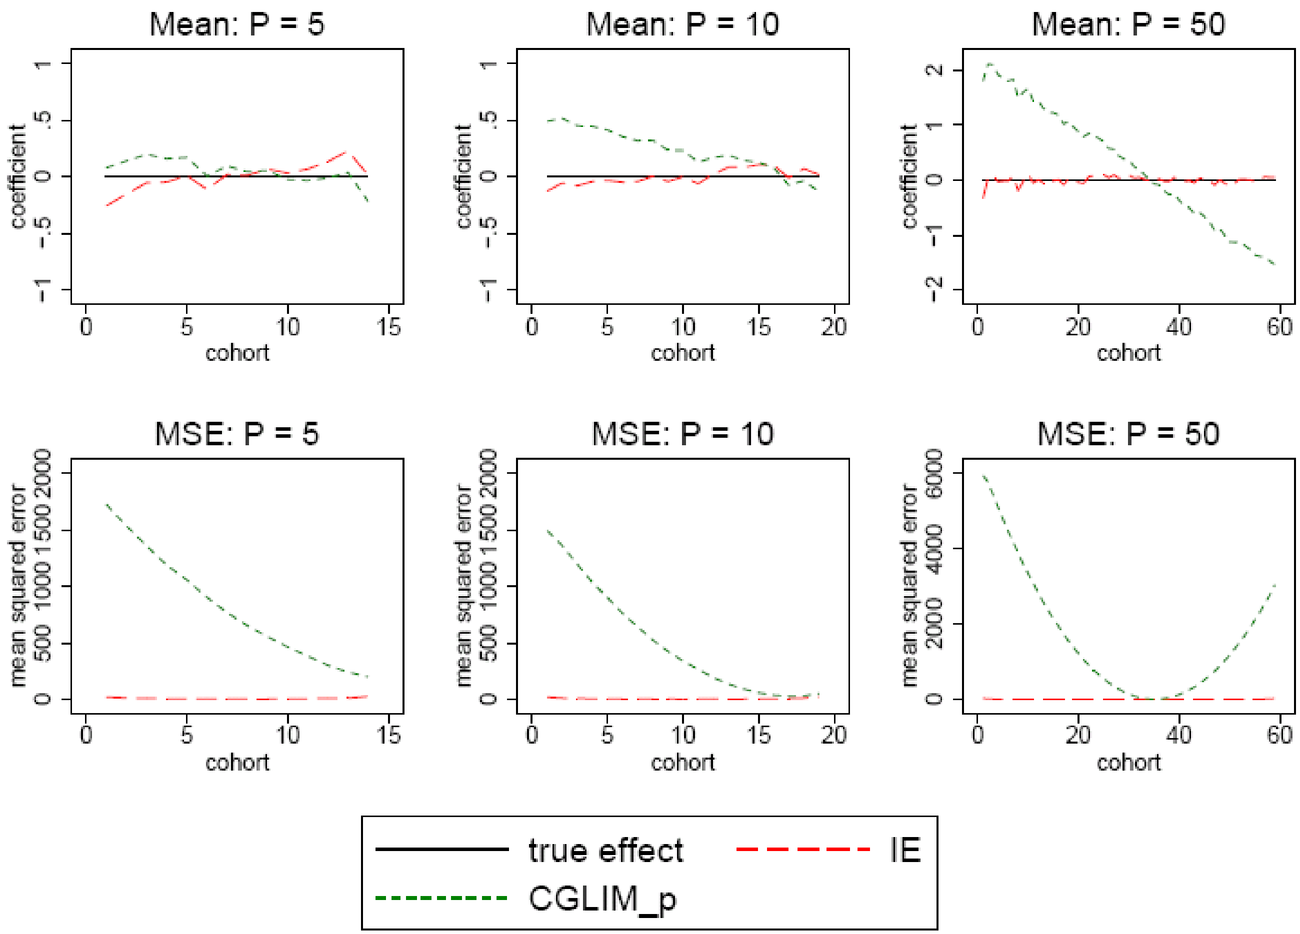

Figure 8. Means of Estimates of Period and Cohort Effects from 1000 Simulations by P: True Age Effects Only