The Intrinsic Estimator for Age-Period-Cohort Analysis: What It Is and How to Use It[1]

Yang Yang
University of Chicago

Sam Schulhofer-Wohl
Princeton University

Wenjiang J. Fu
Michigan State University

Kenneth C. Land
Duke University

A new approach to the statistical estimation of age-period-cohort (APC) accounting models, called the intrinsic estimator (IE), recently has been developed. This article (1) further describes the IE algebraically, geometrically, and verbally, (2) reviews properties of the IE as a statistical estimator, (3) provides model validation evidence for the IE both from an empirical example and from a simulation exercise, (4) relates the coefficients of the IE to those of conventional constrained APC models using formal definitions of statistical estimability, hypothesis testing, and empirical applications that directly address a criticism that often has been lodged at general-purpose methods of APC analysis, and (5) introduces computer software for application of the IE that interested users can readily access. The authors conclude that the IE holds the potential for applications not only to APC analysis but also to similar problems of structural underidentification in sociology.

INTRODUCTION

Age-period-cohort (APC) analysis has played a critical role in studying time-specific phenomena in sociology, demography, and epidemiology for the past 80 years (Mason and Wolfinger 2002). Broadly defined, APC analysis distinguishes three types of time-related variation in the phenomena of interest: *age effects*, or variation associated with different age groups; *period effects*, or variation over time periods that affect all age groups simultaneously; and *cohort effects*, or changes across groups of

[1] Direct correspondence to Yang Yang, Department of Sociology, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637. E-mail: yangy@uchicago.edu

0002-9602/2008/11306-0006$10.00

individuals who experience an initial event such as birth in the same year or years. These distinctions have important implications for measurement and analysis. The considerable regularity of age variations in many social outcomes across time and place reflects the developmental nature of true age changes. In contrast, period and cohort effects reflect the influences of social forces. Period variations often result from shifts in social, historical, and cultural environments. Cohort variations are conceived as the essence of social change and may reflect the effects of early life exposure to socioeconomic, behavioral, and environmental factors that act persistently over time to produce differences in life course outcomes for specific cohorts (Ryder 1965).

The APC accounting/multiple classification model was introduced to sociologists by Mason, Mason, Winsborough, and Poole (1973) and serves as a general methodology for cohort analysis when age, period, and cohort are all potentially of interest. This general methodology focuses on the APC analysis of data in the form of tables of percentages or occurrence/exposure rates of events such as births, deaths, disease incidence, and crimes. In spite of its theoretical merits and conceptual relevance, APC analysis of tabulated data suffers from the "identification problem" induced by the exact linear dependency between age, period, and (birth) cohort: period = age + cohort. This can be viewed as a special case of collinear regressors that produces, in this instance, a singular matrix (of one less than full rank) used in the statistical estimation process. Since a singular matrix produces multiple estimators of the three effects, it is difficult to estimate the unique set of true separate effects.

A number of methodological contributions to the specification and estimation of APC models have occurred in recent decades in a wide variety of disciplines, including social and demographic research (e.g., Glenn 1976; Fienberg and Mason 1978, 1985; Hobcraft, Menken, and Preston 1982; Wilmoth 1990; O'Brien 2000) and biostatistics and epidemiology (e.g., Osmond and Gardner 1982; Clayton and Schifflers 1987; Holford 1992; Tarone and Chu 1992; Robertson and Boyle 1998; Fu 2000). Various analytic approaches have produced ambiguous and inconsistent results. Researchers do not agree on methodological solutions to these problems and conclude that APC analysis is still in its infancy (Kupper et al. 1985; Mason and Wolfinger 2002).

Recent developments in APC methodology in biostatistics have emphasized the utility of estimable functions that are invariant to the selection of constraints on the parameters (Holford 1983, 1991, 1992; Kupper et al. 1983; Kupper et al. 1985; Clayton and Schifflers 1987; Robertson, Gandini, and Boyle 1999). This is the approach applied by Fu (2000) in the derivation of a new APC estimator—termed the *intrinsic estimator*

(IE)—based on estimable functions and the singular value decomposition of matrices.

In earlier work (Yang, Fu, and Land 2004), we compared two approaches to the identification problem in APC accounting models: the IE method and the constrained generalized linear models (CGLIM) estimator that has been conventional among demographers and other social scientists for more than two decades (Fienberg and Mason 1978, 1985; Mason and Smith 1985). Through data analyses of population mortality rates, we illustrated some similarities and differences of these two methods in parameter estimates and model fit. We also discussed some key statistical properties of the IE compared with the CGLIM estimator.

From these results, we concluded that the IE offers a useful alternative to conventional methods for the APC analysis of tables of rates. The conceptual foundations of the IE are sufficiently abstract and difficult to understand, however, that additional exposition and illustration are merited. Accordingly, the objectives of this article are to (1) further describe the IE algebraically, geometrically, and verbally, (2) review properties of the IE as a statistical estimator, (3) report results of model validation assessments of the IE from both an empirical example and a simulation exercise, (4) give some usage advice and show how to relate the coefficients of the IE to those of conventional constrained APC models with applications to U.S. female mortality rates, 1960–99, and (5) introduce computer software that interested users can readily access.

Since the IE is a general-purpose method of APC analysis with potentially wide applicability in the social sciences, it is appropriate to recall the criteria for acceptability of such a general-purpose method, recently articulated by Norval Glenn, a long-time critic of attempts to provide general solutions to the APC analysis problem. Glenn (2005, p. 20) stated that such a method "may prove to be useful . . . if it yields approximately correct estimates 'more often than not,' if researchers carefully assess the credibility of the estimates by using theory and side information, and if they keep their conclusions about the effects tentative." These are strong criteria with which we agree. The purpose of this article is to assess the extent to which the IE satisfies them.

THE ALGEBRA OF THE APC IDENTIFICATION PROBLEM

We focus on the APC analysis of rectangular arrays of demographic rates arranged in conventional fashion, with age intervals defining the rows and time periods defining the columns. Specifically, in table 1 we analyze the same U.S. female mortality rates from 1960 through 1999 that we have previously studied (Yang et al. 2004). The data from which this table

TABLE 1
Deaths per 100,000: U.S. Females, 1960–99

Age	1960–64	1965–69	1970–74	1975–79	1980–84	1985–89	1990–94	1995–99
0	518	444	371	309	259	227	192	160
5	39	36	34	27	23	20	18	16
10	31	30	29	25	22	20	19	18
15	53	58	60	54	48	48	46	43
20	70	73	72	64	57	54	51	48
25	88	86	82	70	64	64	64	59
30	123	123	113	90	80	83	86	82
35	183	185	172	135	115	113	119	119
40	275	282	267	221	186	170	166	170
45	419	422	405	344	304	277	261	252
50	631	625	591	524	483	454	419	395
55	906	905	868	770	741	711	671	631
60	1,414	1,337	1,282	1,156	1,126	1,111	1,053	1,014
65	2,228	2,120	1,967	1,740	1,686	1,656	1,590	1,543
70	3,608	3,400	3,112	2,698	2,606	2,563	2,445	2,428
75	5,886	5,593	5,143	4,380	4,087	3,996	3,796	3,778
80	9,678	9,162	8,395	7,337	6,946	6,641	6,302	6,351
85	15,665	14,993	13,757	12,016	11,487	11,223	10,444	10,934
90	23,635	22,839	21,261	18,991	18,253	18,361	17,395	17,666
95+ ...	33,478	32,491	31,170	28,986	28,613	29,536	28,973	29,796

was constructed were obtained from the Berkeley Human Mortality Database (http://www.mortality.org). As is conventional in demographic and epidemiological analyses of arrays of this type, both age and period are measured in five-year intervals, and the diagonal elements of the matrices correspond to birth cohorts.

For example, the death rate for the 35–39 age group recorded in 1960–64 was 0.001835, or 183 deaths per 100,000 females in the population. By comparison, the mortality rate for the same age group in 1995–99 was 119 per 100,000. The complication that arises from the APC analysis of such arrays of rates is that the former rate (183/100,000) also corresponds to the mortality rate of the cohort of U.S. females born in the years 1925–29 and the latter rate (119/100,000) also corresponds to the mortality rate for the birth cohort of the years 1960–64.

The APC accounting/multiple classification model (Mason et al. 1973) is a conventional tool for the analysis of age-by-time-period arrays of demographic rates. We have already described the algebraic structure of this model (Yang et al. 2004) and, for the sake of expository completeness, we repeat it here. For mortality rates, this model can be written in *linear regression form* as

$$M_{ij} = D_{ij}/P_{ij} = \mu + \alpha_i + \beta_j + \gamma_k + \varepsilon_{ij}, \tag{1}$$

where M_{ij} denotes the observed death rate for the ith age group for $i = 1, \ldots, a$ age groups at the jth time period for $j = 1, \ldots, p$ time periods of observed data; D_{ij} denotes the number of deaths in the ijth group; P_{ij} denotes the size of the estimated population in the ijth group, the population at risk of death; μ denotes the intercept or adjusted mean death rate; α_i denotes the ith row age effect or the coefficient for the ith age group; β_j denotes the jth column period effect or the coefficient for the jth time period; γ_k denotes the kth diagonal cohort effect or the coefficient for the kth cohort for $k = 1, \ldots, (a + p - 1)$ cohorts, with $k = a - i + j$; and ε_{ij} denotes a random error with expectation $E(\varepsilon_{ij}) = 0$.

Conventional APC models as represented in model (1) fall into the class of generalized linear models (GLIM; see McCullagh and Nelder [1989] or McCulloch and Searle [2001] for expositions) that can take various alternative forms. First, model (1) can take a *log-linear regression form*, via a *log link*, as

$$\log(E_{ij}) = \log(P_{ij}) + \mu + \alpha_i + \beta_j + \gamma_k, \tag{2}$$

where E_{ij} denotes the expected number of deaths in cell (i, j) that is assumed to be distributed as a Poisson variate, and $\log(P_{ij})$ is the log of the exposure P_{ij} in model (1) and is called the "offset" or adjustment for the log-linear contingency table model. Models of this type are widely used in demography and epidemiology, where the counts of demographic events such as deaths or the incidence of diseases generally follow Poisson distributions and the rates are estimated through log-linear models (Agresti 1996). A second alternative formulation of the model is to treat the underlying number of deaths as a binomial variate. The canonical link changes from a log link to a *logit link*, which yields a *logistic model*,

$$\theta_{ij} = \log\left(\frac{m_{ij}}{1 - m_{ij}}\right) = \mu + \alpha_i + \beta_j + \gamma_k, \tag{3}$$

where θ_{ij} is the log odds of death and m_{ij} is the probability of death in cell (i, j). This model has been implemented more widely in demographic research (e.g., Mason and Smith 1985).

Regression models (1), (2), and (3) can be treated as *fixed-effects generalized linear models* after a reparameterization to center the parameters:

$$\sum_i \alpha_i = \sum_j \beta_j = \sum_k \gamma_k = 0. \tag{4}$$

After this reparameterization, model (1) can be written in the conventional matrix form of a least squares regression:

$$Y = Xb + \varepsilon, \tag{5}$$

where Y is a vector of mortality rates or log-transformed rates, X is the regression design matrix consisting of "dummy variable" column vectors for the vector (of which $m = 1 + [a - 1] + [p - 1] + [a + p - 2]$ represents the dimension) of model parameters b:

$$b = (\mu, \alpha_1, \ldots, \alpha_{a-1}, \beta_1, \ldots, \beta_{p-1}, \gamma_1, \ldots, \gamma_{a+p-2})^T. \tag{6}$$

The T superscript denotes vector transposition, and the ε in model (5) is a vector of random errors with mean 0. Note that the parameters α_a, β_p, and γ_{a+p-1} are not included in the parameter vector b, because of the constraints in equation (4), and can be uniquely determined by use of (4) in conjunction with each estimator of b. Later in this discussion, we will illustrate by an empirical example that the use of reference categories is equivalent to the translation by a constant of the parameter estimates produced by the constraints in equation (4) and thus of no substantive importance.

The ordinary least squares (OLS) estimator of the matrix regression model (5) is the solution b of the normal equation:

$$\hat{b} = (X^T X)^{-1} X^T Y. \tag{7}$$

But this estimator *does not exist*—there is no uniquely defined vector of coefficient estimates. This is because of the fact that the design matrix X is singular with one less than full column rank (Kupper et al. 1985), on account of the perfect linear relationship between the age, period, and cohort effects:

$$\text{period} - \text{age} = \text{cohort}.$$

Therefore, $(X^T X)^{-1}$ does not exist. This is the *model identification problem* of APC analysis. It implies that there are an infinite number of possible solutions of the matrix equation (7) (i.e., OLS estimators of model [5]), one for each possible linear combination of column vectors that results in a vector identical to one of the columns of X. Thus, it is not possible to separately estimate the effects of cohort, age, and period without imposing at least one constraint on the coefficients in addition to the reparameterization in equation (4).

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 defined in equation (6), namely, constraining the effect coefficients of the first two periods to be equal ($\beta_1 = \beta_2$). With this one additional constraint, the model (5) is just-identified, the matrix (X^TX) becomes nonsingular, and the least squares estimator (eq. [7]) exists (as do related maximum likelihood estimators for models [1], [2], and [3]).[2]

The main problems with this CGLIM approach have been discussed in a large body of methodological literature in demography, epidemiology, and statistics. First, the analyst needs to rely on external or side information to find constraints, but such information often does not exist or cannot easily be verified (Mason and Wolfinger 2002). Second, different choices of identifying constraints can produce widely different estimates of patterns of change across the age, period, and cohort categories. As we have previously demonstrated (Yang et al. 2004; see also, e.g., Mason and Smith 1985), estimates of model effect coefficients are sensitive to the choice of the equality coefficient constraint. Third, all just-identified models will produce the same levels of goodness of fit to the data, making model fit a useless criterion for selecting the best-constrained model. But no other criteria or guidelines have been proposed. These problems often lead to the conclusion that it is impossible to obtain meaningful estimates of the distinct effects of age, period, and cohort in analysis of social change.

DESCRIPTION OF THE IE

So what is new about the IE? The consensus has been that the key problem for APC analysis is to identify an estimable function that uniquely determines the parameter estimates. But there continues to be controversy over whether there exists such an estimable function that solves the identification problem. The conventional wisdom is that only the nonlinear, but not the linear, components of APC models can be estimable (Rodgers 1982*a*; Holford 1983). As noted in Fu (2008), however, there have been only numeric demonstrations, but no rigorous proofs, to support the idea that no estimable function exists. It should also be noted that Kupper et al. (1985) provided a condition for estimable functions and suggested that an estimable function satisfying this condition resolves the identification problem. Subsequent publications have shown in that the IE satisfies this

[2] There are other solutions to APC identification problems. The *proxy variables approach* uses one or more proxy variables as surrogates for the age, period, or cohort coefficients (see, e.g., O'Brien 2000). The *nonlinear parametric transformation approach* defines a nonlinear function of one of the age, period, or cohort variables so that its relationship to others is nonlinear (see, e.g., Fienberg and Mason 1985; Yang and Land 2006). These approaches do not take the form of an APC accounting model and are not directly comparable to the IE in statistical terms. As will be illustrated later, however, comparison of results obtained by different approaches provides a means of model validation.

condition and estimates the unique estimable function, including both the linear and nonlinear components of the parameter vector of the multiple classification model (Fu 2000, 2008; Fu, Hall, and Rohan 2004; Yang et al. 2004).

Within the context of the foregoing description of the algebra of the APC identification problem in conventional linear regression models, we next describe the IE in three ways: algebraically, geometrically, and verbally.

Algebraic Definition

We have shown previously (Yang et al. 2004) that, because the design matrix X is one less than full column rank, the parameter space of the unconstrained APC regression model (5) can be decomposed into the direct sum of two linear subspaces that are perpendicular to each other. One subspace corresponds to the unique zero eigenvalue of the matrix X^TX of equation (7) and is of dimension 1; it is termed the null subspace of the design matrix X. The other, non-null subspace is the complement subspace orthogonal to the null space.

Because of this orthogonal decomposition of the parameter space, each of the infinite number of solutions of the unconstrained APC accounting model (5) can be written as

$$\hat{b} = B + sB_0, \tag{8}$$

where s is a scalar corresponding to a specific solution and B_0 is a unique eigenvector of Euclidean norm or length 1. The eigenvector B_0 does not depend on the observed rates Y, only on the design matrix X, and thus is completely determined by the numbers of age groups and period groups—regardless of the event rates. In other words, B_0 has a specific form that is a function of the design matrix. To construct an explicit representation of B_0, note first that the exact linear dependency between the age, period, and cohort variables in model (5) is mathematically equivalent to

$$XB_0 = 0. \tag{9}$$

This equation expresses the property that X is singular; that is, there exists a linear combination of the columns of the design matrix X that equals a zero vector. Kupper et al. (1985) showed that B_0 has the algebraic form

$$B_0 = \frac{\tilde{B}_0}{\left\|\tilde{B}_0\right\|}, \tag{10}$$

meaning that B_0 is the normalized vector of $\tilde{B}_0$:

$$\tilde{B}_0 = (0, A, P, C)^T, \tag{11}$$

where $A = (1 - \frac{a+1}{2}, \ldots, [a-1] - \frac{a+1}{2})$, $P = (\frac{p+1}{2} - 1, \ldots, \frac{p+1}{2} - [p-1])$, and $C = (1 - \frac{a+p}{2}, \ldots, [a+p-2] - \frac{a+p}{2})$, with a, p, and c denoting the number of age categories, time periods, and cohorts in the age-by-time-period array of rates. It is important to note that the vector B_0 is fixed, or nonrandom, because it is a function solely of the dimension of the design matrix X or 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 the conventional CGLIM approach violates this principle if the scalar s in equation (8) is nonzero.

The idea that B_0 should not affect the results is a key point, as intuition suggests that the eigenvector corresponding to the zero eigenvalue should be an arbitrary vector. And indeed, sB_0 is arbitrary. But B_0 is not arbitrary; it is fixed by the design matrix. Furthermore, with equation (8), any APC estimator obtained by placing any identifying constraint(s) on the design matrix can be written as the linear combination $B + sB_0$, where B is the special estimator termed the IE that lies in the parameter subspace that is orthogonal to the null space and determined by the Moore-Penrose generalized inverse.[3]

One computational algorithm for the IE is a principal components regression method, whereby the user (a) computes the eigenvalues and eigenvectors (principal components) of the matrix X^TX; (b) normalizes them to have unit length; (c) identifies the eigenvector B_0 corresponding to the unique eigenvalue 0; (d) estimates a (principal components) regression model with response vector Y as in model (5) and design matrix U, whose column vectors are the principal components determined by the eigenvectors of nonzero eigenvalues; and then (e) uses the orthonormal matrix of all eigenvectors to transform the coefficients of the principal components regression model to the regression coefficients of the intrinsic estimator B.

Geometric Representation

The parameter space of the unconstrained vector b can be decomposed into two parts that are orthogonal or independent in relation to each other:

$$b = b_0 + sB_0, \tag{12}$$

[3] See, e.g., Searle (1971, pp. 16–19) for a definition of the Moore-Penrose generalized inverse and its properties.

where $b_0 = P_{\text{proj}} b$ is a special parameter vector that is a linear function of b, corresponding to the projection of the unconstrained parameter vector b to the non-null space of X. Specifically, the special parameter vector b_0 corresponding to $s = 0$ satisfies the geometric projection

$$b_0 = (I - B_0 B_0^T) b. \tag{13}$$

This projection is illustrated in figure 1, which shows the projection of two parameterizations, b_1 and b_2, onto the non-null parameter space (the vertical axis in fig. 1), which is independent of the real number s. The geometric representation in figure 1 can be thought of either as a simple parameter space of dimension 2 or as multidimensional, with the vertical axis representing a direction in a multidimensional non-null space. In either case, since the projection of any parameterization of b in figure 1 yields the same parameter vector b_0, the latter is estimable.[4]

Figure 1 also helps to illustrate geometrically that the IE may in fact also be viewed as a constrained estimator. But, in contrast to the equality constraints on two or more *coefficients* of the parameter vector b that are imposed in conventional approaches to the estimation of APC accounting models, the constraint imposed by the IE to identify model (5) is a constraint on the *geometric orientation of the parameter vector* b *in parameter space*. Specifically, the IE imposes the constraint that the direction in parameter space defined by the eigenvector B_0 in the null space of the design matrix X have zero influence on the parameter vector b_0 (i.e., on the specific parameterization of the vector b that is estimated by the IE). Since B_0 is a fixed vector that is a function solely of the design matrix (e.g., the number of time periods of data in an analysis) and does not depend on the observed event rates or frequencies being analyzed, this seems to be a reasonable constraint.

Corresponding to the projection of the parameter vector b onto b_0, we have the following projection of the estimators of equation (13) onto the IE B:

$$B = (I - B_0 B_0^T) \hat{b}. \tag{14}$$

This equation provides another algorithm for computing the IE, one that computes an initial estimator $\hat{b}$ of model (5), using, for instance, an equality constraint on two of the age, period, or cohort parameters, and then geometrically projects $\hat{b}$ to the IE B by removing the component in the B_0 direction.

[4] Fig. 1 provides only a geometric illustration. Full algebraic details of the proof that b_0 is estimable and the *only* estimable function that determines both the linear and nonlinear trends in the age, period, and cohort coefficients are given in Fu et al. (2004).

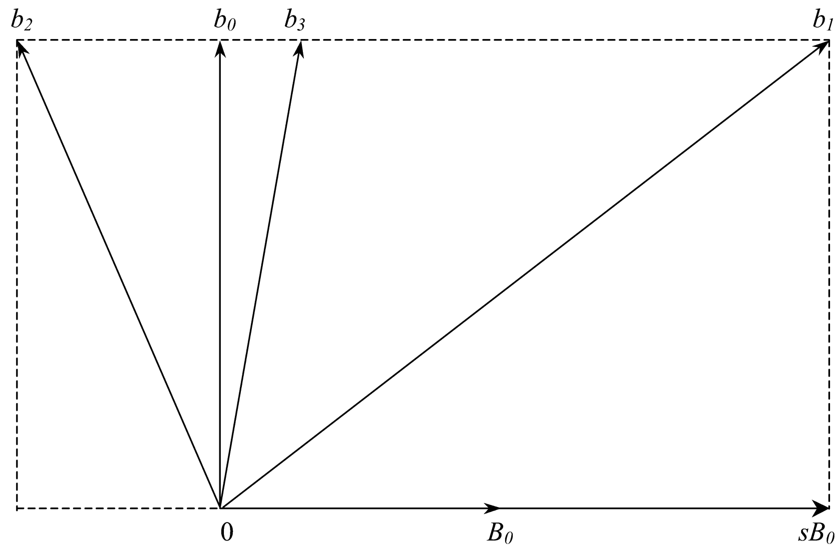

FIG. 1.—Geometric projection of parameter vectors

Verbal Description

Statisticians have known since Kupper et al. (1985) that the dimension of the design matrix (i.e., the numbers of age groups and time periods) in the APC accounting model may affect the estimates obtained from CGLIM estimators. Put in the simplest possible terms, the basic idea of the IE is to remove the influence of the design matrix on coefficient estimates. As noted later herein, this approach produces an estimator that has desirable statistical properties.

The IE also can be viewed as a special form of principal components regression estimator (for a standard exposition of principal components regression, see, e.g., Sen and Srivastava [1990]) that removes the influence of the null space of the design matrix X on the estimator. It specifically estimates a constrained parameter vector b_0 that is a linear function of the parameter vector b of the unconstrained APC accounting model (5). This constrained parameter vector b_0 corresponds to the projection of the unconstrained parameter vector b onto the non-null subspace of the design matrix X.

Since the IE is a principal components estimator, one might well ask, Why not just calculate the eigenvectors of the matrix X^TX by application of principal components, regress the observed rates on the subspace spanned by these eigenvectors, and leave it at that? The answer is that regression coefficient estimates in this subspace are not directly interpretable in terms of age, period, and cohort effects. Therefore, the IE uses the extra step of inverse orthonormal transformation of the coefficient estimates of the principal components regression back to the original space of age, period, and cohort coordinates. The inverse transformation is what makes the IE a special form of principal components estimator. It yields coefficients, as will be illustrated by numerical example below, that are directly interpretable as age, period, and cohort effects and that can be compared to corresponding effects estimated by the conventional imposition of equality constraints on parameters.

STATISTICAL PROPERTIES AND VALIDATION OF THE IE

Statistical Properties

Earlier, we (Yang et al. 2004) stated and proved some properties of the IE as a statistical estimator that we briefly summarize here. For context, consider first the analysis of an APC data set for a finite number of time periods p. That is, suppose that an APC analysis is to be conducted for a fixed matrix of observed rates or event counts. This implies that the corresponding design matrix X is fixed (i.e., X has a fixed number of age groups and time periods). The randomness in the error term ε of model (5) then corresponds to measurement errors in the rates or in the event counts and/or to intrinsic randomness in the rates or counts.

In this context of an age-by-time-period table of population rates with a fixed number of time periods p of data, it has been shown (Yang et al. 2004, p. 101) that the IE satisfies a condition for estimability of linear functions of the parameter vector b that was established by Kupper et al. (1985, app. B) and recently further elaborated by Fu (2008). *Estimable functions* are invariant with respect to whatever solution (see eq. [7] above) to the normal equations is obtained;[5] these functions are desirable as statistical estimators because they are linear functions of the unidentified parameter vector that can be estimated without bias (i.e., they have unbiased estimators).[6] Specifically, the condition for estimability of a con-

[5] See Searle (1971, pp. 180–88) or McCulloch and Searle (2001, pp. 120–21) for an exposition of the concept of estimable functions.

[6] In the history of discussions of the APC accounting model in sociology, Rodgers (1982*a*) was early to argue that analysts should seek estimable functions of the unidentified parameter vector (eq. [6]); see also the comment by Smith, Mason, and Fienberg (1982) and the response by Rodgers (1982*b*). In some respects, the IE can be regarded as providing a practical, easily applicable method for producing estimates of estimable functions from data in the form of age-by-time-period tables of rates, as called for by Rodgers over two decades ago. The estimability referred to by Rodgers, however, essentially means identifiability that can be achieved by any linear constraints, which differs from statistical estimability defined for APC models by Kupper et al. (1985).

straint on the parameter vector that was established algebraically by Kupper et al. (1985) is, in the notation defined above, that $l^T B_0 = 0$, 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 $s = 0$ (i.e., that the arbitrary vector B_0 have zero influence), $l^T = (I - B_0 B_0^T)$ for the IE. Since $B_0^T B_0 = 1$, it follows that $l^T B_0 = (I - B_0 B_0^T) B_0 = B_0 - B_0 B_0^T B_0 = B_0 - B_0 = 0$; the Kupper et al. condition holds for the IE. Note also that the Kupper et al. condition implies that any constrained estimator that is obtained by imposing an equality constraint on the parameter vector b and that contains any nonzero component due to the vector B_0 defined by the design matrix is not estimable; it produces biased estimates of the A, P, and C effect coefficients.

Because the IE B satisfies the estimability condition for APC models, it follows that, for a fixed number of time periods of data, the IE B is an unbiased estimator of the special parameterization (or linear function) b_0 of b defined in equation (12).[7] Thus, a first statistical property of the IE in the context of an APC analysis of demographic rates with a fixed number of time periods of data is that it produces unbiased estimates of the regression coefficients of the projected coefficient vector b_0. Second, we have also shown (Yang et al. 2004, p. 108) that, for a fixed number of time periods of data, the IE is more statistically efficient (has a smaller variance) than any CGLIM estimator that is obtained from a nontrivial equality constraint on the unconstrained regression coefficient estimator b—that is, any equality constraint that does not produce a projection of b onto b_0.

In brief, the IE has nice finite-time-period properties. In addition to finite-time-period properties, the asymptotic properties of APC estimators as the number of time periods p of data increase have been studied by Fu, Hall, and Rohan (2004), Fu and Hall (2006), and Fu (2008). These properties derive largely from the fact that the eigenvector B_0 converges elementwise to zero with increasing numbers of time periods of data. The vector can converge elementwise to zero even though its length is fixed at 1 because the number of elements of the vector grows as we add time periods. Therefore, for any two estimators $\hat{b}_1 = B + s_1 B_0$ and $\hat{b}_2 = B + s_2 B_0$, where s_1 and s_2 are nonzero and correspond to different identifying constraints placed on model (5), as the number of time periods in an APC analysis increases, the difference between these two estimators decreases toward zero, and, in fact, the estimators converge toward the IE B. Suffice it to say that the proof proceeds by demonstrating that the coordinates of B_0 are bounded by a quantity that is a function of the number of age groups and periods and this function converges to zero as $p \rightarrow \infty$. Under

[7] In our earlier work (Yang et al. 2004, p. 107), we also proved this property directly.

suitable regularity conditions on the error term process and a fixed set of age categories with effect coefficients, this feature of B_0 yields a convergence of the IE asymptotically to these "true" effect coefficients. In addition, Fu and Hall (2006) argue that as the number of time periods of data increases, there are definite bounds on the differences between the effect coefficients estimated by the IE and the effect coefficients of the true period and cohort processes. A corollary is that CGLIM estimators $\hat{b}$ with nonzero s in equation (8) of the parameter vector in model (5) are biased in finite-time-period APC analyses. Except under certain conditions specified by Fu, Hall, and Rohan (2004), this bias decreases as the number of time periods in the analysis increases, in which case the CGLIM estimators may converge to the IE estimator.

Remark 1.—These statistical properties are not trivial and merit comment. Both the intrinsic estimator B and any other estimator $\hat{b} = B + sB_0$ with $s \neq 0$ obtained from an equality constraint are asymptotically consistent as the number of time periods of data increases without bound. Therefore, with a large number of periods of data (e.g., 30 or 40), 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 of time periods of observations (e.g., 4 or 5) available for analysis. In these cases, the differences can be substantial.

Remark 2.—As just noted, the IE, by its very definition and construction, satisfies the estimability condition. Other estimators using theoretically derived equality constraints on the parameter vector b may satisfy this condition either exactly or statistically (in a sense defined below). 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.

Remark 3.—These properties provide a means for differentiating among estimators. That is, for tables of rates with a finite number of time periods of data, especially a small number (e.g., 4 or 5), an unbiased estimator should be preferred to a biased estimator, as the latter can be misleading with respect to the estimated trends across the age, time period, and cohort categories.

Remark 4.—In contrast, as has been noted many times over the years in discussions of the APC accounting model (see, e.g., Pullum 1978, 1980; Rodgers 1982*a*, 1982*b*; Smith, Mason, and Fienberg 1982), different just-identified models will generate the same data and yield exactly the same model fit. In particular, linear transformations of the estimated A, P, and C coefficients obtained by the IE (i.e., linear transformations of the elements of the B vector) will fit the observed data just as well as the IE.[8] If, however, such a linear transformation of the A, P, and C coefficients (or any subset thereof) results in coefficients that depart sufficiently far from the coefficients in B that they contain a significant component of the B_0 vector (i.e., a significantly nonzero s coefficient), then the resulting transformed vector will not be estimable—that is, it will not be unbiased.[9] Thus, even though the transformed coefficients will reproduce the data just as well as those obtained by the IE, they will be biased and will give poor indications of the patterns of change across the age, period, and cohort categories used in the analysis. Therefore, goodness of fit to the data (as measured, e.g., by log-likelihood functions or deviance statistics) *cannot* be used as a criterion for selecting among estimators. But estimability can be used for this purpose.

Remark 5.—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. To see the intuitive logic of this statement, recall, for example, the distinction in classical mechanics between the steady-state and general solutions to the ordinary differential equation for Hooke's law on the motion of a displaced spring-mass system subject to an additional forcing motion. This law has the algebraic form $F = -kx + a\cos(\omega t)$, where F denotes acceleration (second derivative of the motion with respect to time) of the mass, x denotes the distance of displacement, k is a constant unique to the particular spring under study, and $a\cos(\omega t)$ is the forcing term (see, e.g., Marion and Thornton 1995, p. 125). If the mass is displaced by, say, a distance of two feet, it will oscillate back and forth with some influence of the length of the initial displacement, but it eventually will settle down to a characteristic pattern of oscillations that depends only on the driving force. By comparison, if the initial displacement of the mass is a distance of four feet, then the mass will display an initial set of larger oscillations that are different from the pattern observed for the two-foot displacement. But after an initial series of oscillations, the driving force will cause the mass to settle into the same set of oscillations as those found after a two-foot displacement. Mathematically, the pattern of oscillations observed after the effect of the initial length of displacements has worn off are termed the *steady-state solution* of the Hooke's law differential equation,

[8] Geometrically, a linear transformation of the coefficient vector corresponds to a rotation of the vector in parameter space. Such a rotation will produce distorted and misleading indications of patterns of change across the age, period, and cohort categories used in the analysis.

[9] This point is illustrated in our empirical analyses and is the basis of the statistical test for estimability derived below.

whereas the *general solution* of the equation consists of the steady-state solution plus a factor that takes into account the initial conditions or displacement of the spring. Because initial conditions can vary from application to application, the general solution of the differential equation can be unique to the application. On the other hand, the steady-state aspect of the solution is invariant and generalizable to the motion of the system regardless of the initial conditions.

Analogously, the IE is essentially a steady-state solution to the APC accounting model estimation problem that factors out the initial conditions of the dimension of the matrix of observed data—namely, the number of age and time period categories that define the design matrix. Because the IE does not allow these "initial conditions" to influence the estimates it produces of the A, P, and C effect coefficients, they will be more invariant to changes in the design matrix, such as additional time period data, than estimates produced by estimators that incorporate such influences. In this sense, the IE removes the aspect of the subjectivity in the estimator that is due to the shape of the data. We will illustrate this feature of the IE below in our model validation analyses.

Validation: An Empirical Example

In brief, the IE has some valuable properties as a statistical estimator. But, 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 effect coefficients that are valid—that reveal the true effects. This is a question of *model validation*: does the identifying constraint imposed by the IE, the projection of the unconstrained APC accounting model vector b onto the non-null space of the design matrix X, as in equation (12), produce estimated coefficients that meaningfully capture the true age, time period, and cohort trends?

One approach to the question of validity is to compare results from an APC analysis of empirical data by application of the IE with results from an analysis of the same empirical data by application of another approach, with a different family of models that do not use the identifying constraint of the IE or CGLIM estimators for the accounting model. As an instance of such an empirical comparison, we next describe analyses of verbal test-score data from 15 cross-sections of the General Social Survey (GSS) from 1974 to 2000.[10] This is an extension of GSS verbal test data for 1974–96; a series of articles on these data in 1999 (see n. 11 below) posed and disputed the existence of an intercohort decline in verbal vocabulary.

[10] The survey years in which verbal ability was tested are 1974, 1976, 1978, 1982, 1984, 1987–90, 1991, 1993, 1994, 1996, 1998, and 2000.

In the GSS tests, a survey respondent's vocabulary knowledge is measured by a composite scale score called WORDSUM, which is constructed by adding the correct answers to 10 verbal test questions and which ranges from 0 to 10. WORDSUM has a distribution that is approximately bell shaped, which has a mean of about 6 and is reported in previous studies to have an internal reliability of .71 (Wilson and Gove 1999*a*, p. 258; 1999*b*).[11] The data include 19,500 respondents who had WORDSUM scores and other covariate measures across all survey years. Respondents' ages in the data pooled across all surveys range from 18 to 89. The average amount of education completed is 12.7 years. Fifty-seven percent of respondents are female, and 15% are black. There are 19 five-year birth cohorts. The oldest cohort member was born in 1890, and the youngest was born in 1980.

We apply the IE to data on verbal test scores grouped into five-year age groups and time periods, shown in table 2. In this age-by-period array, there are 12 five-year age groups from ages 20 to 75+,[12] five five-year period groups from 1976 to 2000, and $12 + 5 - 1 = 16$ cohorts born from 1901 to 1976.[13] This yields 60 degrees of freedom. The event/exposure rates (sample mean proportions) of correct answers to GSS vocabulary questions can be transformed by a log link and modeled by a log-linear regression for which the IE can be obtained. The events are the total number of correct answers for every age and period group and are calculated by multiplying the mean verbal scores by the number of individuals. These are nonnegative counts and can be considered to be distributed as Poisson variates. The population exposure is calculated as the product of the number of people in each cell and the total number of possible correct answers (10). Vocabulary knowledge test scores are not available for every year from 1974 to 2000. For the missing years, we interpolated the mean verbal scores and the numbers of individuals at risk based on the data of neighboring years.

[11] In an item analysis of individual words in WORDSUM, Alwin (1991, p. 628) found that some of the words have become more difficult over time. The general conclusion in the series of articles on the GSS verbal test data (Alwin and McCammon 1999; Glenn 1999; Wilson and Gove 1999*a*, 1999*b*), however, is that word obsolescence does not account for observed changes in the test scores over time.

[12] We excluded ages 18–19 to obtain age groups of equal interval length so that the diagonal elements of the age-by-period matrix refer to cohort members. The results are not influenced by omitting this cell of small sample size. Ages 75–89 are grouped into the 75+ category in order to combine small population exposures for more stable estimates.

[13] The definition of the width of the time intervals used to define birth cohorts in APC analyses is somewhat arbitrary. For consistency with extant articles on trends in GSS vocabulary knowledge data, we use five-year birth cohorts.

TABLE 2
VERBAL TEST CORRECT ANSWER RATES: GSS 1976–2000

Age Group	1976–80	1981–85	1986–90	1991–95	1996–2000
20–24	.554 (8,203)	.531 (8,489)	.526 (5,553)	.548 (4,950)	.555 (5,045)
25–29	.596 (9,671)	.554 (10,653)	.577 (6,946)	.589 (6,480)	.575 (6,780)
30–34	.630 (8,643)	.634 (8,718)	.588 (7,096)	.613 (7,745)	.596 (7,740)
35–39	.625 (6,776)	.643 (7,946)	.632 (6,799)	.634 (7,960)	.622 (8,400)
40–44	.622 (5,619)	.633 (5,996)	.642 (5,548)	.652 (6,940)	.629 (8,355)
45–49	.609 (4,595)	.603 (5,112)	.617 (4,734)	.675 (5,915)	.653 (6,775)
50–54	.617 (5,454)	.597 (4,560)	.594 (3,314)	.617 (4,770)	.661 (5,885)
55–59	.630 (5,439)	.594 (5,335)	.589 (3,207)	.623 (3,590)	.676 (4,225)
60–64	.625 (4,915)	.596 (4,830)	.561 (3,319)	.610 (3,430)	.604 (3,700)
65–69	.605 (4,010)	.576 (498)	.575 (3,633)	.648 (3,000)	.596 (2,985)
70–74	.523 (3,420)	.571 (3,333)	.610 (2,789)	.587 (3,305)	.607 (3,540)
75+	.549 (4,394)	.532 (4,264)	.538 (3,584)	.573 (4,890)	.563 (4,970)

NOTE.—Numbers in parentheses indicate exposure.

Figure 2 shows the results from application of the IE to the data in table 2. Estimated coefficients and their 95% confidence intervals are plotted for successive categories within the age, period, and cohort classifications. Since they indicate changes in correct answer rates from one age group, time period, or cohort to the next, the estimated coefficients represent the temporal trends of vocabulary knowledge along each of these three dimensions, net of the effects of the other two. The age effect coefficients display a concave pattern. This corroborates the quadratic age effect found by Wilson and Gove (1999*a*, 1999*b*): low at youth, rising to a peak in the forties, staying largely flat until the midfifties, and declining gradually into late life. The period effects curve is fairly flat for the first 10 years, from the mid-1970s to the mid-1980s, but this is followed by a jump into the mid-1990s, which flattens out again until 2000. The slightly zigzag shape of the period curve shows some variation in vocabulary knowledge over time for the past 30 years. This is quite revealing, given the absence of direct estimates of period effects in the previous studies. The cohort effects are more complicated and are characterized by a bimodal curve, with the peaks occurring for cohorts born in 1911–16 and 1941–46. The intercohort declines, as Glenn (1994, 1999) and Alwin (1991; see also Alwin and McCammon 1999) insist, are evident for cohorts between and after these periods—that is, among the post–World War I and post–World War II cohorts. But there are also increases in vocabulary knowledge for cohorts born before World War I and for the 1930–50 cohorts. These graphical results give us an overview of the average APC trends that help to inform and specify models for additional APC analyses using individual-level data.

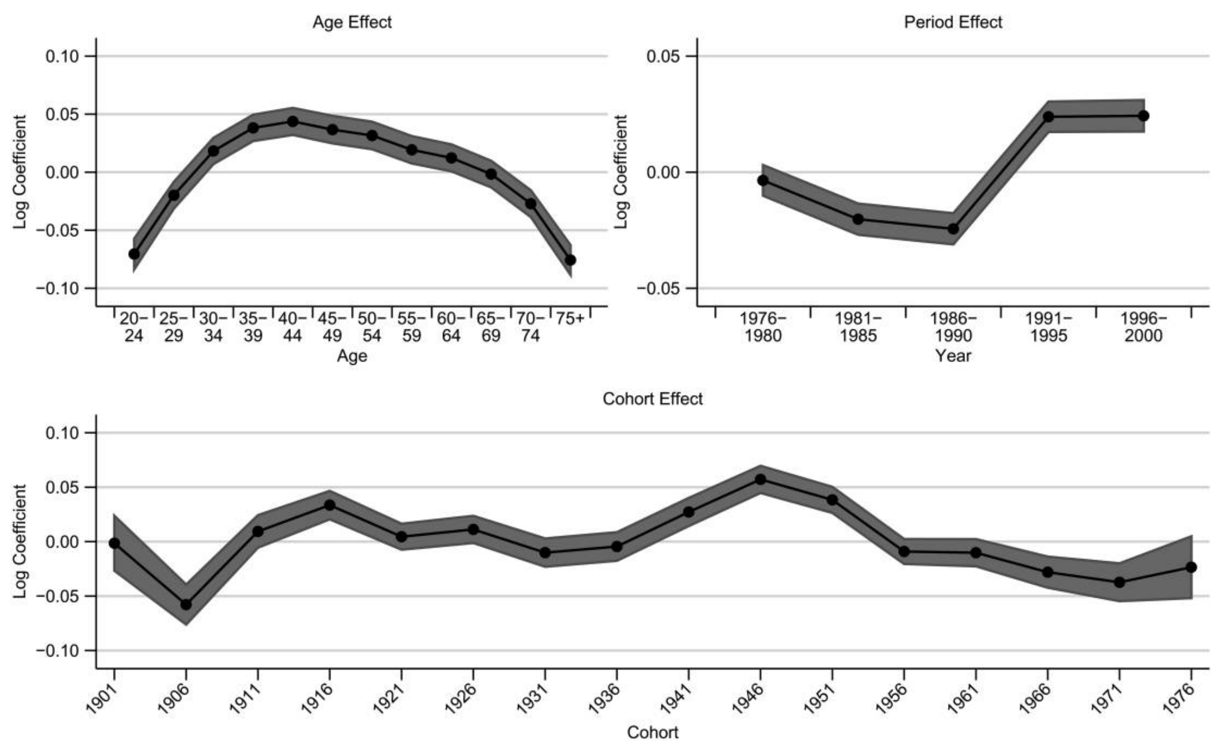

FIG. 2.—IE coefficient estimates and 95% confidence intervals of the age, period, and cohort effects on vocabulary ability using GSS data, 1976–2000.

The GSS data on verbal test scores were also analyzed (Yang 2006; Yang and Land 2006) using a very different approach to APC analysis—namely, a hierarchical age-period-cohort (HAPC) analysis in the form of a cross-classified random-effects regression model. This approach proceeds by building a level-1 fixed-effects regression model at the individual level of analysis and then a random-effects model for cohort and time period effects at level 2. The resulting analysis takes full advantage of the existence of individual-level data on all respondents to the GSS, while at the same time nesting these within the time periods and cohorts to which they correspond. The mixed-effects model does not require the assumption of additive and fixed age, period, and cohort effects used by the conventional linear regression models that cause the identification problem. The resulting estimates of cohort and period effect coefficients estimated in this hierarchical regression models approach are average residual effects of the cohort and period across all time periods and cohorts, respectively, and are not constrained in any way to conform to the IE or any other identifying constraint required by the account model.

The results of the HAPC analysis are shown in figure 3, which is reproduced from Yang's (2006) Bayesian estimates of the HAPC model. While the coefficient metrics and time scales for the age and period effects in figure 3 are in single years rather than the five-year groups of those from the IE analysis in figure 2, the trends of estimated effect coefficients are quite similar across the graphs. That is, both figure 2 and figure 3 exhibit age-effect curves that are quadratic and period-effect curves that show a slight decline from the 1970s into the 1980s and then a slight rise into the 1990s. And the cohort-effect curves are quite similar, showing peaks in the early and mid 20th century, followed by declines.

In brief, this comparison shows that the independent estimates from a hierarchical models analysis corroborate the estimated patterns of change across age, period, and cohort categories that are obtained by imposing the identifying assumption of IE analysis. Of course, this is only one example of comparative analysis, and additional empirical studies are needed before it can be concluded that the IE produces substantively meaningful and empirically valid results under various circumstances. Any statistical model has its limitations and surely will break down under some circumstances. The identification of these limits for the IE is a topic that merits additional research.

Validation: Simulation Results

Following good statistical practice, we now consider another approach to model validation: simulation analyses. It is well known that, given any age-by-period array of rates, it is impossible to know what true effects

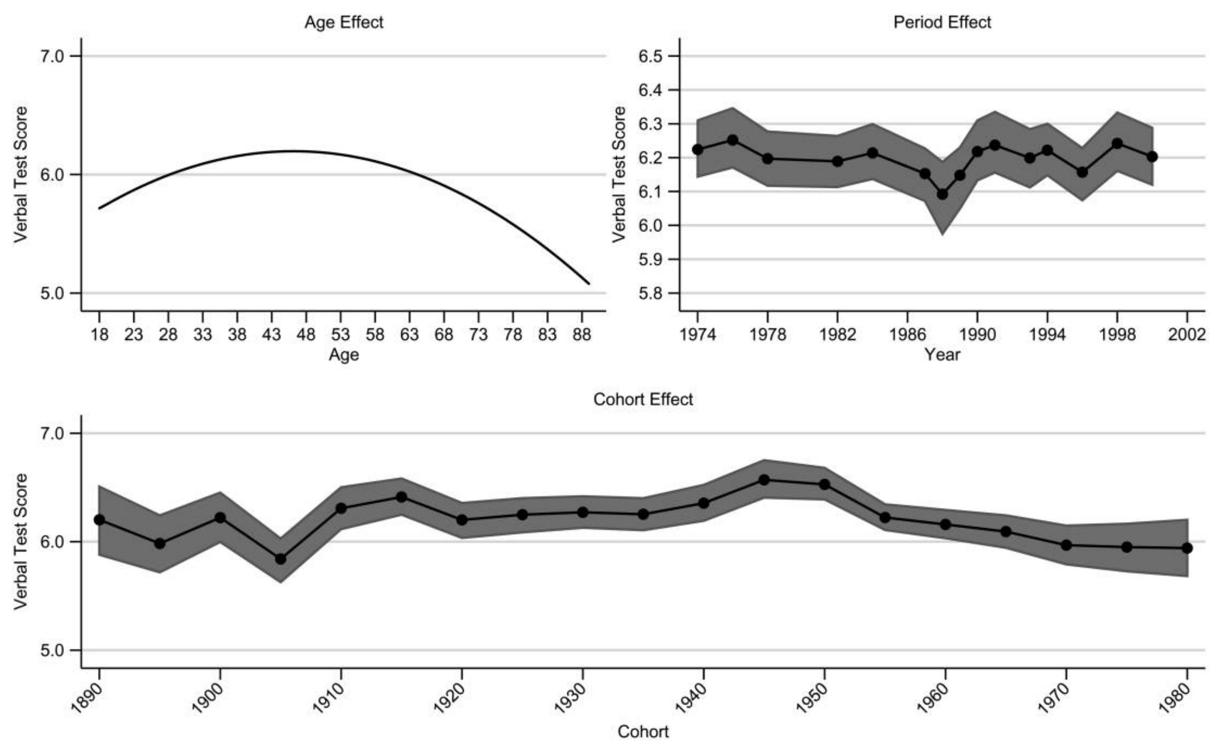

FIG. 3.—Bayes estimates of age, period, and cohort effects on verbal scores. Source: Yang 2006.

generated the data, because any number of just-identified constrained models fit the data equally well. It is therefore important that we apply an estimator such as the IE or CGLIM to artificial data for which we know the true form of the underlying model. We present here the results of some Monte Carlo simulation analyses to determine whether IE and CGLIM estimators recover the true parameters.

We first investigate whether the IE is unbiased in finite samples. The asymptotic properties of the IE apply as the number of periods in the data set goes to infinity. Any data set used in practice has only a finite number of periods. We explore whether the asymptotic results give good approximations of the behavior of the IE in finite samples by simulating data sets with five and 50 periods. The basic asymptotic result we expect 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.[14]

We fix the number of ages in all of our simulations at 9 without loss of generality. For a given number of periods P, we generate 1,000 data sets by Monte Carlo simulation in which the entries in the $P \times 9$ outcome matrix are distributed according to the equation

$$\begin{aligned} y_{ij} \sim \text{Poisson}\Big\{&\exp[0.3 + 0.1(\text{age}_{ij} - 5)^2 + 0.1\sin(\text{period}_{ij}) \\ &+ 0.1\cos(\text{cohort}_{ij}) + 0.1\sin(10 \cdot \text{cohort}_{ij})]\Big\}. \end{aligned}$$

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

Age effect at age a	$.1(a-5)^2$
Period effect in period p	$.1\sin(p)$
Cohort effect in cohort c	$.1\cos(c) + .1\sin(10c)$

We subtract constants from the above values so that the true effects are normalized to have mean 0 in each category according to equation (4), where the constants can be calculated as the mean effects for each category (see n. 15 below). We then estimate age, period, and cohort effects in each simulated data set using the IE and 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).

[14] 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 data set does not add information about the previous periods or about cohorts not present in the period just added.

TABLE 3
SIMULATION RESULTS OF THE IE AND CGLIM ESTIMATORS: AGE EFFECTS (N = 1,000)

		TRUE VALUE	IE		CGLIM_A (A1 = A2)		CGLIM_P (P1 = P2)		CGLIM_C (C1 = C2)	
			p = 5	p = 50	p = 5	p = 50	p = 5	p = 50	p = 5	p = 50
1	Mean	.933	.924	.942	−1.903	−1.876	1.214	1.078	1.714	1.211
	SD		.276	.060	1.265	.348	1.366	1.308	5.294	2.965
	MSE		.076	.004	9.646	8.011	1.942	1.730	28.609	8.862
2	Mean	.233	.217	.238	−1.903	−1.876	.435	.340	.810	.440
	SD		.280	.077	1.265	.348	1.041	.983	3.967	2.223
	MSE		.079	.006	6.164	4.569	1.123	.977	16.053	4.980
3	Mean	−.267	−.303	−.267	−1.717	−1.676	−.158	−.199	.092	−.132
	SD		.343	.094	.798	.209	.742	.653	2.665	1.486
	MSE		.119	.009	2.740	2.028	.562	.431	7.221	2.223
4	Mean	−.567	−.629	−.574	−1.335	−1.279	−.556	−.540	−.431	−.507
	SD		.401	.112	.536	.146	.526	.351	1.384	.751
	MSE		.165	.013	.878	.528	.276	.124	1.932	.567
5	Mean	−.667	−.721	−.674	−.721	−.674	−.721	−.674	−.721	−.674
	SD		.428	.121	.428	.121	.428	.121	.428	.121
	MSE		.186	.015	.186	.015	.186	.015	.186	.015
6	Mean	−.567	−.599	−.576	.108	.128	−.671	−.610	−.796	−.644
	SD		.397	.106	.518	.141	.514	.343	1.387	.749
	MSE		.158	.011	.723	.503	.275	.119	1.974	.566
7	Mean	−.267	−.240	−.266	1.173	1.143	−.386	−.334	−.636	−.401
	SD		.337	.097	.813	.218	.758	.661	2.660	1.484
	MSE		.114	.009	2.733	2.034	.588	.440	7.203	2.217
8	Mean	.233	.304	.237	2.424	2.350	.086	.135	−.289	.035
	SD		.288	.077	1.114	.295	1.076	.981	3.997	2.226
	MSE		.088	.006	6.039	4.569	1.177	.971	16.233	4.989
9	Mean	.933	1.048	.940	3.875	3.758	.757	.804	.257	.671
	SD		.268	.060	1.437	.390	1.366	1.304	5.272	2.962
	MSE		.085	.004	10.716	8.132	1.894	1.716	28.221	8.835

To make the CGLIM results comparable to the IE results, we renormalize, using procedures described in the next section of the article, each set of age, period, and cohort effects to sum to zero for both the CGLIM estimators and the IE.

Table 3 reports the results for age effects. For each age 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 see whether each estimator is unbiased. The standard deviation of the simulated estimates shows how much the estimated parameters vary from sample to sample. The MSE is the average squared difference between the estimated parameter and the true value; this measure of accuracy takes account of both bias and variance.

Three of the four estimators recover the qualitative U-shaped profile of the age effects. The CGLIM estimator that constrains the first two age effects to be equal fails to recover even the qualitative form of the age effects, because the constraint is far from correct: the first two true age effects are not nearly equal. However, even the CGLIM estimators that recover the qualitative shape of the age effects are far off the mark in quantitative terms; only for the IE is the mean of each estimated age effect close to the true value. The IE also exhibits substantially less sampling variation than the CGLIM estimators and has much lower MSE. The IE performs better than the CGLIM whether we use five or 50 periods of data.

Because the numbers of period and cohort effects in our samples are large, we report our results for these effects graphically in figure 4 for the simulations with five time periods of data. The figure shows the mean of the simulated estimates and the MSE for each coefficient and each estimator. Again, the mean for the IE is close to the true value, while the CGLIM estimators are not close, and the MSE is much smaller for the IE than for the CGLIM. It is noteworthy that the cohort effects are particularly poorly estimated by the CGLIM models.

We also investigated how well the estimators perform when there are no true period or cohort effects. Figure 5 shows results from simulations in which the age and period effects were as described above but the cohort effects were all zero. Only CGLIM_p estimates are shown among all CGLIM estimates because they are the closest to the true effects. On average, the IE correctly estimates that there are no cohort effects in the data. The CGLIM estimators incorrectly find cohort effects that are different from zero and change substantially across birth cohorts. Results were similar when we set all of the period effects to zero in the data: only the IE detected that no period effects were truly present.

In simulations not reported here, we found similar results for data sets with 10 and 20 periods of data.[15] The simulations thus show that, regardless of the sample size, the IE is more accurate than the CGLIM. The IE performs well even in data sets 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 50 periods, which would be an unusually large sample in many demographic applications.

[15] Due to space constraints, only a limited set of simulation results are reported here. Additional results, including analyses of other models in the GLM family and formal tests for unbiasedness will be reported in subsequent work.

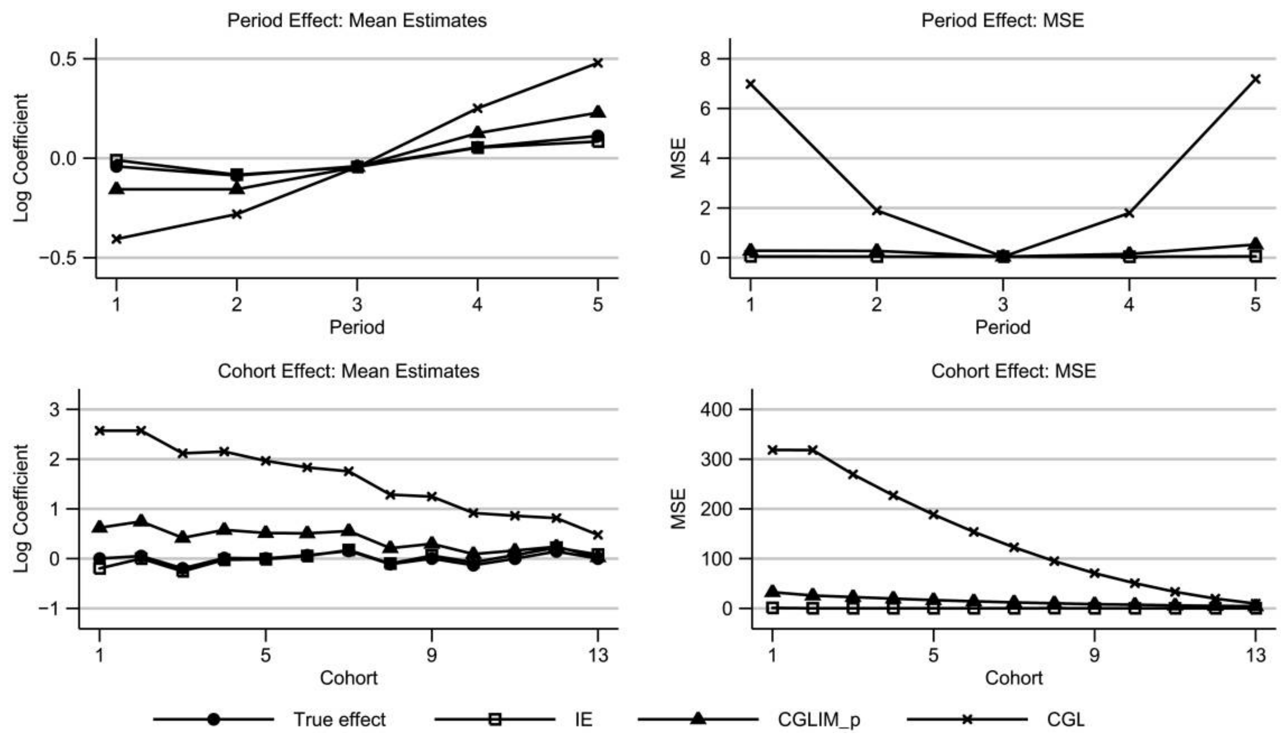

FIG. 4.—Simulation results of the IE and CGLIM estimators: period and cohort effects ($N = 1{,}000$)

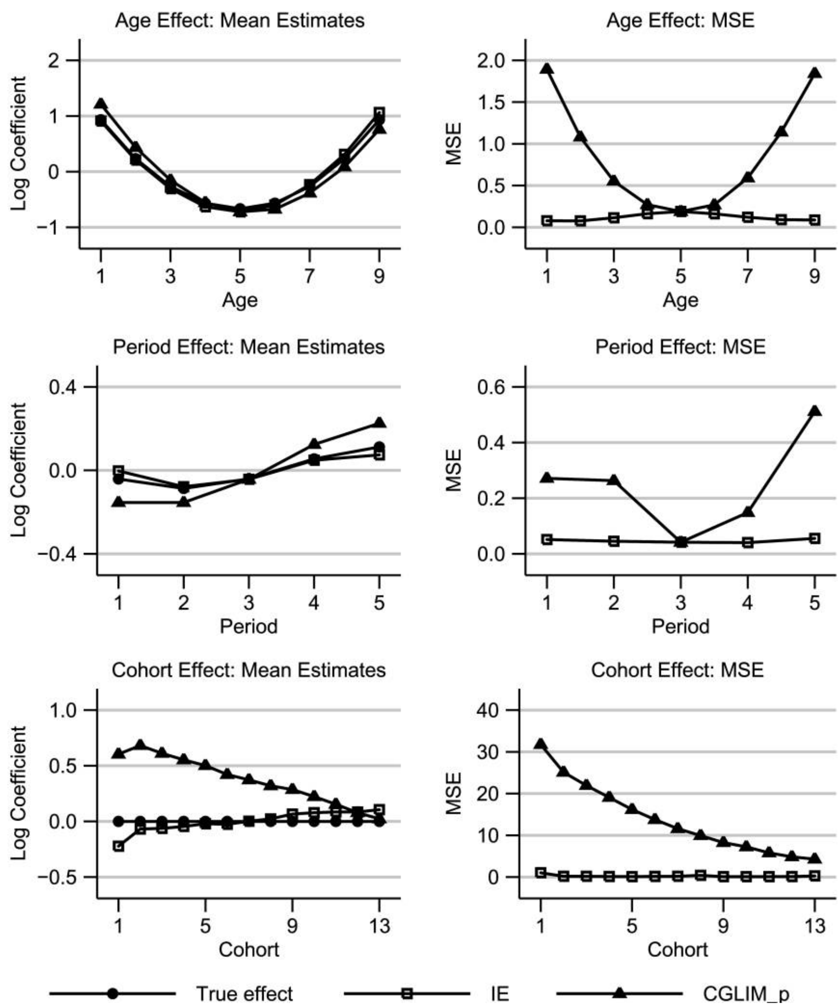

FIG. 5.—Simulation results of the IE and CGLIM estimators; true cohort effect = 0

INTERPRETATION AND USE OF THE IE

Empirical Applications, Interpretation of Model Parameters, and Hypothesis Testing

Given the desirable properties of the IE as a statistical estimator and its evident ability to produce valid estimates of age, period, and cohort effect coefficients, the question becomes one of how to interpret and use this estimator. The question of interpretation arises because the identifying constraint imposed by the IE on the unidentified APC accounting model parameter vector b—namely, projection onto the non-null (column) space of the design matrix X—appears to be a constraint purely of algebraic convenience, devoid of substantive meaning. By contrast, conventional equality-constrained estimators of APC accounting models often are motivated by substantive hypotheses derived from theory or prior studies that indicate that certain coefficients are, say, of the same magnitude and hence can be constrained to be equal. As noted earlier, however, if the equality constraints are not in fact valid, the constraints result in a nonestimable function of b and produce CGLIM estimates of APC effect coefficients that may be wildly off the mark.

To focus the discussion, in table 4 we present the effect coefficient estimates for the mortality data for U.S. females, 1960–99, given in table 1. The first two columns of table 4 report the coefficient estimates and standard errors produced by applying the IE to these data. The next two columns report the corresponding estimates for a CGLIM model, which (for reasons that will be made clear below) we label CGLIM 3, wherein the coefficients for the respective first categories of the age, period, and cohort groups are taken as the reference categories and have effects set to zero, and the identifying constraint is that the second birth cohort (C2), 1870–74, is constrained to have the same effect coefficient (zero) as the first cohort (C1), 1865–69. The note below table 4 gives the overall model fit statistics, which are the same for all models. Earlier we noted (in remark 4) that all just-identified models that incorporate effect coefficients for the full array of age, period, and cohort categories will fit the data equally well. This is evident in table 4. Thus, to reiterate the point made earlier, one cannot use fit statistics to discriminate among just-identified models. Rather, some other criterion must be employed. The criterion applied in this article is that the constrained vector must be estimable.

The numerical estimates in table 4 show that the constraint imposed in the CGLIM 3 model (in which the effect of the 1865 cohort is constrained to equal that of the 1870 cohort) is statistically valid (i.e., within sampling variability) for these data: the estimated effect for the 1865–69 cohort in the IE column is .939 (SE = .060) and the effect for the 1870–74 cohort is .937 (SE = .031). The difference between these two effect coefficient estimates, .002, is well within sampling error and is therefore effectively zero. In other words, this equality constraint produces an estimable function in a statistical sense (to be made precise below). The consequence is that the coefficient vectors for these two models are statistically identical up to a centering or normalizing transformation. This equivalence is demonstrated numerically in the last column of table 4, which gives the numerical values of the corresponding recentered/renormalized CGLIM 3 effect coefficients (i.e., the CGLIM 3 age, period, and cohort coefficients transformed by subtracting their respective group

TABLE 4

IE and CGLIM Estimates, U.S. Female Adult Mortality, 1960–99

	Intrinsic Estimator		CGLIM 3 (C1 = C2)		Centered CGLIM 1 (A1 = A2)	Centered CGLIM 1B (A15 = A16)	Centered CGLIM 2 (P1 = P2)	Centered CGLIM 3 (C1 = C2)
	Effect	SE	Effect	SE	Effect	Effect	Effect	Effect
Intercept	−4.571	.007	−5.803	1.058	−4.571	−4.571	−4.571	−4.571
Age:								
20–24	−2.032	.032	.000	.000	−1.749	.807	−2.410	−2.044
25–29	−1.995	.029	.039	.087	−1.749	.466	−2.322	−2.005
30–34	−1.819	.026	.217	.156	−1.611	.263	−2.096	−1.828
35–39	−1.555	.023	.482	.228	−1.385	.148	−1.782	−1.563
40–44	−1.245	.020	.794	.301	−1.112	.080	−1.421	−1.250
45–49	−.915	.017	1.125	.375	−.821	.031	−1.041	−.919
50–54	−.596	.015	1.446	.449	−.540	−.028	−.672	−.599
55–59	−.295	.013	1.749	.524	−.276	−.105	−.320	−.296
60–64	.026	.011	2.071	.598	.007	−.163	.051	.027
65–69	.349	.010	2.395	.672	.292	−.219	.424	.351
70–74	.701	.009	2.750	.747	.607	−.245	.827	.705
75–79	1.072	.008	3.122	.821	.940	−.253	1.249	1.078
80–84	1.484	.008	3.535	.896	1.313	−.220	1.710	1.491
85–89	1.896	.009	3.949	.970	1.688	−.187	2.173	1.904
90–94	2.273	.011	4.328	1.045	2.027	−.188	2.601	2.284
95+	2.652	.014	4.708	1.118	2.368	−.188	3.030	2.664
Period:								
1960–64	−.115	.009	.000	.000	−.247	−1.440	.062	−.109
1965–69	−.064	.008	.049	.075	−.159	−1.011	.062	−.060
1970–74	−.042	.007	.069	.149	−.099	−.610	.033	−.040
1975–79	−.081	.007	.029	.224	−.100	−.270	−.056	−.080
1980–84	−.034	.007	.074	.298	−.015	.156	−.059	−.034

1985–89	.044	.007	.150	.373	.101	.612	−.032	.041
1990–94	.094	.007	.199	.447	.189	1.041	−.032	.090
1995–99	.197	.008	.301	.522	.330	1.522	.021	.192
Cohort:								
1865	.939	.060	.000	.000	1.355	5.104	.385	.921
1870	.937	.031	.000	.000	1.315	4.723	.433	.921
1875	.900	.020	−.035	.094	1.241	4.308	.447	.886
1880	.834	.016	−.100	.165	1.137	3.863	.431	.821
1885	.748	.014	−.184	.239	1.013	3.398	.395	.737
1890	.653	.012	−.277	.313	.880	2.925	.351	.644
1895	.556	.011	−.373	.387	.746	2.450	.305	.548
1900	.449	.011	−.478	.461	.601	1.964	.248	.443
1905	.331	.011	−.595	.536	.445	1.467	.180	.327
1910	.243	.011	−.682	.610	.318	1.000	.142	.240
1915	.148	.012	−.775	.685	.186	.527	.098	.146
1920	.060	.013	−.861	.759	.060	.060	.060	.060
1925	−.018	.015	−.938	.834	−.056	−.397	.032	−.016
1930	−.124	.017	−1.042	.908	−.200	−.882	−.024	−.121
1935	−.249	.019	−1.165	.983	−.362	−1.385	−.098	−.244
1940	−.415	.022	−1.330	1.057	−.567	−1.930	−.214	−.409
1945	−.567	.024	−1.480	1.132	−.756	−2.460	−.315	−.559
1950	−.678	.026	−1.589	1.207	−.905	−2.950	−.376	−.668
1955	−.732	.030	−1.642	1.281	−.997	−3.383	−.380	−.721
1960	−.805	.035	−1.713	1.356	−1.108	−3.834	−.402	−.792
1965	−.918	.044	−1.825	1.431	−1.258	−4.325	−.464	−.903
1970	−1.063	.060	−1.969	1.506	−1.442	−4.850	−.560	−1.048
1975	−1.230	.103	−2.134	1.583	−1.646	−5.395	−.676	−1.213
t-ratio					.742	23.946	−4.822	−.076
p-value					.458	.000	.000	.939

NOTE.—Deviance = 18,903.2; df = 84; overdispersion = 224.3.

means so that the transformed coefficients sum to zero).[16] The resulting recentered CGLIM 3 and IE effect coefficients generally agree up to two or three digits. Again, this equivalence is due to the fact that the equality constraint imposed in the CGLIM 3 model is statistically valid for this data set. The equivalence is illustrated in figure 6, which shows graphs of the values of the age, period, and cohort effect coefficients for the CGLIM 3 and IE models. The patterns of the respective groups of effects are virtually identical.

Other possible identifying assumptions do not fare as well. To illustrate this, table 4 reports comparable effect coefficient estimates for two alternative CGLIM models: CGLIM 1, which identifies the model by constraining the effect coefficient for the 25–29 age group (A2) to be the same as that of the 20–24 age group (A1); and CGLIM 2, which achieves identification by constraining the effect coefficient for the 1965–69 time period (P2) to equal that for the 1960–64 period (P1). Figure 7 shows graphs of the respective sets of effect coefficients for these two CGLIM models, and, for comparison, shows the corresponding effect coefficients of CGLIM 3, which imposes the constraint C1 = C2, as reported above. For all three sets of effect coefficients, it can be seen that those produced by the CGLIM 1 and CGLIM 2 models bracket those produced by the CGLIM 3 model—that is, the two alternative CGLIM models yield effect coefficients that diverge substantially from those given by the CGLIM 3 model. The divergences are substantial for the age and cohort effects and dramatic for the period effects. The reason for this behavior is that the equality constraints used to produce the CGLIM 1 and CGLIM 2 models do not produce statistically estimable functions and corresponding coefficient estimates.

To proceed more systematically, a procedure is needed for assessing whether two estimated coefficient vectors are within sampling error of being equal. There are several ways to accomplish this. For instance, centered CGLIM effect coefficients could be compared *element by element* with the corresponding estimated IE effect coefficients in the first column

[16] CGLIM coefficients can be subjected to such a normalizing transformation through the following procedures: Under the original CGLIM normalization, let a_i be the estimated age effect for age i, let p_j be the estimated period effect for period j, and let c_k be the estimated cohort effect for cohort k. Let d be the estimated intercept. Transforming these estimates to a different normalization means that we want to find new coefficients a_i', p_j', c_k', and d' such that (1) the predicted value for each data point does not change and (2) the age, period, and cohort effects each sum to zero. The solution is to subtract the mean of the original age effects from each a_i to obtain a new age effect, a_i'—this guarantees that the new age effects sum to zero—but then to add the mean of the original age effects to the intercept, to guarantee that the predicted values do not change. The same process can be used to transform the period and cohort effects.

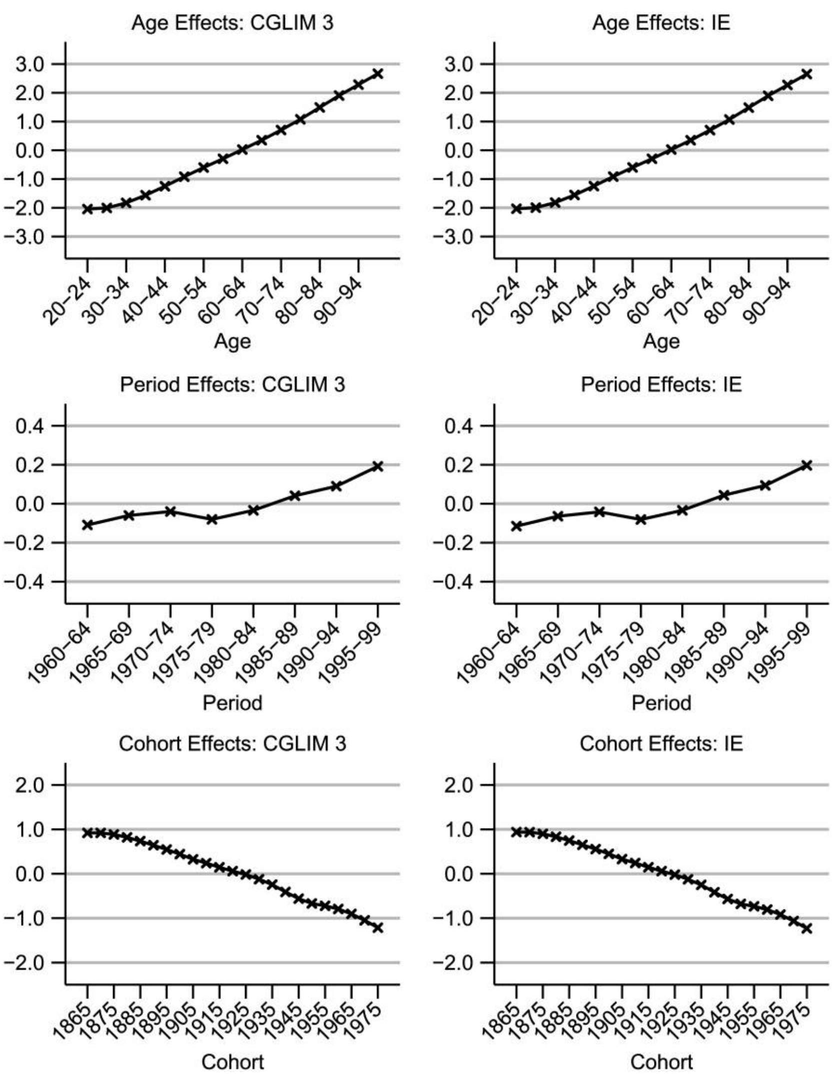

FIG. 6.—Comparison of graphs of age, period, and cohort effect coefficients: CGLIM 3 and IE models.

to determine whether the former are within, say, two standard errors of the latter. Using the standard errors of the IE coefficients, it can be seen that all of the centered effect coefficients for the CGLIM 3 model are within two standard errors of the IE coefficients. By comparison, this is not the case for the centered effect coefficients of the CGLIM 1 and CGLIM 2 models. This confirms the inference stated above, based on

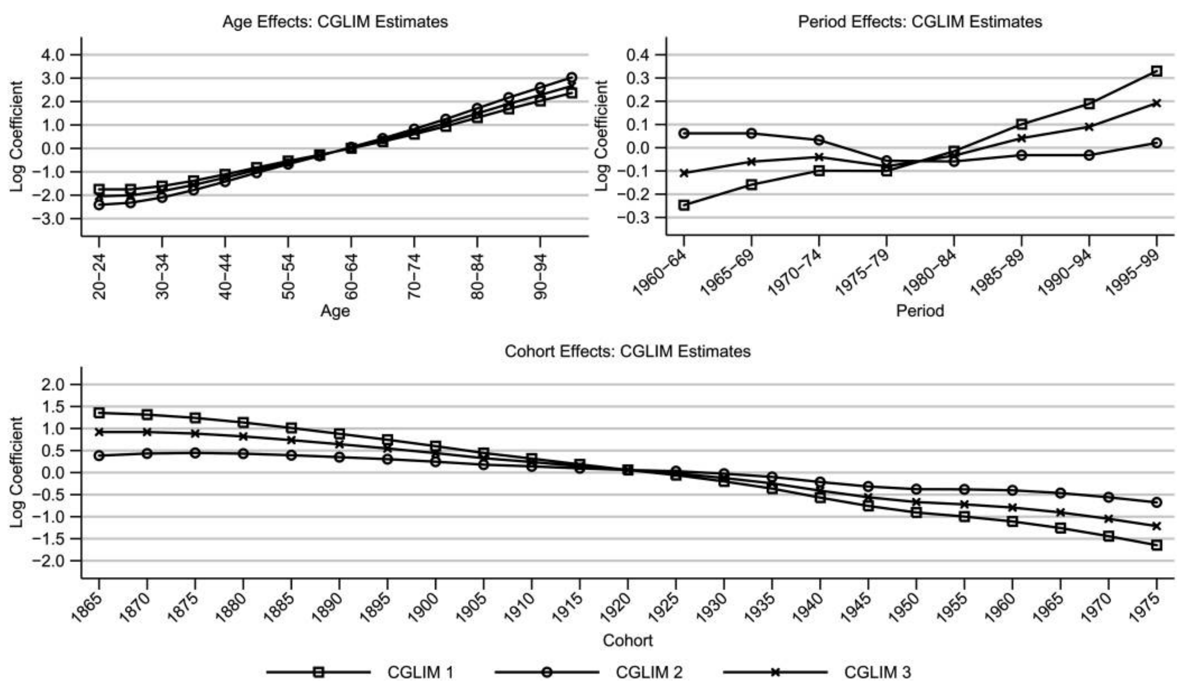

FIG. 7.—Age, period, and cohort effects estimated for three CGLIM models

visual inspection of figures 6 and 7—namely, that the CGLIM 3 model produces estimates of age, period, and cohort effects that are quite consistent with those estimated by the IE and thus is a statistically estimable function, whereas the CGLIM 1 and CGLIM 2 models are not.

An alternative procedure is to define a test statistic based on the *entire vector of coefficients*. To do so, we state the null hypothesis:

$$H_0 : E(\hat{b}B_0) = (B + sB_0)^T B_0 = s = 0.T. \tag{15}$$

In words, the null hypothesis is that the expected value of the product of the estimated and renormalized CGLIM vector $\hat{b}$ and the eigenvector B_0 that is fixed by the design matrix is zero. Because of the orthogonality of the vectors B and B_0, this, as equation (15) indicates, is equivalent to the hypothesis that the expected value of the scalar s is equal to zero. Using the geometric projection illustrated in figure 1, applying this test is equivalent to testing whether the estimated parameter vector under a given set of constraints (b_1 or b_2 or b_3) lies significantly far away from the estimable function, b_0, so that one can infer that its horizontal projection results from a nonzero s.

To specify a test for this null hypothesis, we build upon a well-known asymptotic distribution property of the maximum likelihood estimator (MLE) used to estimate $\hat{b}$: under broad regularity conditions, as sample size or the number of time periods of data increases, the MLE of $\hat{b}$ is consistent and asymptotically normally distributed (McCulloch and Searle 2001, p. 306). This property facilitates the definition of an asymptotic t-test for the null hypothesis (15) as

$$t = \frac{s - 0}{se(s)} = \frac{s}{se(s)}, \tag{16}$$

where $se(s)$ denotes the estimated (asymptotic) standard error of the scalar s.[17] Note that to obtain this test statistic, the numerator s can be computed, as indicated in equation (15), by calculating the product of the vectors $\hat{b}$ and B_0. Then the denominator can be computed by transforming the asymptotic variance-covariance matrix $\hat{\Sigma}^{-1}$ (i.e., the inverse of the Fisher information matrix) that is obtained in the process of estimating b by maximum likelihood to obtain $\hat{b}$. Because the vector B_0 is orthogonal to the design matrix (see eq. [9] above), the variance of the scalar s can be

[17] This asymptotic t-test is based on asymptotic normality properties of the maximum likelihood estimator, which assume large numbers of degrees of freedom—i.e., large APC matrices. A systematic study of the statistical power of the test for small to moderate-sized matrices is beyond the scope of the present article. Even without this, the test can provide useful information for assessing the plausibility of constraints based on theory or side information from other studies and for comparing alternative constraints, as illustrated here.

computed by imposition of the usual quadratic form transformation applied to obtain the variance of the restricted maximum-likelihood estimator from the MLE, namely $B_0^T\hat{\Sigma}^{-1}B_0$.[18] Taking the square root of this transformation yields an estimate of the standard deviation of $s = \hat{b}\, B_0 T$, which, when divided by the degrees of freedom of the model, $df = ap - (1 + [a - 1] + [p - 1] + [a + p - 2])$, produces an estimate of the standard error of s in the denominator of equation (16).

Applied to the three alternative CGLIM models for which the renormalized coefficients are given in table 4, we obtain the following t-ratios: 0.742 ($P = .458$) for CGLIM 1, -4.822 ($P < .001$) for CGLIM 2, and -0.076 ($P = .939$) for CGLIM 3. The results show that the CGLIM 1 and CGLIM 3 models estimates are not statistically different from those of the IE, but the CGLIM 2 model estimates are. Because the degrees of freedom of the APC models are small, it may be appropriate to use a P value larger than the conventional value of .05 for assessing the t-ratios. In this sense, the CGLIM 1 model, with a P value of .5, deviates from the IE model more than the CGLIM 3 model, which has a P value of .9, although both models clearly are acceptable by the conventional .05 criterion. From a substantive point of view, the equality constraint used by the CGLIM 1 model—that the effects be equal for ages 20–24 and 25–29—is consistent with the demographic evidence from life table studies that the mortality risk is generally low and roughly the same in the 20s. Reiterating the point made in remark 4 above, the results of applying the asymptotic t-test to the CGLIM 1, 2, and 3 models also show that estimators that incorporate significant components of the B_0 vector (and thus the design matrix) may give misleading results concerning the patterns of change of the estimated effect coefficients across the age, period, and cohort categories. This is illustrated graphically in figure 7. The figure shows that the patterns of changes of the coefficients for the CGLIM 1 and CGLIM 3 models, for which the foregoing asymptotic t-test is not significant, are quite close to those of the IE and thus lead to similar conclusions. On the other hand, those of the CGLIM 2 model, which has an asymptotic t-ratio that indicates a statistically significant departure

[18] Formally, the asymptotic normality property of the MLE yields $\hat{b} \overset{\text{asy}}{\approx} N(b, \Sigma^{-1})$; i.e., the sampling distribution of the estimated coefficient vector of a model that is identified by imposition of a theoretically motivated equality constraint on two or more coefficients is asymptotically multivariate normal with a mean (expected) parameter vector b and a variance-covariance matrix Σ^{-1}, where Σ^{-1} is estimated by $\hat{\Sigma}^{-1}$, the inverse of the Fisher information matrix. By properties of the MLE (see, e.g., McCulloch and Searle 2001, p. 309), a linear transformation of $\hat{b}$ using B_0 then yields $\hat{b}B_0 \overset{\text{asy}}{\approx} N(0, B_0^T\Sigma^{-1}B_0)T$, where $\hat{b}B_0 = sT$. From this, the asymptotic standard error of s can be computed as the square root of $B_0^T\hat{\Sigma}B_0{}^{-1}$ divided by degrees of freedom, as indicated in the text.

from the IE coefficients, show patterns of change that depart substantially from those of the IE.

To understand how not all constraints based on age effect coefficients produce acceptable *t*-ratios, consider an alternative constraint on the coefficients of the two oldest age groups, a constraint that might be motivated by recent studies of the deceleration of mortality at the oldest ages (see, e.g., Vaupel 1997). To see how this different constraint compares to the current CGLIM 1 model, we estimated a CGLIM 1B model constraining the last two age groups to be equal in coefficients (A15 = A16). The estimates result in a large *t*-ratio of 23.9 ($P < .001$), suggesting that this equality constraint does not produce an estimable function.

In brief, the foregoing test for statistical estimability leads to new avenues for using and interpreting the IE as applied to a particular data set. On the one hand, an analyst can apply the IE in an *exploratory data analysis* approach, in which the objective is to ascertain good estimates of the patterns of age, period, and cohort effects in a table or set of tables of demographic rates. In such an exercise, the analyst does not approach the data with strong prior notions about particular patterns of effects that should be evident, but rather seeks to let the patterns emerge from application of the IE, taking advantage of the fact that the IE is an estimable function and thus has desirable statistical properties.

On the other hand, using the definition of statistical estimability and the test described above, an analyst can use the IE in a *confirmatory data analysis* approach,[19] in which a vector of effect coefficients estimated from application of the IE is used as a benchmark to assess whether a corresponding vector of coefficients estimated from the imposition of one or more theoretically or substantively motivated constraints to achieve model identification is acceptable. In this type of exercise, an analyst approaches a table or tables of demographic rates with a definite hypothesis or set of hypotheses about the underlying age, period, and cohort effects that generated the data. The analyst can use the vector of effect coefficients estimated by the IE to assess the empirical plausibility of the hypotheses. In this way, the definition of statistical estimability and the test described above directly address a criticism that often has been lodged against general-purpose methods of APC analysis, namely, that they provide no avenue for testing specific substantive hypotheses and thus are mere devices of algebraic convenience that may be misleading.

[19] The distinction between "exploratory" and "confirmatory" data analysis dates back in statistics at least to the classic work of Mosteller and Tukey (1977).

Computational Tools

Interested users may desire a computer program for conveniently applying the IE to their data. To date, programs for estimating the IE have been written as add-on files to two commercially available software packages, S-Plus (Venables and Ripley 2000) and Stata (Rabe-Hesketh and Everitt 2004).

The S-Plus program can be obtained by writing to Wenjiang Fu (fuw@epi.msu.edu) and requesting a copy. An add-on file for calculating the IE in Stata may be obtained by typing *ssc install apc* on the Stata command line on any computer connected to the Internet. It can also be downloaded from the Statistical Software Components archive at http://ideas.repec.org/s/boc/bocode.html. The program uses much the same syntax as Stata's *glm* command for generalized linear models. For example, a user whose data set contains a dependent variable *y*, an exposure variable *x*, an age variable *a*, and a period variable *t* can fit a Poisson model with age, period, and cohort effects by typing

```
apc_ie y, exposure(x) family(poisson) link(log) age(a) period(t).
```

Stata will then display coefficient estimates, standard errors, confidence intervals, the log-likelihood value, and a variety of other statistics. The program is documented more fully in a help file that can be read by typing *help apc_ie* in Stata after installation. The package also includes the command *apc_cglim* for calculating CGLIM estimators that impose equality constraints on pairs of coefficients; for details, type *help apc_cglim* in Stata after installation.

CONCLUSIONS

The problem of obtaining reliable estimates of the patterns of change across age groups, time periods, and cohorts has long provided an intriguing challenge in many contexts in the social sciences, demography, and epidemiology. To address this challenge, this article has examined a number of properties and the performance of the intrinsic estimator for the APC accounting model in the context of age-by-time-period tables of rates. As noted above, Glenn (2005, p. 20) has stated several strong criteria for judging the acceptability and utility of a general-purpose method of APC analysis. The IE appears to satisfy these criteria. As shown here, the IE has passed both empirical and simulation tests of validity and can be used to test theoretically motivated hypotheses and to incorporate and test side information from other studies. The IE therefore may provide a useful tool for the accumulation of scientific knowledge about the distinct effects of age, period, and cohort categories in social research.

Indeed, since the APC underidentification problem is an instance of a larger family of such structural issues, the potential range of application for the IE may be even larger. Structural underidentification problems occur when a conceptualization of the effects of structural arrangements leads to an exact linear dependency among the effects. An example is the classical problem in mobility analysis of distinguishing the effects of socioeconomic mobility or distance moved on an outcome variable from the effects of origin and destination statuses (see, e.g, Duncan 1966). Another example pertains to the estimation of the effects of years of labor force experience separate from the effects of current age and age at labor force entry. These and similar problems of structural underidentification occur frequently in sociology and related disciplines. The IE and/or algebraic approaches similar to the IE may prove useful in applications to such problems.

Is the IE then a complete solution to the structural identification problem in APC and similar models? No. Structural identification problems are just that—points of underidentification of parameters resulting from 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 these problems with respect to statistical properties. Because of its desirable properties as a statistical estimator, including its ability, demonstrated above, to produce good estimates of the underlying patterns of change in the age, period, and cohort coefficients in the accounting model with a small number of time periods of data, the IE adds a potentially useful method to the toolkit available for these analyses. Does this mean that researchers should naively apply this method to tables of rates and expect to obtain meaningful results? Again, no. Every statistical model has its limits and will break down under some conditions. 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 an awareness of its many pitfalls. Replications, in particular, are called for when additional data become available.

Last but equally important, the limitations of the IE in theoretical and empirical studies are related to the limitations of the APC accounting model (Smith 2004). This model is based on the assumption of addivity of age, period, and cohort effects that not only incurs the identification problem but also may be a poor approximation of how social change occurs. Additional new models and methods are needed for testing other theories of social change. The recent development of hierarchical APC models is useful for studying the contextual effects of cohort membership and period events on a wide range of social processes (Yang 2006, 2008;

Yang and Land 2006). Other conceptually appealing models, such as the cohort inversion model and the continuously accumulating/evolving cohort effects model (Hobcraft et al. 1982), should also be explored and more fully developed in future research.

REFERENCES

Agresti, A. 1996. *An Introduction to Categorical Data Analysis*. Wiley Series in Probability and Statistics. New York: John Wiley & Sons.

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

Alwin, D. F., and R. J. McCammon. 1999. "Aging versus Cohort Interpretations of Intercohort Differences in GSS Vocabulary Scores." *American Sociological Review* 64:272–86.

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

Duncan, O. D. 1966. "Methodological Issues in the Analysis of Social Mobility." Pp. 51–97 in *Social Structure and Mobility in Economic Development*, edited by N. J. Smelser and S. M. Lipset. Chicago: Aldine.

Fienberg, S. E., and W. 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." Pp. 45–88 in *Cohort Analysis in Social Research*, edited by W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.

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* 29:263–78.

———. 2008. "A Smoothing Cohort Model in Age-Period-Cohort Analysis with Applications to Homicide Arrest Rates and Lung Cancer Mortality Rates." *Sociological Methods and Research* 36:327–61.

Fu, W. J., and P. Hall. 2006. "Asymptotic Properties of Estimators in Age-Period-Cohort Analysis." *Statistics and Probability Letters* 76:1925–29.

Fu, W. J., P. Hall, and T. Rohan. 2004. "Age-Period-Cohort Analysis: Structure of Estimators, Estimability, Sensitivity, and Asymptotics." Technical report. Michigan State University, Department of Epidemiology.

Glenn, N. D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period, and Cohort Effects." *American Sociological Review* 41:900–904.

———. 1994. "Television Watching, Newspaper Reading, and Cohort Differences in Verbal Ability." *Sociology of Education* 67:216–30.

———. 1999. "Further Discussion of the Evidence for an Intercohort Decline in Education-Adjusted Vocabulary." *American Sociological Review* 64:267–71.

———. 2005. *Cohort Analysis*. Thousand Oaks, Calif.: 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* 39:311–24.

———. 1991. "Understanding the Effects of Age, Period, and Cohort on Incidence and Mortality Rates." *Annual Review of Public Health* 12:425–57.

———. 1992. "Analysing the Temporal Effects of Age, Period, and Cohort." *Statistical Methods in Medical Research* 1:317–37.

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: 811–30.

Kupper, L. L., J. M. Janis, I. A. Salama, C. N. Yoshizawa, and B. G. Greenberg. 1983. "Age-Period-Cohort Analysis: An Illustration of the Problems in Assessing Interaction in One Observation per Cell Data." *Communications in Statistics—Theory and Method* 12:2779–2807.

Marion, J. B., and S. T. Thornton. 1995. *Classical Dynamics of Particles and Systems*, 4th ed. Fort Worth, Tex.: Harcourt Brace College Publishers.

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. 2189–94 in *International Encyclopedia of the Social & Behavioral Sciences*. New York: Elsevier.

McCullagh, P., and J. A. Nelder. 1989. *Generalized Linear Models*, 2d ed. New York: Chapman and Hall.

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

Mosteller, F., and J. W. Tukey. 1977. *Exploratory Data Analysis*. Reading, Mass.: Addison-Wesley.

O'Brien, R. M. 2000. "Age Period Cohort Characteristic Models." *Social Science Research* 29:123–39.

Osmond, C., and M. J. Gardner. 1982. "Age, Period, and Cohort Models Applied to Cancer Mortality Rates." *Statistics in Medicine* 1:245–59.

Pullum, T. W. 1978. "Parametrizing Age, Period, and Cohort Effects, with Applications to U.S. Delinquency Rates, 1964–1973." *Sociological Methodology* 9:116–40.

———. 1980. "Separating Age, Period, and Cohort Effects in White U.S. Fertility, 1920–1970." *Social Science Research* 9:25–244.

Rabe-Hesketh, S., and B. Everitt. 2004. *A Handbook of Statistical Analysis Using Stata*, 3d ed. New York: Chapman & Hall.

Robertson, C., and P. Boyle. 1998. "Age-Period-Cohort Analysis of Chronic Disease Rates I: Modeling Approach." *Statistics in Medicine* 17:1305–23.

Robertson, C., S. Gandini, and P. Boyle. 1999. "Age-Period-Cohort Models: A Comparative Study of Available Methodologies." *Journal of Clinical Epidemiology* 52:569–83.

Rodgers, W. L. 1982*a*. "Estimable Functions of Age, Period, and Cohort Effects." *American Sociological Review* 47:774–87.

———. 1982*b*. "Reply to Comments by Smith, Mason, and Fienberg." *American Sociological Review* 47:793–96.

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

Searle, S. R. 1971. *Linear Models*. New York: John Wiley & Sons.

Sen, A., and M. Srivastava. 1990. *Regression Analysis: Theory, Methods, and Applications*. New York: Springer-Verlag.

Smith, H. L. 2004. "Response: Cohort Analysis Redux." Pp. 111–19 in *Sociological Methodology* 34:111–19.

Smith, H. L., W. M. Mason, and S. E. Fienberg. 1982. "More Chimeras of the Age-Period-Cohort Accounting Framework: Comment on Rodgers." *American Sociological Review* 47:787–93.

Tarone, R., and K. C. Chu. 1992. "Implications of Birth Cohort Patterns in Interpreting

Trends in Breast Cancer Rates." *Journal of the National Cancer Institute* 84: 1402–1410.

Vaupel, J. W. 1997. "Trajectories of Mortality in Advanced Ages." Pp. 17–37 in *Between Zeus and the Salmon*, edited by K. W. Wachter and C. E. Finch. Washington, D.C.: National Academy Press.

Venables, W. N., and B. D. Ripley. 2000. *Modern Applied Statistics with S-PLUS*, 3d ed. New York: Springer.

Wilmoth, J. R. 1990. "Variation in Vital Rates by Age, Period, and Cohort." *Sociological Methodology* 20:295–335.

Wilson, J. A., and W. R. Gove. 1999*a*. "The Intercohort Decline in Verbal Ability: Does It Exist?" *American Sociological Review* 64:253–66.

———. 1999*b*. "The Age-Period-Cohort Conundrum and Verbal Ability: Empirical Relationships and Their Interpretation. (Reply to Glenn and to Alwin and McCammon)" *American Sociological Review* 64:287–302.

Yang, Y. 2006. "Bayesian Inference for Hierarchical Age-Period-Cohort Models of Repeated Cross-section Data." *Sociological Methodology* 36:39–74.

Yang, Y. 2008. "Social Inequalities in Happiness in the United States, 1972–2004: An Age-Period-Cohort Analysis." *American Sociological Review* 73:204–226.

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.

Yang, Y., and K. 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.