Sociological Methods & Research
Volume 36 Number 3
February 2008 402-428

10.1177/0049124106290392
http://smr.sagepub.com
hosted at
http://online.sagepub.com

A Mixed Model Estimation of Age, Period, and Cohort Effects

Robert M. O'Brien
University of Oregon, Eugene
Kenneth Hudson
University of South Alabama, Mobile
Jean Stockard
University of Oregon, Eugene

For more than 30 years, sociologists and demographers have struggled to come to terms with the age, period, cohort conundrum: Given the linear dependency between age groups, periods, and cohorts, how can these effects be estimated separately? This article offers a partial solution to this problem. The authors treat cohort effects as random effects and age and period effects as fixed effects in a mixed model. Using this approach, they can (1) assess the amount of variance in the dependent variable that is associated with cohorts while controlling for the age and period dummy variables, (2) model the dependencies that result from the age-period-specific rates for a single cohort being observed multiple times, and (3) assess how much of the variance in observations that is associated with cohorts is explained by differences in the characteristics of cohorts. The authors use empirical data to see how their results compare with other analyses in the literature.

Keywords: *age; period; cohort; mixed models*

Forty years ago, Norman Ryder (1965) eloquently delineated the distinctive impact of birth cohorts on social change and continuity. As he noted,

> Society persists despite the mortality of its individual members, through processes of demographic metabolism and particularly the annual infusion of birth cohorts. These may pose a threat to stability, but they also provide the opportunity for societal transformations. Each birth cohort acquires coherence and continuity from the distinctive development of its constituents and from its own persistent macroanalytic features. (p. 843)

Since Ryder's (1965) statement, researchers have shown that members of particular birth cohorts experience historical and demographic events differently than members of other cohorts and that they carry these experiences throughout their life course (Elder 1974, 1996; Elder and Caspi 1990). Sociologists have documented cohort influences in a wide variety of areas, including antiblack prejudice (Firebaugh and Davis 1988), opinions on democracy and Nazism (Weil 1987), parental values (Alwin 1990), political orientation and voting (Firebaugh and Chen 1995; Alwin and Krosnick 1991), sex role attitudes (Mason and Lu 1988), intellectual skills (Alwin 1991; Wilson and Gove 1999), criminal behavior (O'Brien 1989; O'Brien, Stockard, and Isaacson 1999; Smith 1986), and suicide (Stockard and O'Brien 2002a, 2002b).

Despite this attention to the presence of cohort effects, isolating and estimating the distinctive impact of cohorts has proven difficult because of the linear dependency between age, period, and cohort. Although they were not the first to note this dependency (see Greenberg, Wright, and Sheps 1950; Hall 1971; Sacher 1960), Mason et al.'s (1973) exposition of the age, period, cohort conundrum continues to challenge sociologists and demographers. They noted that if age, period, and cohort were all causally related to an outcome variable, modeling the outcome variable's dependence on only two of the three variables would result in spurious estimates. Using all three of these variables, however, is problematic because age, period, and birth cohort are linearly dependent.[1] Knowing an observation's value on any two of these variables tells us exactly its value on the third. For example, if we have yearly observations on homicide rates by age, knowing the age and year (period) for a specific rate will tell us the cohort to which the rate belongs. Thus, including age, period, and cohort as independent variables in a regression analysis produces perfect multicollinearity in the matrix of explanatory variables. Because this matrix is not of full rank, it is not possible to estimate unique values for the ordinary least squares (OLS) regression coefficients.[2]

There is an extensive literature on the attempts to estimate age, period, and cohort effects. Three traditional strategies (Mason et al. 1973) include (1) transforming at least one of these variables so that its relationship to the others is nonlinear; (2) assuming that two or more age groups, periods, or cohorts have the same effect on the dependent variable; or (3) assuming that the effects of membership in particular cohorts are variable through time rather than fixed. The most widely used of these strategies assumes that two or more age groups, periods, or cohorts have the same effect on the dependent variable.[3]

We focus on a different type of solution that does not assume that one or more coefficients for age groups or periods or cohorts have the same value or that specific types of nonlinear relationships exist across age groups, periods, or cohorts. We measure the age and period effects using fixed-effect dummy variables and assess the effects of cohorts using random-effect dummy variables to represent the cohorts. The use of a mixed model in which the age and period dummy variables are fixed effects and the cohorts are treated as random effects is a major innovation in our approach. It allows us to assess whether the variance associated with cohort membership is statistically significant after controlling for the effects of age groups and periods, and this assessment does not place restrictions on the functional form of the relationship between cohort membership and the dependent variable. We then introduce two theoretically selected "cohort characteristics" as fixed effects to examine how well they account for the differences among cohorts.[4]

The use of cohort characteristics, thought to be associated with the effects of cohorts on the dependent variable, has been used extensively to help solve the age–period–cohort (APC) conundrum. Referring to this class of solutions, Rodgers (1982a) states that "a solution to the dilemma [the APC conundrum] lies in the specification and measurement of the theoretical variables for which age, period, and cohort are indirect indicators" (p. 774). Then this variable (or variables) may be used in the analysis rather than age or period or cohort, eliminating the linear dependency in the set of independent variables. Kahn and Mason (1987) use this approach when they employ the relative size of the cohort as a cohort characteristic posited to be responsible for cohort membership's association with the dependent variable. Farkas (1977) uses the analogous solution but does so by finding a direct indicator for the period effects. He conceptualizes period effects as associated with the business cycle and uses the unemployment rate to model this effect in his study.

In this article, we focus our examination on cohort effects and use cohort characteristics to help explain cohort effects (but the same method can be used to estimate period or age group effects and can be used with period or age group characteristics).[5] A problem with the cohort (or period or age group) characteristic solution is that it only captures cohort effects to the extent that the cohort characteristic captures the entire effect of cohort membership. To more completely capture the cohort effect, we propose a mixed model in which age and period are estimated as fixed effects and cohort membership is estimated as a random effect. The age and period dummy variables do not assume a particular functional form for the relationship between age and the dependent variable or for period and the dependent variable. The same is true for the random effects of cohort membership. This allows us to estimate the variance *uniquely* associated with cohorts while controlling for the effects of age and period, without assuming that we capture this variance with the cohort characteristics. This partitioning of variance is limited by the inevitable problem of partitioning the variance among independent variables when they are correlated and the causal ordering of the independent variables is not specified. We then demonstrate how cohort characteristics can be used to explain the variance that is associated with cohorts and is independent of age and period. Since the typical cohort in an APC model contributes multiple observations to the analysis, we also model the correlation of the residuals within cohorts.

The Traditional APC and APCC Models

Given our focus on the problems inherent in the traditional APC model and our later use of the age–period–cohort-characteristic (APCC) model, we briefly describe these models. The traditional layout for APC data with a rectangular period by age data matrix is depicted in Figure 1a. Each row represents a different period and each column a different age group. Each cell represents an age-period-specific statistic, such as a rate or mean. The diagonals running downward from left to right represent cohorts. For example, in Figure 1a, the cohort born between 1940 and 1944 contributes its first age-period-specific observation when it is 15 to 19 years old in 1960. It is next observed when it is 20 to 24 years old in 1965, and it is last observed when it is 45 to 49 years old in 1990. The cohort born between 1970 and 1974 contributes its first age-period-specific rate in 1990 when it is 15 to 19 years old, its second in 1995 when it is 20 to 24 years old, and its last in 2000. The two examples in Figure 1 represent schematically the two data sets we analyze later in the article.

This traditional APC model is represented by equation (1):

$$Y_{ij} = \mu + \alpha_i + \pi_j + \chi_k + \varepsilon_{ij}, \tag{1}$$

where the effect of the ith age group is given by α_i, the effect of the jth period by π_j, the effect of the kth cohort is χ_k, μ is the intercept, ε_{ij} is the random disturbance, and $i = 1, \ldots, A-1$; $j = 1, \ldots, P-1$; $k = 1, \ldots, (A+P)-1$. A is the number of age groups, and P is the number of periods.[6] The α_i, π_j, and χ_k are the effects of age, period, and cohort dummy variables, and one category of each serves as a reference category. Y_{ij} is

Figure 1
Schematic Period-by-Age Tables for the Homicide Offense and Suicide Data

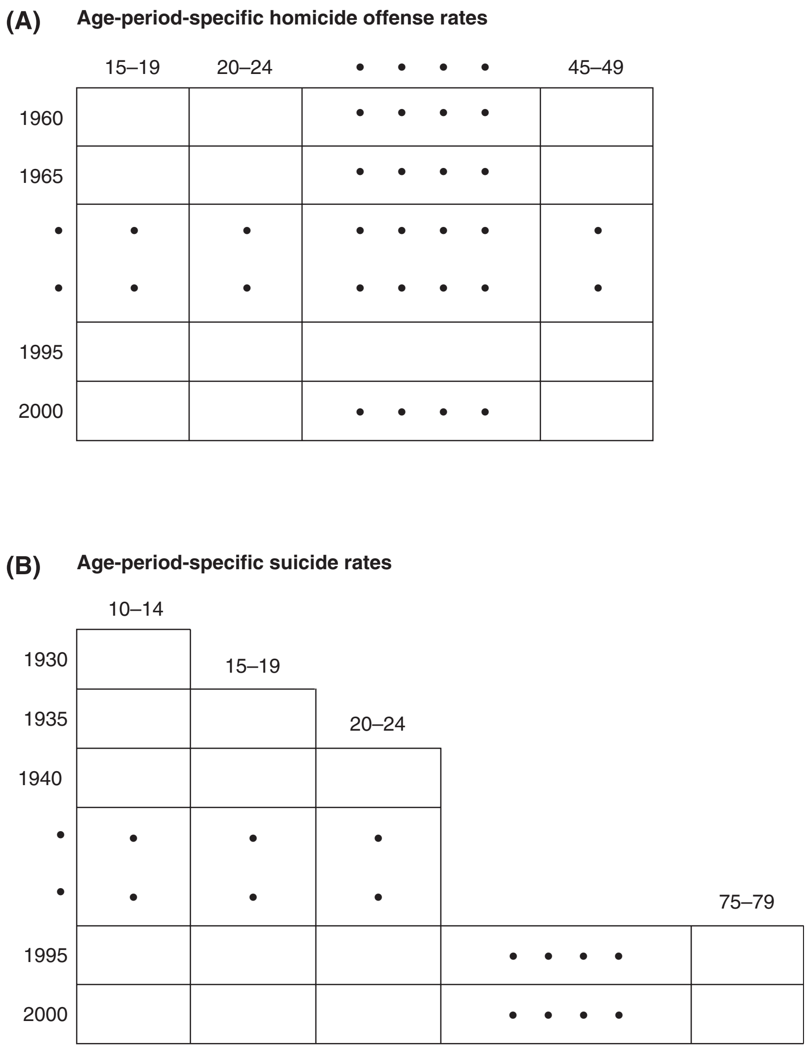

the age-period-specific statistic for the ith age group and jth period, and ε_{ij} is the random disturbance for the ijth observation. When OLS is used to estimate this equation, the ε_{ij} are typically assumed to be independently, identically, and normally distributed. This model of age, period, and cohort effects suffers from a linear dependency among the variables representing age, period, and cohort since knowing the values of any two of these variables allows one to know the value of the third. This linear dependency means that we cannot estimate the $A - 1$ age effects, the $P - 1$ period effects, and the $A + P - 1$ cohort effects in a single OLS model.[7]

Equation (1) also can be used to denote the traditional APCC model, where the effect of the α_i and π_j represents the age and period effects, μ represents the intercept, and ε_{ij} is the random disturbance as in (1). In the APCC model, however, χ_k represents the value of the cohort characteristic for the kth cohort (here we represent only a single cohort characteristic). This formulation eliminates the linear dependency inherent in the traditional APC model since the cohorts' values on the cohort characteristic are unlikely to be perfectly linearly related to the age and period. Using cohort characteristics to adequately measure cohort effects, however, is dependent on choosing appropriate cohort characteristics that capture most of the effects of the cohorts. Without capturing the cohort effects, the coefficients associated with age and period will almost certainly be biased. If the cohort effect is not fully captured, the age and period dummies will be confounded by cohort effects due to the missing cohort characteristics. Unless the missing cohort characteristics are uncorrelated with the included ones, the coefficients for the included ones will be biased due to omitted variable bias. The method that we propose helps gauge the amount of variance that is unique to cohorts that is captured by the cohort characteristics.

The traditional APCC model represents a conservative model for assessing the effects of cohorts. By using dummy variables to control for age effects and period effects, the model captures the relationship between age and the dependent variable and between period and the dependent variable, whether that relationship is curvilinear or not. The price for this control is in the loss of degrees of freedom and parsimony of representation. The fixed-effect dummy variables for age and period in the APCC model control for the effect of all variables related to age and all variables related to period (to the extent that these effects are constant across periods and ages, respectively). In addition, O'Brien et al. (1999) show that, because cohort is linearly dependent on the age and period dummy variables, the APCC model controls for any linear relationship between cohort year of birth and the dependent variable. In effect, this controls for any linear relationship between the cohort characteristic and the "age" of the cohort. Thus, the effects of cohort characteristics estimated in this model are independent of the effects of age group, period, and cohort year of birth. While the control for cohort year of birth may seem extreme, this sort of control is analogous to the control for the linear effects of time in a time-series model.

Both the APC model and the APCC model treat the effects of age, period, and cohort (APC) or cohort characteristics (APCC) as fixed effects. The first model is not identified if we include dummy variables for the cohorts (excluding one for a reference category) or if we use a linear variable to represent cohort year of birth in the model. The second model, which uses a cohort characteristic, is identified since the cohort characteristic is not likely to be a linear function of age and period.

Equation (2a) represents an age-period model, where Y_{ij} is the age-period-specific observation for the ith age group and jth period, μ is the intercept, α_i is the effect of the ith age group, π_j is the effect of the jth period, and ε_{ij} is the random disturbance associated with the ijth observation. We can use the residuals from this model to obtain an OLS analogue to the mixed model random effect for cohorts that we introduce in the next section. Equation (2a) is the age-period model, and we can use the estimated coefficients from this model (2a) to generate the observed residuals, e_{ij}, in (2b).

$$Y_{ij} = \mu + \alpha_i + \pi_j + \varepsilon_{ij}, \tag{2}$$

$$Y_{ij} - \left(\mu + \hat{\alpha}_i + \hat{\pi}_j\right) = e_{ij}. \tag{3}$$

We could then examine the pattern of these residuals to see if they indicate an effect of cohort membership (after controlling for the effects of age and period).

We would proceed in the following manner: (1) run a regression analysis with the age-period-specific statistic as the dependent variable and the age and period dummy variables as the independent variables, (2) construct an age-by-period table of residuals from the analysis, and (3) compute the mean residual for each of the cohorts. As an example, in Figure 1a, these would be the mean of the single residual for the cohort that was 15 to 19 years old in 2000, the mean of the two residuals for the cohort that was 15 to 19 years old in 1995 (they are observed a second time when they are 20 to 24 years old in 2000), and so on. These mean residuals indicate how much the cohort is above or below its expected value, controlling for age and period. We label this the method of cohort residuals. We could, of course, control for age and cohort and examine the mean residuals for each period (the method of period residuals), or we could enter dummy variables for cohort and period and examine the mean residuals for age groups (the method of age group residuals).

To the extent that there are cohort effects (that are statistically independent of age and period effects), we would expect these mean cohort residuals to vary systematically between cohorts. For instance, in cohorts that have a higher than expected propensity toward lethal violence, we would expect that their residuals would tend to be positive. For those with a lower propensity toward lethal violence, we would expect their residuals to be negative. This conceptualization suggests that we examine the mean residuals from the diagonals associated with different cohorts. Later we compare the results of the use of this method to the preferred method proposed in the next section.

This logic suggests that in the OLS framework, we can obtain information about the effects of age, period, and cohort, even though we cannot enter each of these variables as regressors in a traditional OLS model. We can do so by examining the residuals from a model in which we have entered the other two "complete" sets of dummy variables. Although age, period, and cohort are linearly related, the residuals associated with age, period, or cohort are not linearly related to the other two factors. In the next section, we use a mixed model to help us separate the residual variance into two parts: one associated with cohort membership and the other with the residual variance after entering the random-effect variables for cohorts.[8]

Age, Period, Cohort Mixed Model Estimation

By using a mixed model to estimate age, period, and cohort effects, we can (1) assess the amount of variance in the dependent variable that is associated with cohorts while controlling for the age and period dummy variables, (2) model the dependencies that result from the age-period-specific rates for a single cohort being observed multiple times, and (3) assess how much of the variance in observations associated with cohorts is explained by differences in cohort characteristics. We label this model the age, period, cohort mixed model (APCMM).

Using the notation introduced for equation (1), we can write this mixed model (without cohort characteristics) as

$$Y_{ij} = \mu + \alpha_i + \pi_j + \nu_k + \varepsilon_{ij}, \tag{4}$$

where μ is the intercept, α_i represents the fixed effects of the dummy variables for the age groups, π_j represents the fixed effects of the dummy variables for the periods, and ε_{ij} is the random disturbance. The new term, ν_k, represents the random disturbance characterizing the kth cohort. We can view these random disturbances as collections of factors not in the regression model that are specific to particular cohorts. The ν_k and ε_{ij} are assumed to be uncorrelated random variables with zero means and uncorrelated with the fixed-effect independent variables in the analysis. The restricted maximum likelihood technique used in our analysis assumes that these random variables have a multivariate normal distribution. By modeling the cohort effects as random effects, the APCMM provides an estimate of the variation between cohorts while controlling for age and period dummy variables. This measure does not impose a specific functional form on the relationship between cohort membership and the dependent variable and, in this sense, assesses the variance in the dependent variable (age-period-specific rates) that is associated with cohort membership, controlling for age and period.

We can add one or more cohort characteristics to equation (3) to help explain the variation among cohorts and to control the age and period effects for the cohort effects (as represented in the cohort characteristics):

$$Y_{ij} = \mu + \alpha_i + \pi_j + \chi_k + \nu_k + \varepsilon_{ij}. \tag{5}$$

Here, χ_k provides an estimate of the effects of a cohort characteristic, and ν_k provides an estimate of the systematic variance between cohorts that remains unaccounted for after controlling for the age and period effects and differences between cohorts in their values on the cohort characteristic.

The mixed model estimation procedure that we use allows us to model the autocorrelation of observations within cohorts that typically is not modeled in APCC models (see O'Brien et al. 1999). It is reasonable to expect that observations within a given cohort exhibit a degree of dependency. In addition to a cohort exhibiting systematically higher rates of homicide or suicide than other cohorts, the observations within a cohort may be correlated. We model the former by treating the cohorts as a random variable and the latter by including an autoregressive parameter estimate in the model.[9]

The estimation of separate effects for age, period, and cohort is possible in the mixed model because we treat the age and period effects as fixed and the cohort effects as random. The mixed model estimation of the fixed and random effects uses a process that is distinctly different from that employed in OLS estimation. The matrix formulation of the mixed model may be written as

$$\mathbf{Y} = \mathbf{X}\beta + \mathbf{Zu} + \mathbf{e}, \tag{6}$$

where **Y** is a column vector of observations on the dependent variable (the age-period-specific rates), and **X** is the design matrix for the fixed effects. For the APCCMM, **X** consists of one column for the intercept, $P - 1$ columns for the period dummy variables, $A - 1$ columns for the age group dummy variables, and a column for each of the cohort characteristics in the model. There is one row for each of the age-period-specific observations. The first column consists of all 1s, the columns for the dummy variables contain a 1 in the row of the column if the age-period-specific observation in the row corresponds to an observation in the period or age group represented by the column and a 0 if it does not, and each column representing a cohort characteristic contains the value of the cohort characteristic for the case represented by the corresponding row. β is an $(A + P + CC - 1)$ column vector of the corresponding fixed effect parameters (where CC represents the number of cohort characteristics).[10] **Z** is the design matrix for the random effects. It consists of a column for each of the cohorts and a row corresponding to each of the age-period-specific observations. There is a 1 in the row of the column if the age-period-specific observation in the row is an observation that is in the cohort represented by the column and a 0 if it is not. The vector of random effects, **u**, is a column vector of the corresponding random effects with one entry per cohort, and **e** is a column vector of residual errors (one row for each age-period-specific observation). It is assumed that the **u** and **e** are uncorrelated random variables with zero means and that they are uncorrelated with the fixed effects in **X**. For significance testing and for maximum likelihood estimation, **u** and **e** are assumed to come from a population with a multivariate normal distribution. This formulation allows the estimation of the fixed effects using the design matrix **X** and the random effects (the variation in the cohort diagonals) using the **Z** design matrix. The random effects are independent of **e**, and the fixed effects are independent of the random effects and **e**. This means that the random effects of the cohorts are orthogonal to the age and period dummy variable effects (and cohort characteristics when they are included in the **X** matrix) and the "residual errors." In the traditional APCC model, there is no **Z** matrix or random effects (**u**), and the crucial assumption is that **e** and the fixed effects are independent of each other. It is one of the reasons that the cohort effects estimated in the mixed model differ from those estimated using the method of cohort residuals from the OLS model (in the section with empirical results, we will see that the OLS model with cohort characteristics gives us somewhat different estimates of the fixed effect than does the mixed model). That occurs, in part, because in addition to requiring that the residuals be orthogonal to the fixed effects, the mixed model requires that the random effects, the fixed effects, and the residuals be orthogonal to each other. We use SAS (2004) Proc Mixed to estimate this model using restricted maximum likelihood estimation, and SAS implements empirical Bayes estimates for the random effects.[11] We use an AR(1) option to account for possible autocorrelation for the observations within cohorts.

The mixed model framework provides an estimate of the variation between cohorts while controlling for age and period effects. It does not solve the problem that the age and period dummy variables may be controlling for some of the variance that “belongs to” cohorts. This is a standard problem in regression analysis when the independent variables are not orthogonal to one another. Thus, when researchers control for the effects of other variables, they may control for part of the relationship (even if it is a causal relationship) between the independent variable of interest and the outcome variable. Entering cohorts as a random variable in a model with fixed dummy variables for age groups and periods does not control the estimated age and period effects for the impacts of cohorts. To control the age group effects for period and cohort, using our method, one must enter dummy variables for the cohorts and periods and estimate the age group effects as random variables. To control the period effects for age groups and cohorts, one must enter dummy variables for age groups and cohorts and estimate the period effects as random variables.

Since the age and period dummy variables are perfectly related to cohort time of birth, to the extent that steady increases or steady decreases over time in the cohort characteristics cause increases or decreases in the dependent variable, this relationship would be suppressed by controlling for the age and period dummy variables. On the other hand, it is exactly this sort of over-time relationship that researchers often control for in time-series analyses to avoid a spurious relationship. In both the time-series context and in the context of an APCC model, it makes for a conservative test of the effects of the independent variable on the dependent variable.

Estimation of Models With Empirical Data

Figures 1a and 1b are schematic representations of the period-by-age tables that contain the data for the two empirical examples that we analyze using the APCMM. Our empirical examples involve the role of cohort characteristics in explaining shifts in the age distributions of homicides and suicides that occurred during the latter part of the twentieth century. During part of that time, younger birth cohorts became much more prone to both forms of lethal violence. We analyze data on homicide offenses (Federal Bureau of Investigation [FBI], *Crime in the United States*, various years) that were used by O'Brien et al. (1999) and data on suicide (U.S. Department of Health, Education, and Welfare, *Annual Vital Health Statistics Report*, various years) that were used by Stockard and O'Brien (2002a), except that in their analyses, the final period was 1995, while we have updated their analyses to include data for the United States through the year 2000. As is typical of such analyses, the number of years for each of the age groups is equal to the spacing between periods and is equal to the number of years included in each birth cohort. In the table for homicide offenses (Figure 1a), there are an equal number of observations for each of the age groups and for each of the periods. There are multiple observations for most cohorts, but the number of observations per cohort is markedly unbalanced. For example, the cohort born between 1940 and 1944 corresponds to those who were 15 to 19 years old in 1960, and there are seven observations on this cohort before it "ages out" of our data after 1990, when it was 45 to 49 years old. Two of the cohorts in this analysis have only one observation: the cohort born between 1910 and 1914 that was 45 to 49 years old in 1960 and the most recent cohort that was born between 1980 and 1984 and was 15 to 19 years old in 2000.

The data for suicide (Figure 1b) appear in a nearly triangular table. In this table, the oldest cohort contributes the most observations (along with the second oldest cohort). The oldest cohort, born between 1915 and 1919, was 10 to 14 years old in 1930 and "exited" our analysis after 1995, when it reached the age of 75 to 79 years old. Here, the number of observations for periods, age groups, and cohorts is unbalanced. The differences in the dimensions of these two period-by-age tables result from the availability of data for these two analyses (for details, see O'Brien et al. 1999; Stockard and O'Brien 2002a). Again, most of the cohorts contain multiple observations (the only exception, for the suicide data, is the most recent cohort that was born between 1985 and 1989 and was 10 to 14 years old in 2000).

In the analyses reported in Tables 1 and 2, we use the log of the homicide (or suicide) rate per 100,000 as the dependent variable, and we also log the cohort characteristics. Logging these variables reduced the skew in both the homicide and suicide rates and in one of the cohort characteristics (the skew for the other cohort characteristic remained insignificant). These log transformations also provided a model that assessed whether proportionate shifts in the cohort characteristics result in proportionate shifts in the homicide (or suicide) rates. Table 1 presents results from the analysis of the homicide offense data, and Table 2 presents results from the suicide data.

(text continues on page 418)

Table 1
APCC Mixed Models (and OLS APCC Model) for Homicide Offenses, 1960 to 2000

	Mixed Models						OLS Model	
	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*
Fixed effects								
Intercept	1.0963	0.0936	–6.3361	1.7966	–6.4368	1.7886	–5.4160	1.0000
Ages 15-19	1.2228	0.1171	–0.3627	0.2715	–0.3636	0.2714	–0.1986	0.1640
Ages 20-24	1.5175	0.1020	0.2092	0.2232	0.2077	0.2247	0.3407	0.1340
Ages 25-29	1.2826	0.0871	0.2315	0.1788	0.2210	0.1805	0.3096	0.1091
Ages 30-34	0.9869	0.0731	0.1976	0.1370	0.1712	0.1393	0.2420	0.0880
Ages 35-39	0.7204	0.0611	0.1956	0.0991	0.1699	0.1010	0.2251	0.0726
Ages 40-44	0.3820	0.0523	0.1243	0.0663	0.1116	0.0618	0.1430	0.0609
Ages 45-49	0.0000	—	0.0000	—	0.0000	—	0.0000	—
Period 1960	0.2205	0.1526	2.3357	0.3611	2.3236	0.3620	2.1243	0.2183
Period 1965	0.3802	0.1374	2.2393	0.3184	2.2387	0.3194	2.0423	0.1948
Period 1970	0.8243	0.1218	2.4168	0.2780	2.4156	0.2793	2.2354	0.1751
Period 1975	0.9068	0.1064	2.2380	0.2375	2.2295	0.2409	2.0757	0.1558
Period 1980	0.8844	0.0919	1.9515	0.1958	1.9405	0.2017	1.8057	0.1346
Period 1985	0.7342	0.0786	1.5374	0.1542	1.5498	0.1607	1.4168	0.1130
Period 1990	0.8268	0.0673	1.3700	0.1130	1.3835	0.1170	1.2982	0.0901
Period 1995	0.5760	0.0590	0.8579	0.0752	0.8547	0.0702	0.8239	0.0694
Period 2000	0.0000	—	0.0000	—	0.0000	—	0.0000	—
LNNMB			1.9169	0.3178	1.9097	0.3166	1.7222	0.1865
LNRCS			1.6084	0.5067	1.6554	0.5034	1.3855	0.2720

(continued)

Table 1 (continued)

	Mixed Models						OLS Model	
	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*
Random effects								
Cohorts	0.0882	0.0374	0.0093	0.0059	0.0031	0.0107		
Residuals	0.0103	0.0025	0.0104	0.0025	0.0173	0.0120	0.0130	0.0023
AR(1)					0.5443	0.3547		
Model fit								
-2χ Log likelihood	–10.3		–35.1		–37.0			
BIC	–4.9		–29.8		–29.1			

Note: APCC = age-period-cohort-characteristic; OLS = ordinary least squares; LNNMB = log of the percentage of nonmarital births; LNRCS = natural log of relative cohort size; BIC = Bayesian information criterion.

Table 2
APCC Mixed Model (and OLS APCC Model) for Suicides, 1930 to 2000

	Mixed Models						OLS Model	
	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*
Fixed effects								
Intercept	2.5534	0.1320	–7.2348	1.4742	–7.0396	1.5000	–3.4371	0.5743
Ages 10-14	–2.5598	0.1941	–5.6614	0.3825	–5.6558	0.3825	–5.6085	0.2199
Ages 15-19	0.5007	0.1841	–3.3606	0.3527	–3.3553	0.3534	–3.3100	0.2037
Ages 20-24	0.1537	0.1737	–2.4624	0.3236	–2.4667	0.3241	–2.4182	0.188
Ages 25-29	0.2420	0.1636	–2.1349	0.2961	–2.1444	0.2957	–2.0961	0.1741
Ages 30-34	0.2737	0.1538	–1.8688	0.2696	–1.8856	0.2685	–1.8418	0.1610
Ages 35-39	0.2826	0.1446	–1.6218	0.2438	–1.6404	0.2423	–1.5999	0.1490
Ages 40-44	0.2814	0.1360	–1.3872	0.2187	–1.4032	0.2173	–1.3723	0.1380
Ages 45-49	0.2326	0.1284	–1.1977	0.1943	–1.2109	0.1933	–1.1845	0.1280
Ages 50-54	0.1542	0.1220	–1.0398	0.1715	–1.0521	0.1711	–1.0320	0.1198
Ages 55-59	0.0517	0.1169	–0.9031	0.1508	–0.9118	0.1512	–0.8931	0.1135
Ages 60-64	–0.0770	0.1136	–0.7994	0.1341	–0.8068	0.1353	–0.8064	0.1104
Ages 65-69	–0.1349	0.1126	–0.6188	0.1218	–0.6229	0.1212	–0.6266	0.1091
Ages 70-74	–0.0428	0.1150	–0.2828	0.1164	–0.2875	0.1029	–0.2788	0.1112
Ages 75-79	0.0000	—	0.0000	—	0.0000	—	0.0000	—
Period 1930	–1.2244	0.2254	2.1357	0.4223	2.1151	0.4218	2.1537	0.2539
Period 1935	–0.9929	0.1941	2.1088	0.3825	2.0947	0.3825	2.0558	0.2199
Period 1940	–0.9478	0.1751	1.9137	0.3516	1.9174	0.3537	1.8647	0.2020
Period 1945	–0.9584	0.1593	1.6593	0.3224	1.6322	0.3261	1.5985	0.1855

(continued)

Table 2 (continued)

	Mixed Models						OLS Model	
	Model 1		Model 2		Model 3		Model 4	
	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*	Coefficient	*SE*
Period 1950	–0.9483	0.1449	1.4308	0.2946	1.3999	0.2997	1.3709	0.1714
Period 1955	–0.9137	0.1313	1.2332	0.2675	1.2154	0.2731	1.1902	0.1585
Period 1960	–0.5251	0.1183	1.3826	0.2392	1.3521	0.2447	1.3374	0.1431
Period 1965	–0.2832	0.1057	1.3881	0.2112	1.3624	0.2165	1.3483	0.1282
Period 1970	–0.0599	0.0936	1.3731	0.1833	1.3577	0.1883	1.3351	0.1131
Period 1975	0.1277	0.0823	1.3249	0.1560	1.3189	0.1606	1.2938	0.09877
Period 1980	0.0700	0.0718	1.0291	0.1292	1.0383	0.1330	1.0032	0.08476
Period 1985	0.1503	0.0624	0.8750	0.1029	0.8807	0.1054	0.8608	0.07124
Period 1990	0.1626	0.0545	0.6481	0.0777	0.6530	0.0780	0.6393	0.05877
Period 1995	0.1305	0.0486	0.3722	0.0557	0.3749	0.0501	0.3698	0.04885
Period 2000	0.0000	—	0.0000	—	0.0000	—	0.0000	—
LNNMB			1.6114	0.2049	1.5879	0.2098	1.5616	0.1200
LNRCS			1.8495	0.3374	1.8121	0.3401	1.8407	0.1795
Random effects								
Cohorts	0.0446	0.0187	0.0043	0.0028	0.0021	0.0034		
Residuals	0.0152	0.0024	0.0149	0.0024	0.0179	0.0039	0.0144	0.0019
AR(1)					0.3970	0.1319		
Model fit								
-2χ Log likelihood	–34.1		–59.3		–64.4			
BIC	–28.6		–53.9		–56.3			

Note: APCC = age–period–cohort-characteristic; OLS = ordinary least squares; LNNMB = log of the percentage of nonmarital births; LNRCS = natural log of relative cohort size; BIC = Bayesian information criterion.

Using the log of these rates as the dependent variable rather than treating the dependent variable as a count variable is common practice in most of the literature in sociology and criminology when the number of cases (events) is large. In our situation, the unit of analysis is the age-period-specific rates. The mean number of homicides for these units is 2,255.48, while the minimum number was 509 and the maximum number was 6,484. For our analyses of suicides, the mean number of cases on which the age-period-specific rates were based is 1,355.55, while the minimum was 65 and the maximum was 4,425.[12] We used dummy variable coding for the age groups and periods with 45-49 as the reference category for age groups for the homicide data (75-79 for the suicide data) and 2000 as the reference category for periods. The SAS (2004) program that we used to estimate our models allows us to handle binary or count-based dependent variables.[13]

In Model 1, we include the age and period dummy variables as fixed-effects predictors and cohort as a random effect. The crucial information in Model 1, which is not available in the traditional OLS APC model, is the variance estimate for cohorts. The variance among these 15 cohorts is .0882 and is statistically significant ($p < .01$); that is, while controlling for the age and period fixed effects, there is a statistically significant amount of variance between cohorts.[14]

Model 2 incorporates two cohort characteristics to explain the variation between cohorts. Both of these variables are hypothesized, from a Durkheimian framework, to influence the social integration and regulation of birth cohorts (see O'Brien et al. 1999; O'Brien and Stockard 2002; Stockard and O'Brien 2002a, 2002b). One of these characteristics is the natural log of the percentage of nonmarital births in the cohort (LNNMB). NMB is the percentage of live births in a cohort that were born to women who were not married.[15] For example, NMB for the cohort that was born between 1940 and 1944 is the percentage of the live births between 1940 and 1944 in the United States that were born to women who were not married. The second cohort characteristic is the natural log of relative cohort size (LNRCS). In this article, relative cohort size (RCS) was measured as the percentage of the resident U.S. population ages 15 to 64 who were 15 to 19 years old when the cohort was 15 to 19 years old.[16]

These two characteristics are each positively related to the logged age-period-specific homicide offense rate, while controlling for the age and period dummy variables. The coefficient for LNNMB is statistically significant at the .0001 level, while the coefficient associated with LNRCS is significant at the .01 level.[17] Including the two cohort characteristics in Model 2 reduces the variance between cohorts that is not accounted for by the model from .0882 in Model 1 to .0093. Thus, these two cohort characteristics account for nearly 89 percent $(= [(.0882 - .0093)/.0882] \times 100)$ of the variance between cohorts in Model 1. This indicates that these particular cohort characteristics are very effective in accounting for these cohort variations in homicide offending. Not surprisingly, Model 2 fits better than Model 1 according to the Bayesian information criterion (BIC): The lower the value of BIC, the better the model. Using the likelihood ratio chi-square test, we find that the fit of Model 2 is significantly better than the fit of Model 1$(\chi^2 = 24.8[= -10.3 - (-35.1)]$, with 2 degrees of freedom).

Controlling for LNNMB and LNRCS (in Model 2) shifts the age and period effects. For example, in Model 1, the strongest positive period effect is 1975, while in Model 2, it is 1970. In Model 1, the 20 to 24 age group has the strongest positive association with homicide of any of the age groups, while in Model 2, the strongest positive association is with the 25 to 29 age group. According to our cohort-based theory of the effects of cohort characteristics on rates of homicide offending, Model 1 is misspecified.

Model 3 reports the results after adding an autoregressive 1 (AR(1)) component to the model. This adjusts for autocorrelation between the observations within cohorts and estimates a single autocorrelation for all of the observations. Given the limited number of observations within each cohort, this procedure is a reasonable way to proceed. The estimated autocorrelation of .5443 is not statistically significant, and BIC indicates that estimating this autocorrelation does not produce a better fitting model—correcting for parsimony. The likelihood ratio chi-square of 1.9 with 1 degree of freedom is not statistically significant at the .05 level. Not surprisingly, the coefficients and standard errors are affected only slightly by the inclusion of the AR(1) component.

The mixed model approach has allowed us to estimate variance between cohorts, while controlling for the age and period effects, and to estimate the extent to which the cohort characteristics that we introduced in Model 2 can account for that variance. It also permits us to model dependencies between the multiple observations that appear within cohorts, although for these data, the AR(1) component was not statistically significant.

Model 4 shows the results for the traditional APCC model using OLS fixed-effect estimates. There are no random effects except for the residual in this model. In terms of the relationship of the cohort characteristics to the age-period-specific homicide rates, the substantive results are similar for all of the models that contain the cohort characteristics, as are the patterns of the age and period coefficients. The standard errors tend to be less in Model 4 (for all 17 estimates for the fixed effects in comparison to

Model 2). This is to be expected since the OLS model does not take into consideration the random error due to cohorts. The close similarity between the OLS results and the APCCMM results is due to the cohort characteristics used in the estimation. Since these cohort characteristics "explain" much of the random cohort effect, the exclusion of the random cohort effect in this analysis is not as important as it would be otherwise.

For our second empirical example, we use data on suicides from vital statistics beginning in 1930 and ending with 2000. These data are represented schematically in Figure 1b. Again, we logged the age-period-specific rates for suicide. The results of this analysis are reported in Table 2. In Model 1, we note that the variance associated with cohorts (.0446) is statistically significant ($p < .01$).

In Model 2, we find that both LNNMB and LNRCS are strongly related to the logged age-period-specific suicide rates while controlling for the age and period dummy variables. Once we enter LNRCS and LNNMB into the model, the variation between cohorts is greatly reduced. Using a proportionate reduction in variance measure, the variance between cohorts in Model 1 (.0446) has been reduced to .0043 or by 90 percent. The measures of fit indicate that Model 2 fits the data statistically significantly better than Model 1($\chi^2 = 25.2$, with 2 degrees of freedom), and BIC also suggests Model 2 as the better fitting model. These cohort characteristic appear to be very strongly related to differences in rates of suicides for cohorts.

Once more, we note shifts in the age and period effects as we move from Model 1 to Model 2 (or any of the models that contain the cohort characteristics: Models 2, 3, or 4). When we add the cohort characteristics in Model 2, we see that the suicide rates increase monotonically with age. In Model 1, these rates increase to ages 35 to 39 and then decrease until 65 to 69. The patterns of the period effects also differ between Model 1 and models that contain the cohort characteristics. Again, according to our cohort-based theory, Model 1 is misspecified.

In Model 3, we add the AR(1) component and find that it provides a statistically significant improvement to the fit of the model (using the likelihood ratio chi-square test) and that BIC suggests that this is the best fitting of the mixed models in Table 2. The autocorrelation is moderately strong (.3970), but the substantive results for these data are little affected by the addition of this AR(1) component. In Models 2 and 3, the coefficients for the cohort characteristics are quite similar, and their standard errors are close in value.

Model 4 is the traditional OLS APCC model using fixed effects only. Our substantive conclusions would not be altered using this model in terms of the relative magnitude of the cohort characteristics, the direction of their relationships, or their statistical significance. But the APCCMM approach has provided a gauge of how effective our cohort characteristics are in accounting for the cohort effects that remain after controlling for the effects of age groups and periods. The standard errors from the OLS approach tend to be underestimated (29 out of 30 times for the fixed effects in comparison to Model 3). Again, this is expected since this model does not take into consideration the error generated by the random cohort effects. Not taking this random component into account in the OLS approach leads to inaccurate estimates of the standard errors.

We conducted an auxiliary analysis to compare the method of cohort residuals that we discussed earlier to show the possibility of computing age, period, and cohort effects using an OLS model with the results from the mixed model. The OLS method involved regressing age-period-specific rates on age and period and computing a matrix of residuals. The residuals for each cohort were summed and the OLS mean residuals for each cohort calculated. We compared these mean residuals with the mixed model estimates for the cohort effects (in a model containing only the fixed effects for age groups and periods and a random effect for cohorts). SAS (Proc Mixed) produces estimates for the random effect for each cohort. We found that the correlations between these two estimates of cohort effects were .87 for the homicide data and .83 for the suicide data. We are not advocating using the mean OLS residuals to measure cohort effects rather than the mixed model approach. As Raudenbush and Bryk (2002:47) point out, the mixed model Bayes estimates of the random effects (although they are biased toward zero) produce the least mean squared error. These "cohort effects" can be interpreted as the residual variance in the dependent variable that is associated with the cohorts after controlling for age and period effects (of course, the mixed model estimation requires that these random effects, the fixed age and period dummy effects, and the residuals be mutually independent).

It is reasonable to ask why we chose to treat cohort as a random effect and period and age group as fixed-effect dummy variables. Certainly, we could treat cohort and period as fixed effects (using dummy variables to code the cohort and period effects) and treat age group as a random variable. As an example, when we used dummy variables for periods (with one period serving as a reference group) and for cohorts (with one cohort serving as a reference cohort) with the homicide offense data, we found that the mean age group residuals, using our method of residuals, indicated that the two age groups with the greatest tendency toward homicide offending were the age groups 20 to 24 and 25 to 29. The other age groups declined monotonically on both sides of this peak. The correlation between these mean age group residuals and the solutions for the random effects of age groups in the mixed model was .99. For homicide offenders, most criminologists would agree that we know the general relationship between age and homicide over a wide variety of settings (Gottfredson and Hirschi 1990). That is, the homicide rate grows rapidly with age and peaks in the mid-20s and then declines with age. We chose to use cohorts as the random variable in this article mainly because of our research interests. In previous research, our focus has been on shifts in the age distributions of homicides and suicides over time. It is natural to explain these shifts in the age distribution of lethal violence with shifts in the propensity toward lethal violence associated with cohorts.

In the normal OLS age-period model or age–period–cohort characteristic model, the residuals are contained in a single "random error" component, even though we might suspect that they are based both on systematic differences between the cohorts and on a random error component. The APCMM approach allows us to model two random variables. One random variable is specified as a patterned difference between cohorts, and the other random variable represents the random error variance.[18] Treating cohorts as a random variable helps to isolate this variance component from the residual variance.

In the analyses of both homicide (reported in Table 1) and suicide (reported in Table 2), we note a striking pattern with regards to the standard errors in both the mixed models and the OLS model. The estimated standard errors for the age dummy variables become smaller as we move from the youngest to the oldest age group. The same sort of decrease in the estimated standard errors is evident for the period effects, with the largest standard errors for the early periods and the smallest for the later periods. This "fanning out" of the confidence interval (standard errors) from the reference category was noted by Yang, Fu, and Land (2004). In unreported analyses, we changed the reference category and found that the standard errors "fanned out" from the new reference category.

Conclusion

For more than 30 years, sociologists and demographers have struggled to come to terms with the age, period, cohort conundrum. The path is littered with valiant, but not fully adequate, attempts to create analyses and techniques that will allow us to understand how these three important variables affect social outcomes. There are good reasons for believing that these three factors have significant and independent effects on outcomes such as prejudice, parental values, intellectual skills, criminal behavior, suicide, mortality, and fertility. Models that capture only two of these three effects are not satisfactory—they are misspecified.

We offer an approach to this vexing problem that uses mixed models and cohort characteristics to obtain separate estimates of the effects for age, period, cohorts, and cohort characteristics. These estimates suffer from the inevitable problem that occurs when using correlated independent variables, without unambiguous causal ordering: the problem of assigning "credit" to one or the other independent variable for "common" variance explained in the dependent variable. We also note that there is no substitute for finding theoretically important cohort characteristics. Models that contain only the fixed effects for age groups and periods, as well as random effects for cohorts, are not likely to produce accurate estimates for age groups and period effects since such models do not control for cohort effects.

As with any new method applied to empirical data, replications with diverse data sets will be needed to assess how useful this model is in separating age, period, and cohort effects in a wide range of applications. But for the data analyzed in this article, this approach (1) provided estimates of cohort effects after controlling for age and period, (2) allowed the modeling of dependencies of the observations within cohorts by estimating the autocorrelation of residuals within cohorts, and (3) allowed an assessment of how well the cohort characteristics account for the variation among the cohorts.

Notes

1. Mason and Feinberg (1985) note that $C = P - A$, "where C denotes the time of system entry, P denotes system time, and A denotes duration in the system" (p. 3).

2. Mason et al. (1973) cite several situations in which we might expect distinct age, period, and cohort effects (e.g., when political party identification or income is the dependent variable). So it is important not to ignore one of these effects.

3. Glenn (1976) and Rodgers (1982a) critique this solution. They point out that the choice of which two or more effects to specify as the same can greatly change the results of the analysis (for responses to these critiques, see Mason, Mason, and Winsborough 1976; Smith, Mason, and Fienberg 1982; Rodgers 1982b). Sasaki and Suzuki (1987) suggested a search method using a Bayesian criterion for determining the effects of age, period, and

cohort. This method assumes that successive age, period, or cohort parameters change gradually, and they note that Mason et al.'s (1973) constrained effects analysis is a special case of this method. Glenn (1989) argues that Sasaki and Suzuki's method is mechanical and, although it sometimes yields reasonable estimates, if it is used in an automatic way, it will almost inevitably lead to many incorrect solutions. Sasaki and Suzuki (1989) respond with a defense of their method.

4. A *cohort characteristic* is the term used by O'Brien, Stockard, and Isaacson (1999); the term used by Heckman and Robb (1985) is *proxy variable*.

5. A series of recent articles by O'Brien and Stockard (O'Brien and Stockard 2002; O'Brien et al. 1999; O'Brien 1989, 2000; Stockard and O'Brien 2002a, 2002b) has examined the effects of age, period, and cohort on outcome variables by using cohort characteristics to represent the effect of cohorts. In this article, we use the same data sources but extend their analyses to the year 2000 and use mixed models to examine cohort effects.

6. The $(A + P) - 1$ count of the number of cohorts is appropriate for a rectangular matrix such as the one in Figure 1a.

7. Unless, of course, we transform one or more of these independent variables to be nonlinear or the effects of two or more of the coefficients for the dummy variables to be equal, or place some other restriction on these fixed effects for age, period, or cohort.

8. We choose to use age groups and periods as fixed effects and to examine the cohort residuals for two reasons. (1) Our focus is on cohort effects, and we want to control as completely as possible any effects of age and period before concluding that there are effects of cohorts. (2) Including age group and period dummy variables controls for any effects of cohorts that are linearly related to cohort time of birth (O'Brien 2000). This is analogous to detrending the cohort effect for time prior to analyzing its effect.

9. The SAS Proc Mixed procedure, used to estimate the age, period, cohort mixed model (APCMM), allows us to model the dependencies of observations nested within cohorts using the repeated command for autocorrelation within cohorts.

10. Other fixed effects, such as an interaction between an age group and a period, can be added to the $\mathbf{X}$ matrix and estimated in the β vector.

11. We employ a random-effects model, in part because we think it is more intuitive to think of the residuals along, for example, the cohorts when age and period dummy variables have been controlled as the equivalent of blocks of observations that share a "locational" trait in common. It is possible to obtain the same results by treating cohorts as the second level in a hierarchical linear model. In that case, we can treat the intercepts of the cohorts as random and the dummy variables for age and period as fixed effects at Level 1. In SAS, that means specifying the dummy variables for age and period on the model statement and the randomly varying intercepts on the random statement. Specifically, we use "random cohorts/solution;" in SAS to treat the cohorts as random variables. The equivalent model, in terms of results, treats cohorts as a Level 2 variable in a hierarchical linear model. In SAS, we would use the following statement: "random intercept/subject = cohorts solution;."

12. The homicide offenses are based on Uniform Crime Report data (Federal Bureau of Investigation, *Crime in the United States*, various years) and the suicide rates are based on data from the U.S. Department of Health, Education, and Welfare (*Annual Vital Health Statistics Report*, various years).

13. One reviewer pointed out that our log transformation amounts to assuming that the rates are log-normally distributed, which is an approximation to the Poisson distribution. In

auxiliary analyses, we used the GLIMMIX macro in SAS to run Poisson models, with the dependent variable being the age-period-specific *number* of homicides (suicides), and the age-period-specific population was used as an exposure variable. We do not report the results of these analyses—but they are substantively similar to the results that we do report in Tables 1 and 2. The major difference is that the natural log of relative cohort size (LNRCS) is only marginally significant ($p < .10$) in both the homicide and suicide analyses.

14. In general, the null hypothesis tests that the variance of the random effects equals zero may be highly misleading when it is based on the standard errors generated in these mixed models, especially when the sample size is small (Raudenbush and Bryk 2002:64). Thus, in addition to this test, we conduct the likelihood ratio chi-square test by comparing the fit of the model with and without the random cohort component. This test tends to be conservative, decreasing the chances of rejecting a false null hypothesis (Raudenbush and Bryk 2002:284). The statistical significance (significant or not significantly different from zero) of the variances for the cohort effects based on this test are the same as those based on the standard errors for each model that includes the random cohort effects in Tables 1 and 2.

15. These data were drawn from vital statistics (U.S. Bureau of the Census 1946, 1990).

16. These data were drawn from the *Current Population Surveys* (U.S. Bureau of the Census, various years). In one case, for the suicide data, we analyze data for those ages 10 to 14 in 2000. For this cohort, we do not have data on the percentage of those 15 to 64 who are 15 to 19 when the cohort is 15 to 19 since the cohort has not reached that age by the year 2000. We use instead the percentage of those 10 to 59 who are 10 to 14 when this cohort is 10 to 14 to measure relative cohort size (RCS) for this one cohort's single observation.

17. Since both the dependent variable and the cohort characteristic are logged, we can interpret the coefficients for the natural log of the percentage of nonmarital births (LNNMB) and LNRCS as the impact of a percentage increase in these measures on the percent increase in the dependent variable. Technically, these coefficients represent the point elasticity of the dependent variable (homicide) with respect to the independent variable: the instantaneous rate of change in terms of a proportionate change in the independent variable on a proportionate change in the dependent variable.

18. The mixed model estimates the fixed and random effects simultaneously.

References

Alwin, Duane F. 1990. "Cohort Replacement and Changes in Parental Socialization Values." *Journal of Marriage and Family* 52:347-60.

———. 1991. "Family of Origin and Cohort Differences in Verbal Ability." *American Sociological Review* 56:625-38.

Alwin, Duane F. and Jon A. Krosnick. 1991. "Aging, Cohorts, and the Stability of Sociopolitical Orientation Over the Life Span." *American Journal of Sociology* 97:169-95.

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

———. 1996. "Human Lives in Changing Societies: Life Course and Developmental Insights." Pp. 31-62 in *Developmental Science*, edited by Robert B. Cairns, Glen H. Elder Jr., and E. Jane Costello. New York: Cambridge University Press.

Elder, Glen H., Jr. and Avshalom Caspi. 1990. "Studying Children in a Changing Society: Sociological and Personological Explorations." Pp. 201-47 in *Studying Persons and Lives*, edited by A. I. Rabin, Robert Zucker, and Susan Frank. New York: Springer.

Farkas, George. 1977. "Cohort, Age, and Period Effects Upon the Employment of White Females: Evidence for 1957-1968." *Demography* 67:33-42.

Firebaugh, Glenn and Kevin Chen. 1995. "Vote Turnout of Nineteenth Amendment Women: The Enduring Effects of Disenfranchisement." *American Journal of Sociology* 100:972-97.

Firebaugh, Glenn and Kenneth E. Davis. 1988. "Trends in Antiblack Prejudice, 1972-1984: Region and Cohort Effects." *American Journal of Sociology* 94:251-72.

Glenn, Norval D. 1976. "Cohort Analysts' Futile Quest: Statistical Attempts to Separate Age, Period and Cohort Effects." *American Sociological Review* 5:900-1.

———. 1989. "A Caution About Mechanical Solutions to the Identification Problem in Cohort Analysis: Comment on Sasaki and Suzuki." *American Journal of Sociology* 95:754-61.

Gottfredson, Michael R. and Travis Hirschi. 1990. *A General Theory of Crime*. Stanford, CA: Stanford University Press.

Greenberg, B. G., J. J. Wright, and C. G. Sheps. 1950. "A Technique for Analyzing Some Factors Affecting the Incidence of Syphilis." *Journal of the American Statistical Association* 25:373-99.

Hall, R. E. 1971. "The Measurement of Quality Change from Vintage Price Data." Pp. 240-71 in *Price Indexes and Quality Changes*, edited by Z. Griliches. Cambridge, MA: Harvard University Press.

Heckman, James and Richard Robb. 1985. "Using Longitudinal Data to Estimate Age, Period, and Cohort Effects in Earning Equations." Pp. 137-50 in *Cohort Analysis in Social Research: Beyond the Identification Problem*, edited by William M. Mason and Stephen E. Fienberg. New York: Springer-Verlag.

Kahn, Joan R. and William M. Mason. 1987. "Political, Alienation, Cohort Size, and the Easterlin Hypothesis." *American Sociological Review* 52:155-69.

Mason, Karen O. and Yu-Hsia Lu. 1988. "Attitudes Toward Women's Familial Roles: Changes in the United States, 1977-1985." *Gender & Society* 2:39-57.

Mason, Karen O., William M. Mason, and H. H. Winsborough. 1976. "Reply to Glenn." *American Sociological Review* 5:904-5.

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

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

O'Brien, Robert M. 1989. "Relative Cohort Size and Age-Specific Crime Rates: An Age-Period-Relative-Cohort-Size Model." *Criminology* 27:57-78.

———. 2000. "Age Period Cohort Characteristic Models." *Social Science Research* 29:123-39.

O'Brien, Robert M. and Jean Stockard. 2002. "Variations in Age-Specific Death Rates: A Cohort Explanation for Changes in the Age Distribution." *Social Science Research* 31:124-50.

O'Brien, Robert M., Jean Stockard, and Lynne Isaacson. 1999. "The Enduring Effects of Cohort Characteristics on Age-Specific Homicide Rates, 1960-1995." *American Journal of Sociology* 104:1061-95.

Raudenbush, Stephen, W. and Anthony S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed. Thousand Oaks, CA: Sage.

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

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

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

Sacher, G. A. 1960. "Analysis of Life Tables With Secular Terms." Pp. 253-7 in *The Biology of Aging*, edited by J. D. Ebert, H. B. Glass, and N. W. Shock. Washington, DC: American Institute of Biological Sciences.

SAS. 2004. SAS/MIXED Version 8. Cary, NC: SAS Institute Incorporated.

Sasaki, Masamichi and Tatsuzo Suzuki. 1987. "Changes in Religious Commitment in the United States, Holland, and Japan." *American Journal of Sociology* 92:1055-76.

———. 1989. "A Caution About the Data to be Used for Cohort Analysis: Reply to Glenn." *American Journal of Sociology* 95:761-5.

Smith, Dwayne M. 1986. "The Era of Increased Violence in the United States: Age, Period, or Cohort Effect?" *The Sociological Quarterly* 27:239-51.

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

Stockard, Jean and Robert M. O'Brien. 2002a. "Cohort Variations and Changes in Age-Specific Suicide Rates Over-Time: Explaining the Recent Increase in Youth Suicide." *Social Forces* 81:606-42.

———. 2002b. "International Comparisons of Cohort Effects on Suicide Rates: The Role of Contextual Variables." *American Sociological Review* 67:854-72.

U.S. Bureau of the Census. 1946. *Vital Statistics of the United States: Natality*. Washington, DC: Government Printing Office.

———. 1990. *Vital Statistics of the United States: Natality*. Washington, DC: Government Printing Office.

Weil, Frederick D. 1987. "Cohorts, Regimes, and the Legitimation of Democracy: West Germany Since 1945." *American Sociological Review* 52:308-24.

Wilson, James A. and Walter R. Gove. 1999. "The Intercohort Decline in Verbal Ability: Does It Exist?" *American Sociological Review* 64:253-66.

Yang, Yang, Wenjiang J. Fu, and Kenneth C. Land. 2004. "A Methodological Comparison of Age–Period–Cohort Models: The Intrinsic Estimator and Conventional Generalized Linear Models." Pp. 77-110 in *Sociological Methodology 2004*, edited by Ross M. Stolzenberg. New York: American Sociological Association.

Robert M. O'Brien is a professor of sociology and department head at the University of Oregon. His primary research areas are criminology (with a focus on ecological approaches to explaining crime rates) and the development and use of quantitative methods to better understand social phenomena. He has two forthcoming publications with Jean Stockard. One involves the development of a unified explanation of shifts in the age distributions of homicides and suicides over time, and the other involves tracing the effects of cohort characteristics (for cohorts born as early as 1875) on the age distribution of suicide rates across countries.

Kenneth Hudson is an assistant professor of sociology and research associate in the Youth Violence Prevention Program at the University of South Alabama. His past research has focused on poverty, inequality, nonstandard work arrangements, and labor market dualism in the United States. His current projects include an examination of the relationship between neighborhood poverty, schools, and persistent youth offending, as well as a study of the contemporary conflict between modernity and religious fundamentalism. His recent articles include "The Dark Side of the Protestant Ethic: A Comparative Analysis of Welfare Reform" with Andrea Coukos in *Sociological Theory* (2005) and "The New Labor Market Segmentation: Labor Market Dualism in the New Economy," which is forthcoming in *Social Science Research*.

Jean Stockard is a professor in the Department of Planning, Public Policy and Management at the University of Oregon. Much of her work focuses on trends in social phenomena related to the well-being of youth and families. Recent publications have dealt with changing age patterns of lethal violence (with Robert O'Brien), changes in rates of nonmarital births (with Jo Anna Gray, Joe Stone, and Robert O'Brien), and the relationship of the built environment to physical health.