

Can Cohort Replacement Explain Changes in the Relationship Between Age and Homicide Offending? Author(s): Robert M. O'Brien and Jean Stockard Source: Journal of Quantitative Criminology, Vol. 25, No. 1 (March 2009), pp. 79-101 Published by: Springer Stable URL: http://www.jstor.org/stable/23367546 Accessed: 11-08-2016 00:09 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at http://about.jstor.org/terms

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Springer is collaborating with JSTOR to digitize, preserve and extend access to *Journal of Quantitative Criminology*

ORIGINAL PAPER

Can Cohort Replacement Explain Changes in the Relationship Between Age and Homicide Offending?

Robert M. O'Brien · Jean Stockard

Published online: 18 September 2008 © Springer Science+Business Media, LLC 2008

Abstract This paper focuses on shifts in the age distribution of homicide offending in the United States. This distribution remained remarkably stable with small but significant changes over a long period of time. Then between 1985 and 1990 the rates of homicide offending doubled for 15-to-19 year olds and increased nearly 40% for 20-to-24 year olds, while the homicide offending rates decreased for those over 30. In addition to this "epidemic of youth homicide," which lasted through the mid-1990s, there have been systematic changes in the age distribution of homicide in the United States associated with cohort replacement over the past 40 years. We introduce an estimable function approach for estimating the effects of age, period, and cohort. The method allows us to assess simultaneously the impacts of periods and cohorts on the age distribution of homicide offending. We find that although the age curve remains relatively stable, there are shifts in it associated systematically with cohort replacement. Cohort replacement accounts for nearly half of the upturn in youth homicides during the epidemic of youth homicides, but a significant fraction of that upturn is not associated with cohort replacement.

Keywords Age distribution of homicide · Epidemic of youth homicide · Estimable functions · Cohort replacement

Introduction

In a 1983 publication Hirschi and Gottfredson (1983) postulated an "invariant" relationship between age and crime, suggesting that, across cultures and time, the age distribution of crime remains the same—rising from adolescence to a peak in the early twenties and then gradually declining to very low levels by middle age. In the decade that followed

J. Stockard

R. M. O'Brien (\boxtimes)

Department of Sociology, University of Oregon, Eugene, OR 97403, USA e-mail: bobrien@uoregon.edu

Planning, Public Policy and Management, University of Oregon, Eugene, OR, USA

Hirschi and Gottfredson's (1983) initial statement, however, the age distribution of homicide offending, the measure of violent crime generally considered the most reliable, rose sharply among younger age groups, and declined among older age groups. As a result, the age distribution of homicide altered with the highest rates among adolescents rather than among those in their twenties and relatively lower rates in the older age ranges. Clearly the age distribution of crime is not rigidly invariant [we note that Hirschi and Gottfredson (1983) did not consider it to be rigidly invariant].

This research focuses on shifts in the age distribution of homicide offending and not on the absolute levels of homicide offending. We consider two factors that may account for changes in this distribution over the period 1965–2005 in the United States. The first is the most popular explanation for the epidemic of youth homicide (which occurred during the last half of the 1980s through the mid-1990s): the development of crack cocaine drug markets during this period (Blumstein 1995; Blumstein and Cork 1996). The second, cohort replacement, suggests that different birth cohorts have characteristics that predispose them to higher or lower propensities for involvement in homicide offending (O'Brien et al. 1999; Savolainen 2000). This explanation could account for shifts in the age distribution of homicide offending throughout the period from 1965 to 2005. To test the extent to which these two explanations can account for the changes in the age distribution of homicides, we present a new form of analysis for examining age, period, and cohort effects.

Background Literature

The Invariance Thesis

Hirschi and Gottfredson (1983, pp. 553–54) original statement of age invariance was strongly worded: "... In this paper we advance and attempt to defend the following theses: (1) the age distribution of crime is invariant across social and cultural conditions." They bolstered this thesis by presenting evidence that the age distribution of homicide offending has remained remarkably "invariant" across countries and time.¹ Not surprisingly such a provocative and strongly stated thesis created a strong reaction (Greenberg 1985), and this challenge was met with an equally strong response (Hirschi and Gottfredson 1985).²

Although Hirschi and Gottfredson (1983) used the word invariant, a careful reading of their original article indicates that their interpretation is not so rigidly invariant. For instance, they state:

"So, while we may find social conditions in which age does not have as strong an effect as usual, the isolation of such conditions does not lead to the conclusion that age effects may be accounted for by social conditions. On the contrary, it leads to the conclusion that in particular cases the age effect may be to some extent obscured by countervailing social processes" (Hirschi and Gottfredson 1983, p. 561).

² There were many other papers published that discussed the invariance of crime rates and what invariance would look like; see, for example, Britt (1992); Greenberg (1994); Steffensmeier et al. (1989). From these papers and others we would conclude that the age distribution of crime is not strictly invariant.



¹ In the mid-1980s Cook and Laub (1986) pointed to a surprising stability of youth crime rates given major changes in the family structure.

They further note (footnote 9, p. 561) that their invariance hypothesis began as an empirical generalization derived from the relationship between observable variables:

"Defense of such generalizations in the face of variation across indicators, however minor, requires some degree of conceptualization. In the discussion here, we have substituted 'tendency to commit criminal acts' for 'crime,' and our invariance hypothesis has thus become 'the age distribution of the tendency to commit criminal acts is invariant across social and cultural conditions.' This revised hypothesis is not strictly at the mercy of the facts and is therefore not necessarily contradicted by the observation that the relationship between age and various indicators of crime is not precisely the same under all conditions."

We label their argument for a stable age distribution of homicide as the "invariance thesis." It involves an age distribution of crime that is relatively stable across long periods of time though not rigidly invariant. In other words, the pattern of homicide offending that increases from adolescence through the early twenties and then declines is expected to appear across historical eras and cultures, although the strength of these age differences may vary from one situation to another.³

Rigid age invariance for homicide offending in the United States is not consistent with the data. As alluded to above, the age patterning of homicide altered dramatically from the 1985 to early and mid-1990s. For instance, in 1985 the homicide offense rate for those 15-to-19 years old was 16.32 per 100,000 but it rose to 36.52 per 100,000 in 1990—a 115% increase. The homicide rate for the same periods also rose sharply for those who were 20-to-24 years old from 21.11 to 29.10—a 38% increase.⁴ At the same time the rates for those in the 5-year age categories from 30 to 34 and above dropped. The rates for youth remained at these very high levels in 1995 and then dropped dramatically in 2000 both absolutely and in relationship to other age groups (see Fig. 1). We discuss the invariance thesis not to test it in its rigid form, but to point out that there is a patterning to the age distribution of homicides in the United States and to examine relatively small, but significant (e.g., the epidemic of youth homicide), changes in that pattern over the period that we investigate.

The Cohort Replacement Thesis

Cohort replacement focuses on cohort differences in the propensity to engage in criminal activity. For example, if the cohort born between 1940 and 1944 is particularly prone to commit homicide, then in 1960 when they are 15-to-19 years old we can expect that age group to commit homicides at a higher than expected rate. When this same cohort is 20-to-24 years old in 1965, we can expect that age group to commit homicides at a higher

 $^{^3}$ Hirschi and Gottfredson (1983) view the age distribution of homicide as universal and not socially determined. That is, it is likely to be biological/genetic (see Gove 1994, for a similar view on age and violence). In contrast, David Greenberg (1977, 1985) portrays this distribution as socially created by the social position of youth in economically advanced societies. Our data will not be able to adjudicate between these sorts of explanations.

⁴ The UCRs provide the number of those arrested for homicide in the United States in various age categories for those agencies reporting 12 months of data during the year. We used these age-specific arrest counts to calculate these and other rates as described in the methods section of this paper.

	15 10	00.04	05.00	00.04	95.00	40.44	45.40	E0 E4	FF F0	60.64
	12-19	20-24	23-29	30-34	33-39	40-44	43-49	50-54	55-59	00-04
1965	10	9	8	7	6	5	4	3	2	1
	9.07	15.18	14.69	11.70	9.76	7.41	5.56	4.60	3.13	2.38
1970	11	10	9	8	7	6	5	4	3	2
	17.22	23.75	20.09	16.00	13.13	10.10	7.50	5.68	4.38	2.78
1975	12	11	10	9	8	7	6	5	4	3
	17.54	25.62	21.05	15.81	12.83	10.52	7.32	4.91	3.34	2.99
1980	13	12	11	10	9	8	7	6	5	4
	18.00	23.97	18.88	15.23	12.32	8.80	6.76	4.36	3.28	2.16
1985	14	13	12	11	10	9	8	7	6	5
	16.32	21.10	16.79	12.58	9.60	7.50	5.31	4.32	3.31	1.90
1990	15	14	13	12	11	10	9	8	7	6
	35.17	29.10	18.00	12.44	9.38	6.81	5.17	3.38	2.36	1.77
1995	16	15	14	13	12	11	10	9	8	7
	35.08	31.93	16.76	10.05	7.25	5.47	3.67	2.68	2.50	1.39
2000	17	16	15	14	13	12	11	10	9	8
	14.63	18.46	10.90	6.63	5.41	3.74	2.30	1.70	0.89	0.64
2005	18	17	16	15	14	13	12	11	10	9
	13.87	18.70	11.85	6.80	4.69	3.69	3.09	1.74	1.22	0.76

Fig. 1 Period by age table containing the age-period-specific homicide rates per 100,000

than expected rate.⁵ Such shifts associated with cohorts are expected to endure throughout the lifespan of the cohort. In this way it is possible for cohort replacement to create shifts in the age distribution of homicides over time. On the other hand, if cohorts were not associated with homicide offending and we found there was a higher than expected rate of offending by 15-to-19 year olds in 1960, there would be no reason to predict an unusually high rate of offending for those 20-to-24 years old in 1965.

O'Brien et al. (1999), Savolainen (2000) among others have used such an approach and suggested that shifts in the age distribution of homicides are associated with the relative size of birth cohorts (baby booms and busts) and the family structure of cohorts (growing up in single parent families and the percentage of non-marital births). The effect of booms and busts in cohort size on the outcomes for cohorts is often labeled the Easterlin effect (Easterlin 1978, 1987). Such cohorts are likely to have fewer parents per child, fewer teachers per child when they reach school age, fewer adults per child for supervision, and fewer entry level jobs per entry level worker when they hit the job market. O'Brien, et al. (1999, p. 1063) note that this results in "lower levels of supervision and attention from parents, teachers, counselors, and other adults as well as more crowded homes and schools."

With respect to the family structure indicators of cohorts, those growing up in single parent families are four or five times more likely to grow up in poverty (O'Hare 1996), and

 $^{^{5}}$ Our data come from the Uniform Crime Reports and that source reports the number of arrests for those (for example) who are 25–29 in 1970. These reports do not tell us the year in which these people were born. It turns out that someone who is 29 in 1970 could be born in 1940 or 1941, and those who are 25 in 1970 could be born in 1940 or 1941, and those who are 25 in 1970 could be born in 1944 or 1945. So to say that they are from the cohort born between 1940 and 1944 is "shorthand" and no better than saying they were born between 1941 and 1945 or that they were born between 1940 and 1945. This final classification would create overlapping birth cohorts when we classified each of the age groups to their birth cohort. This problem is well illustrated using Lexis diagrams (Vandeschrick 2001).

these children are more likely to grow up in less desirable and safe neighborhoods, have adequate medical care, day care, and after school care (McLanahan and Sandefur 1994). At the macro-level there is much evidence that the proportion of single parent or disrupted families is related to homicide rates (Blau and Blau 1982; Jacobs and Helms 1996; Huff-Corzine et al. 1986; Messner 1983; Messner and Sampson 1991; Sampson 1985, 1986; Williams 1984; Williams and Flewelling 1988). Sampson finds that this relationship is especially strong for juvenile crime (Sampson 1987). And there has been speculation and evidence that cohorts characterized by high rates of single parent or disrupted families have higher rates of homicide For instance, Sampson and Wilson (1995) state that "the roots of urban violence among today's 15-to-21 year old cohort may stem from childhood socialization that took place in the late 1970s and early 1980s" (p. 53), a time when family disruption began increasing sharply. O'Brien et al. (1999) and Savolainen (2000) have shown that cohorts characterized by high percentages of single parent families or families with a high percentage of non-marital births are more likely to produce homicide offenders.

We label the argument that shifts in the age distribution of homicide offending is due to shifts in the propensity of cohorts to be involved in homicide offending the "cohort replacement thesis." The methods introduced in this paper allow us to fully account for any cohort effects on the age-period-specific homicide rates (not just cohort effects associated with specific cohort characteristics, such as, relative cohort size, or measures of family structure). This moves us beyond age-period-cohort characteristic models (O'Brien 2000).

The Crack Cocaine Thesis

The most popular explanation for the dramatic increase in youth homicide from the late 1980s and through the mid 1990s relative to the rates of older age groups suggests that these changes resulted from the development and subsequent decline of the crack cocaine drug market (Blumstein 1995). Crack cocaine became available during the mid 1980s allowing users to purchase a "hit" for 5–10 dollars rather than purchase powdered cocaine at a price that was out of reach of many people (especially the poor).⁶ Youth became dealers in this expanded market selling small quantities of this relatively inexpensive drug. The financial rewards from this illegal work were far better than other work offered to youth in poor neighborhoods. Furthermore, juveniles were not subject to adult prosecution and the strong sentences associated with selling crack cocaine. Given the buyers of crack often had little money, there were many transactions for a limited number of hits. These transactions took place on the street and in neighborhoods and necessitated carrying money. The market was lucrative and juveniles involved in this drug trade needed guns to protect themselves and their markets. Other youth purchased guns for self protection, which increased the potential for violence even more (Blumstein and Cork 1996; Fryer et al. 2005). This lethal mix led to increases in homicide offending among young adults and adolescents.

Additional evidence supports this compelling story. During this period, the mix of gun related homicides increased, especially for youth (Blumstein 1995; Cook and Laub 1998, 2002). Those cities that experienced the crack cocaine epidemic earliest, recorded the earliest onsets of epidemics of youth homicide (Cork 1999). The crack cocaine thesis has also been used to explain the end of the epidemic of youth homicide and the return to a

⁶ "Crack cocaine is a smokeable, solid form of cocaine that is obtained by evaporating a solution of cocaine, sodium bicarbonate (baking soda), and water ... crack was a technological innovation in that it was much easier and less hazardous to produce than other forms of 'freebase,' which typically necessitated the use of flammable ether" (Cork 1999, p. 381).

more typical age distribution of homicides. The argument is that by the late 1990s the crack cocaine market became regularized with less need for violence as markets were established (Baumer et al. 1998; Blumstein and Rosenfeld 1998; Steffensmeier and Harer 1999). As Fryer et al. (2005, p. 7) note, "by the year 2000, we observe little impact of crack ... We hypothesize that the decoupling of crack and violence may be associated with the establishment of property rights and the declining profitability of crack distribution.

We label this explanation for the sharp upturn in the homicide offense rates for youth relative to the rates for older offenders and its return to a "more typical" age distribution for homicides as the "crack cocaine thesis." Both the crack cocaine thesis and the cohort replacement thesis suggest factors that may represent some of the "countervailing social processes" alluded to by Hirschi and Gottfredson (1983, p. 561) that obscure an underlying age pattern.

In the following section we describe the analytic techniques used to examine the hypothesized age, cohort replacement, and crack cocaine effects. The important methodological innovation, which we label the "estimable function approach," allows us to examine the effects of cohorts without relying on proxies (cohort characteristics) for these effects as has been done in prior research (O'Brien et al. 1999 and Savolainen 2000). The estimable function approach can determine whether there is a significant amount of variation in the age-period-specific homicide rates associated with birth cohorts after controlling for age and period. This provides a way of assessing whether cohort replacement can explain shifts in the age distribution of homicides over time and whether it can explain the sharp upturn in age-period-specific homicide rates for youth associated with the epidemic of youth homicide. If it cannot explain this sharp upturn, then there is need for some other explanation.

Analytical Methods

To capture proportional shifts in the level of the age distribution from period to period, we use dummy variables to represent periods. We do this to control for overall increases or decreases in the homicide rate from period to period—since our focus is on changes in the age distribution of homicides and not the overall level of homicide offending. To capture the average shape of the age distribution over time, we use dummy variables for age groups. To capture shifts in this age distribution due to cohort replacement, we use dummy variables for cohorts. To capture unique shifts in the age distribution of homicide offenses during the epidemic of youth homicide offending that are not due to cohort replacement, we use appropriate age by period interactions. Using dummy variables for age, period, and cohort allows the form of the relationship of these variables to the dependent variable to be non-linear. The result is an age-period-cohort model with some specific age by period interactions.

Equation 1 represents a multiple classification model in which the rates in each cell of Fig. 1 constitutes the dependent variable and these cells are associated with particular values for age dummy variables, period dummy variables, and cohort dummy variables:

$$\ln(Y_{ij}) = \mu + \alpha_i + \pi_j + \chi_k + \sum_{\substack{\text{Selected} \\ i \text{ and } j}} (a_i \times \pi_j) + \varepsilon_{ij}, \tag{1}$$

🖉 Springer

where μ is the intercept, the effect of the ith age group is given by α_i , the effect of the jth period is given by π_j , the effect of the kth cohort his given by χ_k , the product term $\alpha_i \times \pi_j$ is a selected interaction between an age group dummy variable and an period dummy variable, and ε_{ij} is the random disturbance term. In this equation i = 1, ..., A - 1; j = 1, ..., P - 1, and the k subscript runs through (A + P) - 2, where A is the number of age groups and P is the number of periods. $Ln(Y_{ij})$ is the natural log of the age-period-specific homicide offending rate for the ith age group and jth period.

To think about this analysis visually, we refer to Fig. 1. We use a dummy variable to represent each row (period) of this table (except that one period is reserved for the reference group). We use dummy variables to code for each of the columns (age groups) of this table (except that one age group is reserved for the reference group). We use dummy variables to code for each of the main diagonals (cohorts) in Fig. 1 (except that one cohort is reserved for the reference group). Note that the effect of each period is by construction⁷ the same across all age groups, the effect of each age group is by construction the same across all periods, and the effect of each cohort is by construction the same throughout its life span. We plan (at times) to add four interaction terms to this model: terms for those 15-to-19 and 20-to-24 in 1990 and 1995.

The statistical problem with such a model is well known: there is a linear dependency between the dummy variables for age, period, and cohort. If one knows the age group and period associated with a given rate, one can determine the birth cohort associated with that rate; if one knows the period and birth cohort, one can determine the age group; and if one knows the age group and birth cohort, the period is known. More formally; Age = Period – Cohort, Period = Age + Cohort, and Cohort = Period – Age.

This linear dependency does not allow the simultaneous estimation of the α_i , π_i , and χ_k coefficients in Eq. 1. This is the classic Age-Period-Cohort conundrum. These coefficients are "not estimable" (Searle 1971), and it is this quandary that has long vexed researchers who attempt to simultaneously study age, period, and cohort effects. Note that it is the specific coefficients (for example the period effects for 1980, 1985, etc. or the effects associated with the cohort born between 1940 and 1944, 1945 and 1950, etc.) that are not estimable in a model that contains all of the age, period, and cohort dummy variables. It is possible, however, to determine whether each of the sets of dummy variables and the interactions between age and period account for variance that is not accounted for by the other sets of dummy variables and age by period interactions in the model. This is possible because the predicted value of the dependent variable in such models is an estimable function (Scheffé 1959; Searle 1971). In our case the estimated values of the age-periodspecific homicide rates are estimable (even though the specific age, period, and cohort, coefficients are not estimable). Additionally, these estimated values of the age-periodspecific homicide rates are best linear unbiased estimates of these rates (Scheffé 1959; Searle 1971).

Given that we can estimate the best linear unbiased values of the age-period-specific homicide rates based on age, period, and cohort and we can predict the values of the ageperiod-specific rate of homicide using the age and period dummy variables only, we can determine whether the predictions are significantly improved by adding cohorts to the model and just how much they are improved. Because the approach allows cohorts to

⁷ This means that these dummy variables are fixed effects so the estimated coefficient for a period is the same across all age groups in that period; the estimated coefficient for an age group is the same across all periods in that age group, the estimated coefficient for a cohort is the same across all combinations of age and period within that cohort.

be fully modeled in the prediction of the age-period-specific homicide rates, we can determine the extent to which cohort replacement can explain the epidemic of youth homicide and the extent to which some other explanation is necessary. We calculate the proportions of variance uniquely associated with the age effects, period effects, cohort effects, and the age by period interaction effects to determine the necessity of each of these factors in predicting the age-period-specific homicide rates. We also graph the best linear unbiased predictors of the age-period-specific homicide rates based on age, period, and cohort effects and compare these to the observed age-period-specific homicide rates as well as the age-period-specific rates predicted on the basis of age and period alone. This graphical approach using estimable functions has not been used in criminological studies nor, as far as we are aware, has it been used in other areas. It is an important interpretive device.

Although the estimable function approach to age-period-cohort modeling that we use is unusual, we note that some authors (Mason et al. 1973; Robertson et al. 1999; Smith 2004) have been aware that in models that are just identified the predicted values of the age-period-specific rates are the same even though the coefficients for the age, period, and cohort effects are likely to be different (often radically different) depending on the constraint used to make the age-period-cohort model just identified (as opposed to being underidentified). For example, a common constraint used to identify the individual age, period, and cohort coefficients has been to make two of the age coefficients equal, or two of the period coefficients produce just identified models (saturated models) that may yield radically different estimates of age, period, and cohort effects, but produce the same fit to the observed data and the same predicted values for the age-period-specific rates.⁹ But these researchers have taken little advantage of the estimability of the age-period-specific values of the dependent variable in substantive research.

Summarizing the estimable function approach—the predicted values produced are the best linear unbiased estimates of the age-period-specific rates based on the age, period, and cohort dummy variables. This allows us to find the best predicted values based on age and period, and on age, period, and cohort, and on age, period, cohort and selected interactions even though the parameter estimates for the specific ages, periods, and cohorts are not uniquely estimable. We use this property to see how these factors affect the fit of the model in terms of predicted values and how each of these factors contribute to the variance associated with the age-period-specific homicide rates that is not associated with the other factors. A researcher might conclude that without uniquely identified values for say the individual effects of each cohort, there is little use in pursuing an analysis. We show that there is much to be learned without such estimates.

⁸ Smith (2004, p. 113) states, "regardless of the identifying constraint, the estimated model has the same degrees of freedom and the same goodness-of-fit. This means that the test for whether an APC decomposition of variance is better than one based on, say, age and period alone is independent of the constraint chosen."

⁹ "[T]he estimates from the full APC models are sensitive to the choice of equality constraints on the parameters of the model. Specifically, different restrictions that equate coefficients of different subsets of adjacent categories can lead to widely different trend estimates of age, period, and cohort effects, all of which fit the data equally well" (Yang et al. 2004, p. 96).

Estimable Functions: Predicted Values and Partitioning the APC Variance

In this section we briefly describe estimable functions.¹⁰ It is well known that the least squares solution leads to unique solutions for the regression coefficients when the design matrix for the independent variables is of full column rank: $\mathbf{b} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$, where \mathbf{X} is the design matrix for the independent variables with a column of ones in the first column followed by columns for each of the independent variables (P - 1 dummy variables of periods, A - 1 dummy variables for age groups, and C - 1 dummy variables for cohorts). **Y** is a column vector of observations on the dependent variable. Both \mathbf{X} and \mathbf{Y} have *n* rows (one for each observation). **b** is a column vector of unique regression coefficients for the independent variables (with the first element representing the intercept).

In the situation described above, however, the individual parameters are non-estimable, because the design matrix, and thus the X'X product matrix of independent variables, is not of full column rank. The regular inverse does not exist for X'X, if we include dummy variables (excluding the reference categories) for each age, period, and cohort. The consequence is that there are an infinite number of solutions for **b**. It is still possible, however, to compute one of the many possible non-unique least squares solutions by employing a generalized inverse (see Scheffé 1959; Searle 1971). In this case, we can estimate the regression coefficients as $\mathbf{b}^0 = (\mathbf{X}'\mathbf{X})^{-}\mathbf{X}'\mathbf{Y}$ where the superscripted minus sign indicates that the inverse is a generalized inverse, and \mathbf{b}^0 indicates that the regression coefficients for this solution are not unique (there are an infinite number of solutions). And this is a problem for which we have no solution, as Searle (1971, p. 160) states: "An investigator having data to be analyzed will clearly have no use for any \mathbf{b}^{0} [the non-unique solutions to the set of normal equations] as it stands, whatever its numerical value. But what about linear functions of the elements of \mathbf{b}^{0} ?"¹¹ This last sentence is somewhat cryptic. What Searle is suggesting is that linear functions of these non-unique estimates may be quite useful. For example, we can use a linear function of these coefficients multiplied by values on the independent variables to provide us with predicted values of the dependent variable. Each set of estimates of the regression coefficients using a different generalized inverse will differ-but each set will yield the same predicted values for the dependent variable. The estimated value of the dependent variable is in this situation estimable, while the individual regression coefficients are not.

In the present situation, the estimated values of the dependent variable as predicted by the age group, period, and cohort dummy variables are estimable—they are least square estimates. (This is also the case when we add some selected interaction terms to the model). The same estimates of these dependent variable scores will result from any of the sets of solutions for the regression coefficients provided by any generalized inverse for **X'X**. Searle (1971, p. 181) notes that in our situation: "(*i*) the expected value of any observation is estimable." Thus, even though the vector \mathbf{b}^0 is not unique, each \mathbf{b}^0 vector resulting from any generalized inverse produces the same estimates of the expected values of the observations, and they are best linear unbiased estimates.

¹⁰ See Searle (1971) for a comprehensive discussion of generalized inverses and estimable functions. A briefer, but adequate, discussion for the presentation in this paper is found in Rawlings (1988).

¹¹ Yang et al. (2004) argues that a particular generalized inverse is the best one to use, and this provides a rationale for a particular \mathbf{b}^0 vector being superior to other such vectors. If they are correct, that justifies the use of the individual regression coefficients in an APC model as the best estimates. For a discussion of the strengths and weaknesses of this approach see (Smith 2004). The procedure we use does not suggest that one \mathbf{b}^0 is superior to another.

We use these estimates of the expected values of the observed rates in many ways. We show graphically the extent to which cohort replacement affects the age distributions of homicide offending even with age and period effects already accounted for. Using partitioning of the variation, we show that the epidemic of youth homicide is only partially accounted for by cohort replacement. We can accomplish these tasks because the least squares *predicted values* (based on models containing, age, period, and cohort dummy variables) can be estimated using the generalized inverse.¹²

Since we can estimate the variation associated with the age and period dummy variables, the age and cohort dummy variables, and the period and cohort dummy variables by running separate regressions with each of these two sets of dummy variables in the equation, we can assess the variation that is uniquely associated with cohorts, periods, or age groups. For example, the variation uniquely associated with cohorts equals the variation associated with age, period, and cohort dummy variables minus that associated with age and period in a regression that includes only those dummy variables. SAS (2004) handles all of this in a single step by producing Type III sums of squares: these are the unique sums of squares (the additional sums of squares added to the model when the set of dummy variables is entered last) and providing estimates of predicted values using the variables in the model. Although SAS provides these sums of squares associated uniquely with each of the sets of dummy variables, its output clearly labels the parameter estimates associated with the individual dummy variables in the full model as non-unique.¹³

Data and Measures

Homicide Arrests

Ideally we would have available an accurate record of homicide offenses by age group over an extended period of time. Unfortunately that is not the case—the Uniform Crime Reports provide the number of homicides known to the police—but not the age of the offenders. Vital statistics provide the ages of homicide victims—but not data on offenders. We rely on homicide arrest data from the Uniform Crime Reports, which are broken down by age. These data are not ideal, but are the best available for the questions addressed in this paper. A major concern is that there have been shifts in the clearance rates for homicides over time. We provide some protection for this problem by including dummy variables for periods in the analysis. To the extent that shifts in the clearance rates are similar across age groups, we control for the effect of clearance rates with the period dummy variables. Importantly, we would not expect shifts in clearance rates by age to correspond to the patterns taken by cohort replacement effects. We, like most researchers in this situation, use Uniform Crime Report arrest data in our analyses.

¹² A generalized inverse is a default in PROC GLM in SAS (2004) for situations in which the X'X matrix is singular.

¹³ In SAS we used the program below to provide the estimates of the age-period-specific homicide offense rates based on age, period, and cohort dummy variables and age by period interactions:

PROC GLM; CLASS cohort agegroup period;

MODEL lnhom1564 = period agegroup cohort age1*period90 age1*period95 age2*period90 age2*period95/SOLUTION PREDICTED; RUN;

The Uniform Crime Reports began collecting data in the 1930s, but the coverage was spotty to begin with. By 1960 arrest data are available for the entire country broken down by single years of age for those 15-to-24 and for 5 year age groups for those 25-to-29, 30-to-34, through those 45-to-49. By 1965 the UCR provided data on those arrested broken down by 5 year age groups to age 60-to-64. We use this data (FBI, 1961, 1966, ..., 2006), which in our analysis will mean the loss of one period (1960) but will add three age groups (50-to-54, 55-to-59, and 60-to-64). We lose seven age-period-specific observations from 1960, but we gain 27 age-period-specific observations for the nine periods from 1965, 1970,..., 2005 for a total of 20 more age-period-specific observations.

The UCR system is voluntary, although the vast majority of law enforcement agencies report homicide arrests broken down by age to the FBI, the percentage of the total United States population covered by the agencies reporting for the periods covered in this paper has varied from 65% in 2000 to 92% in 1980. We correct for this underreporting by dividing the total population of the United States by the number of residents in the areas reporting to the FBI that year. This ratio is multiplied by the number of homicide arrests in each of the age groups. This corrected number is then divided by the number of U.S. residents in a particular age group and multiplied by 100,000 to obtain an estimate of the number of homicide arrests in that age group per 100,000 residents. We obtained data on the number of residents in each age group from the Current Population Surveys: Series P-25 (US Bureau of the Census, 1995, 2000, 2005, a).

We logged the homicide arrest rates (natural log) for two reasons. These age-periodspecific rates exhibited a strong positive skew that was substantially reduced when we logged the variable. But perhaps the most important reason was that we are just as interested in changes in the homicide arrest rates for those age groups with low rates as for age groups with relatively high rates. This transformation helps insure that a doubling of the rates for those 60-to-64 is treated as the same degree of change as a doubling or the rates for those 15-to-19 in our models.

Cohort Characteristics

The analytic technique described above is an important innovation. It can determine whether there is a unique cohort effect (or age or period effect) as defined by the sets of dummy variables for age, period, and cohort. It does this while controlling for the other sets of dummy variables. This is descriptively crucial for criminologists for answering questions such as: Do cohorts differ in their propensity to homicide offending? Can cohort replacement account for the upturn in youth homicides during the latter part of the 1980s through the mid 1990s? The approach outlined above allows us to answer those questions.

A different question is why cohorts might differ in their rates of homicide offending. To determine why the cohorts might differ, we test whether the two characteristics of the cohorts used by O'Brien et al. (1999) can account for some of the cohort variation that occurs in our data. Those two cohort characteristics are relative cohort size and the percentage of non-marital births. We measure relative cohort size (RCS) as the percentage of the resident population age 15-to-64 that is in the cohort when the cohort is age 15-to-19. It is a measure of the relative size of the cohort to the cohorts older than it including the parental generation. As noted earlier, we can conceptualize this variable as an indicator of baby boom and baby bust generations. Theoretically, relatively large cohorts have fewer parents per child, larger class sizes, and more youth (for peer associations) relative to adults (for cross generational associations). When they are young adults these cohort members have fewer entry level jobs per job applicant and (in part) because of this they

delay marriage and child bearing. These factors are likely to lead to a generation that is more prone to committing homicides controlling for changes associated with period and age. The data for computing RCS were obtained from the Current Population Surveys: Series P-25 (US Bureau of the Census, 1995, 2000, 2005, a).

The second cohort characteristic is the percentage of live births that were to nonmarried women during the years in which the cohort members were born.¹⁴ Such families are more likely to be single parent families throughout much of the childhood of the cohort members, have less monitoring and supervision of children, and children in such families are more likely to grow up in poverty. Data for the percentage of non-marital births (births to unwed mothers per 100 live births) were obtained from Vital Statistics of the United States (US Bureau of the Census, 1946, 1990, b; National Center for Health Statistics, 2003).

In our examination, each cohort is assigned a value on each of the two cohort characteristics and the question is the extent to which these variables can explain the variation that is unique to cohorts that we observe in our model. If these characteristics are associated with much of the unique variance due to cohorts, we have a possible mechanism for these differences across cohorts. But again we emphasize that by focusing on the predicted value of the dependent variable, which is an estimable functions, we can ascertain whether cohorts are important without using these cohort characteristics. They are used only to see if we can provide an explicit explanation for the cohort replacement effects.

Results

We now apply the approach outlined in the last two sections of this paper to the questions outlined in the first two sections. Table 1 presents the results from an analysis in which we have ordered the entry of the dummy variable sets for periods, age groups, cohorts, and then the interactions between age and period. We justify this order of entry below, but not too adamantly, since it is just one of several possible orderings. Table 2 will take the more conservative approach of presenting the results when each of these interactions and sets of dummy variables are entered into the equation last (indicating the variance uniquely associated with that factor controlling for all of the other variables in the equation). We will present the rationale for the order of entry of the variables in Table 1 after we describe the results in Table 1.

Table 1 reports the degrees of freedom for each of the sets of dummy variables when it enters the equation and the Type I sums of squares associated with its entry. These are the sums of squares associated with the set over and above the sums of squares accounted for by the variables already in the equation. The total sum of squares (squared deviations from the mean) for the age-period-specific homicide rates is 76.711. When the period dummy variables are added to the equation they account for 11.7% [=(8.943/76.711) × 100] of the variability in these rates. When we add the age dummy variables to the analysis they account for an additional 83.0% [=(63.648/76.711) × 100] of the variability (this, of



🙆 Springer

¹⁴ Savolainen (2000) analyzed Public Use Micro Samples (PUMS) census data to estimate the percentage of those in 5-year birth cohorts who lived in single parent families when they were ages 5-to-9. These data were not available for all years or for all 10 year census periods, so he used interpolation for several of his estimates. Still, the percentage of non-marital births and the percentage of cohort members in single parent families at ages 5-to-9 are highly correlated. Using data based on cohorts born from the 1915–1919 to 1975–1979, the correlation between these two measures is 0.98. For the first differences of the measures, the correlation is 0.90. The data on single parent families were supplied by Jukka Savolainen.

Source	df	Type I SS	Mean square	F value		
Period	8	8.943	1.118	127.450	0.0001	
Age group	9	63.648	7.072	806.270	0.0001	
Cohort	16	3.175	0.198	22.620	0.0001	
Total for interactions	4	0.489	0.122	13.940	0.0100	
Total	89	76.711	-	-	-	
Model	37	76.255	2.061	234.970	0.0001	
Residual	52	0.456	0.009	-	-	

Table 1 Sums of squares associated with period, age groups, cohorts, and four interactions of age and period when entered in the order specified below

Table 2Sums of squares uniquely associated with period, age groups, cohorts, and four interactions of ageand period

Source	df	(Type III SS)	Mean square	F value	<i>p</i> <
Period	7	3.900	0.557	63.520	0.0001
Age group	8	3.564	0.446	56.800	0.0001
Cohort	16	2.263	0.141	16.120	0.0001
Age (15–19) × Per (1990)	1	0.293	0.293	33.430	0.0001
Age (20–24) × Per (1990)	1	0.209	0.209	23.850	0.0001
Age (15–19) × Per (1995)	1	0.053	0.053	6.030	0.0500
Age (20–24) × Per (1995)	1	0.113	0.113	12.940	0.0010
Total for interactions	4	0.489	0.122	13.940	0.0100
Residual	52	0.456	0.009	-	-

course, is with controls for the sums of squares accounted for by the period dummy variables). Controlling for the effects of period and age, the cohort dummy variables account for 4.1% [=(3.175/76.711) × 100] of the sums of squares. Note that each of these sets of dummy variables account for a statistically significant amount of variability (p < 0.0001) when entered into the equation. Each of these factors: period, age, and cohort improve the fit of the model when added to the factors that precede it. When we enter the four interaction terms between age and period into the equation to account for any specific variability due to the two youngest age groups in 1990 and 1995 (the years of the epidemic of youth homicide), we find that they account for 0.64% [=(0.489/76.711) × 100]] of the sums of squares, which is a statistically significant amount of the variability (p < 0.01). There is a statistically significant upturn in youth homicide offending during these two periods that is not fully accounted for by a model containing age, period, and cohort dummy variables.

The rationale for the ordering of the entry of the variables in Table 1 is that our focus is on the age distribution of homicide and the factors that change the shape of that distribution. The period dummy variables do not affect the shape of the age curve for homicides (they proportionately increase or decrease the age-specific rates during different periods) and should therefore be controlled before examining the factors that contribute to changes in the age curve. The age curve and its degree of stability is one focus of our concern. We enter age into the equation next to estimate the amount of variation that is associated with age. This gives an advantage to Hirschi and Gottfredson's (1983) perspective, since it gives age credit for the maximum amount of variability that it can receive (even credit for variability that it jointly accounts for with cohorts and the interaction terms). Even so, the percentage of variability in the age-period-specific rates of homicide offending associated with the age dummy variables (83.0 percent) is quite impressive. This is the solid age curve that Hirschi and Gottfredson (1983) refer to and label as invariant. We already know, given the epidemic of youth homicide, that this curve is not rigidly invariant. Hirschi and Gottfredson (1983) do not claim that it is strictly invariant. We focus on the factors that change the shape of this curve. Are changes in the shape of this age curve over time related to cohort replacement? The cohort dummy variables represent this replacement effect and when they are added to the equation they account for an additional 4.1% of the variability. But even these cohort effects do not account entirely for the strong upturn in homicide offenses for the two youngest age groups in 1990 and 1995 during the epidemic of youth homicides. These four interactions account for 0.64% of the variance and are statistically significant.

These additional amounts of variability associated with cohorts and the age by period interactions may seem trivial, but just how important they are can be seen in the graphs presented later. The age, period, and cohort dummy variables and the four age by period interactions account for 99.4% [= $(76.26/7,671) \times 100$] of the variability in the age-periodspecific homicide offense rates. Even when we correct this R-square value for the number of independent variables used in the analysis its value is 98.7%.

Table 2 provides a different perspective on these results by focusing on the unique sums of squares associated with each of these factors. That is, when each set of dummy variables enter the equation last, how much additional variance is associated with the set (Type III sums of squares). The degrees of freedom differ slightly from those in Table 1, because when period and age groups are entered into the equation last they lose an extra degree of freedom because of the linear dependency when all three sets of dummy variables are in the equation at the same time. Here we see a much more equal "impact" of each of these factors (with the dummy variables for period variables accounting for the largest share of the "unique" variability and cohorts being almost as important as age groups). In this table, we present the interactions both separately and their total "impact" when all four are entered into the equation simultaneously. Again, we note that the period set of dummy variables, the age group set of dummy variables, and the cohort set of dummy variables each account for a statistically significant amount of variation when entered last into the equation (p < 0.0001). Each of the interactions accounts for a statistically significant amount of variation, when entered last into the equation (both separately and when they enter as a group).

Even though the individual coefficients for the age, period, and cohort dummy variables are non-estimable, these findings tell us that there are unique effects of each of these sets of dummy variables. To see more specifically how each of these factors contributes to changes in the age curve of homicide offenses, we show on the same graphs: (1) the observed homicide offense age curve; (2) the age curve based on the age and period effects, which shows the average age curve at different periods¹⁵; and (3) the age curve based on the age and period and cohort effects, which shows how cohort replacement shifts the age curve in different periods. We do not graph the interactions between age and period for the two youngest age groups in 1990 and 1995, but their effects in those 2 years are clear and will be mentioned.

92

¹⁵ The shape of the age curve is proportionately the same—invariant—at each period.

The first set of graphs (Fig. 2) consists completely of periods (1965, 1970, and 1975) before the epidemic of youth homicide. Figure 2 shows how well the strictly invariant age curve and the curve that takes cohort replacement into account track the observed age curve for homicide offenses in the earlier periods for which we have data. Importantly the cohort replacement effects is not based on a proxy (cohort characteristics), but on the full effect of cohorts as measured by a dummy variable for each cohort. The strictly invariant age curve (fitted using the age-period model) has the same general shape as the observed age curve, but does not fit the observed data well for some of the age groups. But how well



D Springer

93

this age curve fits the observed data is relative. It is certainly better than many other curves that might be fit to the data and has the general shape discussed by Hirschi and Gottfredson (1983). These graphs shows that the age curve is not rigidly invariant and that cohort replacement can account for much of this lack of invariance from 1965 to 1975. During this period there is only one statistically significant (p = 0.03) discrepancy between the predicted and observed age-period-specific homicide rates when cohort replacement is taken into account and that is for the youngest age group in 1965.¹⁶

Figure 3 shows the same three age curves for the two periods that preceded the epidemic of youth homicide (1980 and 1985) and for 1990 a period that clearly reflects the epidemic of youth homicide. Note that in each of the figures (1980–1990) the age curves (based on age and period) are the same (proportionately).¹⁷ In 1985 the observed age distribution of homicide offenses is well modeled by the strictly invariant age curve. During the epidemic of youth homicide, which is typically seen as occurring during the last half of the 1980s and through the middle 1990s, we see that the strictly invariant age curve does not fit the observed data well, especially for the two youngest age groups. Most important, for the crack cocaine thesis, the model that includes age, period, and cohorts does not account for the entire sharp upward trend in youth homicide that occurred between 1985 and 1990. It does account for part of this upturn. In 1990 it accounts for 46 percent of the discrepancy for those 15–19, but none of the discrepancy for those who are 20–24.

Of course, when we add the dummy variables for those 15–19 and 20–24 in 1990 they account for all of this variance (they fit the observed data perfectly) for these two age groups. Our statistical tests presented earlier (in Table 2) show that this improvement in fit is statistically significant (not likely just a chance fluctuation). Thus, even taking cohort replacement into consideration, the rates for 15-to-19 and 20-to-24 year olds in 1990 and, we should add 1995, have significantly higher homicide rates than predicted by age, period, and cohort replacement. This is consistent with the crack cocaine thesis.

Figure 4 presents the age curves for the final period (1995) of the epidemic of youth homicide in our data and the two subsequent periods. Again the age period model does not capture the upturn in youth homicide in 1995 and the age-period-cohort model does not do so completely. In 1995 the model that contains the cohort effects accounts for 60 percent of

🖉 Springer

¹⁶ We tested this by entering age by period interactions for the largest discrepancies for each age in each period and assessing whether the interaction significantly improved the fit of the model that contained the age, period, and cohort dummy variables and the four interactions associated with the epidemic of youth homicide. We report the results of this procedure for each of the discrepancies—keeping the interactions that are statistically significant (p < 0.05) in the equation and searching for others to add. We began in the earliest period and proceeded to the most recent. If an interaction became insignificant (p > 0.05) in our model, we eliminated it. This is an informal way of testing for the most important discrepancies. We note that this procedure probably is too generous in including interaction terms. Since there are 86 such significant tests for the interactions (excluding the four associated with the epidemic of youth homicide), the Bonferroni corrections suggests that we use an alpha value of 0.0006.

¹⁷ This estimate for the log of the homicide rate for those 15-to-19 in 1985 is the log of the sum of the intercept—which is the same across all of the periods—plus the age effect for those from 15-to-19—which is the same across all of the periods, and the period effect. It is only the period effect that varies from period to period which increases or decreases the absolute value of the predicted age-period-specific homicide rate for that age group for that period. The same can be said for each of the age groups in any period. If we convert this logged age-period-specific rate predicted from the age-period model to the "raw" age-period-specific rate for say the youngest age group in the third period (1975) (see Fig. 1) by exponentiating it, we have $e^{(u + al + p^3)}$, which we can also write as $e^{u} \times e^{al} \times e^{p^3}$. Clearly this shift is proportional since e^{p^3} is constant for all the age-period-specific rates in period 3 and its effect is multiplicative.



Fig. 3 Age distributions for homicide offenses observed, predicted by age and period, and predicted by age, period and cohort: 1980–1990

the discrepancy between the standard age curve (based on period and age) and the observed age curve for those 15–19 and 74 percent of the discrepancy for those who are 20–24.

By the year 2000 the cohort replacement model is quite consistent with the observed age curve and the strictly invariant age curve does not fit well at both the youngest and oldest age levels. But the age curve for 2005 has interesting deviations for ages 25–29 and 30–34,

🙆 Springer



Fig. 4 Age distributions for homicide offenses observed, predicted by age and period, and predicted by age, period and cohort: 1995–2005

being lower than expected using the cohort replacement model. The deviation is statistically significant, however, only for those age 25–29 (p = 0.01). Could this be a "mini" cohort effect caused by those who went through the epidemic of youth homicide 10 years earlier as 15–19 and 20–24 year olds? If it is, there was no indication of such an effect in

D Springer

2000. The only other significant deviation from the observed age-period-specific rates and those predicted by cohort replacement in Fig. 4 is for those 55–59 in 2000.

Overall, the results are consistent with the crack cocaine thesis, which suggests that during this period the crack cocaine markets became regularized, with less need for violence, and led to a return to more typical rates of youth homicide (Baumer et al. 1998; Blumstein and Rosenfeld 1998; Steffensmeier and Harer 1999). Whether there is a mini cohort effect for cohorts who went through the epidemic of youth homicide as youths will not be known until data not yet generated are available. We also note that the discrepancy between the predicted homicide rates using cohort replacement and the observed rates are substantially greater for those 15–19 year old than for those 20–24 both in 1990 and in 1995. This suggests that the impact of the epidemic was far greater for the youngest age group. This is also consistent with the crack cocaine thesis that emphasizes the participation of youth in this market and that their participation was, in part, because of more lenient sentences for those who were minors.

The results indicate that there is (not surprisingly) a strong relationship between age and homicide rates in the United States from 1965 to 2005 and the general shape of this distribution is that predicted by Hirschi and Gottfredson (1983). There is not, however, a strictly invariant age curve during this time period. The fit of the predicted age curve to the observed age curve when cohort replacement is taken into account is fairly remarkable over most of the period covered. The exception is for youth violence during the 1990 and 1995 periods. These are the years of the epidemic of youth homicide that are associated with the crack cocaine thesis. Although cohort replacement is associated with almost half of the discrepancy between the standard age curve and the observed age curve for these two age groups during these two periods, there remains a significant discrepancy for each of these age groups in each of these periods between the predictions with cohort replacement included and the observed rates. This is consistent with the crack cocaine thesis, which predicts just such a discrepancy.

Our final analysis explores whether the cohort explanation suggested by O'Brien et al. (1999) can provide an explanation for the cohort replacement effects found in our analyses. These authors suggested that relatively large size cohorts and cohorts that have a higher percentage of non-marital births are more prone to homicide offending. We described the measurement of these two variables earlier. To examine how well these two variables (cohort characteristics) can explain the cohort replacement effects found in our analyses, we entered them into the equation, rather than cohort dummy variables, to represent the effects of cohorts. These models are of full column rank, even though they contain the age and period dummy variables and the cohort characteristics, since the values of relative cohort size and the percentage of non-marital births are not linear functions of age and period.

We use these cohort related variables to see if they can account for the variation uniquely associated with cohorts in our model. The process is straightforward. We run a regression (model 1) with age and period dummy variables only and find that the sums of squares associated with these dummy variables is 70.49. We add the cohort dummy variables to this model and find that the sums of squares associated with this model (model 2) is 72.17. We remove the dummy variables associated with cohorts from the model and replace them with the two cohort characteristics and find that the sums of squares associated with this model (model 3) is 71.98. Since we used dummy variable to code the cohorts in the second model, the increase in the sums of squares from model 1 is the maximum amount of variability that is uniquely associated with cohorts: 1.68 (=72.17 - 70.49). Since the cohort characteristics have constant values within cohorts, they represent

🖉 Springer

another way (not optimal) of coding cohorts and the amount of variability associated uniquely with the two cohort characteristics is 1.49 (=71.98 - 70.49). The two cohort characteristics account for $89\% [=(1.49/1.68) \times 100]$ of the sums of squares that are unique to cohorts.¹⁸ Both of the cohort characteristics are statistically significant and substantively strongly related to the age-period-specific homicide rates. Since both the dependent variables and the cohort characteristics are logged, the regression coefficients represent the effect of a one percent change in the independent variable in terms of a percent change in the dependent variable. For the logged relative cohort size, the coefficient is 1.11 and for the logged percent of non-marital births the coefficient is 1.34. Thus a one percent change in the relative cohort size (controlling for all of the other variables in the model) is associated with a 1.11 percent change in the age-period-specific homicide rate. A one percent change in the percentage of non-marital births is associated with a 1.34 percent change in the age-period-specific homicide rate.

Accounting for Changes in the Age Distribution of Homicide

This analysis focuses on changes in the age distribution of homicide offending. It does not address the average level of homicide offending in the United States which for our periods has ranged from 5.1 per 100,000 in 1965 to 10.2 per 100,000 in 1980. These changing levels are modeled using the period dummy variables. We also use dummy variables to estimate an average age curve over the nine periods that we examine. We use the age-period model as our baseline. We do not try to explain the why the age effects are the way they are [Hirschi and Gottfredson (1983) would suggest that this is not the role of the social scientist] nor why the specific period effects are the way they are. The task we have set for ourselves is to account for shifts in the age distribution over time. We assess how well we have accomplished this task.

To make this assessment, we must bear in mind that we are trying to explain shifts in the age distribution of homicide offenses over time. The baseline model for such an attempt is the model that contains age and period. This model provides the same shaped age distribution in each of the periods (the age rates from one period to another differ by only a constant of proportionality based on the period dummy variables). Turning to Table 1 we see that the total sum of squares for the age-period-specific homicide offending rates is 76.711 and the sum of squares accounted for by age and period is 72.591. As noted earlier these two sets of dummy variables account for nearly 95% of the variability in these rates. Our task, however, is to account for shifts in the shape of the age distribution from period to period. The variability that we are interested in explaining is associated with the ageperiod-specific rates around this standard age curve. This sums of squares equals 4.12 (=76.711 - 72.591) and it represents the sum of squared deviations of the observed ageperiod-specific rates from this age curve in different periods. Again turning to Table 1 we find that cohort replacement is associated with 77% [=(3.175/4.12) \times 100] of this variability of the observed age-period-specific rates from the standard age curve. The four interaction terms for the two youngest age groups in 1990 and 1995 are associated with

¹⁸ These sums of squares differ from those in Tables 1 and 2, because we do not have data on the percentage of nonmarital births for the cohorts born from 1910 to 1914 and before. Thus, this data is missing for those age 50–54 and above in 1965; 55–59 and above in 1970; and 60–64 in 1975. The sums of squares reported in this paragraph are based on analyses in which these six cases are not included in any of the analyses.

12% [= $(0.489/4.12) \times 100$] of this variation. Cohort replacement and the interactions for the epidemic of youth homicide account for 89 percent of the variability of the age-period-specific homicide rates from the invariant age distribution as represented by our average age curve.

The other assessment is more informal. It involves comparing the standard age curve and its fit to the data relative to the fit of a model that contains the cohort dummy variables. The question is does the fit appear to be significantly (in a substantive sense) better. The same assessment can be made for the epidemic of youth homicide. Was something significant (again in a substantive sense) going on in addition cohort replacement?

Discussion

This paper focuses on the age distribution of homicide offending in the United States over nearly half a century. While overall rates of homicide have varied by nearly a factor of two over this time span our focus is on the relative rates between age groups within each of the periods: the shape of the age distribution. A strict invariance thesis would maintain that the age distribution is absolutely invariant except for random shocks. Our data indicate that this is not the case—even though there is on average a relatively stable age curve with a peak at the 20-to-24 age group and a decline thereafter. This general pattern in terms of the periods we investigated (1965, 1970,..., 2005) except for 1990 and 1995. This general pattern is consistent with Hirschi and Gottfredson (1983) who noted that "social conditions" at times affect the age curve.

Even though the peak ages of homicide offending are relatively stable, we find important shifts in age distribution throughout the period 1965–2005 that are associated with cohort replacement. Figures 2–4 show the importance of cohort replacement: models containing age, period, and cohort dummy variables fit the observed age curve substantially better than those that contain only age and period. The impact of cohort replacement is also reflected in the statistically significant increase in the sums of squares accounted for in the age-period-specific rates when the cohort dummy variables are entered into models that control for age and period effects. The cohort replacement thesis finds strong support in these data throughout the years 1965–2005. During the epidemic of youth homicide (1990 and 1995) there are statistically significant deviations between the cohort replacement model predictions and the observed homicide rates for the youngest two age groups, but in those two periods it fits better than the model representing the strictly invariant age curve.

While cohort replacement is associated with a substantial amount of the upturn in the homicide offending of youth during 1990 and especially 1995, it does not account completely for this upturn. The graphs indicate that it accounts for slightly less than half of this upturn. There is certainly a strong deviation from the cohort replacement model during the two periods of the epidemic of youth homicide and this deviation is in the direction suggested by the crack cocaine thesis. These are the periods in which researchers suggest that crack cocaine markets were opening up with their associated arming of sellers and then the associated arming of other young people in response. The same researchers note that by the late 1990s these markets had stabilized: as Fryer et al. (2005, p. 7) note in their paper measuring the general impact of crack cocaine, "by the year 2000, we observe little impact of crack." We see this same return to normalcy in Fig. 4 for the year 2000. Here the model predicting the observed age-period-specific homicide rates on the basis of cohort replacement fits the data very well. This is what we would expect if the crack cocaine

🖉 Springer

market had regularized by the end of the 1990s. The 2005 data raise the question of whether there is a continuing impact on the homicide offending rates of those who were youth during the epidemic of youth homicide. It will be interesting to see if this is the case as data become available in the future.

The estimable function approach has allowed us to fully model the cohort replacement effects. This cannot be done by using cohort characteristics, although the characteristic that we use seem to do a good job of modeling cohort effects. Combining this approach with the graphs that we use, allows us to see the impact of cohort replacement and the additional potential impact of the crack cocaine epidemic. The estimable function approach allows us to enter dummy variables for cohorts to a model that contains dummy variable for age groups and periods and produce estimates of the best predicted values of the age-periodspecific homicide rates based on age, period, and cohorts. This allows us to see how well cohort replacement can account for shifts in the age distribution of homicides. Without such a technique one could argue that the model had not fully accounted for the effects of cohorts on homicide rates. For these data, it appears that two cohort characteristics (relatively cohort size and the percentage of non-marital births) account for much of the variability of the age-period-specific homicide rates that are associated uniquely with cohorts. These cohort characteristics provide a potential explanation of the cohort replacement effects found in this research in much the same way that the crack cocaine thesis provides an explanation for the discrepancies found for young homicide offenders' rates in the two periods associated with the epidemic of youth homicide.

References

- Baumer E, Lauritsen JL, Rosenfeld R, Wright R (1998) The influence of crack cocaine on robbery, burglary, and homicide rates: a cross-city, longitudinal analysis. J Res Crim Deling 35:316–340
- Blau JR, Blau PM (1982) The cost of inequality: metropolitan structure and violent crime. Am Sociol Rev 47:114-129

Blumstein A (1995) Youth violence, guns, and the illicit drug industry. J Crim Law Criminol 86:10-36

Blumstein A, Cork D (1996) Linking gun availability to youth gun violence. Law Contemp Probl 59:5–24 Blumstein A, Rosenfeld R (1998) Explaining recent trends in U.S. homicide rates. J Crim Law Criminol

88:1175-1216

Britt CL (1992) Constancy and change in the US age distribution of crime: a test of the 'invariance hypothesis'. J Quant Criminol 8:175-187

Cook PJ, Laub JH (1986) The (surprising) stability of youth crime rates. J Quant Criminol 2:265-276

- Cook PJ, Laub JH (1998) The unprecedented epidemic in youth violence. In: Tonry M (ed) Crime and justice: a review of research, vol 24. University of Chicago Press, Chicago, pp 27-64
- Cook PJ, Laub JH (2002) After the epidemic: recent trends in youth violence in the United States. In: Tonry M (ed) Crime and justice: a review of research, vol 28. University of Chicago Press, Chicago
- Cork D (1999) Examining space-time interaction in city-level homicide data: crack markets and the diffusion of guns among youth. J Quant Criminol 15:379-406
- Easterlin RA (1978) What will 1984 be like? Socioeconomic implications of recent twists in the age structure. Demography 15:397-421
- Easterlin RA (1987) Birth and fortune: the impact of numbers on personal welfare. University of Chicago Press, Chicago
- FBI (1961, 1966,...,2006). Crime in the United States (1960, 1965,...,2005). US Government Printing Office, Washington DC
- Fryer RG, Heaton PS, Levitt SD, Murphy KM (2005) Measuring the impact of crack cocaine. NBER Work Pap Ser 11218:1–65
- Gove WR (1994) Why we do what we do: a biopsychosocial theory of human motivation. Social Forces 73:363-394
- Greenberg DF (1977) Delinquency and the age structure of society. Contemp Crises 1:189-223
- Greenberg DF (1985) Age, crime, and social explanation. Am J Sociol 91:1-27

D Springer

- Greenberg DF (1994) The historical variability of the age-crime relationship. J Quant Criminol 10:361–373 Hirschi T, Gottfredson M (1983) Age and the explanation of crime. Am J Sociol 89:552–584
- Hirschi T, Gottfredson M (1985) Age and crime, logic and scholarship: comment on Greenberg. Am J Sociol 91:22–27
- Huff-Corzine L, Corzine J, Moore DC (1986) Southern exposure: deciphering the South's influence on homicide rates. Social Forces 64:906–924
- Jacobs D, Helms R (1996) Toward a political model of incarceration: a time-series examination of multiple explanations for prison admission rates. Am J Sociol 102:323–357
- Mason KO, Mason WH, Winsborough HH, Poole WK (1973) Some methodological issues in cohort analysis of archival data. Am Sociol Rev 38:242–258
- McLanahan S, Sandefur G (1994) Growing up with a single parent, what hurts, what helps. Harvard University Press, Boston
- Messner SF (1983) Regional and racial effects on the urban homicide rate: the subculture of violence revisited. Am J Sociol 88:997-1007
- Messner SF, Sampson RJ (1991) The sex ratio, family disruption, and rates of violent crime: the paradox of demographic structure. Social Forces 69:693-713
- National Center for Health Statistics (2003) Vital statistics of the United States: natality. US Government Printing Office. Available at: http://www.cdc.gov/nchs/data/statab/natfinal2003.annvol1_17

O'Brien RM (2000) Age period cohort characteristic models. Soc Sci Res 29:123-139

- O'Brien RM, Stockard J, Isaacson L (1999) The enduring effects of cohort characteristics on age-specific homicide rates, 1960–1995. Am J Sociol 104:1061–1095
- O'Hare WP (1996) A new look at poverty in america. Popul Bull 51(2):48. Population Reference Bureau, Washington DC
- Rawlings JO (1988) Applied regression analysis: a research tool. Wadsworth & Brooks/Cole, Pacific Groove, CA
- Robertson C, Gandini S, Boyle P (1999) Age-period-cohort models: a comparative study of available methodologies. J Clin Epidemiol 52:569–583
- Sampson RJ (1985) Race and criminal violence: a demographically disaggregated analysis of urban homicide. Crime and Delinq 31:47-82
- Sampson RJ (1986) Effects of socioeconomic context on official reaction to juvenile delinquency. Am Sociol Rev 51:876-886
- Sampson RJ (1987) Urban black violence: the effects of male joblessness and family disruption. Am J Sociol 93:348-382
- Sampson RJ, Wilson WJ (1995) Toward a theory of race, crime, and urban inequality. In: Hagan J, Peterson RD (eds) Crime and inequality. Stanford University Press, Stanford California, pp 37–54

SAS (2004) SAS/GLM: Version 8. SAS institute incorporated, Cary NC

- Savolainen J (2000) Relative cohort size and age-specific arrest rates: a conditional interpretation of the Easterlin effect. Criminology 38:117–136
- Scheffé H (1959) The analysis of variance. Wiley & Sons, New York
- Searle SR (1971) Linear models. Wiley & Sons, New York
- Smith HL (2004) Response: cohort analysis redux. In: Stolzenberg RM (ed) Sociological methodology 2004. American Sociological Association, New York, pp 111–119
- Steffensmeier DJ, Allan EM, Harer MD, Streifel C (1989) Age and the distribution of crime. Am J Sociol 94:803-831
- Steffensmeier D, Harer MD (1999) Making sense of recent U.S. crime trends, 1980 to 1996/1998: age composition and other explanations. J Res Crime Deling 36:235-274
- US Bureau of the Census (1995, 2000, 2005 a) Current population surveys: series P-25. US Government Printing Office, Washington DC (Numbers 98, 114, 170, 519, 870, 1000, 1022, 1058, 1127, for 1995, 2000, 2005 data). http://www.census.gov/popest/states/tables/NST2007-01.xls and www.census.gov/ popest/archives/EST90INTERCENSAL/US-EST90INT-dataset.html
- US Bureau of the Census (1946, 1990, b) Vital statistics of the United States: natality. US Government Printing Office, Washington DC
- Vandeschrick C (2001) The Lexis diagram, a misnomer. Demogr Res 4:97-124
- Williams KR (1984) Economic sources of homicide: re-estimating the effects of poverty and inequality. Am Sociol Rev 49:283–289
- Williams KR, Flewelling RL (1988) The social production of criminal homicide: a comparative study of disaggregated rates in American cities. Am Sociol Rev 53:421-431
- Yang Y, Fu WJ, Land KC (2004) A methodological comparison of age-period-cohort models: the intrinsic estimator and conventional generalized linear models. In: Stolzenberg RM (ed) Sociological methodology 2004. American Sociological Association, New York, pp 75–110