NBER WORKING PAPER SERIES

ON THE NATURE AND ESTIMATION OF AGE, PERIOD, AND COHORT EFFECTS IN DEMOGRAPHIC DATA

David E. Bloom

Working Paper No. 1700

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1985

This research was supported by NICHD Contract No. N01-HD9-12824. The author is indebted to Said Haidar and David Neumark for helpful assistance in the preparation of this paper and to Mark Watson for helpful comments. The research reported here is part of the NBER's research program in Labor Studies. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #1700
September 1985

On the Nature and Estimation of Age, Period, and Cohort Effects in Demographic Data

ABSTRACT

This paper develops a general procedure for estimating age, period, and cohort effects in demographic data. The procedure involves structuring, mathematically, the effect of cross-cohort changes in the timing and level of a vital event on period rates of occurrence of the event. The procedure is illustrated and tested in an application to data on the first birth rates of American women. Overall, the empirical results provide support for the procedure. The results also provide evidence that period effects are highly age-specific and that the size of cohort effects may be substantially overestimated by models which fail to allow for the age specificity of period effects.

David E. Bloom
Department of Economics
Littauer Center
Harvard University
Cambridge, MA 02138

I. Introduction

Identifying age, period, and cohort patterns in demographic data is perhaps the most time-worn activity of demographers.[1] These patterns are interesting primarily because they reveal aspects of the underlying physiological and behavioral mechanisms which lead to the occurrence of vital events. For example, age is a proxy for the state of an individual's physiology which itself directly influences variables such as fertility and mortality. Age can also be regarded as a proxy for exposure to social forces which may influence other variables such as nuptiality and divorce. Period, on the other hand, is a proxy for changes in a vast set of conditions that can also influence the occurrence of demographic events. Among this set of conditions we may include the state of the economy, the existence of war, factors related to the legal environment in a society, and technological developments such as those related to improved medical care and contraception. Finally, cohort is a proxy for variables whose influence on demographic events depends on the point in time at which some other demographic event is experienced, e.g., birth, completion of schooling, marriage, etc. It has been argued that such timing is important because it determines the degree of competitiveness individuals face throughout their lives in their search for spouses, jobs, housing, etc., which itself may influence the occurrence of demographic events such as marriage and fertility. Cohort effects have also been explained as being age-specific period effects although little empirical evidence has been provided in connection with this proposition.

The most natural approach to estimating the effects of age, period, and

1. Hobcraft, Menken, and Preston (1982) provide an excellent review of age, period, and cohort effects in demography.

cohort factors on some demographic rate (r) begins by specifying the following general model:

(1) $r = r(a,p,c)$

where a, p, and c are variables which measure age, period, and cohort and r is some general function. This expression can be linearized to yield

(2) $r_{apc} = \gamma + A'\alpha + P'\phi + C'\Phi + \varepsilon_{apc}$

where α, ϕ, and Φ are vectors of age, period, and cohort effect parameters, γ is a scaling constant, ε is a remainder term, and A', P', and C' are vectors of indicator variables which have a single corresponding element equal to unity for each observation, with all other elements equal to zero.

At first glance, it appears that the parameters of equation (3) can be estimated by ordinary least squares regression. However, because the age, period, and cohort variables are related by the identity

(3) $a + c = p$

the parameters of equation (3) are not separately identified, i.e., this is a case of perfect multicollinearity. Thus, demographers and statisticians have had to explore alternative approaches to estimation.

One set of alternative approaches involves a procedure in which γ, α and Φ are estimated under the assumption $\phi = 0$ (Ewbank, 1974; Rodgers and Thornton, 1983), with period effects taken to be the residuals from this model. However, because the assumption $\phi = 0$ omits P' from the regression, and because p' is correlated with A' and C' by the identity in equation (3), this procedure guarantees that the estimates of α and Φ will be biased. Thus, the estimates of ϕ computed from the residuals will also be biased.[2]

Another approach to the identification problem involves the imposition of constraints on the parameter vectors to be estimated. More specifically, it has been shown that identification can be achieved by constraining, in a variety of possible ways, the values of at least some of the parameters in equation (2) (see Palmore, 1978; Feinberg and Mason, 1979). Thus, the solution to the identification problem lies in the imposition of mathematical structure on equation (2). Unfortunately, however, little research has been done regarding the nature of an appropriate structure, i.e., a set of reasonable identifying constraints.

In this paper, we seek to remedy this deficiency by specifying a class of assumptions which can be used to identify the parameters of equation (2). We argue that the class of assumptions we propose is *a priori* reasonable and we outline simple tests for choosing an assumption from among this class. We then illustrate and test this technique using data on first birth rates experienced by American women between 1917 and 1980. Two important substantive findings which emerge from this example are that period effects are highly age specific and that the size of cohort effects may be substantially overestimated by models which fail to allow for the age specificity of period effects.[3]

---

2. Indeed, given the identity in equation (3), the residuals will be orthogonal to the true period effects by construction of the linear regression.

3. A third class of solutions to the identification problem involves replacing the age, period, or cohort variables by the underlying variables which they proxy. Provided these underlying variables are not definitionally dependent, this procedure solves the identification problem. Unfortunately, this procedure is totally infeasible when, as is often the case, data on the underlying variables are not available. See Heckman and Robb (1983) for further discussion of this point.

II. The Model

In this section, we propose a method for estimating age, period, and cohort effects in demographic data. To begin with, we discuss a set of separability assumptions which allow equation (1) to be written in empirically more tractable forms. In particular, let us consider the assumption of weak separability between age and period effects. This assumption suggests that equation (1) can be rewritten as

(4) $r(a,p,c) = r_1(a,c)*r_2(a,p)$

where r_1 and r_2 are general functions. In other words, it suggests that there is no interaction between cohort and period effects. Thus, period phenomena may have age-specific effects, but not cohort-specific effects. Observe that weak separability between both cohort and period effects and age and period effects is a special case of this assumption. In this special case equation (1) can be written as

(5) $r(a,p,c) = r_3(a,c)*r_4(p)$

where r_3 and r_4 are general functions. Observe that equation (5) implies that period phenomena affect demographic data by the same multiplicative factor at all ages.

With the stage set in this manner, we shall proceed by suggesting a class of methods for estimating $r_2(a,p)$ and $r_4(p)$ in equations (4) and (5). Using these estimates, we can remove the effects of period factors from the data (e.g., by dividing $r(a,p,c)$ by $r_2(a,p)$ or $r_4(p)$) and then use least squares regression to estimate age and cohort effects from the data that are purged in this way.

Our procedure for estimating $r_2(a,p)$ builds on three findings from earlier research on age, period, and cohort effects in demographic data (see, for example, Ryder, 1964; Pullum, 1978; Hobcraft, Menken, and Preston, 1979; Ryder, 1980; Bloom, 1982a; Bloom, 1982b; Bloom and Trussell, 1984). First, much existing literature indicates that demographic rates exhibit greater fluctuations in age and period dimensions than they do across cohorts. Second, existing literature provides strong evidence that cohort-to-cohort changes in most demographic rates are relatively small and that most series of cohort rates are fairly smooth in nature. Third, Ryder (1964) has proven that period rates may vary either because of genuine period effects or because of cross-cohort changes in the level or timing of the demographic event of interest.

We operationalize these three findings by assuming that the component of period variations that is due to cross-cohort changes in timing or level can be modeled as a low-order polynomial of the period index. Thus, for the case in which period effects are assumed <u>not</u> to be age-specific, we may write

$$(6) \quad r(p) = \sum_{i=0}^{j} \beta_i p^i + \varepsilon_p$$

where $r(p)$ is the period rate $(= \sum_{a=0}^{\omega} r(a,p))$, j is the order of the polynomial, and ε_p is a period effect which we treat as a random variable that has mean zero and is uncorrelated with the cohort effects captured by the polynomial. Observe that the β_i can be estimated by ordinary least squares regression of the period rate on a jth degree polynomial of the period index. Observe also that if the cross-cohort change component of period rate variations is adequately modeled, then the ratio $r(p)/\sum_{i=0}^{j} \hat{\beta}_i p^i$ is an estimate of the proportionate period effect in any period p. Such period effect estimates can be used to convert the raw data r_{apc} into $r_3(a,c)$ (see equation (5)). Standard regression analysis can then be applied to the transformed data to estimate age and cohort effects (i.e., $r_3(a,c) = \gamma + A'\alpha + C'\phi + \mu_{ac}$).

We now consider the case in which period effects are assumed to be age-specific. In this case, we seek to estimate period effects separately for different ages or age groups. Following the method outlined above, a natural approach to this problem involves modeling the cross-cohort component of period variations in the demographic rate at a particular age or in a particular age group as a low-order polynomial of the period index. Thus, we may write

$$r^*(a,p) = \sum_{i=0}^{j} \gamma_i p^i + \varepsilon^*_p$$

where $r^*(a,p) = \sum_{a=a_1}^{a_2} r(a,p)$ and a_1 and a_2 are the endpoints of the age group covered ($a_2 \geq a_1$), and ε^*_p is a random period effect with the same properties as ε_p in equation (6). Observe that γ_i are estimable parameters for the age group and that it will be necessary to estimate more than one set of γ_i, i.e., one set for each age or age group. Observe also that the estimated γ_i can be used to remove period effects from the raw data so that cohort and age effects can be estimated.

Before concluding this section, three additional comments deserve mention. First, the degree of the polynomial (j) used to model cohort effects on period rates is somewhat arbitrary. In this research, we have experimented with linear (j = 1), quadratic (j = 2), and cubic (j = 3) specifications. In most cases we have found the quadratic to be best and to provide results which are little different from those computed using higher-order polynomial specifications. Second, the degree of age-specificity allowed in the period effects is also somewhat arbitrary. However, it is possible to test for the age-specificity of period effects by studying patterns of correlation among the estimated period effects for different ages or age groups. In general, correlations that significantly differ from positive unity suggest that period effects are age specific. Third, goodness-of-fit in the period-adjusted regressions is a useful criterion for choosing among different models (i.e., different degree polynomials and different degrees of age-specificity in period effects) since the removal of period effects should in principle remove unexplained variance from those regressions.

III. An Example

This section illustrates and tests the use of the procedure developed above by applying it to data on the first birth rates of American women at ages fourteen to forty-nine during the years 1917 to 1980. Period effects are estimated using polynomials of varying degree and assuming different degrees of age specificity. Age and cohort effects are then estimated and goodness-of-fit statistics are analyzed to compare the performance of the procedure under different assumptions.

Before presenting the estimation results, it should be noted that the application of the procedure to first birth data provides a reasonable test of the procedure. To begin with, much research has been done recently on the form of cohort distributions of first births by age (see Trussell, Menken, and Coale, 1981; Bloom, 1982a; Bloom, 1982b). In general, this body of research has demonstrated the existence of strong regularities in the age pattern of first births for cohorts of women in many different populations and over a long period of time, i.e., in the absence of major temporal disturbances, these distributions are smooth, unimodal, have little density below age fifteen or above age fifty, and they are skewed to the right. Moreover, these distributions are closely replicated by the Coale-McNeil marriage model, a three-parameter, closed-form expression for the relative frequency of cohort first births at different ages. This expression is presented in equation (7)

$$(7) \quad g(a) = \frac{E}{\sigma} \cdot 1.2813 \exp\left[-1.145 \left(\frac{a-\mu}{\sigma} + .805\right) - \exp\left[-1.896 \left(\frac{a-\mu}{\sigma} + .805\right)\right]\right]$$

where g(a) is the proportion of women having a first birth at age a in the observed population, μ and σ are parameters representing the mean and standard deviation of age at first birth for those women who bear children, and E is a parameter representing the proportion of women who bear children. We shall proceed by assuming that in the absence of period effects the age distribution of cohort first births follows the form of equation (7). Provided then, that the adjustment model yields unbiased estimates of the true period effects, equation (7) should replicate first birth data that have been adjusted for period effects more closely than it replicates unadjusted data. We examine this proposition empirically below. For the moment, however, we wish to stress the fact that a test of this form, in which we utilize independent information from outside the model, is essential since we are dealing with a model that is fundamentally unidentified from the data.

In order to conserve space, we do not present parameter estimates for all of the age, period, and cohort effect estimates we have computed. Instead, we report selected results in Tables I-III. First, Table I presents goodness-of-fit measures for the Coale-McNeil model fit to cohort first birth distributions adjusted using a second-degree polynomial with each of three assumptions about the age-specificity of period effects: (i) period phenomena have identical multiplicative effects at all ages, (ii) period phenomena have different multiplicative effects in the seven age intervals 14-19, 20-24, 25-29, 30-34, 35-39, 40-44, and 45-49, and (iii) period effects are single-year age-specific. For completeness, Table I also presents goodness-of-fit results for period-unadjusted data. The goodness-of-fit statistic reported is the normalized mean absolute deviation (NMAD). This statistic represents the area between the fitted and observed first birth distribution divided by the area under the fitted distribution. Perfect fits are thus indicated by values of NMAD equal to zero with worse fits indicated by successively larger positive values of NMAD. Results are reported in Table I for all cohorts for whom completed first birth data are available (i.e., data for the entire age range 14 to 49 during the years 1917 to 1980).

The statistics reported in Table I are strongly supportive of the procedure developed in Section II. In particular, application of the procedure results in improved goodness-of-fit statistics for every cohort studied and for every successive increase in age-specificity. Moreover, in some cases, the improvement is dramatic. For example, a 16.2 percent error for the 1924 birth cohort is reduced to 2.3 percent error. Thus, the model developed seems to be providing extremely good estimates of the period effects.

The second finding worth noting relates to the age specificity of the period effects. In this connection, Table II presents the correlation matrix for period effects estimated under age-specificity assumptions (i) and (ii) above. Note that the correlations are often substantially less than unity, indicating that period effects do seem to be age-specific.[4] This seems to be the first empirical evidence demonstrating the age-specificity of period

4. Although not reported here, the correlation matrix for single-year age specific period effects also exhibits non-perfect positive correlations although the correlations within the seven age groups are relatively high.

effects.

The third finding worth noting relates to the nature of the parameter estimates for the cohort first birth distributions. These parameter estimates, which are reported in Table III, show substantially less variation across cohorts after the removal of period effects. This finding therefore supports the proposition that much of what appears to be cohort effects is simply age-specific period effects. For example, the age-specific period adjustments reduce the standard deviations of the cross-cohort estimates of μ, σ, and E by 71 percent, 47 percent, and 77 percent, respectively. Although these reductions are not of sufficient magnitude to contradict the existence of cohort effects, they do suggest that the size of cohort effects may be substantially overestimated by models which fail to allow for the age specificity of period effects.

Table I
Normalized Mean Absolute Deviations[1],
1903-1931 Birth Cohorts, Women Aged 14-49

Birth Cohort	NMAD Unadjusted	NMAD, Period - Adjusted: No Period-Age Interaction	Period Interaction with 5 year age groups	Period Interaction with single yr. age groups
1903	0.058	0.048	0.042	0.029
1904	0.057	0.052	0.040	0.028
1905	0.072	0.062	0.039	0.028
1906	0.093	0.068	0.037	0.029
1907	0.107	0.072	0.037	0.028
1908	0.116	0.075	0.034	0.028
1909	0.119	0.078	0.033	0.028
1910	0.129	0.079	0.034	0.028
1911	0.122	0.074	0.031	0.028
1912	0.113	0.071	0.033	0.028
1913	0.108	0.069	0.032	0.028
1914	0.099	0.069	0.035	0.028
1915	0.110	0.070	0.039	0.028
1916	0.122	0.072	0.037	0.027
1917	0.128	0.072	0.037	0.027
1918	0.138	0.073	0.040	0.027
1919	0.135	0.064	0.041	0.026
1920	0.136	0.059	0.036	0.026
1921	0.141	0.066	0.041	0.025
1922	0.153	0.075	0.053	0.024
1923	0.158	0.075	0.048	0.024
1924	0.162	0.073	0.040	0.023
1925	0.158	0.064	0.041	0.022
1926	0.136	0.063	0.045	0.022
1927	0.101	0.068	0.041	0.022
1928	0.088	0.067	0.034	0.021
1929	0.069	0.061	0.030	0.022
1930	0.060	0.045	0.034	0.023
1931	0.063	0.037	0.035	0.024

---

1. $$\text{NMAD} = \frac{\sum_{a=14}^{49} |g(a) - \hat{g}(a)|}{\sum_{a=14}^{49} g(a)}$$

Table II
Estimated Correlation Matrix for Period Effects
Computed for Selected Age Groups

	14-49	14-19	20-24	25-29	30-34	35-39	40-44
14-49	1.00	.75	.95	.64	.50	.61	.78
14-19		1.00	.83	.02	-.14	.06	.42
20-24			1.00	.41	.27	.44	.68
25-29				1.00	.92	.80	.64
30-34					1.00	.93	.71
35-39						1.00	.88
40-44							1.00

Table III

Estimated Parameters with No Period Adjustments (I) and with Single-Year-Age-Specific Period Adjustments (II)

Birth Cohort	Parameter: μ I	μ II	E I	E II	σ I	σ II
1903	22.91	23.83	.79	.80	4.88	5.55
1904	22.95	23.79	.79	.81	4.99	5.53
1905	22.96	23.76	.77	.82	5.02	5.51
1906	22.98	23.71	.76	.82	5.07	5.45
1907	23.09	23.68	.76	.83	5.23	5.42
1908	23.25	23.65	.76	.84	5.45	5.40
1909	23.51	23.61	.77	.84	5.73	5.40
1910	23.81	23.58	.78	.85	6.01	5.35
1911	24.11	23.55	.79	.86	6.27	5.33
1912	24.39	23.52	.81	.86	6.47	5.29
1913	24.63	23.49	.83	.87	6.62	5.28
1914	24.79	23.46	.84	.87	6.58	5.25
1915	24.84	23.44	.85	.87	6.51	5.24
1916	24.83	23.42	.86	.88	6.42	5.22
1917	24.83	23.39	.87	.88	6.36	5.19
1918	24.81	23.36	.88	.89	6.34	5.18
1919	24.69	23.34	.89	.89	6.11	5.17
1920	24.56	23.32	.90	.89	5.97	5.14
1921	24.47	23.30	.92	.89	5.87	5.13
1922	24.37	23.28	.91	.90	5.72	5.11
1923	24.26	23.25	.92	.90	5.57	5.10
1924	24.11	23.24	.93	.90	5.40	5.09
1925	23.92	23.22	.92	.90	5.13	5.07
1926	23.62	23.20	.90	.90	4.75	5.06
1927	23.25	23.20	.88	.90	4.44	5.04
1928	22.96	23.16	.88	.90	4.38	5.03
1929	22.85	23.14	.89	.90	4.50	5.02
1930	22.82	23.12	.91	.90	4.64	5.01
1931	22.79	23.10	.92	.90	4.68	5.00

REFERENCES

1. Bloom, David E. 1982a. "What's Happening to the Age at First Birth in the United States? A Study of Recent Cohorts," Demography 19: 351-370.

2. ___________. 1982b. "Age Patterns of Women at First Birth." Genus XXXVIII: 101-128.

3. ___________ and James Trussell. 1984. "What Are the Determinants of Delayed Childbearing and Permanent Childlessness in the United States?" Demography 21: 591-611.

4. Ewbank, Douglas C. 1974. An Examination of Several Applications of the Standard Pattern of Age at First Marriage, unpublished Ph.D. dissertation, Princeton University.

5. Feinberg, Stephen E. and William M. Mason. 1979. "Identification and Estimation of Age-Period Cohort Models in the Analysis of Discrete Archival Data," in Sociological Methodology, 1979, Karl F. Schuessler, Editor, San Francisco: Jossey-Bass, pp. 1-65.

6. Heckman, James and Richard Robb. 1983. "Using Longitudinal Data to Estimate Age, Period, and Cohort Effects in Earnings Equations." In H. Winsborough and O. Duncan, Eds., Analyzing Longitudinal Data for Age, Period, and Cohort Effects, New York: Academic Press.

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

8. Palmore, Erdman. 1978. "When Can Age, Period, And Cohort be Separated?" Social Forces 57: 282-295.

9. Pullum, Thomas W. 1980. "Separating Age, Period, and Cohort Effects in White U.S. Fertility, 1920-1970." Social Science Research 9: 225-244.

10. Rodgers, William L. and Arland Thornton. 1985. "Changing Patterns of First Marriage in the United States." Demography 22: 265-279.

11. Ryder, Norman B. 1964. "The Process of Demographic Translation." Demography 1: 74-82.

12. ___________. 1980. "Components of Temporal Variations in American Fertility." In Demographic Patterns in Developed Societies, R. W. Hiorns, Editor, London: Taylor and Francis, pp. 11-54.

13. Trussell, James, Jane Menken, and Ansley J. Coale. 1981. "A General Model for Analyzing the Effect of Nuptuality on Fertility." In L. Ruzicka, Editor, Nuptuality and Fertility: Proceedings of a Conference. Liege: Ordina Editions, 1981.