

Demography: Analysis and Synthesis

Graziella CASELLI, Jacques VALLIN, and Guillaume WUNSCH

with contributions by

Daniel COURGEAU, Nico KEILMAN, Eva LELIÈVRE, James VAUPEL,
Anatoli YASHIN, and John WILMOTH

VOLUME I




ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Academic Press is an imprint of Elsevier



Academic Press is an imprint of Elsevier
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
525 B Street, Suite 1900, San Diego, California 92101-4495, USA
84 Theobald's Road, London WC1X 8RR, UK

This book is printed on acid-free paper. 

Copyright © 2006, Elsevier Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the publisher.

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, e-mail: permissions@elsevier.co.uk. You may also complete your request online via the Elsevier homepage (<http://elsevier.com>), by selecting "Customer Support" and then "Obtaining Permissions."

Library of Congress Cataloging-in-Publication Data

Caselli, Graziella.

[Demographie. English]

Demography : analysis and synthesis : a treatise in population studies / Graziella Caselli, Jacques Vallin, and Guillaume Wunsch ; with contributions by Daniel Courgeau ... [et al.].
p. cm.

Includes bibliographical references and index.

ISBN 0-12-765660-X (set : alk. paper)—ISBN 0-12-765661-8 (volume 1 : alk. paper)—
ISBN 0-12-765662-6 (volume 2 : alk. paper)—ISBN 0-12-765663-4 (volume 3 : alk. paper)—
ISBN 0-12-765664-2 (volume 4 : alk. paper) 1. Demography. I. Vallin, Jacques.
II. Wunsch, Guillaume J. III. Title.

HB871.C37513 2005

304.6—dc22

2005019413

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

ISBN 13: 978-0-12-765660-1

ISBN 10: 0-12-765660-X

For all information on all Elsevier Academic Press publications
visit our Web site at www.books.elsevier.com.

Printed in the United States of America

05 06 07 08 09 10 9 8 7 6 5 4 3 2 1

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER

BOOK AID
International

Sabre Foundation

Age-Period-Cohort Models In Demography

JOHN R. WILMOTH

Department of Demography, University of California, Berkeley, California

Demographic events such as birth, death, marriage, and migration mark life's most important transitions and, not surprisingly, are influenced by a multitude of factors. These influences are not constant across time or space, and thus demographic rates display considerable variation. A crucial task of demographic analysis is to identify the factors that account for this variability.

One particular approach to this quite general problem is to study the variation in demographic rates along three critical dimensions: age of the event, year (or period) of its occurrence, and cohort of the individuals involved. Within this framework, *cohort* may refer to an individual's year of birth or to the time of some other important transition in the life course (e.g., a marriage cohort in a study of divorce). Likewise, *age* may mean a person's chronologic age at the time of the event or the time elapsed since another important transition (e.g., duration of employment in a study of retirement).

The underlying assumption of age-period-cohort (APC) analysis is that all factors influencing a demographic event or its rate of occurrence can be grouped meaningfully into these three categories. Of course, there may be interactions or overlap between the three sets of factors, but it is assumed, implicitly, that these influences are less important than the main effects associated with each dimension. The typical analytic strategy of APC studies, therefore, is first to identify

patterns of change in demographic rates that are associated with these three dimensions or sets of causal factors. Once this elaborate description is available, the analyst attempts to identify the specific influences (biologic processes, historic trends, etc.) that are responsible for the observed patterns in terms of age, period, and cohort.

Based on this description, the logic of APC analysis seems simple and relatively straightforward. There is, however, a fundamental problem with this strategy related to the fact that, mathematically,

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

This relationship is illustrated by means of a Lexis diagram in Figure 18-1. In the view of some researchers, the exact mathematic connection between age, period, and cohort renders all forms of APC analysis meaningless, since changes in a demographic process along one of the three dimensions cannot be distinguished statistically from changes along the two remaining dimensions (Glenn, 1976; Goldstein, 1979; Rodgers, 1982). Other researchers have proposed a variety of solutions for overcoming this problem and claim that valid and useful results can still be derived within an APC framework (Fienberg and Mason, 1979; Clogg, 1982; Caselli and Capocaccia, 1989; Wilmoth, 1990). All researchers would agree, however, that the identification problem that plagues APC analysis is a fundamentally difficult and perplexing problem.

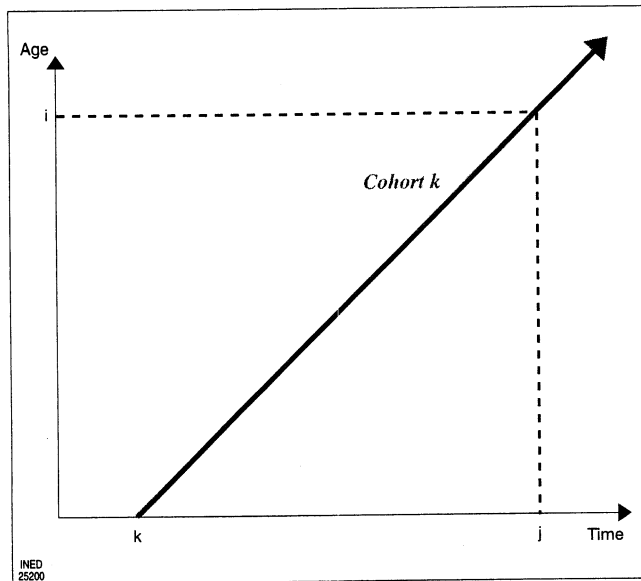


FIGURE 18-1 Lexis diagram showing the relationship between age, period, and cohort of a demographic event. The lifeline of individuals born at time k (cohort) is represented by the diagonal arrow. For such persons, a demographic event at age i occurs at time $j = k + i$ (period).

In this chapter, we first discuss the theoretical motivation of the age-period-cohort framework. We then provide a formal definition and discussion of the identification problem that affects APC modeling. Although we do not provide a complete literature review, we discuss strategies for estimating APC models that have been proposed in particular contexts. We also describe a small number of empirical applications of APC modeling that seem particularly informative. In conclusion, we offer some general remarks about the applicability of the APC framework.

I. THEORETICAL MOTIVATION

We begin our theoretical discussion by asking, "Why is the age-period-cohort framework useful?" The period of an event, for example, has no direct influence on its outcome but is merely a marker for other factors that may affect a demographic process. Periods may be associated with wars, epidemics, political and legal changes, famines, inventions, scientific discoveries, and changing fashions. It is these occurrences or situations, not the period itself, that account for the variation in demographic rates along the period dimension. Likewise, it is not cohort membership and chronologic age themselves that influence demographic rates, but rather the characteristics associated with those two dimensions.

It is worth considering, however, whether all three dimensions should be important in all situations. In the case of the classical demographic events (birth, death, marriage, and migration) it seems obvious that age, or duration, should be associated with sharp differences in outcomes. Due to biologic mechanisms, both the risk of death and the chance of giving birth vary enormously with age. Owing mostly to social and psychological mechanisms, the risk of divorce varies depending on the duration of marriage and the probability of migration changes over the life cycle. The age pattern of marriage is a complex function of both biologic maturation and social convention.

The immediacy of period-specific events makes their connection to demographic processes self-evident. The role of cohort-specific influences, on the other hand, is less obvious. In an influential article, Ryder (1965) laid out the theoretical justification for the use of cohorts in the study of social change. According to Ryder, fresh cohorts enable the process of social change, since young people are less constrained by history and more capable of adapting to or creating new modes of living. Furthermore, individuals experience certain critical events (e.g., infancy and childhood, education, military service, first employment) alongside peers from their particular cohort and the after-effects of these situations may mark them for life. The imprint of life-defining events may be biologic or social, but in either case we should expect cohorts to represent an important dimension of variation in subsequent demographic processes.

The direction and magnitude of the influence of cohort membership on demographic phenomena have been a matter of considerable speculation, although few firmly established generalizations have emerged. Perhaps the most well-known theory about the influence of cohort membership on demographic events is Easterlin's explanation of the American baby boom (Easterlin, 1961, 1978). According to this theory, smaller cohorts enjoy significant advantages as they enter the job market, which facilitates family formation and thus raises fertility levels. With regard to the baby boom, however, it is difficult to distinguish the effects of cohort size from the effects of an expanding postwar economy, both of which may have influenced life chances and family behavior around the 1950s.

In mortality studies, the discussion of the role of cohorts has focused on the after-effects of adverse events in early life on subsequent probabilities of death or survival (for a review, see Elo and Preston, 1992). Theoretically, it is possible that adverse events in early life could result in either higher or lower mortality in later life. On one hand, the survivors of a traumatic

event or situation (e.g., famine, war, poverty) may be weakened by the experience and thus demonstrate unusually high levels of mortality later on. On the other hand, it is possible that the survivors in these situations may be an especially robust subset of the original cohort and thus could display unexpectedly low levels of mortality in subsequent years. These competing explanations are often referred to as processes of *debilitation* and *selection*, respectively. It is possible, of course, that both processes are operating at the same time. In these situations, it is their relative magnitudes that should determine whether an affected cohort experiences unusually high or low mortality in later years.

In summary, there is a sound theoretical basis for believing that age, period, and cohort may all be important dimensions of variation in demographic processes. It is this theoretical interest that has encouraged the development of statistical models incorporating all three components. We now turn our attention to such models.

II. AGE-PERIOD-COHORT MODELS

The standard age-period-cohort model has the following form:

$$f(r_{ijk}) = \mu + \alpha_i + \beta_j + \theta_k + \varepsilon_{ijk}$$

In this formulation, r_{ijk} is an observed demographic rate for some event that occurs at age i in year j for cohort k (thus, $k = j - i$); the function, $f(\cdot)$, is some transformation applied to the observed rates; the parameter, μ , establishes an overall level for $f(r_{ijk})$; the parameters, α_i , β_j , and θ_k , describe patterns of change in $f(r_{ijk})$ by age, period, and cohort, respectively; and the last term, ε_{ijk} , represents error (either in the specification of the model or in the original data) and random fluctuations.

Although the standard APC model is quite simple in form, its parameters are not easily estimated. The central difficulty is that there is no obvious means of identifying a unique set of parameter estimates that provide an optimal fit to the observed data. This identification problem affects not only the statistical estimation of the parameters but also their subsequent interpretation.

1. Identification Problem

The identification problem of the standard APC model is easily illustrated by considering the model in its estimated form:

$$\hat{y}_{ijk} = \hat{\mu} + \hat{\alpha}_i + \hat{\beta}_j + \hat{\theta}_k$$

Thus, the predicted value of the transformed rates, \hat{y}_{ijk} , equals the sum of the four estimated parameters, $\hat{\mu}$, $\hat{\alpha}_i$, $\hat{\beta}_j$, and $\hat{\theta}_k$.

The first task in identifying a unique solution for this model involves constraining the levels of the four parameters. A common approach is to require that

$$\sum_i \hat{\alpha}_i = \sum_j \hat{\beta}_j = \sum_k \hat{\theta}_k = 0$$

With this solution, $\hat{\mu}$ gives some average level of $f(r_{ijk})$, and the other three sets of parameters describe changes in the transformed rates with respect to the average. Alternatively, we could require that $\hat{\alpha}_i = \hat{\beta}_j = \hat{\theta}_k = 0$ for some specific choice of i , j , and k . In this case, $\hat{\mu}$ gives the predicted value of this particular $f(r_{ijk})$, and the other parameters describe changes with respect to this point of reference.¹

Constraints of this sort are common in other statistical applications (analysis of variance, regression, etc.). Obviously, they involve an arbitrary choice about how to obtain a unique solution, but this choice affects only the level of the various parameters and thus does not fundamentally confuse their interpretation. APC models, however, present an additional, more perplexing identification problem involving not only the level of the various parameters but also their slope.

For example, suppose that $\hat{\alpha}_i$, $\hat{\beta}_j$, and $\hat{\theta}_k$ were adjusted such that

$$\begin{aligned}\alpha_i^* &= \alpha_i + \lambda \cdot i \\ \beta_j^* &= \beta_j - \lambda \cdot j \\ \theta_k^* &= \theta_k + \lambda \cdot k,\end{aligned}$$

where λ is some real number. Now, note that the predicted values in the model are the same for any choice of λ , since $k = j - i$:

$$\begin{aligned}\hat{y}_{ijk} &= \hat{\mu} + \alpha_i^* + \beta_j^* + \theta_k^* \\ &= \hat{\mu} + \alpha_i + \lambda \cdot i + \beta_j - \lambda \cdot j + \theta_k + \lambda \cdot k \\ &= \hat{\mu} + \alpha_i + \beta_j + \theta_k + \lambda \cdot (i - j + k) \\ &= \hat{\mu} + \alpha_i + \beta_j + \theta_k.\end{aligned}$$

¹ Perhaps the best choice, for reasons of mathematic convenience, is to require that

$$\sum_i w_i \hat{\alpha}_i = \sum_j w_j \hat{\beta}_j = \sum_k w_k \hat{\theta}_k = 0,$$

where w_i , w_j , and w_k are weights corresponding to the number of observations for each particular age, period, and cohort. In a rectangular age-by-period array, for example, w_i and w_j would be constant (equaling the number of columns and rows, respectively) and thus would fall out of the equation. In this situation, only w_k would vary, equaling the number of elements in a diagonal of the matrix corresponding to an individual cohort.

Since different sets of age, period, and cohort parameters yield identical predicted values, there is no means of choosing between alternative solutions on the basis of goodness-of-fit. As with the earlier problem of uniquely identifying the level of these parameters, a unique solution demands an arbitrary choice. In this case, however, alternative solutions may have vastly different implications, since the speed and even direction of change in the age, period, and cohort components of the process are affected. In fact, by careful manipulation of λ , the linear trend in the three sets of parameters can be altered *ad libitum*.²

Typically, discussions of APC models have focused on the problem of uniquely identifying the linear trend in the age, period, and cohort parameters. Higher order components of these parameters (quadratic, cubic, etc.) and interactions between the three sets of factors have generally been ignored because they make identification of the model even more difficult (Fienberg and Mason, 1985; Clogg, 1982). However, second- and higher order components of the main effects are not uniquely identified if interaction terms are added to the standard APC model (Wilmoth, 1990).

2. Estimation Strategies

In this chapter, we do not attempt to provide a complete review of the literature on APC modeling. For that purpose, we refer the reader to excellent discussions by other authors (Hobcraft, Menken, and Preston, 1982; Mason and Fienberg, 1985). Instead of a review of the literature, we provide an overview of the main strategies that have been proposed for estimating and/or modifying the standard APC model in light of the identification problem discussed earlier. We divide these approaches into three groups: arbitrary assumptions, interaction terms and demographic translation, and direct measurement. See also the excellent discussion of these topics by Holford (1991).

² Note that the transformed parameters, α_i^* , β_j^* , and θ_k^* , as defined here, would not normally sum to zero. Alternatively, define them as follows:

$$\begin{aligned}\alpha_i^* &= \alpha_i + \lambda \cdot (i - \bar{i}) \\ \beta_j^* &= \beta_j - \lambda \cdot (j - \bar{j}) \\ \theta_k^* &= \theta_k + \lambda \cdot (k - \bar{k})\end{aligned}$$

where \bar{i} , \bar{j} , and \bar{k} are defined as the means of i , j , and k . It can be shown that the sums of these transformed parameters equal zero if the data matrix is rectangular. If the data array is not rectangular, a convenient solution is to use weights, as defined in Footnote 1, both for defining the parameter constraints and for computing the means of i , j , and k in the above formulas.

a. Arbitrary Assumptions

The first strategy for identifying an APC model is to make some arbitrary assumption about the linear trend in one of the three dimensions. Specific applications differ, however, depending on whether the assumption is viewed by the analyst(s) as an accurate reflection of an underlying reality or merely as a convenient strategy for estimating an arbitrary statistical model.

Unique estimates of the standard APC model (with three main effects, as described in prior sections) can be obtained by assuming that the parameters pertaining to some pair of cohorts (or ages or periods) are equal. For example, one could assume that $\theta_k = \theta_{k+1}$ for a specific cohort k . This restriction permits estimation of the model, although at the expense of assuming that the cohort-specific effects for cohorts k and $k + 1$ are equal. Arguably, such an assumption is justified in specific cases based on prior knowledge.

In a widely cited paper, Fienberg and Mason (1979) illustrated the application of the standard APC model to the analysis of educational attainment in the United States. Here, the dependent variable consisted of proportions of individuals ages 20 and above who had, at some previous moment in their lives, advanced from one educational level to the next. These proportions were computed using decennial census data from 1940 to 1970. The authors noted correctly that educational attainment should under most circumstances be constant or slightly increasing over the adult life of a cohort (this expectation could be violated only by selective migration or differential mortality). In other words, age effects are of little importance in this model and can reasonably be set equal over some age range. Similarly, period effects can have little impact on the educational attainment of adults (most of whom have completed their schooling prior to the period of observation). Indeed, the authors justify the inclusion of period effects in the model based solely on concerns about changes in the recording of educational data in decennial censuses.

Due to the relative unimportance of these two sets of components (age and period) *a priori*, the choice by Fienberg and Mason to use educational attainment as an illustration of APC modeling is somewhat puzzling. Their resolution of the standard identification problem in this particular case was relatively straightforward and unproblematic, since the authors merely equated age effects for age groups 30 to 34 through 55 to 59 years.³ This example illustrates the sensible use of

³ Obviously, the model would have been fully identified if age effects for only two age groups had been set equal. Fienberg and Mason (1979: 53) argue, however, that there are "*a priori* grounds for interpreting age effects primarily at the tails of the age distribution" and thus prefer to equate the coefficients throughout the middle age range.

prior information to resolve the APC identification problem. It provides little guidance, however, about how to resolve the identification problem in situations where there is no strong basis for asserting the equality of at least one pair of effects *a priori*.

In another application of the APC model, Clogg (1982) examines rates of labor force participation in the United States. Here, there are strong reasons for expecting that the dependent variable is influenced by factors related to age, period, and cohort, but there is only a very weak basis for supposing, *a priori*, a particular relationship between some pair of effects. Nevertheless, the author constrains the last two cohort effects to differ by one or two percent (depending on the specific example) as a means of obtaining a unique set of parameter estimates. The substantive justification of this choice can fairly be termed *weak*, although arguably it has little impact on the conclusions of the study.

In general, even if prior knowledge supports the choice of a specific identifying restriction, it is important to recognize the sensitivity of the results to minor violations of the chosen assumption. If θ_k and θ_{k+1} are close but not equal, for example, the error that results from equating them will be slight for adjacent cohorts, but much larger for distant ones (Clogg, 1982). The importance of this imprecision will depend on the application, and the sensitivity of the results to any particular assumption can often be assessed quite easily. Nevertheless, the validity of results derived from APC models in which this last identifying constraint is, allegedly, a reflection of underlying reality remains quite dubious in most cases.

Alternatively, rather than asserting the reality of identifying constraints, certain parameter restrictions can be adopted as a matter of convenience merely to obtain a unique solution. For example, Pullum (1980) fit the standard APC model by equating two adjacent cohort parameters but avoided interpreting the resulting estimates literally. Pullum's main purpose in fitting such a model was merely to determine whether the fit (in a model of American fertility rates) is improved more by the addition of period or cohort parameters. He concluded that cohort parameters provide less explanatory power (per additional parameter) than period parameters. Since he was interested only in the overall fit of each model, not in the parameter estimates, the choice of identifying constraints was irrelevant.

Pullum also noted, however, that the second (and higher) differences of the parameter estimates in the three-factor APC model are invariant to the choice of an identifying assumption. In our notation, first differences depend on the choice of λ , although second differences do not. For example,

$$\begin{aligned}\Delta\alpha_i^* &= \alpha_{i+1}^* - \alpha_i^* \\ &= \hat{\alpha}_{i+1} + \lambda \cdot (i+1) - [\hat{\alpha}_i + \lambda \cdot i], \\ &= \hat{\alpha}_{i+1} - \hat{\alpha}_i + \lambda\end{aligned}$$

whereas,

$$\begin{aligned}\Delta^2\alpha_i^* &= \Delta\alpha_{i+1}^* - \Delta\alpha_i^* \\ &= \hat{\alpha}_{i+2} - \hat{\alpha}_{i+1} + \lambda - [\hat{\alpha}_{i+1} - \hat{\alpha}_i + \lambda], \\ &= \hat{\alpha}_{i+2} - 2\hat{\alpha}_{i+1} + \hat{\alpha}_i\end{aligned}$$

First differences are related to the linear trend, and thus their dependence on λ is a reminder that the slope of the three sets of parameters cannot be uniquely identified in the standard APC model. Second differences are measures of deviation from the linear trend, and thus these quantities can be accurately estimated in the simple model with three main effects (but no interaction terms).

An analogous solution is to require that the slope of one of the three sets of parameters be constrained to equal zero. Following Wilmoth (1990), a simple three-factor model applied to a rectangular age-by-period array of demographic rates could be estimated by using the following constraints:

$$\sum_i \alpha_i = \sum_j \beta_j = \sum_k w_k \theta_k = \sum_k k w_k \theta_k = 0$$

These constraints on θ_k are equivalent to requiring that both the *level* and the *slope* of a weighted least-squares regression line fitted to the cohort parameters should equal zero.⁴

Thus, one solution to the APC identification problem is to choose, arbitrarily, one dimension whose slope is constrained to equal zero (see also Caselli and Capocaccia, 1989). As a matter of convenience, the chosen dimension may correspond to the diagonals of a rectangular array of demographic rates. The resulting parameter estimates for this third factor have a simple interpretation: They represent deviations from the long-term pattern of change in that dimension. An analysis of such deviations is logically similar to an analysis of second derivatives, since both emphasize deviations from the long-term trend rather than absolute effects.

b. Interaction Terms and Demographic Translation

An obvious shortcoming to the standard APC model, aside from its problems of identifiability, is the absence of interaction terms. To address this inadequacy, Wilmoth (1990) proposed models with the following form:

⁴ Here, as in Footnotes 1 and 2, the weights equal the length of the diagonals corresponding to individual cohorts.

$$f(r_{ijk}) = \mu + \alpha_i + \beta_j + \sum_{m=1}^{\rho} \phi_m \gamma_{im} \delta_{jm} + \theta_k + \varepsilon_{ijk}^5$$

This model differs from the standard APC model by the addition of ρ (usually one or two) multiplicative terms. The use of these interaction terms is motivated by the fact that the pace of change (over time) in demographic rates often differs by age. For example, mortality rates have typically fallen much more rapidly at younger than at older ages. This component of differential change by age cannot be expressed well in a model with no age-period (or age-cohort) interactions.

The above model is appropriate for an analysis of an age-period array of demographic rates, where the rates of change in r_{ijk} differ by age. The multiplicative terms capture these differential trends by age. The cohort parameters, θ_k , constrained to have zero level and slope, express deviations from the overall pattern of change by age and period. In short, the θ_k captures residual patterns that seem truly to lie along diagonals of the age-period array. Without the multiplicative term(s), however, estimates of θ_k may be heavily influenced by age-period interactions. The introduction of interaction terms by no means resolves the identification problem in the standard APC model. Their use is consistent, however, with the philosophy that the proper means of including all three sets of parameters in the model is to focus the description on two dimensions only and to treat the third dimension as a sort of residual.

The multiplicative interaction terms also help us to understand the notion of demographic translation (Ryder, 1964, 1980; see also Chapter 17). In the above model, the cohort term, θ_k , has been reduced to a residual component, and thus the linear or long-term cohort trend has been translated into the age and period dimensions. Even this residual component of the cohort trend, θ_k , can be expressed as age-period interactions, if a sufficient number of multiplicative terms are added to this model (Wilmoth, 1990). In this sense, the choice to assign any portion of the variability in demographic rates to one of the three dimensions is arbitrary, because it is always possible to re-express the pattern in terms of the other two dimensions by the addition of a sufficient number of interaction terms.

The only resolution of this logical dilemma is to acknowledge, as a matter of plausibility, that certain patterns must be due to changes in one specific dimension. For example, the residual cohort effects that emerge in a complete analysis of an age-by-period array (including at least some form of age-period inter-

action) cannot plausibly be attributed to age- and period-specific factors. These residual cohort effects tend to vary with high frequency and thus are not easily translated into the age and period dimensions. It is worth noting that Foster (1990) reached a similar conclusion (i.e., that only high-frequency cohort effects are plausibly identifiable) by an entirely different method of analysis.

c. Direct Measurement

Several analysts have noted that the only true resolution of the identification problem affecting APC analysis would be to measure directly the factors whose effects are summarized by the coefficients of the standard APC model. These factors would not typically be linearly related (although they might be strongly collinear), and thus a unique solution for the model could be found. An obvious difficulty with this strategy is the problem of defining and then measuring the proximate variables for which age, period, and cohort are, admittedly, only rough approximations.

In a study of mortality in Italy, for example, Caselli and Capocaccia (1989) use the probability of death in infancy or during the first 15 years of life as a measure of long-term cohort influences, whose magnitude may vary over the life course. In our notation, their model can be written

$$f(r_{ijk}) = \mu + \alpha_i + \beta_j + \gamma_i \cdot Q_k + \varepsilon_{ijk}$$

where Q_k equals either q_0 or ${}_{15}q_0$. This model requires no special identifying restrictions. Because the estimated γ_i 's are positive below age 45 and negative above this age, the authors conclude that "higher mortality early in life is associated with higher mortality up to age 45 and lower levels at later ages" (Caselli and Capocaccia, 1989, p. 152).

An alternative interpretation of these results, however, is that the age-cohort term in their model is, in reality, capturing the same sorts of patterns that were treated as age-period interactions by Wilmoth and colleagues (Wilmoth *et al.*, 1989; Wilmoth, 1990). Indeed, the multiplicative age-cohort term in the model of Caselli and Capocaccia (1989) is remarkably similar, in both its form and its estimated values, to the multiplicative age-period term in the models by Wilmoth and collaborators. In other words, the age-cohort term in the above model may serve merely to document the relatively more rapid rate of change in mortality at younger than at older ages. Since infant-child mortality has declined almost monotonically for the cohorts in question, its function in the model may be simply to provide a marker for the second dimension of temporal mortality change.

⁵ For a full discussion of the constraints required for fitting this model, see Wilmoth (1990).

The innovative applications and clear explanations of the age-period-cohort model by Caselli and Capocaccia (1989) are certainly valuable. Their attempt to measure age, period, or cohort effects directly, however, illustrates some of the statistical difficulties inherent in the task. As noted by Clogg (1982),

The selection of the proper causal variables to be considered in a modeling procedure is a most difficult theoretical task, one that is at least as difficult as applying the age-period-cohort accounting framework (p. 460).

In other words, although it may seem in principle that direct measurement is an obvious solution for the APC identification problem, the practical application of this technique is perhaps more fraught with difficulties than the problem we are trying to overcome.

3. Empirical Examples

Two specific examples of the application of APC models to demographic data are reviewed here. These examples were chosen to illustrate two kinds of conclusions that can legitimately be drawn from APC models applied to time series of age-specific demographic rates. The first example, Pullum's (1980) study of American white fertility, illustrates the use of the standard APC model to compare the variability by cohort versus period. The second example, Wilmoth's (1990) analysis of French male mortality, employs a more complicated APC model with interaction terms to document the unusually high or low levels of mortality for certain cohorts.

a. American White Fertility, 1920 to 1970

Pullum (1980) analyzed an age-by-period array of age-specific fertility rates for American white women aged 15 to 44 during years 1920 to 1970. He fit all three possible two-factor models and the standard three-

factor APC model. Table 18-1 shows χ^2 measures of fit for each model in four time periods, each spanning two decades (lower values indicate a better fit). The period \times cohort model is clearly inferior to the other two-factor models, reflecting the primary importance of age as a dimension of variation in this case. It is not clear from the comparison in Table 18-1, however, whether age should be paired with period or cohort to obtain the best two-dimensional description of these data.

Table 18-2 reveals the greater importance of periods than cohorts as a dimension of variation in these data by comparing each two-factor model to the full APC model. The values in this table (which can be derived from those in Table 18-1) can be thought of as the improvement in fit when age, period, or cohorts is added as the third and final dimension of the standard APC model or as the loss of fit when one of these dimensions is removed from the full three-factor model. In each case, the improvement or loss of fit is measured relative to the change in degrees of freedom.

Pullum (1980) applied these same four models to an age-by-cohort array of age-specific fertility rates for American white women aged 15 to 44 years born during years 1905 to 1926 (thus including all cohorts whose reproductive lives fell entirely within the observation period) (Table 18-3). He notes that

In contrast to the previous applications to rectangular [age-by-period] data, the inclusion of period effects appears *much* more useful than the inclusion of cohort effects (p. 236, emphasis in original).

For example, using measures like those presented here in Table 18-2, the loss of fit caused by removing all period parameters from the full APC model is 13.1 (increase in chi-square per degree of freedom) compared to only 0.4 if the cohort dimension is eliminated.

In conclusion, Pullum stated (p. 241)

TABLE 18-1 χ^2 Measure of Fit for Four Models Applied to Age-specific Fertility Rates, American White Women, Aged 15-44 Years, 21-year Time Intervals

| Model | df | Interval of time | | | |
|-------------------------------------|-----|------------------|-----------|-----------|-----------|
| | | 1920-1940 | 1930-1950 | 1940-1960 | 1950-1970 |
| Age \times period | 580 | 148.1 | 330.2 | 411.3 | 93.3 |
| Age \times cohort | 551 | 90.2 | 397.8 | 411.6 | 800.2 |
| Period \times cohort | 560 | 13,306.8 | 14,952.4 | 20,603.5 | 21,130.1 |
| Age \times period \times cohort | 532 | 21.6 | 133.1 | 98.0 | 40.8 |

Note: Data for each time interval were in age-by-period format. χ^2 estimate is based on an artificial case base of 630,000 cases for each time interval.

Source: Pullum Thomas W., 1980. Separating age, period, and cohort effects in white U.S. fertility, 1920-1970. *Social Science Research*, vol. 9(3), p. 227-244. New York, Academic Press, Table 2.

TABLE 18-2 Reductions in χ^2 Per *df* When Age, Period, or Cohort is Added as Final Term of Standard APC Model of Age-specific Fertility Rates, American White Women, Aged 15-44 Years, 21-year Time Intervals

| Effect | Interval of time | | | |
|--------|------------------|-----------|-----------|-----------|
| | 1920-1940 | 1930-1950 | 1940-1960 | 1950-1970 |
| Age | 474.5 | 529.3 | 732.3 | 753.2 |
| Period | 3.6 | 13.9 | 16.5 | 40.0 |
| Cohort | 2.6 | 4.1 | 6.5 | 1.1 |

Note: See note for Table 18-1.

Source: Pullum Thomas W., 1980. Separating age, period, and cohort effects in white U.S. fertility, 1920-1970. *Social Science Research*, vol. 9(3), p. 227-244. New York, Academic Press, Table 3.

TABLE 18-3 χ^2 Measures of Fit for Four Models Applied to Age-specific Fertility Rates, American White Women, Aged 15-44 Years, Cohorts Born 1905-1926

| Model | <i>df</i> | χ^2 |
|-----------------------|-----------|----------|
| Age x period | 580 | 96.1 |
| Age x cohort | 609 | 730.4 |
| Period x cohort | 588 | 15,846.8 |
| Age x period x cohort | 560 | 87.6 |

Note: Data were in age-by-cohort format. χ^2 estimate is based on an artificial case base of 660,000 woman-years.

Source: Pullum Thomas W., 1980. Separating age, period, and cohort effects in white U.S. fertility, 1920-1970. *Social Science Research*, vol. 9(3), p. 227-244. New York, Academic Press, Table 4.

Despite the great theoretical appeal of the notion that continuities exist in the behavior of cohorts, we have found that the explanatory gain per cohort parameter is far less than the gain per period parameter. . . . The implications for our understanding of U.S. fertility are that, as a set, changes in those temporal variables which cut across cohorts, such as economic cycles, appear to be more important than changes in those variables which distinguish cohorts, such as shared socializing experiences.

Similar explorations of fertility patterns for other populations and time periods seem warranted. Pullum also suggested that these models could be usefully applied to fertility rates by parity.

b. French Male Mortality, 1946 to 1981

Wilmoth (1990) fit an APC model with two age-period interaction terms to age-specific mortality rates for French males during years 1946 to 1981. The full model, as stated earlier, was as follows:

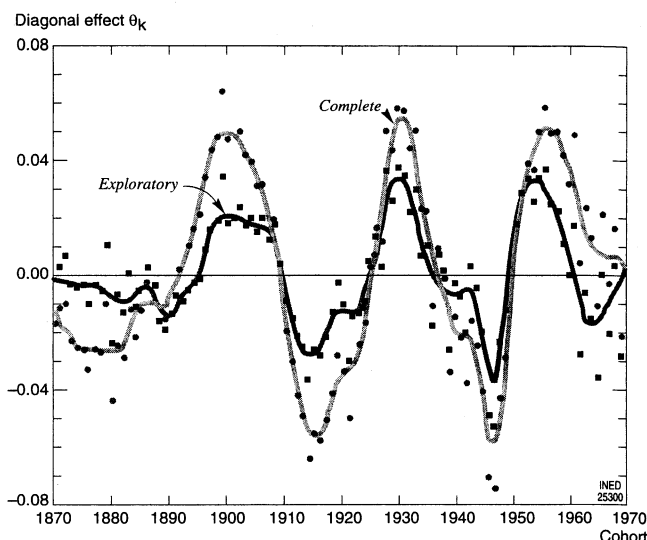


FIGURE 18-2 Residual cohort effects estimated by two methods in age-period-cohort model with interaction terms applied to age-specific mortality rates, French males, aged 0-89, years 1946-1981. Note: Data analyzed were in age-by-period format. The values depicted here are measures of the average level of excess or deficit mortality for each cohort (as a proportion of predicted levels) over the observation period. (Source: Wilmoth [1990], p. 307).

$$f(r_{ijk}) = \mu + \alpha_i + \beta_j + \sum_{m=1}^p \phi_m \gamma_{im} \delta_{jm} + \theta_k + \varepsilon_{ijk}.$$

The interaction terms in this model describe variations in the pattern of mortality decline by age: The first term captures the relatively faster decline at younger compared to older ages, while the second term reflects the pattern of increasing (or slowly decreasing) levels of mortality in late adolescence and early adulthood.

The cohort parameters in this model, θ_k , reflect residual levels of mortality relative to the thorough description by age and period. As shown here in Figure 18-2, Wilmoth's model can be used to document the persistence of relatively high or low levels of mortality for certain cohorts. Two sets of estimates for θ_k were derived using different methods for fitting the model, though the *complete* series is considered more reliable.

Wilmoth demonstrated the presence of substantial levels of excess or deficit mortality for various cohorts, successfully rejecting alternative explanations, for example, that the patterns were artifacts of bad data or of the model itself. Nevertheless, his explanation of their historical or biologic causes was less satisfactory. He acknowledged the inadequacy of existing explanations of these patterns when a comparison of results from different populations is made.

For example, it had been suggested that excess mortality for groups of male cohorts born around 1900 and 1930 may be related to early combat experiences near the end of the World Wars or to nutritional deprivation in adolescence resulting from the social and economic dislocations at these times (see Horiuchi, 1983; Wilmoth *et al.*, 1989). However, these explanations fail to account for the presence of similar patterns of excess mortality among Japanese cohorts, both male and female, born around 1900, since Japan was only nominally involved in World War I (see Wilmoth, 1988, p. 107–113). Furthermore, the presence of similar patterns for females suggests that combat experience in the wars may not be the crucial explanatory variable. Finally, the apparent recurrence of a similar pattern for cohorts born after World War II suggests that the phenomenon may operate through mechanisms at least partially unrelated to the two wars.

Although the failure to explain these findings does not invalidate them, it does illustrate the difficulties of interpreting cohort-specific variations in demographic rates.

CONCLUSION

The studies by Pullum (1980) and Wilmoth (1990) illustrate that valid results *can* be derived within an APC framework if one accepts that there is no magic solution to the *identification problem*. Instead, it is necessary to seek findings that are invariant to the choice of identifying constraints (e.g., measures of fit, second differences) or that acknowledge the fundamental arbitrariness of these constraints (e.g., residual cohort effects with zero slope).

It is important to remember that the identification problems affecting APC analysis are not some statistical aberration. Rather, they reflect a fundamental lack of information in the data being analyzed. The first identification problem, involving the level of the various parameters, is not problematic in most situations, since our purpose in estimating these sorts of models is to analyze patterns of change over age or time, which are unaffected by an arbitrary choice of level. Furthermore, this kind of identification problem is familiar and affects other statistical methods as well.

The unusual aspect of APC analysis is that, even for the simplest model, identifying constraints are required in order to obtain unique estimates of the linear trend in the parameters. This situation arises directly from the relationship linking the three sets of causal factors:

$$\text{cohort} + \text{age} = \text{period}$$

In regression analyses, of course, it is standard practice to avoid using independent variables that are highly collinear (either individually, or in some combination). By their very nature, some variants of APC analysis ignore the sound logic of this practice and attempt, by some statistical trick, to perform the impossible, namely, to extract more information than what is contained in the data. Because of this fundamental lack of information, the results of an APC analysis reflect both the underlying patterns in the data and the assumptions adopted by the analyst. This reality does not invalidate the method in all situations, but it serves as an important reminder of the limitations that affect all models used with the purpose of separating age, period, and cohort effects.

References

- CASELLI Graziella, CAPOCACCIA Riccardo, 1989. Age, period, cohort and early mortality: an analysis of adult mortality in Italy. *Population Studies* vol. 43(1), p. 133–153.
- CLOGG Clifford C., 1982. Cohort analysis of recent trends in labor force participation. *Demography* vol. 19(4), p. 459–480.
- EASTERLIN Richard A., 1961. The American baby boom in historical perspective. *American Economic Review* vol. 51(5), p. 869–911.
- EASTERLIN Richard A., 1978. What will 1984 be like? Socioeconomic implications of recent twists in age structure. *Demography*, vol. 15(4), p. 397–421.
- ELO Irma T., PRESTON Samuel, 1992. Effects of early-life conditions on adult mortality: A review. *Population Index* vol. 58(2), p. 186–212.
- FIENBERG Stephen E., MASON William M., 1979. Identification and estimation of age-period-cohort models in the analysis of discrete archival data, *Sociological Methodology*, vol. 10, San Francisco, Jossey-Bass, p. 1–65.
- FIENBERG Stephen E., MASON William M., 1985. Specification and implementation of age, period, and cohort models, in: William M. Mason and Stephen E. Fienberg (eds.), *Cohort analysis in social research*, p. 45–88. New York, Springer-Verlag, 408 p.
- FOSTER Andrew, 1990. Cohort analysis and demographic translation: a comparative study of recent trends in age-specific fertility rates from Europe and North America. *Population Studies*, vol. 44(2), p. 287–315.
- GLENN Norval D., 1976. Cohort analysts' futile quest: statistical attempts to separate age, period, and cohort effects. *American Sociological Review*, vol. 41, p. 900–904.
- GOLDSTEIN Harvey, 1979. Age, period and cohort effects. A confounded confusion. *Bulletin in Applied Statistics*, vol 6(1), p. 19–24.
- HOBSCRAFT John, MENKEN Jane, PRESTON Samuel H., 1982. Age, period, and cohort effects in demography: a review. *Population Index*, vol. 48(1), p. 4–43.
- HOLFORD Theodore R., 1991. Understanding the effects of age, period, and cohort on incidence and mortality rates. *Annual Review of Public Health* vol. 12, p. 425–457.
- HORIUCHI Shiro, 1983. The long-term impact of war on mortality: old-age mortality of First World War survivors in the Federal Republic of Germany. *United Nations Population Bulletin* vol. 15, p. 80–92.

- MASON William M., FIENBERG Stephen E. (eds.), 1985. *Cohort analysis in social research: beyond the identification problem*. New York, Springer-Verlag, VIII + 400 p.
- PULLUM Thomas W., 1980. Separating age, period, and cohort effects in white U.S. fertility, 1920-1970. *Social Science Research*, vol. 9(3), p. 227-244. New York, Academic Press.
- RYDER Norman B., 1964. The process of demographic translation, *Demography*, vol. 1(1), p. 74-82.
- RYDER Norman B., 1965. The cohort as a concept in the study of social change. *American Sociological Review*, vol. 30(6), p. 843-861.
- RYDER Norman B., 1980. Components of temporal variations in American fertility, in: Robert W. Hiorns (ed.), *Demographic patterns in developed societies*, p. 15-54. Londres, Taylor and Francis, 208 p.
- WILMOTH John R., 1988. *On the statistical analysis of large arrays of demographic rates*. Princeton, Department of Statistics, Princeton University, 158 p. (Ph.D.Thesis).
- WILMOTH John R., 1990. Variation in vital rates by age, period, and cohort, *Sociological methodology*, vol. 20, ed. by Clifford C. Clogg, p. 295-335. Oxford, Basil Blackwell.
- WILMOTH John R., VALLIN Jacques, CASELLI Graziella, 1989. Quand certaines générations ont une mortalité différente de celle que l'on pourrait attendre. *Population* vol 44(2), p. 335-376. (Also in English: When does a cohort's mortality differ from what we might expect?, *Population: An English Selection*, vol. 2, 1990, p. 93-126.)