Event History Modeling of World Fertility Survey Data [1]

Adrian E. Raftery
University of Washington

Steven M. Lewis
University of Washington

Akbar Aghajanian
Fayetteville State University

Michael J. Kahn
St. Olaf College

May 21, 1992
Revised: October 25, 1994

[1]This research was supported by NIH grant no. 5R01HD26330. An earlier version of this paper was presented to the 1992 Annual Meeting of the American Sociological Association in Pittsburgh, August, 1992, and we are grateful to Alberto Palloni for his very useful written comments prepared for that meeting. Andrew Gross did a great deal of the work necessary to bring the original data to a form in which it could be analyzed, and pointed out the good fit of the cubic polynomial to the age effect. Charles Hirschman has provided leadership and insight throughout the project, and useful comments on this paper. We are grateful to Rod Little and Diane Lye for helpful discussions, and to the Editor and two anonymous referees whose comments considerably improved the manuscript.

Abstract

Event history analysis seems ideally suited for the analysis of World Fertility Survey (WFS) data, which consists of full birth histories and related information. However, it has not been much used for this purpose, and most analyses of WFS data have consisted of tabulations of standard fertility rates, and regressions with children ever born as the dependent variable, both of which have disadvantages. We suggest that this is because event history analysis has practical drawbacks for WFS data, even though, in principle, it provides a superior analytic framework. These are the many partial dates, the computational burden of discrete-time event history analysis, the need to take account of five clocks at once (age, period, cohort, time since last event, and parity), and the difficulty of interpreting the coefficients.

We propose a modeling strategy for the event history analysis of WFS data which aims to overcome these problems, and we apply it to the previously unanalyzed WFS data from Iran. This yields estimates of the time of onset of fertility decline and of the extent to which it was due to compositional changes in the population, and it also enables us to determine whether it was a period effect, a cohort effect, or both. These results would have been hard to obtain using other approaches. In addition, the usefulness of ACE as an exploratory tool for determining the best coding of independent variables is illustrated.

Contents

List of Tables

List of Figures

List of Acronyms

ACE	Alternating Conditional Expectations
BIC	Bayesian Information Criterion
CEB	Children Ever Born
EHA	Event History Analysis
IFS	Iran Fertility Survey
PPR	Parity Progression Ratio
TFR	Total Fertility Rate
WFS	World Fertility Survey

1 Introduction

1.1 The World Fertility Survey

The World Fertility Survey (WFS) was a set of individual surveys carried out in 42 developing countries between 1972 and 1984. The data collected included full retrospective pregnancy histories, marital histories, contraceptive use data and socioeconomic characteristics. The WFS is perhaps the largest and certainly one of the most analyzed social surveys ever carried out. It may well be the richest and most studied source of event history data ever collected: the number of studies using the data is huge, exceeding one thousand by 1984 according to Cleland and Verma (1989), who gave an overview of WFS methodology. Initial summaries of the results of the WFS were provided in the volumes edited by Bulatao and Lee (1983) and Cleland and Scott (1987).

Most of the initial analyses, carried out by WFS staff, were descriptive and consisted of tabular displays of standard demographic fertility rates (e.g. Verma, 1980). Many of these were contained in the "First Country Reports" issued by the WFS for each country except Iran.[1] Many later analyses have consisted of individual-level linear regressions with children ever born (CEB) as dependent variable; we refer to this as CEB regression.

The form of the data suggests that event history analysis (EHA) would have advantages over the other two methodologies, but it has been less used in the analysis of WFS data. Why is this? We suggest that, although EHA provides an analytic framework that is in principle superior to the others, it also has practical disadvantages in terms of both implementation and interpretation that have discouraged its use. These include the frequent unavailability or unreliable reporting of the exact times of events, the heavy computational demands of discrete-time EHA, the need to take account of the five clocks that are running in parallel (age, period, cohort, time since last birth and parity)[2], and the relative lack of interpretability of the coefficients that it yields. A further problem is how to code ordered categorical independent variables such as education or size of place of residence; this arises with CEB regression also.

Here we describe a modeling strategy for the event history analysis of WFS data, which aims to resolve these difficulties. We apply it to the analysis of the Iran Fertility Survey (IFS),

[1] A First Country Report for Iran is now being prepared (Aghajanian, 1993).

[2] Here we use the word "period" to mean the current date, namely the time since a fixed baseline that is the same for every woman and every birth. However, according to Pullum (1987), the WFS staff used the word "period" to mean the time elapsed since the last birth. Here we use the word "duration" to mean the time elapsed since the last birth.

which was carried out in 1977 as part of the WFS, but which has not been analyzed before because of the revolutionary situation in Iran. We try to show that our approach yields several results that would have been hard to obtain using other methods. These include good estimates of the time at which fertility decline started, the extent to which it was due to changes in the composition of the population in terms of education and urbanization, and whether it was a period effect, a cohort effect, or both. This provides support for the comments of Singer (1989), who emphasized the potential payoff from improved EHA methods for WFS data.

1.2 Difficulties with Children Ever Born Regression

Secondary analyses of WFS data have focused on the estimation of theoretically derived equations or the testing of specified hypotheses about fertility. Regression with children ever born as the dependent variable has been much used with WFS data to determine the proximate determinants of fertility, especially by Bongaarts (1978, 1982), who used aggregate level decompositions rather than individual level data. CEB regression has also been much used with individual level data to assess the effects of socioeconomic variables on fertility, in the tradition of Easterlin and Crimmins (1985). It was also the basis of the multilevel modeling that has been carried out (Entwisle and Mason, 1985).

There are several difficulties with this much-used method. Perhaps the most sophisticated development of CEB regression is that of Little (1978, 1980, 1981), as illustrated by Little and Perera (1981). This was a major step forward compared with previous uses of the approach, and has been a model for many subsequent applications. However, even if Little's advice is followed, CEB regression still has several drawbacks. These include the facts that (i) the need to deal with censoring can lead to a loss of data or distorted results; (ii) period and cohort effects cannot be separated; and (iii) if the effects of covariates are parity-specific, they may be diluted or suppressed. There are also several other, more minor problems.

We now discuss the difficulties with CEB regression in more detail.

(i) **Censoring:** Many of the birth histories are censored by the interview. There are two main ways to deal with censoring in the context of CEB regression. One is to try to control for exposure by introducing an independent variable to represent it, such as duration of marriage or age. The expected value of CEB at a given duration or age is the integral of the age-specific fertility rates up to the current age. Thus, since fertility varies with age, the dependence of CEB on the exposure variable is nonlinear.[3] In practice, it often happens that the exposure variable explains far more of the variance in the dependent variable, CEB, than any other independent variable. Thus, failure to accurately specify the nonlinear function of exposure to which CEB is linearly related may well mislead.

Little (1981) assumed that this function was quadratic, an assumption that has been adopted by many subsequent authors. This implies that age-specific fertility is a linear function of age.[4] Our results below in Section 3.1 show that this was far from being the case in Iran; see especially Figure 1. Pullum (1987, p. 656) has pointed out several other difficulties with Little's (1981) much-used formulation.

The second way of dealing with censoring is to redefine the dependent variable as children ever born by age x and to restrict the analysis to women aged x or more. This can lead to the loss of much data and yields no information about later fertility (after age x) or younger cohorts (younger than x), which may be of considerable interest.

(ii) **Period and cohort effects:** In CEB regression, it is impossible to separate period and cohort effects. If all women are included and an exposure variable used, then period is a constant equal to the date of the survey, and an estimated cohort effect could be made up of cohort effects alone, cohort and period effects, or period effects alone. If the analysis is restricted to births by age x to women aged x or more at the time of the survey, age is a constant and period and cohort are the same, so that again it is not possible to distinguish between period and cohort effects.

The problem has often been avoided by simply omitting period effects, cohort effects, or both, from the model, as was done, for example, by Hobcraft and Casterline (1983). When this is done, it is impossible to infer anything about either period or cohort effects. Also, estimates of the effects of any other independent variables that are related to period or cohort may be biased as a result of the omitted variables.

(iii) **Effects that vary with parity:** If the effect of an independent variable varies with parity (i.e. by the number of previous births to the woman), its effect may be diluted or even suppressed in CEB regression. A case in point is the period effect in our Iran

[3] This is because the integral of any non-constant function is non-linear.

[4] This is because expected CEB is the integral of age-specific fertility. Thus age-specific fertility is the partial derivative of expected CEB with respect to age. If expected CEB is a quadratic function of age, then age-specific fertility is the derivative of a quadratic function, namely a linear function.

data, which was both strong and in opposite directions at different parities; see Section 4.

(iv) **Other difficulties:** It is not possible to take account of time-varying covariates in the CEB regression framework. Admittedly, this is not a major obstacle here, because time-varying covariates are rare in WFS data. A further problem is that the dependent variable, CEB, is not at all normally distributed. Also, the variance of the error term is not at all constant. This last point can be met by weighting the cases in the regression as recommended by Little (1981), but this seems often not to have been done in practice. Also, Pullum (1987, p. 656) has pointed out several other, more minor, difficulties, with CEB regression in its most used formulation, that of Little (1981).

1.3 The Iran Fertility Survey

We now briefly describe the context of the 1977 Iran Fertility Survey (IFS), and the data it produced. We will use these data in the rest of the article to illustrate our methodology.

Most of the post-World War II period in Iran was characterized by rapid economic growth fuelled by oil exports, and rapid modernization imposed from above (Ebrahamian, 1982). Oil production was the key motor of the economy. It grew steadily during the first half of the century, but the really dramatic growth started only in 1955; by 1974 oil production had reached ten times the highest pre-1955 figure, while total economic activity increased by a factor of four in half a generation.

One would expect several of the accompanying societal changes to have led to a decline in fertility. These include increasing education, measures designed to increase women's status, the state-sponsored family planning program started in 1967, and the decline in infant mortality. Here we attempt to quantify the overall decline in marital fertility and to assess the relative importance of some of these contributing factors. For other discussion of the determinants of fertility in Iran see Aghajanian (1978, 1988), Ajami (1976), Paydarfar (1975) and Goode and Farr (1980) which refer only to data collected for local areas, and Aghajanian (1991) which analyzes census data.

Our analysis is based on data from the IFS, which was carried out in 1977 as part of the WFS. It consists of the full fertility histories of a sample of married women of whom the oldest were born in 1927. It was not fully analyzed at the time it was collected because of the revolutionary situation in Iran.

The IFS was based on a nationally representative sample of ever-married women aged less than 50. Based on a multistage random sampling procedure, 6,056 households were visited and all ever-married women who were less than 50 years old were interviewed. This procedure resulted in interviews with a total of 4,932 women. There were 42 nonresponses for an effective sample size of 4,890. There was a total of 20,346 births in 72,248 ever-married-woman-years.

In general it seems reasonable to expect high quality data from the IFS. However, in data collection in a developing country with a low female literacy rate, errors in the dates of events and the number of children ever born are possible. Estimates of the levels and trends of fertility depend on the extent to which women forget to mention births, misreport the dates of births or misreport their own ages.

The quality of the IFS was comprehensively assessed by Aghajanian, Gross and Lewis (1993), who concluded that the survey data are of good quality, at least as good and in many respects better than that of WFS data from comparable countries. Trussell (1984) indicated that the errors present in WFS data of this quality do not change the substantive conclusions from complex regression analyses. The single biggest problem is that, while respondents were asked both the month and the year in which each event occurred, the month was missing in a majority of cases, and so we have worked only with the year of occurrence.

Otherwise there were few missing data. We have excluded those women who had missing data on any of the quantities of interest to us here; for parities two and above, this led to the exclusion of only 8% of the intervals. Ultimately it would be more satisfactory to use all the data and deal with missing data via multiple imputation (Rubin, 1987). However, the amount of missing data is small and our analysis shows no systematic differences between women with and without missing data. We expect that an analysis using multiple imputation would yield results very similar to those that we report here. We have not used the single imputation methods for missing data that were routinely used by WFS staff in the analysis of other surveys. Note that, in principle, such single imputation methods can lead to the precision of estimates being overstated (e.g. Rubin, 1987), although for the IFS data this overstatement is unlikely to be substantial.

1.4 Event History Analysis for the WFS

EHA was introduced by Cox (1972) as a synthesis of regression and life tables, initially in the context of nonrepeatable events such as death. The key idea is that the dependent variable is the hazard rate, namely the instantaneous probability of an event. This resolves the main problems of linear regression for such data, namely the censoring problem and the non-normality and non-constant variance of the response. It also makes it possible to avoid the difficulties that are more specific to CEB regression for WFS data, namely the inability to separate period and cohort effects, and the dilution of effects that vary with parity. Excellent introductions to EHA for social data which put more emphasis on repeatable events are provided by Carroll (1983), Tuma and Hannan (1984), Allison (1984), Petersen (1991) and Yamaguchi (1991). EHA has been widely applied in demography; Hobcraft and Murphy (1986) provide a review, including references to some studies with WFS data. Many of the relatively few EHA analyses of WFS data to date have been of limited scope; see Section 5.2.

While EHA appears to overcome the inadequacies of CEB regression, it also has several features that have made its application to WFS data difficult in practice. The main difficulties we met in applying EHA to the WFS data were as follows.

Partial Dates: Overall, the IFS data are of good quality and there are few missing responses. The biggest problem with the data is that for almost two-thirds of the respondents, only the year in which a birth occurred was reported; only about one-third reported the month of birth. Standard, continuous-time, EHA depends on full information about the time of events, and we therefore used discrete-time EHA (Allison, 1984) with the (Persian) calendar year as the interval. This led to the additional problem that, contrary to the usual assumption of discrete-time EHA, the occurrence of an event does not end the exposure for that interval, since a woman can give birth more than once in the same calendar year. Our solution to this is described in Section 2.1.

Computer Demands of Discrete-Time EHA: As usually applied, discrete-time EHA requires the creation of a computer file with one record per woman-year of exposure containing the relevant outcomes and covariates. In the IFS data, this contains up to four million numbers and makes estimation computationally expensive. We wrote a special computer program that overcomes this problem.[5]

The Five Clocks: Age, Period, Cohort, Duration and Parity: Each woman-year of exposure is characterized by five time-like variables: age, period, cohort, duration (i.e. time since the last event), and parity. It is essential to take account of all these clocks; the need to do so is a difficulty, but the ability to do so is a strength, since it allows us to disentangle these effects. We are not aware of previous work on age-period-cohort modeling of event

[5]This program, which is written in Fortran, is available on request from the authors by sending e-mail to *raftery@stat.washington.edu*.

history data; the need to deal, in addition, with duration and parity is a feature of event history data, but not of archival data such as that considered by Fienberg and Mason (1978), for example.

We solved the identification problem by parametric modeling of those components for which there is a reasonable basis for assuming smooth effects, namely age, parity (for parities two and above), and duration (for durations three and over). Period and cohort effects were represented by dummy variables without constraints. We allowed the data to suggest the parametric forms, and we found that simple forms fit the data well. This worked well enough that we did not consider the more elaborate approach of representing the age effect by model fertility schedules (Coale and Trussell, 1974; Johnson, 1985). There have been many previous efforts to fit parametric forms to the age component alone (e.g. Farid, 1973; Duchêne and Gillet-de Stefano, 1974).

Our approach to the age-period-cohort problem is "midway" between two others that have been discussed in the literature. That of Mason *et al.* (1973) and Fienberg and Mason (1978) involves imposing the least restrictive single constraint possible to achieve identifiability; this does not exploit the fact that several of the components vary smoothly. The approach of Nakamura (1986), applied by Sasaki and Suzuki (1987), assumes that *all* components vary gradually, although no specific parametric forms are imposed. As pointed out by Glenn (1989a,b), this may well not be the case, and in our data it is important to allow for possible abrupt changes in the period and cohort effects. A different, exploratory, approach, has been described by Wilmoth (1990). For a review of other demographic applications of age-period-cohort modeling, see Hobcraft, Menken and Preston (1982).

Coding Ordered Categorical Variables: Our independent variables include several ordered categorical variables with around 6-8 categories. Should these be coded as interval or nominal variables? We answer this question by a preliminary exploratory application of the ACE technique (Breiman and Friedman, 1985; see DeVeaux, 1989 for an exposition). This suggests transformations of the original coding that fit the data well and are substantively meaningful. These are then used in the formal estimation.

Interpreting EHA parameters: Standard demographic fertility rates are readily interpretable, as are coefficients from CEB regressions. The parameters from EHA models, by contrast, are not, which is a real disadvantage of EHA compared to the other approaches. We suggest a simple way of approximately translating EHA parameters to changes in the Total Fertility Rate. Rodriguez and Cleland (1981) and Hobcraft and Casterline (1983) have done this for subgroups that share the same covariate values, while Palloni (1985) suggested a more complex approach.

2 Event History Modeling

2.1 Discrete Event History Analysis

Our data are in the form of event histories and so event history analysis (e.g. Tuma and Hannan, 1984) is the method of choice. However, we are working with discretized data which include only the calendar year of each event. Since a year is long relative to typical inter-event (mostly inter-birth) intervals, this leads to many intervals being of the same length, and so the standard approximations for dealing with ties in the usual continuous-time event history analysis (e.g. Cox and Oakes, 1984) break down.

We therefore considered a discrete-time event history analysis method (Allison, 1984). In this approach, each woman-year of exposure is treated as a separate case, with the response being the occurrence or not of an event in that year; the baseline hazard rate and the effects of covariates are modeled by logistic regression. However, this approach cannot be used without modification here because a woman may give birth on two separate occasions in the same year, and ignoring this would bias the results, leading to underestimation of both overall fertility and fertility differentials.

Our modified discrete event history analysis was carried out as follows. A file was created in which each woman-calendar year was a single case if she did not give birth that year, two cases if she gave birth once, and three cases if she gave birth on two separate occasions. Duration t was coded as 0 if she had already given birth that year, 1 if the previous birth had occurred in the preceding year, and so on. The first interval for each woman started with her first marriage and for that interval duration was coded as 0 in the year of marriage, 1 in the following year, and so on. We assumed that women were not at risk of giving birth before marriage.

The resulting data set was modeled by the logistic regression model

$$\begin{aligned} \text{logit}(\pi_{ity}) &= \log\left(\frac{\pi_{ity}}{1-\pi_{ity}}\right) \\ &= \beta_0 + \sum_{j=1}^{p} \beta_j x_{jity}, \end{aligned} \tag{1}$$

where π_{ity} is the probability of a birth to the ith woman at duration t in calendar year y, x_{jity} is the jth covariate for the ith woman at duration t in calendar year y, and $\beta_0, \beta_1, \ldots, \beta_p$

are unknown regression coefficients. This model was estimated by maximum likelihood. We used a specially written Fortran program which both created the data set and also did the estimation without actually writing the data set to a file. This avoided some of the computer storage problems associated with the manipulation of such a large data set (typically up to about four million numbers) and made it possible to analyze it using a desktop workstation.

The covariates considered for inclusion in equation (1) include functions of age, duration t, parity, birth cohort, calendar year y (i.e. period), and individual characteristics of the woman such as her education and that of her husband, where she grew up and where she lived at the time of interview, and her age at marriage.

We therefore have the classic demographic problem of modeling age, period and cohort, with, in addition, duration and parity. Thus there are, in a sense, five clocks running simultaneously. We deal with this by exploiting the longitudinal nature of the data, and also by modeling these effects parametrically whenever it makes theoretical sense and is well supported by the data. This breaks the formal identities that can make estimation impossible. Any remaining near collinearities will lead to high correlations between parameter estimates and hence will be readily detectable. They will lead to inflated standard errors, but not biased estimates, and this will be taken account of in inference.

We found that the effects of individual characteristics were quite different for the lower parities than for the higher ones, and so we fitted three separate versions of equation (1), one each for parity 0, parity 1, and parities 2 and higher.

We based model comparison on the BIC statistic (Raftery, 1986a, b), in the form

$$BIC = -\chi^2 + p \log n, \tag{2}$$

where χ^2 is the likelihood ratio test statistic for comparing the null model with no covariates with the model of interest, p is the number of independent variables in the model of interest as defined by equation (1), and n is the sample size, i.e. the number of cases (woman-years) in the logistic regression (1). With this definition, the *smaller* BIC is, the better the model. For the model with no covariates, BIC is zero, so a positive BIC indicates a model that is worse than the null model.

One rule of thumb for the interpretation of BIC is as follows (Raftery, 1993a). Suppose we wish to compare a model M_1 with a larger model M_2 within which it is nested, where BIC_1 and BIC_2 are the BIC values for the two models. If $\Delta = \mathrm{BIC}_2 - \mathrm{BIC}_1$ is positive there is no evidence for the additional effect included in M_2, if Δ is between -5 and 0 there is evidence for the additional effect but it is weak, if Δ is between -10 and -5 the evidence is strong but not conclusive, while if Δ is less than -10 the evidence is very strong. The BIC statistic is justified as an approximation to twice the logarithm of the Bayes factor in favor of the model of interest against the null model; see Raftery (1988, 1993a, b) and Kass and Raftery (1993) for justification and evaluation of the quality of the approximation.

2.2 Translating Event History Parameters into Total Fertility Rates

EHA is very well suited for modeling WFS data and representing theoretical hypotheses about it. However, the magnitudes of the estimated parameters that it yields are not easy to interpret directly in terms of more standard and familiar demographic rates and quantities, such as age-specific fertility rates and the Total Fertility Rate (TFR), although the signs of the parameters do have a direct interpretation. Here we describe a simple approximate way of "translating" event history parameters into quantities that are more familiar to demographers.

An approximate way of translating the event history parameters for our models into Total Fertility Rates (TFRs) is as follows. Let $\bar{\pi}$ be the average value of π_{ity} as defined by equation (1), $\bar{L} = \text{logit}(\bar{\pi})$, and $\bar{f}$ be the fertility rate corresponding to $\bar{\pi}$. Note that $\bar{f}$ and $\bar{\pi}$ are not the same, because years in which there was a birth are counted more than once in the denominator of $\bar{\pi}$, but not of $\bar{f}$. We have $\bar{\pi} = \bar{f}/(1+\bar{f})$, so that $\bar{f} = e^{\bar{L}}$. Now, approximately, for our data,

$$\text{TFR} = 2 + A\bar{f} = 2 + Ae^{\bar{L}}, \tag{3}$$

where $A = (C - B - 1)$, B is the average age at the second birth and the age-specific fertility rate for parities two and higher is approximated by

$$\text{fertility rate (age)} \approx \begin{cases} \bar{f}, & \text{if } D \leq \text{age} \leq C \\ 0, & \text{if not,} \end{cases} \tag{4}$$

(provided that $B \geq D$). From our data we have estimated $D = 15$, $C = 42$, $B = 22$, so that $A = 19$. This corresponds to the rough assumption that all women eventually have at least two children. This is not too unreasonable in the context of pre-1977 Iran, but it would have to be modified for countries with lower fertility rates. Equation (4) approximates the age-specific fertility rate by a constant rate during a given exposure period that starts at age $(B+1)$ and ends at age C.

We describe the method by example rather than by giving general equations. Let us examine the effect of the woman's education on fertility (Table 3 below). Our estimation of model (1) is based on 63,667 cases including 13,099 births, so that $\bar{\pi}$ is $13099/63667 = 0.2057$, corresponding to $\bar{L} = \text{logit}(0.2057) = -1.351$, $\bar{f} = e^{-1.351} = 0.259$, or about $\text{TFR} = 2{+}19\bar{f} = 6.92$. The average educational attainment is 0.5, and the event history parameter for this variable is -0.239. Thus, for education 0, $\bar{L} = -1.351 + (0 - 0.5) \times (-0.239) = -1.2315$, so that $\text{TFR} = 2 + 19e^{-1.2315} = 7.55$. For education 5, $\bar{L} = -1.351 + (5 - 0.5) \times (-0.239) = -2.426$, so that $\text{TFR} = 2 + 19e^{-2.426} = 3.68$. Thus the average difference in TFR between women with some higher education and those with no education was about $7.55{-}3.68 = 3.87$, or about four children.

The average effect on TFR of a change in an independent variable x whose event history parameter is β is about $A\beta e^{\bar{L}}$, by differentiating equation (3) with respect to x. For our data, this is about $19e^{-1.351}\beta = 4.92\beta$, on average. Thus we can "read" Table 3 in a rough way by multiplying the parameter estimates by about 5 to gauge the effect on TFR of a unit change in the independent variable.

3 Coding Variables Using Parametric Models and ACE

3.1 Coding Age and Duration Using Parametric Models

Age: There is a clear pattern to the relationship between fertility and age, with a rapid increase followed by a more gradual decrease; it makes sense to represent it in a parametric form if this is well supported by the data. To do so, we calculated the number of woman-years and the number of births to women of that age in our data set for each year of age below 50. Dividing the number of births by the number of woman-years gives an estimate of the average age-specific marital fertility. This is very well fit by a cubic polynomial function of age, namely

$$\text{fertility} = -.809 + 1.120a - .334a^2 + .0286a^3 \quad (R^2 = .99), \tag{5}$$

where $a = \text{age}/10$.

The raw estimated fertility rates are shown together with the fitted curve given by equation (5) in Figure 1. The good fit of the curve is clear. Not only does it capture the main features, namely the rapid increase followed by the more gradual decrease, but it also captures the slowdown in the decrease after age 40, a relatively subtle feature. This latter feature is represented by the fitted curve having a point of inflection at age 39. Since equation (1) is a linear model of probabilities on the logit scale, we take the logit of the function of age on the right-hand side of equation (5) to be our age effect.

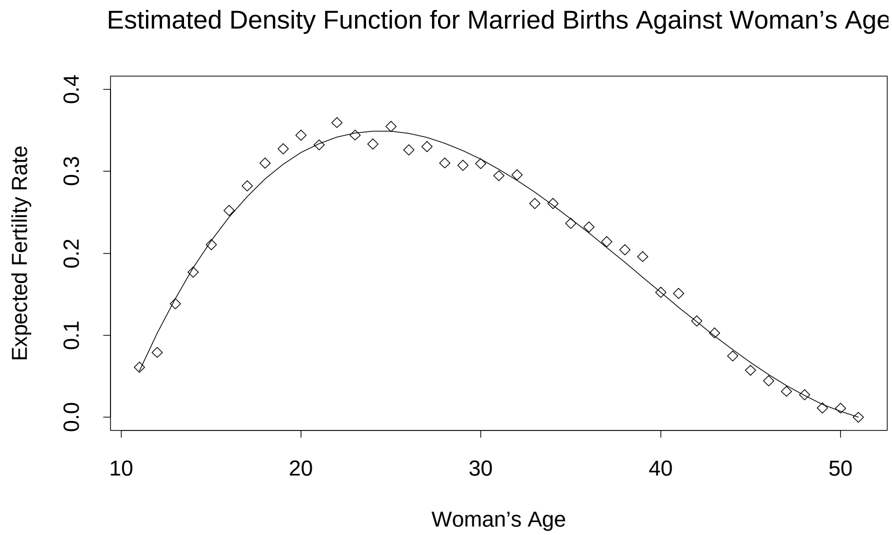

Figure 1: Estimated fertility rate with fitted cubic polynomial.

When the shape of the age effect is estimated separately in this way, it may not be fully separated from the effects of the other clocks. To check that this was not the case, we estimated the age effect directly as part of the model in two different ways. We entered a, a^2 and a^3 as independent variables in model (1). We also fitted several of the models reported later using dummy variables for five-year age categories instead of the single age correction factor reported here. These alternative specifications made almost no difference to the estimates of other parameters or to the model comparisons, and so led to the same conclusions. We therefore chose to report our results in terms of the estimated age effect (5).

What we have done here is in some ways analogous to effect-proportional scaling of categorical variables. We have replaced several polynomial terms by a pre-estimated linear combination. This linear combination takes up only one degree of freedom, rather than three, for model comparison purposes. This does not matter given the sample size here, but it could matter in a smaller data set.

Duration: We now consider the coding of duration t, namely the time since the last event. Calculating the average fertility rate by duration t, using the method just described for age (but using only parities 2 and higher), gives the results shown in Figure 2. For durations 3 and higher, the fertility rate declines almost exactly exponentially ($R^2 = .99$ in a weighted logarithmic regression); the fit is shown in Figure 3. We therefore coded duration using four covariates: one dummy variable each for durations $t = 0$, 1 and 2, and one further variable

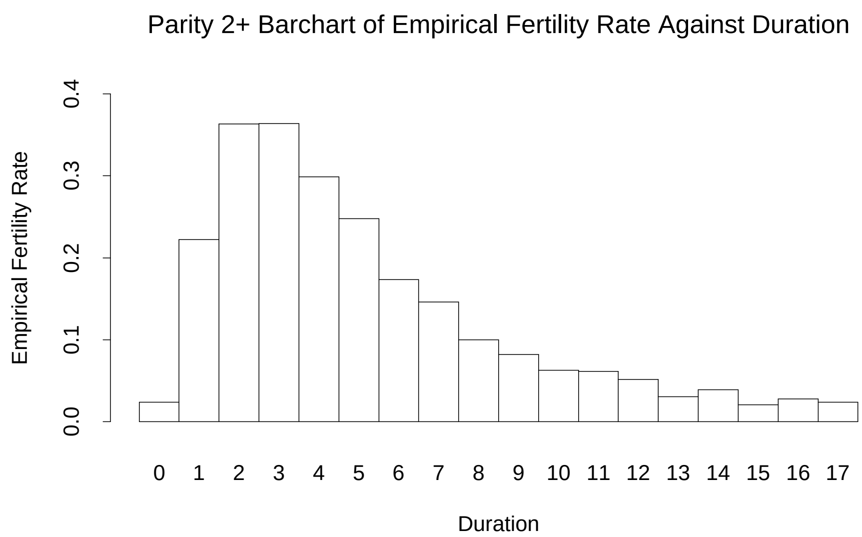

Figure 2: Estimated fertility rate by duration (parity 2+).

equal to

$$\begin{cases} 0 & \text{if } t = 0, 1, 2 \\ \text{logit}(ab^t) & \text{if } t \geq 3, \end{cases} \tag{6}$$

where $a = 0.726$ and $b = 0.795$.

3.2 Coding Other Variables Using ACE

Several of the independent variables considered are categorical but ordered; these include educational attainment and size of place of current residence. Other variables are measured on an interval scale but their effects are (or could be) quite nonlinear; these include age at first marriage. One possibility is to code these as sets of dummy variables, but this is often unparsimonious and computationally prohibitive. Instead we searched for transformations of the variables that were theoretically meaningful and well supported by the data, as follows.

We carried out ordinary least squares linear regression with each birth as a case. The length of the interval ending in that birth was the dependent variable, and the independent variables were the same as those in equation (1), with the exception of duration; calendar year was coded as the year at the beginning of the interval. The variables were not, however, coded linearly. Rather, the ACE technique was used to find the monotonic (nonparametric) transformation of each variable that made the regression as linear as possible.[6] The results

[6]This was done using the *ace* function in S-PLUS.

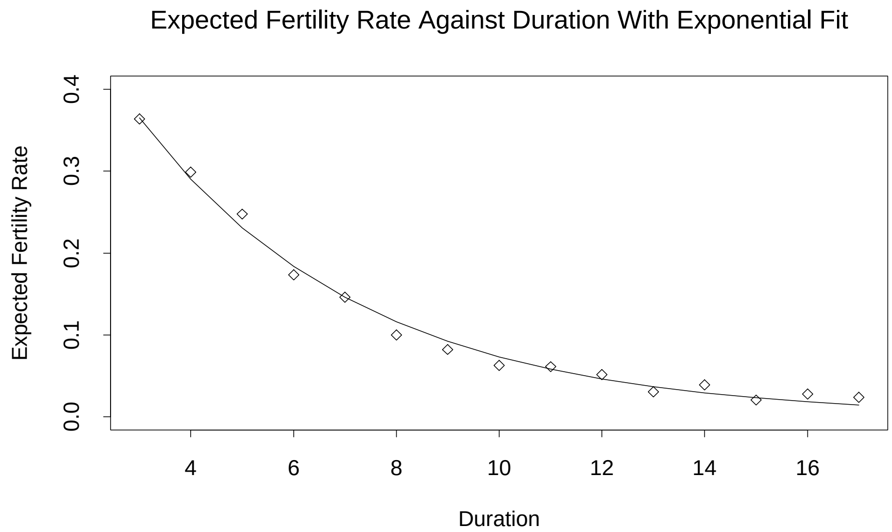

Figure 3: Raw estimate and fitted exponential curve to fertility rate at durations 3 and above (parity 2+).

were used to suggest theoretically meaningful *parametric* transformations for use in equation (1).

Note that this linear regression does not itself give valid results because all the open intervals have been ignored, a source of selection bias, and also because the response is quite nonnormal and so ordinary least squares will not yield valid inferences. However, we found ACE to be a valuable *exploratory* tool, leading to transformations of the variables that were substantively meaningful, and subsequently turned out to provide a much better fit to the data in the event history model.

Age at first marriage: The ACE transformation of age at first marriage is shown in Figure 4 for parity 0. This suggests that the effect is piecewise linear, with a greater effect of very early marriage. We therefore represented age at first marriage by two independent variables, namely

$$\begin{aligned} \text{early age at marriage} &= (K - \text{age at marriage})^+ \\ \text{late age at marriage} &= (\text{age at marriage} - K)^+, \end{aligned}$$

where $(x)^+ = x$ if $x \geq 0$ and 0 if $x \leq 0$, and K is the breakpoint. We found that $K = 12$ fitted better than the alternatives considered.

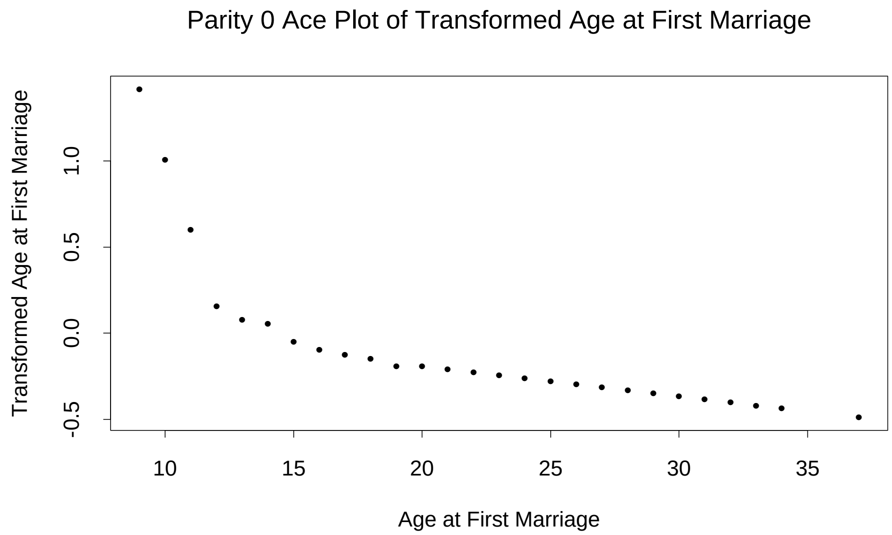

Figure 4: ACE plot for age at first marriage (parity 0).

Education: Educational attainment is coded in six categories: none (0), incomplete primary (1), complete primary (2), incomplete second level (3), complete second level (4), and higher (5). This is now a fairly standard international categorization and has been used, for example, in a set of recent studies of educational stratification in several countries (Raftery and Hout, 1993; Shavit and Bloesfeld, 1992). The ACE transformation of the respondent's education indicated that the effect was linear when coded this way and so no transformation was needed. Husband's education was coded the same way as wife's education.

Place of residence: Size of place of residence was originally coded in eight categories ranging from the largest, Tehran (1), to isolated farm dwellings (8). The ACE transformation is shown in Figure 5. This is fairly linear between 1 and 5, but there is almost no change between 5 and 8. Categories 6–8 are all rural, while category 5 consists of small towns. In fact, most of the small towns in category 5 are more rural than urban, and are really large villages with more than 5,000 inhabitants. We amalgamated categories 5–8 to form a single category. This indicates that there are no systematic differences in marital fertility between rural and small town residents in terms of the size of their place of residence, but that for city dwellers the size of the place where they live does have an effect. Place of childhood residence is coded in two categories: city (1) and village (2).

Descriptive statistics for some of the key variables are shown in Table 1.

Figure 5: ACE plot for size of place of current residence (parity 2+).

Table 1: Means and standard deviations for variables in the study.

variable	*mean*	*standard deviation*
Year of Birth	1945	9.6
Size of Place of Residence	3.8	1.6
Education	0.5	1.1
Husband's Education	0.9	1.3
Age at First Marriage	15.7	3.3
Place of Childhood Residence	1.7	0.5

Table 2: Model fits for parity 2+. The best model is shown in bold.

#	*Model*	χ^2	p	*BIC*
0	Null	0	0	0
1	A D	8994	5	−8939
2	A D P	9047	6	−8981
3	A D P C	9182	12	−9049
4	A D P Y	9308	12	−9175
5	A D P C Y	9327	18	−9127
6	A D P C Y S W H	10079	21	−9847
7	**A D P Y S W H**	**10065**	**15**	−9899

NOTE: The independent variables are as follows: A = age effect defined by equation (5); D = duration into the current interval (see text); P = parity (equal to number of previous births); C = cohort (7 levels); Y = period (7 levels); S = size of the place (5-category) in which the woman resides; W = woman's education (6 categories); H = husband's education (6 categories); The quantities χ^2, p and BIC are defined by equation (2).

4 Results

4.1 Parities 2 and Higher

The modeling results for parities 2 and higher are shown in Table 2. Model 0 is the null model with no covariates, and model 1, which accounts for age and duration, fits much better. Model 2 accounts for parity also; while this fits better, parity did not account for much of the variation.[7] Model 3 includes cohort effects; this provides an improved fit according to BIC, but the corresponding parameter estimates are not significant.[8] Model 4 includes the period effect but not the cohort effect, and this is substantial and significant. When the period effect is in the model, the cohort effect accounts for very little additional variability (model 5).

Model 6 introduces the individual characteristics, namely the woman's education, that of her husband, and the size of the place where she lives; these are all substantial and highly significant.[9] Finally, in model 7, the cohort effect is removed. Model 7 is the preferred one,

[7] Parity is coded linearly as the number of previous births. We also fit a model that coded parity as a set of dummy variables, but the linear coding accounts for almost as much of the variability and fits much better according to BIC.

[8] Cohort could also be coded linearly and this would provide a better fit in model 3 according to BIC. However, our final conclusion is that the cohort effect is not significant, and so we have not reported on our efforts to refine the coding of cohort.

[9] We also tested the effect of age at marriage, place of childhood residence, premarital work and post-

confirming that the fertility decline is a period effect and not a cohort effect.[10]

This finding is not without precedent in the literature. Fertility decline has been found to be a period effect and not a cohort effect elsewhere also: in the U.S. by Pullum (1980) and Isaac *et al.* (1979) as reported by Hobcraft, Menken and Preston (1982), and in Sweden, England and Wales, and Australia by Page (1977).

The parameter estimates for the preferred model are shown in Table 3. The estimated period effect is graphed in Figure 6 along with an approximate confidence interval. Fertility started to decline in the first half of the 1960s, and the decline accelerated markedly at the beginning of the 1970s.[11] From 1961 to 1977 the logit of the probability of giving birth in a given year declined by 0.43; this corresponds to a decline in the marital total fertility rate (TFR) of about 2.1, from 8.1 to 6.0. Using data from the 1966 and 1976 Censuses of Iran, Aghajanian (1991) calculated that the marital TFR declined by 2.0, from 8.3 to 6.3, between 1966 and 1976, which is in good agreement with our results here. Note, however, that the present method and data give a much more precise idea of the *timing* of the decline than is possible with Census data.

The effects of education, husband's education and size of place of current residence are all large and highly significant. Each unit increase in the woman's education decreases the logit of the probability of a birth by 0.24. As we saw in Section 2.2, this implies that the difference in TFR between women with no education and those with some higher education is about four children. The effect of the husband's education was about one-third as large as that of the woman's own education. The effect of place of current residence was also substantial, but less so than that of education. The average TFR for women in rural areas was about 7.4 and that for women in Tehran was about 5.9, a difference of $7.4 - 5.9 = 2.5$.

Is the decline in fertility a compositional effect, due to the increasing proportion of educated women? The answer is no. The effect of the educational composition of the

marital work, but these did not have significant effects. The nonsignificance of age at marriage is not too surprising since here we are considering only parities 2 and higher.

[10]Model 7 is preferred to model 6 according to BIC, but the P-value from a χ^2 likelihood ratio test is about 0.03. Thus the situation is akin to that of Grusky and Hauser (1984), discussed by Raftery (1986b), in which the data set was very large and a model that appeared to fit the data very well in most ways was nevertheless rejected by a conventional likelihood ratio test but favored by BIC. Our decision to reject model 6 is reinforced by the fact that the estimated cohort effects do not differ significantly from one another, and that they exhibit no clear pattern. The sequence of between-cohort fertility changes estimated in model 6 is as follows: down, up, down, down, up, down; we can think of no plausible interpretation of such an erratic pattern. We also fit models that included interactions between period and age, period and duration, and period and parity, but there was no evidence for the existence of any of these interaction effects.

[11]The estimated increase up to the early 1950s is not significant.

Table 3: Parameter estimates for the preferred model for parity 2+.

Variable	β	*Standard error*	t
Intercept	-2.856	–	–
Age Effect	1.836	0.109	16.87
Duration 0	-1.998	0.072	-27.81
Duration 1	0.521	0.053	9.77
Duration 2	1.281	0.053	24.16
Duration 3+	2.534	0.106	23.92
Parity	-0.058	0.006	-10.04
Before 1948	0.114	0.214	0.53
1948-1952	0.320	0.080	3.98
1953-1957	0.410	0.047	8.75
1958-1962	0.426	0.036	11.92
1963-1967	0.298	0.031	9.57
1968-1972	0.301	0.029	10.53
Size of Place of Residence	0.085	0.008	10.60
Woman's Education	-0.239	0.021	-11.40
Husband's Education	-0.086	0.013	-6.52

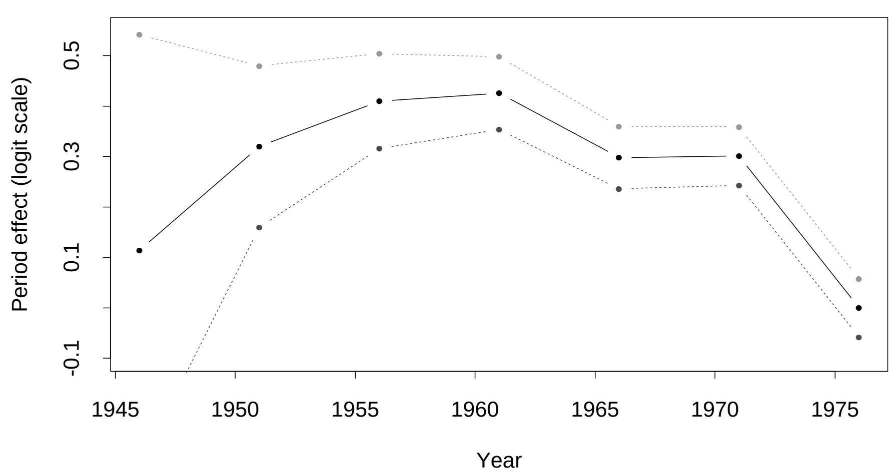

Figure 6: Period effects with approximate 95% confidence intervals, based on the preferred model for parities 2 and higher.

population would manifest itself as a cohort effect and not a period effect. In model 3 of Table 2, which does not include the individual characteristics, the estimated fertility decline was 0.48, compared with 0.43 when the individual characteristics are included. Thus the individual characteristics may account for 10% of the fertility decline, but even this is doubtful since the decline of 0.05 is only one and a half times the standard error of the estimate of the decline. Note that this does not exclude the possibility that the fertility decline was due in part to the *contextual* effect of a more highly educated population; this *would* manifest itself as a period effect.

4.2 Parities 0 and 1

Table 4 shows the models fit for parity 0, i.e. the interval from first marriage to first birth, and Table 5 shows the parameter estimates for the preferred model. The results are rather different from those for parities 2 and above. As before, there is a strong period effect and no cohort effect. However, the period effect is in the direction of a shortening of the interval, which would lead, all other things being equal, to an *increase* in fertility. This is the opposite of what was observed for parities 2 and above. The logit increased by 0.95 over the time span, which means that the average reported time from marriage to conception decreased by over two-thirds, from about 4.5 to 1.3 years.

This is based on a rough calculation which goes as follows. There were 3,590 births for 15,251 woman-years of exposure at parity 0, i.e. an average probability of birth of 0.235, whose logit is -1.18. The logit increased by 0.945, about from -1.653 to -0.708, so that the probability of a birth increased from 0.16 to 0.33. Thus the average interval went down from 6.25 to 3.03 woman-years of exposure. However, both the first and the last year of each interval contribute only about half a year on average to the interval, and the interval includes nine months of gestation, so that these numbers have to be reduced by 1.75, to 4.50 and 1.28 respectively.

We cannot determine the reasons for this dramatic change. It may be partially due to misreporting or measurement error. However, it may also be due in part to a decline in the proportion of arranged marriages, and to the gradual increase over time in the average age at first marriage. Hirschman (1985) obtained changes of a similar magnitude for several Asian societies. Rindfuss and Morgan (1983) argued that the dramatic decline in the first birth interval in Asia was due to the decline of arranged marriages, which led to more frequent intercourse in the early years of marriage. Such an explanation could also be valid for Iran.

Place of childhood residence has an effect, with women raised in cities giving birth about

Table 4: Model fits for parity 0. The best model is shown in bold.

#	*Model*	χ^2	p	BIC
0	Null	0	0	0
1	A D	814	5	−766
2	A D C	992	11	−886
3	A D Y	1007	11	−901
4	A D C Y	1031	17	−867
5	A D C Y S W H E M R	1082	23	−860
6	A D Y S W H E M R	1063	17	−899
7	**A D Y E M R**	**1062**	**14**	−927

NOTE: The independent variables are as follows: A = age correction effect defined by equation (5); D = duration into the current interval (see text); C = cohort (7 levels); Y = period (7 levels); S = size of the place (5-category) in which the woman resides; W = woman's 6-category level of education; H = husband's 6-category level of education; E = negative of woman's age at first marriage before 12 years; M = woman's age at first marriage after 12 years of age; R = woman's type of childhood residence (1=city, 2=village). The quantities χ^2, p and BIC are defined by equation (2).

Table 5: Parameter estimates for preferred model for parity 0.

Variable	β	*Standard error*	t
Intercept	-2.501	–	–
Age Effect	3.320	0.274	12.11
Duration 0	0.076	0.093	0.82
Duration 1	1.124	0.083	13.51
Duration 2	1.233	0.084	14.76
Duration 3+	1.875	0.169	11.09
Before 1948	-0.945	0.118	-8.04
1948-1952	-0.858	0.082	-10.43
1953-1957	-0.488	0.071	-6.88
1958-1962	-0.332	0.067	-4.94
1963-1967	-0.179	0.067	-2.69
1968-1972	-0.043	0.065	-0.66
Early Marriage Age	0.113	0.050	2.26
Late Marriage Age	-0.053	0.009	-6.23
Childhood Residence	-0.156	0.042	-3.69

Table 6: Model fits for parity 1. The best model is shown in bold.

#	*Model*	χ^2	p	BIC
0	Null	0	0	0
1	A D	2977	5	−2930
2	A D C	3006	11	−2901
3	A D Y	3038	11	−2933
4	A D C Y	3079	17	−2917
5	A D C Y S W H E M R	3162	23	−2943
6	A D Y S W H E M R	3124	17	−2962
7	A D Y W H M	3122	14	−2989
8	A D Y W	3117	12	−3002
9	**A D W**	**3078**	**6**	−3021

NOTE: The independent variables are coded as follows: A = age effect defined by equation (5); D = duration into the current interval (see text); C = cohort (7 levels); Y = period (7 levels); S = size of the place (5-category) in which the woman resides; W = woman's 6-category level of education; H = husband's 6-category level of education; E = negative of woman's age at first marriage before 12 years; M = woman's age at first marriage after 12 years of age; R = woman's type of childhood residence (1=city, 2=village). The quantities χ^2, p and BIC are defined by equation (2).

six months sooner after marriage, on average, than women raised in villages. Again this is different from the effect found for parities 2 and higher, where place of childhood residence was not important but place of current residence was. Also, at parity 0 we see city women having higher fertility, whereas at higher parities it was the other way around.

The models and parameter estimates for parity 1 are shown in Tables 6 and 7. The preferred model included only age, duration and woman's education; there is no evidence of a substantial period effect. The effect of woman's education is in the same direction as for higher parities, with more education corresponding to lower fertility, but the effect is weaker. Thus, the model for parity 1 is "between" that for parity 0 and that for parities 2 and higher. The differences between individuals and across time in fertility at parity 1 are small. This probably reflects the fact that once there has been a first child, few efforts are made to avoid or delay the birth of a second child. Important differences do not start to appear until after two children have been born.

Table 7: Parameter estimates for preferred model for parity 1.

Variable	β	*Standard error*	t
Intercept	-2.125	0.139	-15.30
Age Effect	1.360	0.227	5.99
Duration 0	-2.435	0.161	-15.14
Duration 1	0.740	0.110	6.71
Duration 2	1.590	0.112	14.17
Duration 3+	2.694	0.231	11.64
Woman's Education	-0.202	0.021	-9.67

5 Discussion

5.1 Substantive Points

Our methods have allowed us to reach several substantive conclusions. The most important of these is the fact that the decline in marital fertility was a period effect and not a cohort effect. It started in the early 1960s, just after the onset of rapid economic and educational expansion, and coincided with the 1963 "White Revolution", which gave women the vote and introduced a Literacy Corps and a Health Corps, among other modernizing measures. It accelerated in the early 1970s, after the establishment of the official family planning program and the passing of the Family Protection Act which raised the legal status of women, both in 1967.

Other possible causes of the decline include the reduction in infant mortality and the contextual effect of the population being more educated (but not the individual effect of education on fertility). Individual characteristics, namely the woman's education and that of her husband and where she lives, do not explain the overall decline in fertility but they are important predictors of the fertility of individuals.

Previous work based on WFS data has tended to emphasize the effects of individual characteristics that are fixed early in a woman's life, notably her education and her place of childhood residence (e.g. Entwisle and Mason, 1985). This would lead one to expect fertility decline to be a cohort effect produced by the increasingly educated and urban composition of the population. Our finding that fertility decline was a period effect indicates that analyses based on such early-life variables are not very useful for explaining fertility decline in Iran. We also find that it is current rather than childhood residence that is important; this is congruent with the period effect result.

5.2 Methodological Points and Other Approaches

Our work here leads us to make several methodological points. Event history analysis seems preferable to regression with children ever born as the dependent variable, because it allows us to use all the data and avoids biases due to censoring. It is possible to translate the parameters of event history models into more familiar demographic measures such as age-specific fertility rates and the TFR, at least approximately. The event history approach has also enabled us to model the effects of age, period, cohort, parity and duration simultaneously, so that we have five clocks running at once. Parametric modeling of these effects is useful and works well here.

We have found that the effects of the independent variables are quite different for parities 2 and above than for parities 0 and 1, and it is important to model them separately. This indicates another potential problem with regression on children ever born: it effectively considers all parities together and so the effects of covariates may be diluted or even suppressed. This is particularly relevant here since for parities 2 and above there was a decreasing period effect in fertility and for parity 0 there was an increasing period effect. The trend in parities 2 and above would tend to produce a decrease in CEB, while the trend in parity 0 would tend to produce an increase. Thus the period effect would not be as clear in CEB regression as in the EHA analysis, and its parity-related nature would not be revealed at all.

We have found that ACE is a useful exploratory tool for determining the best way to code independent variables.

An alternative approach that has been used in the analysis of fertility is the calculation of parity progression ratios (PPRs) (Henry, 1953; Ryder, 1971, 1982; Feeney, 1983). For a given group or subgroup of women, the PPR at parity i is essentially the proportion of women with an ith birth who go on to have an $(i+1)$th birth. Since we take account explicitly of parity in our event history model, it does not seem that an analysis based on PPRs would provide any additional information.

A difficulty with the PPR approach is that it involves analyzing separately each subgroup created when the sample is cross-classified by the independent variables. In our case, the number of such subgroups is very large, and so, as pointed out by Pullum (1987), very large samples are typically required for the method to give statistically stable results. It seems that it would be hard for the PPR approach to answer the questions that we have considered here with a sample as small as that in the IFS. It has been applied mostly to large groups such as entire cohorts (e.g. Ryder, 1971), rather than to determining the effects of independent variables, which is a key focus in our work here.

The PPR approach has been applied mostly to developed countries and has been little used in the analysis of WFS data, This is in part because of its sample size requirements, but perhaps it could also be because the "stopping" behavior that it aims to capture has played a bigger role in recent fertility patterns in developed countries than in the developing nations surveyed by the WFS. Feeney (1983) provides an example of the application of PPR methods in a developing nation context, but this is not to WFS data and is relatively unusual in the PPR-based literature.

There have been some other EHA analyses of WFS fertility data, but these have been relatively few and of limited scope, restricted, for instance, to a single birth interval. For example, Bumpass, Rindfuss and Palmore (1986) carried out what amounted to discrete-time event history analysis of the Korean WFS, but they analyzed parity 2, parity 3, and parities 4–8 separately, and avoided the difficulties that we met with a more comprehensive approach. One of the more comprehensive EHA analyses of WFS data is that of Gilks (1984), as reported by Pullum (1987). He used continuous-time EHA rather than the discrete-time version used here. His approach differs in that he did not include cohort as an independent variable, so that he made no effort to distinguish between period and cohort effects. Also, while age, parity and period were included in his model, these are all entered linearly and he did not try to find a more appropriate functional form.

EHA has also been successfully applied to WFS child *mortality* data (e.g Trussell and Hammerslough, 1983). This is somewhat easier than the EHA of fertility data because there is only a single event (death) rather than the sequence of events (births) in fertility data, because the "five clocks" problem typically does not arise, and because the computer demands tend to be much smaller.

5.3 Unobserved Heterogeneity

We have not dealt with the very important issue of unobserved heterogeneity. We are currently working on ways of taking account of this by adding an unobserved woman-specific random effect to the right-hand side of equation (1). The resulting model can be estimated using recently developed Gibbs sampling and other Markov chain Monte Carlo methods (Lewis, 1993), and the methodological issues discussed in this paper remain relevant.

The random effect can be specified to have a given parametric distribution, such as normal, t, or log-gamma; this distribution can be allowed to vary with parity or period.

Results with related models have been found to be sensitive to the distribution assumed for the random effect (e.g. Trussell and Richards, 1985). This seems to us to constitute model uncertainty, which should be accounted for explicitly in any inferences (Raftery, 1993a); there is no need to select a single model, and, indeed, it may be misleading to do so.

An alternative is to allow the random effect to have a finite discrete distribution; this can also be estimated in the Gibbs sampling framework, and leads to a discrete-time analogue of the model that underlies the nonparametric maximum likelihood estimator of Heckman and Singer (1984). This has the advantage of allowing explicit modeling of sterility. Specifying a mixture of normal or other distributions for the random effect yields a rich class of models that includes both the Heckman-Singer nonparametric model and the parametric models mentioned above, to any desired degree of approximation. Note that previous modeling of unobserved heterogeneity in fertility has concentrated on continuous-time data (e.g. Heckman and Walker, 1987, and references therein), and so is not directly applicable to our discrete-time context, although some of the basic ideas are generally useful.

5.4 Concluding Remarks

Pullum (1987) commented that "WFS has contributed relative little to the causal modeling of socioeconomic or background variables ... For the most part, the effects of explanatory variables have been analyzed in a rather shallow manner." We have tried to argue here that EHA promises to provide a good framework for such modeling, perhaps better than the other methods used or developed by the WFS. This should not be too surprising since, as Pullum points out, the methods developed by the WFS were not mainly meant for this purpose. For EHA to fulfil this promise, the problems identified in Section 1.4 have to be solved, and we hope that our proposals here will help to do this.

References

Aghajanian, A. (1978). "Fertility and family economy in rural communities of Iran." *Journal of Comparative Family Studies* **9**, 119–128.

Aghajanian, A. (1988). "The value of children in rural and urban Iran: A pilot study." *Journal of Comparative Family Studies* **19**, 85–97.

Aghajanian, A. (1991). "Population change in Iran, 1966-86: A stalled demographic transition?" *Population and Development Review* **17**, 703–716.

Aghajanian, A. (1993). "The Iran Fertility Survey of 1976–1977: First County Report." In preparation.

Aghajanian, A., Gross, A.B. and Lewis, S.M. (1993). "Evaluation of the Iran Fertility Survey". Working Paper, Center for Studies in Demography and Ecology, University of Washington.

Ajami, I. (1976). "Differential fertility in peasant societies." *Population Studies* **30**, 435–464.

Allison, P. (1984). *Event History Analysis.* Beverly Hills: Sage.

Bongaarts, J. (1978). "A framework for analyzing the proximate determinants of fertility." *Population and Development Review* **4**, 105–132.

Bongaarts, J. (1982). "The fertility-inhibiting effects of the intermediate fertility variables." *Studies in Family Planning* **13**, 179–186.

Breiman, L. and Friedman, J.H. (1985). "Estimating optimal transformations for multiple regression and correlation (with Discussion)." *Journal of the American Statistical Association* **80**, 580–619.

Bulatao, R.A. and Lee, R.D., eds. (1983). *Determinants of Fertility in Developing Countries.* New York: Academic Press.

Bumpass, L.L, Rindfuss, R.R. and Palmore, J.A. (1986). "Determinants of Korean birth intervals: The confrontation of theory and data." *Population Studies* **40**, 403–423.

Carroll, G. (1983). "Dynamic models of discrete dependent variables: A didactic essay." *Quality and Quantity* **17**, 425–460.

Cleland, J. and Scott, C. (1987). *The World Fertility Survey: An Assessment.* Oxford University Press.

Cleland, J. and Verma, V. (1989). "The World Fertility Survey: An appraisal of methodology (with Discussion)." *Journal of the American Statistical Association* **84**, 756–774.

Coale, A.J. and Trussell, T.J. (1974). "Model fertility schedules: Variations in the age structure of childbearing in human populations." *Population Index* **40**, 185–258.

Cox, D.R. (1972). "Regression models and life tables (with Discussion)." *Journal of the Royal Statistical Society, series B* **34**, 187–201.

Cox, D.R. and Oakes, D. (1984). *Analysis of Survival Data.* London: Chapman and Hall.

DeVeaux, R.D. (1989). "Finding transformations for regression using the ACE algorithm." *Sociological Methods and Research*, **18**, 327–359.

Duchêne, J. and Gillet-de Stefano, S. (1974). "Ajustement analytique des courbes de fécondité générale." *Population et Famille* **32**, 53–93.

Easterlin, R.A. and Crimmins, E.M. (1985). *The Fertility Revolution.* Chicago: University of Chicago Press.

Ebrahamian, E. (1982). *Iran Between Two Revoluations.* Princeton, N.J.: Princeton University Press.

Entwisle, B. and Mason, W. (1985). "Multilevel effects of socioeconomic development and family planning programs on children ever born." *American Journal of Sociology* **91**, 616–649.

Farid, S.M. (1973). "On the pattern of cohort fertility." *Population Studies* **27**, 159–168.

Feeney, G. (1983). "Population dynamics based on birth intervals and parity progression." *Population Studies* **37**, 75–89.

Fienberg, S.E. and Mason, W.M. (1978). "Identification and estimation of age-period-cohort models in the analysis of discrete archival data." In K.F. Schuessler (ed.) *Sociological Methodology 1979.* San Francisco: Jossey-Bass, pp. 1–65.

Gilks, W.R. (1984). "The relationship between birth history and current fertility." WFS unpublished manuscript.

Glenn, N.D. (1989a). "A caution about mechanical solutions to the identification problem in cohort analysis: Comment on Sasaki and Suzuki." *American Journal of Sociology* **95**, 754–761.

Glenn, N.D. (1989b). "A flawed approach to solving the identification problem in the estimation of mobility effects: A comment on Brody and McRae." *Social Forces* **68**, 789–795.

Goode, M. and Farr, D. (1980). "Social status and fertility: A study of Twon and the villages in North Western Iran." *Population Studies* **34**, 311-319.

Grusky, D.B. and Hauser, R.M. (1984). "Comparative social mobility revisited: Models of convergence and divergence in 16 countries." *American Sociological Review*, **49** 19–38.

Heckman, J. and Singer, B. (1984). "A method for minimizing the impact of distributional assumptions in econometric models for duration data." *Econometrica* **52**, 271–320.

Heckman, J. and Walker, J. (1987). "Using goodness of fit and other criteria to choose among competing duration models: A case study of Hutterite data." In *Sociological Methodology 1987*, C.C. Clogg (ed.). Washington, D.C.: American Sociological Association, pp. 247–307.

Henry, L. (1953). *Fécondité des Mariages. Nouvelle méthode de mesure.* Travaux et Documents, Cahier no. 16, Paris: Presses Universitaires de France.

Hirschman, C. (1985). "Premarital socioeconomic roles and the timing of family formation: A comparison of five Asian societies." *Demography* **22**, 35–59.

Hobcraft, J.N. and Casterline, J.B. (1983). *Speed of Reproduction.* WFS Comparative Study no. 25. Voorburg, Netherlands: International Statistical Institute.

Hobcraft, J.N., Menken, J. and Preston, S.H. (1982). "Age, period, and cohort effects in demography." *Population Index* **48**, 4–43.

Hobcraft, J.N. and Murphy, M. (1986). "Demographic event history analysis: A selective review." *Population Index* **52**, 3–27.

Isaac, L., Cutright, P., Jackson, E. and Kelly, W.R. (1979). "Period effects on race and parity-specific birth probabilities of American women, 1917–1976: A new measure of fertility." Unpublished manuscript, Florida State University.

Johnson, R.A. (1985). "Analysis of age, period and cohort effects in marital fertility." In Mason, W.M. and Fienberg, S.E. (eds.) *Cohort Analysis in Social Research.* New York: Springer-Verlag, pp. 229-258.

Kass, R.E. and Raftery, A.E. (1993). "Bayes factors and model uncertainty." Technical Report no. 254, Department of Statistics, University of Washington.

Lewis, S.M. (1993). "Contribution to the Discussion of three papers on Gibbs sampling and related Markov chain Monte Carlo methods." *Journal of the Royal Statistical Society, series B*, **55**, 79–81.

Little, R.J.A. (1978). "Generalized linear models for cross classified data from the WFS." WFS Technical Bulletin No. 5. Voorburg, Netherlands: International Statistical Institute.

Little, R.J.A. (1980). "Linear models for WFS data." Voorburg, Netherlands: International Statistical Institute.

Little, R.J.A. (1981). "Statistical models for World Fertility Survey data." In *World Fertility Survey Conference 1980: Record of Proceedings.* Voorburg, Netherlands: International Statistical Institute, vol. 3, pp. 189–260.

Little, R.J.A. and Perera, S. (1981). "Illustrative analysis: Socio-economic differentials in cumulative fertility in Sri Lanka — A marriage cohort approach." *WFS Scientific Reports*, no. 12. Voorburg, Netherlands: International Statistical Institute.

Mason, K.O., Mason, W.M., Winsborough, H.H. and Poole, W.K. (1973).. "Some methodological issues in cohort analysis of archival data." *American Sociological Review* **38**, 242–258.

Nakamura, T. (1986). "Bayesian cohort analysis for general cohort table analysis." *Annals of the Institute of Statistical Mathematics* **38**, 353–370.

Page, H.J. (1977). "Patterns underlying fertility schedules: A decomposition by both age and marriage duration." *Population Studies* **31**, 85–106.

Palloni, A. (1985). "Assessing the effects of intermediate variables on birth interval-specific measures of fertility." *Population Index* **50**, 623–657.

Paydarfar, A. (1975). "Socio-cultural correlates of fertility among tribal, rural and urban populations of Iran." *Social Biology* **22**, 151–166.

Petersen, T. (1991). "The statistical analysis of event histories." *Sociological Methods and Research* **19**, 270–323.

Pullum, T.W. (1980). "Separating age, period and cohort effects in white U.S. fertility, 1920–1970." *Social Science Research* **9**, 225–244.

Pullum, T.W. (1987). "Analytic methodology." In J. Cleland and C. Scott (eds.) *The World Fertility Survey: An Assessment.* Oxford University Press, pp. 644–676.

Raftery, A.E. (1986a). "A note on Bayes factors for log-linear contingency table models with vague prior information." *Journal of the Royal Statistical Society, series B* **48**, 249–250.

Raftery, A.E. (1986b). "Choosing models for cross-classifications." *American Sociological Review* **51**, 145–146.

Raftery, A.E. (1988). "Approximate Bayes factors for generalized linear models." Technical Report no. 121, Department of Statistics, University of Washington.

Raftery, A.E. (1993a). "Bayesian model selection for structural equation models." In K.A. Bollen and J.S. Long (eds.) *Model Selection in Structural Equation Models.* Beverly Hills: Sage.

Raftery, A.E. (1993b). "Approximate Bayes factors and accounting for model uncertainty in generalized linear models". Technical Report no. 255, Department of Statistics, University of Washington.

Raftery, A.E. and Hout, M. (1993). "Maximally maintained inequality: Expansion, reform and opportunity in Irish education, 1921–1975." *Sociology of Education*, **66**, 41–62.

Rindfuss, R. and Morgan, P. (1983). "Marriage, sex and the first birth interval: The quiet revolution in Asia." *Population and Development Review* **9**, 259–278.

Rodriguez, G. and Cleland, J. (1981). "Socio-economic determinants of marital fertility in 20 countries: A multivariate analysis." In *World Fertility Survey Conference, 1980: A Record of Proceedings.* Voorburg: International Statistical Institute, vol. 2, pp. 325–414.

Rubin, D.B. (1987). *Multiple Imputation for Non-Response in Surveys.* New York: Wiley.

Ryder, N.B. (1971). "Notes on fertility measurement." *Milbank Memorial Fund Quarterly* **49**, 109–127.

Ryder, N.B. (1982). "Progressive fertility analysis." *WFS Technical Bulletins*, no. 8. Voorburg, Netherlands: International Statistical Institute.

Sasaki, M. and Suzuki, T. (1987). "Changes in religious commitment in the United States, Holland and Japan." *American Journal of Sociology* **92**, 1055–1076.

Shavit, Y. and Bloesfeld, P. (1992) (eds.). *Trends in Educational Opportunity.* Westview Press.

Singer, B. (1989). "Comment." *Journal of the American Statistical Association* **84**, 771–772.

Trussell, J. (1984). "An evaluation of imputation of dates in WFS surveys." Unpublished manuscript, Office of Population Research, Princeton University.

Trussell, J. and Hammerslough, C. (1983). "A hazards-model analysis of the covariates of infant and child mortality in Sri Lanka." *Demography* **20**, 1–26.

Trussell, J. and Richards, T. (1985). "Correcting for unobserved heterogeneity in hazard models: An application of the Heckman-Singer model to demographic data." *Sociological Methodology 1985* (N.B. Tuma, ed.), **15**, 242–276. San Francisco: Jossey-Bass.

Tuma, N.B. and Hannan, M.T. (1984). *Social Dynamics.* Orlando, Fla.: Academic Press.

Verma, V.K. (1980). *Basic Fertility Measures from Retrospective Birth Histories.* WFS Technical Bulletin No. 4. Voorburg, Netherlands: International Statistical Institute.

Wilmoth, J. (1990). "Variation in vital rates by age, period and cohort." *Sociological Methodology 1990* (C.C. Clogg, ed.), **20**, 295–336.

Yamaguchi, K. (1991). *Event History Analysis.* Newbury Park, Calif.: Sage.