On Cohort Forecasts of First Marriage for U.S. Women

P. C. Roger Cheng*

Abstract

This paper develops a simple age-period-cohort framework in completing cohort incomplete first marriage schedules and makes full use of the 1935–1993 U.S. data to obtain robust outcomes. From a theoretical perspective, we show that the period effect represents not only transitory macro shocks that impinge on all cohorts, but also the interaction between the age and cohort effects. Empirically, we indicate that the period effect is the key to transforming a marriage level into a marriage schedule. Accompanied by the smoothed version of the adjusted period measure proposed in Kohler and Philipov (2001), we approximate the cohort marriage schedules fairly well and the estimates of all distributional parameters can be thereby obtained. Our approach is easy to implement and the data requirement is relatively light, indicating that the proposed method is readily applicable to countries whose data lengths are not long enough and would be helpful for further empirical investigation in the relationship between cohort nuptiality behavior and other cohort-specific socioeconomic factors.

1 Introduction

The most comprehensive way to describe the first marriage behavior for women of a birth cohort is to construct their complete marriage (nuptiality) schedule, so that all distribution-related parameters, such as cohort total nuptiality rate (CTNR), mean or median age at first marriage, standard deviation, skewness, kurtosis, and so on, can be directly derived. Such schedules can be very useful as building blocks for the study of interdependence among women's education, labor force participation, marriage, and childbearing.

*Correspondence: P. C. Roger Cheng, Department of Economics, National Central University, Taoyuan 320, Taiwan. Tel: 886-3-422-7151 ext. 66310; Fax: 886-3-422-6903; E-mail: paochih@mgt.ncu.edu.tw.

Measuring the actual marriage behavior for birth cohorts of women, however, demands lengthy series of annual data. With a data set consisting of age-specific first marriage rates (ASFMRs) over a period of consecutive years, one can track a cohort of women over time and observe the changes in their nuptiality rates. If the series of annual data is lengthy enough, one will obtain *complete* marriage schedules for several cohorts. Nevertheless, there are always a few cohorts whose marriage schedules are *incomplete*, no matter how long the series is; data concerning their younger or older experiences are just unavailable. Statistical and behavioral approaches have been developed to help project the past or future experience of women, especially of those who are still young today. In this paper, two popular curve-fitting models proposed in the literature, the Hernes model and the Coale-McNeil model, will be evaluated according to their performance in predicting incomplete marriage schedules.

By contrast, the period total nuptiality rate (PTNR) in a particular year can be easily computed by summing up the ASFMRs in that year. In consideration of the fact that the PTNR is usually affected by changes in the age pattern of first marriage, alternative period indicators such as Bongaarts and Feeney (1998), Kohler and Philipov (2001), and Kohler and Ortega (2002), which are originally designed for fertility rates, can be used to disentangle the tempo (and/or the spread) and the quantum components. While quite a few researchers are skeptical about the usefulness of these adjusted period measures due to their assumptions regarding the shape of the age patterns and/or due to their fluctuating and occasionally absurd values (e.g., Kim and Schoen, 2000; Li and Wu, 2003; Schoen, 2004; van Imhoff and Keilman, 2000; van Imhoff, 2001), some researchers have compared them with the cohort measures (e.g., Bongaarts and Feeney, 2006; Kohler and Ortega, 2002; Ryder, 1990; Schoen, 2004; Smallwood, 2002; Sobotka, 2003; van Imhoff and Keilman, 2000), explicitly or implicitly regarding these period indicators as forecasts of the completed (fertility or nuptiality) levels for cohorts who are still young. To ease the problem of random fluctuations and the occurrence of nonsensical values that confront the primitive values of these adjusted period indicators, we propose a smoothed version of these indicators and then evaluate their performance in predicting cohort nuptiality levels.

Furthermore, this paper develops a framework in completing cohort incomplete marriage schedules which can operate well even when the data length is not too long, broadening the range of its application. As will be shown, this framework outperforms previous methods in forecasting incomplete cohort schedules. Specifically, we propose a simple age-period-cohort (APC) model to decompose ASFMRs into age, period, and cohort effects. From a theoretical perspective, we show that the period effect so derived represents not only transitory macro shocks that impinge on all cohorts, but also the interaction between the age and cohort effects. Empirically, we indicate that such a period effect is the key to transforming a nuptiality level into a marriage schedule. In other words, once the actual nuptiality levels of cohorts who have not finished first marriage were known, their incomplete marriage schedules could have been estimated very well. Notwithstanding the value of actual CTNR is in reality not available, our research provides a possible substitute for it; smoothing the values of adjusted period indicators can removed fluctuations in them and then summarize the trend which becomes a good approximation to the CTNR.

Whatever forecasting approach is adopted, accuracy should be a major criterion for any forecast worthy of the name. To test the predictive power of a model, calculating ex ante projections and then comparing with actual data is more rigorous than just evaluating the ex post fit in a particular sampling period. Moreover, since the results of testing the predictive power of a model may depend on the choice of prediction period, it is suggested that ex ante projections should be evaluated in a few periods rather than in just a particular one (see De Beer, 1985, p.529). This paper makes as full use as possible of the long span in the U.S. data: we set a fixed subperiod length, move the starting point of subperiod along calendar years, compute measures of competing indicators in each subperiod, and then assess their average performance in all subperiods.

The remainder of this paper is organized as follows. Section 2 provides an overview of some U.S. nuptiality data sets, and then evaluates the Hernes model and the Coale-McNeil model. Given a constructed continuum of cohort marriage histories, Section 3 demonstrates the empirical performances of the smoothed version of some adjusted period measures in predicting cohort nuptiality levels. The approach proposed in this paper is introduced in Section 4, and we will discuss its theoretical properties and show that the period effect is the key to successfully completing incomplete marriage schedules. Section 5 presents how the collaboration of the smoothed period indicator with our APC indicator can yield better forecasts, and Section 6 summarizes.

2 Data and Curve Fitting Models

2.1 Data

The marriage history of a particular cohort can be derived from sample survey data, either directly or indirectly. One can construct the history using the answer to the question about the age at first marriage (AFM), or compute the weighted proportion of female samples whose marital status is other than 'Never Married' at each age and then track the proportions for this cohort of women over time. Only a few data sets contain the question about the AFM, including the 1960 census data, the March Current Population Survey (CPS) from 1962 through 1971, and the June 1995 CPS,[1] while the question about the marital status is included in most data sets, such as the March CPS from 1962 through 2008.

As for data sets of the first type, ASFMRs of cohorts are available and the related cumulated marriage levels can be thereby derived. The main problem regarding data of this type comes from subjects' memorial correctness. By contrast, cumulated marriage levels of cohorts are obtained from data sets of the second type and the related ASFMRs can be calculated by substracting cumulated marriage levels at two adjacent ages. Although the memorial correctness problem can be slight in data of this type, nonsensically negative

[1] All these data sets are available from the DataFerrett, U.S. Census Bureau. The sample size is about 1.8 million in the 1960 census and about 140,000 in other CPS data sets.

ASFMRs might arise due to sampling errors.

To evaluate the performance of a model in predicting incomplete marriage schedules, we need some complete cohort marriage histories over ages 15–50. The 1960 census, the March 1971 CPS, and the June 1995 CPS allow us to observe the whole marriage schedules of those born from 1885 through 1909, from 1896 through 1920, and from 1930 through 1944, respectively.[2] The 1962–2008 March CPS gathered cumulated marriage levels over ages 15–50 for cohorts 1947–1958. As one can see, some cohorts between 1885 and 1958 are covered in more than two data sets while some are not covered at all.

2.2 Curve Fitting Models

In this subsection, two popular curve fitting models, developed by Hernes (1972) and by Coale and McNeil (1972), are evaluated by their performance in forecasting the eventual cohort marriage levels.[3] Since complete marriage histories for 77 cohorts are available from the four data sets, one can truncate the latter part of a marriage schedule at any chosen age and see how close the predicted values are to the actual CTNR. The absolute percentage error (APE)

$$\mathrm{APE}_c(a) = \frac{|I_c(a) - \mathrm{CTNR}_c|}{\mathrm{CTNR}_c} \times 100\%$$

is adopted as the criterion of evaluation when the observed schedule of cohort c is limited to ages 15 through a, where $I_c(a)$ is the predicted value of the eventual marriage level, CTNR_c. We select every five ages from 24 as the cut point a and present the results in Figure 1 (note that the maximum scale on the vertical axis in the panel where $a = 24$ is different from those in other panels).

For all cohorts with complete marriage histories, the figure shows that the APE decreases as the observed history is truncated at a higher age and that the Hernes model

[2] The oldest age group we choose from the data is set at 75, while the June 1995 CPS provides related information only up to 65.

[3] For a clear description of these two models, also refer to Goldstein and Kenney (2001).

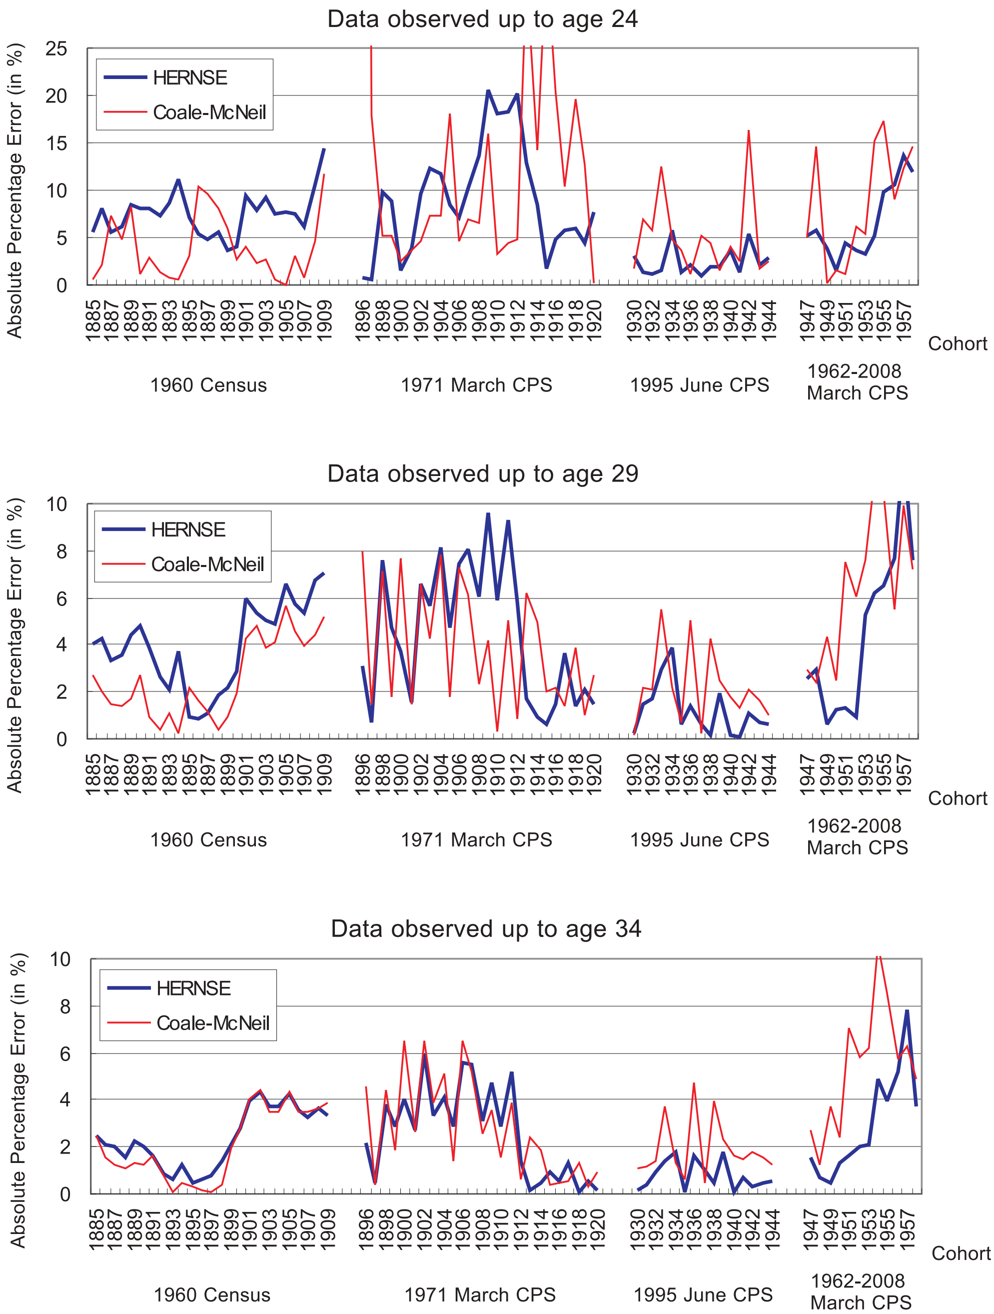

Figure 1. Performance of Curve Fitting Models in Predicting the CTNR with Limited Observations

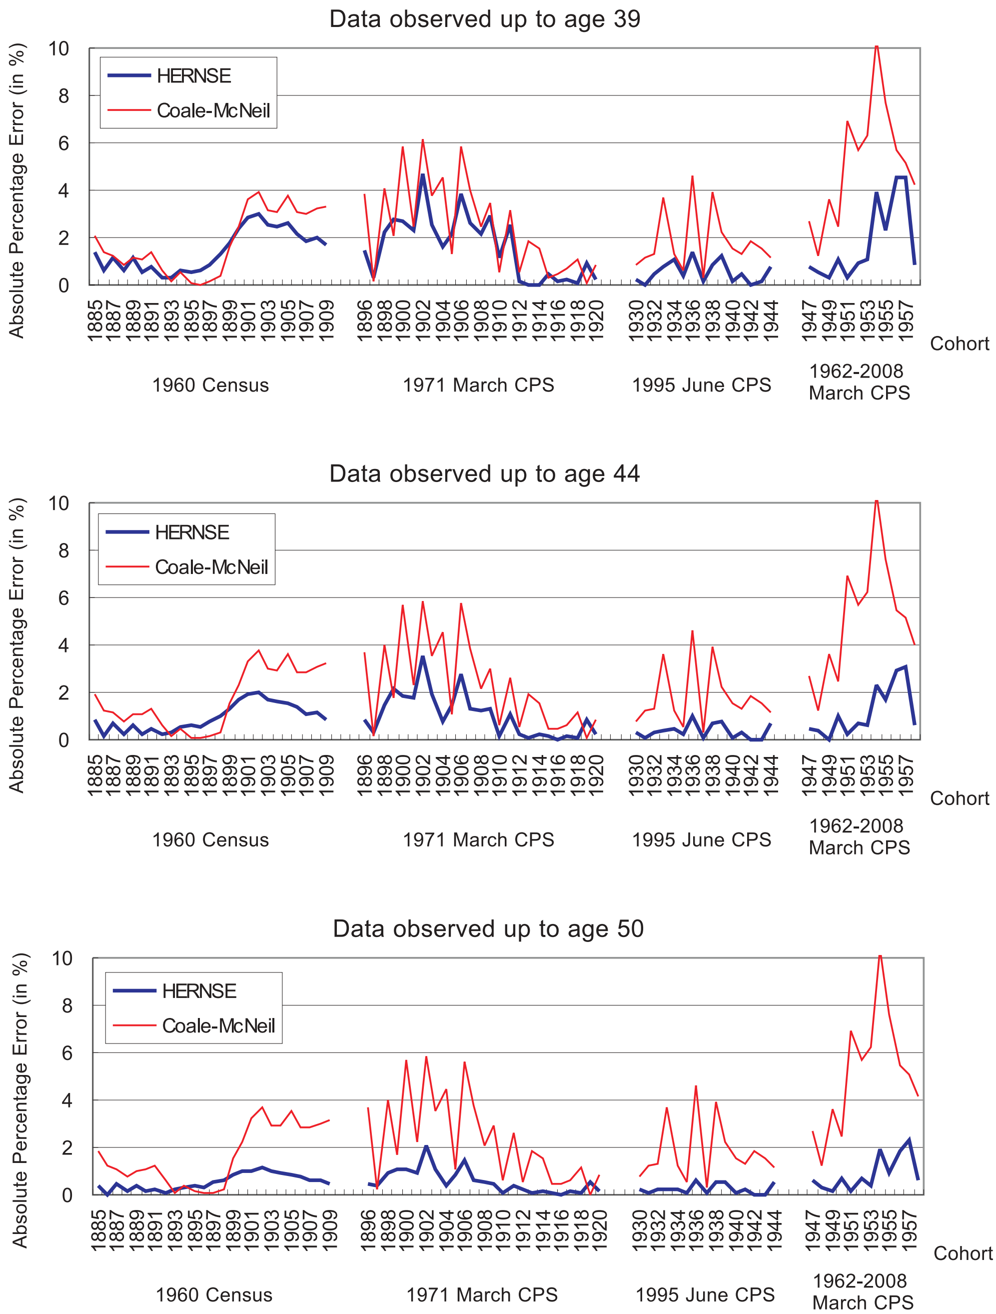

Figure 1. (Continued) Performance of Curve Fitting Models in Predicting the CTNR with Limited Observations

performs better than the Coale-McNeil model in general. Regarding the performance of the Hernes model, the maximum (mean) APE over the 77 cohorts is 20.64% (6.89%) when $a = 24$, 4.70% (1.37%) when $a = 39$, and 2.31% (0.55%) when $a = 50$. If one takes 5% as a standard, then the Hernes model with observations covering from 15 to at least 39 provides reliable projections. This conclusion allows us to construct a continuum of complete cohort marriage histories for cohorts 1885 through 1969. For consistency of this constructed data set, the Hernes fitted values for all cohorts are used so that possible measurement and sampling errors, as well as macroeconomic shocks that impinge on all cohorts, will be removed. As for the cohorts covered in more than two data sets, the 1960 census is prior to the March 1971 CPS and the June 1995 CPS is prior to the 1962–2008 March CPS.

3 Period Nuptiality Indicators

Given this constructed data set of complete marriage histories for cohorts 1885–1969, one can derive a corresponding ASFMR Lexis rectangle describing time and age, covering years 1935–1984 and ages 15–50.[4] Furthermore, in order to proceed the following analysis, we extend the rectangle to year 1993 in which women of cohort 1969 were aged 24 by adding the Hernes fitted values at early ages for cohorts 1970–1978.

Figure 2 shows period total nuptiality rates, PTNRs, and the period mean age at first marriage, PMAFM, for the U.S. from 1935 through 1993. Also presented in the same figure are cohort total nuptiality rates, CTNRs, for women of cohorts aged 23 in corresponding years . It is clear that the CTNR is relatively stable than the PTNR and that the period shown in the figure can be divided into two parts. Before year 1960, the PTNR is higher than the CTNR and the mean age at first marriage declines.[5] By contrast, the PTNR falls

[4] The Lexis rectangle of ASFMRs can also be obtained from official register data, by dividing the number of first marriages in a given age group by the number of all women in that age group.

[5] The above-one rates in the late 1940s and the early 1950s are anomalous because a woman can experience at most one first marriage. The apparent explanation is the declining mean age at first marriage.

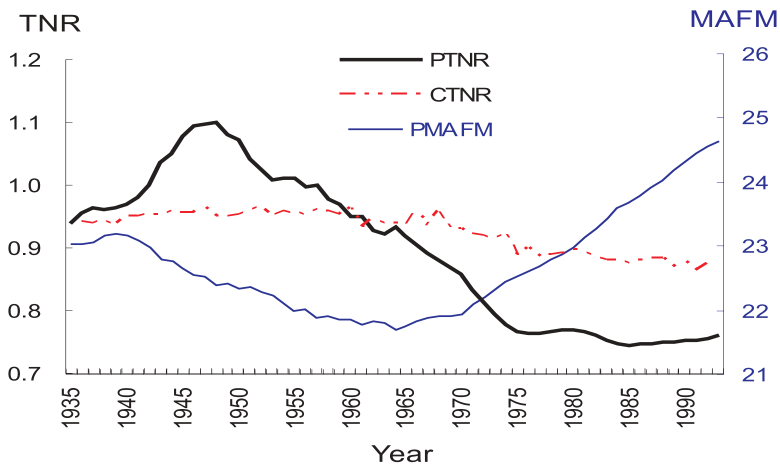

Figure 2. Total Nuptiality Rate and Mean Age at First Marriage,
Female in the U.S.

below the CTNR after 1960, with values around 0.75 toward the end of the period, and the mean age at first marriage rises. This suggests that the period nuptiality measure is significantly distorted by the tempo effect.

To remove distortions caused by the tempo effect (and/or the spread effect), one can compute period adjusted measures and examine their proximity to the cohort measure.[6] Based on our 1935–1993 ASFMR data, we investigate the performance in approximating cohort marriage levels of two representative period indicators, including

(1) tempo-adjusted total nuptiality rates proposed by Bongaarts and Feeney (1998), but the total nuptiality rate and the mean age of the marriage schedule are smoothed before being substituted into the formula, denoted as BFS hereafter.

(2) tempo-adjusted and variance-adjusted total nuptiality rates proposed by Kohler and Philipov (2001), denoted as KP hereafter.

Note that given ASFMR data for T consecutive years, there will be T figures for the PTNR

[6]What we do here is just to compare the estimated values with actual cohort ones, as van Imhoff (2001) indicates: "Any procedure trying to estimate cohort quantum from period quantum is based on simplifying assumptions, the justifiability of which can only be verified empirically" (p. 36).

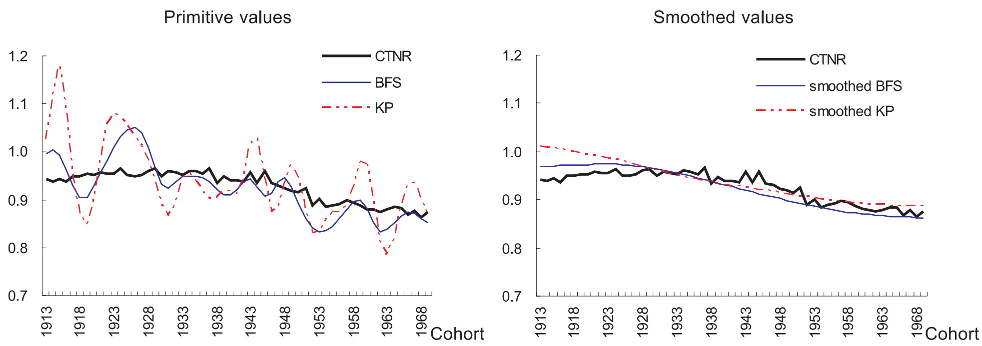

Figure 3. Comparison of Adjusted Period Measures with the CTNR:
Primitive and Smoothed Values

but only $T - 2$ figures for the BFS and the KP; values for the first and the last years will not be available for the calculation reason.

Period adjusted measures in a particular year are compared with the CTNR of women who reached the mean age at first marriage in that year. For example, the BFS and the KP in year 1936 will be compared with the CTNR of cohort 1913 since the 1936 MAFM was 23. The left panel in Figure 3 depicts the comparison of the BFS and the KP (in their primitive values) with the CTNR. As can be seen, period indicators exhibit considerable fluctuations and suffer the problem of the occurrence of impossible above-one values, while the cohort related parameter usually follow a relatively stable pattern. Smoothing the primitive values of period indicators is then a reasonable choice to examine whether there will be any improvement in the approximation,[7] and the right panel in Figure 3 presents the smoothed results. The contrast between these two panels is too salient to deny the usefulness of smoothing. The mean absolute percentage errors (MAPEs) for the primitive BFS and the primitive KP are 3.27% and 5.85%, while they drop to 1.72% and 2.08% after smoothing. If we focus on the predictability of these period adjusted measures and restrict the comparison to the last 15 cohorts (1955–1969), the MAPEs become 2.43%

[7] As is well known, smoothers are by definition devised to remove fluctuations and summarize the trend of data, and there have been a few studies (e.g., Silverman, 1996; Kohler and Philipov, 2001; Currie et al., 2004) trying to smooth the original data before implementing their models.

and 4.45% for the primitive BFS and KP and drop to 1.41% and 1.04% after smoothing.[8]

Instead of using up all observations in 59 years to compare period and cohort indicators only once, we further investigate the average performance in all subperiods of a chosen fixed length, thus making full use of the long span in the data and yielding more robust results. Mark any period of 25 consecutive years during 1935–1993 as a round,[9] and there will thus be a total of 35 rounds, from the first one (1935–1959) to the last (1969–1993). In the mth round ($m = 1, \ldots, 35$), ASFMRs for 25 years from $1934 + m$ through $1958 + m$ are available, and the corresponding 23 values (from $1935 + m$ through $1957 + m$) of the BFS and the KP can be derived. Period nuptiality indicators in a particular year t are compared with the CTNR of the cohort whose year of birth is $c = t - 23$, where 23 is the average of the mean age at first marriage over the entire data span from 1935 through 1993, rounded up or down to the nearest integer.[10] The proximity of period adjusted indicators to the corresponding CTNRs in a particular round can be measured by the MAPE. The average MAPE across all 35 rounds for the smoothed KP is 1.21% for the most recent 15 cohorts in each round, and the maximum of the maximal APE in each round for the smoothed KP is 4.43%.

4 APC Model in Completing Cohort Incomplete Marriage Schedules

Next, this paper develops a framework in completing cohort incomplete marriage schedules which can operate well even when the data length is not too long. Let $f(a, c)$ represent the age specific nuptiality rates at age a for women of birth cohort c. The life-cycle

[8]The reader might want to know the maximum APEs for these two measures. Among the last 15 cohorts, they are 5.31% and 10.65% for the primitive BFS and KP and drop to 2.45% and 2.61% after smoothing.

[9]In this paper we choose 25 as the subperiod length because a 25-year ASFMR data set can be obtained for many countries. Alternatively, subperiod lengths other than 25 could be chosen. All related results for lengths of 20, 30 or 35 are available upon request from the authors.

[10]The period mean age at first marriage in fact ranges from 21.69 to 24.65 during 1935–1993, so that a period measure in year t should be compared with cohort $t - \mathrm{PMAFM}_t$ rather than $t - 23$. We use interpolation and extrapolation to solve this problem.

pattern in timing their first marriage can be viewed as a function of the eventual marriage level throughout their life (up to age 50 in this paper). With a density $h(a, c)$ satisfying $\int h(a, c)da = 1$, the relationship

$$f(a, c) = h(a, c)\ \text{CTNR}(c) \tag{1}$$

must hold by definition. Furthermore, we specify $h(a, c)$ as the product of two components:

$$h(a, c) = h_1(a)\ h_2(a, c), \tag{2}$$

where h_1 contains only the variable a, and h_2 contains nothing else but all interaction terms between a and c.[11] By substituting Equation (2) into Equation (1) and taking logs, we have

$$\ln f(a, c) = \ln h_1(a) + \ln h_2(a, c) + \ln \text{CTNR}(c). \tag{3}$$

The problem that arises here is whether there is a simple and intuitive way to specify the interaction term h_2. Since $f(a, c)$ also denotes the age specific nuptiality rates for women aged a at period $p = a + c$, a function of p (except for the linear term) can in effect capture most interaction between a and c, with a minor limitation that symmetric terms (such as a^2c and ac^2) share a common coefficient. Unless the limitation is severely violated, Equation (3) can be re-expressed by

$$\begin{aligned} \ln f(a, c) &= H_1(a) + H_2(p) + H_3(c) \\ \text{or} \quad f(a, c) &= \exp[H_1(a) + H_2(p) + H_3(c)], \end{aligned} \tag{4}$$

where H_1, H_2, and H_3 denote the age, period, and cohort effects, respectively. Note that there is no one-to-one correspondence for each single RHS item between Equations (3)

[11] It is of no use specifying another component containing only the variable c since $\int h(a, c)da = 1$ must hold for all c.

and (4). Equation (4) can be estimated by regressing the age specific nuptiality rates (plus 1 per thousand to account for the logarithm of zero nuptiality rates otherwise being undefined) against age dummies (A), period dummies (P), and cohort dummies (C) in a nonlinear least-squares way:

$$f(a,c) = \exp(\iota\beta + A\alpha + P\psi + C\gamma), \tag{5}$$

where ι is a vector of units, β denotes the intercept, and vectors α, ψ, and γ are the parameters of the age, period, and cohort effects. To solve the identification problem arising from the linear dependence among these three factors, we follow the normalization method proposed by Deaton and Paxson (1994), requiring that the period dummies sum to zero ($\sum \psi_p = 0$) and are orthogonal to a linear time-trend ($\sum p\,\psi_p = 0$).[12] As a result, any linear time trend in the logarithm of nuptiality rates is attributed to ages and cohorts, rather than to periods. Such derived period effects consist of two components: 1) cyclical fluctuations (or macroeconomic shocks) that impinge on all cohorts and 2) interaction terms between the age and cohort effects. Deaton and Paxson only mentioned the macro shock part, while the second is our emphasis.

As defined in previous section, regressions of Equation (5) can be conducted for all 35 rounds, and there will be totally 900 ASFMRs ($=$ 36 ages $\times$ 25 years) regressed on the corresponding age, period, and cohort dummies. The quality of fit of the APC model can be expressed by R^2 which is defined as $1 -$ SS(Residual)/SS(Total Corrected), and the measure shows that our model performs very well (around 0.99) in explaining existing data in all rounds. Nevertheless, in a way different from other studies using APC models, this paper is focused on completing cohort incomplete marriage schedules rather than on identifying any individual effect of age, period, or cohort. To achieve this forecasting purpose, information regarding the out-of-range period parameters are necessary. Taking round 15

[12]We have explored some other normalizations and found there is almost no difference in completing cohort incomplete marriage schedules, except for occasional absurd situations in which marriage schedules might go up at oldest ages when other normalizations are adopted.

(where data are confined to years 1949–1973) as an example, all marriage schedules for cohorts 1927–1949 are incomplete. Period estimates outside the data range are required to complete the upper-left triangle (nuptiality at early ages for cohorts 1927–1933) and the lower-right trapezoid (nuptiality at older ages for cohorts 1927–1949). For instance, the first marriage schedule of cohort 1949 is incomplete due to the lack of ASFMRs at ages 25–50. Although estimates of cohort dummy 1949 and age dummies 25–50 are available from the regression, one can not yet derive projected values for this missing part without estimating the period dummies 1974–1999.

An assumption concerning how period effects change outside the data range is thus called for. Utilizing time series models based on estimates of existing period effects in the data range is one attractive option, and assuming that out-of-sample period effects equal the estimate for the nearest in-sample period is another. But which alternative is better? Can we have any clue to the answer? In fact we do have actual ASFMR data for these incomplete parts, which helps solve the problem associated with the out-of-sample period effects. Since estimates of the intercept, the age and the cohort effects (i.e., $\hat{\beta}$, $\hat{\alpha}$ and $\hat{\gamma}$) in Equation (5) have been obtained, out-of-range period effects can thus be estimated using the following nonlinear least-squares regression:

$$f(a, c)/\exp(\iota\hat{\beta} + A\hat{\alpha} + C\hat{\gamma}) = \exp(P\psi), \tag{6}$$

where $f(a, c)$ is now the ASFMR outside the data range, P denotes those out-of-range period dummies, and the vector ψ captures the out-of-range period effects.[13] In other words, peeking at these outside data reveals the *unknown* period effects which are depicted in Figure 4 by thin broken lines extended from both ends of the *existing* period effects (expressed by thick lines in the figure). Based on estimates of period effects outside the data range in selected rounds, Figure 4 shows that there appear some sharp turns upward or downward in estimates around the boundary points of the 25-year data range, which

[13] In estimating the out-of-range period effects, observations are discarded if the corresponding age is greater than 39, owing to their very low levels of nuptiality rate.

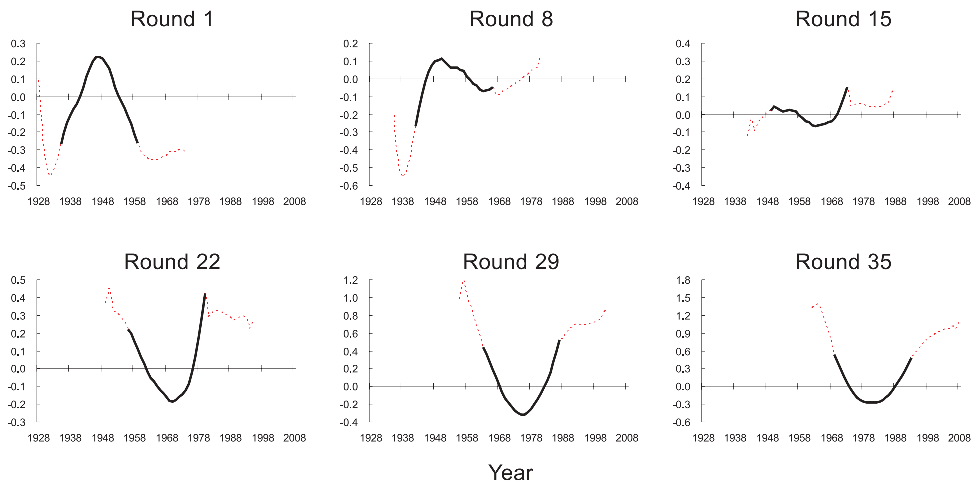

Figure 4. Period Effects Outside the Data Range in Selected Rounds

Note: Thick lines denote the existing period effects in the data range and thin broken lines extended from both ends represent the unknown period effects outside the data range.

might frustrate the ambition of utilizing time-series models to project out-of-range period effects. Besides, the assumption that out-of-sample period effects equal the estimate for the nearest in-sample period obviously fails to capture the reality.

In light of these "revealed" out-of-range period effects, we propose an assumption that:

Assumption. *Out-of-range period effects depart in a specific direction and then move along a path which is constrained to a zero slope some periods beyond the end of the existing data range.*

The main concern of the zero-slope constraint is to prevent possible absurd situations in which marriage schedules might go up rather than down at oldest ages. In addition, this paper specifies the form of the locus in outside period effects ψs as:

$$\psi_x = \begin{cases} \hat{\psi}_{\text{end}} + bx(12-x), & \text{if } x = 1, \dots, 6 \\ \psi_6, & \text{if } x > 6, \end{cases} \tag{7}$$

where x denotes the distance from the end of the existing data range and b is related to the departing direction. In sum, outside period effects move along a half parabola within 6 periods and turn horizontally beyond.

This assumption however presents its own problem, that is, to gauge in which direction the data will go. In short, b needs to be estimated. To estimate the directions along which out-of-range period effects depart, some additional information from outside the APC model are required. Let us temporarily *pretend* that the actual CTNRs of the 23 cohorts are somehow acquired. Since the ASFMRs included within the data range are known and can be summed up for any particular cohort, the discrepancy from its actual CTNR represents the value to which the outside projected nuptiality rates should be summed up. With these target values in hand, one can 1) estimate the departing directions, $\hat{b}$, 2) compute the out-of-range period effects, $\hat{\psi}_x$, 3) get the projected ASFMRs outside the data range, and 4) thus complete the incomplete marriage schedules. The outside projected values based on this assumption plus those actual ASFMRs within the data range construct a curve measure, denoted as APC-TRUE hereafter since it is the collaboration between the APC framework and the actual CTNR.

As a curve measure, the APC-TRUE can yield all distribution-related parameters. Therefore we use cohort total nuptiality rate (CTNR), cohort mean age at first marriage (CMAFM), and standard deviation to represent the quantum, the tempo, and the spread of a marriage schedule, respectively Figure 5 presents an overview of the performance of the APC-TRUE in four selected rounds, demonstrating that the APC-TRUE performs excellently not only in the estimation of the quantum (i.e., the marriage level) but also in the estimation of other nuptiality parameters. As a consequence, it has been approved that the period effect is the key to transforming a marriage level into a marriage schedule.

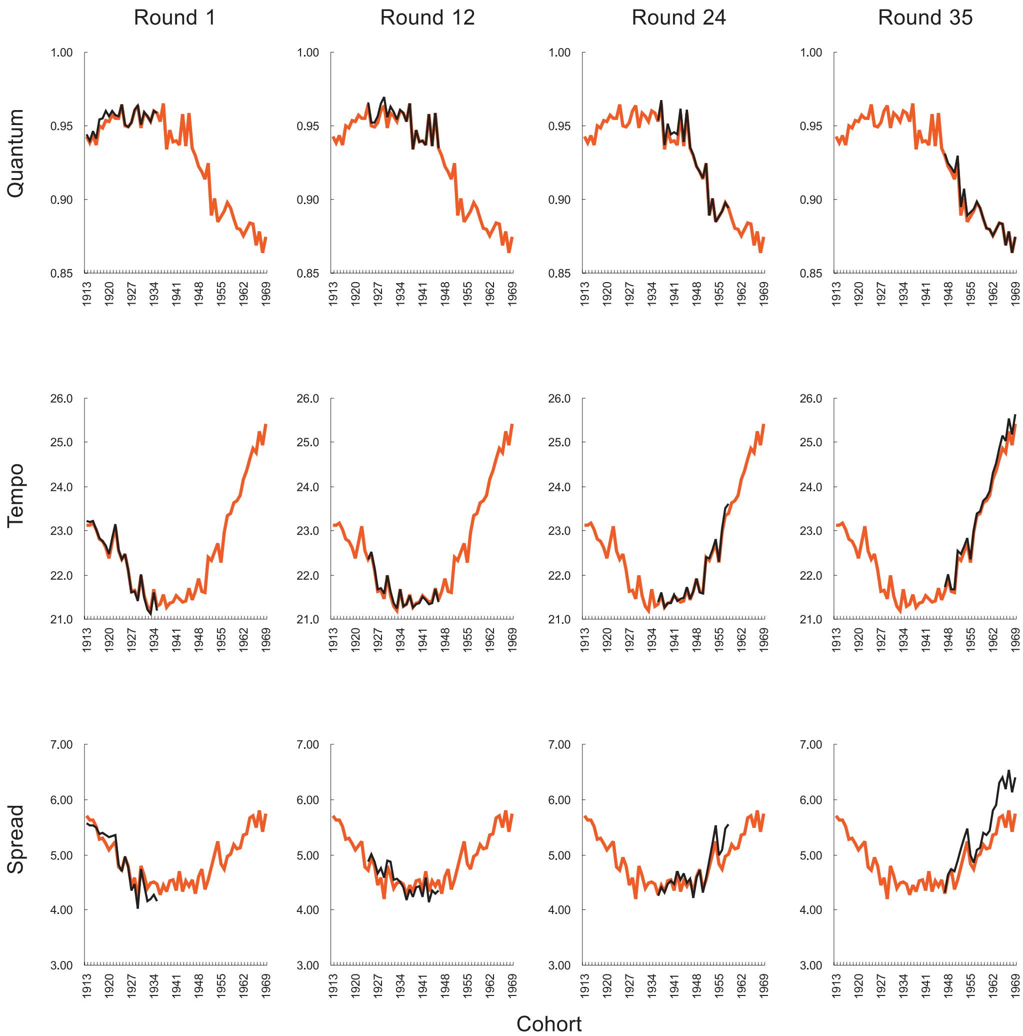

Figure 5. Performance of the APC-TRUE in Selected Rounds: Quantum, Tempo, and Spread

5 APC Model with the Smoothed KP

However excellent, the APC-TRUE is after all a *synthetic* measure since no actual CTNR for an incomplete cohort is available due to the data-range limitation. The next question is: Can we can find a good substitute for the actual CTNR? After investigating the effects of smoothing period nuptiality indicators in Section 3, we suggest that the smoothed KP can be a possible candidate, though not necessarily the best. In fact, our APC framework is ready to collaborate with any better method in forecasting the CTNR.

Denote the APC framework accompanied with the smoothed KP as APC-KP hereafter, Figure 6 presents an overview of its performance in four selected rounds. It is not unexpected that the APC-KP is outperformed by the APC-TRUE. But to fairly evaluate its performance, an appropriate way is to find some rivals to compete with, according to the accuracy of projections in the incomplete age pattern of nuptiality rates. Among methods aiming at forecasting incomplete cohort schedule based on age specific data, we consider the approach developed by Evans (1986) and the Coale-McNeil curve fitting model mentioned in Section 2.

In short, Evans's methodology uses a linear regression model to predict the ASFMRs after age 25 from the cumulated cohort nuptiality through age 24 and the ratio of nuptiality at ages 15–19 to that at ages 20–24. Owing to the estimation reason, however, 25-year data range is not long enough for this approach to work. We grant the rival 10 additional years in each round as a premium. In addition, while the APC-KP can estimate incomplete marriage schedules not only for young cohorts (whose future marriage has not been totally realized yet) but also for old ones (whose early marriage has been finished but unavailable from the data), Evans's approach is limited to forecasting incomplete parts for young cohorts only. The comparison is thus targeted at cohorts whose marriage has not been totally realized given the data range.

To save space, we target at the most recent cohort evaluated in each round. Specifically, the comparison is about how well the projected ASFMRs approximate the actual ones at

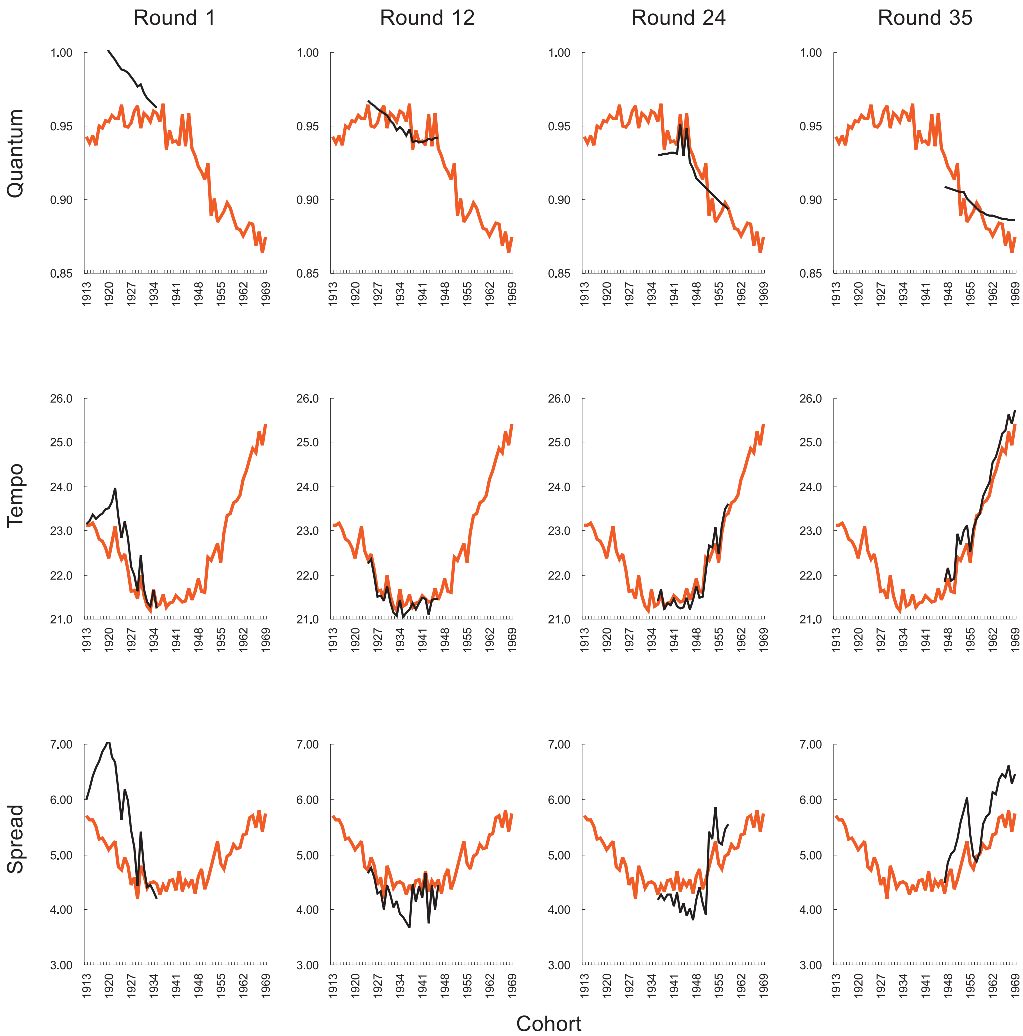

Figure 6. Performance of the APC-TRUE in Selected Rounds: Quantum, Tempo, and Spread

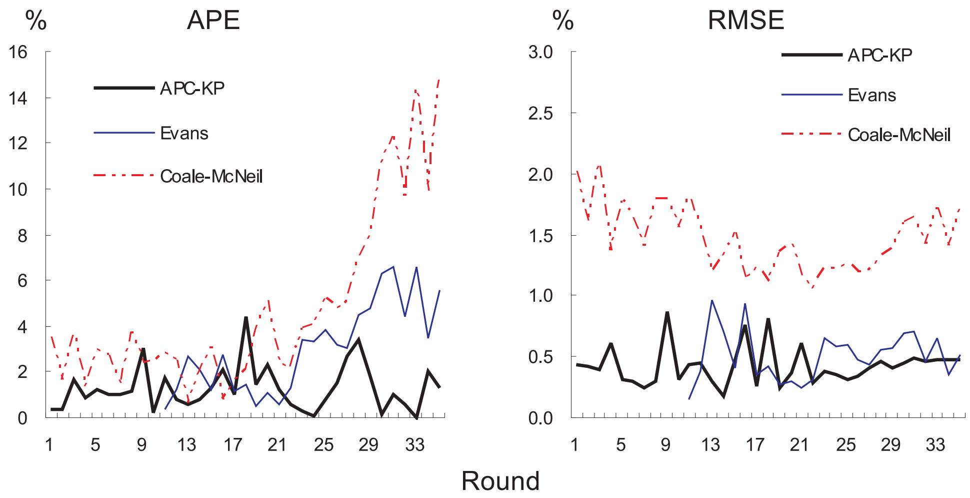

Figure 7. Performance of the APC-KP, Evans (1986)'s Method, and the Coale-McNeil Model in All Rounds

ages 25–49. Besides the APE criterion (of which the focus is on the marriage levels), we adopt the standardized root of mean square error (RMSE)

$$\text{RMSE}_m = \frac{\sqrt{\frac{1}{25}\sum_{a=25}^{49}\left[\hat{f}(a,c) - f(a,c)\right]^2}}{\sum_{a=25}^{49} f(a,c)} \tag{8}$$

as another criterion to compare the proximity of a projected marriage schedule to an actual one. From both panels in Figure 7, it is apparent that the Coale-McNeil model is outperformed by the other two methods. Since the information utilized by the Coale-McNeil model is pretty few, the result will not be unexpected. As for the performance of the APC-KP and the Evans method, it is not that clear to decide which approach is definitely better from Figure 7, even though the APC-KP wins in more rounds. However, since the APC-KP utilizes 10 years less than the Evans method, it merits better evaluations.

6 Summary

Unlike most previous studies which mainly focused on forecasting marriage levels only, this paper tries to accomplish cohort incomplete marriage schedules. We propose a simple age-period-cohort model to decompose the age-specific first marriage rates into age, period, and cohort effects, following the approach developed in Deaton and Paxson (1994) to solve the identification problem arising from the linear dependence of any one factor on the remaining two. From a theoretical perspective, we show that the period effect so derived represents not only macro shocks that impinge on all cohorts, but also the interaction between the age and cohort effects. Empirically, we indicate that the period effect is the key to transforming a marriage level into a marriage schedule, and related nuptiality statistics including the tempo, the spread, and other distributional parameters can be thereby derived. In this paper, we suggest that the smoothed version of the tempo-variance-adjusted period measure proposed in Kohler and Philipov (2001) can provide useful information on marriage level, and the collaboration between the APC framework and the smoothed KP performs better than some rival methods in approximating the cohort incomplete marriage schedules. All empirical results presented in this paper are fairly robust for we make efficient use of the 1935–1993 U.S. data. As a final emphasis, our approach is easy to implement and the data requirement is relatively light, indicating the proposed method is readily applied to countries whose data lengths are not long enough and would be helpful for further empirical investigation in the relationship between cohort nuptiality behavior and other cohort-specific socioeconomic factors.

References

Bongaarts, John and Griffith Feeney (1998), "On the Quantum and Tempo of Fertility", *Population and Development Review*, 24(2), 271–291.

——— (2006), "The Quantum and Tempo of Life-Cycle Events", *Vienna Yearbook of Pop-*

ulation Research 2006, 115–151.

Coale, Ansley J. and Donald R. McNeil (1972), "The Distribution of Age at First Marriage in a Female Cohort", *Journal of the American Statistical Association*, 67, 743–749.

Currie, Iain D., Maria Durban, and Paul H. C. Eilers (2004), "Smoothing and Forecasting Mortality Rates", *Statistical Modelling*, 4, 279–298.

De Beer, Joop (1985), "A Time Series Model for Cohort Data", *Journal of the American Statistical Association*, 80, 525–530.

Deaton, Angus S. and Christina H. Paxson (1994), "Saving, Growth, and Aging in Taiwan", in *Studies in the Economics of Aging*, edited by David A. Wise, Chicago: University of Chicago Press.

Evans, M. D. R. (1986), "American Fertility Patterns: A Comparison of White and Nonwhite Cohorts Born 1903–56", *Population and Development Review*, 12(2), 267–293.

Goldstein, Joshua R. and Catherine T. Kenney (2001), "Marriage Delayed or Marriage Forgone? New Cohort Forecasts of First Marriage for U.S. Women", *American Sociology Review*, 66, 506–519.

Hernes, Gudmund (1972), "The Process of Entry into First Marriage", *American Sociological Review*, 37, 173–182.

Kim, Young J. and Robert Schoen (2000), "On the Quantum and Tempo of Fertility: Limits to the Bongaarts-Feeney Adjustment", *Population and Development Review*, 26(3), 554–559.

Kohler, Hans-Peter and José Antonio Ortega (2002), "Tempo-Adjusted Period Parity Progression Measures, Fertility Postponement and Completed Cohort Fertility", *Demographic Research*, 6, 92–144.

Kohler, Hans-Peter and Dimiter Philipov (2001), "Variance Effects in the Bongaarts-Feeney Formula", *Demography*, 38(1), 1–16.

Li, Nan and Zheng Wu (2003), “Forecasting Cohort Incomplete Fertility: A Method and An Application”, *Population Studies*, 57(3), 303–320.

Ryder, Norman B. (1990), “What Is Going to Happen to American Fertility”, *Population and Development Review*, 16(3), 433–454.

Schoen, Robert (2004), “Timing Effects and the Interpretation of Period Fertility”, *Demography*, 41(4), 801–819.

Silverman, Bernard W. (1996), “Smoothed Functional Principal Components Analysis by Choice of Norm”, *The Annals of Statistics*, 24(1), 1–24.

Smallwood, Steve (2002), “The Effect of Changes in Timing of Childbearing on Measuring Fertility in England and Wales”, *Population Trends*, 109, 36–45.

Sobotka, Tomáš (2003), “Tempo-Quantum and Period-Cohort Interplay in Fertility Changes in Europe: Evidence from the Czech Republic, Italy, the Netherlands and Sweden”, *Demographic Research*, 8, 151–214.

van Imhoff, Evert (2001), “On the Impossibility of Inferring Cohort Fertility Measures from Period Fertility Measures”, *Demographic Research*, 5, 23–64.

van Imhoff, Evert and Nico Keilman (2000), “On the Quantum and Tempo of Fertility: Comment”, *Population and Development Review*, 26(3), 549–553.