Title
A Cohort Model of Fertility Postponement

Journal
Demography, 51(5)

ISSN
0070-3370

Authors
Goldstein, Joshua R
Cassidy, Thomas

Publication Date
2014-10-01

DOI
10.1007/s13524-014-0332-7

Peer reviewed

HHS Public Access
Author manuscript
Demography. Author manuscript; available in PMC 2015 October 01.

Published in final edited form as:
Demography. 2014 October ; 51(5): 1797–1819. doi:10.1007/s13524-014-0332-7.

A Cohort Model of Fertility Postponement

Joshua R. Goldstein and
Max Planck Institute for Demographic Research and Department of Demography, University of California, Berkeley, 2232 Piedmont Avenue, Berkeley CA 94720

Thomas Cassidy
Department of Mathematics, Bucknell University, Lewisburg, PA 17837

Abstract

We introduce a new formal model in which demographic behavior such as fertility is postponed by differing amounts depending only on cohort membership. The cohort-based model shows the effects of cohort shifts on period fertility measures and provides an accompanying tempo adjustment to determine the period fertility that would have occurred without postponement. Cohort-based postponement spans multiple periods and produces "fertility momentum," with implications for future fertility rates. We illustrate several methods for model estimation and apply the model to fertility in several countries. We also compare the fit of period-based and cohort-based shift models to the recent Dutch fertility surface, showing how cohort- and period-based postponement can occur simultaneously.

1 Introduction

One can view much of the demographic change that is taking place in advanced societies as a result of the changing meaning of age. We often hear that 40 is the new 30, or even that 80 is the new 60. Demographers have developed formal models to show that shifting age schedules (or equivalently, shifting meanings of age) can produce dramatic changes in cross-sectional period measures. Most notably, the "tempo-adjusted total fertility rate" introduced by Bongaarts and Feeney (1998) has become an important part of the modern demographic toolkit.

The transformation of the human life cycle is a process that takes place within individual lives and is thus most naturally conceptualized as a cohort process. The magic of the Bongaarts and Feeney (BF) "tempo adjustment" is that only period data are needed to perform the correction: postponement is modeled in period manner, with all ages (and thus all cohorts) postponing their events in the same manner in a given year. This rate of change can vary from period to period, but within a period, all cohorts are treated the same. An enormous advantage of this approach—apart from any degree of realism it may or may not have—is that it produces a simple mathematical model in which changes in period mean ages completely determine the presence and extent of tempo effects.

Criticisms have been leveled at the BF approach. Notably, a number of authors have stated that the uniform postponement across all ages is unrealistic (e.g., Kim and Schoen 2000). However, few authors have shown the consequences of departure from this assumption or have proposed alternatives. Zeng and Land (2001) have shown, using a set of simulations, that violations of the uniform postponement by age assumption matter relatively little. Kohler and Philipov (2001) proposed an important extension to the BF approach, giving a more general framework for age- and time-varying postponement and offering a special case in which postponement differed linearly by age within any period. The Kohler and Philipov results have not been widely used, in part because the study is complex, the assumptions required to estimate the model are somewhat stylized, and the estimation procedure requires strong smoothing in order to produce interpretable results. The BF method itself is known to be subject to "seemingly random year-to-year fluctuations" caused by "sensitivities to small errors" and "deviations from the constant shape assumption" (Bongaarts and Feeney 2006, p. 131).

An alternative to period-based fertility shifts is a cohort-based model in which the past influences future behavior. In this formulation of postponement, each cohort experiences a fixed change in fertility timing. The cohort-shift model allows postponement choices made in the past to contribute to a "fertility momentum" that plays out in terms of evolving fertility rates over the life of cohorts. In contrast, period-based postponement has no necessary implications for an ongoing transformation in fertility rates.

The cohort perspective on fertility postponement harkens back to Ryder's (1964) original formulation of quantum and tempo effects on fertility, which was based on the changing mean ages of cohort age schedules. The ideas of Lesthaeghe and Willems (1999) about postponement and recuperation also have a cohort basis, as does the average completed fertility (ACF) measure developed by Butz and Ward (1979) and analyzed by Schoen (2004).[1] Empirical efforts to disentangle period and cohort effects on fertility were undertaken by Foster (1990) and reviewed by Ní Bhrolcháin (1992), both of whom concluded that period quantum effects appear to dominate. More recently, Bongaarts and Sobotka (2012) used stylized simulation to argue that period postponement is more common than cohort postponement.

The first contribution of the present study is to formalize the cohort postponement model, compare it with other fertility models, derive methods for fitting the model, and present applications. We introduce an adjusted measure of period total fertility, analogous to the BF measure, which can be used to recover the total fertility that would have been observed in the absence of cohort postponement. We show how the cohort shift can be estimated from data and apply our adjustment formula to real data from several countries to see how it compares with the total fertility rate (TFR) and to the BF adjustment (TFR^*). We also present a combined model of fertility postponement that encompasses simultaneous cohort and period shifts, and we offer a method for estimating this model. The combined model is

[1]The $TFR^\dagger$ measure we introduce turns out to be equivalent to ACF in the special case of unchanging period and cohort quantum, which gives ACF a formal behavioral foundation. However, in general, tempo-adjusted fertility under cohort shifts does not equal ACF.

compatible with the cohort-shift model studied in this article and with the period-shift model used elsewhere.

Given any model of fertility behavior, it is natural to ask how closely the model matches reality. One innovation in this study is to introduce goodness of fit as a criterion for comparing models. We demonstrate this criterion by using data from Holland to investigate whether actual fertility changes are better modeled via cohort shifts or period shifts. Our goal in this demonstration is not to perform an exhaustive empirical comparison, but rather to make precise an alternative model of fertility change, describe its qualitative features, and show how data can be used to compare different models of tempo and quantum changes over the Lexis surface.

2 Models of Cohort- and Period-Based Postponement

In this section, we introduce a model of fertility in which timing shifts are cohort driven, and we illustrate some of its distinctive features. We first give a verbal characterization before turning to the formal description. We contrast this model with the period-shift model, and we show how both types of timing changes can be combined in a single model. Each model comes with a tempo-adjustment formula that allows us to recover the fertility rate that would have occurred had there been no shifts. None of these models can completely describe the full variety of fertility rates on the Lexis surface. Rather, a model serves as a simplified description of both fertility decisions and broad trends in fertility. Later in the article, we explore how closely the various models match the observed data.

2.1 Behavioral Interpretations

The cohort-shift model decomposes observed fertility into an interaction of cohort-based decisions and period-based events. Fluctuations in period fertility are then understood as the result of both (1) *cohort* tempo, which refers to changes in period fertility that result from the timing of births in cohorts; and (2) *period* quantum, or level changes in period fertility that result from events that are independent of age and cohort.

We represent cohort tempo as the outcome of shifts that describe how much each cohort may have advanced or delayed its fertility schedule. Different generations will have different plans for the timing of childbearing that plays out over the course of their lives. These underlying schedules of intended fertility then encounter period-driven events or shocks that may ultimately reduce or increase the cohort's total fertility. This model, a mixture of cohort and period influences, captures both the lifetime implications of cohort fertility intentions and the immediate responses to unanticipated period events. It is the interplay of cohort plans and period events that produces variety in the observed fertility surface.

Because modern contraceptive technologies enable women to control when they have children, fertility at any one age is no longer disconnected from fertility at other ages. Rather, lower fertility at young ages tends to be followed by higher fertility at older ages; conversely, higher fertility and young ages tends to be followed by lower fertility at older ages. A decision to postpone childbirth until after completing college or after establishing a career has implications for years beyond the initial decision.

Modeling shifts in fertility as a cohort driven process allows the decision to postpone fertility to have implications that span multiple periods. Because life-course expectations can also evolve over time in response to changing circumstances, it is useful to allow for period specific changes in total fertility. These quantum fluctuations are a function of time only, and so they provide a measure of changing fertility relevant to specific periods.

Fertility can also be modeled entirely in terms of periods, with both postponement shifts and quantum changes responding only to temporal events. A major advantage of such a model is the ease of calculation it affords. Although the practical value of an easily determinable model is great, the behavioral interpretation of such of model is more difficult. A challenge that other authors (Kohler and Philipov 2001) have noted is that different ages may respond to a period shock in different ways: for example, compared with older women, younger women may be more likely to choose to postpone in response to a period shock. An additional challenge is that the period model allows for the complete independence of behavior at different ages across a given cohort. So, for example, postponement at younger ages has no implication for the fertility of the same cohort at older ages. The estimation advantages of basing tempo adjustment on a single period are offset, we argue, by excluding the potentially useful information that comes from surrounding years.

In “tempo models,” cohort fertility targets are not explicit parameters (cf. Lee 1980). Instead, the underlying timing of fertility is emphasized. However, it is possible to reconcile these models with cohort targets by thinking of a cohort’s initial target as some level of period quantum held constant over the cohort’s reproductive span. Period quantum variations can be interpreted as reassessments of cohort intentions in light of changing circumstances. The models are flexible enough to allow recovery of lost fertility (say, because of a recession) by structuring the time series of period quantum values.

2.2 Formalization of Postponement Models

Here, we formalize period postponement, cohort postponement, and a model that combines both types of fertility postponement. Denote the fertility rate at age a and time t by $f(a, t)$, with the fertility rate at age a of the cohort born at time c given by $f(a, c + a)$. Let $f_0(a)$ denote a normalized, standard baseline fertility schedule that sums to 1.

Period-based postponement is represented by $R(t)$, which denotes the total amount by which women in period t have shifted their fertility. We introduce cohort-based postponement through a shift function $S(c)$. The shift $S(c)$ indicates the amount by which women from cohort c have shifted their fertility. For example, if “40 is the new 30” for the cohort of 1960, then $S(1960) = 10$. Simultaneous period and cohort postponement occurs as a combination R and S.

Our first postponement model has been studied extensively in the literature and is the basis for the BF tempo adjusted fertility rates.

Author Manuscript

Author Manuscript

Author Manuscript

The Period-Shift Model of Fertility

$$f(a,t)=f_0(a-R(t))(1-R^{'}(t))q(t). \quad (1)$$

Notice that the period-shift model includes the term $(1 - R'(t))$, which lowers or raises the fertility level depending on whether women have delayed or advanced fertility. This term appears because timing changes that occur within cohorts spread or consolidate period births, an effect that is not found in the cohort counterpart to the period-shift model.

The Cohort-Shift Model of Fertility

$$f(a,t)=f_0(a-S(t-a))q(t). \quad (2)$$

These models are derived as follows. Let $F_0(a)=\int_0^a f_0(x)dx$ be the cumulative fertility for the baseline schedule. We assume that $F_0(\omega) = 1$, where ω is sufficiently large so as to be well beyond the last age of fertility, leaving room for shifts. Under the period-shift model, the shifted cumulative fertility schedule for women born in year c and currently of age a is $F_0(a - R(c + a))$. Differentiating cumulative fertility with respect to age a gives us $f_0(a - R(c + a))(1 - R'(c + a))$. To obtain the observed fertility rate $f(a, t)$, we include period quantum effects $q(t)$ and rewrite cohort in terms of period and age.[2]

The shifted cumulative fertility schedule for the cohort-shift model is $F_0(a - S(c))$. Differentiating cumulative fertility with respect to age a gives us $f_0(a - S(c))$, the fertility rate for cohort c at age a in the absence of any period quantum effects. To obtain the observed fertility rate $f(a, t)$, we include period quantum effects and rewrite in terms of period and age.

Both formulations assume period quantum effects, but in Eq. (2), shifts are cohort-based rather than period-based. The cohort-shift model shares with Eq. (1) the assumption that both shift effects and period-quantum effects and are independent of age. Nonetheless, the cohort model in Eq. (2) allows the shape of both period and cohort fertility schedules to vary, since shifting cohort fertility timing causes period fertility schedules to fluctuate by period, and changing period quantum means that that the shape of the cohort fertility schedules is also not constant. We will see in the next section that the cohort-shift model implies its own tempo-adjusted measure of period fertility.

Cohort and period shifts are not incompatible ideas. Both types of postponement can be encompassed in one model.

The Combined Model of Shifted Fertility

$$f(a,t)=f_0(a-S(t-a)-R(t))(1-R^{'}(t))q(t). \quad (3)$$

[2]See Rodriguez (2006) for a lucid derivation of the period-shift model.

This combined model[3] includes as special cases both the cohort-based model explored in depth in this article (setting $R = 0$) and the period-based model used by Bongaarts and Feeney (setting $S = 0$). As with the previous models, the combined model is obtained by differentiating the cumulative fertility $F_0(a - S(c) - R(c + a))$ with respect to age, rewriting cohort in terms of period and age, and then including period quantum effects.

2.3 Shift-Adjusted Period TFR

In this section, we show how to adjust period TFR to compensate for the tempo distortion caused by timing changes. Shift-adjusted (or tempo-adjusted) total fertility is a measure of the TFR that would have been observed in the absence of timing changes. Because the baseline schedule f_0 has been normalized to 1, this unshifted TFR is exactly the period quantum term $q(t)$. A shift-adjusted fertility rate is therefore a method for obtaining $q(t)$ from the observed fertility rates.

Each of the models presented in the previous section has a corresponding adjustment factor that, when multiplied by the observed fertility rates, enables one to recover $q(t)$. For the period-shift model of Eq. (1), $q(t)$ is obtained from observed rates via the BF adjustment procedure,

$$\mathrm{TFR}^{*}(t)=\frac{\int f(a,t)da}{1-R'(t)}.$$

To see that $\mathrm{TFR}^*(t) = q(t)$, use Eq. (1) to replace $f(a, t)$ with $f_0(a - R(t))(1 - R'(t))q(t)$. Substituting $w = a - R(t)$ and $da = dw$ produces the desired result. We refer to the quantity $1/(1 - R'(t)$ as the adjustment factor for period t.

For the cohort-shift model of Eq. (2), $q(t)$ can be recovered from observed rates by defining the shift-adjusted period TFR as

$$\mathrm{TFR}^{\dagger}(t):=\int_0^{\omega} f(a,t)(1+S'(t-a))da, \quad (4)$$

where $S'(t - a)$ is the derivative of $S(c)$ evaluated at $t - a$, the incremental shift relevant for the cohort aged a at time t. To show that $\mathrm{TFR}^{\dagger}(t) = q(t)$, use Eq. (2) to replace $f(a, t)$ with $f_0(a - S(t - a))q(t)$. Substituting $w = a - S(t - a)$ and $dw = (1 + S'(t - a))da$ gives the desired result. Thus, the shift-adjusted $\mathrm{TFR}^{\dagger}(t)$ recaptures the period TFR that would have been observed in the absence of cohort shifts. We present a method for estimating the adjustment factor $1 + S'$ in the appendix.

Definition (4) works because the age shifts from cohort to cohort are recapitulated in the cross-section from age to age. Although a shift in fertility timing does not compress the fertility experience for any cohort (see Fig. 1a), increasing postponement effectively speeds up the clock in a synthetic cohort within a given period (see Fig. 1b). If postponement

[3]Compare this to Kohler and Philipov's (2001) equation 3.

increases over cohorts, then in a fixed time period, the passage from one age to the next will cover a greater range of the fertility behaviors implied by the baseline schedule. Conversely, if successive cohorts have earlier birth schedules, then within a given period, the synthetic cohort will have more exposure at the relevant age-specific fertility rate (not shown). In order to exactly compensate for the compression or extension of age introduced by the cohort shifts, $\mathrm{TFR}^{\dagger}$ can be interpreted as the result of either inflating or deflating the time spent at each age. The individual weights, $1 + S'(t - a)$, tell us how to inflate the observed fertility in order to account for the cross-sectional compression of age that results from increasing postponement.[4]

For the combined model of period and cohort shifts given in Eq. (3), the adjustment factor needed to recover $q(t)$ is given by

$$\frac{1+S'(t-a)}{1-R'(t)}.$$

We present a method for estimating this adjustment factor from the observed fertility rates in the appendix. With this adjustment factor, we can define shift-adjusted period total fertility as

$$\mathrm{TFR}^{\dagger+*}(t):=\int f(a,t)\frac{1+S'(t-a)}{1-R'(t)}da.$$

That this formula reproduces the desired period quantum $q(t)$ can be verified by replacing $f(a, t)$ with $f_0(a - S(t - a) - R(t))(1 - R'(t))q(t)$ and then substituting $w = a - S(t - a) - R(t)$ and $dw = (1 + S'(t - a))da$ before integrating.

2.4 Fertility Momentum

A feature of the cohort-approach to fertility postponement is that future trends in fertility are determined in part by the timing of cohorts already having children. By analogy with population momentum, we can consider fertility "momentum," defining it within the cohort model as the trajectory of period fertility if postponement were to come to an instant halt. More specifically, freeze cohort postponement by letting $S(c) = S(c_0)$ for all cohorts born after c_0, the youngest cohort having children at some time t_0. Furthermore, freeze period quantum at the level observed at time t_0. Under this scenario, period fertility will converge gradually to $\mathrm{TFR}^{\dagger}(t_0)$.[5]

The cohort formulation of fertility postponement has inherent implications for future fertility rates and therefore can be used to forecast TFR under the assumption of no changes in quantum. For all currently active cohorts, we look at the timing of births to deduce the

[4]Equation (4) can be interpreted in two ways. We can see the adjustment factor as modification to the observed rates—that is, $\mathrm{TFR}^{\dagger}$ is $\int[f(a, t)(1 + S'(t - a))]da$. Alternatively, we can think of $1 + S'$ as a scaling factor applied to each age—that is, $\mathrm{TFR}^{\dagger} = \int f(a, t)[(1 + S'(t - a))da]$.

[5]From Eq. (2), $\mathrm{TFR}(t) = \int f_0(a - S(t - a))q(t_0)\, da$. As $S(t - a)$ becomes constant over all childbearing ages, $\mathrm{TFR}(t) \to q(t_0)$ as $t \to \infty$.

intended fertility schedule for these cohorts. We then assume that the newer cohorts who have not yet begun to give birth will follow the timing of the most-recent cohorts. (This assumption is relatively unimportant because in the short run, the newer cohorts will contribute much less to fertility than will the currently active cohorts.) By projecting fertility schedules for each cohort individually, we can predict TFR in subsequent years. If the currently active cohorts have experienced different schedules, then we will see fertility rates evolve for several years even if quantum and shifts are now fixed.

Figure 2 shows TFR in the Czech republic from 1950 to 2010, along with a forecast of TFR for the next 10 years. This forecast is based on the observed timing of births for cohorts active in 2010, along with the assumption that the timing of births for newer cohorts is the same as that for the most recent cohort. This projection shows fertility rates rising for the next 10 years as result of older cohorts fulfilling postponed fertility plans. Although this forecast assumes no change in period quantum, if reasonable quantum predictions are available, one could build a projection of fertility rates that combines forecasted quantum with observed cohort shifts. An equivalent projection is not possible in the period-shift model because the current timing of births in the present provides no indication of future timing, even for cohorts that are already in the midst of childbearing.

2.5 Comparison With BF Tempo Adjustment

The BF adjustment factor, used to weight the observed fertility rates, is

$$\frac{1}{1-R'(t)}. \quad (5)$$

Under the assumptions of the period-shift model of Eq. (1), $R'(t)$ can be estimated using the time derivative of period mean age at birth. The BF tempo-adjusted TFR is then

$$\mathrm{TFR}^{*}(t)=\frac{\int f(a,t)da}{1-\mu'(t)},$$

where μ' is the derivative of period mean age at birth.

To understand how TFR^{*} and $\mathrm{TFR}^{\dagger}$ can differ, consider the shift pictured in Fig. 1, where we portray a rise in cohort-based postponement that spans multiple cohorts but eventually stabilizes. Because this idealized scenario is purely cohort-based, $\mathrm{TFR}^{\dagger}$ will be constant throughout the transition. In contrast, TFR^{*} will fluctuate. At the onset of postponement, fertility rates change only at the very youngest ages. Because fertility is comparatively low at these ages, TFR is only slightly depressed by this change. However, because these ages are furthest from the mean age at birth, this change in rates has a disproportionate impact on the period mean age at birth; thus, TFR^{*} will overcompensate for the decline in TFR. Similarly, as the end of postponement brings TFR close to its original level, fertility rates change only at the highest ages. This produces an outsized change in mean age at birth, and TFR^{*} again overshoots the ideal tempo-adjusted level.

It is also possible for TFR* and TFR† to produce identical results. Consider the combined model of period and cohort shifts given in Eq. (3). If the cohort shift S is a linear function of c, the derivative of period mean age at birth produces the desired factor $(1 + S')/(1 - R')$. Likewise, if the period shift S is linear in t, the cohort mean age at birth estimator presented in the appendix is also equal to $(1+S')/(1 - R')$, showing that TFR* and TFR† can coincide. See the last subsection of the appendix for more on distinguishing between period and cohort shifts.

3 Applications to Modern Fertility Patterns

To calculate the cohort shift–adjusted TFR from data, we need estimates of the adjustment factor $1 + S'(c)$ for each cohort c. Here, we present a method for deriving $1 + S'(c)$ from data. We then apply these estimation techniques to determine TFR† for Holland, Sweden, the Czech Republic and the United States. Finally, we use numerical techniques to fit the Dutch fertility surface with both the cohort shift and the period-shift models in order to make a goodness-of-fit comparison.

3.1 Estimation

The rate of change in cohort shifts, $S'(c)$, can be estimated in a variety of ways. Our preferred method for incomplete cohorts is to use the rate of change in the truncated mean age at birth, after taking period quantum effects into account.

Our formula for estimation is

$$S'(c)=\frac{\mu'(c)}{1+v(h,c)(\mu(c)-h)+v(l,c)(l-\mu(c))}, \quad (6)$$

where $\mu(c)$ is the truncated mean age of the quantum-adjusted fertility schedule of the cohort born in year c; $\mu'(c)$ is the rate of change in this mean age; v is a truncation adjustment; and l and h are, respectively, the “lowest” and “highest” observed ages. Note that $\mu(c)$ is not the mean age of childbearing of the cohort, which depends on all previous period shocks. The details and derivation of this equation are given in the first subsection of the appendix.

3.2 Applications to Holland, the United States, the Czech Republic, and Sweden

In this section, we show the results of fitting the period and cohort models to a range of populations from developed countries in the Human Fertility Database (n.d.). Holland and the United States are examples of countries that had quite low fertility fairly early and may be nearing the end of their postponement transitions. The Czech Republic exemplifies the sudden period onset of postponement after the fall of Communism and a dramatic decline and recovery of period fertility. Finally, Sweden has experienced quite dramatic period fluctuations in total fertility and has also been used by Kohler and Philipov (2001) as an example of how the BF approach might need improvement.

For the estimates of TFR†, we have applied our estimation to fertility of all parities combined. We use the truncated-means estimator discussed earlier in the text and described in detail in the appendix. For the estimates of TFR^* we use the Bongaarts-Feeney recommendation of combining separate estimates of adjusted total fertility by parity.

$TFR^\dagger$ can also be estimated separately by parity, but we find that this is not necessary in practice. Indeed, a potential advantage of the cohort-postponement approach is that it appears to work quite well or even better when all births are combined, which means that the analysis can be applied in countries where parity-specific data are not available. For a cohort, the link in fertility timing between parities is mechanical because postponement of low parities will be followed by a corresponding postponement of higher parities. This connection need not hold within periods.

In terms of data requirements, the BF approach typically uses a minimum of three annual periods of age-specific fertility rates in order to produce a centered estimate of tempo changes and the corresponding tempo-adjusted fertility rate. In addition, Bongaarts and Feeney recommend (1) estimating tempo effects separately by each parity, and (2) smoothing the resulting TFR in order to remove “seemingly random fluctuations.” (Bongaarts and Feeney 2006, page 131.)

In theory, the cohort model can be estimated from data as limited as those used for the BF method. In practice, however, more periods of observation yield more-reliable estimates of cohort timing. Complete cohort histories are not required because of methods that can estimate shifts for both right- and left-censored cohorts (see the appendix). In practice, a decade or two of period fertility rates appears sufficient to obtain plausible estimates of $TFR^\dagger$. We experimented with the estimation methods by artificially truncating cohorts and found them to be quite robust as long as truncation remains beyond the age of peak fertility. For earlier ages, estimates are less reliable.

Looking across all four populations displayed in Fig. 3, we see several common features. First, from a broader view, the cohort and period approaches produce qualitatively similar stories: they both show that tempo-adjusted fertility is higher than the observed TFR after about 1970 in Holland, the United States, and Sweden, and after 1990 in the Czech Republic. Second, closer examination reveals that the period approach produces more-volatile estimates with more year-to-year change. In contrast, the cohort approach tracks the ups and downs of the TFR, with the magnitude of the tempo adjustment changing slowly over time. The greater volatility in the period estimator occurs because the tempo effect for each period is independently estimated, whereas the smoothness of the cohort approach is a consequence of the gradual change in the composition of cohorts from one period to the next.

Looking at Holland, the cohort approach tells us that fertility quantum has been nearly unchanged since the mid-1970s; by contrast, the period approach suggests a small boom in quantum in the late 1980s and early 1990s, a time of rapid fertility postponement but also one of relatively poor economic conditions in Holland. Finally, the values of tempo-adjusted fertility estimated around 2010 suggest that there is still a considerable tempo effect according to the period model (on the order of 0.1 children); this lies in contrast to the finding of no effect according to the cohort approach.

The United States presents an example similar to that of the Netherlands. The period approach shows a boom in quantum from the early 1990s to the mid-2000s, whereas the cohort approach suggests no such boom.

In the Czech Republic, the period and cohort tempo adjustments diverge most sharply. This is a case in which we know *a priori* that there was a sudden period onset of postponement. Indeed, the period-adjustment procedure gives more plausible results than the cohort-adjustment procedure, which suggests implausibly high period quantum before the fall of the Wall. The cohort estimates have problems because we are estimating a single index of postponement for cohorts living across clearly distinct periods. If we estimate $TFR^{\dagger}$ using only the period before 1990 or only the period after 1990 (not shown), we find that $TFR^{\dagger}$ tracks TFR^{*} fairly closely. In practice, this means that period adjustment is to be preferred in the presence of a known period change in fertility tempo regimes.[6]

Finally, in Sweden, the cohort and period approaches produce quite similar estimates except for the period around 1990, for which the convergence of TFR^{*} and TFR suggests that period postponement came to an complete halt. This is the period analyzed by Kohler and Philipov (2001), whose suggested variance-adjusted procedure produces an estimate nearly identical to the cohort tempo adjustment shown here. A potential advantage of the cohort approach over that introduced by Kohler and Philipov is that it is conceptually and analytically simpler. A second advantage is that it does not require any arbitrary smoothing, which Kohler and Philipov used to reduce volatility. Instead, the smoothing of periods comes part and parcel with the use of the cohort approach.

Overall, we find that the cohort approach is capable of producing tempo-adjusted fertility measures that are less volatile and, in many cases, more plausible than the period procedure. However, as the Czech example shows, there are contexts in which the cohort approach is not appropriate and for which the period model can produce more sensible results.

3.3 Goodness-of-Fit Comparisons of Cohort and Period Shifts

Ever since Bongaarts and Feeney's 1998 paper, demographers have been adjusting period measures and developing models based on the assumption that fertility shifts are period-based. We have presented an alternative formulation based on cohort shifts. It is natural to ask whether one of these formulations is closer to reality. In this section we compare the fit of period-based and cohort-based shift models to recent age-period surfaces. We recognize that neither model can be a perfect description of reality, and our goal is not to make a general statement that tempo changes are always period-based or always cohort-based. Rather, we aim to introduce a method for evaluating such questions and to show that the period formulation fits better for some populations and time intervals but the cohort formulation describes others.

Figure 4 illustrates fit of the period- and cohort-shift models to observed fertility schedules. At four periods in Holland, we see the actual fertility schedule and the best-fit

[6]The same difficulty does not apply to period changes in fertility quantum, which is incorporated in the model for cohort postponement with period quantum (Eq. (2)).

approximations from both the cohort-shift model and the period-shift model. (The fitting was carried out according to the method discussed in the next two paragraphs.) Both the cohort and period models provide good approximations of the observed fertility schedules, with the cohort model fitting somewhat better for two of these periods. One can characterize the overall fit of the model by summing the absolute errors over age for each period. We find the average period error from 1970 to 2008 to be 0.13 children for the optimized period-shift model, compared with an error of 0.10 children for the optimized cohort-shift model.[7] (These can be compared with TFR values of about 2 children.)

Denote the predicted fertility surface from the period-shift model as

$$\hat{f}^{I}(a,t)=\hat{f}_0^{I}(a-\hat{R}(t))(1-\hat{R}^{'}(t))\hat{q}^{I}(t),$$

where the superscript I is used to identify the period-shift parameters and associated predicted values, $\hat{R}(t)$ is estimated cumulative period shift up to time t, and $\hat{R}'(t)$ is this quantity's time derivative. The predicted surface of the cohort-shift model is

$$\hat{f}^{II}(a,t)=\hat{f}_0^{II}(a-\hat{S}(t-a))\hat{q}^{II}(t),$$

with the superscript II identifying the cohort-shift parameters and predicted values, and $\hat{S}(t - a)$ being the estimated cumulative cohort shifts. Note that the baseline schedules in the two models $\hat{f}_0^{I}$ and $\hat{f}_0^{II}$ differ. Likewise the two estimated period quantum $\hat{q}^{I}$ and $\hat{q}^{II}$ can differ.

Numerical optimization can be used to minimize the sum of squared differences between $\hat{f}(a, t)$ in either of the two models and the observed values $f(a, t)$. A computationally intensive approach is to use numerical methods to estimate not only the baseline fertility schedule and resulting inference of period quantum but also the shift parameters. An abridged approach is to estimate the shift parameters $\hat{R}$, $\hat{R}'$ using the methods recommended by Bongaarts and Feeney and the $\hat{S}$ terms using the methods given earlier in the Estimation section. We present the abridged result here. This abridged version is faster to calculate, and it allows a comparison of the models as they are typically applied.[8]

Table 1 presents more-comprehensive results by parity and time interval in terms of the sum of squared errors over the entire Lexis surface. The units here are different, but the models can still be compared directly, with smaller errors indicating better fit.

Comparison of the period and cohort approaches involves several choices. First, because Bongaarts and Feeney clearly recommended that their model be fit by parity, we optimize the fit of the period-shift model separately by parity. Second, we need to chose the time interval covered and the shape of the region on the Lexis surface. A rectangular age-period

[7]We thank an anonymous reviewer for suggesting that we express the error in units that can be directly compared with the TFR.

[8]The results in this section were calculated using the *nlm*() function in R.

region will involve more cohorts than periods; an age-cohort parallelogram will involve more periods than cohorts.

The results of the estimation are shown in Table 1. Using the first line of Table 1 to compare the period and cohort models, we see that the cohort formulation performs better when we do not distinguish parity. This makes sense: even with period-based postponement of first births, the sequential nature of birth order creates a cohort pattern of postponement for higher-order births. This result confirms Bongaarts and Feeney's recommendation of calculating TFR^* by parity.

Considering each parity separately, we see that the period formulation is superior for 1960–2007 for parities 1 and 2. However, the 1960–2007 interval mixes at least three different chapters in the evolution of fertility in Holland: the end of the baby boom, the introduction of the Pill, and the post-1975 trend toward postponing births. The first two of these are understood to be period phenomena, particularly the introduction of the Pill.

If we focus attention on recent postponement, where we think cohort shifts are most likely to be found, it is useful to restrict ourselves to the post-1975 interval. Here, the cohort model indeed tends to fit slightly better than the period model, particularly for first births (see Table 1).

A comprehensive analysis involving more populations and time intervals is needed for more-definitive results. However, we can tentatively conclude at least two things. First, in the absence of parity-specific data, the cohort-postponement approach appears to better describe fertility change than the period-based approach, at least in Holland. Second, with parity-specific data, it is still possible for the cohort model to outperform the period model. In the post-1975 interval, we find that the cohort approach is slightly superior to the period approach.

The goodness-of-fit comparison in Table 1 includes an additional possible source of error: namely, the potential for nonoptimal values of the parameters $R(t)$, $R'(t)$, and $S(t - a)$. A more definitive comparison can be made using the fully optimized approach to estimate these parameters in addition to the baseline fertility schedule.

4 Discussion

In this article, we have introduced a model of cohort fertility postponement combined with period effects on the level of fertility. The model is a cohort analogy to the period perspective underlying Bongaarts and Feeney's tempo adjustment measures.

Presenting the models in mathematical terms reveals the similarities and differences between the period and cohort approaches. In essence, both the strengths and the weaknesses of the period approach come from the localization of change to a particular moment in time. Likewise, the strengths and weaknesses of the cohort approach are due to the spanning of multiple periods by consecutive cohorts.

Our derivation of a tempo-adjusted measure for cohort postponement shows that it is possible to remove the effect of changing cohort timing of fertility, just as the BF measure is capable of removing the effect of period-based timing changes. The cohort-shift model can be estimated through several methods. We also showed that the period and cohort approaches can, in theory, be combined. (We also provided an estimator for the combined model but found it to be somewhat unstable. Here, further research could be useful.)

The introduction of a formal model underlying period and cohort adjustment allows us to introduce a new criterion for judging models of fertility change over age and time: goodness of fit. We found that the cohort model applied to Holland is at least as appealing as the period model in terms of goodness of fit for describing fertility change across the Lexis surface. (Again, further research could be useful here on other countries and time periods.)

The appeal of the cohort model is severalfold. First, it has some behavioral appeal in terms of life-course theory: each cohort has its own intended schedule of births by age, which then is updated by period events. Second, the effect of changes in cohort timing on period fertility can be described as the compression (or expansion) of the passage of time for the synthetic cohort constituting experiencing period fertility rates. Third, the model appears to describe better some of the recent history of fertility change.

The disadvantage of the cohort model in practice is that it is considerably more challenging to estimate. We needed to invent new techniques in order to produce estimates that would be consistent with the model. Even these techniques, however, cannot solve the fundamental challenge of cohort analysis, which is that we observe only a small portion of the lives of cohorts alive at the beginning and end of the portion of the Lexis surface for which we have observations. Although we obtained estimates for all cohorts, the estimates for highly truncated cohorts are unavoidably based on little information. Fortunately, these uncertainties for the youngest cohorts have little influence on the cohort tempo-adjusted TRF because most fertility occurs at older ages.

In practice, we found that the cohort model produces plausible estimates of cohort shifts and shift-adjusted total fertility in all countries we studied. In some cases, the estimates look somewhat better than those of the BF period approach, with less high-frequency volatility as well as less evidence of mini baby booms in quantum (e.g., Holland in the years around 1990).

Our goal in presenting the cohort approach is not to argue that it is necessarily a better approach in all cases. Rather, our hope is to extend the perspectives used by demographers for thinking about and analyzing fertility change. Neither the period nor the cohort fertility postponement approach represents the “true” description of how fertility changes over time, nor how individuals respond to changing conditions.[9] However, we believe that together they provide a more complete picture of how to model fertility change. This is very much in line with the ideas of Burch (2004) and the urgings of Ní Bhrolcháin (2011), emphasizing models as modes of thought and explanation.

[9]See [Lee 1980] for an alternative approach focusing on stocks and flows rather than postponement.

In terms of measurement, the cohort approach has the advantage of builtin smoothing through the process of translation from cohort to period. Multiple cohorts are alive in any given period, and in the cohort approach, each one of them influences its respective age group in a different way. Thus, each period is influenced by a kind of moving filter of cohorts. However, this same process also makes the cohort approach potentially limiting. Because the filter changes slowly, with only one new cohort being introduced and only one cohort exiting each year, the magnitude of the “tempo adjustment” can change only slightly from year to year. Sudden period changes in the tempo adjustment factor are simply not possible in the cohort framework.

We hope that an important contribution of our current efforts is the use of goodness of fit as a criterion for comparing models. The approach used here, and extensions of it, should assist in the evaluation of existing and new models of fertility change over time.

Acknowledgments

This work was partially supported by a grant from the Simons Foundation (Grant No. 210199 to Thomas Cassidy).

References

Bongaarts J, Feeney G. On the quantum and tempo of fertility. Population and Development Review. 1998; 24:271–291.

Bongaarts J, Feeney G. The quantum and tempo of life-cycle events. Vienna Yearbook of Population Research. 2006; 4:115–151.

Bongaarts J, Sobotka T. A demographic explanation for the recent rise in European fertility. Population and Development Review. 2012; 38:83–120. [PubMed: 22833865]

Burch TK. On teaching demography: Some non-traditional guidelines. Canadian Studies in Population. 2004; 31:135–142.

Butz WP, Ward MP. The emergence of countercyclical U.S. fertility. American Economic Review. 1979; 69:318–328.

Foster A. Cohort analysis and demographic translation: A comparative study of recent trends in age specific fertility rates from Europe and North America. Population Studies. 1990; 44:287–315.

Human Fertility Database. Rostock, Germany, and Vienna, Austria: Max Planck Institute for Demographic Research and Vienna Institute of Demography; n.d. Retrieved from http://www.humanfertility.org

Kim YJ, Schoen R. On the quantum and tempo of fertility: Limits to the Bongaarts-Feeney adjustment. Population and Development Review. 2000; 26:554–559.

Kohler HP, Philipov D. Variance effects in the Bongaarts-Feeney formula. Demography. 2001; 38:1–16. [PubMed: 11227839]

Lee RD. Aiming at a moving target: Period fertility and changing reproductive goals. Population Studies. 1980; 34:205–226. [PubMed: 22077121]

Lesthaeghe R, Willems P. Is low fertility a temporary phenomenon in the European Union? Population and Development Review. 1999; 25:211–228.

Ní Bhrolcháin M. Period paramount? A critique of the cohort approach to fertility. Population and Development Review. 1992; 18:599–629.

Ní Bhrolcháin M. Tempo and the TFR. Demography. 2011; 48:841–861. [PubMed: 21643883]

Rodriguez G. Demographic translation and tempo effects. Demographic Research. 2006; 14(article 6): 85–110.10.4054/DemRes.2006.14.6

Ryder NB. The process of demographic translation. Demography. 1964; 1:74–82.

Schoen R. Timing effects and the interpretation of period fertility. Demography. 2004; 41:801–819. [PubMed: 15622955]

Zeng, Yi; Land, KC. A sensitivity analysis of the Bongaarts-Feeney method for adjusting bias in observed period total fertility rates. Demography. 2001; 38:17–28. [PubMed: 11227842]

Zeng, Yi; Land, KC. A sensitivity analysis of the Bongaarts-Feeney method for adjusting bias in observed period total fertility rates. Demography. 2002; 39:269–285. [PubMed: 12048952]

Appendix

In this appendix, we present methods for deriving the adjustment factors $1+S'$ or $(1+S')/(1-R')$ from data. We also explain the inherent difficulty in distinguishing period-based postponement from cohort-based postponement. If shifts are purely period-driven, as modeled in Eq. (1), then the adjustment factor is $1/(1-R'(t))$, and this quantity can be calculated from changes in the period mean age at birth. If shifts are purely cohort-driven, as in Eq. (2), then the adjustment factor is $1+S'(c)$, which can be calculated from changes in the (truncated) mean age at birth for cohorts, as shown below. If shifts are a combination of period and cohort factors (Eq. (3)), the necessary adjustment factor is $(1+S')/(1-R')$. We show, in the second subsection of this appendix, that the ratio of the derivatives of period and cohort fertility with respect to age gives an approximation of the adjustment factor for the combined model. We do not need to estimate either the shift function S or the baseline schedule f_0 because $\text{TFR}^{\dagger}$ is calculated from the observed fertility schedule f.

For both the cohort-shift model and the combined model, it is necessary to control for the influence of changes in the period quantum q. This is achieved by choosing an easily identifiable feature that indicates quantum. We employ the height of the mode for this purpose, but other features could also work. For example, when studying mortality, one might choose the minimum value of the hazards. Our "quantum-normalized fertility rates," $\tilde{f}(a,t)$, are calculated as follows. In each period t there is an age (or ages) with the peak fertility rate, and we denote this rate by $m(t)$. We can easily calculate $m(t)$ from observed data as $\text{Maximum}_a\{f(a,t)\}$. We then define the quantum-normalized rate as

$$\tilde{f}(a,t):=\frac{f(a,t)}{m(t)}.$$

Let m_0 denote the peak value of the baseline schedule f_0. Under the assumptions of purely cohort shifts, $m(t)=m_0q(t)$, and the quantum-normalized rates are

$$\tilde{f}(a,t)=\frac{1}{m_0}f_0(a-S(t-a)). \quad (7)$$

Under the assumptions of combined period and cohort shifts (Eq. (3)), $m(t)=m_0q(t)(1-R'(t))$, and the quantum-normalized rates are

$$\tilde{f}(a,t)=\frac{1}{m_0}f_0(a-S(t-a)-R(t)). \quad (8)$$

Note that although one could also employ the quantum-normalized rates when estimating the adjustment factor for the period-shift model, this will have no impact because raising or lowering the overall level of fertility will not change the mean age at birth.

Estimation Using (Truncated) Mean Age at Birth

The BF adjustment uses changes in period mean age at birth to estimate the factor R'. If cohort c has completed its years of fertility, we can use the change in the cohort quantum-adjusted mean age at birth to estimate $S'(c)$. For cohorts that have not yet completed the fertile years, we can estimate $S'(c)$ using the change in mean age at birth truncated to the latest year of data.

Our formula depends on several quantities. Let $\mu(c)$ be the quantum-adjusted mean age at birth for cohort c as of latest available age; that is,

$$\mu(c)=\frac{\int_l^h x\tilde{f}(x,c+x)\,dx}{\int_l^h \tilde{f}(x,c+x)\,dx}, \quad (9)$$

where l is the lowest age of available fertility data for cohort c, and h is the highest age of available fertility data for cohort c. Because fertility data come in discrete form, we use an interpolating function to calculate integrals of $\tilde{f}$. Let $v(a, c)$ be the proportion of cohort c's truncated fertility that occurs at age a calculated from the quantum-adjusted schedules; that is,

$$v(a,c)=\frac{\tilde{f}(a,c+a)}{\int_l^h \tilde{f}(x,c+x)\,dx}. \quad (10)$$

Let $\mu'(c)$ be the derivative of $\mu(c)$. In practice this can be estimated using

$$\frac{\mu(c+1)-\mu(c-1)}{2},$$

provided the same l and h are used for cohorts $c + 1$ and $c - 1$.

Use Eq. (7) to replace $\tilde{f}(a, t)$ with

$$\frac{1}{m_0}f_0(a-S(t-a))$$

in Eq. (9). We get

$$\mu(c)=\frac{\int_l^h x f_0(x-S(c))\,dx}{\int_l^h f_0(x-S(c))\,dx}.$$

Because *l* and *h* are constants, differentiating this new expression for $\mu(c)$ with respect to *c* and using the rule $(p/g)' = p'/g - (p/g)(g'/g)$ gives us

$$\mu^{'}(c)=\frac{-S^{'}(c)\int_l^h x f_0^{'}(x-S(c))dx}{\int_l^h f_0(x-S(c))dx}-\mu(c)\frac{-S^{'}(c)\int_l^h f_0^{'}(x-S(c))dx}{\int_l^h f_0(x-S(c))dx}. \quad (11)$$

The left-hand fraction in Eq. (11) can be evaluated via integration by parts as

$$\begin{aligned} S^{'}(c)&\frac{l f_0(l-S(c))-h f_0(h-S(c))+\int_l^h f_0(x-S(c))\,dx}{\int_l^h f_0(x-S(c))\,dx}\\ &=S^{'}(c)\cdot(lv(l,c)-hv(h,c)+1)\,. \end{aligned}$$

The right-hand fraction in Eq. (11) is

$$-\mu(c)S^{'}(c)\frac{f_0(h-S(c))-f_0(l-S(c))}{\int_l^h f_0(x-S(c))dx}=-S^{'}(c)\mu(c)\ (v(h,c)-v(l,c))\,.$$

Subtracting this second fraction from the first gives us

$$\begin{aligned}\mu^{'}(c)&=S^{'}(c)\cdot(lv(l,c)-hv(h,c)+1)+S^{'}(c)\mu(c)\,(v(h,c)-v(l,c))\\ &=S^{'}(c)\cdot(v(l,c)(l-\mu(c))-v(h,c)(h-\mu(c))+1)\,,\end{aligned}$$

and solving this for $S'(c)$ yields this estimate of $S'(c)$:

$$S^{'}(c)=\frac{\mu^{'}(c)}{1+v(h,c)(\mu(c)-h)+v(l,c)(l-\mu(c))}. \quad (12)$$

Estimation Using First Derivatives of Period and Cohort Fertility

For the combined model of cohort and period postponement given in Eq. (3), the adjustment factor needed to recover $q(t)$ from the observed fertility rates is $(1+S')/(1-R')$. In this section, we show how to estimate this factor using a ratio of partial derivatives. Because these derivatives tend to be very small, this estimation technique is much more volatile than the method using mean age at birth presented in the previous section.

In the absence of postponement, the derivatives of period fertility and cohort fertility with respect to age should be the same. With postponement, this ratio of derivatives measures the degree to which age has been compressed in a period; thus, multiplying fertility rates by this ratio gives us the appropriate quantity to include in our adjusted total fertility rate. In mathematical terms, this means that

$$\frac{1+S'(t-a)}{1-R'(t)}=\frac{\tilde{f}_a(a,t)}{\tilde{f}_a(a,t)+\tilde{f}_t(a,t)}=\frac{\frac{d}{da}\tilde{f}(a,t)}{\frac{d}{da}\tilde{f}(a,c+a)}. \quad (13)$$

Note that because both derivatives in this ratio can be zero at ages that correspond to peak levels of fertility and also at ages with no fertility, those ages cannot be used to calculate the adjustment factor.[10]

To see why the first equality in Eq. (13) holds, notice that by Eq. (7), $\tilde{f}(a, t) = f_0(a - R(t) - S(t - a))/m_0$. Differentiating gives us

$$\tilde{f}_a(a,t)=f_0'(a-R(t)-S(t-a))(1+S'(t-a))/m_0,$$

and

$$\tilde{f}_t(a,t)=-f_0'(a-R(t)-S(t-a))(S'(t-a)+R'(t))/m_0.$$

Although Eq. (13) suffers from some instability near the ages at peak fertility, it has the advantage of being agnostic toward the origin of shifts; that is, if shifts are purely cohort-based or purely period-based, this equation still approximates the correct adjustment factor. In contrast, Eq. (12) is applicable only if shifts can be modeled entirely as cohort-based; likewise, $1/(1 - \mu'(t))$ provides the appropriate adjustment factor only if shifts can be modeled as entirely period-based.

Distinguishing Period Shifts From Cohort Shifts

The combined model of fertility postponement (Eq. (3)) allows for concurrent cohort-based and period-based shifts, denoted by S and R, respectively. It is natural to wonder whether either of these two factors dominates. In this section, we show that up to linear terms, it is not possible to distinguish between cohort and period shifts from observed data. Nonetheless, it is still possible to estimate the tempo-adjusted TFR without specifically

[10]The second derivative analog to Eq. (13) is

$$\frac{\tilde{f}_{aa}+\tilde{f}_{at}}{\tilde{f}_{aa}+2\tilde{f}_{at}+\tilde{f}_{tt}}.$$

In theory, this equation will produce the same results as Eq. (13) and could be used at ages with peak fertility. However, the estimation of second derivatives from discrete data can be very noisy.

separating cohort and period components, as demonstrated in the previous subsection of this appendix.

We start with the combined postponement model of Eq. (3):

$$f(a,t) = q(t)(1 - R'(t)) f_0(a - R(t) - S(t - a)).$$

Expand R and S via Taylor series as $R(t) = r_0 + r_1 t$ plus terms of degree 2 and higher, and $S(c) = s_0 + s_1 c$ plus terms of degree 2 and higher. We may replace $R(t)$, $S(c)$, and f_0 without changing any values of $f(a, t)$ or $q(t)$ as follows:

$$\tilde{R}(t) := \frac{R(t) - r_0 - r_1 t}{1 - r_1}, \quad \tilde{S}(c) := \frac{S(c) + r_0 + r_1 c}{1 - r_1}, \quad \text{and} \quad \tilde{f}_0(a) := (1 - r_1) f_0(a - r_1 a).$$

Alternatively, we may replace $R(t)$, $S(c)$, and f_0 with

$$\check{R}(t) := \frac{R(t) + s_0 + s_1 t}{1 + s_1}, \quad \check{S}(c) := \frac{S(c) - s_0 - s_1 c}{1 + s_1}, \quad \text{and} \quad \check{f}_0(a) := (1 + s_1) f_0(a + s_1 a),$$

and again f and q remain unchanged at all ages and times.

As an example, suppose that shifts are linear and purely period-based, so that $S(c) = 0$ and $R(t) = r_0 + r_1 t$. This scenario is indistinguishable from a linear cohort shift with

$$R(t) = 0, \quad S(c) = \frac{r_0 + r_1 c}{1 - r_1} \quad \text{and} \quad \tilde{f}_0(a) = (1 - r_1) f_0(a(1 - r_1)).$$

In this case, the BF period-based tempo adjustment and our cohort-based 1 tempo adjustment both produce the same result because $\frac{1}{1 - R'(t)}$ and $1+\tilde{S'}$ are both equal to $1/(1 - r_1)$. This shows that when the period schedule shifts at a rate of r_1, the cohort schedule shifts at a rate of $r_1/(1 - r_1)$, which is the result obtained by Zeng and Land (2002). Similarly, linear cohort shifts can be reinterpreted as linear period shifts; thus, if the cohort schedule shifts at a rate of r_1, the period schedule shifts at a rate of $r_1/(1 + r_1)$, and the same tempo-adjusted TFR is obtained from either adjustment procedure.

Because the linear terms in R and S are interchangeable, there can be no definitive way to disentangle cohorts and periods. This unavoidable ambiguity between R and S suggests that neither the magnitude nor the direction of the shifts is intrinsic to the cohort or period approaches. This inherent undecidability between period and cohort makes efforts such as those by Bongaarts and Sobotka (2012) a considerable challenge.

(a)

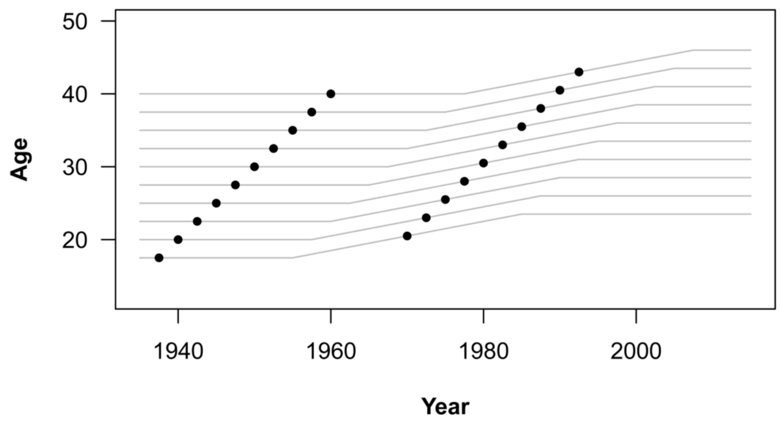

(b)

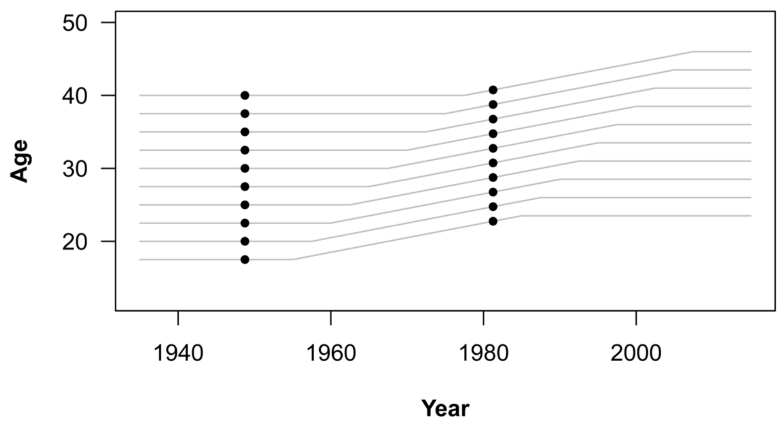

Figure 1.

A hypothetical fertility contour map with cohort shifts. Cohort fertility is unchanged during the shift, but synthetic cohort fertility is compressed during the shift years

(a) The trajectories of two cohorts. The first completes fertility before shifts begin. The second undergoes fertility during the shift years.

(b) The fertility trajectories of two synthetic cohorts, one predating the shifts and one during a period with shifts.

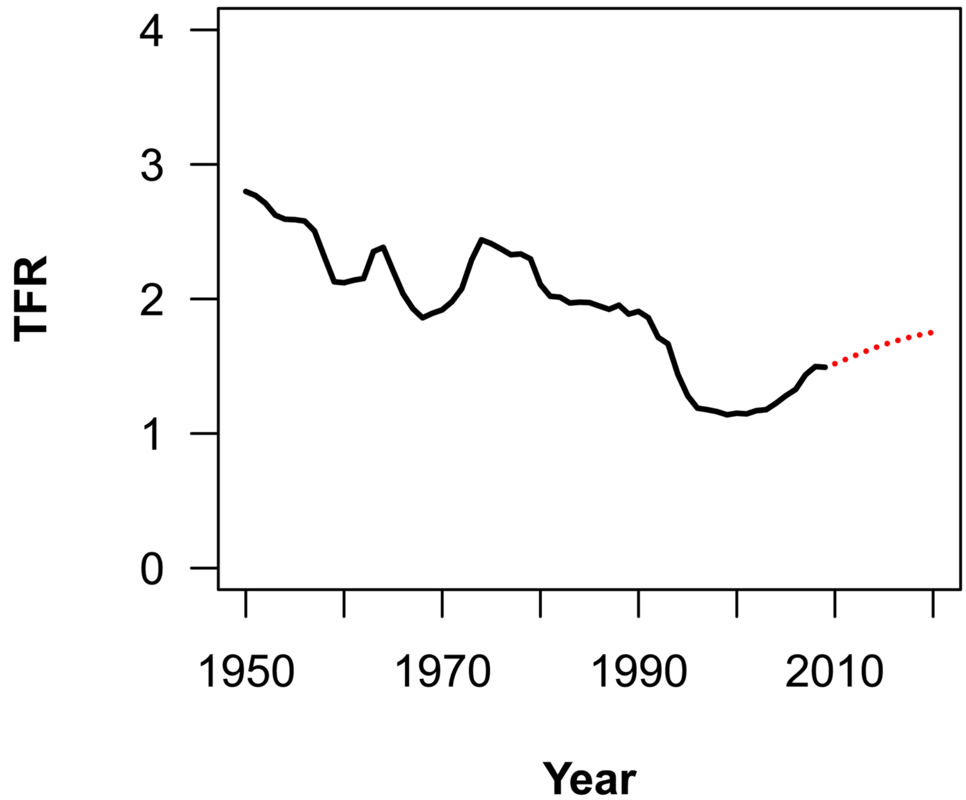

Figure 2.
A cohort-based forecast of TFR in the Czech Republic

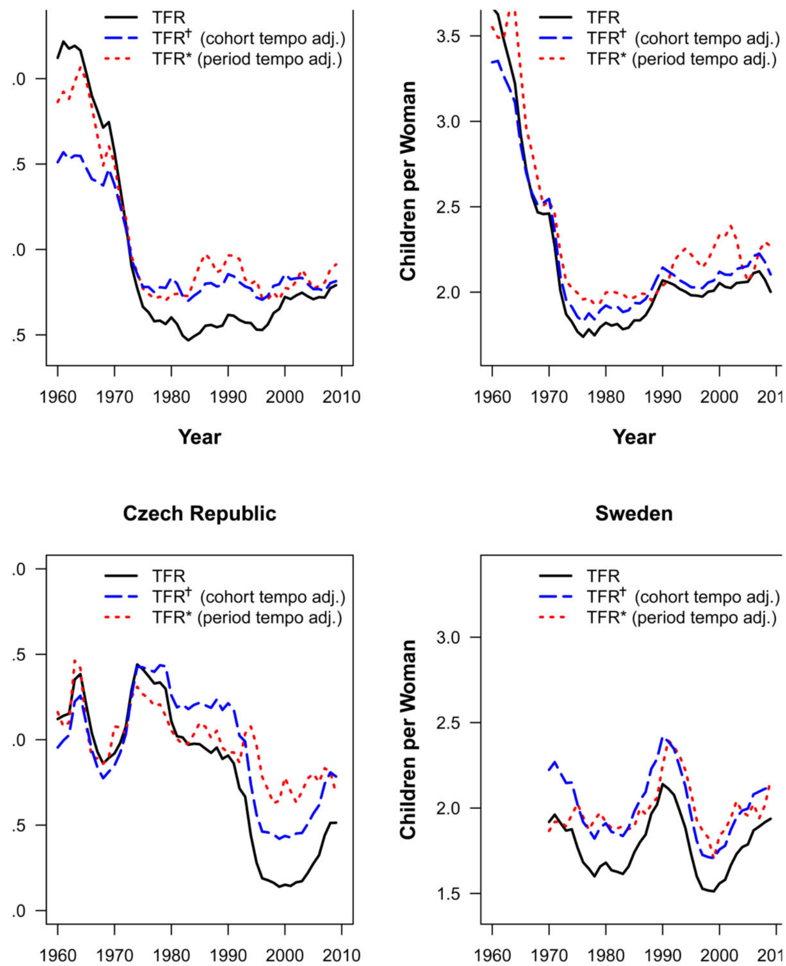

Figure 3.

TFR, TFR†, and TFR* in the Netherlands, the United States, the Czech Republic, and Sweden

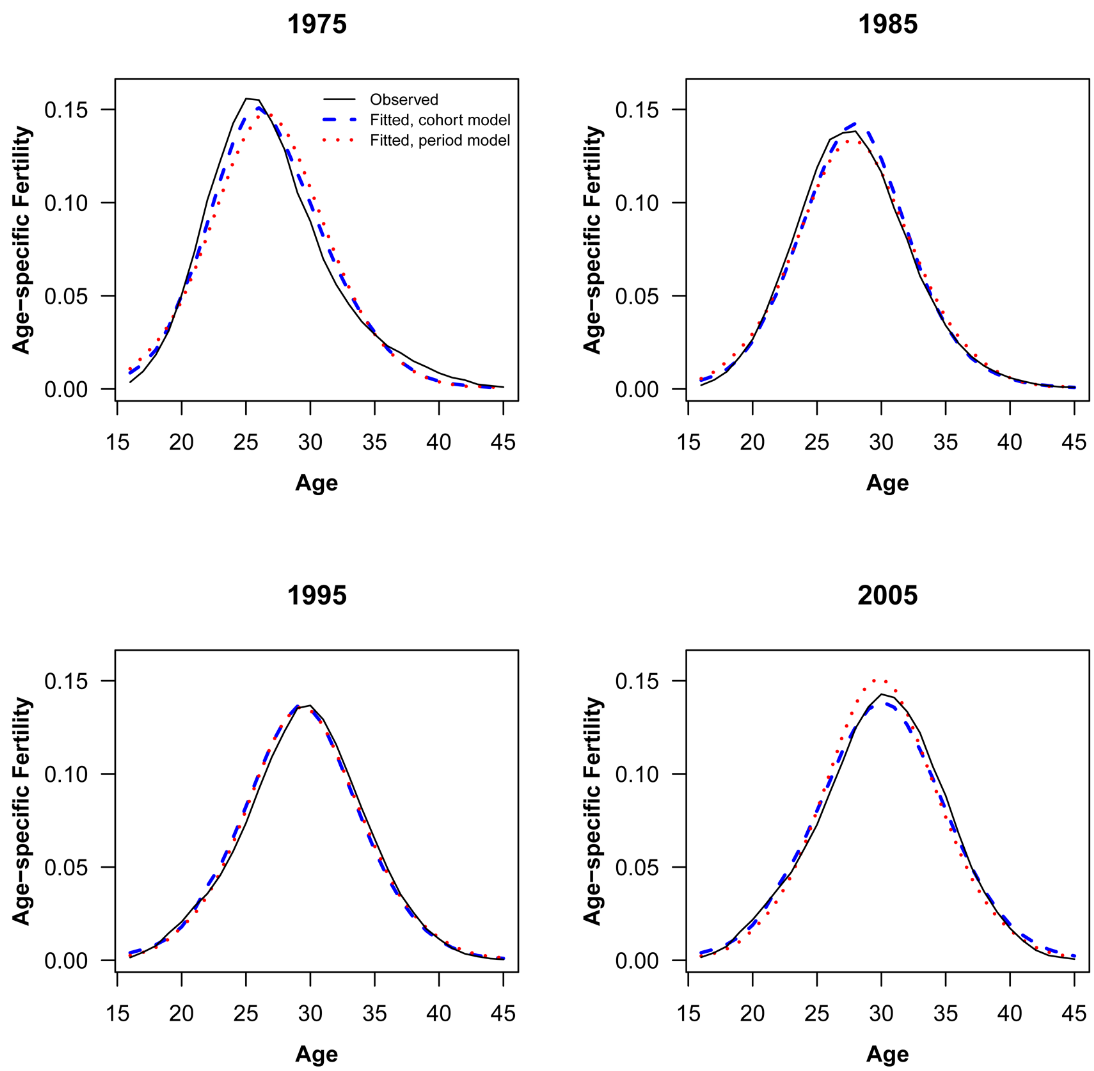

Figure 4.
Selected period fertility age schedules in Holland, optimal fitted schedules from the period- and cohort-shift models, and the observed schedule

Table 1

Goodness-of-fit comparisons over two time intervals in Holland, with the smaller sum of squared errrors (SSE) of fertility rates indicating a closer fit between the model and the observed data

	Holland 1960–2008		Holland 1975–2008	
	Period Fit SSE	**Cohort Fit SSE**	**Period Fit SSE**	**Cohort Fit SSE**
All Parities	0.231	**0.078**	0.039	**0.021**
Parity 1	**0.029**	0.038	0.013	**0.008**
Parity 2	**0.014**	0.018	0.007	**0.005**
Parity 3 +	0.009	**0.003**	0.001	0.001
All Parities, Fit Separately	0.097	**0.093**	0.038	**0.020**

Notes: The smaller residuals are marked in bold, indicating the better fit. "All Parities" refers to the sum of the residuals for fertility rates without regard to parity; that is, $\Sigma_{a,t}\,(f(a,t) - \hat{f}(a,t))^2$, where $\hat{f}(a,t)$ is the optimal fit to the all parity surface. "All Parities, Fit Separately" is calculated by taking the square of the difference between the observed fertility and the fertility schedule made from the sum of the optimal schedules for each parity; that is, $\Sigma_{a,t}(f(a,t) - \Sigma_i \hat{f}_i(a,t))^2$.