Tempo Effects and Demographic Measurement:
Toward a Unified Framework for Analysis and Interpretation

################### or ###################

A Simple Explanation of Tempo Effects in Fertility – and Why Such Issues are Irrelevant for Mortality Analysis

John R. Wilmoth
Department of Demography
University of California, Berkeley

September 15, 2004

Extended Abstract of a Proposed Paper
Session on "Mathematical Demography" (Nico Keilman, chair)
General Conference of the International Union for the Scientific Study of Population
Tours, France, July 2005

Overview of Project and Paper

It goes without saying that the effects of timing changes on standard demographic measures, such as the total fertility rate (TFR) and life expectancy at birth (e0), have become a central preoccupation of methodologically oriented demographers in recent years. Following the publication of two widely-discussed papers by Bongaarts and Feeney, in 1998 and 2003, it became clear that this is one topic that formal demographers care about very much, but about which we seem to disagree almost fiercely.

The paper I am proposing will be nothing less than my third attempt to formulate a coherent argument on this topic. The first paper was intended largely as a verbal critique of the BF tempo-adjusted measure of total fertility, with supplementary modeling and simulations in the second paper. After months and indeed years of work, I abandoned any thought of publishing these unfinished works, because I wanted to approach this topic in a more constructive manner. That was the genesis of the current paper.

Bongaarts and Feeney have proposed tempo-adjusted measures of both fertility and mortality. Despite the fact that the usual focus of my research has been on mortality, I have devoted most of my thinking about tempo issues to an analysis, reformulation, and elaboration of the BF fertility model. I now propose a general model of the effects of timing changes in fertility on standard demographic measures. From my perspective, it is important to emphasize the possibility of multiple interpretations of tempo change, and thus to eschew the notion of an adjusted "tempo-free" measure with general applicability.

Although this new fertility model is less than a week old at this writing, I am cautiously optimistic about its usefulness for clarifying tempo issues in fertility. The BF fertility model is the simplest case of this much more flexible model. In the general model, tempo change is defined by linear shifts in percentiles, with slopes that can vary across the distribution of ages at birth. With the sole assumption that the quantum of cohort fertility is unaltered (both in total and within each percentile segment), it is possible to derive a measure of the effect that such a timing change would have on period fertility (both age-specific and total rates) in the complete absence of changes in cohort quantum.

In other words, the model is used to derive a hypothetical "pure" tempo effect, which would alter the level of period fertility by a certain amount completely on its own. A shortcoming of the current approach is that it does not (yet) address the issue of simultaneous changes in cohort quantum, or interactions between tempo and quantum effects (it may be possible to add such things later on). By formalizing the problem in this general way, we see much more clearly the importance of the key assumption in the BF adjustment procedure concerning the unchanging level of cohort quantum. This assumption is not problematic so long as one adheres to a proper interpretation based on clearly articulated hypothetical situations.

In fact, it is possible to turn around the hypothetical tale of tempo effects, in order to address a fundamentally different question. For example, instead of asking how much period fertility might be suppressed purely as a result of tempo delay (assuming constant cohort quantum), it is also possible (and perhaps equally valuable) to ask how much cohorts in this situation would need to raise their completed fertility above 2.1 (in a low-mortality population) in order to keep the period TFR at replacement level. One insight that emerges from an analysis of tempo effects in fertility is that full population replacement requires that cohorts who delay their fertility must pay a price in terms of a higher level of completed fertility than would otherwise be necessary for this purpose. This is an important point about what the BF adjusted measure does *not* tell us – despite comments in their original article suggesting that tempo-adjustment demonstrates that low-fertility populations today are not as far from replacement level as everyone seems to think. In such a context, the (unadjusted) TFR is the correct measure, and the cohort CFR would need to be tempo adjusted in order to make it comparable to the TFR, whenever there are substantial changes in fertility timing.

In summary, concerning tempo effects in fertility, I am in full agreement that such effects exist and are crucial for understanding the relationship between cohort and period fertility levels. My new model provides, I believe, the most general approach yet proposed for measuring and analyzing such processes. It is worth noting that it seems possible, using this model, to discern tempo effects on measures of fertility timing itself. Admittedly, these are rather small effects and disappear entirely under an assumption of a parallel shift in percentiles of the fertility distribution, as assumed in the original BF model. However, if one accept the formal definition of this concept (I am still undecided), we might conclude from this analysis that, indeed, measures of the timing of demographic events may also be subject to tempo effects. This point is especially pertinent with regard to measures of mortality (which is all about tempo, with a fixed quantum).

At this point in my analysis, I remain skeptical (although not entirely dismissive) of the notion that there are important tempo effects on measures of tempo itself. It is possible to adapt my general fertility model to the case of mortality by focusing on the distribution of events (births or deaths) over the life course. The main difference lies in the relationship between these two distributions and their associated schedules of age-specific rates. In the case of fertility, rates are merely the distribution of (female) ages at birth multiplied by the quantum of total or completed fertility (per woman). For mortality, on the other hand, death rates and the underlying force of mortality equal the probability distribution of deaths divided by an exposure-to-risk based on numbers of survivors.

Thus, whereas the size of the synthetic cohort used in constructing period fertility measures is fixed (the TFR refers to the average fertility of a hypothetical group of women who survive through menopause), the base population to which age-specific death rates are applied when computing a period life expectancy at birth diminishes steadily over the life course. Thus, tempo changes in mortality involve not only a shift in the timing and distribution of events (i.e., deaths) over the life course; they also invoke important secondary effects on implied survivorship. The proper methodological approach and substantive interpretation of this more complicated process are not immediately obvious.

My initial adaptation of the fertility model to mortality produced some results that I still find puzzling. Some of these early results appear to be consistent in form and even magnitude with those of Bongaarts and Feeney. I think this is because both approaches rely on similar but rather questionable assumptions about percentile trends for both periods and cohorts (see discussion below).

Clearly, I need to do much more work to address the issue of tempo effects in mortality adequately. From one perspective, the BF adjusted $\mathring{e}_0$ is merely an incompletely standardized mean age at death. That is, it removes any effects of population age composition due to historic differences in birth cohort size. However, it retains those effects of age structure resulting from past differences in cohort survivorship. Life expectancy at birth, on the other hand, is based on a model of a stationary population in which both of these factors are held constant. It is unclear why we would wish to reject a fully standardized measure of the mean age at death in favor of one that only partially removes the influence of population age composition.

Some Illustrations, Key Insights, and Formal Modeling

To illustrate some of the points made in the preceding overview, in this section I will present assorted pieces of work that I have done so far. Most of it concerns the fertility model, but I will also discuss my attempts to develop a formal mortality model briefly at the end.

Fertility Model

A few simple diagrams help to illustrate some of the key concepts. Figure 1A shows the parallel linear shift (with slope r) of the BF fertility model. However, the model has been modified so that tempo change is episodic rather than continuous. The diagram provides a visual depiction of the lengthening of cohort reproductive intervals that occurs during the shift. Figure 1B presents a model with the same shift at upper ages but no changes at the lower end of the reproductive age range. It also illustrates the lengthening of cohort reproductive lifetimes.

In Figures 2A and 2B, the trend in the span of cohort reproductive intervals is compared to the trend in period total fertility (in the notation used here, a cohort born at time t' attains age α_0 at time t, and thus $t = t' + \alpha_0$). Since the model was specified in terms of a period shift in the fertility age pattern, the trends are discontinuous for the period TFR but gradual for the cohort reproductive interval. It is important to note the exact inverse relationship between these two quantities during the central part of the shift, which includes complete reproductive intervals for some cohorts. In such cases the cohort reproductive interval is longer by a factor of $1/(1-r)$, which matches exactly the reduction in the TFR by a factor of $1-r$.

Figures 3A and 3B depict a pair of more general models. The simple models of Figures 1 and 2 might lead us to believe that the fertility-suppressing effect of tempo delay depends on an overall lengthening of cohort reproductive intervals. In Figure 3A, however, the reproductive age range is held constant at both younger and older ages, and fertility rates during a gradual change of fertility regime are merely weighted averages of a preceding early-fertility pattern and a subsequent late-fertility schedule. Formally, the age pattern of fertility in this model is defined as follows:

$$f_p(x,t) = \begin{cases} f_0(x) & \text{if} \quad t < t_1 \\ \frac{t_4 - t}{t_4 - t_1} f_0(x) + \frac{t - t_1}{t_4 - t_1} f_1(x) & \text{if} \quad t_1 \le t < t_4 \\ f_1(x) & \text{if} \quad t \ge t_4 \end{cases} ,$$

where $TFR_0 = \int f_0(x)dx = \int f_1(x)dx = TFR_1$ (implying a constant quantum of period fertility before and after the regime change), and $\bar{a}_0 = \frac{\int x \cdot f_0(x)dx}{TFR_0} < \frac{\int x \cdot f_1(x)dx}{TFR_1} = \bar{a}_1$ (indicative of a tempo delay).

Although detailed results are not available for the model depicted in Figure 3A, a few moments of reflection should be enough to convince most people that, assuming a constant level of cohort completed fertility, the TFR from time t_0 to t_4 will be consistently lower than its constant level both before and after the shift.[1] This example helps to demonstrate that a lengthening of the total cohort reproductive interval is a sufficient but not necessary condition for producing a tempo effect (assuming constant cohort quantum). In cases where there is a steady upward shift in the percentiles of the cohort distribution of ages at birth, each cohort experiences an increase in the time spent in most percentiles of the age distribution of births (especially within the central age range of relatively high fertility). The direct effect of this lengthening of cohort reproductive lifetimes within fixed percentiles of the age distribution, combined with the assumption of a constant cohort quantum, is that the level of period fertility is reduced during the entire period of the shift (compared to what is observed both before and afterwards).

The general fertility model considered here is depicted in Figure 3B. Its main feature is its flexibility concerning fertility distributions: any two fertility distributions can be used to describe fertility timing before and after the shift, during which cohorts experience linear changes in the percentiles of the age distribution of births. The key result emerging from an analysis of this model is that the total effect of timing changes on period total fertility (assuming no change in cohort quantum) is a reduction in the TFR during the transition period by a factor of

$$1 - \bar{r} = \int \phi_0(y) \cdot \left(1 - r(y)\right) dy \; .$$

This quantity is simly a weighted average across the entire age range of $1 - r(y)$, with weights equal to the baseline probability distribution of ages at birth, $\phi_0(y) = \frac{f_0(y)}{TFR_0}$, and where $r(y)$ is the linear slope of the percentile associated with age y in the baseline fertility schedule.

[1] A simple proof consists of observing that the total volume of period fertility from time t_0 to t_4 equals a similar total for cohorts born from t_0' to t_4', less the two wedges that lie outside the transition period. The assumption of a constant cohort completed fertility rate (CFR) assures that the average CFR for this range of cohorts equals the average TFR for the associated periods. However, it is obvious that the two missing wedges contain relatively high levels of fertility, since each is more concentrated in an age range of relatively high fertility (compared to the other baseline schedule and thus also to the transitional period, which is an average of the two).

This result rests on only two key assumptions: (1) that trends in all percentiles of the series of cohort fertility distributions are linear (over the period of the shift); and (2) that the quantum of a cohort's childbearing between any two percentiles of its fertility age distribution is unaffected by the change of fertility timing. Thus, we are assuming that all births postponed in the present period will still occur eventually.

Formally, we assume that the cumulative distribution of cohort fertility by age during the period of the shift (beginning at time $t = 0$, as depicted in Figure 3A) is related to a baseline distribution as follows:

$$\Phi_c(x,T) = \Phi_0(y)$$

where $y = x - (T + x) \cdot r_0(y)$. Thus, the proportion of its lifetime fertility that a cohort born at time *T* has been completed by age *x*, $\Phi_c(x,T)$, equals the proportion implied by the baseline fertility schedule, $\Phi_0(y)$, at age *y*. We integrate this function to obtain the distribution of cohort fertility:

$$\phi_c(x,T) = \tfrac{d}{dx}\Phi_c(x,T) = \frac{\phi_0(y)\cdot\left[1 - r(y)\right]}{1 + (T + x)\cdot r'(y)} \ ,$$

where $\phi_0(y) = \Phi_0'(y)$. We then multiply by the (fixed) quantum of cohort fertility, which by assumption equals the TFR of the baseline fertility schedule:

$$f_c(x,T) = TFR_0 \cdot \phi_c(x,T) \ .$$

Next, we equate the period fertility rate at age *x* and time *t* with the age-specific rate for the cohort born at time $T = t - x$:

$$f_p(x,t) = f_c(x,t-x) = TFR_0 \cdot \phi_c(x,t-x) = \phi_0(y)\frac{1 - r(y)}{1 + t\cdot r'(y)}$$

where now $y = x - t \cdot r_0(y)$. Integrating this function with respect to age yields the result mentioned earlier:

$$TFR(t) = \int f_p(x,t)\, dx = TFR_0 \cdot (1 - \bar{r}) \ ,$$

where $\bar{r} = \int r(y)\cdot\phi(y)\, dy$.

This model can be used also to investigate the effect of timing changes on measures of fertility tempo, such as the mean age at childbearing. For such purposes, it is convenient to write the formula for the age distribution of period fertility as follows:

$$\phi_p(x,t) = \frac{\phi_0^*(y)}{1 + t\cdot r(y)} \ ,$$

where by definition $\phi_0^*(y) = \phi_0(y)\dfrac{1-r(y)}{1-\bar{r}}$ and is thus a modification of the baseline fertility distribution.

The mean age at childbearing using the baseline fertility schedule is defined as follows:

$$\bar{a}_0 = \int y \cdot \phi_0(y)\, dy \ .$$

Similarly, we obtain the following formula for the period mean age at childbearing implied by this model:

$$\bar{a}_p(t) = \begin{cases} \bar{a}_0 & , \quad t < 0 \\ \bar{a}_0^* + t \cdot \bar{r}^* & , \quad 0 \le t < t^* \\ \bar{a}_0 + t \cdot \bar{r} & , \quad t \ge t^* \end{cases} ,$$

where $\bar{a}_0^*$, and $\bar{r}^*$ are analogous to $\bar{a}_0$ and $\bar{r}$, except that they are computed using $\phi_0^*(y)$ instead of $\phi_0(y)$. In the simple model depicted in Figure 1A, $r(y) = r$ for all y in the reproductive age range, and thus these two distributions are identical. In this case alone, the mean age at birth during the shift lies along a linear trend linking the mean age of childbearing before and after the shift. In all other cases, however, there is a discontinuity in the trend of the mean age of childbearing in this model, and its actual value deviates from the before-after trend. It may be appropriate to interpret this result as a tempo effect on measures of the timing of period fertility. More work is needed on this topic, which is closely related to issues encountered in a discussion about possible tempo effects in mortality.

Mortality Model

It is possible to adapt the above fertility model to the case of mortality, although with some fundamental differences. In both cases, $\phi_0(y)$ describes the distribution of the event of interest (birth or death) over the life course. However, in the case of mortality, the rate – which by definition is the link between the period and cohort age patterns – is based on the population of survivors, which diminishes with age (in contrast, the synthetic cohort of the standard fertility model has a constant size over the reproductive age range). Thus, as with fertility, we define a model in terms of percentiles of cohort age distributions of deaths, which are constant for an initial period, shift linearly over some interval, and then stabilize again at constant levels. This implies a surface of death rates, from which the period age distribution of deaths can be derived.

With such a model, we do indeed recover something that might be interpreted as a tempo effect: discontinuities at both ends of the trends in life expectancy at birth and in age-specific death rates (compared to those of the baseline distribution). The size of the discontinuities is a direct function of the slope of change in percentiles across the age range. Recall that tempo change in fertility had the effect of multiplying birth rates by $1-r(y)$ across the age range. Perhaps not surprisingly, period death rates in the mortality model are affected in an equivalent manner: they equal this same factor multiplied by the death rates of the baseline mortality schedule. Indeed, this result might have been expected, as it follows from the assumption of a linear shift in percentiles of the cohort age distribution of deaths. As with fertility, whenever the time of passage across successive percentiles of the age distribution of an event goes up or down, there is an exact compensatory change in the rate of the event associated with that percentile – and this observation holds true for both kinds of rates.

I am not sure that I understand the BF tempo model of mortality well enough to conclude that it involves a fundamentally similar process. However, the form of the two mathematical results does appear to be somewhat similar, and both present the puzzle of a quite curious discontinuity in period trends. In the case of the model I am describing here, the discontinuity in period trends appears to be due to the utter implausibility of its core assumption, which involves a sudden linear increase in cohort percentiles of mortality. However, trends in cohort percentiles for ages at death are highly dependent on mortality experiences from earlier in life, and it would be quite unusual for a cohort to experience a sudden rise in such a manner (the assumed trend is not even continuous in its first derivative). To achieve this sudden extension of its life course across the full age range, the cohort must drop its mortality rate immediately to a level substantially below its value before the shift. In short, the model I have described is simply a bad model, and the implied discontinuity is meaningless. (I will, of course, attempt to determine if a similar conclusion is warranted for the BF mortality model.)

On the other hand, in the last few hours before finalizing this abstract, I have effectively turned the mortality model around, so that it now involves a linear shift in period percentiles of ages at death. The initial results seem quite plausible, although I have not had time to check them carefully. In the simple case of parallel linear shifts for percentiles of the period age distribution of deaths, life expectancy at birth increases linearly during all years of the shift at a rate of r. If the duration of the shift is long enough to include some complete cohort lifetimes, then the rise in mean lifetime for such cohorts occurs at a rate of $r/(1-r)$.[2] Although period and cohort death rates are equivalent by definition in this model, cumulative hazard rates for cohorts differ from those for the periods in which they live roughly by a factor of $1/(1-r)$. For example, in cases of delayed mortality, the cumulative mortality experience at age x is greater (by a factor of $1/(1-r)$, approximately) for the cohort born in year T, compared to the synthetic cohort of period $t = T + x$ at the same age.

Unlike my first attempt, this model of tempo change in mortality yields results that are entirely plausible. Moreover, they contain no suggestion of anything that might be interpreted as a tempo effect. The relationships between period and cohort measures implied by this model seem (thus far) to present no puzzles. Thus, these new results seem to confirm my previous

[2] These are, in fact, the same slopes, and the two trends would appear as parallel lines if plotted on a Lexis diagram. The difference in their numerical values owes to the fact that the period slope is defined with respect to time of the event, t, whereas the cohort slope is defined in relation to time at birth, $T = t - x$, where x is the cohort's age at time t. For a fixed x, $dT/dt = 1$. However, for the percentile slopes – as well as many measures of timing that can be derived from them – the associated age, x, is rising as a linear function of time, t, with a slope of $r(y)$, where y is the associated age in the baseline mortality schedule. Therefore, $dT/dt = 1 - r(y)$ for a trend at some particular age, and $dT/dt = 1 - \bar{r}$ for a trend in mean values. If $\bar{r}$ is positive, then cohort time runs more slowly than period time. As a result, from the standpoint of cohorts, changes occur more quickly than what is implied by the period trend. Of course, just the opposite occurs when $\bar{r}$ is negative.

belief that there are simply no relevant tempo effects in the case of mortality. This is perhaps explainable by the fact that there is no set of questions analogous to the ones asked in the fertility tempo story. In that case, given an observed rise in the mean age of childbearing, accompanied by a reduced level of period fertility, it is quite reasonable to ask hypothetically whether the lower-than-expected period quantum might be due merely to delayed reproduction, and thus might have no impact on the eventual completed fertility of cohorts.

I can imagine no comparable story for mortality. In fact, I see a very different story associated with these mortality models. Whereas a constant level of cohort completed fertility is a central assumption of the fertility tempo model, the quantum of mortality is fixed automatically in the mortality model. A behavioral model in which cohorts choose to die the same number of times – although sometimes they start dying more slowly in order to live longer – does not seem particularly useful. It seems that the fundamental problem with a tempo model of mortality is that it has causation reversed. It is not tempo delay that causes lower mortality rates, but rather lower mortality rates that cause tempo delay. It should be possible to demonstrate the following points in some formal manner: (1) that any reduction in death rates produces a predictable corresponding increase in average cohort lifetimes (mortality delay), (2) that lower rates and delayed ages at death are merely two ways of depicting the same historical changes, and therefore (3) that tempo changes in no way introduce a bias into our standard measures of mortality and longevity.

Conclusion

The above description offers a taste of the paper I would like to write for this session. Clearly, many puzzles still wait to be solved. To guide my work in this area, I have attempted to identify some key requirements for a unified approach to the analysis of tempo effects in mortality and fertility, and (by extension) in other demographic events as well. Some of the main points are as follows:

- It goes without saying that we need **clear definitions** of concepts, and this will only come through continued **formal development** of this topic. For example, I believe that the fertility model I am proposing helps to clarify a key aspect of the intuition that underlies the BF fertility model. It seems that the central idea is that tempo delay in fertility consists of a postponement of births without their cancellation (i.e., constant quantum). For this reason, the reduction in rates associated with timing delay is inversely related to the increase in time spent by cohorts as they progress between successive percentiles of the age distribution of births. At least in the case of repeatable events, it becomes clear that the rate-suppressing effect of a tempo delay depends on: (1) a key assumption (that the contribution to overall quantum that occurs in each percentile of the distribution is maintained even as the distribution shifts around), and (2) a key fact (rising percentile trends result in a lengthening of cohort life lines within equivalent segments of the distribution, requiring a comparable reduction in rates to achieve the same quantum).

- We need some **common framework** for understanding tempo effects, if any, as they affect fertility, mortality, and other demographic events. One of the key distinctions is between **repeatable** and **non-repeatable** events. However, in some cases the same process can be studied in both ways (e.g., fertility can be modeled using either transition or accretion rates), perhaps offering the opportunity to clarify various issues. A related conceptual distinction concerns tempo effects with regard to **quantum** measures (e.g., the TFR), versus the very different case of **tempo** measures (e.g., life expectancy at birth).

- For various reasons, even our simplest models should be based on **historically relevant** scenarios.[3] For example, the fertility model used by Bongaarts and Feeney to derive their tempo-adjusted TFR consists of a fertility age distribution that shifts upward or downward continuously over time at a constant linear rate (with each percentile of the distribution moving in parallel). They use that model to justify an adjusted fertility measure, invoking arguments about what would happen in the presence or absence of the assumed delay. This simple approach was surely useful for developing some initial insights about the topic. However, I believe that we will clarify many issues if we adhere to models in which tempo change occurs over a finite interval, which can probably be defined in terms of either periods or cohorts for any event, if done properly.[4] This approach offers two advantages: (1) it allows us to think more clearly about historical episodes of tempo change, and (2) by shrinking the interval toward zero, it may allow us to develop a formal definition of localized tempo effects.

- One last observation is that it is important to keep in mind our **interpretive goal** when discussing different approaches to the formalization of tempo effects in the analysis of demographic events. In the case of the TFR, tempo change is relevant when discussing this period measure as a representation of female reproductive lives, since tempo effects clearly alter the relationship between period and cohort indices of fertility quantum. On the other hand, as mentioned earlier, it is important to avoid possible confusion about what a tempo-adjusted fertility measure tells us concerning the implications of current fertility levels for population replacement. In this case, it could be more informative to turn the question around and examine the tempo effect on cohorts of (hypothetically) maintaining a constant replacement level for period fertility.

Finally, we may wish to observe a sort of demographer's Hippocratic oath, and thus to be careful lest tempo effects be misinterpreted by the various consumers of demographic research. Of course, many issues about the nature of tempo effects in relation to demographic measurement remain unresolved within our discipline. Hopefully this can be accomplished without breeding unnecessary confusion in the outside world concerning the meaning of our most basic measures.

---

[3] Realistic scenarios are even better, of course, but relevant ones are more practical (and realistic!) from the modeler's perspective.

[4] Perhaps another valid formulation for the mortality model would be to initiate the shift across the age range as a function of cohorts (i.e., along diagonals of the Lexis diagram).

Figure 1A
Parallel Shift Model (Fertility)

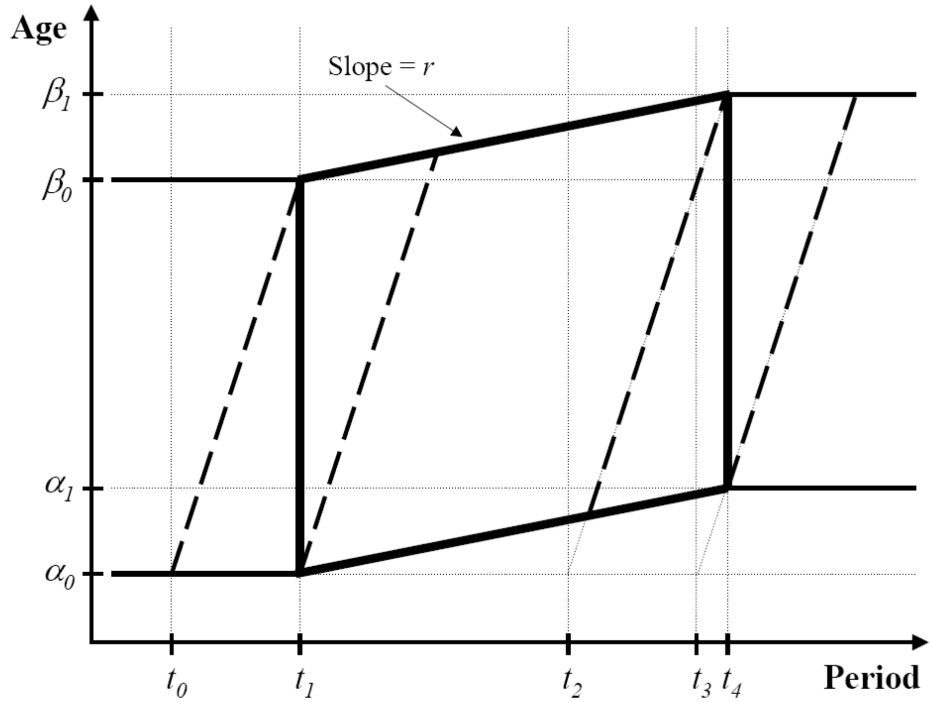

Figure 1B
Proportional Stretch Model (Fertility)

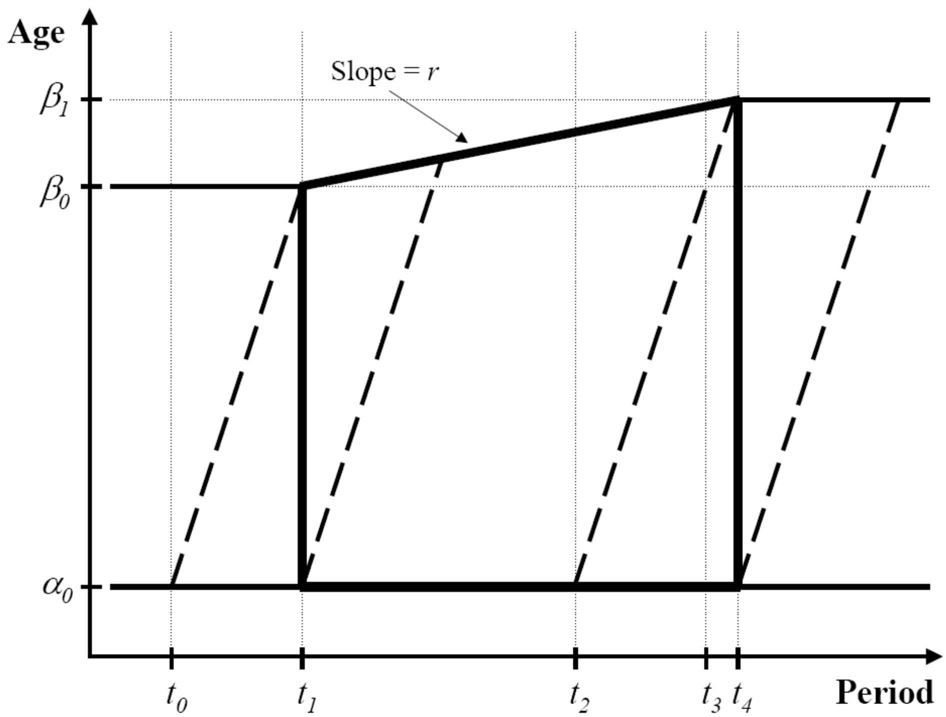

Figure 2A

Reproductive Age Range: Parallel Shift and Proportional Stretch Models

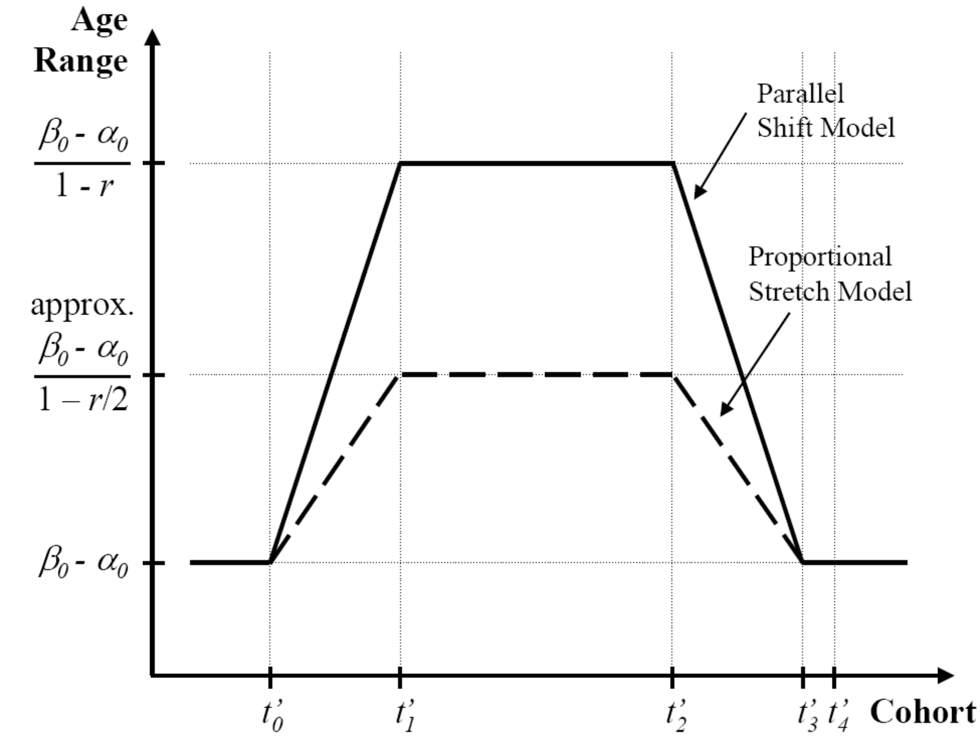

Figure 2B

Total Fertility Rate (TFR): Parallel Shift and Proportional Stretch Models

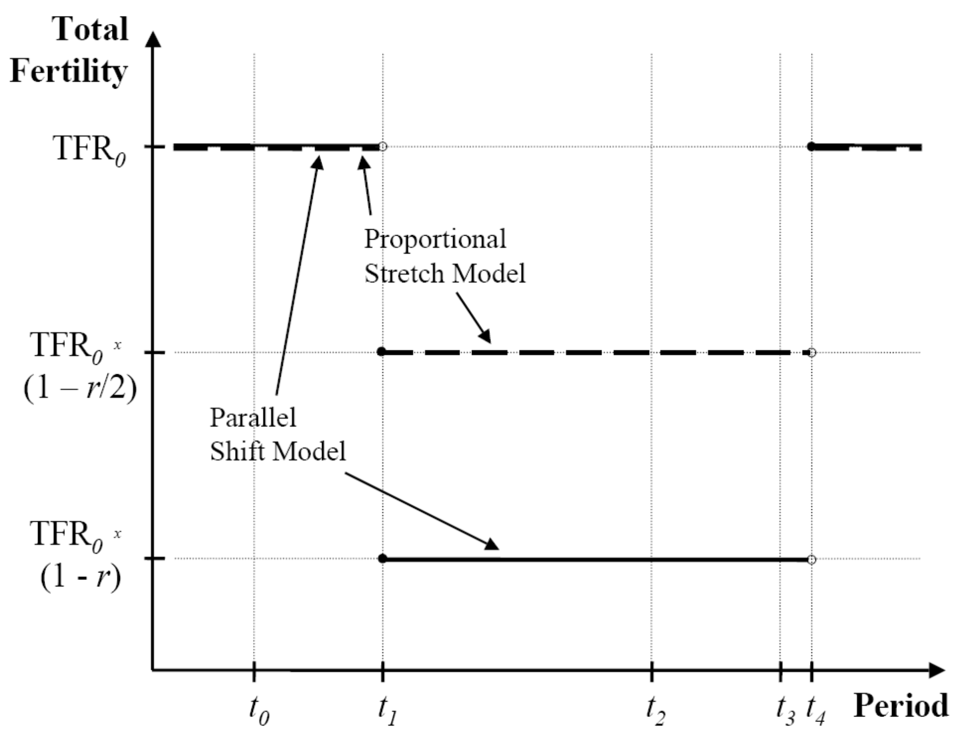

Figure 3A
Regime Change Model (Fertility or Mortality)

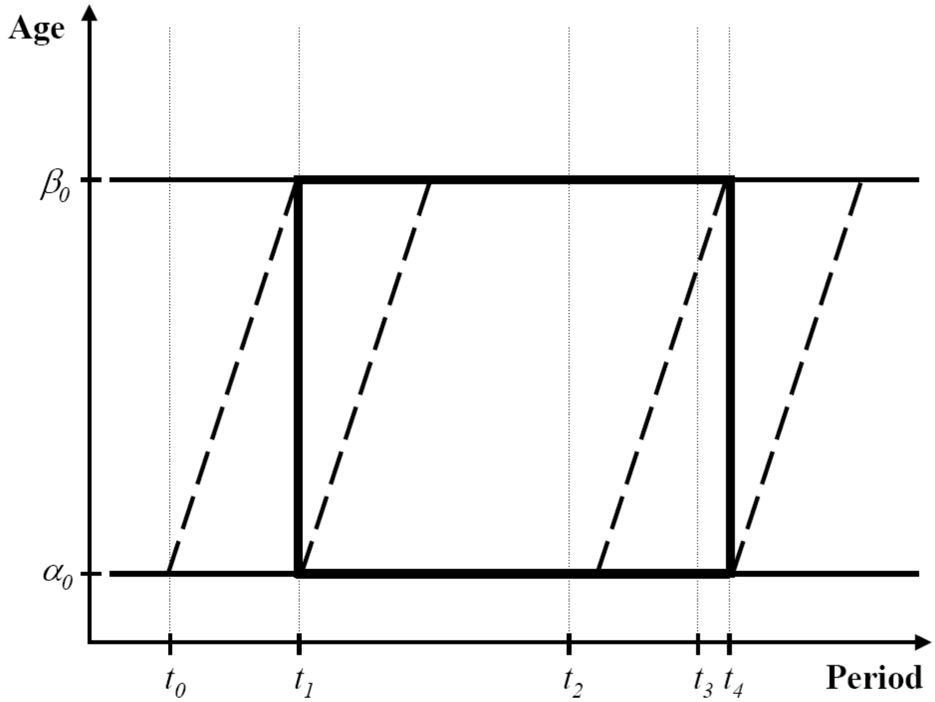

Figure 3B
Percentile Shift Model (Fertility or Mortality)