



Demographic Research a free, expedited, online journal of peer-reviewed research and commentary in the population sciences published by the Max Planck Institute for Demographic Research Konrad-Zuse Str. 1, D-18057 Rostock · GERMANY www.demographic-research.org

DEMOGRAPHIC RESEARCH

**VOLUME 14, ARTICLE 6, PAGES 85-110
PUBLISHED 14 FEBRUARY 2006**

<http://www.demographic-research.org/Volumes/Vol14/6/>
DOI: 10.4054/DemRes.2006.14.6

Research Article

**Demographic translation and tempo effects:
An accelerated failure time perspective**

Germán Rodríguez

© 2006 Max-Planck-Gesellschaft.

Table of Contents

1	Introduction	86
2	Fertility	87
2.1	Translating fertility	88
2.2	Tempo-adjusted fertility	89
2.3	A synthetic cohort interpretation	91
2.4	Cohort and period shifts	94
3	Mortality	95
3.1	Mortality translation	96
3.2	The Bongaarts-Feeney model	97
3.3	Four measures of longevity	99
3.4	Cohort survival	101
3.5	A proportional hazards models	104
4	Discussion	106
5	Acknowledgements	108
	References	109

Demographic translation and tempo effects: An accelerated failure time perspective

Germán Rodríguez¹

Abstract

In this paper I review the concept of tempo effects in demography, focusing on the tempo adjustments proposed by Bongaarts and Feeney and drawing on the work of Ryder and Zeng and Land. I show that the period-shift model that underlies the proposed adjustments can be motivated from an accelerated failure time cohort perspective. I propose alternative measures of tempo under changing fertility and mortality that share a synthetic cohort interpretation with the adjusted measure of quantum. I stress similarities between the results for fertility and mortality, particularly in terms of mean age of childbearing and mean age at death, but also note some important distinctions. I conclude that the fertility adjustments can help distinguish quantum and tempo effects, but argue that in the case of mortality the Bongaarts-Feeney measure of tempo-adjusted life expectancy differs from conventional estimates because it reflects past mortality.

¹Princeton University, E-mail: grodri@Princeton.edu

1. Introduction

How long do we live? According to the U.S. National Center for Health Statistics, “in 2002 the overall expectation of life at birth was 77.3 years”(Arias, 2004). The center makes clear that this measure represents “what would happen to a hypothetical (or synthetic) cohort if it experienced throughout its entire life the mortality conditions of a particular period in time”, in this case 2002. In real life a child born in the U.S. in 2002 would probably live longer than 77.3 years on average, because we expect mortality to improve in the future.

Bongaarts and Feeney (2002, 2003, 2005) have challenged the conventional wisdom, and created quite a stir in the demographic community, by postulating the existence of mortality “tempo effects” that bias standard measures of longevity, such as the period life expectancy, whenever mortality is changing. The measures are believed to be biased upwards when expectation of life is increasing, so we don’t live as long as we think. Bongaarts and Feeney (2003) note that “[e]stimates of the effect for females in three countries with high and rising life expectancy range from 1.6 yr in the U.S. and Sweden to 2.4 yr in France for the period 1980-1995”.

The concept of tempo distortion originated in the field of fertility analysis, where one can draw a clear distinction between quantum and tempo, and refers to the fact that a reduction in period rates could be caused by delays in childbearing without any changes in completed cohort family size. Many demographers have found the extension of these ideas to mortality baffling because a reduction in period mortality rates can only mean that people will die later. With mortality the quantum is fixed, only tempo can change, and no one would mistake one for the other.

It is, of course, possible for cohort and period summaries of age-specific mortality rates to differ. But Bongaarts and Feeney (2003) make the stronger claim that “tempo effects distort both the observed death rates and the corresponding life expectancy”. It is also quite likely that mortality rates are distorted by unobserved heterogeneity, particularly at old ages, but Vaupel (2002) reports that Bongaarts believes that “tempo effects can distort mortality in homogeneous populations”.

Like others I have gone over the underlying mathematical argument and have found no fault. But I come up with a different interpretation of the Bongaarts-Feeney results. I show that working strictly within their framework, one can produce an estimate of expectation of life when mortality is declining that is *higher*, not lower, than the conventional estimate. This differs, of course, from the Bongaarts-Feeney adjustment, and I hope the argument will clarify exactly why this is the case. As Wachter (2005) has noted “every measure measures something”, and we are just measuring different things. Specifically, I will argue that their measure combines the observed force of mortality with features of the age distribution that reflect past rather than current mortality.

Because so much of the work builds upon earlier results on fertility I start with a brief review of Ryder's (1964) famous translation formula. My main goal is to clarify its intent and the conditions under which it is valid. I then review the Bongaarts-Feeney (1998) tempo-adjusted total fertility rate and a synthetic-cohort interpretation due to Zeng and Land (2001, 2002). I show that the period-shift fertility model used by Bongaarts and Feeney can be motivated in terms of a cohort-delay model where the passage of time slows down. I then obtain a measure of mean age of childbearing under changing tempo that complements the Bongaarts-Feeney tempo-adjusted total fertility rate, yet differs from their tempo estimate.

Having laid the groundwork in the field of fertility, where these ideas are less controversial, I move to the field of mortality. I mention briefly why Ryder (1964) didn't pursue a translation formula for mortality, as well as how one might go about it knowing what we know today. I then turn to the Bongaarts-Feeney framework showing how their period-shift mortality model results from a slowing down of time in an accelerated-failure-time framework. I then discuss, and I hope explicate, the various measures of longevity that have been proposed, noting how some of these indices depend on the past via the age structure. I also derive a synthetic cohort measure of life expectancy under changing mortality that provides an exact analog of the measure of fertility tempo derived earlier, yet differs substantially from the Bongaarts-Feeney tempo-adjusted measure of life expectancy.

While most of the paper emphasizes parallels between the analysis of fertility and mortality, in the discussion I return to some of the fundamental differences noted at the outset. In the case of fertility we have recurrent events where a distinction between quantum and tempo is meaningful and, more importantly, adjustments can be useful in determining the extent to which period changes reflect quantum or tempo effects. In the case of mortality trends have an unambiguous interpretation as tempo effects. The fact that the proposed adjusted measures differ from conventional life expectancy is not due to a bias or distortion, but simply to the fact that they measure different things. Specifically, conventional life expectancy depends only on the force of mortality, whereas the adjusted measures are affected by age composition and thus past mortality.

2. Fertility

Let us consider a surface of age-period fertility rates where $f(a, t)$ is the fertility rate at age a and time t . This rate pertains both to period t , and to the cohort born at time $t - a$.

2.1 Translating fertility

Ryder (1964) was interested in the relative strengths and weaknesses of cohort and period summaries of these rates. Useful summaries for the cohort born at time t include the average number of children per woman, $\text{TFR}_c(t)$, a measure of the *quantum* of fertility, and the mean age of childbearing $\mu_c(t)$, a measure of the *tempo* of fertility, defined as

$$\text{TFR}_c(t) = \int f(a, t + a) da \quad \text{and} \quad \mu_c(t) = \int af(a, t + a) da / \text{TFR}_c(t). \quad (1)$$

Together these indices tell us whether women have more or fewer children, and whether they have them earlier or later in life.

The aggregates can also be computed for periods, and are usually interpreted in terms of a synthetic cohort that goes through life bearing children at the current observed rates. The synthetic cohort representing period t has $\text{TFR}_p(t)$ children at an average age of $\mu_p(t)$ where

$$\text{TFR}_p(t) = \int f(a, t) da \quad \text{and} \quad \mu_p(t) = \int af(a, t) da / \text{TFR}_p(t). \quad (2)$$

Ryder's chief concern was that period summaries provide a distorted view of the behavior of cohorts when fertility is changing, and he was able to formalize this view in a remarkable result.

Ryder (1964) assumes that $f(a, t)$ may be expanded in a Taylor series separately for each age. The most useful result is obtained by expanding rates for the cohort which is now at its mean age of childbearing and ignoring terms beyond the first derivative. If the cohort of interest has mean age of childbearing μ , and was thus born at $t - \mu$, we have

$$f(a, t - \mu + a) \approx f(a, t) + (a - \mu)f'(a, t). \quad (3)$$

Under this approximation Ryder obtained the following relationship between cohort and period TFRs:

$$\text{TFR}_c(t - \mu) = \frac{\text{TFR}_p(t)}{1 - r_c}, \quad (4)$$

where r_c is the time derivative or rate of change of *cohort* mean age of childbearing at time $t - \mu$.

This remarkable formula shows that if cohorts postpone childbearing then, to a first order of approximation, the period TFR will fall *below* the cohort TFR (for the cohort at its mean childbearing age) by an amount that depends on how fast the mean age of childbearing is increasing. If mean age of childbearing is decreasing then the period TFR will rise *above* the corresponding cohort TFR. This in fact happened during the baby

boom, when period TFRs rose to levels that exceeded the completed fertility of all active cohorts (Ryder, 1964; Schoen, 2004).

It is important to note that Ryder's result relies solely on a first-order Taylor series approximation to the rates at each age. Contrary to popular belief, there is no assumption that the shape of the period or cohort schedules is constant, or that the cohort and period TFRs are constant. To see this point note that one can generate rates $f(a, t)$ that satisfy the assumption of linearity by interpolating between any two arbitrary age schedules $f(a, 0)$ and $f(a, \tau)$.

Ryder (1964) also considered a translation procedure for mean age of childbearing, introducing a second type of formula with stronger assumptions (which may account for some of the confusion). We will not pursue this development further because it is not central to the argument that follows, except to note Ryder's conclusion that "the period mean is a distorted version of the cohort mean" when quantum is changing, "just as the period sum is a distorted version of the cohort sum" when tempo is changing.

2.2 Tempo-adjusted fertility

Bongaarts and Feeney (1998) proposed a tempo-adjusted total fertility rate, usually denoted TFR^* , based on an expression that looks remarkably like Ryder's translation formula:

$$\text{TFR}^*(t) = \frac{\text{TFR}_p(t)}{1 - r_p(t)}. \quad (5)$$

There are, however, two subtle but important differences. First, the $r_p(t)$ on the right-hand-side is the rate of change in the *period*, not the cohort, mean age of childbearing at time t . This is much easier to calculate from available data. Second, TFR^* is not a cohort rate, but rather a pure-period measure representing tempo-corrected fertility, as we will see presently.

A third difference I should mention is that Bongaarts and Feeney recommend applying their procedure separately by birth order, using rates that divide births of a given order by all women. I ignore this breakdown to keep the argument simple. (I also believe that parity-specific fertility is best analyzed using hazard rates where births of order k are divided by women at parity $k - 1$, but that's an argument best left for another time; see van Imhoff and Keilman (2000) and the rejoinder by Bongaarts and Feeney (2000).)

We will derive the adjustment in Equation 5 considering a situation where all cohorts start delaying fertility at the same time and rate without reducing their completed family size. The situation where quantum is fixed is simpler—and more relevant to the analysis of mortality—than where quantum is changing as well, although the Bongaarts-Feeney adjustment can be applied in both cases. The assumption of a constant rate also simplifies

things, in particular it leads to explicit cohort results, although Equation 5 can also be applied when the rate of change varies over time.

It will be useful to introduce a function $F(a, t)$ representing the cumulative fertility or average parity of women age a at time t (the cohort born at time $t - a$). This schedule can be obtained as a *cohort* integral, by accumulating fertility along a diagonal of the Lexis diagram:

$$F(a, t) = \int_0^a f(x, t - a + x) dx. \quad (6)$$

The age-period specific rates $f(a, t)$ are the *cohort* derivatives of these rates, and can be recovered by differentiating $F(a, c + a)$ with respect to a , i.e. with respect to both age and time.

Let us also introduce a fertility schedule $f_0(a)$ with corresponding cumulative schedule $F_0(a)$, total fertility rate $\text{TFR}_0 = \int f_0(a) da$ and mean age of childbearing $\mu_0 = \int a f_0(a) da / \text{TFR}_0$. This baseline schedule will represent the situation at time zero, so that $F(a, 0) = F_0(a)$. If fertility has been constant for a long time we could view all rates prior to time zero as generated by the baseline schedule, but this assumption is not necessary for the developments that follow. All we need is the assumption that just before time zero women were following the cumulative schedule $F_0(a)$.

Now suppose that at time zero all cohorts *slow down* their pace of childbearing at the same rate r . Let us give this statement a precise meaning. The cohort that has reached average parity $F_0(a)$ at age a and time zero, and would have been expected to reach parity $F_0(a + 1)$ a year later, will instead climb only as far as $F_0(a + 1 - r)$. This is similar to taking a pill that prevents all births (and stops a woman's biological clock) for a fraction r of the year, but I prefer to work in continuous time. The same idea is used in Coale's (1971) classic nuptiality model, where he speeds up or slows down the Swedish schedule of first marriages. The device of accelerating or slowing down the passage of time is also used in survival analysis, as we will see in Section 3.

It turns out that this slowing down of time is exactly equivalent to a period shift in the cumulative fertility schedule, so that

$$F(a, t) = F_0(a - rt), \quad t \geq 0. \quad (7)$$

For example the cohort age a at time zero had parity $F(a, 0) = F_0(a)$ and will now move to $F(a + 1, 1) = F_0(a + 1 - r)$.

If we now take *cohort* derivatives, differentiating with respect to both age and time (which of course vary together for a cohort) we obtain

$$f(a, t) = f_0(a - rt)(1 - r), \quad t \geq 0. \quad (8)$$

This shows that when all cohorts slow down the pace of childbearing at the same rate r the age-specific rates are instantly deflated by a factor $1 - r$ and start shifting to older ages.

The simplest way to prove Equation 8 is to write the period-shift model for a cohort that reaches age a at time $t = c + a > 0$, which is

$$F(a, c + a) = F_0(a - r(c + a)) = F_0(a(1 - r) - rc), \quad (9)$$

and then take derivatives with respect to a for fixed c to obtain

$$f(a, c + a) = f_0(a(1 - r) - rc)(1 - r) = f_0(a - r(c + a))(1 - r). \quad (10)$$

Integrating the period schedule in Equation 8 over a for fixed t we obtain the period TFR, and we can also obtain the period mean age of childbearing. As long as the cumulative schedule continues to shift at a rate r ,

$$\text{TFR}_p(t) = \text{TFR}_0(1 - r) \quad \text{and} \quad \mu_p(t) = \mu_0 + rt. \quad (11)$$

The period TFR declines at time zero by a factor $1 - r$ as a result of the delay. This could be misinterpreted as a change in the quantum of fertility when in fact it is a pure tempo effect. The fact that the derivative of period mean age of childbearing is r provides an ingenious way to recover the baseline TFR simply dividing by $1 - r$, which leads to the Bongaarts-Feeney formula 5. The key assumption required is that all cohorts delay fertility at the same time and rate.

This leads to a direct interpretation of the tempo-adjusted TFR as a *counterfactual* measure; paraphrasing Bongaarts and Feeney (1998), it provides an estimate of what the period TFR would have been if cohorts had not delayed childbearing at time t . Note that this is indeed a pure period measure as claimed; it estimates TFR_0 , which does not correspond to the completed family size of any real cohort unless fertility has been constant for the last thirty five years or so. It can, however, be interpreted as the completed family size of a synthetic cohort, as we will see below.

It is interesting to note that Bongaarts and Feeney adjust the quantum but not the tempo of fertility, considering the mean age of childbearing unaffected by tempo distortions. This can be seen to be the case in the present framework because $\mu_p(0) = \mu_0$, a result that obtains because the factor $1 - r$ appears both in the numerator and the denominator of the mean. Delays affect the mean age of childbearing only after time zero. This point will be quite important when we turn to an analysis of mortality.

2.3 A synthetic cohort interpretation

In the previous section we focused on period measures. Let us now consider what happens to the cohort that starts childbearing at time zero, when the passage of time slows down.

Let a_0 denote the lowest age of childbearing, so the cohort in question was born at time $-a_0$. From Equation 8, we see that this cohort would follow the schedule

$$f^\dagger(a) = f_0(a - r(a - a_0))(1 - r) = f_0(a(1 - r) + ra_0)(1 - r). \quad (12)$$

Integrating this expression over all ages a we find the total fertility rate for this cohort to be

$$\text{TFR}^\dagger = \int f_0(a(1 - r) + ra_0)(1 - r)da = \text{TFR}_0, \quad (13)$$

where the results follows by changing variables from a to $y = a(1 - r) + ra_0$ and noting that the Jacobian $da/dy = 1/(1 - r)$ cancels out the multiplier $1 - r$. This result is due to Zeng and Land (2001), who provide a simplified derivation of the Bongaarts-Feeney adjustment.

Because $\text{TFR}^\dagger = \text{TFR}^*$, the Zeng-Land approach leads to an interesting interpretation of the Bongaarts-Feeney measure in synthetic cohort terms, as the number of children that a cohort would have under current *conditions*, if by that we mean the current rates and the fact that they are shifting to older ages at a constant rate r .


The corresponding mean age of childbearing for this cohort can easily be obtained using the same change of variables technique, but appears to have been overlooked in the literature:

$$\mu^\dagger = \int af_0(a(1 - r) + ra_0)(1 - r)da / \text{TFR}_0 = \frac{\mu_0 - ra_0}{1 - r}. \quad (14)$$

The notation could be streamlined considerably if we measured age from a_0 as done by Zeng and Land (2001), in which case Equation 14 would simplify to $\mu^\dagger = \mu_0/(1 - r)$ and we would have the remarkable result that under a period shift the quantum and tempo of fertility are affected exactly the same way.

Bongaarts and Feeney (1998) argue that TFR^* removes a tempo distortion from TFR , and one could make the point that μ^\dagger removes a tempo distortion from μ . I prefer the more neutral view that the two sets of indices measure different things: TFR (and μ) tell us how many children a synthetic cohort would have (and when) if it followed a *fixed* period fertility schedule with constant shape, quantum and tempo. In contrast, TFR^* (and μ^\dagger) tell us how many children the synthetic cohort would have (and when) if it followed a *shifting* period schedule with constant shape and quantum but changing tempo.

Figure 1 illustrates these ideas with a Coale-Trussell (1974) fertility schedule where 90% of women marry, age at marriage has mean 23 and standard deviation 4, the level of natural fertility (M) is 1 and the control parameter (m) is -1 . Under this schedule the TFR is 4 children per woman and the mean age of childbearing is 29.2. Suppose, however, that women start delaying fertility at the rate of $r = 0.2$ years per year. As


Figure 1: Period and Cohort Rates when Childbearing is Delayed

shown in Equation 8, the period age-specific fertility rates would be instantly reduced by 20%, a necessary consequence of the fact that women have slowed down childbearing. The curve labelled “period” shows the deflated schedule, which has a TFR of 3.2 children per woman but the same mean age of childbearing as the original. The curve labelled “cohort” shows the schedule followed by the cohort just starting its reproductive career, assuming the shift continues indefinitely at the same rate. This cohort would have 4.0 children per woman, on average at age 33.5 given by Equation 14.

Figure 2 shows how a shift in a period schedule leads to a stretched cohort schedule. Here we plot the cumulative schedule $F_0(a)$ in the example at 10 year intervals. We also show in gray the parity schedule for the cohort starting reproductive life when the shift starts, and we mark the points where it “borrows” its cumulative fertility from the three central curves. Note that all schedules lead to a completed family size of four, but the cohort takes longer to climb that far.

To summarize, we have illustrated how a reduction in period fertility from 4.0 to 3.2 can result from delayed childbearing without changes in quantum. Noting that mean age of childbearing increases 0.2 years per year we obtain a TFR* of 4.0. We can interpret this number as a counterfactual estimate of what the period TFR would have been if women had not delayed childbearing, in which case the mean age of childbearing would still be 29.2. We can also interpret it as the number of children that a synthetic cohort would have if the delay continued indefinitely, in which case mean age of childbearing would be 33.5. The last estimate pairs TFR* with μ^\dagger , the estimate of mean age of childbearing under changing tempo proposed here.

Figure 2: How a Period Shift in a Parity Schedule Translates into a Cohort Delay



2.4 Cohort and period shifts

The foregoing results generalize to multiple cohorts if we assume that the cumulative period schedule $F(a, t)$ continues to shift according to Equation 7. For later cohorts this means not only that once childbearing starts it proceeds at a slower pace than before, but also that the start of childbearing itself is delayed. This implication of period-shift models will be of some significance when we turn to mortality, and represents a departure from accelerated failure time models.

Following exactly the same change of variables technique we used for the Zeng-Land cohort, we can show that the cohort born at time t for $t \geq -a_0$ has


$$\text{TFR}_c(t) = \text{TFR}_0 \quad \text{and} \quad \mu_c(t) = \mu^\dagger + r_c(t + a_0) \quad (15)$$

where r_c is the rate of change of cohort mean age of childbearing, and is related to the period derivative by

$$r_c = \frac{r}{1 - r}. \quad (16)$$

Equation 16 is due to Zeng and Land (2002), who noted that period changes in tempo provide a distorted view of cohort changes in tempo. (They use the notation r^* for r_c .) Note that the cohort considered earlier was born at $t = -a_0$, and that evaluating these expressions at that value leads to TFR^\dagger and μ^\dagger .

An interesting implication of these results is that a shift in period fertility schedules generates a parallel shift in cohort fertility schedules, with both moving up the age axis

Figure 3: Shifting Period and Cohort Fertility Schedules

but at slightly different rates r and r_c . Figure 3 illustrates this idea using model Coale-Trussell schedules. The left panel shows a period schedule that is shifting to older ages at the rate of $r = 0.2$ years per year, and the right panel shows the corresponding cohort schedules shifting at the rate of $r_c = 0.25$ years per cohort.

Thus, under a simple linear shift model cohort and period quantum are constant and differ by a factor $1 - r$ at time zero and later. Cohort and period tempo change over time. The period mean age of childbearing increases at the rate of r years per year starting from μ_0 at time zero. Cohort mean age of childbearing varies between μ_0 and μ^\dagger for the active cohorts at time zero, and increases at the rate of r_c years per cohort for cohorts that start their reproductive careers after that. These results provide a way to translate cohort and period quantum and tempo, but the assumptions required are stronger than for a simple counterfactual interpretation of TFR*.

3. Mortality

Let us now turn our attention to mortality, focusing on a surface of age-period specific rates $\mu(a, t)$ representing the force of mortality at age a and time t for the cohort born at $t - a$. The rates along a diagonal can be used to compute a *cohort* life table, but the data required are often not available and the calculation can only be completed after the cohort has died.

More often the mortality rates for fixed t are used to compute a *period* life table, which may be interpreted in terms of a synthetic cohort that goes through life subject to the force of mortality prevailing at time t . Bongaarts and Feeney's concern is that period measures, including the period expectation of life and the rates themselves, may be distorted by a tempo effect.

3.1 Mortality translation

Ryder (1964) noted that “the development of translation procedures has proven more difficult for mortality functions than for fertility functions” because of the multiplicative relationships involved in an attrition process, although he made some headway working with the logarithms of the rates. Keilman (1994) later obtained useful translation formulas for the hazards of non-repeatable events, but these do not lead to simple summary results such as Equation 4.

Further progress can be made working with a *survival* surface where $S(a, t)$ represents the probability that someone born at time $t - a$ will survive to age a at time t ,

$$S(a, t) = \exp\left\{-\int_0^a \mu(x, t - a + x)dx\right\}. \quad (17)$$

A nice feature of this surface is that integrating along a diagonal leads to cohort life expectancy:

$$e_0^{(c)}(t) = \int_0^\infty S(a, t + a)da. \quad (18)$$

Unfortunately, integrating over a for fixed t does *not* lead to period life expectancy unless mortality is constant. It does, however, lead to a meaningful alternative period measure of longevity, the cross-sectional average length of life (CAL) described by Guillot (2003):

$$\text{CAL}(t) = \int_0^\infty S(a, t)da. \quad (19)$$

The survival probabilities $S(a, t)$ for fixed t may be interpreted as the age distribution of a population that has a constant stream of births and is subject to the mortality risks $\mu(a, t)$. Bongaarts and Feeney (2003) call this the *standardized* age distribution. CAL is a function of this age distribution and thus depends on past mortality, a point to which we will return later.

In addition to life expectancy and CAL it will be useful to define $\alpha = \int aS(a)da / \int S(a)da$, the mean age in the stationary population implied by a survival schedule $S(a)$. A straightforward application of Ryder's (1964) translation formula, which would expand

the survival probabilities for the cohort now at its mean stationary age around the current age distribution using a first-order Taylor series, yields

$$e_0^{(c)}(t - \alpha) = \frac{\text{CAL}(t)}{1 - r_c}, \quad (20)$$

where r_c is the rate of change in the cohort mean stationary age. This shows that, to a first order of approximation, CAL falls below cohort life expectancy when mortality is declining, to an extent determined by the speed of the decline, provided we line up cohorts and periods using mean stationary age.

Guillot (2006) applies Ryder's ideas using a somewhat different approach, but reaches essentially the same conclusions. He divides $\text{CAL}(t)$ by an index of distributional distortion to obtain an adjusted measure, which can be interpreted as a weighted average of the life expectancies of all cohorts alive at t . He then notes in an application to France that the result is close to the life expectancy of the cohort born at time $t - A(t)$, where $A(t)$ is the mean age of the stationary population at time t , between 30 and 37 years for France in the twentieth century. Here we divide by $1 - r_c$ instead of the distortion index, and use cohort rather than period mean age. But we both conclude that when mortality declines CAL falls below the life expectancy of the cohort near its mean stationary age. (I later show under different assumptions that CAL equals the life expectancy of the cohort now at its mean age at death.)

One could take this result to mean that CAL provides a distorted view of cohort life expectancy, or is subject to a tempo effect when mortality is declining, in much the same way that the period TFR distorts cohort fertility. I prefer to view it as indicating that when mortality is declining the age structure lags behind the cohort mortality schedule. In other words, it takes a while for a population to forget its past.

I realize that applying a formula developed for the quantum of fertility to the tempo of mortality seems unusual, if not plain wrong, but Ryder's result is quite general. Given any age-period surface, it relates a cohort integral to a period integral and to the rate of change of the first cohort moment. In fertility we applied it to age-specific rates, so the integrals are measures of quantum and the first moment is tempo. In mortality we applied it to survival probabilities (or age distributions), so the integrals are mean survivals and the first moment is mean stationary age.

3.2 The Bongaarts-Feeney model

The Bongaarts-Feeney model of mortality change is formally identical to the fertility model, except that the period schedule that shifts over time is the standardized age distribution $S(a, t)$ rather than the parity schedule $F(a, t)$. In this section we motivate the

model in terms of a slowing down of the passage of time, just as we did for fertility. Later we discuss various period and cohort measures of longevity under the model.

Let $S_0(a)$ denote a survival function and let $d_0(a)$ and $\mu_0(a)$ denote the corresponding density and hazard functions. This could be a conventional period life table or a mathematical model. We will assume that at time zero survival is governed by $S_0(a)$ in the sense that all cohorts are following this schedule. This is equivalent to assuming that the population is stationary with age distribution $S_0(a)$.

Suppose, however, that at time zero all cohorts postpone death at the same rate r . Consider specifically the cohort that has reached age a at time zero, of which a fraction $S_0(a)$ is still alive. We would expect a fraction $S_0(a + 1)$ to be alive a year later at age $a + 1$, but instead we observe that the proportion surviving has increased to $S_0(a + 1 - r)$. It is precisely as if the cohort had aged only $1 - r$ years in one year. This type of model is known in the statistical literature as an accelerated life model, see for example Kalbfleisch and Prentice (2002). The situation is similar to taking a pill that prevents death (and stops aging) for a fraction r of the year, but I prefer to view the process as developing in continuous time.

Remarkably, this model is equivalent for all active cohorts to a period shift in the standardized age distribution, where

$$S(a, t) = \begin{cases} 1 & \text{if } a < rt \\ S_0(a - rt) & \text{if } a \geq rt \end{cases} \quad (21)$$

For example the survival probabilities for the cohort considered in the previous paragraph are $S(a, 0) = S_0(a)$ and $S(a + 1, 1) = S_0(a + 1 - r)$. If we compute a cohort derivative, differentiating Equation 21 with respect to both age and time, and changing sign, we obtain a density reflecting the age distribution of deaths at each time

$$d(a, t) = \begin{cases} 0 & \text{if } a < rt \\ d_0(a - rt)(1 - r) & \text{if } a \geq rt. \end{cases} \quad (22)$$

Note that $d(a, t)$ is a probability density function only for a cohort, i.e. if we consider $d(a, c + a)$ for fixed c . The period profile is not a real density but a collection of densities for various cohorts, and in this model it integrates to $1 - r$, not one. Bongaarts and Feeney (2003) call the integral of $d(a, t)$ for fixed t the total mortality rate (TMR). Watcher (2005) notes that it can be interpreted as a period count of deaths.

If we divide the deaths $d(a, t)$ by the numbers exposed $S(a, t)$ we obtain the age-period specific force of mortality

$$\mu(a, t) = \begin{cases} 0 & \text{if } a < rt \\ \mu_0(a - rt)(1 - r) & \text{if } a \geq rt. \end{cases} \quad (23)$$

This is both a period and a cohort hazard, pertaining to time t and to the cohort born at $t-a$. Note that when all cohorts start delaying death at the same rate the hazard is instantly deflated by a factor $1-r$ and starts shifting to older ages. This is clearly a tempo effect, as it is caused by a delay in death. I don't believe, however, that it is a distortion. The only way that cohorts can delay death is by dying at lower rates, so I view the reduction in hazards as real. The interesting question concerns the implications of this change for longevity.

It will be useful to introduce for completeness two additional functions defined by Bongaarts and Feeney (2003) in (their) Equations 5a and 5b. If we differentiate $S(a, t)$ with respect to time only (as opposed to time and age simultaneously) we obtain the death density

$$d_s(a, t) = d_0(a - rt), \quad (24)$$

and dividing this by the survivors $S(a, t)$ we obtain the hazard

$$\mu_s(a, t) = \mu_0(a - rt). \quad (25)$$

These are proper density and hazard functions for $a \geq rt$ and can best be viewed as inherent features of the standardized age distribution $S(a, t)$, so I will call then the age-distribution density and hazard, respectively. Note that under the period shift model the observed force of mortality $\mu(a, t)$ is proportional to the age distribution hazard $\mu_s(a, t)$, with proportionality factor $1-r$. This is called the *proportionality* assumption in the Bongaarts-Feeney framework.

I should also note that Bongaarts and Feeney consider a more general shift model where the rate of delay is not a constant r but a function of time $r(t)$. I stick to the linear case because it is simpler and leads to explicit results for cohorts.

3.3 Four measures of longevity

Bongaarts and Feeney (2003) consider four measures of longevity, denoted M_1 to M_4 . Three of them are equal under the period-shift model of the previous section. The odd one out is period life expectancy.

The first measure is cohort average length of life (CAL)

$$M_1(t) = \text{CAL}(t) = \int_0^\infty S(a, t) da. \quad (26)$$

This measure is easily computed by integrating the standardized age distribution. From Equation 21 we find that under the period shift model

$$\text{CAL}(t) = \text{CAL}(0) + rt, \quad (27)$$

where $CAL(0)$ is both CAL and the conventional expectation of life in the baseline schedule $S_0(a)$. CAL may be computed as an ordinary mean age at death where deaths are obtained by applying the age-distribution hazard $\mu_s(a, t)$ to the standardized age distribution $S(a, t)$. Interestingly, CAL doesn't change when cohorts start postponing death, but it starts increasing at the rate of r years per year as long as the shift (or slow down of time) continues. This occurs because CAL is based solely on the age structure at time t , and does not respond to changes in mortality until these are reflected in the age structure.

The second measure is standardized mean age at death

$$M_2(t) = \int_0^{\infty} ad(a, t)da / \int_0^{\infty} d(a, t)da, \quad (28)$$

which is based on the standardized age distribution of deaths at time t . The deaths in this index result from applying the current force of mortality $\mu(a, t)$ to the standardized age distribution $S(a, t)$, and may thus be viewed as a measure that depends both on current mortality risks and the current age distribution.

Under the period-shift model the force of mortality $\mu(a, t)$ and the age-distribution hazard $\mu_s(a, t)$ are proportional, with proportionality factor $1 - r$. Because this factor appears both in the numerator and denominator of the mean it cancels out, so $M_2(t) = M_1(t)$ as noted by Bongaarts and Feeney (2003). If the proportionality assumption is not satisfied, however, the two indices will differ.

The third measure is conventional period life expectancy

$$M_3(t) = e_0^{(p)}(t) = \int_0^{\infty} \exp\left\{-\int_0^a \mu(x, t)dx\right\}da. \quad (29)$$

This index may also be viewed as an ordinary mean age at death where deaths result from applying the force of mortality $\mu(a, t)$ to the stationary population implied by that hazard, which is of course the period survival function $\exp\left\{-\int_0^a \mu(x, t)dx\right\}$ (not to be confused with $S(a, t)$). This measure depends on the current force of mortality only.

Under the period shift model the force of mortality $\mu(a, t)$ is proportional to $\mu_s(a, t)$ and therefore the period survival function is a power of the standardized age structure, but there is no simple relationship between $M_3(t)$ and either $M_1(t)$ or $M_2(t)$.

Note that when cohorts start postponing death the conventional expectation of life reacts instantly. Because it depends only on the force of mortality $\mu(a, t)$, which has been deflated by a factor $1 - r$, conventional life expectancy e_0 will increase. This is again a tempo effect, but in my view is not a distortion. Conventional life expectancy is just a summary of age-period specific mortality, and responds appropriately by increasing when the rates decline. In particular, the synthetic cohort interpretation of e_0 as the mean lifetime implied by the current rates continues to be correct.

The fourth measure is the Bongaarts-Feeney tempo-adjusted life expectancy. This index seeks to remove the tempo effect from the force of mortality dividing by $1 - r$ and is therefore defined as

$$M_4(t) = \int_0^{\infty} \exp\left\{-\int_0^a \mu(x, t)/(1 - r) dx\right\} da. \quad (30)$$

Under the period-shift model $\mu(a, t)$ is proportional to $\mu_s(a, t)$ with proportionality factor $(1 - r)$ and therefore $M_4(t) = M_1(t) = M_2(t)$, as noted by Bongaarts and Feeney (2003). In this case the adjusted measure can be viewed as an ingenious way to estimate CAL or mean age at death from the observed hazard. If the model does not hold, however, $M_4(t)$ is a different measure that ostensibly depends only on the current force of mortality and the rate of delay r , but in practice requires knowledge of the standardized age distribution for estimation. Watcher (2005) provides a characterization of $M_4(t)$ that clarifies this issue.

To summarize, when cohorts start delaying death conventional life expectancy reacts instantly, whereas the other three measures react more slowly, increasing only as the changes work their way into the age structure. The fundamental issue is whether this is a bias or distortion in conventional life expectancy. I argue that it is just a reflection of the fact that when mortality declines the age structure lags behind the force of mortality. To further explore this issue we now look at the cohort implications of the period-shift model.


3.4 Cohort survival

Consider again the cohort born at the time the period shift, or the slowing down of the passage of time, starts. This cohort would have been expected to follow the schedule $S_0(a)$ but instead will follow a stretched schedule, where the probability of surviving to age a is

$$S^\dagger(a) = S_0(a(1 - r)). \quad (31)$$

This result follows directly from the period-shift model in Equation 21 and shows that each calendar year the cohort ages only $1 - r$ years.

Figure 4 illustrates how a period shift leads to a cohort delay using a Weibull distribution that is shifting towards higher ages at a rate of 0.2 years per year, an artificially high rate chosen to make the illustration clear. I show the schedule at the start of the process as well as 25, 50, 75 and 100 years later, and superimpose the survival probabilities that would apply to a synthetic cohort undergoing this regime, highlighting the ages where the cohort survival “borrows” its probability from the three central curves. The analogy to Figure 2 for fertility should be obvious.

Figure 4: How a Period Shift in Survival Translates into a Cohort Delay

We can compute the expectation of life under $S^\dagger(a)$ using the same change of variables technique that we used in the case of fertility:

$$e_0^\dagger = \int_0^\infty S_0(a(1-r))da = \int_0^\infty S_0(y) \frac{dy}{1-r} = \frac{e_0}{1-r}. \quad (32)$$

We find that if $r > 0$ the expectation of life under a shifting schedule exceeds the value it would have if the schedule remained fixed. The area under the original curve is e_0 , the shaded area under the stretched curve is e_0^\dagger .

Note by way of illustration that life expectancy in the U.S. today is 77.3 under a fixed mortality schedule, but would be 85.8 if the schedule shifted 0.1 years per year, which is the observed gain in period life expectancy between 2001 and 2002. The value 85.8 is computed simply as $77.3/0.9$.

Let us return to $S^\dagger(a)$, the survival function that applies to our synthetic cohort. Differentiating we find the density to be

$$d^\dagger(a) = \frac{d}{da} S^\dagger(a) = \frac{d}{da} S_0(a(1-r)) = d_0(a(1-r))(1-r). \quad (33)$$

The hazard, computed as the ratio of deaths to survival, is

$$\mu^\dagger(a) = d^\dagger(a)/S^\dagger(a) = \mu_0(a(1-r))(1-r). \quad (34)$$

Thus, if the mortality schedule shifts 0.1 years per year, a 60 year old would be exposed to 90% of the risk that would have applied at age 54 under a static schedule. These results

are consistent with Equations 22 and 23 in the previous section, and thus with equations 8b and 8c in Bongaarts and Feeney (2003). (We showed before that their $\frac{d}{dt}M_2(t) = r$.)

We note again that as soon as time slows down the hazard is deflated by a factor $1 - r$, which is how the cohort manages to live longer. Consider an example where the baseline survival $S_0(a)$ is Weibull with parameters p and λ , so $S_0(a) = \exp\{-(\lambda a)^p\}$. In this case the stretched survival $S^\dagger(a)$ is also Weibull with parameters p and $\lambda^\dagger = \lambda(1 - r)$, so the shift and consequent slowing down of the passage of time translate into a proportionate reduction in the hazard at all ages. Kalbfleisch and Prentice (2002) show that the Weibull is the only distribution where the accelerated life and proportional hazards families coincide.

For an example more relevant to human mortality, at least in adult ages, consider a Gompertz model with parameters α and β , where the baseline hazard $\mu_0(a) = \exp\{\alpha + \beta a\}$ increases exponentially with age. In this case the stretched survival is also Gompertz but with parameters $\alpha^\dagger = \alpha + \log(1 - r)$ and $\beta^\dagger = \beta(1 - r)$, a result that follows directly from the general expression given above. In this case the change in the hazard is not proportional, but relatively larger at older ages. For a country such as the U.S., where adult mortality is roughly Gompertz, a shift of 0.1 years per year starting at age 30 would reduce the hazard by 10% at age 30, 30% at age 60 and 46% at age 90. As a result a 30 year old, who is expected to live another 48.4 years under current conditions, would live on average about 53.8. (These calculations are based on $\alpha = -9.696$ and $\beta = 0.0855$, which implies $\alpha^\dagger = -9.545$ and $\beta^\dagger = 0.07694$. Note that for a shift starting at age a_0 rather than zero $\alpha^\dagger = \alpha + \log(1 - r) + \beta r a_0$. The value of $e_0^\dagger = 53.8$ can be obtained as $48.4/0.9$ or by numerical integration of the Gompertz hazard.)

These results can be extended to multiple cohorts, just as we did in the case of fertility, by assuming that the standardized age distribution continues to shift at a constant rate. Using essentially the same argument as in the previous section, we can show that the cohort born at time $t > 0$ goes through the survival schedule

$$S(a, t + a) = \begin{cases} 1 & \text{if } a < tr/(1 - r) \\ S_0(a - r(t + a)) & \text{otherwise} \end{cases} \quad (35)$$

and thus has life expectancy

$$e_0^{(c)}(t) = e_0^\dagger + r_c t, \quad (36)$$

where e_0^\dagger is the life expectancy of the cohort born at time zero and r_c , the rate of change in cohort life expectancy, is

$$r_c = \frac{r}{1 - r}. \quad (37)$$

The cohort born at time zero experiences just a stretching of the survival function $S_0(a)$, which yields a plausible model for all ages. Subsequent cohorts, however, are assumed to

experience no mortality until they reach age $r_c t$, at which time they join a stretched and shifted schedule. This feature makes the model less realistic in multiple-cohort settings unless one restricts its applicability, as Bongaarts and Feeney do, to the adult ages, say above 30, in low mortality populations.

With these caveats, the foregoing results allow us to relate period CAL or mean age at death to cohort life expectancy. As we noted in the previous section, when mortality declines the age structure lags behind the force of mortality and as a result

$$\text{CAL}(t) < e_0^{(p)}(t) < e_0^{(c)}(t). \quad (38)$$

Under the period-shift model we can be a bit more precise. We can show that the Bongaarts-Feeney measure M_4 , which is then the same as CAL, M_1 and M_2 , is the life expectancy of the cohort now at its mean age at death:

$$\text{CAL}(t + e_0^{(c)}(t)) = e_0^{(c)}(t), \quad (39)$$

a result easily verified by direct substitution, noting that the cohort born at t has mean age at death $(e_0 + r t)/(1 - r)$. Alternatively, one can go back in time and note that the cohort dying today was born at time $(t - e_0^\dagger)/(1 + r_c)$ and has life expectancy $\text{CAL}(t)$.

Goldstein (2006) has also derived the translation formula (39) and has used it to show that under a continuing linear shift the cohort born today would have life expectancy given by equation (32); this provides increased confidence in these results.


To summarize, conventional life expectancy e_0 measures how long a new born would live under current rates. This may not be a realistic estimate if mortality is declining. Under a period-shift model we have shown that a new born would in fact live longer, e_0^\dagger years. On the other hand period CAL, mean age at death and the Bongaarts-Feeney adjusted measure M_4 would all be lower, corresponding to the mean age at death of the cohort now reaching its life expectancy, provided the assumptions underlying the simpler linear shift model are satisfied.

3.5 A proportional hazards model

We now consider an example where the assumption is not quite satisfied, and therefore CAL, M_2 and M_4 differ. Specifically, consider a population with a constant stream of births and no mortality before age 30. Suppose the force of mortality follows a Gompertz function with $\alpha = -9.997$ and $\beta = 0.0855$, which as noted earlier fits very closely the U.S. 2002 life table. Suppose further that mortality has been constant long enough for the population to become stationary. In this case all four measures, CAL, mean age at death, e_0 and the Bongaarts-Feeney tempo-adjusted life expectancy M_4 are 78.45.

Suppose now that at time zero the force of mortality declines 20% at all ages. The conventional period life expectancy, being just a summary of age-specific mortality, would increase instantly to 80.97 to reflect this improvement. One has to be careful not to conclude that all cohorts will live this long, as the calculation applies only to the cohort age 30 at time zero, assuming mortality remains constant thereafter. CAL, on the other hand, doesn't change at time zero but starts increasing immediately afterwards as the decline in mortality is reflected on the standardized age distribution. Eventually the population becomes stationary again and CAL reaches 80.97. Figure 5 shows the trajectory of CAL for this example.

Figure 5: Measures of Longevity After a One-Time Reduction in Hazard



Mean age at death doesn't change instantly either. Although this index depends on the observed force of mortality, which is 20% lower at time zero, the reduction factor appears both in the numerator and denominator and cancels out. It is only as the reduction works its way into the age structure that mean age at death starts to increase, eventually reaching 80.97. Figure 5 shows that the trajectory of mean age at death is very similar to CAL. The Bongaarts-Feeney tempo-adjusted measure depends on the force of mortality and a correction factor based on r , which I estimated using the TMR. (Using a numerical derivative of $M_2(t)$ gives very similar results except for the first two years.) The key result is that M_4 is very similar to the other two measures. It takes them nearly sixty years to fully reflect the instantaneous change in mortality that occurred at time zero.

The figure also shows cohort life expectancy, estimated assuming that mortality was constant both before and after time zero at the specified level. We plot a cohort's life

expectancy on the year when it reaches its mean age at death. We note that the three measures of longevity track the increase in cohort life expectancy, albeit only approximately.

4. Discussion

This paper has emphasized similarities between the analysis of fertility and mortality. I have argued that Ryder's translation formula can be applied quite generally to demographic surfaces. When the surface represents age-specific fertility rates the formula translates period and cohort quantum. When the surface represents survival probabilities the formula translates period and cohort tempo, but using CAL rather than conventional life expectancy. The common theme is that period and cohort demographic summaries can differ in times of change. I believe that labelling these differences a bias or distortion has been unfortunate. Period aggregates provide convenient summaries, while cohort aggregates are often needed to fully understand the underlying process.

I have also stressed the fact that the Bongaarts-Feeney framework is essentially the same for fertility and mortality, postulating a period shift in a cumulative schedule representing average parity or survival probabilities. The shift can be motivated by assuming that all cohorts delay childbearing or postpone death at the same rate, and is closely linked to accelerated failure time models used in survival analysis. The shift results in a proportionate reduction in fertility or mortality rates, which also move to older ages. The model applies to multiple cohorts but requires assuming that later cohorts experience not just a slowing down of time but also a delay in the onset of exposure, an assumption that may be less realistic and, in the case of mortality, requires restricting application to adult ages in low mortality populations. I have also proposed measures of tempo under changing fertility or mortality which complement the Zeng-Land interpretation of the Bongaarts-Feeney adjustment by applying to the same synthetic cohort.

Having stressed similarities between fertility and mortality, it is perhaps appropriate to remind ourselves of some fundamental differences. In the case of fertility a reduction in age-period specific rates could represent changes in the quantum or tempo of fertility: women could be having fewer children or just having them later (or both). By assuming that delays occur at all ages at the same rate the Bongaarts-Feeney framework can ingeniously separate the two types of change. In our illustration we could have misinterpreted a reduction in TFR from 4.0 to 3.2 as a change in completed family size, but because it was accompanied by an annual increase of 0.2 years in mean age of childbearing—which would lead to just such a reduction—we concluded that it was a pure tempo effect. This does not mean, incidentally, that the reduction in period rates is not real. The only way cohorts can still have 4.0 children but over a longer time is by having them at a slower pace. The new measure of tempo introduced here tells us how much longer it would take.

There are two reasons why mortality is different, even if the same period-shift model applies. First, mortality is a pure tempo phenomenon; everyone dies exactly one time and the only question is when. Consequently, a reduction in the period force of mortality can only mean that cohorts are delaying death. There is no risk of misinterpretation, and therefore, one might argue, no need for adjustment. Bongaarts and Feeney implicitly acknowledge this point when they note that mean age at death, which they view as a direct analog of mean age of childbearing, needs no adjustment. They do adjust the force of mortality, of course, but I view this adjustment as merely a device to bring the conventional calculation of life expectancy inline with CAL or mean age at death. I see no bias or distortion in the observed force of mortality, just as I see no bias in age-specific fertility, and the best proof of that is the fact that cohort survival is determined entirely by $\mu(a, t)$, not by its tempo-adjusted version. The question then is whether we should use standardized mean age at death or conventional life expectancy as a measure of longevity.

That brings us directly to the second reason why mortality is different, and it has to do with exposure. In fertility all women are exposed to have a birth, whether they have had one before or not, which makes $f(a, t)$ a true event-exposure rate. Both the cohort and period TFR and mean age of childbearing are summaries of these rates and are not affected by exposure. In the case of mortality only survivors are at risk of dying, which is why analytical interest usually focuses on the force of mortality $\mu(a, t)$, which acts on survivors $S(a, t)$ to produce deaths $d(a, t)$. For a cohort the choice of measure is immaterial because exposure is itself determined by the force of mortality and as a result conventional life expectancy and mean age at death are identical. For a period the two measures can be quite different when mortality is changing. Conventional life expectancy depends only on the period force of mortality $\mu(a, t)$, whereas mean age at death depends also on $S(a, t)$ and thus on the population's past mortality history. We have seen that under the strong assumption of a linear-shift model, mean age at death coincides with the life expectancy of the cohort now reaching its mean age at death.

The question we asked at the outset, 'How long do we live?', can thus be seen to have different answers depending on our precise definition of 'we'. Conventional life expectancy applies to a hypothetical cohort that is exposed to a constant set of rates. It has the great merit of also applying to everyone else when mortality is constant. But when mortality is changing the construction is less useful; why ask how long someone would live subject to these rates if they are changing? We know that they would probably live longer than that, and we can estimate how much longer if we are willing to make strong assumptions about future changes. In particular, a continuing linear shift to older ages leads to e_0^\dagger , the simple measure of life expectancy under changing mortality proposed here. It is also the case that when mortality is declining no cohort has yet lived that long, or even as long as e_0 would imply. The Bongaarts-Feeney measure tells us how long those dying today have lived, standardizing for cohort size, when the proportionality assumption

holds. The fact that those dying today haven't lived as long as today's newborns will probably live, under either fixed or changing rates, is not a bias or distortion; it's just a fact of life.

The foregoing discussion has emphasized the practical interpretation of various measures of longevity while implicitly accepting the conventional view that mortality change is driven by the hazard function. But the Bongaarts-Feeney approach is fundamentally different; it views mortality change as driven by gains in longevity that shift the age distribution. This deflates the hazard by a factor $1 - r$ and shifts it to older ages. Unfortunately, it is difficult to differentiate these frameworks empirically because the age patterns in low-mortality countries are very close to a Gompertz model, where a proportionate reduction in the hazard cannot be distinguished from a shift to older ages. But if mortality were to stop declining we would soon know, because the period-shift model predicts an increase in the hazard as the factor $1 - r$ disappears and our past catches up with us, whereas the conventional view is that the hazard would stay constant. Faced with such choice, one may very well prefer to see hazards continue to decline and live longer with the uncertainty.

5. Acknowledgments

I am grateful to John Bongaarts and James Vaupel for inviting me to participate in the Workshop on Tempo Effects in Mortality held in New York, November 18-19, 2004, sponsored by the Max Planck Institute for Demographic Research and the Population Council. I learned a lot from all the participants. I am also grateful to John Bongaarts for detailed comments on an earlier draft and for a stimulating email exchange.

References

- Arias, E. (2004). United States life tables, 2002. *National Vital Statistics Reports*, 53(6):271–291.
- Bongaarts, J. and Feeney, G. (1998). On the quantum and tempo of fertility. *Population and Development Review*, 24(2):271–291.
- Bongaarts, J. and Feeney, G. (2000). On the quantum and tempo of fertility: Reply. *Population and Development Review*, 26(3):560–564.
- Bongaarts, J. and Feeney, G. (2002). How long do we live? *Population and Development Review*, 28(1):13–29.
- Bongaarts, J. and Feeney, G. (2003). Estimating mean lifetime. *Proceedings of the National Academy of Sciences*, 100(23):13127–13133.
- Bongaarts, J. and Feeney, G. (2005). The quantum and tempo of life-cycle events. *Policy Research Division Working Paper*, New York City: Population Council, 207. <http://www.popcouncil.org/>.
- Coale, A. J. (1971). Age patterns of marriage. *Population Studies*, 25(2):193–214.
- Coale, A. J. and Trussell, T. J. (1974). Model fertility schedules: Variations in the age structure of childbearing in human populations. *Population Index*, 40:185–258.
- Goldstein, J. R. (2006). Found in translation? A cohort perspective on tempo-adjusted life expectancy. *Demographic Research*, 14(5):71–84. <http://www.demographic-research.org/volumes/vol14/5/>.
- Guillot, M. (2003). The cross-sectional average length of life (CAL): A cross-sectional mortality measure that reflects the experience of cohorts. *Population Studies*, 57(1):41–54.
- Guillot, M. (2006). Tempo effects in mortality: An appraisal. *Demographic Research*, 14(1):1–26. <http://www.demographic-research.org/volumes/vol14/1/>.
- Kalbfleisch, J. D. and Prentice, R. L. (2002). *The Statistical Analysis of Failure Time Data*. John Wiley and Sons, New York, 2nd edition.
- Keilman, N. (1994). Translation formulae for non-repeatable events. *Population Studies*,

48(2):341–357.

Ní Bhrolcháin, M. (1992). Period paramount? a critique of the cohort approach to fertility. *Population and Development Review*, 18:599–629.

Ryder, N. B. (1964). The process of demographic translation. *Demography*, 1(1):74–82.

Schoen, R. (2004). Timing effects and the interpretation of period fertility. *Demography*, 41(4):801–819.

van Imhoff, E. and Keilman, N. (2000). On the quantum and tempo of fertility: Comment. *Population and Development Review*, 26:549–553.

Vaupel, J. W. (2002). Life expectancy at current rates vs. current conditions: A reflexion stimulated by Bongaarts and Feeney's "How long do we live?". *Demographic Research*, 7. Article 8.

Wachter, K. W. (2005). Tempo and its Tribulations. *Demographic Research*, 13(9):201–222. <http://www.demographic-research.org/volumes/vol13/9/>.

Zeng Yi and Land, K. C. (2001). A sensitivity analysis of the Bongaarts-Feeney method for adjusting bias in observed period total fertility rates. *Demography*, 38:17–28.

Zeng Yi and Land, K. C. (2002). Adjusting period tempo changes with an extension of Ryder's basic translation equation. *Demography*, 39(2):269–285.