Gaps and Lags: Relationships between Period and Cohort Life Expectancy

Joshua R. Goldstein Kenneth W. Wachter

June 24, 2003

Abstract

The paper provides a new interpretation of period life expectancy as a lagged indicator of the experience of real cohorts in populations experiencing steady mortality improvement. We find that current period life expectancy applies to cohorts born about 50 years ago. Cohorts born today will be 70 years old before the period life table has a life expectancy at birth equal to their own. We illustrate the relationship between period and cohort mortality using historical and forecast life tables from Sweden and the United States. Analytical approximations of the lag are found that provide new insights into the relationship between periods and cohorts. We relate our findings to Ryder's classic results on period-cohort translation as well as the new approach of Bongaarts and Feeney on tempo distortion.

> Development of translation procedures has proven more difficult for mortality functions than for fertility functions ... (Ryder, 1964, p. 81)

1 Introduction

In 1999, the Swedish official statistics agency reported that the period life expectancy for both sexes was 79.4 years. A natural question that arises is who, exactly, is expected to live this long? Period life expectancy is usually thought of as a synthetic measure, referring to a hypothetical cohort that lives its entire life according to the rates of a single period. Real people and real cohorts do not live this way, aging instead through successive periods. In this paper, we explore the correspondence between period and cohort life expectancy using historical and forecast life tables from Sweden and the United States for the 19th, 20th, and 21st centuries. We derive analytical formula relating period and cohort life expectancy that provide insight into the relationship between periods and cohorts. We find that current period life expectancy is about equal that of the cohort born 50 years ago. The lag of period measures has increased over time and is expected to increase further as mortality improves.

The correspondence between period and cohort rates is of natural interest. Period life expectancy is almost all that is ever reported by statistical agencies. The advantage of the synthetic measure is its timeliness. Cohort life expectancy can only be observed once the last member of a cohort dies, and is thus is typically out of date by more than a century. The disadvantage of period measurement is that people cannot use it to gauge the length of their own lives or of others. Individuals, actuaries, and demographic forecasters would choose cohort rather than period measurement of life expectancy, if only it were available.

Recently, period life expectancy has been the target of a new criticism. Bongaarts and Feeney (2002) argue that the period life tables are distorted by "tempo" effects in much the same way that the period total fertility rate is distorted by postponements in the timing of childbearing. This innovative interpretation remains controversial.[1] Fertility tempo and quantum are, at least in theory, distinguishable. For mortality, however, there is no such clear distinction. Since everyone must in the end die, the only kind of mortality decline that can occur is a postponement of the inevitable.

A specific problem with the tempo-based theory of mortality change is the interpretation of "tempo-adjusted life expectancy." It is said to be something like the period life expectancy that would be observed in the absence of "tempo effects." Intuitively, in the absence of changes in the timing of mortality, period mortality should equal cohort mortality. However, "tempo effects" are defined technically not as the product of changes in mortality

[1] Vaupel (2002), for example, writes some "scholars have told me that they think the notion of tempo effects on mortality is a red herring, a pink elephant, a bte noire, a siren's song." (p. 374).

timing but as a consequence of a particular shift in the distribution of age at death from one period to the next. As a result, no relationship between cohort and ”tempo-adjusted” period life expectancy has been asserted or established.

A potentially simple way around this confusion is to explore empirically how observed period life expectancy corresponds to cohort life expectancy. Setting aside the issue of whether mortality improvement is or is not a tempo effect, we can address the question of cohort survival head-on.[2] The empirical translation is possible in all industrialized countries where period and cohort life expectancy are slowly and steadily improving. The steady, monotonic change allows every observed value of period life expectancy to be mapped to a unique cohort with the same life expectancy.

The translation of life expectancy between cohorts and periods has remained an open problem for nearly 40 years. In his seminal paper, Ryder (1964) noted that “development of translation procedures has proven more difficult for mortality functions than for fertility functions” (p. 81).

Our approach adds a new perspective to the literature but is not too distant from Ryder’s original thinking. Whereas Bongaarts and Feeney view period life expectancy as a “distorted” measure of underlying mortality, Ryder considered periods a moving average of underlying cohort processes. View-

[2] It is not our intent here to enter into a critique of the Bongaarts-Feeney approach. Guillot (2002) offers an extensive critique and alternative period measure; Vaupel (2002) provides additional discussion and interpretations of mortality tempo.

ing period measures as lagged indicators of cohort experience, as we do here, is an extension of Ryder's logic using reasoning similar to the mean-value theorem in calculus. If each period combines many cohorts and if cohorts are changing steadily, then there exists a particular cohort with the same summary mortality as the current period.

2 The relationship sketched

Steady mortality improvement has been a feature of 20th century demography throughout the industrialized world. Despite setbacks like the Spanish flu epidemic, wars and reconstruction, economic ups and downs, the invention of antibiotics, chemotherapy, surgery and so forth, age-specific survival rates have improved at a remarkably constant pace (Lee and Carter 1992, Tuljapurkar et al. 2000). Age-specific mortality rates have fallen at about 1 to 2 percent per year.

In concert with the decline in age-specific mortality rates, life expectancy at birth has risen steadily. The improvement in period life expectancy is sketched in Figure 1 along with the associated improvements in the life expectancy of cohorts born in each period. Cohort life expectancy is greater than period life expectancy because the cohort will experience improving mortality rates as they age.

Figure 1 illustrates the three concepts we will use to analyze the relationship between period and cohort life expectancy. We use the term "lag" or λ

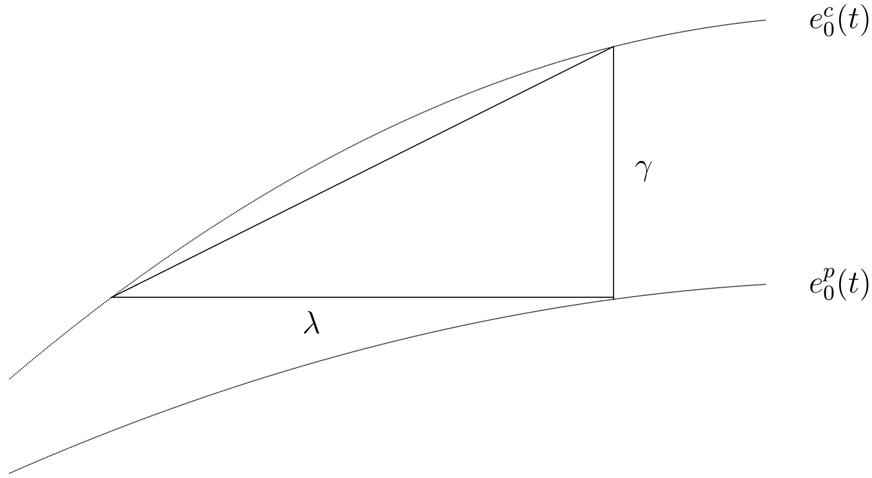

Figure 1: Sketch of gap and lag between cohort and period life expectancy when mortality and entropy are decreasing with time. The slope of the line forming the hypothenuse of the triangle is $\frac{\gamma}{\lambda}$.

to describe the horizontal distance between the cohort and period lines for any given value of life expectancy. The lag tells how out-of-date period life expectancy is as a measure of cohort experience. It tells how far back in time one has to go to find the birth of a cohort with life expectancy equal to that of the current period.

We define the vertical distance between the period and cohort curves in a given year as the "gap" or γ. The magnitude of the gap tells how much today's period life expectancy differs from the life expectancy of the cohort born today. It is the gain in life expectancy that a cohort will benefit from by experiencing future declines in mortality.

Finally, we use the term "slope" to describe derivative of life expectancy with respect to time either curve. The basic relationship between these concepts is that the slope will be approximately equal to the ratio of the gap to the lag, the rise to the run. For example, in 1999 Swedish period life expectancy was increasing at an annual rate of .. years, lagged behind cohort life expectancy by xxx years, and was about yyy years smaller than the forecast life expectancy of the cohort born in 1999. We can see that the slope is about, although not exactly, equal to the ratio of the gap to the lag.

In populations experiencing a steady worsening of mortality, the gap will be negative, but the lag will still be positive. This is because period mortality will still consist of mortality experienced by cohorts born in the past.

3 Observed and Projected Relationships between Cohort and Period Life Expectancy

We are now ready to look at actual human experience. In this paper, we use Swedish and U.S. mortality as illustrations. Sweden has a long time series of accurately-measured mortality rates, is a leader in mortality improvement, and has been used in past papers (Vaupel, Bongaarts and Feeney.) It also happens to be the only country with easily available historical cohort life tables (BMDB). The United States has also experienced steady mortality decline over the last century. We use U.S. data to show that generality of our results among the industrialized nations and because the BMD offers cohort mortality into the 21st century based on forecasts by the U.S. Social Security Administration. This allows us to illustrate the relationship between cohort and period mortality without doing any forecasting ourselves.

Figure 4 shows observed and estimated values of period and cohort life expectancy in Sweden and the United States for females. The data is publicly available at the Berkeley Mortality Data Base.

The Swedish values are based on observed periods and completed cohorts. They illustrate the relationship between period and cohort life expectancy during the demographic transition as life expectancy rises rapidly from under 50 to over 70 years. We see that the vertical gap between period and cohort life expectancy in any year is about 5 years. The horizontal lag averages

Figure 2: Period and Cohort Life Expectancy for females in Sweden and the United States

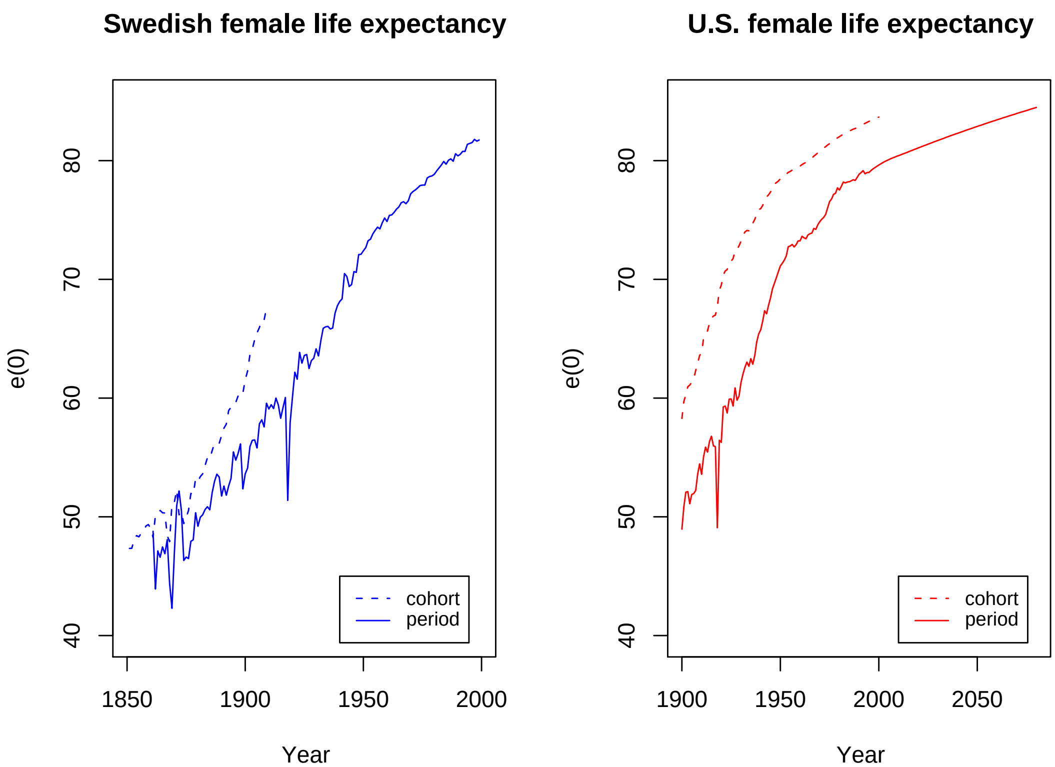

Source: Berkeley Mortality Database. Note: U.S. life expectancy estimates include mortality forecasts from 1999 onwards, based on Social Security Medium projection.

about 10 to 15 years.

The United States values are based on projected from 1999 onwards, based on the Social Security Administration's medium projection. The slowing rate of life expectancy improvement over time is due in part to the increasing entropy of the life table but also to the fact that the projections assume a slowing of age-specific mortality decline from the historic rate of about 1 percent to about 0.5 percent.

This figure illustrates our basic finding, that current period life expectancy is equal to cohort life expectancy lagged by about 50 years. We also see that this lag has grown over time and is expected to grow in the future. In 1900, period life expectancy was out of date by only about 20 years with respect to cohorts. By 2100, however, the lag will have grown to 70 years. The relationship between the lag and the level of life expectancy is also not constant, growing from about less than half of life expectancy to more than three-quarters.

At the same time, the gap between period and life expectancy in terms of expected years of life has not grown so much. In 1870 Sweden, the gap is about 4 years; in 1900 Sweden it is about 5 years, compared to 8 years in the United States at the same time; by 2000 it is about 5 in the United States. It is worth noting that the size of the gap is not a constant relative to life expectancy. As life expectancy has nearly doubled, the gap has grown only slightly.

4 Understanding Gaps and Lags

Having seen how period and cohort life expectancy are empirically related to one another over the past few centuries, we turn to a theoretical explanation for the magnitude of the gaps and lags between period and cohorts.

Our approach is to consider a simple, but not unrealistic, model of steady age-specific mortality change. We use this simple model to approximate cohort life expectancy in terms of the current period life table, the rate of mortality change, and the age-pattern of mortality change. Our cohort life expectancy approximation is then used to obtain explicit expressions for the gap and lag between period and cohort life expectancy.

4.1 An Age-specific Model of Mortality Change

Here first provide an analytic approximation of cohort life expectancy in terms of period mortality and the pattern of mortality improvement. Let the age specific hazard rates at age x and year t be

$$h(x,t) = h(x)e^{-kb(x)t}, \tag{1}$$

where $h(x)$ is an arbitrary baseline age-schedule of mortality, k is the average annual exponential rate of mortality improvement, and $b(x)$ is an arbitrary fixed schedule allowing different rates of improvement at each age. For identifiability, we constrain the $b(x)$ terms to a unit mean.[3]

Take $t = 0$ as the reference period. In this case, the cohort born λ years earlier experiences cohort hazards

$$h(x, x - \lambda) = e^{-kb(x)(x-\lambda)}\, h(x). \tag{2}$$

Setting the reference period to time equals zero simplifies the algebra but allows us to consider any arbitrary time period without loss of generality.

The life expectancy of the cohort born λ years before the reference period is

$$e_0^C(-\lambda) = \int_0^{\omega} l(x, x - \lambda)\, dx = \int_0^{\omega} \exp\left(-\int_0^x h(a, a - \lambda) da\right) dx \tag{3}$$

Period life expectancy at the reference period time $t = 0$ will be simply

$$e_0^P(t = 0) = \int_0^{\omega} l(x) dx = \int_0^{\omega} \exp\left(-\int_0^x h(a) da\right) dx \tag{4}$$

4.2 Approximating cohort life expectancy

To explore the lags and gaps between period and cohort life expectancy, it is useful to write cohort life expectancy as a function of the rate of mortality improvement k as well, again allowing age-schedules of mortality and mor-

---

[3]Our model is analogous to the *expected* forecast of Lee and Carter model of mortality change, with Lee and Carter's k_t following the time path $k \cdot t$ (Lee and Carter, 199, p. 661). It can also be expressed using Vaupel and Romo's notation, replacing $\rho(x)$ with $kb(x)$ and imposing constancy on $\rho(x)$ over time (Vaupel and Romo 2003, p. 202).

tality improvement to take arbitrary forms. Substituting the cohort hazards given by (2) into (3), we see that cohort life expectancy as a function of k is

$$e_0^C(-\lambda, k) = \int_0^{\omega} \exp\left(-\int_0^x e^{-kb(a)(a-\lambda)}h(a)da\right) dx \tag{5}$$

Cohort life expectancy can now be approximated for small k by expanding $e_0^C(-\lambda, k)$ around $k = 0$. To second order,

$$e_0^C(-\lambda, k) \approx e_0^C(\lambda, 0) + k\, \frac{\partial e_0^C}{\partial k}\bigg|_{k=0} + \frac{1}{2}k^2\, \frac{\partial^2 e_0^C}{\partial k^2}\bigg|_{k=0} \tag{6}$$

Some manipulation then gives

$$\begin{aligned} e_0^C(-\lambda, k) \;\approx\;& \int_0^{\infty} l(x)dx + \\ & k\int_0^{\infty}(x-\lambda)b(x)e(x)l(x)h(x)\,dx + \\ & \frac{k^2}{2}\left[\int_0^{\infty} l(x)\left(\int_0^x (a-\lambda)b(a)h(a)\,da\right)^2 dx - \int_0^{\infty} l(x)\int_0^x (a-\lambda)^2 b(a)^2 h(a)\,dadx\right]. \end{aligned} \tag{7}$$

All of the life table functions in this approximation refer to the period life table of the reference period.

4.3 Approximating Lags and Gaps

We can now use equation (7) to approximate the lag between cohort and period life expectancy. Setting period life expectancy equal to cohort life expectancy we can solve for λ, the date of birth of the cohort. Using only the zero'th and first order terms of (7),

$$e_0^P(t=0) = e_0^C(-\lambda) \tag{8}$$

$$\approx \int_0^\infty l(x)dx + k\int_0^\infty (a-\lambda)b(a)e(a)l(a)h(a)da. \tag{9}$$

The integral $\int_0^\infty l(x)dx$ is, however, period life expectancy, and can be subtracted from both sides. This allows one to obtain an estimate for λ by distributing the $(a-\lambda)$ term and and rearranging :

$$\hat{\lambda}_1 = \frac{k\int_0^\infty ab(a)e(a)l(a)h(a)da}{k\int_0^\infty b(a)e(a)l(a)h(a)da}. \tag{10}$$

We call $\hat{\lambda}_1$ the first order estimate of the lag.

We can also use our cohort life expectancy approximation to estimate the gap between period and cohort life expectancy in the arbitrary reference period. The gap, recall, is defined as the difference between period life expectancy and the life expectancy of the cohort born in that period. Formally,

$$\gamma = e_0^C(\lambda=0) - e_0^P(t=0) \tag{11}$$

To first order, applying (7),

$$\hat{\gamma}_1 = \int_0^\omega xb(x)e(x)l(x)h(x)dx. \tag{12}$$

4.4 Interpreting the Approximations

The approximation for the lag has a direct relationship to the sketch shown in Figure 1. The numerator of $\hat{\lambda}_1$ is the gap $\hat{\gamma}_1$, which suggests that the denominator should be equal to the slope of the hypotenuse shown in the figure. The first order approximation linearizes cohort and period life expectancy over time. Indeed, from equations (1) and (4), the derivative of period life expectancy at time with respect to time is $k \int_0^\infty b(x)e(x)l(x)h(x)\,dx$, the numerator of the lag. should

Interestingly, and somewhat surprisingly, this first order approximation of the lag does not depend on the rate of mortality improvement k, since it cancels out of both the numerator and denominator.[4]

The figure illustrates how little the lag varies across a wide range of k, including negative values. Instead, the rate of mortality improvement primarily drives the difference between cohort and period life expectancy at a moment in time (the gap), and the slope. The fact that the gap and the slope both depend, to first order, linearly on k, allows k to cancel out of both the numerator and denominator of the first order lag estimate.

The lag approximation can be interpreted as a mean age of a distribution related to entropy. To gain some insight into the distribution of which the lag is the mean, recall that period life expectancy at birth is equal to

[4]The approximation is only defined for $k \neq 0$. When $k = 0$ the lag and gap are both zero since period and cohort life expectancy coincide.

Figure 3: Cohort e(0) trajectories by rate of mortality decline k

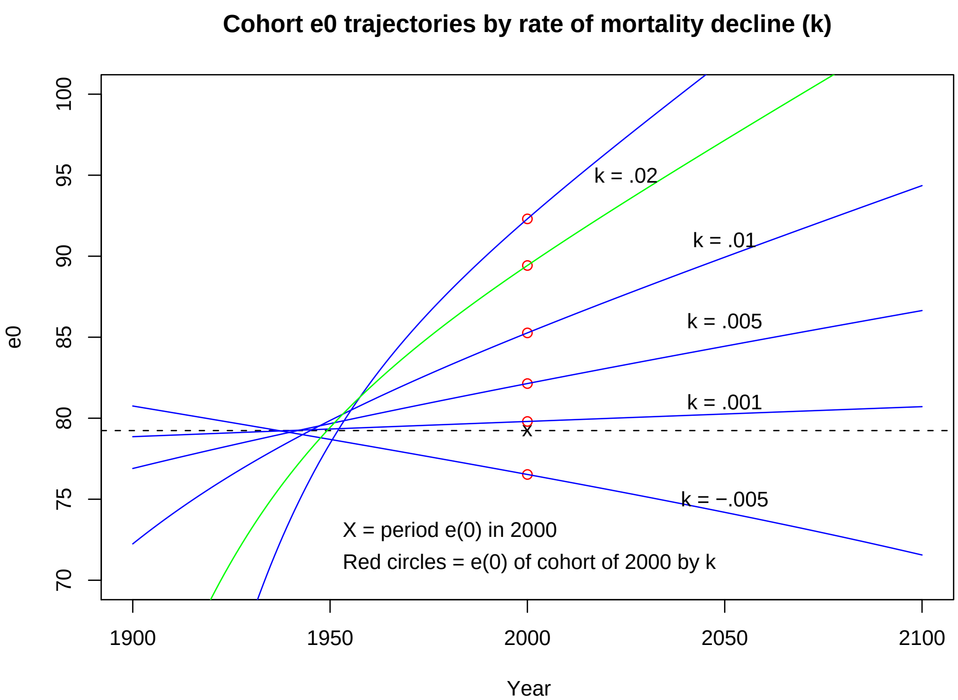

Dashed line gives period e(0) in 2000. Lag can be seen as the horizontal distance between "x" and the intersection of the cohort curve with the horizontal dashed line. The gap is the vertical distance between "x" and "o". The figure shows that the lag is roughly constant over a wide range of mortality decline rates but that the gap is highly dependent on k.

$\int ah(a)l(a)da$ since $h(a)l(a)$ is the density of age at death. The distribution we have is a weighted mean age at death, where the weights are equal to the remaining years of life expectancy itself weighted by the relative rate of mortality improvement.

Further insight into the approximations can be gained by thinking of mortality improvement as the years of additional life that would be granted to someone whose life was "saved." (Vaupel 1986, p. 148). Vaupel notes that the usual measure of entropy H can be rewritten as $\int_0^\infty h(x)l(x)e(x)dx/e_0$. Since $e(x)$ gives the expected years of life of someone whose "first death" is averted at age x, the numerator of this expression is total effect of eliminating the first death of everyone (which would have occured at age $h(x)l(x)$). The denominator makes this total effect relative to life expectancy and gives the proportional effect of "staying the hand of death once" for everyone in the population.

In our case, we have a similar expression to Vaupel's alternative expression of entropy, except that we have inserted the additional weight function $b(a)$ and we have the mean of the distribution rather than its sum.

The gap, or numerator of the lag approximation, has a direct interpretation. If we want to know how many years of life expectancy a cohort will gain by living through future mortality improvements, then $k \int_0^\omega xb(x)e(x)l(x)h(x)dx$ tells how many person years will be gained by "staying the hand of death" with an intensity of $kb(x)x$ at each age x. The factor x in included in this intensity to account for the longer period of mortality decline for the cohort by the time it has reached this age.

The gap itself varies relatively little over, but is instead mostly a function of k. This is because as mortality falls, $h(x)$ gets smaller while $l(x)$ and $e(x)$ get larger, largely canceling each other out.

The approximations also have the intuitive geometric interpretation given in Figure 1. The first order approximation linearizes the curves of life expectancy over time. In this case, $\frac{\gamma}{\lambda}$ should be equal to the slope of life expectancy with respect to time. And, indeed this turns out to be the case. Differentiating,

$$\frac{\partial e_0^P(t)}{\partial t} = k \int_0^\infty b(x)e(x)l(x)h(x)\,dx \tag{13}$$

. The same result is found in Vaupel and Romo's equation (15).

The analytical approximations of the gap and lag thus give us a substantive interpretation of the relationship between cohort and period life expectancy. We see that the gap and lag are related to each other via the rate of change of period life expectancy with respect to time, which is itself a function of the entropy of the period life table.

Furthermore, the approximations show that the lag of period life expectancy will to first order be independent of the rate of mortality change but will depend on level of mortality as expressed by the mean age of entropy of the period life table. The gap, on the other hand, will be very dependent on the rate of mortality change but insensitive to mortality level.

Figure 4: Observed and approximated gaps and lags between period and cohort life expectancy in the United States

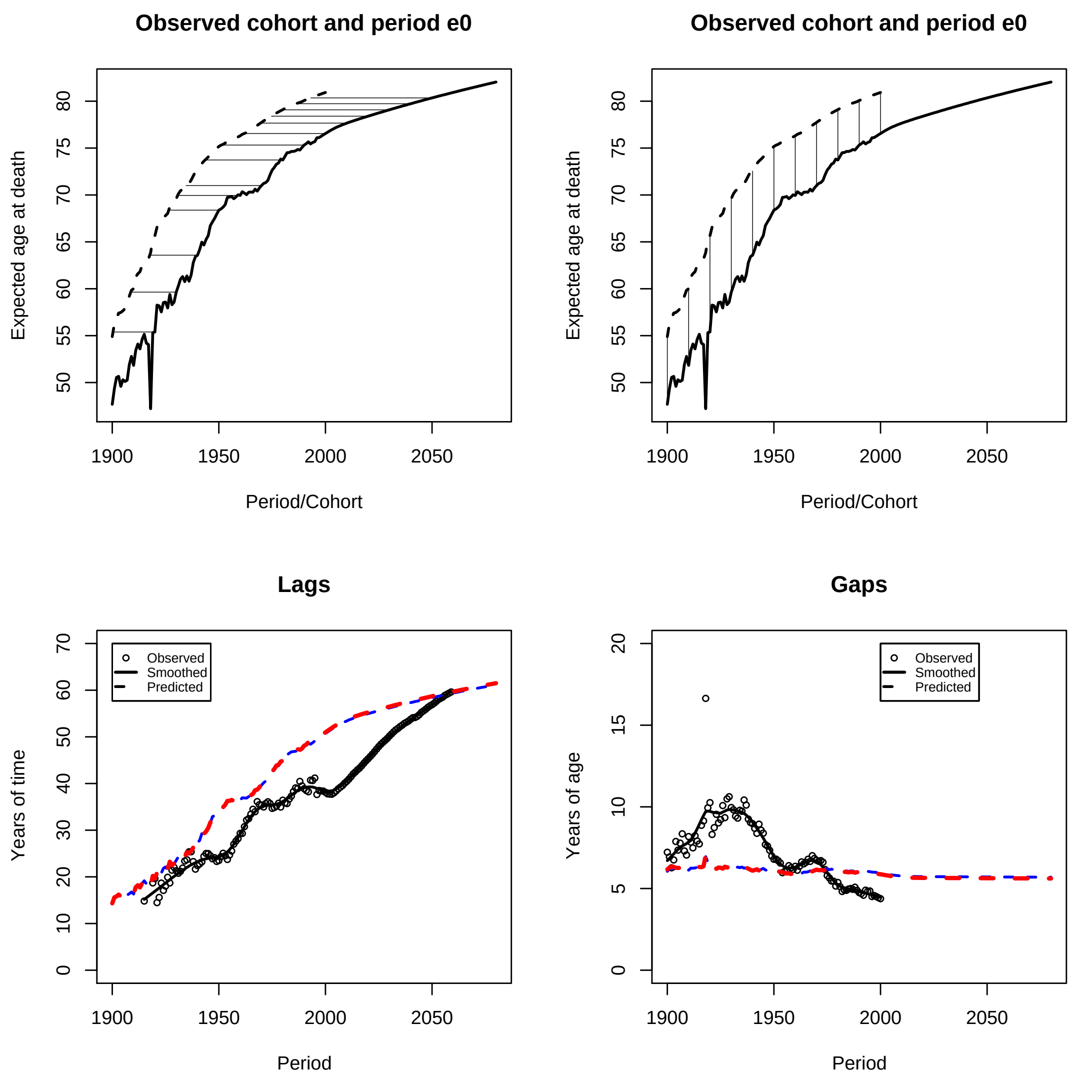

Predicted Gaps and Lags are based on first order approximations

The approximations are compared to the historical and forecast experience in Figure xxx. The approximations differ from the observed gaps and lags both (i) because age-specific mortality decline has fluctuated over time rather than according to the equation xxx and (ii) our approximations are accurate only to first order even when mortality decline has been perfectly steady.[5] Nonetheless, we see that the approximations do a good job in capturing the increasing lag and the relative constancy of the gap.

5 Translation of Other Life Table Measures

It is possible to use our approach to translate other life table quantities in addition to life expectancy at birth.

5.1 Life expectancy at age x

Extending our results to life expectancy at other ages is straightforward. Since, life expectancy at older ages is of great interest for retirement and pension planning, we consider e(65).

The equations generalize easily.

[to be written]

[5]It is possible to solve the second of these problems by including the second order terms in the approximation but there is little payoff to introducing additional complexity since the model of age-specific mortality decline is only approximates what has really happened. Our philosophy is a first order approximation is good enough for a model that is itself only an approximation of historical experience.

5.2 Life expectancy at older ages under Gompertzian mortality

It is also interesting to consider Gompertzian mortality in which mortality decline is constant across all ages. This model, first put forward by Vaupel, is analytically tractable and consistent with 20th century mortality decline in the industrialized world. When mortality is Gompertzian, the gaps and lags take on a particularly simple form. [to be written]

5.3 Survival

We have concentrated for the most part on life expectancy which is the area under the survival curve. The cohort born λ years before a given period will have a survival curve with the same area underneath it but will not have identical chances of survival to each age.

Using the same techniques as above, and denoting the gaps and lags with x subscripts to indicate that they apply to the survival function l_x, we find

$$\hat{\lambda}_x = \frac{\int_0^x ab(a)h(a)\,da}{\int_0^x b(a)h(a)\,da} \tag{14}$$

and

$$\hat{\gamma}_x = kl_x \int_0^x ab(a)h(a)\,da. \tag{15}$$

[discussion for these results to be written]

Table 1: Summary of results for lags and gaps

Period e(0)	Lag
50-60	20
70-80	50
85-95	70

k	Gap
0	0
0.001	0.6
0.005	3.0
0.010	6.0
0.020	12.0

Note: this values are rough and are taken from the earlier figures not from the analytical approximations. They need to be recalculated exactly taking Sweden as baseline.

6 Rules of Thumb

Table 1 gives the essence of our results. For high mortality populations at the beginning stages of mortality decline, period life expectancy lags cohort life expectancy by about 20 years. Current low mortality populations are experiencing period mortality that applies to cohorts born about 50 years ago. The life expectancy of cohorts born today will be seen in period measures some 70 years from now. These rules of thumb are applicable across a wide range of speeds of mortality decline.

The period measure is becoming increasingly out of date from the point of view of how lagged it is. On the other hand, the absolute difference in years of expected life remains rather constant across different levels of mortality, instead depending mostly on the rapidity of mortality decline. In this case, the rule of thumb appears to be that a 1 percent rate of annual decline translates into a gap between cohort and period life expectancy of about 6 years.

7 Conclusions

In populations undergoing steady mortality change, we have found that period life expectancy can be fruitfully interpreted as a lagged measure of underlying cohort experience. Our approach has been to look at the relationship between period and cohort life expectancy in terms of two measures: (1) the lag of λ years following a cohort's birth that it takes for period life expectancy to rise to the level experienced by that cohort and (2) the gain in life expectancy γ that a cohort benefits from by experiencing changing rather than fixed mortality.

We find that as mortality has fallen, the lag between periods and cohorts has increased. Our formal results show that the lag is determined by the level of mortality as expressed by an expression related to life table entropy and the age pattern of mortality decline. The rate of mortality decline is not a major factor.

On the other hand, we find that the gap between period and cohort life expectancy has remained roughly constant across a range of mortality levels. What matters here, it appears, is the rate of mortality decline, more than the absolute level.

Compared to Ryder's translation of total fertility rates, our approach to mortality is both satisfying and unsatisfying. The satisfaction comes from having solved what Ryder himself termed a more difficult problem. On the other hand, even given a simple model of mortality decline, we were able to produce only approximate rather than exact analytical expressions for λ and γ. More accurate approximations can be obtained at the price of introducing the complexity of second order approximations. But even these more exact solutions apply only to a model that represents a simplified pattern of mortality decline. The first order approximations, we feel, provide the appropriate accuracy for working with a model of steady mortality decline that itself is only an approximation of historical experience.

The main import of our results is that in populations experiencing steady mortality decline, period life expectancy can now be equated with the experience of specific cohorts. The rule of thumb for current life expectancy in most low mortality populations is that it applies to cohorts born about half a century ago.