Prof. Joshua Goldstein

Office of Population Research

Wallace Hall

Princeton University

Princeton, New Jersey 08544

U.S.A.

josh@opr.princeton.edu

Gaps and Lags: Relationships between Period and Cohort Life Expectancy

by Joshua R. Goldstein and Kenneth W. Wachter

Princeton University and University of California, Berkeley, U.S.A.

Short Title: Gaps and Lags

Submitted to *Population Studies* 22 March 2005

Revised for resubmission 23 March 2006

Gaps and Lags: Relationships between Period and Cohort Life Expectancy

Joshua R. Goldstein Kenneth W. Wachter

23 March 2006

Abstract

This paper provides an empirical and analytic context for regarding 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 in the industrialized world applies to cohorts born some 40 to 50 years ago. Lags track an average age at which future years of life are being gained, in a sense which we make precise. Our findings augment Ryder's classic results on period-cohort translation.

Keywords: expectation of life, demographic translation, period and cohort, linear shift models.

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

1 Introduction

In 2003, the Swedish official statistics agency reported that the period life expectancy for both sexes was 80.2 years. Who, in fact, would be expected to live this long? The usual answer is "No one": Period life expectancy is thought to be an entirely synthetic measure, referring to an hypothetical cohort living its entire life according to the rates of a single period. Real people do not live out their lives in this way. They age and die as members of cohorts through successive periods subject to ever-changing rates.

In modern industrialized nations, however, these ever-changing rates tend to be changing quite steadily. Thanks to the steadiness of change, usually period life expectancy for each period *is* the life expectancy for a real cohort whose birth date is separated from the period by a systematic lag. Period life expectancy is not as purely hypothetical as the language in demography textbooks might lead us to believe. In this paper we focus attention on the systematic correspondence between periods and cohorts in terms of life expectancy. Empirically, we explore the correspondence in historical and projected life tables from the United States and Sweden for the nineteenth, twentieth, and twenty-first centuries. Analytically, we derive approximations

which yield a simple interpretation and representation of period-to-cohort lags. Historical comparisons and simulation show that the lag approximations capture the magnitude and the pattern of time change.

Under contemporary conditions, period life expectancy is about equal to cohort life expectancy for the cohort born about 40 to 50 years ago. This lag has lengthened over time and it is expected to lengthen further as mortality improves. As we shall see, the longer lags reflect the later ages at which gains in years of life from improvements in mortality have come to be concentrated. The lag between periods and cohorts is not equal, as one might initially suspect, to the mean age of death, but rather to the mean age of mortality improvement.

Interpreting period life expectancy in terms of cohort experience has obvious appeal. Almost anyone would prefer cohort measures, were it not that cohort measures lack timeliness, remaining incomplete for more than a century pending the last cohort member's death. The one main virtue of period mortality rates is their quick availability. Unlike period fertility rates, which drive population renewal by determining the size of each period's new crop of births, period mortality rates do not enter into the demographic Renewal Equation. It is cohort survivorship that shapes the populations at risk.

Our proposal to regard period life expectancy as a systematically lagged cohort indicator follows in the spirit of Norman Ryder's tradition of demographic translation. Ryder (1964) recognized that the approach to the translation of fertility measures which he pioneered did not carry over read-

ily to mortality. His well-known comment is the epigraph to this paper. The translation of life expectancy between periods and cohorts has remained an open problem over the intervening forty years.

If, following Ryder, one wants to view period measures as moving averages of underlying cohort processes, then it is a natural step to think in terms of lags. Consider, by analogy, the mean-value theorem of the calculus which requires the average value of a continuous function over an interval to equal some particular value of the function within the interval. If each period is in effect averaging the present experience of many cohorts born in the past, and if cohorts are changing steadily, then we expect there to be a particular past cohort with the same summary measure of mortality as the period. Hence we expect to be able to represent period life expectancy as a lagged cohort measure. Like Ryder's fertility translation, our translation of periods to cohorts applies to the average lifetime experience – in this case, life expectancy – but not to a full correspondance of the lifetable at every age.

Renewed interest in demographic translation for mortality measures has been stimulated by Bongaarts and Feeney (2002). Their controversial claims about tempo distortions have been assessed by each of us elsewhere, in Goldstein (2006) and Wachter (2005). What matters for present purposes are measures and models which have come to the fore during the ensuing discussion. Instead of starting with period life expectancy and seeking cohort equivalents, Schoen and Canudas-Romo (2005) define a period measure based

directly on cohort information. Their "Average Cohort Life Expectancy" or ACLE is a weighted average of the cohort life expectancies of the cohorts alive in a given period, weighted by the proportional representation these cohorts would have in the period population if initial cohort size had been constant in the past. Current values of ACLE are calculated from forecasts of future mortality along with records of past mortality. As a measure of the current state of overall longevity, the ACLE has a straightforward rationale.

Building on Brouard (1986), Guillot (2003) analyzes the "Cross-sectional Average Length of Life" CAL, equal to the total period population which would be produced from a constant unit stream of past births subject to a set of age and time-specific past survival rates. Since CAL averages over the past survival of the many cohorts composing the standardized period population, it is relatively immune to temporary period shocks that may affect period life expectancy, and, in the face of steady change, it lags behind period life expectancy.

Relationships between the full CAL measure and cohort life expectancy at birth are complex. However, CAL and life expectancy can also be calculated solely over adult ages, conditioning on survival to some age like 30. For these adult measures, under a set of stringent but illuminating special conditions, the period value of CAL comes out to equal the cohort life expectancy of the cohort born CAL years in the past, as shown by Goldstein (2006). In other words, under these special conditions (which constrain the future as well as the past), when a cohort reaches an age equal to its life expectancy, its life

expectancy equals the period value of CAL. This correspondence leads to a simplified formula for the lag between period and cohort adult life expectancy for this special case, with consequences discussed in Section 6.

2 Sketch of the relationship

Steady mortality improvement has been a feature of twentieth century demography throughout the industrialized world. Despite the shocks of wars, the flu epidemic, economic depressions, and new diseases, and despite breakthroughs such as the invention of antibiotics, chemotherapy, and open-heart surgery, age-specific survival has been improving at a remarkably constant pace (Lee and Carter 1992, Tuljapurkar, Li and Boe 2000). Mortality rates at all ages in many developed countries have been falling by between 1 and 2 percent per year.

With the decline in age-specific mortality rates, life expectancies at birth have been rising steadily. The pattern is shown in stylized form in Figure 1. The horizontal axis represents time, the vertical axis years of life. The lower curve stands for period life expectancy $e_0^P(t)$ at time t. The upper curve stands for cohort life expectancy $e_0^C(t)$ for the cohort born at time t. Cohort life expectancy is greater than period life expectancy because the cohorts experience improving mortality rates as they age. The curves are concave downward to reflect the diminishing rate of increase in life expectancy over time often described in terms of the measure called lifetable entropy.

Figure 1 about here.

The straight lines on Figure 1 illustrate the three concepts we use to analyze the relationship between period and cohort life expectancy. We define the vertical distance between the period and cohort curves in a given current year as the "gap" γ. Formally,

$$\gamma(t) = e_0^C(t) - e_0^P(t). \tag{1}$$

The magnitude of the gap tells how much period life expectancy differs from the life expectancy of the cohort born in the current period. It is the "bonus" that a real cohort receives from taking advantage of future mortality improvement. The units of γ are years of life.

We define the horizontal distance between cohort and period curves as the "lag" λ. The lag tells us how far back in time from the current period we have to go to find a cohort with equivalent life expectancy. (If no such cohort exists, the lag is undefined.)

$$\lambda(t) = \min\{\lambda : e_0^C(t-\lambda) = e_0^P(t)\} \tag{2}$$

The units of λ are years of time.[1]

[1]The lag defined (2 is a "backward" lag in the sense that it looks backward for a cohort that matches the current period. One can also define a "forward" lag by the number of years it will take for period life expectancy to reach the current level of cohort life expectancy. The first-order estimates we present below are the same for forward and backward lags, but the second order estimate is not, and in practice, the two can differ substantially.

We use the term "slope" to describe the derivative of life expectancy with respect to time on either curve. Figure 1 contains an approximate triangle with a mildly curving hypotenuse. The ratio γ over λ is the average cohort slope along the upper hypotenuse, not far from the period slope at the vertex.

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

In this paper we ask, first, what is the relationship between current period and cohort life expectancy? This question can be answered empirically by looking at the magnitude of the current gap and lags. We ask, second, how has this relationship been changing? This question can also be answered empirically, using historical data and forecasts. We ask, third, what determines the magnitudes of the gaps and lags? For answers, we turn to mathematical models, obtaining approximate formulas and spelling out the intuitions to which they lead. The appeal of the formulas is not that they enable highly accurate estimates of gaps and lags – which can be obtained directly from observed and forecast life expectancy trajectories – but rather that they provide a qualitative understanding of the magnitude, determinants, and dynamics of gaps and lags.

3 Empirical gaps and lags

We begin with empirical relationships between cohort and period life expectancy seen in documented human experience. For illustration, we use official estimates and projections for females from Sweden and the United States, prepared by Statistics Sweden and by the U.S. Social Security Administration. [2]

Figure 2 about here.

Empirical curves for Swedish and U.S. women are shown in Figure 2. Cohort life expectancies plotted by cohort date of birth for 1900 to 1959 for Sweden and for 1900 to 2000 for the U.S. at the top of each plot depend significantly after about 1915 on the projections. Below them, period life expectancies for Sweden and the United States from 1900 to 2000 reflect observed data. Lags are shown in the middle of the plots. Gaps appear at the bottom.

We see steady improvement in period life expectancy in both countries, with the exception of the 1918 influenza pandemic. Future improvements are expected to be at a slower rate due in part to increases in lifetable entropy (Keyfitz and Caswell 2005, p. 78). In the United States, there is also a

[2]Swedish and U.S. historical period estimates and U.S. cohort projections are available in the Human Mortality Database, *www.mortality.org* and the Berkeley Mortality Database at *www.demog.berkeley.edu/∼bmd.* Swedish projections are described at *www.scb.se.* We thank Hans Lundstrom for making the Swedish age-specific projections available to us. Cohort life expectancy in Sweden was obtained to 1959 by extrapolating the forecast exponential age-specific improvements above age 100 from 2050 to 2070.

slowdown in the pace of age-specific mortality decline built into the Social Security forecasts (Technical Panel on Assumptions and Methods 2003).

With improving survival, cohort life expectancy is consistently higher than period life expectancy, making the observed gaps positive in both countries. In Sweden, gaps rise from about 5 years in the middle of the 19th century to about 12 years in the middle of the 20th century. In the United States, the observed gaps begin at about 10 years in 1900 and fall to about 6 years in 1960. In both countries, the gap spikes during the influenza epidemic.

Lags in both countries rise over time. In Sweden, the observed lags rise from about 20 years in 1900 to about 50 years in 2000. In the United States, there is a similar rise, but at a slower pace.

There is no simple relationship between gaps and lags and the level of period life expectancy. The lags grow from less than half of life expectancy to more than three-quarters. The lack of relationship between the gap and the level of life expectancy is even more striking, with the gap rising in the early part of the 19th century and falling slowly since.

4 Modeling Gaps and Lags

Having seen how period and cohort life expectancy are empirically related to one another, we turn to a theoretical explanation for the magnitude of the gaps and lags. We adopt a popular model of steady age-specific mortality change widely used for modern industrialized countries. Our strategy is to

approximate cohort life expectancy in terms of current period life table values along with parameters specifying an overall rate and age pattern of mortality decline. Within this model we derive closed-form expressions for gaps and lags. Our purpose is not to predict gaps and lags, which after all can be computed directly from the lifetables themselves. Rather, our expressions are intended to promote structural understanding of the magnitudes of the indices and their determinants.

4.1 An Age-specific Model of Mortality Change

Our model is a log-linear specification for age-specific hazard rates $\mu(x,t)$ at age x and year t in terms of a baseline schedule $\mu(x)$:

$$\mu(x,t) = \mu(x)e^{-kb(x)t} \tag{3}$$

For each age group, it is a model of constant proportional change. The age-schedules of mortality $\mu(x)$ and mortality improvement $kb(x)$ can take any desired form. For identifiability, we constrain the $b(x)$ schedule to have unit mean, making k the mean annual exponential rate of mortality improvement over the age range under consideration.

This model of mortality change is a natural choice. It is essentially the Lee-Carter forecasting model (Lee and Carter 1992) currently in use by the United Nations (United Nations 2004). The linear trend factor kt in (3) is the usual Lee-Carter form for forecasts of central tendency. Full Lee-Carter fore-

casts also have stochastic terms and are fitted to historical data with varying factors $k(t)$ in place of kt. The model we are using can also be expressed using the notation of Vaupel and Romo (2003, p. 202), replacing their $\rho(x)$ with $kb(x)$ and imposing constancy on $\rho(x)$ over time. This constancy is the main assumption. The assumption is not perfect over the timescales of lags and cohort lifetimes, but it is a valuable starting point.

For notational convenience, let $t = 0$ represent the period of current interest. We write $l(x)$ for survivorship in the current period lifetable and include k as an argument in life expectancies, so that $e_0^P(t;k)$ is the period life expectancy at birth at time t given the overall rate of mortality improvement k. Period life expectancy at $t = 0$ is

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

The cohort born λ years earlier experiences cohort hazards

$$\mu(x, x-\lambda) = e^{-kb(x)(x-\lambda)}\,\mu(x). \tag{5}$$

Its cohort life expectancy is

$$e_0^C(-\lambda;k) = \int_0^\omega \exp\left(-\int_0^x \mu(a, a-\lambda)da\right)dx \tag{6}$$

4.2 Approximating cohort life expectancy, lags, and gaps

We obtain cohort life expectancy at $t = -\lambda$ as a function of the parameter k by substituting the cohort hazards given by (5) into (6):

$$e_0^C(-\lambda; k) = \int_0^\omega \exp\left(-\int_0^x e^{-kb(a)(a-\lambda)}\mu(a)da\right)dx \tag{7}$$

Cohort life expectancy can now be approximated for small k by expanding $e_0^C(-\lambda; k)$ in a Taylor series around $k = 0$:

$$\begin{aligned} e_0^C(-\lambda_b; k) &= e_0^C(-\lambda_b; 0) + k \left.\frac{d\, e_0^C}{dk}\right|_{k=0} \dots \\ &= \int_0^\infty l(a)da + k\int_0^\infty (a-\lambda)b(a)e(a)l(a)\mu(a)\, da \dots \end{aligned} \tag{8}$$

All of the life table functions in this approximation refer to the period life table of the reference period, a pleasing feature, because the period life table is usually what we have in hand.

We now use the expansion in (8) to approximate the lag between cohort and period life expectancy. Setting lagged cohort life expectancy equal to period life expectancy at $t = 0$, which equals the first term on the right of (8), we solve for λ:

$$\hat{\lambda}_1 = \frac{k \int_0^\infty ab(a)e(a)l(a)\mu(a)da}{k \int_0^\infty b(a)e(a)l(a)\mu(a)da} \tag{9}$$

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

Our Taylor expansion, with λ set to zero, also leads to a first-order estimate of the gap:

$$\hat{\gamma}_1 = k \int_0^\omega ab(a)e(a)l(a)\mu(a)da. \tag{10}$$

4.3 Interpreting the approximations

Equation (10) for the gap gives the the years of life a cohort gains by being able to live into the future when mortality rates have fallen. Inspection of $\hat{\gamma}_1$ shows that it is in fact a sum of age-specific gains. A useful way of thinking about mortality improvement is as additional years of life to those who would have otherwise died (Vaupel and Yashin 1987). Following this logic, the term $l(a)\mu(a)$ is the age distribution of deaths in the absence of mortality decline, and thus the distribution of ages at which deaths are averted. The term $ab(a)k$ gives the magnitude of mortality improvement at each age, with the a term accounting for the longer time it takes for cohort survivors to reach older ages. Finally, $e(a)$ gives the years of life gained for each improvement in mortality. The product of these three terms gives the life expectancy gain at each age; summing over all ages gives the total gain in life expectancy.

Our approximate formula for the lag matches the geometric interpretation of the triangle of Figure 1. The numerator is an approximation for the gap γ

and the denominator is an exact formula for the "slope", the time derivative of the period life expectancy:

$$\dot{e}_0^P = \left.\frac{\partial e_0^P(t;k)}{\partial t}\right|_{t=0} = k\int_0^\infty b(a)e(a)l(a)\mu(a)\,da. \tag{11}$$

(The same result is Equation 15 of Vaupel and Romo (2003).) In words, the first-order estimate of the lag is the number of years it takes for period life expectancy to catch up to cohort life expectancy, assuming constant improvement in life expectancy at the rate observed in the reference period.

A deeper interpretation of Equation (9) shows the lag as the mean age of mortality improvement of a cohort, the average age at which future years of life are gained. This interpretation helps explain why the lag is not, as one might at first suspect, simply equal to the mean age at death. Translated cohort and period fertility are offset by a quantity equal to the mean age of childbearing, but for mortality, it is the mean age of mortality improvement, not the mean age of mortality, that is relevant. We can write the lag as a mean, such that $\lambda_1 = \int ag(a)da / \int g(a)da$. Here, $g(a) = e(a)l(a)\mu(a)kb(a)$ is the years of life gained at age a by mortality improvement. A simple case helps illustrate this relationship. If all mortality improvements occurred only at one age a', with no change at any other ages, then the life-expectancy of a cohort would be equal to the life-expectancy of the period a' years after the cohort was born. When improvements take place over many ages, then the effects of these improvements need to be averaged over these many ages.

The average age of improvement depends on the age-schedule of mortality change, the age distribution of deaths, and on life expectancy by age. For a given age-schedule of mortality change $b(x)$, lower mortality will increase the lag, because deaths $l(x)\mu(x)$ will be concentrated at higher ages. For a given level of mortality, the older the schedule of mortality improvement $b(x)$ the larger the lag λ.

4.4 The unimportance of k

Intuitively, it makes sense that the increased life expectancy that cohorts receive from future declines in mortality should depend on the pace of mortality change. Indeed, we see that our analytical expression for γ is directly proportional to k.

On the other hand, it is somewhat surprising that the first order approximation of the lag does not depend on the rate of mortality improvement k, which cancels out of the numerator and denominator. [3]

Figure 3 about here

Figure 3 illustrates using simulation how little observed lags vary across a wide range of k, including negative values. Instead, the rate parameter k primarily drives the difference between cohort and period life expectancy at

[3]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, but the limit of the lag as k goes to zero does have a non-zero limit.

a moment in time (the gap), and the slope. The fact that the gap and slope both depend to first order linearly on k allows k to cancel out.

4.5 Accounting for Curvature

From the geometric interpretation, we can see in Figure 1 that the use of the current slope in $\hat{\lambda}_1$ will not be exact if there is curvature in the trajectory of period life expectancy over time. Instead, what is needed is the average slope over the period $t - \lambda$ to t, call it $\bar{\dot{e}}$. A second-order approximation is thus

$$\hat{\lambda}_2 = \frac{\hat{\gamma}_1}{\bar{\dot{e}}} \tag{12}$$

We estimate the average slope, using the first order estimate of the lag as guide, letting $\bar{\dot{e}}$ be the slope $\hat{\lambda}_1/2$ years before the time of interest. In simulations, we use the observed value of the slope at this point. In empirical applications, we approximate the slope at this point linearly as

$$\bar{\dot{e}} \approx \dot{e} - \frac{\hat{\lambda}_1}{2}\ddot{e}, \tag{13}$$

where the time derivative of the slope given in equatin (11) is

$$\ddot{e} = k^2 \int_0^\omega l(x) \left\{ \left[\int_0^x \mu(a)b(a)da \right]^2 - \int_0^x \mu(a)b(a)^2 da \right\} dx. \tag{14}$$

5 Validation

The interpretations of the magnitudes of gaps and lags offered in Section 4 make intuitive sense. But they are based on approximations, not exact formulas. We now ask how well the formulas capture the actual magnitudes. The story we have offered is a simplification. How much of a simplification does it turn out to be?

We investigate the validity of our approximations first with simulations and then with empirical series. It is important to bear in mind that the goal of our formulas is not to predict values of gaps and lags. The values can be computed directly from the observed or projected lifetables from which the quantities in our formulas are estimated. The goal is to understand the structure of gaps and lags and to see how much of that structure is uncovered by our model and our approximations.

The simulations allow assessment of our formulas when the assumptions of the model hold, that is, under conditions of proportional mortality improvement at a rate and with an age pattern that are constant over time. The empirical comparisons test how well the logic of our analytic results continues to apply in the face of observable variations in the course of mortality improvement.

5.1 Simulations

We present simulated trajectories of period and cohort life expectancy, and the associated gaps and lags, for two patterns of age-specific mortality improvement. The"classical" simulation provides a stylized version of historical mortality change to date, with faster improvements at younger ages. The Gompertz simulation offers a simple view of what mortality change could look like in the future. Childhood mortality is taken to be negligible and proportional changes in old-age mortality are taken to be uniform over time and age.

For the "classical" simulation we follow Canudas-Romo and Schoen (2005), employing a Siler model as the baseline schedule of mortality. The Siler model adds an exponentially declining hazard, dominating childhood, and an exponentially rising hazard, dominating adulthood, along with a constant background component. We calibrate the baseline schedule to approximate the period mortality rates of Swedish females in 2000. Mortality improvement is modeled with age-specific coefficients $b(x)$, constant over time, which drop linearly with age between ages 0 and 20 and remain steady above age 20. The rate k of mortality decline is chosen to be consistent with long-term drops in Swedish female period life-expectancy from 47.3 years in 1850 to 82.2 years in 2000, and onward, according to U.N. forecasts, to 106.5 years in 2300.[4]

[4]The estimated Siler baseline hazard is $\mu(x) = \alpha_1 \exp(-\beta_1 x) + \alpha_2 \exp(+\beta_2 x) + \alpha_3$, with $\alpha_1 = 0.003$, $\alpha_2 = 0.00001428$, $\alpha_3 = 0.002$, and $\beta_1 = 1$, and $\beta_2 = 0.100$. The estimated pattern of mortality change is given by of $k = 0.0092$, with $b(0) = 3.09$ and $b(20) = 0.84$.

For the Gompertz simulation, we use the adult component of mortality from the Siler model on its own.

The simulated life expectancies, gaps, and lags are shown in Figure 4. In the classical simulation, between 1850 and 2150 period life expectancy rises from about 47 years to 95 years. Cohort life expectancy rises from about 51 years to 102 years. Simulated gaps rise from 4.0 to 7.1 years, and lags lengthen from 7.6 to 90.3 years.

We see that $\hat{\gamma}_1$ is an accurate estimate of the actual gap, with the maximum error being around half a year. The first-order estimate of the lag $\hat{\lambda}_1$ tracks the overall magnitude of the increase in the actual lag but is off by as much as 8.9 years for intermediate values of life expectancy, when the curvature in $e_0^P(t)$ is considerable. The second order estimate of the lag $\hat{\lambda}_2$ accounts for this curvature nearly completely.

In the Gompertz simulation, cohort and period life expectancy become very nearly linear with time, as do gaps and lags. The first order estimates for the gap are quite accurate. They are systematically slightly low by about 0.6 years. The first order estimates of the lag are also accurate, but they are systematically slightly high by about 1 year. These outcomes can be understood in more detail by an approach to be described in Section 6.

Overall, the "classical" simulation shows a setting in which the first order estimate for the gap is highly accurate, the first order estimate for the lag is only approximate, and the second order estimate for the lag based on time derivatives is highly accurate. The Gompertz simulation, which is relevant

for the gaps and lags of life expectancy at higher ages and perhaps for future mortality decline, shows a setting in which the first order for both gaps and lags capture the trends with high precision.

5.2 Observations

We now turn to the trajectories of life expectancy observed and forecast in Sweden and the United States. Our estimated gaps and lags are shown in Figure 5 superimposed on the observed trajectories taken from Figure 2. The estimates are computed from Equations (9), (10), and (13), with the gap and lag for each period estimated from the period lifetable values for that period. Consistent with our model, we estimate patterns of mortality improvement with k and $b(x)$ constant over time, matched to the overall change across the period of observation [5]

We see from the graphs that the estimates of gaps based on unchanging values of k and b succeed in capturing the overall level of the gaps but miss the temporal variation. The estimated gaps represent a highly smoothed version of the observed gaps for both countries. The estimated current value of $\hat{\gamma}_1$ in 2000 of about 4 years is a plausible one and it is roughly consistent with the simulations.

The estimated lags in both Sweden and the United States track observed lags well until about World War II, but afterwards overshoot the observed

[5]Specifically, we let $\rho(x) = -\left[\log {}_1M_x(t_2) - \log {}_1M_x(t_1)\right]/(t_2 - t_1)$ and set k to the mean over all observed ages of $\rho(x)$ and $b(x) = \rho(x)/k$, where $t_1 = 1900$ for both Sweden and the United States and $t_2 = 2050$ for Sweden and $t_2 = 2080$ for the United States.

values. Using $\hat{\lambda}_2$ to account for the convexity of the life expectancy trajectories makes the estimated lags accurate until about 2000 in Sweden. In the United States, $\hat{\lambda}_2$ is an improvement over $\hat{\lambda}_1$ but still does not capture the slowdown in mortality improvement built into the Social Security forecasts that we have already mentioned. This hypothesized slowdown is controversial. The slowdown projected for Sweden is less dramatic.

Comparison between the fits to the simulations and the fits to the empirical series suggest that variations over time in the rate and age pattern of mortality improvement in the empirical series are responsible for the deviations between estimates and observations. Although the assumption of constant k and $b(x)$ is widely adopted and underlies most applications of Lee-Carter forecasting, it is not fully tenable for these long-term series (Lee and Miller 2001).

We could compute a time-varying series of estimates for k and $b(x)$ for use with our formulas, but complications arise. For the gap, future values of $kb(x)$ are relevant, for they determine the life experience of the cohort just born. For the slope, past values of $kb(x)$ are relevant, as we work back in time to the lagged value of cohort life expectancy. For earlier epochs, changes in $b(x)$ are more pronounced, and at later epochs changes in k stand out. If accurate prediction were our goal, more elaborate formulas might be justified. But for the purpose of structural understanding, little is to be gained from complicated expressions. Our approximations do well enough in practice to support the interpretations offered in Section 4.

6 Linear Shift Models

Curvature in graphs of life expectancies at birth is largely induced by reductions in infant and child mortality. Our Gompertz simulations, in which infant mortality is set to negligible levels, produce graphs that are quite straight. More generally, when we restrict attention to adult mortality, trajectories of further life expectancy beyond some age like 30 are typically quite nearly linear. Such linearity can be represented by fitting variants of our proportional change model from Section 4 into the framework of the Linear Shift Model. This special case from a family of models developed by Bongaarts and Feeney (2003) has been studied by Rodriguez (2006) and Goldstein (2006).

6.1 Linear Shifts

Under the Linear Shift Model, the hazard rate at every age x is given by the hazard rate at a younger age $x - rt$, t years before:

$$\mu(x,t) = \mu(x - rt, 0), \tag{15}$$

This model applies to some stretch of years before and after some reference time $t = 0$ and it only applies to adult ages. The convention is to let age $x = 0$ in the formula correspond to some human age like 30, and to condition on survival to that starting age, implicitly setting younger hazard rates (for negative x values) to zero. So far, no assumptions about the shape of the

baseline hazard are being made.

Because the whole period mortality schedule is shifting to higher ages at a constant rate r, period life expectancy under the model increases exactly linearly with slope r. The same is true of the measure CAL, the Cross-Sectional Average Length of Life, described in the Introduction. The value of CAL at time t depends on prior values of the hazards, often as a moving average of prior values of period life expectancy (Wachter 2005). Under the Linear Shift Model, it is the same as lagged period life expectancy, with a lag denoted by $G(r)$:

$$CAL(t) = e_0^p(t - G(r)) \tag{16}$$

This $G(r)$ is constant over time, since $e_0^p(t)$ and $CAL(t)$ have the same slope with respect to time. $G(r)$ depends not only on r but on the underlying hazard schedule.

It is a remarkable fact, proved by Goldstein (2006), that the Linear Shift Model implies

$$e_0^C(t - CAL(t)) = CAL(t) \tag{17}$$

The measure CAL lags behind cohort life-expectancy by a number of years equal to CAL itself.

Manipulation of equations (16) and (17), keeping in mind that both period life expectancy and CAL are increasing at a constant rate r, leads to a

full set of formulas for gaps and lags under Linear Shifts.

$$\lambda(t) = e_0^p(0) - G(r) + rt \tag{18}$$

$$\gamma(t) = \frac{r}{1-r}\lambda(t) = \gamma(0) + \frac{r^2}{1-r}\,t \tag{19}$$

The Linear Shift Model thus has the following features:

i) The lag comes out to be linear with the same slope as period life expectancy.

ii) The lag comes out exactly equal to the ratio of the gap to the slope of cohort life expectancy. In the absence of curvature, the cohort slope rather than the period slope picks up the relevant hypotenuse of Figure 1.

iii) If we approximate the lag by a first-order Taylor expansion in r, repeating our approach from Section 4, we obtain an estimate that differs from the exact lag by an amount which remains constant over time.

iv) Under the model, cohort slope exceeds period slope when r is positive.

6.2 Shifting Gompertz Hazards

Our results for the Linear Shift Model are exact, and they hold regardless of the form of the baseline hazard $\mu(x, 0)$. For numerical calculations, however, we need to specify the baseline hazard in order to fix $G(r)$ and $e_0^P(0)$.

One alternative is to let the baseline hazard be an exponential Gompertz function. Combining this specification with the Linear Shift Model, however, leads to hazards for $t > 0$ that are not fully Gompertzian, since zero hazards for negative ages x get shifted into positive ages as time goes on. Formulas for G are feasible but unwieldy. This alternative preserves the Linear Shift Model but relaxes our proportional change assumptions.

A second alternative, which we used in the simulations above, is to assume Gompertz hazards at all times at all ages over age $x = 0$, that is to say, $\mu(x,t) = \alpha e^{-kt} e^{\beta x}$. Such hazards have the form assumed in Section 4 with the $b(a)$ all set equal to 1. With this specification, when α is small, the Linear Shift Model holds approximately, but only approximately, with a shift parameter $r = k/\beta$. We maintain our proportional change model, but relax the strict conditions on shifts.

With this alternative, our first-order estimates of gap and lag take the following forms:

$$\hat{\gamma}_1 = k \int ae(a)l(a)h(a)da = \frac{k}{\beta}\left(e_0^P(t) - \frac{1}{\beta}(1 - \alpha e_0^P(t))\right). \tag{20}$$

$$\hat{\lambda}_1 = \frac{e_0^P}{1 - \alpha e_0^P} - \frac{1}{\beta} \tag{21}$$

Often in developed societies α is small enough for αe_0^P to be neglected when as $x = 0$ corresponds to some young adult age like 30. Then Equations (20) and (21) agree with the Linear Shift Equations (18) with $G(r) \approx 1/\beta$. For

present-day Sweden or the U.S.A., β is of order 10^{-1} and k is of order 10^{-2}, r of order 10^{-1}, and $G(r)$ of order 10^{1}.

A third alternative yields exact expressions at the cost of some artificiality. We posit Gompertz hazards at all negative and positive values of x, evaluating the mean age at death of the extended Gompertz distribution on the age interval $(-\infty, \infty)$. This device allows closed-form solutions, and the sacrifice in realism is not as great as it might seem. When $x = 0$ stands for an adult age like 30 years, small negative x values stand for sensible human ages, and large negative values make only the tiniest of contributions. Formally speaking, we have a single model which simultaneously satisfies the assumptions of proportional change and of linear shifts.

With this setup, as Pollard and Valkovics (1992) pointed out, the mean age at death is given by

$$e_0^P(0) = \frac{1}{\beta}\left[\log\left(\frac{\beta}{\alpha}\right) - \gamma^E\right]. \tag{22}$$

Here γ^E is Euler's Constant $0.577215\ldots$, produced by integrating the exponential integral (Abramowitz and Stegun 1964). The function $G(r)$ has an exact expression which is a little greater than $1/\beta$ when r is positive:

$$G(r) = \frac{1}{\beta}\log\left(\frac{1}{1-r}\right) = \frac{1}{\beta r}\left(1 + \frac{r}{2} + \frac{r^2}{3} + \ldots\right). \tag{23}$$

Augmented by these formulas, the equations for the Linear Shift Model give easy numerical predictions. With $\beta \approx r \approx 0.10$, $G(r) \approx 10.5$. The lag

stays about 10.5 years less than period life expectancy (that is, than further life expectancy conditional on survival to the adult age chosen to correspond to $x = 0$). The gap is about 1/9 as big as the lag, growing at no more than about 0.01 per year. The CAL measure hovers about 1 year below period life expectancy, and both of them grow by about a year every decade, while cohort life expectancy grows by about a year every nine years. If present trends in countries like Sweden or the U.S.A. continues, the future story of lags and gaps for adult life expectancies could well resemble this scenario.

6.3 Lags and gaps for e_{65}

The approach just described can also be deployed to study remaining life expectancy at an older age such as an age at retirement. For this purpose, the easiest formulas are the first-order Gompertz approximations (20) and (21). For example, for age 65, we set $\alpha = \mu(65)$, keeping in mind that αe_0^P is no longer vanishingly small. To illustrate using values that are close to contemporary conditions in low morality societies, let the annual exponential rate of mortality decline be $k = 0.01$ and the current period Gompertz schedule be determined by $\beta = 0.10$ and $\alpha = 0.01$. These estimates yield $e_{65} = 20.5$. Simulating, we find that the exact value of γ is 1.31, versus 1.26 for $\hat{\gamma}_1$. The exact value of the lag λ is 15.0 years, versus 15.9 for $\hat{\lambda}_1$. Period life expectancy at age 65 in current lifetables actually corresponds to the life expectancy of the cohort that reached age 65 some 15 years ago. Those now aged 65 can expect to live about one-year longer than the current period life

table would indicate, adding around five percentage points to pension costs.

7 Conclusion

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 by which the equivalent period and cohort life expectancies are observed and (2) the gain in life expectancy γ that a cohort benefits from by experiencing improving rather than fixed mortality.

We use a simple model of mortality decline to show that the lag between periods and cohorts is not equal, as might have been suspected, to the mean age of death, but rather to the mean age at which mortality improvement is occurring.

We find that as mortality has fallen, the lag between periods and cohorts has increased. This is largely a function of the greater ages at which deaths are occurring, and thus the greater ages at which mortality improvement is effectively taking place. Period mortality is in this sense becoming a more and more outdated measure of the experience of cohorts.

We find, on the other hand, that the absolute difference between period and cohort life expectancy, the gap, has risen and then fallen over time. Thus, while period mortality has become more "out of date", its divergence

from cohort mortality has shrunken. The reason for this paradox is that the pace of change in period mortality has itself flattened. It takes more years to cover less ground.

We find in our model that the pace of mortality decline plays an important role in determining the magnitude of the gap, but to a first approximation it plays no role in determining the size of the lag. This is because the change in gap is, to a first approximation, exactly offset by a change in the slope. However, the analysis of historical and forecast mortality in the United States and Sweden showed that changes in the pace of mortality decline can cause actual lags to depart from our first-order, and even second-order, approximations.

We have made some progress in translating period and cohort life expectancy. Ryder was correct that the translation of mortality measures would be more difficult than fertility. However, we have seen that interpretable analytical expressions can be derived from a model of steady mortality decline and that these approximations perform well even with data which is not perfectly in accord with the model.

The main import of our results is that in populations experiencing steady mortality decline, period life expectancy can be regarded as a lagged measure of cohort life expectancy. The lag lengthens gradually with declines in mortality. For countries with mortality rates close to those of the present-day United States, today's period life expectancy at birth summarizes the expected longevity of people born about 40 to 50 years ago who are or would be now in the prime of middle age. Current period life expectancy at age 65

matches expected survival beyond 65 for people who celebrated their 65th birthday some 15 years ago. Period life expectancies, defined abstractly through synthetic cohorts, can be associated with the experience of real cohorts. The correspondence makes them more tangible and easier to grasp.

References

Abramowitz, M. and Stegun, I. (eds): 1964, *Handbook of Mathematical Functions*, Applied Mathematics Series 55, National Bureau of Standards, U.S. Department of Commerce, Washington, D.C.

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, *PNAS* **100**(23), 13127–13133.

Brouard, N.: 1986, Structure et dynamique des populations, *Espaces, Populations, Sociétés* **2**(14-15), 157–168.

Canudas-Romo, V. and Schoen, R.: 2005, Age-specific contributions to changes in the period and cohort life expectancy, *Demographic Research* **13**(3), 63–82.

Goldstein, J. R.: 2006, Found in translation? A cohort perspective on tempo-adjusted life expectancy, *Demographic Research* **14**(5), 71–84.

Guillot, M.: 2003, The cross-sectional average length of life (CAL), *Population Studies* **1**(1), 41–54.

Keyfitz, N. and Caswell, H.: 2005, *Applied Mathematical Demography, Third Edition*, Springer Verlag, New York.

Lee, R. D. and Carter, L.: 1992, Modeling and forecasting U.S. mortality, *JASA* **87**(419), 659–671.

Lee, R. D. and Miller, T.: 2001, Evaluating the performance of the Lee-Carter method for forecasting mortality, *Demography* **38**(4), 537–549.

Pollard, J. H. and Valkovics, E. J.: 1992, The Gompertz distribution and its applications, *Genus* **48**(3-4), 15–28.

Rodriguez, G.: 2006, Demographic translation and tempo effects: An accelerated failure time perspective, *Demographic Research* **14**(6), 85–110.

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

Schoen, R. and Canudas-Romo, V.: 2005, Changing mortality and average cohort life expectancy, *Demographic Research* **13**(5), 117–142.

Technical Panel on Assumptions and Methods: 2003, Report to the Social Security Advisory Board, *Technical report*, Social Security Advisory Board, Washington, D.C.

Tuljapurkar, S., Li, N. and Boe, C.: 2000, A universal pattern of mortality decline in the G7 countries, *Nature* **405**, 789–792.

United Nations: 2004, *World Population to 2300*, United Nations, Department of Economic and Social Affairs, Population Division.

Vaupel, J. W. and Romo, V. C.: 2003, Decomposing change in life expectancy: A bouquet of formulas in honor of Nathan Keyfitzs 90th birthday, *Demography* **40**(2), 201–216.

Vaupel, J. W. and Yashin, A. I.: 1987, Repeated resuscitation: How lifesaving alters life tables, *Demography* **24**(1), 123–135.

Wachter, K. W.: 2005, Tempo and its tribulations, *Demographic Research* **13**(9), 201–222.

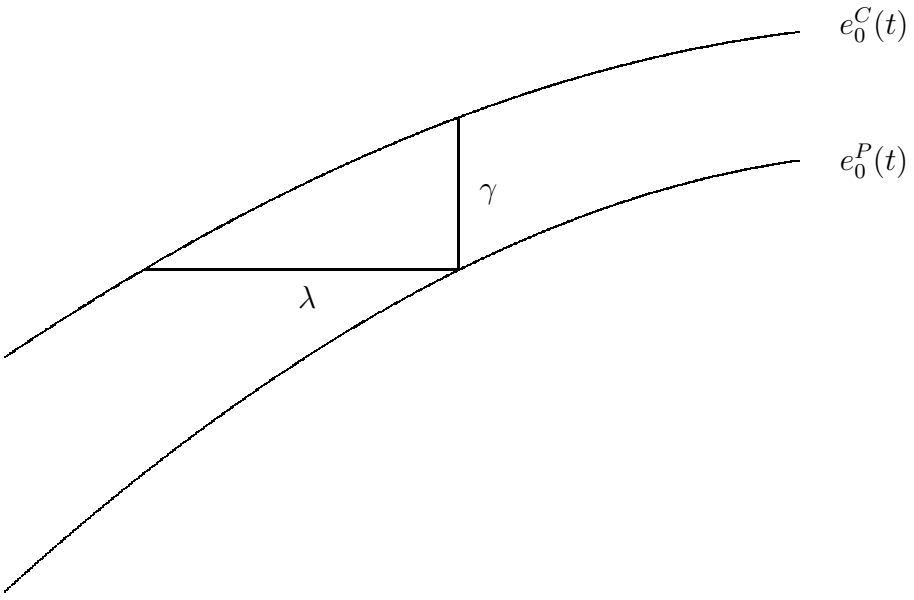

Figure 1: Sketch of gap (γ) and lag (λ) between cohort and period life expectancy when mortality and entropy are decreasing with time.

Figure 2: Observed and forecast cohort and period female life expectancy at birth in Sweden and the United States and accompanying gaps and lags.

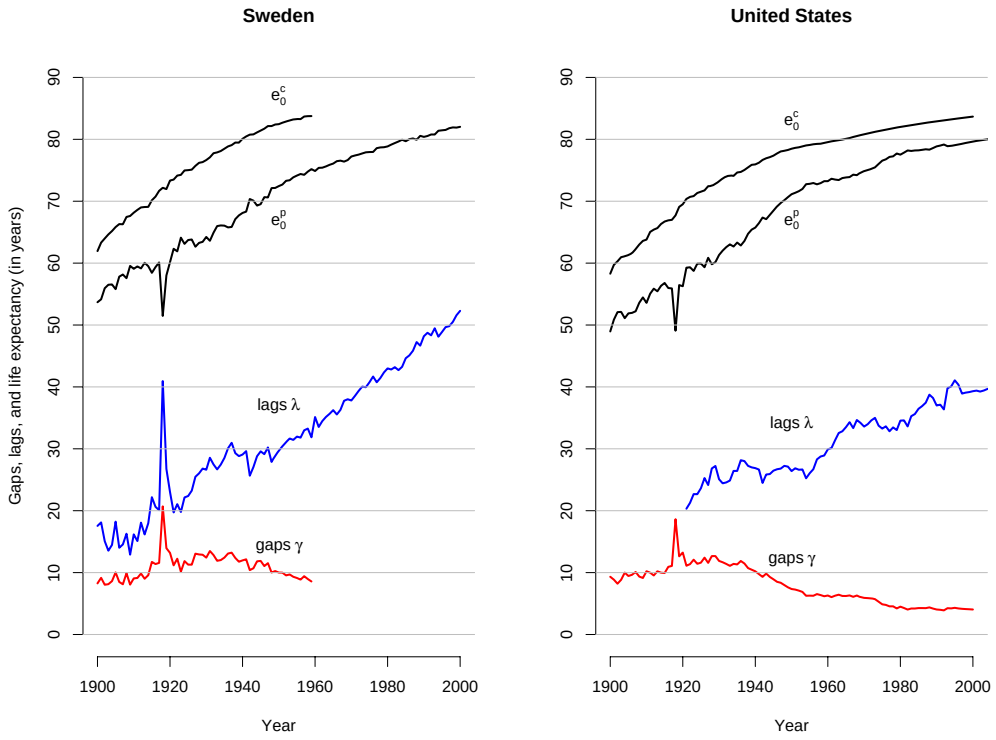

Sources: U.S. Social Security Administration available at Berkeley Mortality Database (www.demog.berkeley.edu/∼bmd); Human Mortality Database (www.mortality.org); Statistics Sweden. Early Swedish lags are based on 19th century e_0^C (not shown).

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

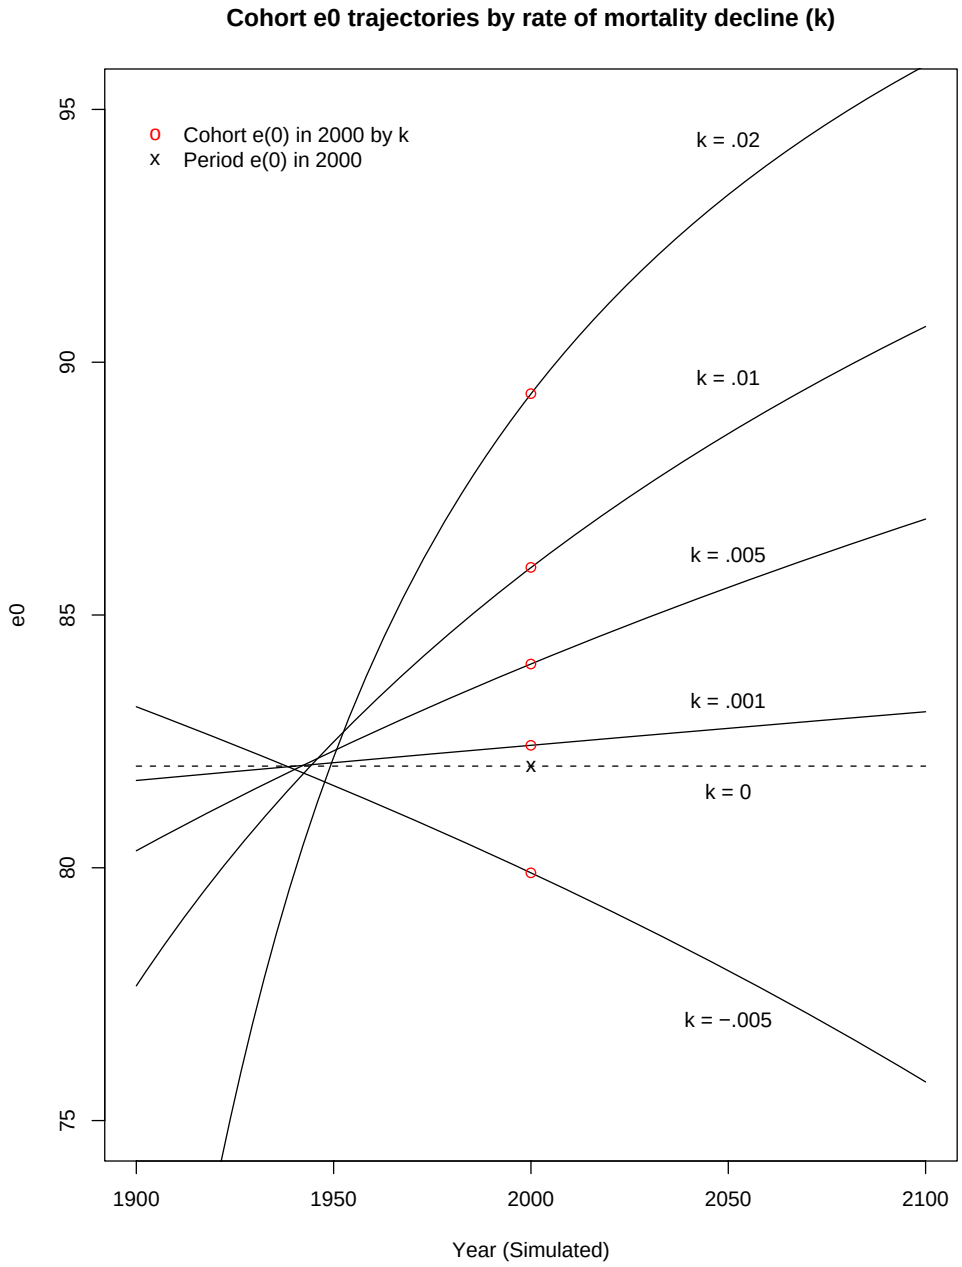

Dashed line gives period e(0) in 2000. Lag can be seen as the horizontal distance between "x" and the intersection of the cohort curve with the horizontal dashed line. The gap is the vertical distance between "x" and "o". The figure shows that the lag is roughly constant over a wide range of mortality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

Figure 4: Simulated trajectories for period and cohort life expectancy with accompanying observed and approximated gaps and lags

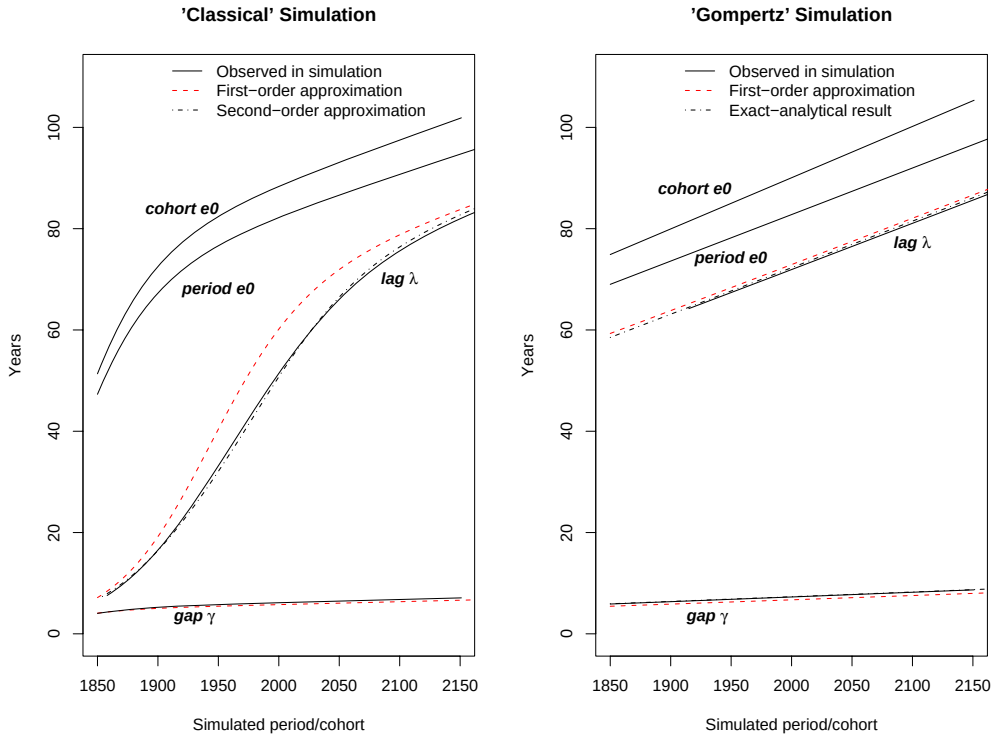

Classical scenario uses Siler baseline mortality for year 2000, with mortality change calibrated to period $e_0(1850) = 47.3$ years, $e_0(2000) = 82.2$ years, and $e_0(2300) = 106.5$ years. Mortality rate decline is faster at youngest ages. Gompertz simulation based on Gompertz baseline $\mu(x)$ and uniform pattern of mortality decline. See text for details.

Figure 5: Comparison of observed gaps and lags for females in Sweden and the United States with analytical estimates

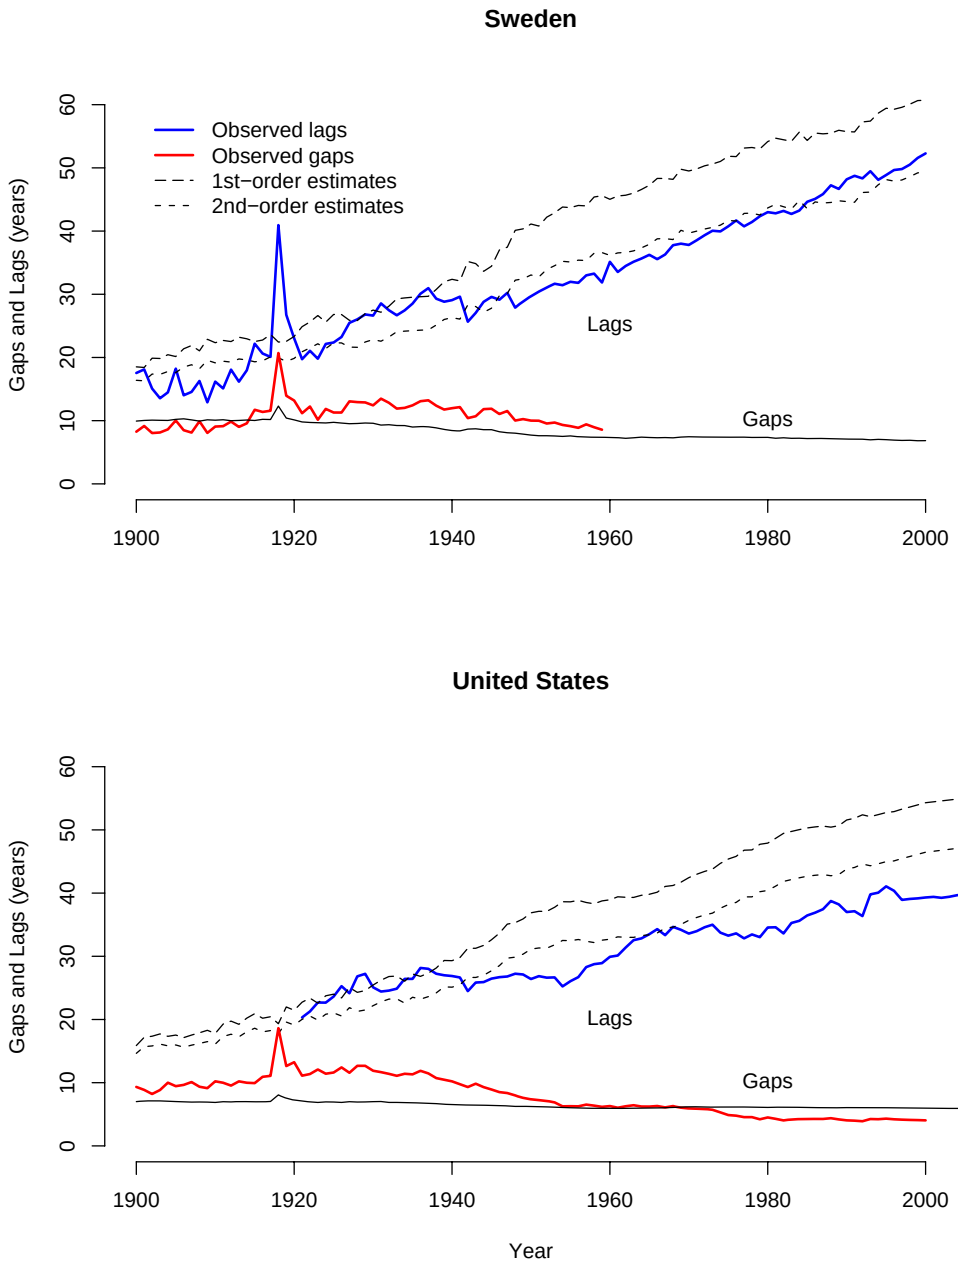

Analytical estimates made assuming constant age-pattern and level of mortality decline. See text for details.