Bayesian Probabilistic Projections of Life Expectancy for All Countries [1]

Jennifer L. Chunn, Adrian E. Raftery, and Patrick Gerland

Working Paper no. 105
Center for Statistics and the Social Sciences
University of Washington.

October 28, 2010

[1]Jennifer L. Chunn is research assistant, Department of Statistics, National University of Singapore; email: jenchunn@gmail.com. Adrian E. Raftery is Professor of Statistics and Sociology, University of Washington, Box 354322, Seattle, WA 98195-4322; e-mail: raftery@u.washington.edu. Patrick Gerland is s a Population Affairs Officer with the United Nations Department of Economic and Social Affairs, Population Division, New York, NY 10017; e-mail: gerland@un.org. This research was supported by NICHD grant R01 HD54511. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the National Institute of Child Health and Human Development. Also, the views expressed in this paper are those of the authors and do not necessarily reflect the views of the United Nations. Its contents have not been formally edited and cleared by the United Nations. The designations employed and the presentation of material in this paper do not imply the expression of any opinion whatsoever on the part of the United Nations concerning the legal status of any country, territory, city or area or of its authorities, or concerning the delimitation of its frontiers or boundaries. The authors are grateful to Leontine Alkema, John Bongaarts, Thomas Buettner, Samuel Clark, Joel Cohen, Gerhard Heilig, Ronald Lee, Nan Li, François Pelletier, Kenneth Wachter and Hania Zlotnik for helpful comments and discussions.

Abstract

We propose a Bayesian hierarchical model for producing probabilistic forecasts of male life expectancy at birth for all the countries of the world from the present to 2100. Such forecasts would be an input to the production of probabilistic population projections for all countries, which is currently being considered by the United Nations. To evaluate the method, we did an out-of-sample cross-validation experiment, fitting the model to the data from 1950–1995, and using the estimated model to forecast for the subsequent ten years. The ten-year predictions had a mean absolute error of about 1 year, about 40% less than the current UN methodology. The probabilistic forecasts were calibrated, in the sense that (for example) the 80% prediction intervals contained the truth about 80% of the time. We illustrate our method with results from Madagascar (a typical country with steadily improving life expectancy), Latvia (a country that has had a mortality crisis), and Japan (a leading country). We also show aggregated results for South Asia, a region with eight countries.

Keywords: Bayesian hierarchical model; Double logistic function; Lee-Carter model; Markov chain Monte Carlo; Mortality; Population projection.

List of Figures

1 Observed five-year gains in life expectancy 4
2 Illustration of the double-logistic function. 5
3 Absolute residuals versus life expectancy with fitted regression spline. 7
4 Life expectancy projections for Madagascar. 10
5 Life expectancy projections for Latvia. 11
6 Life expectancy projections for Japan. 13
7 Life expectancy projections for South Asia for the Bayesian hierarchical model, IIASA and the UN. 16

List of Tables

1 Out of sample cross-validation results for the Bayesian hierarchical model and the current UN methodology. 9

Every two years, the United Nations Population Division (UN) publishes an updated edition of the *World Population Prospects* (WPP), which includes projections of the populations of all the countries in the world to 2050, broken down by age and sex. The UN issues several variants of its population projections, of which the most important is the Medium variant, viewed as the best projection of future population trends. It also issues High and Low variants, obtained by increasing and decreasing the total fertility rate (TFR) by half a child, respectively. The High and Low variants are scenarios designed to indicate what would happen if the assumptions underlying the Medium variant were violated in various ways; they are not probabilistic projections. The UN does not issue variants indicating the likely impact of differences in future mortality rates. Largely because of the need to project climate change over the next century, the UN is considering extending its population projections to 2100.

Fully probabilistic population projections are an alternative to scenarios, which may be preferable because, unlike scenarios, they indicate the likely range of future population outcomes (Bongaarts & Bulatao, 2000). To produce probabilistic population projections for all the world's countries, one needs probabilistic projections of the main demographic processes affecting national populations: fertility, mortality and international migration.

This paper is part of a research program aiming to do this. The UN produces many of its current (deterministic) projections by projecting broad summaries of population processes, and then breaking them down into age-specific rates using model schedules, to yield the age- and sex- specific fertility, mortality and migration rates that are required by the standard cohort-component population projection method. Our research therefore focuses on probabilistic projection of summary population measures. Fully probabilistic projections of the total fertility rate have already been produced (Alkema et al., 2008; Raftery et al., 2009; Alkema et al., 2011).

The best known approach to probabilistic projection of mortality is that of Lee & Carter (1992), which uses past age-specific death rates for each year over time in a country. They used a log-linear model for the age-specific mortality rate at age x in year t, $\log(m_{xt}) = a_x + k_t b_x + \varepsilon_{xt}$, where k_t is the mortality index. They project k_t into the future using a random walk model with constant drift. The Lee-Carter method has been found to perform well for some developed countries (Booth et al., 2005; Bell, 1997). When forecasting a group of countries simultaneously, a common age parameter is fixed to ensure consistent forecasts of multiple countries (Li & Lee, 2005).

However, the Lee-Carter method requires that age-specific death rates be available for at least three time periods, which is not the case for many developing countries (Li et al., 2004). Also, while the Lee-Carter model specifies changes in age-specific mortality rates to

have a constant distribution on the logarithmic scale, White (2002) found that linear changes in life expectancy gave a better fit to data from 21 industrialized countries. Further, Lee & Miller (2001) found that the forecasts of the Lee-Carter method for the United States, Japan, Canada, Sweden, and France systematically underestimated future life expectancy.

Lutz and colleagues at the International Institute for Applied Systems Analysis (IIASA) have taken a different approach, producing probabilistic projections from expert opinion. They addressed data limitation by aggregating countries into regions and forecasting regional life expectancy based on expert-based probabilistic projections (Lutz et al., 2004).

We propose an approach that is adapted to the UN's task of projecting populations for all the countries in the world, many of which have data that are patchy and of variable quality. We project life expectancy directly. We do this using a random walk model with a nonconstant drift. The drift term is a nonlinear function of current life expectancy and reflects the fact that life expectancy tends to improve more slowly for the countries with the lowest and highest life expectancies, and more quickly for the countries in the middle. Also, the overall rate of improvement varies by country. We use a Bayesian hierarchical model (BHM), which allows us to estimate the rate of improvement in life expectancy for a country using past data from that country, and also taking account of the observed past patterns in all other countries.

We develop a one-sex model for males and discuss potential extensions to a two-sex model in the closing discussion section. We use the estimates of historical male life expectancy at birth from the UN *World Population Prospects* (WPP) 2008 Revision from 1950 through 2010 (United Nations, 2009). Because of the significant impact of the HIV/AIDS epidemic on mortality rates, we do not include countries with a generalized HIV/AIDS epidemic in this analysis. We base our results on data from 158 countries.

This paper is organized as follows. We first discuss the UN's current projection methodology. We then describe our proposed model, which is a natural extension of the UN's current practices from deterministic to probabilistic projections. We assess our model by holding back the last ten years of data (1995–2005), reestimating our model using only the data from 1950 to 1995, generating probabilistic projections for 1995 to 2005, and comparing them with what was actually observed.

We then present probabilistic projections of life expectancy to 2100 in three countries with widely different current life expectancies, Madagascar, Latvia, and Japan, each of which presents different forecasting challenges. Life expectancy in Madagascar has been improving steadily but is still in the lowest quartile. Latvia has experienced a mortality crisis in the past generation, with both declining and increasing life expectancy. Japan is a leading country, with one of the highest life expectancies. For each country, we present

out-of-sample projections for 1995-2005. Lastly, we discuss aggregation of country-specific projections. Comparisons are made with the recently updated regional projections for South Asia by IIASA.

METHODOLOGY

Current UN Population Projection Methodology

Currently, the UN projects life expectancy at birth deterministically. The life expectancy at birth, $l_{c,t+1}$ for country, c, in the next five-year period, $t+1$, is projected to be the life expectancy in the current time period, $l_{c,t}$, plus the expected gains in life expectancy, $g(l_{c,t})$. Observed five-year gains in life expectancy for 158 countries from 1950 to 2005 are plotted in Figure 1. This figure highlights the non-constant rate of change in life expectancy. To capture this, the UN has developed models that represent the gains in life expectancy by a double-logistic function of current life expectancy. The five deterministic UN models are shown in Figure 1(b), where models vary by pace of gains in life expectancy.

The double-logistic function (Meyer, 1994) has six parameters, as illustrated in Figure 2. Four of them identify intervals of life expectancy when the rate of life expectancy gains is changing, one describes the approximate maximum gain in life expectancy, and the last parameter gives the asymptotic rate of gains as life expectancy increases. For each country, a UN analyst chooses one of five prescribed choices of the six parameters[1] by assessing the recently observed pace of mortality decline (United Nations, 2009). The model implies that after life expectancy reaches a certain threshold, life expectancy increases at a constant rate. This is consistent with research indicating that there is no evidence of an upper limit to life expectancy (Oeppen & Vaupel, 2002).

The double-logistic function is the sum of two three-parameter logistic growth functions. Demographic transition theory suggests that life expectancy improves slowly at first, and then faster as a country enters the demographic transition. The early rapid gains are largely a result of improvements in infant and child mortality. Once most of the possible gains from reductions in child mortality have been realized, increases slow as mortality improvements shift to older ages.

To summarize, the UN projects life expectancy in the next time period deterministically using the equation

$$l_{c,t+1} = l_{c,t} + g(l_{c,t}). \tag{1}$$

[1]Using data from different percentiles, the UN developed five models for each sex to describe different rates of increase in life expectancy: very fast (based on the upper 10^{th} percentile of countries), fast (upper 25^{th} percentile), medium (based on median), slow (lower 25^{th}), and very slow (lower 10^{th} percentile).

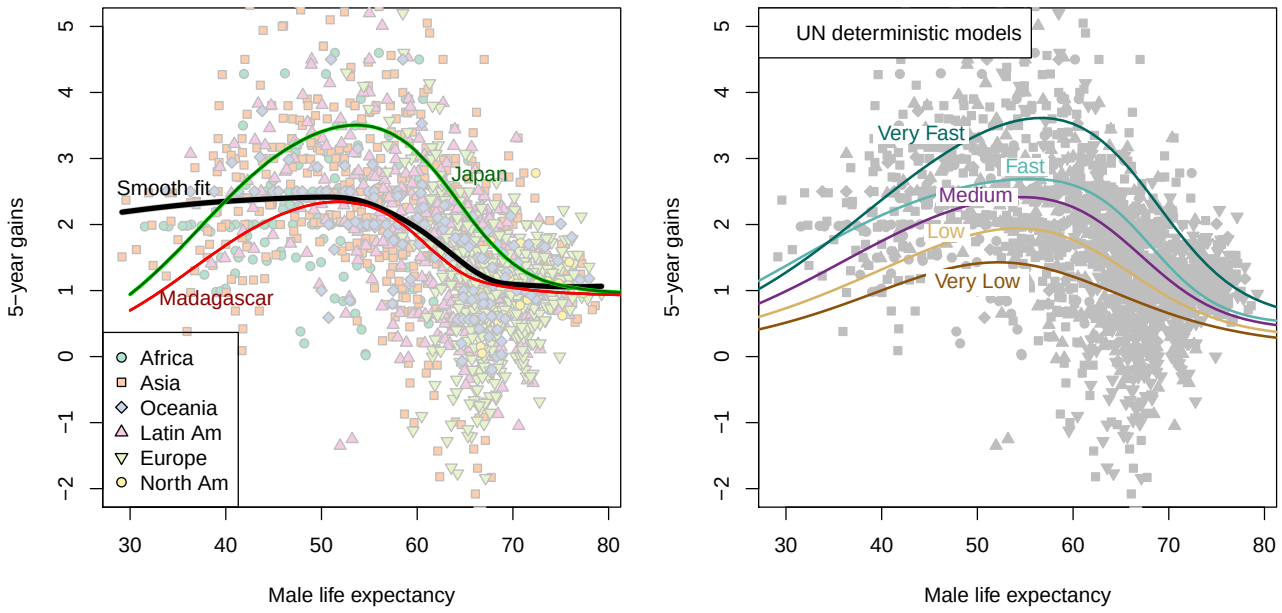

Figure 1: Observed five-year gains in life expectancy, plotted against the life expectancy at the beginning of the five-year period. UN estimates for 158 countries from 1950 to 2005 are included in this figure ($n = 1738$). Each point represents an observed five-year gain in life expectancy within a country. The black line is a locally-weighted polynomial (lowess) regression of the observations, which highlights the non-constant rate of gains in life expectancy. Included in the left plot are the fitted posterior median double-logistic functions for Japan and Madagascar from our model. The UN deterministic models are included in the right plot. (Note: 31 observations (1.8%) are outside the range of the plot and not shown, but were included in the local regression.)

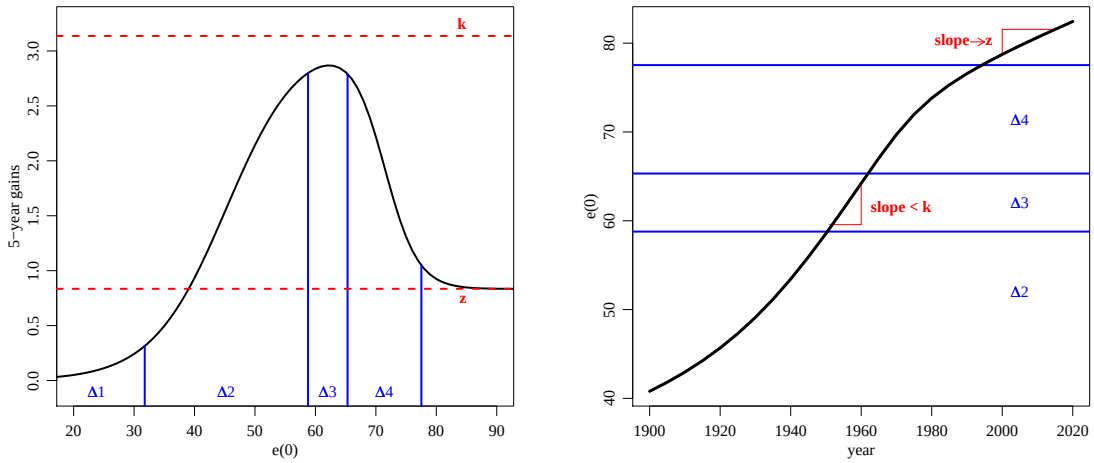

Figure 2: Illustration of the double-logistic function, based on a curve from the posterior distribution for Japan. The left plot illustrates the double-logistic function of 5-year gains in life expectancy. The right plot is a time-series of life expectancy with gains modeled according to the double-logistic function

The expected five-year gain in life expectancy is a function of the current level of life expectancy, namely

$$
\begin{aligned}
g(l_{c,t}|\boldsymbol{\theta}^c) &= \frac{k^c}{1+exp(-\frac{log(9^2)}{\Delta_2^c}(l_{ct}-\Delta_1^c-0.5\Delta_2^c))} + \frac{z^c-k^c}{1+exp(-\frac{log(9^2)}{\Delta_4^c}(l_{ct}-\sum_i \Delta_i^c+0.5\Delta_4^c))}, \\
\boldsymbol{\theta}^c &\in (\boldsymbol{\theta}^{\text{Very Slow}}, \boldsymbol{\theta}^{\text{Slow}}, \boldsymbol{\theta}^{\text{Medium}}, \boldsymbol{\theta}^{\text{Fast}}, \boldsymbol{\theta}^{\text{Very Fast}}), \\
\boldsymbol{\theta}^c &= (\Delta_1^c, \Delta_2^c, \Delta_3^c, \Delta_4^c, k^c, z^c).
\end{aligned}
$$

The parameters $\boldsymbol{\theta}^c$ for country c are chosen by a UN analyst from the five possibilities.

Stochastic Model

The UN projection method is deterministic and does not account for uncertainty. We now extend it to a stochastic model to allow for uncertainty. This involves two extensions. The first allows for stochastic changes within a country by replacing the deterministic model (1) by a stochastic one by adding a random perturbation to (1). It then becomes a random walk with drift, where the drift term is given by the double-logistic function.

The second extension is to allow the parameters of the double-logistic function to vary between countries over a continuous range rather than among the current five UN possibilities.

The resulting hierarchical model is

$$l_{c,t+1} = l_{c,t} + g(l_{c,t}|\boldsymbol{\theta}^{(c)}) + \varepsilon_{c,t+1}, \tag{2}$$

where

$$\begin{aligned}
g(l_{c,t}|\boldsymbol{\theta}^c) &= \text{Double-Logistic function with parameters } \boldsymbol{\theta}^c, \\
\boldsymbol{\theta}^c &= (\Delta_1^c, \Delta_2^c, \Delta_3^c, \Delta_4^c, k^c, z^c), \\
\Delta_i^c|\sigma_{\Delta_i} &\overset{\text{iid}}{\sim} \text{Normal}_{[0,100]}(\Delta_i, \sigma^2_{\Delta_i}), \quad i = 1, \ldots, 4, \\
k^c|\sigma_k &\overset{\text{iid}}{\sim} \text{Normal}_{[0,10]}(k, \sigma_k^2), \\
z^c|\sigma_z &\overset{\text{iid}}{\sim} \text{Normal}_{[0,1.15]}(z, \sigma_z^2),
\end{aligned}$$

where $\text{Normal}_{[a,b]}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and standard deviation σ, truncated to lie between a and b.

This model allows us to pool information about the rates of gains across countries by assuming that each set of country-specific double-logistic parameters is randomly sampled from a common truncated normal distribution. The normal distribution is truncated such that all the double-logistic parameters are positive. The first five parameters clearly must be positive because they correspond to values of life expectancy and maximum gains.

The sixth double-logistic parameter, z^c, is the asymptotic average rate of increase in life expectancy per five-year period. Our prior distribution for this is informed by the results of Oeppen & Vaupel (2002), who found a strong positive linear trend in the "best practices" life expectancy (i.e. the highest life expectancy in a given year) from the mid-19th century through 2000. By assuming that z^c is non-negative, we are assuming that life expectancy will continue to increase on average. In their regression of highest male life expectancy on year, Oeppen & Vaupel (2002) estimated a slope of 1.11 years per five-year period with $R^2 = 0.98$. Because this is the rate of increase for "best practices" countries, we assume that the asymptotic rate of increase for any given country will not exceed the upper bound of the 99.9 percent confidence interval for this estimate, namely 1.15.

To specify the distribution of the random perturbations, $\varepsilon_{c,t}$, we first estimated the model assuming them to be normally distributed with a constant variance, using the estimation method described later. Figure 3 shows the absolute residuals from this fit with a fitted regression spline. The spread of the residuals clearly decreases with increasing life expectancy. To account for this, we modeled the $\varepsilon_{c,t}$ as normally distributed with standard deviation proportional to the regression spline fitted to the absolute residuals shown in Figure 3, so that

$$\varepsilon_{ct} \overset{\text{iid}}{\sim} N(0, (\omega \times f(l_{c,t-1}))^2). \tag{3}$$

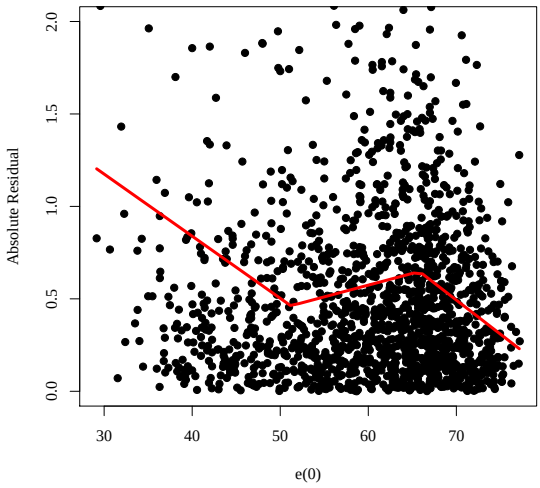

Figure 3: Absolute residuals from the constant variance model plotted against across life expectancy, with fitted regression spline. (Note: 44 (2.8%) of the residuals are outside the range of the plot, but were included in the regression spline fit.)

Parameter Estimation

We adopted a Bayesian approach to estimating our model, making it a Bayesian hierarchical model. This requires specifying prior distributions for the 13 world parameters of the model: $(\Delta_i, \sigma^2_{\Delta_i})$ for $i = 1, \ldots, 4$, k, σ^2_k, z, σ^2_z and ω. We specified prior distributions that are proper but much more diffuse than the posterior distributions.

We set $\Delta_i \sim N_{[0,100]}(a_i, \delta_i^2)$ for $i = 1, \ldots, 4$, $k \sim N_{[0,10]}(a_5, \delta_5^2)$ and $z \sim N_{[0,1.15]}(a_6, \delta_6^2)$. We set $(a_1, \ldots, a_6)$ to the values specifying the UN medium-pace model, namely $(15.77, 40.97, 0.21, 19.82, 2.93, 0.40)$. We set $(\delta_1^2, \ldots, \delta_6^2)$ to the variances of each parameter among the different UN models, namely $(3.56, 3.93, 3.96, 3.80, 0.99, 0.16)$.

For the world variance parameters, $\sigma^2_{\Delta_i} (i = 1, \ldots, 4)$, σ^2_k and σ^2_z, we used inverse-gamma prior distributions with 4 degrees of freedom (i.e. a shape parameter equal to 2). To set the parameters of these priors, we first fit the double-logistic model by least squares to the data from each country individually, and then for each parameter we computed the empirical average squared deviations from the values for the UN medium-pace model. We then set the prior means of the reciprocals of the world variance parameters equal to the reciprocals of these values. This yields rate parameters $(15.6^2, 23.5^2, 14.5^2, 14.7^2, 3.5^2, 0.6^2)$ for the six inverse gamma prior distributions. The resulting prior distributions are guaranteed to be much more spread out than the posterior distribution. Finally, a diffuse Uniform[0,10] prior was used for ω.

Experiments showed that the results were insensitive to changes in these priors, which is to be expected because the resulting prior distribution is much more spread out than the posterior distribution.

The posterior distribution of the parameters was approximated by Markov chain Monte Carlo, using the package Winbugs 1.4 (Lunn et al., 2000), run in R using the R function R2WinBUGS (Sturtz et al., 2005). We used three chains of length 100,000 scans each, with a burn-in of 10,000 scans. Visual inspection of the trace plots, the Raftery-Lewis diagnostic (Raftery & Lewis, 1992) and the Gelman-Rubin statistic (Gelman & Rubin, 1992) all indicated that the chains had converged and had explored the posterior distribution enough to yield good estimates of posterior quantiles of interest.

MODEL VALIDATION

To assess our probabilistic projections, we fit our model to data from 1950 to 1995 (n =1,422). We then used the resulting model to forecast life expectancy for males from 1995 to 2005, and compared the forecasts with what was actually observed. We had 316 out-of-sample predictions.

Comparing probabilistic forecasts with outcomes is not simple because it involves comparing two different kinds of objects : a probability distribution and a single value (Gneiting et al., 2007). A probabilistic forecast should be calibrated, i.e. x% prediction intervals should contain the truth x% of the time. It should also be sharp, i.e. the intervals should be as narrow as possible.

To assess the predictive ability of our model, we examined numerical measures of calibration via the coverage of our prediction intervals, root mean squared error (RMSE), mean absolute error (MAE) and the mean standardized absolute predictive error (SAPE). The standardized absolute predictive error for country c at time t is

$$a_{ct} = \sqrt{\frac{2}{\pi}} * \frac{|l_{ct} - \hat{l}_{ct}|}{\hat{\sigma}_{pred,ct}}.$$

Thus the SAPE is the absolute difference between the actual observed life expectancy (l_{ct}) and the median forecast ($\hat{l}_{ct}$), standardized by the standard deviation of the predictive distribution. When the model is correctly specified, the expected mean SAPE value is equal to one.

These metrics are given in Table 1. Overall, our model was well calibrated with our 95% prediction intervals capturing the actual observations 92% of the time. The empirical coverage of our 80% prediction intervals was 82%. The mean SAPE was 1.04, which is close

to the theoretical mean of 1. The MAE of our median predictions was 1.07 years. Thus our "best guess" was within just over one year of the actual observation, on average.

Table 1: Summary measures for 10 year out-of-sample predictions for the Bayesian hierarchical model and the current UN methodology.

Measure	BHM	UN replication
Root mean squared error (RMSE)	1.64	2.50
Std absolute prediction error (SAPE)	1.04	n/a
Mean absolute error (MAE)	1.07	1.86

Prediction Intervals		
Nominal	Actual	Mean half-length
95%	92.1%	2.54
90%	89.2%	2.13
80%	82.0%	1.66

We compared the predictive ability of our model with the current UN methodology. Replicating our cross-validation methodology by using WPP 2008 data through 1995, a UN analyst computed life expectancy forecasts for 1995-2005 using one of the five prescribed UN models of gains in life expectancy at birth based on levels and trends in the preceding two decades. The RMSE and the MAE of our method were both substantially smaller than those of the current UN method (1), by over 40%.

We assessed the sharpness of our projections using the distribution of prediction interval lengths. For the 1995-2000 time period, the prediction interval half-lengths ranged from 0.7 to 1.9 years with an average half-length of 1.3 years. For the next five-year period, 2000-2005, the interval half-lengths increased to a range of 1.0 to 3.2 years with an average half-length of 1.9 years. Average life expectancy at birth and prediction interval lengths both varied by region. Among the continental regions from 1995 to 2005, Africa had the lowest life expectancy of 59.8 years with an average interval half-length of 2.2 years. With an average life expectancy of 73.5 years and an interval half-length of 1.3 years, North America had both the highest life expectancy and the narrowest prediction intervals.

CASE STUDIES

Typical country: Madagascar

Currently, the UN estimates life expectancy at birth among males in Madagascar to be 58.5 years. Figure 4(a) shows projections of life expectancy starting from 2005–2010. The median

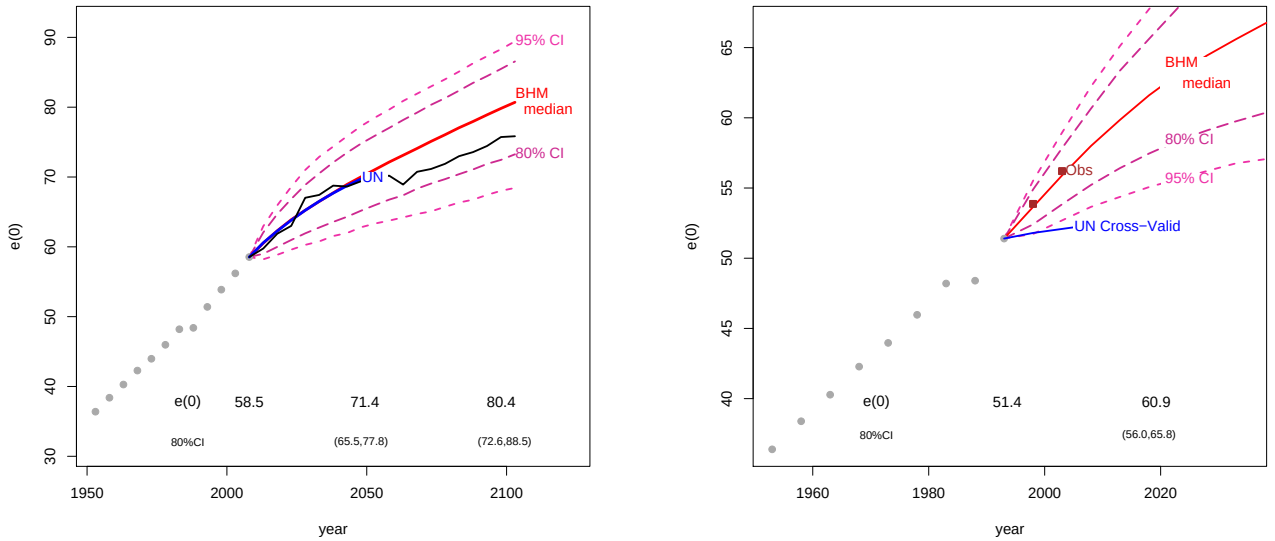

(a) Projections from 2005-2010.

(b) Cross-validation projections from 1990-1995.

Figure 4: Life Expectancy Projections for Males in Madagascar. The above plots include the UN projections and our median projections, with 80% and 95% prediction intervals. The life expectancy values used to estimate our model are indicated by grey circles. (a) Projections from 2005-2010 with a sample BHM trajectory. (b) Cross-validation projections from 1990-1995. Observed life expectancies from 1995-2005 are shown as squares.

forecast from our BHM is similar to WPP 2008 forecasts through 2050. In 2045-2050, the UN projects male life expectancy will be 69.7 years. We project life expectancy will be 71.4 years with an 80% prediction interval of (65.5,77.8). Fifty years after that, in 2095-2100, we project life expectancy to reach 80.4 years with a wider 80% prediction interval of (72.6,88.5). Figure 4(b) shows out-of-sample projections for Madagascar with projections beginning in 1990-1995. UN observed estimates are indicated in the plot by brown squares and are close to our median projection. The exclusion of two time periods results in more uncertainty in our projections for 2095-2100 with a median of 76.6 years and an 80% prediction interval of (64.7,87.8).

Along with quantiles of the projected life expectancy distribution, in Figure 4(a) (as well as Figures 5(a) and 6(a) for Latvia and Japan), we also include a sample stochastic trajectory for Madagascar. We see that unlike the quantiles, the sample trajectory does not follow a smooth path. For this trajectory the mean absolute deviation from the median is equal to the median mean absolute deviation among the posterior sample of projected trajectories. It can be viewed as a trajectory with typical deviation from the projected median.[2]

[2]This choice of a typical stochastic trajectory was suggested by Joel Cohen (personal communication,

Mortality Crisis: Latvia

Figure 5 shows estimated and projected life expectancies for Latvia. Male life expectancy in Latvia increased from 62.5 years in 1950 to 66.3 years in 1965. However, in the subsequent 15 years, male life expectancy in Latvia decreased by 2.2 years, to 64.1 years. Life expectancy increased again until a 3.8 year decline was recorded between 1985-1990 and 1990-1995. Since then, life expectancy in Latvia has been increasing again. Both the UN and our median projections predict a continuous increase in life expectancy, but ours increase more slowly.

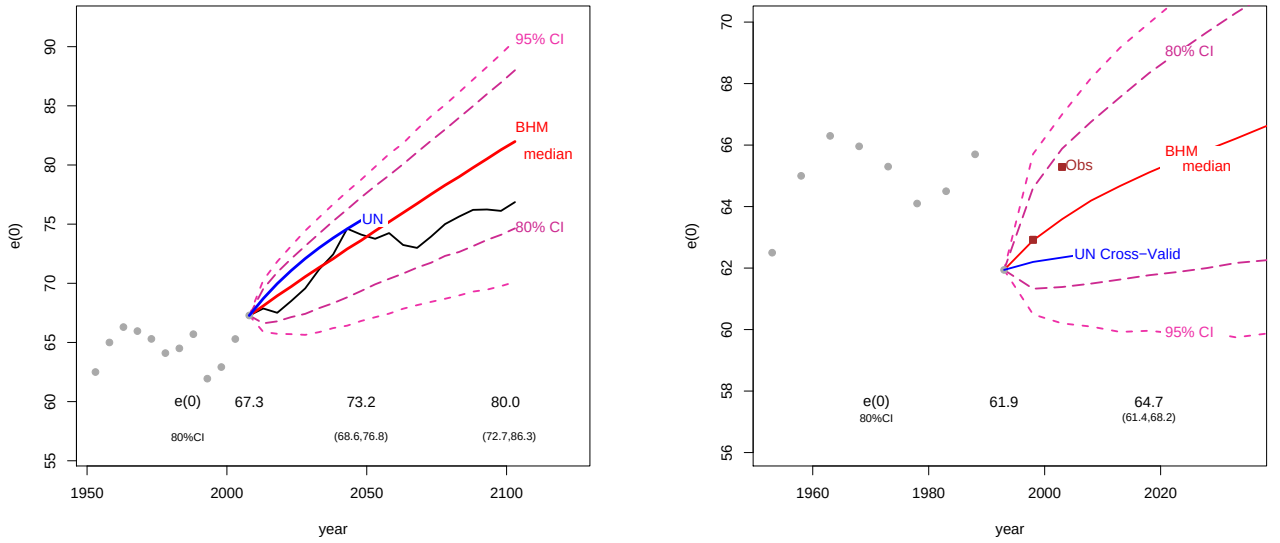

(a) Projections to 2100 starting from 2005-2010.

(b) Cross-Validation projections from 1990-1995.

Figure 5: Life Expectancy Projections for Latvia. The above plots include the UN projections and BHM median projections, with 80% and 95% prediction intervals. The past values of life expectancy values used to estimate our model are shown by grey circles. (a) Projections to 2100 starting from 2005-2010. A typical stochastic trajectory is shown in black, illustrating the non-smoothness of individual projections. (b) Out-of-sample projections starting from 1990-1995. Observed life expectancies from 1995-2005 are shown as squares. By 1995, Latvia had not yet recovered from its mortality crisis; the BHM projection intervals reflect uncertainty about a full recovery.

As can be seen from Figure 5(b), our 80% prediction intervals capture the observed estimates of life expectancy from 1995 to 2005. For the first time period, from 1995-2000, the upper bound of our 80% prediction interval was 64.4 years. Yet the lower bound of our 80% prediction interval, 61.1 years, indicates that life expectancy may continue to decrease.

December, 2009).

In fact, our prediction intervals allow for the possibility of life expectancy not increasing for the following 50 years, reflecting the erratic progress over the previous 40 years.

Leading country: Japan

One of the difficulties with projecting mortality is accurately projecting the country with the highest life expectancy. Historically, "pessimists" believed that life expectancy could not keep rising at historic rates and assumed that there must be a "ceiling" to life expectancy for humans (Fries, 1980, Olshansky et al., 1990, Olshansky et al., 2001, Olshansky et al., 2002, Olshansky et al., 2005). "Optimists", on the other hand, see no evidence of a limit to increases in life expectancy (Oeppen & Vaupel, 2002, Tuljapurkar et al., 2000, Tuljapurkar, 2005). However, past estimates of the "maximum life expectancy" have continually been surpassed (Oeppen & Vaupel, 2002), and old-age mortality rates continue to decline (Vaupel et al., 1998). Oeppen & Vaupel (2002) presented evidence that the world's highest, or "best practices", life expectancy at birth has increased linearly across time and shows no signs of leveling off. They estimated that the "best practices" life expectancy for males has increased at a rate of 1.11 per five-year period. See Bongaarts (2006) for a review of the historical debate on life expectancy limits.

Although Japan does not currently have the highest male life expectancy (that title has belonged to Iceland since 2000), it is the country with the highest overall life expectancy and has been since 1980. Figure 6(a) is a plot of male life expectancy in Japan. Also included in the plot is what the trajectory would be if male life expectancy in Japan increased at the "best practices" rate of 1.11 per five-year period. Vallin & Meslé (2009) updated and expanded the data time period (from 1840-2000 to 1750-2005) for "best practices" life expectancy. They found a segmented line fit the extended time frame better, with the most recent segment (1960-2005) still with a strong positive slope (1.13 years per five-year period for women), and concluded that the Oeppen-Vaupel line may be too optimistic for the long-term future. Our median projection increases more slowly than the Oeppen-Vaupel "best practices" linear projection, reflecting the fact that the "best-practices" lines is for the best country at each time, and thus likely to increase faster than for an individual country. However, the "best practices" trajectory is just within the upper bound of our 80% prediction interval.

Bongaarts (2006) also found the Oeppen-Vaupel "best-practices" rate to be overly optimistic. By decomposing mortality into juvenile, background, and senescent mortality, he observed that historically large gains in life expectancy were due to declines in juvenile mortality. Then as juvenile mortality reached low levels, the rate of gains in life expectancy diminished. The rationale for this decomposition is similar to that for the double-logistic function, where there are periods of high gains in life expectancy (i.e., when juvenile mortal-

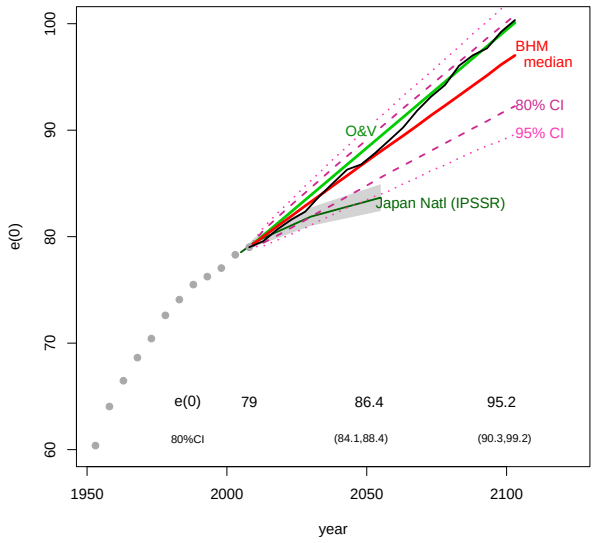

(a) Projections beginning in 2005-2010

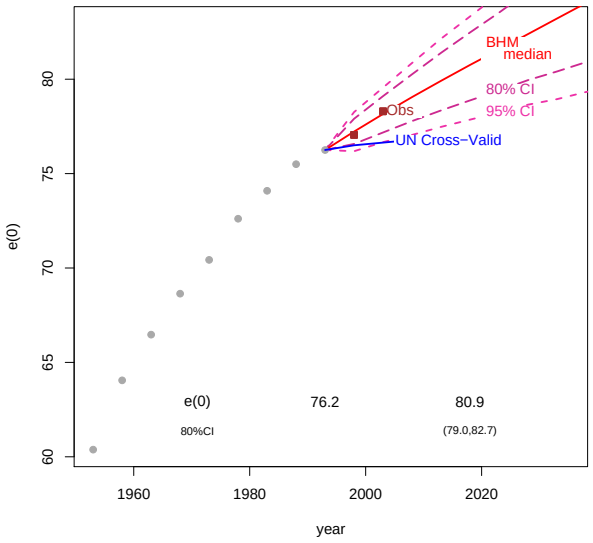

(b) Cross-Validation projections from 1990-1995.

Figure 6: Life Expectancy Projections for Japan. The above plots include the UN projections and BHM median projections, with 80% and 95% prediction intervals. The life expectancy values used to estimate our model are show by grey circles. (a) Projections from 2005-2010 with a sample trajectory. National Institute of Population and Social Security Research (IPSS) medium variant projections are the same as the UN projections with uncertainty bounds indicated in the shaded region. We include a trajectory with a constant increase of 1.11 years per five-year period, as estimated by Oeppen and Vauppel (O & V) (2002) for the "best practices" country. A typical stochastic trajectory is also included to illustrate the non-smoothness of individual BHM projections. (b) Cross-validation projections from 1990-1995. Observed life expectancies from 1995-2005 are shown as squares.

ity is declining) followed by a leveling off of gains (i.e., when the gains in life expectancy are due to incremental declines in senescent mortality). Bongaarts (2006) found that senescent life expectancy in countries with low mortality, on average, increased at a rate of 0.75 years per five-year period. The average asymptotic rate of gains estimated in our model was 0.84 years per five-year period, which is much closer to the Bongaarts (2006) projected gains in senescent life expectancy than the "best practices" rate of increase.

Recently, the Japanese official projections made by the National Institute of Population and Social Security Research (IPSSR) extended the Lee-Carter method to provide more refined estimates of mortality at higher ages. The original Lee-Carter method estimated age-specific mortality rates for five-year age groups, with the last age group aggregating those 85 and older. The IPSSR now uses the shifting logistic model (Bongaarts, 2005) to account for continued increases in life expectancy in Japan. IPSSR projections (low/medium/high rates of mortality decline variants, with the medium variant being equivalent to the UN projections) (Kaneko et al., 2008) are included in Figure 6(a). The IPSSR projections are more conservative and project an earlier leveling off of life expectancy than our projections, but are still within our prediction intervals.

When looking at out-of-sample projections in Figure 6(b), which begin in 1990-1995, we see that the UN out-of-sample projections suggested an immediate leveling off of life expectancy, unlike our projections. In fact, the observed life expectancy in 1995-2005 did not level off and instead continued to increase.

AGGREGATION TO REGIONAL PROJECTIONS: SOUTH ASIA

Researchers at the International Institute for Applied Systems Analysis (IIASA) (Lutz et al., 2004) produced regional probabilistic prediction intervals for life expectancy using Delphi-type methods. A group of experts were asked to give 90% prediction intervals for future life expectancy in each of 13 specified regions. Linear paths were then drawn from a normal distribution to produce probabilistic predictive distributions. This method uses demographic knowledge as an input whereas time-series methods rely on past trends.

Country-specific projections allow regional projections to be made regardless of how the region is defined. To compare our projections with those of IIASA for South Asia, we aggregated UN estimates and projections and our projections to be proportional to the regional male average populations from 2005 to 2010.[3] The countries included in the IIASA-

[3]Regional aggregation by weighted average using total male population is an approximation to the true regional aggregation. The true computation is more complex and requires the survival ratios by age and

defined region of South Asia (with population percentage) are: India (75.1%), Pakistin (10.5%), Bangladesh (9.9%), Nepal (1.7%), Afghanistan (1.5%), Sri Lanka (1.3%), Bhutan (0.04%), Maldives (0.02%).

Our model takes no account of correlations between the random perturbations in life expectancy gains in different countries. Previous work (Alho, 2008) has suggested that cross-country correlations are non-zero and should be modeled as such. Within South Asia, in the past 60 years, life expectancies were indeed highly correlated, as follows:

$$\begin{pmatrix}
\text{Afghanistan} & 0.91 & 0.91 & 0.99 & 0.98 & 0.94 & 0.98 & 0.99 \\
0.91 & \text{Bangladesh} & 0.99 & 0.92 & 0.96 & 0.99 & 0.97 & 0.89 \\
0.91 & 0.99 & \text{Bhutan} & 0.92 & 0.96 & 0.99 & 0.96 & 0.89 \\
0.99 & 0.92 & 0.92 & \text{India} & 0.99 & 0.95 & 0.99 & 0.99 \\
0.98 & 0.96 & 0.96 & 0.99 & \text{Maldives} & 0.98 & 0.99 & 0.98 \\
0.94 & 0.99 & 0.99 & 0.95 & 0.98 & \text{Nepal} & 0.98 & 0.92 \\
0.98 & 0.97 & 0.96 & 0.99 & 0.99 & 0.98 & \text{Pakistan} & 0.97 \\
0.99 & 0.89 & 0.89 & 0.99 & 0.98 & 0.92 & 0.97 & \text{Sri Lanka}
\end{pmatrix}.$$

However, for projection purposes, what matters are the cross-country correlations between the random perturbations (effectively the forecast errors) rather than between the life expectancies themselves. The hierarchical modeling of the change in life expectancy allows for between-country correlation in life expectancy. We are interested in the residual correlations in life expectancies gains, ρ_{c_i,c_j}, between countries c_i and c_j, namely:

$$(\hat{\rho}_{c_i,c_j}) = \begin{pmatrix}
\text{Afghanistan} & -0.28 & -0.23 & 0.09 & 0 & -0.07 & -0.08 & 0.19 \\
-0.28 & \text{Bangladesh} & 0.05 & -0.03 & 0.03 & 0 & 0.05 & -0.19 \\
-0.23 & 0.05 & \text{Bhutan} & -0.03 & 0.05 & 0.09 & 0.07 & 0.2 \\
0.09 & -0.03 & -0.03 & \text{India} & 0.16 & 0.03 & 0.03 & 0 \\
0 & 0.03 & 0.05 & 0.16 & \text{Maldives} & 0.24 & 0.25 & 0.34 \\
-0.07 & 0 & 0.09 & 0.03 & 0.24 & \text{Nepal} & 0.08 & 0.05 \\
-0.08 & 0.05 & 0.07 & 0.03 & 0.25 & 0.08 & \text{Pakistan} & -0.01 \\
0.19 & -0.19 & 0.2 & 0 & 0.34 & 0.05 & -0.01 & \text{Sri Lanka}
\end{pmatrix}.$$

These correlations are much smaller than those between the life expectancies themselves, and in most cases are consistent with the absence of any correlation. Our projection model has accounted for all or most of the between-country correlation in South Asia.

Country-specific projections were then made by sampling the vector of random perturbations, $(\delta_{c_i,t})$, from a multivariate normal distribution whose covariance matrix incorporates the residual correlation matrix, $\left(\hat{\rho}_{c_i,c_j}\right)$. The aggregated projections for the South Asian region are shown in Figure 7. The 2007 IIASA projections available on their website (Lutz et al., 2008) are also depicted. We found that our median projections were consistent with

sex weighted by their respective population by age and sex. In addition, future weights depend on future fertility by country, and therefore future weights will change over time.

IIASA's median projections, but our intervals are much sharper, ranging from 36–72% narrower than those of IIASA.

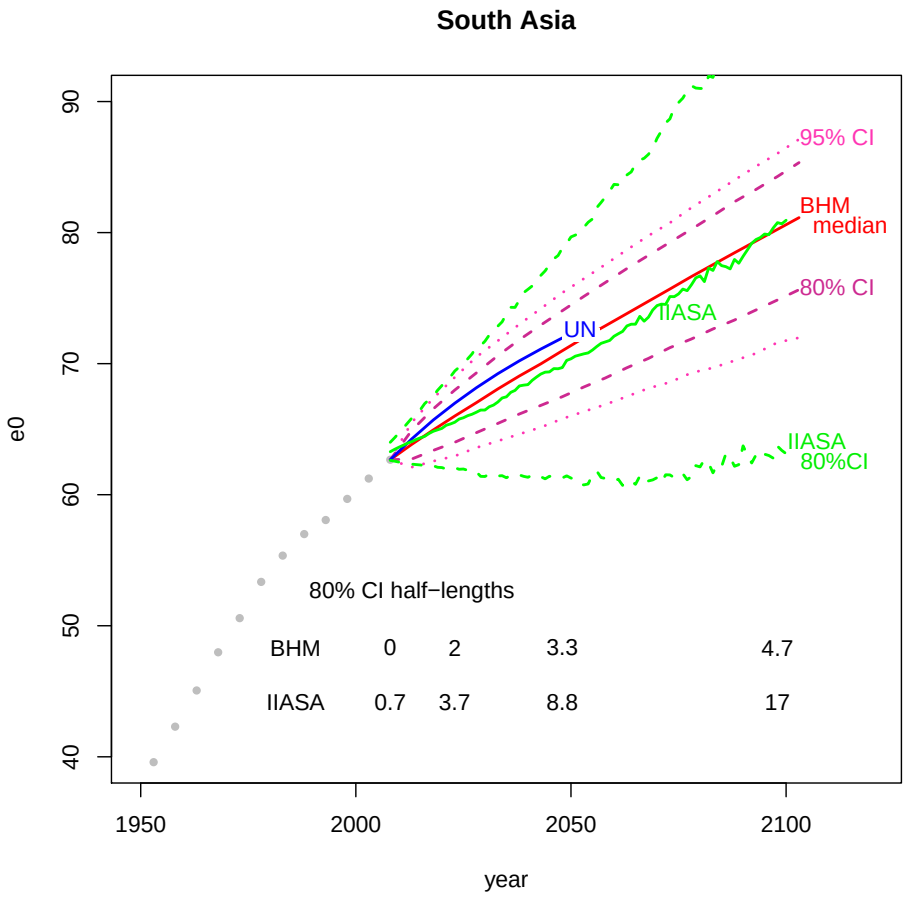

Figure 7: Life expectancy projections for South Asia (IIASA-defined) for our BHM model, IIASA and the UN. The median projections for BHM and IIASA are similar with each other, but the IIASA 80% intervals are wider than the BHM 80% interval.

Our approach yields stochastic trajectories that can fluctuate around the projected life expectancy. By contrast, IIASA samples random linear trajectories, which change at a constant rate throughout the projection period, with no fluctuations.

DISCUSSION

We have proposed a Bayesian way to produce probabilistic projections of life expectancy for all the countries of the world. One possible use of these projections is as inputs to the United Nations estimation and projections. We have shown by out-of-sample cross-validation

that the method gives well-calibrated and sharp prediction intervals, and also gave better forecasts than the current UN method.

We have restricted our analysis to countries without generalized HIV/AIDS epidemics because of the singular nature of their demographic impact, predominantly sexually active adults. To explore the possible applicability of our method to the excluded countries, we loosened the exclusion rule and fitted our model to all countries with a 2000-2005 HIV/AIDS prevalence rate of less than 4% (n =179). In the 10-year cross-validation of these countries, the 80% prediction intervals included the observed life expectancy 84% of the time. Including the additional countries did not greatly degrade sharpness: the mean absolute error was 1.2 years and the average 80% prediction interval half-length was 1.9 years, compared to 1.0 and 1.7 years respectively without the HIV/AIDS countries. Further research is needed to generalize our model to countries with a generalized HIV/AIDS epidemic while properly accounting for the uncertainty in AIDS mortality, but these results suggest that the current model could provide reasonable approximate results.

Our method projects male life expectancy only. Further research is needed to apply this model to life expectancy at birth among males and females simultaneously while ensuring that trajectories by sex do not diverge or cross. Lee & Carter (1992) suggested doing this by modeling the two sexes independently and introducing a new parameter to ensure that the stochastic trajectories do not cross or diverge. Such a model could also be made more complex by allowing the double-logistic parameters for the two sexes to be correlated. An alternative approach would be to combine two stochastic models: one model for one sex, as here, or for average life expectancy, and a second model for the gap between the sexes.

Our model assumes that the rate of increase of life expectancy will decline in the future, but that gains will continue, asymptotically at a linear rate on average. It would be possible to modify the model to allow the rate of increase to decline to zero over time, but experience to date does not provide much support for this (Oeppen & Vaupel, 2002).

It could also be dangerous. Olshansky et al. (2009) argued that current official forecasts of life expectancy by the U.S. Social Security Administration and the U.S. Census Bureau may be underestimated, by about three years for males and eight years for females. They estimated that this discrepancy could cost as much as \$3 to \$8 trillion more than currently projected for Medicare and Social Security. The current official governement forecasts of U.S. male life expectancy in 2050 are 80−81 years, whereas Olshansky et al. (2009) projected 83−86 years under their two main scenarios, with a range of 78−113 years for their more extreme scenarios. Our BHM projections are close to theirs, with a median of 84.5 years and an 80% prediction interval of 82.2−86.4.

Much other research has been done on the forecasting of mortality (Booth, 2006; Dowd

et al., 2010). However, efforts have focused on developed countries where reliable age-specific data are available. The best known time-series method for forecasting age-specific mortality rates is the Lee-Carter method and its various parallels (e.g., Renshaw & Haberman, 2006), generalizations (e.g., de Jong & Tickle, 2006; Hyndman & Ullah, 2007; Pedroza, 2006; Koissi et al., 2006; Brouhns et al., 2002) and extensions (e.g., Li & Lee, 2005; Li et al., 2004; Ishii, 2008). As discussed earlier, the empirical evidence about the Lee-Carter method's performance for forecasting life expectancy is mixed.

There have been other time-series methods to estimate and project age-specific mortality rates. The Brass relational method fits a two-parameter model where the age-specific mortality rates are assumed to be given by a linear function of a user-chosen model life table on a logit scale (Brass, 1971). The Heligman-Pollard model is an eight-parameter model with three parts describing mortality in childhood, young adulthood, and later life (Heligman & Pollard, 1980). Although both models have fitted mortality data well (Keyfitz, 1991; Hartmann, 1987), difficulties may arise in projecting the parameters (Keyfitz, 1981). Like the Heligman-Pollard model, the Bongaarts shifting logistic model (Bongaarts, 2005) differentiates mortality at different ages by fitting a three-parameter logistic model and fixing the slope parameter across time while allowing the other two parameters to vary with time. However, the shifting logistic model focuses on senescent mortality and may be most relevant when infant/child and adult mortality are already negligible. Other models have focused on senescent mortality in terms of biological and evolutionary phenomena that incorporate the idea of heterogeneity in the population and frailty (Steinsaltz & Wachter, 2006, Yashin et al., 2000). These approaches require age-specific data.

Instead of modeling age-specific mortality rates, Gage (1993) used a competing hazards model developed by Siler (1979), which was developed for animal mortality. The five-parameter model describes the hazard function as made up of three components, immature (or childhood), residual (or background), and sensecent mortalities. Bongaarts (2006) also suggested decomposing mortality into these three components with future projections focusing on the senescent mortality component. Both approaches require cause of death data, and Gage (1993) acknowledged that this is a limitation due to issues of data quality and cause of death classification.

In addition to time-series approaches, there are two other main approaches to developing probabilistic projections (Lee, 1998). As previously discussed, expert-based probabilistic projections have been produced by Lutz and colleagues at the IIASA (Lutz et al., 1998, 2004, 2008). However, this method does not explicitly rely on the use of available data, instead relying on a collection of experts and their ability to specify probabilistic bounds, which may or may not be accurate (Alho, 2005). The other alternative to time-series methods is ex-post

analysis of previous projections (Keyfitz, 1981; Stoto, 1983; Smith & Sincich, 1990). In this method, previous forecast errors are used to create probabilistic errors on future projections.

Girosi & King (2008) recently proposed a Bayesian method that incorporates covariates in a linear regression model. However, their approach requires additional data that may not be reliable or even available in many countries. Pedroza (2006) proposed a Bayesian approach to the Lee-Carter model by accounting for the uncertainty in the age parameters as well as the mortality index usually forecasted. Czado et al. (2005) presented a Bayesian approach to the Poisson log-bilinear formulation of the Lee-Carter model. While the latter two approaches account for uncertainty in the Lee-Carter model, their generalization to all countries is hindered by the non-availability of age-specific mortality rates.

References

Alho, J. (2008). Aggregation across countries in stochastic population forecasts. *International Journal of Forecasting*, *24*, 343–353.

Alho, J. M. (2005). *Statistical Demography and Forecasting*. New York: Springer.

Alkema, L., Raftery, A. E., Gerland, P., Clark, S. J., & Pelletier, F. (2008). Assessing uncertainty in fertility estimates and projections. In *Proceedings of the 2008 Annual Meeting of the Population Association of America*, New Orleans, La., USA. http://paa2008.princeton.edu/abstractViewer.aspx?submissionId=80115.

Alkema, L., Raftery, A. E., Gerland, P., Clark, S. J., Pelletier, F., Buettner, T., & Heilig, G. (2011). Probabilistic projections of the total fertility rate for all countries. *Demography*, *48*, in press.

Bell, W. (1997). Comparing and assessing time series methods for forecasting age-specific fertility and mortality rates. *Journal of Official Statistics*, *13*, 279–202.

Bongaarts, J. (2005). Long-range trends in adult mortality: Models and projection methods. *Demography*, *42*, 23–49.

Bongaarts, J. (2006). How long will we live? *Population and Development Review*, *32*, 605–628.

Bongaarts, J. & Bulatao, R. A. (2000). *Beyond Six Billion: Forecasting the World's Population*. Washington, D.C.: National Academy Press.

Booth, H. (2006). Demographic forecasting: 1980 to 2005 in review. *International Journal of Forecasting*, *22*, 547–581.

Booth, H., Tickle, L., & Smith, L. (2005). Evaluation of the variants of the Lee-Carter method of forecasting mortality: A multicountry comparison. *New Zealand Population Review, Special Issue on Stochastic Population Projections*, *31*, 13–37.

Brass, W. (1971). On the scale of mortality. In Brass, W. (Ed.), *Biological Aspects of Demography*, (pp. 69–110)., London. Taylor and Francis.

Brouhns, N., Denuit, M., & Vermunt, J. (2002). A Poisson log-bilinear regression approach to the construction of projected life tables. *Insurance: Mathematics and Economics*, *31*, 373–397.

Czado, C., Delwarde, A., & Denuit, M. (2005). Bayesian Poisson log-bilinear mortality projections. *Insurance: Mathematics and Economics*, *36*, 260 – 284.

de Jong, P. & Tickle, L. (2006). Extending Lee-Carter mortality forecasting. *Mathematical Population Studies*, *13*, 1–18.

Dowd, K., Blake, D., & Cairns, A. J. G. (2010). Facing up to the uncertain life expectancy: The longevity fan charts. *Demography*, *47*, 67–78.

Fries, J. F. (1980). Aging, natural death, and the compression of morbidity. *New England Journal of Medicine*, *202*, 130–135.

Gage, T. B. (1993). The decline of mortality in England and Wales 1861 to 1964: Decomposition by cause of death and component of mortality. *Population Studies*, *47*, 47–66.

Gelman, A. & Rubin, D. (1992). Inference from iterative simulation using multiple sequences (with discussion). *Statistical Science*, *7*, 457–511.

Girosi, F. & King, G. (2008). *Demographic Forecasting*. Princeton, NJ: Princeton University Press.

Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society, Series B*, *69*, 243–268.

Hartmann, M. (1987). Past and recent attempts to model mortality at all ages. *Journal of Official Statistics*, *3*, 19–36.

Heligman, L. & Pollard, J. H. (1980). The age pattern of mortality. *Journal of the Institute of Actuaries*, *107*, 49–80.

Hyndman, R. J. & Ullah, S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics and Data Analysis*, *51*, 4942–4956.

Ishii, F. (2008). Mortality projection model for Japan with age-shifting structure. Paper presented at the annual meeting of The Population Association of America, April 17-19, New Orleans, Louisiana.

Kaneko, R., Ishikawa, A., Ishii, F., Sasai, T., Iwasawa, M., Mita, F., & Moriizumi, R. (2008). Population projections for Japan: 2006-2055 outline of results, methods, and assumptions. *The Japanese Journal of Population*, *6*, 76–114.

Keyfitz, N. (1981). The limits of population forecasting. *Population and Development Review*, *7*, 579–593.

Keyfitz, N. (1991). Experiments in the projection of mortality. *Canadian Studies in Population*, *18*, 1–17.

Koissi, M. C., Shapiro, A. F., & Högnäs, G. (2006). Evaluating and extending the Lee-Carter model for mortality forecasting: Bootstrap confidence interval. *Insurance: Mathematics and Economics*, *38*, 1–20.

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

Lee, R. D. (1998). Probabilistic approaches to population forecasting. *Population and Development Review*, *24*, 156–190.

Lee, R. D. & Carter, L. R. (1992). Modeling and forecasting U. S. mortality. *Journal of the American Statistical Association*, *87*, 659–671.

Li, N. & Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography*, *42*, 575–594.

Li, N., Lee, R., & Tuljapurkar, S. (2004). Using the Lee-Carter method to forecast mortality for populations with limited data. *International Statistical Review*, *72*, 19–36.

Lunn, D., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS v1.4.3 – a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, *10*, 325–337.

Lutz, W., Sanderson, W. C., & Scherbov, S. (1998). Expert-based probabilistic population projections. *Population and Development Review*, *24*, 139–155.

Lutz, W., Sanderson, W. C., & Scherbov, S. (2004). *The End of World Population Growth in the 21st century: New Challenges for Human Capital Formation and Sustainable Development*. Sterling, VA: Earthscan.

Lutz, W., Sanderson, W. C., & Scherbov, S. (2008). IIASA's 2007 probabilistic world population projections, IIASA world population program online data base of results. Retrieved November 6, 2009 (http://www.iiasa.ac.at/Research/POP/proj07/index.html?sb=5).

Meyer, P. (1994). Bi-logistic growth. *Technological Forecasting and Social Change*, *47*, 89–102.

Oeppen, J. & Vaupel, J. (2002). Broken limits to life expectancy. *Science*, *296*, 1029–1031.

Olshansky, S., Carnes, B., & Cassel, C. (1990). In search of Methuselah: Estimating the upper limits to human longevity. *Science*, *250*, 634–640.

Olshansky, S., Passaro, D., Hershow, R., Layden, J., Carnes, B., Brody, J., Hayflick, L., Butler, R., Allison, D., & Ludwig, D. (2005). A potential decline in life expectancy in the United States in the 21st century. *New England Journal of Medicine*, *352*, 1103–1110.

Olshansky, S. J., Carnes, B., & Brody, J. (2002). A biodemographic interpretation of the life span. *Population and Development Review*, *28*, 501–513.

Olshansky, S. J., Carnes, B. A., & Désesquelles, A. (2001). Prospects for human longevity. *Science*, *291*, 1491–1492.

Olshansky, S. J., Goldman, D. P., Zheng, Y., & Rowe, J. W. (2009). Aging in America in the twenty-first century: Demographic forecasts from the MacArthur Foundation Research Network on an Aging Society. *The Milbank Quarterly*, *87*, 842–862.

Pedroza, C. (2006). A Bayesian forecasting model: Predicting U.S. male mortality. *Biostatistics*, *7*, 530–550.

Raftery, A. & Lewis, S. (1992). One long run with diagnostics: Implementation strategies for Markov chain Monte Carlo. *Statistical Science*, *7*, 493–497.

Raftery, A. E., Alkema, L., Gerland, P., Clark, S. J., Pelletier, F., Buettner, T., Heilig, G., Li, N., & Ševčíková, H. (2009). White paper: Probabilistic projections of the total fertility rate for all countries for the 2010 *world population prospect*s. In *United Nations Expert Group Meeting on Recent and Future Trends in Fertility*. United Nations, Department of Economic and Social Affairs, Population Division. Available at http://www.un.org/esa/population/meetings/EGM-Fertility2009/egm-fertility2009.html.

Renshaw, A. E. & Haberman, S. (2006). A cohort-based extension to the Lee-Carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, *38*, 556–570.

Siler, W. (1979). A competing-risk model for animal mortality. *Ecology*, *60*, 750–757.

Smith, S. K. & Sincich, T. (1990). On the relationship between length of base period and population forecast errors. *Journal of the American Statistical Association*, *85*, 367–375.

Steinsaltz, D. R. & Wachter, K. W. (2006). Understanding mortality rate deceleration and heterogeneity. *Mathematical Population Studies*, *13*, 19–37.

Stoto, M. A. (1983). The accuracy of population projections. *Journal of the American Statistical Association*, *78*, 13–20.

Sturtz, S., Ligges, U., & Gelman, A. (2005). R2winbugs: A package for running winbugs from r. *Journal of Statistical Software*, *12*(3), 1–16.

Tuljapurkar, S. (2005). Future mortality: A bumpy road to Shangri-La? *Science of Aging Knowledge Environment*, *14*, pe9.

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

United Nations (2009). *World population prospects: The 2008 Revision*. New York, NY: United Nations.

Vallin, J. & Meslé, F. (2009). The segmented trend line of highest life expectancies. *Population and Development Review*, *35*, 159–187.

Vaupel, J., Carey, J., Christensen, K., Johnson, T., Yashin, A., Holm, V., Iachine, I., Kannisto, V., Khazaeli, A., Liedo, P., Longo, V., Zeng, Y., Manton, K., & Curtsinger, J. (1998). Biodemographic trajectories of longevity. *Science*, *280*, 855–860.

White, K. M. (2002). Longevity advances in high-income countries, 1955-96. *Population and Development Review*, *28*, 59–76.

Yashin, A. I., Iachine, I. A., & Begun, A. S. (2000). Mortality modeling: a review. *Mathematical Population Studies*, *8*, 305–332.