

Bayesian Population Projections for the United Nations*

Adrian E. Raftery

Leontine Alkema

University of Washington

National University of Singapore

Patrick Gerland

United Nations

October 2, 2012; 2nd revision February 4, 2013

Abstract

The United Nations regularly publishes projections of the populations of all the world's countries broken down by age and sex. These projections are the de facto standard and are widely used by international organizations, governments and researchers. Like almost all other population projections, they are produced using the standard deterministic cohort-component projection method and do not yield statements of uncertainty. We describe a Bayesian method for producing probabilistic population projections for most countries that the United Nations could use. It has at its core Bayesian hierarchical models for the total fertility rate and life expectancy at birth. We illustrate the method and show how it can be extended to address concerns about the UN's current assumptions about the long-term distribution of fertility. The method is implemented in the R packages `bayesTFR`, `bayesLife`, `bayesPop` and `bayesDem`.

*Adrian E. Raftery, Department of Statistics, Box 354322, University of Washington, Seattle, WA 98195-4322. Email: raftery@uw.edu; Web: www.stat.washington.edu/raftery. Leontine Alkema, Department of Statistics and Applied Probability and Saw Swee Hock School of Public Health, National University of Singapore, Singapore 117546; Email: alkema@nus.edu.sg. Patrick Gerland, Population Estimates and Projections Section, United Nations Population Division, New York, NY 10017; Email: gerland@un.org.

Key words and phrases: Bayesian hierarchical model; cohort component projection method; double logistic function; Leslie matrix; life expectancy; total fertility rate.

1 Introduction

The United Nations (UN) publishes projections of the populations of all countries broken down by age and sex, updated every two years in a publication called the *World Population Prospects* (WPP). It is the only organization to do so. These projections are used by researchers, international organizations and governments, particularly with less developed statistical systems, and researchers. They are used for planning, social and health research, monitoring development goals, and as inputs to other forecasting models such as those used for predicting climate change and its impacts (Intergovernmental Panel on Climate Change, 2007; Seto et al., 2012). They are the de facto standard (Lutz and Samir, 2010).

Like almost all other population projections, the UN’s projections are produced using the standard cohort-component projection method (Whelpton, 1936; Leslie, 1945; Preston et al., 2001). This is a deterministic method based on an age-structured version of the basic demographic identity that the number of people in a country at time $t + 1$ is equal to the number at time t plus the number of births, minus the number of deaths, plus the number of immigrants, minus the number of emigrants.

The UN projections are based on assumptions about future fertility, mortality and international migration rates; given these rates the UN produces the “Medium” projection, a single value of each future population number with no statement of uncertainty. The UN also produces “Low” and “High” projections using total fertility rates (the average number of children per woman) that are, respectively, half a child lower and half a child higher than the Medium projections. These are alternative scenarios that also have no probabilistic interpretation.

Scientists, including researchers working on climate change, have long expressed interest

in UN population projections that would include statistical uncertainty intervals. This was first expressed in 1986 by a call to incorporate a probabilistic element in UN projections and to probabilistically specify the range of error (El-Badry and Kono, 1986). Independent evaluations of UN projections (National Research Council, 2000; Keilman et al., 2002) and expert-based probabilistic projections for the world and major regions (Lutz et al., 1998, 2004, 2008) have further highlighted the desirability of uncertainty bounds.

Responding to the call for the inclusion of uncertainty in populations projections, the UN is interested in producing probabilistic population projections for all countries; here we describe the current state of an ongoing effort to develop a methodology for doing so. Our method builds on previous work on time series methods for probabilistic population projections (National Research Council, 2000), particularly the work of Ronald D. Lee and his collaborators (Lee and Carter, 1992; Lee and Tuljapurkar, 1994; Lee, 2011).

In Section 2 we summarize the current UN approach and in Section 3 we describe our probabilistic approach. In Section 4 we consider how a modification to the method could accommodate disagreement about the long-term behavior of fertility assumed in the model, and in Section 5 we discuss the contribution of Bayesian thinking to the method.

2 Current UN Population Projection Methodology

We now outline the UN’s current (deterministic) population projection method, as used in the *World Population Prospects 2008* (United Nations, 2009) and described by United Nations (2006). The most recent UN projections published in the *World Population Prospects 2010* (United Nations, 2011a) incorporate some aspects of the new methods we will describe here. Thus we will refer to the 2008 WPP method as the “current” method.

2.1 Cohort Component Projection Method

At the heart of the UN's current population projection method lies the cohort component projection or Leslie matrix method. To fix ideas we describe a simplified version here. We consider one sex (female) and divide the population into N k -year age groups; those in the x -th age group are aged from $k(x - 1)$ years to $(kx - 1)$ years. The projection is done by k -year time periods, where k is typically 5 or 1 (in our work we use $k = 5$), and the beginning of the t -th time period will be referred to as time t .

We let $n_{x,t}$ be the number of females in the x -th age group at time t . We let $S_{x,t}$ be the survival ratio for the x -th age group in the t -th period, i.e. the proportion of the females in the x -th age group at time t who are still alive at time $t + 1$. We let $B_{x,t}$ be the number of female offspring of females in the x -th age group at time t who are born in the t -th period and survive to time $t + 1$, divided by $n_{x,t}$. Finally, we denote net migration by $m_{x,t}$, equal to the number of immigrants during the t -th period who were in the x -th age group at time t and are still in the population at time $t + 1$, minus the corresponding number of emigrants. For the highest age group, N , we assume that those who survive stay in the same age group at the next time point.

Then the model is simply a series of deterministic accounting identities:

$n_{1,t+1} = \sum_{x=1}^N B_{x,t}n_{x,t} + m_{1,t}$, $n_{x+1,t+1} = S_{x,t}n_{x,t} + m_{x+1,t}$ for $x = 1, \dots, N - 2$, and $n_{N,t+1} = S_{N-1,t}n_{N-1,t} + S_{N,t}n_{N,t} + m_{N,t}$. If we define the projection matrix for period t , \mathbf{P}_t , by

$$\mathbf{P}_t = \begin{bmatrix} B_{1,t} & B_{2,t} & \cdots & B_{N-1,t} & B_{N,t} \\ S_{1,t} & 0 & \cdots & 0 & 0 \\ 0 & S_{2,t} & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & S_{N-1,t} & S_{N,t} \end{bmatrix},$$

then the model can be rewritten in matrix form as

$$\mathbf{n}_{t+1} = \mathbf{P}_t \mathbf{n}_t + \mathbf{m}_t, \quad (1)$$

where $\mathbf{n}_t = (n_{1,t}, \dots, n_{N,t})^T$, and $\mathbf{m}_t = (m_{1,t}, \dots, m_{N,t})^T$. This can be applied recursively to obtain population projections. It can be extended in a fairly straightforward way to project two-sex populations.

The formulation (1) is due to Leslie (1945). The deterministic analysis of (1) and its use for population projections are the subject of classical mathematical or formal demography; see for example Preston et al. (2001), Keyfitz and Caswell (2005) and Caswell (2006).

The use of (1) for population projections requires that values of future age-specific mortality, fertility and migration be specified for each future time period to be projected. This is the hard part, and most of the uncertainty about future population is due to uncertainty about these future quantities.

2.2 Projecting mortality and fertility rates

The UN's current method generates assumptions about future age-specific fertility and mortality rates for most countries by projecting forward the overall level of future fertility or mortality, and then converting the overall levels to age-specific rates.

The UN's current method for projecting life expectancy at birth (hereafter just referred to as life expectancy) for most countries is as follows. Five-year gains in life expectancy for country c in time period t , $\ell_{c,t}$, are projected using a deterministic double logistic function, namely

$$\ell_{c,t+1} = \ell_{c,t} + g(\ell_{c,t} | \boldsymbol{\theta}^c), \quad (2)$$

where five-year gains $g(\ell_{c,t} | \boldsymbol{\theta}^c)$ are given by

$$g(\ell_{c,t} | \boldsymbol{\theta}^c) = \frac{k^c}{1 + \exp(-\frac{A_1}{\Delta_2^c}(\ell_{ct} - \Delta_1^c - A_2 \Delta_2^c))} + \frac{z^c - k^c}{1 + \exp(-\frac{A_1}{\Delta_4^c}(\ell_{ct} - \sum_{i=1}^3 \Delta_i^c - A_2 \Delta_4^c))}. \quad (3)$$

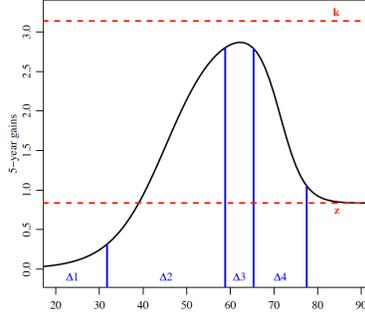


Figure 1: Example of a double logistic function used for projecting the five-year gains in life expectancy, plotted against life expectancy.

In (3), $\theta^c = (\Delta_1^c, \Delta_2^c, \Delta_3^c, \Delta_4^c, k^c, z^c)$ are the six parameters of the double logistic function for country c , whose meaning is illustrated in Figure 1, and A_1 and A_2 are constants. The parameters to be used for a given country are chosen by the UN analyst for that country from a list of five pre-determined patterns that represent different rates of improvement in life expectancy.

The most used measure of the overall level of fertility is the total fertility rate (TFR) for country c at time t , defined as $f_{c,t} = k \sum_{x=1}^{N-1} F_{c,x,t}$, where $F_{c,x,t}$ is the fertility rate in country c for age group x at time t . The TFR is the average number of children a woman would bear in her life if exposed to the age-specific fertility rates prevalent at time t .

The UN's current method for projecting TFR takes account of several empirical regularities. The past century has been dominated by the fertility transition, a shift from high fertility and high mortality to low fertility and low mortality, that started in Europe and North America in the late 19th century and in East Asia in the mid 20th century. It has now started in almost all countries and is complete in many (Hirschman, 1994). The patterns of change in different countries have been similar. The TFR starts from a high level that differs among countries but is typically between 4 and 8, and then starts to decline slowly. The pace of decline reaches a peak about half way through the transition. Then the pace of decline slows, stopping some time after the TFR goes below the replacement level of about 2.1 children per woman. In several low-fertility countries, a slow increase has been observed

after this point.

The UN has projected five-year decrements in the TFR using a double logistic function, similar to the function used for projecting gains in life expectancy. The parameters to be used for a given country are chosen by the UN analyst for that country from a list of three pre-determined patterns. In the projections, the TFR is held constant once it reaches 1.85 children, which represents the deterministic ultimate fertility level or asymptote. For countries where the TFR is below 1.85 at the start of the projection, the TFR is projected to increase by 0.05 children per 5-year period until the ultimate level is reached.

Projected values of the total fertility rate and life expectancy are converted to age-specific fertility and mortality rates based on past patterns or model life tables, and population projections are produced using the cohort component method. Finally, high and low variants are produced by increasing or decreasing the total fertility rate in each future period by half a child.

3 Bayesian Population Projections

The UN's current projection method does not yield an assessment of uncertainty about future population quantities. It is somewhat subjective because the double logistic functions used have been selected by the analyst from a small number of predetermined possibilities rather than estimated from the data. It is also somewhat rigid in that the set of double logistic functions used is small and may not cover a full range of realistic future possibilities.

To address these issues we have developed a Bayesian probabilistic population projection method. This involves building Bayesian hierarchical models to project the total fertility rate and life expectancy, each of which produces a large number of possible future trajectories from the posterior predictive distribution. These are then input to the cohort component projection method to provide a posterior predictive distribution of any future population quantity of interest. We now briefly describe the method.

The resulting model is as follows:

$$f_{c,t+1} = f_{c,t} - r(f_{c,t}|\boldsymbol{\delta}^c) + a_{c,t}, \quad (4)$$

where the five-year decrement $r(f_{c,t}|\boldsymbol{\delta}^c)$ is given by

$$r(f_{c,t}|\boldsymbol{\delta}^c) = \frac{-d^c}{1 + \exp\left(-\frac{2\ln(9)}{\nabla_1^c}(f_{c,t} - \sum_{i=2}^4 \nabla_i^c + 0.5\nabla_1^c)\right)} + \frac{d^c}{1 + \exp\left(-\frac{2\ln(9)}{\nabla_3^c}(f_{c,t} - \nabla_4^c - 0.5\nabla_3^c)\right)}, \quad (5)$$

with $\boldsymbol{\delta}^c = (\nabla_1^c, \nabla_2^c, \Delta_3^c, \nabla_4^c, d^c)$ being a vector of country-specific parameters and $a_{c,t} \stackrel{\text{iid}}{\sim} N(0, \sigma(t, f_{c,t})^2)$, where $\sigma(t, f_{c,t})$ is a function that describes how the error standard deviation changes with fertility level and time period.

The country-specific parameters, $\boldsymbol{\delta}_c$, are assumed to be drawn from a world distribution whose parameters (or hyperparameters) themselves have a diffuse prior distribution, namely $\boldsymbol{\delta}_c \stackrel{\text{iid}}{\sim} h(\cdot, \boldsymbol{\phi})$, where $\boldsymbol{\phi} \sim p(\boldsymbol{\phi})$. The resulting Bayesian hierarchical model is estimated using Markov chain Monte Carlo.

We define a country as having entered Phase III once two consecutive five-year increases below a TFR of 2 children have occurred. By this definition 21 countries had entered Phase III by 2010: 19 European countries, the USA and Singapore. For these countries, TFR has tended to increase back towards replacement level after they entered Phase II, reversing the secular trend of fertility decline. This is by now a well-documented trend (Myrskylä et al., 2009).

To model this we used a single first-order autoregressive model with long-term mean μ equal to the approximate replacement fertility level of 2.1 for all countries in Phase III, namely

$$f_{c,t+1} - \mu = \rho(f_{c,t} - \mu) + b_{c,t}, \quad (6)$$

where $b_{c,t} \stackrel{\text{iid}}{\sim} N(0, \sigma_b^2)$. The parameters ρ and σ_b were estimated by maximum likelihood from the 54 time periods observed in the 21 countries that have entered Phase III, yielding

$\hat{\rho} = 0.89$ and $\hat{\sigma}_b = 0.10$. The estimated value of ρ gives expected increases that are similar to the current UN increments of 0.05 children for each five-year period until the TFR equals 1.85.

3.2 Bayesian Life Expectancy Projection Model

We model female life expectancy similarly to Phase II total fertility. We use the UN's double logistic function to project expected gains, but we add a heteroscedastic error term, we allow the parameters of the country-specific double logistic functions to vary continuously among countries rather than being restricted to five pre-assigned possibilities, and we assume that the double logistic parameters are draws from a common "world" distribution (Raftery et al 2012a).

The resulting Bayesian model is as follows:

$$\ell_{c,t+1} = \ell_{c,t} + g(\ell_{c,t}|\boldsymbol{\theta}^{(c)}) + e_{c,t}, \quad (7)$$

where

$$g(\ell_{c,t}|\boldsymbol{\theta}^{(c)}) = \frac{k^c}{1 + \exp(-\frac{A_1}{\Delta_2^c}(\ell_{ct} - \Delta_1^c - A_2\Delta_2^c))} + \frac{z^c - k^c}{1 + \exp(-\frac{A_1}{\Delta_4^c}(\ell_{ct} - \sum_{i=1}^3 \Delta_i^c - A_2\Delta_4^c))}. \quad (8)$$

In (8), $\boldsymbol{\theta}^{(c)} = (\Delta_1^c, \Delta_2^c, \Delta_3^c, \Delta_4^c, k^c, z^c)$, $e_{c,t} \stackrel{\text{ind}}{\sim} N(0, \omega(\ell_{c,t})^2)$, where $\omega(\cdot)$ is a smooth function representing how the error standard deviation depends on the current level of life expectancy, and A_1 and A_2 are constants. The country-specific parameters are assumed to be drawn from world distributions, as follows:

$$\Delta_i^c \stackrel{\text{iid}}{\sim} \text{TN}_{[0,100]}(\Delta_i, \sigma_{\Delta_i}^2), \quad i = 1, \dots, 4, \quad (9)$$

$$k^c \stackrel{\text{iid}}{\sim} \text{TN}_{[0,10]}(k, \sigma_k^2), \quad (10)$$

$$z^c \stackrel{\text{iid}}{\sim} \text{TN}_{[0,1.15]}(z, \sigma_z^2), \quad (11)$$

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

The world hyperparameters on the right-hand sides of equations (9)–(11) are given diffuse prior distributions, with one notable exception, namely z . The country-specific parameter z^c is the asymptotic linear increase of life expectancy in country c per five-year period, and this is restricted to be less than 1.15, which is highly informative. This is based on the empirical fact that over the past 170 years the maximum country-specific life expectancy in year y has been increasing highly linearly with y (Oeppen and Vaupel, 2002); 1.15 years per five-year period is the upper bound of a 99.9% confidence interval for the rate of increase. Accordingly, the prior distribution of z is also bounded above by 1.15.

Male life expectancy is highly correlated with female life expectancy and is almost invariably lower. We therefore project female and male life expectancy jointly, by first projecting female life expectancy using the Bayesian hierarchical model described above and then projecting the gap between them. On average the gap tends to increase as a function of female life expectancy as female life expectancy increases up to about 75 years, and then tends to decrease. It also has extreme values, often corresponding to conflicts when male life expectancy is more affected than female. We represent this using the regression model of Lalic (2011) with t -distributed errors for $G_{c,t}$, the gap in country c at time period t .

The model is as follows:

$$G_{c,t+1} = \min\{(G_{c,t+1}^*)_+, 18\},$$

where

$$G_{c,t+1}^* = \begin{cases} \beta_0 + \beta_1 \ell_{c,1950-1955} + \beta_2 G_{c,t} + \beta_3 \ell_{c,t} + \beta_4 (\ell_{c,t} - 75)_+ + \epsilon_{c,t}, & \text{if } \ell_{c,t} \leq M, \\ \gamma_1 G_{c,t} + \epsilon_{c,t}, & \text{if } \ell_{c,t} > M, \end{cases}$$

$$\epsilon_{c,t} \stackrel{\text{iid}}{\sim} t(\mu = 0, \sigma^2 = 0.0665, \nu = 2),$$

and $M = 86.2$ years, the highest life expectancy recorded in the WPP 2010. The gap is restricted to be no more than 18 years, which is slightly above the highest value observed to date. The model was estimated by maximum likelihood.

3.3 Bayesian Population Projections

The methods just described formed the basis for probabilistic population projections for 159 countries, comprising just under 90% of the world's population in 2010. The 38 countries with generalized HIV/AIDS epidemics were not included because they have very different mortality patterns and require special treatment. Thirty small countries or areas with populations under 100,000 were also excluded.

To produce the probabilistic projections, 2,000 trajectories of the total fertility rate for each five-year period from 2010 to 2100, and 2,000 joint trajectories of female and male life expectancy were simulated from their posterior predictive distributions. These were then converted to age-specific fertility and age- and sex-specific mortality rates using established UN methods, and input to the cohort-component method. Current UN assumptions about future international migration were used. This yielded joint probabilistic projections of any future population quantity of interest.

The method was described in more detail by Raftery et al (2012b), which also reported the results of a study assessing the model by estimating it from data for 1950–1990 and using it to project population for all 159 countries in the 20-year period 1990–2010. The projections of total population, total fertility rate, and female and male life expectancy were reasonably accurate and the projection intervals were reasonably well calibrated.

The results for Egypt are shown in Figure 3. For the total fertility rate, the UN high and low variants turn out to be similar to the limits of the Bayesian pointwise 80% projection intervals. For life expectancy, the current UN projections do not provide any assessment of uncertainty or even scenarios. The Bayesian approach suggests higher future life expectancy, but the current UN projection is within the Bayesian 95% interval for most years. The median

Bayesian projection of total population in 2100 is about 10% higher than the UN's WPP 2010 projection, but this has to be seen in the context of the considerable uncertainty about Egypt's total population at century's end. The Bayesian 80% interval for Egypt's population in 2100 ranges from 96 to 184 million.

Perhaps the most striking result is the projected trend in the potential support ratio, equal to the number of people aged 20–64 per person aged 65 or over. This can be roughly interpreted as the number of workers per retiree and is important, for example for planning old-age social security systems. In Egypt this is currently 10.7, but is projected to decline dramatically to 1.4 by the end of the century, with an 80% projection interval 1.0–2.0. For context, in the U.S. this is currently 4.6 and is projected to decline to 1.8 by 2100.

This trend is well known for the U.S. and other developed countries (Lee, 2011) and features in political debate and policy making there. What is perhaps surprising is that the same trend is projected for developing countries with young populations and currently high potential support ratios like Egypt. Indeed the decline is likely to be even steeper in many developing countries than in developed countries. Egypt in 2100 may well have an older population than the U.S. or any other country in the world does now. The projection intervals show that this overall trend is essentially inevitable, even if there is some uncertainty about the extent of the eventual decline.

4 Accommodating Controversy About Ultimate Fertility Level via Bayes

Our method for projecting the total fertility rate for all countries was discussed during a three-day Expert Group Meeting convened by the UN in December 2009 (United Nations Population Division, 2009) and was favorably assessed. The predictive median from our method was then used as the UN's (deterministic) projection of TFR in the WPP 2010 (United Nations, 2011a). Apart from that, the UN used the same deterministic projection

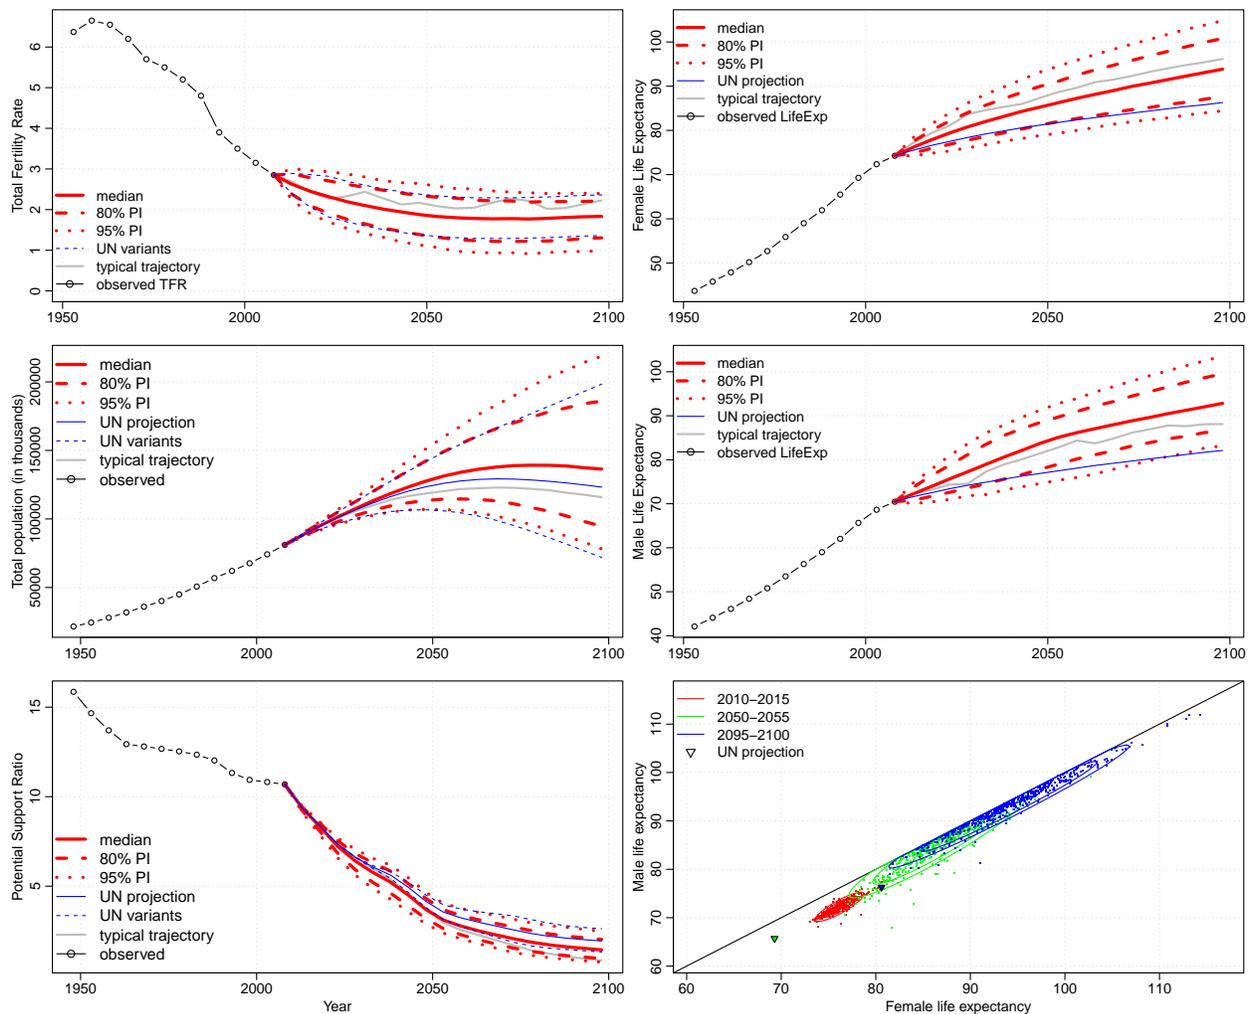


Figure 3: Bayesian probabilistic population projections for Egypt, 2010–2100: major population indicators. Left, top to bottom: Total fertility rate; Total population; Potential support ratio (20–64 population/65+ population). Right, top to bottom: Female life expectancy; Male life expectancy; Joint predictive distribution of female and male life expectancy for 2010–2015, 2050–2055 and 2095–2100. The Bayesian predictive distributions are shown in red: median—solid; 80% projection interval—dashed; 95% projection interval—dotted. The UN WPP 2010 medium projection is shown as a solid blue line and the UN WPP 2010 high and low projections are shown as dashed blue lines. The solid grey line represents a typical trajectory.

method in WPP 2010 as in WPP 2008 and previous projections, but is considering making future projections probabilistic.

Substantively this led to projections of a slower decline in fertility in Africa than had previously been expected. It also led to projections of a slow increase in fertility in Europe, which had previously been projected to remain at the sub-replacement level of 1.85 child per woman once it reached this threshold. The previous UN projection in WPP 2008 went up only to 2050 and projected a world population of 9 billion. The new WPP 2010 projection went up to 2100 for the first time and projected a world population of 10 billion, a billion more (although for a longer time horizon).

Overall, the new projections were well received. However, there was one critique, relating not to the statistical method, but to the assumption in the model for TFR in Phase III that asymptotically TFR oscillates around the approximate replacement rate of 2.1, namely that $\mu = 2.1$ in (6).

Basten et al. (2012) argued that the UN's assumption of an eventual recovery of fertility towards replacement is not justified for five advanced East Asian economies (Korea, Japan, Hong Kong, Singapore and Taiwan). They pointed out that the national statistical agencies of these countries project lower fertility rates than does the UN, that the relevant scientific literature does not suggest an increase in fertility in the short term, that a recent unpublished survey of experts concluded that fertility would not increase as markedly as the UN predicts, and that current evidence about fertility intentions does not suggest an immediate appetite for more children in these countries. They also argued that the UN assumption is based largely on European experience and that there is no reason to assume that it will carry over to East Asia. We interpret their arguments as implying that μ in (6) may differ between East Asian countries and others, and that for the five East Asian countries they consider, μ should be less than the replacement level of 2.1.

We do not necessarily accept this critique, and there are counterarguments. However, here we suggest a possible way to relax the assumptions underlying the Phase III fertility

model (6) so as to accommodate the critique and make the model more fully data-based. Instead of requiring every country in Phase III to follow the same model (6), we allow both μ and ρ to vary between countries following a Bayesian hierarchical model:

$$f_{c,t+1} - \mu_c = \rho_c(f_{c,t} - \mu_c) + \varepsilon_{c,t}, \quad (12)$$

where $\varepsilon_{c,t} \stackrel{\text{iid}}{\sim} N(0, \sigma_\varepsilon^2)$ and

$$\mu_c \sim TN_{[0,\infty)}(\bar{\mu}, \sigma_\mu^2); \quad \rho_c \sim TN_{[0,1]}(\bar{\rho}, \sigma_\rho^2). \quad (13)$$

The prior distributions for the hyperparameters are as follows:

$$\bar{\mu} \sim U[0, 2.1]; \quad \sigma_\mu \sim U[0, 0.318]; \quad \bar{\rho} \sim U[0, 1]; \quad \sigma_\rho \sim U[0, 0.289]; \quad \sigma_\varepsilon \sim U[0, 0.5]. \quad (14)$$

The priors are chosen to be diffuse, except for the prior on $\bar{\mu}$, the world mean of the country-specific asymptotes, which is restricted to be no greater than the replacement level of 2.1. Since the only critiques to date suggest that the UN's current choice of 2.1 may be too high, this truncation seems to accommodate current expert opinion.

Figure 4 shows some of the results of fitting the Bayesian hierarchical model (12)–(14) to the data from the 21 countries that have entered Phase III. The posterior distribution of $\bar{\mu}$, the world mean of the country-specific TFR asymptotes, essentially excludes values below 1.6. For most of the 21 countries, the projections are similar to those from the WPP 2010 model (6) of Alkema et al. (2011) with fixed $\mu = 2.1$; this can be seen, for example, for the Netherlands in the middle panel of Figure 4.

The only one of Basten et al. (2012)'s five advanced East Asian economies that has entered Phase III is Singapore, for which the projections are shown in the left panel of Figure 4. Remarkably, it is also the only one of the 21 Phase III countries for which the projections differ substantially between the $\mu = 2.1$ model and the Bayesian hierarchical

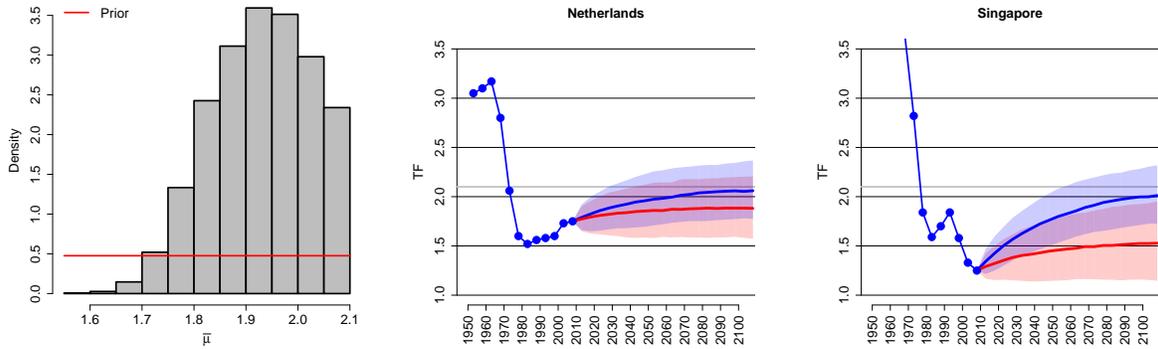


Figure 4: Results for the Bayesian hierarchical model for Phase III fertility (12)–(14): Left: Posterior distribution of $\bar{\mu}$, the world mean of the country-specific asymptotes. Center and right: Projections of the TFR with 80% projection intervals for the Netherlands and Singapore under the WPP 2010 model (6) in blue and the Bayesian hierarchical model for Phase III in red.

model with country-specific μ_c . The projection under the Bayesian hierarchical model (12)–(14) is much lower than under the previous model (6), as Basten et al argued it should be. It asymptotes at 1.5 instead of 2.1. This suggests that the data provide some support for Basten et al’s contention, and also that the proposed Bayesian hierarchical model for Phase III can accommodate differences of this kind among countries.

5 Discussion

We have developed a method for probabilistic population projections for possible use by the United Nations in its biennial projections of the populations of all countries. We have developed Bayesian hierarchical models for projecting future overall levels of fertility and mortality for each country, as measured by the total fertility rate and life expectancy for females and males. Samples from the resulting posterior predictive distributions are converted into age-specific fertility and mortality rates. These are then used as inputs to the cohort component projection method, yielding samples from the posterior predictive distributions of any future population quantity of interest. The method has performed reasonably well in terms of out of sample predictive performance.

The method focuses on the overall levels of fertility and mortality, traditionally viewed as the biggest sources of error in population projections. However, there are other sources of uncertainty not taken into account in our method. These include uncertainty about the age structure of fertility and mortality, and also about net international migration, the latter being increasingly important (National Research Council, 2000). Not including these does not seem to have led to substantial miscalibration of the method, but future work should aim to incorporate these other sources of uncertainty. Also, the method is not currently adapted to the countries with generalized HIV/AIDS epidemics which comprise about 10% of the world's population, and it will be important to extend the method to these countries.

The UN projects populations not only for all countries, but also for regions and other sets of countries including trading blocs and economic, political and ecological groupings (United Nations, 2011b). Aggregation of deterministic projections is straightforward: just add up the projections for the component countries. For probabilistic projections it is not so simple, because between-country correlations have to be taken into account. Our approach treats countries as exchangeable rather than independent, but it does not model the additional correlation that may exist between, for example, contiguous countries. For life expectancy, work to date suggests that our model does account for most of this correlation; there is little correlation between forecast errors (Raftery et al 2012a). For fertility there may be some excess between-country correlation, and this needs to be accounted for in future work.

Our approach is largely Bayesian, and Bayesian thinking was essential in overcoming the statistical challenges, including the limited amount of data for each country and the differences and similarities among countries in the way fertility and life expectancy have evolved. The Bayesian approach allowed us to borrow strength from other countries through the hierarchical model, thus avoiding instabilities in estimation. It also gave us a way to combine the posterior predictive distributions of fertility and mortality with the cohort component model in a natural way, and to incorporate external information about the asymptotic rate of increase in life expectancy through its prior distribution. In addition, it provided the basis

for a way to accommodate Basten et al. (2012)’s critique of the UN’s assumption about the equilibrium distribution of total fertility rate. Fienberg (2011) described several other major governmental and policy problems for which Bayesian thinking proved useful.

We have adopted a fully Bayesian approach, however, only when it led to an improved solution. For example, our model for the female-male gap in life expectancy is not Bayesian, because there are enough data to estimate the model reliably via maximum likelihood. Fully Bayesian estimation of this model would give similar results and would be more complicated, and so did not seem worth doing.

Several other Bayesian approaches to population projection have been proposed, but these have focused mainly on projecting mortality. Girosi and King (2008) proposed a Bayesian method for forecasting age-specific mortality that can incorporate covariates. They showed that for short-term forecasts, their method outperformed the widely used time series method of Lee and Carter (1992) (without covariates) on average for 48 countries with better mortality data. This result, obtained when covariates were used, requires additional data that may not be reliable or even available in many countries. The reliability of their approach remains unproven for medium- or long-term projections due to the reliance on covariates and the difficulties in predicting them beyond a few decades. They did not give probabilistic forecasts, although their method may in principle be able to provide them.

Czado et al. (2005) developed a Bayesian method for estimating the Poisson log-bilinear formulation of Brouhns et al. (2002)’s version of the Lee-Carter model. 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. 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. Daponte et al. (1997) developed a Bayesian approach to the problem of reconstructing past populations, which is different from the problem of projecting future populations that we address here.

Expert-based probabilistic population projections have been produced by Lutz and col-

leagues (Lutz et al., 1998, 2004, 2008). However, this method is not explicitly based on available data, and instead relies on a collection of experts and their ability to specify probabilistic bounds, that may or may not be accurate (Alho, 2005). Thus, while like Bayesian approaches this method uses expert knowledge, it does not update it formally using data, and so is not Bayesian in the usual sense.

Our model for the total fertility rate has been adopted by the UN as the basis for its deterministic projections in WPP 2010 (United Nations, 2011a). The UN also issued probabilistic projections using the methods described here on an experimental basis in November 2012, at <http://esa.un.org/unpd/ppp>. The UN is considering issuing official probabilistic projections for the first time in WPP 2014, using our methods.

The methods described here are implemented in the R packages `bayesTFR` (Ševčíková et al., 2011), `bayesLife` (Ševčíková and Raftery, 2011), `bayesPop` (Ševčíková and Raftery, 2012) and `bayesDem` (Ševčíková, 2011).

Acknowledgements

This work was supported by grants R01 HD054511 and R01 HD070936 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) and a research grant from the National University of Singapore. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of NICHD. Also, the views expressed in this article 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 article 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 John Bongaarts, Thomas Buettner, Samuel Clark, Joel Cohen, Gerhard Heilig,

Ronald Lee, Nan Li, Sharon McGrayne, Hana Ševčíková, Hania Zlotnik, participants in the Bay Area Colloquium in Population (BACPOP), the editor, the associate editor and two referees for very helpful comments and discussions.

References

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

Alkema, L., A. E. Raftery, P. Gerland, S. J. Clark, F. Pelletier, T. Buettner, and G. K. Heilig (2011). Probabilistic projections of the total fertility rate for all countries. *Demography* 48, 815–839.

Basten, S. A., D. A. Coleman, and B. Gu (2012). Re-examining the fertility assumptions in the UN’s 2010 World Population Prospects: Intentions and fertility recovery in East Asia? Presented at the Annual Meeting of the Population Association of America, San Francisco. <http://paa2012.princeton.edu/sessionViewer.aspx?SessionId=112>.

Brouhns, N., M. Denuit, and J. K. Vermunt (2002). A Poisson log-bilinear approach to the construction of projected lifetables. *Insurance: Mathematics and Economics* 31, 373–393.

Caswell, H. (2006). *Matrix Population Models: Construction, Analysis and Interpretation*. Sunderland, Mass.: Sinaurer Associates Inc.

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

Daponte, B. O., J. B. Kadane, and L. J. Wolfson (1997). Bayesian demography: Projecting the Iraqi Kurdish population, 1977-1990. *Journal of the American Statistical Association* 92, 1256–1267.

El-Badry, M. A. and S. Kono (1986). Demographic estimates and projections. *Population Bulletin of the United Nations* 19–20, 35–43.

- Fienberg, S. E. (2011). Bayesian models and methods in public policy and government settings (with discussion). *Statistical Science* 26, 212–239.
- Giroi, F. and G. King (2008). *Demographic Forecasting*. Princeton, NJ: Princeton University Press.
- Hirschman, C. (1994). Why fertility changes. *Annual Review of Sociology* 20, 203–233.
- Intergovernmental Panel on Climate Change (2007). *Climate Change 2007: Synthesis Report*. Geneva, Switzerland: IPCC.
- Keilman, N., D. Q. Pham, and A. Hetland (2002). Why population forecasts should be probabilistic – illustrated by the case of Norway. *Demographic Research* 6, 409–454.
- Keyfitz, N. and H. Caswell (2005). *Applied Mathematical Demography* (3rd ed.). New York: Springer.
- Lalic, N. (2011). Joint probabilistic projection of female and male life expectancy. Master’s thesis, Department of Statistics, University of Washington, Seattle, Wash.
- Lee, R. D. (2011). The outlook for population growth. *Science* 333, 569–573.
- Lee, R. D. and L. Carter (1992). Modeling and forecasting the time series of US mortality. *Journal of the American Statistical Association* 87, 659–671.
- Lee, R. D. and S. Tuljapurkar (1994). Stochastic population forecasts for the United States: Beyond high, medium, and low. *Journal of the American Statistical Association* 89, 1175–1189.
- Leslie, P. H. (1945). On the use of matrices in certain population dynamics. *Biometrika* 33, 183–212.

- Lutz, W. and K. C. Samir (2010). Dimensions of global population projections: what do we know about future population trends and structures? *Philosophical Transactions of the Royal Society B* 365, 2779–2791.
- Lutz, W., W. C. Sanderson, and S. Scherbov (1998). Expert-based probabilistic population projections. *Population and Development Review* 24, 139–155.
- Lutz, W., W. C. Sanderson, and S. Scherbov (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., W. C. Sanderson, and S. Scherbov (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>).
- Myrskylä, M., H. P. Kohler, and F. C. Billari (2009). Advances in development reverse fertility declines. *Nature* 460, 741–743.
- National Research Council (2000). *Beyond Six Billion: Forecasting the World’s Population*. Washington, D.C.: National Academy Press.
- Oeppen, J. and J. W. Vaupel (2002). Broken limits to life expectancy. *Science* 296, 1029–1031.
- Pedroza, C. (2006). A Bayesian forecasting model: Predicting U.S. male mortality. *Biostatistics* 7, 530–550.
- Preston, S. H., P. Heuveline, and M. Guillot (2001). *Demography: Measuring and Modeling Population Processes*. Malden, Mass.: Blackwell.
- Raftery, A. E., J. L. Chunn, P. Gerland, and H. Ševčíková (2012a). Bayesian probabilistic projections of life expectancy for all countries. *Demography* 49, in press.

- Raftery, A. E., N. Li, H. Ševčíková, P. Gerland, and G. K. Heilig (2012b). Bayesian probabilistic population projections for all countries. *Proceedings of the National Academy of Sciences* 109, 13915–13921.
- Seto, K. C., B. Güneral, and L. R. Hutyrá (2012). Global forecasts of urban expansion to 2030 and direct impacts on biodiversity and carbon pools. *Proceedings of the National Academy of Sciences* 109, 16083–16088.
- United Nations (2006). *World Population Prospects: The 2004 Revision. Volume III: Analytical Report*. New York, NY: United Nations.
- United Nations (2009). *World Population Prospects: The 2008 Revision*. New York, NY: United Nations.
- United Nations (2011a). *World Population Prospects: The 2010 Revision*. New York, NY: United Nations.
- United Nations (2011b). *World Population Prospects: The 2010 Revision. Special Aggregates – DVD-ROM Edition – Dataset in Excel and ASCII format*. United Nations Publication ST/ESA/SER.A/311, Population Division, Department of Economic and Social Affairs, United Nations. http://esa.un.org/unpd/wpp/Other-Information/WPP2010_Special%20Aggregates%20-%20list%20of%20groupings.pdf.
- United Nations Population Division (2009). Expert group meeting on recent and future trends in fertility. <http://www.un.org/esa/population/meetings/EGM-Fertility2009/egm-fertility2009.html>. Downloaded August 29, 2012.
- Ševčíková, H. (2011). *bayesDem: Graphical User Interface for bayesTFR, bayesLife and bayesPop*. R package version 1.6-0.
- Ševčíková, H., L. Alkema, and A. E. Raftery (2011). bayesTFR: An R package for probabilistic projections of the total fertility rate. *Journal of Statistical Software* 43, 1–29.

Ševčíková, H. and A. E. Raftery (2011). *bayesLife: Bayesian Projection of Life Expectancy*.
R package version 0.4-0.

Ševčíková, H. and A. E. Raftery (2012). *bayesPop: Probabilistic Population Projection*. R
package version 1.0-3.

Whelpton, P. K. (1936). An empirical method for calculating future population. *Journal of
the American Statistical Association* 31, 457–473.