Probabilistic Projections of the Total Fertility Rate

Leontine Alkema, Adrian E. Raftery,
Patrick Gerland, Samuel J. Clark, François Pelletier *

August 13, 2009

*Leontine Alkema, Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546; Email: alkema@nus.edu.sg. Adrian E. Raftery, Departments of Statistics and Sociology, University of Washington, Seattle, WA 98195-4320; Email: raftery@u.washington.edu. Patrick Gerland, Population Estimates and Projections Section, United Nations Population Division, New York, NY 10017; Email:gerland@un.org. Samuel J. Clark, Department of Sociology, University of Washington, Seattle, WA 98195-3340; Email: samclark@u.washington.edu. François Pelletier, Mortality Section, United Nations Population Division, New York, NY 10017; Email: pelletierf@un.org. The project described was partially supported by Grant Number 1 R01 HD054511 01 A1 from the National Institute of Child Health and Human Development. 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 or those of the United Nations. Its contents have not been formally edited and cleared by the United Nations. The authors are grateful to Thomas Buettner, Gerhard Heilig and Taeke Gjaltema for helpful discussions and insightful comments. Alkema thanks the United Nations Population Division for hospitality.

Abstract

This paper discusses a Bayesian projection model to construct country-specific probabilistic projections of the total fertility rate (TFR), for all countries in the world. Our methodology for high fertility countries builds onto the one currently used by the United Nations Population Division, which assumes that fertility will eventually fall below replacement level. In the model, the pace of the fertility decline is decomposed into a systematic decline with distortion terms added to it. The pace of the systematic decline is modeled as a function of its level, based on the UN methodology. We propose a Bayesian hierarchical model to estimate the parameters of the decline function.

For low fertility countries, and to project what will happen after the fertility transition has ended, we assume that the TFR will fluctuate around 2.1 children in long term projections. This is modeled with a first order autoregressive time series model.

The result of the Bayesian projection model is a set of future TFR trajectories for each country. The "best" projection for the TFR is given by the median TFR outcome in each period. The "low" and "high" variant are given by the 10th and 90th percentile of the sample of TFR outcomes in each period, the lower and upper bounds of the 80% prediction intervals.

This new approach provides valuable insights about future fertility trends worldwide. The prediction intervals for future fertility levels vary by country. The intervals are wider in high-fertility countries than those currently inferred with the low and high variants of the official UN population projections. The projected TFRs and the corresponding prediction intervals will shed new light on future population dynamics, including on dependency ratios and on the pace of population ageing.

1 Introduction

Population forecasts predict the future size and composition of populations, based on predictions of fertility, mortality and migration. They are used for many purposes, including for predicting the demand for food, water, education, medical services, labor markets, pension systems, and predicting future impact on the environment. It is important for decision makers to not only have a point forecast that states the most likely scenario of a future population, but also to know the uncertainty around it, that is, the possible future values of an outcome, and how likely each set of possible future values is.

Fertility is a key driver of the size and composition of the population. Fertility decline has been a primary determinant of population ageing and projected levels of fertility have important implications on the age structure of future populations, including on the pace of population ageing. The total fertility rate (TFR) is one of the key components in population projections; it is the average number of children a woman would bear if she survived through the end of the reproductive age span, and experienced at each age the age-specific fertility rates of that period. The UN Population Division produces projections of the total fertility rate for 196 countries that are revised every two years and published in the World Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division ming). For countries with above-replacement fertility, a demographic transition model is used to project the decline in the total fertility rate and assumes that fertility will eventually fall below replacement level. Three sets of parameter values describe three different trajectories of future declines, from which the UN analyst chooses one which seems most appropriate for the country of interest. The UN projections for countries that are currently experiencing below-replacement fertility are constructed based on the assumption that fertility will increase again towards replacement level, to stabilize at 1.85 children. Fertility is assumed to increase linearly at a maximum rate of 0.05 children per woman per quinquennium.

While using the cohort-component method, the TFR projection, together with projections of mortality and international migration, provide the so-called *Medium* variant of the official United Nations population projections. The effect of lower or higher fertility when projecting populations is illustrated with the *Low* and *High* variants of the projections. In the high variant, half a child is added to the medium variant in order to examine the influence of a slower fertility decline on the population projections. Similarly, for the low variant, half a child is subtracted from the medium variant.

Though useful to highlight the sensitivity that one child difference makes on demographic outcomes, the drawback of the variants is that they do not assess the uncertainty in future fertility levels (Bongaarts and Bulatao 2000), and to what extent the low or high fertility variants are more likely. Future levels of fertility will be more uncertain in countries where the fertility transition has only just started than in countries where fertility is close to replacement level. A shortcoming with the current projection methodology is that the rate of change implied in the projections is not sufficiently country-specific; only three options for modeling the future rate of change as a function of the fertility level are considered, from which one is chosen for each country. This means that the current approach works well for capturing the average experience of groups of countries which experience a similar pace of decline at the same fertility level, but it is less adequate to depict much slower or faster declines deviating from the typical group average experiences.

In this paper we develop methodology to construct probabilistic projections of the TFR for all countries in the world. Our methodology builds onto the one currently used by the United Nations Population Division for projecting the TFR. For countries that are going through the fertility transition from high fertility towards replacement fertility, the pace of the fertility decline is decomposed into a systematic decline, with distortion terms added to it. The pace of the systematic decline in TFR is modeled as a function of its level, based on the UN methodology. We propose a Bayesian hierarchical model to estimate the parameters of the decline function. A time series model is used for projecting trends in fertility after reaching replacement level, assuming that in long-term projections the TFR will fluctuate around replacement level fertility. The results are country-specific projections that are reproducible and take into account past trends.

This new approach provides valuable insights about future fertility trends worldwide. The prediction intervals for future fertility levels vary by country. The intervals are wider in most high-fertility countries than those currently inferred with the low and high variants of the official UN population projections. The projected TFRs and the corresponding prediction intervals will shed new light on future population dynamics, including on dependency ratios and on the pace of population ageing.

This paper is organized as follows; In Section 2.1 we discuss the model as used by the UN Population Division to project the total fertility rate, Section 2.2 explains the projection model to construct probabilistic, country-specific projections. In Section 3 we present the results of this model and Section 4 discusses possible improvements on the methodology.

In this paper UN estimates and projections of the TFR are taken from the 2008 revision of the UN World Population Prospects (United Nations, Department of Economic and Social Affairs, Population Division ming).

2 Methods

2.1 UN methodology for projecting fertility

The UN Population Division estimates and projects the TFR for five-year time periods from 1950 until 2050 (in the most recent revision). Five year intervals are chosen such that the estimates and projections can be used as input in the cohort-component projections, which are based on 5-year age groups.

A demographic transition model is used to project a fertility decline for countries in which the TFR is above 2.1 children for each woman (which is equal to replacement level fertility for countries with low mortality rates). In this model, the TFR is predicted to decline because of decreasing child mortality and economic development. The UN projects that fertility will decline towards 1.85 children for each woman. This assumption is based on what has been observed in countries that have gone through their fertility transition. The pace of the future fertility decline is modeled as a function of the level of the TFR, also based on what has been observed in countries that have gone through (most of) their fertility transition. This is illustrated for Thailand and India in Figure 1. The plot on the left shows the 5-year UN estimates for Thailand and India over time, with $f_{c,t}$ the TFR for country c, 5-year period t. Thailand went through its fertility transition relatively fast compared to other countries. The fertility transition in India has not been completed yet; its TFR has decreased from around 6 to 3 children and is still declining. The pace of the fertility decline during the transition is modeled as a function of its level in terms of 5-year decrements, which are the decreases in TFR in a 5-year period. The 5-year decrements as observed in Thailand and India are plotted against TFR in Figure 1(b). The TFR decreases on the horizontal axis such that when examining a fertility decline, the decline curve as given by the 5-year decrements is followed from left to right over time. The decline curves in Thailand and India shows the typical pattern of a fertility decline that starts slowly at high TFR values. The pace increases and is at its maximum around a TFR of 5 children per woman, and then slows down again towards the end of the transition.

The UN uses a parametric function to project the next 5-year decrement given a certain level of fertility, whose shape is similar to the curves as observed in Thailand and India (United Nations, Department of Economic and Social Affairs, Population Division 2006). The UN projection model is given by:

$$f_{c,t+1} = f_{c,t} - d(\boldsymbol{\theta}, f_{c,t}), \tag{1}$$

with $d(\cdot,\cdot)$ the parametric decline function to model the fertility transition. This function gives a 5-year decrement (decrease) as a function of the current level of the TFR and parameter vector $\boldsymbol{\theta}$. The decline function itself is given by the sum of two logistic functions, a double logistic function (Meyer 1994). The first logistic function describes a high pace of decline at high total fertility rates decreasing towards a slower pace for lower fertility. The second function describes the opposite effect to slow down the pace of fertility decline at the beginning of the transition. The sum of the two is a parametric function with 6 parameters that describes a decline in fertility that starts with a slow pace at a high TFR, peaks around

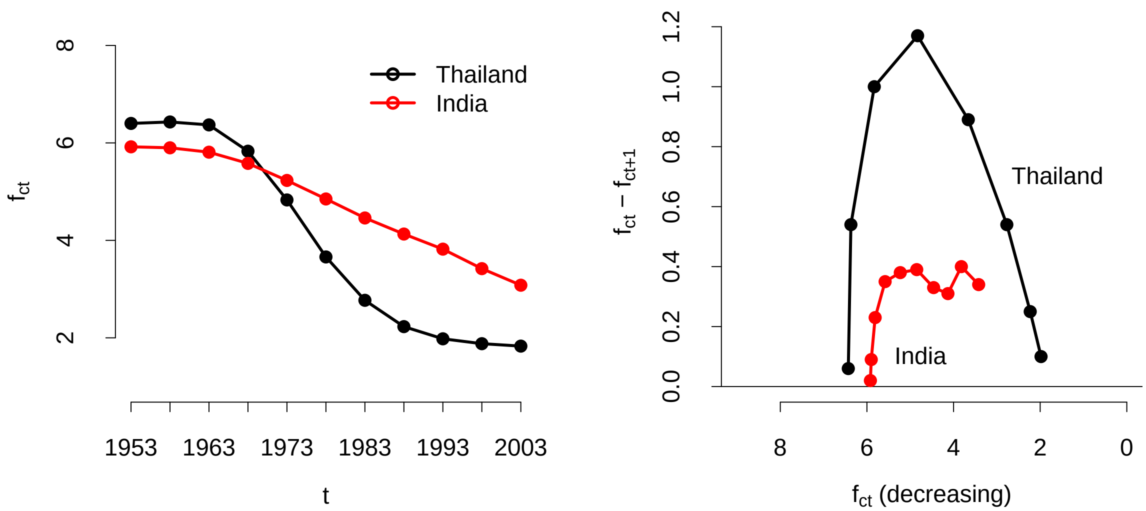

Figure 1: UN estimates of the fertility transition in Thailand and India: (a) Total fertility rate $f_{c,t}$ versus 5-year period t, (b) 5-year decrements $f_{c,t} - f_{c,t+1}$ versus total fertility rate $f_{c,t}$.

a TFR of 5 and slows down again at lower TFR outcomes. In the UN projections, the parameter vector is chosen from a set of 3 different vectors, with each parameter vector giving a different paced fertility decline: $\boldsymbol{\theta} \in \{\boldsymbol{\theta}_{SS}, \boldsymbol{\theta}_{FS}, \boldsymbol{\theta}_{FF}\}$. These parameter vectors have been estimated based on fertility declines in countries that have completed the fertility transition (United Nations, Department of Economic and Social Affairs, Population Division 2006). The subscripts of $\boldsymbol{\theta}$ refer to the pace at the start and the end of the fertility decline, with "S" meaning slow, and "F" meaning fast. The decline functions, as given by these three parameter vectors are shown in Figure 2. In the figure, the Fast/Slow decline curve is given by the solid line. Compared to the Fast/Slow decline curve, the Slow/Slow decline curve gives a slower-paced decline at the start of the transition, the Fast/Fast trajectory a faster pace at the end. For all three projected declines, the TFR is kept constant after it reaches 1.85 children.

For each country, the UN analyst chooses the decline curve that seems most reasonable for the future fertility decline in that country, based on what has been observed in that country or region so far, or based on expert knowledge about the country. Generally, projected fertility paths yielded by the UN models are checked against recent trends in fertility for each country. When a country's recent fertility trends deviate considerably from the standard decline curves, fertility is projected over an initial period of 5 or 10 years in such a way that it follows recent experience. The model projection takes over after that transition period. For instance, in countries where fertility has stalled or where there is no evidence of fertility decline, fertility is projected to remain constant for several more years before a declining path sets in.

Note that the double logistic model does not predict the onset of the fertility transition,

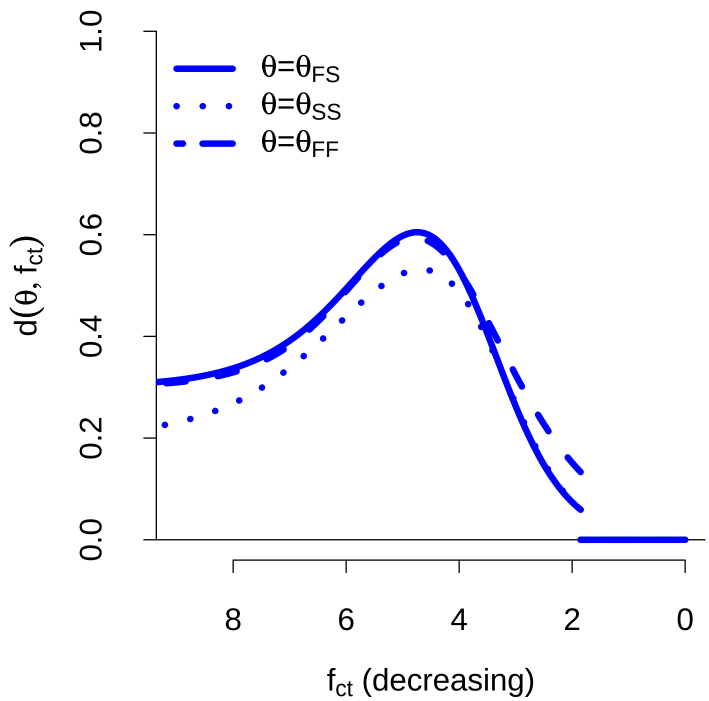

Figure 2: The UN decline curves that underlie the fertility projections for countries with above-replacement fertility. Each curve is given by the double logistic decline function with one choice of the parameter vector $\boldsymbol{\theta}$.

it gives the pace of the decline after its onset. In order to predict future fertility levels in countries for which a decline has not yet been observed, additional assumptions are needed about the timing of the onset of the decline, e.g. the decline takes off in the next five or ten years.

Several countries (e.g. in Europe and Asia) are currently experiencing below-replacement fertility. The UN projections for these countries are constructed based on the assumption that fertility will increase again and will stabilize at 1.85 children. For these countries it is assumed that over the first 5 or 10 years of the projection period fertility will follow the recently observed trends. After that transition period, fertility is assumed to increase linearly at a rate of 0.05 children per woman per quinquennium until it reaches 1.85 children per woman. For countries with very low fertility, replacement does not need to be reached by 2050.

2.2 Bayesian projection model

There are some drawbacks of the UN projection model. It is a deterministic model, thus there is no uncertainty assessment of the projections. Secondly, the projections for the high fertility countries are based on choosing the parameter vector $\boldsymbol{\theta}$ of the decline function from a set of three vectors. This results in projections that are not country-specific. Moreover, the three sets of parameter values do not capture the variation in the past. This is illustrated in Figure 3. This figure shows the decrement curves as observed in Thailand and India, with the outcomes of the UN decline function for the three parameter vectors. The UN decline curves do not differ much at all compared to the observed decrements in Thailand and India.

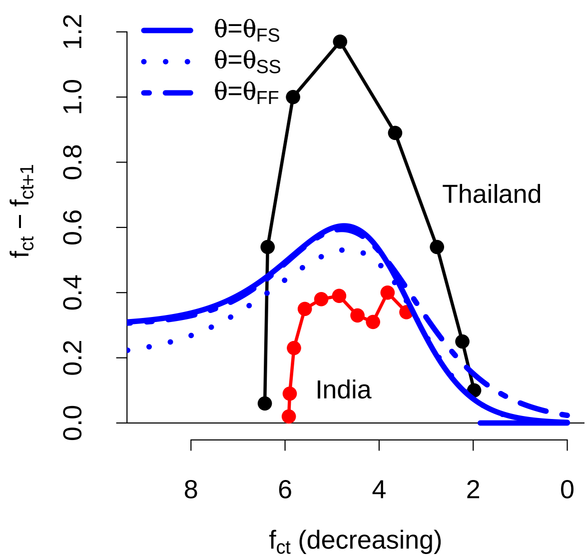

Figure 3: Comparison of the UN decline curves with the observed decrements in Thailand and India.

Our objective is to construct country-specific probabilistic projections of the TFR; during the fertility transition from high towards replacement fertility and after replacement fertility has been reached. The projection model for the fertility transition is based on the UN model, with modifications to overcome its drawbacks. The UN model is modified as follows: (i) instead of only having 3 options for parameter vector $\boldsymbol{\theta}$ of the decline function, it will be estimated for each country separately, (ii) an uncertainty assessment is included by allowing for random distortions from the parametric decline curve, and by assessing the uncertainty in $\boldsymbol{\theta}$ for each country. For low fertility projections the main assumption is that the TFR will converge towards and fluctuate around replacement level fertility.

The next sections discuss the methodology for high fertility and low fertility projections. The result of the Bayesian projection model (BPM) is a set of future TFR trajectories for each country. The "best" projection for the TFR is given by the median TFR outcome in each period. The "low" and "high" variant are given by the 10th and 90th percentile of the sample of TFR outcomes in each period, the lower and upper bounds of the 80% prediction intervals.

2.2.1 Fertility transition

For countries that are currently going through the fertility transition, the 5-year decrements are decomposed into a systematic decline, with a distortion term added to it. More formally, the TFR for 5-year periods is modeled with a random walk model with drift, defined by:

$$\begin{aligned} f_{c,t+1} &= f_{c,t} - d_{c,t} + \varepsilon_{c,t}, \qquad (2)\\ \varepsilon_{c,t} &\sim N(0, \sigma^2_{c,t}), \qquad (3) \end{aligned}$$

with $d_{c,t}$ the drift term which models the systematic decline during the fertility transition, and $\varepsilon_{c,t}$ the random distortions, which model the derivations from the systematic decline. The expression for the standard deviation $\sigma_{c,t}$ is based on examination of the absolute distortions as a function of the TFR level (which showed a higher variance around a TFR of 4-5) and over time (which showed a higher variance before 1975, fertility transitions have become more predictable since 1975 which is possibly explained by more attention being devoted to family planning programs). The standard deviation function is given in the Appendix.

The drift term $d_{c,t}$ gives the 5-year decrement during the fertility transition. A slightly modified version of the double logistic function, as used by the UN, is chosen as the decline function to model the decrements. The decrements are given by $d_{c,t} = d(\boldsymbol{\theta}_c, \lambda_c, \tau_c, f_{c,t})$, with

$$d(\boldsymbol{\theta}_c, \lambda_c, \tau_c, f_{c,t}) = \begin{cases} g(\boldsymbol{\theta}_c, f_{c,t}) & \text{for } \tau_c \leq t \leq \lambda_c; \\ 0 & \text{otherwise,} \end{cases} \qquad (4)$$

with τ_c the start period of the fertility transition, λ_c its end period and $g(\boldsymbol{\theta}_c, f_{c,t})$ the double logistic function with country-specific parameter vector $\boldsymbol{\theta}_c = (\triangle_{c1}, \triangle_{c2}, \triangle_{c3}, \triangle_{c4}, d_c)$, given by

$$\frac{-d_c}{1 + \exp\left(-\frac{2\ln(9)}{\triangle_{c1}}(f_{c,t} - \sum_i \triangle_{ci} + 0.5\triangle_{c1})\right)} + \frac{d_c}{1 + \exp\left(-\frac{2\ln(9)}{\triangle_{c3}}(f_{c,t} - \triangle_{c4} - 0.5\triangle_{c3})\right)}.$$

Figure 4 illustrates the parametrization of the double logistic function. The 5-year decrements as given by the decline function are plotted against TFR. The maximum pace of the decline (the maximum 5-year decrement) is given by d_c. Note that the actual attained maximum pace tends to be slightly smaller than d_c, it depends on the four $\triangle_{ci}$'s, which describe the ranges of the TFR in which the pace of the fertility decline changes. The decline takes off at TFR level $U_c = \sum_{i=1}^4 \triangle_i$, where the pace is around 10% of its maximum pace ($0.1d_c$). During $\triangle_{c1}$, or more correctly, from TFR level U_c to $U_c - \triangle_{c1}$, the pace of the decline increases from around $0.1d_c$ to over $0.8d_c$. During the TFR range denoted by $\triangle_{c2}$, the TFR is declining at a higher pace than during the rest of the transition; its 5-year decrements range between $0.8d_c$ and d_c. In $\triangle_{c3}$ the pace of the fertility decline decreases further to below $0.1d_c$. The asymptotic level of the TFR is given by $\triangle_{c4}$, at which the drift term becomes

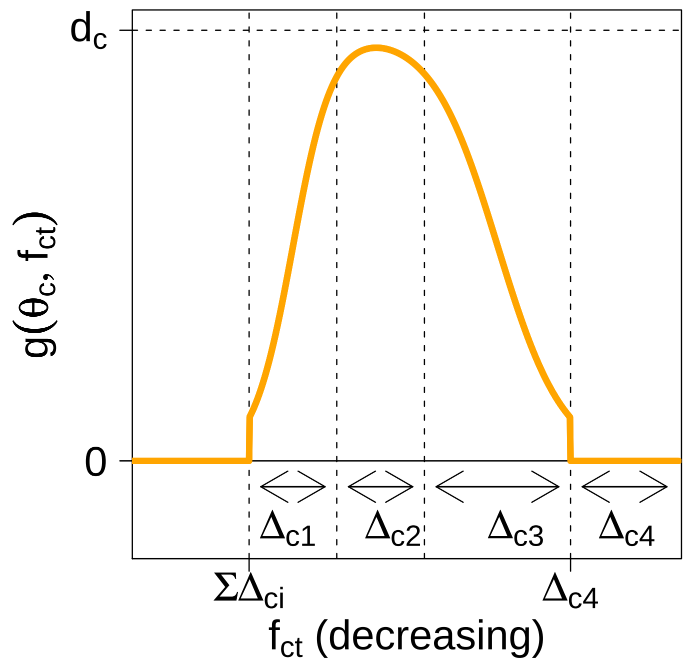

Figure 4: 5-year decrements as given by the double logistic function plotted against the TFR. Note that the horizontal TFR axis is negatively oriented (i.e. decreasing).

zero. With this parametrization, the fertility transition starts in period τ_c at level U_c and ends in λ_c at level $\triangle_{c4}$.

The double logistic function is chosen to project the pace of the fertility decline during the fertility transition because of (i) the straightforward interpretation of its parameters, (ii) to keep consistency with the current UN methodology, and (iii) because of its ability to represent various declines by varying the values of the maximum decrement, and the $\triangle_{ci}$'s. This is illustrated in Figure 5, which shows the observed decrements in Thailand and India, as discussed earlier. The orange line gives the least-squares fit of the parameters of the decline function as described above to these decrements to illustrate the flexibility to model various declines.

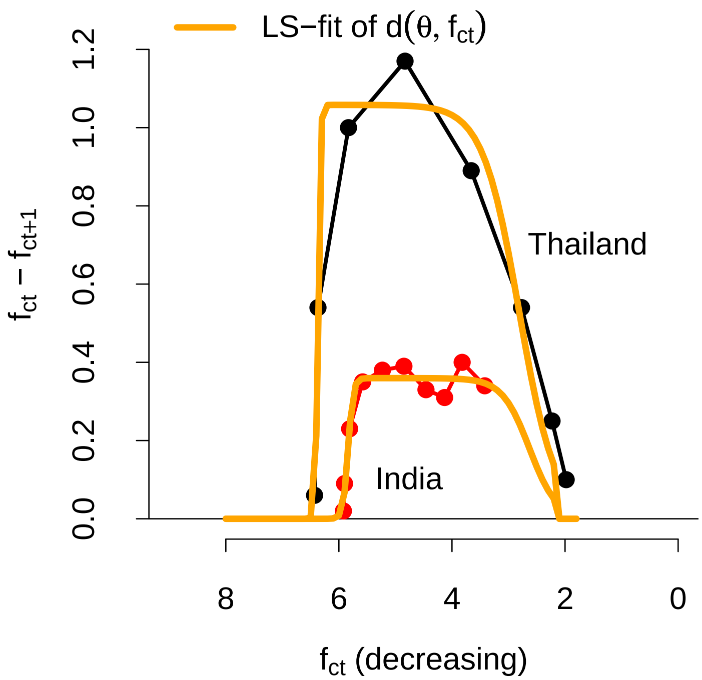

Figure 5: Observed 5-year decrements in Thailand and India, with least-squares fits of the double logistic decline function that is used in the time series projection model.

The parameters of the decline function will be estimated for each country. The start period of the transition, τ_c, is given by the period in which the TFR starts declining and its outcome can be determined before fitting the projection model. For countries in which the fertility decline starts after 1950, its start level U_c is observed and given by:

$$U_c = f_{c,\tau} \text{ if } \tau_c \geq 1950. \quad (5)$$

For countries in which the decline started before 1950, the start level U_c is added as a parameter to the model (for details on τ_c and U_c, see Appendix).

The fertility transition ends at a TFR level that is equal to parameter $\triangle_{c4}$. Based on the UN estimates, 50 countries have experienced a minimum below 2.1 children, a turn-around point such that the TFR is increasing again. For these countries $\triangle_{c4}$ is equal to the TFR at the turn-around point. The mean of the observed $\triangle_{c4}$'s is 1.49 (the median 1.45). There are an additional 7 countries in which no minimum has been observed yet, but for which the current TFR is lower than 1.5 (so these countries would bring down the mean/median of the $\triangle_{c4}$'s even more). Based on the current mean of the minima, we set $\triangle_{c4} = 1.5$ for all countries in which no minimum has been observed yet, which gives:

$$\begin{aligned} \triangle_{c4} &= \min_t f_{c,t}, \text{ for } c \in S_T && (6) \\ &= 1.5 \text{ otherwise,} && (7) \end{aligned}$$

with S_T the set of countries in which a turn-around point below 2.1 children has been observed. The end period λ_c is given by the first period in which the TFR decreases below $\triangle_{c4}$:

$$\lambda_c = \min\{t : f_{c,t} < \triangle_{c4}\} \quad (8)$$

The additional 4 parameters in the double logistic function that determine the pace of the decline are given by $(\triangle_{c1}, \triangle_{c2}, \triangle_{c3}, d_c)$; the ranges of TFR values in which the pace is at its maximum and decreases from its maximum to zero, and the maximum pace of the fertility decline. To estimate these parameters in high-fertility countries (esp. in countries where not so much of the decline has been observed yet), and assess their uncertainty, we assume that these parameters are exchangeable between countries and use a Bayesian hierarchical model to derive the country-specific distributions (Gelman et al. 2004). This means that the predicted systematic part of the fertility decline in a country is based on its observed decline so far, as well observed declines in all other countries. The details on the hierarchical model are discussed in the Appendix.

Ideally, empirical data would be used to fit the projection model to and estimate the parameters of the systematic trend as given by the drift term, and the variance of the distortion terms. However, empirical data are not available in a standard format for most countries. Also, for most developing countries, issues with data quantity and quality require extra attention. To overcome this problem, the 5-year UN estimates are used as the data set of TFR observations. Using the UN estimates allows for constructing prediction intervals for all countries, based on the declines and trends that have been observed so far in all countries. We assume that the UN estimates for period t, denoted by $u_{c,t}$, are equal to the TFR:

$$u_{c,t} = f_{c,t}. \quad (9)$$

In reality the UN estimates are measured/estimated with error. No sampling model for the estimates is included here, because the error variance of the UN estimates cannot be estimated based on single 5-year estimates for each country. This means that the prediction intervals as constructed in this paper are expected to be more narrow than intervals in which the additional error variance would have been accounted for.

2.2.2 Low fertility projections

The model as discussed above is used for modeling and projecting the TFR during the fertility transition. After the fertility transition has been completed we assume that in long-term projections the TFR will converge towards and fluctuate around replacement level fertility (around 2.1 children). This is modeled with a first order autoregressive time series model, an AR(1) model, with its mean fixed at replacement fertility $\mu = 2.1$. This model is given by:

$$f_{c,t} \sim N(\mu + \rho(f_{c,t-1} - \mu), s^2) \text{ for } t > \lambda_c \tag{10}$$

with ρ the autoregressive parameter with $|\rho| < 1$, and s^2 the variance of the random errors. In this model the expected increase (decrease) towards 2.1 is larger if the current TFR is further away from 2.1, and depends on ρ. For example, at a TFR of 1.5, the expected next TFR is $2.1 - 0.6\rho$; a smaller ρ will give a larger expected increase. The smaller ρ, the faster the TFR will increase towards replacement level fertility.

The autoregressive parameter ρ is estimated based on the UN estimates after a turn-around point has been observed (100 pairs $(f_{c,t-1}, f_{c,t})$ for 50 countries). Based on simple linear regression through the origin, the least-squares estimate for ρ is 0.906. The fitted regression line is shown in Figure 6(a). The fitted regression line is very close to the loess smoother and fits the data very well. This outcome of ρ is similar to the current UN methodology of an expected increase in the TFR of 0.05 children at a TFR of 1.5.

In the AR(1) model, the asymptotic 80% prediction interval is given by

$$\left(2.1 - 1.28\frac{s}{\sqrt{1-\rho^2}}, 2.1 + 1.28\frac{s}{\sqrt{1-\rho^2}}\right),$$

with $\rho = 0.906$. The estimate for s in the regression model is 0.07, this would give the prediction interval (1.89, 2.31) which seems too narrow as little is known about the future range of possible outcomes of the TFR, especially within one to two future generations. Alternatively, assuming that the 100 points are roughly a random sample from the marginal distribution of the AR(1) process, we estimate that $s = 0.19$ (given the mean $\mu = 2.1$, the estimates range from 1.2 to 2.1, so standard deviation is $0.9/2 = 0.45 = s/\sqrt{1-\rho^2}$). Also, when fitting the AR(1) model to all data points after the TFR has decreased below 2.1 we find that $s = 0.2$. Based on $\mu = 2.1$, $\rho = 0.906$ and $s = 0.2$, the asymptotic 80% prediction interval is given by: 2.1 +/- 0.60 children. This approximate estimate of s seems more in line with the current (lack of) knowledge about future levels of the TFR, and is used in the model.

Figure 7 illustrates the projections with the AR(1) model for Finland and the Netherlands, starting at their turn-around outcome. The AR(1) projection matches well with the observed trend.

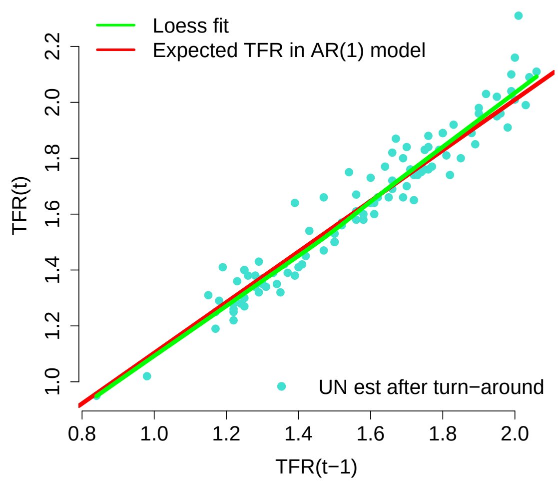

Figure 6: Observed UN estimates after turn-around point, with loess smoother and fitted regression (through the origin) line for AR(1) model.

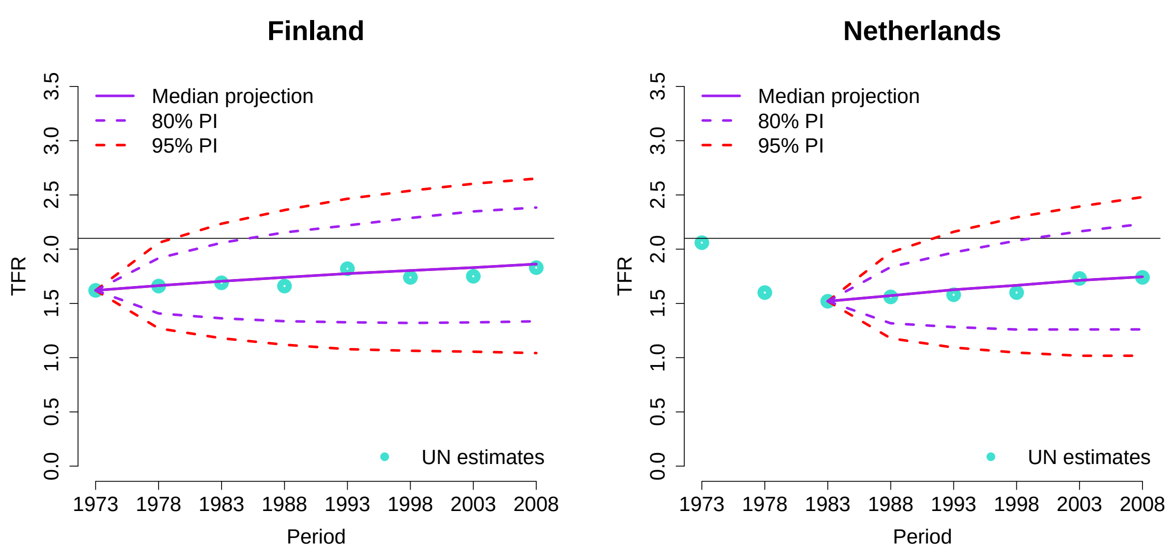

Figure 7: In-sample low-fertility prediction intervals for (a) Finland and (b) the Netherlands.

3 Results

3.1 Probabilistic TFR projections

In this section projections for several countries will be discussed to illustrate the results of the Bayesian projection model.

Figure 8 shows the prediction intervals for future TFR in (a) India, (b) Bolivia, (c) Mozambique and (d) Uganda, which are currently going through the fertility transition. The most recent estimates for 2005-2010 and projections for 2045-2050 for these countries are given Table 1. The median projections with the Bayesian projection model (BPM) show a slower decline in all four countries than given by the UN predictions, based on what has been observed in these countries so far, and in other countries that have (partly) completed the fertility transition. The end level in 2045-2050 is slightly lower for India with the BPM, because the fertility transition is not projected to end at 1.85 as explained in the previous section. Comparing the widths of the prediction intervals of the four countries shows that the higher the current level of the TFR, the more uncertainty in its projections, with the width of the 80% prediction interval for 2045-2050 ranging from 0.77 children for India to 2.08 children for Uganda.

Table 1: Projection results for 2045-2050 for selected countries, ordered by increasing TFR in 2005-2010: UN estimate for 2005-2010 and projection for 2045-2050, median projection for 2045-2050 with Bayesian projection model (BPM), and 80% prediction interval (PI).

Country	2005-2010	UN 2045-2050	BPM 2045-2050			
			Median	80% PI		Width 80% PI
Italy	1.38	1.74	1.77	1.23	2.31	1.08
Netherlands	1.74	1.85	1.94	1.39	2.47	1.09
India	2.76	1.85	1.76	1.41	2.17	0.77
Bolivia	3.50	1.85	2.08	1.61	2.60	0.99
Mozambique	5.11	2.41	2.50	1.77	3.27	1.50
Uganda	6.38	2.62	3.03	1.99	4.07	2.08

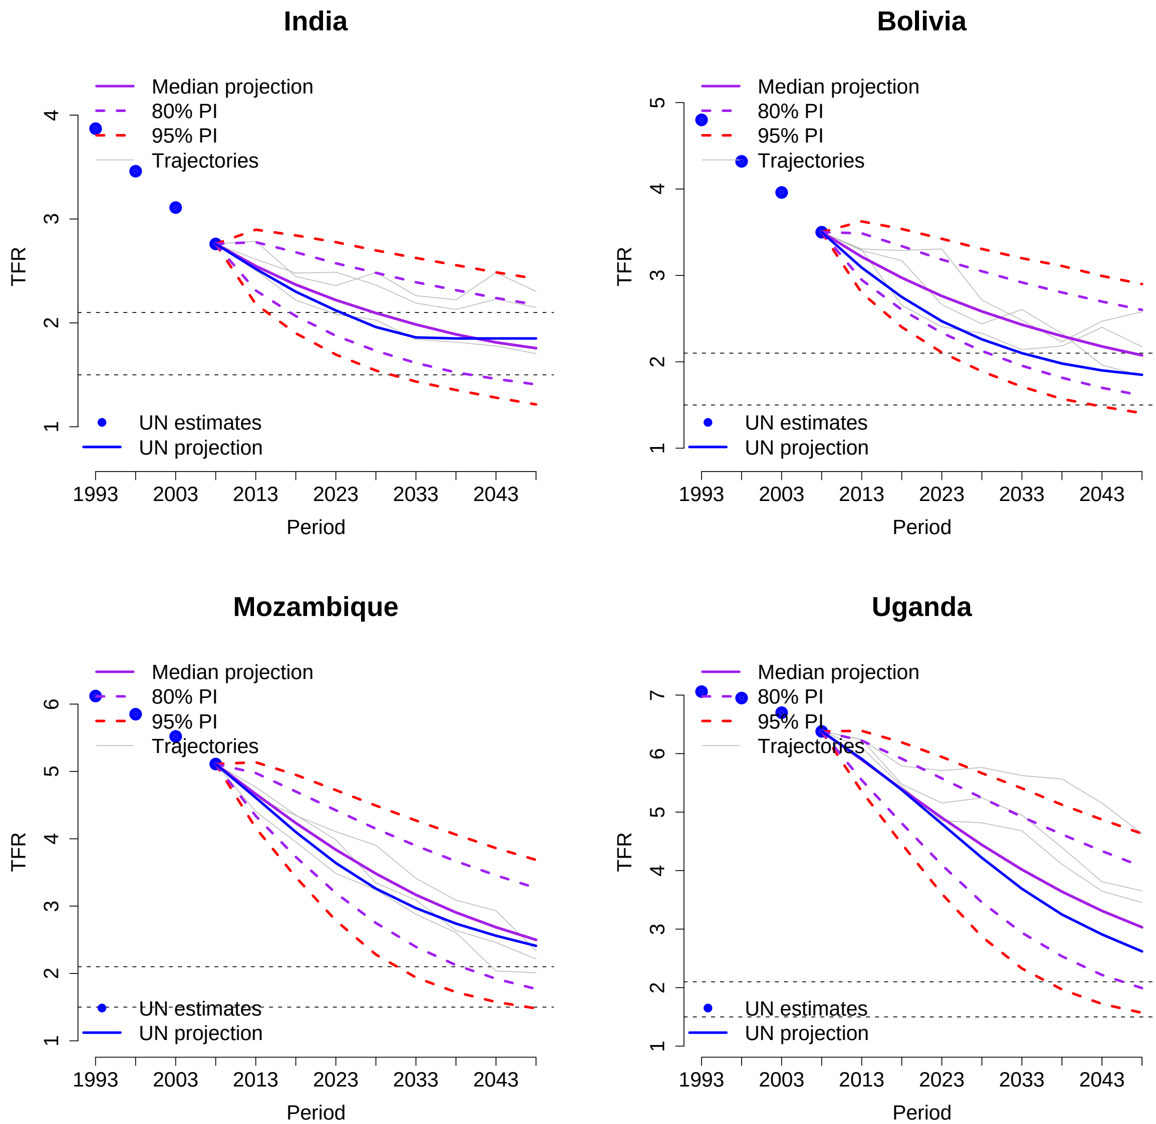

Figure 8: Prediction intervals for (a) India, (b) Bolivia, (c) Mozambique and (d) Uganda, which are currently going through the fertility transition. The solid red line is the median projection with 95% and 80% prediction intervals (red and purple dashed lines). A random sample of 3 future trajectories are shown in grey. The UN estimates and projections (2008 revision) are shown in blue.

Figure 9 shows the prediction intervals for future TFR in two countries with currently below-replacement fertility: (a) Italy and (b) the Netherlands. Table 1 gives the outcomes of the estimates and projections in these countries. In the Netherlands current fertility is 1.74. The median BPM projection for 2048 is 1.94 child, 0.09 child higher than the UN prediction. The current total fertility rate in Italy even lower, 1.38 children. The projection is 1.74 child for 2045-2050, similar to the UN projection. The uncertainty as given by the 80% prediction intervals is around 1 child for both countries.

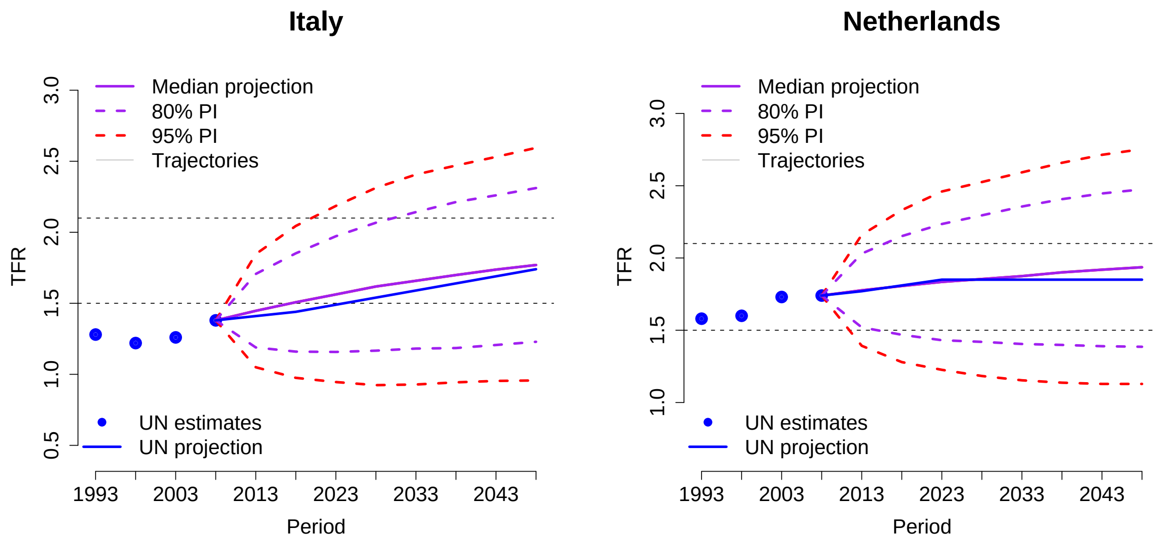

Figure 9: Prediction intervals for (a) Italy and (b) the Netherlands, both countries are currently experiencing below-replacement fertility. The solid red line is the median projection with 95% and 80% prediction intervals (red and purple dashed lines). A random sample of future trajectories is shown in grey. The UN estimates and projections (2008 revision) are shown in blue.

Figure 10(a) shows the widths of the 80% prediction intervals in 2045-2050, plotted by region against the current TFR level (2005-2010). The prediction intervals are widest for countries with currently high TFR levels, and smallest for countries with a TFR between 2 and 3. Figure 10(a) shows the ratios of the width of the lower half of the 80% prediction interval, over its total width. This figure shows that the prediction intervals are approximately symmetric around the median projection, as the ratios are scattered around 0.5. There is slightly more uncertainty towards higher outcomes of the TFR for projections in countries where the TFR is currently around three to four children.

Figure 11 shows the outcomes of the decline curves in India and Thailand together with the UN fast/fast and slow/slow scenarios. As expected, the median decrements as given by the BPM are quite different for the two countries, with larger decrements for Thailand.

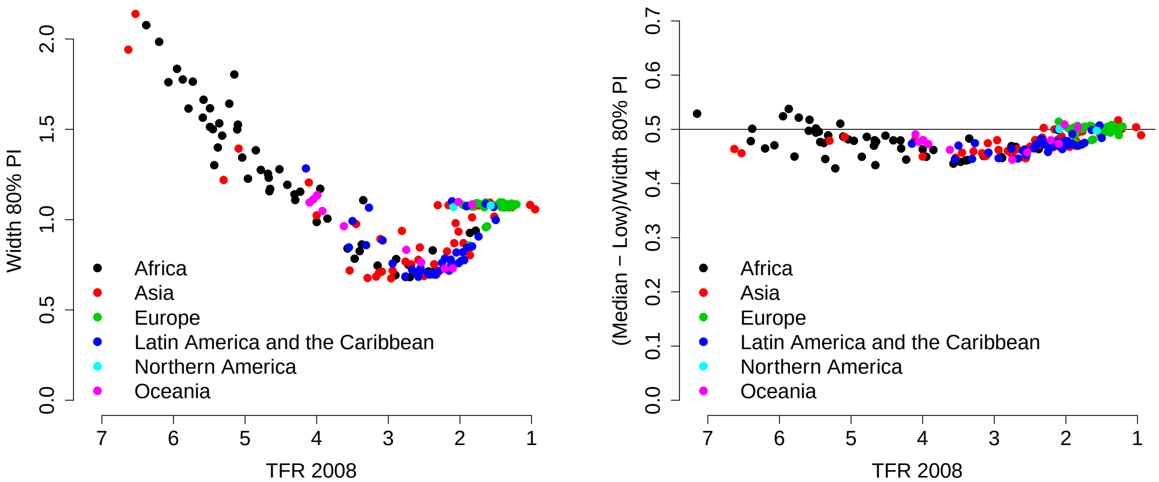

Figure 10: (a) Widths of 80% prediction intervals for 2045-2050, plotted against TFR in 2005-2010 (decreasing). The width is largest at high TFR, and at its minimum for countries with a TFR between 2 and 3. (b) Ratios of the width of the lower half of the 80% prediction interval, over its total width.

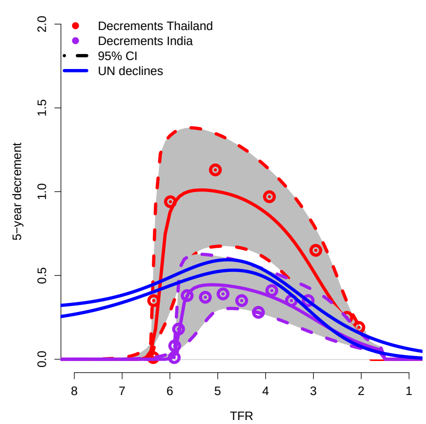

Figure 11: Outcome of the double logistic function for Thailand and India, together with the fast/fast and slow/slow UN scenarios, to illustrate the difference in the outcomes of the double logistic function between countries in the Bayesian projection model.

3.2 Model validation

Modeling assumption are validated by carrying out out-of-sample projections. In the out-of-sample projections, the BPM is used to construct projections based on the UN estimates up to and including the 5-year period 1995-2000. The projections are compared to the UN estimates in 2000-2005 and 2005-2010. The calibration of the prediction intervals is evaluated by calculating the proportion of left-out UN estimates that fall outside their prediction intervals. If modeling assumptions hold, we expect 10%/2.5% of the estimates to fall above/below the 80%/95% intervals. The results are given in Table 2. The results show that the prediction intervals are reasonable for high fertility countries (in which the TFR is currently above 2 children), but too wide for the low fertility countries. The latter is as expected because the model uses a larger error variance in the projections around replacement fertility than what has been observed so far (as explained in the methods section), to allow for extra uncertainty in future trends around replacement fertility levels.

Table 2: Model validation results: the proportion of left-out UN estimates that falls below or above their 80% and 95% prediction intervals in 2000-2005 and 2005-2010, when projecting from 1995-2000. We expect 10%/2.5% of the estimates to fall above/below the 80%/95% intervals.

TFR in 1995-2000	Period	# Countries	Proportion of obs.			
			<80%PI	>80%PI	<95%PI	>95%PI
<2	2000-2005	58	0.08	0.00	0.00	0.00
>2	2000-2005	138	0.07	0.06	0.02	0.01
Total	2000-2005	196	0.07	0.05	0.01	0.01
<2	2005-2010	58	0.03	0.00	0.00	0.00
>2	2005-2010	138	0.04	0.12	0.01	0.02
Total	2005-2010	196	0.04	0.09	0.01	0.01

4 Discussion

In this paper we propose a new unified approach to project the total fertility rate in all countries of the world; in countries at any stage of the fertility transition or post fertility transition at below-placement fertility. The projection model incorporates the worldwide experience of observed fertility transitions during the last 60 years. The results are presented in term of median predictions and 80% prediction intervals. In general for policy planning purposes, the 80% prediction interval may often be sufficient to get enough insight into future scenarios. If needed, 95% prediction interval would allow for a more conservative analysis of future scenarios. The prediction intervals as presented here are based on UN estimates. Using the UN estimates as the true TFR will give narrower prediction intervals as it ignores the uncertainty in the UN estimates.

The goal of the Bayesian projection model is to construct projections that are roughly calibrated for all countries combined. The out-of-sample predictions in 1995-2005 based on 1950-1995 data show that for all countries combined the projected trend matches up well with what has been observed in 1995-2000 and 2000-2005. The prediction intervals are slightly conservative in the sense that they overestimate the uncertainty.

The TFR projections, together with projections of mortality and migration, give population projections. Upper bounds of the TFR projections for high-fertility countries like Uganda could possibly yield unfeasible population levels. Expert information on regional or country-specific demographic issues (like carrying capacity, population pressure on the environment, basic services and infrastructure) could be incorporated, e.g. in a Bayesian melding framework, to exclude unrealistic population outcomes and the TFR trajectories that underlie these outcomes.

Possibly a better prediction can be constructed if information on current and future trends in fertility-related outcomes, for example child mortality and economic development, is included in the model. However, projecting these outcomes is difficult and adds another level of complexity to the model. Also, the UN restriction of having to apply the same model to a broad range of countries imposes difficulties with respect to availability of data and applicability of more detailed models to a range of different countries.

References

Bongaarts, J. and R. Bulatao (2000). *Beyond six billion: forecasting the world's population.* National Research Council (U.S.), Washington, D.C.: National Academy Press.

Gelfand, A. and A. F. M. Smith (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association 85*, 398–409.

Gelman, A., F. Bois, and J. Jiang (1996). Physiological pharmacokinetic analysis using population modeling and informative prior distributions. *Journal of the American Statistical Association 91*, 1400–1412.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (2004). Bayesian Data Analysis, Second Edition. *Boca Raton: Chapman & Hall/CRC Chapter 5.*

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

Neal, R. M. (2003). Slice sampling. *The Annals of Statistics 31, No. 3*, 705–767.

United Nations, Department of Economic and Social Affairs, Population Division (2006). *World Population Prospects. The 2004 Revision, Vol. III, Chapter VI. Methodology of the United Nations population estimates and projections*, pp. 100–104. http://www.un.org/esa/population/publications/WPP2004/WPP2004 Volume3.htm.

United Nations, Department of Economic and Social Affairs, Population Division (2006). *World Population Prospects. The 2006 Revision, Vol. III, Analytical report.* United Nations publication, http://www.un.org/esa/population/publications/wpp2006/wpp2006.htm.

United Nations, Department of Economic and Social Affairs, Population Division (2009, forthcoming). *World Population Prospects. The 2008 Revision.* CD-ROM Edition - Extended Dataset in Excel and ASCII formats, United Nations publication, http://www.un.org/esa/population/publications/wpp2008/.

5 Appendix

5.1 Standard deviation of the random distortions

The standard deviation $\sigma_{c,t}$ of the random distortions, $\varepsilon_{c,t}$, is given by:

$$\sigma_{c,t} = c_{1975}(t) \left(\sigma_0 + (f_{c,t} - S) \left(-aI_{f_{c,t}>S} + bI_{f_{c,t}<S}\right)\right), \tag{11}$$

with parameters σ_0, the maximum standard deviation of the distortions, attained at TFR $f_{c,t} = S$, a and b are multipliers of the standard deviation, to model the linear decrease for larger and smaller outcomes of the TFR. The constant $c_{1975}(t)$ is added to model the higher error variance of the distortions before 1975, and is given by:

$$c_{1975}(t) = \begin{cases} c, & t \in [1950 - 1955, 1970 - 1975]; \\ 1, & t \in [1975 - 1980, \infty) \end{cases} \tag{12}$$

The variance function is illustrated in Figure 12.

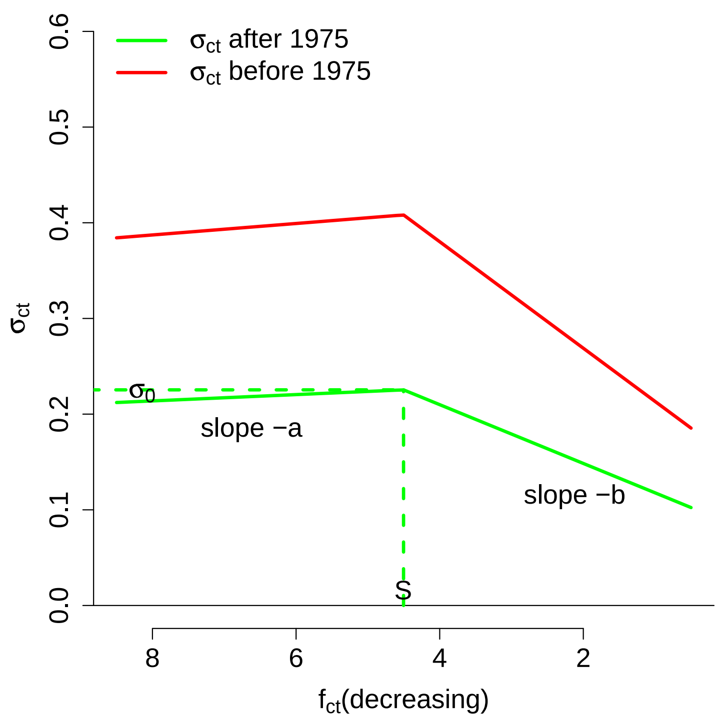

Figure 12: Standard deviation of the distortion terms

5.2 Start period and start level

The start period of the transition, τ_c, is given by the period in which the TFR starts declining and its outcome can be determined before fitting the projection model.

To exclude temporary dips in the TFR before the start of the fertility transitions, the start period is given by the period with the most recent local maximum within 0.5 children of the global maximum of the TFR. For some countries the decline had started before 1950 and so the start of the decline was not observed in our observation period. These countries are identified by a maximum TFR smaller than 5.5 children in the observation period. The cut-off of 5.5 children was chosen after visual inspection of the start periods for all countries based on different cut-off values. Using 5.5 children best identified the countries in which the decline had possibly already started before 1950. With this definition the start period is defined by:

$$\tau_c = \max\{t : (M_c - L_{c,t}) < 0.5\}, \text{ if } L_{c,t} > 5.5, \tag{13}$$

$$< 1950 \text{ otherwise.} \tag{14}$$

with global maximum $M_c = \max_t f_{c,t}$, and local maxima denoted by $L_{c,t}$.

In some countries an increase in the TFR has been observed before the onset of the fertility transition, for example in countries in western Africa after independence. The distribution of the distortion terms during this period is possibly different from the distribution during the fertility transition. As the period before the onset of the decline is not important for projecting the TFR, the UN estimates during this observation period are left out of the analysis.

For countries in which the fertility decline starts after 1950, its start level U_c is observed and given by:

$$U_c = f_{c,\tau} \text{ if } \tau_c \geq 1950. \tag{15}$$

For countries in which the decline started before 1950, the start level U_c is added as a parameter to the model, with prior distribution:

$$U_c = U[5.5, 8.8] \text{ for } \tau_c \geq 1950, \tag{16}$$

The upper bound of the prior distribution on starting value $U_c = \sum \triangle_{ci}$ is based on the observed maximum in the UN estimates (8.7). Its lower bound of 5.5 children is the same lower bound used to define the start level of the decline. This number is based on examining decline curves, the minimum level at which the decline starts is slightly under 6.

5.3 Bayesian hierarchical model for parameters of drift term

A Bayesian hierarchical model is used to estimate the decline parameters $(\triangle_{c1}, \triangle_{c2}, \triangle_{c3}, d_c)$ for each country.

A logit transform is used to restrict the maximum decrement d_c to be between 0.25 and 2.5 children decrease per time period (e.g. the maximum pace of fertility decline observed in the past is around 2 children/5-year period in China). Its hierarchical model is given by:

$$\phi_c \sim N(\chi, \psi^2), \tag{17}$$

with ϕ_c the logit-transform of $d_c/5$:

$$\phi_c = \log\left(\frac{d_c/5 - l_d}{u_d - d_c/5}\right), \tag{18}$$

and $[l_d, u_d] = [0.05, 0.5]$, χ the hierarchical mean of the logit-transformed maximum decline for all countries in which (part of) the fertility transition has been observed and ψ^2 the error variance. (The model is implemented in terms of $d_c/5$ because it was originally fitted to 1-year decrement data).

Given $\triangle_{c4}$ and the start level $U_c = \sum_i \triangle_{ci}$, the other three TFR ranges $(\triangle_{c1}, \triangle_{c2}, \triangle_{c3})$ can be expressed as proportions of $U_c - \triangle_{c4}$. Define:

$$p_{ci} = \frac{\triangle_{ci}}{U_c - \triangle_{c4}} \text{ for } i = 1, 2, 3, \tag{19}$$

such that $\sum_{i=1}^{3} p_{ci} = 1$. We assume that the proportions are exchangeable between countries. For the purpose of computation, we transform the model in terms of a new set of parameters γ_{ci} and then define the p_{ci}'s as follows (Gelman et al. 1996):

$$p_{ci} = \frac{\exp(\gamma_{ci})}{\sum_j \exp(\gamma_{cj})}. \tag{20}$$

The hierarchical model for the γ_{ci}'s is given by:

$$\gamma_{ci} \sim N(\alpha_i, \delta_i^2), \tag{21}$$

with α_i the hierarchical mean of the γ_{ci}'s and δ_i^2 their variance.

A Markov Chain Monte Carlo sampling procedure is used to get samples of the posterior distributions of each of the parameters (Gelfand and Smith 1990). This procedure is a combination of Gibbs, Metropolis-Hastings and slice sampling steps (Neal 2003). The MCMC sampling algorithm is implemented in the statistical package `R`.