

Estimating Child Mortality Using a Bayesian Hierarchical Time Series Model

Leontine Alkema*, Wei Ling Ann

Department of Statistics and Applied Probability, National University of Singapore, Singapore.

*E-mail: alkema@nus.edu.sg

Abstract

Background: Millennium Development Goal 4 calls for a reduction in the under-five mortality rate by two-thirds between 1990 and 2015, which corresponds to an annual rate of decline of 4.4%. The United Nations Inter-Agency Group for Child Mortality Estimation estimates child mortality in every country to measure progress. For the majority of countries, the estimates within a country are based on the assumption of a piece-wise constant rate of decline.

Methods and Findings: This paper proposes an alternative method to estimate child mortality, such that the underlying rate of change is allowed to vary smoothly over time using a time series model. Information about the average rate of decline and changes therein is exchanged between countries using a Bayesian hierarchical model. Cross-validation exercises suggest that the proposed model provides U5MR confidence bounds which are reasonably well calibrated during the observation period. The alternative estimates suggest smoother trends in under-five mortality and give new insights into changes in the rate of decline within countries.

Conclusions: The proposed model offers an alternative modeling approach for obtaining estimates of child mortality which removes the restriction of a piece-wise linear rate of decline and introduces hierarchy to exchange information between countries. The newly proposed estimates of the rate of decline in under-5 mortality and the uncertainty assessments would help to monitor progress towards reducing child mortality.

Introduction

Millennium Development Goal 4 (MDG4) calls for a reduction in the under-five mortality rate by two-thirds between 1990 and 2015. The United Nations Inter-Agency Group for Child Mortality Estimation (IGME) estimates child mortality in every country to measure progress. Its estimates are revised every year; the most recent estimates were published in September 2010 [1, 2]. Measurements of the under-five mortality rate (U5MR) are available from different sources, including vital registration systems and sample surveys, and are published by IGME in the data base *CME Info* (www.childmortality.org). A set of standard weights that quantify the relative accuracy of different types of observations has been developed [3, 4]. In this weighting system, observations that are deemed to be of high quality are assigned a higher weight than observations of lesser quality. Observations from sources that are known to be over- or underestimating child mortality are assigned a zero weight.

For the majority of countries, IGME uses a linear spline regression estimation approach for estimating U5MR [2]. In short, the observation period is broken up into different periods and a constant rate of decline is estimated within each period. The number of periods and their start and end points depend on the availability of observations and the weight that has been assigned to the observation. The estimated rate of decline within the period is estimated using weighted least-squares estimation.

The spline estimation approach is illustrated in Figure 1 for Pakistan. The first plot shows the observations of U5MR expressed as the number of deaths per 1,000 live births (black dots) and the spline estimates. The vertical lines represent the time points, which are called the knots, at which the rate of decline is allowed to change. The second plot shows the annual rate of decline, given by $\log(U5MR(t-1)/U5MR(t))$, which is held constant between the knots.




Figure 1. Illustration of the spline estimation approach for Pakistan. Left: U5MR against time. Observations are represented by black dots, spline estimates are represented by the blue line. Right: Underlying annual rate of decline of the spline estimates against time. The vertical lines in both plots represent the time points (knots) at which the rate of decline changes.

The drawback of the spline approach is that it leads to inaccuracies in the estimated rate of decline over time by forcing it to be constant between the knots, especially when there is a large difference between the rate of decline from one period to the next, or when the rate of decline is kept constant during a long observation period. This is illustrated in Figure 2 for the Latvia. Child mortality might have increased in early 1990s, but there is no evidence in the data for the “elbows” in the U5MR estimates in 1989 and 1994: the abrupt changes in the rate of decline are caused by the placements of the knots. If the knots would have been placed a few years earlier or later, a different trend in U5MR would have

been estimated.

The spline estimates for Bhutan are illustrated in the second plot of Figure 2. Because of the limited number of observations, the rate of decline is kept constant throughout the observation period in Bhutan. As a result of this approach, the U5MR estimates do not follow the trend in the data in the early 1990s very well. Moreover, the estimation approach gives relatively narrow confidence bounds for U5MR during the 1980s, even though data during that period are lacking.

To overcome the drawbacks of using a piece-wise constant model for the rate of decline in the spline estimation approach, we developed an alternative method for estimating under-5 mortality for all countries, which allows the rate of decline to change continuously over time. A Bayesian hierarchical model is used to estimate the model parameters for each country, such that trends within countries are estimated based on the country's own history as well as observed trends in all other countries.

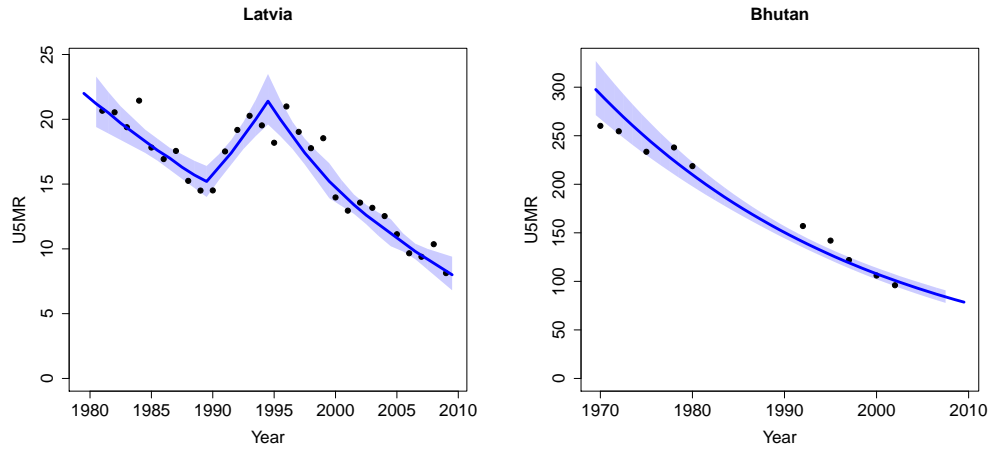


Figure 2. Spline estimates for Latvia and Bhutan. Spline estimates of U5MR in blue, observations in black, 95% confidence intervals are represented by the blue area.

Methods

Time series model: The rate of decline in U5MR is modeled with a time series model, such that changes are not limited to occur at predefined time points. For country c , the rate of decline from year $t - 1$ to year t is given by $r_{c,t} = \log\left(\frac{u_{c,t-1}}{u_{c,t}}\right)$, where $u_{c,t}$ is the U5MR in year t . We assume that the rate of decline fluctuates around a country-specific average annual change using an autoregressive time series model of order one (an AR(1) model). For each country, the parameters of the time series model are (i) average rate of decline β_c , (ii) parameter ρ_c that models the autocorrelation in the rate of decline, and (iii) a variance parameter δ_c^2 which determines the smoothness of its trajectory. The AR(1) model is given by:

$$r_{c,t} = \beta_c + \rho_c(r_{c,t-1} - \beta_c) + \varepsilon_{c,t}, \quad (1)$$

$$\varepsilon_{c,t} \sim N(0, \delta_c^2), \quad (2)$$

where $\varepsilon_{c,t}$ are the distortion terms that model the random fluctuations in $r_{c,t}$. The model behavior is illustrated in Figure 3. In the top row the black curve represents one simulated trajectory of the annual rate of decline for $\beta_c = 0.03$, $\rho_c = 0.5$ and $\delta_c = 0.016$, the bottom row illustrates the corresponding trajectory of U5MR based on a start level of 200 deaths per 1,000 births in 1970. The parameters are changed individually in each plot to illustrate their influence on the rate of decline and corresponding U5MR trajectory (the standardized distortion terms $\varepsilon_{c,t}/\delta_c$ in the time series model are held constant). The influence of the average decline β_c is shown in the first column: a larger rate of decline gives a faster decline in U5MR. For outcomes of ρ_c that are closer to one, the rates of decline are more correlated over time and the trajectory of $r_{c,t}$ can “stay away from” the average decline during a longer period. This effect is illustrated in the second column of Figure 3 where ρ_c is increased from 0.5 to 0.85. The resulting rate of decline in red is larger than β_c from 1981 until 2003, leading to a faster decline of U5MR during that period. Finally, decreasing variance parameter δ_c^2 gives a smoother trend in the rate of decline and U5MR, as illustrated in the third column. The simulations illustrate the interpretation of the parameters and the flexibility of the model to represent various trends in the rate of decline and U5MR.

Data model: The sampling model for the data is taken from the IGME spline estimation approach such that the U5MR estimates from the two modeling approaches are directly comparable. The sampling model is given by:

$$\log(y_{c,t,s}) \sim N(\log(u_{c,t}), \sigma_c^2/w_{c,t,s}),$$

where $y_{c,t,s}$ is the observed U5MR for country c , year t , observation s . Its variance on the log-scale is proportional to the inverse of its weight $w_{c,t,s}$. The country-specific error variance is denoted by σ_c^2 . Observation years are rounded and observations with zero weight are excluded.

Bayesian hierarchical model: Estimating the country-specific parameters β_c and ρ_c presents a challenge for some countries because of a limited number of observations and/or “noisy data” (large sampling and non-sampling errors). We use a Bayesian hierarchical model [5, 6] to estimate these parameters in each country, such that the estimates are based on the observations in the country of interest, as well as on the global experience. A hierarchical approach to estimating and projecting demographic outcomes for a number of countries is a natural way to exchange information between countries while constructing country-specific estimates and projections, that include an uncertainty assessment. The fewer the number of observations in the country of interest, the more its estimates and projections are driven by the experience of other countries, while in countries with many observations the estimates will be driven more by its observed trends.

In the Bayesian hierarchical model for average rate of decline β_c we assume that for all countries β_c is drawn from a probability distribution that represents the range of outcomes of the average annual growth across all countries. For β_c in a specific country, its probability distribution based on the world-level

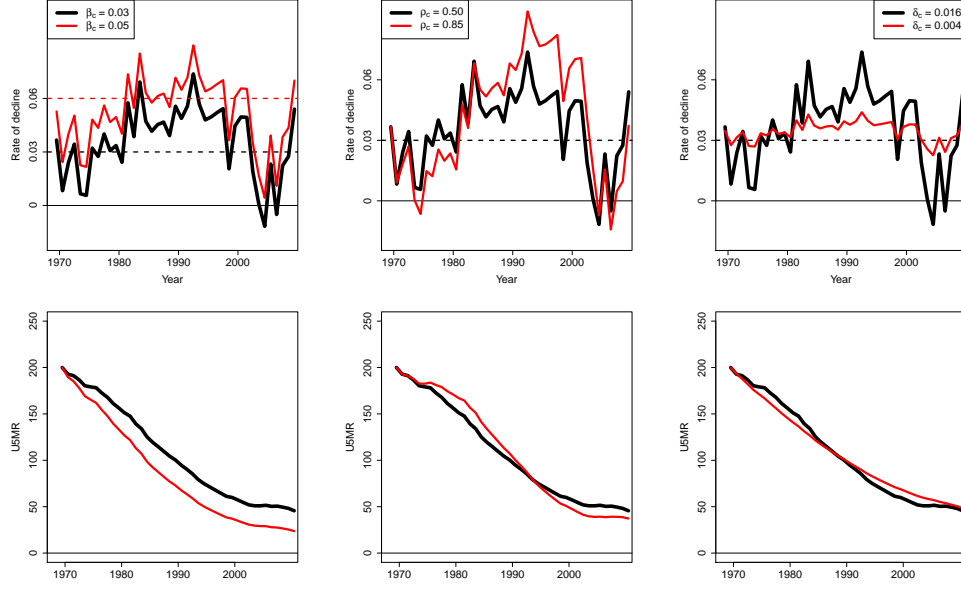


Figure 3. Model illustration. Simulations of the annual rate of decline and the corresponding trajectories of U5MR for different values of average rate of decline β_c , autoregressive parameter ρ_c and variance parameter δ_c^2 . The start level of U5MR is set at 200 deaths per 1,000 live births in 1970. In all plots the black curve represents one simulated trajectory for $\beta_c = 0.03$, $\rho_c = 0.5$ and $\delta_c = 0.016$. The red curve illustrates the effect of changes in one of the parameters. Column 1: A larger average decline β_c gives a faster decline in U5MR. Column 2: For outcomes of ρ_c that are closer to one, the rates of decline are more correlated over time and the trajectory of $r_{c,t}$ can “stay away from” the average decline during a longer period. Column 3: Decreasing variance parameter δ_c^2 gives a smoother trend in the rate of decline and U5MR.

experience is then updated using Bayes’ theorem with the observed trend in the country, which results in the posterior distribution for β_c . The resulting estimates (draws from the posterior distribution) can be viewed as weighted averages of a “world pattern” and information from the country data. The hierarchical distribution for β_c is given by:

$$\log(\beta_c) \sim N(m_\beta, \sigma_\beta^2),$$

with hierarchical mean and variance m_β and σ_β^2 . The log-transform is applied to restrict the average rate of decline to be positive. The same approach is used for the autoregressive parameter ρ_c . Its hierarchical distribution is given by:

$$\rho_c \sim N(m_\rho, \sigma_\rho^2),$$

with m_ρ the hierarchical mean and σ_ρ^2 the variance of the ρ_c ’s.

Variance parameter δ_c^2 determines the smoothness of the trajectories: a larger variance allows trajectories from the time series model to follow trends in the data more closely. Because error variance is estimated within each country separately, an upper bound on δ_c^2 is needed to avoid highly fluctuating U5MR trajectories that follow the data closely in countries where large measurement errors are present. Smoothness of trajectories is guaranteed by restricting the expected absolute outcome of the annual distortion term to be smaller than the average rate of decline β_c , which corresponds to $\delta_c = \lambda \cdot \beta_c \cdot \sqrt{\pi/2}$, with $0 \leq \lambda \leq 1$.

Estimation: The under-5 mortality rate was estimated with the Bayesian hierarchical time series model for 165 countries within their observation periods, and projected for an additional five years after the most recent observation year in each country. All model parameters were estimated in a Bayesian framework. Diffuse prior distributions were assigned to the additional model parameters. The complete model and assigned prior distributions are given in the Appendix. A Markov Chain Monte Carlo (MCMC) algorithm was used to sample from the posterior distribution of the parameters using Winbugs software [7, 8]. Results were obtained from eight chains; the total number of iterations in each chain was 600,000, keeping every 500th iteration, and discarding the first 100,000 iterations as burn-in. Convergence of all model parameters and $\log(\text{U5MR})$ was assessed visually using trace plots and with the run length diagnostic of Raftery and Lewis [9, 10]. This resulted in a posterior sample of U5MR trajectories for each country. The “best” estimates during the observation period were given by the median outcome, and 95% confidence intervals by the 2.5% and 97.5% percentiles of the set of outcomes in each year. Similarly for the projection period, the “best” projection was given by the median outcome, and the 2.5% and 97.5% percentiles were used to construct 95% projection intervals.

Model validation: Cross-validation was used to validate model performance, in which some observations are excluded while the method is applied to the remaining observations (the training set). Calibration was assessed based on the predictive distributions for the excluded observations by checking that the prediction intervals contain the left-out observations the right proportion of the time. The proportion of excluded observations that fall outside their 80% and 95% prediction intervals was used as a measure of calibration. These proportions are expected be around 20% and 5%, respectively. Higher proportions suggest that the confidence and projection intervals for U5MR are too narrow, while smaller proportions suggest that the confidence and projection intervals for U5MR are too wide.

When leaving out observations, the structure of the data should be taken into account, e.g. by leaving out all observations from the same survey, such that the observations in the test set are independent of the observations in the training set. Unfortunately the data on *IGME info* are not yet in a format which allows for this exercise to be carried out in an automated fashion. Instead, we carried out two validation exercises. In the first exercise, 20% of the observations within each country are left out at random, such that the training set consists of 80% of all observations. In the second validation exercise, all observations in the last five years of the observation period within each country were left out. That is, if the last observation year is T_c in country c , all observations in years $T_c - 4, \dots, T_c$ were excluded from the training set for that country. The motivation for this second exercise is to check the predictive performance of the model: for many countries the last observation year is well before 2010, thus the “2010 estimates” are in fact projections.

To assess how confident one can be about the current estimates and projections, we also compared the current estimates and projections of U5MR based on the full data set to the confidence and projection intervals for U5MR based on the training set. The goal is to check that additional data has not changed the U5MR estimates and projections significantly: as more data become available, we expect the current (updated) estimates to lie well within the previously constructed confidence and projection intervals. The smaller the proportion of estimates and projections that fall outside their respective confidence and projection intervals, the better (contrary to the previously described measures of calibration).

Results

U5MR estimates: The under-5 mortality rate was estimated with the Bayesian hierarchical time series model (BM) for 165 countries within their observation periods, and projected for an additional five years after the most recent observation year in each country. For many countries, the differences between the spline estimates and the median estimates of the BM are small. This is illustrated in Figure 4 for Pakistan. The median estimates from the BM are given by the red line, and the red shaded area represents the 95% confidence intervals. The IGME estimates are shown in blue, with the 95% confidence and projection intervals represented by the blue shaded area. For Pakistan, the confidence intervals overlap and point estimates are similar.

As discussed in the introduction, the limited number of knots in the spline approach introduce abrupt changes in the rate of decline in U5MR in Latvia. The BM estimates differ from the IGME estimates, as shown in Figure 4: U5MR changes more smoothly in the BM estimates. For Bhutan, a country with a small number of observations, the spline regression approach gives relatively narrow confidence bounds for U5MR during the 1980s, even though data during that period are lacking. The BM confidence bounds differ from the spline confidence bounds and show greater uncertainty about the U5MR in the 1980s.

The results for Australia, India and Nigeria are shown in Figure 4 on the second row. Australia is a country with a high quality vital registration system; the BM estimates follow the observations closely and there is little uncertainty about the level of U5MR. Based on the splines estimates, India and Nigeria together account for nearly a third of all under-five deaths worldwide in 2009 [1]. The BM estimates are similar but slightly more uncertain for both countries in 2009.

The results for all countries are given in the Appendix, in Figure 10.

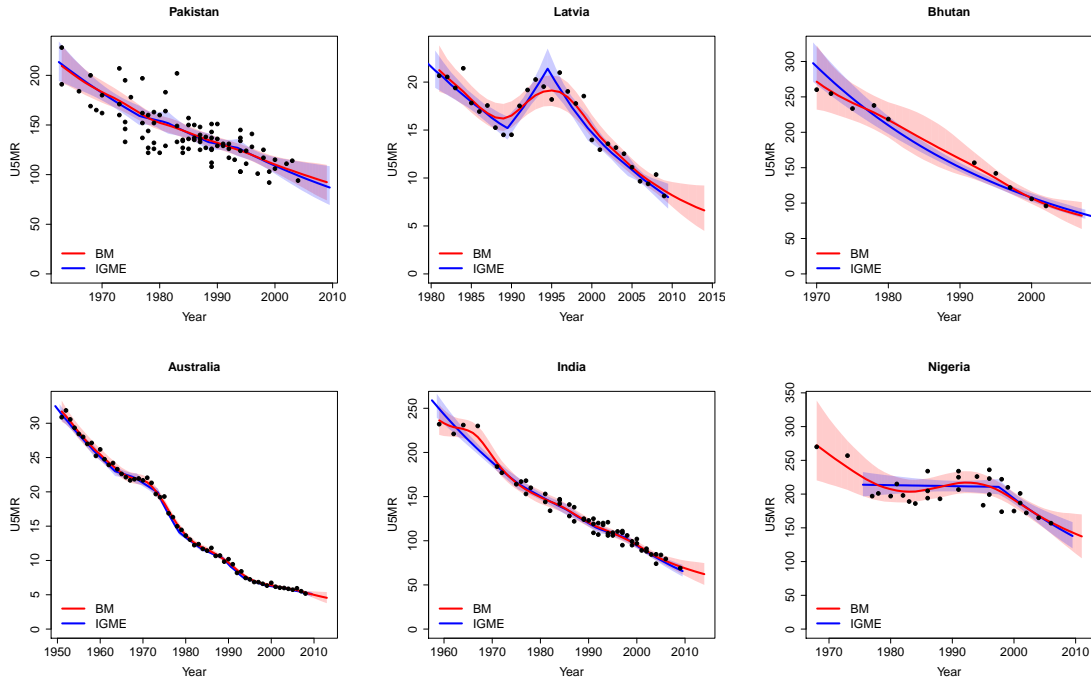


Figure 4. Estimates of U5MR for Pakistan, Latvia, Bhutan, Australia, India and Nigeria. The median U5MR estimates from the Bayesian model (BM) are shown in red, and the IGME estimates are shown in blue. The 95% confidence and projection intervals are represented by the red area for the Bayesian model, and the blue area for the IGME estimates. Observations are represented by the black dots.

Rate of decline: Millennium Development Goal 4 calls for a reduction in the under-five mortality rate by two-thirds between 1990 and 2015, which corresponds to an annual rate of decline of 4.4%. This target has increased the attention for monitoring the rate of decline in U5MR to assess whether countries are reducing U5MR at the 4.4% rate a year, and whether the rate of decline has increased recently in countries where progress has been slow. The average annual rate of decline, measured over 5-year periods, is plotted for Pakistan, Latvia, Bhutan, Australia, India and Nigeria in Figure 5. The green line is drawn at the target rate of 4.4%. The median estimates for the rate of decline have been above 4.4% for a short period in Bhutan, and for low-mortality countries Australia and Latvia; on average, the rate of decline in the countries with higher mortality rates tends to be much lower.

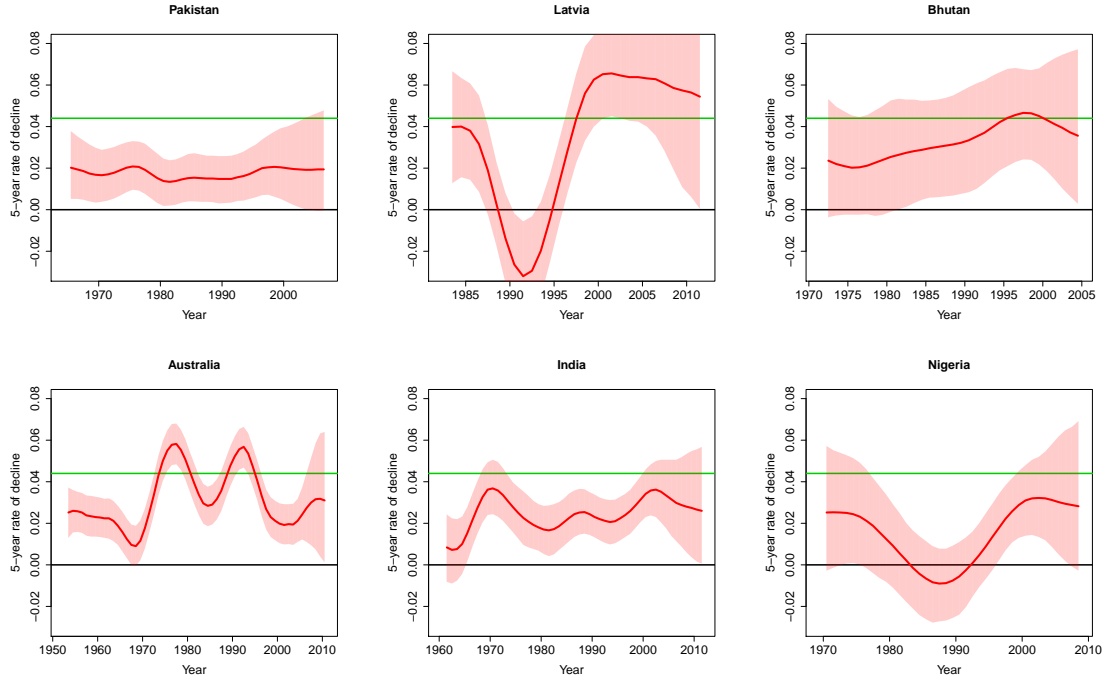


Figure 5. Estimates of the annual rate of decline for Pakistan, Latvia, Bhutan, Australia, India and Nigeria. The average rate of decline over a 5-year period is plotted against the midpoint of the 5-year period under consideration. The median estimates are shown in red, and the 95% confidence and projection intervals are represented by the red area. The green line is drawn at a rate of decline of 4.4%, which corresponds to Millennium Development Goal 4.

The observed rate of decline in the last five observation years is given in Figure 6 for selected high mortality countries, where high mortality countries are defined as those countries with at least 40 deaths per 1,000 births (as in [1], 55 countries in total). The selected high mortality countries are the six countries with the lowest, and the six countries with the highest estimates of the 5-year rate of decline among all high mortality countries. The highest rates were estimated in Pakistan, Bolivia, Nepal, Mongolia, Timor-Leste and Liberia, while Comoros, Burundi, Mauritania, Chad, Latvia and Congo DR had the lowest rates of decline.

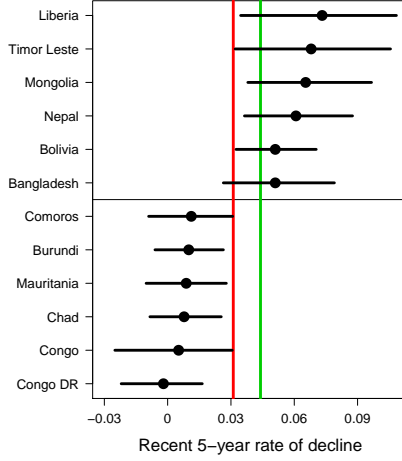


Figure 6. Estimates of the recent 5-year rate of decline for the six high mortality countries with the highest/lowest outcomes. Median estimates (dots) and 95% confidence intervals (horizontal line) for the rate of decline in the last five observation years in the six countries with the highest, and the six countries with the lowest median estimate. The red vertical line is the median of all country estimates, which is at 3.5%. The green line is drawn at a rate of decline of 4.4%, which corresponds to Millennium Development Goal 4.

Bayesian inference allows for calculating the posterior probability (or the posterior odds) that the recent 5-year rate of decline was above the 4.4% goal, which gives more insights into a country's recent progress towards reducing child mortality at the target rate. Figure 7(a) shows the probability that the recent rate of decline was higher than 4.4%, plotted against the median estimate of the rate of decline, for all high mortality countries. The most recent rate of decline was higher than 4.4% with probability 0.8 or higher for only four countries: Liberia, Mongolia, Nepal and Timor-Leste (plotted in green). For almost half of the countries (46%), the posterior probability that the recent rate of decline was less than 4.4% is at least 0.8. This corresponds to posterior odds of at least four to one that the recent rate of decline was less than 4.4% for those countries.

The difference between the most recent 5-year rate of decline and the preceding 5-year period can be used to assess evidence of acceleration or deceleration in child mortality reduction. The difference between the most recent 5-year rate of decline and its preceding 5-year period are shown in Figure 7(b) for all high mortality countries. The probability of a positive difference, of an increase in the rate of decline, is shown on the vertical axis. For eight countries the posterior odds of a positive increase are at least four to one. These countries are Liberia, Ghana, Tajikistan, Senegal, Uzbekistan, Cambodia, Mali and Chad. There are no countries in which the odds of a negative change in the rate of decline are more than four to one.




Figure 7. Posterior probability that the recent 5-year rate of decline was larger than 4.4% and of an increase in the 5-year rate of decline for all high mortality countries. (a)

Probability that the recent 5-year rate of decline in the observation period within each country is more than 4.4%, plotted against the median estimate of its recent rate of decline. The four selected countries in green have a posterior probability of more than 0.8 that the difference was positive. The countries that are plotted in pink have a posterior probability of less than 0.2 that the rate of decline was more than 4.4%. (b) Probability of a positive difference between the most recent 5-year increase and the preceding 5-year increase in the last ten years of the observation period within each country is plotted against the median estimate of that difference. The eight selected countries have a posterior probability of more than 0.8 that the difference was positive.

Other model parameters: Histograms of the posterior samples of the non-country specific parameters λ , m_ρ , σ_ρ , m_β and σ_β are given in Figure 9 in the Appendix. The autocorrelation parameter of the autoregressive process (ρ_c) does not vary much between countries: all 95% confidence intervals for ρ_c are roughly within 0.5 and 0.9. With ρ_c smaller than one for all countries, the rate of decline converges towards to average rate of decline β_c in long term projections, as shown for Pakistan in Figure 5.

Figure 8 shows the median estimates and 95% confidence intervals for the average rate of decline β_c in the six countries with the highest, and the six countries with the lowest median estimates. The red vertical line is the median of all country estimates, which is at 3.4%. The average rates of decline were lowest in Sierra Leone, Burundi, Burkina Faso, Pakistan, Togo and Afghanistan.

Model validation: To assess model performance we carried out the cross-validation exercises as discussed in the previous section. Countries with less than four observations in the training set were excluded (Korea DPR was excluded in the first exercise in which 20% of the observations were left out at random, and Grenada, Korea DPR, Liechtenstein, Montserrat and Somalia were excluded in the second exercise, where observations during the last five years of the observation period were left out). The results are summarized in Tables 1 to 6 in the Appendix.

When leaving out 20% of the observations at random, the results based on the predictive distributions for the left-out observations show that slightly fewer observations fall outside their respective 80% and 95% prediction intervals than expected (13% instead of 20% for the 80% interval, and 4% instead of 5% for the 95% interval), which suggests that the confidence bounds for U5MR are somewhat conservative (too wide), see Table 1.

As more data become available, we expect the current (updated) estimates to lie well within the previously constructed confidence bounds. This holds true when leaving out 20% of the observations:

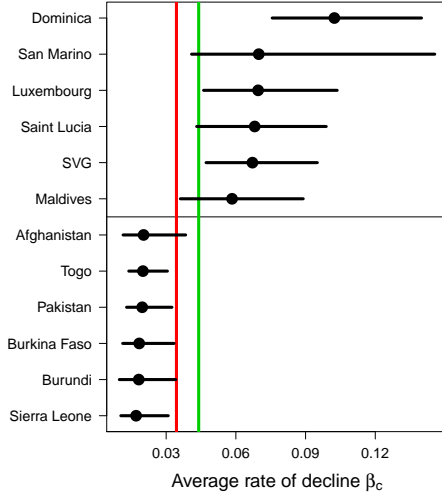


Figure 8. Estimates of the average rate of decline for the six countries with the highest/lowest estimates. Median estimates (dots) and 95% confidence intervals (horizontal line) for average rate of decline β_c in the six countries with the highest, and the six countries with the lowest median estimate. The red vertical line is the median of all country estimates, which is at 3.4%. The green line is drawn at a rate of decline of 4.4%, which corresponds to Millennium Development Goal 4. (*SVG = St Vincent and the Grenadines)

only 3% of the current estimates fall outside the 80% confidence bounds that were constructed based on the training set for the last observation year (Table 2). Similarly, only 3% of the current projections fall outside the 80% projection intervals that were constructed based on the training set (Table 3).

When leaving out the observations in the last five observation years within each country, the results for all countries combined suggest that the model is reasonably well calibrated (19% of the left-out observations fall outside their 80% prediction intervals, and 6% of the observations fall outside their 95% prediction intervals, Table 4). The calibration of the confidence interval for the last observation year within the second training set is also satisfactory: few updated estimates fall outside the previously constructed 80% confidence intervals for that year (Table 5). We expect that at most around 20% of the current estimates are outside the previously constructed projection intervals when projecting ahead for five years beyond the last observation year, which holds true for all countries combined: 18% of the estimates fall outside the previously constructed 80% projection intervals (see Table 6).

Broken down by U5MR level, countries with current U5MR between 40 and 100 deaths per 1,000 live births have more observations falling below their respective 80% prediction intervals than expected (19% instead of 10%, see Table 4), which suggests the rate of decline was under-predicted for some countries in this group. This is confirmed when comparing the current estimates for the last observation year to the projection intervals constructed when excluding the last 5 years of data (Table 6): the current estimates are below the 80% projection intervals for seven out of thirty countries with current U5MR between 40 and 100 deaths per 1,000 live births. This percentage is higher than expected and suggests that the rate of decline was under-predicted in these seven countries, which are Nepal, Tajikistan, Mongolia, Bangladesh, Bolivia, Ghana and Azerbaijan. The list of countries is not very surprising: the first six countries were highlighted earlier because of either a high most recent 5-year rate of decline or evidence of an increase in the 5-year rate of decline. For all countries, the current estimates are within the previously constructed 95% projection intervals.

Discussion

We proposed a Bayesian hierarchical time series model for estimating U5MR and constructing short-term projections, as an alternative to the spline regression modeling approach. An autoregressive model for the annual rate of decline in U5MR captures changes more smoothly than a piece-wise linear regression model. The hierarchical approach for estimating the model parameters has the advantage of sharing information about parameter estimates between countries, without restricting these parameters to be equal.

Cross-validation exercises suggest that the proposed model provides U5MR confidence bounds which are reasonably well calibrated during the observation period. The out-of-sample performance of the model for all countries combined was satisfactory as well, although the rate of decline was slightly underestimated for countries with a recent acceleration in the rate of decline. This is to be expected: projections quantify “our best guesses” and uncertainty therein based on observed trends in the past. If future trends are significantly different from the past, e.g. through interventions which lead to acceleration in the rate of decline that have not been observed in the past, projection intervals are no longer calibrated. This highlights the need to distinguish between U5MR estimates and projections when analyzing progress towards reducing U5MR: estimates represent observed trends, while projections represent the expected future trends based on the past and model assumptions.

The sampling model for the data in the Bayesian model was taken from the IGME approach such that the U5MR estimates from the two modeling approaches are directly comparable. The observations are assumed unbiased on the log-scale, with variance inversely proportional to an observation-specific weight which has been fixed a priori. Biased observations, such as observations from incomplete vital registration systems, are excluded or down-weighted in the current approach. If biased observations are informative of the trend in U5MR, inclusion of the observation and estimating their biases could improve the estimation of U5MR. The sampling model can also potentially be improved upon by including error variance parameters for different data types which are estimated from the data, and by using a hierarchical model to exchange information between countries about measurement errors. Sharing information about measurement errors between countries would give more informed estimates of error variance in countries with fewer observations, and it would remove the need for the restriction on the variance of the distortion terms in the current proposed model.

For countries with high HIV/AIDS prevalence and/or extreme events, IGME uses a loess smoother to model U5MR [2, 11–13]. The main difficulties in the loess smoother approach are the choice of the smoothing parameter and the construction of confidence intervals. In IGME’s 2010 revision of the U5MR estimates, the smoothing parameter was determined based on the number of independent data series within a country [2]. Confidence bounds were not included for the set of countries in which the loess smoother was used. The Bayesian modeling approach could also be used for these countries, although adjusted modeling procedures might be needed for the periods with extreme events or to model reversal in trends in high HIV prevalence countries.

Gaussian process regression (GPR) has been proposed as an alternative modeling approach for estimating U5MR for all countries [14]. In this modeling approach, trajectories of U5MR on the log(10)-scale are draws from a multivariate normal distribution. The prior mean is given by a loess smoother fit to the data. The U5MR estimates from our proposed model cannot be compared directly to the GPR’s estimates because the data set that was used in [14] differs from the data set from *IGME info*. In-sample and out-of-sample validation checks were used to assess model performance of the GPR model, but the reported measures of calibration only assessed how close the U5MR estimates based on the training set were to the left-out data from the test set. The widths of the reported uncertainty bounds from the GPR approach have not been validated. We found that uncertainty in U5MR estimates and its rate of decline is non-negligible for many countries and needs to be taken into account when analyzing progress towards MDG4. For this reason, validation exercises of the confidence bounds of any U5MR modeling approach should be included in model validation exercises, such that results can be compared across models.

Acknowledgments

The authors are grateful to François Pelletier, Patrick Gerland, Cheryl Sawyer, Maria Wong, Pei Rong Seah, Dingxuan Zhang and the members of the UN Inter-Agency Group for Child Mortality Estimation for helpful discussions and insightful comments, to Danzhen You for assistance with the data, and to Jennifer Chunn and Alex Cook for helpful comments on the manuscript.

References

1. UN Inter-agency Group for Child Mortality Estimation (2010) Levels & Trends in Child Mortality. http://www.childmortality.org/stock/documents/Child_Mortality_Report_2010.pdf.
2. UN Inter-agency Group for Child Mortality Estimation (2010) Estimation methods used by the UN inter-agency group for child mortality estimation. Technical report. [Http://www.childmortality.org/stock/documents/Methods for Estimating Child Mortality_2010.pdf](http://www.childmortality.org/stock/documents/Methods_for_Estimating_Child_Mortality_2010.pdf).
3. Hill K, Pande R, Mahy M, Jones G (1998) Trends in child mortality in the developing world: 1960 to 1996. Technical report, UNICEF, New York.
4. UNICEF, the World Health Organization (WHO), World Bank, United Nations Population Division (UNPD) (2007) Levels and trends of child mortality in 2006 estimates developed by the interagency group for child mortality estimation. Working paper http://www.childinfo.org/files/infant_child_mortality_2006.pdf.
5. Lindley DV, Smith AFM (1972) Bayes estimates for the linear model. *Journal of the Royal Statistical Society, Series B* 34: 1–41.
6. Gelman A, Carlin JB, Stern HS, Rubin DB (2004) *Bayesian Data Analysis*. Boca Raton, FL: Chapman & Hall/CRC, 2nd edition.
7. Gelfand A, Smith AFM (1990) Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85: 398–409.
8. Lunn D, Thomas A, Best N, Spiegelhalter D (2000) WinBUGS - A Bayesian modeling framework: Concepts, structure and extensibility. *Statistics and Computing* 10(4): 325–337.
9. Raftery AE, Lewis SM (1992) How many iterations in the Gibbs sampler? In: J M Bernardo et al, editor, *Bayesian Statistics 4*, Oxford University Press. pp. 763–773.
10. Raftery AE, Lewis SM (1996) Implementing MCMC. In: Gilks WR, Spiegelhalter DJ, Richardson S, editors, *Markov Chain Monte Carlo in Practice*, London: Chapman and Hall. pp. 115–130.
11. Cleveland WS, Devlin SJ (1988) Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association* 83: 596–610.
12. Cleveland W, Grosse E, Shyu W (1991) Local regression models. In: Chambers JM, Hastie T, editors, *Statistical Models in S*. Chapman & Hall, New York.
13. Murray CJL, Loakso T, Shibuya K, Hill K, Lopez AD (2007) Can we achieve Millennium Development Goal 4? New analysis of country trends and forecasts of under-5 mortality to 2015. *Lancet* 370: 1040–1054.
14. Rajaratnam JK, Marcus J, Flaxman A, Wang H, Levin-Rector A, et al. (2010) Neonatal, post-neonatal, childhood, and under-5 mortality for 187 countries, 1970-2010: a systematic analysis of progress towards millennium development goal 4. *The Lancet* 375 Issue 9730: 1988–2008.

Appendix

The Bayesian hierarchical time series model is given by (with S_c the first observation year):

$$\begin{aligned}
r_{c,t} &= \log \left(\frac{u_{c,t-1}}{u_{c,t}} \right) \\
r_{c,t} &= \beta_c + \rho_c(r_{c,t-1} - \beta_c) + \varepsilon_{c,t} \text{ for } t \geq S_c, \\
\varepsilon_{c,t} &\sim N(0, \delta_c^2), \\
\log(u_{c,S_c-1}) &\sim N(\log(y_{c,S_c,1}), 1), \\
r_{c,S_c-1} &\sim N \left(\beta_c, \frac{\delta_c^2}{1 - \rho_c^2} \right), \\
\log(\beta_c) &\sim N(m_\beta, \sigma_\beta^2), \\
m_\beta &\sim N(0, 10^2), \\
1/\sigma_\beta^2 &\sim \text{Gamma}(0.01, 0.001), \\
\rho_c &\sim N(m_\rho, \sigma_\rho^2), \\
m_\rho &\sim N(0, 10^2), \\
1/\sigma_\rho^2 &\sim \text{Gamma}(0.01, 0.01), \\
\delta_c &= \lambda \cdot \beta_c \cdot \sqrt{\pi/2}, \\
\lambda &\sim U(0, 1), \\
\log(y_{c,t,s}) &\sim N(\log(u_{c,t}), \sigma_c^2/w_{c,t,s}), \\
1/\sigma_c^2 &\sim \text{Gamma}(0.01, 0.01).
\end{aligned}$$

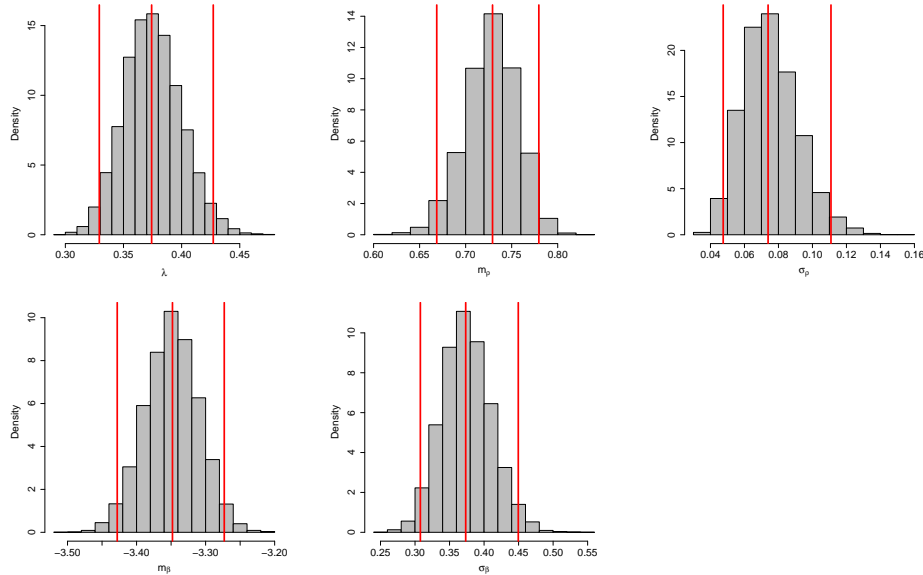


Figure 9. Histograms of the posterior samples of λ , m_ρ , σ_ρ , m_β and σ_β . The 95% confidence intervals and median estimates are represented by the red lines.

Table 1. Cross-validation results when leaving out 20% of the observations.

Observations	# Obs	Above median	Below 80% PI	Above 80% PI	Below 95% PI	Above 95% PI
All	1310	0.46	0.06	0.07	0.02	0.02
U5MR < 40	931	0.46	0.07	0.08	0.02	0.02
U5MR \geq 40	379	0.45	0.06	0.07	0.02	0.01
40 \leq U5MR <100	235	0.44	0.06	0.07	0.02	0.01
U5MR \geq 100	144	0.47	0.04	0.06	0.01	0.01
Expected proportions		0.50	0.10	0.10	0.025	0.025

The proportion of excluded observations that fall above the median U5MR estimate and outside their 80% and 95% prediction intervals (PI), when leaving out 20% of the observations at random. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country (based on the complete data set).

Table 2. Change in U5MR estimate in last observation year when leaving out 20% of the observations.

Countries	# Countries	Above median	Below 80% CI	Above 80% CI	Below 95% CI	Above 95% CI
All	164	0.48	0.02	0.01	0.01	0.00
U5MR < 40	110	0.42	0.03	0.02	0.02	0.00
U5MR \geq 40	54	0.59	0.02	0.00	0.00	0.00
40 \leq U5MR <100	30	0.57	0.03	0.00	0.00	0.00
U5MR \geq 100	24	0.62	0.00	0.00	0.00	0.00
Expected proportions		0.50	\leq 0.10	\leq 0.10	\leq 0.025	\leq 0.025

The proportion of countries in which the median U5MR estimate for the last observation year based on the full data set falls above the median estimate and outside the 80% and 95% confidence intervals (CI), that were constructed based on a training data set in which 20% of the observations were left out. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country based on the complete data set. Smaller outcomes suggest better calibration for the expected proportions that are denoted with \leq 0.10 and 0.025.

Table 3. Change in 5-year U5MR projection when leaving out 20% of the observations.

Countries	# Countries	Above median	Below 80% PI	Above 80% PI	Below 95% PI	Above 95% PI
All	164	0.54	0.02	0.01	0.00	0.00
U5MR < 40	110	0.44	0.03	0.01	0.00	0.00
U5MR ≥ 40	54	0.74	0.02	0.00	0.00	0.00
40 ≤ U5MR < 100	30	0.73	0.03	0.00	0.00	0.00
U5MR ≥ 100	24	0.75	0.00	0.00	0.00	0.00
Expected proportions		0.50	≤ 0.10	≤ 0.10	≤ 0.025	≤ 0.025

The proportion of countries in which the 5-year median U5MR projection based on the full data set falls above the median projection and outside the 80% and 95% projection intervals (PI), which were constructed based on a training data set in which 20% of the observations were left out. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country based on the complete data set. Smaller outcomes suggest better calibration for the expected proportions that are denoted with ≤ 0.10 and 0.025 .

Table 4. Cross-validation results when leaving out the last five observation years.

Observations	# Obs	Above median	Below 80% PI	Above 80% PI	Below 95% PI	Above 95% PI
All	631	0.48	0.10	0.09	0.03	0.03
U5MR < 40	461	0.51	0.08	0.10	0.03	0.03
U5MR ≥ 40	170	0.42	0.14	0.05	0.02	0.02
40 ≤ U5MR < 100	98	0.38	0.19	0.03	0.03	0.01
U5MR ≥ 100	72	0.49	0.07	0.08	0.01	0.03
Expected proportions		0.50	0.10	0.10	0.025	0.025

The proportion of excluded observations that fall above the median U5MR estimate and outside their 80% and 95% prediction intervals (PI), when leaving out the most recent five years of the observation period. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country (based on the complete data set).

Table 5. Change in U5MR estimate in last observation year when leaving out the last five observation years.

Countries	# Countries	Above median	Below 80% CI	Above 80% CI	Below 95% CI	Above 95% CI
All	160	0.45	0.06	0.04	0.00	0.01
U5MR < 40	107	0.50	0.07	0.06	0.00	0.01
U5MR ≥ 40	53	0.36	0.04	0.00	0.00	0.00
40 ≤ U5MR < 100	30	0.27	0.03	0.00	0.00	0.00
U5MR ≥ 100	23	0.48	0.04	0.00	0.00	0.00
Expected proportions		0.50	≤ 0.10	≤ 0.10	≤ 0.025	≤ 0.025

The proportion of countries in which the median U5MR estimate based on the full data set falls above the median estimate and outside the 80% and 95% confidence intervals (CI) in the last observation year in the training set, where the confidence intervals were constructed based on the training set in which the most recent five years of the observation period were left out. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country based on the complete data set. Smaller outcomes suggest better calibration for the expected proportions that are denoted with ≤ 0.10 and 0.025.


Table 6. Change in 5-year U5MR projection when leaving out the last five observation years.

Countries	# Countries	Above median	Below 80% PI	Above 80% PI	Below 95% PI	Above 95% PI
All	160	0.46	0.11	0.07	0.04	0.01
U5MR < 40	107	0.51	0.09	0.10	0.05	0.02
U5MR ≥ 40	53	0.34	0.15	0.02	0.02	0.00
40 ≤ U5MR < 100	30	0.27	0.23	0.00	0.00	0.00
U5MR ≥ 100	23	0.43	0.04	0.04	0.04	0.00
Expected proportions		0.50	≤ 0.10	≤ 0.10	≤ 0.025	≤ 0.025


The proportion of countries in which the median U5MR estimate for the last observation year based on the full data set falls above the median projection and outside the 80% and 95% projection intervals (PI), that were constructed based on a training data set in which observation in the last five observation years were left out. The results are broken down by the median estimate of the U5MR level in the most recent observation year within the country based on the complete data set. Smaller outcomes suggest better calibration for the expected proportions that are denoted with ≤ 0.10 and 0.025.

Figure 10. U5MR estimates for all countries (in alphabetical order). The median estimates from the Bayesian model (BM) are shown in red, and the IGME estimates are shown in blue. The 95% confidence intervals are represented by the red area for the Bayesian model, and the blue area for the IGME estimates. Observations are represented by the black dots.


Afghanistan




Algeria




Andorra




Angola




Antigua & Barbuda



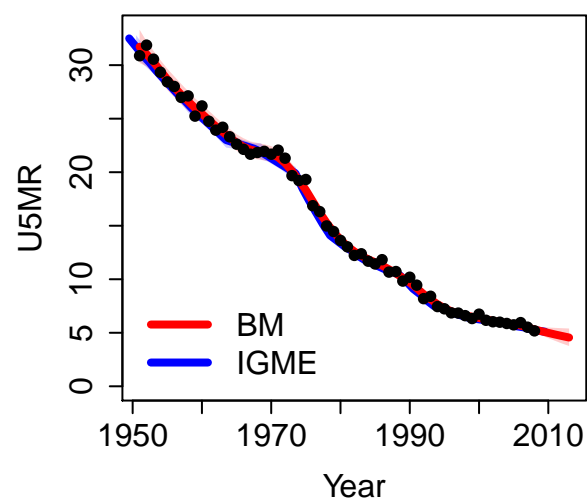
Argentina




Armenia




Australia




Austria




Azerbaijan




Bahamas




Bahrain




Bangladesh




Barbados




Belarus




Belgium




Belize




Benin




Bhutan




Bolivia




Bosnia & Herzegovina




Brazil




Brunei




Bulgaria




Burkina Faso




Burundi




Cambodia




Canada




Cape Verde




Chad




Chile




China




Colombia




Comoros




Congo DR




Congo




Costa Rica




Croatia




Cuba




Cyprus




Czech Republic




Denmark




Djibouti




Dominica




Dominican Republic




Ecuador




Egypt




El Salvador




Equatorial Guinea




Eritrea




Estonia




Ethiopia




Federated States of Micronesia




Finland




France




Gambia The




Georgia




Germany




Ghana




Greece




Grenada




Guatemala




Guinea




Guinea-Bissau




Guyana




Haiti




Honduras




Hungary




Iceland




India




Indonesia




Iran




Iraq




Ireland




Israel




Italy




Jamaica




Japan




Jordan




Kazakhstan




Kiribati




Korea DPR




Korea Rep




Kuwait




Kyrgyzstan




Latvia




Lebanon




Liberia




Libya




Liechtenstein




Lithuania




Luxembourg




Macedonia




Madagascar




Malaysia




Maldives




Mali




Malta




Marshall Islands




Mauritania




Mauritius




Moldova




Monaco




Mongolia




Montenegro




Montserrat




Morocco




Myanmar




Nepal




Netherlands




New Zealand




Nicaragua




Niger




Nigeria




Norway




Oman




OPT




Pakistan




Panama




Papua New Guinea




Paraguay




Peru




Philippines




Poland




Portugal




Qatar




Romania




Russian Federation




Saint Kitts & Nevis




Saint Lucia




Samoa




San Marino




Sao Tome & Principe




Senegal




Serbia




Sierra Leone




Singapore




Slovakia




Slovenia




Solomon Islands




Somalia




Spain




Sri Lanka




St Vincent & the Grenadines




Suriname




Sweden




Switzerland




Tajikistan




Thailand




Timor Leste




Togo




Tonga




Trinidad & Tobago




Turkey




Turkmenistan




Ukraine




United Arab Emirates




United Kingdom




United States of America




Uruguay




Uzbekistan




Vanuatu



Venezuela



Vietnam



Yemen

