"Global population projections by the United Nations"
John Wilmoth, Population Association of America, San Diego, 30 April 2015

Revised 5 July 2015

[Slide 1] Let me begin by thanking Wolfgang Lutz for reaching out to me about this topic last year and initiating the discussion that we are about to have. I also want to thank Jane Menken for agreeing to chair this event, and our two discussants, Mike Teitelbaum and Brian O'Neill, for agreeing to share their insights on this topic.

In the time that I have, I will briefly describe some key aspects of the global population projections produced by the United Nations. The UN's work in this area has changed and, I think, improved over time. I hope that by the end of my presentation you will understand better the background and the meaning of the UN's projections.

[Slide 2] I will start with a bit of history. The UN has produced 23 sets of global population projections, beginning in 1951. Early projections were for the world as a whole and for large regions. Projections for individual countries began in 1968, and the latest set, the 2012 Revision, includes full information for 201 countries or areas.

[Slide 3] This next slide offers a quick overview of the method used by the United Nations to project the populations of countries and the world. All calculations are made within a cohort-component framework. When projecting the population of an individual country, assumptions about future trends in fertility and mortality are derived mostly from information about historical trends for the country itself. Nevertheless, the model of demographic change for an individual country is a specific case of a more general global model, and the available data for a single country are often supplemented by information from other countries, as necessary.

Let me just say quickly that I will focus here on how we model future trends in fertility and mortality. For international migration, the UN assumes, quite simply, that current levels will be maintained for several decades, and then gradually decline toward zero in the distant future. This assumption can be a topic for

discussion on another day. For the rest of this talk, I will focus on fertility and mortality as components of the UN projections.

Traditionally, the UN medium projection has been accompanied by a small number of alternative future trends, derived by varying key assumptions, in particular fertility, or by constructing hypothetical scenarios, such as an instant shift to replacement fertility or to zero migration. More recently, a new method based on a formal statistical model has made it possible to assess the probability of observing a range of plausible future trends.

The median trajectory of this probabilistic model has the same assumptions about future trends in fertility, mortality and migration as the traditional medium-variant projection. Thus, the difference between the two approaches involves only the way of describing alternative future trends, not the central trend itself.

[Slide 4] Projected trends in fertility and mortality are derived from historical information about trends in those same variables. As you all know, observed regularities in historical trends have led to various theories of demographic change. In building a projection model, the UN used existing theories of demographic change to guide the specification of a general model, which was then fitted to the experience of individual countries.

To estimate the model for a given country, we used the data for that country and for other countries as well. When data are plentiful for a given country, the estimation procedure gives most of the weight to data for that country when choosing the parameters of the country-specific model. When data are sparse, the method gives more weight to data for other countries.

[Slide 5] Now, some of you may be saying, there is not a single theory of demographic change that explains *all* changes in *all* countries, or that is commonly accepted by *all* professional demographers. Fair enough. However, the theories that exist about long-term trends in fertility and mortality do share certain common features, which are reflected in the UN's projection model.

For example, in the case of fertility, the model depicts a transition from high to low levels of TFR, while for mortality, the model depicts a transition from low to high levels of life expectancy at birth. As Paul Demeny wrote in 1968, "In traditional societies fertility and mortality are high. In modern societies fertility and mortality are low. In between, there is a demographic transition." [cited in Coale (1973)] This is the most basic and relatively uncontroversial aspect of demographic transition theory: that a transition *has* occurred, *is* occurring or *will* occur. The models of fertility and mortality change that underlie the UN's projection method give form to a general theory of transition.

There is less theoretical consensus about what happens after these major transitions have occurred. Will life expectancy continue to increase? Will fertility remain well below replacement level for a long time? In each case, the UN projections reflect a view of the future that is widespread but probably not universal in the same way that there is almost universal recognition of transition theory. For fertility, it is assumed that post-transition trends will fluctuate and that countries with very low levels of fertility will see at least a modest recovery toward higher levels – not necessarily to the replacement level but in that direction. For mortality, it is assumed that the future increase of life expectancy at birth will be roughly linear, with a rate of change that converges toward the observed increase of the maximum age at death.

[Slide 6] The next slide illustrates the complete fertility model in three phases, corresponding to the periods before, during and after the transition. Note that the rapid drop of the TFR during the transition is followed by a period of fluctuation and modest increase.

[Slide 7] The next slide illustrates the shape of the trend in life expectancy at birth that is embedded in the model. It is a kind of S-curve, with a future increase that decelerates and eventually becomes linear, with a slope that must be chosen based on a reasoned argument. In this case, it was argued that the rate of increase in life expectancy at birth is likely to decline in the future, since in the long run the rate of

increase in the mean of a distribution cannot exceed the rate of increase in its maximum, which has been moving up more slowly.

[Slide 8] The next three slides focus on phase II of the fertility model, or the period of major decline. The model of decline depicts the rate of change in the TFR as a function of the TFR itself. The general form of this relationship is an inverted U, which is modeled using a double-logistic curve. The inverted U-shape implies that the transition begins slowly, proceeds more quickly in the middle, and then slows down at the end. We use Bayesian statistical methods to estimate the double-logistic model for each country, yielding information about both the country-specific parameter estimates and their uncertainty.

[Slide 9] Here is an illustration of the fitted model for India. In the left panel, the solid red line depicts the estimated curve, depicting the rate of change in the TFR as a function of its current level. There is uncertainty around that estimate, as depicted both by the individual simulated trajectories (in grey) used to derive the estimate, and by the percentile bands (in dashed and dotted red) corresponding to the 80- and 95-percent prediction intervals.

In the panel on the right, rates of change from the panel on the left are transformed into projected future changes in the TFR. Again, the uncertainty bands (in red) correspond to 80- and 95-percent prediction intervals. The individual trajectories of the future TFR (in grey) were generated by Monte Carlo simulation, taking into account both the uncertainty of parameter estimates (on the left) and the uncertainty of future annual changes, as described by a random walk. In short, this is a model of a random walk with variable drift, which takes into account the uncertainty about the drift parameter in assessing the uncertainty of future trends.

[Slide 10] Briefly, this slide illustrates how the Bayesian hierarchical model yields estimates of the rate of decline in the TFR for countries with different historical patterns, and also for countries that differ greatly in terms of data availability and their current stage of transition.

[Slide 11] Again briefly, this slide illustrates the third phase of the fertility model, which starts by definition after there has been an increase over two consecutive 5-year time intervals for a country with a TFR that is already below 2. The empirical base used for estimating this portion of the fertility model now includes 25 countries or areas, mostly in Europe but also in North America, East Asia and elsewhere.

[Slide 12] It goes without saying that there is considerable uncertainty about future population trends. A major challenge for demographers has been how to describe and communicate that uncertainty to everyday consumers of population projections. Although the traditional variants and scenarios have no probabilistic interpretation, they remind users of the possibility that future trends may differ from the central forecast. Thus they emphasize the *existence* of uncertainty, even if they do not help very much in understanding the *extent* or the *magnitude* of that uncertainty. For a quantitative assessment of the uncertainty of population projections, one needs a statistical model.

[Slide 13] It is instructive to compare the two means of depicting alternative future population trends in some specific cases. Here, for Nigeria, the medium variant is shown as a solid red line. The other red lines depict the distribution of future trends according to the probabilistic model, while the blue lines depict the high- and low-fertility assumptions (plus or minus half a child).

In the case of a high-fertility country like Nigeria, the high- and low-fertility assumptions are not particularly extreme and do not cover the full range of uncertainty revealed by the probabilistic analysis. In short, if our notion of a "high" and "low" scenario derives from a TFR that varies by plus or minus half a child, we will understate the uncertainty of future trends in this context, as seen on the left for the TFR and on the right for total population.

[Slide 14] The situation is just the opposite for low-fertility countries, such as the Russian Federation. Although future trends are expected to fluctuate, those changes are likely to take place within a fairly small range, which is easily covered by plus or

minus half a child. In this case, defining a "high" and "low" scenario by this traditional method leads to a modest overestimation of the future uncertainty of population trends.

[Slide 15] For large regions or the world as a whole, the high- and low-fertility variants grossly overstate the uncertainty of future trends. This is because the high- and low-fertility assumptions are applied uniformly across *all* countries and *all* future time periods. The uniform application of this assumption yields an exaggerated depiction of the range of plausible future trends. Although some countries will have higher levels of fertility than implied by the medium scenario, others will have lower levels. Likewise, some countries will have fertility levels that are higher than expected in some time periods but lower in others. The impact of these differences will tend to cancel out.

For the UN's projection of world population in 2100, the high- and low-fertility variants differ by around 10 billion, spanning a range from roughly 6 and half to 16 and a half billion. By contrast, the 95-percent projection interval for 2100 runs from 9 to around 13 billion, thus covering a span of roughly 4 billion.

[Slide 16] What have we learned from the new probabilistic projections? We have learned, as mentioned already, that the traditional high- and low-fertility variants tend to exaggerate the uncertainty of future trends in some cases while understating the uncertainty in others. We have also learned that stabilization of the world population within this century seems somewhat unlikely. However, stabilization in that time frame is not impossible or even implausible according to this analysis: we have estimated that there is a 30-percent chance that the world population will either level off or begin to decline by the end of the 21st century.

[Slide 17] The Population Division acknowledges that the assessment of the uncertainty of our global population projections using the probabilistic approach is still a work in progress. The current method does not incorporate several forms of uncertainty as listed here, which may be of major or minor importance compared to the uncertainty already reflected in the model.

I will briefly mention just two items on this list. First, the current method does not incorporate uncertainty related to future trends in international migration. Ongoing research suggests that this omission may lead to a substantial understatement of the uncertainty of future population trends for a country like Germany, where international migration is now the major driver of population change. Second, the current method does not take account of uncertainty related to our knowledge of the past and present demographic situation of a country.

[Slide 18] This graph, for example, shows available data on levels of the TFR in Nigeria since 1970. Obviously, the various sources of information do not provide a clear and consistent signal. Between the 2010 and the 2012 Revisions, the UN revised upward its TFR estimates for Nigeria in the latest time periods, since newly available data suggested a higher level. However, this upward revision was soon contradicted by the next available data point, which became available within a year of publishing the 2012 Revision and suggested a TFR value clearly below the revised estimate.

[Slide 19] To summarize my key points:

- Population projections by the United Nations are derived from models of future trends in the components of change, in particular fertility and mortality.
- The UN's method of population projection has a strong basis in demographic theory. For each country, the models are calibrated using data from the country itself and, especially when data are sparse, from other countries as well.
- The uncertainty of future population trends has traditionally been depicted using variants and scenarios. The 2012 Revision introduced a new method that yields probabilistic statements about plausible future trends.
- Lastly, work on the probabilistic assessment of uncertainty by the UN is ongoing and could benefit from further efforts to incorporate additional sources of uncertainty.

[Slide 20] I would like to acknowledge the efforts of our fine colleagues in the United Nations Population Division, who do the hard work of producing the

estimates and projections of world population that so many of you use in your own work. I also want to thank Professor Adrian Raftery and his collaborators for an extremely fruitful and also very enjoyable collaboration over the past decade.

[Slide 21] Finally, I want to point out that there are many freely available R packages for implementing the UN's probabilistic projection method.

[Slide 22-25] And for those of you who would like more detailed information on the probabilistic projection method, the last four slides of my presentation contain further information about the suite of R programs, as well as an extensive bibliography.

You will be able to obtain a copy of this presentation either from the website of the Population Division or by sending me an email.

Thank you for your attention.