International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236 807 342
Fax: +43 2236 71313
E-mail: publications@iiasa.ac.at
Web: www.iiasa.ac.at

Interim Report **IR-04-049**

Putting Oeppen and Vaupel to Work: On the Road to New Stochastic Mortality Forecasts

Warren Sanderson (wsanderson@notes.cc.sunysb.edu)
Sergei Scherbov (scherbov@iiasa.ac.at)

Approved by

Wolfgang Lutz (lutz@iiasa.ac.at)
Leader, POP Program

September 22, 2004

Contents

Abstract

Oeppen and Vaupel (2002) revolutionized the field of human mortality forecasting by showing that best-practice life expectancy has risen almost linearly from the mid-nineteenth century to the present. In this paper, we present a methodology that makes use of that finding. We show that among a set of 14 low mortality countries, the distribution of life expectancies in the last 40 years has had almost perfectly linear mean and median values. We use this observation to estimate the parameters of models that include both trend error and idiosyncratic error. We compare the outcomes of the new procedure with the United Nations (2003) forecasts for Germany, Japan, and the U.S., where only mortality rates differ. The projections are most similar for Japan and most different for the U.S.

Acknowledgments

This paper was presented at the 24th International Symposium on Forecasting, Sydney, Australia, July 5, 2004.

About the Authors

Warren Sanderson is from the Departments of Economics and History at the State University of New York at Stony Brook, USA.

Sergei Scherbov works at IIASA and at the Vienna Institute of Demography of the Austrian Academy of Sciences.

Putting Oeppen and Vaupel to Work: On the Road to New Stochastic Mortality Forecasts

Warren Sanderson and Sergei Scherbov

1. Introduction

Oeppen and Vaupel (2002) revolutionized the field of human mortality forecasting. They demonstrated that best-practice life expectancy has risen almost linearly from the 1850s to 2000. To this they added a long list of forecasts of limits to human life expectancy that have consistently been proven wrong. While there is certainly no guarantee that the linear rise in best-practice life expectancy will continue into the 21st century, neither is there a compelling argument that it will not. In particular, there is certainly no reason to believe that the year 2000 marked a watershed between a 150-year period of rapid increases in best-practice life expectancy and a subsequent period of much slower gains. The most plausible belief, based on Oeppen and Vaupel (2002), is that, at least in the short run, the remarkable linearity in best-practice life expectancy will remain with us. This paper demonstrates one way to apply the Oeppen and Vaupel (2002) finding in mortality forecasting.

Oeppen and Vaupel (2002) concentrated on best-practice life expectancy because their point was that there is no indication that human life expectancy is nearing a limit. Read narrowly, their work pertains only to a particular order statistic of historical life expectancy distributions. Read more broadly, Oeppen and Vaupel (2002) show that interesting properties emerge when we consider the distribution of life expectancy histories across countries – properties that we would not see if we are considering each country separately. We began looking for other potentially interesting and useful properties by asking the question: Is the linearity that Oeppen and Vaupel observed a characteristic only of best-practice life expectancy, or is it a more general characteristic of the whole group of countries that are experiencing low mortality? To answer this question and to put the Oeppen and Vaupel (2002) finding into perspective, we took the most recent 40 years of data from 14 countries with high life expectancies in the Human Mortality Database (2004). We excluded countries of the former Warsaw Pact because their recent mortality histories could have been affected by the significant political and economic changes they experienced. The countries are: Austria, Canada, Denmark, England and Wales, Finland, France, (Western) Germany, Italy, Japan, the Netherlands, Norway, Sweden, Switzerland, and the United States. Their recent life expectancy histories (for females) can be seen in Figure 1.

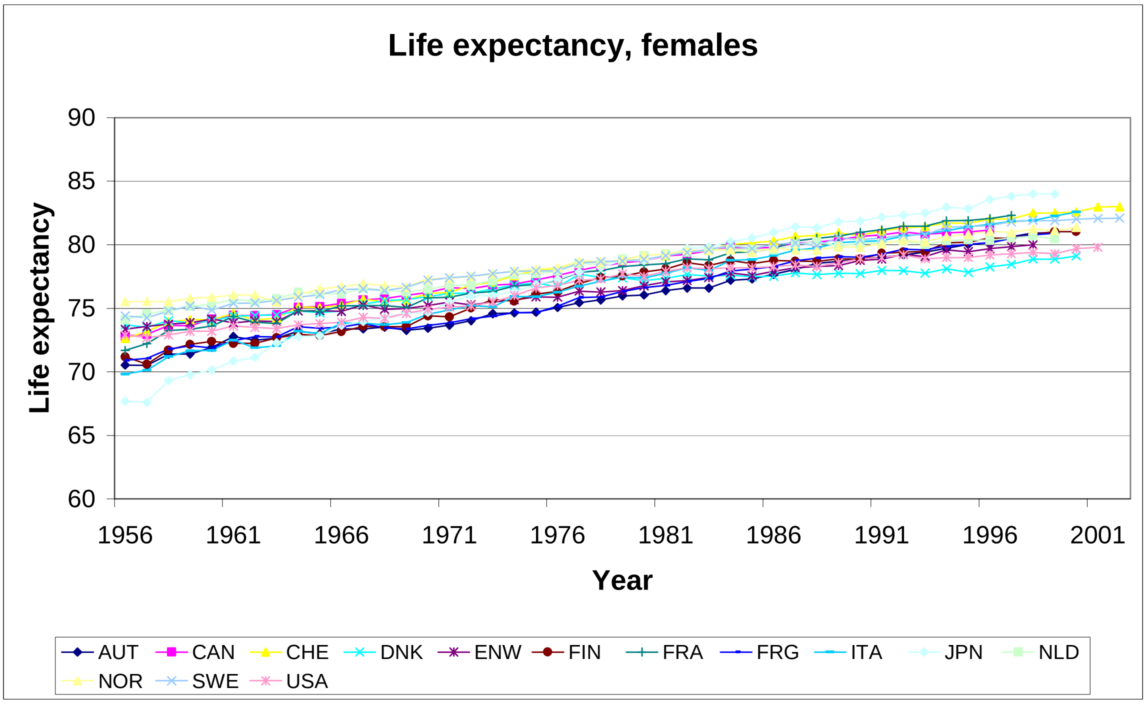

Figure 1. Life expectancy of females in 14 countries with high life expectancies in the Human Mortality Database (2004).

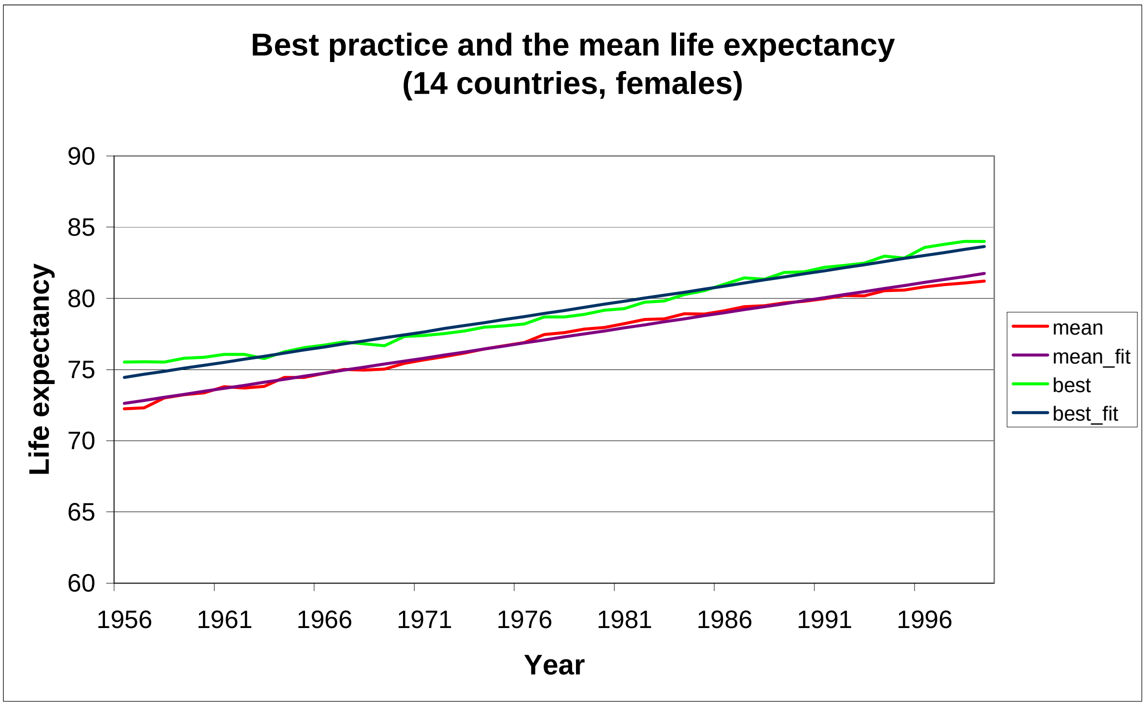

Figure 2. Best-practice and the mean life expectancy of females in 14 countries with high life expectancies in the Human Mortality Database (2004).

Figure 2 clearly answers the question. The graph showing best-practice life expectancy is nearly linear, of course, but so is the graph of the mean of the life expectancy distributions. Indeed, the rate of increase of mean life expectancy and best-practice life expectancy are practically identical, around 2.1 years per decade for the mean and 2.2 years for best practice, and these are statistically indistinguishable from one another. In other words, the linearity that Oeppen and Vaupel observed is not only a characteristic of best-practice life expectancy, but, at least recently, it has also been a general feature of the means of the distributions of life expectancies in low mortality countries.

Fortunately, we do not lack methods for making stochastic mortality forecasts. But do they put Oeppen and Vaupel (2002) to work? The most famous and most widely used method of making stochastic mortality forecasts is the Lee-Carter method (Lee and Carter 1992). Table 1 shows the predicted life expectancies (both sexes combined) for seven low mortality countries using the Lee-Carter method (Tuljapurkar et al. 2000) and the corresponding observed increases in life expectancies over the last two decades.

Table 1. Observed and predicted increases in life expectancies for both sexes combined. Sources: Forecasts for the Lee-Carter model are from Tuljapurkar et al. (2000); observed life expectancies from the Human Mortality Database (2004).

Country	Observed increase of life expectancy per decade (time period)	Lee-Carter predicted increase in median life expectancy per decade: 2000-2020	*Lee-Carter predicted increase in median life expectancy per decade: 2020-2050*
Canada	1.78 (1976-1996)	1.41	*1.20*
France	2.21 (1981-2001)	1.74	*1.5*
Germany	2.29[a] (1979-1999)	1.28	*1.14*
Italy	2.60 (1980-2000)	1.63	*1.35*
Japan	2.62 (1979-1999)	2.27	*1.63*
UK	1.85[b] (1978-1988)	1.32	*1.22*
US	0.87 (1979-1999)	1.35	*1.22*
Median	*2.21*	*1.41*	1.22
Average	2.03	1.57	1.32

[a] Former Federal Republic of Germany (West Germany)
[b] England and Wales

Clearly, the answer is that the Lee-Carter method does not put Oeppen and Vaupel (2002) to work. First, note that the mean observed increase in life expectancy over the last 20 years of available data was 2.0 years per decade. That is consistent with the 2.1 years per decade increase that we observe over a nearly 40-year period in our 14-country sample shown in Figure 2. The predictions for the first two decades of the 21st century using the Lee-Carter model clearly are very different from the change in life expectancies observed in the last two decades. The mean increase falls from 2.0 to 1.6 years. If we were to believe the Lee-Carter results, we would have to believe that the year 2000 saw a downward shift in the trend of life expectancy increases. There is no evidence so far that this has been the case.

Second, note that the Lee-Carter forecasts show the continuing decreases in additional years of life expectancy per decade. In contrast, Oeppen and Vaupel (2002) show a remarkable consistency in the increase in life expectancy for the best-practice case, and Figure 2 shows that, at least in the last four decades of the 20th century, the same was true for the mean of 14 low mortality countries. One can, of course, argue that around the beginning of the 21st century the pattern of life expectancy increase changed – that the future of life expectancy increases will no longer be like the past. But again to make this argument convincing, its proponents would need some evidence. As Oeppen and Vaupel (2002) show, the past is littered with incorrect forecasts based on the presumption that we were nearing the limit to human life expectancy. The Lee-Carter forecasts do not assume such a limit, but they do show a remarkable and continuing slowing in life expectancy gains. Third, the Lee-Carter approach forecasts continuing and increasing gaps in life expectancy between the countries, although history shows that countries may lag behind and after a while become leaders. Clearly, the Lee-Carter method does not put Oeppen and Vaupel (2002) to work and therefore, we must continue looking for an approach that does.

Keilman et al. (2002) use a different methodology for making probabilistic mortality forecasts for Norway. They use annual life expectancies at birth in Norway from 1945 to 1995 separately for men and women, and estimate a multivariate ARIMA (2,0,0) model. However, they modify their estimated coefficients so that median forecasted life expectancies for men and women in 2050 match those assumed in the 1999 forecast made by Statistics Norway. That agency assumed that male life expectancy would rise to 80.0 years and female life expectancy to 84.5 years in 2050. Since in 2000, life expectancy for males was 75.95 years and for females was 81.37 years, the implied increase is a paltry 0.81 years per decade for men and 0.63 years per decade for women. For comparison, between 1950 and 2000, the life expectancy of males and females increased by 1.21 and 1.62 years per decade, respectively. Clearly, the small increases in life expectancy forecasted by Statistics Norway do not put Oeppen and Vaupel (2002) to work. Perhaps the ARIMA (2,0,0) model unconstrained to match the forecast of Statistics Norway would have.

In general, adding uncertainty to the medium variant forecasts of official organizations does not put Oeppen and Vaupel (2002) to work because those forecasts typically show much less life expectancy gain. This can be seen from Table 2, where we reproduce forecasts gathered in Lundström (2003). Indeed, most of the official mortality forecasts shown there assume increases in life expectancy of less than one year per decade over the period 2000 to 2050, less than half the rate of the most recently observed changes.

Table 2. Average life expectancy for women, 2000-2050. Forecasts in different countries. Source: Lundström (2003).

Country	2000	2010	2020	2030	2040	2050
France	83	84.8	86.5	87.9	89.2	90.4
Belgium	81.1	82.3	83.6	85	86.5	88.1
Austria	81.3	82.8	84.2	85.5	86.2	87
Switzerland	83	84.4	85.6	86	86.4	86.9
USA	79.9	81.4	82.9	84.2	85.4	86.6
Sweden	82.1	83.4	84.4	85.3	86	86.5
Japan	84.1	85.1	85.6	86	86.3	86.5
Finland	80.9	82	83.1	84	84.8	85.5
Italy	82.3	83.5	84.7	84.7	84.7	84.7
Norway	81.5	82.7	83.5	84.1	84.4	84.5
UK	80.1	81.5	82.6	83.2	83.6	83.8
Netherlands	80.6	81.1	81.6	82.2	82.7	83
Denmark	78.5	78.6	78.6	78.6	78.6	78.6

Lutz, Sanderson, and Scherbov (2004, 2001, 1996) have devised two methods for stochastic mortality forecasting that, in different ways, are closer to Oeppen and Vaupel (2002). We call our two approaches LSS1 and LSS2. They were devised to be applied to large regions of the world, not individual countries. Nevertheless, they are relevant to our discussion here. In both, the mean value of life expectancy increases linearly by two years per decade, and so both are broadly consistent with recent observations. In LSS1 (see Lutz et al. 1996) stochastic realizations of life expectancy paths are straight lines originating from observed life expectancies in 1995. The trends are drawn from a normal distribution with a mean increase of two years of life expectancy per decade and with the standard deviation determined so that there would always be a five percent chance that life expectancy would be less than it was in 1995. LSS1, in other words, had only trend uncertainty; all that uncertainty was immediately realized in 1995. Although on average LSS1 captures the observed trends, the individual paths do not seem fully consistent with the expectation that the observed trend in life expectancy increase would persist, at least in the short run.

LSS2 (see Lutz et al. 2001, 2004), in contrast, has no trend uncertainty at all and has only idiosyncratic uncertainty around a life expectancy trend that increases by two years per decade. LSS2, then, is potentially a candidate for putting Oeppen and Vaupel (2002) to work. But when we estimated the parameters of the LSS2 model for the 14 countries in our low mortality sample, we found that the standard deviation of life expectancies at the end of the century was much lower than we expected.

Thus, we had to set off to find a new model that was both consistent with Oeppen and Vaupel (2002) and consistent with the empirical data on life expectancy changes in our 14 sample countries. It is probably not surprising that we found the most consistent model to be one that combined the features of both LSS1 and LSS2. The model contains both trend uncertainty and idiosyncratic variation around the uncertain trends. LSS1 and LSS2 are both special cases of model presented here. The combination of the features of both is required to fit the data and to make Oeppen and Vaupel work.

We proceed as follows. In Section 2, we present the new framework and talk about how we estimated its parameters for three countries, Germany, Japan, and the U.S. The third section contains a comparison of our forecasts for those three countries and those of the United Nations (2003). The comparison is constructed so that the forecasts differ only in mortality assumptions. We present some concluding thoughts about putting Oeppen and Vaupel (2002) to work in Section 4.

2. The Model and Its Estimation

We chose the 14 countries from which we took the data based on three criteria: (1) Data had to be available for at least 40 years after 1955 in the Human Mortality Database (2004), (2) there could be no significant economic or social changes that could have caused a break in the long-run trend of life expectancy changes, and (3) all the countries had to have high life expectancies. The linearity of the mean life expectancy that we saw in Figure 1 and the idea, based on Oeppen and Vaupel (2002), that we can learn about mortality change by studying distributions of outcomes, suggested to us that we could treat the latest four decades of life expectancy observations for the 14 countries as being generated from a common statistical model. That model is:

$$e_0(t,c) = \tau(t) + \mu(t,c), \quad (1)$$

where $e_0(t,c)$ is female life expectancy at time t in country c, $\tau(t)$ is the mean value of the set of 14 female life expectancies at time t, as shown in Figure 1, and $\mu(t,c)$ is the deviation from the mean at time t in country c.

As in LSS2, we specify that $\mu(t)$ has a moving average representation

$$\mu(t,c) = \frac{\alpha \sum_{i=t-n+1}^{t} x_{i,c}}{n}, \quad (2)$$

where n is the order of the moving average, the $x_{i,c}$ are independently distributed realizations from a standard normal distribution, and the standard deviation of $\mu(t,c)$ depends on α according to the formula:

$$sd\left(\mu(t,c)\right) = \frac{\alpha}{\sqrt{n}}. \quad (3)$$

The moving average specification in Eq. (2) is unusual because the coefficient on each of the moving average terms is constant (= α). We also estimated a model where the coefficients varied, but the results were so similar to the constant coefficient version in Eq. (2) that we chose the simpler approach. First, we discuss how we estimate the two parameters, α and n and then move on to how we use those estimates in our forecasts.

Since both the life expectancies and $\tau(t)$, the trend, are observed, we can calculate the $\mu(t,c)$ from Eq. (1). The $\mu(t,c)$ are a function of the two parameters, α and n. We estimate these two parameters using the 560 (14 countries times 40 years) observations on the μ's. The autocorrelation function depends only on n and we used this fact to find the value of *n* that best fit it. We computed the autocorrelation function for correlations between the $\mu(t,c)$ from one year apart to 30 years apart. For autocorrelations *i* years apart we had $(40-i)\cdot 14$ pairs of observations.

To find the value of *n* that best fit the data, we defined the statistic *S* as:

$$S(n)=\sum_{i=1}^{30}\left(\frac{\left[\gamma(i)-\bar{g}(i,n)\right]}{sd\left[g(i,n)\right]}\right)^{2}, \tag{4}$$

where $\gamma(i)$ is the observed autocorrelation between the $\mu(t,c)$ that are *i* years apart (using the $(40-i)\cdot 14$ relevant pairs of observations), $g(i,n)$ is a single simulated measurement of the autocorrelation of terms that are *i* years apart (based on $(40-i)\cdot 14$ pairs of simulated values), conditional on there being *n* terms in the moving average; $\bar{g}(i,n)$ is the expected value of the $g(i,n)$ and $sd[g(i,n)]$ is the standard deviation of the autocorrelation, computed again assuming the same number of observations used in calculating the $\gamma(i)$.

The $\bar{g}(i,n)$ and the $sd[g(i,n)]$ were computed using a bootstrap procedure. We did 1,000 iterations. In each iteration we took 14 time series from our 14 countries with replacement. In other words, in a single iteration, we could include France twice, Norway twice, and have no observations for the U.S. or Japan. The probability that a single country would be appear only once among the 14 time series in a given iteration is 2.7 percent (100*(1/14)*(13/14)^13). We found that the minimum value of $S(n)$ occurred at *n* equals 44. We also computed the $\bar{g}(i,n)$ and the $sd[g(i,n)]$ on the assumption that the moving average specification was exactly correct. When we minimized $S(n)$, in this case, we obtained an *n* of 34. The difference between a value of *n* of 44 and 34 makes only a trivial difference in the mortality forecasts. We decided to use an *n* of 44 because it kept us closer to the observed data.

Using $n = 44$, we performed the analogous calculation, on the standard deviations of life expectancy differences over various spans from one year to 30 years in order to compute α. We defined the test statistic $D(n)$ as:

$$D(\alpha)=\sum_{i=1}^{30}\left(\frac{\left[\delta(i)-\bar{d}(i,\alpha)\right]}{sd\left[d(i,\alpha)\right]}\right)^{2}, \tag{5}$$

where $\delta(i)$ is the observed standard deviation of differences in $\mu(t,c)$ values *i* years apart; $\bar{d}(i,\alpha)$ is the mean value of those standard deviations when we are considering differences in $\mu(t,c)$ values *i* years apart, for a specific value of α; and $sd[d(i,\alpha)]$ is the standard deviation of $d(i,\alpha)$. We followed an analogous bootstrap procedure to the one described above. We found that $D(\alpha)$ was minimized at $\alpha = 7.3$, implying a standard deviation of the $\mu(t,c)$ of 1.1 years of life expectancy.

The model above is almost identical to the one in Lutz et al. (2001). The main difference is that here we estimated a value of α that fit the data for high life expectancy countries that did not have a trend shift within the last 40 years of observations. In Lutz et al. (2001) we let α increase gradually to reflect the increasing uncertainty of life expectancies with the passage of time that we derived from expert opinion.

The difference between these two views provides us with both a challenge and an opportunity. The annual variability of observed life expectancy paths does not increase over time. Nevertheless, our uncertainty about the future certainly increases over time. We bridge the gap by adding an element that was crucial in Lutz et al. (1996), trend uncertainty.

Oeppen and Vaupel (2002) make a convincing case that best-practice life expectancy is not slowing down. We show that the same has been true of the mean life expectancy in our 14 country sample. There is certainly no plausible forecasting methodology, which we know of, that suggests that life expectancy trends should change at the instant that a forecast begins. Instead, it is more likely that current life expectancy trends will continue for a while. But for how long?

Here we assume that there will be a trend change sometime during the 21st century. We assume that the size of the trend shift is more likely to be small than large, and that it will occur in a random year between 2002 and the end of the century. The probability of a trend shift happening in any given year is 1/97. The first trend shift can happen in 2002 and the last in 2099, so each of the 97 years has an equal probability of being the one in which the trend shift occurs. We assume that the change in the trend, when it happens, has a normal distribution with mean zero.

Our full model, then, is:

$$e_0(t,c) = \tau(t) + \theta(t,c) \cdot \eta(c) \cdot \left(t - t^*\right) + \mu(t,c),$$

where $\theta(t,c)$ is zero for all the years prior to the trend shift and is unity thereafter, *n(c)*is the value of the trend shift, and t^* is the year in which the trend shift occurs. $\eta(c)$ is normally distributed with mean zero and standard deviation ν. This leaves us with the problem of estimating the standard deviation of the trend change distribution, v.

In keeping with the spirit of Oeppen and Vaupel (2002), we decided to estimate v on the basis of observed life expectancies. We used a bootstrap procedure, similar to the one we used in the estimation of α, with the life expectancy data for six countries that had long time series, England and Wales, Norway, and Sweden (dating back to 1847), Switzerland (beginning 1876), Italy (beginning 1872), and France (beginning 1899). The resulting bootstrap measure of the standard deviation of life expectancy differences using observations between 80 and 100 years apart was 5.3 years of life expectancy. In this paper, we rounded the standard deviation up to 6.0 years of life expectancy for predictions 100 years in the future. Our estimate of v was set so that the standard deviation of life expectancy differences 100 years apart is 6.0 years.

Our incorporation of trend change is consistent with the available time series data on life expectancy. Figure 3 shows the long-term evolution of female life expectancy in five of the countries that we used in our analysis. Italy was omitted because its life expectancy in the late 19th century was considerably lower than the other five. Three periods of stable but distinct trends are clearly evident: Prior to 1900, 1900-1960, and 1960 onwards. We constructed our 14 country sample on the condition that their most recent 40-year histories did not include a trend change. This allowed us to estimate n and α first and then use those estimates in computing v from the longer time series data.

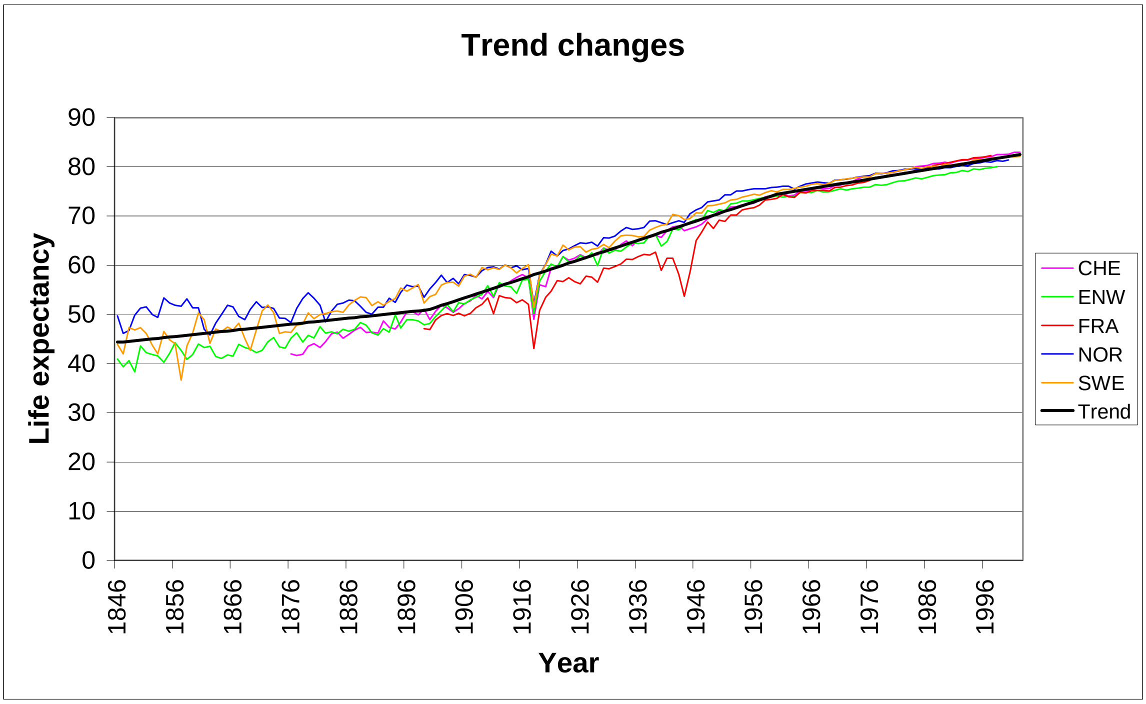

Figure 3. Long-term evolution of female life expectancy in five countries used in the analysis.

3. The Results

What difference does it make if we put Oeppen and Vaupel (2002) to work? In order to answer this question, we need suitable standards of comparison. The obvious candidates here are the UN forecasts. Since our interest is in assessing the effects of putting Oeppen and Vaupel (2002) to work, we have produced probabilistic forecasts using fertility and migration assumptions that are as close as possible to the United Nations (2003) long-run projections. To keep our presentation manageable, we show the results only for Germany, Japan, and the United States.

Figure 4 shows our life expectancy forecasts for Japanese women and the UN forecast. The two agree well up to 2050. The UN forecasts a nine year increase in life expectancy from 2000 to 2050, a 1.8 year increase per decade. We predict about the same, because since Japan has a well above average life expectancy in 2000, our methodology requires that its median value will be closer to the trend line in 2050. After 2050, the UN forecasts a significant slowing in Japan's life expectancy growth. Female life expectancy is projected to grow by three years in the second half of the century, from 93 years to 96 years, at a rate of 0.6 years increase per decade. We do not assume such a slow down, on average (although half of our life expectancy paths do exhibit a slowdown), and show a median life expectancy for Japanese women in 2100 of around 105 years.

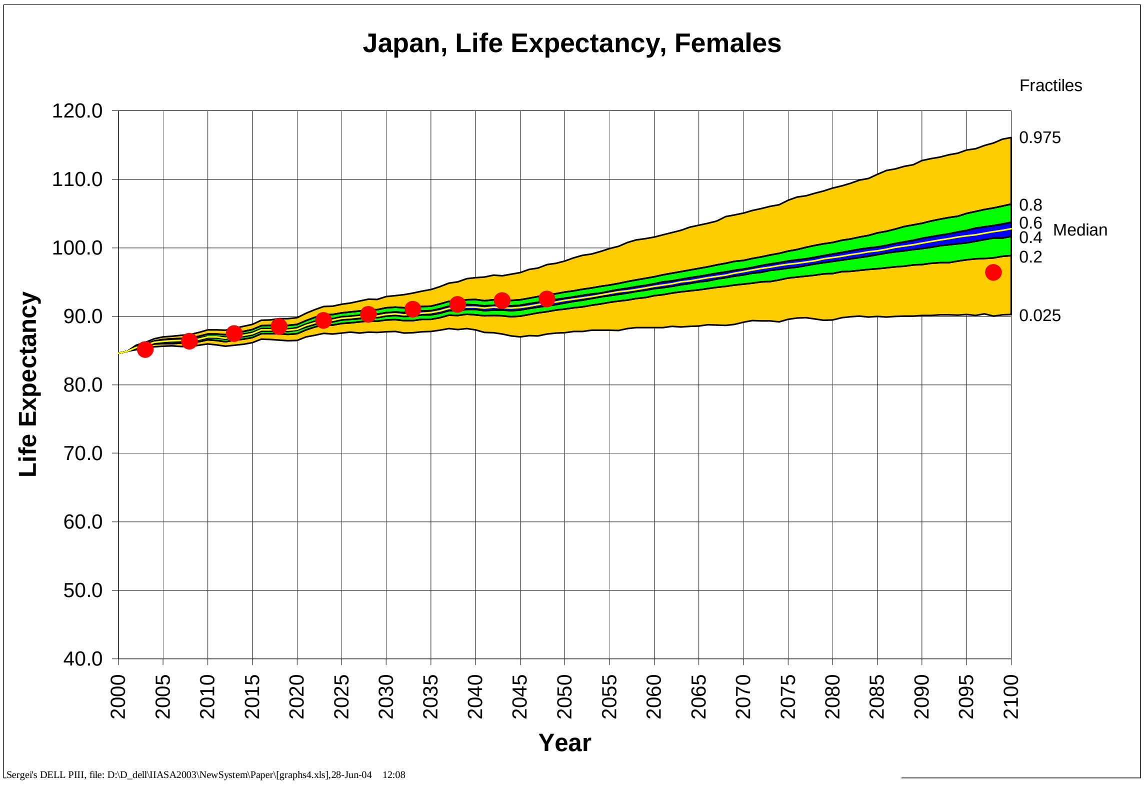

Figure 4. Probabilistic forecast of life expectancy and the UN medium variant, Japan, females.

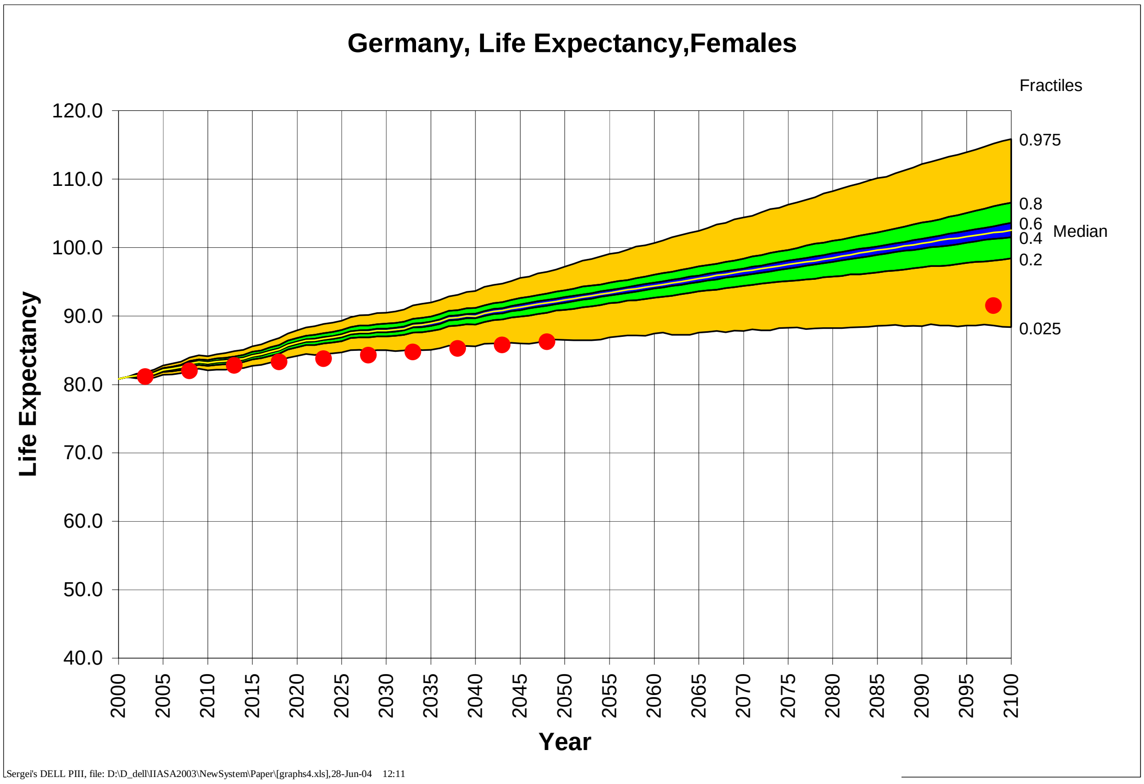

Figure 5. Probabilistic forecast of life expectancy and the UN medium variant, Germany, females.

Figure 5 shows a similar graph for female life expectancy in Germany. The UN forecasts a slower increase in life expectancy there than in Japan in the first half of the century, an increase of 1.2 years per decade, and a faster increase in the second half of the century, again 1.2 years per decade. After 2020, the forecasted life expectancy for German women remains around the 0.025 fractile of our distribution of life expectancies. In this case, putting Oeppen and Vaupel (2002) to work implies that the UN life expectancy path for German women is rather unlikely.

Figure 6 shows the corresponding graph for women in the United States. The UN life expectancy forecast increases by one year per decade in the 2050 period, and only by 1.2 years per decade in the second half of the century. The U.S. is rather far below the median of the 14 country sample in 2000, and in this sense is the mirror image of the Japanese case. Our methodology assumes that, on average, the U.S. female life expectancy will eventually approach the forecasted median line more closely and therefore, will on average (in the long run) have faster life expectancy increases than Japan. The UN life expectancy forecasts are below the 0.025 fractile of our life expectancy distributions throughout the first half of the century. By 2100, the UN forecast is just barely above the 0.025 fractile. Again, from our perspective, the UN forecasts seem rather low.

Population sizes are shown in Figures 7, 8, and 9, and median ages in Figures 10, 11, and 12. Since mortality paths are the only population change component that varies between the UN projections and the new ones given here, the life expectancy changes are closely reflected in the population sizes and median ages. The UN projections and ours are closest for Japan and most different for the United States, with Germany always being an intermediate case. It is easy to misjudge the differences based on these graphs. Using the UN forecast, the U.S. population grows on average by 0.4 percent per year over the century. Putting Oeppen and Vaupel (2002) to work yields a growth rate of 0.5 percent. The U.S. population is only around 12 percent higher in 2100 using the new projections than using the UN's. There would be much more variability in measures like the proportion of the population at age 80 and above, but those numbers have not yet been published by the UN for 2100, so we cannot make the comparison at this time.

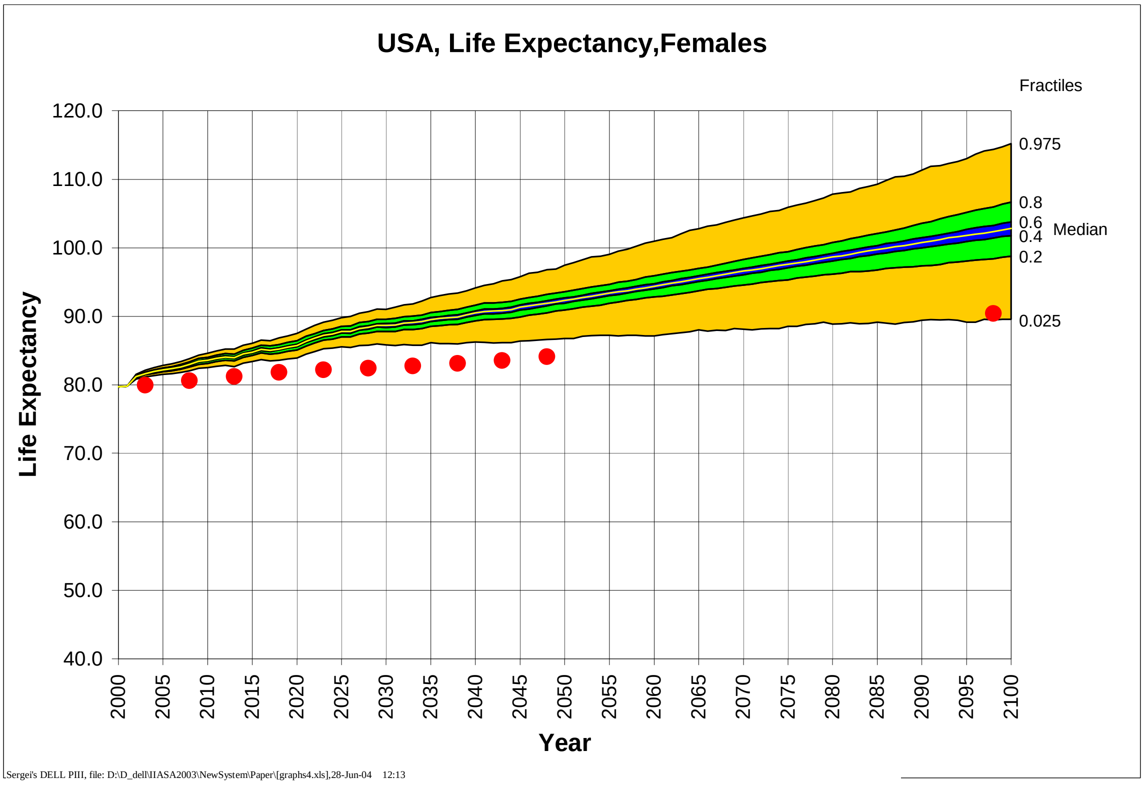

Figure 6. Probabilistic forecast of life expectancy and the UN medium variant, USA, females.

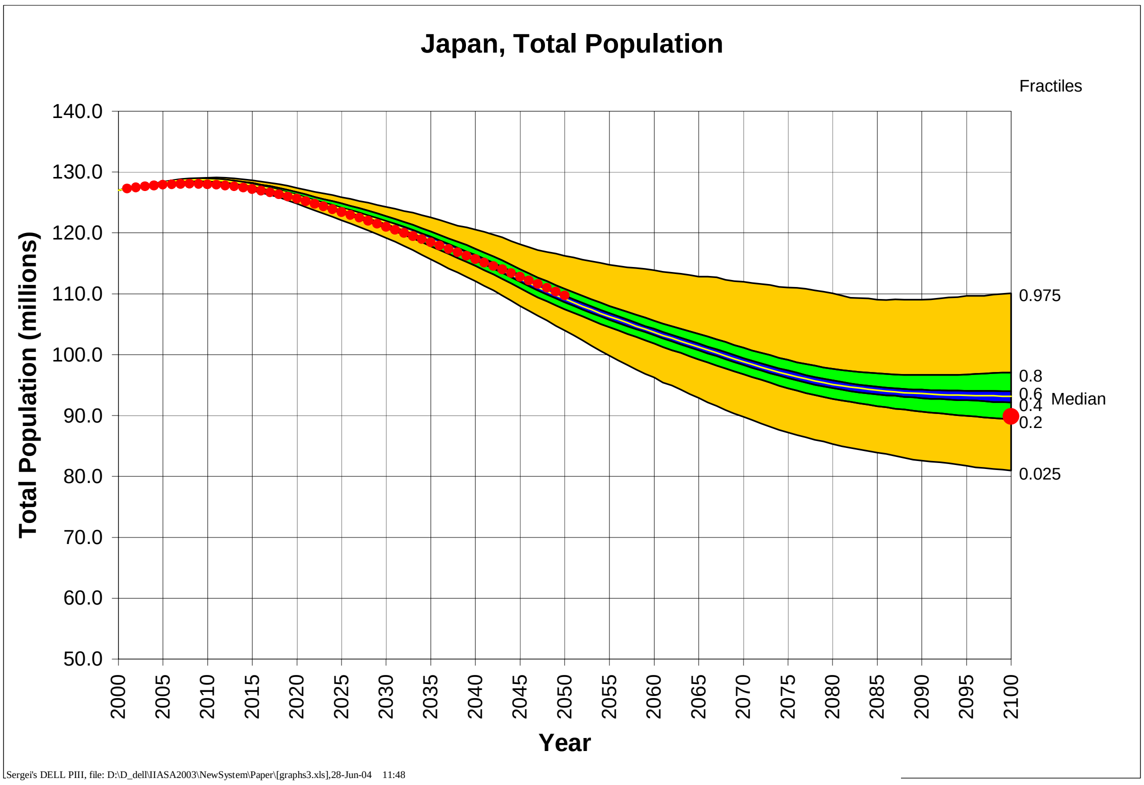

Figure 7. Probabilistic forecast of population size and the UN medium variant, Japan.

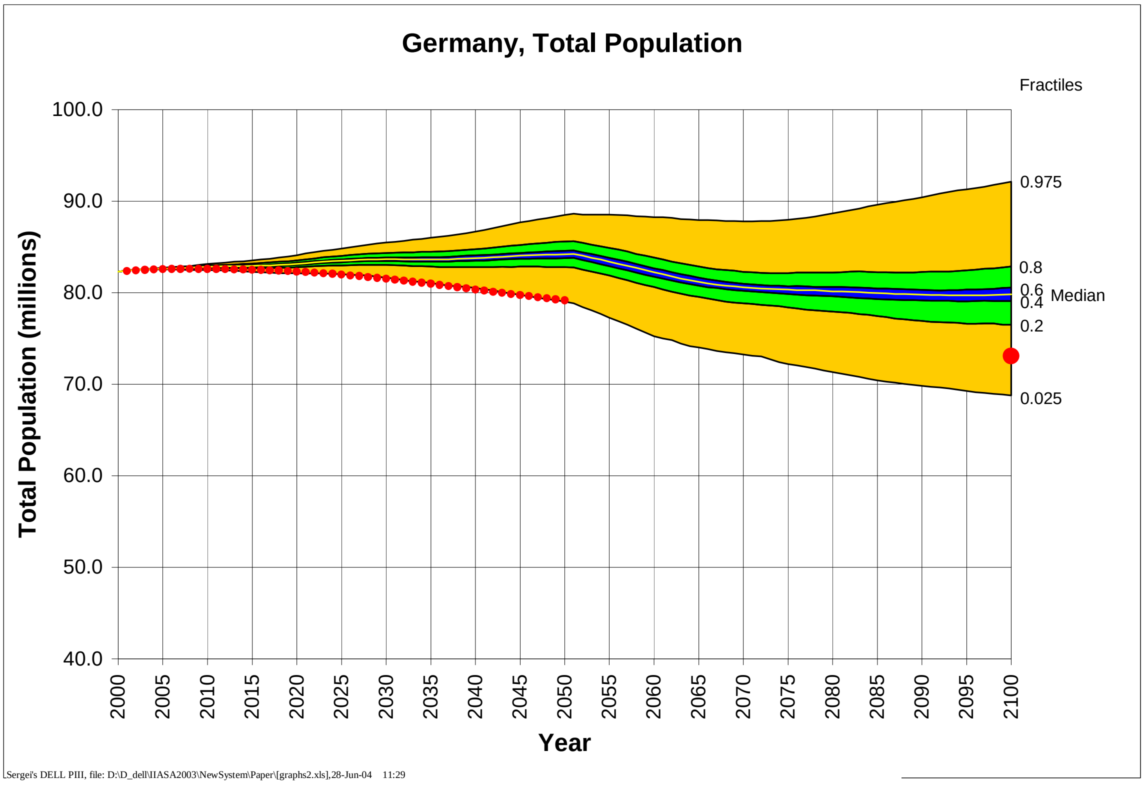

Figure 8. Probabilistic forecast of population size and the UN medium variant, Germany.

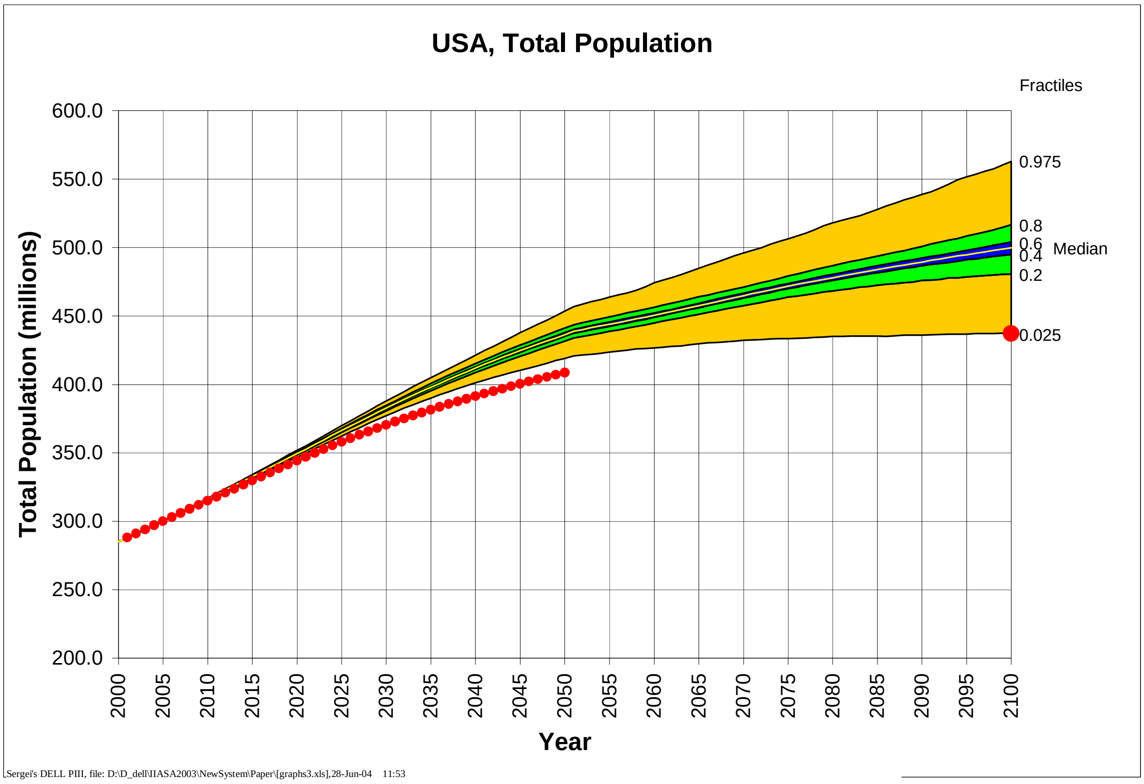

Figure 9. Probabilistic forecast of population size and the UN medium variant, USA.

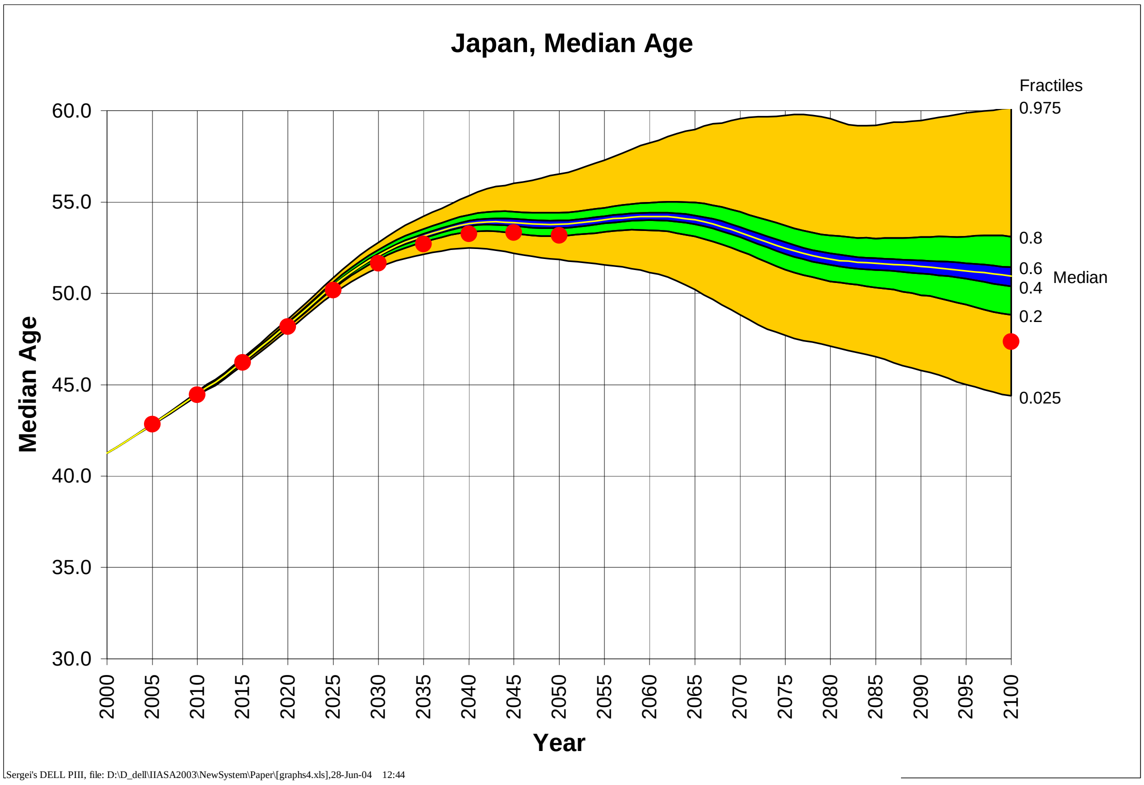

Figure 10. Probabilistic forecast of median age and the UN medium variant, Japan.

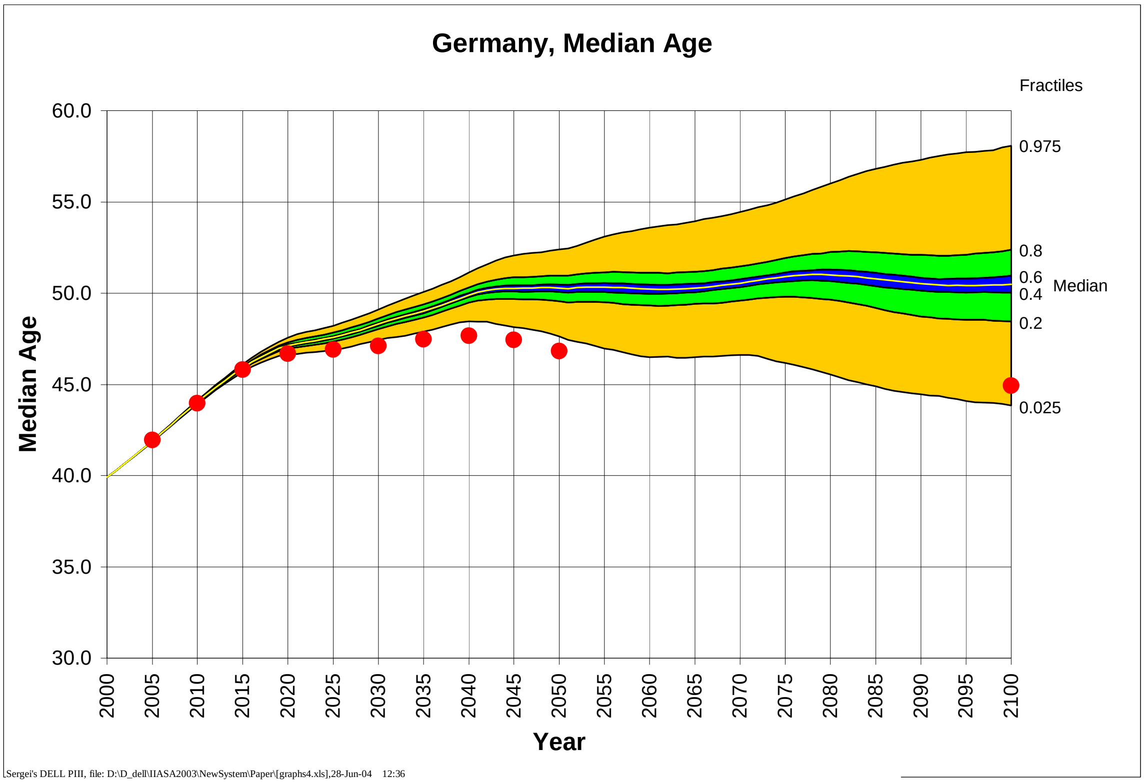

Figure 11. Probabilistic forecast of median age and the UN medium variant, Germany.

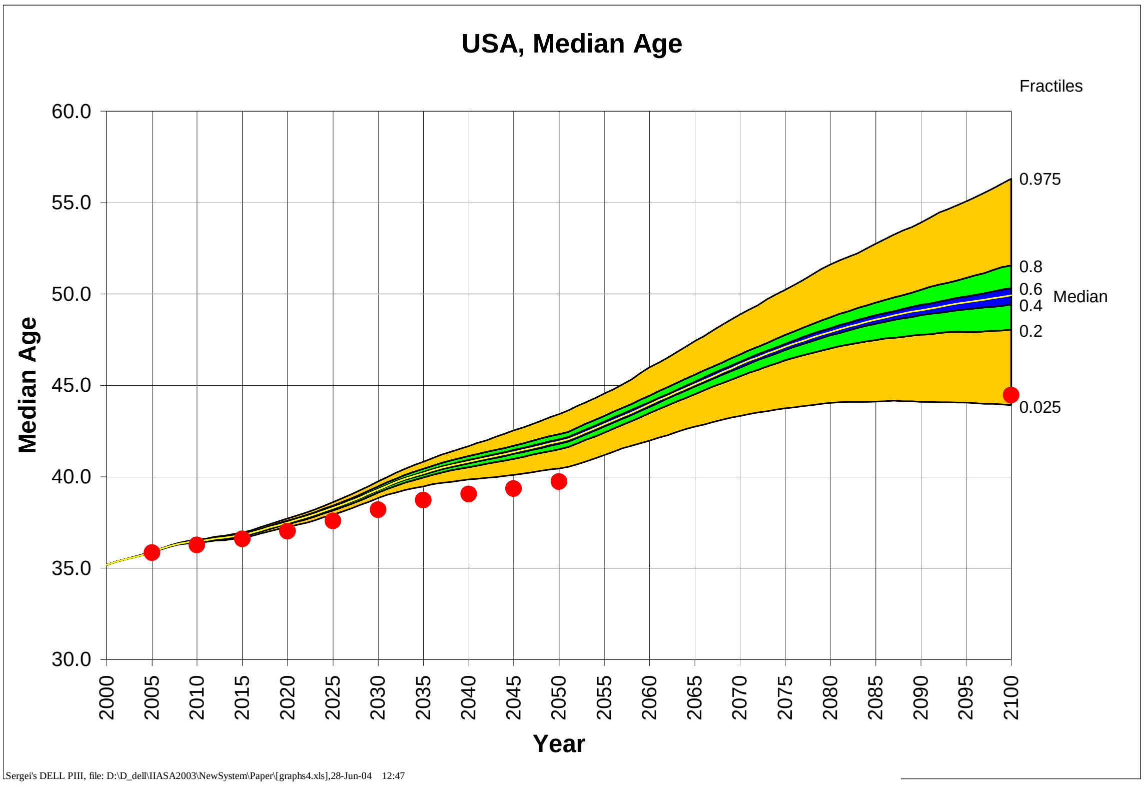

Figure 12. Probabilistic forecast of median age and the UN medium variant, USA.

4. Conclusions

In this paper we have put Oeppen and Vaupel (2002) to work. The result is a new methodology for forecasting life expectancy changes in low mortality countries. Our methodology has applied three key lessons that we have learned from Oeppen and Vaupel (2002). First, our methodology produces life expectancy forecasts in which the immediate future exhibits a continuation of the trends that we have seen in the recent past. It was surprising to us to see how few mortality forecasts do this. Second, our methodology produces sets of life expectancies in which one country does not remain the life expectancy leader into the indefinite future. History shows that the country with the best-practice life expectancy changes. Our procedure is consistent with this history, while forecasts that assume that the leader today will always be the leader are not. Third, our forecasts do not show that life expectancy increases must slow down. We allow for the possibility that life expectancy increases could slow down as well as for the possibility that they may speed up.

We have also followed in the spirit of Oeppen and Vaupel (2002). They showed that the assessment of many experts on the limitations of human life expectancy have consistently been proven wrong. This does not mean that expert opinion on mortality is worthless. Indeed, we remain strong believers in the value of expert opinion in demographic forecasting generally, but only when that opinion is consistent with observations. Expert opinions that invariably fail that test need to be rejected and alternatives need to be found. Oeppen and Vaupel (2002) show the power of insightful analysis of empirical life expectancy data. In our view, they have produced the best approach to understanding the evolution of life expectancies and we have followed them by basing our model and our parameter estimates on the same sorts of data.

Our work on putting Oeppen and Vaupel (2002) to work is not complete. The methodology that we offer here is appropriate only for low mortality countries. An analysis that can be used for countries that are catching up to the leaders remains on our research agenda. Finally, ours is, of course, only one of the many possible ways of putting Oeppen and Vaupel (2002) to work. Others are undoubtedly developing ways to do this as well. Certainly, Oeppen and Vaupel (2002) will never be unemployed.

5. References

Human Mortality Database. 2004. J.R. Wilmoth and V. Shkolnikkov, directors. http://www.mortality.org, accessed March 2004.

Keilman, N., D.Q. Pham, and A. Hetland. 2002. Why population forecasts should be probabilistic – illustrated by the case of Norway. *Demographic Research* 6(15): 407-453. http://www.demographic-research.org.

Lee, R.D. and L. Carter. 1992. Modeling and forecasting the time series of U.S. mortality. *Journal of the American Statistical Association* 87(419): 659-671.

Lundström, H. 2003. Mortality assumptions for Sweden. The 2000-2050 population projections. Pages 59-73 in T. Bengtsson and N. Keilman (eds.), *Perspectives on Mortality Forecasting*. Social Insurance Studies Nr. 1. Stockholm: Swedish National Social Insurance Board.

Lutz, W., W.C. Sanderson, and S. Scherbov. 2004. The end of world population growth. Pages 17-83 in W. Lutz, W.C. Sanderson, and S. Scherbov (eds.), *The End of World Population Growth in the 21st Century: New Challenges for Human Capital Formation and Sustainable Development*. London: Earthscan. London.

Lutz, W., W.C. Sanderson, and S. Scherbov. 2001. The end of world population growth. *Nature*. 412: 543-545.

Lutz, W., W.C. Sanderson, and S. Scherbov. 1996. Probabilistic population projections based on expert opinion. Pages 397-385 in W. Lutz (ed.), *The Future Population of the World: What Can We Assume Today?* London: Earthscan, revised edition.

Oeppen, J. and J.W. Vaupel. 2002. Broken limits to life expectancy. *Science* 296: 1029-1031.

Tuljapurkar, S., N. Li, and C. Boe. 2000. A universal pattern of mortality decline in G7 countries. *Nature* 405: 789-792.

United Nations. 2003. *World Population Prospects: The 2002 Revision Population Database*. New York: United Nations. http://esa.un.org/unpp.