A Discussion of Three Papers on Mortality "Laws" and Models

Thomas P. Edwalds, FSA, ACAS, MAAA
Munich American Reassurance Company

Presented at The Living to 100 and Beyond Symposium
Sponsored by the Society of Actuaries

Orlando, Fla.

January 12–14, 2005

Copyright 2005 by the Society of Actuaries.
All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

The three researchers who prepared and presented papers for this session have done an excellent job of exploring some alternative mathematical models of mortality. I'd like to begin by making some technical comments about each of these papers, then discuss the implications of these models for estimating advanced age mortality.

Technical Comments

In her paper, "'Makeham-Type' Mortality Models," Marie Redina Mumpar-Victoria illustrates a possible approach to estimating mortality rates when data is sparse. While the problem of sparse data is particularly severe in small countries such as the Philippines, this problem is universal at the most advanced ages.

Ms. Mumpar-Victoria offers two alternatives to the traditional Makeham model—the Inverse-Makeham model and the Modified-Makeham Select model—that sketch out possible trajectories for advanced age mortality given a mortality data set lacking advanced age data. Not surprisingly, the Modified-Makeham Select model was a better fit for her data than the Inverse-Makeham model. This would be expected, because the Modified-Makeham Select model has seven parameters, while the Inverse-Makeham has only three. In addition, the Modified-Makeham Select model takes advantage of the select-and-ultimate structure of the data, while the Inverse-Makeham could only be fit to ultimate mortality experience in the data.

What is surprising is that these models predict remarkably low mortality rates over age 90. Of course, Ms. Mumpar-Victoria is careful to point out that extrapolations become less and less reliable the farther they are from the range of the fitted data, so these predicted mortality rates should be viewed with a skeptical eye. Nonetheless, it is worthy to note that mortality rates below 0.15 or even 0.1 at ages up to 96 are plausibly consistent with the experience data at the working ages.

Kathryn Robertson has presented an excellent exposition of Extreme Value Theory (EVT) and applied it to the question of advanced age mortality in her paper, "An Extreme Value Analysis of Advanced Age Mortality Data." This approach is quite promising, because the theory is robust and its application to this question is intuitively appealing. EVT techniques can be used without making any assumption about the underlying distribution, and meaningful results can be obtained with relatively few data points.

I like the fact that Ms. Robertson chose to use birth cohort data for her analysis. I believe that this approach is the most consistent with the underlying assumptions of EVT, in that we are trying to estimate the distribution of the most extreme values of

survival time from birth. This also means that the trends over time from one birth cohort to the next are not confounded by dependence relationships.

Of course, when using EVT techniques, the treatment of outliers is critical. The Japanese data Ms. Robertson used contained one extreme observation of death at age 120, which she chose to reject as an outlier and replace with a more plausible value. She clearly justifies her approach and her choice of a replacement value, but other choices could also be reasonably justified. Inclusion of that observation without modification would have radically altered her results. Alternative choices for replacement values would also have significantly affected her results.

Jerry Han also uses EVT techniques in his paper, "Living to 100 and Beyond: An Extreme Value Study." Notwithstanding my previous comments about the virtues of birth cohort data, his choice to use a full century of U.S. Life Tables provides lots of data points for examining trends. Using the top 15 ages (instead of the top 15 observations) from each year also means there are lots of "extreme" observations in each year in his data set. Using EVT techniques, he calculates the maximum attainable age implied by each year's data, then studies the trend in that age over time. Not surprisingly, the implied maximum attainable age in any year appears to be strongly influenced by the actual maximum age attained in that year.

His paper shows that, while the estimated limiting age does not increase monotonically, its upward trend is clear and relentless. He attempts to quantify this upward trend by simple linear regression, which is useful, but visually it appears that a straight line is not a particularly good fit for the pattern of the increase.

Estimates of ω

We can understand these papers better by considering how the models and calculations in them address three big questions: First, is there a maximum attainable age, ω? If so, what is it? And what can we reasonably conjecture about mortality rates at very high ages?

Let us take a look at the variety of estimates of ω offered by these papers and other sources. By her choice of models, Ms. Mumpar-Victoria asserts that there is no finite limiting age. The traditional Makeham model and the two variations on it presented in her paper all have that assumption embedded in them. In other words, you can use these models to calculate the probability of surviving to age 275. The answer may be an infinitesimal probability, but it won't be exactly zero.

The EVT techniques used by Ms. Robertson may or may not yield an estimate of a finite maximum attainable age, depending on the fitted value of the shape parameter. The data she used consistently yielded a negative value for ξ, implying a finite value for ω, except for the data for Japanese men with the outlier included. Using a parametric bootstrap on the Japanese data, she estimated ω to be 124.21 for females and 118.57 for males (when the outlier was replaced by a more reasonable value). However, the 95 percent confidence interval that she calculated using this parametric bootstrap had no upper bound, so the notion that there is no finite value for ω remains plausibly consistent with her data. Furthermore, she states in her paper that she intuitively prefers to believe that there is no fixed finite value for ω.

The largest value I could find for ω in the graphs presented by Mr. Han was about 116, but his analysis clearly showed that the estimates of ω were increasing over time. His regression lines would eventually extend to some very large values. If you were hoping for solid evidence of a limit on longevity risk, there is not much comfort to be found in these papers.

I'd also like to mention two other estimates of ω from noteworthy sources. The Torah or Bible pegs the value of ω at 120 (Gen 6:3). Jeanne Calment of France asserted that ω is 122, and made that assertion by actually living to 122 but not to 123. I guess there are a lot of ways to offer a dissenting opinion.(SIDE NOTE TO EDITOR: Jeanne Calment was not a researcher; she is simply the person who died at the oldest documented age. The fact that Mme. Calment died at age 122 does not prove that omega = 122; it does prove that omega > 121. My point was that her documented longevity contradicts most of the other estimates for omega.)

Mathematical Models of Mortality

The question of what we can reasonably conjecture about mortality rates at very high ages leads us into the use of mathematical models. Ms. Mumpar-Victoria presents a concise list of the advantages of using a mathematical model to describe the relationship between mortality rates and age across the full range of ages. First, the notion that there "ought to be" a concise relationship between age and mortality rate is intuitively appealing. Second, fewer parameters are needed both to fit a model to a set of raw data and to communicate that model to others. Finally, and most importantly for the purpose of estimating mortality rates at very advanced ages, a mathematical model makes it possible to estimate mortality rates for ages lacking data, based on what data there is at other ages.

The model proposed by Gompertz in 1825, $\mu_x = B \cdot c^x$, has proven to be quite useful and is still in use today. In this model, the natural logarithm of the force of mortality, $\ln(\mu_x)$, is linear in x. In 1860, Makeham proposed a refinement of the Gompertz model: $\mu_x = A + B \cdot c^x$. In Makeham's model, the natural logarithm of the first difference in the force of mortality, $\ln(\mu_{x+1} - \mu_x)$, is linear in x. The models used by the authors of the papers presented in this session have radically different shapes. Both Ms. Robertson and Mr. Han use the Generalized Pareto Distribution to model mortality rates at advanced ages, since EVT demonstrates that any distribution approaches this family of distributions at extreme values. Using this distribution implies that for $x \geq u$, $\mu_x = \frac{1}{\tilde{\sigma} + \xi(x-u)}$, where u is the threshold value beyond which values of x are considered extreme. In this model, $1/\mu_x$ is linear in x. Ms. Mumpar-Victoria, on the other hand, has devised a model called the Inverse-Makeham model, in which $\mu_x = \frac{z}{\sigma \cdot (e^z - 1)} + A$, where $z = e^{-(x-m)/\sigma}$. This model is highly nonlinear, so there is no apparent transformation of μ_x that is linear in x. This makes it very difficult to fit approximately using back-of-the-envelope methods.

Shape of Mortality Curves at High Ages

Let's take a look at the shape of these models for advanced age mortality. In all of the following figures, $\ln(\mu_x)$ is plotted against attained age. Based on the commonly observed pattern where mortality increases sharply by attained age, this scaling is generally an informative way to look at mortality data and models. Gompertz predicted that such a graph would appear as a straight line, so each figure indirectly provides a comparison to the Gompertz model.

Figure 1
Shape of Models

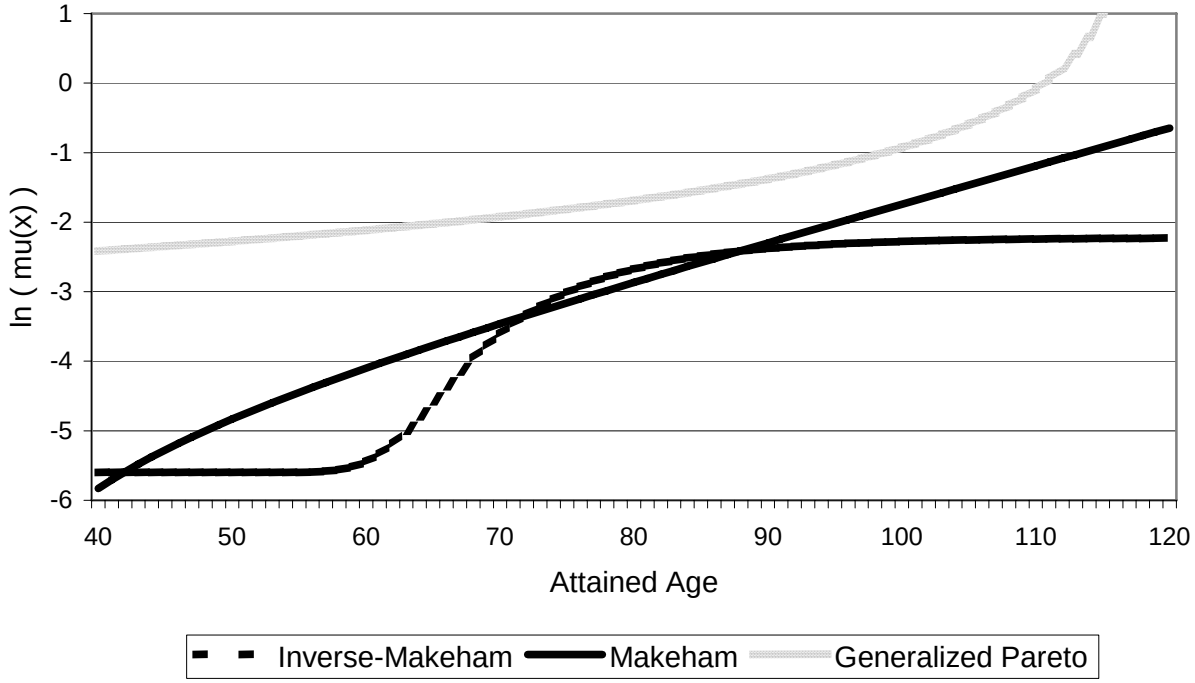

Figure 2
Shape of Models at High Ages

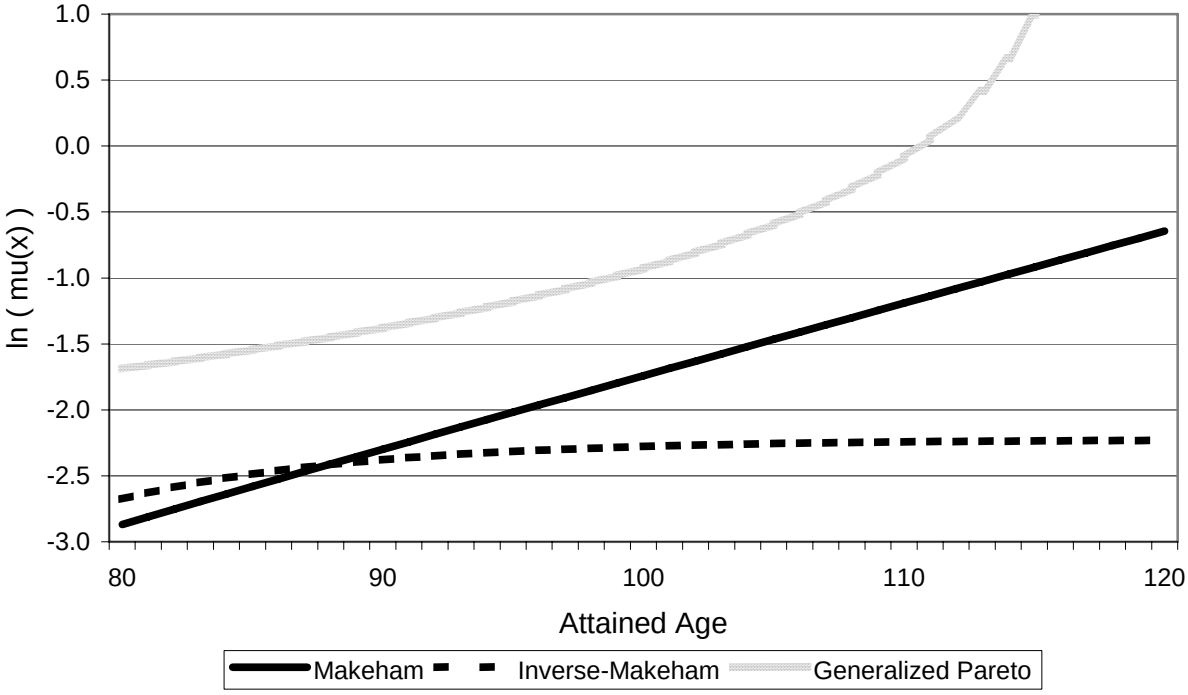

Figure 1 traces three mortality models from age 40 to age 120. The Generalized Pareto curve shown utilizes the parameters calculated by Ms. Robertson as the best fit to the data for females in cohort 1 of the Canadian data with $u = 95$. The extrapolation of this curve to ages under 95 is spurious, since the model specifically applies only to ages above 95, but I thought that the shape would be more evident by extending the curve to the entire range of the graph. Both the Makeham and Inverse-Makeham curves shown utilize the parameters calculated by Ms. Mumpar-Victoria as the best fit to the data from the Philippine Intercompany Mortality Table.

We can see from Figure 1 that the Inverse-Makeham model has a very peculiar S-shape, with a steep increase in mortality between the ages of 60 and 80 and surprising stability outside of that range. We can also see the Generalized Pareto curve taking off toward infinity as it approaches the limiting age, ω. Now let's look at these curves just at the high end of the range, from age 80 to age 120. Figure 2 shows how divergent these models are at the high ages. Incidentally, Figure 2 also is a reasonably good representation of the three types of curves in the Generalized Pareto family. The Inverse-Makeham curve is similar to a Generalized Pareto with $\xi > 0$, the Makeham curve is similar to a Generalized Pareto with $\xi = 0$, and the Generalized Pareto curve shown has $\xi < 0$.

Figure 3
Model Extensions of US 2002 - Male

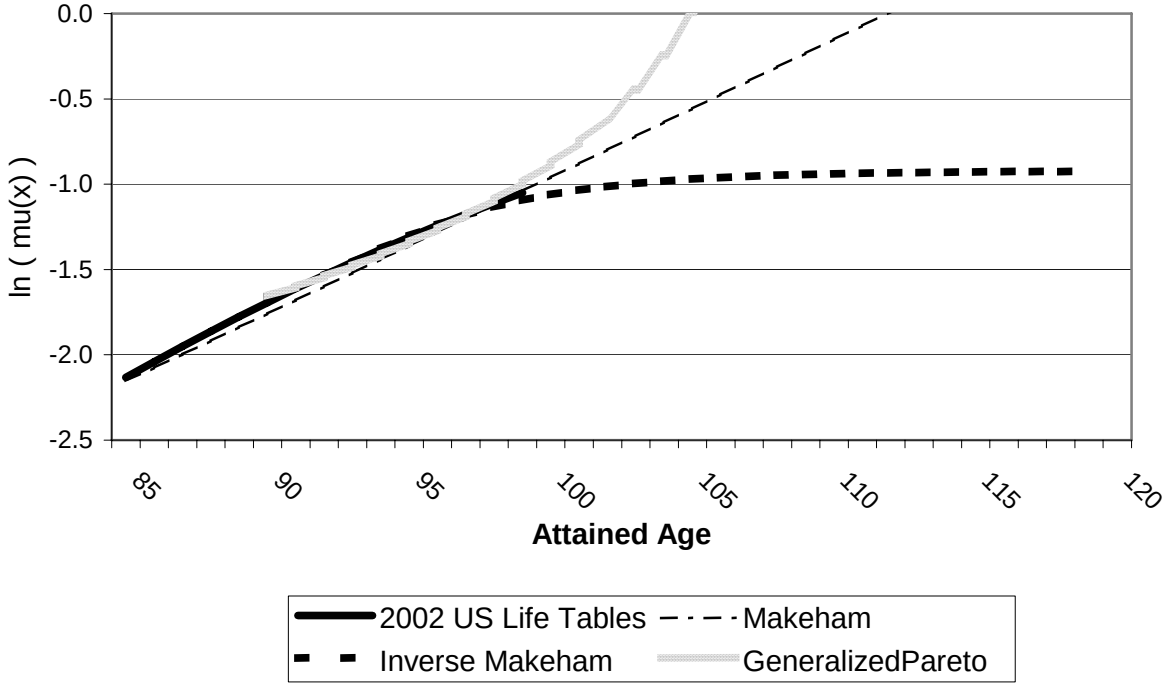

Figure 4
Model Extensions of US 2002 - Female

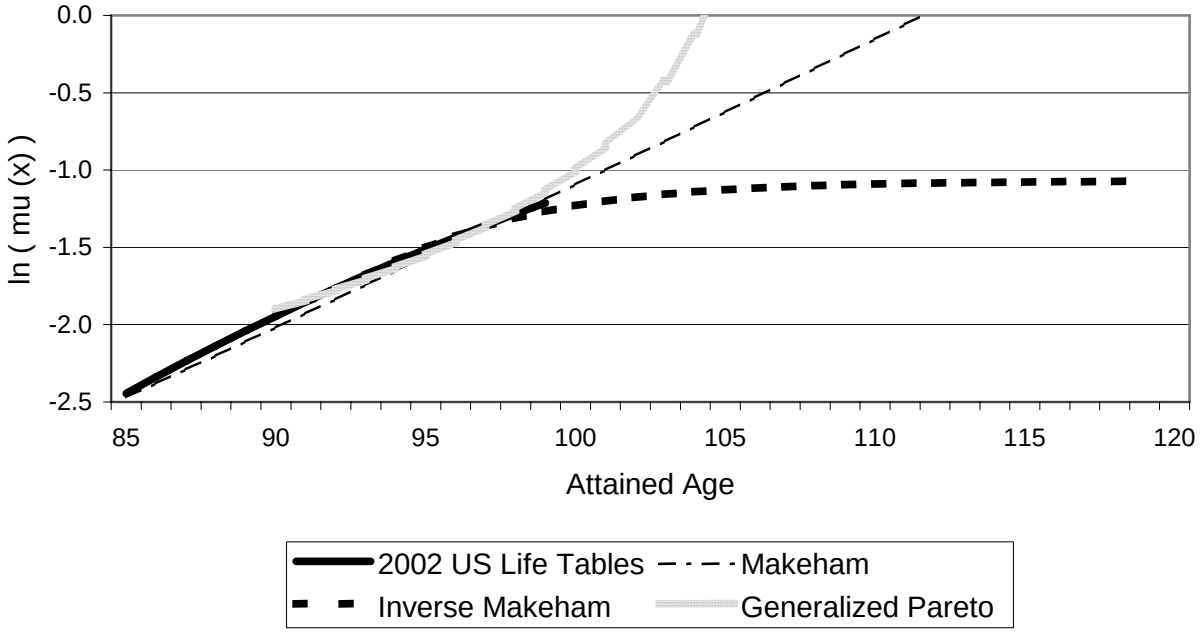

The authors of these three papers used very different data sources for their analysis, so we would expect the graphs of the models fitted in the papers to differ substantially, even if all of the authors were fitting the same model. It is illuminating to compare these models as if they were all used to extend the new 2002 U.S. Life Tables. This is the table prepared by the Division of Vital Statistics of the National Center for Health Statistics. It arbitrarily ends at age 99 without a value for ω. I calculated a regression line on the first difference of the force of mortality between ages 65 and 99 to get parameters for a Makeham curve. I also calculated a regression line on the reciprocal of the force of mortality between ages 90 and 99 to get parameters for a Generalized Pareto curve. The Inverse-Makeham curve was fit heuristically by trial and error to get a curve that looked reasonably close between ages 90 and 99. Figure 3 shows these extensions for males, and Figure 4 shows them for females. These figures show that the divergence of the models is even more dramatic when identical data is used to parameterize the models.

To get a sense of how the advanced age mortality trajectories predicted by these models compare with current mortality tables being used by practicing actuaries, let's take a look at the graphs of a few current tables for ages 40 to 120. Figure 5 shows the 2002 U.S. Life Tables separately for males and females. These tables are firmly based in the data for the U.S. population without extrapolation. The graphs look quite straight with only a couple of minor hiccups, but they do appear to taper off slightly at the very end, possibly suggesting a tail similar to the Inverse-Makeham or Generalized Pareto with $\xi > 0$.

Figure 5
2002 US Life Tables

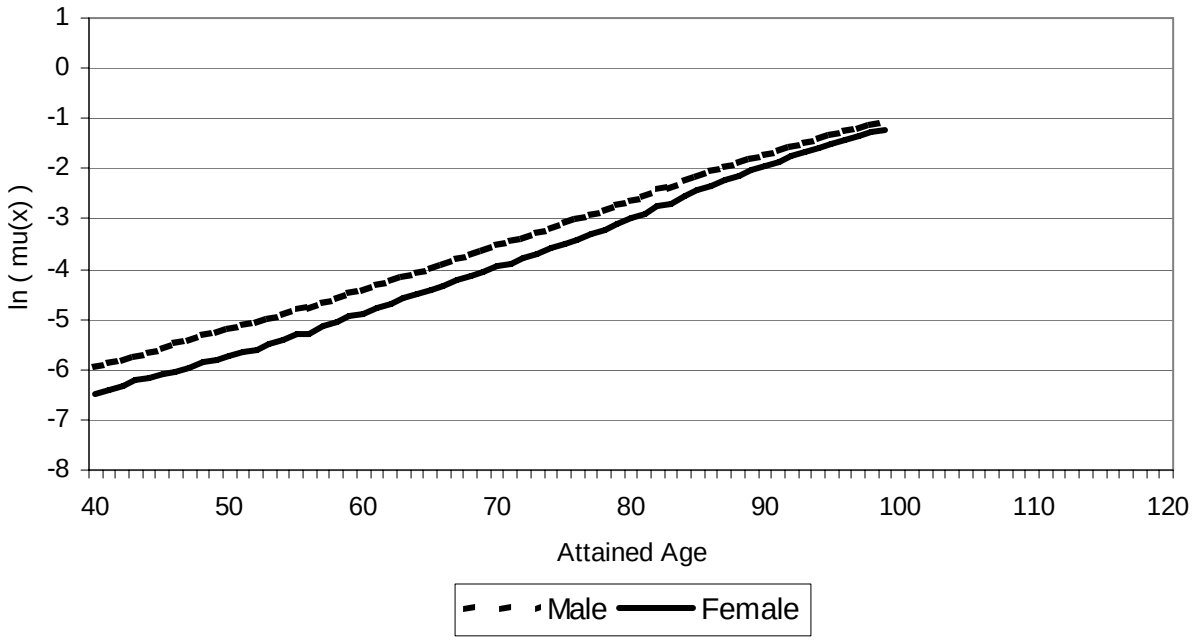

Figure 6
2001 VBT Nonsmoker Ultimate

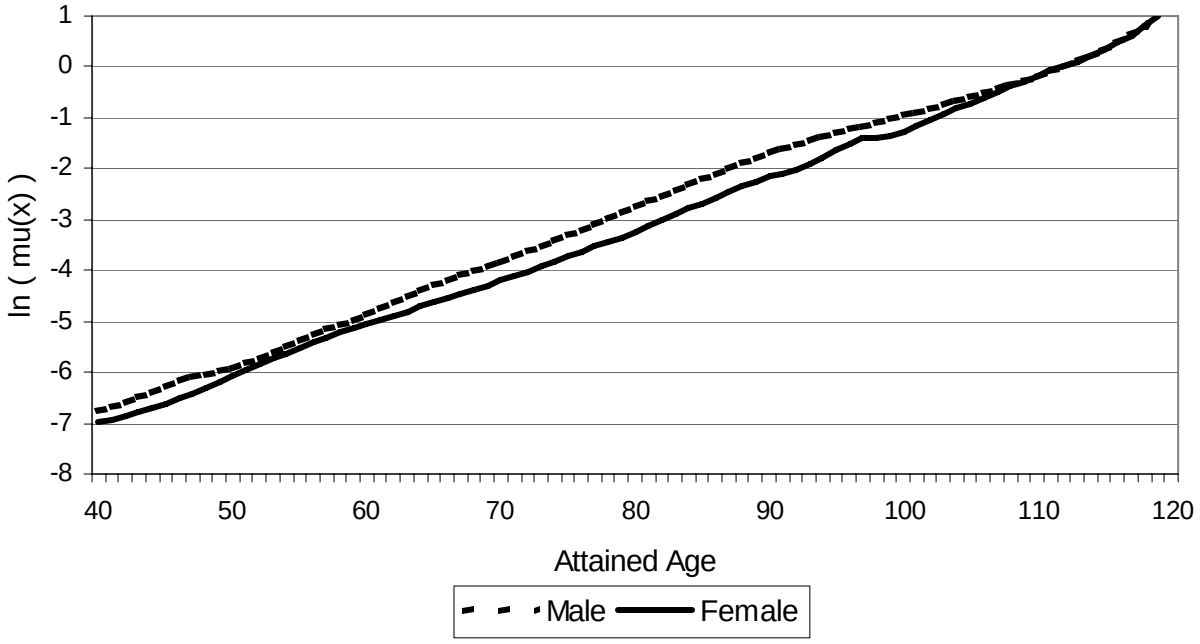

Figure 7
A-2000 Basic

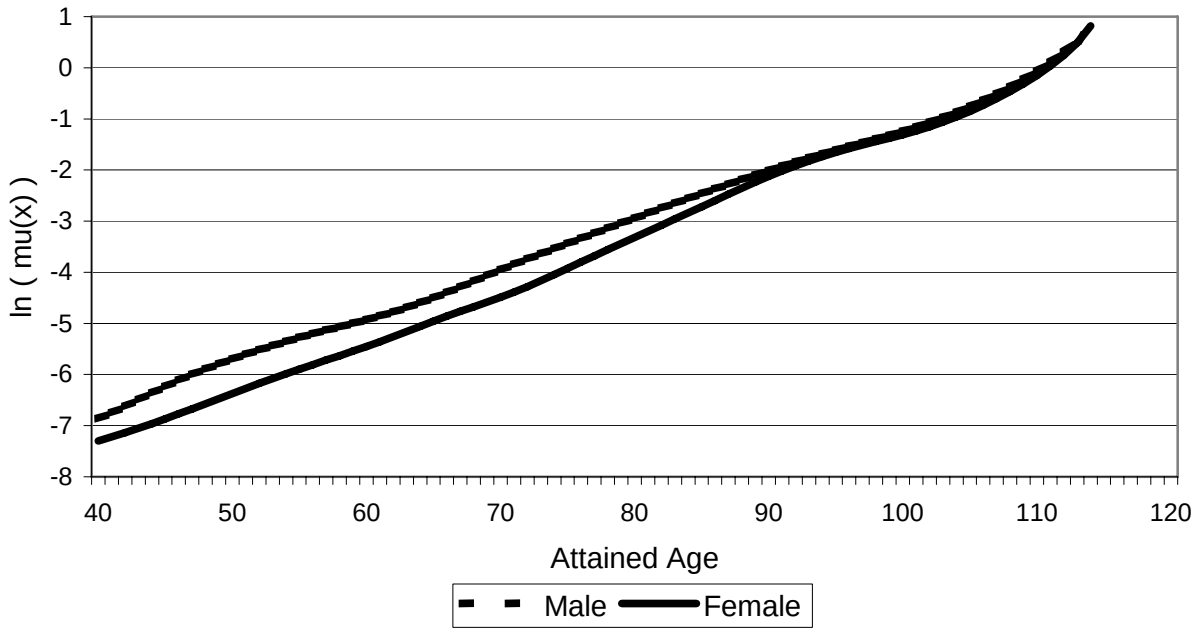

Figure 8
RP-2000 Healthy Annuitant

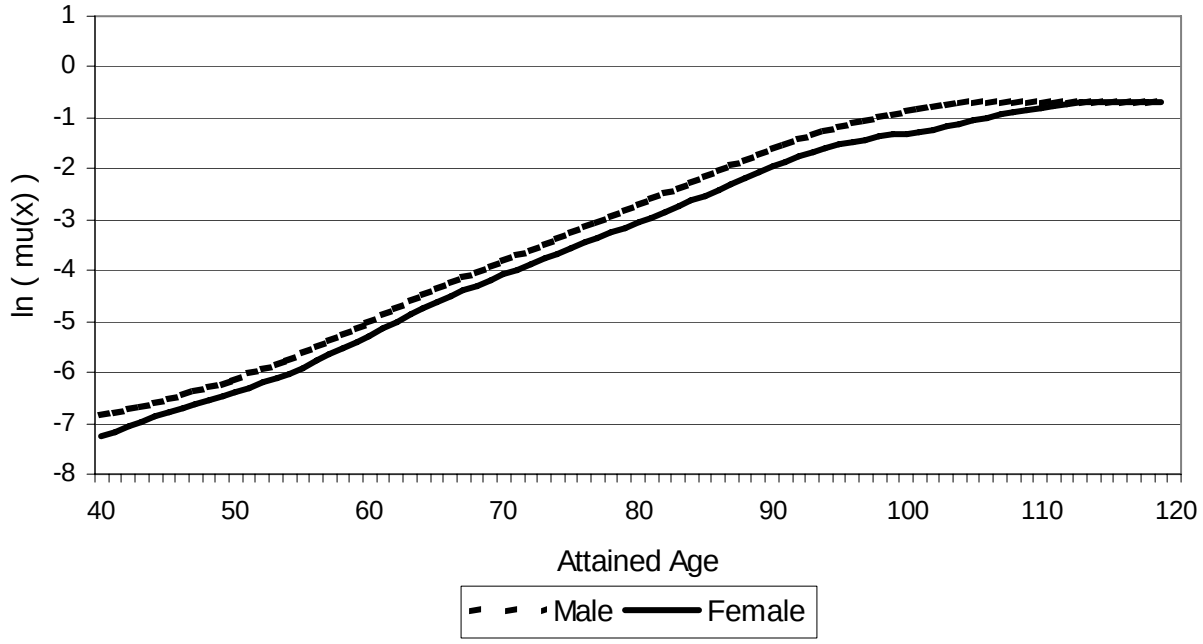

Figure 6 shows the 2001 Valuation Basic Table (VBT) ultimate mortality for male and female nonsmokers. This table is based on recent life insurance industry experience. The current valuation standard for life insurance, the 2001 CSO Table, was created by adding valuation margins to the 2001 VBT. The lines on this graph are still fairly straight, although not quite as straight as the 2002 U.S. Life Tables. These lines appear to bend up slightly at the highest ages.

The current valuation standard for individual annuities is the A-2000 Table, shown in Figure 7. The lines on this graph clearly bend upward at the highest ages, consistent with the Generalized Pareto distribution with $\xi < 0$. On the other hand, Figure 8 traces the mortality pattern for healthy annuitants from the RP-2000, the most recent table based on data from large pension plans. The lines on this graph clearly flatten out in an even more pronounced manner than the Inverse-Makeham or the Generalized Pareto with $\xi > 0$.

Of course, the mortality rates at the highest ages for all of these tables currently used by actuaries were established by mathematical extrapolations that the developers considered aesthetically pleasing. Credible data was available for use in developing these tables only up to ages in the mid-90s, so the shapes of the bends in the graphs at the highest ages are more a reflection of the opinions of the developers than of any indication from the underlying data. In order to understand what the indications of the underlying data really are, let's compare these four tables for ages 80 to 100 only. Figure 9 shows these four tables for males, and Figure 10 shows these tables for females. In general, there seems to be a very slight downward bend in these graphs.

Figure 9
Male Mortality Tables

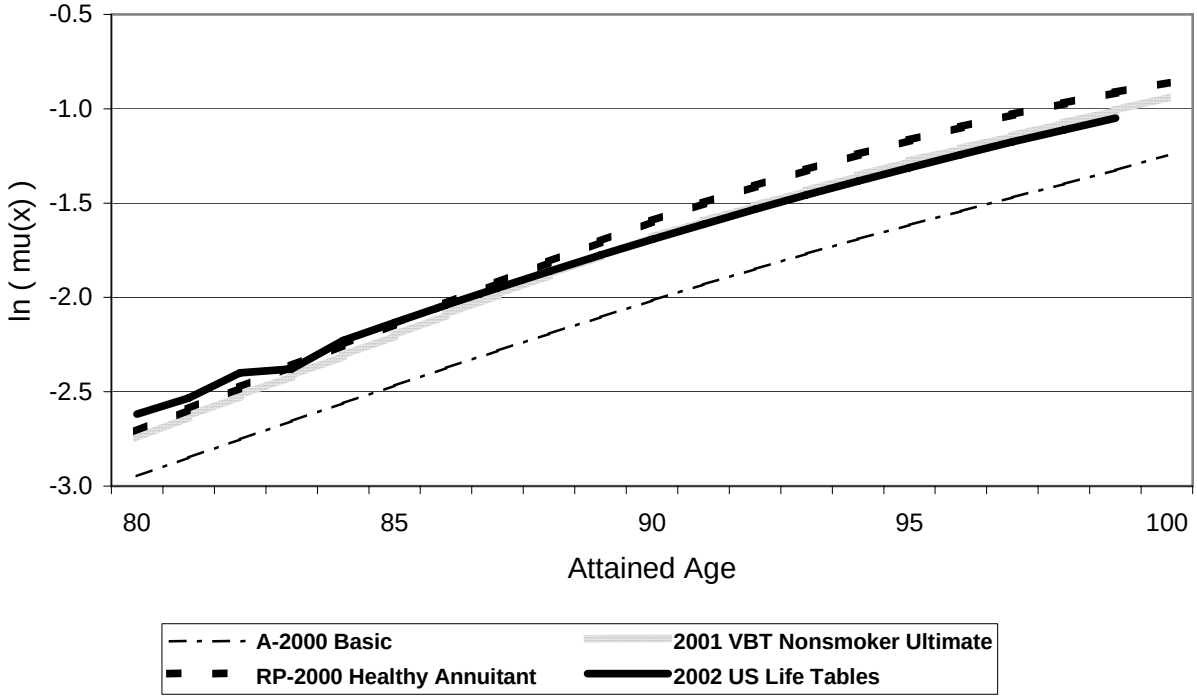

Figure 10
Female Mortality Tables

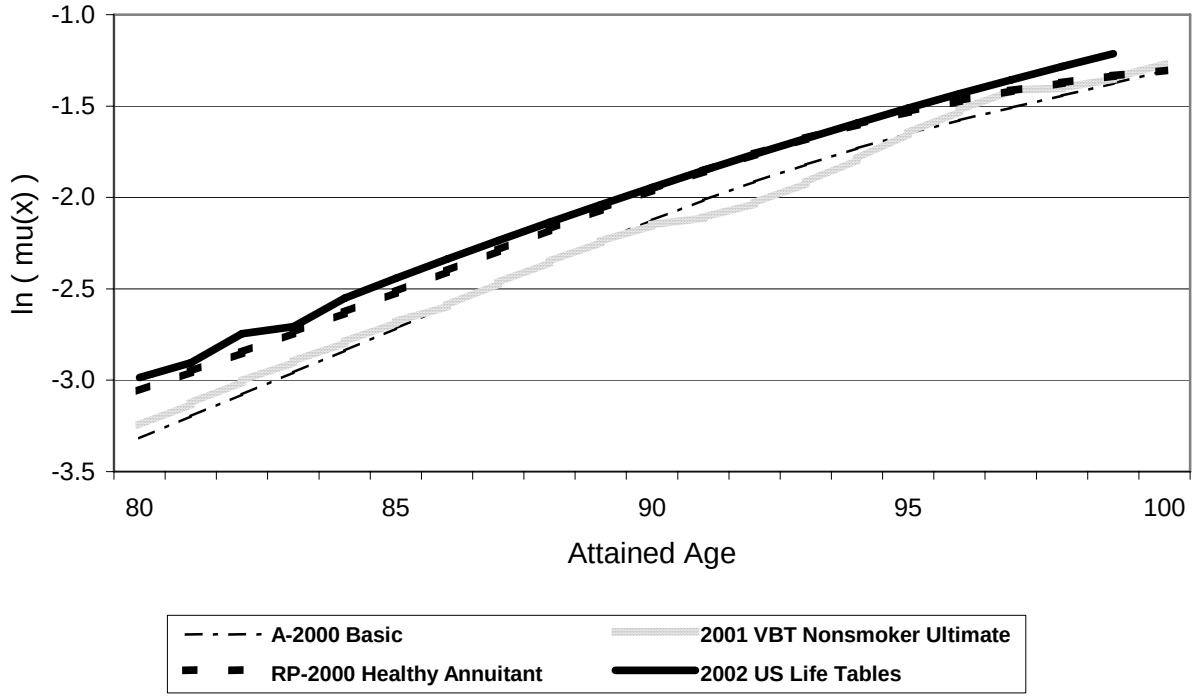

In conclusion, I'd like to commend these authors once more for their excellent papers proposing new mathematical models of mortality.