A Quick Look at Cohort Effects in US Male Mortality

Summary: Leave-one-out cross fitting analysis in MCMC provides a newly efficient way to select the optimal degree of parameter shrinkage. Applying this to mortality trends finds a correspondence between mortality cohorts and birth cohorts, e.g., the baby boomers. In particular it appears that mortality increases and decreases with the birth rate four years earlier. Links of mortality trends to HIV and perhaps medical insurance show up as well.

Basic Findings

The mortality trend model discussed below in more detail combines a base mortality table by age, a mortality trend across years, and a smaller cohort effect by year of birth. The cohort impacts are in Figure 1.

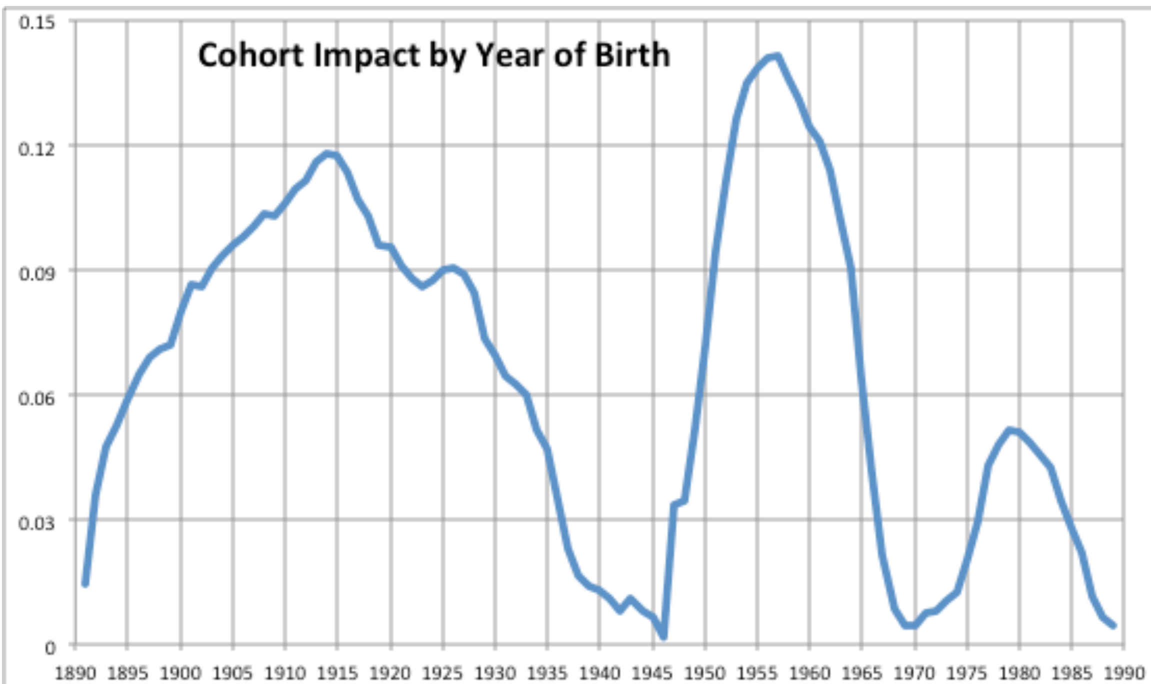

Figure 1: Cohort Impacts: Increase from Base Mortality

Cohort effects in mortality tend to be difficult to explain. Often strings of coincidences are invoked – what age they were when smoking reduced, when heart disease treatments improved, etc. One possibility suggested by this pattern is a relationship to the size of the cohort. Being in a bigger age group, especially when society had to adjust to cope, could itself stress that population, producing higher mortality rates. Also the rate of change in cohort size could have related effects.

Looking at the birth rates supports this, but with a lag – see Figure 2. It may be that having a big population three to five years older could produce more stress – less employment and advancement opportunities, for example, while a smaller slightly older population could produce more opportunities. Here I should note that in the mortality data, year of birth is defined by subtracting age at death from year of death. Due to truncation etc. there is some imprecision in this calculation. For comparison, birth rates were thus averaged over the two years ending with each year shown.

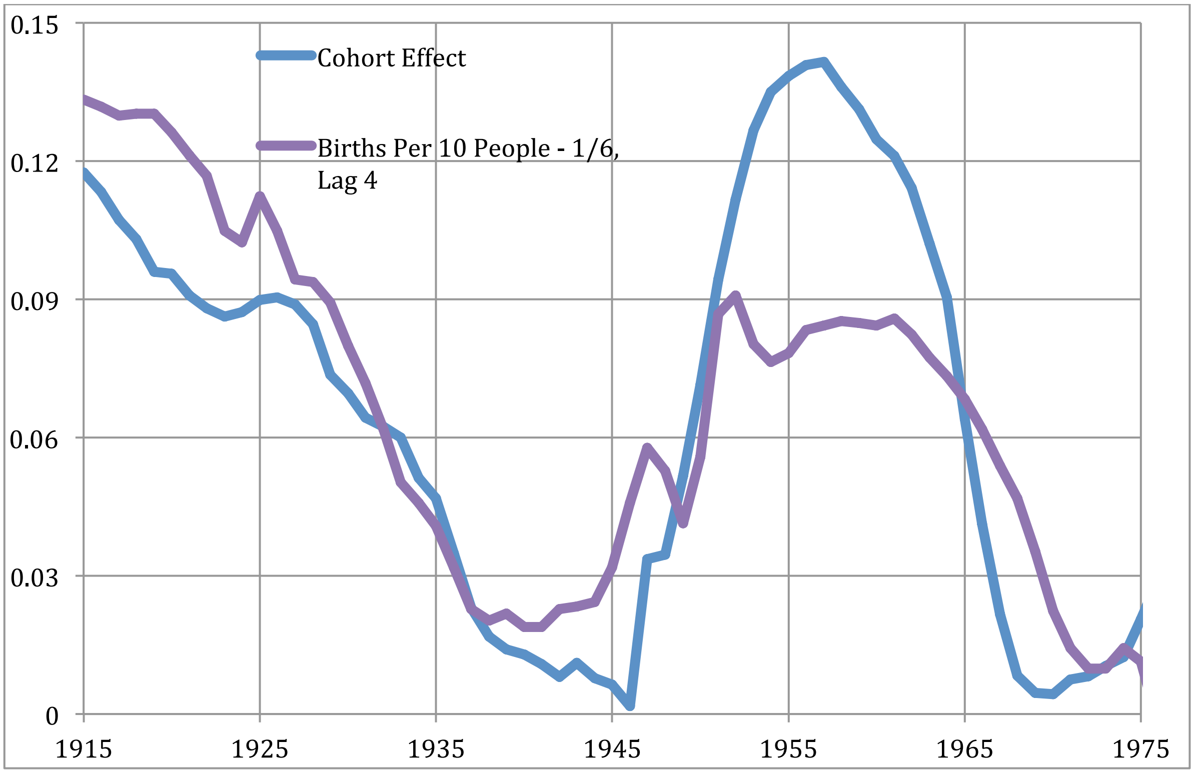

Figure 2: Mortality Cohort Effects and Shifted Birth Rates

The lag may have changed in later groups, and the slope, or change in birth rates, may also have an effect. Another possible driver here is military service. The boomer male mortality is higher than the cohort-size effect would suggest, which could be Vietnam related – both in service and post-service. The blip in Figure 1 at 1980 does not correspond with cohort size but could be related to Mideast service. The data for that group only goes through ages in the twenties, where mortality is usually low, so military deaths could show up as a population effect for these males. Possibly adjusting the earlier birth data for infant mortality would shift the WWII cohort towards a similar effect, but a much longer view of mortality is there for this group. They may also have enjoyed better veterans' benefits than later cohorts.

Other Trends

The model starts off with a base mortality rate, shown on a log scale in Figure 3. This is pretty linear after a bump at age 20. To this is applied a mortality trend over time, as in Figure 4. The trend is pretty steady at 1% to 2% improvement per year except for 1985 – 1995, which is apparently HIV related. The trend immediately catches up in 1996, when HIV deaths slow dramatically. More detail on HIV mortality and the annual trend rates are in Figures 5 and 6. However there is also a reduction in the improvement trend in 2012 –3.

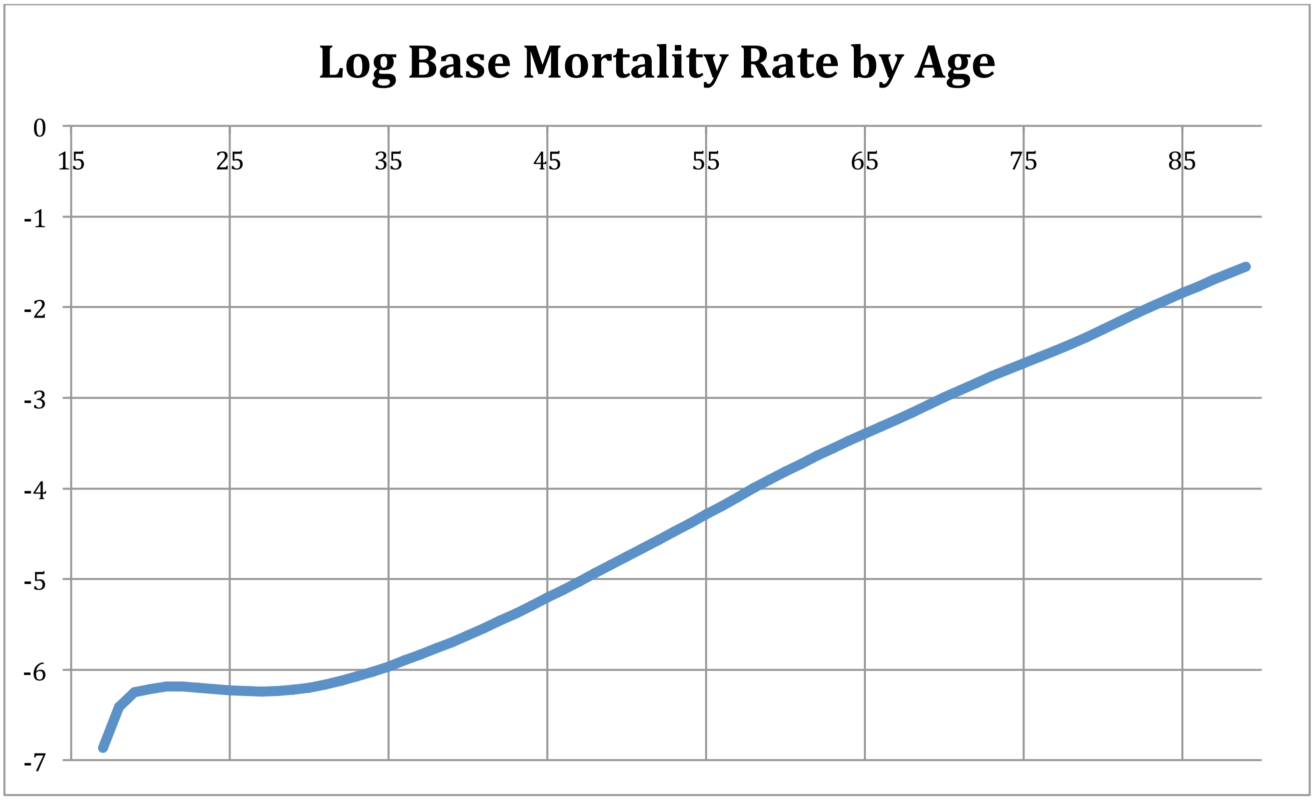

Figure 3: Modeled Base Mortality

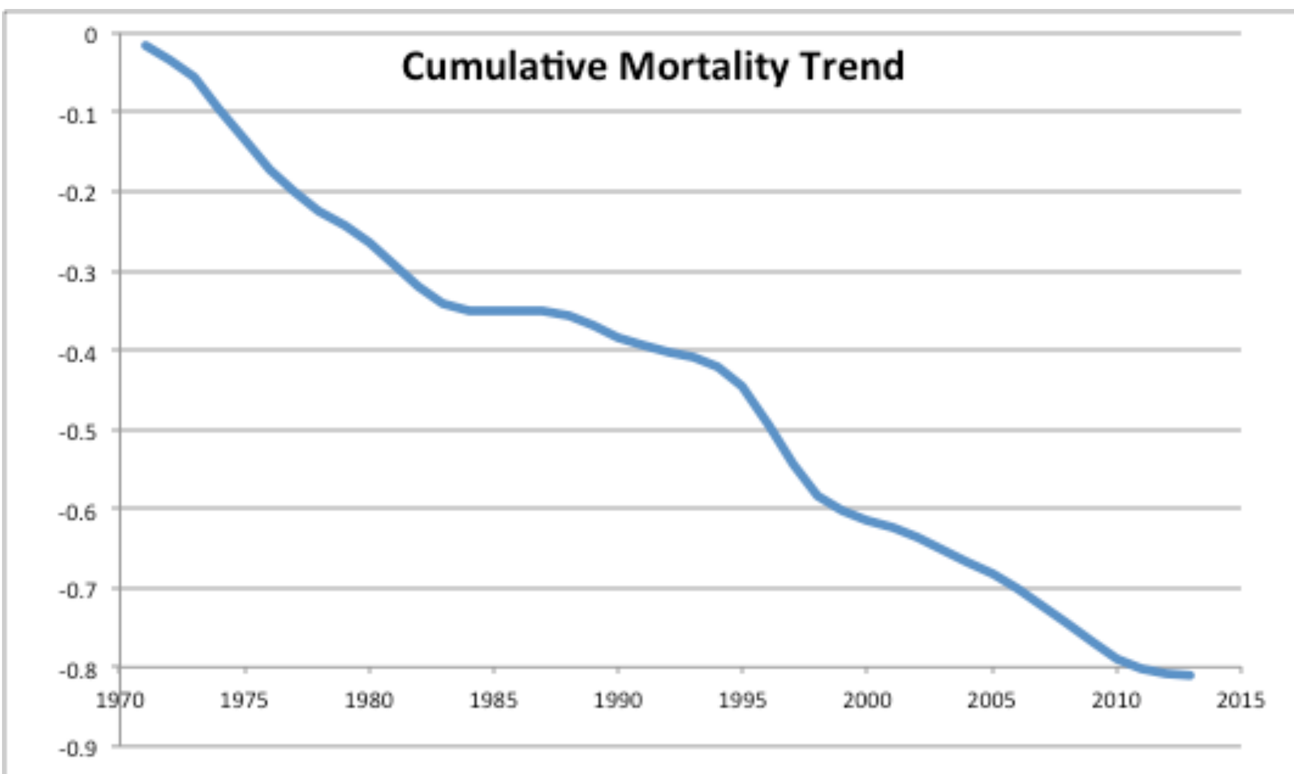

Figure 4: Modeled Mortality Trend

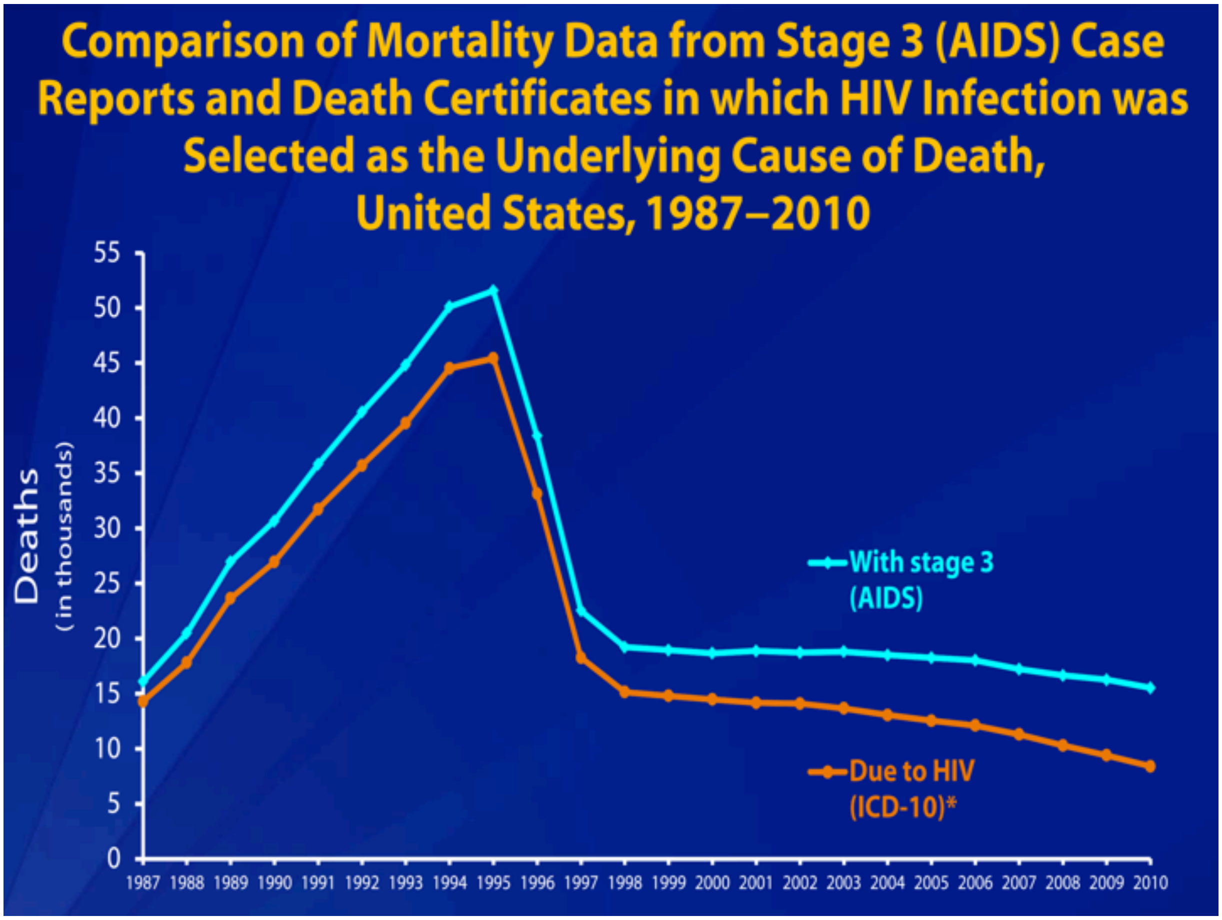

Figure 5: HIV Mortality (from CDC)

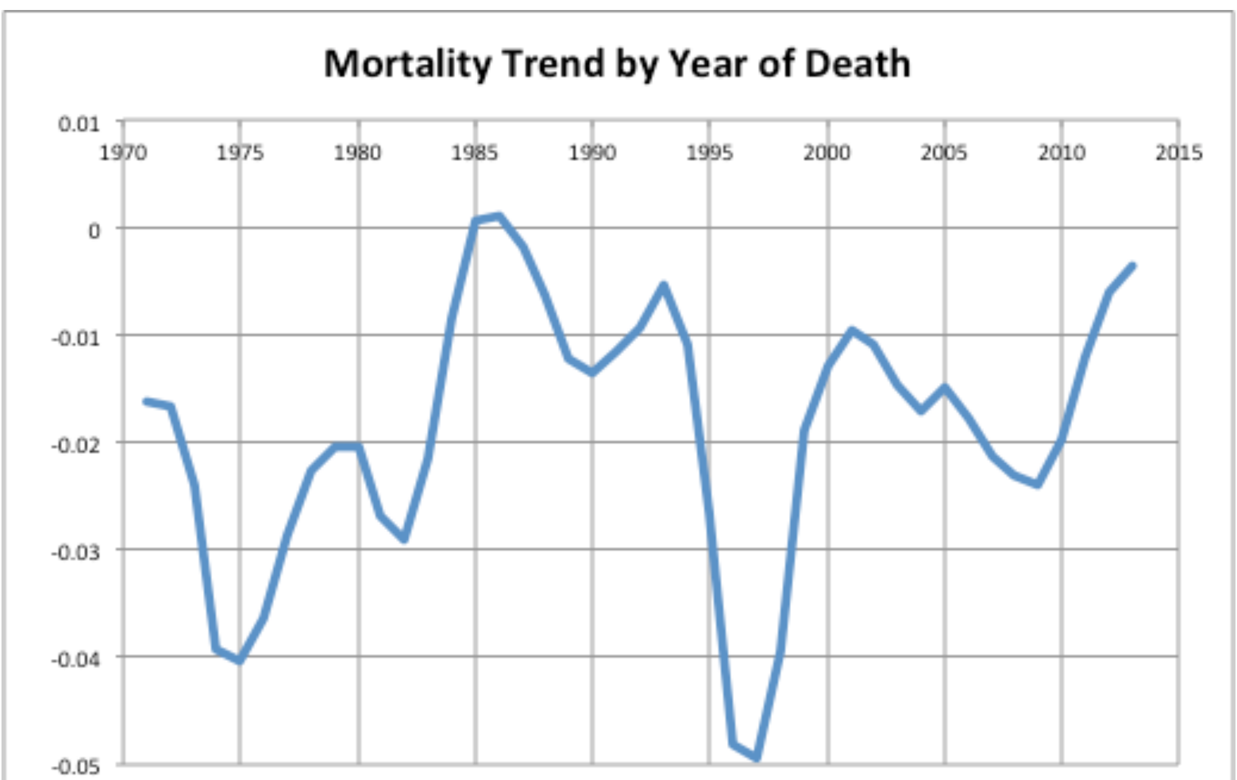

Figure 6: Modeled Mortality Annual Trend Changes

The reduced trend for 2012 – 3 could actually be an artifact of the fitting. The methodology for parameter shrinkage is pretty good at eliminating random fluctuations when there is subsequent data showing that in fact they were not new trends, but this mechanism does not work so well for the last few observations in a data set. However another possible impact here is reduced medical insurance coverage – see Figure 7. This may have a cumulative impact so may take a few years to respond to the latest reduction in uninsureds.

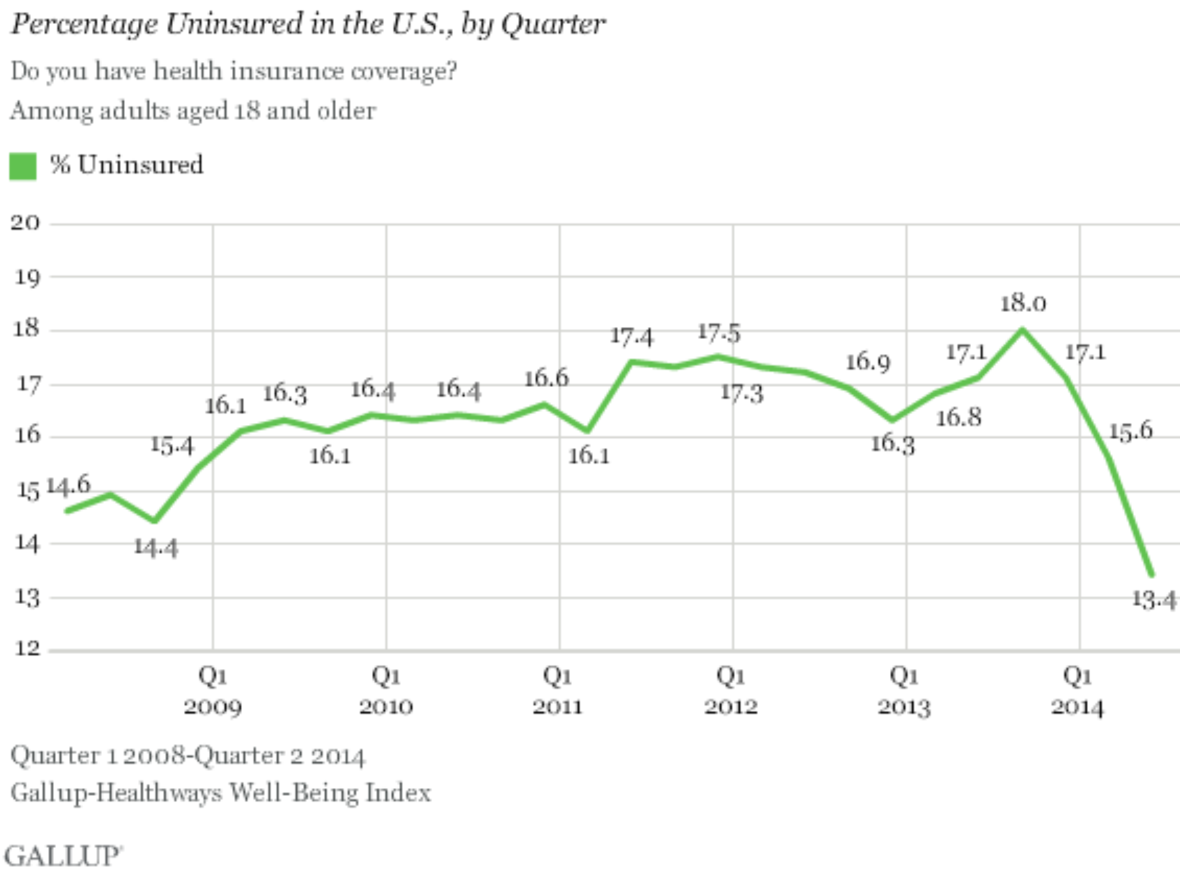

Figure 7: Percentage without Medical Insurance

Finally, the mortality trends appear to have greater impact at older ages, reversing a bit for the very oldest. The factors in Figure 8, normalized to a maximum of unity, are applied to the trends.

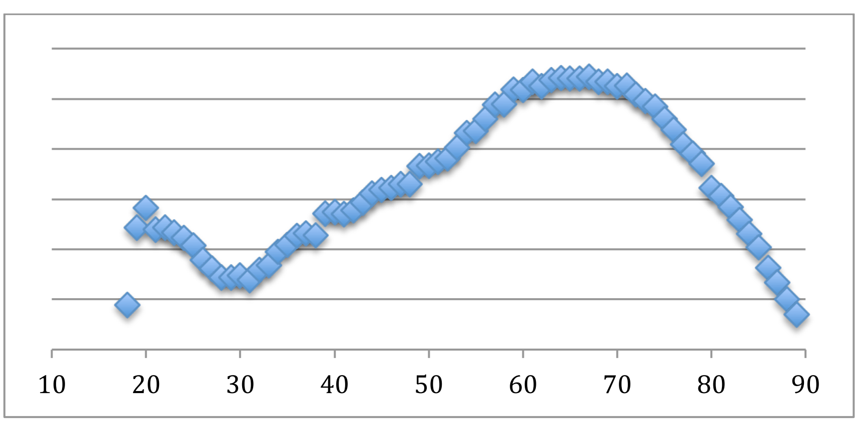

Figure 8: Age Adjustment to Trend

Fitted Values

Although mortality has trended downward over virtually the entire observed period, the cohorts having higher mortality could see observed mortality rates at some ages temporarily increase. The trend being not totally picked up at lower ages contributes to this. Figure 9 combines the individual factors to show the implied mortality rates that would be observed at various years, by age. Ages around 50-55 especially show some temporary reversals in the last decade, but at older ages the trend asserts itself more normally.

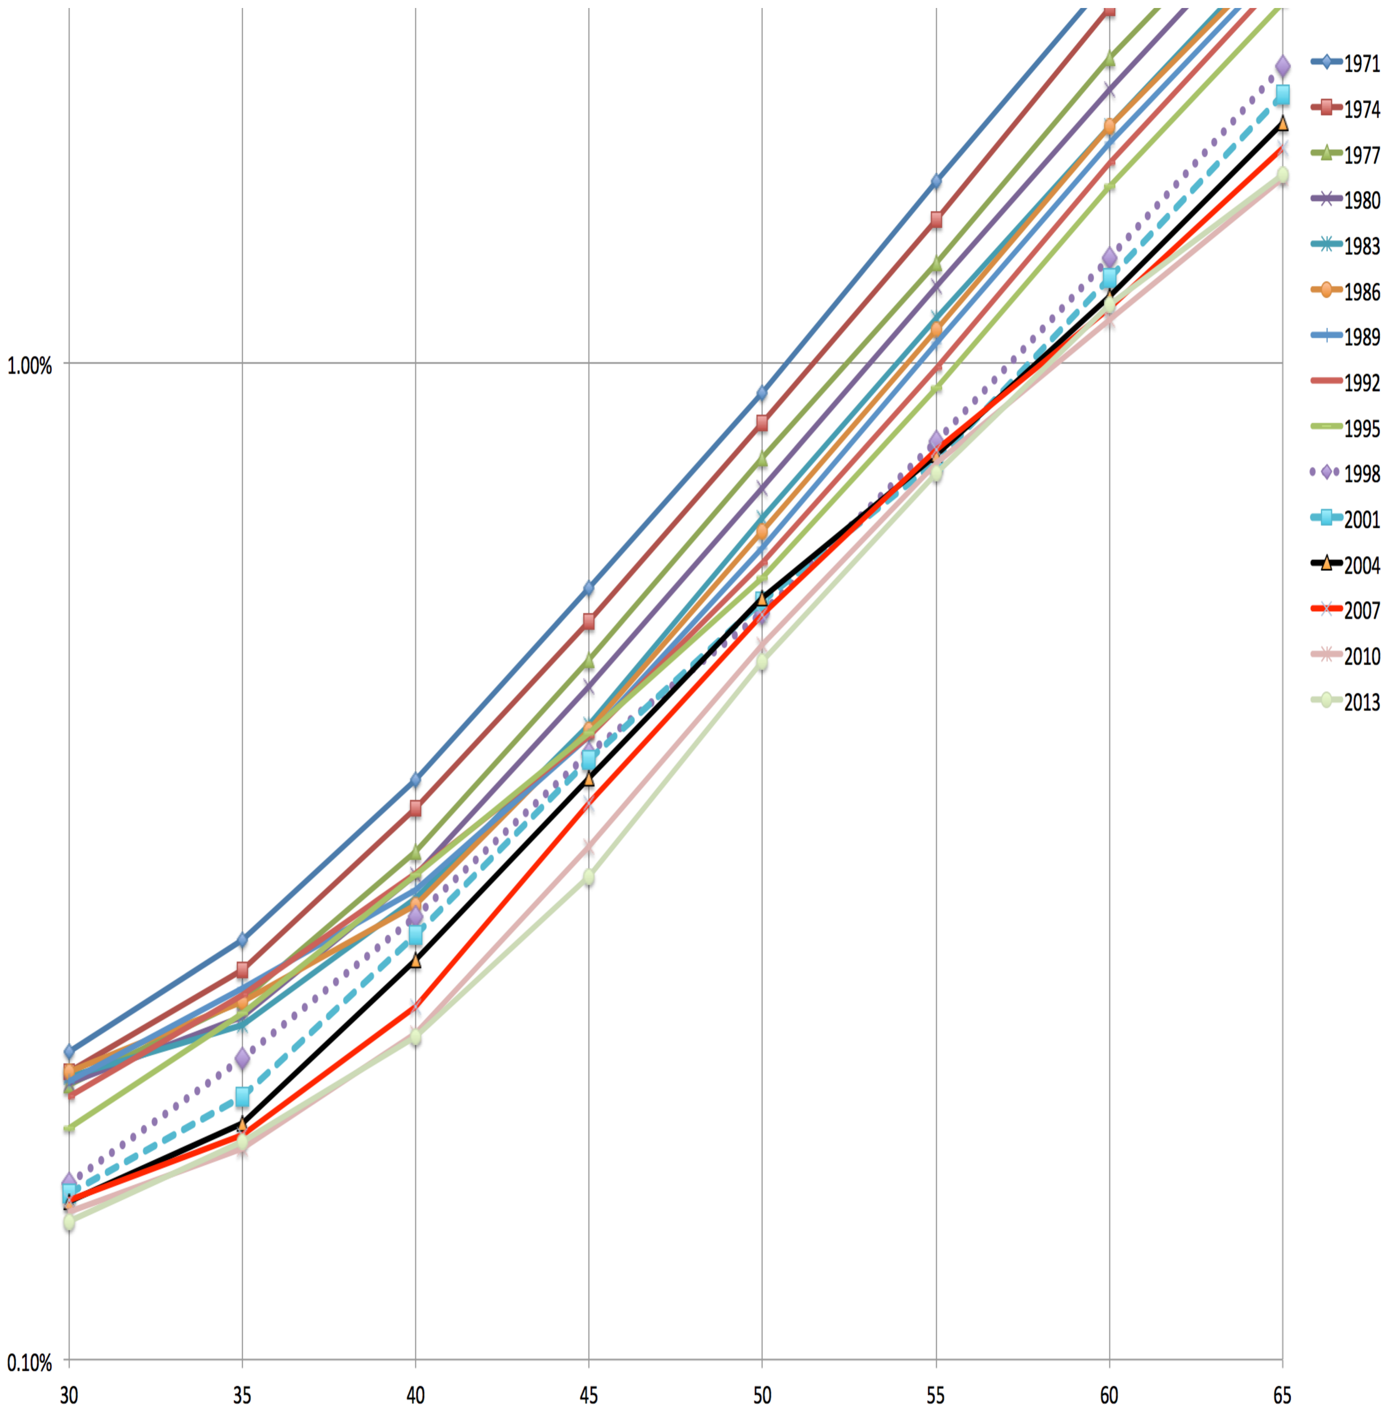

Figure 9: Model Implied Mortality Rates by Age and Year of Death

The Model

For someone born in year n, 1891-1989, the log of the probability of dying at age u, 17-89, so in year n+u, 1971-2013 is assumed to be normally distributed following:

$$y[n,u] \sim normal(p[n] + q[u] + w[n+u]*r[u], sigma\_y)$$

Here p is a vector with 99 elements, one for each year, while q has 73 elements and w and r each have 43. A few initial values are set to zero for specificity, but still there are over 250 parameters. Here p, q, w and r are actually transformed parameters. The original parameters are the second differences of these, or the changes in slope in the graphs, and those are what are shrunk towards zero Lasso-like, then summed up to give the transformed parameters. Specifically, each 2^{nd} difference was given a double exponential prior in (0, 0.1*sigma_y), where the 0.1 came from trying a number of values and using LOO to determine the best in terms of this form of out-of-sample prediction.

The cohort effects are thus the p parameters, base mortality rates are given by q, w is the annual trend, and r (which is normalized to be in [0,1]) expresses how the trend is picked up by age. This model is a generalization of the traditional Lee-Carter model, which does not have cohort effects, and is a special case of the Renshaw-Haberman model, in which the cohort effects have an age interaction like the trend has. Earlier work by Roman Gutkovich, submitted for publication, found that the cohort-age interaction disappeared quickly under classical Lasso, so it was not included here. Other modelers have noted the need for parameter shrinkage, but this has been most typically done by using cubic splines across the transformed parameters rather than by this kink of shrinkage.

Model Issues

The age-trend interaction is probably the most problematic part of the model. Most of the medical advances have directly reduced mortality at more advanced ages, but this is not always the case. In particular, HIV affected younger people more strongly, as did its treatment. This fact probably shifted the r parameters somewhat towards lower ages, but that makes them less accurate for other periods without really picking up the mid-1990s very well. In fact the strong divergence between 1995 and 1998 fitted mortality at older ages seen in Figure 9 is probably misstated. Some dummy variables may be called for to pick up the HIV effect.

Another problem with that interaction is its estimation. Traditionally an iterative process is used to estimate w with r assumed fixed, then these roles are reversed, etc. until they both converge. This has the advantage of making the w values reasonable trends. Sometimes (as in Venter, G. G. 2011 "Mortality Trend Risk" http://www.casact.org/pubs/forum/11wforumpt2/Venter.pdf) MLE is used to estimate r and w simultaneously. But even though there are numerous local maxima, most MLE routines will converge to one or another of these, which can be directed somewhat to still produce reasonable trends. With MCMC, however, there appear to be many local maxima where the trends and percentage impacts do not look like they are meant to. This seems to forestall convergence, and the parameter means can come out looking like huge changes over ages with only fluctuations over time in this model with two age effects, namely q and r. Here we simply took r values from a previous classical Lasso model. Probably some greater clarity in model specification can constrain the trends to look like they are meant to, but more research would be needed for this.

Summary

The modeling so far has focused on finding the degree of parameter shrinkage to use. The resulting cohort parameters informally seem to relate to demographic trends in birth rates and possibly to war and veterans issues. The trends in mortality appear fairly steady except for an HIV effect and a slowdown in the last two years, possibly a random fluctuation or possibly a result of declining medical insurance. However formal regression for the parameter movements have yet to be carried out. Other forms of parameter shrinkage could also be explored, like other priors, for instance.