Appendix: A Simple Model of Uncertain Mortality

This model provides a simple framework for looking at the uncertainty in mortality rates around a given path for expected mortality. Taking the projected mortality rates, $m*_t$ for the cohort which was sixty-five in 2012, assume that actual mortality rates, m_t, follow the following random process.

$$logm_t = logm_t^* + u_t, \text{ if } logm_t^* + u_t < \log(0.8);\ t>65$$
$$\log m_t = \log(0.8) \text{ if } logm_t^* + u_t \geq \log(0.8)$$
$$\log m_t = \log(0.8) \text{ if } \log m_{t-1} = \log(0.8)$$
$$u_t = u_{t-1} + \theta_t + \vartheta_{t-1} + \varepsilon_t$$
$$\vartheta_t = \vartheta_{t-1} + \omega_t$$

Here $\varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right)$ and $\vartheta_t \sim N\left(-\frac{\sigma_\vartheta^2}{2}, \sigma_\vartheta^2\right)$

These expressions set out mortality rates as a random process defined with reference to an exogenous path for mortality rates, m_t^*, which I take to be the cohort mortality rate projected by the Office for National Statistics. The two random terms are assumed to be drawn from distributions with means below zero to correct for the fact that if the logarithm of a variable is normally distributed, the mean value of that variable is greater than zero.

This is not, however, enough to ensure that, for each period $E(m_t) = m_t^*$ for two reasons. First, there is an element of drift in u_t which is not fully addressed by the offset in the distributions of ε_t and ϑ_t. Note that

$$u_t = \sum_{i=65}^{t} \varepsilon_i + \sum_{i=65}^{t} (t-i)\vartheta_i$$
$$E(u_t) = E(u_{t-1}) + \theta_t - \frac{\sigma_\varepsilon^2}{2} - (t-65)\frac{\sigma_\vartheta^2}{2}$$

$Var(u_t) = (t-64)\sigma_\varepsilon^2 + \sigma_\vartheta^2 \sum_{i=1}^{t}(t-i)^2$ and

$$Var(u_t) = Var(u_{t-1}) + \sigma_\varepsilon^2 + (t-65)^2\sigma_\vartheta^2$$

For $E(e^{u_t}) = 1$ we require, since u_t is normally distributed, $E(u_t)=-Var(u_t)/2$ giving $\theta_t = \frac{\sigma_\vartheta^2\{(t-65)-(t-65)^2\}}{2}$

Secondly, we have imposed maximum mortality rate of 0.8. There is substantial debate about whether there is a maximum mortality rate of less than one; since very few very old people are observed, it is not possible to explore this satisfactorily statistically, but many life tables are constructed on the assumption of a maximum mortality rate lower than this[1]. A maximum mortality rate means that, even if the underlying distribution were symmetric around m_t^* the expected value generated by the censored distribution would differ from this.

The approach we have adopted here offers a means of representing shocks to mortality rates for a particular cohort with the assumption that there is powerful drift in them. In some sense this represents a structure which is very unfavourable for

[1] If the mortality rate is 0.5 the fewer than one in a thousand of those reaching their hundredth birthday would live to a hundred and ten.

annuity contracts. Not only are changes in the level of mortality assumed to persist in the cohort being studied, but changes in the trend rate of change of mortality with age are assumed to persist. Thus, if log mortality rates are rising by one per cent a year less than had been forecast, that change is assumed to persist.