A value-at-risk framework for longevity trend risk

Opening comments from Stephen Richards to Faculty of Actuaries sessional meeting

19th November 2012

Longevity trend risk is most commonly found in annuity portfolios and defined-benefit pension schemes. The risk is that mortality rates improve faster than expected, thus causing the insurer or pension scheme to pay pensions for longer than anticipated and making a loss. The paper before you tonight is about answering two related questions posed by the holders of this risk:

1. How might expectations of future mortality improvements change over a single year?
2. What financial impact could these changes have?

The above two questions might be the definition of a one-year, value-at-risk assessment of longevity trend risk. However, this particular risk unfolds over a long period of time, specifically the lifetime of the annuitants. A one-year, value-at-risk approach is therefore not a natural way to view this risk. Despite this, regulatory requirements in the shape of ICA in the UK and Solvency II in the EU are clear in pushing insurers to consider all their risks through a one-year prism, not least because many other risks ***are*** sensibly viewed in this way. This paper therefore grew out of the regulatory need to put longevity trend risk into a one-year, value-at-risk framework, specifically to estimate the capital required to cover at least 99.5% of scenarios arising in one year.

We will start first with why we believe we need stochastic projection models for this task. We are not alone in this view. Writing about the solvency capital requirement, Börger (2010) stated:

SLIDE 3

> "the computation of the SCR for longevity risk via the VaR approach obviously requires stochastic modelling of mortality"

On the same subject, Plat (2011) wrote:

SLIDE 4

> "naturally this requires stochastic mortality rates"

It is worth illustrating why these authors — and the authors of tonight's paper — believe this. If we take the mortality-experience data from 1961 to 1992 we can fit a Lee-Carter model, which would have been the state-of-the-art in mortality projections in 1992. We can then plot the central Lee-Carter projection, together with the equivalent contemporaneous CMI projection.

SLIDE 6

There is a substantial agreement between the two methods, at least at age 70 shown in the graph. The CMI projection is also on the more prudent side, as one would expect of an actuarial projection! We can now add in the actual mortality rates experienced since 1992 to see how these two projections fared.

SLIDE 7

We can see that both projections were wrong: at the point they were made, nobody was properly aware of the emerging "cohort effect", as introduced to the actuarial profession by Willets (1999). However, although both models were wrong, the statistical model still has value in the form of confidence intervals, which we can also add to the graph:

SLIDE 8

We can see that while improvements were faster than the Lee-Carter model's central assumption, the confidence intervals would have correctly alerted actuaries to the possibility of improvements being faster than this. This, then, is why we believe stochastic models are useful for this kind of work. They alert us to the possibilities which we think unlikely, and the purpose of insurance reserves is to guard against the unlikely.

Longevity trend risk, in either an annuity portfolio or a defined-benefit pension scheme, is a long-term risk. An adverse trend will unfold over a number of years, and so a natural approach would be to use a stressed-trend in reserve valuation.

SLIDE 10

An important point to note is that different models will produce different results, as shown in Figure 2 in the paper.

SLIDE 11

It is easy to see that model risk exists, but it is hard to quantify. The allowance for model risk is therefore an area requiring actuarial judgment.

However, Figures 1 and 2 in the paper are the result of a stressed-trend approach to capital requirements, whereas modern regulation of insurance companies is more orientated around a value-at-risk approach. By this we mean that reserves should be adequate to cover all events occurring over the coming year, barring those with a probability lower than one in 200. The question is then how to take a long-term risk like an adverse longevity trend and put it into a one-year view?

The framework described in the paper does this in a relatively simple way:

- First, we use a stochastic mortality model to simulate mortality rates for the coming year.
- Second, we use those rates to simulate the mortality experience of a chosen population.
- Third, we use the new mortality-experience “data” to refit the mortality projection model.
- Fourth, we calculate a financial measure using the updated projection, such as an annuity factor or a portfolio valuation.

SLIDE 13

Figure 5 in the paper illustrates how central projections under a Lee-Carter (1992) model vary when fitted and refitted with one year’s new data. The different projections are shown, and specimen annuity factors are calculated using these updated projections. The process is repeated many times to obtain a set of annuity factors. This set can then be used to estimate the value-at-risk, for example estimating the 99.5^{th} percentile and thus the minimum capital requirement for longevity trend risk under Solvency II. This is shown in Table 5, which again shows the importance of model risk. Figure 6 shows how the VaR capital requirement varies by age.

SLIDE 15

An important point to remember is that capital requirements in a VaR analysis are quantiles, i.e. they are order statistics. As such, they have uncertainty over their value, and this uncertainty can be quantified — Figure 6 in the paper shows the 95% confidence envelope for the capital requirement for the Lee-Carter model.

SLIDE 16

An interesting question is how many simulations are required, and we can see that increasing the number of simulations from 1,000 to 10,000 materially reduces the uncertainty, although the central estimate is not hugely changed. However, while reducing the uncertainty over the tail estimate is clearly desirable, it is important not to get distracted with an ever-increasing number of simulations. The first reason is that models are only approximations, and this is particularly true when estimating the tails. The second reason is that model risk is material, and different models needs to be run. We view it as more important to run VaRs for five different models with 1,000 simulations each than it is to run a single VaR with 10,000 simulations.

Another practical point is that VaR calculations are dependent on the discount function or yield curve used. VaR calculations will have to be run as frequently as there are major shifts in the yield curve.

SLIDE 18

We have developed the method in the paper as an answer to the problem of putting longevity trend risk into a one-year view. However, the capital requirements produced will only be minimum floor values for two main reasons. The first reason is that we think a one-year, value-at-risk view will under-estimate the true capital requirement for longevity trend risk. This is because most models will not immediately respond to the beginning of a new longevity trend. This is a desirable property of a model: we do not want it to immediately and fully respond to what might be a mere random fluctuation. Most fitted models are therefore more heavily influenced by the existing data than the newly simulated experience, and so will only respond partially to new data. This is why the VaR capital requirements in Table 5 are usually lower than their stressed-trend equivalents in Table 3.

A second reason is that a data-driven approach is unable to incorporate external events which have no precedent in the data. An example of an external event is the recently announced revision to recent population estimates arising from the 2011 Census, and Willets (2012) points out that mortality rates above age 90 are disproportionately affected by this. There have been several retrospective revisions to population estimates from the ONS in recent years, and this is only one example of the sort of thing which can change stated mortality rates and yet is not captured by a statistical mortality model.

In practice therefore we expect that both practitioners and regulators will view one-year VaR capital requirements as a floor for longevity trend risk. Nevertheless, we feel that the framework in the paper is still a very useful tool — being able to set a minimum value for the capital requirement for longevity trend risk is itself a step forward. In addition, the framework offers two further practical benefits. The first is that the framework allows users to explore the oft-overlooked subject of model risk; it can be used with any stochastic model capable of generating sample paths, and any projection model capable of being fitted to data. When projecting mortality rates, or calculating capital requirements for longevity trend risk, it is essential never to rely on any one single model. The advantage of the framework in this paper is that it leaves the user free to specify which models to investigate.

The other use of this framework is to test models for robustness. For practical reasons, a life office will often want to focus on one or two models for everyday use. But how does a life office ensure that a model selected today won't cause problems tomorrow? One answer is to use the framework in this paper to test a model's reaction to new data, and to ensure the resulting projections are stable. As the example in Section 10 shows, not every model is robust to the addition of new data, and it is useful to know this well in advance of investing resources in that model. Cairns et al (2009) described a series of tests for selecting a mortality projection model, to which we would suggest the addition of this framework as a test for model robustness.

While the framework itself is complete, refinements are still possible. For example, the method in the paper uses specimen annuity factors, when it could be usefully extended to perform an entire portfolio valuation instead. This would allow for the distribution of liabilities by age and size. Similarly, the liabilities could be calculated allowing for dependent spouses' benefits instead of just single-life benefits. There is also the possibility of including a portfolio's own mortality experience. We think that the framework outlined in this paper is a useful step in gaining a one-year, value-at-risk view of longevity trend risk, but now we are interested in hearing the audience's opinions!

References

BOOTH, H. & TICKLE, L. (2008) Mortality modelling and forecasting: a review of methods, *Annals of Actuarial Science*, Volume 3, Parts 1 and 2, page 8.

BÖRGER, M. (2010) Deterministic shock vs. stochastic value-at-risk: An analysis of the Solvency II standard model approach to longevity risk, *Blätter DGVFM*, **31**, 225–259.

BOX, G. E. P. & DRAPER, N. R. (1987) Empirical Model-Building and Response Surfaces, *Wiley*, ISBN 0-471-81033-9.

CAIRNS, A. J. G., BLAKE, D. & DOWD, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration, *Journal of Risk and Insurance*, **73**, 687–718.

CAIRNS, A. J. G, BLAKE, D., DOWD, K., COUGHLAN, G. D., EPSTEIN, D., ONG, A. & BALEVICH, I. (2009) A quantitative comparison of stochastic mortality models using data from England and Wales and the United States, *North American Actuarial Journal*, **13(1)**, 1–35.

CARRIERE, J. F. (1994) Dependent decrement theory, *Transactions of the Society of Actuaries*, Volume XLVI, 1–21.

CMIB (CONTINUOUS MORTALITY INVESTIGATION BUREAU) (1999) Report Number 17, Institute and Faculty of Actuaries.

LEE, R. D. AND CARTER, L. (1992) Modelling and forecasting the time series of US mortality, *Journal of the American Statistical Association*, **87**, 659–671.

OEPPEN, J. AND VAUPEL, J. W. (2002) Broken limits to life expectancy, *Science*, Volume 296, 10 May 2002, 1029–1031.

PLAT, R. (2011). One-year Value-at-Risk for longevity and mortality, *Insurance: Mathematics and Economics*, **49(3)**, 462–470.

RICHARDS, S. J. & CURRIE, I. D. (2009) Longevity risk and annuity pricing with the Lee-Carter model, *British Actuarial Journal*, Volume 15, Part II, No. 65, 317–365 (including discussion).

RICHARDS, S. J. (2010) Selected Issues in Modelling Mortality by Cause and in Small Populations, *British Actuarial Journal*, 15 Supplement, pages 267–283.

RICHARDS, S. J., CURRIE, I. D. & RITCHIE, G. P. (2012) A value-at-risk framework for longevity trend risk, *presented to the Faculty of Actuaries in Edinburgh on 19th November 2012*.

WILLETS, R. C. (1999) Mortality in the next Millennium, *Staple Inn Actuarial Society*, London.

WILLETS, R. C. (2012) Longevity and the 2011 Census, http://www.longevitas.co.uk/site/informationmatrix/longevityandthe2011census.html [Accessed on 12th November 2012].