Pensions Institute

DISCUSSION PAPER PI-0703

Facing up to uncertain Life Expectancy: The Longevity Fan Charts

Kevin Dowd, David Blake, and Andrew J.G. Cairns

February 2010

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

FACING UP TO UNCERTAIN LIFE EXPECTANCY: THE LONGEVITY FAN CHARTS*

KEVIN DOWD, DAVID BLAKE, AND ANDREW J.G. CAIRNS

This article uses longevity fan charts to represent the uncertainty in projections of future life expectancy. These fan charts are based on a mortality model calibrated on mortality data for English and Welsh males. The fan charts indicate strong upward sloping trends in future life expectancy. Their widths indicate the extent of uncertainty in these projections, and this uncertainty increases as the forecast horizon lengthens. Allowing for uncertainty in the parameter values of the model adds further to uncertainty in life expectancy projections. The article also illustrates how longevity fan charts can be used to stress-test longevity outcomes.

It has become increasingly clear over the past few years that life expectancy has not only been rising but has been rising at a much faster rate than was previously anticipated. Future life expectancies are therefore uncertain, and this uncertainty is no longer seriously disputed.[1]

This uncertainty has major implications for those providing services to the elderly, such as health and long-term care or pensions. For example, if people are living longer than previously anticipated, pension providers will be paying out for longer periods and someone has to bear the resulting higher costs. Uncertain longevity exposes pension funds, life companies, and the state itself to longevity risk, and their exposure to this risk ought to be managed.

The phenomenon of rising but uncertain longevity has major public policy implications. As Dr. Mervyn King, the Governor of the Bank of England, stated in a recent lecture to the British Academy, "We cannot avoid taking decisions, so we must accept the need to analyse the uncertainty that inevitably surrounds them" (King 2004). And yet, as Dr. King explained, "policy debates continue to be permeated by [an] 'illusion of certainty'" that refuses to acknowledge the uncertainty intrinsic to any forecasts of the future. This refusal to acknowledge risks makes it difficult to manage them and effectively undermines much public policy debate.

Some indication of the scale of the longevity problem can be seen in the revisions made to expected longevity forecasts over the past 20 years. For instance, in 1980, the U.K. CMI anticipated that a British man who reached 60 in 1999 could expect to live another 21 years; however, by 1999, that forecast was revised upward to 26 years (CMI 2006a). Over the course of almost 20 years, the expected remaining lifetime of a 60-year-old man had increased by 5 years. There can therefore be little doubt that forecasts of expected longevity are highly uncertain.

---

*Kevin Dowd, Pensions Institute, Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, United Kingdom; e-mail: Kevin.Dowd@hotmail.co.uk. David Blake, Pensions Institute, Cass Business School, City University, United Kingdom. Andrew J. G. Cairns, Maxwell Institute for Mathematical Sciences, Edinburgh, and Department of Actuarial Mathematics and Statistics, Heriot-Watt University, United Kingdom.
The authors thank Chris O'Brien; the Editor, Kenneth Land; and three anonymous referees for helpful comments that have much improved the paper. The usual caveat applies.

1. To illustrate, the uncertainty surrounding expected future lifetimes was formally acknowledged by the U.K. Continuous Mortality Investigation (CMI) when it published its proposed new mortality tables (the "00" series tables) in September 2005 (CMI 2006b); the CMI now advises practicing actuaries not to rely on a single mortality projection, but instead to rely on a range of scenarios reflecting the uncertainty in its projections.

ASSESSING LONGEVITY RISK: ALTERNATIVE APPROACHES

One approach to this problem is to consider alternative expert views about the impact of potential biomedical factors on future longevity. However, the range of alternative views is vast: on the one hand, "pessimists," led by Jay Olshansky (e.g., Loladze 2002; Mizuno et al. 2004; Olshansky, Carnes, and Cassel 1990; Olshansky, Carnes, and Désesquelles 2001; Olshansky et al. 2005), have suggested that future life expectancy might level off or even decline because of factors such as obesity and decreased food-derived health benefits associated with higher levels of atmospheric CO_2. On the other, "optimists," led by James Vaupel (e.g., Oeppen and Vaupel 2002; Tuljapurkar 2005; Tuljapurkar, Li, and Boe 2000; Vaupel et al. 1998), have argued that there is no natural upper limit to the length of human life. Moreover, even demographers critical of the extrapolative forecasting approach adopted by Vaupel and other "optimists" have still accepted the possibility that scientific advances and the sociopolitical responses to them might lead to substantial increases in life expectancy over the next century (e.g., de Grey 2006).

An alternative is to carry out stress tests. For example, the U.K. Pensions Commission and the Government Actuary's Department (GAD) recently carried out a stress test in which they hypothesized that mortality rates might have a 1% error that compounds over time. Using this approach, they found that the life expectancy of 65-year-old males (as of 2005) could lie anywhere between 17.7 and 20.5 years, with a baseline principal projection estimate of 19 years.[2] By 2040, their approach suggested a range between 17.2 and 26.7 around a base case of 21.3 years (Pensions Commission 2005). These and other results from subjective stress-testing exercises reported in the Pensions Commission Report (and elsewhere) are shown in Figure 1 and reveal that there is no consensus surrounding projections of life expectancy in the United Kingdom.[3]

STOCHASTIC MORTALITY MODELS

Empirical observations about changing trends have been accompanied by a growing effort to model future changes in mortality and life expectancy. So-called extrapolative methods using stochastic models date back at least to the landmark paper by Lee and Carter (1992). The nature of extrapolative methods, and their advantages and dangers, are discussed by Wilmoth (1998) and Booth (2006). Since 1992, numerous studies have sought to develop the Lee-Carter approach. Some authors have sought to improve the estimation methodology (see, e.g., Booth, Maindonald, and Smith 2002; Brouhns, Denuit, and Vermunt 2002; De Jong and Tickle 2006; and the discussion in Andreev and Vaupel 2006). A criticism of the Lee-Carter model is that its dynamics are dependent on a single source of risk, a property that does not sit comfortably with the observation of, for example, Wilmoth (1998) and Cairns, Blake, and Dowd (2008) that there have been different patterns of mortality decline at different ages. Other authors have therefore sought to develop new models that incorporate additional sources of risk (see, e.g., Cairns, Blake and Dowd 2006; Cairns et al. 2009; and Renshaw and Haberman 2003), while other studies have sought to incorporate the effects of parameter uncertainty (Cairns et al. 2006; Koissi, Shapiro, and Högnäs 2006; Li and Lee 2005) and model risk (Cairns et al. 2009). These studies have typically found that

2. Strictly speaking, these projections refer to the so-called cohort life expectancy, which is estimated using the 2005 value of the mortality rate of 65-year-olds, the 2006 projection of the mortality rate of 66-year-olds, the 2007 projection of the mortality rate of 67-year-olds, and so forth. This is to be distinguished from period life expectancy, which is estimated using 2005 values of the mortality rates of all ages.

3. We can also compare actual mortality outcomes against earlier forecasts. Using this approach, the Pensions Commission found that the then latest available forecast of the mortality rate for 65-year-old males in 2004 was some 41% lower than anticipated in GAD forecasts made in 1984 (Pensions Commission 2005).

Figure 1. Recent Projections of Life Expectancy for 65-Year-Old English and Welsh Males

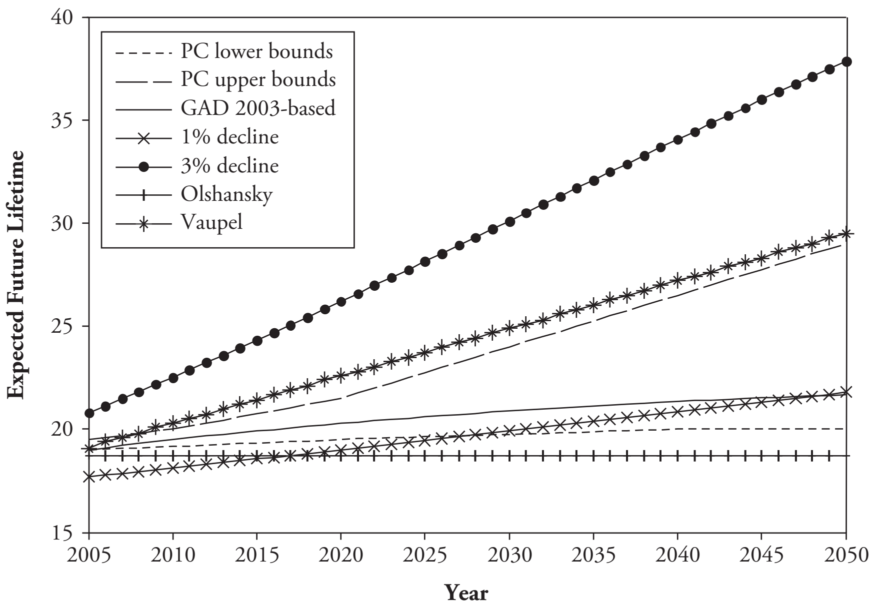

Notes: "PC lower bound" and "PC upper bound" refer to the bounds of the putative 90% prediction interval taken from Figure E.11 of the Pensions Commission (2005) report; "GAD 2003-based" refers to the GAD principal projection based on 2003 data obtained from the same source; "1% decline" and "3% decline" refer to the Pension Commission projections based on 1% and 3% declines in mortality, and are taken from Figure E.4 of the Pensions Commission (2005) report; "Olshansky" and "Vaupel" refer to the putative Olshansky and Vaupel projections given in Figure E.5 of the Pensions Commission (2005) report. For more details concerning the assumptions underlying these projections, see Pensions Commission (2005: Appendix E).

parameter uncertainty in a model can have a significant impact on the level of uncertainty in forecast mortality rates and life expectancy.[4]

QUANTIFYING AND ILLUSTRATING LONGEVITY RISK WITH FAN CHARTS

In this section, we quantify and illustrate longevity risk using fan charts: these are charts showing some central projection (such as the median, mode, or mean) of the variable of interest—in our case, expected future lifetime (EFL)—surrounded by a set of probability bounds. These are shaded darkest around the most likely central projection and become lighter as we move outward toward less likely outcomes. Hence, the fan charts are shaded so that the degree of shading reflects the forecasted probability of the outcome.[5] The longevity fan charts themselves are based on a specified stochastic mortality model. The model chosen is the two-factor CBD model (Cairns et al. 2006), which is known to provide a good fit to English and Welsh male mortality data. This model comes in two versions: a version

4. There are thus three different types of uncertainty to bear in mind: model uncertainty (i.e., we do not know the true mortality model), parameter uncertainty (i.e., whatever mortality model we use, we do not know the true values of its parameters), and forecast uncertainty (i.e., the uncertainty of future mortality rates given any particular model and its calibration).

5. We can also think of fan chart forecasts in a related way: each fan chart gives probability density forecasts for each of $t = 1,2,\dots,T$ periods ahead. Consequently, for any given t, the fan chart gives us a density forecast of the EFL for that future period.

in which the values of the model's parameters are assumed to be known with certainty, and a version that allows for uncertainty in the values of the model's parameters. More details of this model are provided in the appendix.

The earliest fan charts appear to be the Bank of England's inflation fan charts, which were first published in the February 1996 issue of the Bank's *Inflation Report* (Bank of England 1996), and which have been published in every *Inflation Report* since. The first life expectancy fan chart was published in Dr. King's lecture and gave prediction intervals for female life expectancy at birth from 2004 out to a little past 2055: this fan chart had life expectancy starting out at 81; its central projection then rose to about 86.5 years, with a 90% prediction interval running from a little less than 84 to almost 90 years. However, King's longevity fan charts were essentially illustrative, and he did not disclose the underlying statistical model. Other life-expectancy fan charts were produced by Sanderson and Scherbov (2004), utilizing data for 14 countries, but these were based on a rather simple projection model.

We now produce a series of fully calibrated EFL fan charts and use them to illustrate the uncertainty in current and past longevity forecasts. Unlike earlier studies, our fan charts are based on an explicit mortality model, the CBD model, that takes account of possible uncertainty in estimates of the model's parameters, can be scientifically replicated, *and* generates forecasts that are falsifiable. They are also estimated over alternative sample periods to illustrate the sensitivity of results to sample observations.

It is important to emphasize here that these forecasts are based on the assumption that the future will be like the recent past; that is, one is assuming that the historical sample used to calibrate the parameters of the model will give accurate forecasts of future longevity over the forecast horizon. This is a controversial assumption, and the user must interpret the results accordingly. Thus, if a user expected future longevity to improve even more strongly than in the past—for instance, because he/she believed that a cure for cancer was near—then he/she would regard these forecasts as unduly pessimistic. Conversely, a user who expected longevity improvements to slow down or reverse would regard these forecasts as unduly optimistic. However, even if one believed that these forecasts were biased on one side or the other, they nonetheless provide a useful benchmark.[6]

Figure 2 shows the longevity fan chart forecasts for 65-year-old males estimated using mortality data for the period 1961–1980, calibrated on LifeMetrics data for English and Welsh males.[7] Superimposed on the fan charts are the risk bounds—that is, the bounds of the 90% prediction intervals—marked by the two black lines indicated with circles. We can regard these charts as the forecasts a modeler would have been able to make back in 1980, given the model and data up to that year. The forecasts have a maximum horizon of 26 years extending to 2006. Panel a shows the fan chart on the assumption that the parameters of the model are known with certainty, and Panel b shows the corresponding fan chart when we make allowance for possible uncertainty in these parameters. Both fan charts have an initial (as of 1980) EFL of about 13.6 years. The most likely outcomes given by the dark paths in the center of each fan chart are projected to rise gradually over the horizon period to 15.2 years for the parameter certain (PC) case and to 14.9 years for the parameter uncertain (PU) case. The widths of the fan chart intervals show that the projections of life expectancy are quite uncertain and that this uncertainty increases as the forecast horizon increases. A comparison of the two fan charts also shows that the bounds of the parameter uncertain forecasts are considerably wider than those of the parameter certain forecasts—a finding

6. Moreover, as we show in the next section, the fan chart approach can also be adapted to carry out stochastic stress tests of hypothetical mortality scenarios, as well as to provide probability forecasts.

7, These data are available online at http://www.lifemetrics.com.

Figure 2. Longevity Fan Charts for 65-Year-Old English and Welsh Males: 1980–2006

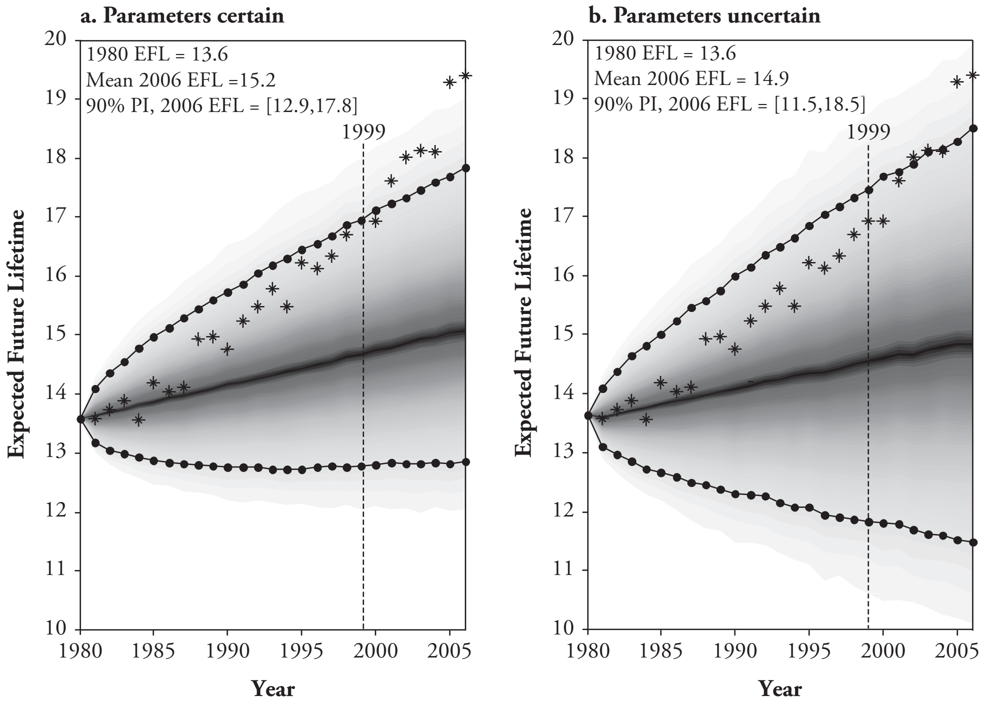

Notes: The charts show the period-by-period forecasted probability densities of future cohort EFL for forecasts made as of 1980 based on LifeMetrics England and Wales mortality data for ages 64–84 and years 1961–1980, and the black lines marked by circles are the bounds of the 90% prediction intervals (PI). Panel a gives the forecasts based on the assumption that the parameters of the model are known with certainty, whereas Panel b gives forecasts based on the assumption that these parameters are estimated with possible error. Forecasts are based on the M5 representation of the Cairns-Blake-Dowd mortality model (see Cairns et al. 2006; and Cairns et al. 2009). The starred points are estimates of subsequently realized EFL for each year over the forecast horizon.

that is consistent with the earlier studies we cited that have also examined the impact of parameter uncertainty on forecasts.[8]

The fan charts also show estimates of subsequently realized values for EFLs for 65-year-old males over the period 1980 to 2006, indicated by the stars in the chart.[9] These are quite close to the central projections of the fan charts for the first seven years, but then gradually move upward relative to the fan chart forecasts and eventually breach and then exceed the upper bound of the 90% prediction interval. This indicates that after seven years

8. The principal reason for this increased width is uncertainty in the underlying trend rather than in the volatility of mortality rates. As our time horizon increases, uncertainty in the trend dominates all other sources of risk in influencing the width of the right side fan chart. This confirms that longevity risk is above all a trend risk: getting the trend right is the key to successful forecasting.

9. The realized EFL for a male aged 65 in year T measures the EFL based on a model estimated using data up to and including year T, making full allowance for forecasted future stochastic mortality improvements. This contrasts with the forecast EFL for a male aged 65 in future year T, which measures the EFL based on a data sample whose latest observation antecedes T, but again making full allowance for forecasted future stochastic mortality improvements.

Table 1. Longevity Forecast Errors Over Long Horizons

	Parameters Certain			Parameters Uncertain		
Horizon	Mean Forecast	Outcome	Error	Mean Forecast	Outcome	Error
20 Years	14.8	16.9	2.1	14.6	16.9	2.3
26 Years	15.2	19.4	4.2	14.9	19.4	4.5

Notes: The table shows the mean forecasts, realized outcomes, and associated forecast errors for forecasted probability densities of future EFL for forecasts made as of 1980, based on LifeMetrics English and Welsh mortality data for ages 64–84 and years 1961–1980. Forecasts are derived from the M5 representation of the Cairns-Blake-Dowd mortality model (see Cairns et al. 2006; and Cairns et al. 2009).

or so, the fan chart projections start to underestimate future EFLs, and the degree of underestimation tends to rise thereafter with the length of the forecast horizon.

In addition, because the PU fan charts are wider than the PC ones, the realized EFL outcomes associated with the former are more likely to be true (in the sense that they usually have higher *p* values) than those associated with the latter. To illustrate, the charts show the vertical lines that highlight the projections and realized outcomes for 1999. In the PC case, the realized value sits astride the upper risk bound, indicating that this outcome has a *p* value under the null of about 5%; by contrast, the realized value for the PU case lies comfortably below the upper risk bound, indicating that this outcome has a *p* value of well above 5%.

The plots of the risk bounds also allow us to test the performance of the fan chart forecasts. Under the null hypothesis that the forecasts are adequate, the probability of any given realized outcome breaching either bound is 5%. We then see that for the PC case, the last six realized outcomes are around or above the upper risk bound; if we take 5% as our significance level, then we would say that the PC forecasts are broadly adequate by this criterion up to about 2000. For the PU case, only the last two realized outcomes are above the upper risk bound, so the PU fan chart forecasts are adequate up to about 2004. These results give us some reason to be confident in the model's forecasting ability, except for very long forecast horizons, and more so for the PU than the PC case.

We can also assess the performance of the fan chart forecasts by comparing their central projections against realized outcomes—that is, we can assess their forecast errors. Table 1 presents some forecast errors over horizons of 20 and 26 years. For the earlier horizon, the forecast errors are 2.1 years and 2.3 years for the PC and PU cases, respectively; the corresponding forecast errors for the 26-year forecast horizon are 4.2 and 4.5. These are not ideal, but we have just seen that the former are not statistically significant, and both compare favorably with the forecast error of 5 years for the CMI's 20-year projections for 60-year-olds mentioned earlier.

Moving ahead in time, Figure 3 gives more contemporary fan chart forecasts based on a sample covering the years 1987–2006. These forecasts start at 2006 and extend out over a 50-year horizon to 2056.[10] The EFL now starts at about 19.7 or 19.8 years, and the most likely EFL for 65-year-old males in 2056 is projected to be 27.2 years for the PC fan chart and 27.4 for the PU fan chart. Again we find the same pattern of a rising trend and increasing uncertainty as the forecast horizon lengthens. For this later sample period, the PU fan charts are very wide indeed and are considerably wider than their PC equivalents. For example, the PU fan chart projects a 90% prediction interval for the EFL of 65-year-old

10. Considerable caution needs to be taken when projecting as far ahead as 50 years with only 20 years of historical data (see, e.g., Wilmoth 1998). This concern is addressed in part through the inclusion of parameter uncertainty. In the 50-year fan charts, we can see that the differences between the PC and PU cases are relatively small initially but get substantially larger as the time horizon increases.

Figure 3. Longevity Fan Charts for 65-Year-Old English and Welsh Males: 2006–2056

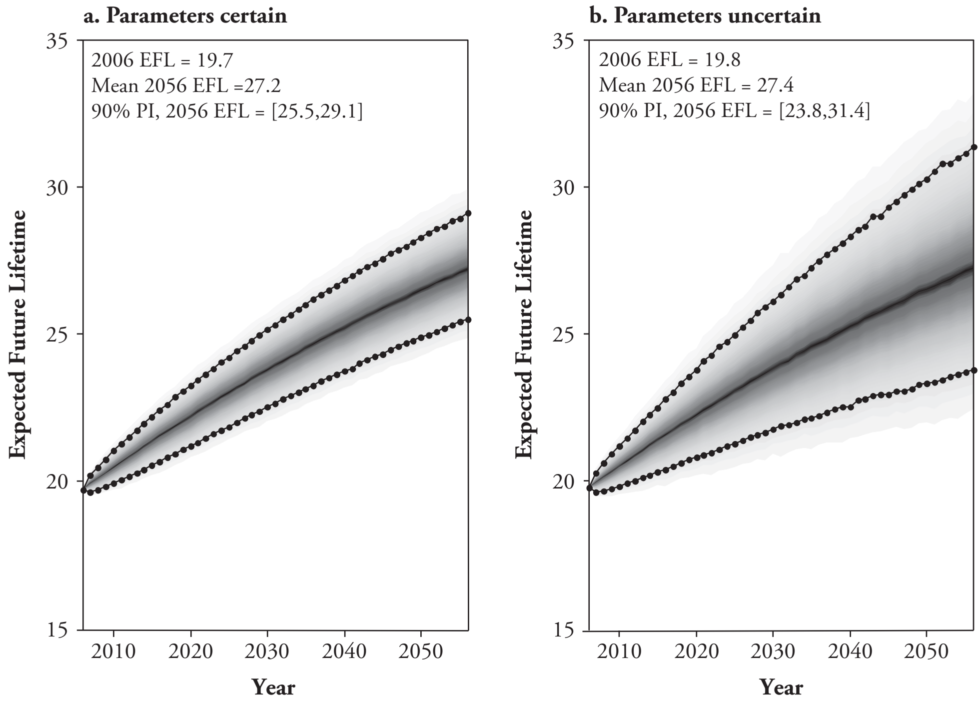

Notes: The charts show the period-by-period forecasted probability densities of future cohort EFL for forecasts made as of 1980 based on LifeMetrics England and Wales mortality data for ages 64–84 and years 1986–2005, and the black lines marked by circles are the bounds of the 90% prediction intervals (PI). Panel a gives forecasts based on the assumption that the parameters of the model are known with certainty, whereas Panel b gives forecasts based on the assumption that these parameters are estimated with possible error. Forecasts are based on the M5 representation of the Cairns-Blake-Dowd mortality model (see Cairns et al. 2006; and Cairns et al. 2009).

males in 2056 that stretches from 23.8 to 31.4 years, whereas the corresponding prediction interval for the PC fan chart is only 25.5 to 29.1. Future life expectancy as viewed from 2006 would therefore appear to be very uncertain.

STRESS TESTING LONGEVITY RISK WITH FAN CHARTS

We can also adapt the fan chart approach to carry out stress tests against specified what-if scenarios. To give a simple example, we might ask, what if future mortality rates fall by x% relative to those projected by the mortality model?[11] To carry out such an exercise, we would reduce the mortality rates in our earlier forecasts by x% and recalculate the fan charts.[12]

11. Say, as a result of a biomedical marker, such as a cure for cancer.

12. We can, of course, imagine all manner of other possible scenarios. For instance, we might ask, what if there were a sudden change in mortality rates at some future time T? The change might be temporary or permanent, anticipated or unanticipated, or might affect some ages more than others, and so forth. Alternatively, we might specify a possible event (e.g., a cure for cancer) and then hypothesize how this might affect mortality rates. All these scenarios can be modeled using the fan chart approach provided that we are prepared to make appropriate assumptions.

An example is given by Figure 4, which shows the fan charts associated with the hypothetical scenario in which mortality rates fall by 3% relative to those projected by our mortality model. The main results of this stress-test exercise are presented in Table 2. We see that a hypothetical fall in our mortality rates of 3% relative to those underlying Figure 3 would lead life expectancies as of 2006 to rise by about 0.22 to 0.25 of a year, or by about 1.1% to 1.3%. For its part, projected life expectancy in 2056 would rise by about 0.7%.

It is interesting to compare these results with the projections in Figure 1. The mean projections for 2056 from our stress tests turn out to be quite close to those of Vaupel and the upper bound of the 90% prediction interval hypothesized by the Pensions Commission, but are well below the 3% reduction in the mortality rate also hypothesized by the Pensions Commission.

CONCLUSIONS

We believe that longevity fan charts are a useful way of showing the stochastic nature of improvements in life expectancy over time and of projecting into the future the longevity uncertainty inherent in recent mortality data. They can also be adapted to carry out

Figure 4. Stress-Test Longevity Fan Charts for 65-Year-Old English and Welsh Males: 2006–2056

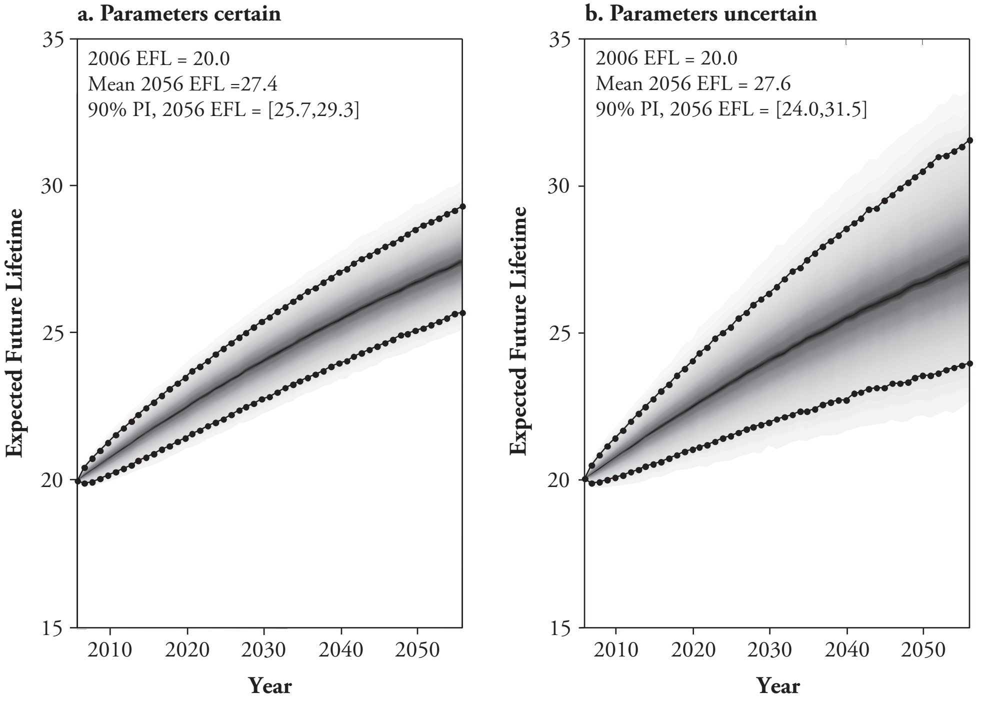

Notes: The charts show the stress-tested period-by-period forecasted probability densities of future cohort EFL for forecasts made as of 1980 based on LifeMetrics England and Wales mortality data for ages 64–84 and years 1986–2005, and the black lines marked by circles are the bounds of the 90% prediction intervals (PI). Panel a gives the forecasts based on the assumption that the parameters of the model are known with certainty, whereas Panel b gives forecasts based on the assumption that these parameters are estimated with possible error. Forecasts are based on the M5 representation of the Cairns-Blake-Dowd mortality model (see Cairns et al. 2006; and Cairns et al. 2009). The stress test assumes that mortality rates are 3% lower than the predictions of the M5 mortality model.

Table 2. Expected Future Lifetimes: Forecasts Versus Stress-Test Outcomes

Variable	Parameters Certain				Parameters Uncertain			
	2006	2056	90% Prediction Interval		2006	2056	90% Prediction Interval	
Forecast (Figure 2)	19.724	27.237	25.500	29.106	19.784	27.358	23.792	31.364
Stress Test (Figure 3)	19.972	27.436	25.692	29.314	20.003	27.553	23.980	31.543
Change	0.25	0.20	0.19	0.21	0.22	0.20	0.19	0.18
% Change	1.3	0.7	0.6	0.7	1.1	0.7	0.8	0.6

Notes: The table shows the forecasts and stress-test outcomes of the following: the 2006 EFLs, the mean EFLs for 2056, and the 90% prediction intervals for the 2056 EFLs. Estimates are made as of 2006 based on LifeMetrics English and Welsh mortality data for ages 64–84 and years 1987–2005, and are obtained using the M5 representation of the Cairns-Blake-Dowd mortality model (see Cairns et al. 2006; and Cairns et al. 2009). The stress test assumes that mortality rates are 3% lower than the forecasts.

stochastic mortality stress tests. Our main findings confirm that expected future lifetimes are projected to increase strongly, but we also find that these projections are highly uncertain and become more so as the forecast horizon lengthens. These findings are bad news for those with obligations to pay pensions or otherwise provide for the elderly: it forces them to anticipate large numbers of people living to very old ages while also raising the question of how to manage the financial risks created by the uncertainty of future longevity. We also find that allowing for parameter uncertainty makes the fan charts even wider—that is, it makes us even more uncertain about future longevity. This finding makes the bad news even worse.

We would stress that although our results are based on a model calibrated on English and Welsh male mortality experience, we have every reason to expect that similar findings would be obtained for both males and females for comparable countries. Our findings therefore have potentially disturbing implications for the health, pensions, and life insurance industries in many countries, and for public policy generally. Nevertheless, policy makers have recently begun to acknowledge these issues, and the capital markets are currently developing financial instruments that can be used to hedge longevity risk once fan charts have been used to help quantify it (see, e.g., Coughlan et al. 2007; and Loeys, Panigirtzoglou, and Ribeiro 2007).

APPENDIX: A MODEL OF EXPECTED LONGEVITY

The longevity fan charts are based on the M5 representation of the mortality model set out by Cairns et al. (2006; see also Cairns et al. 2009). Let $q(t,x)$ be the realized mortality rate in year $t+1$ (that is, from time t to time $t+1$) of a cohort aged x at time 0. We assume that the logit of $q(t,x)$ is governed by a two-factor Perks stochastic process (Perks 1932):

$$\text{logit } q(t,x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x}), \tag{1}$$

where $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ are stochastic processes that are measurable at time $t+1$, and where $\bar{x}$ is the mean of the ages used to estimate the parameters of the model. Now let $\boldsymbol{\kappa}_t = (\kappa_t^{(1)}, \kappa_t^{(2)})'$, and assume that $\boldsymbol{\kappa}_t$ is a random walk with drift:

$$\boldsymbol{\kappa}_t = \boldsymbol{\kappa}_{t-1} + \boldsymbol{\mu} + \mathbf{CZ(t)}, \tag{2}$$

where $\boldsymbol{\mu}$ is a constant 2×1 vector of drift parameters, $\mathbf{C}$ is a constant 2×2 lower triangular matrix reflecting volatilities and correlations, and $\mathbf{Z(t)}$ is a 2×1 vector of independent standard normal variables. Cairns et al. (2006) showed that this model provides a good fit to data for English and Welsh males over the period 1961–2002.

For future time t, we wish to estimate the expected future lifetime for a male aged 65 at that time, which we denote by $EFL(t)$. $EFL(t)$ is a function of $\boldsymbol{\kappa}_t$ and is given by

$$EFL(t) \approx \frac{1}{2} + \sum_{u=t+1}^{\infty} E[S(u)/S(t)|\boldsymbol{\kappa}_t] \tag{3}$$

where $E[S(u)/S(t)|\boldsymbol{\kappa}_t]$ is the expected value of $S(u)/S(t)$ conditional on the values of $\boldsymbol{\kappa}_t$, and $S(u)$ is the proportion of males aged 65 at time t who are still alive at time u.[13]

Estimating $EFL(t)$ is a challenging task. For the model defined in Eqs. (1) and (2), $EFL(t)$ does not have a closed-form solution and hence cannot be evaluated analytically. At the same time, the problem does not lend itself to standard Monte Carlo simulation because the $\boldsymbol{\kappa}_t$ are themselves random. So although we can use stochastic simulation to estimate $EFL(t)$ *given* $\boldsymbol{\kappa}_t$, estimating $EFL(t)$ for a *random future* $\boldsymbol{\kappa}_t$ is altogether more difficult. For example, if we want to obtain m estimates of $EFL(t)$, and if each of these estimates is based on a typically different random $\boldsymbol{\kappa}_t$, then under conventional stochastic simulation, we would run m^2 sets of simulation trials: that is, we would simulate m sets of $\boldsymbol{\kappa}_t$, and for each of these, we would carry out m further simulation trials to obtain one estimate of $EFL(t)$. For the values of m needed to get accurate results, this simulation-within-simulation approach would be very time-intensive even for fast, modern computers.

To avoid these difficulties, we resort to the following procedure, which nests a Taylor series approximation within a stochastic simulation approach. More specifically, we first run a preliminary stochastic simulation exercise that estimates the parameters of the second-order Taylor series approximation for each $E[S(u)/S(t)|\boldsymbol{\kappa}_t]$ term, where the approximation is centered on the expected value of $\boldsymbol{\kappa}_t$, which is itself easily obtained from (2).[14] We then simulate m sets of $\boldsymbol{\kappa}_t$ values, and for each simulated $\boldsymbol{\kappa}_t$ pair, we estimate $EFL(t)$ using the Taylor series approximations of the $E[S(u)/S(t)|\boldsymbol{\kappa}_t]$. For any given t, this approach gives us m estimates of $EFL(t)$, and we estimate the median and 90% prediction bounds for $EFL(t)$ from their order statistics.

Where the $\boldsymbol{\mu}$ and $\mathbf{C}$ parameters are simulated under the PU case, the simulation approach used is a standard Bayesian one suitable for simulating uncertain parameters, bearing in mind that $\boldsymbol{\kappa}_t$ is subjected to multivariate normal shocks with mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{V}$. In the absence of any clear beliefs about the value of $\boldsymbol{\mu}$ and $\mathbf{V}$, we seek a noninformative prior, and a natural choice is the Jeffreys prior:

$$p(\boldsymbol{\mu},\mathbf{V}) \propto |\mathbf{V}|^{-3/2}, \tag{4}$$

where $|\mathbf{V}|$ is the determinant of $\mathbf{V}$. If we let $\mathbf{D}_t = \boldsymbol{\kappa}_t - \boldsymbol{\kappa}_{t-1}$, then the posterior distribution for $\boldsymbol{\mu}$ and $\mathbf{V}$, given $\mathbf{D}$, is

$$\mathbf{V}^{-1}|\mathbf{D} \sim Wishart(n-1, n^{-1}\hat{\mathbf{V}}^{-1}) \tag{5}$$

$$\boldsymbol{\mu}|\mathbf{V},\mathbf{D} \sim WVN(\hat{\mu}, n^{-1}\hat{\mathbf{V}}^{-1}), \tag{6}$$

where $\hat{\boldsymbol{\mu}}$ and $\hat{\mathbf{V}}$ are sample estimates of $\boldsymbol{\mu}$ and $\mathbf{V}$, and n is the size of the $\mathbf{D}$ sample. Eqs. (5) and (6) can be programmed, and simulated values of $\mathbf{C}$ can be obtained by running

13. The "1/2" term in (3) is added because the rest of the right side of (3) refers to the curtate expectation of life—that is, the expected value of the completed years survived. Adding half a year gives a more accurate estimate of expected future lifetime when one wishes to take account of partial years as well as completed years.

14. In fact, the approximation was applied to the probits of the $E[S(u) / S(t)| \boldsymbol{\kappa}_t]$, which is more accurate than applying the approximation to the expected survivor rates themselves. For more details, see Cairns (2007).

the simulated **V** matrix through a lower triangular Choleski decomposition. For further details, see Cairns et al. (2006:695–97).

The realized EFL series in Figure 2 was estimated using the same M5 model but with the most recent available data for the year concerned; that is, the 1981-realized EFL was estimated using data for 1962–1981, and so forth, and the 2006-realized EFL was estimated using data for 1987–2006.[15]

The results reported in this article were based on the following sets of parameters. For Figures 2–4, we used m = 10,000, and models were estimated for ages 64–84 to give us $\bar{x} = 74$. Figure 2 was then estimated over years 1961–1980 using the following parameter estimates obtained by maximum likelihood: $\kappa_t^{(1)} = -2.611$, $\kappa_t^{(2)} = 0.092823$, $\hat{\boldsymbol{\mu}} = [-0.0085624\ \ 0.00013585]^T$, and $\hat{\mathbf{V}}$ =[0.0014766 3.7613e-005; 3.7613e-005 1.2909e-005]. Figure 3 was estimated over the years 1987–2006 using these ML parameter estimates: $\kappa_t^{(1)} = -3.2907$, $\kappa_t^{(2)} = 0.11132$, $\hat{\boldsymbol{\mu}} = [-0.029247\ \ 0.00089111]^T$, and $\hat{\mathbf{V}}$ = [0.00042056 4.3708e-006; 4.3708e-006 9.3843e-007]. Figure 4 was estimated using the same parameters as Figure 3, but on the assumption that mortality rates are 3% lower than those forecasted by the M5 model.[16]

REFERENCES

Andreev, K.F. and J. Vaupel. 2006. "Forecasts of Cohort Mortality After Age 50." MPIDR Working Paper WP 2006-012. Max Planck Institute for Demographic Research, Rostock, Germany.

Bank of England. 1996. *Inflation Report*. London: Bank of England.

Booth, H. 2006. "Demographic Forecasting: 1980 to 2005 in Review." *International Journal of Forecasting* 22:547–81.

Booth, H., J. Maindonald, and L. Smith. 2002. "Applying Lee-Carter Under Conditions of Variable Mortality Decline." *Population Studies* 56:325–36.

Brouhns, N., M. Denuit, and J.K. Vermunt. 2002. "A Poisson Log-Bilinear Regression Approach to the Construction of Projected Life Tables." *Insurance: Mathematics and Economics* 31:373–93.

Cairns, A.J.G. 2007. "The CBD 2-Factor Model: A Note on Approximations to Survivor and Financial Functions." Unpublished technical report. Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh, Scotland.

Cairns, A.J.G., D. Blake, and K. Dowd. 2006. "A Two-Factor Model for Stochastic Mortality With Parameter Uncertainty: Theory and Calibration." *Journal of Risk and Insurance* 73:687–718.

———. 2008. "Modelling and Management of Mortality Risk: A Review." *Scandinavian Actuarial Journal* 108:79–113.

Cairns, A.J.G., D. Blake, K. Dowd, G.D. Coughlan, D. Epstein, A. Ong, and I. Balevich. 2009. "A Quantitative Comparison of Stochastic Mortality Models Using Data From England and Wales and the United States." *North American Actuarial Journal* 13:1–35.

Continuous Mortality Investigation (CMI). 2006a. "The Graduation of the CMI 1999–2002 Mortality Experience: Final '00' Series Mortality Tables—Assured Lives." Working Paper 21. Institute and Faculty of Actuaries, London and Edinburgh.

———. 2006b. "The Graduation of the CMI 1999–2002 Mortality Experience: Final '00' Series Mortality Tables—Annuitants and Pensioners." Working Paper 22. Institute and Faculty of Actuaries, London and Edinburgh.

15. This way of estimating realized EFL ensures a consistent comparison between the forecasted and realized EFLs. At the same time, it still allows for a valid test of the forecasts because it is always possible—as we see in Figure 2—for estimates of realized EFL to fall well outside the predicted ranges and therefore to "fail" a test of forecast adequacy. Thus, forecasts are still falsifiable, even though we use the same model to estimate realized EFL as we use to forecast future EFL.

16. The calculations were made using functions that were specially written in MATLAB. These will be made available on request to bona fide academic researchers three years after the paper is published.

Coughlan, G., D. Epstein, A. Sinha, and P. Honig. 2007. “q-Forwards: Derivatives for Transferring Longevity and Mortality Risks.” JPMorgan Pension Advisory Group, London, July 2. Available online at http://www.lifemetrics.com.

De Jong, P. and L. Tickle. 2006. “Extending Lee-Carter Mortality Forecasting.” *Mathematical Population Studies* 13:1–18.

de Grey, A. D. N. J. 2006. “Extrapolaholics Anonymous: Why Demographers’ Rejections of a Huge Rise in Cohort Life Expectancy in This Century Are Overconfident.” *Annals of the New York Academy of Science* 1067:83–93.

King, M.A. 2004. “What Fates Impose: Facing Up to Uncertainty.” *Eighth British Academy Annual Lecture*. Delivered to the British Academy, London, on December 1, 2004.

Koissi, M.C., A.F. Shapiro, and G. Högnäs. 2006. “Evaluating and Extending the Lee-Carter Model for Mortality Forecasting: Bootstrap Confidence Interval.” *Insurance: Mathematics and Economics* 38:1–20.

Lee, R.D. and L.R. Carter. 1992. “Modeling and Forecasting U.S. Mortality.” *Journal of the American Statistical Association* 87:659–75.

Li, N. and R. Lee. 2005. “Coherent Mortality Forecasts for a Group of Populations: An Extension of the Lee-Carter Method.” *Demography* 42:575–94.

Loeys, J., N. Panigirtzoglou, and R.M. Ribeiro. 2007. “Longevity: A Market in the Making.” J.P. Morgan Securities Ltd., London, July 2. Available online at http://www.lifemetrics.com.

Loladze, I. 2002. “Rising Atmospheric CO_2 and Human Nutrition: Toward Globally Imbalanced Plant Stoichiometry?” *Trends in Ecology & Evolution* 17:457–61.

Mizuno, T., I.-W. Shu, H. Makimura, and C. Mobbs. 2004. “Obesity Over the Life Course.” *Science of Aging Knowledge Environment* 24:re4.

Oeppen, J. and J.W. Vaupel. 2002. “Broken Limits to Life Expectancy.” S*cience* 296:1029–31.

Olshansky, S.J., B.A. Carnes, and C. Cassel. 1990. “In Search of Methuselah: Estimating the Upper Limits to Human Longevity.” *Science* 250:634–40.

Olshansky, S.J., B.A. Carnes, and A. Désesquelles. 2001. “Prospects for Human Longevity.” *Science* 291:1491–92.

Olshansky, S.J., D. Passaro, R. Hershow, J. Layden, B.A. Carnes, J. Brody, L. Hayflick, R.N. Butler, D.B. Allison, and D.S. Ludwig. 2005. “A Potential Decline in Life Expectancy in the United States in the 21st Century.” *New England Journal of Medicine* 352:1103–10.

Pensions Commission. 2005. *A New Pension Settlement for the Twenty-First Century: Appendix E: Uncertainties in Life Expectancy Projections*. Norwich: The Stationery Office.

Perks, W. 1932. “On Some Experiments in the Graduation of Mortality Statistics.” *Journal of the Institute of Actuaries* 63:12–57.

Renshaw, A.E. and S. Haberman. 2003. “Lee-Carter Mortality Forecasting With Age-Specific Enhancement.” *Insurance: Mathematics and Economics* 33:255–72.

Sanderson, W. and S. Scherbov. 2004. “Putting Oeppen and Vaupel to Work: On the Road to New Stochastic Mortality Forecasts.” Working paper. International Institute for Applied Systems Analysis, Austria.

Tuljapurkar, S. 2005. “Future Mortality: A Bumpy Road to Shangri-La?” *Science of Aging Knowledge Environment* 14:pe9.

Tuljapurkar, S., N. Li, and C. Boe. 2000. “A Universal Pattern of Mortality Decline in the G7 Countries.” *Nature* 405:789–92.

Vaupel, J., J. Carey, K. Christensen, T. Johnson, A. Yashin, V. Holm, I. Iachine, V. Kannisto, A. Khazaeli, P. Liedo, V. Longo, Y. Zeng, K. Manton and J. Curtsinger. 1998. “Biodemographic Trajectories of Longevity.” *Science* 280:855–60.

Wilmoth, J.R. 1998. “The Future of Human Longevity: A Demographer’s Perspective.” *Science* 280:395–97.