

# A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty

Andrew J.G. Cairns<sup>ab</sup>, David Blake<sup>c</sup>, Kevin Dowd<sup>d</sup>

First version: February 2005

This version: May 11, 2005

## Abstract

In this paper we consider the evolution of the post-60 mortality curve in the UK and its impact on the pricing of the risk associated with mortality improvements over time: so-called longevity risk. We introduce a two-factor stochastic model for the development of this curve through time. The first factor affects mortality-rate dynamics at all ages in the same way, whereas the second factor affects mortality-rate dynamics at higher ages much more than at lower ages.

The paper then examines the pricing of longevity bonds with different terms to maturity referenced to different cohorts. We find that longevity risk over relatively-short time horizons is very low, but at horizons in excess of 10 years it begins to pick up very rapidly. We propose a method for risk-adjusting the market price of a longevity bond which includes an allowance for parameter risk in the stochastic mortality model. We utilise the pricing information contained in the November 2004 European Investment Bank longevity bond to make inferences about the likely market prices of the risks in the model. Based on these, we investigate how future issues might be priced to ensure an absence of arbitrage between bonds with different characteristics.

## Keywords

Longevity risk; Perks model; longevity bond; market price of longevity risk; market price of parameter risk.

## 1 Introduction

Recently it has become clear that mortality is a stochastic process: longevity has not only been improving, but it has been improving in an unpredictable way. These (largely unanticipated) improvements have also been greatest at higher ages, and

---

<sup>a</sup>Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.

<sup>b</sup>Corresponding author: E-mail A.Cairns@ma.hw.ac.uk

<sup>c</sup>Pensions Institute, Cass Business School, City University, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom.

<sup>d</sup>Centre for Risk & Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham, NG8 1BB, United Kingdom.

have caused life offices (and pension plan sponsors in the case where the plan provides the pension) to incur losses on their life annuity business. The problem is that pensioners are living much longer than was anticipated, say, 20 years ago. As a result life offices are paying out for much longer than was anticipated, and their profit margins are being eroded in the process. The insurance industry is therefore bearing the costs of unexpectedly higher longevity. Looking forward, possible changes in lifestyle, medical advances and new discoveries in genetics are likely to make future improvements to life expectancy very unpredictable as well. This, in turn, will lead to smaller books of life annuity business, smaller profit margins or both.

A recent capital market innovation, the longevity bond, provides life offices and pension plans with an instrument to hedge the aggregate mortality risks that they face. The idea for longevity bonds was first published in the *Journal of Risk and Insurance* in 2001<sup>1</sup>. Longevity bonds are annuity bonds whose coupons are not fixed over time, but fall in line with a given survivor index<sup>2</sup>. For example, the survivor index might be based on the population of 65-year-olds alive on the issue date of the bond. Each year the coupon payments received by the life office or pension plan decrease by the percentage of the population who have died that year. If, after the first year, 1.5% of the population of what are now 66-year olds have died, then the coupon payable at the end of that first year will fall to 98.5% of the nominal coupon rate. But this is exactly what the life office or pension plan wants, since only 98.5% of their own 66-year old annuitants (assuming these are representative of reference population) will be alive after one year too, so they do not have to pay out so much.

The issue by the European Investment Bank (EIB) of the world's first longevity bond was announced in November 2004 by its structurer and manager BNP Paribas (BNP). The bond had an initial market value of about £540m and a maturity of 25 years. Its coupon payments are linked to a survivor index based on the realised mortality experience of a cohort of males from England & Wales aged 65 in 2003 as published by the UK Office for National Statistics (ONS). The longevity risk, however, is not borne by the EIB. Instead, BNP arranged for the the longevity risk to be reinsured with Bermuda-based PartnerRe. The intended main investors were UK pension funds and life offices. For further details of the EIB/BNP longevity bond see Appendix A.

In this paper we propose a stochastic mortality model that we fit to UK mortality data and show how the calibrated model can be used to price longevity-dependent financial instruments such as the EIB/BNP longevity bond. The model involves two stochastic factors. The first affects mortality at all ages in an equal manner, whereas the second has an effect on mortality that is proportional to age. We present empirical evidence that indicates that both these factors are needed to achieve a satisfactory empirical fit over the mortality term structure (that is, to model adequately historical mortality improvements at different ages). The resulting model dynamics allow us to simulate cohort survivor rates, thereby enabling us to model longevity risk, and other mortality indices underlying alternative mortality-linked securities.

To price a mortality-linked security we adopt the risk-adjusted (or 'risk-neutral')

---

<sup>1</sup>Blake and Burrows (2001). See also Cox *et al*, 2000.

<sup>2</sup>For this reason they are also known as *survivor* bonds (e.g. Blake and Burrows, 2001).

approach to pricing. Given the relative dearth of market information, we propose a simple method for making the adjustment between real and risk-adjusted probabilities. The magnitude of this adjustment is established by estimating the market prices of longevity and parameter risk implied by the market price of the EIB/BNP longevity bond.

The layout of this paper is as follows. Section 2 outlines the model. Section 3 fits the model to English and Welsh mortality data, and discusses the plausibility of the fit. Section 4 then presents some simulation results for the survivor rate based on the calibrated model. Two alternative sets of simulation results are presented: first, results that do not take account of parameter uncertainty, and, second, results that do take account of such uncertainty. Section 5 discusses the price of longevity risk – that is to say, it discusses the premium that a life office or pension plan might be prepared to pay to lay off such risk – and uses this to show how the EIB/BNP bond might be priced in a risk-adjusted framework. It also presents some illustrative pricing results. Section 6 discusses the risk premium on new issues, and shows how the earlier results might be used to price new longevity bonds with different terms to maturity and following different cohorts. Section 7 concludes.

## 2 Model specification

In Cairns, Blake and Dowd (2004) we introduced notation for *forward survival probabilities*

$$p(t, T_0, T_1, x) = \begin{array}{l} \text{probability as measured at } t \text{ that} \\ \text{an individual aged } x \text{ at time } 0 \text{ and still alive at } T_0 \\ \text{survives until time } T_1 > T_0. \end{array}$$

Let  $I(u)$  represent the indicator process that is equal to 1 at time  $u$  if the life aged  $x$  at time 0 is still alive at time  $u$ . Furthermore let  $\mathcal{M}_u$  be the filtration generated by the development of the mortality curve up to time  $u$ .<sup>3</sup> Then

$$p(t, T_0, T_1, x) = Pr(I(T_1) = 1 | I(T_0) = 1, \mathcal{M}_t).$$

Note that  $p(t, T_0, T_1, x) = p(T_1, T_0, T_1, x)$  for all  $t \geq T_1$ , since the observation period  $(T_0, T_1]$  is then past and not subject to any further uncertainty.

For simplicity in this exposition we will define  $\tilde{p}(t, x) = p(t + 1, t, t + 1, x)$  to be the *realised* survival probability for the cohort aged  $x$  at time 0. Additionally define the realised mortality rate  $\tilde{q}(t, x) = 1 - \tilde{p}(t, x)$ .

In this paper we will restrict ourselves to the following model<sup>4</sup> for the mortality curve:

$$\tilde{q}(t, x) = 1 - p(t + 1, t, t + 1, x) = \frac{e^{A_1(t+1) + A_2(t+1)(x+t)}}{1 + e^{A_1(t+1) + A_2(t+1)(x+t)}}. \quad (1)$$

<sup>3</sup>That is,  $\mathcal{M}_u$  represents the history of the mortality curve up to time  $u$ .

<sup>4</sup>This is a special case of what are known as Perks models: see, for example, Perks (1932) or Benjamin and Pollard (1993).

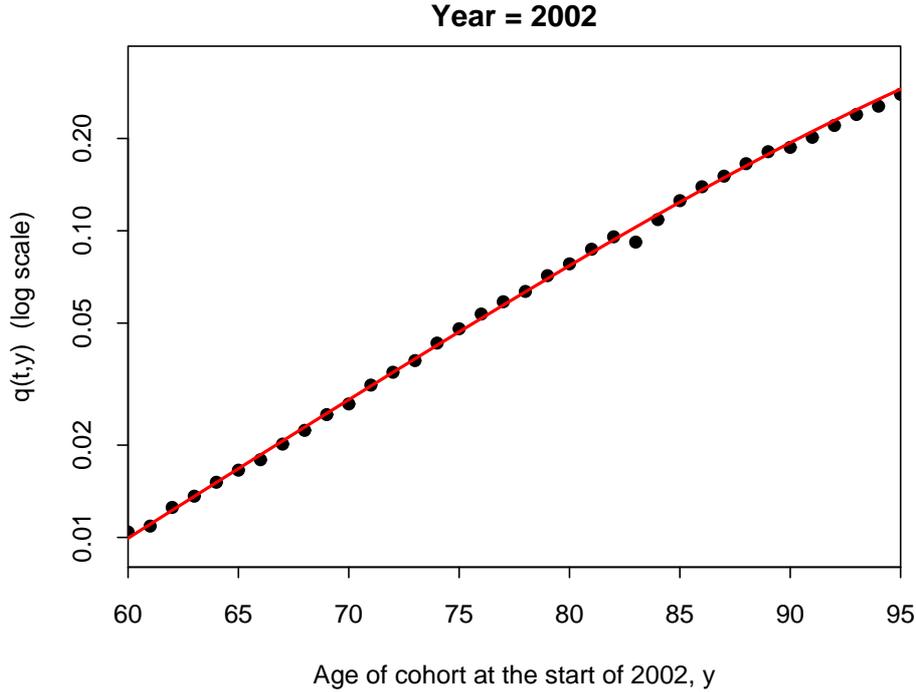


Figure 1: Ungraduated mortality rates above the age of 60 for England & Wales males for the year 2002 (dots) and fitted curve  $e^{A_1+A_2y}/(1+e^{A_1+A_2y})$  for  $A_1 = -10.95$  and  $A_2 = 0.1058$ .

In this equation,  $A_1(u)$  and  $A_2(u)$  are stochastic processes that are assumed to be measurable at time  $u$ . An example of a mortality curve is given in Figure 1. This graph shows the ungraduated mortality rates above the age of 60 for England & Wales males in 2002<sup>5</sup> along with the fitted curve (fitted using least squares applied to (1)). The fit is clearly very good. Simpler parametric curves can also be fitted (for example,  $q_y = a^{A_1+A_2y}$ ) but the chosen curve gives a significantly better fit, especially for higher ages.

### 3 Stochastic mortality

Estimated values for  $A_1(t)$  and  $A_2(t)$  for the years 1961 to 2002 are plotted in Figure 2.<sup>6</sup> These results show a clear trend in both series. The downward trend in  $A_1(t)$  reflects general improvements in mortality over time at all ages. The increasing trend in  $A_2(t)$  means that the curve is getting slightly steeper over time: that is, mortality improvements have been greater at lower ages. There were also changes in the trend and in the volatility of both series. To make forecasts of the future distribution of  $A(t) = (A_1(t), A_2(t))'$ , we will model  $A(t)$  as a two-dimensional random walk with

<sup>5</sup>Available from the Government Actuary's Department website, [www.gad.gov.uk](http://www.gad.gov.uk).

<sup>6</sup>For each  $t$ ,  $A_1$  and  $A_2$  were estimated using least squares by transforming the ungraduated mortality rates from  $q_y$  to  $\log q_y/p_y = A_1 + A_2y + \text{error}$ .

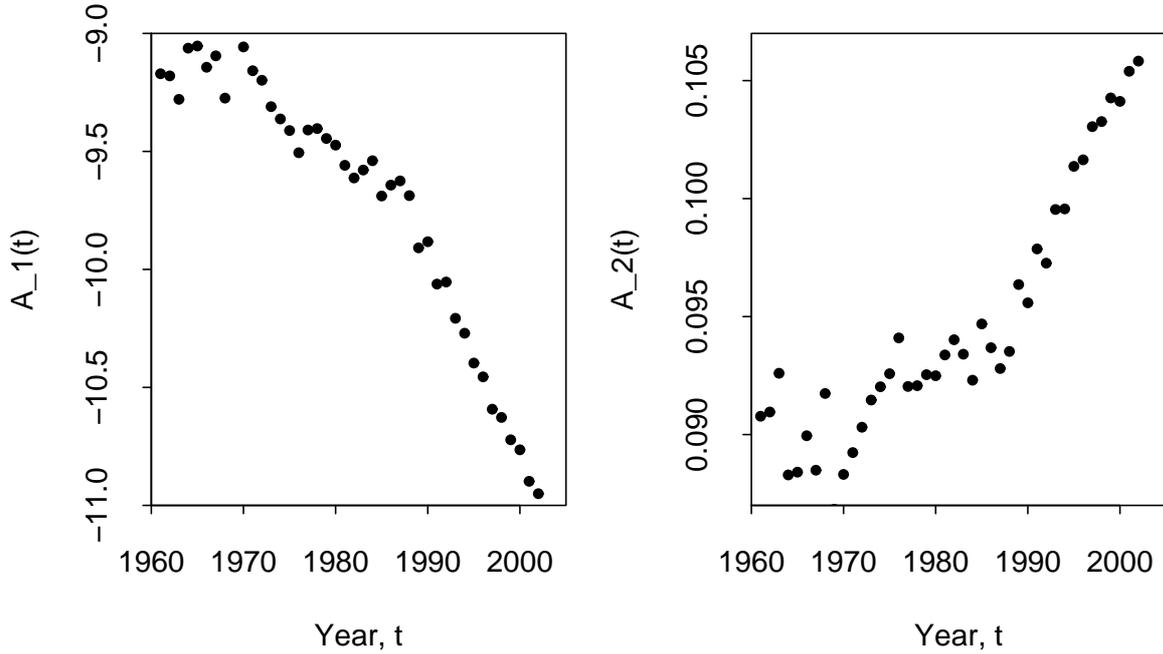


Figure 2: Estimated values of  $A_1(t)$  (left) and  $A_2(t)$  (right) in equation (1) from 1961 to 2002 for England and Wales males.

drift. Specifically

$$A(t+1) = A(t) + \mu + CZ(t+1) \quad (2)$$

where  $\mu$  is a constant  $2 \times 1$  vector,  $C$  is a constant  $2 \times 2$  upper triangular matrix<sup>7</sup> and  $Z(t)$  is a 2-dimensional standard normal random variable. If we use data from 1961 to 2002 (41 observations of the differences) we find that

$$\hat{\mu} = \begin{pmatrix} -0.043 & 4 \\ 0.000 & 367 \end{pmatrix}, \quad \text{and} \quad \hat{V} = \hat{C}\hat{C}' = \begin{pmatrix} 0.010 & 67 & -0.000 & 161 & 7 \\ -0.000 & 161 & 7 & 0.000 & 002 & 590 \end{pmatrix}. \quad (3)$$

If, on the other hand, we use data from 1982 to 2002 only (20 observations) then we find that

$$\hat{\mu} = \begin{pmatrix} -0.066 & 9 \\ 0.000 & 590 \end{pmatrix}, \quad \text{and} \quad \hat{V} = \hat{C}\hat{C}' = \begin{pmatrix} 0.006 & 11 & -0.000 & 093 & 9 \\ -0.000 & 093 & 9 & 0.000 & 001 & 509 \end{pmatrix}. \quad (4)$$

These results show a steepening of trends after 1982, with  $\mu_1$  and  $\mu_2$  both becoming larger in magnitude. They also show that the volatilities in the later period were notably smaller than in the earlier period.

<sup>7</sup>There are infinitely many matrices  $C$  that satisfy  $V = CC'$ , but the choice of  $C$  makes no difference to our analysis. Provided the entries of  $C$  are all real valued,  $CC'$  is always positive semidefinite. The restriction of  $C$  to an upper-triangular triangular form means that  $C$  is straightforward to derive from  $V$  and this (Cholesky) decomposition is unique.

An important criterion for a good mortality model (see, Cairns *et al.*, 2004, for a discussion) requires the model and its parameter values to be *biologically* reasonable. The negative value for  $\mu_1$  indicates generally improving mortality, with this improvement strengthening after 1982. The positive value for  $\mu_2$  means that mortality rates at higher ages are improving at a slower rate. Indeed, above the very high age of 113, the model predicts deteriorating mortality.<sup>8</sup> This might be perceived to be an undesirable feature of our model, but because this crossover point is at such a high age it is not felt to be a serious problem here as the number of lives involved is very low.

An additional criterion for biological reasonableness is that we should normally see, in any given year in the future, mortality rates for older cohorts that are higher than those for younger cohorts (that is, for fixed  $t$ ,  $\tilde{q}(t, x)$  should be an increasing function of  $x$ ). This criterion requires  $A_2(t)$  to remain positive. In our model  $A_2(t)$  could, theoretically, become negative, but the positive value for  $\mu_2$  and the initial value for  $A_2$  in 2002 of 0.1058 means that  $A_2(t)$  is very unlikely to do so. So the possibility of a negative  $A_2(t)$  is of little practical significance.

## 4 Simulation results for the survivor index $S(t)$

A longevity bond of the type issued by the EIB/BNP indexes coupon payments in line with a survivor index  $S(t)$  for a specified cohort of individuals.<sup>9</sup>

We now wish to determine the distribution for  $S(t)$  for the times  $t = 1, 2, \dots, 25$  that are relevant for the EIB/BNP longevity bond. Even though the functional form for  $\tilde{q}(t, x)$  is relatively simple, its distribution for  $t > 2$  is not analytically tractable, so we resort to Monte Carlo simulation and obtain the simulated  $\tilde{q}(t, x)$  and  $S(t)$  from simulations of the underlying process  $A(t)$ .

### 4.1 Results with no allowance for parameter uncertainty

In our first experiment, we simulated the  $A(t)$  according to equation (2) using estimates for  $\mu$  and  $V$  based on data from 1961-2002 and 1982-2002. These parameter estimates were treated as if they were the true parameter values, implying that, to begin with, we ignore parameter uncertainty. The results are plotted in Figure 3. We can make the following observations:

- The solid curves plot the expected values of  $S(t)$ . Measured at time 0, these represent the *ex ante* probabilities of survival from time 0 to time  $t$ ,  $p(0, 0, t, 65)$  (which we refer to as *spot survival probabilities*). The mean trajectory based on data from 1982-2002 (bottom plot) is slightly higher than that in the upper plot (based on 1961-2002 data). This is because steepening trends in  $A_1(t)$

<sup>8</sup>In other words the mortality rate at ages  $> 113$  is rising over time rather than lower.

<sup>9</sup>In the case of the EIB/BNP bond the reference cohort is the set of all England and Wales males aged 65 in 2003. The method used to calculate  $S(t)$  for this cohort is given in Appendix A.

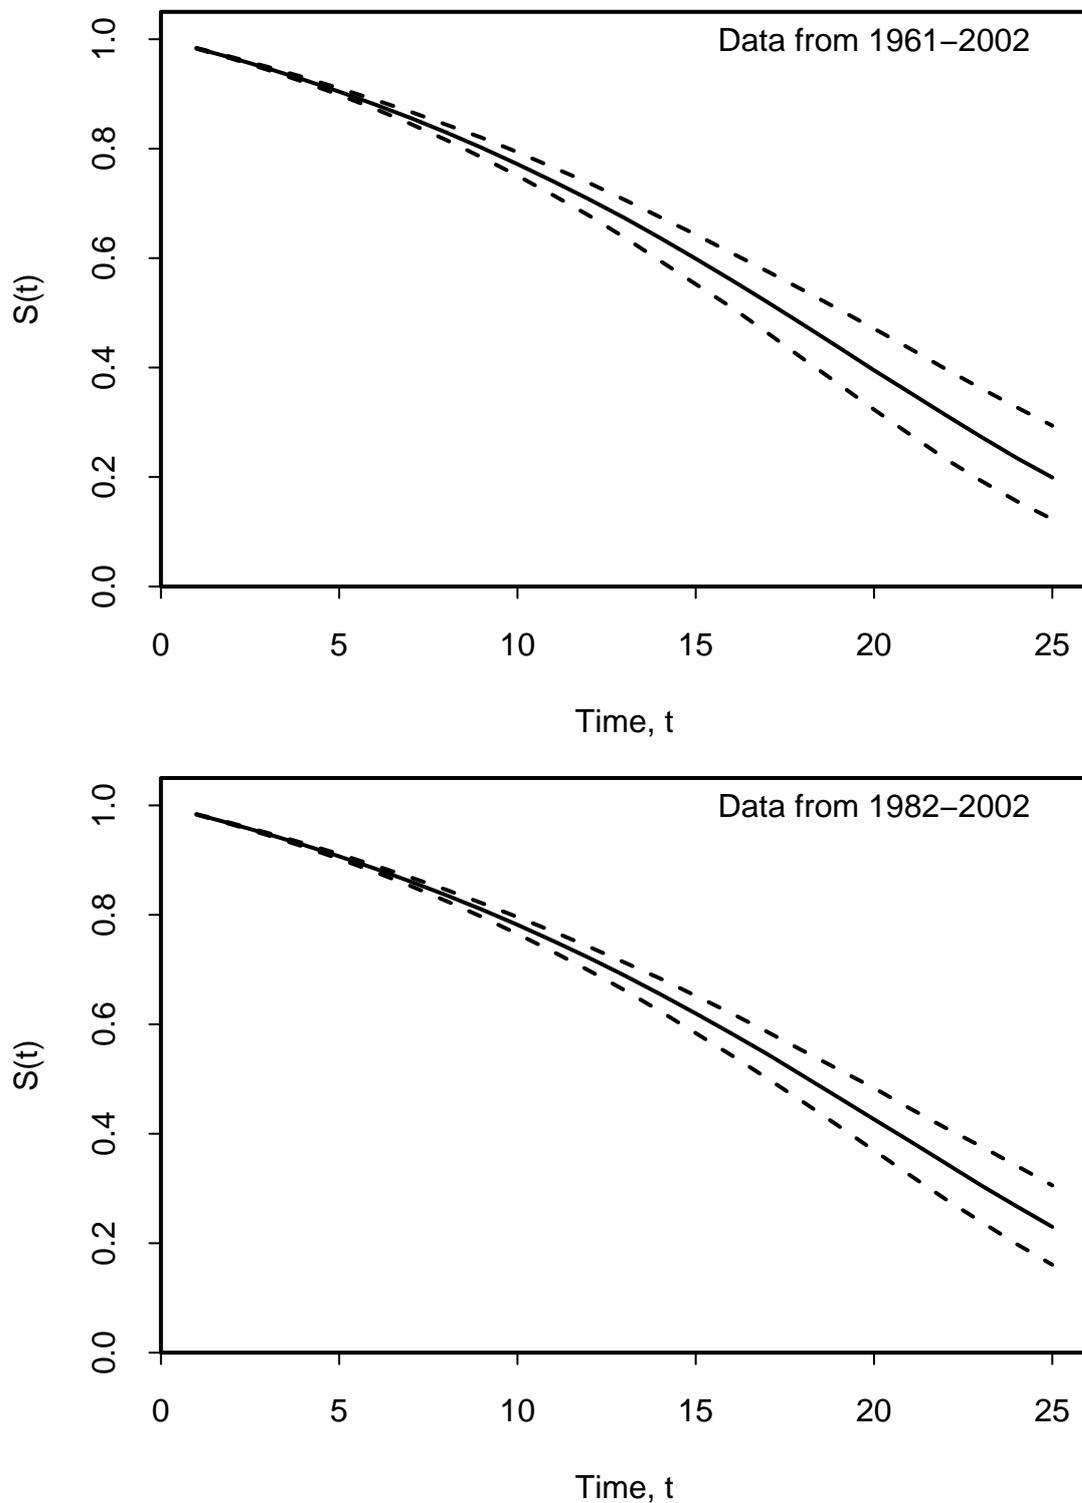


Figure 3: Mean and confidence intervals for projected survival probabilities based on data from 1961-2002 (top) or 1982-2002 (bottom). Each plot shows the mean (solid curve) and the 5th and 95th percentiles (dashed curves) of the simulated distribution of the reference index,  $S(t)$ , with no allowance for parameter uncertainty.

and  $A_2(t)$  in the 1982-2002 data (Figure 2) signal greater improvements in the future.

- The dashed curves in each plot show the 5th and the 95th percentiles of the distribution of  $S(t)$ . We can observe that the resulting 90% confidence interval is initially quite narrow but becomes quite wide by the 25-year time horizon (which is the maturity of the EIB/BNP longevity bond). We can also see that the confidence interval based on 1982-2002 data is a little narrower, reflecting the smaller values on the diagonal of  $V$ .
- The confidence interval for  $S(t)$  grows in quite a different way from, say, that associated with an investment in equities. This point is best illustrated by looking at the variance of the logarithm of  $S(t)$ , as illustrated in Figure 4. We can see that this is very low in the early years indicating that we can predict with reasonable certainty what mortality rates will be over the near future. However, after time 10 the variance starts to grow very rapidly (almost ‘exponentially’). This contrasts with equities where we would expect to see linear, rather than ‘exponential’, growth in the variance if the price process follows geometric Brownian motion.

The explanation for this variance growth is that the longer-term survival probabilities incorporate the compounding of year-by-year mortality shocks: the survival probability for year  $t$  depends on shocks applied to mortality rates in each of the years 1 to  $t$ , and each individual shock affects survival probabilities in all subsequent years.<sup>10</sup>

## 4.2 Results with parameter uncertainty

We consider next the impact of parameter uncertainty. It is clear that we have a limited amount of data and so the parameter estimates above must inevitably be subject to some degree of uncertainty. We will analyse this using standard Bayesian methods.<sup>11</sup>

Recall that we have assumed that the process  $A(t)$  is subject to i.i.d. multivariate normal shocks with mean  $\mu$  and covariance matrix  $V$ . In the absence of any clear prior beliefs about the values of  $\mu$  and  $V$  we will use a non-informative prior distribution. A common prior for the multivariate normal distribution in which both  $\mu$  and  $V$  are unknown is the Jeffreys prior

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

where  $|V|$  is the determinant of the matrix  $V$ . With this prior and with  $n$  observations  $D = \{D(1), \dots, D(n)\}$  (where  $D(t) = A(t) - A(t-1)$ ), it is known that (see,

<sup>10</sup>For further intuition, see Appendix C.

<sup>11</sup>For a general discussion of model and parameter risk, see Cairns, 2000.

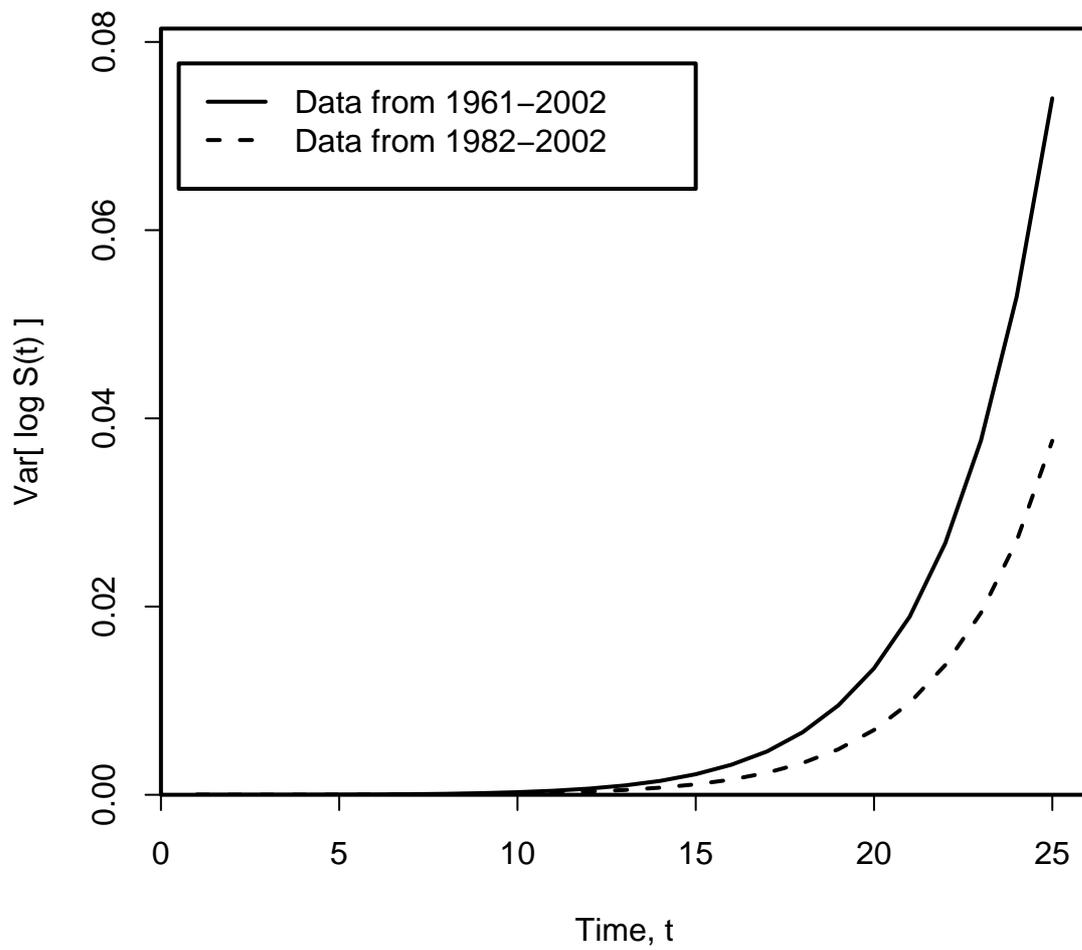


Figure 4: Plot of the variance of  $\log S(t)$  using data from 1961-2002 (solid curve) and from 1982-2002 (dashed curve), with no allowance for parameter uncertainty.

for example, Gelman *et al.*, 1995) the posterior distribution<sup>12</sup> for  $\mu, V|D$  is:

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

$$\mu|V, D \sim \text{MVN}(\hat{\mu}, n^{-1}V) \quad (6)$$

$$\text{where } \hat{\mu} = \frac{1}{n} \sum_{t=1}^n D(t)$$

$$\text{and } \hat{V} = \frac{1}{n} \sum_{t=1}^n (D(t) - \hat{\mu})(D(t) - \hat{\mu})'.$$

In what follows, we will restrict ourselves to an analysis based on data from 1982 to 2002.

For each simulated sample path of  $A(t)$  we simulate first  $\mu$  and  $V$  from the Normal-inverse-Wishart distribution and use these values for the whole of that sample path. The results of these simulations can be seen in Figures 5 and 6. In Figure 5 we can see the impact of parameter uncertainty on the confidence interval: specifically that parameter uncertainty becomes much more significant as a source of uncertainty in  $S(t)$  as  $t$  increases. We can see that 25 years ahead parameter uncertainty accounts for about half of the uncertainty in  $S(t)$ <sup>13</sup>. In Figure 6 we plot the variance of  $\log S(t)$ , and the use here of a log scale allows us to see clearly that for smaller values of  $t$  parameter uncertainty is much less important (that is the difference between the two curves is quite small).

## 5 The price of longevity risk

Now consider the price that a life office or pension fund might be prepared to pay to lay off its exposure to longevity risk. From Figures 3 to 6 we can infer that if premiums are to be paid in respect of each future year, the premium will be much larger for the 25-year payment than, say, the 10-year payment. Furthermore, a reasonable proportion of this premium might be in respect of the desire to eliminate exposure to parameter uncertainty.

### 5.1 Pricing using risk-adjusted probability measures

We propose to specify the dynamics under a risk-adjusted *pricing measure*  $Q$  that is equivalent to, in the probabilistic sense, the current real-world measure (which we shall refer to as  $P$ ).<sup>14</sup> The measure  $Q$  is also commonly referred to as the *risk-neutral*

<sup>12</sup>The Wishart distribution is a multivariate version of the Gamma or Chi-squared distribution. For details on how to simulate the joint Normal-Inverse-Wishart distribution see Appendix B.

<sup>13</sup>That is, at  $t = 25$  the size of the gap on the log scale between the solid and dashed lines equates to a ratio of about 2 between the variances.

<sup>14</sup>An alternative way of generating risk-adjusted measures is to use the Wang transform (Wang, 2000, 2002, 2003). These distort the distributions of each of the  $S(t)$  random variables. Examples of its application to longevity bonds and other mortality-linked securities include Lin and Cox (2004), Denuit, Devolder and Goderniaux (2004), Cox and Wang (2005), and Dowd *et al.* (2005).

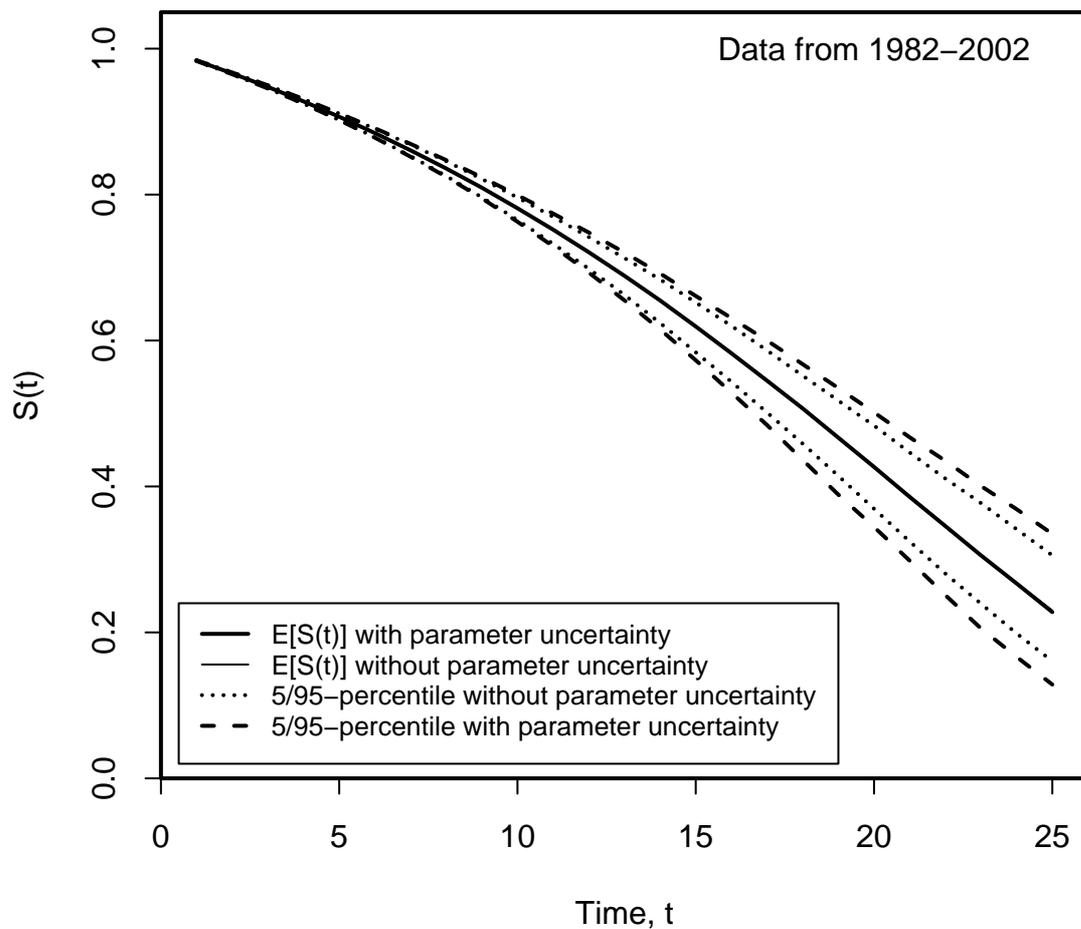


Figure 5: Confidence intervals for projected survival probabilities based on data from 1982-2002. Confidence intervals are shown excluding parameter uncertainty (thin dotted curves) and including parameter uncertainty (thick dashed curves). The mean trajectories (thin and thick solid curves) for the two cases are overlapping.

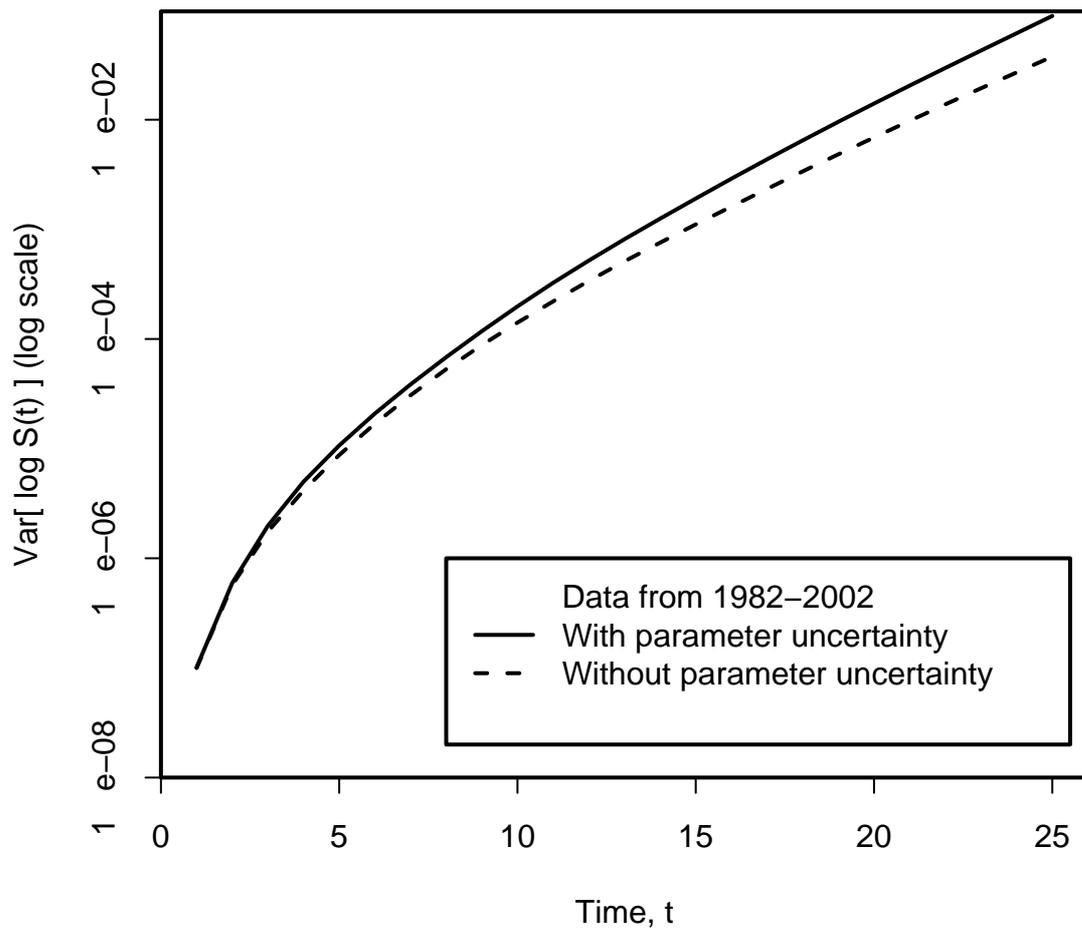


Figure 6: Plot of the variance (on a log scale) of  $\log S(t)$  using data from 1982-2002. The variance has been calculated excluding parameter uncertainty (dashed curves) and including parameter uncertainty (solid curves).

measure<sup>15</sup> or as an *equivalent-martingale* measure.

Recall that we worked earlier with the following dynamics under  $P$ :

$$A(t+1) = A(t) + \mu + CZ(t+1)$$

where  $Z(t+1)$  is a standard 2-dimensional normal random variable under  $P$ .

Under the risk-adjusted measure  $Q(\lambda)$  we propose<sup>16</sup> that:

$$\begin{aligned} A(t+1) &= A(t) + \mu + C(\tilde{Z}(t+1) - \lambda) \\ &= A(t) + \tilde{\mu} + C\tilde{Z}(t+1) \end{aligned}$$

where  $\tilde{\mu} = \mu - C\lambda$ .

In this equation  $\tilde{Z}(t+1)$  is a standard 2-dimensional normal random variable<sup>17</sup> under  $Q$ . The vector  $\lambda = (\lambda_1, \lambda_2)$  represents the *market prices of longevity risk* associated with the processes  $Z_1(t)$  and  $Z_2(t)$  respectively. Under the chosen decomposition for the matrix  $C$  (upper triangular),  $\lambda_1$  is associated with level shifts in mortality (specifically  $\log q/p$ ), while  $\lambda_2$  is associated with a tilt in  $\log q/p$ . We assume (as part of our model) that  $\lambda$  is constant rather than time dependent: indeed it is difficult to propose anything more sophisticated for  $\lambda$  in the absence of any market price data.

We can make the following points about  $Q(\lambda)$ :

- Complete market models such as the Black-Scholes option-pricing model force upon us a unique choice of measure  $Q$ . In contrast, here we have an incomplete market, and a range of possibilities for  $Q(\lambda)$ .
- If there exists some form of market in mortality-linked securities then the choice of  $Q(\lambda)$  needs to be consistent with this (limited) market information (so that theoretical prices under  $Q(\lambda)$  match observed market prices).
- Beyond these restrictions, the choice of  $Q(\lambda)$  becomes a modelling assumption. Thus, here we have postulated that the market price of risk,  $\lambda$ , might be constant over time (in the same way that the market price of risk is normally assumed to be constant in the Black-Scholes model).<sup>18</sup>

<sup>15</sup>In an incomplete market, the term *risk-neutral* is vague, but is used to convey the point that expected returns over the short term under  $Q$  are equal to the short-term risk-free rate of interest. At the present time we are very far from having a complete market in which all contingent claims can be replicated using dynamic hedging strategies. This means that the risk-adjusted measure  $Q$  is not unique. Rather the choice of which measure  $Q$  to use becomes one of the modelling assumptions.

<sup>16</sup>Modelling in discrete time means that there are infinitely-many equivalent measures with different means, variances and covariances for  $Z(t+1)$ . However, we choose to restrict ourselves to ones that have a constant market price of risk, that preserve the variance-covariance structure of  $Z(t+1)$ , and that preserve the assumption of bivariate normality. The latter assumptions lead to consistency between the discrete-time model and the continuous-time diffusion model.

<sup>17</sup>That is,  $\tilde{Z}(t+1) = (\tilde{Z}_1(t+1), \tilde{Z}_2(t+1))'$  where  $\tilde{Z}_1(1), \tilde{Z}_1(2), \dots$  and  $\tilde{Z}_2(1), \tilde{Z}_2(2), \dots$  are independent sequences of i.i.d.  $N(0, 1)$  random variables under  $Q(\lambda)$ .

<sup>18</sup>In the case of both the present model and the Black-Scholes model, it would seem appropriate that the market price be allowed to vary in a stochastic fashion over the long term, 25 years,

## 5.2 Example: the EIB/BNP longevity bond

As an example, consider the 25-year longevity bond announced by the European Investment Bank (EIB) in November 2004 with an issue price based on a yield of 35 basis points below LIBOR. The appropriate starting point is the EIB curve for conventional fixed-interest bonds issued typically at 15 basis points below LIBOR. This means that the new longevity bond was priced at 20 basis points below standard EIB rates. This spread below standard EIB rates will be denoted by  $\delta$  in the equations that follow. We will now make the following assumptions:

1. The projected survival rates used in the pricing of the bond (in the case of the EIB bond this is the projection made by the UK Government Actuary's Department) are unbiased estimates at time 0 under the real-world measure  $P$  of the survival rates.
2. The spread of 20 basis points below the standard EIB curve is accounted for entirely by the market price of longevity risk.
3. The development of mortality rates over time is independent of the dynamics of the interest-rate term structure over time.<sup>19</sup>

We will refer to  $\hat{S}(0, T)$  as the survivor index based upon the latest GAD projections available at time 0.<sup>20</sup> Assumption 1 implies that  $\hat{S}(0, T) = E_P[S(T)|\mathcal{M}_0]$ .

Next, let us refer to  $P(0, T)$  as the price at time 0 of a fixed-interest zero-coupon bond that pays 1 at time  $T$  issued by the EIB. The basis declared by the EIB and BNP for the initial price of the bond was  $V(0) = \sum_{T=1}^{25} P(0, T)e^{\delta T}\hat{S}(0, T)$ , where  $\delta$  is the spread (expressed as a continuously compounding rate) below the EIB curve used in pricing the bond.<sup>21</sup> Given assumption 1 this is equivalent to<sup>22</sup>

$$V(0) = \sum_{T=1}^{25} P(0, T)e^{\delta T} E_P[S(T)|\mathcal{M}_0]. \quad (7)$$

of the contract. However, in the equity-modelling literature there seems to be little consensus on the dynamics of the market price of risk. If we combine this observation with the absence of any historical market data on mortality-linked securities we conclude that it is inappropriate to attempt to model the market-price of risk as a dynamic process.

<sup>19</sup>This is a very useful simplifying assumption which we believe to be a reasonable one for relatively short horizons under normal conditions. However, we recognise that over the very long run the term structure of interest rates will be influenced by the relative size of the capital stock to that of the population and the latter might be influenced by mortality (as well as fertility) dynamics. Also in the short run, we recognise that a catastrophe that affects the size of the population (such as nuclear war) will also affect interest rates.

<sup>20</sup>Values for the  $\hat{S}(0, T)$  are specified in the offer document issued by BNP Paribas.

<sup>21</sup>We do not know what the pricing convention for the bond is after issue. However, it seems plausible that it will also be of the form (for integer  $t$ )  $V(t) = \sum_{T=t+1}^{25} P(t, T)e^{\delta(t)(T-t)}\hat{S}(0, T)$ : that is, still with reference to the *initial* estimate  $\hat{S}(0, T)$ , and with reference to an easily-observable zero-coupon curve at time  $t$ .

<sup>22</sup>Strictly, the  $P(0, T)$  in equation (7) should be LIBOR-implied discount factors,  $P_L(0, T)$ , in combination with  $\delta = 0.0035$  as stated in the contract, while the  $P(0, T)$  in equation (8) should be the EIB-implied discount factors,  $P_E(0, T)$ . However, as stated at the beginning of Section 5.2, the approximate relationship between the two is  $P_E(0, T) = P_L(0, T)e^{0.0015T}$ , which leads us to the given form in equation (7) with  $\delta = 0.0020$ .

Values for  $E_P[S(T)|\mathcal{M}_0]$  based upon parameters in (4) and without parameter uncertainty are given in Table 1, column 1. For example, if we assume that implied EIB zero-coupon prices are given by  $P(0, T) = 1.04^{-T}$  and if we set  $\delta = 0.0020$  and  $\lambda = (0, 0)'$  then we find that the price at issue of the bond on this contractual basis (equation 7) is  $V(0) = 11.442$ .

The risk-adjusted approach to pricing assumes that (exploiting assumption 3 above)

$$V_\lambda(0) = \sum_{T=1}^{25} P(0, T) E_{Q(\lambda)}[S(T)|\mathcal{M}_0]. \quad (8)$$

A comparison of equations (7) and (8) shows that  $\delta$  can be interpreted as an average risk premium per annum. We shall see later (Section 6) that this risk premium will depend upon the term of the bond and on the initial age of the cohort being tracked.

We can now ask the question: *what values for the market prices of risk  $\lambda_1$  and  $\lambda_2$  satisfy  $V_\lambda(0) = V(0)$ ?* Put another way, under what circumstances does the risk-adjusted price (equation 8) match the issue price quoted in the contract (equation 7)?

With no parameter uncertainty, and  $\delta = 0$  we found that we could obtain  $V_\lambda(0) = 11.438$  with  $(\lambda_1, \lambda_2) = (0.375, 0)$  and  $(0, 0.316)$ . For these two values for  $\lambda$  the values for  $E_{Q(\lambda)}[S(t)|\mathcal{M}_0]$  are given in Table 1 columns 3 and 4. In column 5, we have also given an intermediate value for  $\lambda$  between these two extremes.<sup>23</sup> Here we can achieve  $V_\lambda(0) = V(0)$  with  $\lambda_1 = \lambda_2 = 0.175$ .<sup>24 25</sup>

We next introduce parameter uncertainty into the analysis. We first simulate under  $P$  with full parameter uncertainty and the values for  $E_{Q(\lambda)}[S(t)|\mathcal{M}_0]$  are given in Table 1, in column 2 with  $\lambda = (0, 0, 0, 0)'$ . We have seen in Figure 6 that parameter uncertainty presents a significant risk to annuity providers. It follows that they will be prepared to pay a premium to reduce this risk in the same way that they are prepared to pay to reduce the impact of longevity risk.

In this analysis we concentrate on introducing a market price of risk for the mean  $\mu$ . In the case of no parameter uncertainty we assume that  $\mu = \hat{\mu}$ . With parameter uncertainty, the basic model (see Appendix B) simulates first  $V$ , then calculates the upper-triangular matrix  $C$  that satisfies  $V = CC'$ , and then simulates

$$\mu = \hat{\mu} + n^{-1/2} CZ_\mu$$

where  $Z_\mu$  is a standard bivariate normal random variable. Once again we have to identify possible equivalent measures. We propose here a similar restriction that,

<sup>23</sup>This can be found by fixing first the value for  $\lambda_1$  and then solving for  $\lambda_2$ .

<sup>24</sup>In fact, the set of values for  $(\lambda_1, \lambda_2)$  that gives a price of 11.438 is approximately linear running from  $(0.375, 0)$  to  $(0, 0.316)$ . A straight line between the two end points would pass through  $(0.171, 0.171)$  rather than  $(0.175, 0.175)$ .

<sup>25</sup>It is tempting to think that  $(\lambda_1, \lambda_2) = (0.375, 0)$  and  $(0, 0.316)$  represent the extreme values for the market price of longevity risk. This would be true for  $(\lambda_1, \lambda_2) = (0, 0.316)$  if the demand for such assets is coming from annuity providers. However, if the market is dominated by life offices hedging longevity or mortality risk in their term-assurance portfolios then  $\lambda_1$  might, in fact, be negative. Similarly,  $\lambda_2$  might be negative if longevity risk at ages below 60 present the greatest risk to annuity providers.

Column reference:	1	2	3	4	5	6	7
Parameter uncertainty:	N	Y	N	N	N	Y	Y
$\lambda_1$	0	0	0.375	0	0.175	0	0
$\lambda_2$	0	0	0	0.316	0.175	0	0
$\lambda_3$	–	0	–	–	–	1.684	0
$\lambda_4$	–	0	–	–	–	0	1.419
$t$	$E_{Q(\lambda)}[S(t) \mathcal{M}_0]$						
1	0.9836	0.9836	0.9837	0.9836	0.9836	0.9837	0.9836
2	0.9661	0.9661	0.9664	0.9662	0.9663	0.9664	0.9662
3	0.9475	0.9475	0.9482	0.9477	0.9479	0.9482	0.9477
4	0.9278	0.9278	0.9289	0.9281	0.9285	0.9289	0.9281
5	0.9068	0.9068	0.9086	0.9074	0.908	0.9086	0.9074
6	0.8845	0.8845	0.8872	0.8856	0.8863	0.8872	0.8856
7	0.861	0.8609	0.8646	0.8626	0.8635	0.8646	0.8626
8	0.836	0.8359	0.8408	0.8384	0.8395	0.8407	0.8383
9	0.8095	0.8095	0.8157	0.8129	0.8142	0.8156	0.8129
10	0.7816	0.7815	0.7893	0.7862	0.7877	0.7892	0.7861
11	0.7522	0.752	0.7616	0.7583	0.7599	0.7615	0.7582
12	0.7213	0.721	0.7326	0.7292	0.7308	0.7325	0.729
13	0.6888	0.6885	0.7023	0.6989	0.7004	0.7021	0.6987
14	0.6548	0.6545	0.6707	0.6675	0.6689	0.6704	0.6672
15	0.6195	0.6191	0.6378	0.635	0.6362	0.6374	0.6346
16	0.5828	0.5823	0.6036	0.6015	0.6024	0.6032	0.6011
17	0.5448	0.5443	0.5684	0.5672	0.5676	0.5679	0.5667
18	0.5059	0.5052	0.5321	0.5321	0.532	0.5315	0.5316
19	0.4661	0.4654	0.495	0.4965	0.4957	0.4944	0.4959
20	0.4258	0.4251	0.4573	0.4606	0.459	0.4566	0.4599
21	0.3853	0.3847	0.4191	0.4245	0.422	0.4185	0.4238
22	0.345	0.3445	0.3809	0.3885	0.3851	0.3803	0.3879
23	0.3054	0.305	0.3428	0.353	0.3486	0.3424	0.3524
24	0.2667	0.2668	0.3054	0.318	0.3128	0.3052	0.3177
25	0.2297	0.2302	0.2689	0.2841	0.278	0.269	0.284
Price							
$\delta = 0$	11.240	11.237	<b>11.442</b>	<b>11.442</b>	<b>11.442</b>	<u>11.439</u>	<u>11.439</u>
$\delta = 0.0020$	<b>11.442</b>	<u>11.439</u>	–	–	–	–	

Table 1: Longevity bond expected cashflows under the risk-neutral measure  $Q(\lambda)$ :  $E_{Q(\lambda)}[S(t)|\mathcal{M}_0]$  and market prices under various assumptions for the market prices of longevity and parameter risk.

under the equivalent measure,  $Z_\mu$  is still bivariate normal with unit variances but a shifted mean. Thus we now simulate  $\mu$  with reference to two additional market prices of risk  $\lambda_3$  and  $\lambda_4$ :

$$\mu = \hat{\mu} + n^{-1/2}C(\tilde{Z}_\mu - \lambda_\mu) = \tilde{\mu} + n^{-1/2}C\tilde{Z}_\mu$$

where  $\lambda_\mu = (\lambda_3, \lambda_4)'$  and  $\tilde{\mu} = \hat{\mu} - n^{-1/2}C\lambda_\mu$ .

We now have four market prices of risk to play with to match the single price derived by discounting expected cashflows under  $P$  at EIB minus 20 basis points. With parameter uncertainty included, the expected cashflows under  $P$  change very slightly (see Table 1, column 2), as does the price of  $V(0) = 11.439$ . The values for  $\lambda_1$  and  $\lambda_2$  required to match this price are essentially unchanged from the values that were determined before (Table 1 columns 3 and 4) and are consequently not repeated in the table. The required values for  $\lambda_3$  and  $\lambda_4$  were, respectively, 1.684 and 1.419, with the corresponding values for  $E_{Q(\lambda)}[S(t)|\mathcal{M}_0]$  quoted in Table 1, columns 6 and 7.

For the various cases presented in Table 1 we have plotted in Figure 7 the expected value under  $P$  or  $Q(\lambda)$  of  $S(t)$  for  $t = 1, \dots, 25$ . This plot helps us to analyse the impact of using the different measures and, in particular, to see where most of the *additional* value in the longevity bond resides. The expected values in the upper plot show us two things. First, the inclusion of parameter uncertainty has almost no effect on the expected values under  $P$ . Second, the expected values under the different  $Q(\lambda)$  measures look similar, and all show up the largest differences compared with the  $P$  measure near  $t = 25$ .

The lower plot in Figure 7 allows us to differentiate more easily between the different  $Q(\lambda)$  measures. Here we have plotted the expected risk premium per annum on a zero-coupon longevity bond that is held from time 0 through to maturity. This is calculated by converting the ratio of two expected values into an additional rate of return per annum as follows:

$$\frac{1}{t} \log \left( \frac{E_{Q(\lambda)}[S(t)|\mathcal{M}_0]}{E_P[S(t)|\mathcal{M}_0]} \right).$$

From this lower plot we can see that the level of the risk premium depends to some extent on the choice of  $Q(\lambda)$ . However, we can note from both Figure 7 and Table 1 that cases 3 and 6 produce very similar results and that cases 4 and 7 also produce very similar results.<sup>26</sup> The reason for these similarities is not, in fact, difficult to explain. Recall that we have

$$\begin{aligned} \mu &= \hat{\mu} - n^{-1/2}C\lambda_\mu + n^{-1/2}C\tilde{Z}_\mu \\ \text{and } A(t+1) &= A(t) + \mu - C\lambda + C\tilde{Z}(t+1) \\ &= A(t) + \hat{\mu} - n^{-1/2}C\lambda_\mu - C\lambda + n^{-1/2}C\tilde{Z}_\mu + C\tilde{Z}(t+1) \\ &= A(t) + \hat{\mu} - C(n^{-1/2}\lambda_\mu + \lambda) + n^{-1/2}C\tilde{Z}_\mu + C\tilde{Z}(t+1). \end{aligned}$$

Thus, relative to simulation under  $P$  with parameter uncertainty, when we simulate under  $Q(\lambda)$ ,  $\lambda_1$  will have the same effect as  $n^{-1/2}\lambda_3$  and  $\lambda_2$  will have the same effect

<sup>26</sup>If we repeat cases 3 and 4 incorporating parameter uncertainty then the remaining small differences between the cases 3 and 6 and between cases 4 and 7 essentially disappear.

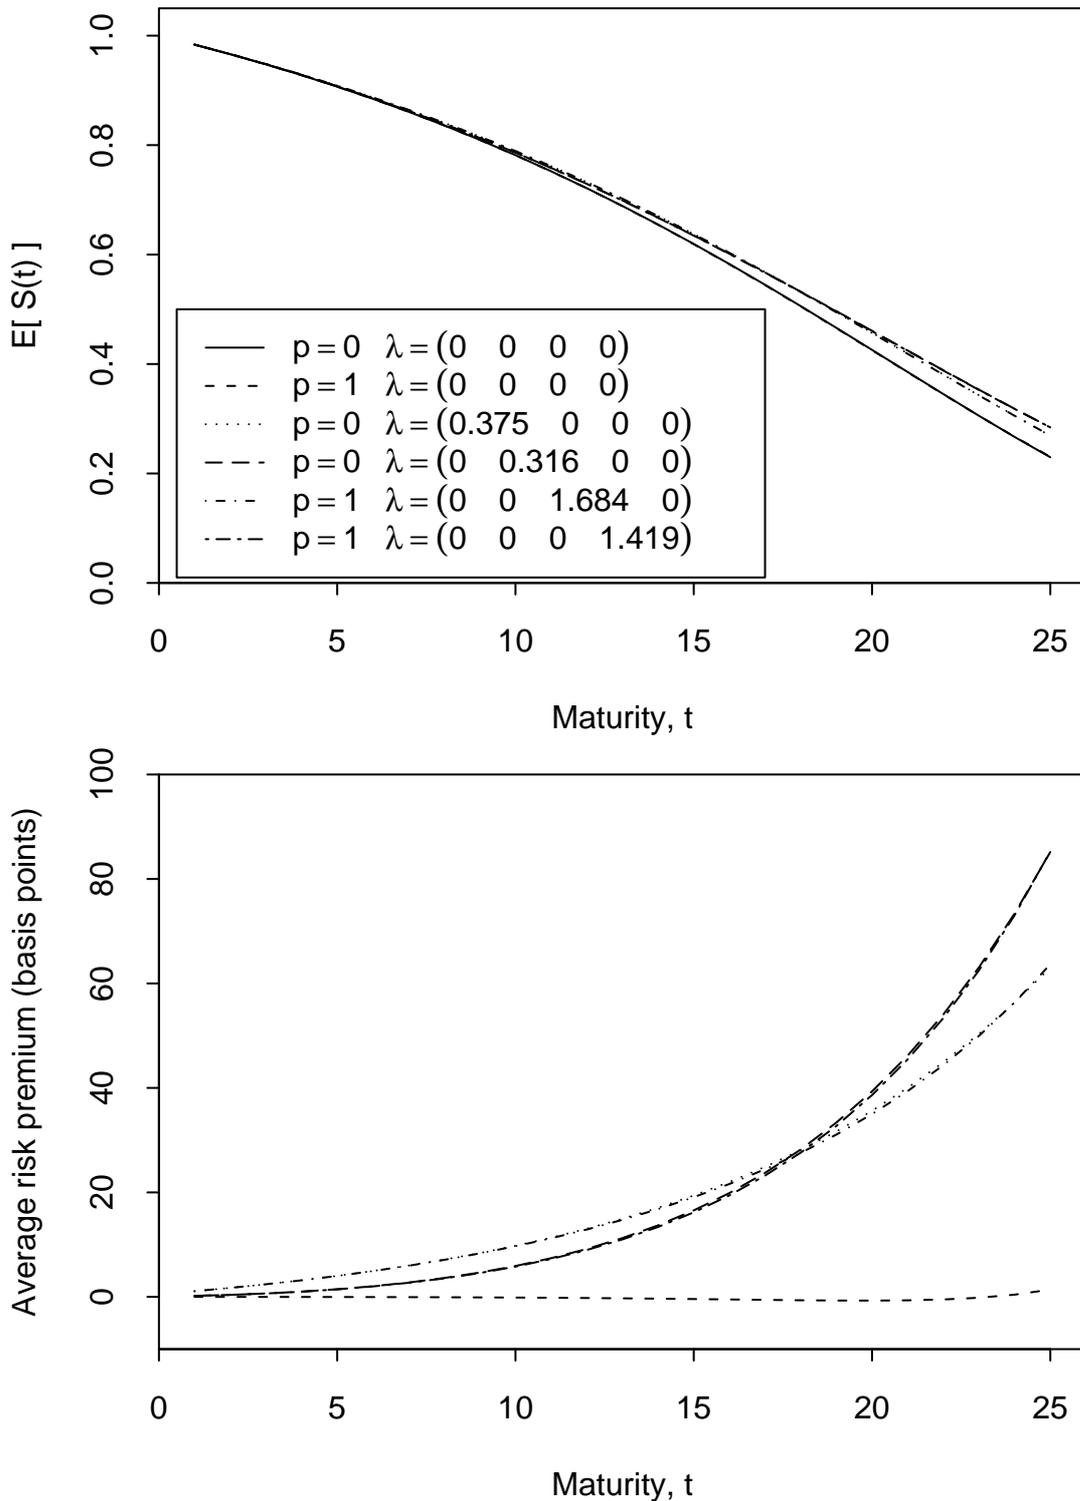


Figure 7: Top: Expected value of  $S(t)$  under different probability measures. Values for the market prices of risk in the six cases considered are given in the legend;  $p = 0$  means without parameter uncertainty and  $p = 1$  means with parameter uncertainty in both  $\mu$  and  $V$ . Bottom: Average risk premium per annum is defined as  $\log\{E_{Q(\lambda)}[S(t)]/E_P[S(t)]\}/t$  on a zero-coupon longevity bond over the full term to maturity. Different line types are defined in the top plot.

as  $n^{-1/2}\lambda_4$ . We can check this by comparing the values of  $\lambda_1$  to  $\lambda_4$  in Table 1. Projections have been made on the basis of  $n = 20$  observations of  $A(t+1) - A(t)$ , and we can see that the ratios of  $\lambda_3$  to  $\lambda_1$  and  $\lambda_4$  to  $\lambda_2$  are both close to  $\sqrt{20}$  as predicted.

Now return to the results presented in Table 1. Why are the required values of  $\lambda_1$  to  $\lambda_4$  positive? And why does the average risk premium per annum plotted in Figure 7 differ in the way that it does for  $\lambda_1$  and  $\lambda_2$  (curves A and B)?

The first thing to do is to analyse the impact on the mortality curve of changes in  $A_1(t)$  and  $A_2(t)$ . Recall that we have

$$\frac{\tilde{q}(t, x)}{\tilde{p}(t, x)} = e^{A_1(t) + A_2(t)(x+t)}$$

Now if we replace  $A_1(t) + A_2(t)(x+t)$  by  $\bar{A}_1(t) + \bar{A}_2(t)(x+t-x_0)$  where  $\bar{A}_1(t) = A_1(t) + A_2(t)x_0$  and  $\bar{A}_2(t) = A_2(t)$ , then for a suitable choice of  $x_0$  (specifically,  $x_0 = \hat{V}_{21}/\hat{V}_{22} \approx 62.2$ ) the processes  $\bar{A}_1(t)$  and  $\bar{A}_2(t)$  become independent random walks with drift.  $\bar{A}_1(t)$  has  $Z_1(t)$  as its driver with market price of risk  $\lambda_1$  and  $\bar{A}_2(t)$  has  $Z_2(t)$  as its driver with market price of risk  $\lambda_2$ . Under this transformation we have

$$\begin{aligned} \bar{A}(t+1) &= \bar{A}(t) + \bar{\mu} + \bar{C}Z(t) \\ \text{where } \bar{\mu} &= \begin{pmatrix} -0.0302 \\ 0.000590 \end{pmatrix} \\ \text{and } \bar{C} &= \begin{pmatrix} 0.01645 & 0 \\ 0 & 0.001229 \end{pmatrix}. \end{aligned}$$

We can now see that, since  $\bar{C}$  is positive, a positive shock  $Z_1(t)$  will produce a level shift in  $\tilde{q}(t, x)/\tilde{p}(t, x)$  over all ages  $x$ : that is, an unanticipated deterioration in longevity. A positive value of  $\lambda_1$ , in contrast, causes  $\bar{A}_1(t)$  to be pushed downwards over time thereby *enhancing* improvements in longevity. So  $\lambda_1 > 0$  is required to produce a positive risk premium (that is, higher expected values of  $S(t)$  under  $Q(\lambda)$ ).

We have to be slightly more careful when we analyse the impact of positive shocks  $Z_2(t)$ . Specifically, for  $x+t$  above age  $x_0 = 62.2$ , a positive value for  $Z_2(t)$  will increase  $\tilde{q}(t, x)/\tilde{p}(t, x)$  (particularly so at high ages). However, the same positive value for  $Z_2(t)$  will cause  $\tilde{q}(t, x)/\tilde{p}(t, x)$  to fall for values of  $x+t$  less than 62.2. Now in our analysis we are considering a cohort who are all aged 65 at time 0, so that  $S(T)$  is constructed from the experienced mortality rates  $\tilde{q}(0, 65), \tilde{q}(1, 65), \dots, \tilde{q}(T-1, 65)$ . Since the minimum age is 65 a positive shock in  $Z_2(t)$  will cause an increase in each of  $\tilde{q}(t-1, 65), \tilde{q}(t, 65), \dots, \tilde{q}(T-1, 65)$ , everything else being equal. Thus, we infer that  $\lambda_2$  must also be positive to produce a positive risk premium.

This discussion also helps us to explain the difference between the curves corresponding to  $(\lambda_1, \lambda_2) = (0.375, 0)$  and  $(\lambda_1, \lambda_2) = (0, 0.316)$  in the lower half of Figure 7 (curves B and A respectively). Specifically risk adjustments to the dynamics of  $A_2(t)$  through the use of  $\lambda_2$  have proportionately a much greater effect on higher-age mortality than adjustments to  $A_1(t)$  through  $\lambda_1$ . This means that the impact on the survivorship probability to higher ages is much more sensitive to  $\lambda_2$  than to  $\lambda_1$ . Thus

we see that curve A in Figure 7 corresponding to  $\lambda_2$  is flatter than curve B initially but then picks up at a much faster rate, ending up at a higher level.

## 6 The risk premium on new issues

The announcement in 2004 of the 25-year EIB longevity bond will, we hope, be followed by other issues with different maturity dates and which will follow different cohorts.

Recall that the 25-year bond following the age-65 cohort (we will refer to this as the  $(T = 25, x = 65)$  bond), had a 20 basis-point risk premium per annum. The question now is: *what risk premiums are appropriate for bonds with different terms to maturity or that follow older or younger cohorts?* It is important to address this question to be sure that possible future bonds are priced in a consistent fashion.

This question can be answered in a relatively straightforward fashion. The key is that the market prices of risk  $\lambda_1$  and  $\lambda_2$  used in pricing the  $(T, x)$  bond must be the same as those used in pricing the  $(25, 65)$  bond. Thus for each  $(T, x)$  we calculate the price of the bond by determining expectations under  $Q(\lambda)$  and then discounting at EIB rates as before. We then calculate the price of the bond using expectations under  $P$  but then discounting at EIB rates minus the risk premium  $\delta$  as in equation (7). We then need to find the value of  $\delta$  that equates the two prices under  $P$  and  $Q(\lambda)$ .<sup>27</sup>

Recall that the only longevity bond so far in existence does not allow us to determine  $\lambda$  uniquely. Instead, for any given proposed bond  $(T, x)$ , the risk premium  $\delta(T, x, \lambda)$  will depend on  $\lambda$ .

Risk premia on  $(T, x)$  bonds are given in Tables 2, 3 and 4.<sup>28</sup> We can make the following observations:

- In each table we see that older cohorts attract a higher risk premium. As we take younger and younger cohorts, the mortality rates get closer to zero, so even if we introduce a market price of risk, the probability of survival will still be close to 1. In contrast at higher ages, the market price of longevity risk will have a proportionally greater impact on the survival probability. These differences between ages 60, 65 and 70 are illustrated by the  $Var[\log S(t)]$  plot in Figure 8. We can see that the longevity risk for the age 60 cohort is much lower than the age 65 and 70 cohorts. Consequently a lower risk premium is appropriate.

<sup>27</sup>We have not made any allowance in these calculations for parameter risk. We have commented earlier that the impact of this is minimal for the  $(25, 65)$  bond priced with a 20 basis-point risk premium.

<sup>28</sup>We concentrate on 20, 25 and 30-year bonds, but, for completeness we have included infinite-maturity longevity bonds (which Blake and Burrows, 2001, called Survivor Bonds). However, we note the practical difficulties associated with such bonds in dealing with the small numbers of survivors at very high ages as well as lack of reliability in mortality statistics at these ages.

		Initial age of cohort, $x$		
		60	65	70
Bond	20	8.9	14.7	23.1
Maturity	25	12.7	<b>20.0</b>	28.7
$T$	30	16.9	24.3	31.5
	$\infty$	22.9	27.2	32.2

Table 2: Longevity bond risk premium in basis points per annum as a function of term to maturity and initial age of cohort. Market price of longevity risk assumed to be  $\lambda = (0.375, 0)$ .

		Initial age of cohort, $x$		
		60	65	70
Bond	20	4.8	12.4	26.1
Maturity	25	9.2	<b>20.0</b>	36.1
$T$	30	15.0	27.6	42.3
	$\infty$	27.1	34.8	44.7

Table 3: Longevity bond risk premium in basis points per annum as a function of term to maturity and initial age of cohort. Market price of longevity risk assumed to be  $\lambda = (0, 0.316)$ .

		Initial age of cohort, $x$		
		60	65	70
Bond	20	6.8	13.4	25.1
Maturity	25	11.0	<b>20.0</b>	33.3
$T$	30	16.2	26.6	37.9
	$\infty$	25.5	33.7	39.6

Table 4: Longevity bond risk premium in basis points per annum as a function of term to maturity and initial age of cohort. Market price of longevity risk assumed to be  $\lambda = (0.175, 0.175)$ .

		Initial age of cohort, $x$		
		60	65	70
		$e_P(x, T)$		
Maximum	20	16.95	15.15	12.74
Years	25	19.59	16.78	13.45
$T$	30	21.30	17.53	13.64
	$\infty$	22.43	17.79	13.66
		$e_{Q(\lambda)}(x, T) - e_P(x, T)$		
Maximum	20	0.12	0.20	0.28
Years	25	0.28	0.40	0.47
$T$	30	0.54	0.65	0.60
	$\infty$	1.22	1.02	0.66

Table 5: Effect of the change of measure from  $P$  to  $Q(\lambda)$  for  $\lambda = (0.175, 0.175)$  on expected future, truncated lifetimes:  $e(x, T) = \int_0^T E[S(t)]dt$ . Upper part of table shows  $e(x, T)$  under the real-world measure  $P$  and the lower part of the table shows the increase in  $e(x, T)$  when we change to the risk-neutral measure  $Q(\lambda)$ .

- In each table we see that the longer the maturity of the bond, the greater the risk premium. This reflects our earlier observations (for example, Figure 7, bottom) that longer-dated cashflows have a higher risk premium per annum.
- In each table consider the diagonal running from cohort 60, term 30 up to cohort 70, term 20. In each case the terminal age is 90. As we move up the diagonal, there are two conflicting trends influencing the risk premium. The shortening maturity serves to push the risk premium down,<sup>29</sup> while the increasing initial age serves to push up the risk premium up. However, we can see from the Table that the latter trend dominates and the risk premium increases as we move up the diagonal.

Compare, for example, one cohort currently aged 60 with another currently aged 70 and consider the contracted cashflows at age 90. This cashflow is clearly subject to greater uncertainty for the age-60 cohort. However, the observation above indicates that the overall impact of this greater uncertainty on the 30-year longevity bond is much reduced by the effect of discounting.

- The risk premium  $\delta(T, x, \lambda)$  varies most with  $(T, x)$  when  $\lambda = (0, 0.315)$  (Table 3). The greater variation with  $T$  reflects the development of the risk premium illustrated in Figure 7, bottom. The greater variation with  $x$  reflects the fact that  $Z_2(t)$  affects mortality rates in different ways at different ages. The market price of risk  $\lambda_2$  has a positive effect on mortality at higher ages and a negative effect at lower ages.
- From Table 4 with the intermediate  $\lambda = (0.175, 0.175)$  we see that the risk premia lie between those given in Tables 2 and 3.

<sup>29</sup>See Table 5 for this trend.

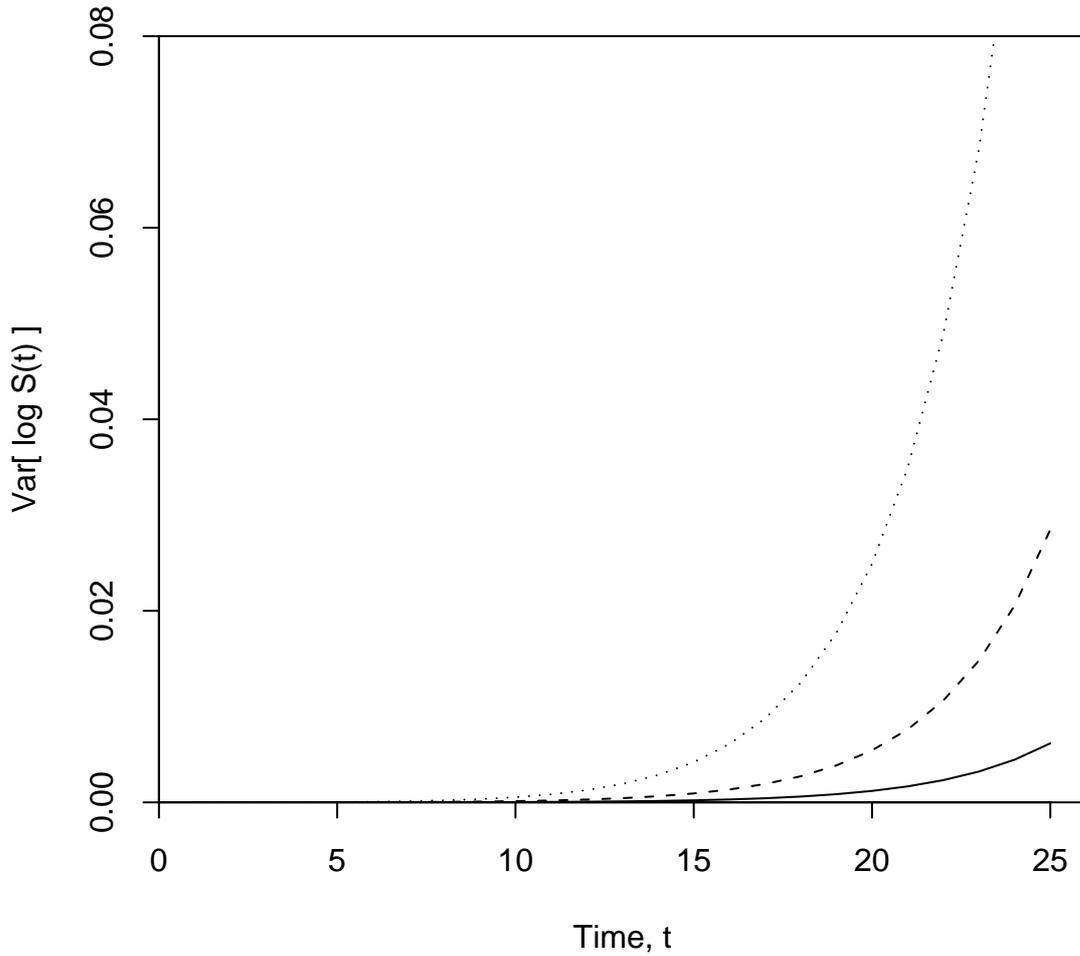


Figure 8: Plot of the variances of  $\log S(t)$  for the age 60 (solid line), 65 (dashed) and 70 (dotted) cohorts, based on data from 1982-2002, with no allowance for parameter uncertainty.

Table 5 shows the impact in the truncated expected future lifetime  $e(x, T)$  when we move from the real-world measure,  $P$ , to the risk-neutral measure  $Q(\lambda)$  when  $\lambda = (0.175, 0.175)$ .<sup>30</sup> The trends in this table match those in Table 4 with the exception of the trend along the diagonal from  $(x, T) = (60, 30)$  to  $(70, 20)$  where the trend is reversed. As we move upwards along the diagonal we have the same two factors working in opposite directions as before: decreasing term and increasing age. In Tables 2 to 4 the impact of discounting was sufficient to allow the increasing-age effect to dominate. In Table 5 the absence of discounting means that the decreasing-term effect is dominant.

<sup>30</sup>If  $\tau_x$  is the random future lifetime of an individual aged  $x$ , then  $e(x, T) = E[\min\{\tau_x, T\}] = \int_0^T E[S(t)]dt$ .

## 7 The impact of the EIB interest rate

We can also investigate the impact of a change in interest rates. Specifically, let us take  $\lambda$  as given but change the EIB interest rate from 4% to 5% per annum. In this case we find that the impact on the risk premium is relatively small. Specifically, if  $\lambda = (0.375, 0)$  then  $\delta(25, 65) = 19.1$  basis points and if  $\lambda = (0, 0.315)$  results in  $\delta(25, 65) = 18.9$  basis points. This reduction in the risk premium reflects the relative lowering, in present-value terms, of the later, more-uncertain cashflows under the bond.

## 8 Conclusions

In this paper we have used a simple two-factor model for the development of the mortality curve over time that seems, nevertheless, to fit the data well. The model allows us to simulate the distribution of a survivor index over various time horizons under both the real-world probability measure and under a variety of possible risk-adjusted measures. By taking expectations under the latter measure, this model enables us to price the longevity risk inherent in longevity bonds, given the known longevity risk premium (of 20 basis points) contained in the world's first longevity bond, namely the November 2004 EIB 25-year bond designed by BNP Paribas with a reference cohort of 65-year old English and Welsh males. The chosen model is well suited to pricing longevity bonds. For other types of contract that involve, for example, derivative characteristics on future mortality rates (such as guaranteed annuity options) models formulated within the forward-mortality-model (such as the Smith-Olivier model described in Olivier and Jeffrey, 2004) or the mortality-market-model (Cairns *et al*, 2004) frameworks are likely to prove more efficient to implement.

We find that the premium increases with both term and the initial age of the reference cohort. In the latter case, this is caused by the greater volatility that is associated with the higher mortality rates of older people compared with younger people. For example, in the worst-case scenario considered (where the entire longevity risk premium is associated with the second (ie volatility) factor), the premium for a 30-year bond with a reference cohort aged 70 is 42.3 basis points.

Another key finding of the paper is that the reference cohort's initial age is more important for determining the premium than the bond's maturity. To illustrate, again in the context of the worst-case scenario, the premium for a 20-year bond with a reference cohort aged 70 is 26.1 basis points, whereas the premium for a 30-year bond with a reference cohort aged 60 is 15.0 basis points. This shows that the greater uncertainty in death rates at higher ages dominates the greater discounting of the more distant cash flows of longer maturing bonds.

These findings suggest that open-ended survivor bonds that continue to pay out so long as members of the reference cohort are still alive would not have an excessively high longevity risk premium. However, they might be unattractive in other respects, such as the administrative inconvenience associated with paying very small coupons

50 years or so after the bond was issued. So fixed-term longevity bonds might well dominate for practical considerations. Our results also suggest that fixed-term longevity bonds might also be favoured by investors wishing to avoid the dominating effects of parameter risk over the basic longevity risk.

We propose in future research to investigate alternatives to the random walk model with drift used here. Possibilities include models drawn from the *ARIMA* class of time series models.<sup>31</sup> By taking this approach we will be investigating the important issue of model risk in addition to the parameter risk considered in this paper.

In this paper we have assumed that the fitted values of  $A(t)$  are known with certainty. A further line of research is to relax this assumption and to use instead filtering approaches or Markov Chain Monte Carlo (MCMC) methods to estimate jointly the posterior distribution of the parameters and of the current values of  $A(t)$ .

## Acknowledgements

AC wishes to thank the Isaac Newton Institute in Cambridge, where he was a visitor during the preparation of this paper. This research was conducted under research grant RES-000-23-1036 from the UK Economic and Social Research Council.

---

<sup>31</sup>The model used here is classified as an *ARIMA*(0, 1, 0) model. Elsewhere (Cairns, *et al.*, 2005) we have used an *ARIMA*(1, 1, 0) model with similar results to those in this paper.

## Appendices

### A The EIB/BNP longevity bond

The EIB/BNP Paribas longevity bond makes reference to a cohort (aged  $x$  at time 0) index that is calculated along the following lines. We start, for convenience, by letting  $t = 0$  correspond to the beginning of 2003 and set the reference index  $S(0) = 1$ . Changes in the reference index,  $S(t)$ , from one year to the next are determined by national mortality rates which are made publicly available. Thus, for  $t = 1, 2, \dots$ ,  $S(t+1) = S(t)(1 - \tilde{m}(t, x))$  where  $\tilde{m}(t, x)$  is the mortality rate for individuals aged  $x + t$  in year  $t$  (that is, age 65 in 2003, age 66 in 2004, and so on) published by the UK Office for National Statistics (ONS).

It is important to note, though, that the ONS calculates its mortality rates in a different way from the UK Government Actuary's Department (GAD) (for example, as with those for 2002 plotted in Figure 1). The reference population (England and Wales, males) is the same in both cases but the GAD publishes mortality rates  $\tilde{q}(t, x)$ , whereas the ONS publishes so-called central mortality rates<sup>32</sup>  $\tilde{m}(t, x)$ . An approximate relationship between the two rates is given by

$$\tilde{m}(t, x) \approx \frac{\tilde{q}(t, x)}{1 - \frac{1}{2}\tilde{q}(t, x)}.$$

### B Simulation of the Normal-Inverse-Wishart distribution

Equation (5) requires simulation of  $V|D$  using its posterior distribution, the Wishart( $n-1, n^{-1}\hat{V}^{-1}$ ) distribution. It is more instructive to show how to simulate from the Wishart distribution than it is to write down its density function. Thus:

- Let  $S$  be the upper triangular matrix that satisfies  $SS' = n^{-1}\hat{V}^{-1}$ .
- Now simulate  $n - 1$  i.i.d. vectors  $\alpha_1, \dots, \alpha_{n-1} \sim MVN(0, SS')$ : that is, let  $\alpha_i = SZ_i$  where  $Z_i$  is a standard  $n$ -dimensional normal random variable (that is, the individual elements of each  $Z_i$  are independent normal random variables with mean 0 and variance 1).
- Let  $X = \sum_{i=1}^{n-1} \alpha_i \alpha_i'$ .
- Then  $X$  has a Wishart( $n - 1, n^{-1}\hat{V}^{-1}$ ) distribution.
- Our final step is to invert  $X$  to get our simulated covariance matrix: that is,  $V = X^{-1}$ .

---

<sup>32</sup>In standard actuarial notation, the mortality rate  $q_y$  represents the proportion of the population aged  $y$  at the start of the year that die during the following 12 months. The central mortality rate  $m_y$  takes the cohort aged  $y$  at the start of the year and divides the number of deaths during the year by the average size of the cohort over the year.

Note that  $E[X] = \frac{n-1}{n}\hat{V}^{-1}$ . Thus the distribution of the simulated matrices  $V = X^{-1}$  will be centred close to  $\hat{V}$  itself.

The second step of simulating from the Normal-Inverse-Wishart posterior distribution is to take the simulated  $V$  from the steps above and then sample  $\mu$  from a multivariate normal distribution with mean  $\hat{\mu}$  and covariance matrix  $n^{-1}V$ . This can be simulated in the usual way.

## C Illustration of accumulated variance

Consider a random walk in which  $W(0) = 0$  and, for each  $t$ ,  $W(t+1) = W(t) + Z(t+1)$  where  $Z(1), Z(2), \dots$  is a sequence of i.i.d. standard normal random variables.

The way in which  $Var(\log S(t))$  builds up is similar to  $Y(T) = \sum_{t=1}^T W(t)$ . Since  $W(t) = Z(1) + \dots + Z(t)$  we have

$$Y(T) = TZ(1) + (T-1)Z(2) + \dots + 2Z(T-1) + Z(T).$$

This has variance  $\sum_{t=1}^T t^2 = \frac{1}{6}T(T+1)(2T+1)$ . In contrast  $Var W(T) = T$ .

If we introduce a drift,  $\lambda$ , into the process (that is, the  $Z(t)$  are i.i.d.  $N(\lambda, 1)$ ) (thereby having the same effect as a market price of risk) then

$$Y(T) = \frac{1}{2}T(T+1)\lambda + TZ(1) + (T-1)Z(2) + \dots + 2Z(T-1) + Z(T).$$

This does not affect the variance if  $\lambda$  is known. However, if  $\lambda$  is unknown with mean  $\hat{\lambda}$  and variance  $\sigma^2$ , then

$$Var Y(T) = \frac{1}{6}T(T+1)(2T+1) + \frac{1}{4}T^2(T+1)^2\sigma^2.$$

As a consequence, uncertainty in  $\lambda$  causes greater uncertainty in  $Y(T)$  than does the underlying volatility in  $W(T)$  as  $T$  gets larger (that is, the  $\frac{1}{4}T^2(T+1)^2\sigma^2$  term dominates as  $T$  gets large).

## References

- Benjamin, B., and Pollard, J.H. (1993) *The Analysis of Mortality and Other Actuarial Statistics, 3rd Edition*. Institute of Actuaries: London.
- Blake, D., and Burrows, W. (2001) "Survivor Bonds: Helping to Hedge Mortality Risk," *Journal of Risk and Insurance* 68, 339-348.
- Cairns, A.J.G. (2000) "A discussion of parameter and model uncertainty in insurance," *Insurance: Mathematics and Economics* 27, 313-330.
- Cairns, A.J.G., Blake, D., and Dowd, K. (2004) "Pricing death: Frameworks for the valuation and securitization of mortality risk," Preprint, Heriot-Watt University.
- Cairns, A.J.G., Blake, D., Dowd, K., and Dawson, P. (2005) "Pricing the risk on longevity bonds," Preprint, Heriot-Watt University.
- Cox, S., Fairchild, J., and Pedersen, H. (2000) "Economic Aspects of Securitisation of Risk," *ASTIN Bulletin* 30, 157-193.
- Cox, S., Lin, Y., and Wang, S. (2005) "Pricing longevity bonds," Presentation at the 1st International Conference on Longevity Risk and Capital Market Solutions, Cass Business School, London, February, 2005.
- Denuit, M., Devolder, P., and Goderniaux, A-C. (2004) "Securitization of longevity risk: pricing survival bonds with Wang transform in the Lee-Carter framework," Preprint, Université Catholique de Louvain.
- Dowd, K., Blake, D., Cairns, A., and Dawson, P. (2005) "Survivor swaps," Preprint, Centre for Risk and Insurance Studies, Nottingham University.
- Lee, R.D., and Carter, L.R. (1992) "Modeling and forecasting U.S. mortality," *Journal of the American Statistical Association* 87, 659-675.
- Lin, Y., and Cox, S.H. (2004) "Securitization of mortality risks in life annuities," Working paper, Georgia State University.
- Olivier, P., and Jeffery, T. (2004) "Stochastic mortality models," Presentation to the Society of Actuaries of Ireland.  
See [http://www.actuaries.ie/Resources/events\\_papers/PastCalendarListing.htm](http://www.actuaries.ie/Resources/events_papers/PastCalendarListing.htm)
- Perks, W. (1932) "On some experiments in the graduation of mortality statistics," *Journal of the Institute of Actuaries* 63, 12-57.
- Wang, S.S. (2000) "A class of distortion operations for pricing financial and insurance risks," *Journal of Risk and Insurance* 67, 15-36.
- Wang, S.S. (2002) "A universal framework for pricing financial and insurance risks," *ASTIN Bulletin* 32, 213-234.
- Wang, S.S. (2003) "Equilibrium pricing transforms: new results using Buhlmann's 1980 economic model," *ASTIN Bulletin* 33, 57-73.