A VALUE-AT-RISK APPROACH

TO MIS-ESTIMATION RISK

By S. J. Richards

LONGEVITAS™

A value-at-risk approach to mis-estimation risk

Richards, Stephen J.*

June 22, 2020

Abstract

Parametric mortality models, such as survival models, permit detailed analysis of risk factors for actuarial work. However, finite data volumes lead to uncertainty over parameter estimates, which in turn gives rise to mis-estimation risk of financial liabilities. Mis-estimation risk can be assessed on a run-off basis by valuing the liabilities with alternative parameter vectors consistent with the covariance matrix. This run-off approach is especially suitable for tasks like pricing portfolio transactions, such as bulk annuities, longevity swaps or reinsurance treaties. However, a run-off approach does not fully meet the requirements of regulatory regimes that view capital requirements through the prism of a finite horizon, such as Solvency II's one-year approach. This paper presents a methodology for viewing mis-estimation risk over a fixed time-frame and results are given for a specimen portfolio. As expected, we find that time-limited mis-estimation capital requirements increase as the horizon is lengthened or the discount rate is reduced. However, we find that much of the so-called mis-estimation risk in a one-year value-at-risk assessment can actually be driven by idiosyncratic variation, rather than parameter uncertainty. This counter-intuitive result stems from trying to view a long-term risk through an short-term window. We also find that parsimonious models tend to produce lower mis-estimation risk than less-parsimonious ones.

Keywords: mis-estimation risk, level risk, annuities, longevity risk, survival model, Solvency II.

1 Introduction and motivation

When pricing or reserving for a block of insurance contracts, mortality assumptions are commonly divided into two separate items: (i) the current level of mortality rates and (ii) projection of future trends. For each basis element there are risks in getting the assumption wrong. For example, for future trends there is no way of knowing if the chosen projection model is correct. This model risk in forecasting is discussed elsewhere; see for example Cairns [1998] and Richards et al. [2020]. This paper is concerned with the first basis element, i.e. the current level of mortality rates in a portfolio and the estimation risk thereof.

In deriving an assumption for current mortality rates, a model must be proposed and calibrated to the available experience data for the portfolio concerned. Using experience data unrelated to the portfolio should be avoided as far as possible, since this could introduce bias from lives with different mortality characteristics, i.e. basis risk. In this paper we will assume that a best-estimate basis for a portfolio can be derived from its own experience data. Many approaches exist: from non-parametric methods [Kaplan and Meier, 1958] to semi-parametric models [Macdonald et al., 2018] to fully parametric models [Richards, 2012]. In each case there is a risk that the true underlying mortality rates are different from the estimated rates due to sampling error. This risk is further compounded by the tendency for liabilities to be concentrated in a relatively small sub-group of lives (see Table 5, where over half of the pension payments are made to just a fifth of pensioners). These risks — sampling error and concentration of liabilities — combine to produce uncertainty over the

*stephen@longevitas.co.uk, Longevitas Ltd, 24a Ainslie Place, Edinburgh, EH3 6AJ. www.longevitas.co.uk

current mortality rates and a magnified impact on the value of the liabilities. This uncertainty is variously labelled mis-estimation risk or level risk.

When pricing block transfers of risk, such as bulk annuities or longevity swaps, an insurer is interested in the financial impact of mis-estimation risk over the entire lifetime of the portfolio. The methodology in Richards [2016] provides the run-off view of mis-estimation risk required for such tasks. However, regulatory frameworks like Solvency II view risk over a one-year horizon, not in run-off. This paper adapts the pricing mis-estimation methodology of Richards [2016] to frame mis-estimation risk in a short time horizon like 1–5 years. To illustrate the results we use the records of a UK pension scheme described in detail in Appendix A.

The plan of the rest of paper is as follows: Section 2 defines various terms used; Section 3 describes the methodology for assessing mis-estimation, together with some basic validity conditions; Section 4 describes how this methodology is adapted to view mis-estimation over a limited time horizon; Section 5 outlines how models are structured, while Section 6 looks at specimen results over various horizons; Section 7 considers the sensitivity to discount rate, while Section 8 looks at variation with model type; Section 9 concludes.

2 Definitions

Denote by $\hat{\boldsymbol{\theta}}$ the maximum-likelihood estimate of a parameter vector $\boldsymbol{\theta}$. Let $\hat{\boldsymbol{\theta}}^{(n)}$ be the maximum-likelihood estimate of $\boldsymbol{\theta}$ from the addition of n years of further experience data. Denote by $V(\boldsymbol{\theta}, y)$ the value of life-contingent liabilities in-force at time y using the mortality rates effective at time y according to the model specified by the parameter vector $\boldsymbol{\theta}$. We assume that all basis elements for the calculation of $V(\boldsymbol{\theta}, y)$ are known apart from the current mortality rates at y, i.e. that $V(\boldsymbol{\theta}, y)$ is a deterministic function of $\boldsymbol{\theta}$, but uncertainty is introduced through uncertainty over $\boldsymbol{\theta}$. Throughout this paper mortality improvements will be modelled up to y, but no future mortality improvements after y will be assumed — the mortality rates used will be those applying at y with no allowance for future changes in time apart from the ageing of the lives assured. We are concerned in this paper with a value-at-risk assessment of the mis-estimation risk in $V(\hat{\boldsymbol{\theta}}, y)$ only; for a value-at-risk approach to longevity trend risk after time y, see Börger [2010], Plat [2011] or Richards et al. [2020].

The best-estimate of the liability at time y is taken to be the expected value, $\mathbb{E}[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]$, which for an unbiased model will equal $\mathbb{E}[V(\hat{\boldsymbol{\theta}}, y)], \forall n$. We value liabilities with $\hat{\boldsymbol{\theta}}^{(n)}$ at time y, rather than at time $y+n$, is because we are interested in the potential impact on current liabilities of n additional years of experience after time y (a risk against which we can hold additional capital).

The risk measure of interest is $\mathrm{VaR}_\alpha[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]$ defined in equation (1), i.e. the proportion of the best-estimate needed to cover a proportion α of losses that might occur due to a change in the best-estimate assumption caused by an additional n years of experience data after time y. $\mathrm{Q}_\alpha[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]$ is the α-quantile of the distribution of liability $V(\hat{\boldsymbol{\theta}}^{(n)}, y)$, which we will estimate according to Harrell and Davis [1982].

$$\mathrm{VaR}_\alpha[V(\hat{\boldsymbol{\theta}}^{(n)}, y)] = \frac{\mathrm{Q}_\alpha[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]}{\mathbb{E}[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]} - 1 \quad (1)$$

$$\Pr[V(\hat{\boldsymbol{\theta}}^{(n)}, y) \leq \mathrm{Q}_\alpha[V(\hat{\boldsymbol{\theta}}^{(n)}, y)]] = \alpha \quad (2)$$

In the European Union (EU) the Solvency II regime for insurer solvency calculations is based on $n = 1$ and $\alpha = 99.5\%$, i.e. $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(1)}, y)]$. In this paper V will be the reserve for a portfolio of continuously paid single-life annuities, as defined in equations (3) and (4). In equation (4) w_i is the level annuity

$$V(\boldsymbol{\theta}, y) = \sum_i a(i, y, \boldsymbol{\theta}) \quad (3)$$

$$a(i, y, \boldsymbol{\theta}) = w_i \int_0^\infty {}_tp_{x_i, y, \boldsymbol{\theta}} v^t dt \quad (4)$$

paid to life i aged x_i at time y and ${}_tp_{x,y,\boldsymbol{\theta}}$ denotes the t-year survival probability at outset age x at time y. v^t is a discount function, which can be adapted to allow for escalating benefit payments if necessary. In this paper we will mainly discount at a constant net rate of 0.75% per annum, applied continuously, although Section 7 considers the impact of different discount rates.

3 Parameter risk and mis-estimation

We broadly recapitulate the mis-estimation methodology of Richards [2016]. We assume that we have a log-likelihood function, $\ell(\boldsymbol{\theta})$, that is twice differentiable so that a Hessian matrix may be calculated [McCullagh and Nelder, 1989, page 6]. The true underlying value of $\boldsymbol{\theta}$ is unknown and is denoted $\boldsymbol{\theta}^*$. The maximum-likelihood estimate of $\boldsymbol{\theta}^*$ is $\hat{\boldsymbol{\theta}}$. Under the maximum-likelihood theorem $\hat{\boldsymbol{\theta}} \approx \mathcal{N}(\boldsymbol{\theta}^*, \mathcal{I}^{-1})$, where $\mathcal{I}$ is the Fisher Information [Cox and Hinkley, 1996, Chapter 9]. In practice we replace the unknown $\boldsymbol{\theta}^*$ with $\hat{\boldsymbol{\theta}}$, i.e. $\hat{\boldsymbol{\theta}} \approx \mathcal{N}(\hat{\boldsymbol{\theta}}, \mathcal{I}^{-1})$. Parameter uncertainty is summarised in $\mathcal{I}^{-1}$, i.e. not just the parameter variances along the leading diagonal but also the covariances between parameter estimates [Richards et al., 2013, Table 14]. To explore parameter risk we can generate an alternative parameter vector, $\boldsymbol{\theta}'$, that is consistent with the data and model from $\boldsymbol{\theta}' = \hat{\boldsymbol{\theta}} + \boldsymbol{A}\boldsymbol{z}$, where $\boldsymbol{A}$ is the Cholesky decomposition of $\mathcal{I}^{-1}$ and $\boldsymbol{z}$ is a vector of independent $\mathcal{N}(0,1)$ variates of the same length as $\boldsymbol{\theta}$.

For assessing mis-estimation risk in run-off Richards [2016] looked at the variation in $V(\boldsymbol{\theta}', y)$ from repeated simulation of $\boldsymbol{z}$. The alternative valuations were normalised by dividing by the mean of the values for V and various quantiles computed to express mis-estimation risk as a percentage of the expected reserve. There are several important conditions for a model to be suitable for this kind of approach, and a summary overview is given in Table 1.

Table 1: Assumption checklist for mis-estimation methodology.

Assumption	Potential problem	Solution or check
Independent lifetimes in model-fitting and simulation.	Multiple benefit records per individual, leading to failure of independence assumption and under-estimation of parameter variance.	Deduplication of benefit records; see Macdonald et al. [2018, Section 2.5].
Parameter estimates have multivariate Normal distribution.	Parameters not distributed as multivariate Normal.	Plotting of marginal log-likelihoods to check inverted quadratic shape; see Figure 1.
Static mortality, or time trends allowed for.	Using multi-year data without allowing for a time trend leads to false confidence in estimate of current mortality.	Tests of fit by calendar year [Macdonald et al., 2018, Section 6.5], inclusion of time-trend parameter in model (Figure 9).
Model suitable for financial purposes.	Liabilities disproportionately concentrated in small sub-group with different mortality characteristics; See Table 5.	Tests of fit by pension size [Macdonald et al., 2018, Section 6.5] or "bootstrapping" [Richards et al., 2013, Section 8.3].

4 A value-at-risk approach to mis-estimation

The methodology in Section 3 was used in Richards [2016] for what might be called pricing mis-estimation, such as when transfering a block of liabilities. The experience data of the block up to time y is used to calibrate a model for best-estimate purposes, and the pricing risk is represented by the range of values for V consistent with $\hat{\boldsymbol{\theta}}$ and the estimated covariance matrix, $\mathcal{I}^{-1}$.

Table 2: Parameter options for simulating lifetimes.

Parameter vector	Parameter risk	Description
$\hat{\boldsymbol{\theta}}$	No	Maximum-likelihood estimate (MLE).
$\boldsymbol{\theta}' = \hat{\boldsymbol{\theta}} + \boldsymbol{A}\boldsymbol{z}$	Yes	Perturbed MLE.

However, regulatory frameworks like Solvency II view risk over a fixed horizon, a feature that is absent from the run-off approach in Section 3. We can however adapt the methodology to an n-year view of mis-estimation risk in two stages. In the first stage, we use the model to simulate the future lifetimes of the survivors and truncate additional time lived for those surviving more than n years; Richards [2012, Table 7] lists formulae for simulating future lifetimes under various mortality models. As per Table 2, we have two options for simulation: we can either use the best-estimate parameters, $\hat{\boldsymbol{\theta}}$, or else we can use the perturbed parameters, $\boldsymbol{\theta}'$. Using the latter will increase the variability in the survival times, and corresponds to what one would understand to drive mis-estimation. The only point in this paper where we will not include parameter risk is in Section 6, where we will switch it off to quantify the relevant contributions of parameter risk and idiosyncratic risk.

In the second stage, we take the real experience data up to time y, add the additional n years of simulated pseudo-experience and refit the model; this will yield an alternative parameter vector, $\hat{\boldsymbol{\theta}}^{(n)}$. $\hat{\boldsymbol{\theta}}^{(n)}$ can be viewed as the response of the parameter estimates to n years of new experience data, and we use it to value the liabilities at time y. We repeat this process of simulating lifetimes, refitting the model and revaluing the liabilities to collect (say) 10,000 realisations of the liability value, $V_j, j = 1, \ldots, 10,000$. The approach of repeatedly refitting models and revaluing liabilities is necessarily computationally intensive, so we use parallel processing over 63 threads to reduce run-times [Butenhof, 1997].

5 Model structure

Following Macdonald et al. [2018] we use survival models for individual lifetimes. We maximise the log-likelihood shown in equation (5), where $\mu_{x,y}$ is the mortality hazard at exact age x and calendar time y. Life i enters observation at exact age x_i at time y_i and is observed for t_i years. d_i is an indicator variable taking the value 1 if life i is dead at observation time t_i, and zero otherwise. $H_{x,y}(t)$ is the integrated hazard function, defined in equation (6). We have a wide choice of functional forms for $\mu_{x,y}$ for post-retirement mortality; [Richards, 2012] reviews seventeen such models applied to the mortality of UK annuitants. However, we start with the simple and familiar model of Gompertz [1825] in equation (7), where the offset of -2000 is to keep the parameters well-scaled. δ represents the portfolio-specific time trend; here it is common to all lives, but it could be interacted with any risk factor (such as with gender to estimate separate time trends for males and females).

$$\ell = -\sum_i H_{x_i,y_i}(t_i) + \sum_i d_i \log \mu_{x_i+t_i,y_i+t_i} \quad (5)$$

$$H_{x,y}(t) = \int_0^t \mu_{x+s,y+s} ds \quad (6)$$

$$\mu_{x,y} = e^{\alpha_i+\beta_i x+\delta(y-2000)} \quad (7)$$

α_i and β_i are parameters for life i structured as in equations (8) and (9), where $\alpha^{(j)}$ is the main effect of risk factor j and $\beta^{(j)}$ is the interaction of the j^{th} risk factor with age. $z_i^{(j)}$ is an indicator variable taking the value 1 if life i has risk factor j and zero otherwise. Using the data in Appendix A we fit a model with age, gender (male or female), early-retirement status (pension commencing before or after age 55), widow(er) status, pension size (below £5,385 p.a., £5,385-12,560 p.a. or above £12,560 p.a.) and calendar time as explanatory variables. For multi-level factors we adopt a policy of making the most numerous level the reference value, i.e. the baseline case is a male first life retiring after age 55 with a small pension. We therefore have parameters for those retiring early, widow(er)s, females and those with medium or large pensions. We estimate these parameter values by maximising the log-likelihood in equation (5), with the results shown in Table 3. Of particular note are the different mortality characteristics of those with the largest pensions. Uncertainty over the mortality of this small sub-group has a disproportionately large impact on uncertainty over the liability value, V.

$$\alpha_i = \alpha_0 + \sum_{j=1}^{m} \alpha^{(j)} z_i^{(j)} \tag{8}$$

$$\beta_i = \beta_0 + \sum_{j=1}^{m} \beta^{(j)} z_i^{(j)} \tag{9}$$

Table 3: Parameter estimates under the Gompertz [1825] model. The second column is $\hat{\boldsymbol{\theta}}$ in the sense of Sections 3 and 4, while the third column is the square root of the leading diagonal of $\mathcal{I}^{-1}$. Source: own calculations fitting model in equations (5)-(9) to the data in Appendix A.

Parameter	Estimate	Standard error	Z-value	p-value	Contributors: Lives	Deaths	Years lived
Age (β_0)	0.10097	0.0020	49.81	0	44,616	10,663	260,374.0
EarlyRetirement	1.1306	0.2572	4.40	0	8,848	1,305	49,681.6
EarlyRetirement:Age	-0.011712	0.0035	-3.37	0.0007	8,848	1,305	49,681.6
Widow(er)	0.903570	0.2384	3.79	0.0002	9,183	3,285	51,643.6
Widow(er):Age	-0.0098666	0.0029	-3.38	0.0007	9,183	3,285	51,643.6
Female	-1.6377	0.2117	-7.74	0	25,541	5,693	150,089.0
Female:Age	0.014363	0.0027	5.35	0	25,541	5,693	150,089.0
Intercept (α_0)	-10.390	0.1618	-64.22	0	44,616	10,663	260,374.0
Medium pension	-0.76801	0.2266	-3.39	0.0007	8,924	1,889	51,846.4
Large pension	-2.7358	0.4985	-5.49	0	2,226	329	12,546.5
Medium pension:Age	0.0072229	0.0028	2.56	0.0104	8,924	1,889	51,846.4
Large pension:Age	0.028330	0.0062	4.60	0	2,226	329	12,546.5
Time (δ)	-0.039991	0.0037	-10.72	0	44,616	10,663	260,374.0

We need to check the validity of our assumptions using the checklist in Table 1 before performing any mis-estimation assessments. We have already deduplicated the data, as described in Appendix A, so the independence assumption holds true.

Regarding the assumption of a multivariate normal distribution for the parameter estimates, we can see in Figure 1 that all profile log-likelihoods are suitably quadratic around the maximum-likelihood estimates. Note that the choice of horizontal scale is important, as it would be possible to find a quadratic-like shape for even poor models by selecting a suitably narrow range; in contrast, in Figure 1 the horizontal range is determined by the estimated standard error of the parameter.

Since we are only interested in detecting exceptions to the quadratic shape, we can create a space-saving "signature" for the profile log-likelihoods in Figure 1 by plotting them one after another without labels to see if there are any that do not have a

Figure 2: "Signature" formed from profile log-likelihoods in Figure 1.

Figure 1: Profile log-likelihoods around estimates in Table 3. The horizontal scale is determined by two standard errors on either side of the joint maximum-likelihood estimate of each parameter.

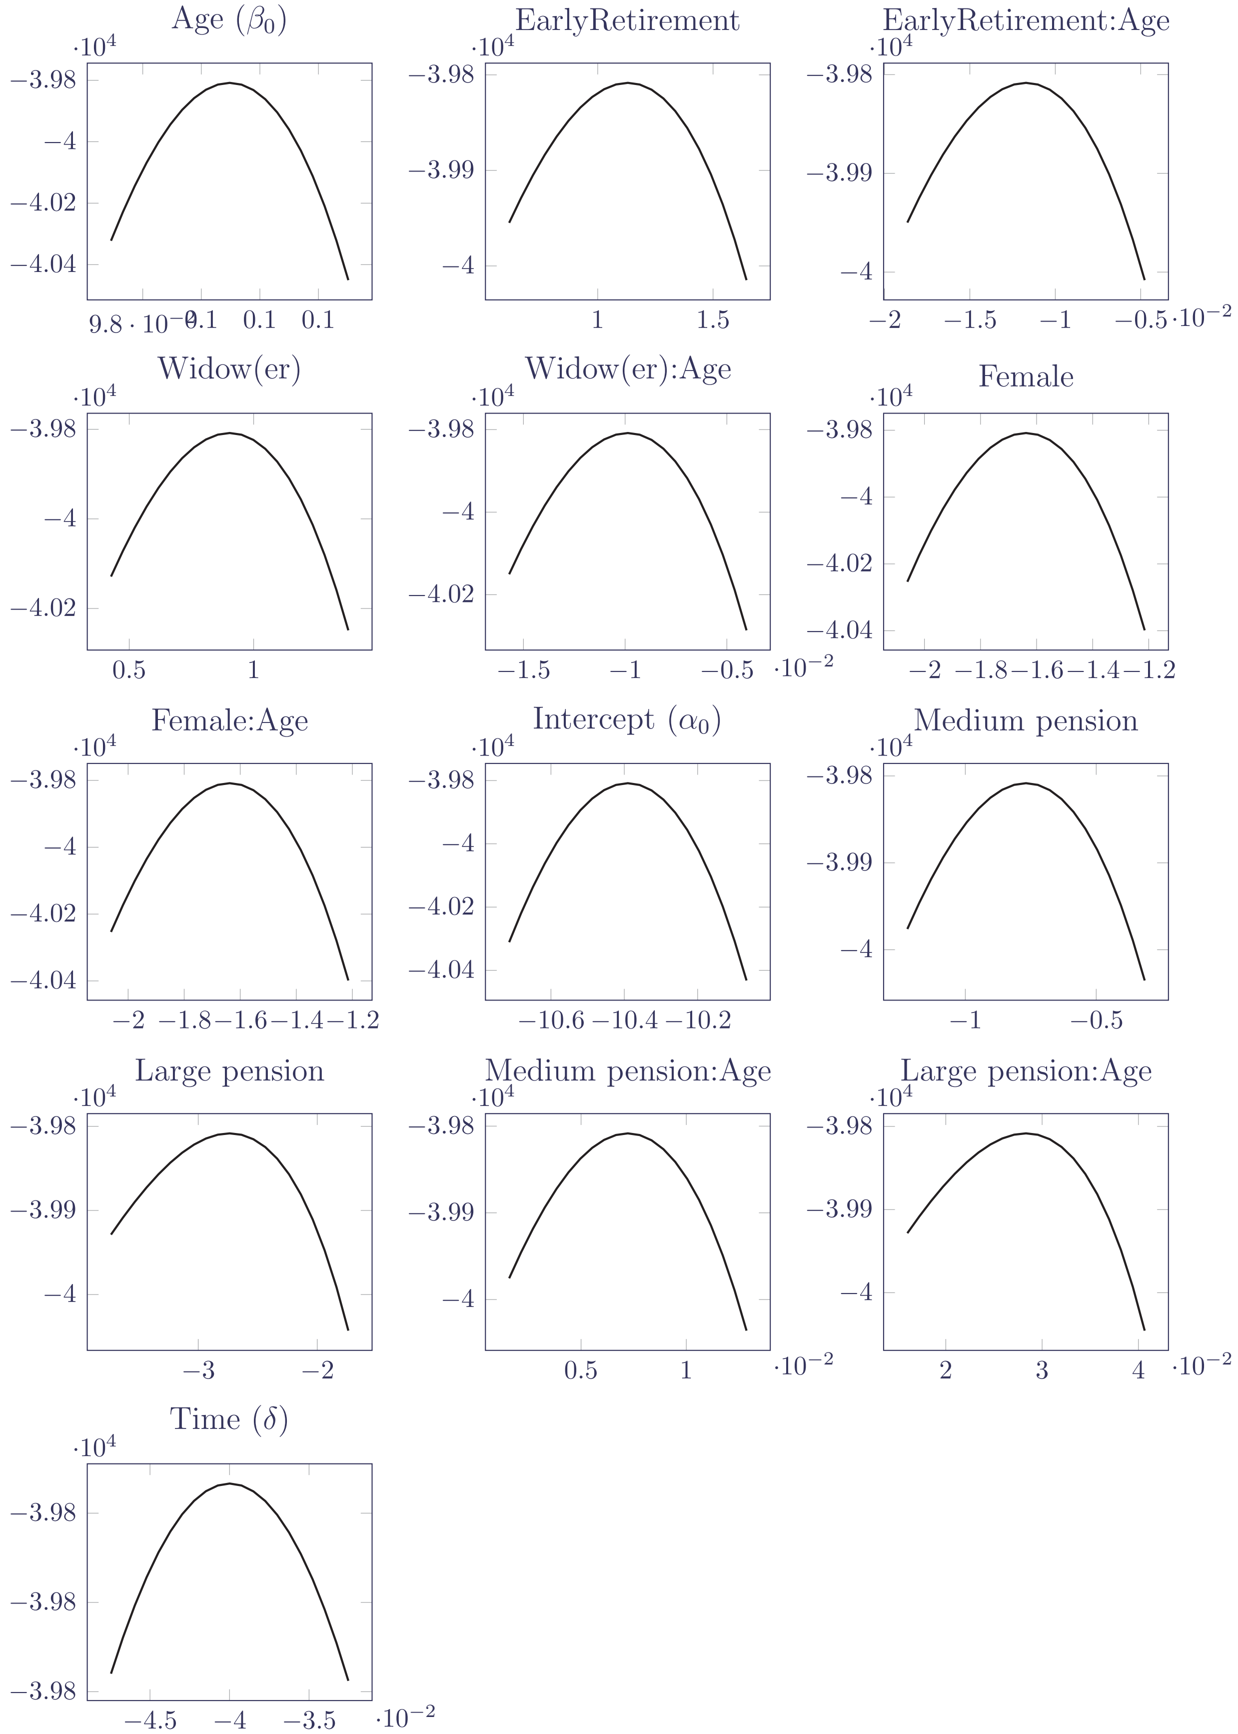

cleanly inverted U-shape. Figure 2 shows how this approach summarises the shapes in Figure 1 in a single short "signature". This approach becomes particularly useful when the number of parameters grows large. For an example where a parameter fails this quadratic test, see the signature for the Makeham-Perks model in Table 4.

The model in Table 3 has a time-trend parameter, so the only remaining item to check is financial suitability. Figure 3 shows that the residuals by pension size-band are plausibly drawn from a N(0,1) distribution [Macdonald et al., 2018, Section 6.5]; the apparent outlier is of minimal financial significance. Furthermore, the bootstrapping procedure of Richards [2016, Section 8.3] shows that on average the model in Table 3 predicts 100.1% of lives-based mortality and 98.9% of amounts-based mortality. The model is therefore broadly suitable for financial purposes, and so all four validity criteria in Table 1 are fulfilled.

Figure 3: Deviance residuals by pension size-band ($1 \equiv 5\%$ of lives with smallest pensions, $20 \equiv 5\%$ of lives with largest pensions).

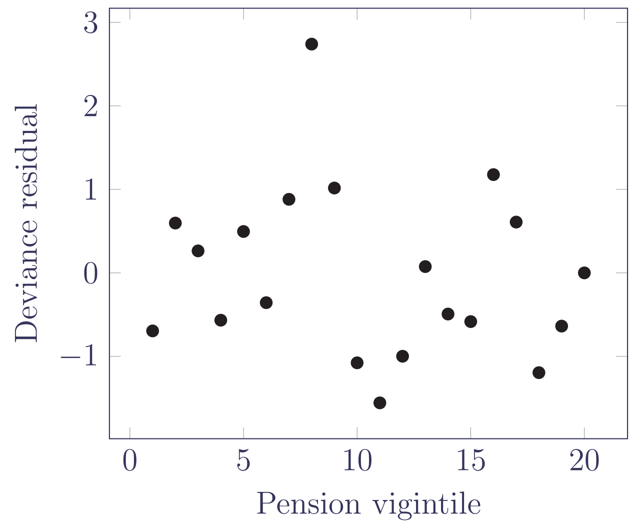

Before considering the value-at-risk approach to mis-estimation of Section 4, we first consider the run-off approach of Section 3. Figure 4 shows the distribution of reserves with 10,000 simulations of the parameter column in Table 3. The 99.5% point of the distribution is 2.61% above the mean, with a 95% confidence interval of 2.55–2.67%. On a run-off basis, the mis-estimation capital requirement would therefore be around 2.6%.

Figure 4: Distribution of 10,000 simulations of $V(\boldsymbol{\theta}', 2010)$ for model in Table 3 applied to survivors at 1st January 2010 for portfolio in Appendix A. Mortality rates are at 1st January 2010 with no further improvements.

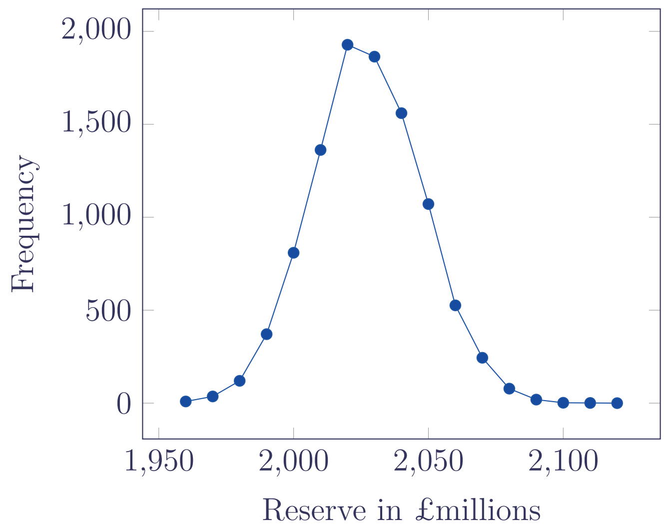

The run-off approach to mis-estimation is useful for setting a confidence interval on the pricing basis for a bulk annuity or longevity swap. For a best-estimate basis, the mean of the distribution in Figure 4 can be back-solved to a given percentage of a chosen table (or done separately for the reserves for males and females). For example, if we use the S2PA table [CMI Ltd, 2014] the equivalent best-estimate percentages are 109.7% for males and 100.1% for females. We can further use the 2.5% and 97.5% points of the distributions to form a 95% confidence interval for this basis: back-solving leads to 95% confidence intervals of 103.6–116.3% for males and 95.3–105.0% for females. Note that the confidence interval is not symmetric around the central estimate in part because the distribution of reserves is not normal — the p-value of the test statistic from Jarque and Bera [1987] is 0.2581 for males, but 0.0009158 for females. This means that one cannot automatically assume normally-distributed risk for mis-estimation.

6 The role of VaR horizon and parameter risk

With the model from Section 5 we turn to the question of the n-year value-at-risk capital assessment. Figure 5 shows the mis-estimation capital requirements for the portfolio described in Appendix A. The impact of horizon, n, is shown both with parameter risk (simulating lifetimes with $\boldsymbol{\theta}'$ as per Table 2) and without (simulating lifetimes with $\hat{\boldsymbol{\theta}}$ only). The capital requirements without parameter risk are fairly flat, as the simulated experience is similar to the real data. In contrast, the capital requirements including parameter risk rise with increasing horizon; most of this difference is driven by the uncertainty over the experienced time trend, which makes the estimate of mortality levels at 1st January 2010 more uncertain.

A comparison of the two series in Figure 5 shows the relative role of parameter risk over a short horizon and reveals an oddity: the value of $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(1)}, 2010)]$ would be used for the purposes of Solvency II mis-estimation capital, yet most of the capital requirement is clearly driven by the idiosyncratic variation in the simulated experience, not parameter risk: we have 1.27% with parameter risk, but we still have 1.08% without it. The run-off or pricing mis-estimation assessment at the end of Section 5 was driven solely by estimation error in $\boldsymbol{\theta}$, but 85% of the one-year VaR mis-estimation capital stems from idiosyncratic risk. This counter-intuitive aspect of the value-at-risk approach is not an outlier: Kleinow and Richards [2016, Table 5] found that most of the value-at-risk capital for longevity trend risk was similarly driven by the simulated experience, not parameter risk.

Figure 5: Mis-estimation $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(n)}, 2010)]$ capital requirements with and without parameter risk in simulation of additional n years of experience data. 95% confidence intervals are marked with −. Source: 10,000 simulations of model valuing single-life immediate-annuity cash-flows discounted at 0.75% p.a.

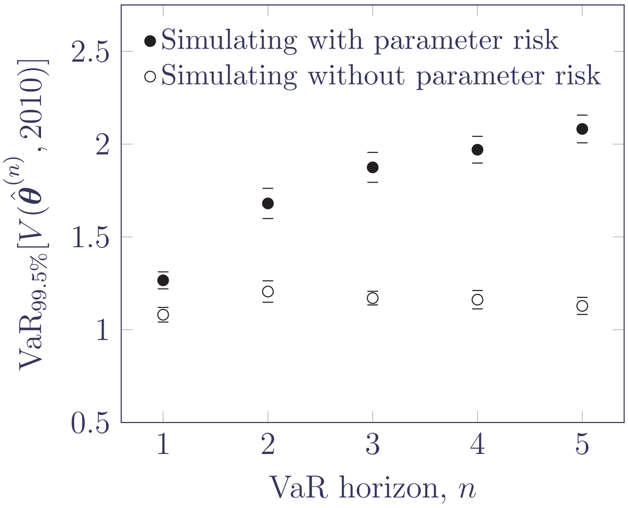

Figure 6: Distribution of 10,000 simulations of $V(\hat{\boldsymbol{\theta}}^{(1)}, 2010)$ for model in Table 3 applied to survivors at 1st January 2010 for portfolio in Appendix A. Mortality rates are at 1st January 2010 with no further improvements.

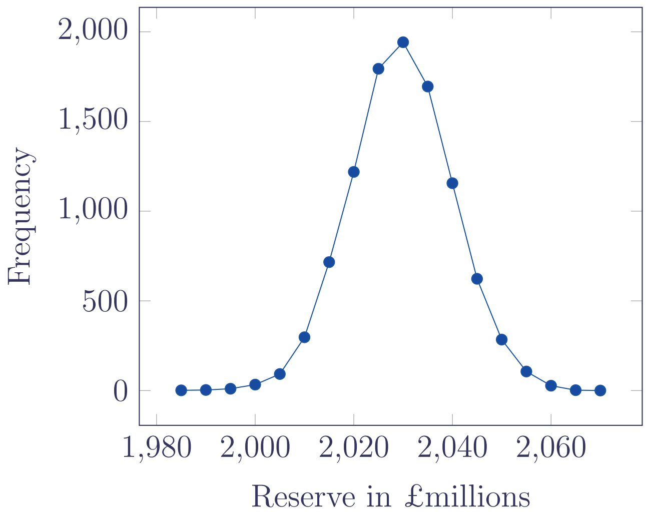

Figure 6 shows the distribution of $V(\hat{\boldsymbol{\theta}}^{(1)}, 2010)$ from which $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(1)}, 2010)]$ was calculated. As with Figure 4, we can use percentiles to back-solve to a percentage of a standard table: the 99.5% reserves for males and females equate to 106.1% of S2PA for males and 97.0% for females. Compared with the central estimates in Section 5, a shorthand 99.5% stress for Solvency II mis-estimation risk would then be -3.6% of S2PA for males and -3.1% for females (larger portfolios would likely have smaller mis-estimation stresses).

7 The role of discount rate

Figure 7 shows the VaR mis-estimation capital requirements using various discount rates. For a given horizon, risk capital increases as the discount rate falls. Mis-estimation assessments clearly need to be regularly updated as the shape or level of the yield curve changes.

With historically low interest rates and inflation-proofed defined-benefit pensions, annuity liabilities for bulk buy-outs will at times have to be valued using a net negative discount rate, implying higher mis-estimation capital requirements than those in Figure 7.

Figure 7: Mis-estimation $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(n)}, 2010)]$ capital requirements at various discount rates. Source: 10,000 simulations with parameter risk of model fitted to data for UK pensioner liabilities in Appendix A.

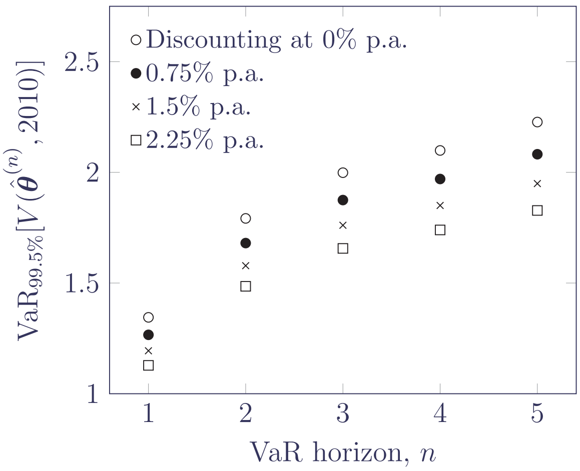

8 The role of mortality law

The results in this paper have so far all been based on the Gompertz [1825] mortality law. In this section we explore some alternative mortality laws, starting with the simplified logistic model from Perks [1932] in equation (10). It has the same number of parameters as the Gompertz law, but captures the tendency for log(mortality) to increase less slowly than linearly at advanced ages; see Barbi et al. [2018] and Newman [2018] for the ongoing debate as to the validity of this phenomenon. Individual mortality differentials are handled in the same way as the Gompertz model with equations (8) and (9). A variation on the logistic Perks law is the model from Beard [1959] shown in equation (11). The Beard and Gompertz laws are linked: if individual mortality follows a Gompertz law, but there is also Beta-distributed heterogeneity in α, then observed mortality will follow the Beard law; see Horiuchi and Coale [1990]. Another variation on equation (10) is to add a constant Makeham-like term [Makeham, 1860], as in equation (12).

$$\mu_{x,y} = \frac{e^{\alpha+\beta x+\delta(y-2000)}}{1+e^{\alpha+\beta x+\delta(y-2000)}} \tag{10}$$

$$\mu_{x,y} = \frac{e^{\alpha+\beta x+\delta(y-2000)}}{1+e^{\alpha+\beta x+\rho+\delta(y-2000)}} \tag{11}$$

$$\mu_{x,y} = \frac{e^{\epsilon}+e^{\alpha+\beta x+\delta(y-2000)}}{1+e^{\alpha+\beta x+\delta(y-2000)}} \tag{12}$$

A more recent option is the Hermite-spline model of Richards [2020] in equation (13), where the Hermite basis-spline h functions are shown in Figure 8. Hermite splines are defined for $t \in [0,1]$ and so we map age x onto $[0,1]$ with $t = (x-x_0)/(x_1-x_0)$ with pre-defined values of $x_0 = 50$ and $x_1 = 105$. We assume $\mu_x = \mu_{x_0}, x \le x_0$ and $\mu_x = \mu_{x_1}, x \ge x_1$.

$$\log \mu_{x,y} = \alpha h_{00}(t) + m_0 h_{10}(t) + \omega h_{01}(t) \tag{13}$$

Figure 8: Hermite basis splines for $t \in [0,1]$.

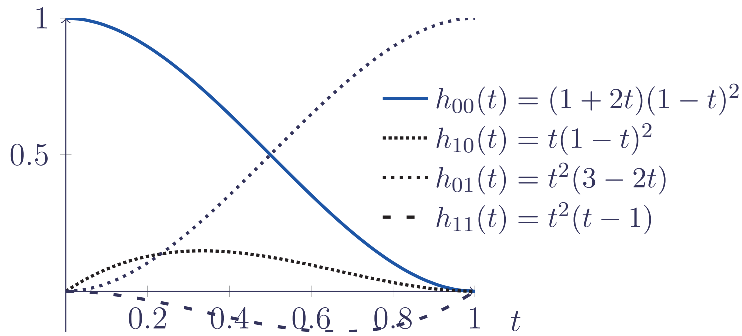

This new class of Hermite-spline models is highly parsimonious. It was specifically designed for modelling post-retirement mortality such that mortality differentials automatically narrow with increasing age. This reduces the number of parameters compared to the other four models, and further avoids the crossing-over of fitted mortality rates at advanced ages; see Richards [2020, Section 1]. Individual mortality differentials are therefore handled with just equation (8) — narrowing age differentials are handled automatically and so there is no need for equation (9) with the Hermite family.

Richards [2020] modelled age-related mortality changes with a peak improvement at an age to be estimated from the data. Here we instead extend equation (13) for time variation as follows:

$$\log \mu_{x,y} = (\alpha + \delta(y-2000))h_{00}(t) + (m_0 + m_0^{\text{trend}}(y-2000))h_{10}(t) + \omega h_{01}(t) \tag{14}$$

where δ plays a similar role to equations (10)-(12) by changing the level of mortality in time and m_0^{trend} changes the shape at younger ages. A common feature to both parameters is the automatic reduction in influence with age when multiplying by the Hermite spline functions h_{00} and h_{10}. We find that δ in equation (14) does not improve the fit for the data set in Appendix A, so we use this simpler version:

$$\log \mu_{x,y} = \alpha h_{00}(t) + (m_0 + m_0^{\text{trend}}(y-2000))h_{10}(t) + \omega h_{01}(t) \tag{15}$$

Equations (7), (10), (11), (12) and (15) are therefore all one-parameter approaches to changes in mortality level. Figure 9 shows the implied annual mortality improvements for the baseline combination of risk factors in Table 3 (the implied improvements under the Beard model are not shown as they are indistinguishable from the Perks model). Each model has a single-parameter allowance for mortality change, but clearly some models have a more reasonable shape for improvements by age than others. At one extreme the Gompertz model in equation (7) has an unrealistic constant rate of improvement at all ages, while perhaps the most realistic shape of all is from the Hermite model of equation (15).

Figure 9: Modelled percentage mortality improvements per annum by age. Source: own calculations of $100\% \times (1 - \mu_{x,2001}/\mu_{x,2000})$ for male first lives with the smallest 75% of pensions who retired after age 55. The period covered by the data is 2001–2009.

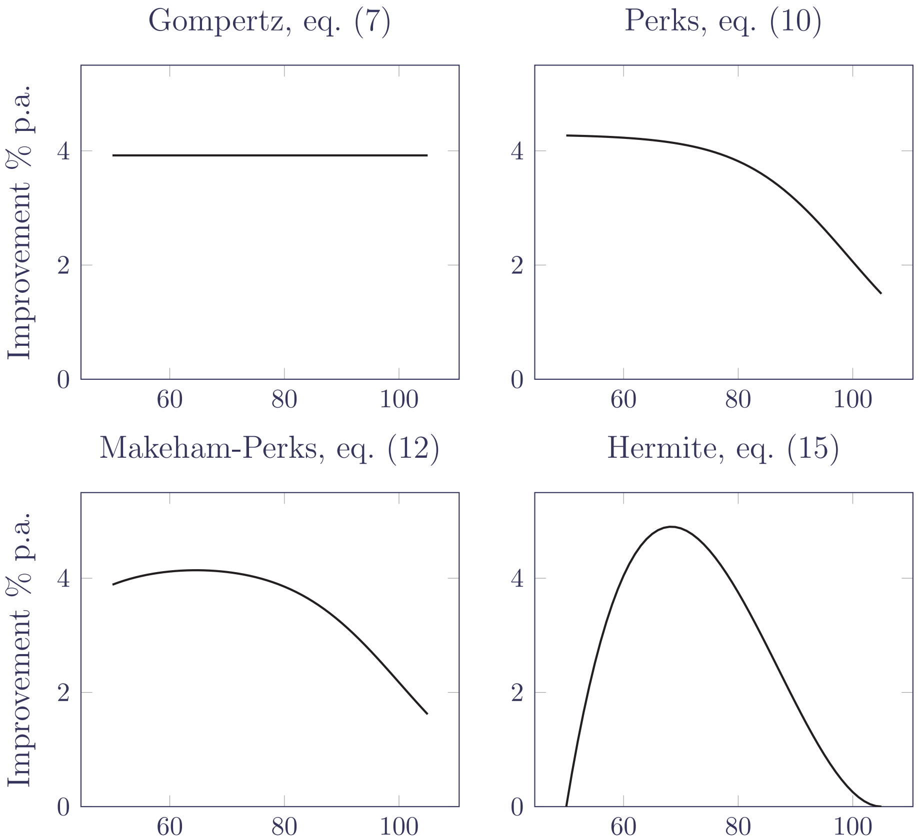

Figure 10 shows the mis-estimation capital requirements for each of the five models. We see that the parsimony of the Hermite model in equation (15), and the more-realistic allowance for improvements in Figure 9, result in less mis-estimation risk and thus less capital. Table 4 shows that there are no unwelcome compromises: the Hermite model has the lowest AIC of the five, while having an ability to predict lives- and pension-weighted variation as good as any of the other models.

Figure 10: Mis-estimation $\mathrm{VaR}_{99.5\%}[V(\hat{\boldsymbol{\theta}}^{(n)}, 2010)]$ capital requirements for various mortality laws. Source: 10,000 simulations with parameter risk of model fitted to data for UK pensioners, single-life immediate-annuity cashflows discounted at 0.75% p.a.

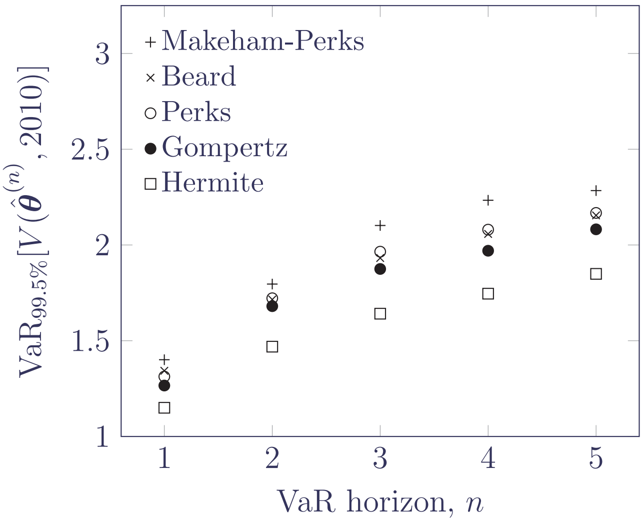

Table 4: Summary of model fits. Note that one of the Makeham-Perks parameters, ϵ, does not have a properly quadratic profile in the log-likelihood signature, although the impact is minimal. Source: own calculations fitting to data in Appendix A; bootstrap percentages are the mean ratio of actual deaths v. model-predicted deaths from 10,000 samples of 10,000 lives (sampling with replacement).

Mortality law	Equation number	Parameter count	Log-likelihood profile signature	AIC	Mean bootstrap percentage: (a) lives	(b) pensions
Gompertz	(7)	13		79,642.7	100.1%	98.9%
Perks	(10)	13		79,637.6	100.0%	99.3%
Beard	(11)	14		79,625.6	100.1%	99.0%
Makeham-Perks	(12)	14		79,624.8	100.0%	99.0%
Hermite	(15)	10		79,623.4	100.1%	99.2%

9 Conclusions

A value-at-risk approach to mis-estimation risk can be obtained by (i) fitting a suitable mortality model, (ii) repeated simulation of additional experience with this model, and (iii) refitting the model and valuing the liabilities with the updated parameters. Quantiles can be calculated from the liability distribution, and can be used to back-solve stress tests expressed in terms of a standard table. We find that the resulting capital requirements for short time horizons are very different from a run-off approach that might be used for pricing. At the shortest horizon of one year, much of the capital requirement stems from the idiosyncratic variation in portfolio simulation, not the parameter risk underlying the original model. This counter-intuitive result arises from the somewhat artifical regulatory need to view risk through a one-year prism — a risk defined in terms of parameter uncertainty ends up being quantified in a manner where parameter uncertainty plays a surprisingly small role. We find that parsimonious models with realistic allowance for trends in the data tend to produce lower mis-estimation capital requirements.

Acknowledgments

The author thanks Gavin Ritchie, Kai Kaufhold, Stefan Ramonat and Torsten Kleinow for helpful comments on earlier drafts. Any errors or omissions remain the responsibility of the author. All models, simulations and mis-estimation calculations were performed with Longevitas, while quantiles were estimated using the `Hmisc` package in R Core Team [2017]. Graphs were done in tikz and pgfplots, while typesetting was done in LaTeX.

References

E. Barbi, F. Lagona, M. Marsili, J. W. Vaupel, and K.W. Wachter. The plateau of human mortality: Demography of longevity pioneers. *Science*, 2018;360:1459–1461, 2018. ISSN 0036-8075. doi: 10.1126/science.aat3119.

R. E. Beard. Note on some mathematical mortality models. In G. E. W. Wolstenholme and M. O'Connor, editors, *The Lifespan of Animals*, pages 302–311. Little, Brown, Boston, 1959.

M. Börger. Deterministic shock vs. stochastic value-at-risk: An analysis of the Solvency II standard model approach to longevity risk. *Blätter DGVFM*, 31:225–259, 2010. doi: https://doi.org/10.1007/s11857-010-0125-z.

D. R. Butenhof. *Programming with POSIX Threads*. Addison-Wesley, Boston, 1997. ISBN 978-0-201-63392-4.

A. J. G. Cairns. Descriptive bond-yield and forward-rate models for the British government securities' market. *British Actuarial Journal*, 4(2):265–321 and 350–383, 1998.

CMI Ltd. *Graduations of the CMI SAPS 2004–2011 mortality experience based on data collected by 30 June 2012 — Final "S2" Series of Mortality Tables*. CMI Ltd, 2014.

D. R. Cox and D. V. Hinkley. *Theoretical Statistics*. Chapman and Hall, 1996. ISBN 0-412-16160-5.

B. Gompertz. The nature of the function expressive of the law of human mortality. *Philosophical Transactions of the Royal Society*, 115:513–585, 1825.

F. E. Harrell and C. E. Davis. A new distribution-free quantile estimator. *Biometrika*, 69:635–640, 1982. ISSN 00063444. doi: https://doi.org/10.1093/biomet/69.3.635. URL `http://www.jstor.org/stable/2335999`.

S. Horiuchi and A. J. Coale. Age patterns of mortality for older women: an analysis using the age-specific rate of mortality change with age. *Mathematical Population Studies*, 2(4):245–267, 1990.

C. M. Jarque and A. K. Bera. A test for normality of observations and regression residuals. *International Statistical Review*, 55(2):163–172, 1987.

E. L. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53:457–481, 1958.

T. Kleinow and S. J. Richards. Parameter risk in time-series mortality forecasts. *Scandinavian Actuarial Journal*, 2016(10):1–25, 2016. doi: https://doi.org/10.1080/03461238.2016.1255655.

A. S. Macdonald, S. J. Richards, and I. D. Currie. *Modelling Mortality with Actuarial Applications*. Cambridge University Press, Cambridge, 2018. ISBN 978-1-107-04541-5.

W. M. Makeham. On the law of mortality and the construction of annuity tables. *Journal of the Institute of Actuaries and Assurance Magazine*, 8:301–310, 1860.

P. McCullagh and J. A. Nelder. *Generalized Linear Models*, volume 37 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, London, second edition edition, 1989. ISBN 0-412-31760-5.

S. J. Newman. Errors as a primary cause of late-life mortality deceleration and plateaus. *PLoS Biology*, 16(12):e2006776, 2018. doi: https://doi.org/10.1371/journal.pbio.2006776.

W. Perks. On some experiments in the graduation of mortality statistics. *Journal of the Institute of Actuaries*, 63:12–40, 1932.

R. Plat. One-year value-at-risk for longevity and mortality. *Insurance: Mathematics and Economics*, 49(3):462–470, 2011. doi: https://doi.org/10.1016/j.insmatheco.2011.07.002.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2017. URL `https://www.R-project.org/`.

S. J. Richards. A handbook of parametric survival models for actuarial use. *Scandinavian Actuarial Journal*, 2012 (4):233–257, 2012. doi: https://doi.org/10.1080/03461238.2010.506688.

S. J. Richards. Mis-estimation risk: measurement and impact. *British Actuarial Journal*, 21 (3): 429–457, 2016. doi: 10.1017/S1357321716000040.

S. J. Richards. A Hermite-spline model of post-retirement mortality. *Scandinavian Actuarial Journal*, 2020:2:110–127, 2020. doi: 10.1080/03461238.2019.1642239.

S. J. Richards, K. Kaufhold, and S. Rosenbusch. Creating portfolio-specific mortality tables: a case study. *European Actuarial Journal*, 3 (2):295–319, 2013. doi: https://doi.org/10.1007/s13385-013-0076-6.

S. J. Richards, I. D. Currie, T. Kleinow, and G. P. Ritchie. Longevity trend risk over limited time horizons. *Annals of Actuarial Science*, 24:1–16, 2020. doi: 10.1017/S174849952000007X.

Appendices

A Description of portfolio and data preparation

We have individual records for survivors and deaths in a local-authority pension scheme in England & Wales, and we follow the data-preparation steps outlined in Macdonald et al. [2018, Chapter 2]. The data fields available for each record are as follows: date of birth, gender, commencement date, total annual pension (either at death or at the date of extract), end date, postcode, National Insurance (NI) number, employer sub-group and whether the pensioner was a child, main life or widow(er) (C, M or W, respectively). The end date was determined differently for deaths, temporary pensions, trivial commutations and survivors to the extract date. For deaths, the end date was the date of death. For children's pensions and trivial commutations the end date was the date the pension ceased or was commuted. For the other survivors, the end date was the date of extract at the end of April 2010. To avoid bias due to delays in reporting of deaths, only the experience data to end-2009 was used. A check of death counts suggested that the earliest usable start date for the experience data would be late spring 2000. However, in order to balance the seasons, we start the exposure period on 1st January 2001 and end on 31st December 2009. To use an exposure period with unequal representation of each season we would need to incorporate a seasonal term in the model [Richards, 2020, Section 8].

There were 55,169 benefit records available before deduplication, of which 21 were rejected due to corrupted dates. Of the remaining 55,148 records, 12,832 were marked as deaths. However, life-office annuitants often have multiple benefits and this phenomenon is also present in pension schemes. In both cases it is necessary to identify the individual lives behind the benefit records to ensure the validity of the independence assumption for statistical modelling. For this we need a process of deduplication [Macdonald et al., 2018, Section 2.5], and for this portfolio we used two composite deduplication keys: the first was a combination of date of birth, gender and postcode (which identified 1,814 duplicates) and the second was a combination of date of birth, gender and National Insurance number (which identified a further 191 duplicates). The highest number of records for a single individual was 7. A particular business benefit of deduplication lies in creating a more accurate picture of the liability for each life: during deduplication the total pension across linked records is summed. There were no instances where the alive-dead status was in conflict among merged records. After deduplication we had 53,143 lives, of which 14 had zero exposure due to ending on the commencement date. This gave 53,129 lives, of which 12,510 were deaths (53,129=55,169-21-1,814-191-14). The resulting data volumes are shown in Figure 11.

Figure 11: Distribution of deaths (top) and time lived (bottom) for 2001–2009 after data validation and deduplication.

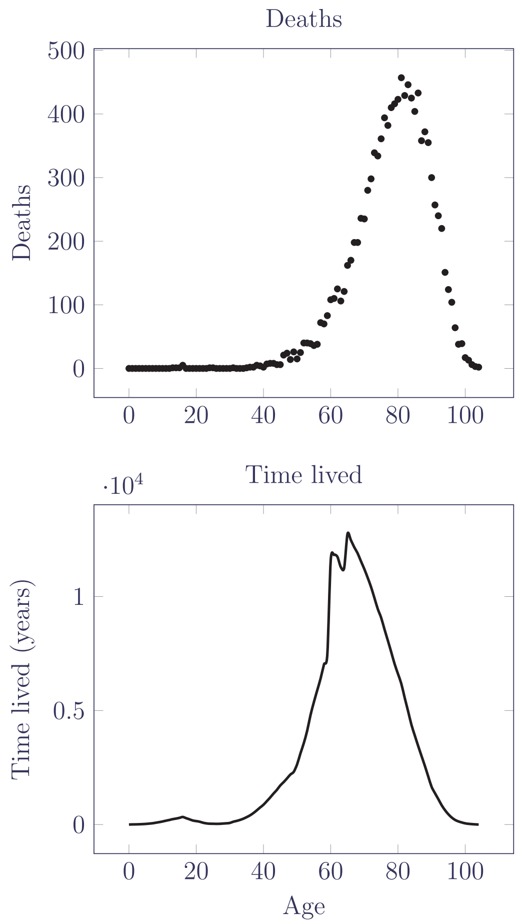

Table 5: Data by pension decile. Pensions to early terminations are revalued at 2.5% p.a. to the end of 2009. The impact of trivial commutations can be seen in the reduced exposure time for the decile of the smallest pensions, S01.

SizeBand	Revalued pension p.a. (£) From...	...to	Lives	Deaths	Exposure (years)	Pensions (£million)	Percentage of total scheme pension
S01	0.00	537.87	5,314	1,090	52,332.8	1.6	0.7%
S02	537.87	963.68	5,313	1,595	70,624.0	4.0	1.8%
S03	963.68	1,464.66	5,313	1,538	70,053.0	6.4	3.0%
S04	1,464.66	2,063.94	5,313	1,507	73,249.5	9.3	4.3%
S05	2,063.94	2,763.69	5,313	1,431	75,705.0	12.7	5.9%
S06	2,763.69	3,602.49	5,313	1,388	77,726.7	16.8	7.8%
S07	3,602.49	4,649.55	5,313	1,174	77,333.2	21.8	10.1%
S08	4,649.55	6,202.61	5,313	1,106	75,465.2	28.5	13.2%
S09	6,202.61	9,009.18	5,313	959	71,680.9	39.4	18.2%
S10	9,009.18	104,751.71	5,311	722	67,329.9	75.3	34.9%
Total			53,129	12,510	711,500.2	215.9	100.0%

Pension-scheme benefits in the UK are increased from year to year. This creates a potential bias problem for cases which terminated in the more-distant past, i.e. deaths and temporary pensions. To put all pension values on the same footing, we need to revalue the pension amounts for earlier terminations. Unfortunately, the formula is exceptionally complex and affects different tranches of benefit accumulated at different times. We therefore opted for a broad-brush approach and revalued early terminations by 2.5% per annum from the date of termination to the end of the period of observation (the Retail Prices Index RPIJ increased by a geometric average of 2.49% over this period).

To check for corruption of records related to paying benefits to surviving spouses, Macdonald et al. [2018, Section 2.10] recommend plotting the Kaplan-Meier survival curves for males and females. Such corruption often goes undetected by traditional actuarial comparisons against a standard table, and Macdonald et al. [2018, Figure 2.8] give an example of a UK annuity portfolio that demonstrates this kind of problem (it is also known to occur in occupational pension schemes). However, Figure 12 shows a clean separation of curves with the expected shape, so there is no such issue for the records of this pension scheme.

Figure 12: Kaplan-Meier survival curves from age 60 using formula from Richards [2012, Section 11]. Experience data 2001–2009.

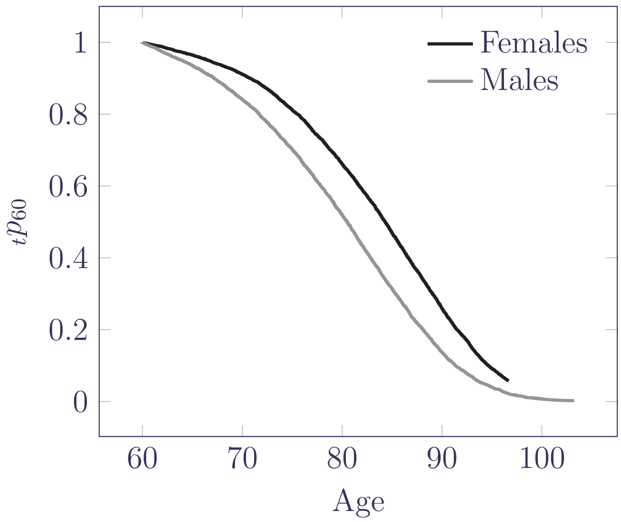

Contact

More information including case studies, latest features, technical documentation and demonstration videos can be found on our website at www.longevitas.co.uk

24a Ainslie Place, Edinburgh, EH3 6AJ
Telephone 0131 315 4470
Email info@longevitas.co.uk

LONGEVITAS™