


ALLOWING FOR SHOCKS IN PORTFOLIO MORTALITY MODELS

By S. J. Richards

Mortality levels in time for UK annuity portfolio.



Allowing for shocks in portfolio mortality models

Richards, Stephen J.*

August 2, 2021

Abstract

The covid-19 pandemic creates a challenge for actuaries analysing experience data that includes mortality shocks. Without sufficient local flexibility in the time dimension, any analysis will be biased by the temporarily higher mortality. Also, depending on where the shocks sit in the exposure period, any attempt to identify mortality trends will be distorted. We present a methodology for analysing portfolio mortality data that offers local flexibility in the time dimension. The approach permits the identification of seasonal variation, mortality shocks and late-reported deaths (OBNR). The methodology also allows actuaries to measure portfolio-specific mortality improvements. Finally, the method assists actuaries in determining a representative mortality level for long-term applications like reserving and pricing, even in the presence of mortality shocks. Results are given for a mature annuity portfolio in the UK, which suggest that the Bayesian Information Criterion (BIC) is better for actuarial model selection than the AIC.

Keywords: mortality trends, seasonal variation, covid-19, B -splines, OBNR.

1 Introduction

The covid-19 pandemic [The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team, 2020] creates the need to allow for mortality shocks in experience analysis performed by actuaries. The intense nature of the repeated covid-19 mortality shocks in many countries means that traditional methods based around annual q_x -style mortality rates are inadequate: mixing periods of shock and non-shock mortality under-states the true intensity of mortality spikes. Furthermore, mortality shocks in the experience data may lead to bias in bases for reserving and pricing. A continuous-time methodology that copes with sharp fluctuations in time is required.

This paper covers the modelling and analysis of portfolio experience data only. The subject of future mortality trends is out of scope, although Section 9 briefly considers the topic of year-on-year improvements. In our models we make extensive use of splines, which are flexible mathematical functions. There are numerous different kinds of spline, each with different properties and thus suitable for different purposes. In this paper we use two kinds of spline: Hermite splines, which span the interval $[0, 1]$, and Schoenberg [1964] splines, which are piecewise local polynomials on the real line. In modern literature, such as de Boor [2001] and Eilers and Marx [2021], references to B -splines are synonymous with Schoenberg splines. Since this paper uses two different kinds of spline, we use the term Schoenberg spline, rather than B -spline, to distinguish from Hermite splines.

Past mortality modelling using splines in calendar time, such as Currie et al. [2006] and Eilers et al. [2008], uses stratified grouped counts. In contrast, we use individual records in a survival model so that we can include covariates like gender, pension size and other factors typically available to

*stephen@longevity.co.uk, Longevity Ltd, Edinburgh. www.longevity.co.uk. <http://orcid.org/0000-0001-6859-6347>


actuaries. We use Hermite splines to model age- and income-related mortality and we use Schoenberg [1964] splines to model mortality in time. Here we compress the knot spacing for splines in time to half a year (or less) to model sharp swings in mortality.

The plan of the rest of this paper is as follows: Section 2 presents important features of covid-19 mortality in the U.K. and the resulting need for continuous-time methods in place of annual q_x rates; Section 3 describes the data set used in the paper; Section 4 describes the use of Hermite splines for modelling mortality by age, while Section 5 describes how to also use Hermite splines for modelling mortality by annuity amount; Section 6 recaps the application of Schoenberg splines for modelling mortality levels over time. Section 7 considers the ability of the method to identify seasonal variation, while Section 8 looks at the additional requirements for modelling the covid-19 mortality shocks in the UK. Section 9 looks at further insights that can be derived, such as portfolio-specific improvement rates. Section 10 considers the use of the methodology to allow for the impact of occurred-but-not-reported (OBNR) deaths. Section 11 considers the conditions under which Schoenberg [1964] splines may — and may not — be used for mortality patterns by age. Section 12 discusses the use and limitations of information criteria, and the role of actuarial judgement in selecting models; Section 13 concludes.

2 Covid-19 and other mortality shocks

Covid-19 is a new viral disease [The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team, 2020] whose arrival in the UK in 2020 caused deaths to surge to levels not seen since the global influenza pandemic of 1918–1920 [Spreeuwenberg et al., 2018]. Figure 1 shows that England & Wales had had sharply higher numbers of deaths in 1918 and 2020 compared to preceding years.

Figure 1: Numbers of deaths in England & Wales (2020 count is provisional). Source: ONS.



In many places the 1918–1920 influenza pandemic appeared as a series of three spikes in mortality, as shown in Figure 2 for Scotland, where the second and third spikes were the most severe.

As with influenza in 1918–1920, a key feature of Covid-19 mortality in many countries is that it too takes the form of relatively sharp peaks. At the time of writing there have been two such peaks in the UK, as shown in Figure 3. These peaks form and subside over a period of weeks or months, and an annualised approach to mortality will under-state such mortality surges. For this reason, it is important to use continuous-time methods like the mortality hazard, μ_x , rather than annual q_x -style rates.

Figure 2: Weekly deaths in Scotland in 1918–1919 as percentage of 1913–1917 average. Source: Craufurd Dunlop and Watt [1915, 1916a,b, 1918, 1919, 1920a,b].


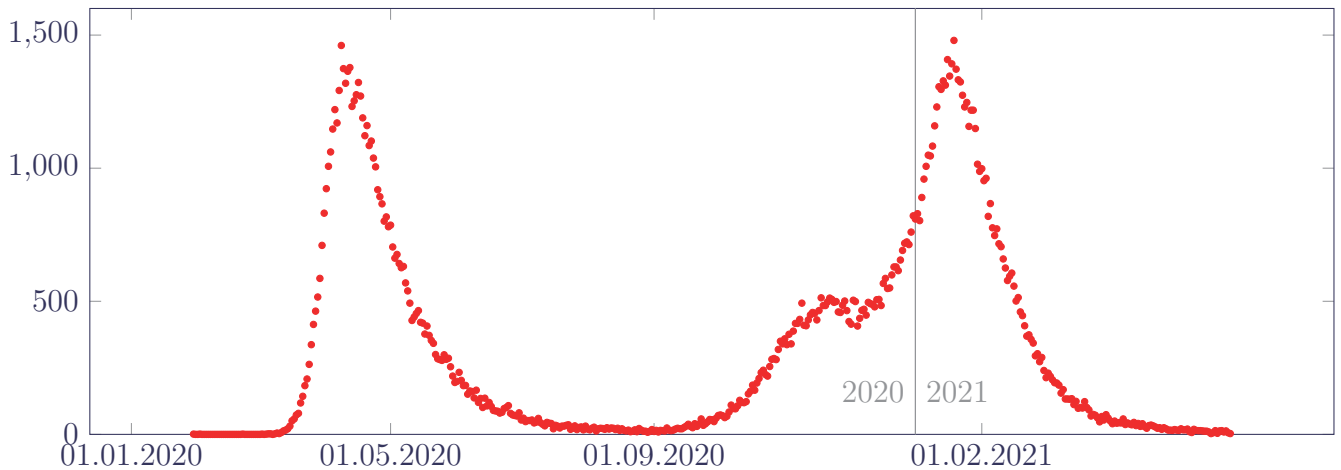


Figure 3: Daily deaths in UK where the death certificate mentions covid-19 as one of the causes. Note that the first lockdown legally commenced in England on 26th March 2020 [Hancock, 2020]. Source: ONS [2021].



Covid-19 made its presence felt in the mortality of annuity portfolios and pension schemes in many countries; Richards [2021] demonstrated this for annuity portfolios in France, the UK and the USA. Covid-19 therefore creates a problem for actuaries analysing experience data to set bases: the extra mortality of 2020–2021 risks either under-reserving (in the case of liabilities already on the balance sheet) or under-pricing (in the case of insurers writing bulk annuities and longevity swaps). Unfortunately, portfolio administrators seldom record the cause of death, so data like Figure 3 are typically unavailable. Excluding periods of covid-affected data is also an unsatisfactory approach — pension schemes looking to transact bulk annuities or longevity swaps may only have experience data for the most recent 3–5 years, and discarding one or more of those years’ data is an unaffordable luxury for a pricing actuary. Actuaries therefore require a methodology that works with all-cause mortality data, but which flexibly tracks mortality levels over time to allow for covid-19 spikes. This way all available experience data can be used, and the actuary can then exercise judgement as to what point in time is most representative for calibrating a mortality basis.

3 Data description

The data set used in this paper comprises individual records from a mature annuity portfolio of a UK insurer. At the end of June 2021 a direct extract was made from the administration systems, as recommended by Macdonald et al. [2018, Section 2.2] to get the most up-to-date data. Policy records were validated using the checks described in Macdonald et al. [2018, Sections 2.3 and 2.4].

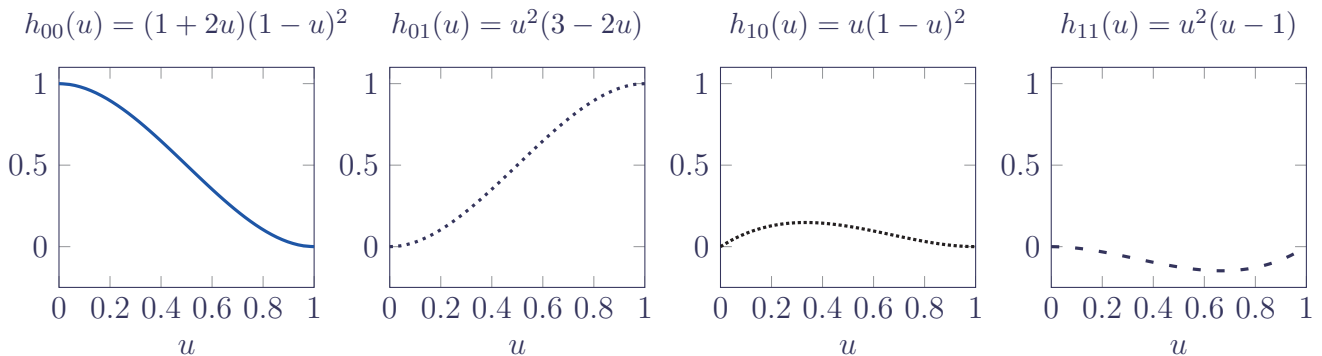
Insurer annuity portfolios often contain two or more policies per person, and so deduplication is required for the independence assumption that underpins statistical modelling; see Macdonald et al. [2018, Section 2.5] for discussion of various approaches to deduplication. As is common for UK annuity portfolios, there are many such duplicates — of 351,947 records passing validation, 124,420 records were found to be someone’s second or third annuity, and there were three people with 29 annuities each. The tendency to have multiple annuities is correlated with wealth and socio-economic status, so deduplication is an essential step in building a statistical model for actuarial purposes. There were 729 annuities where the annuitant was marked alive on one annuity and dead on another; these annuities were excluded from the data for modelling.

The data set used in this paper is an updated extract of the UK3 data set in Richards [2021]. In this paper we will continue to refer to it as UK3 for continuity.

4 Hermite splines for mortality by age

A basis of Hermite splines in one dimension [Kreyszig, 1999, p868] is a collection of four cubic polynomial functions, as shown in Figure 4. A basis of cubic Hermite splines will produce the same fitted curve as a basis of cubic Bézier curves or Bernstein polynomials of degree 3, albeit with different coefficients.

Figure 4: Four Hermite basis splines for $u \in [0, 1]$. Note that $h_{00}(u) + h_{01}(u) = 1$.



Richards [2020b] proposed using Hermite splines for modelling mortality by mapping an age range $[x_0, x_1]$ onto $[0, 1]$ and forming the log hazard as a linear combination of the spline functions in Figure 4. The model for the mortality hazard at age x , μ_x , is given in equation (1):

$$\begin{aligned} \log \mu_x &= \alpha h_{00}(u) + \omega h_{01}(u) + m_0 h_{10}(u) + m_1 h_{11}(u) \\ u &= \frac{x - x_0}{x_1 - x_0} \end{aligned} \tag{1}$$

where the intermediate variable u maps age $x \in [x_0, x_1]$ onto $[0, 1]$ and α , ω , m_0 and m_1 are parameters to be estimated. In practice the h_{11} spline is seldom needed for mortality work and so we often set $m_1 = 0$. Doing so forces the mortality hazard in equation (1) to be monotonic at advanced ages,

a topic we will return to in Section 11. We can further set $m_0 = 0$ for a simple two-parameter alternative to the Gompertz [1825] model that is strictly monotonic at all ages, as per equation (2):

$$\log \mu_x = \alpha h_{00}(u) + \omega h_{01}(u) \quad (2)$$

Actuarial survival data are typically left-truncated; see Macdonald et al. [2018, Section 1.9] for a discussion of the differences between actuarial survival models and those in medical research, such as described by Collett [2003]. To fit a survival model for n individual lifetimes we use the log-likelihood, ℓ , in equation (3):

$$\ell = - \sum_{i=1}^n H_{x_i}(t_i) + \sum_{i=1}^n d_i \log \mu_{x_i+t_i} \quad (3)$$

$$H_{x_i}(t_i) = \int_0^{t_i} \mu_{x_i+s} ds$$

where x_i is the exact age when life $i = 1, \dots, n$ enters observation, t_i is the time in years that life i is observed and d_i is an indicator variable taking the value 1 if life i is dead at age $x_i + t_i$ and zero otherwise; $H_{x_i}(t_i)$ is the integrated hazard function. The log-likelihood in equation (3) is maximised to find the MLEs of the parameters underlying μ_x ; see Appendix A for the technical details of implementation.

Figure 5(a) shows the simplified Hermite-spline model of equation (2) fitted to the UK3 data set. Figure 5(b) shows the deviance residuals [McCullagh and Nelder, 1989, p39], indicating that Hermite splines are an effective approach to modelling post-retirement mortality, especially where the linear assumption of Gompertz [1825] does not hold.

Figure 5: (a) Crude mortality hazard (\circ) for UK3 data set with fitted curve ($-$) from equation (2). (b) Deviance residuals (\bullet) from fit, which are broadly consistent with random $N(0,1)$ i.i.d. variates apart from two negative residuals below -3 . Ages 55–105, 2015–2020.

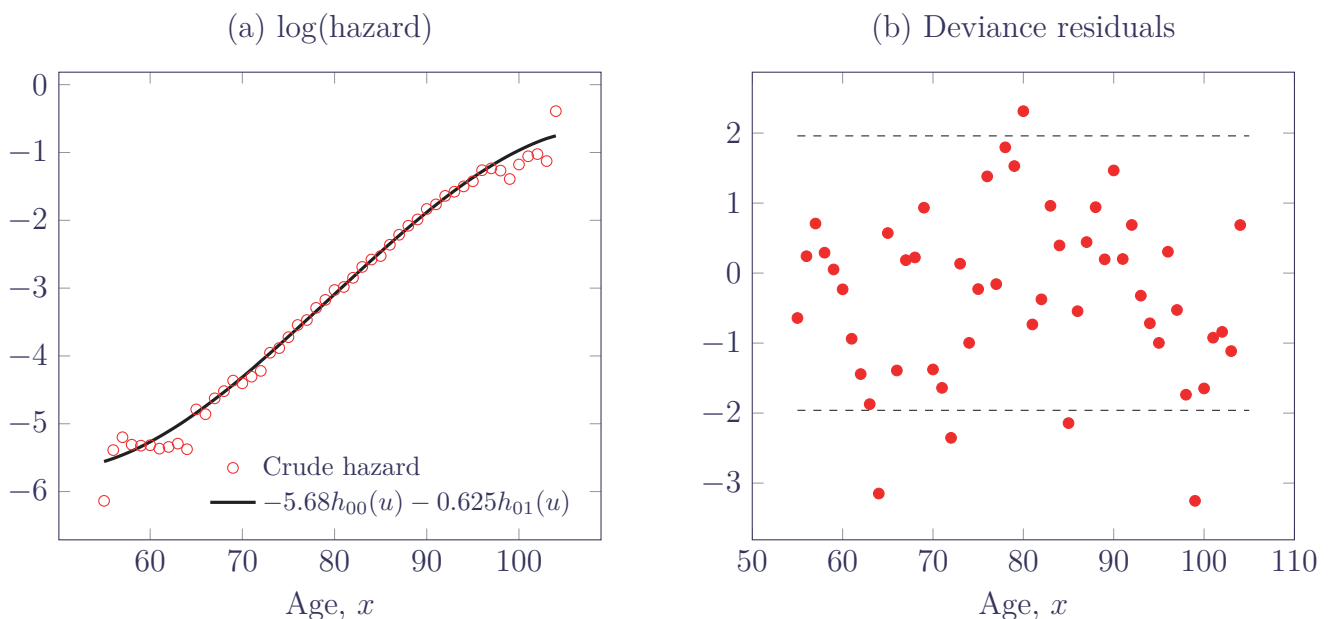


Figure 5 shows that equation (2) does an acceptable job of explaining mortality variation by age. We therefore extend equation (2) to include gender as a risk factor as shown in equation (4), where $z_{i,\text{female}}$ is an indicator variable taking the value 1 if the annuitant is female and zero if male. α_0 , α_{female} , ω_0 and ω_{female} are parameters estimated by maximum likelihood using equation (3).

$$\begin{aligned}\log \mu_{x_i} &= \alpha_i h_{00}(u) + \omega_i h_{01}(u) \\ \alpha_i &= \alpha_0 + \alpha_{\text{female}} z_{i,\text{female}} \\ \omega_i &= \omega_0 + \omega_{\text{female}} z_{i,\text{female}}\end{aligned}\quad (4)$$

5 Hermite splines for mortality by annuity amount

Figure 5 shows that even just using two of the Hermite basis splines works well when fitting a mortality curve by age. This is because mortality rates tend to gradually and monotonically increase with age after retirement. Figure 6 shows the deviance residuals by size-band, which suggest a more complicated relationship between annuity level and mortality. However, the lower-than-expected mortality for the smallest size-bands is actuarially irrelevant: the five size-bands covering the decile of lives with the smallest annuities account for less than 0.5% of total annuity payments. We are therefore only interested in the general downward trend in mortality with increasing annuity amount, and especially the sharply lower mortality of those receiving the very largest annuities. As with age, we can again consider mortality by annuity amount to be monotonic without material loss of actuarial applicability.

Figure 6: Deviance residuals by size-band for the model specified by equation (4), i.e. after accounting for variation by age and gender. Each of the fifty size-bands contains around 2% of the lives in the UK3 portfolio. The dashed lines show the 95% confidence limits for $N(0,1)$ variates.


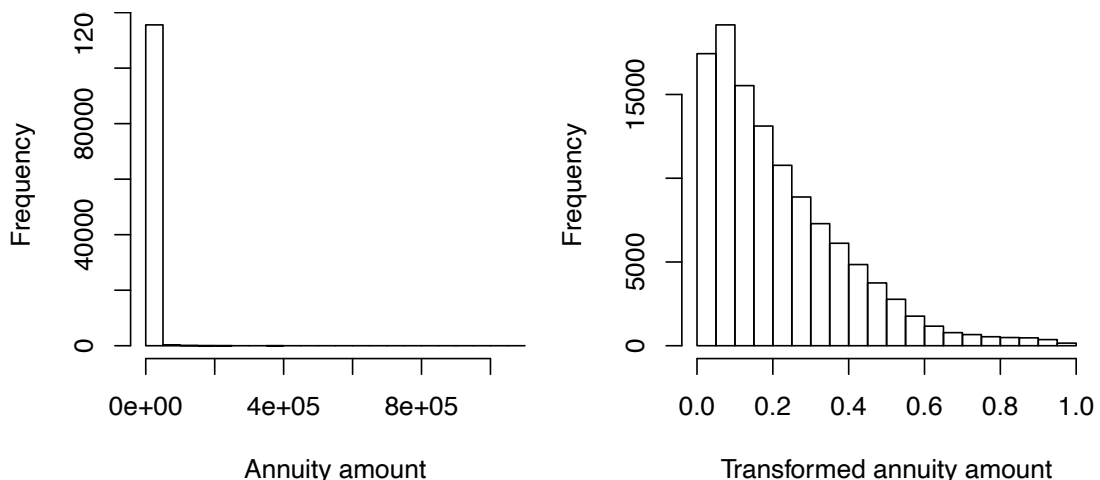


Figure 7: Histograms of annuity amounts. The left panel shows the untransformed amounts, a_i , displaying extreme kurtosis. The right panel shows the amounts transformed by $a_i e^{-8.58082} / (1 + a_i e^{-8.58082})$, showing a marked reduction in kurtosis. The value of $\hat{\lambda}_0 = -8.58082$ comes from the model in Table 1, where it was estimated with reference to the mortality characteristics of the UK3 portfolio.



We adopt a two-part approach: (i) a monotonic transform of the annuity amount onto the interval $[0, 1]$, and (ii) modelling mortality by the transformed amount. For the first step we transform the

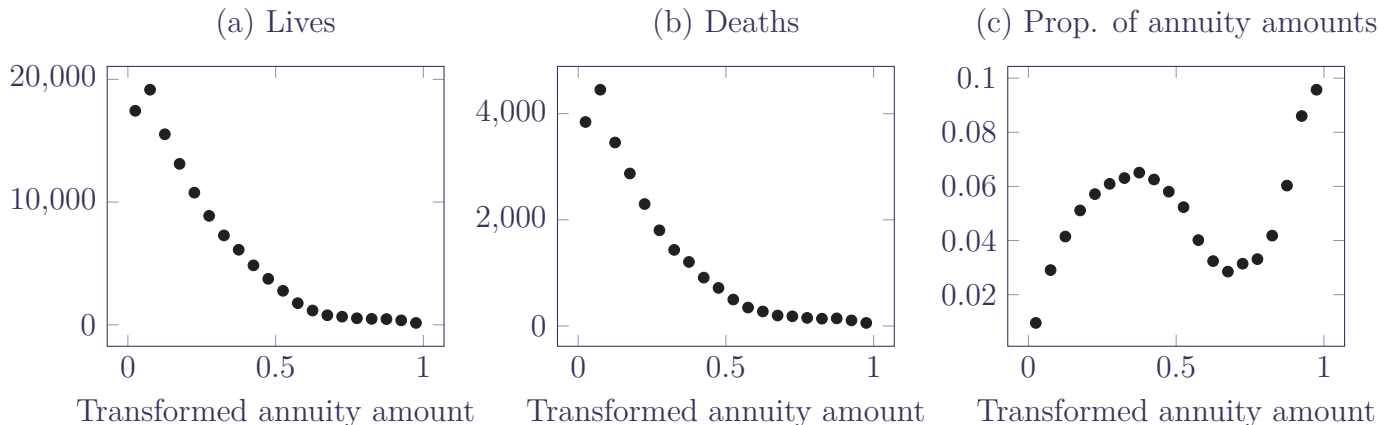
annuity amount, $a \geq 0$ onto the interval $[0, 1]$ so that we can use Hermite splines in the second step. We use a parameterised transform function, $\tau(a, \lambda_0)$, defined in equation (5):

$$\tau(a, \lambda_0) = \frac{ae^{\lambda_0}}{1 + ae^{\lambda_0}} \quad (5)$$

The operation of the transform function in equation (5) is demonstrated in Figure 7, where the parameter λ_0 allows the transform function to adapt simultaneously to the kurtosis of the portfolio and the mortality effect of annuity amount.

The transformation in the right panel of Figure 7 leads to a deliberately unequal distribution of lives and deaths, as shown in Figures 8(a) and (b). However, a useful actuarial consequence is that the unequal distribution focuses on the financially significant annuities: the interval $[0.95, 1]$ contains the 155 largest annuities with 9.57% of total annuity amounts, whereas the similar-sized interval $[0, 0.05]$ has the 18,437 smallest annuities with just 0.96% of the total annuity amounts.

Figure 8: Lives, deaths and proportion of annuity amounts by transformed annuity amount.



We treat a zero annuity amount as the baseline and use the single h_{01} Hermite spline to model the change (reduction) of mortality with increasing annuity amount on the transformed scale. We thus further extend the model in equation (4) to include a continuous allowance for annuity amount, as in equation (6):

$$\begin{aligned} \log \mu_{x_i} &= \alpha_i h_{00}(u) + \omega_i h_{01}(u) \\ \alpha_i &= \alpha_0 + \alpha_{\text{female}} z_{i,\text{female}} + \omega_{\text{amount}} h_{01}(\tau(a_i, \lambda_0)) \\ \omega_i &= \omega_0 + \omega_{\text{female}} z_{i,\text{female}} \end{aligned} \quad (6)$$

The parameter ω_{amount} represents the maximum mortality reduction from a near-infinite annuity income relative to zero income. In equation (6) τ is the logistic transform in equation (5), although Richards [2020a] considers some other transform options.

The above approach to mortality varying by annuity amount means that the amounts effect can be handled statistically without discrete size-bands. This avoids discretisation error and also means that we can extrapolate mortality effects for annuity amounts above the upper limit in the calibrating data set. This latter aspect is particularly useful when calibrating a pricing basis, since an insurer may encounter pension sizes during quotation that are far above the largest pension amount in its own experience data. Table 1 shows the parameter estimates for the UK3 portfolio; it is a parsimonious model accounting for three age-varying risk factors with just $p = 6$ parameters for $n = 116,056$ lives.

The information criterion from Akaike [1987] is shown ($AIC = -2\ell + 2p$), along with the Bayesian Information Criterion [Schwarz, 1978] ($BIC = -2\ell + p \log n$). For further details of this continuous approach to the amounts effect, see Richards [2020a]. Alternatively, see van Berkum et al. [2020] for a different approach using thin-plate splines.

Table 1: Parameter estimates for UK3 portfolio, together with numbers of contributing lives and deaths. The model is specified in equation (6). Ages 55–105, 1st January 2015 to 31st December 2020. AIC=187,693, BIC=187,751.

Parameter	Name	Estimate	Std. err.	Z-value	Lives	Deaths
α_0	Intercept	-5.45242	0.03244	-168.08	116,056	25,081
α_{female}	Gender.F	-0.60228	0.05370	-11.22	34,875	6,429
ω_{female}	Gender.F:Oldest	-0.20897	0.03787	-5.52	34,875	6,429
ω_0	Oldest	-0.53007	0.02021	-26.23	116,056	25,081
ω_{amount}	AmountUltimate	-0.66319	0.11206	-5.92	116,056	25,081
λ_0	AmountTransformParameter	-8.58082	0.27120	-31.64	116,056	25,081

Figure 9 shows the deviance residuals by transformed annuity amount before and after fitting the continuous amounts factor. The overall fit of the model has improved, as measured by the AIC, BIC and χ^2 test statistic by size-band. However, we have made an important trade-off in quality of fit: improved fit for the larger annuity amounts comes at the cost of a worsened fit for the smallest annuity amounts. Actuarially we are comfortable with this, as Figure 8(c) shows that the worsened fit accounts for just 0.96% of the annuity amounts, and is thus a tiny proportion of overall liabilities. Table 2 shows the development of model fit from adding age, gender and annuity amount as mortality risk factors.

Figure 9: Deviance residuals, $\{r_m : m = 1, \dots, 20\}$, by transformed annuity amount. The deviance residuals are calculated for twenty intervals each of length 0.05, which have very different numbers of lives and deaths, as shown in Figure 8. Note that the grouping here was performed purely for the purpose of residual calculation, and that the underlying mortality model by annuity amount is fully continuous. The dashed lines show the 95% confidence limits for $N(0,1)$ variates.

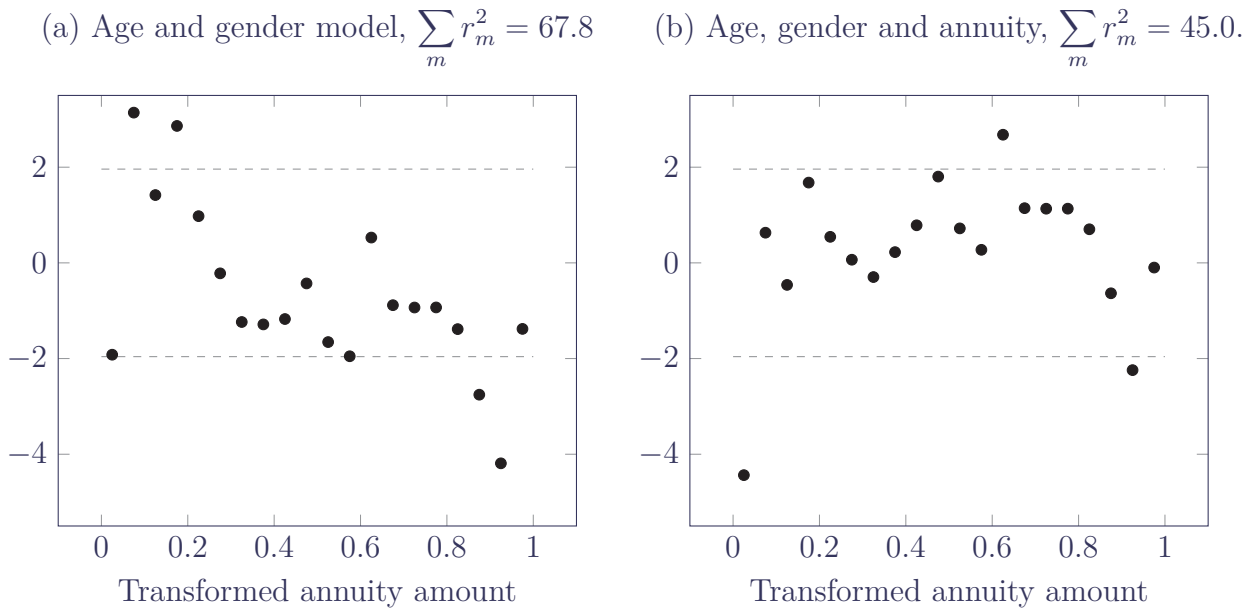


Table 2: Information criteria (ICs) from stepwise inclusion of age, gender and annuity amount as risk factors.

Risk factors	Equation number	Parameter count	AIC	BIC	$\sum_m r_m^2$
Age only	(2)	2	188,413	188,433	60.7
Age and gender	(4)	4	187,774	187,813	67.8
Age, gender and annuity amount	(6)	6	187,693	187,751	45.0

Mortality levels tend to vary monotonically by age. They also can be approximated as such for annuity amount, even if the fit for the very smallest annuities is not ideal. Thus, mortality by age and annuity amount can both be modelled using Hermite splines once each continuous variable is mapped onto the real interval $[0, 1]$. However, modelling mortality variation in calendar time is not monotonic, for which we need another kind of spline.

6 Schoenberg splines for mortality levels over time

Mortality levels can fluctuate sharply over short periods of time, as shown in Figure 10, where seasonal variation is evident along with the Covid-19 shock in April 2020. Unlike mortality by age and annuity amount, patterns of mortality in time are not monotonic. We therefore require a method with sufficient local flexibility to reflect the rise and fall of mortality levels over time. For this we use the splines of Schoenberg [1964]. Schoenberg’s splines are based on polynomials of degree $bdeg$ spanning $bdeg + 2$ knot points. The knot points are unique values on the real line and in many applications they are uniformly spaced. However, uniform knot spacing is not mandatory [Kaishev et al., 2016] and Section 8 will demonstrate an application of unequally spaced knots.

Outside of a spline’s starting and end knots the spline takes the value zero, making it a purely local function. Figure 11 shows four Schoenberg splines of varying degrees, and more detail on such splines can be found in de Boor [2001] and Eilers and Marx [2021].

Schoenberg splines are not new in mortality modelling: Eilers et al. [2004] applied them to modelling mortality trends in grouped counts. This is done by forming a basis of cubic Schoenberg splines (B -splines) in calendar time, as shown in Figure 12.

We define $B_j(y)$ as the j^{th} basis spline evaluated at time y and define $\mu_{x,y}$ as the mortality hazard at exact age x and calendar time y . We can then use a B -spline basis as in Figure 12 to form an age-period model for $\mu_{x,y}$ as per equation (7):

$$\log \mu_{x,y} = \log \mu_x + \sum_{j \geq 1} \kappa_{0,j} B_j(y) \quad (7)$$

Figure 10: Mortality level in time for a mature US annuity portfolio using a semi-parametric estimator. The vertical dotted line indicates 1st April 2020. Source: Richards [2021, Figure 3(c)].




Figure 11: Schoenberg [1964] splines of degree 0, 1, 2 and 3 with one-year spacing between knot points (marked \bullet). In this specimen example, each spline is zero before the knot at 2015 and zero above the knot $bdeg + 1$ knots to the right. The area under each spline is 1.





Figure 12: A basis of nine equally-spaced cubic B -splines spanning 1st January 2015 to end-2020, indexed $j = 0, 1, \dots, 8$. Splines in solid black lie completely within the period, while splines in dashed grey are edge splines that only partly lie in the period being spanned. Knot points are marked \bullet . Note that at any time-point $y \in [2015, 2021]$, there are always four non-zero splines that sum to 1. This is not true outside $[2015, 2021]$, meaning the method cannot extrapolate outside the calibration interval.



The parameter $\kappa_{0,j}$ corresponds to the j^{th} B -spline, and summation is from $j = 1$ as the spline $j = 0$ is absorbed into the baseline hazard. To estimate the $\kappa_{0,j}$ along with any other parameters, we maximise the log-likelihood in equation (8):

$$\ell = -\sum_{i=1}^n H_{x_i, y_i}(t_i) + \sum_{i=1}^n d_i \log \mu_{x_i+t_i, y_i+t_i} \quad (8)$$

$$H_{x_i, y_i}(t_i) = \int_0^{t_i} \mu_{x_i+s, y_i+s} ds$$


where life i enters observation at exact age x_i at exact year y_i , and is observed for exactly t_i years; $H_{x_i, y_i}(t_i)$ is the integrated hazard that now depends on both age at entry and year of entry. Some example estimates of $\kappa_{0,j}$ are given in Table 3, which are then applied to the B_j in Figure 13 to show the local flexibility. Finally, the various $\hat{\kappa}_{0,j} B_j$ products are summed in Figure 14 to show how the estimated mortality level varies over time.

Note that absorbing $\kappa_{0,0}$ into the baseline as in equation (7) is just one way of specifying the necessary identifiability constraint, in this case $\kappa_{0,0} = 0$. That a constraint is required comes from the identity in equation (9):

Table 3: Estimates of $\kappa_{0,j}$ for $j = 1, 2, \dots, 8$ for UK3 portfolio, ages 55–105, 2015 to end-2020. $\kappa_{0,0} = 0$ by construction because it is absorbed into the baseline hazard.

j	1	2	3	4	5	6	7	8
$\hat{\kappa}_{0,j}$	-4.30805	-3.73912	-3.91987	-3.69781	-4.0887	-3.76673	-3.98538	-4.23435

Figure 13: $\hat{\kappa}_{0,j}B_j(y)$ using the nine basis splines in Figure 12 and the estimates in Table 3. Knot points are marked \bullet .




$$\begin{aligned}
 \log \mu_{x,y} &= \alpha_0 h_{00}(u) + \omega_0 h_{01}(u) + \sum_{j \geq 0} (\kappa_{0,j} - c) B_j(y) \\
 &= \alpha_0 h_{00}(u) + \omega_0 h_{01}(u) - c + \sum_{j \geq 0} \kappa_{0,j} B_j(y), \quad \text{as } \sum_{j \geq 0} c B_j(y) = c, \forall c \in \mathbb{R} \\
 &= (\alpha_0 - c) h_{00}(u) + (\omega_0 - c) h_{01}(u) + \sum_{j \geq 0} \kappa_{0,j} B_j(y), \quad \text{as } h_{00}(u) + h_{01}(u) = 1
 \end{aligned} \tag{9}$$

Without an identifiability constraint there is an infinite choice of real-valued c that can be deducted from each $\kappa_{0,j}$ and added to α_0 and ω_0 in equations (1) or (2) and still yield the same fit. As a consequence, the vertical scale in Figure 14(a) is somewhat arbitrary, and depends in large part on the mortality experience covered by the right-hand part of spline B_0 (see for example the spline $j = 0$ in Figure 12 relative to the period from 1st January 2015).

We can use equation (9) to normalise the Schoenberg spline function to take the value zero at a particular point in calendar time. This is useful if we regard that point in time as having a meaning and we want to compare subsequent mortality levels (as is done in Figure 17, for example). For example, 1st October 2019 (2019.75 decimalised) represents the last mid-point between a summer trough and a winter peak before the covid-19 pandemic, and so might be regarded as the most recent suitable timepoint unaffected by seasonal swings, pandemic shocks and unreported deaths. If we define $c_{2019.75} = \sum_{j \geq 1} \hat{\kappa}_{0,j} B_j(2019.75)$ we can then deduct this value from each $\hat{\kappa}_{0,j}$ and normalise the Schoenberg spline function without distorting the model fit as long as we also add $c_{2019.75}$ to the estimates $\hat{\alpha}_0$ and $\hat{\omega}_0$ ($c_{2019.75}$ of course needs to be recalculated whenever the knot points in the spline basis change). Figure 14(b) shows the resulting normalised Schoenberg spline.

Figure 14: Schoenberg time-spline function spanning 1st January 2015 to end-2020 using the basis splines in Figure 12 and the $\hat{\kappa}_{0,j}$ in Table 3. Panel (a) shows the unadjusted spline function, while panel (b) shows the function normalised at zero at 2019.75, the last mid-point between a summer trough and winter peak before the covid-19 pandemic. Note that summation in panel (b) is from $j = 0$ because the coefficient of B_0 is no longer zero.



7 Modelling seasonal variation and covid-19 shocks

Mortality has long been known to have a major seasonal component; Rau [2007, Chapter 2] gives a comprehensive introduction, covering both historical and modern aspects. Seasonal variation is also pronounced in pensioner mortality, as demonstrated by Richards et al. [2020] with a recurring annual cosine term. Such an approach estimates an average seasonal effect within each year, with peak mortality in winter and low mortality in summer. However, this approach cannot account for years with heavier-than-average winter mortality, nor will it account for slight shifts in the timing of the winter peak.

We can address these issues by increasing the number of knots per year. This increases the number of splines needed to span 2015–2020, with the resulting mortality level shown in Figure 15 for two, four and ten knots per year.

Figure 15: $\sum_{j \geq 0} (\hat{\kappa}_{0,j} - c_{2019.75}) B_j(y)$ for y spanning 1st January 2015 to end-2020 with equally-spaced knots.




Figure 15(a) shows that two knots per year does a good job of picking up the seasonal variation with the winter peaks around January of each year 2015–2019. The exception is 2020, where the expected winter peak in January has merged with the covid-19 mortality spike of April and May 2020, resulting in an aggregate peak shifted more towards the time of the first covid-19 shock. Another feature of Figure 15(a) is the deep trough leading up to the end of 2020; this could be either a real feature of the data, say due to the covid-19 mortality being partly due to brought-forward deaths, or else an artefact from the B -spline basis having insufficient flexibility to cope with the short-term intensity of the two covid-19 shocks in Figure 3. To resolve this, we can increase the knot density further, as in Figure 15(b) and (c), but we face a difficult balancing act — increasing the knot density allows the nature of the covid-19 spike to come through, but at the expense of weakening the signal at other times.

Table 4 summarises the model fits using various equally spaced knot densities. According to the AIC, the model with ten knot points per year is the best fit, which is rather contradicted by the fluctuations in Figure 15(c). In contrast, the BIC indicates that two knot points per year is the best fit. The conflicting messages are likely due to the large number of parameters not being properly penalised in the AIC. In mainstream statistical work a small-sample correction to the AIC is often used because of this [Macdonald et al., 2018, page 98]. However, in actuarial work the sample size is typically not the problem due to there being tens of thousands of data points ($n = 116,056$ for the UK3 portfolio here). Rather, the issue is a large number of parameters for a given risk factor — with four or more knot points per year, many of the $\kappa_{0,j}$ parameters prior to the mortality shock are not explaining enough variation to justify their inclusion. (Ye [1998, p120] notes that “flexibility often leads to substantial overfitting”.) A possible explanation for the failure of the AIC here is given by Owen [1991, pp102–103], albeit in the context of a different model: where a statistically significant parameter is being estimated, the estimation costs one degree of freedom; however, if the parameter is statistically insignificant, the estimation costs more than one degree of freedom. Thus, where unnecessary spline parameters are introduced, the standard definition of the AIC may fail to properly penalise this. We will consider this drawback of the AIC again in Sections 11 and 12.

Table 4: Information criteria (ICs) for various knot densities with equal spacing. UK3 portfolio for 1st January 2015 to end-2020. ICs and parameter counts can be compared with Table 2 as the underlying data are identical.

Knots per year	Parameter count	AIC	BIC
1	14	187,594	187,729
2	20	187,412	187,605
4	32	187,324	187,634
10	68	187,244	187,901

8 Increasing the knot density around shock times

Figure 15(c) shows that increasing the density of the knot points allows the covid-19 mortality shock to be clearly identified in terms of height and timing. However, the development from Figure 15(a) to (c) also suggests the introduction of random variation for the non-covid period. It is undesirable to add knot points where flexibility is not required, so we address this by providing extra knots only where they are needed. One approach is to use half-yearly-spaced knots for all years, but to add extra knots around the time of the first covid-19 spike in April and May 2020, as shown in Figure 16. We justify this from our *a priori* knowledge of the population covid-19 mortality in Figure 3.

Figure 17 shows the resulting mortality levels in time, with the expected seasonal fluctuation and the first covid-19 spike around April 2020. The mixed approach with two knots per year for pre-covid years and additional knots for the pandemic shock works well in capturing the salient features in time. Figure 17(b) further shows the usefulness of normalising the Schoenberg splines at zero at a

Figure 16: Part of a basis of nineteen variably-spaced cubic B -splines spanning 1st January 2015 to end-2020, indexed $j = 0, 1, \dots, 18$ (only splines $j = 10, \dots, 18$ are actually shown). Splines in solid black lie completely within the period, while splines in dashed grey are edge splines that only partly lie in the period being spanned. Knot points are marked \bullet .


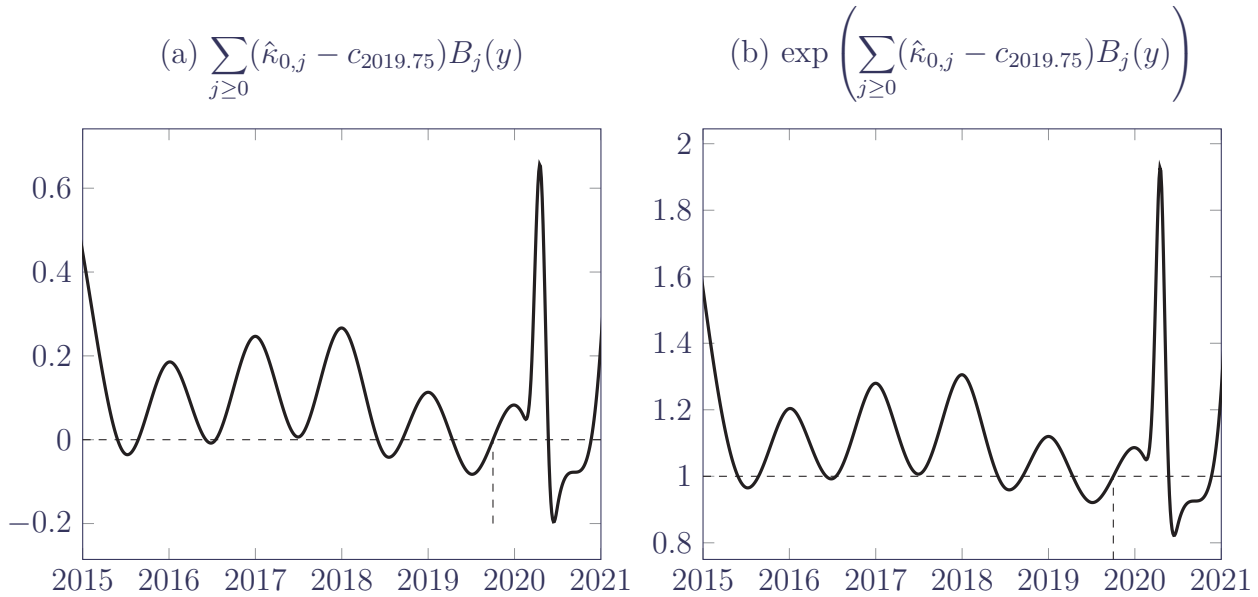


Figure 17: Schoenberg time-spline function for y spanning 1st January 2015 to end-2020 using the spline basis depicted in Figure 16. Panel (a) shows the addition to $\log(\text{mortality})$, while panel (b) shows the multiplier of the mortality hazard.



point in time — when converted to the hazard scale the multiplier is 1 at the reference point (2019.75 in this case) and we can see that the peak of the covid-19 spike is nearly double the reference level of mortality. Following the shock in April 2020 there is also an unusually deep summer trough in 2020, which could be a result of “harvesting” due to brought-forward deaths of the frail. Figure 17(b) also shows that the pre-pandemic seasonal trough-to-peak variation varies between around 15–30%, which is consistent with the results in Richards et al. [2020, Table 2] for a variety of international annuity and pension portfolios.

The BIC for the model behind Figure 17 is 187,489, which is considerably lower than any of the BICs in Table 4. This suggests that the basis of variably spaced knots in Figure 16 has provided flexibility only where it is needed. The time signal is relatively strong: dropping the factors gender and annuity amount from the model leaves Figure 17 largely unchanged, which justifies the “age plus period” nature of equation (7). However, the BIC for the model behind Figure 17 does not take into account the fact that the knots were selected with reference to data and judgement (specifically,

a comparison of the plots in Figure 15 and the distribution of covid-19 deaths in Figure 3). The question of defining and using information criteria is discussed further in Section 12.

There are at least two other alternative approaches that would be possible. The first is the optimisation of the number and position of knots by an algorithm targeting a measure of fit; Kaishev et al. [2016] propose a knot-addition algorithm and also give a comprehensive overview of related knot-optimisation research. The second approach is to deliberately use too many knots and use a tuning parameter to minimise the variability exhibited in Figure 15(c); Eilers and Marx [1996] give details. However, both of these methods have huge computational cost when applied to actuarial data sets, where the number of observations often numbers hundreds of thousands. For example, the UK3 data set in this paper has $n = 116,056$ records and the model behind Figure 17 has 19 Schoenberg splines, giving a knot-to-data ratio of around 1:6000. In contrast, the data set from Kimber et al. [2009] reworked by Kaishev et al. [2016] involves just 1,151 observations and has 227 knots, giving a knot-to-data ratio of 1:5. The model behind Figure 16 took 1.5 hours to fit using parallel processing over 63 threads [Butenhof, 1997], so adding lots more splines — as per Eilers and Marx [1996] — would result in much longer run-times. Similarly, the repeated refitting with new knots — as per Kaishev et al. [2016] — would also be time-consuming when applied to actuarial data sets.

9 Estimating portfolio-specific improvements

An advantage of equation (7) is that it can be used to estimate the portfolio-specific mortality improvement (*PSMI*). This can be estimated by selecting either two winter peaks or two summer troughs at time points $y_1 < y_2$ and using equation (10):

$$PSMI = \left[1 - \exp \left(\frac{\sum_{j \geq 1} \kappa_{0,j} (B_j(y_2) - B_j(y_1))}{y_2 - y_1} \right) \right] \times 100\% \quad (10)$$

Our preference is to use summer troughs, due to the tendency for “sharp peaks in winter and relatively flat troughs in summer” [Marx et al., 2010], meaning that winter peaks can be more extreme and variable. From Figure 17(a) the period 2015.5–2019.5 seems to be the most suitable period for estimating the portfolio-specific improvement rate, as the trough in the summer of 2020 may be unusually deep due to brought-forward deaths. With $y_1 = 2015.5$ and $y_2 = 2019.5$, equation (10) gives $PSMI = \left[1 - \exp \left(\frac{-1.00971 - (-0.96219)}{2019.5 - 2015.5} \right) \right] \times 100\% = 1.18\%$ p.a. This is an age-independent mortality-improvement rate, as equation (7) is an “age effect plus period effect” model.

Table 5 shows the improvement rates for various combinations of summer troughs in Figure 17. Of interest are the strong improvement rates ending in summer 2020 (2020.5), indicating the depth of the mortality dip following the first covid-19 shock in April 2020. (This phenomenon is unrelated to late-reported deaths, which is the subject of the next section.) However, it is an open question what value such volatile improvement rates have when summer troughs can be almost as variable as winter peaks.

Table 5: Annual improvement rates (*PSMI*) between summer troughs in Figure 17.

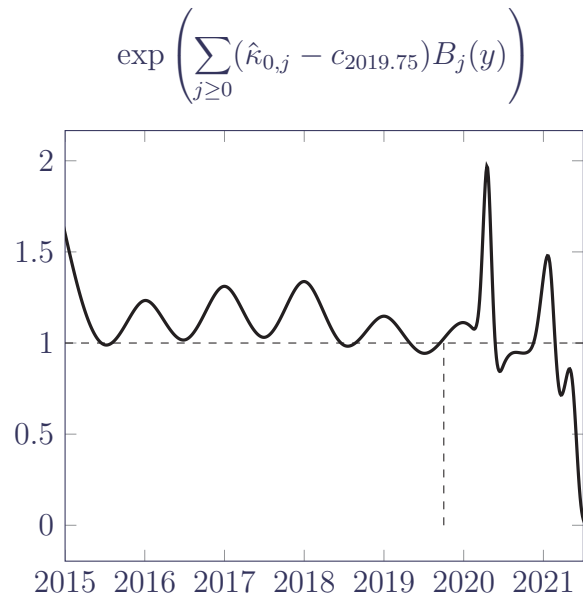
	2016.5	2017.5	2018.5	2019.5	2020.5
2015.5	-2.8%	-2.1%	0.1%	1.2%	2.4%
2016.5		-1.4%	1.5%	2.5%	3.6%
2017.5			4.3%	4.3%	5.2%
2018.5				4.4%	5.7%
2019.5					7.0%

10 Occurred-but-not-reported deaths (OBNR)

Although the UK3 portfolio data includes experience to end-June 2021, we have so far only used the experience to end-2020 to avoid any impact from delays in death reporting. Richards [2021, Figure 7] examined the reporting delays for the same portfolio and found minimal impact on mortality experience quarter of a year or more before the extract date, so discarding the most recent half-year of experience should be more than enough to eliminate OBNR effects. However, discarding experience data is undesirable if it can be avoided, and Richards [2021] proposed a parametric model for late-reported deaths as a means of using all available data without distorting the final results.

However, the reduction in reported mortality leading up to the extract date is just another pattern in time, which means that the Schoenberg spline function in equation (7) is also an alternative means of allowing for late-reported deaths. We can therefore extend the spline basis and include all available experience data, including the period most affected by late-reported deaths. The result is presented in Figure 18, which shows not only a multiplicative factor falling to zero at the extract date at end-June 2021, but also the second Covid-19 spike of January 2021. Figure 3 suggests that the second covid-19 spike should be around the same height as the first, and so it is possible that the lower spike in January 2021 in Figure 18 is due to some OBNR deaths. We could in theory use the forecasting method of Richards [2021, Section 7] to adjust for this, but a parametric OBNR model will not easily fit due to the high correlation with the last spline parameter.

Figure 18: Schoenberg time-spline function for y spanning 1st January 2015 to end-June 2021 using the spline basis depicted in Figure 16 and extended with more knots to cover Q1 2021. The figure shows the time-based multiplier for the mortality hazard. The identifiability constraint is that at 1st October 2019 (2019.75 decimalised) the multiplier is normalised at 1.



11 Schoenberg splines for mortality by age

The demonstrable effectiveness of Schoenberg splines when applied to time-varying mortality raises the question whether we should not also use them for mortality by age. This is not a new idea — McCutcheon [1979] used cubic splines by age when graduating mortality tables, for example. However, the approach of McCutcheon [1979] was designed for stratified grouped data, and not the inclusion of additional covariates per life.


For application to a survival model for individual lives we therefore use equation (11), where α_0 is the baseline, $B_k(x)$ is the k^{th} B -spline evaluated at exact age x and $\alpha_{0,k}$ is the corresponding spline parameter. As before, summation is from $k = 1$ because the spline beginning at $k = 0$ is absorbed into the baseline.

$$\log \mu_x = \alpha_0 + \sum_{k \geq 1} \alpha_{0,k} B_k(x) \quad (11)$$

Figure 19 shows the AIC and BIC for various knot spacings for $B_k(x)$. As with Table 4, the two

information criteria give conflicting guidance on which model to pick: the AIC indicates that the best knot spacing in age is four years, while the BIC suggests that the optimum knot spacing is 25 years or more.

Figure 19: AIC, BIC and number of spline parameters for UK3 mortality rates using model in equation (11) and various equidistant knot spacings. Ages 55–105, 1st January 2015 to end-2020.



One reason to reject the AIC-recommended four-year knot spacing is shown in Figure 20. The model is over-fitted at both the youngest and oldest ages: we have no reason to believe that mortality rates decline at the oldest ages, and the single death at age 55 is not grounds enough for the fitted shape. This is then an advantage of the Hermite-spline approach of equation (1), as it has fewer parameters and less flexibility, thus only permitting slow and stable changes with age. Alternatively, a basis of widely spaced Schoenberg splines will behave similarly. However, we note that the Hermite-spline approach of equation (6) is better suited to the modelling of post-retirement mortality differentials due to automatic convergence with increasing age [Richards, 2020b, Section 3].

In Figure 20 we see that excessive flexibility from closely spaced Schoenberg splines in age is undesirable — the fit might be superficially better quantitatively, but is poorer qualitatively when one takes into account data quality and reasonable prior expectations. However, there will be occasions when the greater flexibility of closely spaced Schoenberg splines is useful, such as where younger ages are included. Figure 21 shows an example of a non-monotonic pattern with age, for which Hermite splines will be unsuitable and for which greater flexibility is required.

Figure 20: Crude mortality hazards (\circ) and fitted rates ($-$) using equation (11) and a four-year knot spacing, showing over-fitting. Source: UK3 data.

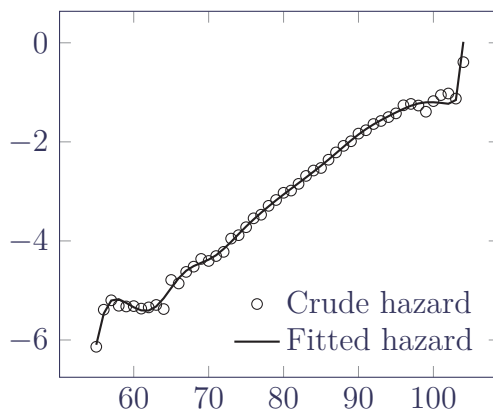



Figure 21: Graduated mortality rates for Australian males, showing “accident hump” around age 20. Source: Heligman and Pollard [1980, Figure 12].



12 Information criteria and actuarial judgement

In this paper we have seen two instances where the AIC led to a qualitatively poorer choice of knot spacing for Schoenberg splines for mortality: first for a time-varying function — Figure 15(a) — and again for an age-varying function (Figure 20). In both cases the AIC led to an excessive number of parameters and over-fitting. Using the small-sample correction for the AIC [Hurvich and Tsai, 1989] would not have led to a different outcome due to the large number of observations [Macdonald et al., 2018, Table 6.1].

The definitions of the AIC and BIC used in this paper are the usual ones given in the titles of Figures 19(a) and (b). However, it is worth noting that more rigorous definitions are given in equation (12), where ℓ is the log-likelihood from equation (3), n is the number of lives and df is the number of degrees of freedom used. In equation (12) it is common to use $df = p$, where p is the number of parameters; for a simple linear regression model the two are synonymous. However, Owen [1991, p103] noted that it is possible for a non-linear model to have $df > p$, while Ye [1998, p122] notes that even a linear smoother can have fewer degrees of freedom than there are parameters ($df < p$). In the field of mortality modelling, Macdonald et al. [2018, Section 11.6] give an example of a penalised GLM where the effective degrees of freedom are $df = 9.0$ with $p = 20$ regression coefficients. Assuming $df = p$ is a simplifying assumption only, and actuaries should not automatically pick a model solely because of its low information criterion.

This problem of the AIC leading to over-fitting may be restricted to relatively large numbers of parameters for a given continuous risk factor. However, relying only on the BIC as an alternative is not a complete solution either for actuarial work. For example, there is a distinction between risk factors that are statistically significant and those that are financially significant; Richards [2020b, Section 9] gives several contrasting examples. In Figure 6 the large negative residuals for the lowest and highest size-bands are of equal significance statistically, yet actuaries will tolerate a poorer fit for smaller amounts if it means better explaining the mortality of those with the largest annuity amounts. While the BIC might be the better information criterion, it cannot be the sole arbiter of model selection for actuarial purposes.

Leaving aside the complexities of model selection, there is also the practical question of how to turn a model into a mortality basis. In the bulk-annuity and longevity-swap markets, specialist valuation software is required to handle the complexities of UK pensioner benefits. However, such valuation systems are seldom capable of handling multi-factor models, and usually a table of mortality rates by age and gender is all that can be accommodated. A useful conversion tool is the equivalent-reserve method of Willets [1999], which expresses one mortality basis in terms of another via the medium of the liabilities in question. For the sake of example, assume that we want to express current mortality levels in terms of the equivalent percentages of some third-party table. We are interested in the mortality rates applying at a point in time, so we ignore future mortality improvements, as these are usually a separate basis item. The approach of Willets [1999] is to solve equation (13) for males and females separately:

$$\sum_i a_i \int_0^\infty v^t {}_t p_x^{\text{model}} dt = \sum_i a_i \int_0^\infty v^t {}_t p_x^{\text{table}} dt \quad (13)$$

where a_i is the annuity amount for life i , v^t is the net discount function to apply to a payment in t years, ${}_t p_x^{\text{model}}$ is the survival probability according to the model and ${}_t p_x^{\text{table}}$ is the survival probability according to the published table. We do not need to worry about the distorting effect of ignoring

mortality improvements, as long as they are ignored on both sides of equation (13). In general, ${}_t p_x = \exp\left(-\int_0^t \mu_{x+s} ds\right)$, and we use the definition of the mortality hazard in equation (14):

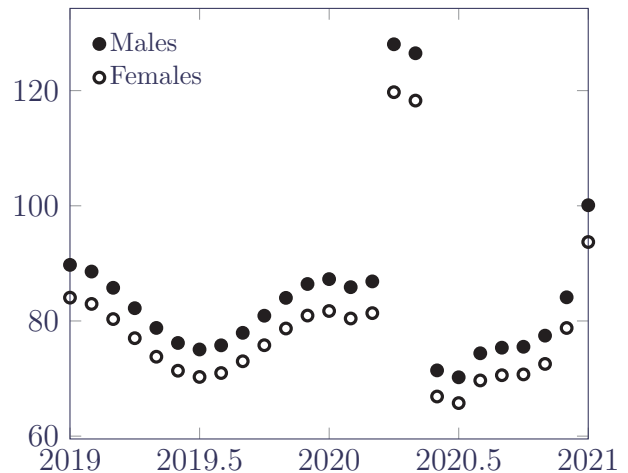
$$\mu_x^{\text{table}} = f \mu_x^{\text{S2PA}} \quad (14)$$

where μ_x^{S2PA} is the mortality hazard according to the table S2PA [CMI Ltd, 2014], and f is the percentage of that table to be solved for in equation (13). Figure 22 shows the results using the modelled mortality levels applying at various points in time over 2019–2020, assuming — somewhat unrealistically — that those mortality levels continue indefinitely, i.e. for the calculation using mortality levels at outset time y the integrated hazard is $H_x(t) = \int_0^\infty \mu_{x+s,y}$.

Figure 22 shows the extent of the challenge for actuaries in setting a best-estimate mortality basis for pricing bulk annuities or longevity swaps. Under normal circumstances one might pick the most recent mid-point between a summer trough and a winter peak; $75\frac{1}{2}\%$ for males and $70\frac{3}{4}\%$ for females in October 2020, say. As it happens, mortality at this point in time was largely unaffected by covid-19 (see Figure 3), but the equivalent percentages for April and May 2020 show the discontinuities possible due to pandemic mortality. The equivalent percentages for January 2020 show an emerging second discontinuity due to the second covid-19 shock shown in Figures 3 and 18.

Of course, no actuary would use equivalent percentages like those for April and May 2020 in Figure 22. However, the importance of allowing for the mortality shock can be seen from imagining what the results would be like without it: without the flexibility of the Schoenberg splines in time, the shock points would be moved down and all the other percentages in Figure 22 would be shifted up; any resulting basis would thus be imprudent for pricing a bulk annuity or longevity swap. The value of the methodology lies in accommodating the mortality spikes so that they cannot drive bias at other points in time.

Figure 22: Percentages of S2PA implied by mortality levels over 2019–2020. Liabilities for UK3 at 1st January 2021 equated as per equation (13) using a net discount rate of 0% p.a. ($v^t = 1, \forall t$).



13 Conclusions

For continuous variables where mortality varies either monotonically or with simple shape variation, a basis of Hermite splines usually provides all the flexibility that is needed. Examples include mortality by age and annuity amount. For actuarial purposes, quality of fit for the smallest annuity amounts is less important than the fit for the largest amounts.

In contrast, continuous variables where mortality fluctuates a lot, or where fluctuations are sharp and extreme, are better handled by a basis of local splines. Using cubic Schoenberg splines — a.k.a. *B*-splines — we can model mortality levels in time for annuity portfolios and pension funds. With two knot points per year we can identify seasonal variation in mortality, and we can estimate portfolio-specific mortality improvements from the change in mortality levels between summer troughs. Where the number of observations is large, we find that the Bayesian Information Criterion (BIC) is materially better for selecting the number of knot points than the AIC.

To handle mortality shocks like covid-19, we add additional knot points from our *a priori* knowledge of the timing of mortality spikes in the population. This mixed approach of regular half-yearly knots and hand-placed additional knots allows the salient mortality features to be identified, i.e. both seasonal variation and mortality shocks. A benchmark time-point can be selected to use in setting a basis, safe in the knowledge that the model’s other parameters are not unduly biased due to the presence of shocks because their effects are explicitly modelled.

The flexibility of the local Schoenberg splines further allows the modelling of the impact of late-reported deaths. This removes the need to discard the most recent experience data and thus permits the use of all available data for analysis and basis-setting.

Acknowledgements

The author thanks Anil Gandhi and Caroline Roberts for their support in providing data. The author also thanks Gavin Ritchie, Torsten Kleinow, Stefan Ramonat and Prof. Angus Macdonald for helpful comments on earlier drafts of this paper. Any errors or omissions remain the responsibility of the author. Calculations were performed using the Longevitas survival-modelling system. Graphs were produced in tikz, pgfplots and R Core Team [2017], while typesetting was done in L^AT_EX.

References

- H. Akaike. Factor analysis and AIC. *Psychometrika*, 52:317–333, 1987. ISSN 0033–3123. doi: <https://doi.org/10.1007/BF02294359>.
- D. R. Butenhof. *Programming with POSIX Threads*. Addison-Wesley, Boston, 1997. ISBN 978-0-201-63392-4.
- 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. Collett. *Modelling Survival Data in Medical Research*. Chapman and Hall, Boca Raton, second edition, 2003. ISBN 1-58488-325-1.
- J. C. Craufurd Dunlop and A. Watt. *Fifty-ninth annual report of the Registrar General for Scotland*, volume 59. H.M. Stationery Office, Glasgow, 1915.
- J. C. Craufurd Dunlop and A. Watt. *Sixtieth annual report of the Registrar General for Scotland*, volume 60. H.M. Stationery Office, Glasgow, 1916a.
- J. C. Craufurd Dunlop and A. Watt. *Sixty-first annual report of the Registrar General for Scotland*, volume 61. H.M. Stationery Office, Glasgow, 1916b.
- J. C. Craufurd Dunlop and A. Watt. *Sixty-second annual report of the Registrar General for Scotland*, volume 62. H.M. Stationery Office, Edinburgh, 1918.
- J. C. Craufurd Dunlop and A. Watt. *Sixty-third annual report of the Registrar General for Scotland*, volume 63. H.M. Stationery Office, Edinburgh, 1919.
- J. C. Craufurd Dunlop and A. Watt. *Sixty-fourth annual report of the Registrar General for Scotland*, volume 64. H.M. Stationery Office, Edinburgh, 1920a.
- J. C. Craufurd Dunlop and A. Watt. *Sixty-fifth annual report of the Registrar General for Scotland*, volume 65. H.M. Stationery Office, Edinburgh, 1920b.
- I. D. Currie, M. Durban, and P. H. C. Eilers. Generalized linear array models with applications to multidimensional smoothing. *Journal of the Royal Statistical Society, Series B*, 68:259–280, 2006.

- C. de Boor. *A practical guide to splines*, volume 27 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2001. ISBN 978-0-387-95366-3.
- P. H. C. Eilers and B. D. Marx. Flexible smoothing with B -splines and penalties. *Statistical Science*, 11(2):89–121, 1996. ISSN 08834237. URL <http://www.jstor.org/stable/2246049>.
- P. H. C. Eilers and B. D. Marx. *Practical Smoothing: the joys of P-splines*. Cambridge University Press, 2021. ISBN 978-1-108-48295-0.
- P. H. C. Eilers, I. D. Currie, and M. Durban. Fast and compact smoothing on large multidimensional grids. *Computational Statistics and Data Analysis*, 50:61–76, 2004.
- P. H. C. Eilers, J. Gampe, B. D. Marx, and R. Rau. Modulation models for seasonal time series and incidence tables. *Statistics in Medicine*, 27(17):3430–3441, 2008. doi: 10.1002/sim.3188.
- B. Gompertz. On the nature of the function expressive of the law of human mortality. *Philosophical Transactions of the Royal Society*, 115:513–585, 1825.
- M. J. D. Hancock. The Health Protection (Coronavirus, Restrictions) (England) Regulations 2020, 2020. URL <https://www.legislation.gov.uk/uksi/2020/350/contents/made>.
- L. Heligman and J. H. Pollard. The age pattern of mortality. *Journal of the Institute of Actuaries*, 107, Part I, No. 434:49–80, 1980.
- C. M. Hurvich and C. L. Tsai. Regression and time-series model selection in small samples. *Biometrika*, 76(2):297–307, 1989.
- William Kahan. Further remarks on reducing truncation errors. *Communications of the ACM*, 8(I):40, 1965.
- V. K. Kaishev, D. S. Dimitrova, S. Haberman, and R. J. Verrall. Geometrically designed, variable knot regression splines. *Computational Statistics*, 31(3):1079–1105, 2016. doi: 10.1007/s00180-015-0621-7.
- S. A. J. Kimber, A. Kreyssig, Y. Z. Zhang, H. O. Jeschke, R. Valenti, F. Yokaichiya, E. Colom-bier, J. Yan, T. C. Hansen, T. Chatterji, R. J. McQueeney, P. C. Canfield, A. I. Goldman, and D. N. Argyriou. Similarities between structural distortions under pressure and chemical doping in superconducting BaFe₂As₂. *Nature Materials*, 8:471–475, 2009.
- E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley and Sons, eighth edition, 1999. ISBN 0-471-33328-X.
- 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.
- B. D. Marx, P. H. C. Eilers, J. Gampe, and R. Rau. Bilinear modulation models for seasonal tables of counts. *Statistics and Computing*, 20(2):191–202, 2010. doi: 10.1007/s11222-009-9144-9.
- P. McCullagh and J. A. Nelder. *Generalized Linear Models*, volume 37 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, London, second edition, 1989. ISBN 0-412-31760-5.
- J. J. McCutcheon. Some remarks on splines. *Transactions of the Faculty of Actuaries*, 37:421–438, 1979. doi: 10.1017/S0071368600008727.
- ONS. Daily deaths with covid-19 on the death certificate by date of death. Technical report, 2021. URL <https://coronavirus.data.gov.uk/details/deaths>. Accessed on 2021-06-15.
- A. Owen. Discussion: Multivariate adaptive regression splines. *The Annals of Statistics*, 19(1):102–112, 1991. doi: 10.1214/aos/1176347969.

- W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C++: The Art of Scientific Computing*. Cambridge University Press, 2005. ISBN 0-521-75033-4.
- 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/>.
- R. Rau. *Seasonality in human mortality: a demographic approach*. Demographic Research Monographs 03 XVI. Springer Verlag, 2007. ISBN 3-540-44900-0.
- S. J. Richards. Modelling mortality by continuous benefit amount. *Longevity Ltd*, 2020a.
- S. J. Richards. A Hermite-spline model of post-retirement mortality. *Scandinavian Actuarial Journal*, 2020:2:110–127, 2020b. doi: 10.1080/03461238.2019.1642239.
- S. J. Richards. Mortality shocks and reporting delays in portfolio data. *Longevity Ltd*, 2021.
- S. J. Richards, S. J. Ramonat, G. Vesper, and T. Kleinow. Modelling seasonal mortality with individual data. *Scandinavian Actuarial Journal*, pages 1–15, 2020. doi: 10.1080/03461238.2020.1777194.
- I. J. Schoenberg. Spline functions and the problem of graduation. *Proceedings of the American Mathematical Society*, 52:947–950, 1964.
- G. E. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6 (2):461–464, 1978.
- P. Spreuwenberg, M. Kroneman, and J. Paget. Reassessing the Global Mortality Burden of the 1918 Influenza Pandemic. *American Journal of Epidemiology*, 187(12):2561–2567, 2018. ISSN 0002-9262. doi: 10.1093/aje/kwy191.
- The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team. The epidemiological characteristics of an outbreak of 2019 novel coronavirus diseases (COVID-19) — China, 2020. *China CDC Weekly*, 2:113, 2020. ISSN 2096-7071. URL <http://weekly.chinacdc.cn//article/id/e53946e2-c6c4-41e9-9a9b-fea8db1a8f51>.
- F. van Berkum, K. Antonio, and M. Vellekoop. Quantifying longevity gaps using micro-level lifetime data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 184(2):548–570, 2020. doi: 10.1111/rssa.12631.
- R. C. Willets. Mortality in the next Millennium. *Staple Inn Actuarial Society, London*, 1999.
- J. Ye. On measuring and correcting the effects of data mining and model selection. *Journal of the American Statistical Association*, 93(441):120–131, 1998. doi: 10.1080/01621459.1998.10474094.

Appendices

A Implementation

A general approach to fitting left-truncated survival models is presented in Richards [2020b, Appendix B]. The starting point is the log-likelihood, ℓ , defined in equation (8). The contribution of a single life i to the log-likelihood is given in equation (15):

$$\ell_i = - \left[\int_0^{t_i} \mu_{x_i+s, y_i+s} ds \right] + d_i \log \mu_{x_i+t_i, y_i+t_i} \quad (15)$$

For maximising the log-likelihood function we require the gradient with respect to each parameter, and for estimating standard errors we require the second pure and cross derivatives for all parameters. General formulae are given in Richards et al. [2020, Appendix A].

The nature of the mortality hazards in equations (1), (2), (4), (7) and (11) means that closed-form expressions for the integrated hazard in equation (15) do not exist. For the results in this paper we have used the Romberg method of numerical integration, but Clenshaw-Curtis integration is also an option. Richards et al. [2020, Appendix A] discuss some practical tests for ensuring that numerical integration has taken place sufficiently accurately. Due to the computationally intensive nature of the calculations, we process work in parallel across 63 threads [Butenhof, 1997]. For the summation of lots of potentially small contributions like equation (15) we use the floating-point error-correction algorithm of Kahan [1965].

Where a material OBNR effect exists, as in Figure 18, the value of $\hat{\kappa}_j$ can be large and negative for the rightmost spline; in the model behind Figure 18, for example, $\hat{\kappa}_j \in (-5, 4), \forall j \in \{1, \dots, 22\}$, but $\hat{\kappa}_{23} = -66.5$. This can lead to a large number of iterations where the algorithm for optimising the log-likelihood places a limit on the size of the change in parameter estimates, e.g. the `parscale` control parameter in R Core Team [2017, `optim` function] or the `STPMX` variable in Press et al. [2005, `dfpmin` function]. To reduce the number of iterations, the limit on the step size between successive parameters can be increased, or else an adaptive step size used when the number of iterations proves large. In our implementation we double the maximum step size after each 50 iterations — this has little impact on most models with parameters in the range $(-20, +20)$, but adapts to cope with models with extreme parameter values caused by OBNR.

B Model parameters

There are two groups of model parameters to be set for the mortality models in this paper: (i) configuration settings, whose values are decided in advance by the analyst, and (ii) parameters whose values are estimated from the data.

B.1 Model settings decided by the analyst

Table 6 lists the configuration settings that need to be decided in advance by the analyst, i.e. they are not estimated from the data. Where appropriate, the values used in the main body of the paper are given.

Table 6: Configuration parameters.

Parameter	Value	Description and role
x_0	50	Age below which $\log \mu_x$ is unmodelled. See equation (1).
x_1	110	Age above which $\log \mu_x$ is unmodelled. See equation (1).
$\tau(a, \lambda_0)$	$\frac{ae^{\lambda_0}}{1 + ae^{\lambda_0}}$	Transform function mapping annuity amount $a \geq 0$ onto $[0, 1]$. See equation (6) and Richards [2020a] for other transforms. The value of λ_0 is estimated from the data.
<i>bdeg</i>	3	Order of Schoenberg splines used in equation (7).

Other analyst decisions include the identifiability constraint on the $\{\kappa_{0,j}\}$ parameters (we use $\kappa_{0,0} = 0$) and the knot spacing.

B.2 Parameters estimated from the data

Table 7 sets out the parameters whose values are estimated from the data by maximising the log-likelihood in equation (3) or (8).

Table 7: Overview of parameters.

Parameter	Name	Description and role of parameter
α_0	Intercept	$\log \mu_{x_0}$.
α_{female}	Gender.F	Addition to α_0 for females.
m_0	AgeGradientYoungest	Initial rate of change of $\log(\text{mortality})$ to the right of x_0 .
m_1	AgeGradientOldest	Rate of change of $\log(\text{mortality})$ approaching age x_1 from below.
ω_0	Oldest	$\log \mu_{x_1}$, the logarithm of the limiting hazard above x_1 .
ω_{female}	Gender.F:Oldest	Addition to ω for females.
ω_{amount}	AmountUltimate	Maximum mortality reduction on a logarithmic scale from having a near-infinite annuity amount; see equation (6).
λ_0	AmountTransformParameter	Parameter used in mapping annuity amount from $[0, \infty)$ to $[0, 1]$; see equation (6).
$\kappa_{0,j}$	TimeSpline. <i>j</i>	Coefficient of time spline <i>j</i> . Splines are numbered from zero, but the zero'th spline is absorbed into the baseline hazard.

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 is a registered trademark for Longevitas Ltd in the UK (registration number 2434941), throughout the European Union (registration number 5854518), and the USA (Trade Mark Registration No. 3707314).

 **LONGEVITAS**TM