B.A.J. **15**, Supplement, 249-265 (2009)

LONGEVITY: A 'SIMPLE' STOCHASTIC MODELLING OF MORTALITY

By B. Browne, J. Duchassaing and F. Suter

ABSTRACT

All UK insurers exposed to longevity risk need to perform stress tests for their Individual Capital Assessment (ICA). Some have put in place deterministic models which are arguably too simple; others have developed stochastic models that can be demanding and complex.

This paper presents a simple model to turn any deterministic mortality scenario into a stochastic model. We propose a simple model of stochastic variation that is easy to explain and to implement, which could be an alternative to and/or complete some of the well known models. The model can be applied to any best estimates of future mortality rates, as it aims to describe how longevity behaves around the projected expected values.

The paper proposes a possible calibration on the England and Wales population mortality that produces a minimum indication of possible future variation and uses the results to validate the model's assumptions. Using sample portfolios and the stochastic model, we can simulate cash flows to determine the distribution of the net present values (NPV) of annuity outgo.

KEYWORDS

Longevity Risk; Mortality Projection; Annuities; Model Risk; Stochastic Mortality; Deterministic Scenario

CONTACT ADDRESS

Juliette Duchassaing, Partner Reinsurance Europe Limited, 153 rue de Courcelles, F-75817 Paris Cedex 17. Tel: +33 (0) 1 44 01 69 89; Fax: +33 (0)1 44 01 17 82;
E-mail: juliette.duchassaing@partnerre.com

1. Introduction

1.1 In this paper the authors present a simple model to turn any deterministic mortality scenario into a stochastic model. Note that simple does not means simplistic. However the paper aims to share a pragmatic solution that could for example be useful in the context of risk management or Solvency II.

1.2 The paper commences by explaining the context and the basic idea. We then describe the model in detail and discuss some calibration issues before validating the assumptions. Numerical examples are shown. Finally we present the benefits of the model when determining capital requirements and briefly some other possible applications before concluding.

1.3 *Why Do We Even Need Stochastic Mortality?*

1.3.1 In the longevity business, each of us has developed our own views on future mortality. We most likely base our views on the results of mathematical models, but we also take into account medical studies. In addition we may obtain expert opinions and attempt to incorporate their implications. We then apply judgement to determine our own view of future mortality for our given purpose, which is usually translated into a single deterministic mortality scenario. This scenario may be used as a basis for risk management, Solvency II or for pricing purposes.

1.3.2 For risk management and Solvency II needs, we must determine the capital to be retained to avoid the risk of ruin. When pricing, we load the premium in order to achieve a certain return on the capital needed to face risks, in particular the risk of deviation of future mortality from our expectations. Both cases quickly become complex and a deterministic methodology is no longer sufficient.

1.3.3 Ideally, defining the capital requires the expected distribution of the results, since it is usually expressed as a percentile of loss. A stochastic model is then necessary to determine the full range of possible results. But herein lies the issue: our best estimate future mortality consists so far of a deterministic scenario.

1.4 *What Are Some of the Options for Stochastic Mortality?*

1.4.1 One option is to use common stochastic models, such as Lee–Carter, or the Cairns–Blake–Dowd family (Cairns *et al*., 2009). Another is to model the future mortality using so-called 'forward mortality models', see, for example, Bauer *et al*. (2007) for further references.

1.4.2 However these methods impose a certain structure on expected mortality that may not be consistent with our own views of future mortality. In addition, their associated volatility corresponds often only to uncertainty in parameter estimation, not to overall model uncertainty and hence they do not capture the full potential deviation of mortality.

1.4.3 This paper proposes a pragmatic solution to make a deterministic scenario stochastic while avoiding these issues and yet being only as complex as necessary.

1.4.4 While this model works well for the authors' needs, some adjustments or refinements may be necessary to be applicable for other practitioners.

2. The Problem

2.1 Let us first identify the problem we wish to address more closely. At the same time we will explain the basic idea behind the model.

2.2 Only the future will tell what the mortality experience will really be.

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

The actual future mortality can differ from our own (deterministic) views through a number of possible sources of variation or risk, including:

— [Random] Fluctuation risk, i.e. we have the right guess on average but the source of observed variation is only the random fluctuation around the deterministic line. A hot summer or a mild winter are unpredictable.
— Pandemic risk: this can be interpreted as a temporary distortion or shock to mortality rates across ages and genders that disappears after a certain period of time.
— Trend risk, i.e. we systemically mis-estimated future mortality improvements, leading to losses or gains in our annuity book that can often be severe. This could arise following the introduction of a new and unexpected medical treatment for example.

2.3 For illustration purposes we assume that our best estimate assumption is the following deterministic scenario:

— PCMA00 as base level in 2000.
— Medium Cohort as mortality improvements from 2001 onwards as published in [CMI 2009a] and [CMI 2009b].

2.4 Figure 1 shows the probability of death by age until age 100 as well as the corresponding survival curve for an insured life aged 65 in 2009. Note that for each age there is a corresponding calendar year as shown in the horizontal scale. We have inserted the effects of examples of the random fluctuation, trend and pandemic risks, albeit heavily exaggerated.

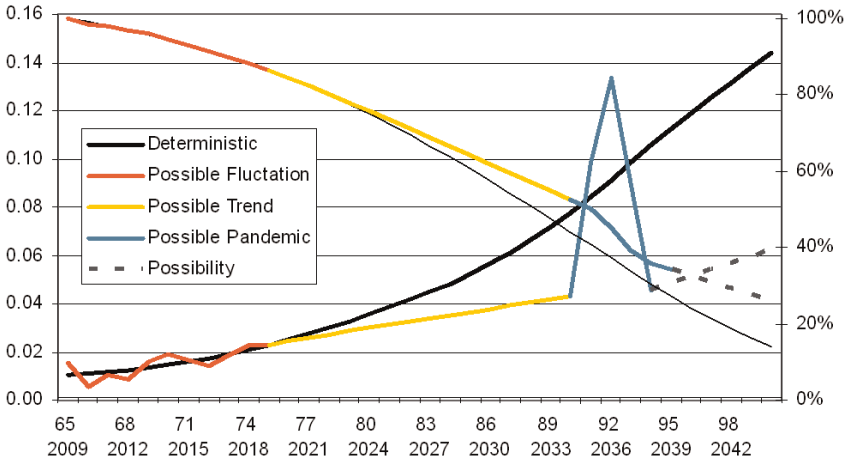

Figure 1. Probabilities of death and corresponding survival curves for a male aged 65 in 2009

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

2.5 Overall various sources of variation from expected mortality can be modelled. This paper presents a model oriented principally towards estimation of the risk of longevity, i.e. the risk that the mortality trend diverges from expectation due to the uncertainty around the central estimate.

2.6 The intention is to model the deviation over time of the observed mortality compared to that expected. Our model focuses on the trend risk only, but pandemic risk could easily be included and random fluctuation must be modelled using the number of lives in the actual portfolio modelled.

3. The Model and Some Features

3.1 We now turn to the model itself and the necessary underlying assumptions.

3.2 The aim of the model is to describe how mortality rates may behave around any chosen deterministic best estimate basis. Therefore the paper proposes a multiplicative model, where the potential future observed mortality is defined as the expected mortality times a stochastic process and which is depending on time only, i.e. independent of age. The model can be presented simply as:

$$\hat{q}_{x,t} = q_{x,t} \times C_t + \varepsilon_{x,t} \tag{1}$$

where for age x at time t:

$q_{x,t}$ is the expected mortality,
$\hat{q}_{x,t}$ is the observed mortality (restricted to [0,1]),
C_t is the stochastic process, and
$\varepsilon_{x,t}$ is the random noise.

3.3 The diversifiable risk, i.e. the fluctuation risk, is represented by the additive component ($\varepsilon_{x,t}$) that we have chosen to leave aside. Neglecting the noise component is appropriate in the context of the large portfolios usually under study by insurers and reinsurers.

3.4 Modelling the systemic risk equates to defining the stochastic process C_t. The model assumes:

$$\begin{cases} C_t = \exp(X_t) \times C_{t-1} \quad \forall t > 0 \\ C_0 = 1 \\ (X_i)_{i \in N} \sim \mathrm{N}(\mu, \sigma^2) \text{ iid} \\ \mathrm{E}[\exp(X_t)] = 1 \\ (\varepsilon_{x,i})_{i \in N} \text{ iid}, \quad \mathrm{E}[\varepsilon_{x,i}] = 0 \ \forall i \\ (C, \varepsilon) \text{ independent.} \end{cases}$$

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

3.5 *Implied Assumptions*

3.5.1 The assumption that the mean of the Lognormal Law should be equal to one implies that the mean of the C_t is also equal to one, which ensures that the mortality simulated with the C_t process will be centred on the chosen best estimate mortality.

$$\mathrm{E}[\exp(X_t)] = 1 \quad \Rightarrow \quad \mathrm{E}[C_t] = 1 \quad \Rightarrow \quad \mathrm{E}\left[\hat{q}_{x,t}\right] = q_{x,t} \quad \text{as } \mathrm{E}[\varepsilon_{x,i}] = 0.$$

3.5.2 The assumption that the mean of the exponential of X_t is equal to one is equivalent to the fact that the parameters μ and σ are linked. Indeed, using the calculation of the mean of a Lognormal Law, we obtain:

$$\mathrm{E}[\exp(X_t)] = \exp\left(\mu + \tfrac{1}{2}\sigma^2\right) = 1 \quad \Leftrightarrow \quad \mu + \tfrac{1}{2}\sigma^2 = 0.$$

Hence, the whole model could be calibrated by defining only one parameter.

3.6 As the very parameter we are interested in, i.e. the one we are trying to understand better, is the potential variation around the best estimate, we have chosen to calibrate the model using the parameter σ, i.e. the volatility.

3.7 The assumption of independence from age may appear to be quite strong, and so this, along with all the other assumptions associated with the model must be validated.

3.8 *An Observation*

3.8.1 Before moving to the next section however, we shall observe that the model described in subsections 3.2 and 3.4 (ignoring the random noise) can be restated as follows:

$$\ln(\hat{q}_{x,t}) - \ln(q_{x,t}) = \sum_{i=1}^{t} X_i$$

where

$$(X_i)_{i \in N} \sim \mathrm{N}\left(-\tfrac{1}{2}\sigma^2, \sigma^2\right) \text{ iid.}$$

3.8.2 As the expected value of the X_i distribution is not zero, the difference between the natural logarithms of the expected and observed mortality rates consequently does not vanish.

3.8.3 This may seem surprising or even unexpected, but it is related to the fact that the difference between the expected values of the natural logarithm of the expected and observed mortality rates is not equal to the natural logarithm of the difference between the expected values themselves, as shown below:

$$\mathrm{E}\left\lfloor \ln(\hat{q}_{x,t}) \right\rfloor - \mathrm{E}\left\lfloor \ln(q_{x,t}) \right\rfloor \neq \ln\left\lfloor \mathrm{E}(\hat{q}_{x,t}) - \mathrm{E}(q_{x,t}) \right\rfloor.$$

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

Section 6 presents a more general observation related to the characteristics of multiplicative models that occurs when projecting annuitant survival.

4. Model Calibration Approach

4.1 Prior to validating the model with historical data in Section 5, we discuss in this section the means by which we can determine the volatility parameter.

4.2 To calibrate the model, we have shown that we only need to determine the σ parameter, i.e. the standard deviation of the Normal Law followed by X_t. In fact the X_t can be calculated from the C_t's, which in turn are defined from the observed and expected mortality rates:

$$X_t = \ln\left(\frac{C_t}{C_{t-1}}\right) \quad \text{where } C_t = \frac{\hat{q}_{x,t}}{q_{x,t}}.$$

4.3 *What is the Expected Mortality?*

4.3.1 Usually the expected mortality is derived using past observed mortality.

4.3.2 From a methodological point of view, we should split the available historical period into two equal periods, use the first set to project the expected mortality for the second period and use the projected and observed mortalities over the second period to determine the volatility.

4.3.3 This means that we would need at least as many years of observation for each of the two periods as we expect to project the mortality in the future.

4.3.4 In practice, this is not always reasonable if we want to project for at least 50 years. Indeed, we would need more than 100 years of history and the data for such a long period does not always exist, or if available, may well suffer from quality problems or not be representative of the future mortality.

4.4 For the purposes of this paper we decided to use 'smoothing' methods on the whole available historic observation period, and to pay attention to the back-testing of those methods. There are various approaches to smoothing a historical mortality surface, but it should be noted that best estimates are often based on a composite of these methods and other information, and may not be a 'pure' application of any one method.

4.5 In this paper we have restricted our choices to some methods described in LifeMetrics, the publicly available JP Morgan toolkit for measuring and managing longevity and mortality risk. We will refer to them by the naming convention proposed in the technical document JPMorgan (2007):

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

— M1: the Lee–Carter method.
— M3: an extension of Lee–Carter with a simplified cohort effect.
— M4: the P-Splines method.
— M7: a model from Cairns, Blake and Dowd including a 'quadratic' cohort effect (Cairns *et al.*, 2009).

4.6 In the next section concerning the determination and the validation of the assumptions we will also briefly comment on these models.

5. Validation of the Model

5.1 An important step when proposing a model is to demonstrate its fit to the observations. In this section we determine the volatility parameter required to generate the X_t distribution and show that $(X_i)_{i\in N} \sim \mathrm{N}\left(-\frac{1}{2}\sigma^2, \sigma^2\right)$, are independent of age.

5.2 *Sample Case*

5.2.1 Here we present an example based on the population of England and Wales. The historical dataset is publicly available from the Human Mortality Database (HMD, 2008). We restricted ourselves to age range 60 to 89 years over the period 1955 to 2005. Hence we worked with two surfaces (one for males, one for females) of 30 ages each with 51 calendar year 'observations'. We also worked with two smaller surfaces (30 ages times 30 years), which produced very similar results.

5.3 *Choice of the 'Smoothing' Method*

5.3.1 The Lee–Carter model was discarded because it does not capture the cohort effect observed in the England & Wales data. We further considered that using P-Splines would produce a residual volatility which would depend too heavily on the chosen balance between smoothness and goodness-of-fit — producing a sort of self-fulfilling prophecy. For this reason we discarded this method also.

5.3.2 M3 and M7 produced robust results from back-testing on the England & Wales data as shown in Cairns *et al.* (2009). Therefore we worked with both of those in our example.

5.3.3 Using two different methods will demonstrate the independence of the results from the underlying model chosen to develop the best estimate 'expected' mortality surface.

5.3.4 On the other hand it is important to note that the structural form of these models have originally been developed and extensively tested on different past periods and here have been parameterised to provide a good fit to the historical data studied, hence the residual volatility that we identify can be considered as a minimum level when using the model to project future potential variation.

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

5.4 As outlined previously, the model implies some assumptions that we wish to validate, as follows:

5.4.1 *The parameters μ and σ are independent of age*

5.4.1.1 In Figures 2 and 3 we show the calculated parameters μ and σ by age for both males and females.

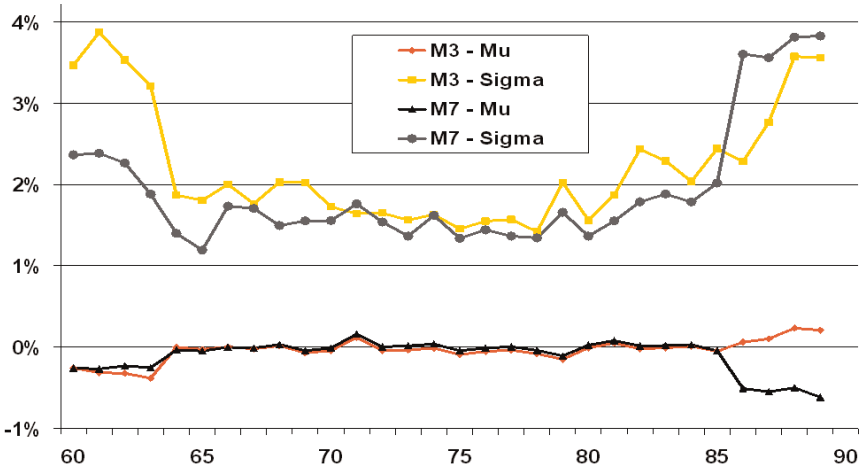

Figure 2. μ and σ by age for males

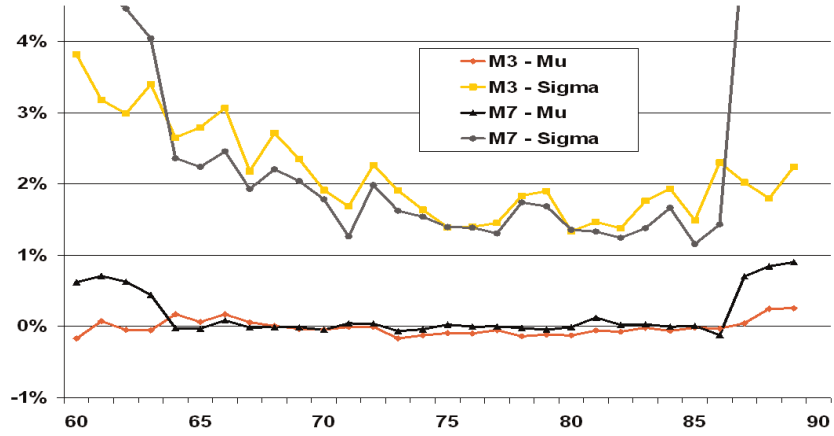

Figure 3. μ and σ by age for females

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

5.4.1.2 Except at the extremes, we observe that the stability of μ and σ, consistent with the assumption that they should be constant across ages.

5.4.1.3 Our studies have shown that in fact this assumption may not hold for younger ages. So it should be considered that this model may not be appropriate without adjustment across all ages from birth to old older ages.

5.4.1.4 Note that it would be relatively straightforward to introduce age-dependent C_t's and introduce an allowance for correlation between successive ages. This is a refinement that should be considered.

5.4.2 *The X_t's follow a Normal distribution*

5.4.2.1 Figure 4 shows the histogram of the 51 years of X_t's for a sample age of 70.

5.4.2.2 To ensure the validity of the Normality assumption we have used the Jarque–Bera statistical test described in Jarque & Bera (1981) for all ages and show the results in Figures 5 and 6.

5.4.2.3 The p-values of the Jarque–Bera test for each age are higher than 5% almost everywhere except at the extreme ages, so we can reasonably accept the Normality assumption for central ages.

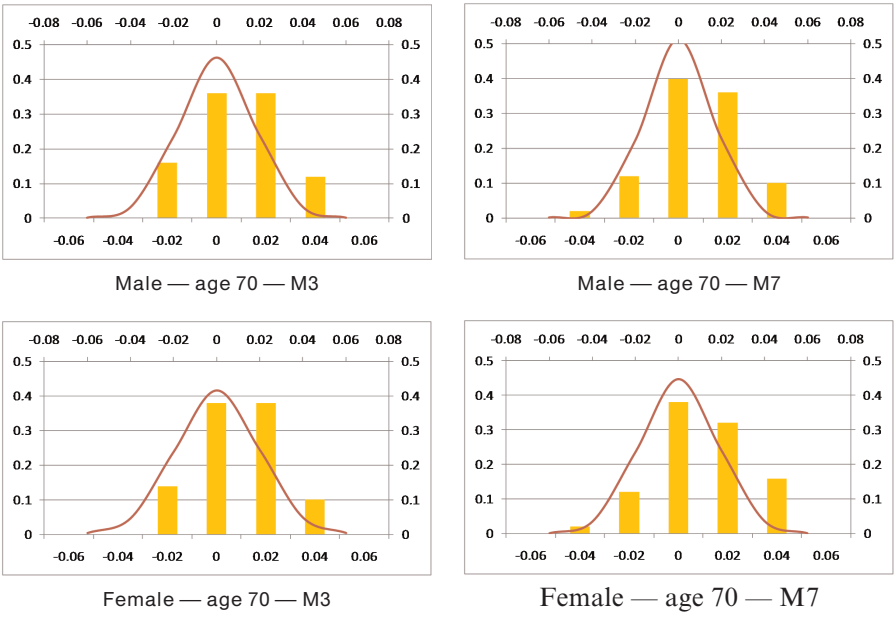

Figure 4. Sample histograms of X_t's

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

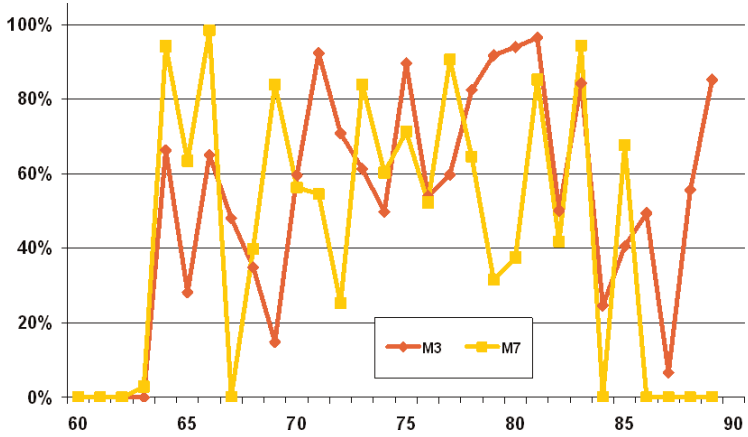

Figure 5. p-values of the Jarque–Bera test for males

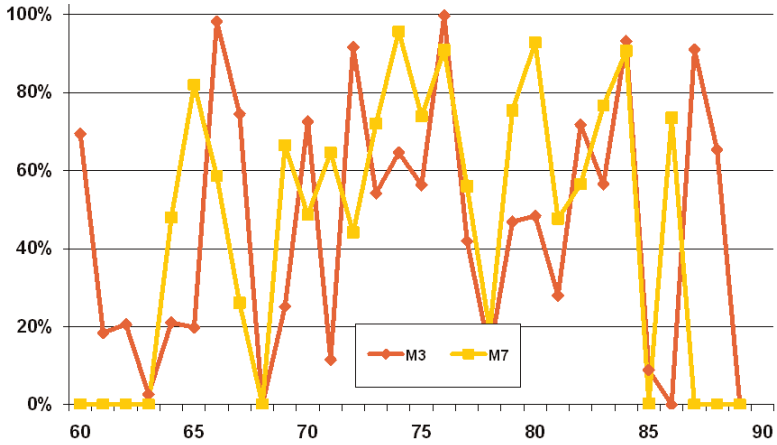

Figure 6. p-values of the Jarque–Bera test for females

5.4.3 'E[exp(X_t)] = 1', *i.e. the distribution is centred around the expected mortality*

5.4.3.1 In Figure 7, we can observe that the mean of e^{Xt} is very close to 1, except for extreme ages.

5.4.3.2 This is consistent with the preceding observation. Taking a pragmatic view we will restrict the observed range to the ages 64 to 84 years, where the assumptions of the model hold.

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

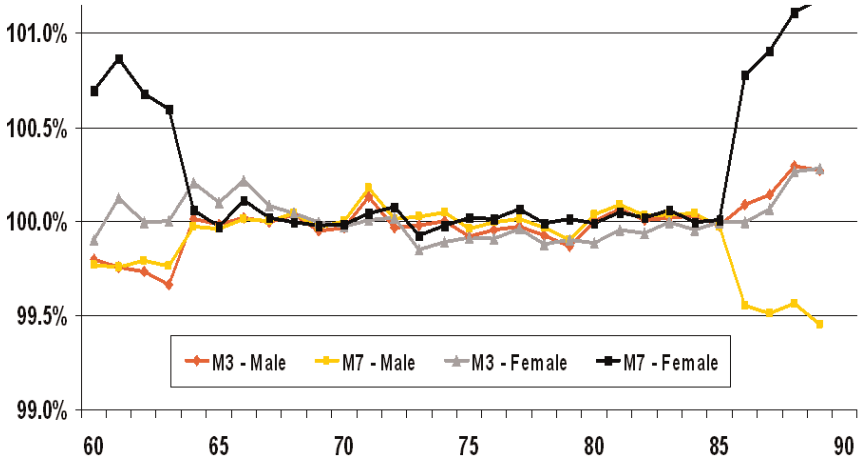

Figure 7. Mean of $\exp(X_t)$ — methods M3 and M7

5.4.3.3 We may examine the same assumption in another way by plotting $\mu + \frac{\sigma^2}{2}$ as shown in Figure 8. Again the assumption that the sum is 0 is validated within the considered age range as the values lie very close to 0 throughout the age range 64 to 84.

5.5 *Determining the Volatility*

5.5.1 Having validated the various assumptions of the model or if necessary defined a restricted area where they can be considered as valid, the

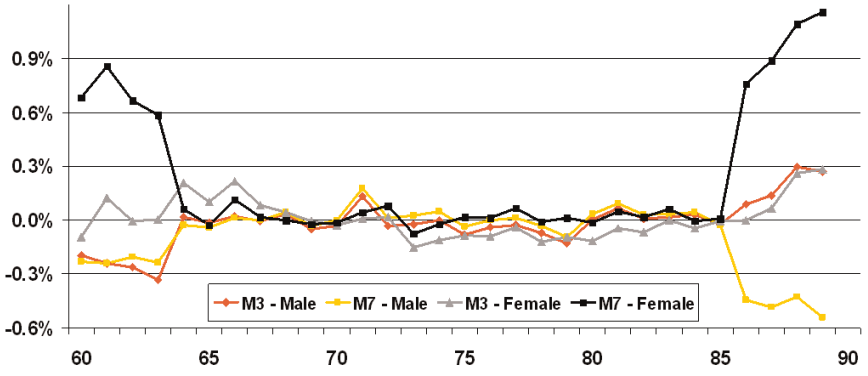

Figure 8. $\mu + \frac{\sigma^2}{2}$ — methods M3 and M7

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

value of the volatility σ can be determined based on the observation of the past. Figures 2 and 3 help to find a suitable value based on the past 51 years and the restricted age range 64 to 84 years.

5.5.2 Of course, as for nearly any model, this one has to be parameterised on past data, and clearly one must question whether the future is likely to only reflect the past and to what extent one would want to use the past volatility to set one's best estimate of the future volatility. Effectively, although a deterministic best estimate has been set, the fact that we are moving from observations on the past to projections about the future clearly introduces an additional level of uncertainty into the projections that can be partially allowed for by choosing a value for the volatility higher than that observed in the past. In any case the final choice for projection purposes remains for each user's judgment but we suggest the volatility should not be less than 2%.

6. Some Numerical Examples

6.1 As outlined above, defining the volatility is the single requirement to use this model and produce stochastic simulations around a deterministic scenario of the future mortality. In the following, for the purposes of this paper, we have assumed an arbitrary volatility of 4%.

6.2 *Sample Paths for C_t*

6.2.1 Figure 9 shows some sample paths over the next 50 years for the cumulative deviation C_t between the deterministic mortality scenario and the stochastic projection.

6.2.2 It can be observed that the deviations are roughly up to −60% and +160% away from the best estimate and positively skewed at the end of the period. This reflects the fact that the deviations follow a log-normal distribution by calendar year.

6.3 *The Corresponding Cash Flows*

6.3.1 We now show the impact of these same C_t factors on some annuity cash flows — a typical problem facing life assurance companies.

6.3.2 Figure 10 displays the corresponding sample cash flows for monthly annuities for a male aged 65 in 2009 based on each of the sample paths and on our best estimate mortality assumption described in section 2.3. We can observe that the cash flows significantly different from best estimate occur when deviations start earlier.

6.4 *'Standardised' Distributions of Present Values*

6.4.1 Simulating a sufficient number of C_t paths, we can calculate the distribution of the present values of cash flows.

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

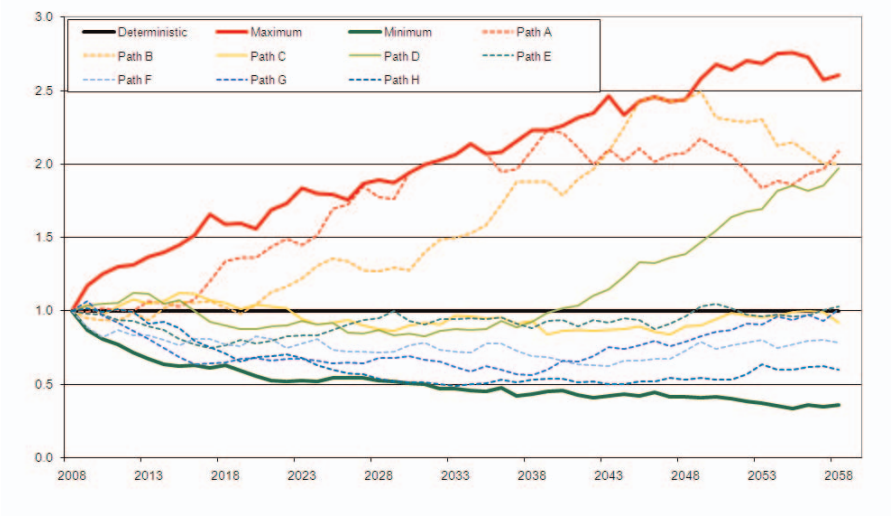

Figure 9. Sample paths for C_t

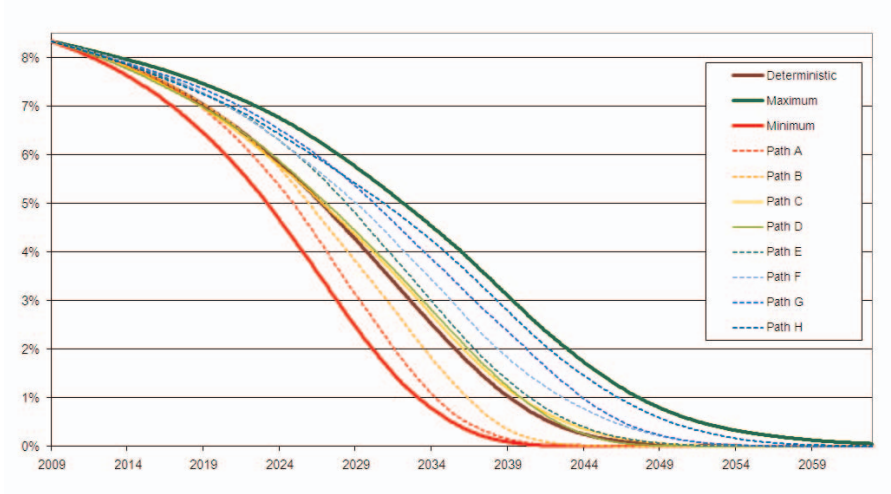

Figure 10 Sample cash flows for monthly annuities for a male aged 65 under scenarios set out in Figure 9

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

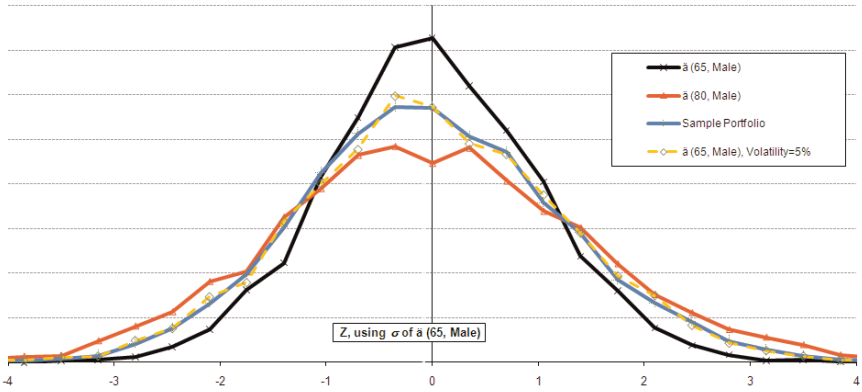

Figure 11 Distributions of the present values (discount rate of 5%) scaled to the standard deviation of ä (65, male)

6.4.2 Figure 11 shows some standardised distributions of present values as a function of the standard deviation of ä (65, male) which highlight the impact of the portfolio composition. The various shapes reflect the age distributions, but also the product features of the portfolio.

6.4.3 Note that for standardisation we keep the original mean of each distribution, but scale its standard deviation to that of the deviation of ä (65, male) using the ratio between the respective coefficients of variations in order to show the relative dispersion.

6.4.4 For clarity we note that the odd shape of ä(80) is due to on the one hand random features of the runs and on the other the creation process of the histogram.

6.4.5 As expected, increasing the assumed volatility from 4% to 5% increases the dispersion of the distribution of ä (65, male), showing its sensitivity to this parameter.

6.5 These distributions can be used to directly determine the percentiles of loss referred to earlier. However, before turning to this, we would like to highlight the following point.

6.6 In addition to the observation made in subsection 3.8, we note here that an issue that one must consider when accepting that mortality has stochastic characteristics is the fact that the average of the present values of the stochastic annuity cash flows is different from the present value of the deterministic scenario alone. The reason is the multiplicative nature of the annual survival mortality within the survival curve which leads to a skewed distribution.

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

7. Use for Capital Requirements

7.1 We mentioned in the introduction that risk management and Solvency II purposes require the determination of the capital needed to face the risk of ruin.

7.2 *UK Individual Capital Assessment*

7.2.1 The usual capital requirement prior to any diversification in the UK ICA regime (FRC, 2006) corresponds to the 99.5% percentile, which is also consistent with Solvency II.

7.2.2 For this model under the previous assumption of a 4% volatility when applied to an annuity for a male aged 65 in 2009, the annuity outgo deviation corresponds to 2.7 times the standard deviation, or an increase of 8.4% in the expected value as shown in Table 1. Using the Tail Value at Risk measure would require an additional 15% increase to 9.6%. The results when increasing the volatility to 5% can also be found in Table 1.

7.3 *Solvency II*

7.3.1 If we now turn to Solvency II's QIS4 Technical Specification (CEIOPS, 2009) for an example of a specific scenario, the 99.5th percentile 'equivalent' scenario is described as a permanent drop of 25% in the mortality.

7.3.2 In our model, using a volatility of 4% on this particular example, we would consider this 'equivalent' scenario to be only at the 98.9th percentile. A possible scenario which could be consistent with the 99.5th percentile could be a scenario based on a combination of a permanent drop in mortality and increased improvements as in the example displayed in Table 2. Should the volatility be increased to 5%, an even more extreme scenario should be chosen.

7.4 We have shown here that the presented model can help to quantify the probability of any deterministic scenario.

Table 1. Numerical results of examples for ä (65, male) under UK ICA

Risk measure	Probability	In standard deviations		Increase from mean	
Volatility		4%	5%	4%	5%
VaR	99.5%	2.70	3.34	+8.4%	+10.4%
	97.5%	1.99	2.49	+6.2%	+7.7%
TVaR	99.5%	3.10	2.42	+9.6%	+12.1%
	97.5%	2.42	3.03	+7.5%	+9.4%

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

Table 2. Numerical results of examples for ä (65, male) under Solvency II

Scenario	Probability		In standard deviations		Increase from mean	
Volatility	4%	5%	4%	5%	4%	5%
Permanent drop of 25% of the base mortality level	98.9%	97.2%	2.35	2.41	+7.3%	+7.5%
Combination of a −15% drop with additional improvements of '+1%'	99.5%	98.7%	2.68	2.84	+8.3%	+8.8%
Combination of a −20% drop with additional improvements of '+1%'	99.8%	99.6%	3.23	3.44	+10.0%	+10.6%

8. Other Potential Uses

8.1 In this penultimate section we would like to present a brief overview of two potential real world applications for the model, to stimulate further reflection around the uses to which such a model could be put.

8.2 We may wish to know the effect on annuity pricing of putting a corridor or any limit around the annuity outgo. While this feature may seem unusual to life insurance practitioners it will become increasingly frequent as banking institutions continue their interest in securitising and often tranching the risk. Pricing this requires the percentiles of the distribution of the cash flows so the impact on capital cost for any specific limit can be determined.

8.3 Secondly, the relatively simple structure of the presented model easily allows integration of a "shock"-based pandemic model into longevity scenarios.

8.4 Each user may find further applications.

9. Conclusion

9.1 In conclusion we have presented a model which:

— can be applied to any deterministic future mortality,
— independently of the underlying structure, with the advantage of being simple, as it depends only on the volatility,
— which was nonetheless validated by data,
— facilitates the assessment of the probability of extreme scenarios,
— can be further extended with other risks, e.g. pandemic.

9.2 In practice this model presents some draw-backs, mainly:

— the calibration issues, as well as
— the computing time,

https://doi.org/10.1017/S1357321700005596 Published online by Cambridge University Press

but these are both common in stochastic modelling.

9.3 We trust that this pragmatic model, which can be further refined as needed, may be considered an option when mortality needs to be stochastically simulated.

Acknowledgements

The authors acknowledge the ground work laid by Michael Koller with Elodie Fontana and Stanislas Boyer (Boyer & Fontana, 2009) and extended in association with Daniel Dubischar and Cédric Fétiveau allowing them to further validate and work with this model.

References

Bauer, D., Börger, M., Russ, J. & Zwiesler, H.-J. (2007). *The volatility of mortality*. Available at http://www.ifa-ulm.de/downloads/Vola_of_mortality.pdf

Boyer, S. & Fontana, E. (2009). RAC (Risk adjusted Capital) calculation for the non diversifiable part of the longevity and mortality risk. Thesis, Private communication.

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A. & Balevich, I. (2009). A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal*, **13**(1), 1-35.
http://www.jpmorgan.com/directdoc/LM_NAAJ2009 Quant Comparison.pdf

CEIOPS (2009). Draft CEIOPS' *Advice for Level 2 Implementing Measures on Solvency II: Standard formula SCR — Article 109 c. Life underwriting risk*. Consultation Paper No. 49, CEIOPS-CP-49/09.

CMI (2009a). *Continuous Mortality Investigation Reports: No. 23*. The Institute of Actuaries and the Faculty of Actuaries.
http://www.actuaries.org.uk/knowledge/cmi/cmi_reports/cmir23

CMI (2009b). *Continuous Mortality Investigation Working Paper: Number 37. Version 1.1 of the CMI Library of Mortality Projections.* The Institute of Actuaries and the Faculty of Actuaries.
http://www.actuaries.org.uk/_data/assets/pdf_file/0017/148211/WP37_v3.pdf

FRC (2006). *GN46 Individual Capital Assessment V2.0 (BAS Amendment 1)*. Financial Reporting Council/Board of Actuarial Standard. http://www.frc.org.uk/documents/pagemanager/frc/GN46 Individual Capital Assessment V2.0 (BAS Amendment 1).pdf

HMD (2008). The Human Mortality Database, as accessed in 2008. http://www.mortality.org/

Jarque, C.M. & Bera, A.K. (1981). Efficient tests for normality, homoscedasticity and serial independence of regression residuals: Monte Carlo evidence. *Economics Letters*, **7**(4), 313-318.

JPMorgan (2007). LifeMetrics: A toolkit for measuring and managing longevity and mortality risks. http://www.jpmorgan.com/directdoc/lifemetrics_technical.pdf