

Robust Hedging of Longevity Risk

Andrew J.G. Cairns

Maxwell Institute for Mathematical Sciences, Edinburgh, and
Department of Actuarial Mathematics and Statistics
Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.

WWW: <http://www.ma.hw.ac.uk/~andrewc>

E-mail: A.J.G.Cairns@hw.ac.uk

First version: August 2011

This version: January 30, 2013

Abstract

We consider situations where a pension plan has opted to hedge its longevity risk using an index-based longevity hedging instrument such as a q -forward or deferred longevity swap. The use of index-based hedges gives rise to basis risk, but benefits, potentially, from lower costs to the hedger and greater liquidity. We focus on quantification of optimal hedge ratios and hedge effectiveness and investigate how robust these quantities are relative to inclusion of recalibration risk, parameter uncertainty and Poisson risk. We find that strategies are robust relative to the inclusion of parameter uncertainty and Poisson risk. In contrast, single-instrument hedging strategies are found to lack robustness relative to the inclusion of recalibration risk at the future valuation date, although we also demonstrate that some hedging instruments are more robust than others. To address this problem, we develop multi-instrument hedging strategies that are robust relative to recalibration risk.

Keywords: Robust hedging, recalibration risk, hedge ratios, hedge effectiveness, Delta hedging, Nuga hedging.

1 Introduction

Pension plans and annuity providers have recently begun to focus greater attention on the measurement and management of longevity risk. This reflects a variety of factors including low risk-free interest rates, changes in accounting regulations, changes in the regulatory environment (e.g. Solvency II; see, for example, Olivieri and Pitacco, 2009), and the rise of enterprise risk management in general terms.

In this paper, we will consider some of the options available to a pension plan or annuity provider (hereafter referred to as the ‘hedger’) that has decided to reduce its exposure to longevity risk.

If the hedger has liabilities exceeding around £100 million then it has the possibility to arrange a customised longevity swap that indemnifies them against unexpected changes in their own mortality experience. Index based longevity hedging instruments such as q -forwards (see www.llma.com) provide an alternative to a customised transaction. They provide one of the more limited range of options for smaller hedgers, but they will also be of interest to larger hedgers who are prepared to accept some degree of basis risk in return for a standardised contract that potentially offers better value for money.

This paper builds on the recent work of Cairns et al. (2013) and will address various issues relating to robustness of various quantities connected to a given financial transaction that are derived from a stochastic or other model for mortality and longevity risk. For a customised longevity swap, the only quantity that might involve use of a model is the price of the transaction (that is, the specification of the fixed leg of the swap). In contrast, with an index based hedge, the measurement of hedge effectiveness, the determination of the appropriate quantities of hedging instruments, and the prices for these instruments are all quantities that are dependent on the underlying stochastic mortality model. We seek to investigate how robust these quantities are relative to various uncertainties in the underlying stochastic model and parameterisation. In other words: how confident are we in a stated hedge effectiveness, hedge ratio or price?

1.1 Structure of the paper

In Section 2 we introduce the preliminaries, the model, the data, and the potential hedging objectives, and summarise recent work on this topic. Section 3 outlines the questions about robustness that underpin the present paper. Section 4 provides some answers to the robustness questions through the use of empirical results and theoretical reasoning. Section 5 shows how strategies can be made more robust through the use of the new concept of Nuga hedging. Section 6 adds additional remarks about lesser risk factors affecting robustness. Section 7 concludes.

2 Hedging

2.1 Data and notation

- There are two, potentially overlapping, populations: $k = 1$ representing England and Wales (EW) males mortality; $k = 2$ representing the UK Continuous Mortality Investigation (CMI) assured lives males mortality. The CMI population is used as a proxy for a typical pension plan.
- Time $t = 0$ represents the end of 2005, the last year for which we have data for both the EW and CMI populations. We use data for 1961 to 2005 inclusive over the ages 50 to 89.
- $m_k(t, x)$ = death rate in year t for males aged x last birthday at the date of death and belonging to population k .
- $q_k(t, x)$ = corresponding population k mortality rate (probability of death) for individuals aged x exactly at the start of year t .
- We assume that $q_k(t, x) = 1 - \exp(-m_k(t, x))$.
- $a_k(T, x)$ = value at time T of a level annuity of 1 per annum payable annually in arrears to an individual aged x at time T in population k .

2.2 The static value hedging problem

We follow Cairns et al. (2013) in considering a simple setting in which the hedger (a) puts in place a static hedge and (b) wishes to assess the financial outcome of the hedge at a specific future valuation date, T (a so-called value hedge).

Other approaches to hedging that complement that in the present paper include Cairns et al. (2008) (static cashflow hedging), Dahl et al. (2008) (who employ value hedging in a dynamic, parameters-certain context to develop a dynamic Delta-hedging strategy using a continuous-time two-population mortality model), Cairns (2011) (discrete-time dynamic Delta-hedging), Li and Luo (2012) (who develop the concept of q -duration – Coughlan et al., 2007 – for estimating hedge quantities), and Dowd et al. (2011c) (static value hedging under a different two-population model).

Keeping things simple, the hedger has a liability at time $T = 10$ years (i.e. the end of 2015) that is equal to $a_2(T, x_0)$ where $x_0 = 65$. The value of the liability is dependent, therefore, on the cohort life table that is in use for population 2 (CMI) at the end of 2015 that will be used for valuing the hedger's liability.

We consider two potential hedging instruments, both index hedges with EW males as the reference population. The first is a *deferred longevity swap* (as employed by

Cairns et al., 2013) in which the floating leg of the swap is linked to the survivorship of EW males aged 65 at the end of 2015, and the fixed leg is set at the end of 2005 and might be an estimate at time 0 of the future cashflows. The second is a simplified version of a q -forward that swaps a fixed estimate of a mortality rate, set at time 0, for a floating mortality rate linked to year T mortality rates and a specific age x . As a variant on this we will also consider use of a q -forward that has a longer-dated maturity date of $T + U$ for some $U > 0$. To distinguish between these we will refer to the maturity T q -forward simply as a q -forward, and to the latter as a *long maturity* q -forward.

The values at T of these three contracts are:

$$a_1(T, x) - \hat{a}_1^F(0, T, x) \quad (1)$$

$$q_1(T, x) - q_1^F(0, T, x) \quad (2)$$

$$e^{-rU} (E_Q[q_1(T + U, x)|\mathcal{M}_T] - q_1^F(0, T + U, x)) \quad (3)$$

where $\hat{a}_1^F(0, T, x)$ is the value at T of the fixed leg of the deferred longevity swap, $q_1^F(0, T, x)$ is the fixed leg of the q -forward, and E_Q refers to expectations taken under a risk-neutral pricing measure.

In equation (3), we have assumed, for simplicity, that interest rates are constant with r being the continuously compounding risk free rate to allow us to focus on hedging longevity risk. This can easily be extended to allow for stochastic interest rates and time-dependent yield curves. Also in equation (3) (and implicit in the annuity value, $a_1(T, x)$, equation (1)) is the use of \mathcal{M}_T which represents the information provided by the mortality data up to time T .

The valuation model underpinning equations (1) and (3) is outlined in Section 2.4.2.

2.3 Hedging objective

Let $L = a_2(T, x)$ represent the value of the liability at T and H represent the value at T of the index-linked derivative (assuming, initially, a single hedging instrument is used). Again, keeping things simple, the only other asset class available to the hedger between 0 and T is a cash account that pays r per annum, continuously compounding.

As our hedging objective, we will assume that the hedger wishes to minimise the variance of $L - hH$, where h is the hedge ratio, or number of units held of the hedging instrument.¹ Solving this is straightforward (see, for example, Coughlan et

¹The use of variance as our objective function is consistent with recent transactions in the UK and, perhaps, elsewhere, where defined benefit pension plans are attempting to derisk: indicating that they have no appetite for risk if it can be avoided. Alternatives to the minimisation of variance, of course, might be used by the hedger such as the minimisation of the standard deviation, value-at-risk or expected shortfall (see, for example, Dowd et al., 2006), or expected utility (see Appendix A).

al., 2004, or Cairns et al., 2013) and we find that the optimal hedge ratio is

$$h = \frac{\text{Cov}(L, H)}{\text{Var}(H)} = \rho \frac{\text{S.D.}(L)}{\text{S.D.}(H)},$$

where ρ is the correlation between L and H , and the optimal hedge effectiveness is

$$1 - \frac{\text{Var}(L - hH)}{\text{Var}(L)} = \rho^2.$$

Hedgers might have alternative, and less extreme attitudes to risk and seek to optimise alternative objectives. By way of example, therefore, we present in Appendix A a stylised account of how the annuity hedging problem might be tackled if the hedger has exponential utility. We see that the minimum variance portfolio plays a key role even in the situation where the hedger has some appetite for risk. Utility-based optimisation is developed further in Zhou et al. (2011) in a longevity context and Cairns (2001) in a more general insurance context. We leave the issue of robustness in a utility-based setting for further work.

In some settings we might seek to construct a static hedge using $n > 1$ hedging instruments with values H_1, \dots, H_n at time T . The minimum variance problem then becomes minimise $\text{Var}(L - \sum_{i=1}^n h_i H_i)$ over h_1, \dots, h_n .

2.4 The stochastic mortality model

A key element of the work of Coughlan et al. (2011) and Cairns et al. (2013) is the distinction between the model that is used to simulate mortality experience between times 0 and T and the model that is used to value the assets and liabilities at time T given the observable mortality experience both before time 0 and between times 0 and T .

Specifically, Coughlan et al. (2011) and Cairns et al. (2013) recognised that valuers only have access to the observable data and not the underlying simulation model (which, in any event, is only an approximation to a more complex reality). As a consequence, they argue that the simulation and valuation models might differ. And even if the true model were known, the way that the valuation model is calibrated at T might be different from the way that the simulation model is calibrated at time 0. Finally, the use of the same model and calibration for both simulation and valuation might lead to an over-optimistic assessment of the benefits of hedging.² We do not claim that the approach taken here is perfect. However, it is a step in the direction

²For example, in equity derivative pricing and hedging, the use of the Black-Scholes model for both pricing equity derivatives and the assessment of the effectiveness of the Delta-hedging strategy suggests that we can eliminate risk entirely. However, we know that the model and its assumptions are, at best, only approximately correct and that complete elimination of risk through dynamic hedging with a single traded underlying is not possible.

of acknowledging that we do not have perfect information about the model or its parameters, and by doing so we reduce the bias (in an optimistic direction) that results from using identical simulation and valuation models.

We will illustrate the methodology for assessment of robustness using the two-population mortality model proposed by Cairns et al. (2011a), a two-population variant of the model “M3” investigated by Cairns et al. (2009). The underlying death rates in populations $k = 1, 2$ are given by

$$m_k(t, x) = \exp [\beta^{(k)}(x) + n_a^{-1} \kappa^{(k)}(t) + n_a^{-1} \gamma^{(k)}(t - x)]$$

where:

- $\beta^{(1)}(x)$ and $\beta^{(2)}(x)$ are the population 1 and 2 age effects;
- $\kappa^{(1)}(t)$ and $\kappa^{(2)}(t)$ are period effects;
- $\gamma^{(1)}(c)$ and $\gamma^{(2)}(c)$ are cohort effects; and
- n_a is the number of ages included in the historical dataset.

Other two or multi-population models have been proposed by Li and Lee (2005), Li and Hardy (2011), Plat (2009), Jarner and Kryger (2011), Dowd et al. (2011a), and Börger et al. (2012) amongst others.

2.4.1 The simulation model

In the simulation model (see Cairns et al., 2013, for more detail):

- $\kappa^{(1)}(t)$ is a random walk with drift ν_κ ;
- the spread between the period effects, $S_\kappa(t) = \kappa^{(1)}(t) - \kappa^{(2)}(t)$, is modelled as an AR(1) time series model with innovations that are correlated with the $\kappa^{(1)}(t)$ innovations;
- $\gamma^{(1)}(c)$ is modelled as an AR(2) process that is reverting to a linear trend;³
- the spread between the cohort effects, $S_\gamma(c) = \gamma^{(1)}(c) - \gamma^{(2)}(c)$ is modelled as an AR(2) time series model with innovations that are correlated with the $\gamma^{(1)}(c)$ innovations.
- No assumptions are made about relationship between the age effects, $\beta^{(1)}(x)$ and $\beta^{(2)}(x)$, at different ages or between populations.

³Numerical examples later in this paper focus on existing cohorts at time 0, suggesting that simulation of $\gamma^{(1)}(c)$ and $\gamma^{(2)}(c)$ is not necessary. However, recalibration of the valuation model at time T requires information about younger cohorts, so simulated cohort effects will have some influence (possibly small) on the cohorts that feature prominently in our study.

2.4.2 The time- T valuation model

Recall, now, the approximate method of Cairns et al. (2013) for calculating at time T simulated annuity prices and long-maturity q -forward prices.⁴ The underlying principle of the valuation model is that stochastic simulation beyond time T is replaced by a single deterministic projection that potentially depends on the sample path up to T . The most important element of this approach is that $\hat{\kappa}^{(1)}(T+u) = \kappa^{(1)}(T) + \nu_{\kappa}u$ and $\hat{\kappa}^{(2)}(T+u) = \kappa^{(2)}(T) + \nu_{\kappa}u$, with the risk-neutral drift, ν_{κ} , being specified at time T . For annuities, we proceed as follows. Let x_0 be the age at time T of a cohort for whom we wish to price an annuity. We next take $\beta^{(k)}(x)$ as given, along with $\kappa^{(k)}(T)$ and $\gamma^{(k)}(T-x_0+1)$. Our deterministic approach for calculating annuity prices then proceeds as follows:

- The cohort will be aged x_0 in year $T+1$, x_0+1 in year $T+2$, ...
- Median estimates for the cohort's death rates in years $T+1, T+2, \dots$ are

$$\log \hat{m}^{(k)}(T+t, x_0+t-1) = \beta^{(k)}(x_0+t-1) + n_a^{-1} (\kappa^{(k)}(T) + \nu_{\kappa}t) + n_a^{-1} \gamma^{(k)}(T-x_0+1).$$

- Define $\hat{p}^{(k)}(T, T+u, x_0) = \exp \left[- \sum_{t=1}^u \hat{m}^{(k)}(T+t, x_0+t-1) \right]$, representing the risk neutral probability that an individual in population k aged x_0 at time T survives to time $T+u$.
- The value at T of a temporary annuity payable annually in arrears for a maximum of U years is

$$a_k(T, U, x_0) = \sum_{u=1}^U e^{-ru} \hat{p}^{(k)}(T, T+u, x_0), \quad (4)$$

where r is the valuation interest rate.

The equivalent valuation at T of $q_k(T+U, x)$ is

$$\mathbb{E}_Q[q_k(T+U, x) | \mathcal{M}_T] = 1 - \exp(-\hat{m}_k(T+U, x))$$

which is required for a long-maturity q -forward.

As we will discuss later, the random-walk drift, ν_{κ} , plays a central role in the valuation at T of assets and liabilities with cashflows that fall after time T .

⁴Alternative approaches for approximating mortality contingent functions have been proposed by Denuit et al. (2010), Dowd et al. (2010a, 2011b), Plat (2010) and Cairns (2011).

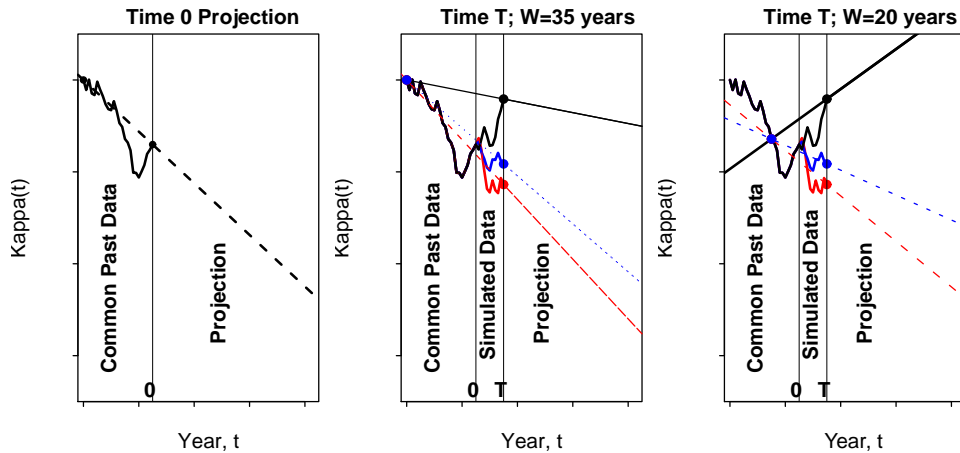


Figure 1: Recalibration risk for a random-walk sample path. Left: historical data up to time 0 and median forecast (solid and dashed lines respectively). Middle: three simulated sample paths up to time 10 (solid lines) plus revised median forecasts for each (dashed lines) assuming a 35-year lookback window. Right: as the middle plot but using a 20-year lookback window.

2.5 Recalibration risk

A key element of the difference between the simulation and valuation models concerns recalibration at T of the model parameters. Consider, for example, the random-walk model for $\kappa^{(1)}(t)$. Initially, we have data up to time 0 (25 observations, Figure 1, left). The median forecast for a random walk is a linear extrapolation of the sample path that connects the start and the end points of the observed path up to time 0. In Figure 1, middle and right hand plots, we have added in three simulated paths that take the random walk from time 0 to time T . At time T we wish to make fresh forecasts of the random walk. The middle plot shows what the median forecasts would be for these three simulations if we used all data up to time $T = 10$ (35 observations in total) to estimate the model parameters, and we can see that the gradient of the forecasts is very much dependent on the simulated path taken between times 0 and T . Additionally, each forecast would be quite different if we left the drift parameter unchanged from its time 0 estimate. These observations echo those of Li and Chan (2005), who pointed out that the principal trend in mortality forecasts is critically dependent on the first and last observations of $\kappa^{(1)}(t)$. Additionally, they point out that forecasts are exposed to potential errors if either the start or end point is considered to be an outlier for some reason.

In Figure 1, right, we show what the median forecasts look like if we use only the last 20 observations rather than all 35 to calibrate the parameters of the random-walk model. Comparing the middle and the right-hand plots, we can see that the choice

of 20 or 35 years for the calibration window can make a big difference in the forecast beyond time T , adding to the concerns expressed by Li and Chan (2005).

The visualisations portrayed in Figure 1 are also reminiscent of Figure 9 in Zhou and Li (2010). However, there, the different forecasts are related to model risk and the treatment of potential structural changes, for example, in the drift of the random walk.

But why would you use only 20 years rather than the full set of observations? In a mortality setting, there is no clear initial observation date: for some countries we have data that extends as far back as the 18th century, but we would not consider using such a long run to calibrate a mortality model. Instead, a more subjective judgment needs to be made to decide what length of recalibration window is appropriate in order to capture most closely what we think is the current trend and/or volatility. An alternative, and more objective, approach is proposed by Booth et al. (2002) by seeking to choose the window which has the best linear fit.

This line of thinking led Cairns et al. (2013) to argue that the parameters of the valuation model be included as sources of risk in a risk assessment. However, we can add that this might be regarded as a source of Knightian uncertainty (that is, a risk that cannot easily be quantified statistically) because the length of the recalibration window is not known, nor is it something that we can estimate from historical data. The approach of Booth et al. (2002) to choosing the length of the lookback window might be seen to reduce the amount of Knightian uncertainty, but their method is not without criticism and other approaches can be justified. Additionally, not all people involved in valuation act in such an objective way, and so we need to keep in mind the possibility that different people will use different recalibration windows.

2.6 Initial conclusions of Cairns et al. (2013)

A principal finding of Cairns et al. (2013) was that inclusion of recalibration risk was essential in the assessment of a longevity hedging strategy. They found that, although total risk is much larger with the inclusion of recalibration risk, the correlation between the liability and an index-linked hedging instrument also rose significantly, and, hence, the hedge effectiveness. This was explained by the common dependence of forecast values of $\kappa^{(1)}(T + s)$ and $\kappa^{(2)}(T + s)$ on the random-walk drift, ν_κ .

Applying this to the present study of robustness, we can conclude that correlation and hedge effectiveness are not robust relative to the inclusion or not of recalibration risk. However, Cairns et al. (2013) made no comment on the robustness of the corresponding hedge ratios or prices.

3 Robustness

We now generalise our discussion of robustness to pose a range of questions.

How robust are estimates of

Q1S optimal static hedge ratios h_1, \dots, h_n (assuming n hedging instruments), or

Q1D optimal dynamic hedge ratios $H_1(t), \dots, H_n(t)$,

Q2 hedge effectiveness, and

Q3 initial (i.e. time 0) prices of hedging instruments $\pi(H_1), \dots, \pi(H_n)$

relative to

R1 the treatment of parameter risk,

R2 the treatment of population basis risk,

R3 the treatment of recalibration risk in the valuation model (or, more generally, the choice of valuation model and its calibration),

R4 the treatment of Poisson risk (sampling variation) in the simulated numbers of deaths between times 0 and T ,

R5 the use of the most recent EW data (2006-2008), where corresponding CMI data is not yet available, and

R6 the choice of simulation model and its calibration?

This paper will focus on quantities Q1S and Q2 and risks R1, R3, R4 and R5.

In addition to looking at a given quantity and its robustness, we can ask what are the consequences if the quantity is not robust? For example, if our estimate of the hedge ratio is not robust then there is a risk that we might buy too many or too few units of the hedging instrument to achieve the targeted level of hedge effectiveness. In particular, if we buy too much then the hedger might be paying a substantial risk premium for excess units that are actually dragging down the hedge effectiveness, perhaps down to zero. Additionally, (a problem that we do not consider here) a lack of robustness in the combined assessment of hedge ratio, hedge effectiveness and price might lead the hedger to make an inappropriate choice of hedging instrument or reference population.

3.1 Modelling variants

Drawing from Cairns et al. (2013) we will work our way through four variants of the model that incorporate differing types of risk. The simulation model (Section 2.4.1) is based on the stochastic two-population model proposed by Cairns et al. (2011a). Model parameters are estimated under the Bayesian paradigm using Markov chain Monte Carlo (MCMC) methods. The output from this includes information about the joint posterior distribution for both the process parameters and the underlying latent state variables (that is, the age, period and cohort effects).

PC: Full parameter certainty: In this case, simulations are conducted using the same point estimate (the mean of the posterior distribution) for the process parameters and latent state variables. In addition, the parameters of the valuation model are calibrated using the historical data and fitted model up to 2005 (time 0), and not recalibrated at the end of 2015 (time $T = 10$). There is no allowance for sampling variation in simulated death counts between 2005 and 2015.

PC-R: Full parameter certainty but with recalibration in 2015: This case is the same as the full PC variant except that the valuation model is recalibrated at the end of 2015. Thus, the simulation component of the exercise is unchanged from the full PC variant, but the calibration of the valuation model at time T now depends on the simulated sample path between times 0 and T (as well as the single historical sample path up to time 0). There is no allowance for sampling variation in simulated death counts between 2005 and 2015.

PU: Full parameter uncertainty with recalibration: The simulation model now incorporates full parameter uncertainty, drawing values for the process parameters and the latent state variables up to the end of 2005 from the posterior distribution. The valuation model is calibrated at the end of 2015 using simulated data up to the end of 2015. There is no allowance for sampling variation in simulated death counts between 2005 and 2015.

PU-Poi: Full PU with recalibration + Poisson risk: This case is the same as the PU variant except that death counts between 2005 and 2015 now include levels of sampling variation consistent with the historical sizes of the two populations and with the Poisson assumption (see Cairns et al., 2011a).⁵ Specifically, population 1 exposures are at the same level as EW in 2005, while population 2 exposures are 1% of the CMI exposures in 2005. This means that population 2 has approximately 5900 lives across ages 50 to 89: typical of a medium/large sized pension plan that might be considering hedging options.

The variants are summarised in Table 1.

⁵For alternatives to the Poisson assumption, see, for example, Li, Hardy and Tan (2009).

Variant	Simulation Model	Valuation Model	Poisson Risk
PC	PC	calibrated in 2005	No
PC-R	PC	recalibrated in 2015	No
PU	PU	recalibrated in 2015	No
PU-Poi	PU	recalibrated in 2015	Yes

Table 1: Inclusion of risk factors in the four modelling variants.

4 Numerical results

Recall from Section 2.2 that we have three options for our choice of hedging instrument: deferred longevity swaps, q -forwards maturing at T , and q -forwards maturing at some later date $T + U$. Initially, we will just consider the first two, but we will then demonstrate that the use of longer-dated q -forwards offers some advantages over contracts that mature at time T itself.

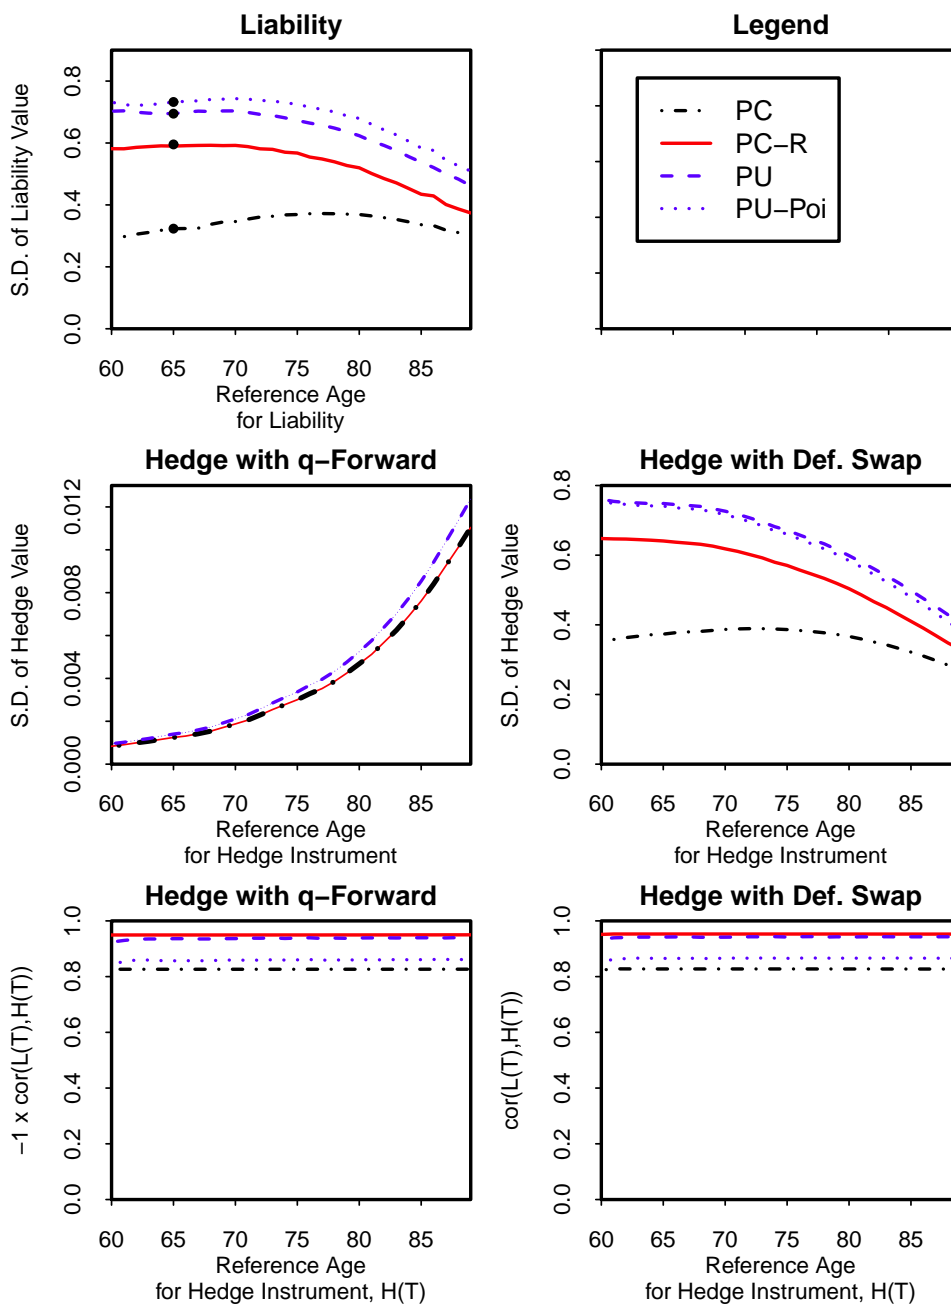


Figure 2: Key statistics for $L(T) = a_2(T, x)$ and $H(T) = q_1(T, y) - q_1^F(0, T, y)$ or $a_1(T, y) - a_1^F(0, T, y)$. Top left: $S.D.(L(T))$ against reference age x under the PC, PC-R, PU and PU-Poi variants: dots highlight the key liability reference age of $x_0 = 65$. Middle row: $S.D.(H(T))$ against reference age y for (left) q -forwards and (right) deferred longevity swaps. Bottom row: $\text{cor}(L(T), H(T))$ for the fixed liability reference age $x_0 = 65$.




Figure 3: Optimal hedge ratios using a single hedging instrument under the PC, PC-R, PU and PU-Poi modelling variants plotted against the reference age of the instrument. Left: q -forwards maturing at time $T = 10$. Right: deferred longevity swaps.

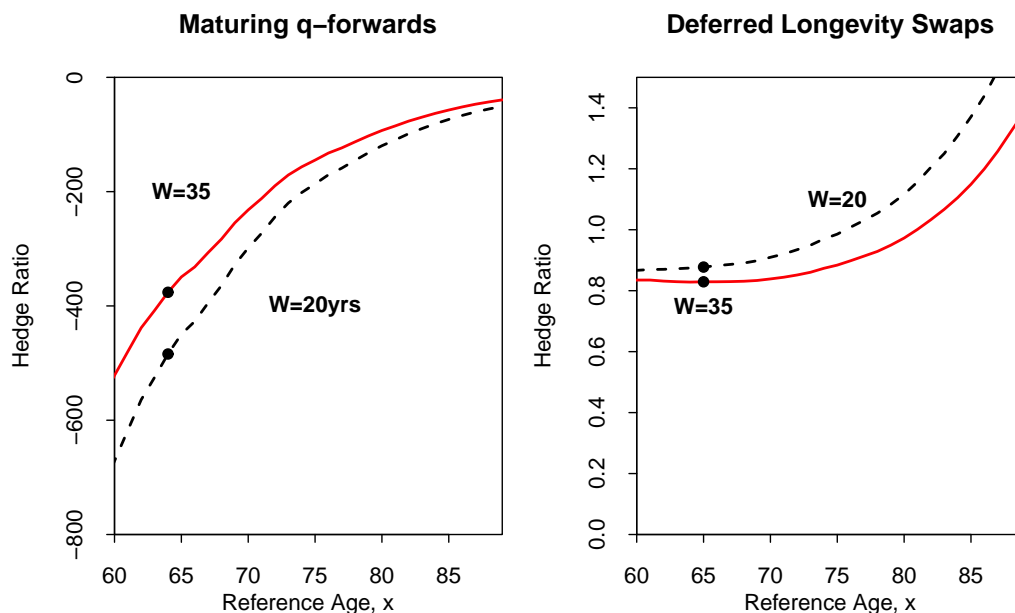


Figure 4: Optimal hedge ratios using a single hedging instrument under the PC-R modelling variant, and assuming a lookback window of $W = 20$ (dashed line) or $W = 35$ (solid line) years plotted against the reference age of the instrument. Left: q -forwards maturing at time $T = 10$. Right: deferred longevity swaps.

4.1 Single-asset static strategies

In Figure 2 we have plotted some statistics that contribute to the hedge ratios plotted in Figure 3. The top left plot shows how the standard deviation of the liability at time $T = 10$ varies with the reference age x . The larger dots highlight age $x_0 = 65$ which is used for the liability in the remainder of the paper. These standard deviations can be compared to a mean liability of around 15 (dependent on the modelling variant). It can be seen that there is some variation with age, as we might expect. More interesting, we can see that ignoring elements such as recalibration risk, parameter uncertainty and even Poisson risk can significantly understate the level of risk being carried by a pension plan, although, of these, Poisson risk is the least important.⁶ Similar differences exist with the deferred annuity as a hedging instrument (middle row, right). q -forward payoffs are not at all dependent on recalibration at T so the PC and PC-R curves are identical (middle row, left). Additionally, the PU and PU-Poi curves are indistinguishable because of the large size of the EW population. The bottom row of Figure 2 shows correlations under the four variants. q -forward correlations have been multiplied by -1 to aid comparison.⁷ We can see that these curves are fairly flat, reflecting the fact that the Age-Period-Cohort model has a single random period effect for each population. The PC-R correlation (solid line) is much higher than the PC correlation (lower dotted line) and this results from an additional common source of risk (reflected most obviously in the top left and middle right plots), namely the recalibrated improvement rate (for further discussion, see Cairns et al. 2013).⁸ The addition of Poisson risk can be seen to reduce correlations significantly (the PU dashed line drops to the PU-Poi line (upper dotted line)). However, the results in Figure 3 indicate that, although Poisson risk causes $\text{cor}(L, H)$ to fall, the covariance, $\text{Cov}(L, H)$, remains almost unchanged.⁹

In Figure 3 we have plotted the estimated hedge ratio ($h = -\text{Cov}(L, H)/\text{Var}(H)$) under the four variants PC, PC-R, PU, PU-Poi. In both plots, hedge ratios have been calculated for a range of reference ages for the chosen hedging instrument. The left-hand plot of Figure 3 shows hedge ratios using q -forwards, while the right-hand

⁶For further discussion of the impact of Poisson risk on run-off surplus, see Richards and Currie (2009).

⁷The two sets of correlations in the bottom row of Figure 2 look almost identical but have small differences. The similarity follows from the fact that uncertainty in the payoffs of the two contracts is driven primarily by uncertainty in $\kappa^{(1)}(T)$. The correlations differ because (a) there is a slight non-linearity in the relationship between the q -forward payoff and the deferred longevity swap payoff, and (b) differing impacts of parameter uncertainty and Poisson risk.

⁸A stylised example is the following. Let $X_1 = \mu + Z_1$ and $X_2 = \mu + Z_2$ where Z_1 and Z_2 are independent standard normal random variables. If μ is assumed to be known, then $\text{cor}(X_1, X_2) = 0$. If we treat μ as an additional source of risk that is independent of Z_1 and Z_2 , then both the variances of X_1 and X_2 are higher. However, additionally, $\text{Cov}(X_1, X_2) = \text{Var}(\mu) > 0$ and so $\text{cor}(X_1, X_2)$ is increased.

⁹We abbreviate $L(T)$ by L and $H(T)$ by H wherever this will not cause confusion.

plot shows corresponding results for deferred longevity swaps.

We can see that, under all variants, there is significant variation using q -forwards (left-hand plot) in the hedge ratio over the range of reference ages. Primarily, this reflects the magnitude of the expected value of $q_1(T, x)$, whereas the estimates of hedge effectiveness (see the correlation plots in Figure 2 as a proxy for hedge effectiveness) are relatively insensitive to the choice of reference age. This reflects the dependence in this stochastic model of future mortality on a single period effect.

For the deferred longevity swap, competing factors come into play that balance each other out to some extent. For younger reference ages, the annuity price at T values payments over a longer period of time and uncertainty in $\kappa^{(1)}(T)$ feeds through to greater variation in these distant cashflows. For a higher reference age, the shorter term cashflows become subject to greater uncertainty, but this is balanced by longer term cashflows being smaller and less significant in terms of their contribution to the risk profile.

Both plots reveal that there is a significant change in the magnitude of the hedge ratio linked to the inclusion of recalibration risk. However, the further inclusion of other forms of parameter uncertainty and Poisson risk (the PC-R, PU and PU-Poi curves are hard to distinguish) have only a small impact on the estimated hedge ratio.

We can conclude from this that, for this model, the hedge ratio is robust relative to the inclusion of parameter uncertainty and Poisson risk, but not robust relative to the inclusion of recalibration risk. But, concerning the latter point, we can also remark that deferred longevity swaps would be preferred to q -forwards maturing at T since the the difference between the PC and PC-R/PU/PU-Poi hedge ratios is smaller.

The conclusion that the hedge ratio is robust relative to the inclusion of parameter uncertainty and Poisson risk contrasts with estimates of the corresponding hedge effectiveness illustrated in Cairns et al. (2013). There, it was demonstrated that the inclusion of parameter uncertainty and Poisson risk could have a moderate impact on the hedge effectiveness (but less than the inclusion of recalibration risk unless population 2 is much smaller than that considered here).

Within the constraints of an index-based hedge, this would not be a problem. However, if alternatives to the index-based hedge are also being considered, then a comparison of the hedge effectivenesses is appropriate, and so inclusion of parameter uncertainty and Poisson risk is important, especially if the hedger's population is relatively small.

4.2 Theoretical support

The observation, based on these numerical results, that optimal hedge ratios are robust relative to the inclusion of parameter uncertainty and Poisson risk can be supported by developing a theoretical argument.

The liability is approximately linear in the $\beta^{(2)}(y)$ (for $y \geq x$), $\kappa^{(2)}(T)$, $\gamma^{(2)}(T-x+1)$ and ν_κ (see Cairns et al., 2013) within the likely range of these parameters.

Start with a given liability, L , and hedge instrument value, H . The optimal hedge ratio is $h = \text{Cov}(L, H)/\text{Var}(H)$. Now suppose we replace L by $\tilde{L} = L + Z_L$ where Z_L is additional noise that is uncorrelated with L and H . This lack of correlation means that $\text{Cov}(\tilde{L}, H) = \text{Cov}(L, H)$ and so the optimal hedge ratio $\tilde{h} = h$ is unchanged. On the other hand, if H is replaced by $\tilde{H} = H + Z_H$ where Z_H is uncorrelated with L , H and Z_L , $\text{Cov}(\tilde{L}, \tilde{H})$ will still equal $\text{Cov}(L, H)$ but $\text{Var}(\tilde{H})$ will be larger and so the optimal hedge ratio will be smaller.

The inclusion of parameter uncertainty in simulations will have a greater impact on population 2: this population is much smaller and so Poisson variation in the historical data up to 2005 will result in greater uncertainty in estimates of the latent state variables.¹⁰ But, arguably, this additional parameter uncertainty will be largely uncorrelated with the future uncertainty in $\kappa^{(1)}(T)$ and $\kappa^{(2)}(T)$. Thus, we conclude that additional noise resulting from the inclusion of parameter uncertainty contributes to approximately uncorrelated additional risks Z_L and Z_H to L and H respectively. As remarked above, Z_L has no impact on the optimal hedge ratio, while the empirical results suggest that Z_H is insignificant given the very small (but nevertheless observable) decrease in the magnitude of the PU curve relative to the PC-R curve (Figure 3).

A similar argument applies when we incorporate Poisson risk in simulated death counts in 2006 to 2015. So long as population 1 is relatively large, the resulting Z_H additional risk will be insignificant. So long as Poisson risk results in additional uncorrelated randomness in L , the size of population 2 does not play a role in the setting of the hedge ratio (other than scaling up to account for the total amount annuity payable).

Key points concerning the use of this theoretical approximation are as follows. First, it helps us to develop an intuitive feel for the more complex numerical results. Second, we emphasize that similar theoretical approximations can be developed for other stochastic mortality models to help develop hedging strategies in these other circumstances.

¹⁰See, for example, Cairns et al., 2011a, Figure 4.

4.3 Robustness relative to the recalibration window

Figure 4 investigates robustness relative to the choice of the recalibration window: $W = 20$ or 35 years under the PC-R variant (Figure 3 uses $W = 20$). Again we see that the estimated hedge ratio for both hedging instruments lacks robustness relative to the choice of W . And, again, we see that although the deferred longevity swap strategy lacks robustness, it is still to be preferred to a q -forward strategy, where the q -forward matures at the valuation date T .

These differences between the robustness of q -forwards and deferred longevity swaps are linked to their dependence, and that of the underlying liability, $L = -a_2(T, 65)$, on $\kappa^{(1)}(T)$ and its drift parameter, ν_κ .

The q -forward payoff, $q(T, x)$, is dependent on $\beta^{(1)}(x)$, $\kappa^{(1)}(T)$ and $\gamma^{(1)}(T - x)$. In the PC and PC-R variants, $\beta_x^{(1)}$, and $\gamma^{(1)}(T - x)$ are constant. In the PU and PU-Poi variants they are subject to only modest amounts of uncertainty relative to uncertainty in $\kappa^{(1)}(T)$. So, as a first approximation, uncertainty in the q -forward payoff depends solely on uncertainty in $\kappa^{(1)}(T)$. A similar argument (see Section 2.4.2), reveals that uncertainty in the value of the deferred longevity swap at T depends upon $\kappa^{(1)}(T)$ and the recalibrated drift, ν_κ . In Figure 4, the large relative change in the q -forward hedge ratio when W changes from 20 to 35 reflects the q -forward's lack of dependence on ν_κ and the fact that ν_κ (which affects the liability valuation) depends directly on $\kappa^{(1)}(T)$. The more modest change in the deferred longevity swap hedge ratio reflects the different relative dependence of the liability and the hedging instrument on $\kappa^{(1)}(T)$ and ν_κ .¹¹

4.4 Using longer-dated q -forwards

In order to address the perceived deficiencies of the q -forward maturing at T relative to a deferred longevity swap, we now consider the use of q -forwards that mature at later dates than T , since their value at T will depend now on ν_κ . Specifically, the revised hedging instrument will be the $(T + U, 54 + T + U)$ q -forward: that is, as we extend the maturity date we continue to refer to the same cohort aged $54 + T$ in year T .

In Figure 5 we investigate how the hedge ratio varies with outstanding term to maturity, U , when the recalibration window W is 20 or 35 years. As in Figure 3, the decreases in magnitude of the hedge ratio simply reflect the increasing magnitude and riskiness in absolute terms of the q -forward as the reference age increases. However, we also see that the ratio between the $W = 35$ and $W = 20$ curves narrows as U increases (an extra 28% when $U = 0$, compared with an extra 11% when

¹¹The greater sensitivity of the deferred longevity swap hedge ratio to W for higher reference ages reflects the fact that the underlying annuity price becomes relatively more sensitive to changes in $\kappa^{(1)}(T)$ and less sensitive to ν_κ .

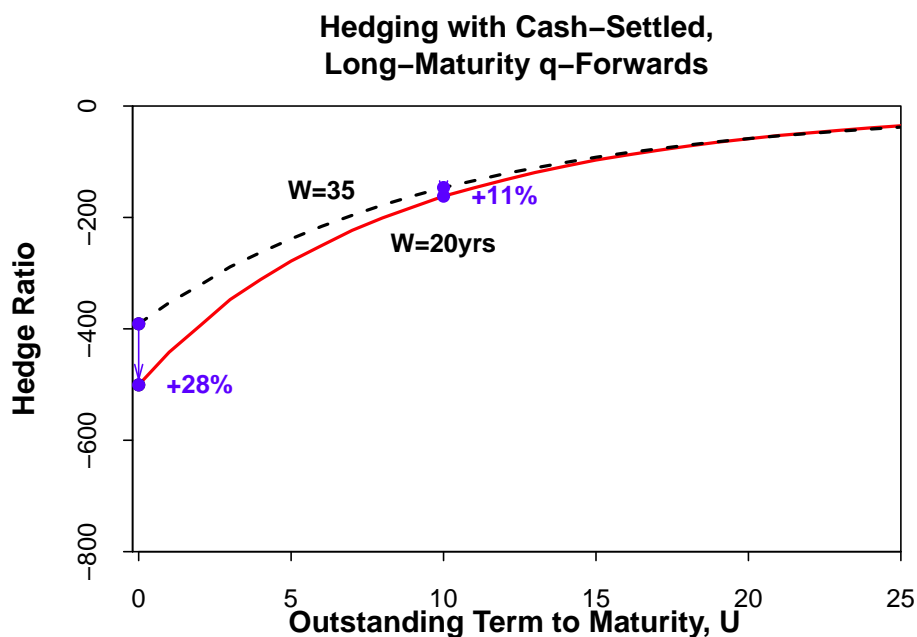


Figure 5: Optimal hedge ratios using a single long-maturity q -forward as the hedging instrument under the PC-R modelling variant. Lookback window is $W = 20$ (dashed line) or $W = 35$ (solid line) years. Plotted shows hedge ratio against the outstanding maturity, U , after time $T = 10$ of the q -forward. The $U = 0$ dots correspond to the $x = 65$ dots in Figure 4 (left-hand plot).

$U = 10$), and this reflects the increasing relative dependence of the value at time T of the q -forward on ν_κ .

Thus, for the liability $L = a_2(T, 65)$, a q -forward with a maturity of more than $T + U = 20$ years seems to be comparable with a deferred longevity swap.¹² For example, it might be linked in some way to the *duration* of the liabilities. This is helpful, since the q -forward contract is a simpler contract and much more likely to find a place in the index-based longevity market.

¹²We leave for future work an investigation of whether a general rule of thumb can be developed for choosing the most appropriate value of U .

5 Improving robustness: Nuga hedging

5.1 Motivation

In the model with recalibration at T , ν_κ is equal to $(\kappa^{(1)}(T) - \hat{\kappa}^{(1)}(T - W))/W$. It follows that if $T - W$ is in the past (that is, not after time 0), and if we know the value of W , then uncertainty in ν_κ can simply be represented by a linear function of $\kappa^{(1)}(T)$. The liability at T can then be decomposed into an approximate linear function of $\kappa^{(1)}(T)$ plus additional uncorrelated risks, and the hedge instrument value at T will, approximately, be primarily a linear function of $\kappa^{(1)}(T)$. So, if W is known, then it will be sufficient in the context of this stochastic model, to use a single hedging instrument and to calculate hedge ratios using a methodology that incorporates recalibration risk.

In reality, W is not known, and the valuation ν_κ might even be calculated using other unknown methods. Thus, W and ν_κ can be regarded as being exposed to sources of Knightian uncertainty, since there is no data that we can use to make the calculation of ν_κ precise: it is a matter of judgment in the hands of those who value and price the liabilities and hedging instruments. So we have an additional source of risk on top of $\kappa^{(1)}(T)$, suggesting that the use of a second hedging instrument might help us to find a more robust solution to the hedging problem.

We will use the expression *Nuga hedging*¹³ to refer to strategies that aim to minimise our exposure to risk associated with the choice of W and the method used to recalibrate the random-walk drift ν_κ . The idea behind Nuga hedging is similar in concept to Vega hedging in equity derivatives. Specifically, we aim to reduce the risk associated with recalibration of a model parameter that is supposed to be constant. Since Vega hedging refers to a volatility parameter rather than a drift parameter we have chosen to refer to hedging against the risk in ν_κ using a different name.

5.2 Delta and Nuga

The hedging strategies employed in Section 4 can be regarded as Delta hedging strategies on the assumption that ν_κ is an explicit linear function of $\kappa^{(1)}(T)$. We will now introduce Delta hedging in its proper form where we aim to hedge against uncertainty in the state variable $\kappa^{(1)}(T)$ on the assumption that all other parameters and state variables in the model remain unchanged.

The valuation model outlined in Section 2.4.2 allows us to simulate annuity prices at future dates. Specifically, let x_0 be the age at time T of a cohort for whom we wish to price an annuity.

Now suppose that we have data up to time 0 and that we seek to hedge the risk at

¹³The name Nuga is derived from the drift parameter, ν_κ .

time T . Assume, also, that we already have estimates for $\gamma^{(k)}(T - x_0 + 1)$: that is, our dataset up to time 0 includes age $x_0 - T - 1$. This means that, at least in a parameters certain world, there is no future uncertainty in the value of the cohort effect for the cohort aged x_0 at time T . For hedging purposes, therefore, $a_k(T, U, x_0)$ is a function of $\kappa^{(k)}(T)$ and ν_κ .

Using the deterministic approach to valuation and formulae in Section 2.4.2, we have:

$$\begin{aligned} \frac{\partial \hat{p}^{(k)}(T, T + u, x_0)}{\partial \kappa^{(k)}(T)} &= -\hat{p}^{(k)}(T, T + u, x_0) \sum_{t=1}^u \hat{m}^{(k)}(T + t, x_0 + t - 1) / n_a \\ &\text{and} \\ \frac{\partial \hat{p}^{(k)}(T, T + u, x_0)}{\partial \nu_\kappa} &= -\hat{p}^{(k)}(T, T + u, x_0) \sum_{t=1}^u t \hat{m}^{(k)}(T + t, x_0 + t - 1) / n_a \\ \Rightarrow \frac{\partial a_k(T, U, x_0)}{\partial \kappa^{(k)}(T)} &= -\sum_{u=1}^U e^{-ru} \hat{p}^{(k)}(T, T + u, x_0) \sum_{t=1}^u \hat{m}^{(k)}(T + t, x_0 + t - 1) / n_a \\ &\text{and} \\ \frac{\partial a_k(T, U, x_0)}{\partial \nu_\kappa} &= -\sum_{u=1}^U e^{-ru} \hat{p}^{(k)}(T, T + u, x_0) \sum_{t=1}^u t \hat{m}^{(k)}(T + t, x_0 + t - 1) / n_a. \end{aligned}$$

We now use the first order Taylor approximation to write

$$L = a_2(T, x) \approx L_0 + \Delta_L(\kappa^{(2)}(T) - \hat{\kappa}^{(2)}(T)) + \mathcal{N}_L(\nu_\kappa - \hat{\nu}_\kappa), \quad (5)$$

where

$$\Delta_L = \frac{\partial a_2(T, U, x_0)}{\partial \kappa^{(2)}(T)} \text{ and } \mathcal{N}_L = \frac{\partial a_2(T, U, x_0)}{\partial \nu_\kappa}.$$

All Deltas and Nugas are evaluated at the best estimates of $\kappa^{(1)}(T)$, $\kappa^{(2)}(T)$ and ν_κ using data available at time 0, since the static hedging strategy must be determined at time 0.

Similarly, we have two hedging instruments with value at T ,

$$H_i \approx H_{i0} + \Delta_{H_i}(\kappa^{(i)}(T) - \hat{\kappa}^{(i)}(T)) + \mathcal{N}_{H_i}(\nu_\kappa - \hat{\nu}_\kappa). \quad (6)$$

For notational convenience, write $\epsilon_i = \kappa^{(i)}(T) - \hat{\kappa}^{(i)}(T)$ and $\epsilon_\nu = \nu_\kappa - \hat{\nu}_\kappa$.

5.3 Delta hedging only

Simple Delta hedging ignores the risk associated with ϵ_ν meaning that hedging can be done with a single asset, in other words mimicking the PC variant in Section 4. Thus, hedging with hedging instrument 1, we adopt a hedge ratio of

$$h = \frac{\Delta_L}{\Delta_{H1}} \frac{\text{Cov}(\epsilon_1, \epsilon_2)}{\text{Var}(\epsilon_1)}.$$

In this way $\text{Var}(L - hH_1)$ would be minimised under the PC variant.

5.4 Delta and Nuga hedging

Next we seek to choose h_1 and h_2 to minimise $\text{Var}(L - h_1H_1 - h_2H_2)$ without making any assumptions about the distribution of ν_κ (i.e. ϵ_ν) or its dependence on ϵ_1 and ϵ_2 .

First, we choose to write $\epsilon_2 = \alpha\epsilon_1 + \tilde{\epsilon}_2$, where the specific choice of $\alpha = \text{Cov}(\epsilon_1, \epsilon_2)/\text{Var}(\epsilon_1)$ results in $\text{Cov}(\epsilon_1, \tilde{\epsilon}_2) = 0$. We will also make the assumption that $\text{Cov}(\tilde{\epsilon}_2, \epsilon_\nu) = 0$ (consistent with the assumption that, whatever the method of recalibration is, ν_κ is set with reference to population 1 only). Then,

$$\begin{aligned} L &\approx L_0 + \Delta_L\alpha\epsilon_1 + \Delta_L\tilde{\epsilon}_2 + \mathcal{N}_L\epsilon_\nu, \\ H_1 &\approx H_{10} + \Delta_{H1}\epsilon_1 + \mathcal{N}_{H1}\epsilon_\nu, \\ H_2 &\approx H_{20} + \Delta_{H2}\epsilon_1 + \mathcal{N}_{H2}\epsilon_\nu. \end{aligned} \tag{7}$$

For Delta and Nuga hedging we recognise that nothing can be done about the residual risk $\tilde{\epsilon}_2$, so we seek to neutralise $\Delta_L\alpha\epsilon_1$ and $\mathcal{N}_L\epsilon_\nu$ through holding h_1 and h_2 units respectively of H_1 and H_2 . This is achieved through solving:

$$h_1\Delta_{H1} + h_2\Delta_{H2} = \alpha\Delta_L \tag{8}$$

$$\text{and } h_1\mathcal{N}_{H1} + h_2\mathcal{N}_{H2} = \mathcal{N}_L. \tag{9}$$

It then follows that the revised deficit at T is

$$P = L - h_1H_1 - h_2H_2 \approx L_0 + \Delta_L\tilde{\epsilon}_2.$$

In other words, the only risk we are left with is the uncorrelated population 2 specific risk, $\Delta_L\tilde{\epsilon}_2$.

Remark 5.4.1

Assume exact, rather than approximate, linearity in equations (7). Additionally, assume that $\nu_\kappa = (\kappa^{(1)}(T) - \kappa^{(1)}(T - W))/W$ for a known value of W , implying that $\epsilon_\nu = \epsilon_1/W$, $\text{cor}(\epsilon_1, \epsilon_\nu) = 1$ and $\text{Cov}(\epsilon_\nu, \tilde{\epsilon}_2) = 0$. Since W is known, one hedging instrument would be sufficient. Here we have two, and so there are many different combinations of (h_1, h_2) that minimise risk. Furthermore, these pairs of risk minimising (h_1, h_2) lie on a straight line that depends upon the known W . Different W 's will result in different straight lines, but all of these lines will intersect at the same point: the Delta and Nuga hedged portfolio.

Remark 5.4.2

With the same assumptions as Remark 5.4.1, the minimum residual risk will always be $\Delta_L^2\text{Var}(\tilde{\epsilon}_2)$ and the optimal hedge effectiveness

$$R^2 = 1 - \frac{\Delta_L^2\text{Var}(\tilde{\epsilon}_2)}{\text{Var}(L)}.$$

However, note that $\text{Var}(L)$ and, hence, the optimal R^2 depend on the method of calibrating ν_κ .

5.5 Numerical example 1: hedging with deferred longevity swaps

In our first example, we assume that the hedger wishes to use deferred longevity swaps that are linked to two cohorts aged 65 and 85 at time T , payable until some high age (we have assumed age 130). So

$$H_1 = a_1(T, 65) - a_1^F(0, T, 65), \quad \text{and} \quad H_2 = a_1(T, 85) - a_1^F(0, T, 85).$$

We compare four strategies using these two hedging instruments.

- Strategy A: no hedging.
- Strategy B: $h_1 = 0.8775$, $h_2 = 0$. This strategy is optimal if we know that ν_κ is recalibrated using a $W = 20$ -year window.
- Strategy C: $h_1 = 0.8291$, $h_2 = 0$. This strategy is optimal if we know that ν_κ is recalibrated using a $W = 35$ -year window.
- Strategy D: $h_1 = 1.3376$ and $h_2 = -0.7199$. This is the Delta-Nuga hedge solving equations (8) and (9).

The calculation of all of these hedge ratios uses simulated data and, therefore, is subject to simulation error. For strategy B the standard error of the estimate of h_1 is 0.0088 and strategy C 0.0110. For strategy D, simulation error only affects the solution to equations (8) and (9) through the parameter $\alpha = \text{Cov}(\epsilon_1, \epsilon_2) / \text{Var}(\epsilon_1) = \text{Cov}(\kappa^{(1)}(T), \kappa^{(2)}(T)) / \text{Var}(\kappa^{(1)}(T))$. This leads, in strategy D, to standard errors of 0.0103 and 0.0299 for the estimates for h_1 and h_2 respectively. In all cases, potential errors in the hedge ratios have been verified to have no material impact on the conclusions that we draw below.

We now look at the performance of these three strategies if, in reality, $W = 20$ or $W = 35$. Results are given in Table 2.

These results are based on 1000 simulated scenarios in each of the two cases for W . For $W = 20$ (upper half of Table 2), strategy B is the best for these 1000 scenarios and so C and D cannot be any better. Theoretically, strategy D would be just as good, but because it has been determined using a numerical method not involving the simulated scenarios, and because of slight non-linearities with respect to the $\kappa^{(k)}(T)$ and ν_κ in the annuity functions, strategy D is not quite as effective as strategy B. However, we see that although strategy C is worse, the differences

Strategy	H_1 h_1	H_2 h_2	Var(Surplus)	Hedge Effectiveness	Rank
$W = 20$					
A	0	0	0.3481	0	
B	0.8775	0	0.03202	0.9080	(1)
C	0.8291	0	0.03298	0.9052	(3)
D	1.3376	-0.7199	0.03209	0.9078	(2)
$W = 35$					
A	0	0	0.2233	0	
B	0.8775	0	0.03353	0.8498	(3)
C	0.8291	0	0.03289	0.8527	(1)
D	1.3376	-0.7199	0.03298	0.8523	(2)

Table 2: Description and outcome of hedging strategies A, B, C and D, using deferred longevity swaps, under the PC-R variant with lookback windows of $W = 20$ (upper half of table) or $W = 35$ (lower half) years. $H_1 = a_1(T, 65) - a_1^F(0, T, 65)$, and $H_2 = a_1(T, 85) - a_1^F(0, T, 85)$.

between the hedge effectiveness of strategies B, C and D are all very small. So, in this case, the benefits of Delta-Nuga hedging are not that great.

When $W = 35$ (lower half of Table 2), strategy C is now best by definition, but again the differences between strategies B, C and D are minimal.

The reason for these small differences is linked to the results displayed in Figure 4. The moderate difference in the hedge ratios 0.8775 and 0.8291 converts into a much smaller reduction in hedge effectiveness (which is quadratic in the hedge ratio). The key point is that for strategies B and C, if ν_κ decreases, then the values of both the liability and the hedge instrument will react in similar ways. This contrasts with the use of a q -forward maturing at time T which we will see in our next example. In this case, changing ν_κ will have an impact on the liability, but there will be no impact on the q -forward payoff at T .

5.6 Numerical example 2: hedging with q -forwards

In our second example we switch to the use of q -forwards for hedging. Instrument 1 is a q -forward maturing at time $T = 10$ linked to age 64, and instrument 2 is a q -forward maturing at time 35 linked to age 89 (so both are linked to the same cohort, but mature 25 years apart).

We compare four strategies using these two hedging instruments.

- Strategy A: no hedging.

- Strategy B: $h_1 = -500.7$ (S.E.=5.2), $h_2 = 0$. This strategy is optimal if we know that ν_κ is recalibrated using a $W = 20$ -year window.
- Strategy C: $h_1 = -391.0$ (S.E.=5.2), $h_2 = 0$. This strategy is optimal if we know that ν_κ is recalibrated using a $W = 35$ -year window.
- Strategy D: $h_1 = -33.2$ (S.E.= 5.1) and $h_2 = -34.8 \exp(25r) = -94.5$ (S.E.= 0^{14}) (where $r = 0.04$ is the discount rate). This is the Delta-Nuga strategy.

Strategy	H_1 h_1	H_2 h_2	Var(Deficit)	Hedge Eff.
$W = 20$				
A	0	0	0.3481	0
B	-500.7	0	0.03435	0.9013 (1)
C	-391.0	0	0.04940	0.8581 (3)
D	-33.2	-94.5	0.04012	0.8848 (2)
$W = 35$				
A	0	0	0.2233	0
B	-500.7	0	0.04882	0.7714 (3)
C	-391.0	0	0.03392	0.8481 (1)
D	-33.2	-94.5	0.03516	0.8426 (2)

Table 3: Description and outcome of hedging strategies A, B, C and D, using deferred longevity swaps, under the PC-R variant with lookback windows of $W = 20$ (upper half of table) or $W = 35$ (lower half) years. $H_1 = q_1(T, 64) - q_1^F(0, T, 64)$, and $H_2 = \exp(-25r) \{E_Q[q_1(T + 25, 89)|\mathcal{M}_T] - q_1^F(0, T + 25, 89)\}$.

Table 3 shares many basic characteristics with Table 2. In a perfectly linear world, strategy B in the $W = 20$ case would deliver the same optimal hedge effectiveness in both Tables 2 and 3. The result in Table 3 is slightly less good, indicating that the minor non-linearities in $a_2(T, 65)$ are slightly better matched by the non-linearities in $a_1(T, 65)$ than $q_1(T, 64)$. In contrast to Table 2, we see that, when $W = 20$, strategy C has a significantly lower hedge effectiveness than strategy B, and, when $W = 35$, B is much worse than C. However, in both cases, strategy D, the Delta-Nuga hedging strategy, produces results that are robust in the sense that its hedge effectiveness is very close to the optimal value.

Again, in all cases, potential errors in the hedge ratios have been verified to have no material impact on the conclusions that have been drawn, including in Strategy D where the standard error of h_1 seems quite large.

¹⁴Standard errors for strategy D depend only on uncertainty in our estimate of the parameter α in the solutions to equations (8) and (9). For our chosen hedging instruments, $\mathcal{N}_{H1} = 0$ and so $h_2 = -\mathcal{N}_L/\mathcal{N}_{H2}$, which is not dependent on α . So the standard error of h_2 is zero.

5.7 Further remarks

In the second numerical example, we also considered hedging with 10-year and 20-year q -forwards (instead of 10 and 35 years). In this case the Delta-Nuga hedge required $h_1 = 279.6$ and $h_2 = -382.5$ units. A long-short position like this might give rise to additional counterparty risk and also greater exposure to model risk.

The 10 and 35 year maturities for the q -forwards were chosen to straddle the relative sensitivity of the liability to uncertainty in $\kappa_1(T)$ and ν_κ .¹⁵ The use of 10 and 20-year maturities does not satisfy this requirement.

We have assumed here that valuation models are calibrated using a version of the Age-Period-Cohort model. In this case, Delta-Nuga hedging will deal with a range of mechanisms for calibrating the ν_κ parameter. We leave for further work an analysis of robustness relative to the choice of valuation model. However, one might conjecture that hedging instruments that share fewer characteristics with the underlying liability (such as q -forwards relative to deferred longevity swaps) are less robust relative to model risk.

6 Robustness relative to other factors

The paper has focused mainly on the inclusion of recalibration risk. However, we have also discussed how results for hedge ratios appear to be robust relative to the inclusion of other forms of parameter risk and Poisson risk. At the same time, we remarked earlier that hedge effectiveness is slightly less robust relative to these same two risk factors.

6.1 Poisson risk

Consider Poisson risk a bit further. Intuition might suggest that the smaller the pension plan, the less useful an index-linked hedging instrument might be. But we have demonstrated empirically with support from general reasoning, that the optimal hedging strategy should not be affected *disproportionately* by the size of the pension plan, if we take it as given that an index-based hedge is to be used. Additionally, in the context of exponential utility, we provide an argument in Appendix A that the price that a plan should be prepared to pay per unit of hedging instrument should not be sensitive to the size of the plan, provided large and small plans have the same exponential utility function at time T .

¹⁵A helpful analogy lies in interest-rate risk management. If we are hedging against parallel shifts in the yield curve and have a positive liability with a duration of D_L years, then if we hedge with two zero-coupon bonds, one with a maturity less than D_L and one with a higher duration, then the duration matched position will have positive holdings of both of the zero-coupon bonds.

On the other hand, the results of Cairns et al. (2013) and Li and Luo (2012) both point to a reduction in hedge effectiveness when we move from a large pension plan to a small plan. From our remarks above on the robustness of the index-based hedging strategy itself, we conclude that a small plan should only consider adoption of a different strategy if there is an alternative to non-index-based hedging that results in a higher hedge effectiveness or at lower cost. As remarked at the beginning of this paper, customised longevity swaps are only available to large pension plans. But small pension plans do have access to the individual annuities market. These will be significantly more expensive than a longevity swap or an index-based hedge, but would result in a much higher reduction in risk (e.g. variance) for a small pension plan compared with an index-based hedge. So there will be a tradeoff, and, even though the index-based hedge is robust relative to plan size, we would argue that there will be a threshold plan size, below which individual annuitisation would be preferred to an index-based hedge.

6.2 Using the latest EW data

The analysis so far has used EW and CMI data up to 2005, this being the latest year for which we have CMI data. But EW data is available up to 2008 (www.mortality.org).

The Bayesian methodology of Cairns et al. (2011a) allows the 2-population model to be fitted using EW data up to 2008 while treating CMI data for 2006-2008 as missing data. Once implemented, it allowed us to refine estimates of the joint distribution of all relevant state variables at the end of 2015 (especially the distribution of $(\kappa^{(1)}(T), \kappa^{(2)}(T))$). We found empirically that a single instrument hedge (assuming $W = 20$) was robust relative to the incorporation of EW data for 2006-2008. If $(\kappa^{(1)}(t), \kappa^{(2)}(t))$ was a bivariate random-walk with correlated innovations, it can be demonstrated that additional knowledge of $\kappa^{(1)}(t)$ for, say, $t = 1, 2, 3$, would not alter the optimal hedge ratio,¹⁶ while the optimal hedge effectiveness would fall. In our model, however, $(\kappa^{(1)}(t), \kappa^{(2)}(t))$ is not a bivariate random walk and so there is not a theoretical argument to support our empirical observation of robustness.¹⁷ However, it is probably the slow rate of mean reversion of the spread between $\kappa^{(1)}(t)$ and $\kappa^{(2)}(t)$ that results in something close to a bivariate random walk up to time T , and, hence, empirically we see that the hedging strategy is robust.

Finally, we note that, although we have not evaluated it, the price would change to reflect the altered mean and covariance matrix for $(\kappa^{(1)}(T), \kappa^{(2)}(T))$.

¹⁶Individually, $\text{Cov}(L, H)$ and $\text{Var}(H)$ are both scaled by a factor of $(10 - 3)/10 = 0.7$, so the hedge ratio $h = -\text{Cov}(L, H)/\text{Var}(H)$ remains unchanged.

¹⁷That is, $\text{Cov}(L, H)$ and $\text{Var}(H)$ are scaled by different factors.

7 Conclusions

In this paper we reinforce the earlier conclusions of Cairns et al. (2013) in finding that simple hedging strategies are not robust relative to the inclusion of recalibration risk. In contrast, hedging strategies are found to be robust relative to the inclusion of parameter uncertainty, Poisson risk and the latest EW data.

We have found that the use of index-based q -forwards that mature on our target valuation date, T , are less robust than contracts (q -forwards or deferred longevity swaps) that have longer maturities. Since a market in q -forwards is more likely than deferred longevity swaps, we recommend that participants in the market give consideration to the issuance of long-dated q -forwards.

We have introduced the notion of Nuga hedging, and we have demonstrated how the use of Delta-Nuga hedging results in strategies that are robust relative to uncertainties in the calibration of the key drift parameter, ν_κ .

As remarked earlier, these are just the first steps in the direction of developing robust hedges in the presence of (Knightian) recalibration methodology risk and model risk.¹⁸ We leave the latter for more detailed analysis in further work.

An additional direction for further research would involve a comparison of dynamic hedging with the results for static hedging in this paper. Dynamic hedging should, in general, be better than static hedging, as, for example, demonstrated in Cairns (2011). However, for the problem being considered here (that is, a path-*independent* value hedging problem where the outcome at time T depends only on the mortality curve at T) the linearisation outlined in equations (5) and (6) turns out to be highly accurate as an approximation to the unlinearised output from the valuation model. As a consequence, we conjecture that, in this case, dynamic hedging would result in only a modest improvement over static hedging.

8 Acknowledgements

I wish to thank David Blake, Kevin Dowd and Guy Coughlan for stimulating discussions on modelling and hedging that provided the motivation for this piece of work. I would also like to thank three anonymous referees for their helpful suggestions on an earlier version. I also acknowledge the Human Mortality Database (www.mortality.org) and the Continuous Mortality Investigation in the UK for providing the data used in the modelling work.

¹⁸See, for example, Cairns et al. (2009) (different models), Richards and Currie (2009) (different calibration methods) and Currie (2011) (extrapolation to higher ages).

A Optimal hedge ratios under exponential utility

We consider here a very simple model to highlight some basic ideas in deriving a hedge ratio.

- It is a one-period model.
- Initial wealth at time 0 is w_0 .
- The hedger has an existing liability to pay an amount $L(T)$ at time T (e.g. $L(T) = a_2(T, 65)$).
- The hedger can invest from time 0 to time 1 in cash earning a risk-free rate of interest of r (continuously compounding), or in a hedging instrument that costs $H(0)$ to buy per unit at time 0 and returns $H(T)$ at time T .

If the hedger buys h units of the hedging instrument at time 0 then her wealth at time 1 is $W_T(h) = w_0e^r - L(T) + h(H(T) - H(0)e^r)$.

- $(H(T), L(T))$ is assumed to be bivariate normal and we write it as

$$\begin{aligned} H(T) &= \mu_1 + \sigma_{11}Z_1 \\ L(T) &= \mu_2 + \sigma_{21}Z_1 + \sigma_{22}Z_2 \end{aligned}$$

where Z_1 and Z_2 are independent standard normal random variables, $\sigma_{11}^2 = \text{Var}(H(T))$, $\sigma_{11}\sigma_{21} = \text{Cov}(H(T), L(T))$ and $\sigma_{21}^2 + \sigma_{22}^2 = \text{Var}(L(T))$.

Assume (without loss of generality) that $\sigma_{11} > 0$, $\sigma_{21} > 0$ and $\sigma_{22} > 0$.

- The hedger has an exponential utility function:

$$u(w) = -\exp(-\gamma w), \quad \text{where } \gamma > 0.$$

It follows that

$$E[u(W_1(h))] = -\exp\left[-\gamma(w_0e^r - \mu_2 - h(\mu_1 - H(0)e^r)) + \frac{1}{2}\gamma^2((h\sigma_{11} - \sigma_{21})^2 + \sigma_{22}^2)\right].$$

The hedger wishes to maximise her expected utility, and this is equivalent to minimising

$$f(h) = h(H(0)e^r - \mu_1) + \frac{1}{2}\gamma(h\sigma_{11} - \sigma_{21})^2.$$

The optimal hedge ratio, h^* , is then¹⁹

$$h^* = \frac{\sigma_{21}}{\sigma_{11}} \left(1 - \frac{(H(0)e^r - \mu_1)}{\gamma\sigma_{11}\sigma_{21}}\right).$$

¹⁹This matches equation (2.3) of Zhou et al. (2011). Zhou et al. develop this idea further to derive the equilibrium price, x_1 , at which supply equals demand. In a non-mortality context, and in a more general setting with multiple investors and additional risky assets, the problem has also been considered by Cairns (2001).

Thus h^* consists of two parts. σ_{21}/σ_{11} is the hedge ratio that minimises the variance of the hedger's wealth at time 1. However, this is scaled back depending on:

- The hedger's risk aversion parameter, γ . The higher γ is, the higher is the hedger's risk aversion and the closer the hedge ratio will be to the variance minimising σ_{21}/σ_{11} .
- The risk premium demanded by the hedge provider, $H(0)e^r - \mu_1$. The higher the risk premium, the smaller the optimal hedge ratio: that is, the hedger will wish to buy less.
- If the covariance, $\sigma_{11}\sigma_{21}$, is larger, then the hedge ratio will be closer to the variance minimising σ_{21}/σ_{11} . Thus for a given price, $H(0)$, the hedger will be prepared to buy more units of $H(T)$ if the *correlation* between $H(T)$ and $L(T)$ is larger.

We can also see that the hedger will only adopt the minimum risk portfolio (i.e. minimum variance) if either the risk premium is zero ($H(0) = \mu_1 e^{-r}$) or the hedger is infinitely risk averse.

The marginal price at which the hedger will begin to go long in the hedging asset is the value of $H(0)$ at which the optimal h^* is equal to 0: that is,

$$H(0)^M = e^{-r} (\mu_1 + \gamma \sigma_{11} \sigma_{21}).$$

If the actual price is below $H(0)^M$, then the hedger will have a positive position ($h^* > 0$). Thus, the hedger will be willing to remain long up to a higher upper limit, $H(0)^M$, if either the hedger is more risk averse or if the covariance between $H(T)$ and $L(T)$ is higher (equivalently, correlation, if the variances are fixed) so that the benefits of hedging are greater.

Important points (assuming exponential utility):

- h^* does not depend on the extent of the unhedgeable risk $\sigma_{22}Z_2$, nor does the minimum variance hedge ratio, σ_{21}/σ_{11} .
- The marginal price below which hedgers will take a long position, $H(0)^M$, is also not dependent on the extent of the unhedgeable risk $\sigma_{22}Z_2$.

In other words, if $H(T)$ is the only available asset for hedging and, in the context of exponential utility, if γ is fixed, then the optimal strategy for a small hedger with more unhedgeable risk would be the same as for a large hedger.

References

- Blake, D., Cairns, A.J.G., and Dowd, K. (2006) Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial Journal*, 12: 153-197.
- Börger, M., Fleischer, D., and Kuksin. N.2012 Modeling mortality trend under modern solvency regimes. Working paper, University of Ulm.
- Booth, H., Maindonald, J., and Smith, L. (2002) Applying Lee-Carter under conditions of variable mortality decline. *Population Studies*, 56: 325-336.
- Cairns, A.J.G. (2001) From financial economics to fair valuation. In *Proceedings of the 11th International AFIR Colloquium*, Toronto, September 2001, Volume 1, pp135-166.
- Cairns, A.J.G. (2011) Modelling and management of longevity risk: approximations to survival functions and dynamic hedging. *Insurance: Mathematics and Economics*, 49: 438-453.
- Cairns, A.J.G., Blake, D., and Dowd, K. (2008) Modelling and management of mortality risk: A review. *Scandinavian Actuarial Journal*, 108: 79-113.
- Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A., and Balevich, I. (2009) A quantitative comparison of stochastic mortality models using data from England & Wales and the United States. *North American Actuarial Journal*, 13: 1-35.
- Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., and Khalaf-Allah, M. (2011a) Bayesian stochastic mortality modelling for two populations, *ASTIN Bulletin*, 41: 29-59.
- Cairns, A.J.G., Blake, D., Dowd, K., and Coughlan, G.D. (2013) Longevity hedge effectiveness: A decomposition. To appear in *Quantitative Finance*.
- Coughlan, G.D., Emery, S. and Kolb, J. (2004) HEAT (Hedge Effectiveness Analysis Toolkit): A consistent framework for assessing hedge effectiveness under IAS 39 and FAS 133, *Journal of Derivatives Accounting*, 1(2): 221-272.
- Coughlan, G., Epstein, D., Sinha, A., and Honig, P. (2007) q -forwards: Derivatives for transferring longevity and mortality risk. Available at www.lifemetrics.com.
- Coughlan, G.D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A.J.G., Blake, D. and Dowd, K., (2011) Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. *North American Actuarial Journal*, 15: 150-176.
- Currie, I.D. (2011) Modelling and forecasting the mortality of the very old. *ASTIN Bulletin*, 41: 419-427.
- Dahl, M., Melchior, M., and Møller, T. (2008) On systematic mortality risk and risk minimisation with survivor swaps. *Scandinavian Actuarial Journal*, 2008(2-3):

114-146.

Dahl, M., Glar, S., and Møller, T. (2009) Mixed dynamic and static risk minimization with an application to survivor swaps. 19th International AFIR Colloquium, Munich, September 2009.

Denuit, M., Haberman, S., and Renshaw, A.E. (2010) Comonotonic approximations to quantiles of life annuity conditional expected present values: Extensions to general ARIMA models and comparison with the bootstrap. *ASTIN Bulletin*, 40: 331-349.

Dowd, K., Blake, D., and Cairns, A.J.G. (2006) Mortality-dependent financial risk measures. *Insurance: Mathematics and Economics*, 38: 427-440.

Dowd, K., Blake, D., and Cairns, A.J.G. (2010a) Facing up to uncertain life expectancy: The longevity fan charts. *Demography*, 47: 67-78.

Dowd, K., Cairns, A.J.G., Blake, D., Coughlan, G.D., and Khalaf-Allah, M. (2011a) A gravity model of mortality rates for two related populations. *North American Actuarial Journal*, 15: 334-356.

Dowd, K., Blake, D., and Cairns, A.J.G. (2011b) A computationally efficient algorithm for estimating the distribution of future annuity values under interest-rate and longevity risks. *North American Actuarial Journal*, 15: 237-247.

Dowd, K., Blake, D., Cairns, A.J.G., Coughlan, G.D. (2011c). Hedging pension risks with the age-period-cohort two-population gravity model. *Seventh International Longevity Risk and Capital Markets Solutions Conference*, Frankfurt, September 2011.

Jarner, S.F., and Kryger, E.M. (2011) Modelling adult mortality in small populations: The SAINT model. *ASTIN Bulletin*, 41: 377-418.

Li, S.H., and Chan, W.S. (2005) Outlier analysis and mortality forecasting: the United Kingdom and Scandinavian countries. *Scandinavian Actuarial Journal*, 2005(3): 187-211.

Li, J.S.-H., Hardy, M.R., and Tan, K.S. (2009) Uncertainty in model forecasting: An extension to the classic Lee-Carter approach. *ASTIN Bulletin*, 39: 137-164.

Li, J.S.-H., and Luo, A. (2012) Key q-duration: A framework for hedging longevity risk. *ASTIN Bulletin*, 42: 413-452.

Li, J.S.-H., and Hardy, M.R. (2011) Measuring basis risk in longevity hedges. *North American Actuarial Journal*, 15: 177-200.

Li, N., and Lee, R. (2005) Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography*, 42(3): 575-594.

Olivieri, A. and Pitacco, E. (2009) Stochastic mortality: The impact on target capital. *ASTIN Bulletin*, 39: 541-563.

Plat, R. (2009) Stochastic portfolio specific mortality and the quantification of mor-

tality basis risk. *Insurance: Mathematics and Economics*, 45: 123-132.

Plat, R. (2010) One-year value-at-risk for longevity and mortality. Pensions Institute Working Paper PI-1015.

Richards, S.J., and Currie, I.D. (2009) Longevity risk and annuity pricing with the Lee-Carter model. *British Actuarial Journal*, 15: 317-343.

Zhou, R., and Li, J.S.H. (2010) A cautionary note on pricing longevity index swaps. To appear in *Scandiavian Actuarial Journal*.

Zhou, R., Li, J.S.H., and Tan, K.S. (2011) Economic pricing of mortality-linked securities in the presence of population basis risk. *Geneva Papers on Risk and Insurance: Issues and Practice*, 36: 544-566.