Pensions Institute

DISCUSSION PAPER PI-1109

Key q-duration: A Framework for Hedging Longevity Risk

Johnny Siu-Hang Li, and Ancheng Luo

July 2011

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

Key q-duration: A Framework for Hedging Longevity Risk

Johnny Siu-Hang Li,* and Ancheng Luo

Abstract

When hedging longevity risk with standardized contracts, the hedger needs to calibrate the hedge carefully so that it can effectively reduce the risk. In this article, we present a calibration method that is based on matching mortality rate sensitivities. Specifically, we introduce a measure called key q-duration, which allows us to estimate the price sensitivity of a life-contingent liability to each portion of the underlying mortality curve. Given this measure, one can easily construct a longevity hedge with a handful of mortality forwards. Our empirical results indicate that using key q-durations, a hedge effectiveness of more than 97% can be attained with only five mortaltiy forwards. We also investigate other important issues that are related to standardized longevity hedges, including the adaptation needed for hedging multiple birth cohorts, and the quantification of sampling risk and basis risk.

Keywords: Cairns-Blake-Dowd model; Mortality forwards; Securitization.

1 Introduction

Pension plans and insurers selling life annuities are subject to longevity risk, the risk that individuals are living longer than expected. Although the risk may not pose as severe a short-term threat as large falls in asset values, it may possibly undermine the long-term sustainability of a portfolio, and thus requires careful management.

*Corresponding author. Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada. Email: shli@uwaterloo.ca.

In recent years, longevity risk has become a high profile risk, partly because of the current low yield environment, and partly because of changes in regulatory regimes. For instance, Solvency II, which is scheduled to come into effect in 2013, requires insurers operating in the European Union to hold longevity risk solvency capital that is based on either a prescribed stress test or an approved internal risk model.

Longevity risk is systematic, so there is a limit to how much longevity risk an entity can take, given its capital base and risk objectives. Pension plans and annuity writers may transfer their longevity risk exposure to capital markets. For example, a pension plan can take a long position of a contract which pays an amount that increases with its realized survival rate to offset the unexpected increase in its liability. Some investors including hedge funds are interested in acquiring an exposure to longevity risk for earning a risk premium, because the risk has no obvious correlation with typical market risk factors such as stock prices, interest rates and foreign exchange rates.

Mortality-linked contracts typically take either an indemnity (bespoke) or standardized form. Indemnity style contracts are based on the actual mortality experience of the hedger's own portfolio. An example is the survivor swap agreed between Babcock International and Credit Suisse in 2009. Under the terms of the contract, Babcock's pension plan will swap pre-agreed monthly payments to Credit Suisse in return for monthly payments dependent on the longevity of the plan's own members. Indemnity style contracts fully mitigate the hedger's longevity risk exposure, but they have limited liquidity. The counterparties may find it difficult to unwind the deal after it had been done.

Standardized contracts, by contrast, are based on the mortality calculated by reference to a national population index. An example is the 25-year longevity bond jointly announced by BNP Paribas and the European Investment Bank in 2004. This bond makes coupon payments that are proportional to the realized survival rates of English and Welsh males who were aged 65 in 2002. Standardized contracts have lower initial and ongoing data requirements. Also, because they are more transparent to investors, they provide quicker execution and, in theory, greater liquidity. However, in relying on standardized contracts, the hedger needs to calibrate the longevity hedge carefully. Specifically, given a certain mortality-linked instrument available in the market, the hedger needs to decide how many units of the instrument it should acquire so that the maximum hedge effectiveness can be attained.

Recently, there has been a wave of work on the pricing of standardized mortality-linked instruments. For example, Cairns et al. (2006) priced longevity bonds with a risk-adjusted mortality process; Li et al. (2011) valued mortality forwards by using a method called canonical valuation; Zhou and Li (2010) discussed various issues related to the pricing of longevity index swaps. However, relatively little research has been devoted to the calibration of longevity hedges.

In a continuous-time setting, Dahl et al. (2008) derived risk-minimizing strategies in markets where survivor swaps are available. Barbarin (2008) performed a similar derivation by considering longevity bonds instead. The results in both studies are built upon some assumed stochastic processes for the evolution of mortality, interest rates and net payments to individuals in the portfolio. For example, in the work of Dahl et al. (2008), it was assumed that the underlying mortality process is driven by a Cox-Ingersoll-Ross (CIR) model and that the net payments to all individuals in the portfolio are identical. These assumptions may be too stringent to fit the actual situation that a hedger is facing.

In this paper, we approach the hedging problem from a different angle by considering a portfolio's price sensitivity to the assumed mortality curve. This idea is first proposed as a concept by Coughlan et al. (2007) and then described in more detail by Coughlan (2009). However, in previous work, the measurement of mortality rate sensitivity has not been studied extensively, and in particular, it is unclear about important properties such as age dependence (see, e.g, Wills and Sherris (2008)) can be incorporated into the sensitivity measures. In this paper, we formalize the setup, and perform a deeper investigation of various issues that are related to the measurement of mortality rate sensitivity.

What we propose in this paper is largely analogous to Ho's (1992) method of key rate durations, which has been extensively used by traders to measure and manage interest rate risk. Key rate durations measure a portfolio's price sensitivity to different segments of an interest rate yield curve, and an interest rate hedge can be constructed using zero-coupon bonds with maturities equal to the key rate durations of the portfolio being hedged. We propose an analogous measure called key q-duration, which allows us to measure the sensitivity of a life-contingent liability to each portion of a mortality curve. With key q-durations, one can create a longevity hedge with mortality forwards, which have been available in the market since they were first launched by JP Morgan in 2007.

We name it key q-duration because, as we demonstrate empirically, the evolution of mortality is driven primarily by a few key mortality rates. This property enables us to create a longevity hedge with a small number of mortality forwards. By requiring fewer contracts, the method we propose can help the longevity market concentrate liquidity on a restricted number of instruments.

We emphasize that the validity of the hedge created by key q-durations does not depend on a specific stochastic mortality model.[1] Since no model (and hence simulation) is involved, the calibration of the hedge is quick and manageable even for a complicated real life pension plan. The ease of implementation requires little sacrifice of hedge effectiveness. Our illustrations indicate that a hedge created by key q-durations is almost equally effective as the corresponding hedge that is optimized by a computationally intensive method.

There are several problems that make longevity risk more difficult to hedge than interest rate risk. Because a typical pension plan contains members who were born in different years, when we construct a longevity hedge, a group of mortality curves have to be considered simultaneously. This means that, in practice, what we need to hedge is the uncertainty arising from the evolution of a two-dimensional surface, rather than just a one-dimensional curve. To overcome this challenge, we generalize key q-durations to a two-dimensional set-up, permitting us to hedge longevity risk associated with multiple birth cohorts using only a handful of mortality forward contracts.

Basis risk and sampling risk are two extra challenges to the hedging work. Basis risk arises from the differences between the mortality experience of the hedger's portfolio and the national population to which the standardized instrument is linked. It affects not only hedge effectiveness but also the calculation of key q-durations. Sampling risk, on the other hand, results from the random variations between the lifetimes of individuals in the same portfolio. Although sampling risk is diversifiable, it may still be significant for smaller pension plans. The issues of basis risk and sampling risk are acknowledged and are analyzed in this paper.

The remainder of this paper is organized as follows. In Section 2, we describe the

[1] In the absence of basis risk, the calculation of key q-durations does not require a mortality model. However, when basis risk is present, a model may be needed to estimate an adjustment factor for the difference between the mortality of the two populations involved. Further details are provided in Section 7.2.

sources of data. In Section 3, we define key q-durations. In Section 4, we explain how a longevity hedge can be constructed with key q-durations. In Section 5, we use a synthetic pension plan to illustrate the use of key q-durations. The performance of the hedge we construct is analyzed from both ex post and ex ante perspectives. In Section 6, we extend key q-durations to a two-dimensional set-up, which can be applied to portfolios involving multiple birth cohorts. In Section 7, we discuss the additional challenges, including basis risk and sampling risk. Finally, in Section 8, we conclude the paper with some suggestions for further research.

2 Sources of Data

Most of our illustrations are based on the LifeMetrics index, on which JP Morgan's q-forwards are written. In particular, we use the historic index data for the population of English and Welsh males. The index data are composed of graduated death probabilities at ages 20 to 90 for the years 1961 to 2009. They are are obtained from the LifeMetrics website (www.jpmorgan.com/lifemetrics).

We illustrate the impact of population basis risk with several populations, including Canadian males, French males, UK males and Scottish males. The required data for these populations are obtained from the Human Mortality Database (2011).

3 Key q-duration

3.1 Key mortality rates

The evolution of a mortality curve is not simple. In Figure 1 we illustrate the shifts of the mortality curve for English and Welsh males over different periods of time. It is clear from the illustration that, just like interest rate yield curves, the mortality curve is subject to non-parallel shifts over time. For this reason, when we measure a portfolio's price sensitivity to the underlying mortality curve, we need to use a vector of numbers rather than just a single number, which by itself is not able to take account of movements other than parallel shifts.

More specifically, we measure a portfolio's price sensitivity to shifts at selected key points on the mortality curves. We let n be the number of key points and let

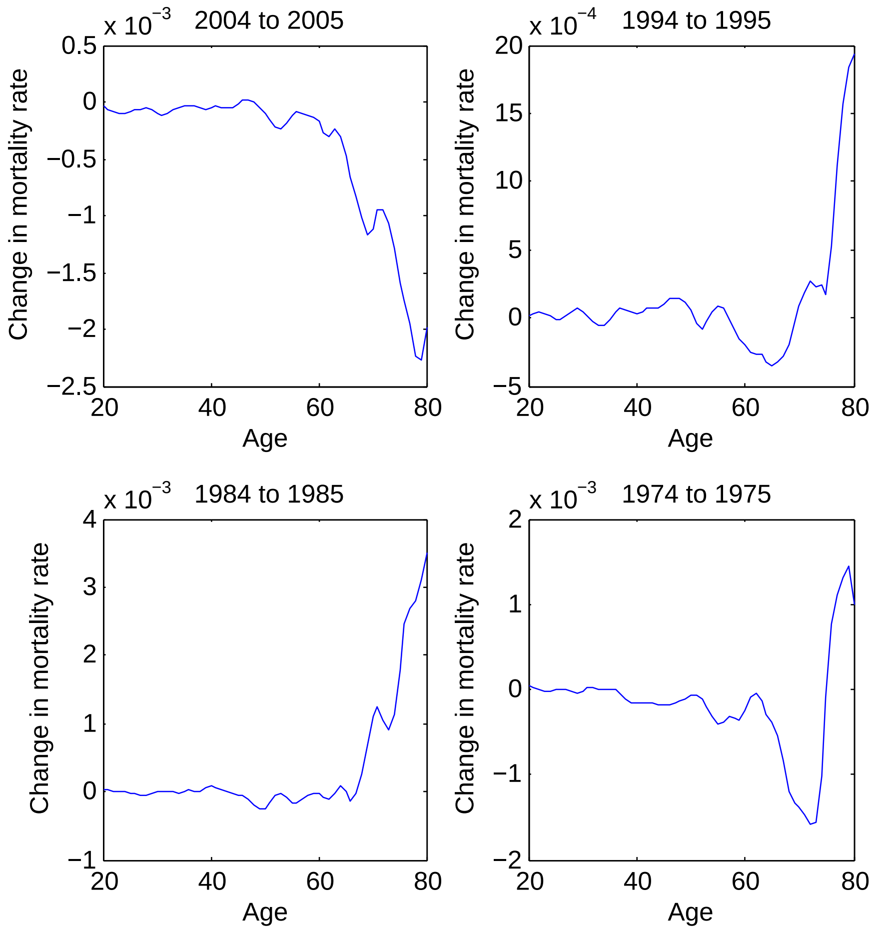

Figure 1: The shifts of the mortality curve for English and Welsh males over different periods of time.

ages $x_1, x_2, \ldots, x_n$ (from smallest to largest) be the key points on a mortality curve. The mortality rates at ages $x_1, x_2, \ldots, x_n$ are called the key mortality rates.

Is it necessary to treat every single point on a mortality curve as a key rate? To answer this question, let us perform a factor analysis of the historic death probabilities for English and Welsh males. In this analysis, we divide the entire age range, 20 to 90, into j consecutive age groups, X_i, $i = 1, \ldots, j$, of (approximately) equal size. For each $j = 1, 2, \ldots$, the following model is then fitted to the data:

$$\ln(q(x,t)) = \alpha_x + \sum_{i=1}^{j} \kappa_t^{(i)} I(x \in X_i) + \epsilon(x,t),$$

where $q(x,t)$ is the probability that an individual aged x at the beginning of year t

dies during year t, given that the individual has survived to age x; α_x is the average of $\ln(q(x,t))$ over the sample period (1961 to 2009); $\kappa_t^{(i)}$ is a time-varying stochastic factor for the ith age group; I is an indicator function; and $\epsilon(x,t)$ is the error term. We can easily fit the model by the method of least squares.

Following Li and Lee (2005), we measure the goodness-of-fit for each fitted model with an explanation ratio, ER, which can be expressed as

$$ER = 1 - \frac{\sum_{x,t} \epsilon(x,t)^2}{\sum_{x,t}(\ln(q(x,t)) - \alpha_x)^2}.$$

We can interpret ER as the proportion of variance in historic $q(x,t)$ explained. A higher value of ER thus indicates a better fit to data, and $ER = 1$ indicates a perfect fit.

The values of ER for $j = 1, 2, \ldots$ are shown graphically in Figure 2. We observe from the diagram that over 98% of the variability in historical mortality rates can be explained with only five factors, and that the benefit from introducing additional factors is very marginal. This factor analysis suggests that we may not need a large number of key mortality rates for an adequate representation. In particular, we may reasonably represent the evoluation of a mortality curve with say five key mortality rates, each of which corresponds to a boarder age group. This property can simplify the measurement of sensitivity, and more importantly, it can reduce the number of instruments needed for a meaningful longevity hedge.

3.2 Modeling key rate shifts

The next step is to model the impact of a shift in a key mortality rate on other rates on the mortality curve. This step is necessary because of the property of age dependence, which we now demonstrate using data from English and Welsh males population.[2]

Let

$$R(x,t) = 1 - \frac{q(x,t+1)}{q_(x,t)}$$

[2]The data from LifeMetrics have been graduated by cubic splines (see Coughlan et al., 2007). The graduation removes sampling fluctuations, which may conceal the statistical properties of the mortality rates.

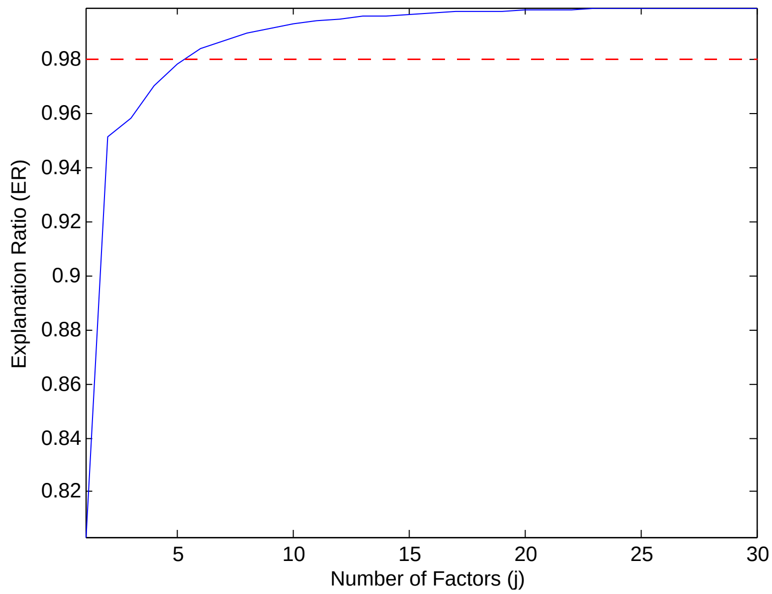

Figure 2: Values of ER for different values of j.

be the mortality reduction factor at age x and in year t, and let $\bar{R}(x)$ be the mean of $R(x,t)$ over the sample period (1961 to 2009). We calculate the sample correlation between the reduction factors at ages x and y using the following formula:

$$\rho(x,y) = \frac{\sum_t \left(R(x,t) - \bar{R}(x)\right)\left(R(y,t) - \bar{R}(y)\right)}{\sqrt{\sum_t \left(R(x,t) - \bar{R}(x)\right)^2}\sqrt{\sum_t \left(R(y,t) - \bar{R}(y)\right)^2}}.$$

In Figure 3 we display the calculated values of $\rho(x,y)$. For example, the right panel shows the values of $\rho(x,80)$ for different values of x. The analysis indicates that a change in a mortality rate is significantly correlated with changes in mortality rates at neighboring ages. The strength of the dependence, however, diminishes as the age gap becomes wider.

To capture age dependence, we assume that a change in the jth key mortality rate is accompanied with changes in mortality rates at ages that are close enough to x_j. Let $s(x,j,\delta(j))$ be the shift at age x associated with a change $\delta(j)$ in the jth key mortality rate. The function s, which may be viewed as an analog to a kernel in density estimation, can take different forms. We suggest the following specification, which uses a linear interpolation to approximate the diminishing dependence between two mortality rates as the age gap becomes wider.

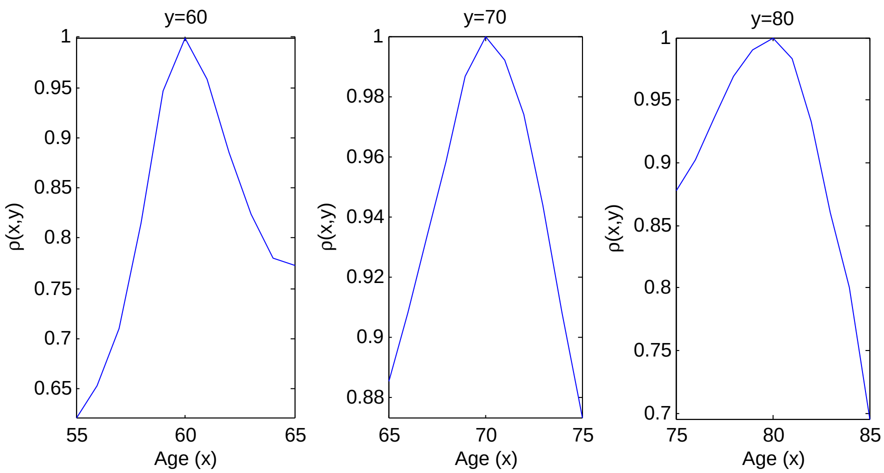

Figure 3: Values of $\rho(x, y)$ for $y = 60, 70, 80$.

For $2 < j < n - 1$,

$$s(x, j, \delta(j)) = \begin{cases} 0 & x \leq x_{j-1} \\ \frac{\delta(j)(x - x_{j-1})}{x_j - x_{j-1}} & x_{j-1} < x \leq x_j \\ \frac{\delta(j)(x_{j+1} - x)}{x_{j+1} - x_j} & x_j < x < x_{j+1} \\ 0 & x \geq x_{j+1} \end{cases}.$$

Note that the impact of $\delta(j)$ is fully absorbed by the time when x reaches the next or the previous key age (i.e., x_{j-1} or x_{j+1}). The specification of s is slightly different for $j = 1$ and $j = n$. For $j = 1$, we change s to $\delta(1)$ when $x \leq x_1$, and for $j = n$, we change s to $\delta(n)$ when $x > x_n$. A similar specification is also used by Ho (1992) to model shifts in an interest rate yield curve.

In effect, the shift in the whole mortality curve is approximated by the sum of $s(x, j, \delta(j))$, $j = 1, 2, \ldots, n$. In Figure 4 we display the overall shift resulting from arbitrary shifts in four key mortality rates, which are located at ages $x_1 = 65$, $x_2 = 70$, $x_3 = 75$, and $x_4 = 80$. The diagrams illustrate how parallel and non-parallel shifts can be approximated by the above specification.

3.3 Calculating key q-durations

A portfolio's price sensitivity to a shift in a key mortality rate is referred to as a key q-duration.

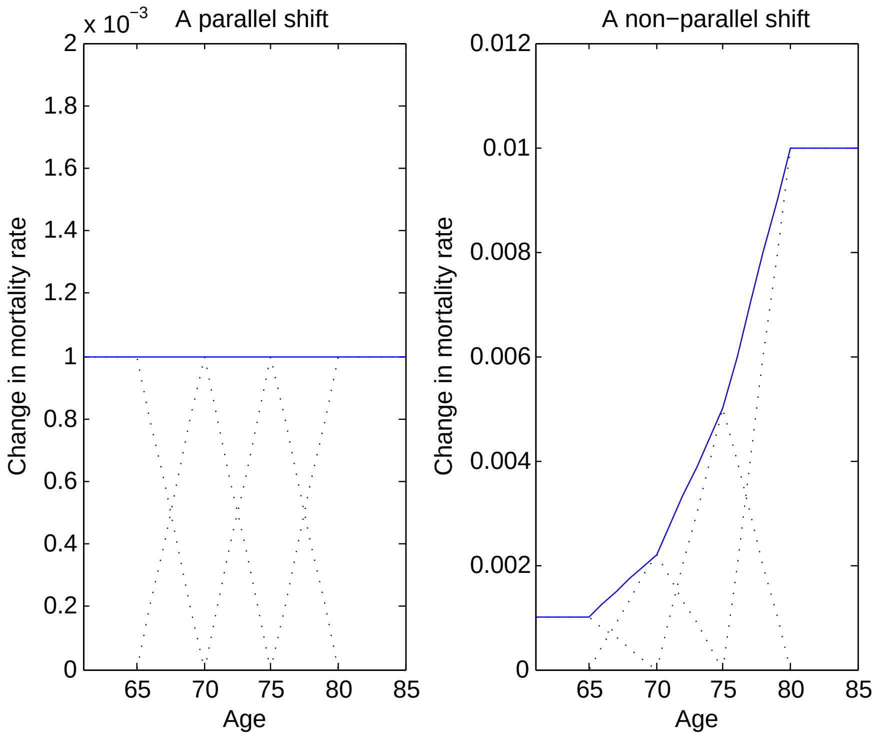

Figure 4: Approximation of the shift in a mortality curve as a sum of $s(x, j, \delta(j))$, $j = 1, 2, 3, 4$. The values shown are arbitrary.

Let $\mathbf{q}$ be the original mortality curve, and let $\tilde{\mathbf{q}}(j)$ by the mortality curve affected by $\delta(j)$, according to the way in which $s(x, j, \delta(j))$ is defined. The jth key q-duration of a portfolio is given by

$$KQD(P(\mathbf{q}), j) = \frac{P(\tilde{\mathbf{q}}(j)) - P(\mathbf{q})}{\delta(j)}, \tag{1}$$

where $P(\mathbf{q})$ is the value of the portfolio on the basis of the mortality curve $\mathbf{q}$. The number $KQD(P(\mathbf{q}), j)$ measures the portfolio's price sensitivity to the jth key rate, and the vector $\{KQD(P(\mathbf{q}), j); j = 1, 2, \ldots, n\}$ as a whole measures the portfolio's price sensitivity to the entire mortality curve.

In most cases, it is difficult to analytically calculate a key q-duration. For practical purposes, we may estimate $KQD(P(\mathbf{q}), j)$ as follows:

1. take $\mathbf{q}$ as the best estimate of the underlying mortality curve;

2. assuming $\delta(j)$ is 10 basis points, calculate $\tilde{\mathbf{q}}$;

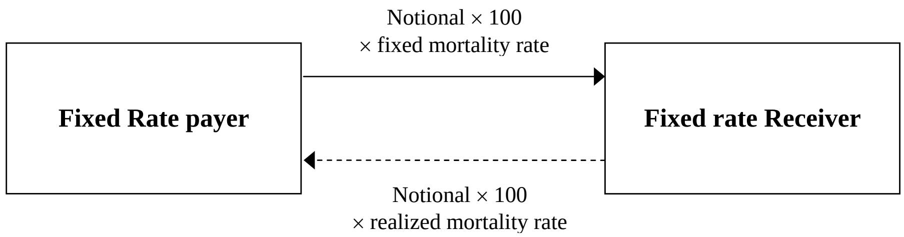

Figure 5: Settlement of a mortality forward at maturity.

3. compute $KQD(P(\mathbf{q}), j)$ with equation (1).

4 Building a Longevity Hedge

4.1 Hedging Instruments

Given the key q-durations of a portfolio, it is easy to construct a longevity hedge for the portfolio with a series of mortality forwards. A mortality forward is a zero-coupon swap that exchanges on the maturity date a fixed amount, determined at time 0, for a random amount that is proportional to an age-specific mortality rate for a certain population (the reference population) in some future time (the reference year). Because there is a lag in the availability of the index data, the reference year may be slightly earlier than the maturity date.

The fixed payment is proportional to the so-called forward mortality rate for the reference population. This rate is chosen so that no payment exchanges hands at inception. However, at maturity, a net payment will be made by one counterparty or the other. The settlement that takes place at maturity is illustrated diagrammatically in Figure 5. The fixed receiver, for example, receives at maturity 100 times the notional amount times the difference between the fixed and realized mortality rates.

An entity wishing to hedge longevity risk could enter into a portfolio of mortality forwards in which it receives fixed mortality rates and pays realized mortality rates. At maturity, the mortality forwards will pay out to the hedger an amount that increases as mortality rates fall to offset the unexpected increase in the hedger's liability. Provided that the weight on each mortality forward is calibrated properly, the resulting longevity hedge can stabilize the hedger's liability with respect to changes

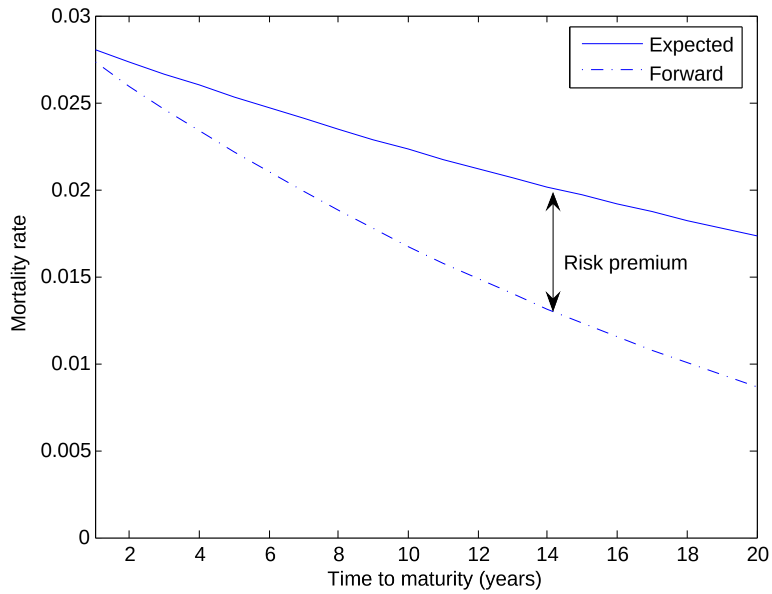

Figure 6: Illustrative expected and forward mortality curves.

in mortality rates.

The fixed payer can be an investor wishing to take longevity risk for a risk premium. To attract investors (fixed rate payers), the forward mortality rate must be smaller than the corresponding expected mortality rate, so that on average (i.e., if mortality is realized as expected), the investor will be paid. The difference between expected and forward rates, therefore, indicates the expected risk premium to the investor. In Figure 6 we illustrate the relationship between forward and expected mortality rates at a certain age for different times to maturity. A widening divergence is expected, because investors demand a higher risk premium from a longer term contract, which involves a prediction further into the future.

4.2 Determining the weight on each mortality forward

To make the hedge effective, the portfolio being hedged and the portfolio of mortality forwards must have similar price sensitivities to the underlying mortality curve. The sensitivities for both portfolios can be measured by key q-durations.

The key q-durations for a mortality forward can be computed analytically. Consider a mortality forward that is linked to jth key mortality rate. Assuming a notional amount of \$1, the present value of the mortality forward (from a fixed receiver's viewpoint) can be written as

$$F_j(\mathbf{q}) = 100(1+r)^{-T_j}(q^f(x_j,t_j) - q(x_j,t_j)),$$

where T_j is the time to maturity (in years), r is the interest rate at which cash flows are discounted, $q^f(x_j,t_j)$ is the forward mortality rate for reference age x_j and reference year t_j. The difference $t_j - x_j$ is the year of birth for the cohort of individuals in question. It is obvious that

$$KQD(F_j(\mathbf{q}), j) = -100(1+r)^{-T_j}. \tag{2}$$

Furthermore, it is noteworthy that according to the way we specify the key rate shifts, $F_j(\mathbf{q})$ is unaffected by changes in other key rates, which means $KQD(F_j(\mathbf{q}), i)$ must be zero for any $i \neq j$. This important property allows us to determine the appropriate weight on each mortality forward independently. It also makes mortality forwards an analog to zero-coupon bonds in Ho's (1992) framework for hedging interest rate risk.[3]

Let V be the time-0 value of the portfolio being hedged. We can construct a hedge with mortality forwards that are linked to the n key mortality rates. Specifically, to make V and the portfolio of mortality forwards have similar sensitivities to $\mathbf{q}$, we need a notional amount of

$$w(j) = \frac{KQD(V(\mathbf{q}), j)}{KQD(F_j(\mathbf{q}), j)}$$

on the mortality forward linked to the jth key mortality rate.

5 An Illustration

5.1 The hedge

We now illustrate the hedging framework with a simple synthetic pension plan, which pays a pensioner \$1 at the beginning of each year until the pensioner dies or reaches age 91, whichever is earlier. In the illustration, the following assumptions are made.

[3]The price of a zero-coupon bond is sensitive only to a shift in the corresponding key interest rate but not other key rates on the yield curve. This property makes it is straightforward to construct an interest rate hedge with zero-coupon bonds, when key rate durations of the portfolio being hedged are known.

1. The pensioner is exactly 60 years old at time 0.[4]
2. The pensioner's mortality experience is exactly the same as that of English and Welsh males with the same year of birth. This assumption will be relaxed in Section 5 when basis risk and sampling risk are studied.
3. The best estimate of the underlying mortality curve is based on the central projection made by the Cairns-Blake-Dowd model (Cairns et al., 2006), fitted to the data from the population of English and Welsh males.
4. The interest rate at which cash flows are discounted is $r = 3\%$.
5. Mortality forwards linked to ages 65, 70, 75, 80 and 85 for the same birth cohort are available.
6. For each mortality forward, the maturity date (the date when the payment is settled) is one year after the reference year.
7. The forward mortality rates are the same as the corresponding best estimate mortality rates. This assumption, which implies zero risk premium, would affect the cost but not the performance of the longevity hedge.

In our baseline calculations, we make use of all five available mortality forwards. We set the key mortality rates to mortality rates at ages 65, 70, 75, 80, and 85. The key q-durations for the mortality forwards are calculated analytically using equation (2), while those for the pension liability are calculated numerically using the procedure described in Section 3.3. The calculated key q-durations are displayed in Table 1. Also shown in the table are the required notional amounts of the mortality forwards in the replicating portfolio.

5.2 Hedge effectiveness

We can evaluate hedge effectiveness by examining the variability in the present value of unexpected cash flows. If a pension plan is unhedged, the present value of the unexpected cash flows from the plan can be expressed as

$$X = V(\mathbf{q}) - V(\mathrm{E}(\mathbf{q})),$$

[4] We set time 0 to the beginning of year 2010, because 2009 is the latest year for which English and Welsh male mortality data are available.

	$j = 1$	$j = 2$	$j = 3$	$j = 4$	$j = 5$
x_j	65	70	75	80	85
$KQD(F_j(\mathbf{q}), j)$	-83.7484	-72.2421	-62.3167	-53.7549	-47.7606
$KQD(V(\mathbf{q}), j)$	-99.9761	-38.4377	-24.0862	-12.4439	-7.5237
$w(j)$	1.1938	0.5321	0.3865	0.2315	0.1575

Table 1: Key q-durations and the required notional amounts.

where V is the time-0 value of the pension liability, and $\mathrm{E}(\mathbf{q})$ is the best estimate of $\mathbf{q}$. By contrast, if there is a longevity hedge for the pension plan, then the present value of the unexpected cash flows from the plan can be written as

$$X^* = V(\mathbf{q}) - V(\mathrm{E}(\mathbf{q})) - H(\mathbf{q}) + H(\mathrm{E}(\mathbf{q})),$$

where H is the present value of all payoffs from the instruments in the hedge portfolio. Of course, both X and X^* are random quantities as they depend on a random vector $\mathbf{q}$. The longevity hedge is effective if X^* is significantly less variable than X. As such, we can measure hedge effectiveness in terms of the amount of risk reduction, R, which is defined by

$$R = 1 - \frac{\sigma^2(X^*)}{\sigma^2(X)},$$

where $\sigma^2(X)$ and $\sigma^2(X^*)$ are the variances of X and X^*, respectively. A higher value of R indicates a better hedge effectiveness.

We simulate 5,000 realizations of $\mathbf{q}$ on the basis of the Cairns-Blake-Dowd model (Cairns et al., 2006) fitted to data from English and Welsh male population. For each realization of $\mathbf{q}$, we calculate the values of X and X^*. This create empirical distributions of X and X^*, from which the amount of risk reduction can be calculated.

The portfolio of five mortality forwards gives a risk reduction of $R = 97.2\%$. The reduction in risk can be visualized from Figure 7, in which the empirical distributions of X and X^* are displayed.

To examine how the amount of risk reduction is related to the number of mortality forwards used, we repeat the analysis by using four mortality forwards (linked to ages 65, 70, 75 and 80) and three mortality forwards (linked to ages 65, 70, 75). The resulting empirical distributions of X^* are shown in Figure 7. The hedge effectiveness

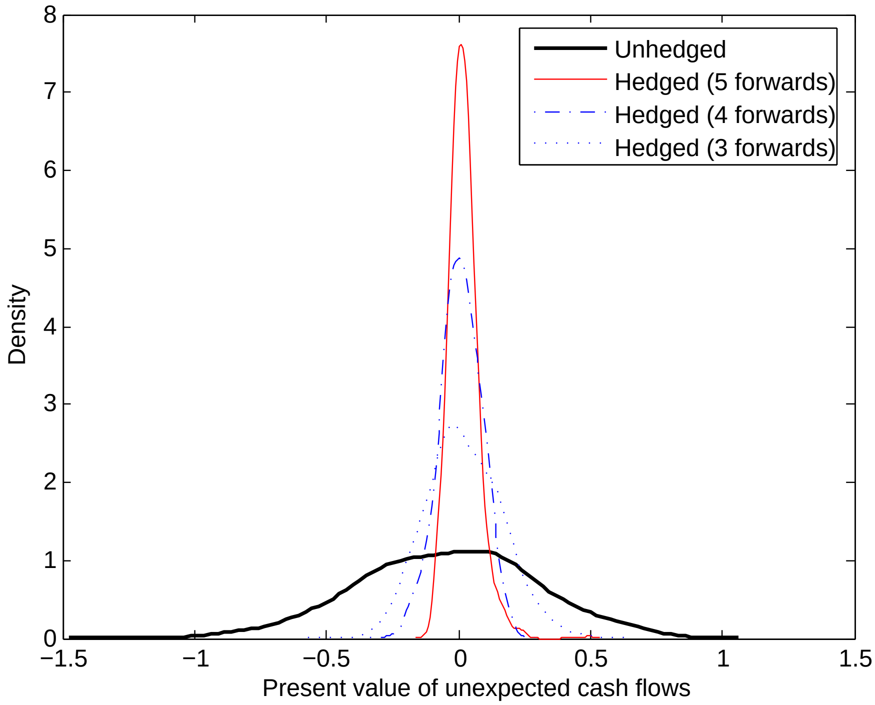

Figure 7: Simulated distributions of X and X^*, based on 3, 4 and 5 mortality forwards.

is reduced to 94.2% and 77.5% when four and three mortality forwards are used, respectively.

5.3 Impact on solvency capital requirement

Solvency II represents a new set of regulatory requirements for insurance firms that operate in the European Union. It is scheduled to come into effect in 2013, replacing Solvency I, which has been in place since the early 1970s.

Under Solvency II, longevity risk is split into two components. The allowance for expected future improvements in mortality is included as Technical Provisions (TP), while the uncertainty concerning the actual future evolution of mortality is covered by the Solvency Capital Requirement (SCR). In principal, SCR is determined as the value-at-risk over a one-year horizon with a 99.5% confidence level.

Insurance companies may compute longevity SCR with an approved internal risk model or, more conveniently, by a standard formula prescribed by the Committee of European Insurance and Occupational Pensions Supervisors (CEIOPS). The setup

	$s = 0.2$	$s = 0.25$
SCR	0.7076	0.8958
SCR^* [key ages: 65, 70, 75, 80, 85]	0.0021	0.0140
SCR^* [key ages: 65, 70, 75, 80]	0.0311	0.0501
SCR^* [key ages: 65, 70, 75]	0.0946	0.1296

Table 2: Longevity SCR at stress levels 20% and 25%.

and calibration of the standard formula is currently being established by a series of Quantitative Impact Studies (QIS). In QIS4, the proposal for the standard formula for longevity SCR is to stress-test using an immediate and permanent 25% fall in mortality rates (CEIOPS, 2008). In QIS5, however, it was decided to reduce the stress to 20%, due possibly to the criticism from the insurance industry that a flat shock of 25% is too stringent (CEIOPS, 2010). We refer interested readers to Börger (2010) for a recent study on longevity SCR.

Let s be the flat shock prescribed by CEIOPS. In the absence of a longevity hedge, the longevity SCR can be expressed as

$$SCR = V((1-s)\mathrm{E}(\mathbf{q})) - V(\mathrm{E}(\mathbf{q})).$$

If a longevity hedge is in place, the longevity SCR is given by

$$SCR^* = V((1-s)\mathrm{E}(\mathbf{q})) - V(\mathrm{E}(\mathbf{q})) - H((1-s)\mathrm{E}(\mathbf{q})) + H(\mathrm{E}(\mathbf{q})).$$

We now study how the proposed hedging strategy may reduce longevity SCR. In Table 2 we display the longevity SCR for the synthetic pension plan, with and without a longevity hedge. At the stress levels of 20% and 25%, the hedge with five mortality forwards can reduce the SCR for the pension plan by 99% and 98%, respectively. The reductions in SCR are smaller if fewer mortality forwards are used, but they are still significant.

5.4 Proximity to optimality

The longevity hedge based on key q-durations is effective, but may not be the most effective, as key q-durations are only approximate measurements of the sensitivity to $\mathbf{q}$. Here we study if a more effective hedge can be formed with the same mortality forwards, and if it can be formed, how much more effective it would be.

Suppose that our sole objective is to maximize the amount of risk reduction, R, produced by the hedge. To achieve this objective, we can gradually adjust the weight on each of the five mortality forwards until the maximum value of R is attained. The value of R in each iteration can be estimated by simulating from a stochastic mortality model. In our illustration, the Cairns-Blake-Dowd model fitted to English and Welsh male data is used.

We found that, on the basis of the five mortality forwards we consider, the optimal amount of risk reduction is 98.5%, which is 1.3% higher that achieved by using key q-duration. The result indicates that key q-durations can yield an amount of risk reduction that is not too far from the maximum attainable value.

Note that the optimization procedure, which involves simulations and a maximization of a multivariable function, demands a lot of computational resources. It works for the synthetic pension plan we consider, but may not be feasible for a sophisticated pension plan in real life. Key q-durations, in contrast, are much easier to calculate and are not dependent on a specific simulation model.

5.5 Choosing key mortality rates

Our hedging strategy requires mortality forwards that are linked to the key mortality rates chosen. In the early stages of the market's development, it is expected that transactions are restricted to a limited number of contracts in which liquidity can be concentrated. Therefore, the choice of key mortality rates is exogenous, depending on the instruments that are available in the market. In our baseline calculations, we assumed that mortality forwards linked to ages 65, 70, 75, 80 and 85 for the cohort in question are available, and chose the key mortality rates accordingly.

Now suppose that mortality forwards linked to all rates on $\mathbf{q}$ are available. In this case, which five key mortality rates (mortality forwards) should we choose? We answer this question by comparing the hedge effectiveness provided by all combinations of key mortality rates. Since the synthetic pension plan involves 31 mortality rates (from ages 60 to 90), there are altogether 169 911 possible combinations.

We found that the maximum amount of risk reduction is attained if the key mortality rates are located at ages 62, 67, 73, 79 and 85. These ages are roughly evenly spaced, and are close to the key mortality rates used in the baseline calculation. This combination of key mortality rates yields a hedge effectiveness of 98.3%, which is

1.1% higher than that achieved in the baseline calculation. Our findings suggest that, to concentrate liquidity and to attract demand from pension plans, it makes sense for investment banks to issue mortality forwards that are linked to rates representing different segments of a mortality curve.

5.6 Comparing with an alternative strategy

In this subsection, we compare our hedging strategy with that proposed by Cairns et al. (2008). As a shorthand, we call their strategy the CBD (Cairns-Blake-Dowd) strategy.

Let $S(x,t)$ be the probability that a person currently aged x will survive t years more. The CBD strategy is based on the following first order approximation of $S(x,t)$:

$$\begin{aligned} S(x,t) &= (1-q(x,t_0))\,(1-q(x+1,t_0+1))\ldots(1-q(x+t-1,t_0+t-1)) \\ &= \prod_{i=0}^{t-1}\big(1-q^f(x+i,t_0+i)-\Delta(x+i,t_0+i)\big) \\ &\approx \prod_{i=0}^{t-1}\big(1-q^f(x+i,t_0+i)\big) \\ &\quad -\sum_{i=0}^{t-1}\Delta(x+i,t_0+i)\prod_{j=0,j\neq i}^{t-1}\big(1-q^f(x+j,t_0+j)\big)\,, \end{aligned}$$

where $\Delta(x+i,t_0+i) = q(x+i,t_0+i) - q^f(x+i,t_0+i)$.

The first term in the approximation is non-random, while the the second term is just a linear combination of the random mortality rates. Hence, the coefficients in the second term tell us the amounts of mortality forwards needed to hedge a liability that depends on $S(x,t)$. In particular, to hedge a liability that pays a random amount of $S(x,t)$ at time t_0+t, the hedger needs to be the fixed receiver of t mortality forwards, and the required notional amount of the ith, $i = 1,2,\ldots,t$, mortality forward is given by

$$\frac{1}{100(1+r)^{t-T_i}}\prod_{j=0,j\neq i}^{t-1}\big(1-q^f(x+j,t_0+j)\big)\,.$$

The hedge for the synthetic pension plan, which has a present value of

$$\sum_{t=0}^{30}\frac{S(60,t)}{(1+r)^t},$$

can be derived accordingly.

Similar to our strategy, the CBD strategy has the advantage of being model free. Nevertheless, it completely ignores the property of age dependence, and as we now demonstrate, the ignorance of age dependence can be a serious problem.

Recall that the synthetic pension plan involves 31 random mortality rates. We first implement both our strategy and the CBD strategy with 31 contracts, each of which is linked to a different point on the underlying mortality curve. Using 31 contracts, both strategies yield an almost perfect hedge effectiveness of 99.9%, and in fact, the required notional amounts calculated from both methods are identical. The identical results are not surprising, because in the extreme case when all rates in the mortality curve are key rates, our model for key rate shifts would imply no age dependence, which makes both methods the same in principle.

However, when we use less contracts, the difference between the two strategies would become apparent. Now we consider the five mortality forwards used in our baseline calculations. Using these five contracts, our strategy produces a hedge effectiveness of 97.2%, but the CBD strategy gives only 35.0%. The CBD strategy does not perform well in this case, because, as it ignores age dependence, it underestimates the price sensitivity of the pension liability to the five mortality rates, and hence understates the required notional amounts.

This comparison points to the conclusion that it is crucially important to take age dependence into account when measuring mortality rate sensitivities. To ameliorate the problem of the CBD strategy, Cairns et al. (2008) suggest that we may consider average (or 'bucketed') mortality rates instead. However, as they point out, using average mortality introduces some basis risk depending on the specific age-structure of the population being hedged. It is also unclear about how mortality rates should be 'bucketed,' particularly when the pension plan involves more than one birth cohorts.

6 Hedging Multiple Cohorts

In Section 3, we presented key q-durations in a one-dimensional setting. This simple setting is easy to understand, and is adequate for hedging the longevity risk arising from one birth cohort. However, in practice, pension plans involve multiple birth cohorts, and therefore when we construct a longevity hedge, we may need to consider a group of mortality curves simultaneously. In this section, we generalize key

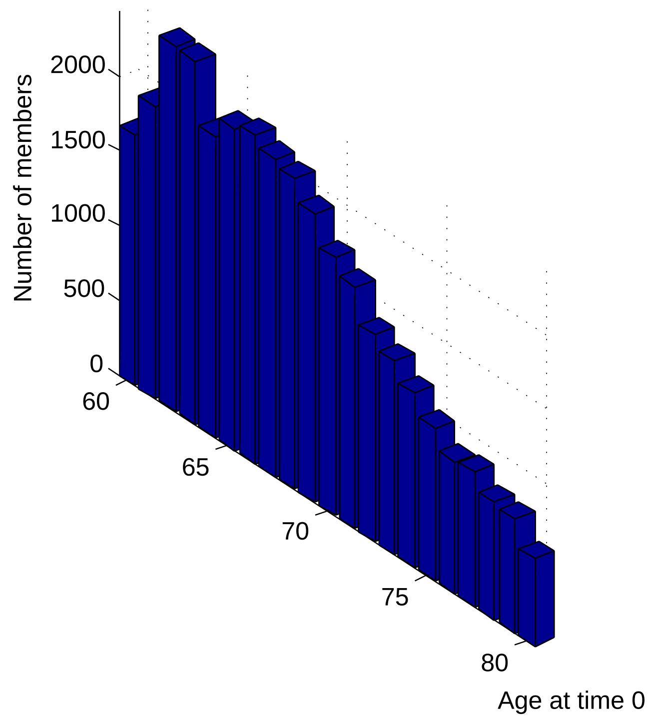

Figure 8: The initial demographic structure of the illustrative multi-cohort pension plan.

q-durations to a two-dimensional setting, which can then be used to hedge the uncertainty arising from the evolution of a 'surface' that is composed of more than one mortality curves.

We illustrate the two-dimensional generalization using a synthetic pension plan with members ranging from age 60 to 80 at time 0 (the beginning of year 2010). The bar chart in Figure 8 summarizes the initial demographic structure of the plan. We assume again that the plan pays each pensioner \$1 at the beginning of each year until the pensioner dies or reaches age 91, whichever is earlier. For simplicity, we assume further that the plan is closed, that is, there is no new entrants to the plan after time 0.

The plan involves 21 birth cohorts, with the youngest born in 1950 and the oldest in 1930. In total, it is subject to the uncertainty associated with 441 future death rates. Assuming that the maximum allowable separation between two key mortality rates is five years of age, our strategy in a one-dimensional setting would require 76 contracts. It is unlikely that the longevity market can provide such a wide variety of standardized contracts. Even if the required contracts are available, the complexity of the hedge would make it costly and difficult to manage in the future.

To mitigate this problem, we attempt to consider the potential dependence along the year of birth dimension. Let us consider the group of individuals who were born in year c. We let

$$R^*(x,c) = 1 - \frac{q(x, c+x+1)}{q_(x, c+x)}$$

be their mortality reduction factor at age x, and let $\bar{R}^*(c)$ be the mean of $R^*(x,c)$ over the sample age range for that cohort. To examine the dependence across cohorts, we calculate the sample correlation between the reduction factors for year of births c and d with the following formula:

$$\rho^*(c,d) = \frac{\sum_x \left(R^*(x,c) - \bar{R}^*(c)\right)\left(R^*(x,d) - \bar{R}^*(d)\right)}{\sqrt{\sum_x \left(R^*(x,c) - \bar{R}^*(c)\right)^2}\sqrt{\sum_x \left(R^*(x,d) - \bar{R}^*(d)\right)^2}},$$

where the summations are taken over the common sample age range for both cohorts.

The calculated values of $\rho^*(c,d)$ for $d = 1910, 1930, 1950$ and for $c = d-5, d-4, \ldots, d+5$ are displayed in Figure 9.[5] The patterns of $\rho^*(c,d)$ indicates that mortality improvement rates of neighboring cohorts are significantly correlated with one another, but the strength of the correlation tapers off as the birth cohorts are wider apart. The analysis of $\rho^*(c,d)$ motivates us to generalize the way in which a key rate shift is modeled. Specifically, in what follows, we assume that a shift in a key mortality rate for a birth cohort would affect mortality rates for not only the same but also the neighboring cohorts.

In the two-dimensional setting, a mortality surface is represented by the key mortality rates for m key cohorts, who were born in years $c_1, \ldots, c_m$, respectively. The mortality curve for key cohort k, $k = 1, \ldots, m$, contains n_k key mortality rates, which are located at ages $x_{1,k}, x_{2,k}, \ldots, x_{n_k,k}$ (from smallest to largest). We identify the jth key mortality rate on the kth key cohort as the (j,k)th key mortality rate.

To hedge the synthetic pension plan, we consider $k = 4$ key cohorts, with $c_1 = 1947$, $c_2 = 1943$, $c_3 = 1939$ and $c_4 = 1935$. Each key cohort contains key mortality

[5] Before calculating $\rho^*(c,d)$, the data have been graduated by cubic splines to remove sampling fluctuations, which may conceal the statistical properties of the mortality rates.

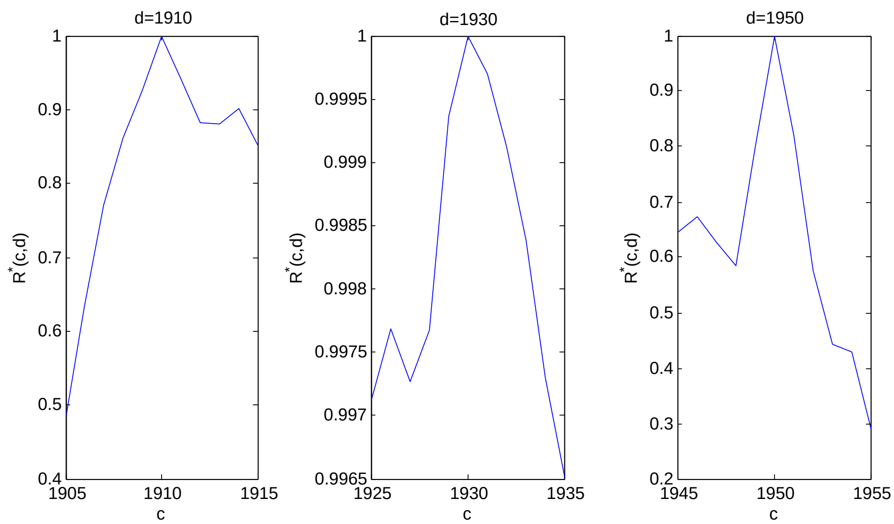

Figure 9: Values of $\rho^*(c,d)$ for $d = 1910, 1930, 1950$.

k	n_k	Locations of key mortality rates
1	5	Ages 65, 70, 75, 80, 85
2	4	Ages 70, 75, 80, 85
3	3	Ages 75, 80, 85
4	2	Ages 80, 85

Table 3: Locations of the key mortality rates used in hedging the multi-cohort synthetic pension plan.

rates that are no more than 5 years of age apart. The strategy we use can be visualized from the lexis diagram in Figure 10. The locations of the 14 key mortality rates used are shown in Table 3. We assume that mortality forwards linked to the key mortality rates are available in the market.

Let $\delta(j,k)$ be the shift in the (j,k)th key mortality rate. Our goal is to model $s(x,c,(j,k),\delta(j,k))$, the impact of $\delta(j,k)$ on the mortality rate for age x and for year of birth c. For $j = 2,\ldots,n_k - 1$ and for $k = 2,\ldots,m-1$, we set

$$s(x,c,(j,k),\delta(j,k)) = \delta(j,k)\,\alpha(x,j,k)\,\beta(c,j,k),$$

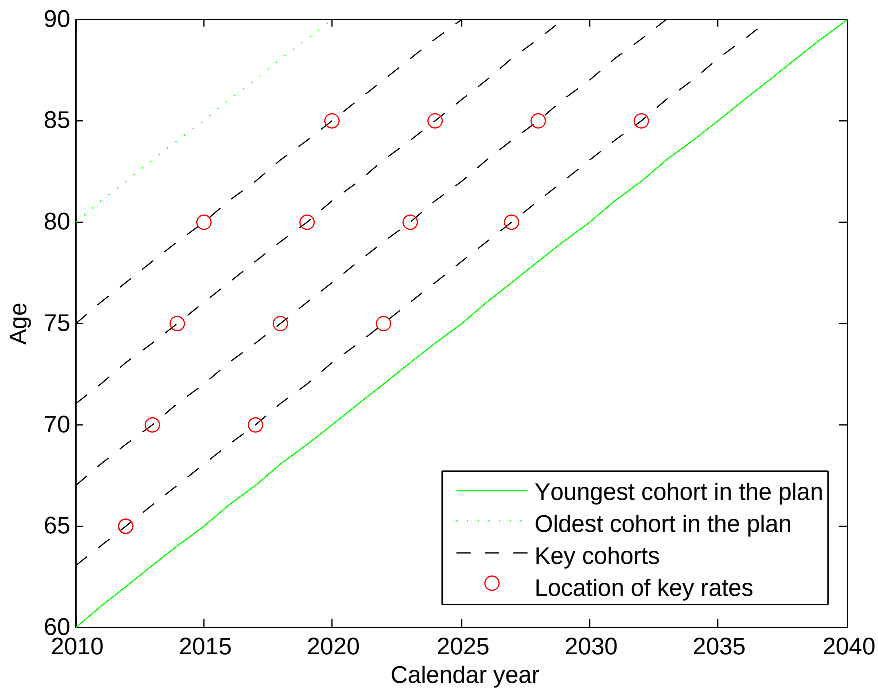

Figure 10: The key cohorts used in the hedging strategy for the multi-cohort synthetic pension plan.

where

$$\alpha(x,j,k) = \begin{cases} 0 & x \le x_{j-1,k} \\ \frac{x - x_{j-1,k}}{x_{j,k} - x_{j-1,k}} & x_{j-1,k} < x \le x_{j,k} \\ \frac{x_{j+1,k} - x}{x_{j+1,k} - x_{j,k}} & x_{j,k} < x < x_{j+1,k} \\ 0 & x \ge x_{j+1,k} \end{cases},$$

and

$$\beta(c,j,k) = \begin{cases} 0 & c \le c_{j,k-1} \\ \frac{c - c_{j,k-1}}{c_{j,k} - x_{j,k-1}} & c_{j,k-1} < c \le c_{k,j} \\ \frac{c_{k+1,j} - c}{c_{k+1,j} - c_{k,j}} & c_{k,j} < c < c_{k+1,j} \\ 0 & c \ge c_{k+1,j} \end{cases}.$$

Similar to the one-dimensional setting, the effect of $\delta(j,k)$ diminishes as x is farther way from x_j, and is fully absorbed when x_{j-1} or x_{j+1} is reached. The dependence along the year of birth dimension is also modeled in a similar manner. In particular, the impact of $\delta(j,k)$ reduces linearly with the distance between c and c_k, and is reduced

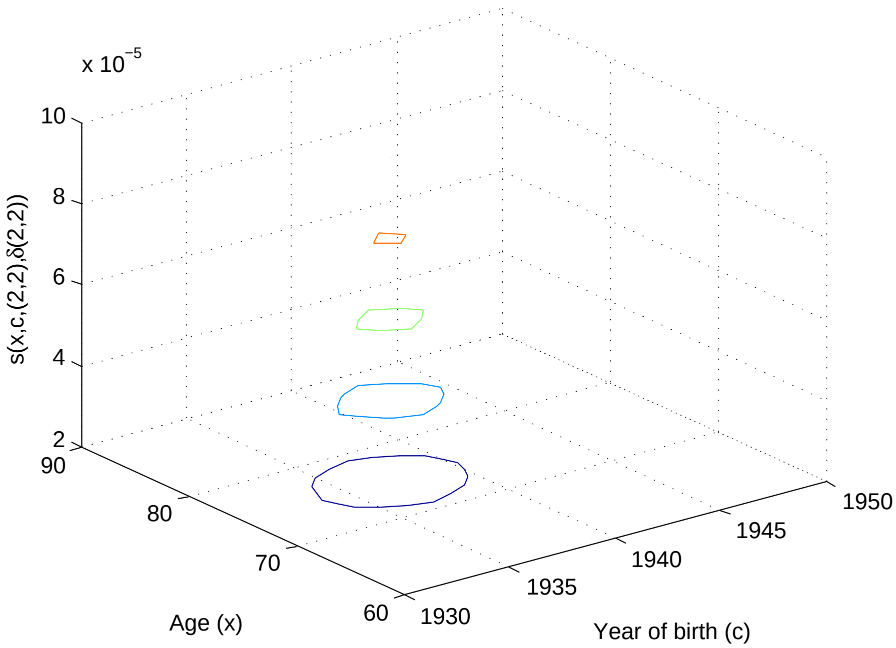

Figure 11: The impact of $\delta(2,2) = 0.0001$ on other mortality rates that are involved in the synthetic multi-cohort pension plan.

to zero when c_{k-1} or c_k is reached. As an example, in Figure 11 we demonstrate how we model the impact of $\delta(2,2) = 0.0001$ on other mortality rates that are involved in the synthetic multi-cohort pension plan.

The specification of α is slightly different for the first and last key mortality rates in a key cohort. For $j = 1$, we set α to 1 when $x \leq x_{1,k}$, and for $j = n_k$, we set α to 1 when $x > x_{n_k,k}$. Also, for $k = 1$, we set β to 1 when $c \leq c_{j,1}$, and for $k = m$, we set β to 1 when $c > c_{j,m}$. It is easy to see that the way we specify s permits a mortality surface to shift in both parallel and non-parallel fashions.[6]

Given $s(x, c, (j, k), \delta(j, k))$, it is straightforward to calculate of the key q-duration associated with the (j, k)th key mortality rate. Let $\mathbf{Q}$ and $\tilde{\mathbf{Q}}(j, k)$ be original mortality surface and the mortality surface affected by $\delta(j, k)$, respectively. Then the key q-duration with respect to the (j, k)th key mortality rate is given by

$$KQD(P(\mathbf{Q}), (j, k)) = \frac{P(\tilde{\mathbf{Q}}(j, k)) - P(\mathbf{Q})}{\delta(j, k)},$$

where $P(\mathbf{Q})$ is the value of the portfolio on the basis of the mortality surface $\mathbf{Q}$.

[6]Setting $\delta(j, k)$ for all j and k to the same value would imply a parallel shift.

The set of key q-durations $\{KQD(P(\mathbf{Q}),(j,k));\ j=1,\ldots,n_k;\ k=1,\ldots,m\}$ as a whole measures the portfolio's price sensitivity to the entire mortality surface. We can adapt the algorithm in Section 3.3 accordingly to estimate the key q-durations for a multi-cohort pension plan.

Let $F_{j,k}(\mathbf{Q},(j.k))$ and $T_{j,k}$ be the present value (from the fixed receiver's viewpoint) and the time to maturity (in years) of a unit-notional mortality forward written on the (j,k)th key mortality rate, respectively. As in the one-dimensional setting, $KQD(F_{j,k}(\mathbf{Q}),(h,i))$ is $-100(1+r)^{-T_{j,k}}$ if $h=j$ and $i=k$, and is zero otherwise. This means that we can determine the weights on the mortality forwards independently. The appropriate notional amount of the mortality forward linked to the (j,k)th key mortality rate is

$$w(j,k)=\frac{KQD(V(\mathbf{Q}),(j,k))}{KQD(F_{j,k}(\mathbf{Q}),(j,k))}.$$

Assuming that assumptions 2, 3, 4, 6 and 7 in Section 5.1 still hold, we calculate the required notional amount of each mortality forward in the hedge portfolio. The details of the mortality forwards are summarized in Table 4.

As before, we evaluate the effectiveness of the hedge by examining X and X^*. Figure 12 shows the simulated distributions of X and X^*, which are based on 5,000 realizations of $\mathbf{Q}$ simulated from the fitted Cairns-Blake-Dowd model. The amount of risk reduction, R, provided by the portfolio of 14 mortality forwards is 95.8%. The results indicate that the two-dimensional extension can help us create a highly effective hedge for multi-cohort pension plans with a relatively small number of hedging instruments.

7 Other Issues

7.1 Small sample risk

In previous illustrations, we assume that there is no small sample risk (or sampling risk), that is, the risk that the realized mortality experience is different from the true mortality rate. Small sample risk is diversifiable, so it does not matter much if the pension plan is sufficiently large. However, for smaller plans, small sample risk may be significant and may affect the hedge's effectiveness.

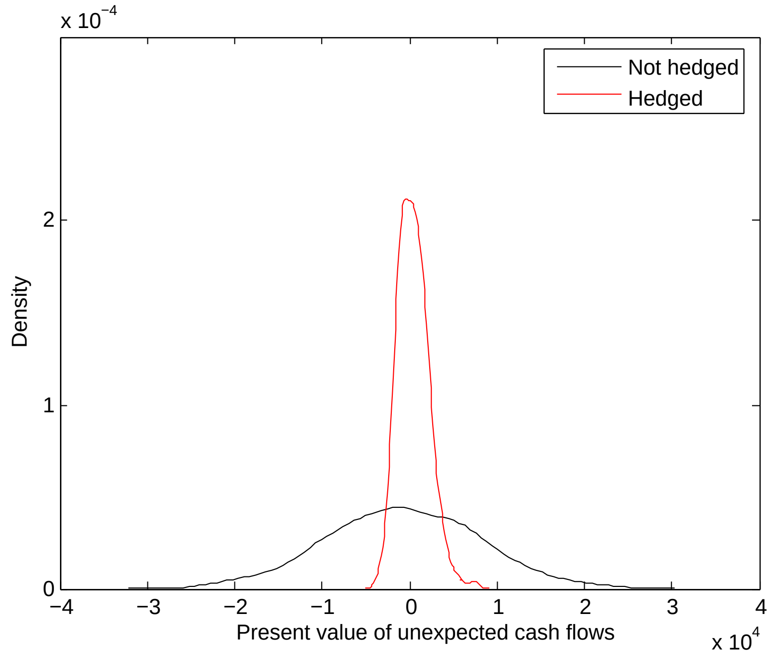

Figure 12: Simulated distributions of X and X^* for the multi-cohort synthetic pension plan.

Let us consider again the synthetic pension plan described in Section 5.1 and suppose that all assumptions (except assumption 2) continue to hold. The plan involves a single cohort of individuals who were born in year 1950 (aged 60 at time 0). We let $l(60)$ be the initial number of pensioners, and $l(x)$ be the number of pensioners who will survive to age x, $x = 61, 62, \ldots$. We incorporate small sample risk by treating the cohort of pensioners as a random survivorship group. This means that $l(x)$ for $x > 60$ is still a random variable even if the mortality curve $\mathbf{q}$ for the pensioners is completely known.

We model small sample risk with the following binomial death process:

$$l(x+1) \sim \text{Binomial}(l(x), 1 - q(x, 1950 + x)), \quad x = 60, 61, \ldots, 90,$$

which is then incorporated into the procedure for simulating unexpected cash flows as follows:

1. simulate a mortality curve, $\mathbf{q}$ using the fitted Cairns-Blake-Dowd model;

2. for each simulated $\mathbf{q}$, simulate the number of survivors $l(x)$, $x = 61, \ldots, 91$;

j	k	Time to maturity, $T_{j,k}$	Ref. age	Ref. year	Notional amt., $w(j,k)$
1	1	3	65	2012	8259.0
2	1	8	70	2017	5747.6
3	1	13	75	2022	4019.9
4	1	18	80	2027	2425.1
5	1	23	85	2032	1301.9
1	2	4	70	2013	5601.5
2	2	9	75	2018	3178.8
3	2	14	80	2023	1946.7
4	2	19	85	2028	1067.1
1	3	5	75	2014	4270.1
2	3	10	80	2019	1645.5
3	3	15	85	2024	917.3
1	4	6	80	2015	3331.6
2	4	11	85	2020	1251.7

Table 4: Reference ages, reference years and notional amounts of the mortality forwards in the longevity hedge for the multi-cohort synthetic pension plan.

3. calculate cash flows on the basis of the simulated $l(x)$;
4. repeat steps 1-3 to derive empirical distributions of the unexpected cash flows.

As in Section 5.1, mortality forwards with reference ages 65, 70, 75, 80 and 85 are used. Note that there is no change to the key q-durations.

In Figure 13 we depict the simulated distributions of the hedged and unhedged cash flows for different values of $l(60)$. When $l(60) = 10,000$, the hedge appears to be equally effective as that when small sample risk is not taken into account. However, for smaller values of $l(60)$, small sample risk has a significant effect on the reduction in volatility. To quantify the effect of small sample risk, we calculate the amount of risk reduction, R, for different values of $l(60)$ (see Table 5). The amount of risk reduction is reduced to only 66% when the initial number of pensioners is 500. This analysis suggests that for a small pension plan with say fewer than 500 members, a bespoke hedge may be a better alternative to one that is based on standardized instruments.

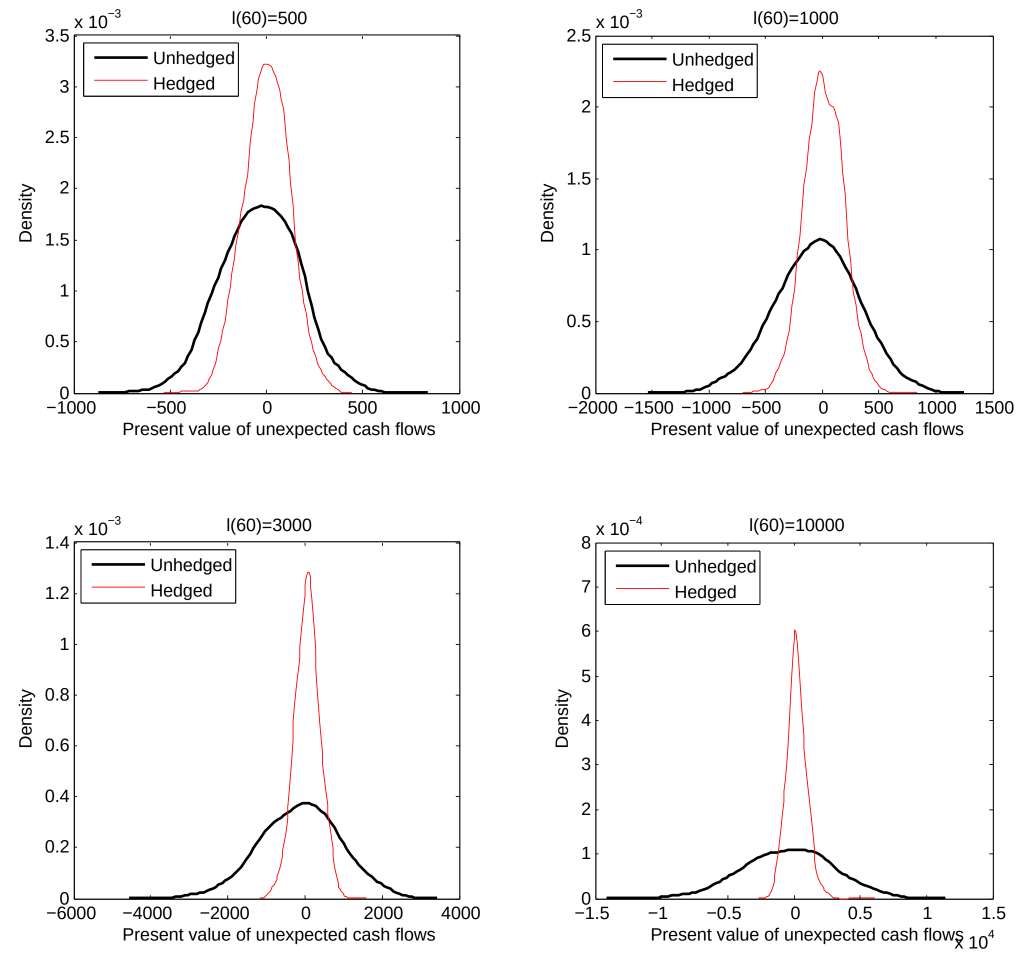

Figure 13: Simulated distributions of X and X^* for different values of $l(60)$.

7.2 Basis risk

When a hedger relies on standardized mortality forwards to hedge its longevity risk exposure, it is inevitably subject to the risk associated with the difference in the mortality experience between the hedger's population and the population to which the mortality forwards are linked. This risk, which is referred to as basis risk, has recently been studied by researchers including Cairns et al. (2011) and Coughlan et al. (2010).

Let $\mathbf{q}_1$ and $\mathbf{q}_2$ be the mortality curve for the hedger's population (population 1) and the reference population of the mortality forwards (population 2), respectively. In general, $\mathbf{q}_1$ and $\mathbf{q}_2$ are different, because of, for example, differing profiles of so-

$l(60)$	R
$+\infty$	97.20%
$10,000$	95.06%
$3,000$	90.20%
$1,000$	77.86%
500	66.01%

Table 5: Values of R for different values of $l(60)$.

cioeconomic group, lifestyle and geography.

The existence of basis risk makes a difference to the calculation of the weights on the mortality forwards. With basis risk, the appropriate weight on the mortality forward linked to the jth key mortality rate is given by

$$w(j) = \frac{KQD(V(\mathbf{q}_1), j)}{KQD(F_j(\mathbf{q}_1), j)} = \frac{KQD(V(\mathbf{q}_1), j)}{KQD(F_j(\mathbf{q}_2), j)} \frac{\partial q(x_j, t_j, 1)}{\partial q(x_j, t_j, 2)}, \tag{3}$$

where $q(x_j, t_j, i)$, $i = 1, 2$, is the jth key mortality rates for population i. The calculation of the adjustment factor $\frac{\partial q(x_j,t_j,1)}{\partial q(x_j,t_j,2)}$ is not straightforward, and may require a two population mortality model.

We consider the augmented common factor model (Li and Lee, 2005), which is defined as follows:

$$\begin{aligned}
\ln(m(x,t,i)) &= a(x,i) + B(x)K(t) + b(x,i)k(t,i) + \epsilon(x,t,i), \quad i = 1,2, \\
K(t) &= c + K(t-1) + \xi(t), \\
k(t,i) &= \phi_0(i) + \phi_1(i)k(t-1,i) + \zeta(t,i), \quad i = 1,2,
\end{aligned}$$

where $m(x,t,i)$ is the central death rate (at age x and year t) for population i; $\epsilon(x,t,i)$ is the error term; c, $\phi_0(i)$, and $\phi_1(i)$ are constants; and $\xi(t)$ and $\zeta(t,i)$ are innovation terms, i.i.d. normal with zero mean. The model assumes mortality rates for both populations are driven by a common stochastic factor, $K(t)$, which follows a random walk with drift, and a population specific stochastic factor, $k(t,i)$, which follows a first order autoregressive model. Parameters $B(x)$ and $b(x,i)$ are the sensitivities to $K(t)$ and $k(t,i)$, respectively.

In a parallel study, Li and Hardy (2011) proved that if the augmented common factor model is assumed, the adjustment factor in equation (3) can be calculated as

Hedger's population	With basis risk	Without basis risk
Canadian males	85.30%	95.18%
French males	87.56%	94.58%
Scottish males	87.48%	94.69%

Table 6: Values of R when population basis risk is present and absent.

follows:

$$\frac{\partial q(x,t,1)}{\partial q(x,t,2)} = \frac{q(x,t,1)(1+0.5m(x,t,2))^2 A(x,t,1)}{q(x,t,2)(1+0.5m(x,t,1))^2 A(x,t,2)}, \tag{4}$$

where $A(x,t,i) = B(x)c + b(x,i)\phi_1(i)^{t-t_0}(\phi_0(i) + (\phi_1(i)-1)k(t_0,i))$ for $i = 1, 2$, and t_0 denotes the year when the hedge is set up.

Let us revisit the synthetic pension plan described in Section 5.1. Suppose that all assumptions except assumption 2 still hold. We now allow the hedger's population to be different from the population to which the hedging instruments are linked. Specifically, we assume that the mortality forwards in the hedge portfolio are linked to the UK male population, while the hedger's population is either Canadian males, French males or Scottish males.

For each case, we calculate the appropriate weights on the five mortality forwards with equations (3) and (4). We then simulate the distributions of the unexpected cash flows, hedged and unhedged, on the basis of the augmented common factor model that is fitted to the corresponding pair of populations. The simulated distributions are displayed in Figure 14. The results indicate that a properly weighted portfolio of mortality forwards can reduce a significant amount of risk, even if basis risk exists.

Also shown in Figure 14 are the corresponding distributions when basis risk is absent, that is, when the hedge is created with mortality forwards that are linked to the hedger's own population. We observe that the reductions in variability are less remarkable when basis risk is present. The impact of basis risk can also be seen from Table 6, which shows the values of R for all cases we consider. On average, basis risk reduces the hedge effectiveness by 8%. The impact of basis risk is the highest when the hedger's population is linked to the population of Canadian males, which is possibly the least related to the population on which the mortality forwards are based.

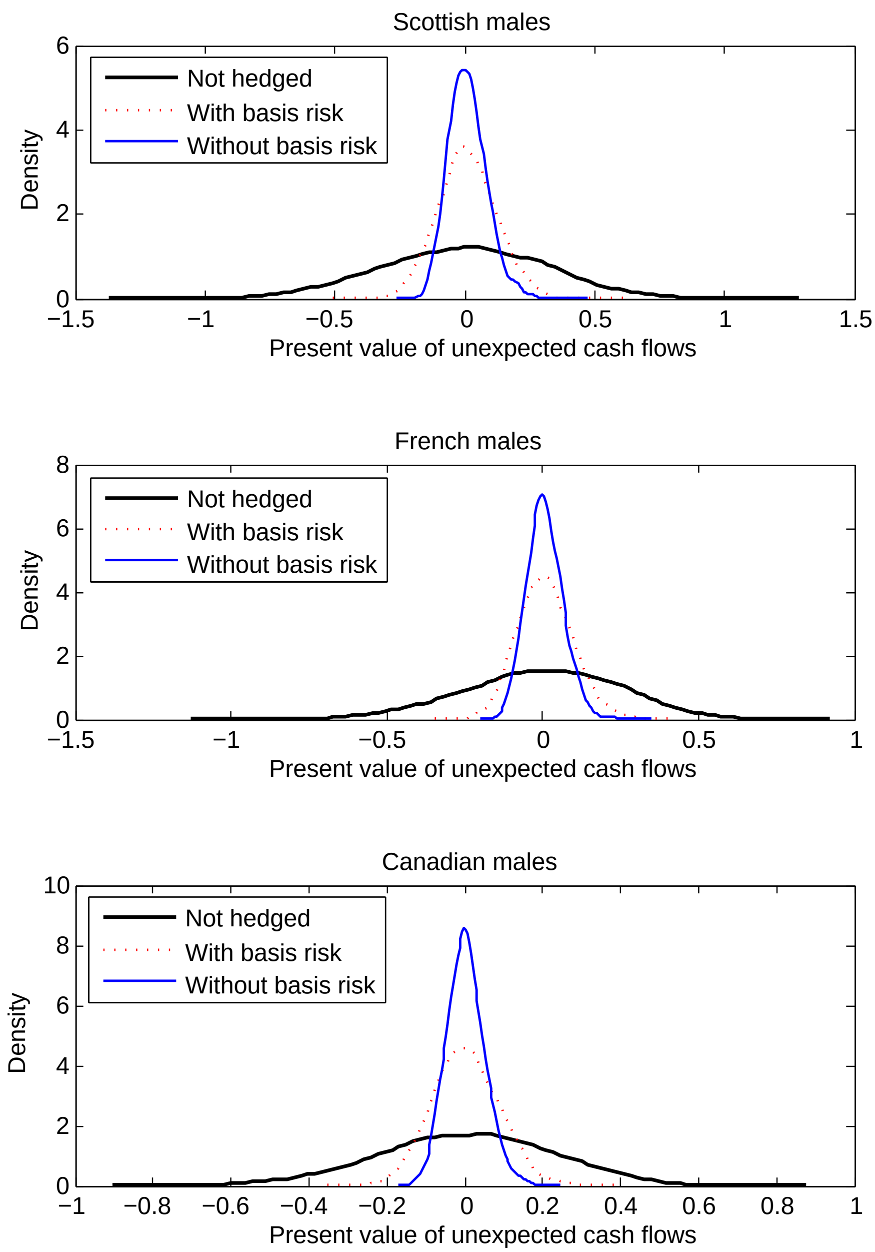

Figure 14: Simulated distributions of X and X^*, with and without basis risk.

8 Concluding Remarks

In this paper, we introduced a measure called key q-duration, which enables us to estimate the price sensitivity of a life-contingent liability to the underlying mortality curve. Given the key q-durations of a portfolio, one can easily hedge the longevity risk associated with the portfolio with only a handful of mortality forwards. The method we propose is easy to implement, yet a hedge calibrated with our method is almost equally effective as one that is calibrated with a computationally intensive optimization.

As Solvency II is rapidly approaching, insurers are increasingly concerned about the additional solvency capital requirement arising from their longevity risk exposures. We demonstrated that our hedging strategy can substantially reduce the longevity risk SCR of a pension portfolio. We emphasize that, however, hedgers must be aware of the sampling risk and basis risk involved in a standardized longevity hedge. In our illustration, basis risk can reduce a hedge effectiveness by 7-10%, while sampling risk can reduce a hedge effectiveness by 2-30%, depending on the size of the portfolio being hedged.

In practice, most pension plans involve a spectrum of birth cohorts. To improve applicability, we extended key q-durations to a two-dimensional setting, which is applicable to a mortality surface containing mortality curves for different birth cohorts. We illustrated the two-dimensional extension with a synthetic pension plan that involves 21 birth cohorts. The hedge we created could reduce the synthetic plan's longevity risk by 95.8%, while keeping the number of hedging instruments to a manageable level. A method requiring a modest number of hedging instruments is desirable, because when the market is still in its infancy, transactions are likely to be restricted to a limited number of instruments in which liquidity can be concentrated.

In the absence of any pricing information, it is assumed in our calculations that the forward mortality rates are the same as (not lower than) the corresponding best estimate mortality rates, which equivalently means that the hedge is costless. However, when a hedger is given a quote of the relevant forward mortality rates, it can easily incorporate the quoted forward rates into the simulations and calculate the cost of the hedge. The hedger can then compare it with the costs associated with other options, including bespoke longevity swaps and buy-ins, that is, the seeking of a financial institution to insure (lock-in) its liabilities. A financially optimal decision can then be made.

When deriving our hedging strategy, it is assumed that the hedger intends to eliminate as much longevity risk as possible. Some entities, however, may only want to transfer a portion of the risk to capital markets. For instance, a life insurer may only want to transfer to capital markets the residual longevity risk that cannot be naturally hedged with its life insurance books. In future research, it would be interesting to adapt the framework of key q-durations so that other hedging objectives, such as hedging 50% of the total risk, can be used.

The calculation of key q-durations requires a central estimate of future mortality and an assumed rate of interest. After the inception of the hedge, both quantities may change as new information is unfolded. This means that key q-durations and hence the optimal hedging strategy may also vary over time. With varying key q-durations, we may be able to achieve a better hedge effectiveness by dynamically adjusting the hedge portfolio. The benefit from dynamic hedging was not studied in this paper, but certainly deserves an investigation when sufficient information about liquidity and transaction costs becomes available.

In the absence of basis risk, the calculation of key q-durations does not require a specific stochastic mortality model. However, when basis risk is present, we need a two-population mortality model to estimate the adjustment term in equation (3). This means that the resulting hedging strategy is inevitably subject to some model uncertainty. Another avenue for future research to investigate how the adjustment term may change if a different two-population mortality model, for example, the models proposed by Cairns et al. (2011), is used. It is also warranted to validate the resulting hedging strategy by non-parametric means such as the block bootstrap.

References

Barbarin, J. (2008). Heath-Jarrow-Morton Modelling of Longevity Bonds and the Risk Minimization of Life Insurance Portfolios. *Insurance: Mathematics and Economics*, **43**, 41-55.

Börger, M. (2010). Deterministic Shock vs. Stochastic Value-at-Risk – An Analysis of the Solvency II Standard Model Approach to Longevity Risk. *Blätter DGVFM*, **31**, 225-259.

Cairns, A.J.G., Blake, D., and Dowd, K. (2006). A Two-Factor Model for Stochastic

Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance*, **73**, 687-718.

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. and Khalaf-Allah, M. (2011). Bayesian Stochastic Mortality Modelling for Two Populations. *ASTIN Bulletin*, to appear.

CEIOPS (2008). QIS4 Technical Specifications. Available at http://www.ceiops.eu/media/docman/Technical%20Specifications%20QIS4.doc.

CEIOPS (2010). Solvency II Calibration Paper. Available at http://www.ceiops.eu/media/files/publications/submissionstotheec/CEIOPS-Calibration-paper-Solvency-II.pdf.

Coughlan, G. (2009). Longevity Risk Transfer: Indices and Capital Market Solutions. In Barrieu, P.M. and Albertini, L. (eds.) The Handbook of Insurance Linked Securities. London: Wiley.

Coughlan, G.D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A.J.G., Blake, D. and Dowd, K. (2010). Longevity Hedging: A Framework for Longevity Basis Risk Analysis and Hedge Effectiveness. Pensions Institute Discussion Paper PAPER PI-1013. Available at http://www.pensions-institute.org/workingpapers/wp1013.html.

Coughlan, G., Epstein, D., Ong, A., Sinha, A., Hevia-Portocarrero, J., Gingrich, E., Khalaf-Allah, M., Joseph, P. (2007). LifeMetrics: A Toolkit for Measuring and Managing Longevity and Mortality Risks. Available at http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics/library.

Dahl, M., Melchior, M. and Moller, T. (2008). On Systematic Mortality Risk and Risk-Minimization with Survivor Swaps. *Scandinavian Actuarial Journal*, **2**, 114-146.

Ho, T.S.Y. (1992). Key Rate Durations: Measurers of Interest Rate Risks. *Journal of Fixed Income*, **2**, 29-44.

Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de (data downloaded on 20 January 2011).

Li, N. and Lee, R. (2005). Coherent Mortality Forecasts for a Group of Population: An Extension of the Lee-Carter Method. *Demography*, *42*, 575-594.

Li, J.S.-H. and Hardy, M.R. (2011). Measuring Basis Risk in Longevity Hedges. *North American Actuarial Journal*, in press.

Li, J.S.-H., Ng, A.C.Y. and Chan, W.S. (2011). On the Calibration of Mortality Forward Curves. *Journal of Futures Markets*, in press.

Wills, S. and Sherris, M. (2008). Integrating Financial and Demographic Longevity Risk Models: An Australian Model for Financial Applications. UNSW Australian School of Business Research Paper No. 2008ACTL05.

Zhou, R. and Li, J.S.-H. (2010). A Cautionary Note on Pricing Longevity Index Swaps. *Scandinavian Actuarial Journal*, first published on: 04 September 2010 (iFirst).