Parametric Mortality Indexes: From Index Construction to Hedging Strategies

Chong It Tan[a,*], Jackie Li[a], Johnny Siu-Hang Li[b], Uditha Balasooriya[a]

[a]*Nanyang Business School, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798.*

[b]*Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1.*

Abstract

In this paper, we investigate the construction of mortality indexes using the time-varying parameters in common stochastic mortality models. We first study how existing models can be adapted to satisfy the new-data-invariant property, a property that is required to ensure the resulting mortality indexes are tractable by market participants. Among the collection of adapted models, we find that the adapted Model M7 (the Cairns-Blake-Dowd model with cohort and quadratic age effects) is the most suitable for constructing mortality indexes, partly because it gives the best fit to the majority of the data sets we consider and partly because the three time-varying parameters in it are highly interpretable and rich in information content. Based on the three indexes created from this model, one can write a standardized mortality derivative called K-forward, which can be used to hedge longevity risk exposures. Another contribution of this paper is a method called key K-duration that permits one to calibrate a longevity hedge formed by K-forward contracts. Our numerical illustrations indicate that a K-forward hedge has a potential to outperform a q-forward hedge in terms of the number of hedging instruments required.

Keywords: Cairns-Blake-Dowd model, securitization, longevity risk reduction

1. Introduction

Pension plan sponsors can mitigate their longevity risk exposures by trading securities that are linked to future realized mortality. To date, the market for such securities is still in its infancy and has yet to overcome a number of challenges. As Blake et al. (2013) pointed out, one of these challenges is the creation of homogeneous and transparent instruments, which allow the market to concentrate liquidity. A significant step in overcoming this challenge is to develop tractable mortality indexes, upon which standardized mortality-linked securities can be written.

*Corresponding author. Tel: +65 9793 8939.
Email addresses: citan1@e.ntu.edu.sg (Chong It Tan), jackieli@ntu.edu.sg (Jackie Li), shli@uwaterloo.ca (Johnny Siu-Hang Li), auditha@ntu.edu.sg (Uditha Balasooriya)

Existing mortality indexes produced by investment banks such as J.P. Morgan, Deutsche Börse and Credit Suisse are constructed without using any model. One disadvantage of this construction approach is that the information content of each index is limited, which means the information reflected by an index is either highly aggregate (e.g., the life expectancy at birth) or specific (e.g., the death probability at a certain age).[1] It follows that a large number of such indexes would be needed to effectively hedge longevity risk, which arises from complex and non-parallel shifts in the underlying mortality curve. This problem hinders market development, because liquidity would be diluted across the large spectrum of indexes.

To improve the information content of a mortality index, one may use a model-based construction method, in which mortality indexes are developed from the time-varying parameters in a stochastic mortality model, such as the Lee-Carter model (Lee and Carter, 1992) and the Cairns-Blake-Dowd (CBD) model (Cairns et al., 2006). The model-based approach was first studied by Chan et al. (2014), who argued that the model on which index construction is based must satisfy the new-data-invariant property, which means that when the model is updated with new mortality data, the mortality indexes (time-varying parameters) for the previous years would not be affected. This property is crucially important, because once made public, an index value cannot (and should not) be changed. Chan et al. (2014) found that among the six stochastic models documented by Dowd et al. (2010), the original CBD model (also called Model M5) is the only model that satisfies the new-data-invariant property.

The first objective of this paper is to further investigate the construction of mortality indexes from stochastic mortality models. We begin with a discussion on how models other than the original CBD model can be adapted to satisfy the new-data-invariant property. We then evaluate the fit of the adapted models to the mortality data from 10 national populations on the basis of the Bayesian Information Criterion (BIC). Among the collection of adapted models, we find that the adapted Model M7 (the Cairns-Blake-Dowd model with cohort and quadratic age effects) yields the best BIC values for the majority of the data sets we consider. On top of that, the time-varying parameters in the model are highly interpretable and are able to reflect the varying age pattern of mortality improvement. For these reasons, we propose to use the three time-varying parameters in the adapted Model M7 jointly as mortality indexes. We call these indexes the 3-factor CBD mortality indexes.

Our second objective is to develop standardized securities written on the 3-factor CBD mortality indexes. In particular, we explain how a security called K-forward, proposed as a concept by Chan et al. (2014), can be used to hedge the longevity risk exposures of a pension plan. The structure of a K-forward is identical to that of a q-forward (see, e.g., Coughlan, 2009), except that the reference rates to which the contracts are linked are different. In more detail, the payoff from a q-forward depends on the realized death probability at a reference age in a reference year, whereas that from a K-forward depends on a realized 3-factor CBD mortality index in a reference year. Compared to a q-forward, a K-forward is an even simpler building block, because its reference rate contains only one parameter (the reference year)

[1]The mortality index of Credit Suisse is based on the life expectancy at birth of the US population, while that of JP Morgan is based on the death probabilities of four national populations.

instead of two.

To ensure that a hedge formed by standardized instruments is effective, there is a need to calibrate it. Following the lines of Cairns et al. (2008), Cairns (2011), Coughlan et al. (2007), Plat (2009, 2010), Li and Luo (2012), Lin and Tsai (2013), Tsai et al. (2010) and Wang et al. (2010), we contribute a measure called key K-duration, which measures a liability's price sensitivity to a specific segment in the time trend of a 3-factor CBD mortality index. The required notional amounts of K-forwards can be determined readily by equating the key K-durations of the portfolio of K-forwards and the liability being hedged. The key K-duration measure is parallel to Li and Luo's (2012) key q-duration, which measures the change in a liability's value due to a small change in a death probability. It also has a close resemblance to Cairns' (2011) approximate deltas (with respect to the time-varying parameters in the original CBD model) and to Plat's (2009) minimum variance hedge ratios that are derived by considering the shifts of the two time-varying parameters in the original CBD model.

The calculation of key K-durations does not require any simulation, so the execution of the proposed hedging strategy is quick and requires minimal computational effort. Our numerical illustrations indicate that the proposed hedging strategy is effective in reducing a portfolio's longevity risk exposure, even if parameter uncertainty and sampling risk are taken into account. Furthermore, although the mortality indexes are derived from the adapted Model M7, our proposed hedging strategy also works well under scenarios simulated from other stochastic mortality models, suggesting that the success in hedging is largely independent of the simulation model and is likely to be achievable in practice.

The advantage of our proposed hedging strategy is particularly apparent when the portfolio being hedged involves individuals who were born in different years. First, a K-forward hedge is easier to execute in comparison to a q-forward hedge, which requires the hedger to determine the key cohorts in the portfolio. Second, our numerical illustrations suggest that compared to a q-forward hedge, a K-forward hedge giving a comparable hedge effectiveness involves a smaller number of securities. This helps the market to concentrate liquidity, thereby facilitating market development.

The remainder of this paper is organized as follows. Section 2 discusses the construction of mortality indexes using common stochastic mortality models. Section 3 specifies a K-forward contract, defines the key K-duration measure and details the proposed hedging strategy. Section 4 illustrates the proposed methods with a hypothetical pension plan involving one single cohort and investigates important issues such as sampling risk. Section 5 presents the generalization to multiple birth cohorts. Section 6 concludes the paper.

2. Constructing Mortality Indexes

In this section, we revisit the problem of model-based mortality index construction, which was previously studied by Chan et al. (2014). The conventions below are used throughout the discussion:

- $m_{x,t} = \frac{D_{x,t}}{E_{x,t}}$ is the central death rate at age x in year t;

Table 1
Specifications of the six stochastic mortality models under consideration.

Model M1: The Lee-Carter Model	
$\ln(m_{x,t}) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)}$	(2 constraints)
Model M2: The Renshaw-Haberman Model	
$\ln(m_{x,t}) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)} + \beta_x^{(3)}\gamma_{t-x}^{(3)}$	(4 constraints)
Model M3: The Age-Period-Cohort Model	
$\ln(m_{x,t}) = \beta_x^{(1)} + n_a^{-1}\kappa_t^{(2)} + n_a^{-1}\gamma_{t-x}^{(3)}$	(3 constraints)
Model M5: The Original Cairns-Blake-Dowd Model	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x})$	(No constraint)
Model M6: The Cairns-Blake-Dowd Model with a Cohort Effect Term	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \gamma_{t-x}^{(3)}$	(2 constraints)
Model M7: The Cairns-Blake-Dowd Model with Cohort Effect and Quadratic Age Effect Terms	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \kappa_t^{(3)}((x-\bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)}$	(3 constraints)

- $D_{x,t}$ is observed number of deaths at age x in year t;
- $E_{x,t}$ is the matching exposures at age x in year t;
- $q_{x,t}$ is the probability that a person aged x at time t will die between time t and $t+1$;
- $\beta_x^{(i)}, i = 1, 2, 3$, are age-specific parameters;
- $\kappa_t^{(i)}, i = 1, 2, 3$, are time-varying parameters;
- $\gamma_c^{(i)}, i = 3, 4$, where $c = t - x$ denotes year of birth, are cohort-related parameters;
- n_a is the number of ages covered in the sample age range;
- $\bar{x}$ is the mean age over the sample age range;
- $\hat{\sigma}_x^2$ is the mean of $(x-\bar{x})^2$ over the sample age range.

The candidate models under consideration are the six stochastic mortality models discussed by Dowd et al. (2010). The specifications of the six models are displayed in Table 1. In fitting these models (except Model M5), we need to impose identifiability constraints to stipulate parameter uniqueness. The number of identifiability constraints needed for each model is displayed in parentheses in Table 1. We refer interested readers to Cairns et al. (2009) for a deeper discussion on the identifiability constraints.

2.1. The New-Data-Invariant Property

Chan et al. (2014) argued that the model on which index construction is based must satisfy the new-data-invariant property. In more detail, the time-varying parameters in the model must remain unchanged when the model is fitted to additional years of data. This property is crucial to the tractability of the mortality indexes, because without it historical index values will need to be revised every time when new mortality data becomes available. Furthermore, if an index value at a particular time point is not fixed, it would be difficult, if not impossible, to define the payoff from an instrument written on the index.

Chan et al. (2014) provided two sufficient conditions for the new-data-invariant property to hold. The first condition is related to the log-likelihood. Assuming $D_{x,t} \sim \text{Poisson}(E_{x,t}m_{x,t})$, the log-likelihood of a stochastic mortality model can be expressed as

$$l = \sum_{t=t_{\text{start}}}^{t_{\text{end}}} \sum_{x=x_0}^{x_1} D_{x,t} \ln(E_{x,t}m_{x,t}) - E_{x,t}m_{x,t} - \ln(D_{x,t}!) = \sum_{t=t_{\text{start}}}^{t_{\text{end}}} \lambda(t),$$

where $[x_0, x_1]$ is the sample age range, $[t_{\text{start}}, t_{\text{end}}]$ is the sample period and $\lambda(t)$ represents the contribution to the log-likelihood from data of year t. We require the log-likelihood to be separable, that is, for $t \neq s$, $\lambda(t)$ and $\lambda(s)$ do not contain any common free parameters. This condition implies that we can estimate the mortality indexes in each year independently. The second condition is associated with the identifiability constraints. In particular, we require that the estimation process involves no identifiability constraint, because otherwise when the model is fitted to additional years of data, parameters for the previous years would be changed to conform to the constraint imposed on them. Among the six candidate models we consider, Model M5 is the only model that meets these two conditions.

In what follows, we discuss how the other five models can be adapted to satisfy the new-data-invariant property. The way we adapt a model depends on whether the model contains age-specific parameters, cohort effect parameters, or both.

To make Model M1 (which contains age-specific parameters) meet the new-data-invariant property, we could keep the age-specific parameters $\beta_x^{(i)}$, $i = 1, 2$, fixed when we update the model. In using this adaptation, the log-likelihood becomes separable and identifiability constraints are no longer needed. We denote the adapted version of Model M1 by Model M1*.

For Models M6 and M7 (which contain cohort effect parameters), we could first estimate the time-varying parameters $\kappa_t^{(i)}$, where $i = 1, 2$ for Model M6 and $i = 1, 2, 3$ for Model M7, and then the cohort effect parameters $\gamma_{t-x}^{(i)}$, where $i = 3$ for Model M6 and $i = 4$ for Model M7, from the residuals. In this way, the estimation of the time-varying parameters would be based on a separable log-likelihood and requires no identifiability constraint. The time-varying parameters in the adapted models thus possess the new-data-invariant property. We denote the adapted versions of Models M6 and M7 by Models M6* and M7*, respectively.

For Models M2 and M3 (which contain both age-specific and cohort effect parameters), we need both of the previously mentioned adaptations. Every time we update the models, we

Table 2
Specifications of the stochastic mortality models that are adapted to satisfy the new-data-invariant property.

Model M1*: The adapted version of model M1	
$\ln(m_{x,t}) = \beta_x^{(1)*} + \beta_x^{(2)*}\kappa_t^{(2)}$	(No constraint)
Model M2*: The adapted version of model M2	
$\ln(m_{x,t}) = \beta_x^{(1)*} + \beta_x^{(2)*}\kappa_t^{(2)} + \beta_x^{(3)*}\gamma_{t-x}^{(3)*}$	(No constraint)
Model M3*: The adapted version of model M3	
$\ln(m_{x,t}) = \beta_x^{(1)*} + n_a^{-1}\kappa_t^{(2)} + n_a^{-1}\gamma_{t-x}^{(3)*}$	(No constraint)
Model M6*: The adapted version of model M6	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \gamma_{t-x}^{(3)*}$	(No constraint)
Model M7*: The adapted version of model M7	
$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \kappa_t^{(3)}((x-\bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)*}$	(No constraint)

keep the age-specific parameters fixed, estimate the time-varying parameters first and then the cohort effect parameters from the residuals. In this way, the time-varying parameters in the adapted model would meet the new-data-invariant property. We denote the adapted versions of Models M2 and M3 by Models M2* and M3*, respectively.

The five adapted candidate models are specified in Table 2, where $\beta_x^{(i)*}$, $i = 1, 2, 3$, and $\gamma_{t-x}^{(i)*}$, $i = 3, 4$, are the parameters that are subject to the proposed adaptations.

2.2. The Impact of the Proposed Adaptations

An adapted model yields a reduced goodness-of-fit in comparison to the corresponding original model, because the estimation of an adapted model is conditioned on some pre-estimated parameters. Specifically, the estimation of the time-varying parameters in Models M1*, M2* and M3* is conditioned on the fixed age-specific parameters, which, in practice, can be determined by data over a restricted sample period, say $[t_{\text{start}}, t_{\text{mid}}]$, where $t_{\text{mid}} < t_{\text{end}}$.[2] On the other hand, the estimation of the cohort effect parameters in Model M2*, M3*, M6* and M7* is conditioned on the estimated time-varying parameters.

To evaluate the impact of the proposed adaptations, we compare the goodness-of-fit to the data over the period of $[t_{\text{mid}} + 1, t_{\text{end}}]$ produced by the original models and the adapted models (whose fixed age-specific parameters, if any, are determined by the data over the period of $[t_{\text{start}}, t_{\text{mid}}]$). The metric we use is $\hat{l}_r - \hat{l}_{r^*}$, the difference between the maximized log-likelihoods for model r and its adapted version r^*.

[2]Parameter estimation over the restricted sample period of $[t_{\text{start}}, t_{\text{mid}}]$ is still subject to the usual identifiability constraints.

Table 3
The data sample periods for the 10 populations under consideration.

Population	Data sample period
Australia, Canada, England and Wales, France, Japan, Norway, Sweden, United States	1950-2009
New Zealand	1950-2008
Taiwan	1970-2010

Table 4
The reductions in log-likelihood values due to the proposed adaptations.

	Males					Females				
Population	M1*	M2*	M3*	M6*	M7*	M1*	M2*	M3*	M6*	M7*
Australia	3734	3125	1624	335	19	1045	6249	1431	1150	24
Canada	6948	24836	1404	203	15	1060	6809	805	1313	194
England and Wales	13886	2362	2600	90	83	9566	1633	2353	1252	27
France	2051	2282	3466	4539	1820	7847	10668	4107	14868	545
Japan	24212	1805	4118	2655	537	88416	2707	3853	13185	261
New Zealand	972	830	399	32	1	555	538	389	191	38
Norway	359	2351	437	54	4	94	267	377	480	27
Sweden	979	364	902	190	12	361	206	892	943	76
Taiwan	8033	6151	951	910	5	3415	1886	718	505	15
United States	53358	14229	10554	8860	1766	20218	11028	7202	14371	2224

We consider gender-specific mortality data from 10 populations across different geographical regions, including Australasia (Australia, New Zealand), East Asia (Japan, Taiwan), the Nordic region (Norway, Sweden), Western Europe (England and Wales, France) and North America (Canada, United States). All data are obtained from the Human Mortality Database (2013). The sample period for each data set is shown in Table 3. We consider an age range of 40-90, because models in the CBD family may not fit the accident hump at younger ages well and the data beyond age 90 are subject to reliability problems. With this age range, we have $\bar{x} = 65$ and $\hat{\sigma}_x^2 = \frac{650}{3}$. Furthermore, in fitting the models that are built for $q_{x,t}$ (Models M5, M6, M6*, M7 and M7*), we assume a constant force of mortality between integer ages, which implies that $m_{x,t} = -\ln(1 - q_{x,t})$.

In evaluating the goodness-of-fit, we set t_{mid} to $t_{\text{end}} - 15$, which means the evaluation is based on the fit to the most recent 15 years of mortality data. The results are tabulated in Table 4. It can be seen that Model M7* consistently gives the smallest reductions in log-likelihood values, indicating that the performances of Models M7* and M7 are similar. One possible explanation is that the three time-varying parameters in Model M7 capture the majority of the historical variations in mortality, so that whether or not the cohort effect parameters are conditionally estimated has little impact on the overall goodness-of-fit. However, other adapted models yield a significantly worse fit in comparison to their original versions. The reductions in log-likelihood values are the most prominent for Models M1*, M2* and M3*, possibly because fixing age-specific parameters makes the model structures a lot more stringent.

Table 5
Values of the Bayesian Information Criterion (BIC) produced by the five adapted models and Model M5.

	Males						Females					
Population	M1*	M2*	M3*	M5	M6*	M7*	M1*	M2*	M3*	M5	M6*	M7*
Australia	15766	14937	11419	11942	8894	8199	10051	20938	10810	17545	10452	7965
Canada	22776	58868	11543	10697	9298	8781	10677	22564	10066	17673	11453	8752
England and Wales	39538	14465	14497	16137	9887	9454	30549	12765	13763	24662	12298	9097
France	15899	14421	16359	56935	23612	13372	26739	30858	17351	121168	41037	11041
Japan	62618	14144	19169	40456	19438	12627	187997	15358	17880	132618	38947	12164
New Zealand	8852	9100	7752	6833	6902	6848	7880	8375	7605	7869	7213	6801
Norway	7801	12354	7989	7338	7057	7021	7061	7922	7621	9977	7716	6849
Sweden	9668	8935	9499	9175	8035	7652	8269	8460	9273	14262	9405	7589
Taiwan	24720	21321	10446	15448	10333	8580	14829	12271	9449	12044	8820	7888
United States	123233	40388	34770	102543	40562	17818	54680	33198	25935	122063	52332	16553

2.3. Finding the Best-Fitting Adapted Model

In this subsection we compare the performances of the five adapted models and Model M5 (which requires no adaptation). The metric we use is the Bayesian Information Criterion (BIC), defined by $\text{BIC}_r = -2\hat{l}_r + v_r \ln n_d$. The BIC indicates a model's goodness-of-fit (reflected by the maximized log-likelihood $\hat{l}_r$), taking into account the number of observations n_d and the effective number of free parameters v_r in the model. We prefer a model with a lower BIC value.

The BIC values produced by the five adapted models and Model M5 are depicted in Table 5. We observe that Model M7* gives the best BIC values for all populations except New Zealand males.[3] This is possibly because, as previously mentioned, the adaptation involved in Model M7* does not have much impact on the model's performance. Another possible explanation is that the quadratic age effect term (which is unique to Model M7*) has a significant explanatory power.

Judging from the evaluation results, Model M7* is the best candidate model for constructing mortality indexes. Furthermore, as we are going to discuss in the next subsection, the time-varying parameters in Model M7* are easy to interpret and are able to explain the varying age pattern of mortality improvement. For these reasons, we propose to use the three time-varying parameters, $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and $\kappa_t^{(3)}$, in Model M7* as mortality indexes. We call these indexes the 3-factor CBD mortality indexes.

2.4. Interpretations of the 3-factor CBD Mortality Indexes

To recap, Model M7* can be written as

$$\ln\left(\frac{q_{x,t}}{1-q_{x,t}}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \kappa_t^{(3)}((x-\bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)*}, \tag{1}$$

[3] For New Zealand males, Model M5 gives the smallest BIC value. The fact that the BIC value produced by Model M5 is less than those produced by Models M6* and M7* implies that cohort effects are not significant in this population. However, the quadratic age effect term is significant, because Model M7* yields a smaller BIC value compared to Model M6*.

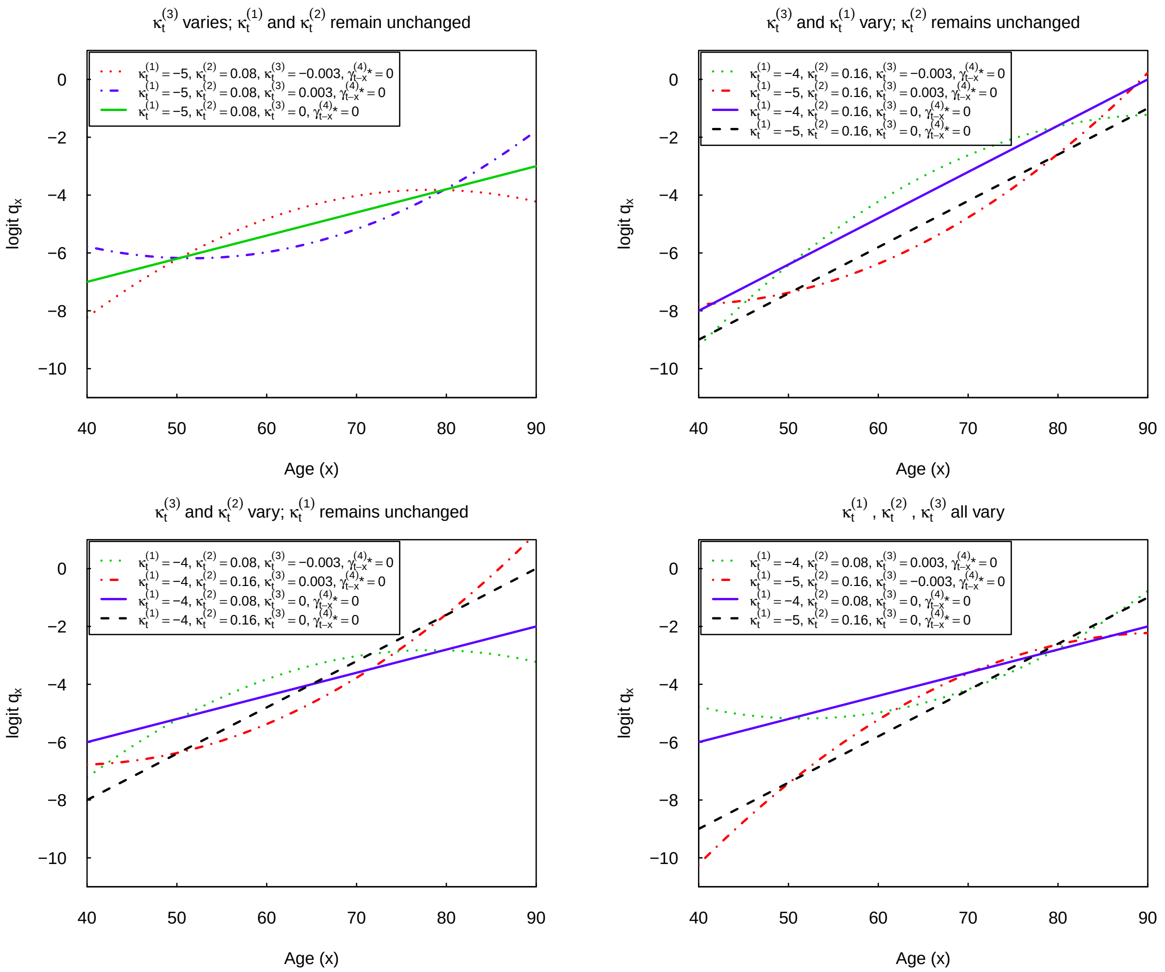

Fig. 1. Changes in the logit-transformed mortality curve under different combinations of the 3-factor CBD mortality indexes.

which means that at a given time t, the logit-transformed death probability is a quadratic function of $(x - \bar{x})$, with coefficients being the mortality indexes $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and $\kappa_t^{(3)}$. This simple structure permits us to interpret the three mortality indexes straightforwardly.

The first index, $\kappa_t^{(1)}$, captures the level of the mortality curve in logit scale. A reduction in $\kappa_t^{(1)}$ signifies an overall mortality improvement across all ages. The impact of a change in this mortality index is illustrated in the top right panel of Fig. 1 (the solid and dashed lines).

The second index, $\kappa_t^{(2)}$, represents the slope of the logit-transformed $q_{x,t}$ with respect to age x. As demonstrated in the bottom left panel of Fig. 1 (the solid and dashed lines), an increase in $\kappa_t^{(2)}$ leads to a steeper logit-transformed mortality curve, which implies that mortality at younger ages (below the mean age $\bar{x}$) declines more rapidly than that at older ages (above the mean age $\bar{x}$).

The third index, $\kappa_t^{(3)}$, measures the curvature of the logit-transformed mortality curve. A higher value of $\kappa_t^{(3)}$ implies an increase in $q_{x,t}$ over the age ranges of 40-50 and 80-90 but

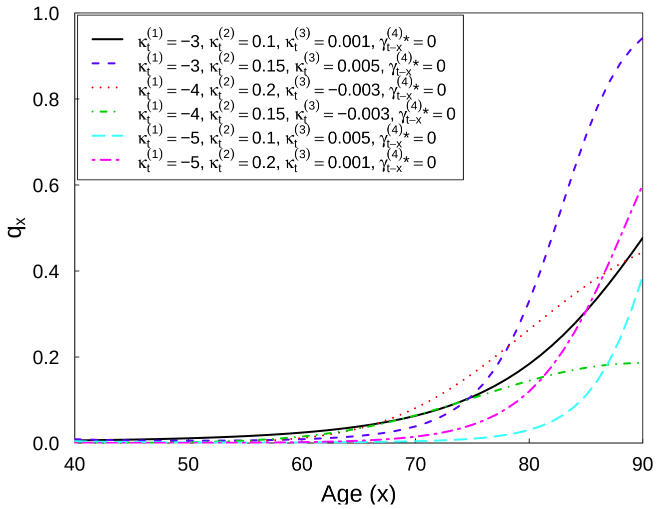

Fig. 2. Mortality curves implied by six different combinations of the 3-factor CBD mortality indexes.

a decrease in $q_{x,t}$ over the age range of 51-79. This effect is due to the term $((x - \bar{x})^2 - \hat{\sigma}_x^2)$, which is positive when $51 \leq x \leq 79$ and negative otherwise. The impact of $\kappa_t^{(3)}$ for fixed combinations of $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ is shown in the top left panel of Fig. 1. The interaction between $\kappa_t^{(3)}$ and the other two indexes is demonstrated in the other panels of Fig. 1.

Besides being easy to interpret, desirable mortality indexes should be able to capture a wide range of mortality age patterns, so that securities written on the indexes can effectively hedge longevity risk, which fundamentally arises from random non-parallel shifts in the underlying mortality curve. The ability of the 3-factor CBD mortality indexes to represent different age patterns of mortality can be observed from Fig. 1 (in logit scale) and Fig. 2 (in the original scale).

Let us illustrate the interpretations of the mortality indexes by considering a portfolio of 3-year term life insurance contracts that are issued to persons aged 40. The present value of this portfolio is higher than expected when:

- the future values of $\kappa_t^{(1)}$ are higher than expected (which implies the overall mortality improves slower than expected),
- the future values of $\kappa_t^{(2)}$ are lower than expected (which implies mortality at younger ages (below age 65) improves slower than expected), and
- the future values of $\kappa_t^{(3)}$ are higher than expected (which implies mortality over the age ranges of 40-50 and 80-90 improves slower than expected).

To further illustrate, we consider a closed pension plan, which makes pension payments to persons aged 65 to 91. The following statements apply to this pension plan:

- The present value is higher than expected if the future values of $\kappa_t^{(1)}$ are lower than expected (which implies the overall mortality improves faster than expected).
- The present value is higher than expected if the future values of $\kappa_t^{(2)}$ are lower than

expected (which implies mortality at older ages (above age 65) improves faster than expected).

- The impact of $\kappa_t^{(3)}$ on the present value is not clear-cut, depending very much on the demographics of the plan. For instance, if the majority of the pensioners are now aged above 80, then lower future values of $\kappa_t^{(3)}$ are likely to yield a higher present value.

3. Hedging Longevity Risk with K-forwards

With three mortality indexes capturing different aspects of mortality patterns, we can split longevity risk into three constituent risks, namely "K1 risk", "K2 risk" and "K3 risk". K1 risk refers to the uncertainty associated with overall mortality improvements, while K2 and K3 risks reflect the uncertainty surrounding the first and second order age effects, respectively.

Chan et al. (2014) proposed, as a concept, a simple security called K-forward that can be used to hedge each of the three constituent risks. In this section, we detail the mechanism of a K-forward and explain how an effective longevity hedge can be built using a portfolio of K-forward contracts.

3.1. Specification of K-forwards

A K-forward contract is a zero-coupon swap that exchanges on the maturity date a fixed amount, determined at time 0 (when the contract is established), for a floating amount that is proportional to a CBD mortality index (the reference rate) for a certain population (the reference population) in some future time (the reference year).

For a K-forward contract with a reference year t^*, the reference rate is $\kappa_{t^*}^{(i)}$, an unknown at time 0. The fixed amount is proportional to the corresponding forward mortality index, denoted as $\tilde{\kappa}_{t^*}^{(i)}$. This forward value is determined in such a way that no payment exchanges hands at time 0. Mathematically, the net payoff to the fixed rate receiver of a K-forward contract can be written as

$$Y \times (\tilde{\kappa}_{t^*}^{(i)} - \kappa_{t^*}^{(i)}), \quad i = 1, 2, 3,$$

where Y is the notional amount. The settlement that takes place on the maturity date T^* is illustrated in Fig. 3. In practice, the maturity date T^* may be slightly later than the reference year t^* due to the time lag in the availability of the mortality index data. For simplicity, we assume in this paper that the maturity date T^* is last day of the reference year t^*.

Pension plans and life insurers can hedge their longevity risk exposures with K-forward contracts. As an illustration, we consider a closed pension plan that promises pension payments up to age 91. At present, the majority of the pensioners in the plan are aged 65-79. Since the pension payouts are negatively associated with the future values of $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$, the pension plan sponsor should write K1- and K2-forwards (K-forwards linked to the first

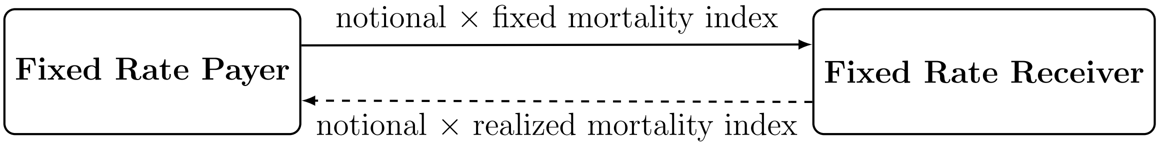

Fig. 3. Settlement of a K-forward contract at maturity.

two CBD mortality indexes) as a fixed rate receiver to offset the adverse outcome when the future values of $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ turn out to be lower than expected. As the majority of the pensioners are currently younger than age 80, higher than expected values of $\kappa_t^{(3)}$ in the first few years will result in a larger aggregate payout. After a sufficiently long time period (say, 15 years), the age profile of the pensioners will be concentrated in the age range of 80-90, which means lower than expected values of $\kappa_t^{(3)}$ in later years will lead to more pension payments. To hedge its exposure to K3-risk, the pension plan sponsor should write short-dated K3-forwards as a fixed rate payer and long-dated K3-forwards as a fixed rate receiver.

The mechanism of K-forwards is similar to that of q-forwards, which were transacted in deals such as the one between Lucida PLC and J.P. Morgan in January 2008. The only difference between the two instruments is that the floating and fixed legs of a q-forward are respectively determined by q_{x^*,t^*}, the death probability at the reference age x^* in the reference year t^*, and $q^f_{x^*,t^*}$, the corresponding forward mortality rate.

Compared to a q-forward, a K-forward is an even simpler security, because it is categorized by only one parameter, the reference year. A K-forward longevity hedge may be easier to implement, because it does not require the hedger to make decisions about the reference ages. This advantage, as we are going to demonstrate in Section 5, is prominent when the portfolio being hedged covers multiple birth cohorts, involving many possible combinations of reference ages and years. From another viewpoint, K-forwards may be more conductive to the development of liquidity, because the market can focus on contracts written on three indexes only instead of the full range of discrete ages.

3.2. Key K-duration

To ensure a K-forward longevity hedge is effective, there is a need to calibrate it. We propose a calibration method called key K-duration (KKD), which is largely parallel to the key q-duration method recently proposed by Li and Luo (2012).

To explain the method of KKD, we need to define the notion of key K-indexes, the mortality indexes $\kappa_t^{(i)}$, $i = 1, 2, 3$, in future years $t = t_1, t_2, \ldots, t_n$ (the key years). It is assumed that K-forwards with reference years equal to the n key years are traded in the market and that no other K-forwards are available.

A KKD measures the change in the value of a liability with respect to a small change in a key K-index. Mathematically, the KKD associated with the ith CBD mortality index and the jth key year can be expressed as

$$KKD_i(P(\boldsymbol{\kappa}), j) = \frac{\partial P(\boldsymbol{\kappa})}{\partial \kappa_{t_j}^{(i)}}, \tag{2}$$

where $P(\boldsymbol{\kappa})$ denotes the portfolio value as a function of the vector of the mortality indexes $\boldsymbol{\kappa}$. The idea behind the KKD calibration method is to split the future time line into parts by the n key years. Then the longevity hedge is calibrated by matching the price sensitivities (measured by KKDs) of the pension plan and hedge portfolio over each subperiod.

In practice, it is generally difficult to compute a KKD by analytical means. To solve this problem, we propose a numerical estimation method, which involves the following two working assumptions:

- A shock in the key K-index $\kappa_{t_j}^{(i)}$ by an amount of $\delta^{(i)}(j)$ is accompanied with a level shift in $\kappa_t^{(i)}$ over the period of $t_j \leq t < t_{j+1}$ by the same amount.[4] The shock $\delta^{(i)}(j)$ is assumed to have no impact on the trajectory of $\kappa_t^{(i)}$ beyond $t_j \leq t < t_{j+1}$.
- The shock $\delta^{(i)}(j)$ has no impact on $\kappa_t^{(h)}$ for all $h \neq i$ and t.

Admittedly, these working assumptions are rather arbitrary, but they still lead to satisfying hedging results (see Sections 4 and 5).

Let $\breve{\boldsymbol{\kappa}}(\delta^{(i)}(j))$ be the vector of mortality indexes that is subject to the shift $\delta^{(i)}(j)$. The proposed estimation procedure can be summarized as follows:

1. take $\boldsymbol{\kappa}$ as the best estimate of mortality indexes;
2. assuming the shift $\delta^{(i)}(j)$ is 0.1% of the best estimate, calculate $\breve{\boldsymbol{\kappa}}(\delta^{(i)}(j))$;
3. estimate the KKD associated with the ith CBD mortality index and the jth key year as $KKD_i(P(\boldsymbol{\kappa}), j) \approx \frac{P(\breve{\boldsymbol{\kappa}}(\delta^{(i)}(j))) - P(\boldsymbol{\kappa})}{\delta^{(i)}(j)}$.

3.3. Building a Longevity Hedge using Key K-durations

We now explain how key K-durations can be used to construct a static longevity hedge for a life-contingent liability at time 0. The hedge we build is a value hedge, which aims to reduce the variability of the liability's present value.

We set time 0 to be the beginning of the current year t_0. Suppose that K1-, K2- and K3-forwards with reference years $t_1, t_2, \ldots, t_n$ are available in the market. From the fixed rate receiver's viewpoint, the random present values of the cash flows from the K-forwards with reference year t_j can be expressed as

$$F_j^{(i)}(\boldsymbol{\kappa}) = (1+r)^{-(T_j - t_0)}(\tilde{\kappa}_{t_j}^{(i)} - \kappa_{t_j}^{(i)}),\ i = 1, 2, 3, \tag{3}$$

[4] We define t_{n+1} as $+\infty$.

where r is the interest rate at which the cash flows are discounted.[5] It can be seen that

$$KKD_i(F_j^{(i)}(\boldsymbol{\kappa}), j) = -(1+r)^{-(T_j - t_0)}, \tag{4a}$$

$$KKD_i(F_j^{(i)}(\boldsymbol{\kappa}), h) = 0, \qquad j \neq h, \tag{4b}$$

$$KKD_h(F_j^{(i)}(\boldsymbol{\kappa}), j) = 0, \qquad i \neq h. \tag{4c}$$

Eq. (4b) and (4c) arise from the two working assumptions in Section 3.2.

We let $w^{(1)}(j)$, $w^{(2)}(j)$ and $w^{(3)}(j)$ be the notional amounts of the K1-, K2- and K3-forwards with reference year t_j, respectively. The present value of hedge portfolio at time 0 is

$$H(\boldsymbol{\kappa}) = \sum_{h=1}^{n} \sum_{g=1}^{3} w^{(g)}(h) F_h^{(g)}(\boldsymbol{\kappa}).$$

Let $L(\boldsymbol{\kappa})$ be the random present value of liability at time 0. To achieve effective hedging, the KKDs of hedge portfolio must match the KKDs of liability. This is equivalent to solving the following set of conditions for $i = 1, 2, 3$ and $j = 1, 2, \ldots, n$:

$$\begin{aligned} KKD_i(L(\boldsymbol{\kappa}), j) &= KKD_i(H(\boldsymbol{\kappa}), j) \\ &= KKD_i\Big(\sum_{h=1}^{n} \sum_{g=1}^{3} w^{(g)}(h) F_h^{(g)}(\boldsymbol{\kappa}), j\Big) \\ &= w^{(i)}(j) KKD_i(F_j^{(i)}(\boldsymbol{\kappa}), j). \end{aligned} \tag{5}$$

The simplification, which follows from Eq. (4b) and (4c), allows us to determine the required notional amount of each K-forward separately without solving a system of equations. Furthermore, by using Eq. (4a), we obtain

$$w^{(i)}(j) = -(1+r)^{T_j - t_0} KKD_i(L(\boldsymbol{\kappa}), j), \tag{6}$$

where $KKD_i(L(\boldsymbol{\kappa}), j)$ can be numerically estimated by the algorithm described in Section 3.2. A longevity hedge can be constructed readily by writing K-forwards with the notional amounts specified by Eq. (6).

4. Illustrations of a K-Forward Hedge: Single Cohort

In this section, we illustrate a K-forward hedge for a hypothetical pension plan involving only a single birth cohort. The following assumptions about the pension plan and the hedging instruments are made:

[5] The settlement date is T_j, which is generally later than the reference year t_j. As previously mentioned, to simplify our calculations, we assume that the settlement date T_j is the last day of the reference year t_j.

- Assumptions about the hypothetical pension plan:
  1. The pension plan contains one pensioner. It pays the pensioner $1 at the beginning of each year starting from age 65 until the pensioner dies or attains age 91.
  2. There is no basis risk, that is, the pensioner's mortality experience is the same as the experience of the population to which the K-forwards are linked.
  3. All cash flows are discounted at an interest rate of 3%.
- Assumptions about the hedging instruments:
  1. The 3-factor CBD mortality indexes are derived from Model M7* that is fitted to data over the age range of 40-90.
  2. The current year t_0 (i.e. time 0) is taken as the beginning of the year immediately after the end of the data sample period (see Table 3). For instance, in the illustration that is based on data from English and Welsh population, time 0 is set to the beginning of year 2010.
  3. K-forwards with reference years 2015, 2020, 2025 and 2030 are available in the market.
  4. The forward mortality indexes, $\tilde{\kappa}_{t_j}^{(i)}$, are identical to the best estimates of the corresponding mortality indexes. This assumption, which implies a zero risk premium, affects the cost of the hedge but not the performance of hedge.
  5. The best estimates of the future mortality indexes required in Assumption 4 and the calculation of KKD values are obtained by extrapolating the time-varying parameters in Model M7* (fitted to the age range of 40-90 and the entire data sample period) through a Vector Autoregressive Integrated Moving Average (VARIMA) process.[6]

We evaluate the effectiveness of a longevity hedge by the amount of longevity risk reduction, R, defined as

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

where $\sigma^2(.)$ is the variance function,

$$X = L(\boldsymbol{\kappa}) - \mathbb{E}(L(\boldsymbol{\kappa}))$$

is the random present value of unexpected cash flows (the random present value of future cash flows minus the expected present value of future cash flows) when the plan is not hedged, and

$$X^* = L(\boldsymbol{\kappa}) - \mathbb{E}(L(\boldsymbol{\kappa})) - H(\boldsymbol{\kappa}) + \mathbb{E}(H(\boldsymbol{\kappa}))$$

[6]The modeling of the CBD time-varying parameters with VARIMA processes is described by Chan et al. (2014). We also refer interested readers to Tiao and Box (1981) and Wei (2006) for further details about VARIMA modeling.

is the random present value of unexpected cash flows when a hedge is in place. A longevity hedge is considered to be effective if the variability of X^* is significantly smaller than that of X, so a higher value of R implies a better hedging performance.

Model M7* plays two roles in the longevity hedge. First, the mortality indexes to which the instruments are linked are derived from Model M7*. Second, in computing KKDs, the best estimates of the future mortality indexes are derived from a Model M7* that is fitted to historical data. So, a natural question to ask is whether the longevity hedge still works if future mortality does not follow Model M7*. To address this concern, we evaluate hedge effectiveness on the basis of mortality scenarios simulated from five different stochastic mortality models:

- Model M7*: the adapted version of model M7
- Model M5: the original CBD model
- Model M3: the original Age-Period-Cohort model
- Model M2: the original Renshaw-Haberman model
- Model MRW: the multivariate random walk model[7] (Bell, 1997)

If the hedge is effective regardless of what simulation model is used, then a high hedge effectiveness is likely to be achievable in reality. The simulation procedure is summarized as follows:

1. Specify the stochastic processes for the time-varying and cohort effect parameters in the simulation models. For Model M7*, the three time-varying parameters are modeled by a VARIMA process, whereas the cohort effect parameter is modeled by an AR(1) process. For Model M5, the two time-varying parameters are modeled by a VARIMA process. For Models M2 and M3, the time-varying parameter is modeled by a random walk with drift, whereas the cohort effect parameter is modeled by an ARIMA(1,1,0) process.[8] For Model MRW, the evolution of the log death rates at different ages is modeled by a multivariate random walk with drift.
2. For each model, simulate 5,000 realizations of future death probabilities and transform them to realizations of the 3-factor CBD mortality indexes.[9] Parameter uncertainty is incorporated into the simulations via the parametric bootstrap approach (see Brouhns et al., 2005).
3. For each realization, calculate the values of X and X^*. Then for each simulation model, compute the variances of X and X^* and finally the amount of longevity risk reduction

[7] It models the log death rates at different ages directly with a multivariate random walk with drift. No time-varying or age-specific parameters are involved.

[8] The ARIMA orders used for the cohort effect parameters in Models M7*, M2 and M3 are taken from the paper by Cairns et al. (2011b).

[9] The step is necessary, because we need the realized 3-factor CBD mortality indexes to calculate the simulated payoffs from the K-forwards. This step is not needed when the simulation model is Model M7* or when the hedging instruments used are q-forwards.

Table 6
The calculated key K-durations and notional amounts of K-forwards. The pension plan and K-forwards are linked to the mortality of English and Welsh males.

	$j=1$	$j=2$	$j=3$	$j=4$
Key year	2015	2020	2025	2030
$KKD_i(F_j^{(i)}(\boldsymbol{\kappa}), j)$	−0.8375	−0.7224	−0.6232	−0.5375
$KKD_1(L(\boldsymbol{\kappa}), j)$	−0.8373	−0.7540	−0.5825	−0.3155
$KKD_2(L(\boldsymbol{\kappa}), j)$	−5.8541	−9.0194	−9.8330	−6.8644
$KKD_3(L(\boldsymbol{\kappa}), j)$	139.2352	54.5539	−40.2955	−80.3121
$w^{(1)}(j)$	0.9997	1.0437	0.9347	0.5869
$w^{(2)}(j)$	6.9901	12.4850	15.7791	12.7698
$w^{(3)}(j)$	−166.2541	−75.5153	64.6624	151.2644

R.

4.1. The Baseline Results

In this subsection, we present the baseline results for English and Welsh male population.[10] The calculated KKDs and notional amounts are tabulated in Table 6. Note that a positive notional amount means the pension plan participates as a fixed rate receiver, whereas a negative notional amount means the pension plan participates as a fixed rate payer. It can be seen that the signs of the calculated notional amounts agree with the arguments presented in Section 3.1.

The longevity risk reductions calculated from different simulation models are displayed in Table 7, and the density functions for X and X^* simulated from Model M7* are shown in Fig. 4. Using a portfolio of K1-, K2- and K3-forwards, the amounts of longevity risk reduction obtained range from 94.0% to 96.0%. The results depend minimally on the simulation model used, indicating that the hedge effectiveness is achievable even if future mortality does not follow Model M7*. The results also indicate that a high hedge effectiveness can be achieved by using K-forwards associated with four key years only.

Table 7 also shows the amounts of longevity risk reduction resulting from different subsets of the available K-forwards. It can be seen that K1 risk is the most crucial whereas K3 risk has the smallest impact. In particular, when the simulation model is Model M5, the use of K3-forwards seems to have no marginal benefit. This result can be explained using the following equation:

$$X^*_{\text{with K3-forward}} = X^*_{\text{without K3-forward}} - \sum_{j=1}^{4} w^{(3)}(j)(1+r)^{-(T_j-t_0)}(\mathbb{E}(\kappa_{t_j}^{(3)}) - \kappa_{t_j}^{(3)}).$$

Because of the nature of Model M5, which assumes that future mortality is driven entirely by $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$, the values of $\kappa_t^{(3)}$ simulated from Model M5 are practically zero. It follows

[10]The numerical results for the other populations are available on request.

Table 7
The amounts of longevity risk reduction (R) resulting from different combinations of K-forwards and different simulation models. The pension plan and K-forwards are linked to the mortality of English and Welsh males. The VARIMA orders for the time-varying parameters in Models M7* and M5 are (1,1,0) and (5,1,0), respectively.

Simulation model	K1 only	K2 only	K3 only	K1&K2	K1&K3	K2&K3	K1&K2&K3
M7*	68.8%	43.5%	0.5%	93.7%	69.4%	44.4%	94.7%
M5	89.5%	3.3%	0.0%	96.0%	89.5%	3.3%	96.0%
M3	95.6%	4.6%	0.2%	95.7%	95.5%	4.7%	95.6%
M2	93.7%	16.1%	0.0%	95.1%	93.7%	16.2%	95.2%
MRW	77.2%	39.3%	0.6%	93.5%	77.9%	39.9%	94.0%

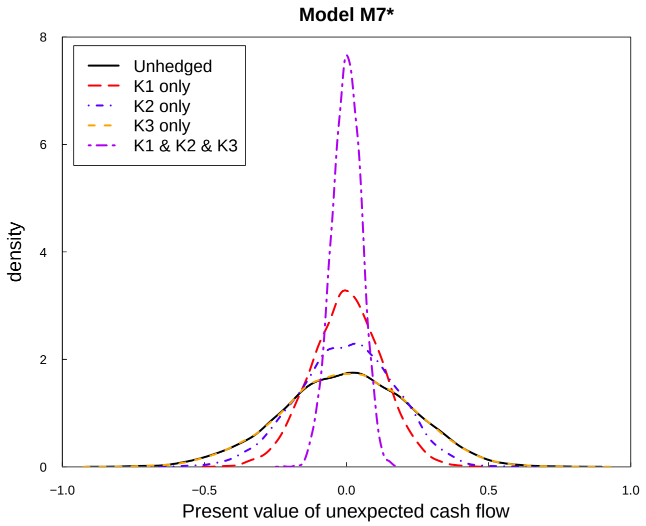

Fig. 4. The distributions of X and X^* simulated from Model M7*. It is assumed that the underlying population is English and Welsh males and that the pension plan consists of a single pensioner currently aged 65.

that the variability in the second term of the right-hand-side of the equation is close to zero. Therefore, $X^*_{\text{with K3-forward}}$ and $X^*_{\text{without K3-forward}}$ have almost the same variability, which implies that the impact of the K3-forwards on the hedge effectiveness is negligible.

In general, the third CBD mortality index may take values that are significantly different from zero.[11] In this case, adding K3-forwards to the hedge portfolio may improve hedge effectiveness. For example, when Model M7* is assumed, adding K3-forwards to a portfolio of K1- and K2-forwards increases the value of R from 93.7% to 94.7%. We also found that the benefit of adding K3-forwards is more apparent for some other populations, including Taiwanese females and U.S. males.

In principle, given a collection of simulated mortality scenarios, we can gradually adjust

[11]For all populations under consideration, Model M7* results in better BIC values compared to Model M6 (see Table 5). This supports the existence of a quadratic age effect, which implies that the third CBD mortality index generally takes non-zero values.

the notional amounts of the K-forwards in the hedge portfolio to compute the highest attainable value of R. The value of R obtained by matching KKDs is generally less than its highest attainable value, because KKDs are only approximate measurements of the sensitivity to the underlying mortality indexes. We found that the maximum achievable R values for a portfolio of K1-, K2- and K3-forwards are 97.3% (for Model M7*), 99.1% (for Model M5), 96.6% (for Model M3), 95.8% (for Model M2) and 96.4% (for Model MRW). These values are very close to those obtained by matching KKDs, providing a strong case for the KKD method. Note that the computation of the highest attainable R value requires heavy computational resources, because it involves simulations and a maximization of a multi-variable function. By contrast, as discussed in Section 3.3, the method of KKD is rather easy to execute.

4.2. Sampling Risk

Sampling risk, or small-sample risk, refers to the risk that the actual mortality experience of a population turns out to be different from the 'true' underlying mortality rates. In this subsection, we study the impact of sampling risk on hedge effectiveness. Again, we use the mortality data from English and Welsh males population to illustrate.

We retain all assumptions specified in Section 4, except that we now assume the pension plan has a cohort of pensioners aged 65 at time 0 instead of only one single pensioner. We let $l(65)$ be the initial number of pensioners, and $l(x)$ be the number of pensioners who survive to age x, where $65 < x \leq 91$. To incorporate sampling risk, we treat the cohort of pensioners as a random survivorship group. Specifically, we model $l(x)$ for $65 < x \leq 91$ by a binomial process:

$$l(x+1) \sim \text{Binomial}(l(x), 1 - q_{x,1945+x}), \qquad x = 65, 66, \ldots, 90.$$

The procedure for simulating liability cash flows is adapted as follows:

1. Simulate a future mortality curve from the assumed mortality model.
2. Given the simulated mortality curve, simulate the number of survivors $l(x)$, for $x = 66, 67, \ldots, 91$, from the binomial process.
3. Calculate the liability cash flows based on the simulated values of $l(x)$.
4. Repeat the steps above to obtain 5,000 scenarios of liability cash flows.

The incorporation of sampling risk does not affect the method by which the KKDs are calculated. Therefore, in this illustration, the values of $KKD_i(F_j^{(i)}(\boldsymbol{\kappa}), j)$ are the same as those in Table 6, whereas the values of $KKD_i(L(\boldsymbol{\kappa}), j)$ are equal to $l(65)$ multiplied by the corresponding values in Table 6 (because the initial number of pensioners is $l(65)$ instead of one). It follows that the required notional amounts are simply $l(65)$ multiplied by the notional amounts shown in Table 6.

The amounts of longevity risk reduction for different values of $l(65)$ are shown in Table 8. The row with $l(65) = \infty$ shows the values of R when there is no sampling risk. It can be seen that the hedge effectiveness drops as $l(65)$ decreases. When $l(65) = 1000$, the value of

Table 8
The amounts of risk reduction (R) for pension plans of different initial sizes. The pension plans and K-forwards are linked to the mortality of English and Welsh males. The simulation model is Model M7* with a VARIMA(1,1,0) process for the time-varying parameters.

$l(65)$	K1 only	K2 only	K3 only	K1&K2	K1&K3	K2&K3	K1&K2&K3
∞	68.8%	43.5%	0.5%	93.7%	69.4%	44.4%	94.7%
10000	65.6%	41.3%	0.6%	89.3%	66.3%	42.3%	90.3%
5000	63.0%	39.8%	0.5%	85.8%	63.6%	40.6%	86.8%
3000	59.6%	37.3%	0.6%	80.8%	60.2%	38.2%	81.9%
1000	46.5%	29.0%	0.4%	63.0%	46.9%	29.6%	63.7%

Table 9
The amounts of risk reduction (R) under different interest rate assumptions. The pension plan and K-forwards are linked to the mortality of English and Welsh males. The simulation model is Model M7* with a VARIMA(1,1,0) process for the time-varying parameters.

Interest Rate	$\sigma(X)$	K1 only	K2 only	K3 only	K1&K2	K1&K3	K2&K3	K1&K2&K3
1%	0.325	68.7%	45.0%	0.3%	94.3%	69.0%	45.8%	95.2%
2%	0.269	68.7%	44.3%	0.4%	94.0%	69.2%	45.1%	95.0%
3%	0.223	68.8%	43.5%	0.5%	93.7%	69.4%	44.4%	94.7%
4%	0.187	68.8%	42.8%	0.6%	93.3%	69.6%	43.7%	94.4%
5%	0.157	68.7%	41.9%	0.8%	92.8%	69.8%	42.9%	94.1%

R resulting from the hedge built from all available K-forwards is only 63.7%. This finding suggests that a small pension plan should consider alternative de-risking solutions, such as pension buy-ins (see Coughlan et al., 2013), instead of using standardized K-forwards.

4.3. Sensitivity Tests

In this subsection, we sensitivity test several assumptions used in estimating hedge effectiveness.

4.3.1. Interest Rate

In the previous illustrations, an interest rate of 3% per annum was used. We now examine the relationship between the hedging results and the assumed interest rate. The results are summarized in Table 9. A higher interest rate produces a lower standard deviation of X, because the more distant and so more uncertain cash flows are discounted by a larger extent. For the same reason, the required notional amounts are lower in magnitude when the assumed interest rate is higher. Nevertheless, the amounts of risk reduction under different interest rate assumptions are similar, indicating that the effectiveness of a K-forward hedge is not very sensitive to the interest rate assumption.

4.3.2. Availability of K-forwards

In the KKD framework, the key years are the reference years for which K-forwards are available. It is assumed in the baseline calculations that the separation between two adjacent

Table 10
The amounts of risk reduction (R) resulting from different spacings between two adjacent key years. The pension plan and K-forwards are linked to the mortality of English and Welsh males. The simulation model is Model M7* with a VARIMA(1,1,0) process for the time-varying parameters.

Separation between two adjacent key years (total number of key years)	K1 only	K2 only	K3 only	K1&K2	K1&K3	K2&K3	K1&K2&K3
8 (3)	53.3%	32.5%	0.5%	77.5%	53.9%	33.2%	78.3%
7 (3)	59.1%	36.6%	0.4%	84.5%	59.6%	37.3%	85.3%
6 (4)	62.8%	39.2%	0.5%	88.5%	63.5%	40.0%	89.4%
5 (4)	68.8%	43.5%	0.5%	93.7%	69.4%	44.4%	94.7%
5 (3)	67.6%	42.4%	0.3%	91.5%	68.2%	43.0%	92.3%
5 (2)	61.1%	36.7%	0.1%	81.4%	61.5%	36.8%	81.7%

key (reference) years is 5 years. Here we examine how the hedging results may change when different separations are assumed. Throughout the analysis, the earliest key year is fixed to 2015. The hedging results are tabulated in Table 10. It can be seen that no matter 3 or 4 key years are used, a smaller separation produces a (slightly) higher amount of risk reduction. This result may possibly be explained by the fact that closer key years (from 2015) yield better hedging results for the earlier random cash flows but worse hedging results for the later ones, which are discounted more heavily and hence less important to the overall hedging performance. It can also be seen that, as expected, the amount of risk reduction increases with the number of key years (i.e., the number of instruments used).

4.3.3. Age Range

In our baseline calculations, we use an age range of 40-90 to define the 3-factor CBD mortality indexes. We now study how the age range over which the indexes are defined may affect the hedging results. In Table 11 we compare the baseline hedging results with the hedging results that are based on two alternative age ranges, 50-90 and 60-90. It can be observed that the choice of age range does not have a significant impact on the amount of risk reduction produced by the hedge with all three types of K-forwards. However, the R values for other combinations of K-forwards vary significantly with the age range used. In particular, the performances of the K1-forward only and K3-forward only hedges improve with the mean age in the age range, but the reverse is true for the K2-forward only hedge. This pattern suggests that if the 3-factor CBD mortality indexes are defined over an older age range, the first and third CBD indexes would capture a larger proportion of the overall longevity risk, while the second CBD index would capture less.

4.4. Advanced Ages

Here, we examine how a K-forward hedge may perform if we extend the coverage of the hypothetical pension plan to age 101. While the age coverage of the pension plan is extended, the definition of the 3-factor CBD mortality indexes remains unchanged; that is, the indexes are still derived by fitting Model M7* to data over the age range of 40-90. We use

Table 11
The amounts of risk reduction (R) for different age ranges over which the 3-factor CBD mortality indexes are defined. The pension plan and K-forwards are linked to the mortality of English and Welsh males. The simulation model is Model M7* (fitted to the corresponding age range) with a VARIMA(1,1,0) process for the time-varying parameters.

Age range	$\bar{x}$	K1 only	K2 only	K3 only	K1&K2	K1&K3	K2&K3	K1&K2&K3
40-90	65	68.8%	43.5%	0.5%	93.7%	69.4%	44.4%	94.7%
50-90	70	79.8%	22.3%	3.8%	91.1%	81.1%	27.8%	94.1%
60-90	75	87.9%	7.6%	5.4%	90.2%	89.9%	14.4%	93.5%

the following procedure to calculate the amount of risk reduction produced by the K-forward hedge for the extended pension plan.

1. Calibrate the K-forward hedge as follows:
   (a) Compute the best estimates of future mortality indexes by extrapolating from a VARIMA process that is fitted to the historical index values.
   (b) Calculate the best estimates of the death probabilities over the age range of 65-101 using Eq. (1) with $\bar{x} = 65$ and $\hat{\sigma}_x^2 = \frac{650}{3}$.
   (c) On the basis of the values from Steps (a) and (b), calculate $P(\boldsymbol{\kappa})$, $P(\breve{\boldsymbol{\kappa}}(\delta^{(i)}(j)))$ and $KKD_i(P(\boldsymbol{\kappa}), j)$ for $i = 1, 2, 3$ and $j = 1, \ldots, n$.
   (d) Using the KKD values computed in Step (c), compute the notional amounts of the K-forwards in the hedge portfolio.
2. Simulate 5000 realizations of future death probabilities from a stochastic mortality model that is fitted to historical data over the age range of 40-100. As before, the simulation models considered are Models M7*, M5, M3, M2 and MRW.
3. Calculate the simulated present values of pension payments from age 65 to 101.
4. Transform the simulated future death probabilities over the age range of 40-90 to realizations of the 3-factor CBD mortality indexes, which are then used to determine the simulated payoffs from the K-forwards.
5. Compute X, X^* and finally the amount of risk reduction R.

We use the same key years despite the coverage of this pension plan is 10 years longer. With this setup, the amounts of longevity risk reduction produced by a calibrated portfolio of K1-, K2- and K3-forwards are 93.9% (for Model M7*), 93.4% (for Model M5), 94.3% (for Model M3), 95.1% (for Model M2) and 93.9% (for Model MRW). These findings suggest that the KKD strategy still produces a satisfactory hedge effectiveness when the pension plan is extended to include more advanced ages.

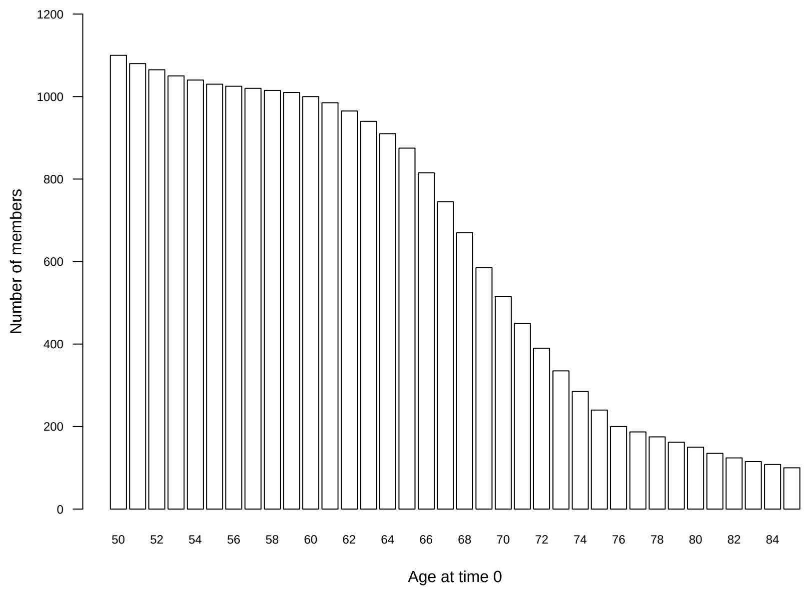

Fig. 5. The initial demographic structure of the multi-cohort closed pension plan.

5. Generalization to Multiple Cohorts

In real life, individuals in a pension plan are generally born in different years. For a multi-cohort pension plan, longevity risk is associated with a two-dimensional mortality surface, composed by a collection of mortality curves for various birth cohorts. The illustrations in the previous section may be viewed as a special case, whereby the longevity risk exposure arises from a particular diagonal over the mortality surface.

In this section, we use a more realistic hypothetical pension plan to illustrate how the proposed method can be applied to situations involving multiple birth cohorts. We retain all assumptions stated in Section 4, except those about the pensionable age and the plan's demographics. Here, we assume that individuals in the plan are currently aged 50 to 85, following the age distribution displayed in Fig. 5. The pensionable age is 60, so the individuals who are currently aged 50 to 59 are now active members (who will receive pension payments when they reach age 60) while those who are currently aged 60 or above are now retirement pensioners (who are currently receiving pension payments). It is further assumed that the plan is closed to new entrants and that all active members remain in their jobs until retirement.

We consider two hedging scenarios:

1. A hedge for the current retirement pensioners only

   This hedge aims to reduce the variability of the present values of the pension payments to the current retirement pensioners in the next 31 years.

2. A hedge for both current active members and retirement pensioners

This hedge aims to reduce the variability of the present values of the pension payments to the current retirement pensioners and the deferred pension payments to the current active members in the next 31 years.

As the market has started to see static longevity hedges for active members, the results for the second hedging scenario may be of many practitioners' interest.[12]

Another objective of this section is to compare the K-forward longevity hedges with the corresponding longevity hedges that are composed of q-forwards, a standardized instrument that has been promoted by the Life and Longevity Markets Association (LLMA). We calibrate q-forward hedges for the aforementioned two hedging scenarios using the method of key q-duration (KQD), proposed by Li and Luo (2012).

As before, all results in this section are computed using mortality data from English and Welsh males population.

5.1. Building K-Forward Hedges

In this illustration, we assume that K-forwards with reference years 2015, 2020, 2025, 2030 and 2035 are available. An additional reference year (2035) is considered, because the hedging horizon here is longer than that for the single-cohort pension plan, which makes at most 26 years of payments (from age 65 to 91). The KKD with respect to each index and reference year can be estimated by using the algorithm described in Section 3.2. Given the KKD values, the required notional amounts (shown in Table 12) can be computed accordingly.

It can be observed that the signs of the notional amounts are completely in line with the arguments presented in Section 3.1. It is interesting to note that the pension plan's positions in the K3-forward with reference year 2025 are different in the two hedging scenarios. For the hedge involving current retirement pensioners only, the plan's position in that K3-forward should be a fixed rate receiver, because in 2025 the individuals associated with the hedge will attain ages 75-90, such that the total pension payments at that time will be negatively related to the index value to which that K3-forward is linked.

It is also interesting to compare the required notional amounts of K1- and K2-forwards in the two hedging scenarios. Since more risk is being hedged, the notional amounts of K1- and K2-forwards in the larger hedge covering both current active members and retirement pensioners are generally higher. The only exception is the K2-forward with reference year 2015. This is because from 2015 to 2019, all individuals associated with the smaller hedge (covering the current retirement pensioners only) will be at least 65 years old, but some individuals associated with the larger hedge will still be younger than 65. As explained in Section 2.4, pension payments below and above age $\bar{x} = 65$ respond to the second CBD mortality index in opposite directions. The offsetting effects thus reduce the amount of K2 risk that the larger hedge has to mitigate.

[12] In January 2011, the Pall (UK) pension fund executed a £70 million deal with J.P. Morgan to hedge the longevity risk associated with its active and deferred members (see Blake et al., 2014).

Table 12
The required notional amounts of K-forwards, calibrated by KKD. The pension plan and K-forwards are linked to the mortality of English and Welsh males.

Reference year t_j	Type	Notional amount	
		Retirement pensioners only	Active members and retirement pensioners
2015	K1	11821	16887
	K2	110873	100389
	K3	−1174133	−2220565
2020	K1	10736	17719
	K2	137888	151606
	K3	−357008	−1775855
2025	K1	8495	16116
	K2	138762	191452
	K3	518404	−696304
2030	K1	5058	12385
	K2	99205	185632
	K3	874180	373324
2035	K1	1588	6730
	K2	35931	121384
	K3	469823	822578

Table 13
Locations of the key mortality rates of the chosen key cohorts.

Key cohort k	Year of birth c_k	Number of key mortality rates n_k	Location of key mortality rates
1	1930	1	Age 85
2	1935	2	Ages 80, 85
3	1940	3	Ages 75, 80, 85
4	1945	4	Ages 70, 75, 80, 85
5	1950	5	Ages 65, 70, 75, 80, 85
6	1955	5	Ages 65, 70, 75, 80, 85
7	1960	4	Ages 65, 70, 75, 80

5.2. Building q-Forward Hedges

We now use the KQD method to develop the corresponding q-forward hedges. The first step in the KQD method is to identify the key cohorts in the pension portfolio. The pension plan involves 36 birth cohorts in total, with the oldest born in year 1925 and the youngest in year 1960. In the smaller hedge for the current retirement pensioners only, we consider five key cohorts, which were born in 1930, 1935, 1940, 1945 and 1950. In the larger hedge that covers additionally the current active members, we use two extra key cohorts, which were born in 1955 and 1960.

The next step is to identify the key mortality rates in each key cohort. Following Li and Luo (2012), the key mortality rates are chosen in such a way that they are no more than five ages apart (see Table 13). It is assumed that the q-forwards linked to the chosen key mortality rates are available in the market.

We then use the q-forwards that are linked to the key mortality rates to build longevity hedges for the two hedging scenarios. The reference rate of the q-forward that corresponds to

Table 14
The required notional amounts of q-forwards, calibrated by KQD. The pension plan and K-forwards are linked to the mortality of English and Welsh males.

Key cohort k	Reference age $x_{l,k}$	Reference year $t_{l,k}$	Notional amount	
			Retirement pensioners only	Active members and retirement pensioners
1	85	2015	35107	35107
2	80	2015	75556	75556
	85	2020	19215	19215
3	75	2015	210216	210216
	80	2020	64683	64683
	85	2025	36228	36228
4	70	2015	418885	418885
	75	2020	162233	162233
	80	2025	102002	102002
	85	2030	58076	58076
5	65	2015	283558	500145
	70	2020	159431	255112
	75	2025	115337	183150
	80	2030	73061	113965
	85	2035	41998	60440
6	65	2020		564347
	70	2025		242804
	75	2030		163699
	80	2035		85929
	85	2040		23716
7	65	2025		341518
	70	2030		138863
	75	2035		83317
	80	2040		30373

the lth key mortality rate of the kth key cohort (which was born in year c_k) is $q_{x_{l,k},t_{l,k}}$, where $x_{l,k}$ and $t_{l,k} = c_k + x_{l,k}$ are the q-forward's reference age and reference year, respectively. For instance, the reference rates of the five q-forwards associated with the key cohort that was born in 1950 are $q_{65,2015}, q_{70,2020}, q_{75,2025}, q_{80,2030}$ and $q_{85,2035}$. Note that the smaller hedge consists of 15 q-forwards, while the larger hedge contains 9 more q-forwards, as two extra key cohorts are introduced.

The notional amounts are determined by matching the KQD values of the pension plan and the portfolio of q-forwards. The KQD values are estimated using the procedure detailed in Section 5 of Li and Luo (2012). The required notional amounts of q-forwards for each hedging scenario are tabulated in Table 14. As a consequence of how the KQD values are estimated, the required notional amounts of the 10 q-forwards corresponding to the first four key cohorts are identical in both hedging scenarios. Also, as expected, the pension plan acts as a fixed rate receiver in all q-forwards, because the pension payments are negatively associated with the realized mortality rates.

5.3. Hedging Results

The hedging results are reported in Table 15 and illustrated graphically in Fig. 6. For the hedging scenario involving the current retirement pensioners only, the amounts of risk

Table 15
The amounts of risk reduction (R) resulting from the K-forward and q-forward hedges for the hypothetical multi-cohort pension plan. The pension plan and hedging instruments are linked to the mortality of English and Welsh males. The simulation model is Model M7* with a VARIMA(1,1,0) process for the time-varying parameters.

Hedging scenarios	K1-, K2- and K3-forwards	q-forwards
Retirement pensioners only	94.5%	91.6%
Active members and retirement pensioners	96.0%	95.5%

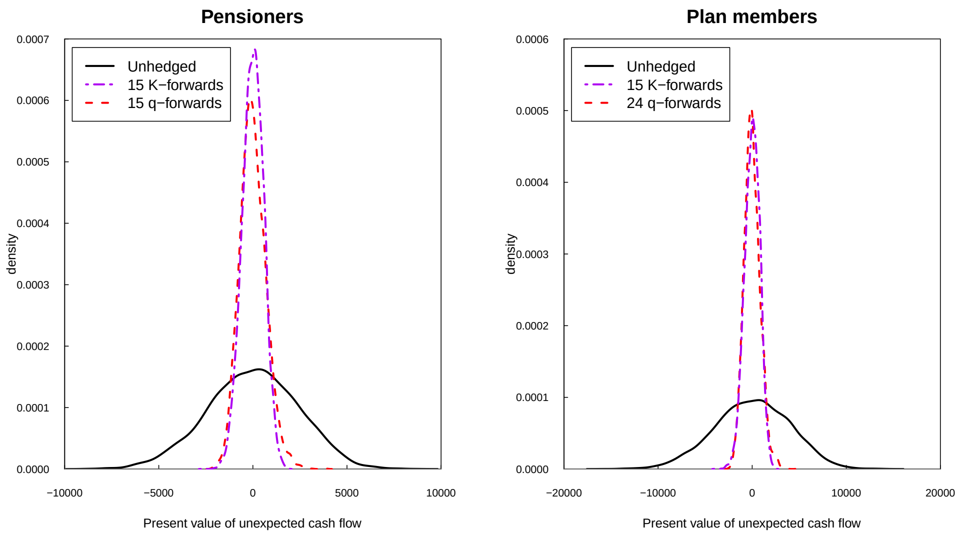

Fig. 6. The distributions of X and X^* for the hedges that are constructed for both current active members and retirement pensioners in the hypothetical multi-cohort pension plan. The pension plan and hedging instruments are linked to the mortality of English and Welsh males. The simulation model is Model M7* with a VARIMA(1,1,0) process for the time-varying parameters.

reduction produced by the K-forward hedge and the q-forward hedge are 94.5% and 91.6%, respectively. For the other scenario that includes additionally the current active members, the amounts of risk reduction resulting from the K-forward hedge and the q-forward hedge are 96.0% and 95.5%, respectively. The simulation results indicate that both types of instruments can yield highly satisfactory hedging results, even if the portfolio being hedged covers multiple birth cohorts. They also indicate that a K-forward hedge may yield a higher amount of risk reduction compared to a q-forward hedge.

What distinguishes the two instruments is the number of securities needed to achieve a given level of hedge effectiveness. For the hedging scenario involving both current retirement pensioners and active members, although both hedges lead to similar R values, the q-forward hedge (calibrated by the KQD method) requires 9 more securities in comparison to the K-forward hedge (calibrated by the KKD method). It follows that when the number of cohorts covered is large, a K-forward hedge is more compact and thus more preferred.

Technically speaking, the KQD measure relies heavily on certain assumptions about how a shift in a key mortality rate affects the mortality rates in the neighbouring cohorts. Li and Luo (2012) showed empirically that these assumptions are reasonably accurate when the separation between two key cohorts is small, but are less likely to hold if the key cohorts are farther apart. Therefore, in implementing the KQD method, additional key cohorts and hence q-forward contracts have to be introduced when the portfolio being hedged covers more cohorts of individuals. Otherwise, the estimated KQD values may not accurately reflect the portfolio's sensitivity to changes in the underlying mortality curve, thereby leading to a less satisfactory hedging result.

By contrast, the KKD value for a particular CBD mortality index depends only on time, having no relationship with the individuals' years of birth. It follows that there is no need to alter the choices of the key K-indexes when more cohorts of individuals are included in the portfolio being hedged. In our illustration, if the pension plan covers even more birth cohorts (say those who are aged 40-49 at time 0), the number of key K-indexes and hence the number of K-forwards will still be 15. To hedge the risk associated with these additional cohorts, all the hedger needs to do is to recalculate the notional amounts.

The benefit of K-forwards to market development has now become clear. With a finite number of K-forwards, the hedging needs of pension plans with different demographic structures can be met. However, the collection of q-forwards that are applicable to one pension plan may not be suitable and/or sufficient for another pension plan that has a different age profile. Hence, relative to q-forwards, K-forwards are easier to attract demand from hedgers and are potentially more liquid.

6. Concluding Remarks

In this paper, we studied the construction of mortality indexes using stochastic mortality models. We recommend using the three time-varying parameters in Model M7* (i.e., Model M7 that is adapted to satisfy the new-data-invariant property) to construct mortality indexes. One basis of our recommendation is that among the collection of six appropriately adapted models, Model M7* provides the best goodness-of-fit to the majority of the data sets under consideration. Another basis is that the three time-varying parameters, which we call the 3-factor CBD mortality indexes, are easy to interpret and are able to reflect the varying age pattern of mortality.

We also investigated the potential applications of the 3-factor CBD mortality indexes to securitization. In particular, we further developed a standardized security called K-forward, the payoff of which is linked to a 3-factor CBD mortality index. We also explained how and why a portfolio of K-forward contracts can be used to hedge the longevity risk associated with a life-contingent portfolio. In terms of structure, the proposed K-forwards are simpler than the existing q-forwards, because each K-forward contract is characterized by one single parameter (the reference year) only.

The proposed KKD measure permits one to calibrate a longevity hedge that is formed by K-forward contracts. From the results of the simulation studies we conducted, the following conclusions can be drawn. First, a calibrated K-forward hedge can substantially reduce the variability in the values of a pension portfolio, even if parameter uncertainty and sampling risk are taken into account. Second, the effectiveness of a K-forward hedge does not vary significantly with the assumed simulation model, indicating that the success in hedging is likely to be achievable in real life. Finally, compared to a q-forward hedge, a K-forward hedge giving a comparable hedge effectiveness requires a smaller number of securities, thereby helping the market to concentrate liquidity.

The proposed KKD method may be improved in future research in a number of ways. For instance, similar to the recent contributions by Lin and Tsai (2013), Tsai et al. (2010) and Wang et al. (2010), a convexity component can be added to improve the approximation of a portfolio's sensitivity to the 3-factor CBD mortality indexes. Also, following the footsteps of Cairns (2011), it would be interesting to extend the KKD hedging framework from static to dynamic, permitting hedgers to periodically adjust their K-forward portfolios. Periodic adjustments allow a K-forward portfolio to reflect the changes in KKD values over time due to, for example, fluctuations in expected interest and mortality rates, thereby possibly leading to even better hedging results.

One important issue that is not addressed in this paper is the pricing of K-forwards. Our illustrations are based on the assumption that the K-forwards are costless, so they provide no information about how a K-forward hedge may cost the hedger. It is warranted in future research to study how a K-forward contract may be priced by no-arbitrage approaches (see, e.g., Cairns et al., 2006; Denuit et al., 2007; Dowd et al., 2006; Li et al., 2011) or economic methods (see, e.g., Zhou et al., 2011, 2013a,b). Given that K-forwards have a very simple structure, it may be possible to derive close-form pricing formulas for them.

Another concerning issue in using a standardized K-forward hedge is population basis risk, which arises from the differences between the mortality experience of the hedger's population and the population to which the K-forwards are linked. Recently, a number of multi-population stochastic mortality models have been developed by researchers including Cairns et al. (2011a), Dowd et al. (2011), Jarner and Kryger (2011), Li and Lee (2005), Yang and Wang (2013) and Zhou et al. (2013b, 2014) to address the issue of population basis risk. In addition, Coughlan et al. (2011) and Li and Hardy (2011) proposed frameworks for quantifying population basis risk. Further research is required to find out how population basis risk can be incorporated in the KKD hedging strategy and how it affects the effectiveness of a K-forward hedge.

Acknowledgements

The first author gratefully acknowledges financial support from Aon Benfield. The third author acknowledges the financial support from the Global Risk Institute, the Center of Actuarial Excellence Program of the Society of Actuaries and the Natural Science and Engineering Research Council of Canada.

References

Bell, W., 1997. Comparing and assessing time series methods for forecasting fertility and mortality rates. Journal of Official Statistics 13, 279-303.

Blake, D., Cairns, A., Coughlan, G., Dowd, K., MacMinn, R. (2013). The new life market. Journal of Risk and Insurance 80, 501-557.

Blake, D., MacMinn, R. Li, J.S.-H., Hardy, M.R. (2014). Longevity risk and capital markets: The 2012-13 update. North American Actuarial Journal, in press.

Brouhns, N., Denuit, M., Keilegom, I.V., 2005. Bootstrapping the Poisson log-bilinear model for mortality forecasting. Scandinavian Actuarial Journal 3, 212-224.

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., Dowd, K., 2006. A two-factor model for stochastic mortality with parametric uncertainty: Theory and calibration. Journal of Risk and Insurance 73, 687-718.

Cairns, A.J.G., Blake, D., Dowd, K., 2008. Modelling and management of mortality risk: a review. Scandinavian Actuarial Journal 2-3, 79-113.

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A., Balevich, I., 2009. A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. North American Actuarial Journal 13, 1-35.

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Khalaf-Allah, M., 2011a. Bayesian stochastic mortality modelling for two populations. ASTIN Bulletin 41, 29-55.

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Khalaf-Allah, M., 2011b. Mortality density forecasts: An analysis of six stochastic mortality models. Insurance: Mathematics and Economics 48, 355-367.

Chan, W.S., Li, J.S.-H., Li, J., 2014. The CBD mortality indexes: Modeling and applications. North American Actuarial Journal (forthcoming).

Coughlan, G.D., 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., Blake, D., MacMinn, R., Cairns, A.J.G., 2013. Longevity Risk and Hedging Solutions. In Handbook of Insurance (2nd Edition), ed. Dionne, G. New York: Springer.

Coughlan, G.D., 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.

Coughlan, G.D., Khalaf-Allah, M., Ye, Y., Kumar, S., Cairns, A.J.G., Blake, D., Dowd, K., 2011. Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. North American Actuarial Journal 15, 114-146.

Denuit, M., Devolder, P., Goderniaux, A.-C., 2007. Securitization of longevity risk: Pricing survivor bonds with Wang transform in the Lee-Carter framework. Journal of Risk and Insurance 74, 87-113.

Dowd, K., Blake, D., Cairns, A.J.G., Dawson, P., 2006. Survivor swaps. Journal of Risk and Insurance 73, 1-17.

Dowd, K., Cairns, A.J.G., Blake, D., Coughlan, G.D., Epstein, D., Khalaf-Allah, M., 2010. Backtesting stochastic mortality models: An ex-post evaluation of multi-period-ahead density forecasts. North American Actuarial Journal 14, 281-298.

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

Human Mortality Database. University of California, Berkeley (USA), and Max-Planck Institute of Demographic Research (Germany). Available at: www.mortality.org (data downloaded on 24 February 2013).

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

Lee, R.D., Carter, L., 1992. Modeling and forecasting the time series of U.S. mortality. Journal of the American Statistical Association 87, 659-671.

Li, N., 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., Hardy, M.R., 2011. Measuring basis risk in longevity hedges. North American Actuarial Journal 15, 177-200.

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

Lin, T., Tsai, C.C.-L., 2013. On the mortality/longevity risk hedging with mortality immunization. Insurance: Mathematics and Economics 53, 580-596.

Li, J.S.-H., Ng, A.C.Y., Chan, W.S., 2011. On the calibration of mortality forward curves. Journal of Futures Markets 31, 941-970.

Plat, R., 2009. Stochastic portfolio specific mortality and the quantification of mortality basis risk. Insurance: Mathematics and Economics 45, 123-132.

Plat, R., 2010. One-year Value-at-Risk for longevity and mortality. Insurance: Mathematics and Economics 49, 462-470.

Tiao, G.C., Box, G.E.P., 1981. Modelling multiple time series with applications. Journal of the American Statistical Association 76, 802-816.

Tsai, J.T., Wang, J.L., Tzeng, L.Y., 2010. On the optimal product mix in life insurance companies using conditional value at risk. Insurance: Mathematics and Economics 46, 235-241.

Wang, J.L., Huang, H.C., Yang, S.S., Tsai, J.T., 2010. An optimal product mix for hedging longevity risk in life insurance companies: the immunization theory approach. Journal of Risk and Insurance 77, 473-497.

Wei, W.W.S., 2006. Time series analysis: Univariate and multivariate Method. Second edition. Addison Wesley/Pearson.

Yang, S.S., Wang, C.W., 2013. Pricing and securitization of multi-country longevity risk with mortality dependence. Insurance: Mathematics and Economics 52, 157-169.

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

Zhou, R., Li, J.S.-H., Tan, K.S., 2013a. Economic pricing of mortality-linked securities: A Tâtonnement approach. Journal of Risk and Insurance, doi: 10.1111/j.1539-6975.2013.12008.x.

Zhou, R., Li, J.S.-H., Tan, K.S., 2013b. Pricing standardized mortality securitizations: a two-population model with transitory jump effects. Journal of Risk and Insurance 80, 733-774.

Zhou, R., Wang, Y., Kaufhold, K., Li, J.S.-H., Tan, K.S., 2014. Modeling period effects in multi-population mortality models: Applications to Solvency II. North American Actuarial Journal, in press.