



Pensions
Institute

DISCUSSION PAPER PI-1105

Modelling and Management of Longevity Risk:
Approximations to Survivor Functions and
Dynamic Hedging

Andrew J.G. Cairns

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/>

Modelling and management of longevity risk: approximations to survivor functions and dynamic hedging

Andrew J.G. Cairns
Maxwell Institute, Edinburgh, and
Department of Actuarial Mathematics and Statistics
Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.
WWW: <http://www.ma.hw.ac.uk/~andrewc>
E-mail: A.Cairns@ma.hw.ac.uk

This version: July 4, 2011

Abstract

This paper looks at the development of dynamic hedging strategies for typical pension plan liabilities using longevity-linked hedging instruments. Progress in this area has been hindered by the lack of closed-form formulas for the valuation of mortality-linked liabilities and assets, and the consequent requirement for simulations within simulations. We propose use of the probit function along with a Taylor expansion to approximate longevity-contingent values. This makes it possible to develop and implement computationally efficient, discrete-time Delta hedging strategies using q -forwards as hedging instruments.

The methods are tested using the model proposed by Cairns, Blake and Dowd (2006a) (CBD). We find that the probit approximations are generally very accurate, and that the discrete-time hedging strategy is very effective at reducing risk.

Keywords: longevity risk, dynamic hedging, Delta hedging, probit-Taylor approximation, CBD model, q -forward, Solvency II.

Accepted for publication in *Insurance: Mathematics and Economics*

Suggested citation:

Cairns, A.J.G. (2011) Modelling and management of longevity risk: approximations to survivor functions and dynamic hedging. To appear in *Insurance: Mathematics and Economics*.

1 Introduction

This paper considers the question of how to hedge a portfolio of pension liabilities where cashflows are exposed to longevity risk: that is, contingent on the development of uncertain aggregate mortality rates over a period of years. In many countries, pension plans are increasingly opting to hedge their pension liabilities. This might be done by simply paying a premium to an insurer in order to transfer their pension liabilities. Alternatively, the pension plan can implement a comprehensive programme of asset-liability management. It has been possible for many years to manage interest rate risk through the use, for example, of interest-rate swaps. In contrast, it has only recently become possible to hedge the plan's exposure to longevity risk through the use of customised longevity swaps and index-based hedging instruments such as q -forwards (see, for example, Blake et al., 2006, Dahl et al., 2008, Cairns et al., 2008, 2011b, Li and Hardy, 2011, Dowd et al., 2011b).

With the exception of Dahl et al., these previous studies have focused on the assessment of static hedging strategies using longevity-linked hedging instruments. In part, this reflects the realities of a market that is in its infancy. However, it also reflects the fact that for realistic, discrete-time mortality models, it is difficult to model how the values of longevity-contingent contracts evolve over time without resorting to the use of simulations within simulations: a difficulty that has hindered the development of dynamic hedging strategies. Theoretically, this can be avoided through the use of market models (see, for example, Cairns et al., 2006b, 2008, Cairns, 2007, and Zhu and Bauer, 2010) but such models are less realistic, at present, than the more popular class of "short-rate" models such as those considered in Cairns et al. (2009). However, if technical difficulties can be overcome and more realistic market models developed, the market-model approach provides a good setting for the development of hedging strategies.

The groundbreaking paper by Dahl et al. (2008) is the exception to this. By using a simplified, continuous-time stochastic mortality model, they are able to derive closed-form solutions for longevity-contingent contracts and derive delta-hedging strategies. Additionally, this is achieved within the context of a two-population model, so the authors are able to assess the effectiveness of a dynamic hedging strategy in the presence of population basis risk.

Here, we attempt – at least in some regards – to make Dahl et al.'s analysis more realistic by proposing an approach that can be applied to a wide range of more realistic, discrete-time models, and with a requirement only for annual rather than continuous rebalancing of the hedge portfolio. In other regards we are less ambitious: for example, by leaving an analysis of the impact of dynamic hedging in the presence of population basis risk for future work.

1.1 Connection to Solvency II

The impending introduction of Solvency II allows for the use of internal models that include stochastic modelling of mortality rates (see, for example, Olivieri and Pitacco, 2009; Plat, 2010; and Börger, 2010). A key use of an internal model is the calculation of solvency capital requirements (SCR). The methodology described in this paper can be used in two ways. First, the avoidance of simulations within simulations will make calculation of the contribution of mortality risk to the SCR feasible. Second, the annual hedging methodology in the paper can be used to manage the magnitude of the SCR.

1.2 Outline of paper

In Section 2 we outline the Cairns, Blake and Dowd (2006a) (CBD) model and calibration as well as other notation that will be used throughout the paper. In Sections 3 to 5 we present the analytical functions that we propose as approximations for the valuation of q -forward contracts and annuity-type liabilities and provide numerical examples to demonstrate their accuracy. In Section 6 we then use these to derive approximations for the so-called “deltas” of the various liabilities and hedging instruments. These are then deployed in a numerical experiment that we report on in Section 7. Finally, in Section 8, we provide some further discussion of possible extensions of this work.

2 Model and notation

In order to illustrate both the approximations and the Delta-hedging strategy proposed in this paper, we will use the CBD 2-factor model (Cairns et al., 2006a) as an example. We define $q(t, x)$ to be the probability, as measured at time $t + 1$, that an individual aged x at time t survives to time $t + 1$. Under the CBD model we have

$$q(t, x) = \frac{\exp [K_1(t + 1) + K_2(t + 1)(x - \bar{x})]}{1 + \exp [K_1(t + 1) + K_2(t + 1)(x - \bar{x})]}$$

(or $\text{logit } q(t, x) = K_1(t + 1) + K_2(t + 1)(x - \bar{x})$) where $K(t) = (K_1(t), K_2(t))'$ is a two-dimensional random walk with drift:

$$K(t + 1) = K(t) + \nu + CZ(t + 1).$$

$Z(t + 1) = (Z_1(t + 1), Z_2(t + 1))'$ is a standard bivariate normal random vector under the real-world probability measure \mathbb{P} . Using England and Wales males data, ages 60 to 89 from 1981 to 2008, and the methodology described in Cairns et al. (2009)

estimate that (taking the end of 2008 as time $t = 0$) $K(0) = (-3.2717, 0.1079)'$,

$$\nu = \begin{pmatrix} -0.02534 \\ 0.0004604 \end{pmatrix}, \text{ and } V = CC' = \begin{pmatrix} 0.0004538 & 0.00001585 \\ 0.00001585 & 0.000001256 \end{pmatrix}. \quad (1)$$

The survivor index is defined as

$$S(T, x) = (1 - q(0, x)) \times (1 - q(1, x)) \times \dots \times (1 - q(T - 1, x)), \quad (2)$$

and represents the *ex post* probability that an individual aged x at time 0 would have survived to time T .

A number of survivor and financial functions are non-linear functions of the future path of $K(t)$ that can only be computed using simulation. This includes spot survival probabilities (see, for example, Cairns et al., 2006b), $p(\tau, \tau + T, x - \tau, k)$: the probability that an individual aged $x - \tau$ at time 0, and still alive at time τ (age x), survives until time $\tau + T$, based on the information about aggregate mortality at time τ , as summarised by the vector $K(\tau) = k$. We have

$$p(\tau, \tau + T, x - \tau, k) = E_{\mathbb{P}} \left[\frac{S(\tau + T, x - \tau)}{S(\tau, x - \tau)} \mid K(\tau) = k \right]. \quad (3)$$

(In a traditional actuarial context, with no mortality improvements, $p(\tau, \tau + T, x - \tau, k)$ can be replaced by the standard actuarial function ${}_T p_x$.) The non-linear dependence of $S(u, x)$ on the $q(t, x)$ means that this expectation can, as mentioned above, only be calculated by simulation or other numerical methods. Expectations in (3) are taken under \mathbb{P} , but, where we are concerned with pricing and valuation, expectations will be taken under the risk-neutral pricing measure, \mathbb{Q} , as proposed in Cairns et al. (2006a) (CBD). As suggested in CBD, we will assume that, under \mathbb{Q} , $K(t + 1) = K(t) + \tilde{\nu} + C\tilde{Z}(t + 1)$, where $\tilde{\nu}$ is the risk adjusted drift, and $\tilde{Z}(t + 1)$ is a standard bivariate normal random vector under \mathbb{Q} . Risk neutral spot survival probabilities calculated under \mathbb{Q} (i.e. $E_{\mathbb{Q}}$ in (3) instead of $E_{\mathbb{P}}$) will be denoted by $p_{\mathbb{Q}}(\tau, \tau + T, x - \tau, k)$.

3 Approximating spot survival probabilities

Many financial functions at time τ derive their value from the spot survival probabilities at time τ evaluated under \mathbb{Q} , $p_{\mathbb{Q}}(\tau, \tau + T, x - \tau, k)$ (since these probabilities determine the mortality table, incorporating mortality projections, that will be in use at time τ for pricing annuities). In the paragraphs that follow, we therefore propose a method for approximating accurately the full range of spot survival probabilities, $p_{\mathbb{Q}}(\tau, \tau + T, x - \tau, k)$. In doing so, we provide ourselves with a flexible framework that can be easily applied *without recalibration* to any portfolio of survivorship-linked cashflows and under any given model for the term structure of

interest rates. This contrasts with the q -duration or q -duration-convexity approximations proposed by Li and Luo (2011) and Plat (2010) respectively which are potentially faster computationally, but which are also dependent on the structure of the portfolio of liabilities.

If our problem is based on a known starting point $K(0) = k$ then simulation under \mathbb{Q} or \mathbb{P} is not a problem. However, we might wish to ask the question: *what is the distribution of $p(\tau, \tau+T, x-\tau, K(\tau))$: the probability that an individual alive and aged x at time τ survives until time $\tau+T$?* Now this depends upon the simulated value of $K(\tau)$ at time τ , so, for each simulated $K(\tau)$ we need to conduct further simulations to establish the value of $p(\tau, \tau+T, x-\tau, K(\tau))$ or $p_{\mathbb{Q}}(\tau, \tau+T, x-\tau, K(\tau))$. This requirement for simulations within simulations is computationally very expensive and points, therefore, to the need for some simple numerical approximations.

The Markov, time-homogeneous nature of the random walk, $K(t)$, and the dependence of $q(t, x)$ on $K(t+1)$, means that

$$p(\tau, \tau+T, x-\tau, k) = p(0, T, x, k).$$

Thus, although our usual starting point $K(0)$ is known, there is no reason why we cannot evaluate spot survival probabilities at time 0 using other initial conditions.

Now $p(0, T, x, k)$ lies between 0 and 1, and so it is common in the statistical literature to transform this onto the interval $(-\infty, \infty)$ using either the logistic or probit transforms before making any approximations. Here we choose to apply the probit transform as our first step: that is, we define $f(T, x, k) = \Phi^{-1}(p(0, T, x, k))$, where Φ^{-1} is the inverse of the standard normal distribution function. We propose the following approximations to $f(T, x, k)$. Let $\hat{k} = (\hat{k}_1, \hat{k}_2)' = E[K(\tau)]$. Then, a linear approximation (based upon the Taylor expansion around \hat{k}) is

$$f(T, x, k) \approx D_0(T, x) + D_1(T, x)'(k - \hat{k})$$

or, if a more accurate quadratic approximation is required,

$$f(T, x, k) \approx D_0(T, x) + D_1(T, x)'(k - \hat{k}) + \frac{1}{2}(k - \hat{k})'D_2(T, x)(k - \hat{k})$$

where $D_0(T, x)$ is a scalar function of (T, x) , $D_1(T, x)$ is a 2×1 vector of first derivatives, and $D_2(T, x)$ is a 2×2 matrix of second derivatives. Specifically, with $k = (k_1, k_2)'$:

$$D_0(T, x) = f(T, x, \hat{k}),$$

$$D_{1,i}(T, x) = \left. \frac{\partial f}{\partial k_i}(T, x, k) \right|_{k=\hat{k}}, \quad \text{and} \quad D_{2,ij}(T, x) = \left. \frac{\partial^2 f}{\partial k_i \partial k_j}(T, x, k) \right|_{k=\hat{k}}.$$

The functions D_0 , D_1 and D_2 will depend on whether we wish to calculate spot survival probabilities under \mathbb{P} or \mathbb{Q} .

The probit transformation was chosen here for two reasons. First and foremost, the probit transformation, in combination with the first and second-order Taylor approximations, was found to produce accurate estimates of the spot survival probabilities that compare favourably with alternative transformations such as the log or logistic transforms (see Appendix A), or no transform at all (as used by Plat, 2010, at the portfolio level).¹ Second, we exploit, later in this paper, a distributional property of the first-order probit-Taylor approximation (in contrast to the logistic transform) to derive closed-form expressions for hedging strategies. The logit and log transforms were considered as alternatives to the use of the probit transform (see Appendix A). The log transform was relatively poor. The logit and probit transforms produced reasonably similar qualities of approximation, the probit being better over a wide range of values for $K(\tau)$ at approximating longer maturity spot survival probabilities and annuity values, while the logit was better for shorter maturity spot survival probabilities. On balance the probit was considered to be better suited for the problems being considered in this paper including dynamic hedging.

Other approximations have been proposed by Denuit and Dhaene (2007), Denuit et al. (2010) for single-factor mortality models such as Lee-Carter, and by Cairns et al. (2011b).

This approximation is computed by making N simulations of $S(T, x)$ given $K(0) = \hat{k}$, and then repeating (after first having reset the random seed) for $K(0) = \hat{k} + (h_1, 0)'$ and $K(0) = \hat{k} + (0, h_2)'$ for small h_1 and h_2 . For the second and third sets of simulations we subtract the expected value for $S(T, x)$ from the baseline ($K(0) = \hat{k}$) and divide by h_1 and h_2 respectively to get the first derivatives. Additional values are required for the second derivatives, but the principle is the same. (See appendix B for further details.)

In Table 1 we present, for initial age $x = 65$ and $\hat{k} = (-3.7785, 0.11699)'$ = $E_{\mathbb{P}}[K(20)]$, values for $D_0(T, x)$, $D_1(T, x)$ and $D_2(T, x)$ for a selection of values for T . All spot probabilities used in the evaluation of the D 's have been calculated under \mathbb{P} . We can comment as follows:

- The $D_0(T, x)$ become more negative with T reflecting the gradually lower probability of survival to T .
- The $D_{1,1}(T, x)$ are all negative, indicating that a higher value of $K_1(\tau)$ means mortality rates will generally be higher at all future ages and therefore survival rates will be lower. Since all future mortality rates will be higher then there will be a proportionately bigger negative impact on survival probabilities to higher ages.

¹The use of first and second-order Taylor expansions is reminiscent of duration and convexity matching in interest-rate risk management (see Plat, 2010, in a mortality context). However, since the approximations here are made to the probit transform of the spot survivor probabilities, the analogy is, at best, a loose one.

T	$D_0(T, x)$	$D_{1,1}(T, x)$	$D_{1,2}(T, x)$	$D_{2,11}(T, x)$	$D_{2,12}(T, x)$	$D_{2,22}(T, x)$
1	2.445	-0.3581	3.4016	-0.039426	0.37466	-3.5577
2	2.1676	-0.39201	3.519	-0.050061	0.4497	-4.1327
5	1.7188	-0.45808	3.3537	-0.073481	0.53907	-4.8488
10	1.2436	-0.5449	2.32	-0.10796	0.46474	-6.269
15	0.83732	-0.63316	0.51431	-0.14525	0.1348	-10.753
20	0.42457	-0.73464	-2.2025	-0.18847	-0.51742	-21.677
30	-0.54931	-1.0009	-11.47	-0.28633	-2.9862	-80.113
40	-1.8594	-1.3244	-26.925	-0.32443	-5.3987	-173.34
55	-4.3429	-1.6401	-53.802	-0.32269	-7.3065	-284.13

Table 1: $D_i(T, x)$ functions, for $x = 65$, for use in the approximation $\Phi^{-1}(p(0, T, x, a)) \approx D_0(T, x) + D_{1,1}(T, x)(k_1 - \hat{k}_1) + D_{1,2}(T, x)(k_2 - \hat{k}_2)$. Parameter estimates as in equation (1). $\hat{k} = (-3.7785, 0.11699)'$. Simulations used in the calculations exclude parameter uncertainty.

- The $D_{1,2}(T, x)$ change sign. A higher than expected value of $K_2(\tau)$ means that mortality rates up to age $\bar{x} = 74.5$ will be lower than anticipated and mortality rates above 74.5 will be higher than anticipated. So for small values of T , the survival probability will be increased, while for larger values of T the raised mortality rates above age 74.5 eventually dominate.

A typical risk measurement problem then involves first simulating future values of $K(\tau)$ and then, under each scenario, calculating liability values using the approximation

$$\begin{aligned}
 p(\tau, \tau + T, x - \tau, K(\tau)) &\approx \hat{p}(\tau, \tau + T, x - \tau, K(\tau)) \\
 &= \Phi \left[D_0(T, x) + \sum_{i=1}^2 D_{1,i}(T, x)(K_i(\tau) - \hat{k}_i) \right].
 \end{aligned}$$

3.1 Approximating financial functions

Now we can see that the approximation, $\hat{p}(\tau, \tau + T, x - \tau, K(\tau))$, is an analytical formula, which means that $\hat{p}(\tau, \tau + T, x - \tau, K(\tau))$ can be observed directly at time τ without resorting to simulations within simulations.

Other life and financial functions can then be approximated as follows:

- Complete expectation of life at age x :

$$EFL_x(\tau) \approx 0.5 + \sum_{T=1}^{\infty} p(\tau, \tau + T, x - \tau, K(\tau)) \approx 0.5 + \sum_{T=1}^{\infty} \hat{p}(\tau, \tau + T, x - \tau, K(\tau)).$$

- Annuity payable for life annually in arrears:

Let $P(\tau, \tau + T)$ be the (zero-coupon-bond) price at τ for 1 payable at time $\tau + T$. Then the value at τ of 1 payable for life annually in arrears to a life aged x at time τ is

$$\begin{aligned} a_x(\tau) &= \sum_{T=1}^{\infty} P(\tau, \tau + T) p(\tau, \tau + T, x - \tau, K(\tau)) \\ &\approx \sum_{T=1}^{\infty} P(\tau, \tau + T) \hat{p}(\tau, \tau + T, x - \tau, K(\tau)). \end{aligned}$$

Note that this expression allows us to combine the probit-Taylor approximation with any stochastic interest-rate model of our choosing. At the other extreme, if we assume that the rate of interest is a constant rate of i then

$$a_x(\tau) \approx \sum_{T=1}^{\infty} (1 + i)^{-T} \hat{p}(\tau, \tau + T, x - \tau, K(\tau)).$$

4 Numerical illustrations

The accuracy of the quadratic and linear approximations is illustrated in Figures 1 to 4.

In Figure 1 we have plotted the probit transform of the spot survival probabilities, $p(0, T, x, k)$, for $x = 65$, and $T = 1, 10$ and 30 .² The plots show contours of $\Phi^{-1}(p(0, T, x, k))$ over a range of values of $k = (K_1, K_2)'$. In the left-hand plots we can see that contours for the quadratic approximation (dashed lines) are almost indistinguishable from the true values (solid lines). The linear approximation (right hand plots) looks reasonable although it is clearly not nearly as good as the quadratic approximation, and the nature of the approximation results in evenly spaced, linear contours. The broad orientation of the contours reflects the sign of $D_{1,1}(T, x)$.

The same comparisons are illustrated further in Figure 2. Here we plot the ratio of the approximation to the true spot survival probability. A value greater than 1 means the approximation is higher than the true value. By construction the approximation is exact (ratio= 1) in all four cases at the centre of the plot, $k = (-3.7785, 0.11699)'$. This confirms the superiority of the quadratic approximation. Both approximations are less good for $T = 30$ compared to $T = 10$, and for $T = 10$ versus $T = 1$, although the deterioration is less important in absolute terms.

²Note that a contour plot of the untransformed spot survival probabilities reveals contours that are much less uniformly spaced than the six plots in Figure 1. This gave a clear indication that better approximations to the spot survival probabilities could be achieved by first transforming the data.

These plots also include ‘clouds’ of dots. These show simulated realisations of $K(\tau)$ for $\tau = 20$. By design that the approximations will be best when the approximation is centred on the expected value of $K(\tau)$ for the correct time horizon, τ . However (not illustrated), even if the time horizon were, say, 10 years or 40 years, the approximation based on $\tau = 20$ would still be good. Thus, if multiple time horizons are relevant, then we can see that the quadratic approximation is generally preferable, the linear approximation might also be adequate. In this figure we can also see that the linear approximation tends to overestimate the spot survival probability.

In Figure 3 we look at the expected future lifetime at age 65 as a function of $K(\tau)$. In the left-hand plots, we show contours of the absolute value of $EFL_{65}(\tau)$: true values (solid black lines); approximate values (dashed red lines). In the right-hand plots, we show contours of the ratio of the approximation to the true value of the expected future lifetime. Again we can see that the quadratic approximation is rather better than the linear approximation. As before, we can note that the linear approximation tends to overestimate the expected future lifetime. However, we can see that the linear approximation is almost always within 1% of its true value and, in many instances, this quality of approximation will be perfectly adequate.

Figure 4 shows equivalent results for the fair price of a life annuity payable annually in arrears, $a_{65}(\tau)$, assuming a fixed rate of interest of 4% per annum effective. The results are generally similar to those for $EFL_{65}(\tau)$, although the approximation errors are rather smaller. For a time horizon of $\tau = 20$ years from 2008, the approximate annuity price is well within 0.05% using the quadratic approximation and significantly less than 0.5% using the linear approximation.

5 Linear approximation: further remarks

5.1 Valuing futures on spot survival probabilities: Result 1

Let \mathcal{M}_t represent the history of $K(s)$ up to time t . This informs us about underlying mortality rates up to time t , but does not allow us to make statements about the history of a given individual.

Let $p^{FUT}(s, t, t+T, x-t, K(0)) = E_{\mathbb{Q}}[p(t, t+T, x-t, K(t)) | \mathcal{M}_s]$ be the futures price at time $s < t$ for the spot survival probability $p(t, t+T, x-t, K(t))$ (also evaluated under \mathbb{Q}) due at time t .

Result 1

We claim that if, using the linear approximation,

$$p(t, t+T, x-t, K(t)) \approx \Phi \left(D_0(T, x) + D_1(T, x)'(K(t) - \hat{k}) \right)$$

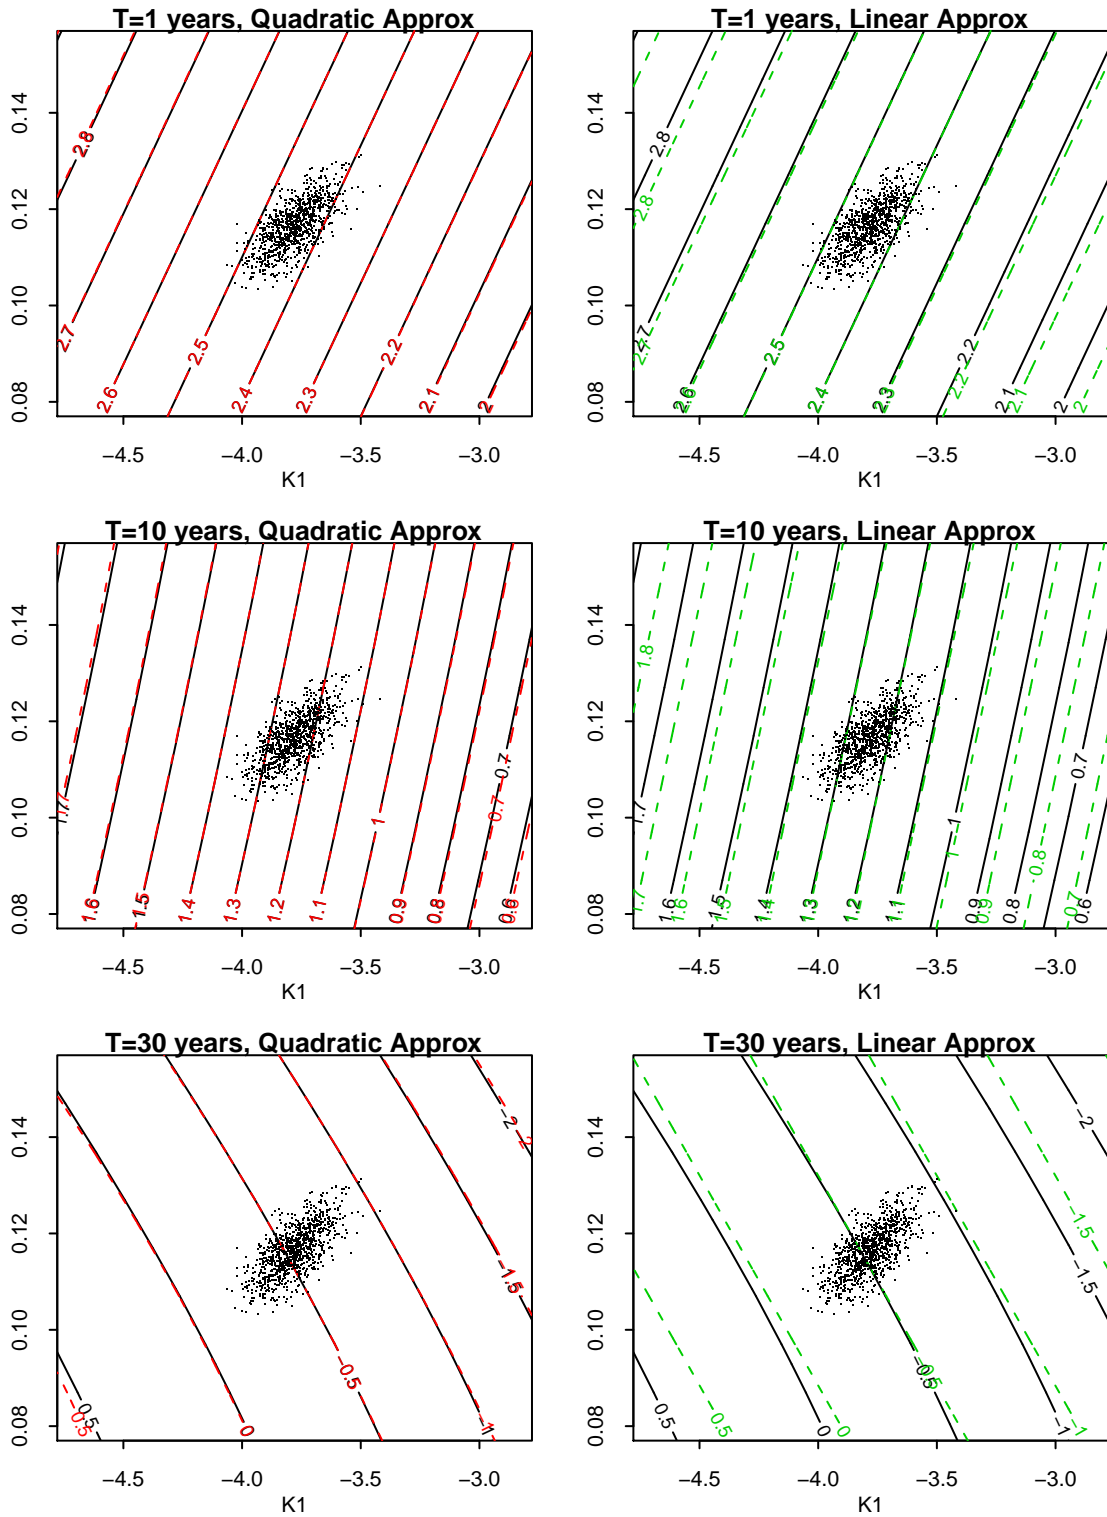


Figure 1: Plots show the probit transform of the spot survival probability, $\Phi^{-1}(p(0, T, x, k))$ for $k = (K_1, K_2)'$. All plots show contours of the true $\Phi^{-1}(p(0, T, x, k))$ (solid black lines) and the approximation $\Phi^{-1}(\hat{p}(0, T, x, k))$ (dashed lines). $x = 65$ throughout; $T = 1$ (upper plots), $T = 10$ (middle plots) and $T = 30$ (lower plots). Left-hand plots use the quadratic approximation. Right-hand plots use the linear approximation. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.

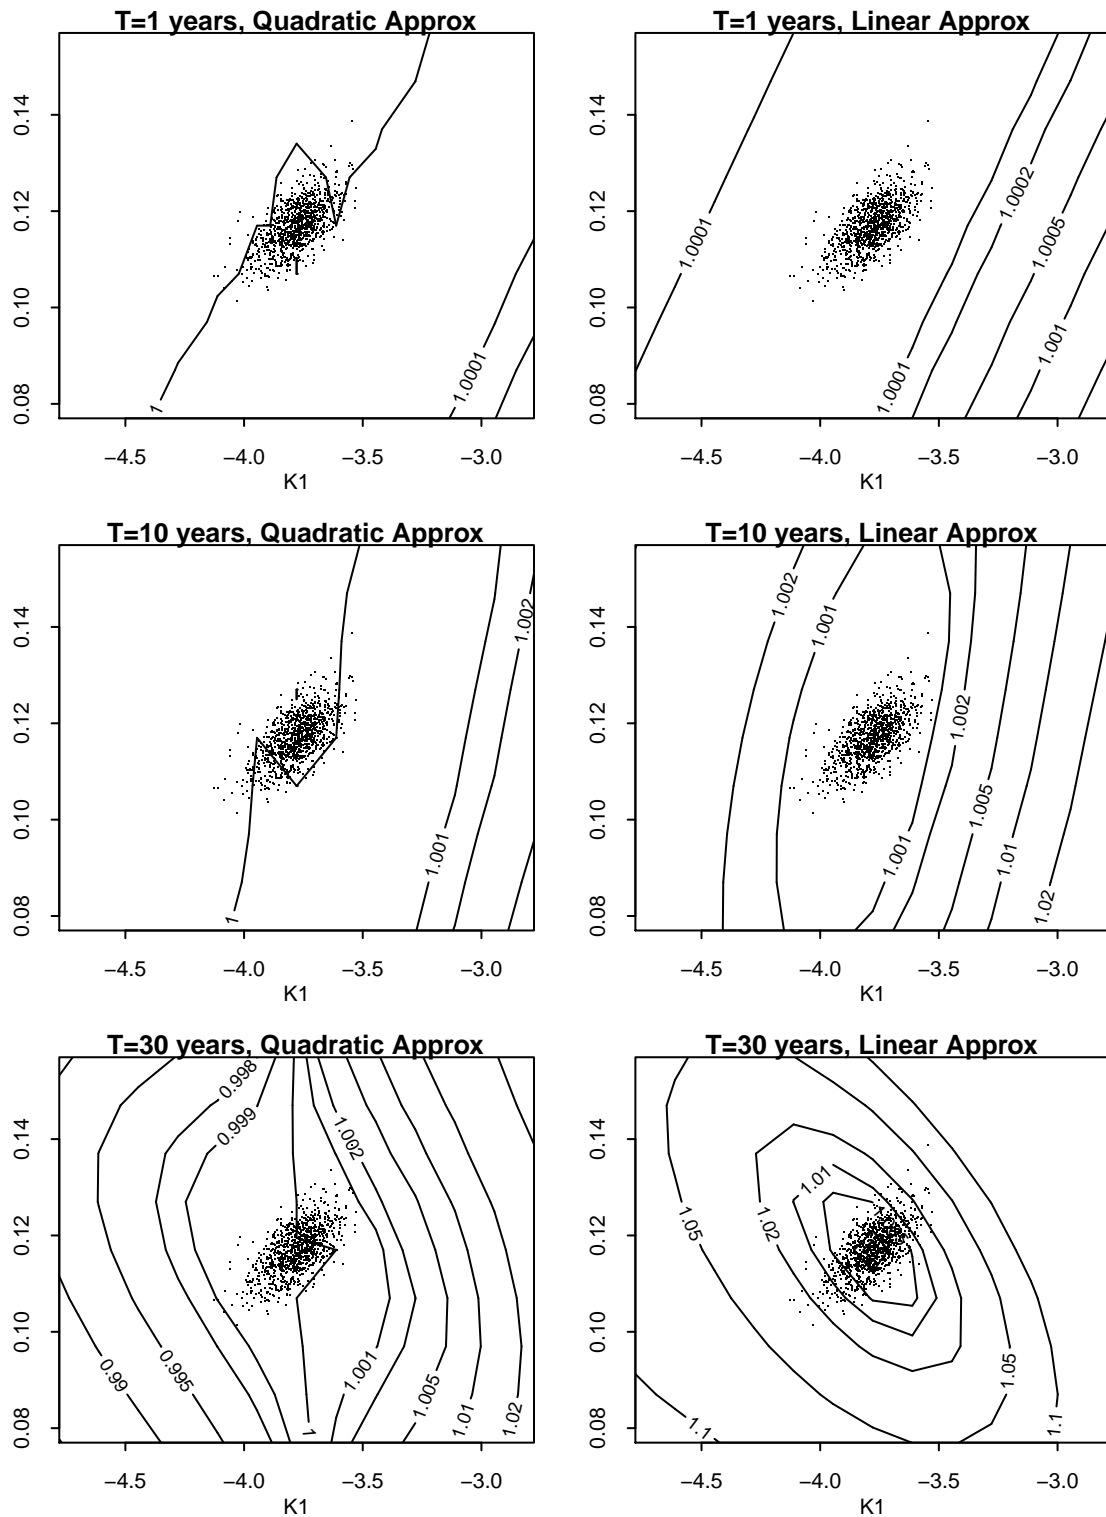


Figure 2: As Figure 1 except that the single set of contours in each plot show the ratio of the approximation to the true spot survival probability, $\hat{p}(0, T, x, k)/p(0, T, x, k)$. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.




Figure 3: Actual versus approximate values for the complete expectation of life, $EFL_{65}(k)$. Left-hand plots: contours of the absolute values for $EFL_{65}(k)$: true values (solid black lines) and approximate values (dashed lines). Right-hand plots show contours of the ratio of the approximate to the true expected future lifetime. Upper plots use the quadratic approximation. Lower plots use the linear approximation. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.

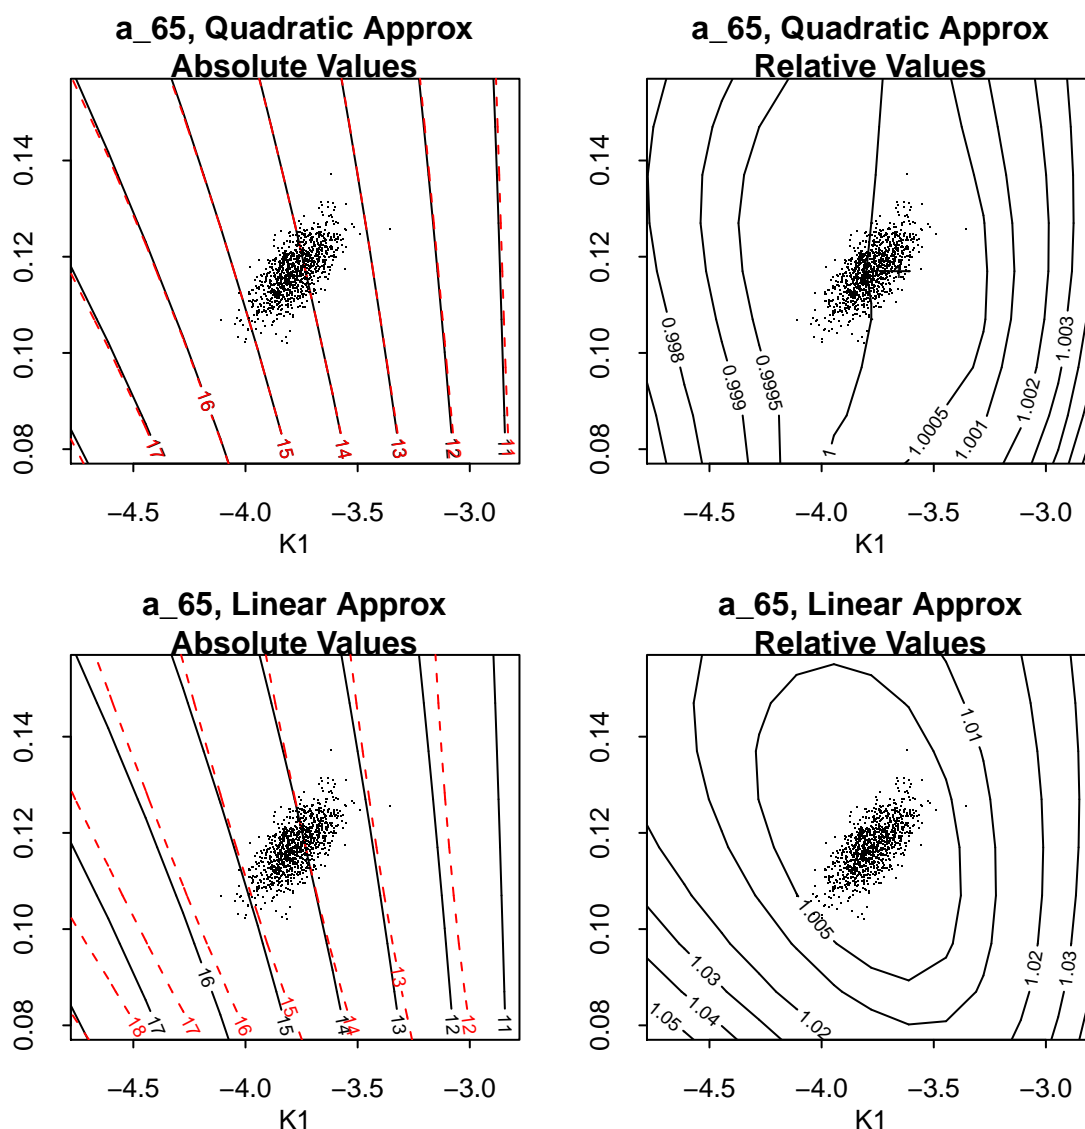


Figure 4: Actual versus approximate values for the annuity function, $a_{65}(k)$. Left-hand plots: contours of the absolute values for $a_{65}(k)$: true values (solid black lines) and approximate values (dashed red lines). Right-hand plots show contours of the ratio of the approximate to the true expected future lifetime. Upper plots use the quadratic approximation. Lower plots use the linear approximation. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.

then

$$p^{FUT}(0, t, t + T, x - t, K(0)) \approx \Phi \left(\frac{D_0(T, x) + D_1(T, x)'(K(0) + \tilde{\nu}t - \hat{k})}{\sqrt{1 + D_1(T, x)'VD_1(T, x)t}} \right) \quad (4)$$

where $\tilde{\nu}$ and V , recall, are the drift vector (under \mathbb{Q}) and the variance-covariance matrix of the random walk process for $K(t)$.

For a proof of (4) see Appendix C.

In other words, the probit-Taylor approximation for a spot survival probability can also be used, in a straightforward manner, to price (approximately) futures contracts.

5.2 Pricing q -forward contracts: Result 2

We now consider q -forward contracts (see, for example, www.llma.com and Coughlan et al., 2007). A simplified version of the q -forward contract is as follows. The crude mortality rate, $q(t, x)$, represents the probability, measured retrospectively, that an individual aged x at time t would have survived to time $t + 1$. We assume that $q(t, x)$ is known at time $t + 1$ but not before. Under the (t, x) q -forward contract:

- the forward price $q^F(0, t, x)$ is set at time 0;
- the contract has zero value at time 0;
- no money exchanges hands between times 0 and $t + 1$ (for example, we assume there are no margin payments between times 0 and $t + 1$);
- at time $t + 1$ the holder of the long position pays a fixed amount of $q^F(0, t, x)$ and receives the floating $q(t, x)$;
- $q^F(s, t, x) = E_{\mathbb{Q}}[q(t, x) | \mathcal{M}_s]$ represents the forward price at time s , $0 \leq s \leq t$.

No arbitrage dictates that, with one year remaining, the q -forward price must equal 1 minus the 1-year-ahead spot survival probability: that is,

$$q^F(t, t, x) = 1 - p(t, t + 1, x - t; K(t)) \approx 1 - \Phi \left(D_0(1, x) + D_1(1, x)'(K(t) - \hat{k}) \right). \quad (5)$$

In Figures 1 and 2 we noted that the approximation for a 1-year-ahead spot survival probability was the most accurate of the linear approximations for various maturities. However, we can remark further that the *retrospective* mortality rate, $q(t, x)$, can be evaluated accurately using the probit transform as an approximation to CBD's use of the logistic function. See Appendix D for further details.

Result 2

Using the probit-Taylor approximation at time t , it follows that

$$q^F(0, t, x; K(0)) \approx 1 - \Phi \left(\frac{D_0(1, x) + D_1(1, x)'(K(0) + \tilde{\nu}t - \hat{k})}{\sqrt{1 + D_1(1, x)'VD_1(1, x)t}} \right). \quad (6)$$

Remark

Appendix D, additionally, discusses how $D_0(1, x)$ and $D_1(1, x)$ can be calculated analytically as an alternative to the simulation approach used in the generation of Table 1.

5.3 The approximation as a market model

Superficially, having a sequence of spot and forward survival probabilities and q -forward prices suggests that we have what might be termed a *market model*. However, this is not true in the same theoretical sense of Olivier and Jeffery (2004) and Smith (2005) which would require the exact relationship

$$p_Q(t, t + T, x - t, K(t)) = E_{\mathbb{Q}} [(1 - q(t, x)) p(t + 1, t + T, x - t, K(t + 1)) | K(t)]$$

to hold. While this result is approximately correct, it is not true as a theoretical result. Even if it were true, the approximations can, in extreme scenarios, give rise to biologically unacceptable results (for example, for extreme values of $K(t)$, the spot survival curve $p(t, t + T, x - t, K(t))$ might not be a decreasing function of T). This contrasts with the underlying CBD model which has no such problems.

Finally, a true market model would need the initial spot survival probabilities to be calibrated to market data rather than be derived from an underlying spot-rate model such as CBD. The currently methodology could be easily adapted to satisfy this requirement.

6 Calculating approximate Deltas

In this section we develop the key quantities required for dynamic Delta hedging of a portfolio of liabilities exposed to longevity risk. We will focus on the use of q -forwards as the hedging instrument, and use these to hedge liabilities linked to the survivor index.

6.1 Deltas for q -forward prices

First we consider the partial derivatives of q -forward prices with respect to the stochastic state variables K_1 and K_2 , which can then be used to develop Delta

hedging strategies for longevity-linked liabilities. Specifically,

$$\Delta_1 = \frac{\partial q^F(0, t, x; k)}{\partial k_1} \quad \text{and} \quad \Delta_2 = \frac{\partial q^F(0, t, x; k)}{\partial k_2}.$$

Using Result 2, equation (6), it follows that, with $k = K(0)$,

$$\begin{aligned} \Delta_1 &= \frac{\partial q^F(0, t, x; K(0))}{\partial k_1} = -\frac{D_{1,1}(1, x)}{\sqrt{1 + D_1(1, x)'VD_1(1, x)t}}\phi(z), \\ \Delta_2 &= \frac{\partial q^F(0, t, x; K(0))}{\partial k_2} = -\frac{D_{1,2}(1, x)}{\sqrt{1 + D_1(1, x)'VD_1(1, x)t}}\phi(z), \\ \text{where } z &= \frac{D_0(1, x) + D_1(1, x)'(K(0) + \tilde{\nu}t - \hat{k})}{\sqrt{1 + D_1(1, x)'VD_1(1, x)t}} \end{aligned}$$

and $\phi(z)$ is the density function of the standard normal.

Alternative approximations for $q^F(s, t, x)$ and its Deltas are presented in Appendix E. Although these are not used in our subsequently analysis, the underlying methodologies been used elsewhere and might be found to be useful in a variety of other circumstances.

6.2 Deltas for spot survival probabilities

Here we calculate the Deltas using the single approximation

$$p_Q(\tau, \tau + T, x, k) = \Phi(z)$$

$$\text{where } z = D_0(\tau, T, x + \tau, \hat{k}(\tau)) + D_1(\tau, T, x + \tau, \hat{k}(\tau))'(k - \hat{k}(\tau)).$$

In this equation, we have augmented our notation so that we properly record the fact that D_0 and D_1 depend on the future valuation time τ , the age of the cohort at time τ , $x + \tau$, the maturity, T , of the payment after time τ , and the value $\hat{k}(\tau)$ around which the approximation is centred. We assume that $\hat{k}(\tau) = E_{\mathbb{P}}[K(\tau)|K(0)]$, and, therefore, can be calculated once in advance of any simulation exercise.

It is straightforward to see that

$$\begin{aligned} \Delta_1 &= \frac{\partial p(\tau, \tau + T, x, k)}{\partial k_1} = D_{1,1}(\tau, T, x + \tau, \hat{k}(\tau))\phi(z) \\ \Delta_2 &= \frac{\partial p(\tau, \tau + T, x, k)}{\partial k_2} = D_{1,2}(\tau, T, x + \tau, \hat{k}(\tau))\phi(z), \end{aligned}$$

and $\phi(z)$ is the density of the standard normal.

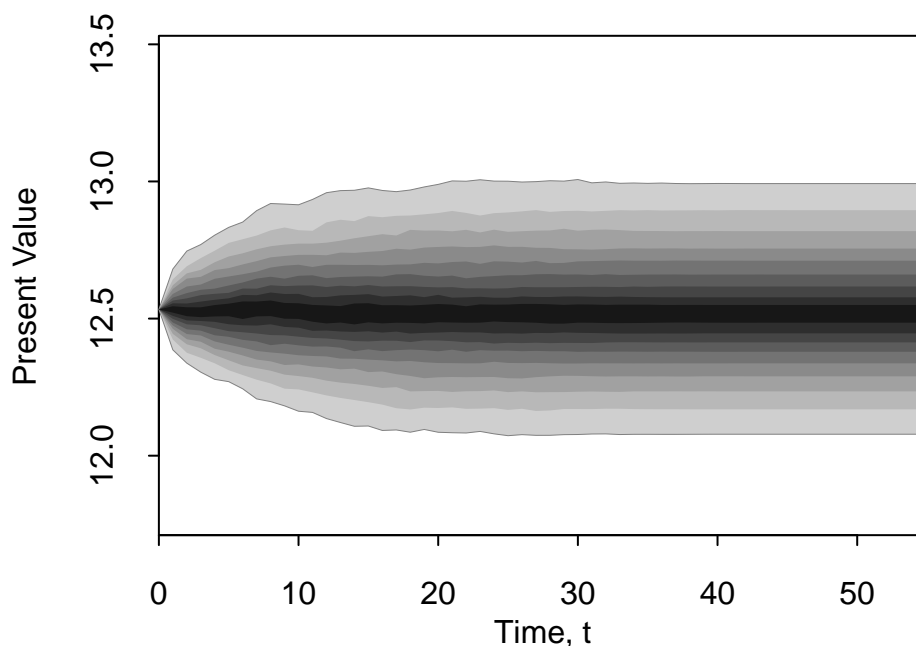


Figure 5: Fan chart of the present value, discounted to time 0, of the annuity portfolio as a function of mortality information available up to time t : $PV(t)$.

7 Example: Delta hedging of an annuity portfolio

Suppose we have a liability to pay $S(t, 65)$ (equation 2) at times $t = 1, 2, \dots, 55$.³ In order to focus on the hedging of longevity risk we will assume that interest rates are constant and that the bank account pays 4% per annum.

In Figure 5 we have plotted a fan chart⁴ showing how the present value of this liability evolves over time, $PV(t)$, based on 1000 simulated scenarios of $K(t)$. Under each scenario, at time t we value future cashflows after time t and then discount these along with all known cashflows at times 1 to t back to time 0. In our simulations we have, for simplicity assumed that the real-world and risk-neutral probability measures are the same. If there were no approximations then $PV(t)$ would be a martingale. Figure 5 appears to be consistent with this property, and further checks on the detail of the simulated $PV(t) - PV(t - 1)$ bear this out, adding weight to the usefulness of the probit-Taylor approximations. After about 35 years, all sample

³ $t = 0$ corresponds to the end of 2008 and age 65. Time 55 corresponds to age 120 which we use as a high cutoff age for a portfolio of life annuities.

⁴The shaded part of the fan chart covers the 5% to 95% quantile range of the distribution at t , and is divided up into 5% quantile bands.

paths of $PV(t)$ remain almost constant as the outstanding payments are usually very small after age 100. The distribution of $PV(55)$ is the same as the distribution of $\sum_{t=1}^{55} S(t, 65)/1.04^t$.

We now show one approach to how the uncertainty in $PV(t)$ can be hedged.

The key points are as follows:

- We will carry out Delta hedging with rebalancing at annual intervals.
- At each date t , three financial instruments can be used: cash paying 4% interest; and two q -forwards both with 10-years to maturity, with reference ages 65 and 75 and zero value at time t , so that the forward rate on each is $q^F(t, t+9, x)$.
- Deltas are calculated using the probit-Taylor approximations. They are based on immediate (i.e. still at time t) small changes in $K_1(t)$ and $K_2(t)$ and measure the impact of these changes on the value of cashflows that fall after t .
- At the end of each period $[t, t+1)$, any q -forward positions set up at time t are closed out and return $1.04^{-9}(q^F(t+1, t+9, x) - q^F(t, t+9, x))$ per unit to be reinvested in the cash account.⁵ q -forward prices are calculated using the probit-Taylor approximations.

Deltas for the liability values, $PV(t)$, are plotted in Figure 6. Δ_1 (top left plot) can be seen to decline steadily in magnitude, albeit with growing uncertainty over the first 20 years. However, this decline is entirely attributable to the gradually falling value of the *outstanding* liabilities after t . If we normalise Δ_1 by dividing by the value of the outstanding liabilities (bottom left plot) then we see that, in relative terms, the outstanding liabilities become more sensitive over time to changes in $K_1(t)$. This is explained by the fact that we are valuing survivorship-linked payments rather than mortality-linked payments. For younger ages and smaller t , survivorship will always be close to 1 regardless of the likely variation in underlying mortality rates, whereas, at higher ages, a 1% relative change (say) in a mortality rate has a relatively bigger impact on short-term survivorship. The lower plot reveals that, for values of t up to about 35, the relative values of Δ_1 do not contain significant levels of uncertainty around a deterministic trend.

Absolute and normalised values for Δ_2 are given in the top right and bottom right plots in Figure 6. The pattern for the absolute values is more interesting than Δ_1 , but still benefits from being normalised by the outstanding present value. The normalised value of Δ_2 (bottom right plot) indicates that, initially, $K_2(t)$ has relatively

⁵This assumes that the contract operates as a forward contract rather than a futures contract. Under the latter, there would be a margin payment of $q^F(t+1, t+9, x) - q^F(t, t+9, x)$ immediately at time $t+1$. Under a forward contract $q^F(t+1, t+9, x) - q^F(t, t+9, x)$ represents the change in the expected payout at time $t+10$.

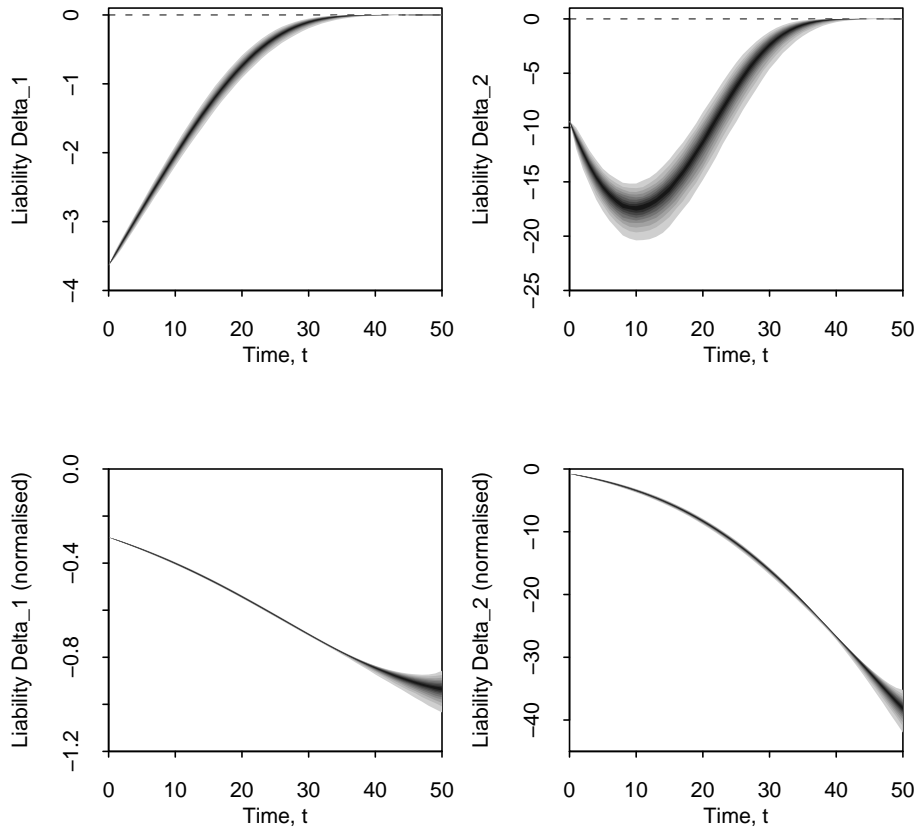


Figure 6: Top row: Fan charts for $\Delta_1(t)$ and $\Delta_2(t)$ respectively for the annuity portfolio, $PV(t)$. Bottom row: Fan charts for $\Delta_1(t)$ and $\Delta_2(t)$ divided by the present value of the outstanding liabilities at time t respectively.

little impact suggesting that the average time to payment is about 9.5 years (i.e. the difference between the pivotal age $\bar{x} = 74.5$ years and the initial $x = 65$). However, as t increases, the average time to payment plus the current time, t , increases steadily above 9.5, and so changes in $K_2(t)$ become, relatively, more and more important. Again (bottom right plot), over the first 40 years the normalised value of Δ_2 follows its deterministic trend very closely.

In Figure 7 we have plotted:

- Fan charts for the age 65 and age 75 q -forward prices, $q^F(t, t + 9, x)$ (top left and top right respectively);
- Fan charts for the age 65 and age 75 q -forward relative Δ_1 's, $\Delta_1(x)/q^F(t, t + 9, x)$ (middle left and middle right respectively);

- Fan charts for the age 65 and age 75 q -forward relative Δ_2 's, $\Delta_2(x)/q^F(t, t + 9, x)$ (middle left and middle right respectively).

The upper plots show how the q -forward prices steadily improve over time (improvement rates of 2% to 3% per annum depending on reference age x). The relative Delta plots look reasonably stable over time. This should not be surprising. If we approximate $q(t, x)$ by $\exp[K_1(t + 1) + K_2(t + 1)(x - \bar{x})]$ then the relative Deltas would be 1 and $x - \bar{x}$, and these levels are, approximately, what we see in the relative Delta plots: namely approximately 1 in the middle and bottom left plots, $65 - \bar{x} = -9.5$ in the middle right plot, and $75 - \bar{x} = 0.5$ in the bottom right plot.

Under the Delta-hedging strategy, the optimal hedge ratios are $u_{65}(t)$ and $u_{75}(t)$ for the age 65 and 75 q -forwards (i.e. the number of units of each q forward to be held from time t to time $t + 1$ that allow the asset and liability Deltas to be matched).⁶

The hedge ratios are plotted in Figure 8.

Both $u_{65}(t)$ and $u_{75}(t)$ become more uncertain over time as we might expect, before declining to zero as the value of outstanding liabilities declines to 0. Figure 7 revealed that the age 65 q -forward would be relatively much better at heading the $K_2(t)$ risk. Thus, as the portfolio ages and the Δ_2 for the liability increases in magnitude (Figure 6, right), we need increasing quantities of the age-65 q -forward. The pattern taken by $u_{75}(t)$ then reflects the balancing act in terms of matching the $\Delta_1(t)$ values.

Finally, Figure 9 demonstrates how successful our hedging strategy has been. We let $A(t)$ represent the value of our assets at time t discounted back to time 0. To compare this with $PV(t)$ (which includes the liabilities already paid at times 1 to t before also discounting to time 0), $A(t)$ values not just the assets in hand at t , but needs to add in the value of the liabilities already paid. Thus, $A(t)$ equals its initial value, $A(0)$ (which we take to be equal to $PV(0)$), plus the value in the gains and losses on the q -forward contracts up to time t , with the liabilities $S(u, 65)$ for $u = 1, \dots, t$ paid out but then added back in. The left-hand plot in Figure 9 shows how the surplus $A(t) - PV(t)$ develops over time. The outer fan shows results when we do not hedge using q -forwards and simply use the cash account. Essentially this matches what we saw in Figure 5, but centred on 0. The inner and much narrower grey fan shows how the surplus evolves over time when the delta hedging strategy has been implemented. We can see a very substantial reduction in risk. Specifically, the standard deviation of the surplus at time 55 discounted to time 0 is reduced from 0.283 to 0.00804: representing a hedge effectiveness of about 97%. This is confirmed when we look at the right hand plot in Figure 9 where we plot $PV(55)$ versus the value, $A(55)$, of the assets at time 55 discounted to time 0. We can see a very high correlation between the two: $\text{cor}(A(55), PV(55)) = 0.9996$.⁷

⁶That is we solve the two linear equations $\Delta_1^{(L)}(t) = u_{65}(t)\Delta_1^{(QF,65)}(t) + u_{75}(t)\Delta_1^{(QF,75)}(t)$ and $\Delta_2^{(L)}(t) = u_{65}(t)\Delta_2^{(QF,65)}(t) + u_{75}(t)\Delta_2^{(QF,75)}(t)$ in the unknown $u_{65}(t)$ and $u_{75}(t)$.

⁷The hedge effectiveness using standard deviation as a risk measure is $1 - \sqrt{1 - \rho^2} = 0.9717$.

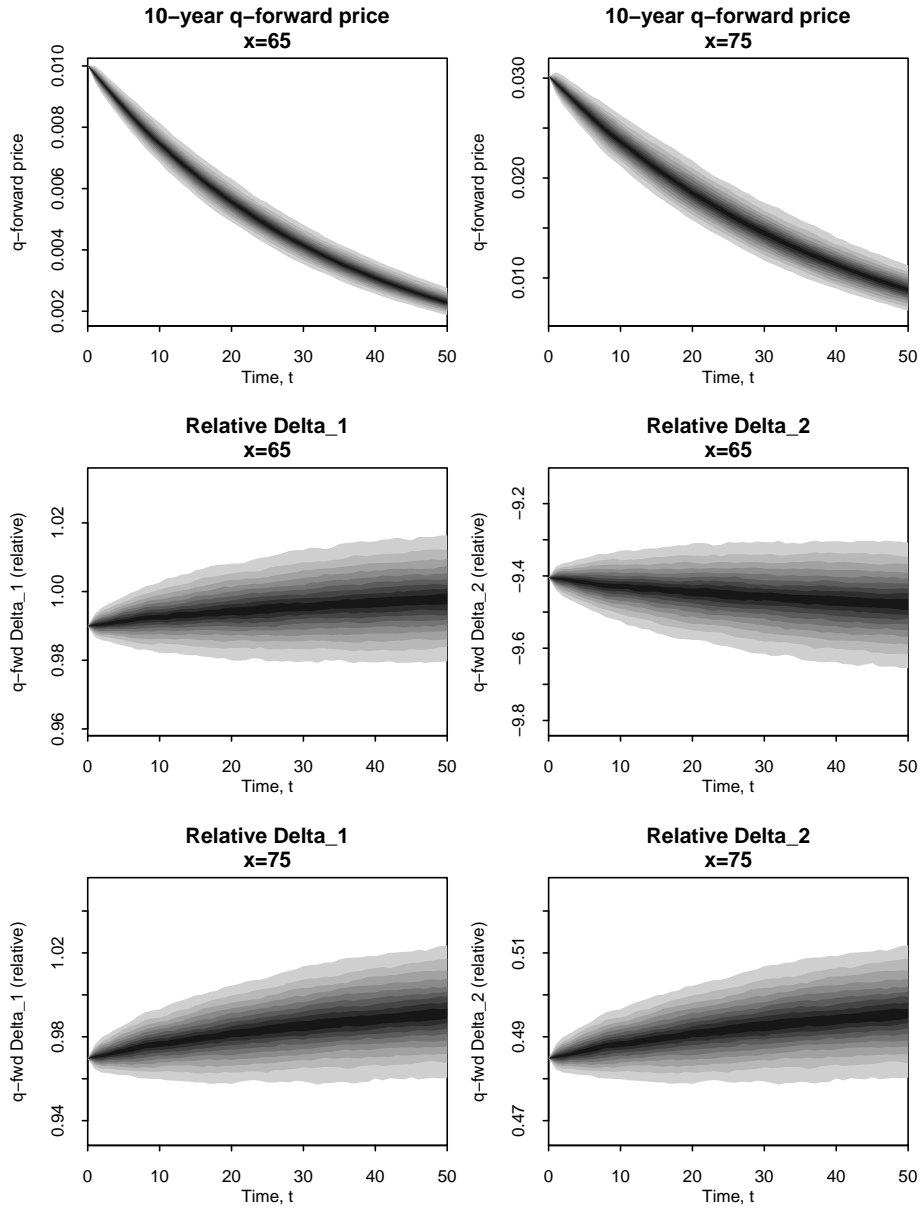


Figure 7: Top row: Fan charts for $q^F((t, t + 9, 65))$ and $q^F(t, t + 9, 75)$ respectively. Middle row: $\Delta_1(t)$ and $\Delta_2(t)$ for $q^F(t, t + 9, 65)$ relative to the q -forward price, $q^F(t, t + 9, 65)$. Bottom row: $\Delta_1(t)$ and $\Delta_2(t)$ for $q^F(t, t + 9, 75)$ relative to $q^F(t, t + 9, 75)$.

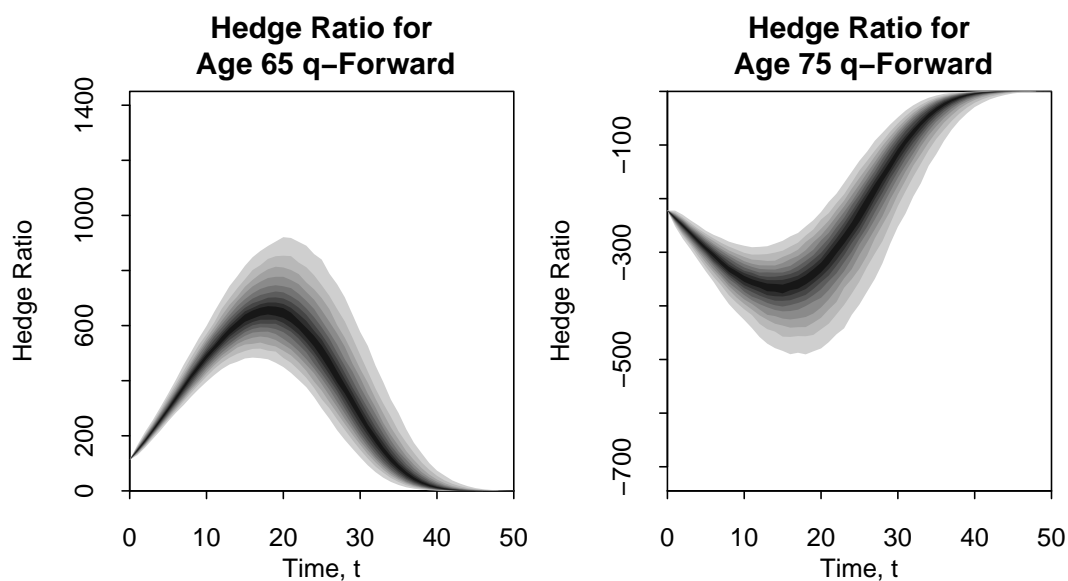


Figure 8: Fan charts for the hedge ratios $u_{65}(t)$ (left) and $u_{75}(t)$ (right).




Figure 9: Left: fan charts for the present value at time 0 of the unhedged surplus, $A(t) - PV(t)$ (wider grey fan) and for the hedged surplus process (narrow central fan). Right: scatter plot of the present value at time 0 of the ultimate liability at time 55 against the value of the hedged portfolio of assets at time 55; the correlation is 0.9996.

7.1 Comparison with static hedging

Although we do not report results in detail here, we compared our dynamic hedging results with what could be achieved using static hedging. Full static hedging used the model-free hedge proposed by Cairns et al. (2008), which, for a 55-year annuity requires the use of 55 distinct q -forwards with maturities in 1 year up to 55 years, all linked to the same cohort currently aged 65. We also looked at a static hedge of the same portfolio using (A) only seven out of the 55 q -forwards maturing in 5, 10, 15, 20, 25, 30 and 35 years and (B) two out of the 55 maturing in 11 and 22 years (these being the optimal pairing out of all possibilities). In Table 2 we report on the standard deviation of the present value at time 0 of the surplus emerging at time 55 under the four hedging strategies, and the resulting hedge effectiveness relative to the case with no hedging. We can see that dynamic hedging with only two q -forwards (hedge effectiveness equal to 0.9716) significantly outperforms even the case with full static hedging (0.8911).⁸ The seven q -forwards used in hedge (A) is a subjective choice that covers the financially more important ages. The optimal hedge is now model dependent with hedge ratios chosen to minimise the standard deviation of the present value of the surplus. The reduction from 55 to 7 q -forwards seems drastic, but only results in a modest, but non-trivial, reduction in hedge effectiveness. Finally, we picked out the two best q -forwards out of the original 55, and found that a static hedge with two q -forwards is substantially worse than a dynamic hedge with two q -forwards.

Hedge	St. Dev. PV surplus	Hedge Effectiveness
No hedging	0.2829	0
Dynamic	0.0080	0.9716
Full model-free static	0.0308	0.8911
(A) static, 7 q -fwds	0.0385	0.8639
(B) static, 2 q -fwds	0.0637	0.7747

Table 2: Hedge effectiveness under five hedging strategies.

There is not necessarily a clear conclusion from this, despite what the table suggests. The reality is that implementation of both static and dynamic hedges are difficult in practice, because of the limited number of contracts that we are likely to see emerging in a traded market, a possible cap of 10 years on the duration of q -forwards, and a lack of liquidity impeding dynamic hedging.

⁸For term annuities, the longer the term of the annuity, the more non-linearities in $S(T, x)$ as a function of the $q(t, x + t)$ cause errors in the model-free hedge. If we wish to hedge a term annuity with a shorter term, then the full model-free static hedge might outperform the dynamic hedge.

8 Further discussion

The success of the hedging strategy comes in spite of the fact that we are using approximations for the q -forward prices and hedging in discrete time. We can conclude from this two things. First, discrete hedging works well because the year on year randomness in mortality rates is relatively small and, therefore, changes in asset and liability values are approximately linear in $K(t)$. Second, we conclude that the probit-Taylor approximation for the liabilities works well. As a tentative, third point, we might also conclude that the approximation also works well for the q -forward prices, but we hesitate on that count because the subsequent gains and losses are determined using the same approximation.

The numerical results that are presented are, of course, the best that we might expect, and, for a variety of reasons, we would not expect hedge effectiveness not to be quite so high. We now discuss these issues and their implications.

A key factor not considered here is the effect of population basis risk: the risk associated with the fact that the population being hedged is different from the reference population underpinning the hedging instrument. While this is a significant consideration, it has been demonstrated elsewhere (Cairns et al., 2011b, Li and Hardy, 2011) that, at least over medium and longer-term horizons, hedge effectiveness is still high even using static hedging. For further discussion of multipopulation modelling, see Cairns et al. (2011a,b), Dowd et al. (2011a,b), Dahl et al. (2008), Li and Lee (2005), Plat (2009), Jarner and Kryger (2011), Li and Hardy (2011). For these models, hedge effectiveness will depend on to what extent the non-hedgeable component of short-term mortality shocks in the hedger's own population persist over time.

In most cases, these models will be amenable to the use of the probit-Taylor approximation, which, in turn, will allow the development of delta-hedging strategies.

The delta hedging strategy assumes a liquid market in which q -forward positions can be closed out at the end of each year at no cost. Primarily this is for computational convenience. If we assume, alternatively, that q -forwards must be held to maturity, then a rolling programme of investment in q -forwards that continues existing positions and takes new positions in 10-year-maturity, age 65 and 75 q -forwards. The new investments will neutralise the Δ_1 's and Δ_2 's taking account of the liability Deltas and also the Deltas of the ongoing q -forwards. Assuming the initial q -forward prices are fair, then such a strategy in an illiquid market should be just as effective as the approach described in this paper. Computationally, the additional complexity would be the requirement to keep track of up to 20 q -forwards rather than just two. We leave for further work the possibility of including some form of transaction cost which might introduce significant additional costs and reduce hedge effectiveness.

The general approach used here – both the probit-Taylor approximation and the discrete-time delta hedging – should be applicable to a wide range of stochastic

mortality models, including all of those considered in Cairns et al. (2009) (other than their model M4), and the multipopulation models mentioned above. Specifically, the approach can be adapted to models with any number of random period and cohort effects, and models that work with the log of the death rate, $\log m(t, x)$, (e.g. Lee and Carter, 1992) rather than logit $q(t, x)$ as in Cairns et al. (2006a).⁹ Additionally, most of the commonly used stochastic mortality models share the characteristic with the CBD two-factor model that randomness in mortality rates builds up gradually over time meaning that discrete-time hedging should work well in most cases. For a given model, the probit-Taylor approximation would need to reflect the number of random factors in the mortality model and a hedging strategy would need cash plus as many hedging instruments as there are random factors. The accuracy of the approximation and of the hedging strategy might be better or worse than we have found for the CBD model, and it is difficult to anticipate which way this might go without analysing each model individually.

We have focused in this paper on hedging “vanilla” annuity payments proportional to $S(t, x)$ which can be valued easily using the probit-Taylor approximation. In contrast, suppose we seek to value an option on an annuity price, $\max\{a(T, x) - g, 0\}$, where $a(T, x) = \sum_{s=1}^{\infty} (1+i)^{-s} p(T, T+s, x-T, K(T))$. The expected value of $a(T, x)$ given $K(t)$ where $t < T$ is straight forward to calculate using the probit-Taylor approximation, but the expected value of $\max\{a(T, x) - g, 0\}$ is not. So it will be necessary to develop further approximations to deal with the non-linearity of the payoff at T .

We also leave for further work consideration of robustness of the Delta hedging strategy. For example, what if the parameters of the CBD model have been miscalibrated, or what if the model itself is wrong? Are there alternatives to the recurring use of age 65 and 75 q -forwards that are more robust under these uncertainties?

Acknowledgements

I would like to thank Kevin Dowd and David Blake for their comments on earlier versions of this paper.

References

- Blake, D., Cairns, A.J.G., and Dowd, K. (2006) Living with mortality: longevity bonds and other mortality-linked securities *British Actuarial Journal*, 12: 153-197.
- Börger, M. (2010) Deterministic shock vs. stochastic value-at-risk – An analysis of the Solvency II standard model approach to longevity risk. *Blätter DGVMF*, 31:

⁹Indeed, the probit-Taylor approximation would cease to be an approximation for q -forward pricing if we choose to model $\Phi^{-1}(q(t, x))$ as being normally distributed instead of the usual assumptions of normality for $\log m(t, x)$ or logit $q(t, x)$.

225-259.

Cairns, A.J.G., Blake, D., and Dowd, K., (2006a) 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., (2006b) Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. *ASTIN Bulletin*, 36: 79-120.

Cairns, A.J.G. (2007) A multifactor generalisation of the Olivier-Smith model for stochastic mortality. Proceedings of the 1st IAA-Life Colloquium, Stockholm, June 2007.

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

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

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., and Khalaf-Allah, M. (2011a) Bayesian Stochastic Mortality Modelling for Two Populations. *ASTIN Bulletin* 41: 29-59.

Cairns, A.J.G., Blake, D., Dowd, K., and Coughlan, G.D., (2011b) Longevity Hedge Effectiveness: A Decomposition. Working paper, Heriot-Watt University.

Coughlan, G., Epstein, D., Sinha, A., and Honig, P. (2007) q-Forwards: Derivatives for transferring longevity and mortality risk. Available at www.lifemetrics.com. (Accessed 17/5/2011.)

Dahl, M., Melchior, M., and Møller, T. (2008) On systematic mortality risk and risk minimization with mortality swaps. *Scandinavian Actuarial Journal*, 2008(2-3): 114-146.

Denuit, M., and Dhaene, J. (2007) Comonotonic bounds on the survival probability in the Lee-Carter model for mortality projection. *Computational and Applied Mathematics*, 203: 169-176.

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

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

Dowd, K., Blake, D., Cairns, A.J.G., and Coughlan, G.D. (2011b) Hedging pension risks with the age-period-cohort two-population gravity model. Seventh In-

ternational Longevity Risk and Capital Markets Solutions Conference, Frankfurt, September 2011.

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

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

Li, J.S.-H., and Luo, A. (2011) Key q-duration: A framework for hedging longevity risk. To appear in *ASTIN Bulletin*.

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

Olivier, P., and Jeffery, T. (2004) Stochastic mortality models. Presentation to the Society of Actuaries of Ireland.

See <https://web.actuaries.ie/events/past> (accessed 25/5/2011).

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

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

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

Smith, A.D. (2005) Stochastic mortality modelling. Workshop on the Interface between Quantitative Finance and Insurance, International Centre for the Mathematical Sciences, Edinburgh.

See <http://www.icms.org.uk/archive/meetings/2005/quantfinance/> (accessed 25/5/2011).

Zhu, N., and Bauer, D. (2010) Gaussian forward mortality factor models: specification, calibration and application. Sixth International Longevity Risk and Capital Markets Solutions Conference, Sydney, September 2010.

A Probit versus logistic and log transforms

The linear and quadratic approximations were each considered in conjunction with the probit, logistic and log transforms. A sample of results are presented in Figures 10 and 11. Figure 10 plots contours of the difference between the approximate and actual spot survival probabilities. In all cases the log transform produces worse results than the probit and logistic transforms. In general, the probit and logistic transforms produce similar results. For lower terms to maturity, the logistic transform produces slightly more accurate results with the linear approximation,

whereas the convexe is true for higher maturities. However, for lower maturities, the approximation errors are typically quite small (e.g. mostly less than 0.0001 for $T = 5$) whereas, for higher maturities, the approximation error tends to be larger (e.g. mostly less than 0.01 for $T = 30$). This might tip the balance slightly in favour of using the probit transform although the decision is marginal based simply on the accuracy of the approximation.

Figure 11 plots contours for the ratio of the approximate to actual annuity values for a male aged 65 under the three transforms using both the linear and quadratic approximations. Again the log transform is relatively poor. The probit transform produces slightly better results than the logistic transform for both the linear and quadratic approximations.

These observations plus the additional analytical properties associated with the probit transform led us to prefer the probit linear approximation for developing a delta hedging strategy.

B Numerical approximation of the derivatives of a function

For a general, twice-differentiable function $g(x_1, x_2)$, for small h_1 and h_2 :

$$\frac{\partial g}{\partial x_1}(x_1, x_2) \approx (g(x_1 + h_1, x_2) - g(x_1, x_2)) / h_1$$

$$\frac{\partial g}{\partial x_2}(x_1, x_2) \approx (g(x_1, x_2 + h_2) - g(x_1, x_2)) / h_2$$

$$\frac{\partial^2 g}{\partial x_1^2}(x_1, x_2) \approx (g(x_1 + h_1, x_2) - 2g(x_1, x_2) + g(x_1 - h_1, x_2)) / h_1^2$$

$$\frac{\partial^2 g}{\partial x_2^2}(x_1, x_2) \approx (g(x_1, x_2 + h_2) - 2g(x_1, x_2) + g(x_1, x_2 - h_2)) / h_2^2$$

$$\text{and } \frac{\partial^2 g}{\partial x_1 \partial x_2}(x_1, x_2) \approx \left(g(x_1 + h_1, x_2 + h_2) - g(x_1 - h_1, x_2 + h_2) - g(x_1 + h_1, x_2 - h_2) + g(x_1 - h_1, x_2 - h_2) \right) / 4h_1h_2.$$

As approximations these formulae are not unique, but they are known to be effective in a variety of circumstances.

C Proof of equation (4)

Let $Z \sim N(0, 1)$ be a standard normal random variable that is independent of the process $K(t)$.




Figure 10: Approximate versus actual values for the spot survival probabilities: $\hat{p}(0, T, x, k) - p(0, T, x, k)$. Top row: $T = 5$ years. Middle row: $T = 10$ years. Bottom row: $T = 30$ years. Left-hand column: Probit transformation, $\Phi^{-1}(x)$. Middle column: Logistic transformation, $\log(x/(1 - x))$. Right-hand column: log transformation, $\log(x)$. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.

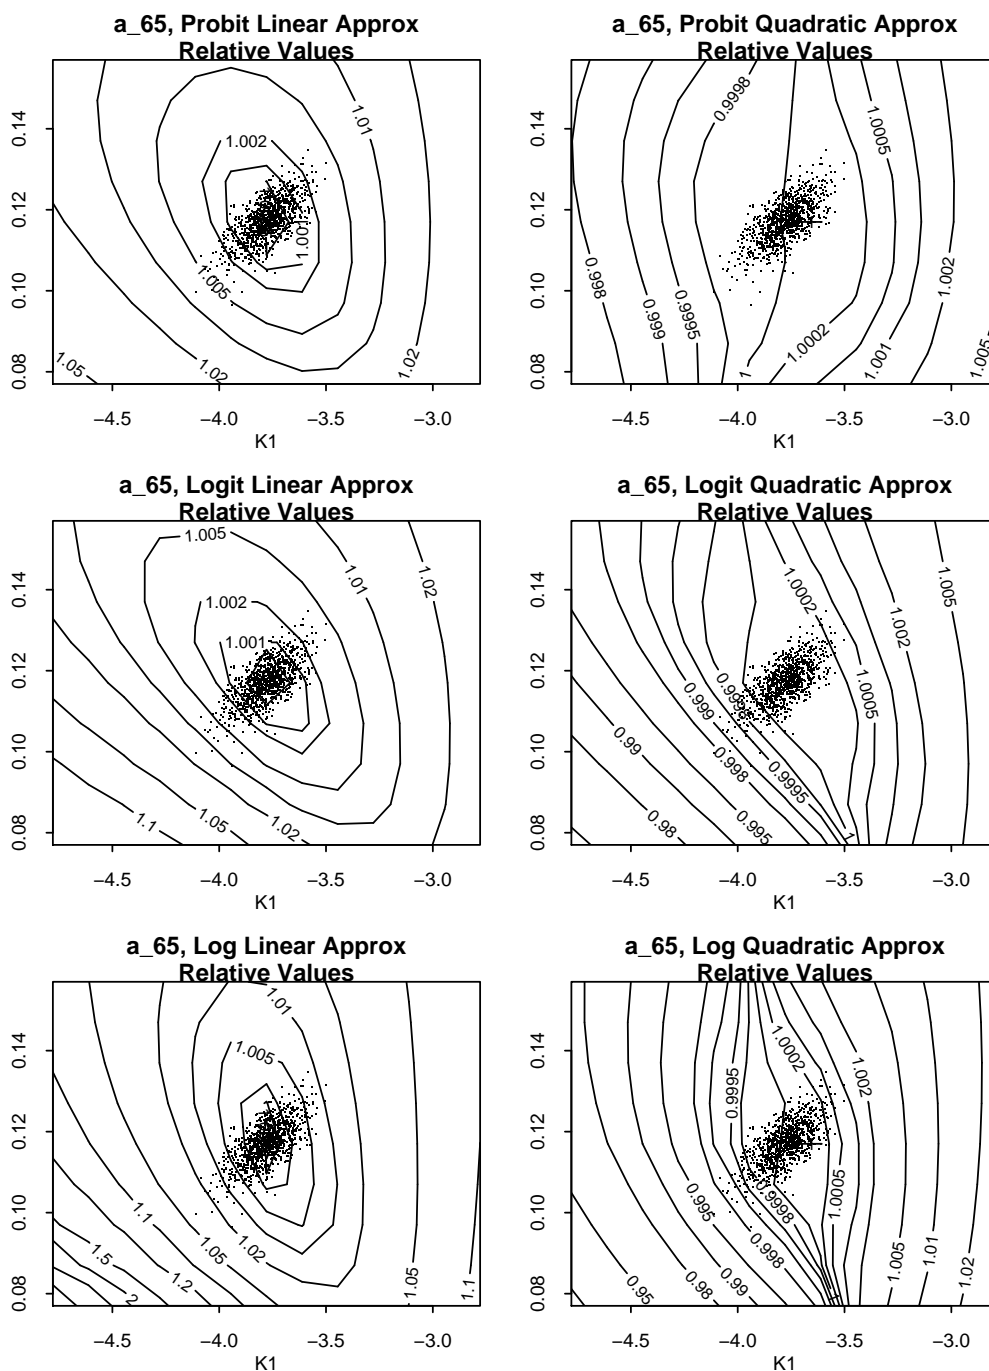


Figure 11: Approximate versus actual values for the annuity function: $\hat{a}_{65}(k)/a_{65}(k)$. Left-hand plots: Linear approximation. Right-hand plots: Quadratic approximation. Top row: Probit transformation, $\Phi^{-1}(x)$. Middle row: Logistic transformation, $\log(x/(1-x))$. Bottom row: log transformation, $\log(x)$. The cloud of dots shows 1000 simulated values of the pair $(K_1(\tau), K_2(\tau))'$ for $\tau = 20$ years after 2008.

Then

$$\begin{aligned}
 p(t, t+T, x-t, K(t)) &\approx \hat{p}(t, t+T, x-t, K(t)) \\
 &= \Phi \left(D_0(T) + D_1(T)'(K(t) - \hat{k}) \right) \\
 &= Pr \left(Z \leq D_0(T) + D_1(T)'(K(t) - \hat{k}) \right) \\
 &= E \left[I_{Z \leq D_0(T) + D_1(T)'(K(t) - \hat{k})} \mid \mathcal{M}_t \right]
 \end{aligned}$$

where $I_{Z \leq D_0(T) + D_1(T)'(K(t) - \hat{k})}$ is equal to 1 if $Z \leq D_0(T) + D_1(T)'(K(t) - \hat{k})$ and 0 otherwise, and \mathcal{M}_t represents the history of the mortality process up to time t , which here can be replaced by $K(t)$ because of the Markov nature of the model.

We wish to calculate

$$\begin{aligned}
 p^{FUT}(0, t, t+T, x-t, K(0)) &= E [p(t, t+T, x-t, K(t)) \mid \mathcal{M}_0] \\
 &\approx E [\hat{p}(t, t+T, x-t, K(t)) \mid \mathcal{M}_0] \\
 &= E \left[\Phi \left(D_0(T) + D_1(T)'(K(t) - \hat{k}) \right) \right] \\
 &= E \left[I_{Z \leq D_0(T) + D_1(T)'(K(t) - \hat{k})} \mid \mathcal{M}_0 \right] \\
 &= Pr \left(Z - D_0(T) - D_1(T)'(K(0) - \hat{k}) \leq 0 \mid \mathcal{M}_0 \right).
 \end{aligned}$$

Now $K(t)$ has the same distribution as $K(0) + \nu t + \sqrt{t}CW$ where the 2×1 vector W is independent of Z and has a standard bivariate normal distribution. Hence

$$\begin{aligned}
 E \left[Z - D_0(T) - D_1(T)'(K(0) - \hat{k}) \right] &= -D_0(T) - D_1(T)'(K(0) - \hat{k}) \\
 \text{and } Var \left[Z - D_0(T) - D_1(T)'(K(0) - \hat{k}) \right] &= 1 + D_1(T)'VD_1(T)t \\
 \Rightarrow p^{FUT}(0, t, t+T, x-t, K(0)) &\approx \Phi \left(\frac{D_0(T) + D_1(T)'(K(0) - \hat{k})}{\sqrt{1 + D_1(T)'VD_1(T)t}} \right).
 \end{aligned}$$

D Probit approximation to logistic mortality rates

The CBD model states that $\text{logit } q(t, x) = K_1(t+1) + K_2(t+1)(x - \bar{x})$. Viewed from a specific start date (say time 0) and for a specific t and x , this can be written as

$$q(t, x) \mid \mathcal{M}_0 \equiv \frac{e^{\alpha + \beta Z}}{e^{\alpha + \beta Z} + 1}$$

where $Z \sim N(0, 1)$, $\alpha = E[\text{logit } q(t, x)]$ and $\beta = \sqrt{Var[\text{logit } q(t, x)]}$.

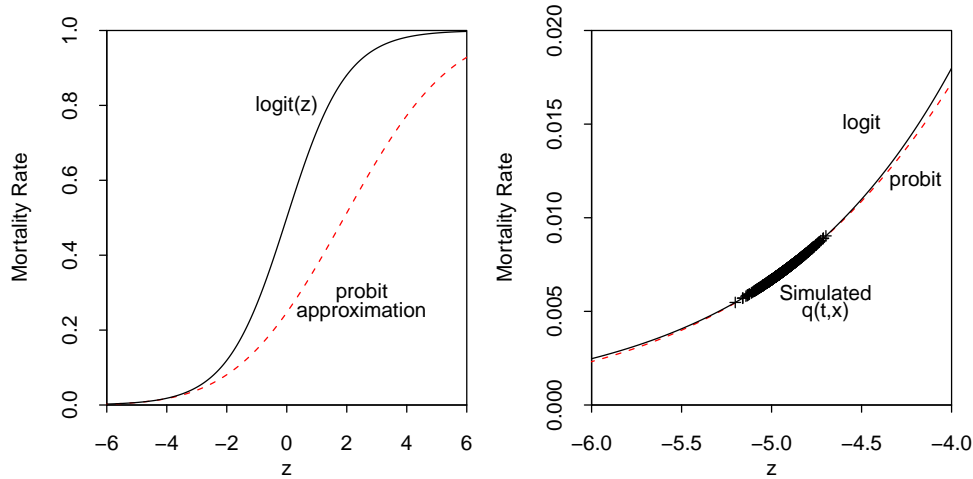


Figure 12: Logistic function $e^z/(e^z + 1)$ and a probit approximation. Probit approximation matches the value and gradient of the logistic function at the expected value for $q(t, x)$ for $x = 65$ and $t = 20$ (i.e. retrospective mortality rate in 2029 for age 65). Left hand plot: logistic and probit functions over a full range of z . Right-hand plot: detail of left hand plot with 1000 simulated values of $(K_1(21) + K_2(21) \times (65 - \bar{x}), q(20, 65))$ (crosses) given $K(0)$.

Now let $f(z) = \Phi(a + bz)$, and define a and b to match $q(t, x; z)$ and $\partial q(t, x; z)/\partial z$ at $z = 0$: that is,

$$a = \Phi^{-1}\left(\frac{e^\alpha}{e^\alpha + 1}\right)$$

$$\text{and } b = \frac{\beta}{\phi(a)} \frac{e^\alpha}{(e^\alpha + 1)^2}.$$

The approximation is illustrated in Figure 12. The probit approximation has been calibrated to match the expected value of $q(t, x)$ in the year 2029 ($t = 20$) given the information available at the end of 2008. From the left hand plot we can see that the probit function does not appear to be an especially good approximation. However, from the right hand plot where we zoom in on the left-hand tail of the functions, we see that the approximation is very accurate in the region of interest: that is, where we find typical simulated values of $q(t, x)$ generated by $K(t + 1)$ for $t = 20$.

Finally, it can be noted that the values of $D_0(T, x)$, $D_{1,1}(T, x)$ and $D_{1,2}(T, x)$ for $T = 1$ and $x = 65$ reported in Table 1 can be verified analytically and independently by deriving appropriate values for a and b as described above and then applying Result 1 in the main body of the paper to move from time 21 back to time 20.

E Alternative approximations for $q^F(s, t, x)$

E.1 Approximation 2: deterministic extrapolation

Let $K(0) = k = (k_1, k_2)'$.

The Deltas require the calculation of the partial derivatives of $q^F(0, t, x; k)$ with respect to k_1 and k_2 the components of $k = (k_1, k_2)$. If we assume that $K(u)$ equals $k + \nu u$ instead of being stochastic, then

$$q^F(0, t, x; K(0) = k) = \frac{\exp(a'K(t+1))}{1 + \exp(a'K(t+1))} \approx \frac{\exp(a'(k + (t+1)\nu))}{1 + \exp(a'(k + (t+1)\nu))} = \tilde{q}(t, x)$$

with $a' = (a_1, a_2)$ and $a_1 = 1$ and $a_2 = (x - \bar{x})$. Note that this approximation is the median of the true distribution at $t + 1$ of $q(t, x)$. Using this approximation is straightforward to calculate the approximate derivatives of $q(0, t, x; k)$: that is,

$$\begin{aligned} \Delta_1 = \frac{\partial q}{\partial k_1} &= \frac{a_1 \exp[a'(k + (t+1)\nu)]}{\left(1 + \exp[a'(k + (t+1)\nu)]\right)^2} = a_1 \tilde{q}(t, x)(1 - \tilde{q}(t, x)) \\ \text{and } \Delta_2 = \frac{\partial q}{\partial k_2} &= \frac{a_2 \exp[a'(k + (t+1)\nu)]}{\left(1 + \exp[a'(k + (t+1)\nu)]\right)^2} = a_2 \tilde{q}(t, x)(1 - \tilde{q}(t, x)). \end{aligned}$$

Note that the ratio of Δ_1 to Δ_2 under this approximation (that is, 1 to $x - \bar{x}$) will be the same as Approximation 1.

E.2 Approximation 3: series expansion

Given $K(0)$, $K(t+1) = K(0) + \nu(t+1) + \sqrt{t+1}CZ$ where $V = CC'$ is the annual variance-covariance matrix of the random walk model for $K(t)$, and Z is a standard bivariate normal random vector.

Now $\text{logit } q(t, x) = a'K(t+1)$ where $a' = (1, x - \bar{x})$. Therefore,

$$\text{logit } q(t, x) \sim N(m(0, t, x), s^2(0, t, x))$$

where $m(0, t, x) = a'(K(0) + \nu(t+1))$ and $s^2(0, t, x) = (t+1)a'Va$. We also have the series expansion $e^w/(1+e^w) = \sum_{j=1}^{\infty} (-1)^{j-1} e^{jw}$. Thus

$$\begin{aligned} E[q(t, x)|K(0)] &= \sum_{j=1}^{\infty} (-1)^{j-1} \exp[jm(0, t, x) + \frac{1}{2}j^2s^2(0, t, x)] \\ &\approx \sum_{j=1}^{2M} (-1)^{j-1} \exp[jm(0, t, x) + \frac{1}{2}j^2s^2(0, t, x)]. \end{aligned} \quad (7)$$

Often $M = 2$ or $M = 3$ will give accurate approximations, but using, for example, $M = 5$ imposes very little computational burden.¹⁰

We then have

$$\begin{aligned} \frac{\partial}{\partial m(0, t, x)} E[q(t, x)|K(0)] &= \sum_{j=1}^{\infty} (-1)^{j-1} j \exp[jm(0, t, x) + \frac{1}{2}j^2 s^2(0, t, x)] \\ &\approx \sum_{j=1}^M (-1)^{j-1} j \exp[jm(0, t, x) + \frac{1}{2}j^2 s^2(0, t, x)] \end{aligned}$$

and this then, with $K(0) = (k_1, k_2)'$, allows us to compute

$$\begin{aligned} \Delta_1 = \frac{\partial}{\partial k_1} E[q(t, x)|K(0)] &= \frac{\partial E[q(t, x)|K(0)]}{\partial m(0, t, x)} \frac{\partial m(0, t, x)}{\partial k_1} \\ &= a_1 \frac{\partial E[q(t, x)|K(0)]}{\partial m(0, t, x)} \\ &= \frac{\partial E[q(t, x)|K(0)]}{\partial m(0, t, x)}, \end{aligned}$$

since $a_1 = 1$. Similarly,

$$\Delta_2 = \frac{\partial}{\partial k_2} E[q(t, x)|K(0)] = (x - \bar{x}) \frac{\partial E[q(t, x)|K(0)]}{\partial m(0, t, x)}.$$

(Again we can note that the ratio of Δ_1 to Δ_2 is 1 to $(x - \bar{x})$.)

Approximation 3 seems to be the preferred approach given that the forward price and its Greeks can be calculated accurately and rapidly.

¹⁰Lower and upper bounds for the approximation in equation (7) are $\sum_{j=1}^{2M} (-1)^{j-1} \exp[jm(0, t, x) + \frac{1}{2}j^2 s^2(0, t, x)]$ and $\sum_{j=1}^{2M-1} (-1)^{j-1} \exp[jm(0, t, x) + \frac{1}{2}j^2 s^2(0, t, x)]$ respectively and can be used to check the accuracy of the approximation. The best approximation can be achieved when $2Mm(0, t, x) + 2M^2 s^2(0, t, x)$ is minimised.