Pricing q-Forward Contracts: An evaluation of estimation window and pricing method under different mortality models

Pauline M. Barrieu *
London School of Economics and Political Science

Luitgard A. M. Veraart †
London School of Economics and Political Science

April 2014

This is an Accepted Manuscript of an article published by Taylor & Francis Group in Scandinavian Actuarial Journal on 05/06/2014, available online: http://dx.doi.org/10.1080/03461238.2014.916228

Abstract

The aim of this paper is to study the impact of various sources of uncertainty on the pricing of a special longevity–based instrument: a q-forward contract. At the expiry of a q-forward contract, the realized mortality rate for a given population is exchanged in return for a fixed (mortality) rate that is agreed at the initiation of the contract. Pricing a q-forward involves determining this fixed rate. In our study, we disentangle three main sources of uncertainty and consider their impact on pricing: model choice for the underlying mortality rate, time-window used for estimation and the pricing method itself.

Key words: longevity risk, q-forward, model uncertainty, estimation window, pricing method.

1 Introduction

Sustained improvements in longevity have been producing a considerable number of new challenges at multiple societal levels. Pension funds are one of the most highly publicised industries to be impacted by rising longevity. Also the insurance industry is facing a number of specific challenges related to longevity.

Under the new European regulatory environment for the insurance industry, Solvency II, it becomes a

*London School of Economics and Political Science, Department of Statistics, Houghton Street, London WC2A 2AE, UK, p.m.barrieu@lse.ac.uk

†London School of Economics and Political Science, Department of Mathematics, Houghton Street, London WC2A 2AE, UK, l.veraart@lse.ac.uk

requirement to measure and evaluate longevity risks. As a consequence, the level of capital required for longevity is increasing and this creates the need for some longevity risk management solutions, such as the transfer of part of the longevity risk to reinsurers and the financial markets.

Several attempts have been made to transfer longevity risk to institutional investors, see Lane Clark & Peacock LLP (LCP) (2013). A very successful transaction of this type so far occurred in February 2012 between AEGON and Deutsche Bank, see Amori et al. (2012) for more details. Such transactions remain however exceptional at this stage. In the absence of both theoretical consensus and established industry practice, the transfer of longevity risk is a difficult process to understand and therefore to manage. Indeed, due to the long-term nature of the inherent risks, accurate longevity projections are unreliable, and modelling the embedded risks (such as interest rate risk and credit risk) remains challenging.

Focusing simply on the longevity modelling, the difficulty comes from various aspects. On the one hand, there is a multiplicity of models for the underlying longevity risk. There is no particular consensus either on the right model to use or on the right data set to be used for the calibration of the model. Moreover, in the case of longevity, more data, i.e. longer data, does not necessarily mean more reliable projections. The underlying risk is a trend risk.

On the other hand, the pricing of a given instrument is also challenging given the nature of the underlying risk, i.e., related to a mortality rate, which is not traded, and the immaturity of the market, which is highly illiquid. The question of the ”right” valuation formula remains open.

In this paper we aim at understanding the impact of various sources of uncertainty on the pricing of longevity-based instruments. To do so, we focus on one of the simplest type of longevity contracts: the q-forward, first introduced by JP Morgan in 2007, see The Life & Longevity Markets Association (2010)[1]. Such a contract specifies a maturity date at which the realized mortality rate for a given population is exchanged in return for a fixed (mortality) rate that is agreed at the initiation of the contract.

We discuss the influence of the choice of the model for mortality rates, that of the corresponding estimation window and that of the pricing rule on the fixed rate of the q-forward. Our objective is not to come up with the pricing method for q-forwards but more to disentangle the impact of the various uncertain aspects (model, estimation window, pricing method) on the price itself to see what is important or what matters the most. In an empirical study we find that the estimation window has a crucial effect on the overall price and the difference in price coming from different choice in the estimation window can even be more severe than the difference between some alternative stochastic mortality models or different actuarial pricing rules.

The rest of the paper is structured as follows: In Section 2 we summarize the main characteristics of a q-forward, in Section 3 we discuss several approaches that can be used for pricing it. Section 4 reviews three classical approaches for modelling mortality. In Section 5 we demonstrate in an empirical study on mortality data from England and Wales the influence of different pricing methods, estimation windows and mortality models on the price of a q-forward. Finally, Section 6 concludes.

[1]The Life and Longevity Markets Association (LLMA in short) has been established as an association in order to promote a liquid traded market in longevity and mortality-related risk. The LLMAs members are key players in the field: AVIVA, AXA, Deutsche Bank, J.P. Morgan, Legal & General, Morgan Stanley, Munich Re, Pension Corporation, Prudential PLC, RBS and Swiss Re.

2 Notation and characteristics of q-forward contracts

Let us start by introducing some notation, recalling the main characteristics of a q-forward contract and discussing how these products can be used for hedging purposes.

Following Cairns et al. (2009), we denote by $m(t, x)$ the crude death rate for individuals of age x in calendar year t, i.e.,

$$m(t, x) = \frac{\#\text{deaths during calendar year } t \text{ of individuals aged } x \text{ at last birthday}}{\text{average population during calendar year } t \text{ aged } x \text{ at last birthday}}.$$

Furthermore, for the mortality rate, denoted by $q(t, x)$, we assume as in Cairns et al. (2009) that

$$q(t, x) = 1 - e^{-m(t,x)}.$$

This mortality rate is the probability that someone aged x at time t will die between times t and $t + 1$. This rate is typically used as underlying risk for longevity-based instruments and in particular for q-*forwards*. These were first introduced by JP Morgan who has been particularly active in trying to establish a benchmark for the longevity market[2]. These contracts are based upon a given population or often a standardized mortality/survival index that draws upon on either the death probability or survival rate as quoted now on the website of the LLMA[3]. For this paper, we will focus on q-forwards based on death probabilities.

Definition 2.1. *"A q-forward is an agreement between two counterparties to exchange at a future date (the maturity of the contract) an amount equal to the realized mortality rate of a given population at that future date (the floating leg), in return for a fixed mortality rate agreed upon at the inception of the contract (the fixed leg)", The Life & Longevity Markets Association (2010).*

Hence, we can interpret the q-forward as a zero coupon swap that exchanges fixed mortality for realized (i.e., floating) mortality of a specified population at a maturity date T. The fixed mortality rate K is determined at the signing time 0 as to make the contract fair for both counterparties. Only at maturity of the contract there is an exchange of cash flow. Formally, the Net Payoff Amount (NPA) of a q-forward contract is given by

$$\text{NPA}(T) = z(q(T) - K), \tag{1}$$

where z denotes the notional agreed at time 0 and $q(T)$ is the realised mortality at time T of the reference population. Note that the NPA can take positive or negative values. The buyer of the q-forward pays the fixed rate and the seller of the q-forward pays the realized rate.

Following the discussion in Coughlan et al. (2007), we can identify various market participants who might be interested in trading these q-forward contracts. Having a better understanding of the potential market will give further insights for a realistic pricing mechanism. Note however that the current state of

[2]JP Morgan developed the LifeMetrics longevity risk analysis platform in 2007.
[3]See www.llma.org.

the market is still embryonic and crucially lacks of liquidity.
One of the characteristics of the q-forward contracts if that they can be used for hedging both mortality and longevity risk. Mortality risk refers to the risk that mortality rates are higher than expected whereas longevity risk refers to the risk that mortality rates are lower than expected.
One the one hand, consider a life insurance company. It can hedge its mortality risk by buying a q-forward contract and hence by paying the fixed (mortality) rate and receiving the realized mortality rate. If the realized mortality rate is higher than the fixed rate at maturity of the contract, it gets a positive payment which can be used to settle the higher death benefits arising from the life insurance contracts.
On the other hand, consider a pension fund that needs to hedge their longevity risk. This can be done by entering into a contract in which it receives a realized survival rate and pays a fixed survival rate. If realized survival rates are higher than the fixed rate, the pension fund receives a positive payment which can be used to settle the higher pension payments that are necessary. Since the survival rate is just 1-q, where q is the mortality rate, receiving realized survival rate and paying fixed survival rate is equivalent to paying realized mortality rate and receiving fixed mortality rate. Hence, this hedge involves selling q-forward contracts.

In principle there are hedgers who are interested in both sides of the trade of a q-forward contract. The main problem, however, is that the magnitude of liabilities that has to be hedged by pension funds etc. is usually much higher than the corresponding one by life insurance companies, see Loeys et al. (2007). Hence there are not enough natural buyers for q-forward contracts. Moreover, the exposure to longevity risk and mortality risk typically applies to different age bands and therefore creates an asymmetry between the needs of potential hedgers. This will certainly affect their market prices if they are ever traded on an established financial market.

Given the fact that no cash flow is exchanged between the buyer and the seller of a q-forward at the signing date 0, the value of the q-forward, denoted by $\mathrm{Value}_0(\cdot)$, should be equal to 0 at that particular time to be fair for both parties, and this fact does not depend on the valuation rule. So, in all generality, the following relationship prevails:

$$\mathrm{Value}_0(\mathrm{NPA}(T)) = \mathrm{Value}_0(z(q(T) - K)) = 0.$$

The main question is then the choice of the valuation rule, i.e., the choice of $\mathrm{Value}_0(\cdot)$ which will then imply the fair value for the fixed mortality rate K which we refer to as forward price.

3 Some pricing methods for q-forward contracts

In this section, we review some of the classical financial and actuarial pricing formulae and look at how they can be applied to the valuation of q-forward contracts. Our aim is certainly not to be exhaustive here but more to introduce some benchmarks for the valuation.
In the following, uncertainty will be described by a standard probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\mathbb{P}$ is the reference probability measure (typically the historical or physical probability measure).

3.1 The role of interest rates

Longevity-linked securities are typically characterised by very long maturities (40 years and beyond). The vast majority of such contracts has an embedded interest rate risk. An appropriate model for interest rates is therefore crucial for pricing and hedging longevity-based instruments. This has been discussed in detail in Barrieu et al. (2012).

In this paper, however, our objective is more focused. We are interested in studying the impact of various sources of uncertainty of the pricing. To do so, we trim the framework as much as we can as to focus on the model uncertainty and uncertainty of the pricing rule. For this purpose, we have chosen a particular type of instrument, the q-forward, whose cash flow structure is such that, for any of the pricing rules we introduce below, the role of interest rates is limited. Obviously, this will not be the case for more complex longevity based products. Having a good understanding of interest rates dynamics, especially in the long term, is then essential and is very complex.

Without loss of generality in the computation, we assume that there is a riskless asset available in the market, whose time-t price is given by e^{rt} with deterministic and constant interest rate $r \geq 0$.

One could include stochastic and dynamic interest rates and then change probability measure (e.g. from the risk-neutral to the forward measure) in the computations of the core expected value. This would then also involve modelling zero-coupon bond prices. To keep the analysis focused, however, we only consider the deterministic and constant case.

3.2 Arbitrage-free valuation and risk-neutral pricing

Since the structure of the q-forward and its payoff (1) are similar to those of a classical forward contract written on a standard underlying asset (such as a stock or a commodity), a natural approach for the valuation might seem to be standard risk-neutral pricing. For this one would take a risk-neutral pricing measure $\mathbb{Q}$ which is equivalent to $\mathbb{P}$ and one would evaluate the forward price by considering

$$\mathrm{Value}_0(\mathrm{NPA}(T)) = \mathbb{E}_{\mathbb{Q}}\left[\exp(-rT)(z(q(T) - K))\right] = 0.$$

This would imply that

$$K = \mathbb{E}_{\mathbb{Q}}\left[q(T)\right].$$

As discussed in many papers, however, see e.g. Cairns et al. (2006a)), in the case of the q-forward, the underlying mortality rate is not tradable and hence the standard arbitrage-free valuation formula for a forward price does not hold true any longer. In particular, $(e^{-rt}q(t))$ does not have to be a $\mathbb{Q}$-martingale. More precisely, as the longevity market is an immature market, based on a non-financial risk, the classical methodology of risk-neutral pricing cannot be used mindlessly. Indeed, the lack of liquidity in the market induces incompleteness, as it is typically the case when non-hedgeable and non-tradable claims exist. Thus, to price the q-forwards, a classical arbitrage-free pricing methodology is inapplicable as it relies upon the idea of risk replication. The replication technique is only possible for markets with high liquidity and for deeply traded assets. It induces a unique price for the contingent claim, which is the cost of the replicating portfolio hedging away the market risk. Hence, in a complete market, the price of the contingent claim is the expected future discounted cash-flows, calculated by the unique risk-neutral probability measure. In contrast, in an incomplete market, such as a longevity-linked securities market, there will be no universal pricing probability measure, making the choice of pricing probability measure crucial. Some papers have been investigating this question in depth, looking at the appropriate characterisation of the market risk premium for the pricing measure, see in particular Bauer et al. (2010), under the assumption that the longevity market is sufficiently mature and that the considered longevity derivative can be replicated using some traded instruments in the market. However, under the current market conditions, it is expected that the historical probability measure will play a key role, as we will see in the following subsection. For this reason, we will not explore the risk-neutral valuation formula in Section 5.

3.3 Some classical actuarial pricing methods

In this section we follow the seminal papers Deprez & Gerber (1985) and Kremer (1986) and introduce some standard actuarial pricing methods to evaluate the NPA at time 0 and to determine the forward price K. Most of the classical approaches in actuarial pricing involve the expected value under the historical probability measure. The simplest case is to use this expectation directly for premium calculations.

Fair premium or net premium principle: It assumes that the time-0 value of the NPA can be derived by considering the expectation of the NPA under the physical probability measure, hence

$$\mathrm{Value}_0(\mathrm{NPA}(T)) = \mathbb{E}_{\mathbb{P}}\left[\exp(-rT)z(q(T)-K)\right] = 0,$$

which implies that

$$K = \mathbb{E}_{\mathbb{P}}\left[q(T)\right].$$

Using this expectation directly as a price can be justified in some situations by appealing to the Strong Law of Large Numbers, see e.g. Mikosch (2009), which in its classical form states that the sample mean $\frac{1}{n}\sum_{i=1}^{n} X_i$ of i.i.d. random variable $X_1, X_2, \dots$ with $\mathbb{E}[X] = \mu$ converges to their expectation μ almost surely as $n \to \infty$. Note that the independence assumption can be weakened and also the assumption on identical distribution can be relaxed. In the case of q-forward contracts this would mean that a large number of q-forward contracts needs to be involved in order for such a limit result to be applicable. This is a limitation at a time where a market is not yet established. Still it seems a less restrictive assumption than e.g. having to assume replicability of these derivatives.

The fair premium is often used as a core element of the price. Usually a risk premium is then added. It may be proportional to the standard deviation which then leads to the standard deviation principle. A general discussion on the risk premium, and in particular on its sign, is provided in the next subsection.

Standard deviation principle: This principle assumes that the time-0 value of the NPA is the sum of the expectation of the NPA under the physical probability measure plus a constant λ times the standard deviation of the NPA. Hence,

$$\text{Value}_0(\text{NPA}(T)) = \mathbb{E}_{\mathbb{P}}\left[\exp(-rT)z(q(T)-K)\right] + \lambda\sqrt{\mathbb{V}_{\mathbb{P}}\left[\exp(-rT)z(q(T)-K)\right]} = 0,$$

where $\mathbb{V}_{\mathbb{P}}(\cdot)$ denotes the variance under the probability measure $\mathbb{P}$ and hence $\sqrt{\mathbb{V}_{\mathbb{P}}(\cdot)}$ is the corresponding standard deviation. This implies that

$$K = \mathbb{E}_{\mathbb{P}}\left[q(T)\right] + \lambda\sqrt{\mathbb{V}_{\mathbb{P}}\left[q(T)\right]}.$$

The standard deviation approach can be re-interpreted in terms of a market risk premium as λ can be related to a measure of the Sharpe ratio of the risky mortality rate $q(T)$ as we will see in the next subsection. The choice of λ can be related to various factors, including some external ones such as the regulation and the imposed safely loadings, or to some internal ones such as the degree of basis risk associated with the choice of this particular mortality index versus the existing portfolio in case of a hedging strategy using q-forwards.

Principle of zero utility: Another standard principle of premium calculation is the so-called principle of zero utility. It is also referred to as indifference pricing in a financial context. Indeed, in an incomplete market framework, where perfect replication is no longer possible, a more appropriate strategy involves utility maximisation, see Carmona (2009) for further details. According to this, the maximum price that an agent is willing to pay depends on their individual preference and it is chosen such that the agent is indifferent between paying this price and obtaining the asset or not engaging in this transaction at all. The obtained price gives an upper bound to the price the buyer is ready to pay for a given instrument. A transaction might or might not occur at this price depending on the preferences of the other market participants. More precisely, for the buyer, given a utility function U and an initial wealth[4] of W_0, the indifference buyer price K should be such that the following relationship prevails:

$$\mathbb{E}_{\mathbb{P}}\big(U(W_0\exp(rT) + z(q(T)-K)\big) = \mathbb{E}_{\mathbb{P}}\big(U(W_0\exp(rT))\big). \tag{2}$$

Hence, if we assume the utility functions to be an exponential utility[5], i.e., $U(y) = -\exp(-\gamma y)$ where $\gamma > 0$ is the constant coefficient of absolute risk aversion and we set the initial wealth to be constant we can solve (2) for K and obtain:

$$K = -\frac{1}{\gamma z}\log\left(\mathbb{E}_{\mathbb{P}}[\exp(-\gamma z q(T))]\right).$$

Note that the forward price K, obtained here as an indifference buyer price, can also be reinter-

[4] We consider capitalised initial wealth (up to time T) as to avoid time inconsistency in terms of cash flows.

[5] Exponential utility functions have been widely used in the financial and economic literature. Several facts may justify their relative importance compared to other utility functions, in particular the absence of constraints on the sign of the considered cash flows or the possibility to obtain analytical expressions for indifference prices. They do have limitations, however, one of which is the fact that they correspond to agents with constant absolute risk aversion.

preted in terms of financial risk measures. Indeed it corresponds to (-1) times the entropic risk measure of the future mortality rate $q(T)$, with a risk tolerance coefficient dependent on the notional amount of the forward contract. More details about entropic risk measures can be found in Föllmer & Schied (2002) for instance.

Remark 3.1 (Applying the principle of zero utility)**.** In the description of the principle of zero utility we have considered the q-forward contract as a stand-alone contract. This view corresponds to the view of an investor in the financial market who does not have any particular exposure to longevity or mortality risk. As discussed earlier, q-forwards can be used for hedging purposes and then their valuation would need to be considered together with the underlying risk that is being hedged. We have seen that the principle of zero utility leads to a forward price that depends on the initial wealth. Similarly such a price would also depend on the underlying risk that is being hedged when these risks are considered simultaneously when the q-forward is priced. Hence a hedger would arrive at an indifference price that is different from that of an investor. For ease of exposition we only consider the investor's perspective here.

Remark 3.2 (Comparison of the actuarial pricing rules)**.** When comparing the three actuarial pricing rules we see that for the fair premium and the standard deviation principle the forward price K is independent of the notional z. This is no longer the case for the principle of zero utility. There K depends on the notional amount in a non-linear way. Note that for exponential utility with parameter γ the forward price K depends on the risk aversion parameter γ and the notional z only via their product γz. Hence if we consider two market participants with different risk aversion parameters $\gamma_1 < \gamma_2$, we see that as long as their corresponding notional z_1, z_2 are such that $\gamma_1 z_1 = \gamma_2 z_2$ they both arrive at the same purchasing price but for different quantities. The more risk averse market participant with the larger parameter γ_2 would buy a smaller quantity than the less risk averse market participant at the same price.

3.4 Discussion on the risk premium

The different pricing principles we have discussed so far are classical and can be used for various types of financial products. To decide which one is more relevant for pricing q-forward contracts we need to account for the overall structure of a possible q-forward market.

As discussed by Loeys et al. (2007), in general, the market is net short longevity, which means that significantly more market participants are interested in selling q-forward contracts than in buying them. This will imply that prices for q-forward contracts will have to include a (negative) risk premium which makes it attractive for investors to buy them. In particular, it is unlikely that q-forwards will be traded at the fair premium price, i.e., at the expected mortality rate under the physical probability measure, but it is to be expected that their prices will be less than this. The difference between the expected mortality rate and the market price is referred to as *risk premium*. We denote it by R and then

$$K = \mathbb{E}_{\mathbb{P}}\left[q(T)\right] + R,$$

where one would expect that $R < 0$.

Note that the standard deviation principle has already this overall structure and the sign of the corresponding risk premium depends solely on the sign of λ.

The risk premium has been discussed in the literature already. In The Life & Longevity Markets Association (2012), the risk premium is modelled explicitly by assuming that mortality improvements are higher than expected, and hence these higher improvement rates are used to account for the (negative) risk premium. Still it is not clear according to which criterion this higher improvement rate should be chosen. Loeys et al. (2007) also argue that one needs $R < 0$ to make investors interested in these products. Furthermore, he states that a sufficient return to risk, i.e., a certain Sharpe ratio, is necessary for investors to buy these products. Hence, according to such an approach one could fix a certain Sharpe ratio S and solve

$$S = \frac{\mathbb{E}_{\mathbb{P}}[q(T)] - K}{\sqrt{\mathbb{V}_{\mathbb{P}}(q(T))}}$$

for the forward price K. This then yields

$$K = \mathbb{E}_{\mathbb{P}}[q(T)] - S\sqrt{\mathbb{V}_{\mathbb{P}}(q(T))}.$$

Hence we see that this approach is related to the standard deviation principle in which the parameter λ is given by $-S$. This gives a natural framework for choosing the parameter λ in the standard deviation principle.

4 Classical models for mortality

A plethora of mortality models is available in the literature. On the one hand, there are the so-called causal mortality models. They use explanatory variables to provide causes for deaths or mortality improvements. In principle, they can be set up to a significant level of detail and include e.g. genetic information, information about lifestyle, income, medical advances etc..

On the other hand there are data-driven extrapolative models. These models only rely on data of the size of a given population and the corresponding number of deaths for given time intervals and given age groups. No explanatory variables are included. For the purpose of this paper we will look at some classical data-driven models only: First we consider the Lee-Carter model (Lee & Carter, 1992) with a random walk model for the time-dependent function, second the Lee-Carter model with a general ARIMA(p, d, q) model and third the Cairns-Blake-Dowd model (Cairns et al., 2006b) again with a random walk. For more details on the modelling of mortality, we refer to Cairns et al. (2009), who provide a detailed analysis of several stochastic mortality models and discuss how well they can describe empirical data from England, Wales and the USA. In line with these results, we have chosen these three models because they usually provide a good fit to empirical data despite their rather simple structure.

4.1 The Lee-Carter model

In the Lee-Carter model, see Lee & Carter (1992), the log death rate is given by

$$\log(m(t,x)) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)}.$$

Note, that this can be written as $\log(m(t,x)) = \beta_x^{(1)}\kappa_t^{(1)} + \beta_x^{(2)}\kappa_t^{(2)}$, where $\kappa_t^{(1)} \equiv 1\ \forall t$.
Hence, there are two functions $\beta_x^{(1)}$ and $\beta_x^{(2)}$ which model the influence of age on the death rate and the function $\kappa_t^{(2)}$ models time-related effects. One usually observes in empirical data that $\kappa_t^{(2)}$ is decreasing with time, showing that death rates in general decrease over time. Furthermore, $\beta_x^{(1)}$ is typically increasing with age and hence death rates increase with age. $\beta_x^{(2)}$ is usually decreasing with age which implies that death rate improvements tend to be higher at lower ages on the log-scale.
In order to estimate the model one needs to impose further parameter constraints to ensure identifiability. As in Cairns et al. (2009) we set

$$\sum_t \kappa_t^{(2)} = 0, \qquad \sum_x \beta_x^{(2)} = 1.$$

The only time-dependent parameter in this model is $\kappa_t^{(2)}$. To be able to price q-forward contracts based on the Lee-Carter model we need to propose a model for the time-dependent function $\kappa_t^{(2)}$. In the literature this is usually being done by first fitting all three functions occurring in the Lee-Carter model. Then the fitted values for $\hat{\kappa}^{(2)}(t)$ are considered as observations from a time series and then a time series model is fitted to describe the dynamics of the $\hat{\kappa}_t^{(2)}$, see Lee & Carter (1992).

4.1.1 Lee-Carter model with random walk (LC1)

For data from England and Wales Cairns et al. (2009) showed that using a simple random walk model with drift, provides a good empirical fit to the time-dependent function $\hat{\kappa}_t^{(2)}$. Hence, the underlying assumption is that

$$\hat{\kappa}_t^{(2)} = \hat{\kappa}_{t-1}^{(2)} + \mu_{LC} + \sigma_{LC} Z_t,$$

where Z_t are i.i.d. standard normally distributed random variables and $\mu_{LC} \in \mathbb{R}, \sigma_{LC} \geq 0$ are parameters that need to be estimated. In our empirical analysis we will refer to the Lee-Carter model with random walk as LC1-model.

We consider $n+1 \in \mathbb{N}$ years of data labeled $t_0, \dots, t_n$ to fit the random walk model for the $\hat{\kappa}_t^{(2)}$. We consider only yearly data, hence $t_i - t_{i-1} = 1$. The corresponding maximum likelihood estimators

(MLE) for μ_{LC} and σ^2_{LC} are then given by

$$\begin{aligned}\hat{\mu}_{LC} &= \frac{1}{n}\sum_{i=1}^{n}(\hat{\kappa}^{(2)}_{t_i} - \hat{\kappa}^{(2)}_{t_{i-1}}) = \frac{\hat{\kappa}^{(2)}_{t_n} - \hat{\kappa}^{(2)}_{t_0}}{n},\\ \hat{\sigma}^2_{LC} &= \frac{1}{n}\sum_{i=1}^{n}\left(\hat{\kappa}^{(2)}_{t_i} - \hat{\kappa}^{(2)}_{t_{i-1}} - \hat{\mu}_{LC}\right)^2.\end{aligned} \tag{3}$$

Note that the MLE $\hat{\mu}_{LC}$ depends only on the first and the last observation of the time-series. It is clear from this definition already that the choice of the estimation window for the time series will be crucial.

Suppose $T = t_n + \Delta t$ for some $\Delta t > 0$ where usually t_n refers to the current year. Then we have

$$\hat{\kappa}^{(2)}_{t_n+\Delta t}|\hat{\kappa}^{(2)}_{t_n} \sim \mathcal{N}_1(\hat{\kappa}^{(2)}_{t_n} + \mu_{LC}\Delta t, \sigma^2_{LC}\Delta t).$$

Hence, based on this random walk assumption, for fixed time T and age x also the mortality rate $q(T, x)$ is a random variable. We can generate samples from its distribution by generating samples $Y_1, Y_2, \ldots,$ from the $\mathcal{N}_1(\hat{\kappa}^{(2)}_{t_n} + \hat{\mu}_{LC}\Delta t, \hat{\sigma}^2_{LC}\Delta t)$ distribution and setting

$$\begin{aligned}m_i &= \exp\left(\hat{\beta}^{(1)}_x + \hat{\beta}^{(2)}_x Y_i\right),\\ q_i &= 1 - \exp(-m_i).\end{aligned}$$

We can then evaluate all previously discussed pricing methods using Monte Carlo simulations as we will see in the next section.

4.1.2 Lee-Carter model with ARIMA (LC2)

In order to allow some analysis of the effect of the time series model used for the time-dependent functions, we consider a general autoregressive integrated moving average (ARIMA(p, d, q)) time series model for the time series $(\hat{\kappa}^{(2)}_{t_i})_{i=0,\ldots,n}$ in the Lee-Carter model. We refer to this model as LC2-model.

Let $p, q, d \in \mathbb{N}_0$. Recall that an ARIMA(p, d, q) model for a time series (X_t), $t \in \mathbb{Z}$, is defined by

$$(1 - \sum_{i=1}^{p}\phi_i B^i)(1-B)^d X_t = c + (1 + \sum_{i=1}^{q}\theta_i B^i)\epsilon_t, \tag{4}$$

where B is the backshift operator, i.e., $B^i X_t = X_{t-i}$ for all $i \in \mathbb{N}_0$, ϕ_i, $i = 1, \ldots, p$ are the parameters corresponding to the autoregressive part of the model and θ_i, $i = 1, \ldots, q$ are the parameters corresponding to the moving average part of the model and c is a constant. It is further assumed that the polynomials $\phi(z) = 1 - \sum_{i=1}^{p}\phi_i z^i$ and $\theta(z) = 1 + \sum_{i=1}^{q}\theta_i z^i$ do not have any roots for $|z| < 1$, see Hyndman & Khandakar (2008); Hyndman & Athanasopoulos (2013). Furthermore, it is assumed that the error term (ϵ_t) is a white noise process with mean 0 and variance σ^2. Note that a random walk is just a special cases, namely an ARIMA(0, 1, 0) model. For details on ARIMA models, their estimation, forecasting and the selection of appropriate orders p, q and d we refer to Hyndman & Khandakar (2008), Hyndman

& Athanasopoulos (2013) and Brockwell & Davis (2009).

4.2 The Cairns-Blake-Dowd model

More recently, Cairns, Blake and Dowd propose in Cairns et al. (2006b) to model logit $(q(t,x))$ directly rather than $m(t,x)$. Their model is given by

$$\log\left(\frac{q(t,x)}{1-q(t,x)}\right) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}).$$

To make comparison with the Lee-Carter model easier, we rewrite this as

$$\log\left(\frac{q(t,x)}{1-q(t,x)}\right) = \beta_x^{(1)}\kappa_t^{(1)} + \beta_x^{(2)}\kappa_t^{(2)},$$

where $\beta_x^{(1)} \equiv 1$ and $\beta_x^{(2)} = (x-\bar{x})$, where $\bar{x}$ denotes the mean age in the sample. There are no identifiability problems in the model and hence there are no further constraints on the functions used.
This model contains two functions depending on time $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$. Cairns et al. (2006b) model their dynamics over time by considering a two-dimensional random walk model with drift. As in the Lee-Carter model one first estimates the functions $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ from empirical data to obtain $\hat{\kappa}_t^{(1)}$ and $\hat{\kappa}_t^{(2)}$. These estimates are then considered as observations from a two-dimensional time series model which is assumed to be given by

$$(\hat{\kappa}_t^{(1)}, \hat{\kappa}_t^{(2)})^\top = (\hat{\kappa}_{t-1}^{(1)}, \hat{\kappa}_{t-1}^{(2)})^\top + \mu_{CBD} + \sigma_{CBD} Z_t,$$

where Z_t are i.i.d. random vectors from the two-dimensional standard normal distribution and the vector $\mu_{CBD} \in \mathbb{R}^2$ and the matrix $\sigma_{CBD} \in \mathbb{R}^{2\times 2}$ need to be estimated. For the estimation we can again use Maximum-Likelihood estimators which are given by the multivariate version of the formulae provided in (3).
We can proceed as before, with the only difference that we now consider a two-dimensional random walk model. Hence, to forecast $(\hat{\kappa}_{t_n+\Delta t}^{(1)}, \hat{\kappa}_{t_n+\Delta t}^{(2)})^\top$ we use that

$$(\hat{\kappa}_{t_n+\Delta t}^{(1)}, \hat{\kappa}_{t_n+\Delta t}^{(2)})^\top | (\hat{\kappa}_{t_n}^{(1)} \hat{\kappa}_{t_n}^{(2)})^\top \sim \mathcal{N}_2((\hat{\kappa}_{t_n}^{(1)}, \hat{\kappa}_{t_n}^{(2)})^\top + \mu_{CBD}\Delta t, \sigma_{CBD}\sigma_{CBD}^\top \Delta t).$$

As before we can generate (two-dimensional) samples $Y_1, Y_2, \dots$ from the $\mathcal{N}_2((\hat{\kappa}_{t_n}^{(1)}, \hat{\kappa}_{t_n}^{(2)})^\top + \hat{\mu}_{CBD}\Delta t, \hat{\sigma}_{CBD}\hat{\sigma}_{CBD}^\top \Delta t)$ distribution and set

$$q_i = \frac{\exp(Y_i^{(1)} + Y_i^{(2)}(x-\bar{x}))}{1+\exp(Y_i^{(1)} + Y_i^{(2)}(x-\bar{x}))}.$$

As for the Lee-Carter model, this can be used for computing prices via Monte Carlo simulations.

5 Comparison of different valuation methods and impact of uncertainty on pricing

For our empirical study on the impact of various sources of uncertainty on the pricing of q-forward contracts, we use mortality and population data from England and Wales retrieved from the Human Mortality Database[6]. We use yearly data from 1961 until 2009 for males aged between 60 years and 89 years in those years. From those raw data, we compute yearly mortality rates. Figure 1 shows the empirical data. As expected, we clearly see that mortality rates are decreasing over time, and are higher for higher ages. In addition to these clear deterministic patterns we can observe that changes in mortality rates also have a stochastic component.

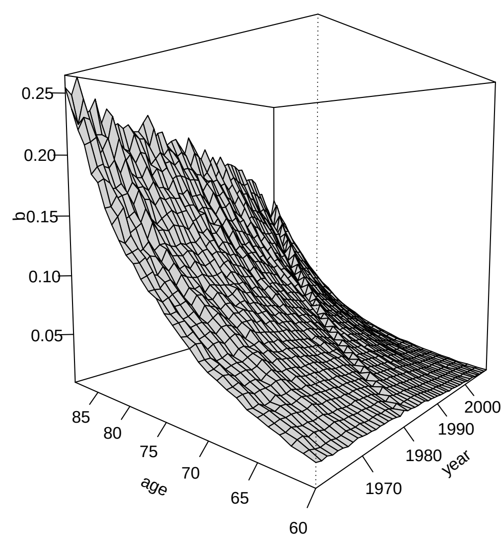

Figure 1: Observed death rates $q(t, x)$ for males with age $x \in \{60, \ldots, 89\}$ in years $t \in \{1961, \ldots, 2009\}$ in England and Wales.

5.1 Fitting the mortality models and uncertainty related to the estimation window

First, we fit the Lee-Carter model to the data from the Human Mortality Database described above. Figure 2 shows the corresponding parameter estimates. The obtained estimates are consistent with empirical results reported in the literature, see Cairns et al. (2009). In particular, $\hat{\beta}_x^{(1)}$ is increasing with age x, corresponding to higher mortality rates at higher ages. $\hat{\kappa}_t^{(2)}$ is decreasing over time t, indicating that mortality rates are generally decreasing over time. Finally, $\hat{\beta}_x^{(2)}$ is decreasing with age x, which implies that mortality rates are improving more at lower ages than at higher ages.

[6] See www.mortality.org.

We then use the estimated parameters $\hat{\kappa}_t^{(2)}$ for some t and consider them as observations from a time series model. As described previously, we first fit a random walk model with drift parameter μ_{LC} and standard deviation parameter σ_{LC}.

Note that the choice of the estimation window appears to be essential and this introduces some additional uncertainty. Indeed, the parameter estimates are crucially dependent on the time period we choose to fit the random walk model. We have chosen two different estimation windows for our analysis: The first is very short and uses only the 6 most recent observations in our data set, i.e., data from 2004-2009. The second is significantly longer and accounts for 21 yearly observations from the years 1989-2009.

If we consider 6 yearly observations from the years 2004-2009 and fit a random walk to the Lee-Carter estimates we obtain $\hat{\mu}_{LC} = -1.0342$, $\hat{\sigma}^2_{LC} = 0.1062$. If we use 21 yearly observations from the years 1989-2009 to fit the random walk we obtain $\hat{\mu}_{LC} = -0.8722$, $\hat{\sigma}^2_{LC} = 0.3053$. Hence, we see that the variance $\hat{\sigma}^2_{LC}$ increases, whereas the trend parameter $\hat{\mu}_{LC}$ is less pronounced for the longer time horizon.

Second we fit a general ARIMA(p, d, q) model to the $\hat{\kappa}_t^{(2)}$ from the Lee-Carter model using the selection method outlined in Hyndman & Khandakar (2008) to find appropriate orders p, q and d. For 6 yearly observations, i.e., from 2004 - 2009, we obtain an ARIMA(2,1,1) model with drift and the parameters, see Definition 4, are estimated to be $\phi_1 = -0.8058$, $\phi_2 = -0.9009$, $\theta_1 = 0.0647$, $c = -0.9801$ and $\sigma^2 = 0.01631$.

For 21 yearly observations, i.e., from 1989 - 2009, we obtain an ARIMA(1, 1, 0) model with drift with parameter estimates $\phi_1 = -0.4545$, $c = -0.8623$ and $\sigma^2 = 0.2406$.

Note that the parameters p, q and d of the ARIMA model are selected by the automatic model selection mechanism outlined in Hyndman & Khandakar (2008), which uses the Akaike Information Criterion (AIC). For 6 yearly observation the chosen model does contain a large number of parameters compared to the available observations. In both cases differencing once is sufficient, i.e., $d = 1$ is selected. For the short time series two autoregressive components and one moving average component is selected. For the longer time series only one autoregressive component and no moving average component is selected. The drift is in both cases strongly negative and it behaves similarly as in the random walk case, i.e. for a shorter estimation window the drift is steeper. In both cases we see that a simple random walk is not selected indicating that a more complicated autocorrelation structure is present in the data.

Third, we fit the Cairns-Blake-Dowd model to our data. Figure 3 shows the parameter estimates. The results are consistent with the results reported in Cairns et al. (2009). Since the age dependent parameter $\beta_x^{(2)}$ is only the difference between x and the sample mean of the considered ages it is increasing in x. (Recall that we consider ages $x \in \{60, 61, \dots, 89\}$ and $\bar{x} = 74.5$.) In contrast to the Lee-Carter model there are now two time-dependent parameters. $\hat{\kappa}_t^{(1)}$ is decreasing with time whereas $\hat{\kappa}_t^{(2)}$ is increasing.

Again we find that the choice of the time window for fitting the (now two-dimensional) random walk model has a crucial effect on the parameter estimates. Considering 6 yearly data between 2004-2009 results in parameter estimates

$$\hat{\mu}_{CBD} = (-0.0353, 0.0009)^\top, \qquad \hat{\sigma}_{CBD}\hat{\sigma}_{CBD}^\top = \begin{pmatrix} 0.0001 & -0.00000007 \\ -0.00000007 & 0.0000005 \end{pmatrix},$$

whereas if we consider the longer time window of 21 observations between 1989 - 2009 we find that

$$\hat{\mu}_{CBD} = (-0.0301, 0.0006)^\top, \qquad \hat{\sigma}_{CBD}\hat{\sigma}_{CBD}^\top = \begin{pmatrix} 0.0004 & 0.00001 \\ 0.00001 & 0.0000009 \end{pmatrix}.$$

Again, the variance increases for the longer time horizon and the trend is slightly weaker.

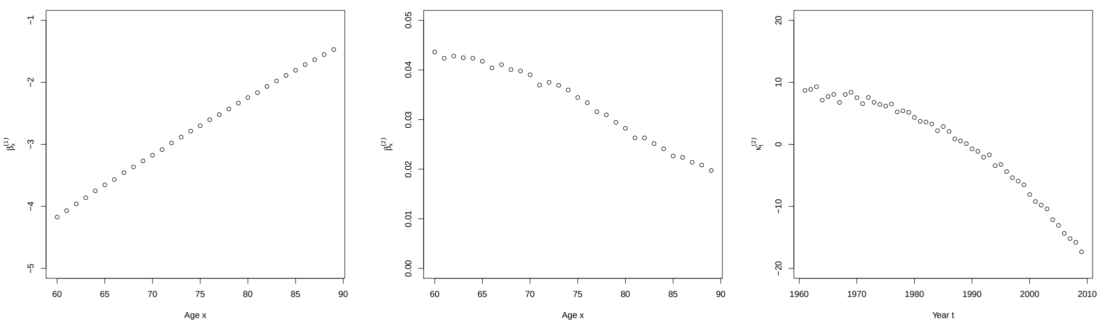

(a) $\hat{\beta}_x^{(1)}$ (b) $\hat{\beta}_x^{(2)}$ (c) $\hat{\kappa}_t^{(2)}$

Figure 2: Parameter estimates for the Lee-Carter model for England and Wales (males) using yearly data from 1961-2009.

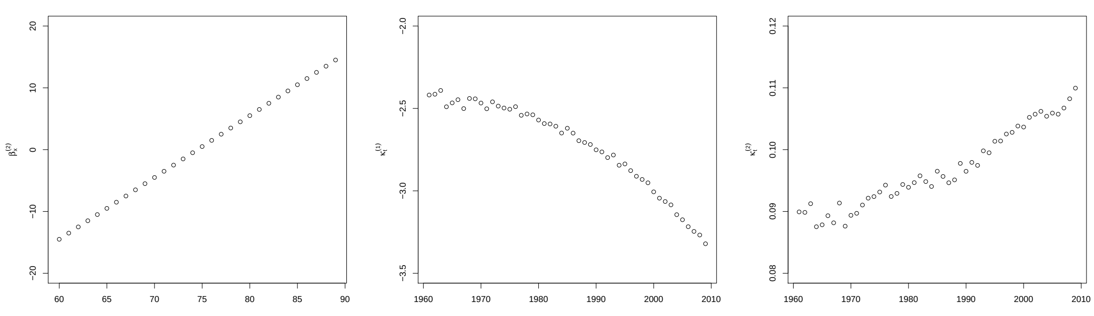

(a) $\beta_x^{(2)}$ (b) $\hat{\kappa}_t^{(1)}$ (c) $\hat{\kappa}_t^{(2)}$

Figure 3: Parameters estimates for the Cairns-Blake-Dowd model for England and Wales (males) using yearly data from 1961-2009.

5.2 Valuing q-forward contracts

We are now able to compute prices of q-forward contracts and assess the impact of different sources of uncertainty. We use:

- Four different actuarial pricing rules:
  - Pricing rule 1 (PR1): Fair premium principle.
  - Pricing rule 2 (PR2): Standard deviation principle with $\lambda = -0.1$.
  - Pricing rule 3 (PR3): Principle of zero utility with $\gamma z = 1$.
  - Pricing rule 4 (PR4): Principle of zero utility with $\gamma z = 10,000$.
- Three different mortality models:
  - Lee-Carter with random walk (LC1),
  - Lee-Carter with *optimal* ARIMA (LC2),
  - Cairns-Blake-Dowd with random walk (CBD).
- Two different estimation windows for the underlying time series model:
  - 6 years (2004 - 2009),
  - 21 years (1989 - 2009).

We consider q-forwards with two different maturities $T \in \{10, 30\}$ and two different underlying ages $x \in \{60, 70\}$. All prices are computed using Monte Carlo simulations which rely on several estimated models. We use a bootstrap method to obtain confidence intervals for the prices, the details of which are outlined in Appendix A.

5.2.1 General comments, age $x = 60$ versus $x = 70$, maturity $T = 10$ versus $T = 30$

All results for age $x = 60$ are presented in Figure 4 and all results for age $x = 70$ are presented in Figure 5. The figures show 95%- confidence intervals for the q-forward prices (drawn as vertical lines with horizontal bars indicating the 2.5%- and the 97.5 % - quantile) and estimated prices (represented by stars) corresponding to the four pricing rules, the three models and the two estimation windows for the two different maturities considered.

Note that Figures 4 and 5 use different scales due to the significant difference in overall price level for q-forwards written on age $x = 60$ compared to age $x = 70$. In general, q-forward prices written on age $x = 60$ are much lower than those written on age $x = 70$ because all mortality models used predict lower mortality rates for the younger age $x = 60$ than for $x = 70$.

Both figures show that across all pricing rules and estimation windows used, prices for maturity $T = 10$ are much higher than prices for maturity $T = 30$. This was expected as all time series models that we fitted to the time-dependent functions in the mortality models included negative drifts. Therefore mortality rates further in the future (here $T = 30$) are predicted to be in general lower than those closer to the starting date (here $T = 10$).

5.2.2 Results for q-forwards written on age $x = 60$

We now look at the prices for age $x = 60$ in Figure 4 in more detail. For all four pricing rules we see that for fixed maturity and fixed estimation window the confidence intervals corresponding to the different mortality models all overlap. In particular, the LC1 model and the CBD model, which both use a random walk to model the time dependent functions, produce particularly similar prices. The corresponding confidence intervals almost coincide for pricing rules PR1, PR2 and PR3.

The confidence intervals are rather wide since they account for several uncertainties related to the estimation of the functions in the mortality model and the estimation and forecasting in the time-series model. In the case of the LC2 model they also account for the uncertainty in the choice of suitable orders p, q and d in the ARIMA(p, d, q) model, since for every bootstrap iteration we determine the optimal parameters p, q, d and do not keep them fixed. The confidence intervals corresponding to the LC2 model are therefore also accounting for model uncertainty within the class of ARIMA models. It is not surprising that they tend to be wider. Furthermore, the ARIMA models selected tend to have more parameters than the random walk model which also explains why corresponding confidence intervals are wider.

Impact of the estimation window

We now look at the influence of the choice of estimation window in more detail. For pricing rules PR1, PR2 and PR3 and the models LC1 and CBD (which both use a random walk model) we find that the confidence intervals corresponding to a 6 years estimation interval or 21 years interval are completely disjoint. Hence, we see that the choice of the estimation window is absolutely crucial for the overall price level. In general prices based on a 21 years estimation window tend to be higher and the corresponding confidence intervals tend to be tighter. For the LC2 model (wich uses a general ARIMA model) we see the same tendency but the confidence intervals for the two different estimation periods still overlap. If we study Figures 2(c), 3(b) and 3(c) we see the reason for the higher price level for an estimation window of 21 years. The time-dependent functions have much steeper slopes when considered over 6 years than over 21 years. Hence when the shorter estimation window is used a stronger decline in mortality rates is predicted which leads to lower prices. The changes in the width of the confidence intervals can be explained by the fact that more available observations will decrease the width of the confidence interval. Still one would have expected that the confidence intervals for 21 observations are a subset of those based on the most recent 6 observations. This is not the case here.

Impact of the pricing rule

Note that some pricing rules depend on parameter(s) which are related to investor's preferences or the notional amount. For the standard deviation principle PR2, we have already seen that $-\lambda$ can be interpreted as the corresponding Sharpe ratio that can be obtained by entering this contract. We have chosen $\lambda = -0.1$ which would represent a Sharpe ratio of 0.1. This seems to be a reasonable assumption compared to empirical findings on products with long maturities.

For the principle of zero utility we have chosen $\lambda z = 1$ for PR3 and $\lambda z = 10,000$ for PR4. Here λ is the constant of absolute risk aversion and z is the notional amount. As discussed in Remark 3.2, the price obtained using the principle of zero utility PR3 and PR4 does depend on the notional amount which is not the case for PR1 or PR2. Suppose the parameter of absolute risk aversion is given by $\gamma = 0.01$, then

PR3 and PR4 would correspond to a notional of 100 and 1,000,000 respectively.
When comparing the four different pricing rules we find very similar prices for PR1, PR2 and PR3 but PR4 yields much wider confidence intervals for the LC2 model and sometimes for the CBD model.

5.2.3 Results for q-forwards written on age $x = 70$

When we now look at the corresponding results for age $x = 70$ in Figure 5 we see that most of the results we obtained for age $x = 60$ carry over but on a higher overall price level. One difference which is striking, however, is that for age $x = 70$ the CBD model has a tendency to yield higher prices than the LC1 and LC2 model. For age $x = 60$ this is not the case. There the LC2 tends to produce the highest prices whereas LC1 and CBD are comparable.

6 Conclusion

In this paper, we have studied the impact of various sources of uncertainty on the pricing of q-forward contracts, disentangling three main sources of uncertainty: model choice for the underlying mortality rate, estimation window and the pricing method itself.

For some classical stochastic mortality models and some classical actuarial pricing rules we find that the choice of the estimation window for the time-dependent functions has a very strong effect on the q-forward price and this choice can even dominate the influence of the choice of the stochastic mortality model or the choice of the pricing rule.

To put these findings into context we must not forget that the stochastic mortality models considered were of rather similar type. We only considered extrapolative, data-driven models which do not attempt to include explanatory variables to explain causes for death or mortality improvements. Furthermore, the time series models used for the extrapolation often included a constant trend. The estimated value of this trend then plays a crucial role in the mortality forecasts. A big risk here is obviously that either the trend is not estimated accurately or that the trend does change over time.

Our findings ultimately highlight the importance of trend risk in pricing longevity products. The fact that the confidence intervals based on the observations of the most recent 21 years or the most recent 6 years are completely disjoint, shows that the time-series model used does not seem to describe the underlying data well. One could argue that either a linear forecasting model is inappropriate for the data or one would need to allow for random changes in this linear trend. Building a stochastic model that adequately deals with these random changes is beyond the scope of this paper. Some attempts have already been made in this direction, see e.g. Sweeting (2011). Nevertheless, it will remain very difficult to estimate such changes from historical data since one would need to use a rather long time-series for this and corresponding data are not very reliable.

It is clear that pricing q-forwards is very delicate due to the long time horizon involved. This means that small errors in the model choice or estimation can be amplified significantly. To reduce the exposure to one model or one estimation horizon, one could further investigate model averaging approaches.

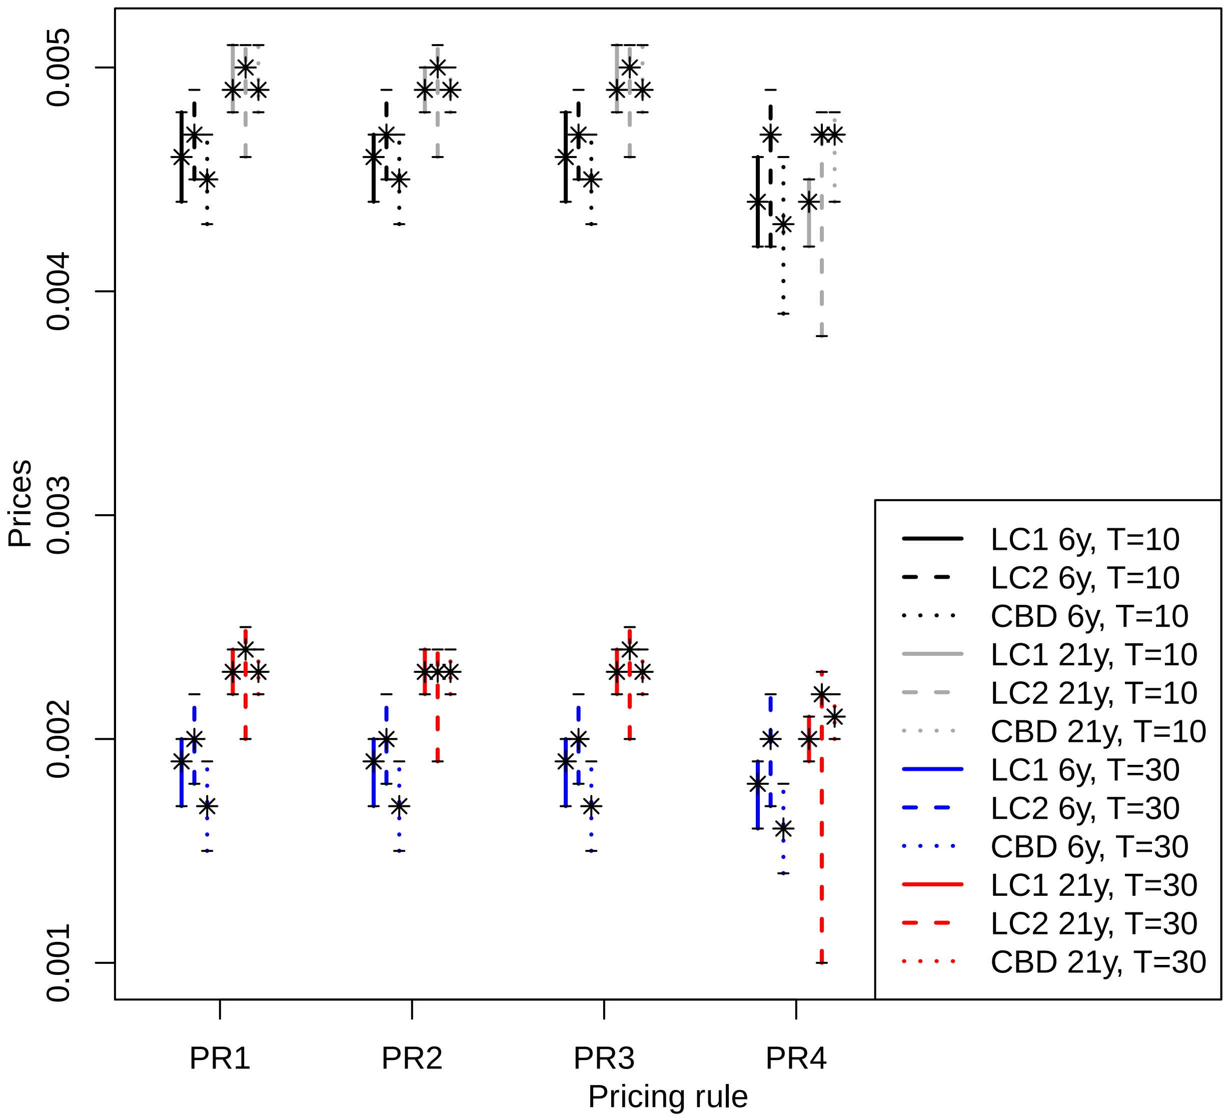

Figure 4: q-forward prices (indicated by stars) and corresponding 95% confidence intervals for age $x = 60$ and maturities $T \in \{10, 30\}$.

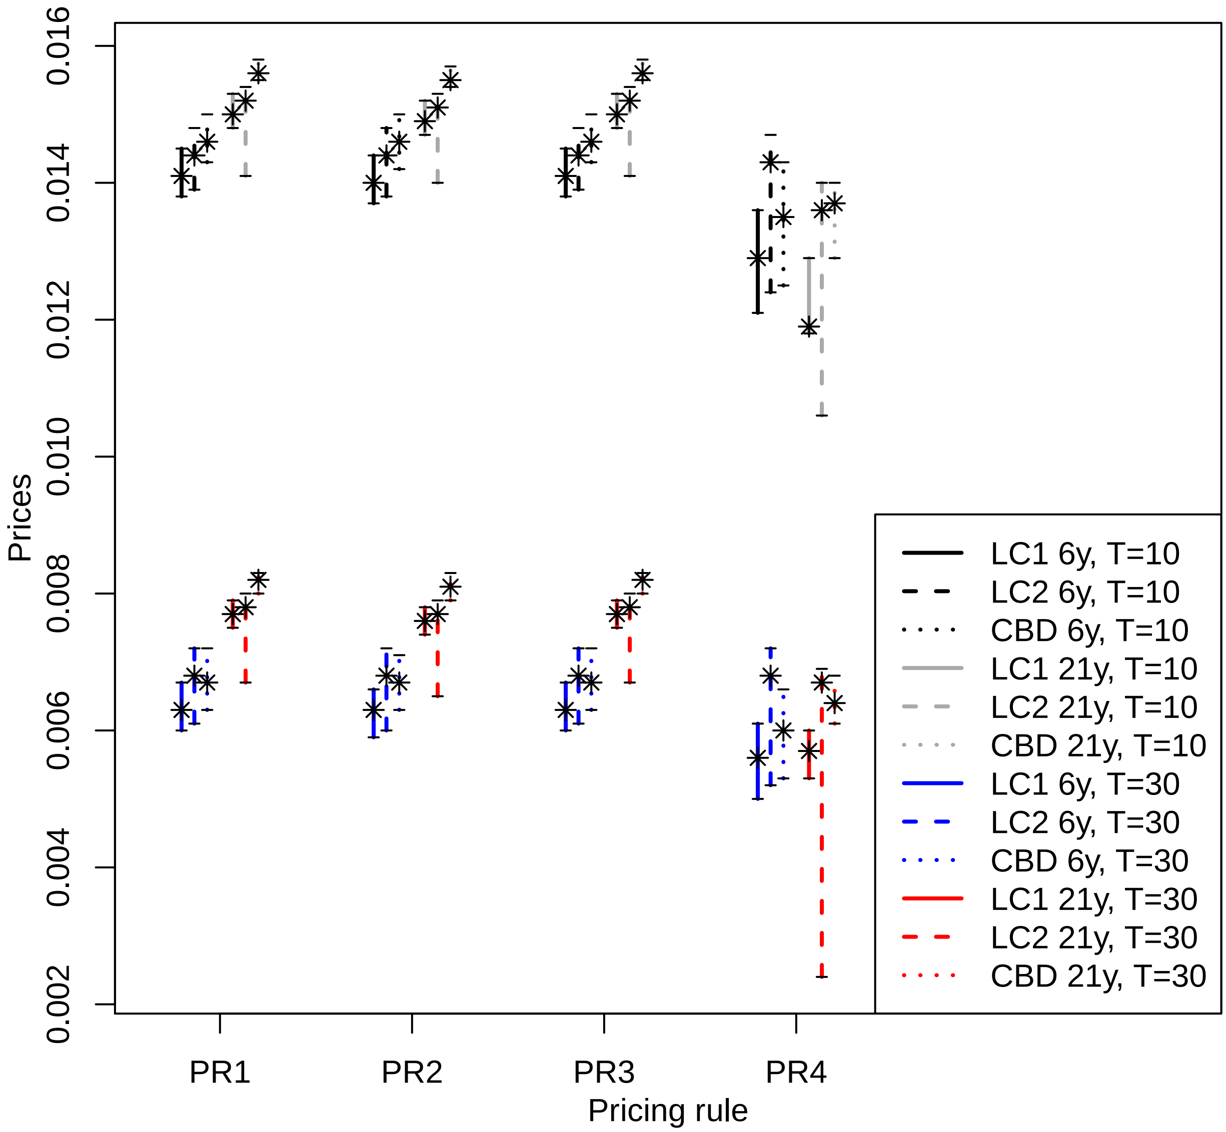

Figure 5: q-forward prices (indicated by stars) and corresponding 95% confidence intervals for age $x = 70$ and maturities $T \in \{10, 30\}$.

A Bootstrap method

In the following we briefly outline the bootstrap method that was used to derive the confidence intervals shown in Figures 4 and 5 for the q-forward prices.

We denote by $D(t,x)$ and $E(t,x)$ the number of deaths and exposures to risk at time t and age x. We assume that the corresponding number of deaths has a Poisson distribution with parameter $\lambda(t,x)$, i.e.,

$$\mathbb{P}(D(t,x) = d) = \frac{\lambda(t,x)^d e^{-\lambda(t,x)}}{d!}$$

and $\lambda(t,x) = \mathbb{E}\left[D(t,x)\right] = E(t,x)m(t,x)$. The exposure-to-risk is observable from empirical data and we estimate $m(t,x)$ using either the Lee-Carter model or the Cairns-Blake-Dowd model. Then we use the following bootstrapping method which is related to the approaches proposed in Li & Luo (2012) and Brouhns et al. (2005).

1. We first fit the corresponding model for the crude death rate or mortality rates (i.e., we compute $\hat{m}(t,x)$ in the Lee-Carter model or $\hat{q}(t,x)$ in the Cairns-Blake-Dowd model) using empirical data.
2. Then we fit the corresponding time series model to the estimates $\hat{\kappa}^{(2)}$ in the Lee-Carter model and $(\hat{\kappa}^{(1)}, \hat{\kappa}^{(2)})^\top$ in the Cairns-Blake-Dowd model.
3. We use the fitted time series model to forecast 10,000 paths of m or q at terminal time T. We use these paths to evaluate the Monte Carlo estimator.
4. Then we start the parametric bootstrap: We assume that the number of deaths has a Poisson distribution with parameter $\hat{\lambda}(t,x) = \hat{m}(t,x)E(t,x)$.
   (a) For each x and each t in the data set we sample the number of deaths $D(t,x)$ from the Poisson distribution with parameter $\hat{\lambda}(t,x)$ and treat those as the new observations.
   (b) We use these new observations and go back to 1.) and refit the mortality model, continue with 2.) and refit the time series model and go to 3.) to compute the new prices based on Monte Carlo.
   (c) Finally we take the 2.5% and the 97.5% quantile of all prices obtained (we used 1,000 bootstrap repetitions) and use them as lower and upper bounds of our confidence intervals.

References

Amori, M., Sagoo, P. & Uncu, O. (2012). Longevity risk transfer to institutional investors: a new era. *Institutional Investors Journal - Special Issue: Longevity Risk Management* **2012**, 25–32.

Barrieu, P., Bensusan, H., El Karoui, N., Hillairet, C., Loisel, S., Ravanelli, C. & Salhi, Y. (2012). Understanding, modeling and managing longevity risk: Key issues and main challenges. *Scandinavian Actuarial Journal* **3**, 203–231.

Bauer, D., Borger, M. & Russ, J. (2010). On the pricing of longevity-linked securities. *Insurance: Mathematics and Economics* **46**, 139–149.

Brockwell, P. J. & Davis, R. A. (2009). *Time series: theory and methods*. Springer.

Brouhns, N., Denuit, M. & Van Keilegom, I. (2005). Bootstrapping the Poisson log-bilinear model for mortality forecasting. *Scandinavian Actuarial Journal* **2005**, 212–224.

Cairns, A., Blake, D. & Dowd, K. (2006a). Pricing death: Frameworks for the valuation and securitization of mortality risk. *Astin Bulletin* **36**, 79.

Cairns, A., Blake, D. & Dowd, K. (2006b). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance* **73**, 687–718.

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

Carmona, R. (ed.) (2009). *Indifference Pricing - Theory and Applications*. Princeton University Press.

Coughlan, G., Epstein, D., Sinha, A. & Honig, P. (2007). q-forwards: Derivatives for transferring longevity and mortality risk. JP Morgan Pension Advisory Group.

Deprez, O. & Gerber, H. (1985). On convex principles of premium calculation. *Insurance: Mathematics and Economics* **4**, 179–189.

Föllmer, H. & Schied, A. (2002). *Stochastic Finance - an Introduction in Discrete Time*. de Gruyter Studies in Mathematics.

Hyndman, R. J. & Athanasopoulos, G. (2013). *Forecasting: principles and practice*. http://otexts.org/fpp/. Accessed on 11/01/2014.

Hyndman, R. J. & Khandakar, Y. (2008). Automatic time series for forecasting: the forecast package for R. *Journal of Statistical Software* **27**.

Kremer, E. (1986). On robust premium principles. *Insurance: Mathematics and Economics* **5**, 271–274.

Lane Clark & Peacock LLP (LCP) (2013). Lcp pension buy-ins, buy-outs and longevity swaps 2013. Available from www.lcp.uk.com.

Lee, R. D. & Carter, L. R. (1992). Modelling and forecasting U.S. mortality. *Journal of the American Statistical Association* **87**, 659–675.

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

Loeys, J., Panigirtzoglou, N. & Ribeiro, R. (2007). Longevity: A market in the making. J.P. Morgans Global Market Strategy.

Mikosch, T. (2009). *Non-Life Insurance Mathematics - An Introduction with the Poisson Process*. Second ed., Springer, Heidelberg, Germany.

Sweeting, P. (2011). A trend-change extension of the Cairns-Blake-Dowd model. *Annals of Actuarial Science* **5**, 143–162.

The Life & Longevity Markets Association (2010). Technical note - q-forward. Available from www.llma.org/publications.html.

The Life & Longevity Markets Association (2012). Longevity pricing framewrok. Available from www.llma.org/publications.html.