Canonical Valuation of Mortality-linked Securities

Johnny S.H. Li*

Abstract

A fundamental question in the study of mortality-linked securities is how to place a value on them. This is still an open question, partly because there is a lack of liquidly traded longevity indexes or securities from which we can infer the market price of risk. This paper develops a framework for pricing mortality-linked securities, on the basis of the theory of canonical valuation. This framework is largely non-parametric, helping us avoid parameter and model risk, which may be significant in other pricing methods. The framework is then applied to a mortality-linked security, and the results are compared against those derived from the Wang transform and some model-based methods.

Keywords: Entropy; Longevity risk; Non-parametric methods; Securitization

*Corresponding author. Address: Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1.

1 Introduction

1.1 Background

Thanks to the combination of better health care and other factors, human mortality in developed countries has been improving steadily for many decades. While improved longevity is generally perceived as a social achievement, it can be a serious problem for actuaries, particularly when it is unanticipated. Longevity risk, that is, the risk that future mortality improvement deviates from today's assumptions, has significantly contributed to the pension crisis that has enveloped many public and corporate pension plans on both sides of the Atlantic.

Actuaries did, of course, take the possibility that people would live longer into account when valuing pensions and annuities. However, what was missed was the pace of mortality reduction. For instance, mortality reduction factors that have been widely used in Britain are found to understate the decline of UK male pensioners' mortality considerably (see Continuous Mortality Investigation Bureau, 1999, 2002). Such an error, which will lead to unforeseen pension and annuity liabilities in the future, cannot be mitigated by selling a large number of contracts, simply because it affects the entire portfolio. Although the risk may be hedged by selling life insurance to the same lives that are buying life annuities, the hedge, as Cox and Lin (2007) pointed out, is cost prohibitive and may not even be practical in many circumstances.

Securitization is seen as a solution to the problem. By securitization we mean laying off mortality or longevity risk exposures with securities that have payoffs tied to a certain mortality or longevity index. There are two main types of mortality-linked security. The first type, for example, the Swiss Re deal in 2003, aims to hedge against the catastrophic loss of insured lives that might result from natural or man-made disasters. The second type, which is the focus of this paper, allows participants to mitigate longevity risk. A well-known example of this type is the 25-year longevity bond announced by BNP-Paribas and European Investment Bank in November 2004. This bond is an annuity bond which pays coupons that are proportional to the survival rates of English and Welsh males who were aged 65 in 2002. Another example is the QxX index swap launched by Goldman Sachs in December 2007. In this swap, the random cash flows are linked to the QxX index, a longevity index for a representative sample of the US senior insured population. We refer readers to Blake and Burrows (2001), Blake et al. (2006a) and Blake et al. (2006b) for deeper discussions on mortality-linked bonds and swaps.

A fundamental question in the study of mortality-linked securities is how to place a value on them. This is still an open question, partly because, as in valuing over-the-counter traded options, there is a lack of liquidly traded longevity indexes or securities from which we can infer the market price of risk, a crucial element in the pricing process. The difficulty can also be seen from another viewpoint by considering the creation of a replicating hedge. If the index on which the mortality-linked security is based is liquidly traded, then the security can be replicated by a portfolio of bonds and the index. Given the principle of no arbitrage, the price of the security is just the value of its replicating portfolio. However, in the absence of a liquidly traded index, we are not able to price the security in this way as a replicating portfolio cannot be formed. Financial engineers call such a situation market incompleteness. In an incomplete market, pricing must rely on some other assumptions.

1.2 Previous work on pricing mortality-linked securities

Various methods have been proposed to approximate the prices of mortality-linked securities in an incomplete market. These methods may be divided into the following three categories:

- *The Wang transform*
  Prices are based on a 'distorted' survival distribution, which is obtained by applying the Wang transform (Wang, 1996, 2000, 2002) to a survival distribution in the real world probability measure. This method is proposed by Lin and Cox (2005), and subsequently extended by other researchers including Dowd et al. (2006), Denuit et al. (2007), and Lin and Cox (2008).

- *Instantaneous Sharpe ratio*
  This method, proposed by Milevsky et al. (2005), assumes that a party who takes non-diversifiable longevity risk should be rewarded a risk premium, which is a multiple (the instantaneous Sharpe ratio) of the standard deviation of the party's portfolio, after all small sample risk has been diversified away. The standard deviation is derived from an assumed process for the evolution of mortality. This approach has also been considered by Young (2008) and Bayraktar et al. (2009).

- *Risk-neutral dynamics of death/survival rates*
  This method is based on a stochastic mortality model, which is, at the very beginning, defined in the real world measure and fitted to past data. For example, Cairns et al. (2006) consider a two-factor model; Bauer et al. (2008) use a model that is parallel to the HJM model for interest rates. The model is then calibrated to market prices, for example, annuity quotes, yielding a risk-neutral mortality process from which security prices are derived.

The Wang transform has some economic justifications. Specifically, it has been shown that the market price of risk in the Wang transform coincides with that implied by the classical capital asset pricing model (CAPM). Nevertheless, the Wang transform has been criticized by a few researchers including Ruhm (2003) and Pelsser (2008), who point out that the Wang transform may not lead to a price consistent with the arbitrage-free price for general stochastic processes. Bauer et al. (2008) have also expressed some concerns about the Wang transform in the context of pricing longevity risk.

Other than the Wang transform, the methods above are heavily dependent on a stochastic process for mortality dynamics. As a result, on top of the uncertainty about the market price of risk, the prices resulting from these methods are subject to two pieces of uncertainty. First, assuming that the stochastic process is correct, parameters in the process may be wrong since they are merely estimates from a finite data sample. This risk, which we call parameter risk, is unavoidable in any model-based approach. The significance of parameter risk in pricing longevity bonds has been demonstrated by Cairns et al. (2006) through Markov Chain Monte Carlo (MCMC).

Second, all methods that involve a stochastic process are affected by model risk, as the process itself may be inaccurate or even incorrect. This situation happens when, for example, the true dynamics of mortality are driven by more factors than assumed in the process. The impact of model risk on pricing is best illustrated by the problem of a 'volatility skew,' the inverse relationship of implied volatility to exercise price, in valuing equity options. Given a volatility skew, the Black-Scholes model, which assumes a constant volatility for all exercise prices, may significantly underestimate values of out-of-the-money puts and in-the-money calls. Although model risk may be reduced by considering a less stringent mortality model, for example, the P-splines regression proposed by Currie et al. (2004), the change of probability measure from real world to risk-neutral under such a model is often difficult, if not impossible.

1.3 Our idea

The problems above can be avoided by considering an alternative pricing method known as 'canonical valuation,' developed by Stutzer (1996). This approach is largely non-parametric, thus reducing parameter and model risk substantially. Another advantage of canonical valuation is that it does not strictly require the use of security prices to predict other security prices.[1] This advantage is especially important when we have only a handful of mortality-linked securities available in the market. Nevertheless, in the future when the market becomes more mature, the method can be modified easily to incorporate more market prices in estimating the risk-neutral density.

Empirical findings indicate that canonical valuation performs well in pricing options on equity indexes. Stutzer (1996) reports that, in a simulated market governed by the Black-Scholes assumptions, canonical valuation produces prices close to Black-Scholes prices, even without using any of the simulated market prices in the valuation process. Grey and Newman (2005) show that, in a stochastic volatility environment, canonical valuation clearly outperforms the historical-volatility-based Black-Scholes estimator for most combinations of moneyness and maturity. Canonical valuation has also been applied to different derivative securities including soybean futures options (Foster and Whiteman, 1999) and bond futures options (Stutzer and Chowdhury, 1999). Results suggest that canonical valuation has merits in both applications.

The primary objective of this paper is to develop a framework for pricing mortality-linked securities, using the theory of canonical valuation. To achieve this objective, we first develop a non-parametric method which allows us to generate a distribution of future mortality rates in the real world probability measure. Then we transform the real-world distribution into its risk-neutral counterpart, by using the maximum entropy principle, which may be regarded as the core of canonical valuation. Finally, we can price a mortality-linked security by discounting its expected payoff, derived from the risk-neutral distribution of future mortality rates, at the risk-free interest rate.

[1] Canonical valuation does not require option price data in pricing options on stocks or equity indexes, and, as we will demonstrate in Section 3, it requires only one market price when valuing longevity securities.

The rest of this article is organized as follows: Section 2 presents the theory of canonical valuation and its economic intuitions; Section 3 set up a non-parametric method for forecasting mortality; Section 4 details how the maximum entropy principle is used to transform the distribution of future death rates from the real world to a risk-neutral measure; Section 5 applies the theoretical results to a mortality-linked security, and compares our framework with the Wang transform and a model-based approach. Finally, Section 6 discusses the limitations of our framework and concludes the paper.

2 The Theory of Canonical Valuation

2.1 A General Set-up

Let us consider a market in which there are m distinct primary securities, whose values evolve according to the state of nature ω. We assume that the ith security, where $i = 1, 2, \ldots, m$, has a time-zero price of F_i and, at the risk-free interest rate, a random discounted payoff of $f_i(\omega)$. Let P be the objective probability measure and $\mathcal{Q}$ be the set of all measures equivalent to P and satisfying

$$\mathbb{E}^Q[f_i(\omega)] = F_i, \quad i = 1, 2, \ldots, m \tag{1}$$

for any Q in $\mathcal{Q}$. That is, $\mathcal{Q}$ is the set of all equivalent martingale measures. Assume further that there are a finite number N of states of nature. If $m = N$, then we say the market is complete. In a complete market, the equivalent martingale measure is unique. However, if $m < N$, which happens when there are only a few securities trading in the market, then we say the market is incomplete. Market incompleteness implies there are infinitely many equivalent martingale measures. To price a derivative in an incomplete market, we need to choose an equivalent martingale measure that is justifiable. This important step may be accomplished by using the principle of canonical valuation.

The principle of canonical valuation is heavily based on the Kullbuck-Leibler information criterion (Kullback and Leibler, 1951). Denote by

$$D(Q, P) = \mathbb{E}^P\left[\frac{\mathrm{d}Q}{\mathrm{d}P} \ln \frac{\mathrm{d}Q}{\mathrm{d}P}\right]$$

the Kullback-Leibler information criterion of measure Q from measure P. Under the principle of canonical valuation, we should choose the equivalent martingale measure Q_0 that minimizes the Kullbuck-Leibler information criterion, that is,

$$Q_0 = \arg\min_{Q \in \mathcal{Q}} D(Q, P),$$

subject to the constraints specified in equation (1). We call Q_0 the canonical measure.

The set-up above is equivalent to the maximization of the Shannon entropy in physics. Therefore, this principle is sometimes referred to as the principle of maximum entropy. Interested readers are referred to Jaynes (1957) and Kapur (1989) for applications of this principle in physical science.

2.2 Intuitions behind the Theory

The principle of canonical valuation can be justified from different angles.

In statistics, the Kullbuck-Leibler information criterion $D(Q, P)$ represents the information gained by moving from measure P to measure Q. From a Bayesian viewpoint, we may regard the objective probability measure P as the prior distribution. In the absence of any information about market prices, the objective probability measure P is the only measure we can use. Given the prices of the m primary securities, we can update the prior assessment by incorporating the information contained in equation (1). However, no information other than equation (1) should be incorporated in the update. As a result, we should choose a measure such that the resulting gain in information is minimal. Equivalently, we choose the measure that minimizes $D(Q, P)$, subject to the price constraints in equation (1).

Geometrically speaking, the Kullbuck-Leibler information criterion $D(Q, P)$ can be considered as a measure of the distance between P and Q, since it is non-negative and is zero if and only if $Q = P$. The geometric interpretation of the principle can be seen from Figure 1. The sheet in Figure 1 is the set of all measures equivalent to P. Of course, P must be a point on the sheet. The line in the sheet represents $\mathcal{Q}$, the set of all measures equivalent to P satisfying the constraints in equation (1). The canonical measure Q_0 is the point on the line and has a shortest distance to the point P.

Furthermore, the principle of canonical valuation is closely related to the expected utility hypothesis. Rittelli (2000) proved the equivalence between the maximization of expected exponential utility and the minimization of Kullbuck-Leibler information criterion. The equivalence holds true in not only the single-period model but also the multi-period model, provided that the optimal solution of the utility maximization problem exists. The dual representation of the canonical measure provides a very clear and explicit financial interpretation of it. Rittelli's results also imply linkages between the principle of canonical valuation and the Esscher transform (Gerber and Shiu, 1994), which has been widely used in actuarial science.

2.3 Implementing the Theory

To implement canonical valuation, we are required to generate a number of scenarios with equal probability. In practice, this can be accomplished by the bootstrap, which generates realizations of a random variable by drawing with replacement from the associated data sample. The scenarios generated may be regarded as a collection of all states of nature. As a result, if N scenarios are generated, then the probability mass function for the state of nature ω under the real-world probability measure P is given by

$$\Pr(\omega = \omega_j) = \pi_j = \frac{1}{N}, \quad j = 1, 2, \ldots, N.$$

The above is often called the empirical probability distribution or the ungrouped histogram of ω. Let π_j^*, $j = 1, 2, \ldots, N$, be the probability distribution of ω under an equivalent martingale measure Q. We can rewrite the constraints in equation (1) as

$$\sum_{j=1}^{N} f_i(\omega_j)\pi_j^* = F_i, \quad i = 1, 2, \ldots, m, \tag{2}$$

and the Kullbuck-Leibler information criterion as $\sum_{j=1}^{N} \pi_j^* \ln \frac{\pi_j^*}{\pi_j}$. As such, to find the canonical measure Q_0, we solve the following constrained minimization problem:

$$Q_0 = \underset{\pi_j^*}{\arg\min} \sum_{j=1}^{N} \pi_j^* \ln \frac{\pi_j^*}{\pi_j} \quad \text{such that } \textstyle\sum_{j=1}^{N} \pi_j^* = 1 \text{ and (2) holds.}$$

Given Q_0, it is straightforward to place a value on a derivative security. Let us consider a security that has a payoff, discounted to time-zero at the risk-free interest rate, of $g(\omega_j)$ in scenario j. The price of this security is simply $\sum_{j=1}^{N} g(\omega_j)\tilde{\pi}_j^*$, where $\tilde{\pi}_j^*$, $j = 1, 2, \ldots, N$, is the probability distribution of ω under Q_0.

2.4 Stutzer's Example

In the original work of Stutzer (1996), canonical valuation is applied to a European option expiring T years from now. The option is written on a single underlying asset, which pays no dividends and has a price of S_t at time t.

Given a time-series of H past prices, $S_{-1}, S_{-2}, \ldots, S_{-H}$, one can generate N possible values of S_T using the bootstrap, which is as follows:

(i) calculate all of the realized single-period gross returns, that is, S_{-i}/S_{-i-1}, $i = 1, 2, \ldots, H-1$;[2]

(ii) draw, with replacement, T values from the $H-1$ realized single-period returns;

(iii) compute a possible value of S_T by multiplying the T returns drawn successively;

(iv) repeat steps (ii) and (iii) N times to generate N possible values of S_T: $S_T(\omega_j)$, $j = 1, \ldots, N$.

These N possible values are equally probable, so under P measure, the probability π_j associated with $S_T(\omega_j)$ is simply $1/N$. The next procedure is to transform the empirical probabilities, π_j, $j = 1, 2, \ldots, N$, into their corresponding risk-neutral (martingale) probabilities, π_j^*, $j = 1, 2, \ldots, N$. To keep the illustration simple, Stutzer considers only one primary asset, the underlying asset itself. Given this assumption, we can rewrite the constraints in equation (2) as

$$S_0 = \sum_{j=1}^{N} B(0,T) S_T(\omega_j) \pi_j^*, \tag{3}$$

where $B(0,T)$ is the price at time 0 of a risk-free zero-coupon bond maturing for 1 at time T. We then derive the risk-neutral probabilities by minimizing the Kullbuck-Leibler information criterion, $\sum_{j=1}^{N} \pi_j^* \ln \frac{\pi_j^*}{\pi_j}$, subject to $\sum_{j=1}^{N} \pi_j^* = 1$ and equation (3).

Using the Lagrange multiplier method, the solution to the minimization problem is as follows:

$$\tilde{\pi}_j^* = \frac{\exp(\gamma^* B(0,T) S_T(\omega_j))}{\sum_{j=1}^{N} \exp(\gamma^* B(0,T) S_T(\omega_j))}, \qquad j = 1, 2, \ldots N,$$

[2] Alternatively, we can generate possible values of S_T by considering the realized T-period gross returns.

where the Lagrange multiplier γ^* is given by

$$\gamma^* = \arg\min_{\gamma} \sum_{j=1}^{N} \exp(\gamma(B(0,T)S_T(\omega_j) - 1)).$$

Given the canonical measure $\tilde{\pi}_j^*$, $j = 1, 2, \ldots, N$, the value C of a European call option with exercise price X expiring T years from now can be expressed as

$$C = \sum_{j=1}^{N} B(0,T) \max[S_T(\omega_j) - X, 0]\tilde{\pi}_j^*.$$

3 Non-Parametric Mortality Forecasting

An important feature of canonical valuation is that it does not require an assumption of a stochastic process for the asset or index to which the derivative security is linked. All we need is to generate, by the bootstrap, an empirical distribution of the security's payoff from a time-series of past asset or index values. This section explores how we may obtain such a distribution for valuing longevity securities like the BNP/EIB bond. The idea is illustrated with the mortality data for the English and Welsh male population from year 1950 to 2005.[3] We focus on ages 65 to 90 only as most longevity securities are unrelated to death rates at younger ages.

3.1 Age and Time Dependency in the Data

The bootstrap in this application is not as straightforward as that in Stutzer's example, since the data we use involve two dimensions, age and time, with potential dependence over both dimensions. Age dependency, as Wills and Sherris (2008) point out, is significant and is a critical factor in pricing mortality-linked securities, particularly when the security has a tranche structure similar to that used in the collateralized debt obligation (CDO) market. To retain age dependency in the bootstrap, we consider mortality rates at different ages jointly by treating them as a vector. That is, we view the data as a multivariate time-series of $\mathbf{m}_t = (m_{65,t}, m_{66,t}, ..., m_{90,t})'$, where $m_{x,t}$ is the central death rate at age x and in year t, and $\mathbf{a}'$ denotes the transpose of $\mathbf{a}$.

[3]The data (historical central death rates) we use are obtained from the Human Mortality Database (2009).

Now we investigate the time (serial) dependency in the vector time-series. In applying the bootstrap, we require the time-series to be weakly stationary.[4] However, the time-series of $\mathbf{m}_t$, as shown in Figure 2, has a clear downward trend, which suggests that it is not weakly stationary. To solve this problem, we consider the transformation of $r_{x,t} = \frac{m_{x,t+1}}{m_{x,t}}$, which may be interpreted as the one-year mortality reduction factor at age x and in year t. Given 56 years of central death rates, we have 55 realized values of $r_{x,t}$ for each age. In Figure 3 we observe no systematic change in $r_{x,t}$ over time, suggesting that it is reasonable to assume that the time-series of $\mathbf{r}_t = (r_{65,t}, r_{66,t}, ..., r_{90,t})'$ is weakly stationary.

To check if the vector $\mathbf{r}_t$ is serially correlated, we examine the cross-correlation matrix (CCM) constructed from the single-period mortality reduction factors at 5 representative ages: 70, 75, 80, 85, and 90.[5] The resulting sample CCMs are shown in Table 1. To better understand the significance of the cross-correlations, in Table 1 we also show the simplified sample CCMs, which consists of three symbols "+," "−," and "·," where

1. "+" means that the corresponding correlation coefficient is greater than or equal to $2/\sqrt{55}$,
2. "−" means that the corresponding correlation coefficient is less than or equal to $-2/\sqrt{55}$,
3. "·" means that the corresponding correlation coefficient is in between $-2/\sqrt{55}$ and $2/\sqrt{55}$.

Note that $2/\sqrt{55}$ is the asymptotic 5% critical value of the sample correlation under the assumption that the series of $(r_{70,t}, r_{75,t}, r_{80,t}, r_{85,t}, r_{90,t})'$ is a white noise series. It is easily seen that significant cross-correlations at the approximate 5% level appear mainly at lag 1. The diagonal entries in the sample CCM at lag-1 indicate that the components $r_{70,t}$, $r_{75,t}$, $r_{80,t}$, $r_{85,t}$, and $r_{90,t}$ demonstrate significant lag 1 autocorrelation. The off-diagonal entries tell us how the components depend on one another. For

[4]Let $\mathbf{y}_t = (y_{1,t}, ..., y_{k,t})'$. We say $\mathbf{y}_t$ is weakly stationary if its mean vector, $\boldsymbol{\mu} = \mathbb{E}(\mathbf{y}_t)$, and its covariance matrix $\mathbb{E}[(\mathbf{y}_t - \boldsymbol{\mu})(\mathbf{y}_t - \boldsymbol{\mu})']$ are constant over time. From an intuitive viewpoint, a time-series is said to be weakly stationary if there is no systematic change in mean (i.e., no trend), no systematic change in variance, and no periodic variations.

[5]Given the data $\{\mathbf{y}_t | t = 1, ..., T\}$, the lag-$l$ cross-correlation matrix $\boldsymbol{\rho}_l$ is estimated by $\hat{\boldsymbol{\rho}}_l = \hat{\mathbf{D}}^{-1}\hat{\boldsymbol{\Gamma}}_l\hat{\mathbf{D}}^{-1}$, where $\hat{\boldsymbol{\Gamma}}_l = \frac{1}{T}\sum_{t=l+1}^{T}(\mathbf{y}_t - \bar{\mathbf{y}})(\mathbf{y}_{y-l} - \bar{\mathbf{y}})'$, $\bar{\mathbf{y}} = \frac{1}{T}\sum_{t=1}^{T}\mathbf{y}_t$, and $\hat{\mathbf{D}}$ is the $k \times k$ diagonal matrix of the sample standard deviations of the component series.

instance, the (2,5)th element in the lag-1 CCM indicates that the reduction factor for age 75 at time t is significantly dependent on that for age 90 at time $t-1$.

3.2 The Bootstrap Procedure

For a sequence with sample CCMs like those in Figure 1, simply drawing with replacement (i.e., the naïve bootstrap) is inappropriate, as it will lose the serial dependency in the data. To retain serial dependency, we can create pseudo-samples by the block bootstrap method, which was first introduced by Carlstein (1986) and further developed by Künsch (1989). The key idea behind the block bootstrap method is that, for a stationary time-series, successive observations are correlated but observations separated by a large time gap are (nearly) uncorrelated. This phenomenon can be seen from the sample CCMs for the series of $(r_{70,t}, r_{75,t}, r_{80,t}, r_{85,t}, r_{90,t})'$ – the cross-correlations taper off as the lag l increases and they all become insignificant beyond lag 3. As a result, individual blocks of observations that are separated far enough in time will be approximately uncorrelated and can be treated as exchangeable. By drawing blocks of data rather than individual values, we can create pseudo-samples that preserve the serial dependence in the original data sequence.

The block bootstrap method can be implemented in different ways. The simplest version divides the data into nonoverlapping blocks of equal size. Assuming a block size of 5, this resampling scheme yields 11 blocks, $(\mathbf{r}_{1950}, \mathbf{r}_{1951}, \mathbf{r}_{1952}, \mathbf{r}_{1953}, \mathbf{r}_{1954})$, $(\mathbf{r}_{1955}, \mathbf{r}_{1956}, \mathbf{r}_{1957}, \mathbf{r}_{1958}, \mathbf{r}_{1959})$, ..., $(\mathbf{r}_{2000}, \mathbf{r}_{2001}, \mathbf{r}_{2002}, \mathbf{r}_{2003}, \mathbf{r}_{2004})$. A variant of this resampling plan is to permit the blocks to overlap. Assuming again a block size of 5, allowing the blocks to overlap will give us 51 blocks, $(\mathbf{r}_{1950}, \mathbf{r}_{1951}, \mathbf{r}_{1952}, \mathbf{r}_{1953}, \mathbf{r}_{1954})$, $(\mathbf{r}_{1951}, \mathbf{r}_{1952}, \mathbf{r}_{1953}, \mathbf{r}_{1954}, \mathbf{r}_{1955})$, ..., $(\mathbf{r}_{2000}, \mathbf{r}_{2001}, \mathbf{r}_{2002}, \mathbf{r}_{2003}, \mathbf{r}_{2004})$. In subsequent calculations, we use the latter resampling plan as it allows for more blocks.[6] To obtain a pseudo-sample, we simply draw blocks with replacement from the original sample and paste the blocks drawn end-to-end to form a new series.

The optimal block size is not always evident. If the blocks are too short, serial dependency in the original sample will be lost. However, using a longer length will effectively reduce the sample size. Hall et al. (1995) show that the optimal block size

[6] This incurs 'end effects,' as the first and last 4 of the original observations appear in fewer blocks then the rest. Such effects can be removed by wrapping the data around a circle, adding the blocks $(\mathbf{r}_{2004}, \mathbf{r}_{1950}, \mathbf{r}_{1951}, \mathbf{r}_{1952}, \mathbf{r}_{1953})$, ..., $(\mathbf{r}_{2001}, \mathbf{r}_{2002}, \mathbf{r}_{2003}, \mathbf{r}_{2004}, \mathbf{r}_{1950})$. This adjustment ensures that each of the original observations has an equal chance of appearing in the simulated series.

depends significantly on the context. In estimating a two-sided distribution function, a block size of $n^{1/5}$, where n is the effective sample size, is optimal. Given this rule, we use a block size of 2 ($55^{1/5} = 2.23 \approx 2$).

Researchers have proposed several ways to improve the block bootstrap procedure, for example, post-blackening, blocks of blocks, and stationary bootstrap. These methods, which are detailed in Davison and Hinkley (1997) and Lahiri (2003), can be incorporated easily into the algorithm we described.

3.3 Making a Mortality Forecast

Let us suppose that the forecast horizon is 30 years. Using a block-size of 2 and the resampling plan that allows the blocks to overlap, we generate 10,000 pseudo-samples of 30 one-year mortality reduction factors. Given these 10,000 pseudo-samples, we can make forecasts of various death and survival probabilities. As an example, we consider the central death rate at age 90 in year 2035 (30 years from 2005). Let $\mathbf{M}$ be a pseudo-sample and $\mathbf{M}(i, j)$, $i = 1, 2, ..., 26$, $j = 1, 2, ..., 30$, be the (i, j)th element in $\mathbf{M}$.[7] On the basis of $\mathbf{M}$, an estimate of $m_{90,2035}$ is given by the product of the base year central death rate, $m_{90,2005}$, and the simulated 30-year reduction factor $\prod_{j=1}^{30} \mathbf{M}(26, j)$ for age 90. With 10,000 pseudo-samples, we can construct an empirical distribution from which we can obtain a central estimate and a confidence interval for $m_{90,2035}$.

With a forecast of cohort death rates, that is, $m_{x,2006}, m_{x+1,2007}, ...$, we can then make a forecast of survival probabilities for different birth cohorts. In Figure 4 we show the empirical distributions of the 10-year, 15-year, 20-year and 25-year survival probabilities for the cohort aged 65 in year 2005.[8] From the means and percentiles of the simulated distributions, we obtain a central estimate and a confidence interval for each of the survival probabilities (see Table 2).

To examine the robustness of the bootstrap relative to the historical data used, we base the bootstrap on three different sample periods:

1. 1960–2005 (46 years; dot-dash line in Figure 4);

[7] Note that $\mathbf{r}_t$ is a 26×1 vector which consists of reduction factors for 26 different ages. Therefore, a pseudo-sample $\mathbf{M}$ of 30 one-year mortality reduction factors would be a 26×30 matrix.

[8] We let ${}_tp_x$ be the probability that a person who was aged x in the base year (year 2005) survives to age $x + t$.

2. 1950–2005 (56 years; solid line in Figure 4);

3. 1940–2005 (66 years; dotted line in Figure 4).

An increase or decrease in the sample period by 10 years seems to have little influence on the central tendency of the simulated distributions, indicating that the bootstrap is reasonably robust relative to how much historical data is used. However, when more years of data are used, the resulting empirical distributions are more dispersed. This observation reflects the greater volatility in mortality rates that can be seen in earlier years. A similar observation is also made in a model-based simulations study conducted by Cairns et al. (2009).

Finally, we compare our non-parametric projection with the projections derived from two parametric mortality models: (1) the Lee-Carter model (Lee and Carter, 1992) and (2) the two-factor model (Cairns et al., 2006). The comparison (see Table 3) indicates that the projections are fairly close to one another.

4 An Equivalent Martingale Measure

Recall that in Stutzer's example, the canonical measure is derived by minimizing the Kullbuck-Leibler information criterion, subject to a constraint (equation (3)) that is based on the asset to which the derivative security is linked.

However, we are unable to derive the canonical measure for longevity securities in this way, as they are linked to either death or survival rates that are not traded in the market. Without a price for the underlying, a constraint similar to that in Stutzer's example cannot be formed. To solve this problem, we need to rely on one or more security that is linked to the relevant death or survival rates and is traded in the capital market at a price we know. We use the BNP/EIB longevity bond in 2004 to illustrate.

4.1 The BNP/EIB Longevity Bond

Before we proceed to the derivation of the canonical measure, let us briefly review the BNP/EIB longevity bond. This bond is a 25-year amortising bond (i.e., a bond without principal repayment) with coupon payments that are linked to a survivor index, which is based on the realized mortality rates of English and Welsh males aged

65 in 2002. The index $I(t)$ on which the coupon payments are based is defined as follows:

$$I(t) = I(t-1)(1 - m_{64+t,2002+t}), \qquad t = 1, 2, \ldots, 25,$$

where $I(0) = 1$, and $m_{x,t}$ is the crude central death rate at age x and in year t. In each year t, $t = 1, 2, ..., 25$, the bond pays a coupon of £50 $\times I(t)$ million.

The issue price was determined by discounting at LIBOR minus 35 basis points the anticipated coupon payments, £50 $\times \mathbb{E}_P[I(t)|\mathcal{F}_0]$ million, $t = 1, 2, \ldots, 25$, where $\mathcal{F}_t$ is the filtration generated by the development of the mortality curve up to time t. Assuming that the evolution of mortality rates over time is independent of the dynamics of the interest rate term-structure over time, the issue price quoted in the contract can be written as

$$£50 \times \sum_{t=1}^{25} B(0,t) \exp(-\delta t) \mathbb{E}_P[I(t)|\mathcal{F}_0],$$

where δ is the longevity risk premium, and $B(0,t)$ is the time-0 price of a risk-free zero-coupon bond that pays 1 at time t (in years).[9] As the EIB curve typically stands about 15 basis points below the LIBOR curve, the risk premium δ is approximately 20 basis points.

Using the non-parametric bootstrap procedure we detailed in Section 3, we calculate the value of $\mathbb{E}^P[I(t)|\mathcal{F}_0]$ for $t = 1, 2, \ldots, 25$. Assuming that the EIB interest rate is 4% per annum, the estimated market price of the bond at $t = 0$ is £561 (£50 $\times$ 11.22) million, which is boardly in line with that derived by Cairns et al. (2006) on the basis of their two-factor mortality model.[10] With this market price, we can formulate a constraint for use in the derivation of the canonical measure.

[9] We define here a risk-free bond by a bond that is free of longevity risk; that is, its payoff is the same regardless of what mortality scenario it turns out to be. On this basis of our definition, a risk-free bond may be subject to other types of risk, for example, counterparty default risk. So such a bond could be one that is issued by EIB (or an institution with a similar credit rating) and is not mortality-linked. In the rest of this article, the risk-free rate refers to the interest rate on such a bond.

[10] Under the same assumption on the EIB interest rate, Cairns et al. (2006) find that the price at issue of the BNP/EIB longevity bond is £572.1 (£50 $\times$ 11.442).

4.2 Deriving the Canonical Measure

Recall that the derivation of the canonical measure involves two steps. The first step is to generate a number, say N, of equally probable mortality scenarios using the non-parametric bootstrap we introduced in Section 3. In each scenario, we have an array of future central death rates from which we can calculate the value of the longevity index $I(t)$ at $t = 1, 2, \ldots, 25$. Let $I(t, \omega_j)$ be the value of the longevity index at time t in the jth scenario. In the jth scenario, the cash flows, discounted to time zero at the risk-free interest rate, from the longevity bond is given by

$$v(\omega_j) = 50 \times \sum_{t=1}^{25} B(0,t) I(t, \omega_j).$$

Under the objective probability measure P, the probability of having a discounted payoff of $v(\omega_j)$ from the longevity bond is $\pi_j = 1/N$, for $j = 1, 2, \ldots, N$. The distribution of $v(\omega)$ under P is shown graphically in the upper panel of Figure 5.

The next step is to perform a constrained minimization of the Kullbuck-Leibler information criterion. We let π_j^* be the probability associated with $v(\omega_j)$ (i.e., the jth scenario) under an equivalent martingale measure Q. Under Q, the expectation of $v(\omega)$ must be the same as the market price of the longevity bond at time zero. In other words, the following constraint must be satisfied:

$$\sum_{j=1}^{N} v(\omega_j) \pi_j^* = 561. \tag{4}$$

The canonical measure is then chosen by minimizing the Kullbuck-Leibler information criterion, subject to $\sum_{j=1}^{N} \pi_j^* = 1$ and the constraint in equation (4). We solve this problem with the method of Lagrange multipliers, which says the constrained minimization is equivalent to minimizing

$$L = \sum_{j=1}^{N} \pi_j^* \ln \pi_j^* - \lambda_0 \left(\sum_{j=1}^{N} \pi_j^* - 1 \right) - \lambda_1 \sum_{j=1}^{N} \left(v(\omega_j) \pi_j^* - 561 \right).$$

Let $\tilde{\pi}_j^*$, $j = 1, 2, \ldots, N$, be the solution, that is, the canonical measure Q_0. We require it to satisfy the first-order conditions:

$$\ln \tilde{\pi}_j^* + 1 - \lambda_0 - \lambda_1 v(\omega_j) = 0, \quad j = 1, 2, \ldots, N,$$

or equivalently,

$$\tilde{\pi}_j^* = \exp(\lambda_0 + \lambda_1 v(\omega_j) - 1), \quad j = 1, 2, \ldots, N,$$

which means π_j^* is proportional to $\exp(\lambda_1 v(\omega_j))$. It follows from $\sum_{j=1}^N \pi_j^* = 1$ that

$$\tilde{\pi}_j^* = \frac{\exp(\lambda_1 v(\omega_j))}{\sum_{j=1}^N \exp(\lambda_1 v(\omega_j))}, \quad j = 1, 2, \ldots, N. \tag{5}$$

What remains is the Lagrange multiplier λ_1, which can be determined by substituting (5) into (4) or by the following expression:

$$\lambda_1 = \underset{\pi_j^*}{\arg\min} \sum_{j=1}^N \exp(\gamma(v(\omega_j) - 561).$$

Using the procedure above, we obtain an estimate of the canonical measure Q_0, which is depicted graphically in the lower panel of Figure 5.

4.3 Additional Primary Securities

In deriving the canonical measure shown in Figure 5, only one primary security, the BNP/EIB longevity bond, is considered. What if there is in the market more than one security that is linked to the mortality of the same reference population? How can we ensure that all these securities are correctly priced under the canonical measure?

We can easily extend the method to incorporate additional primary securities. Suppose that there are $m > 1$ such securities and that the ith, $i = 1, 2, \ldots, m$, security has a price of V_i at time zero and a discounted payoff of $v_i(\omega_j)$ in the jth mortality scenario, $j = 1, 2, \ldots, N$. To ensure correct pricing of these m securities, the following conditions must be satisfied:

$$\sum_{j=1}^N v_i(\omega_j)\pi_j^* = V_i, \quad i = 1, 2, \ldots, m. \tag{6}$$

Therefore, with $m > 1$ primary securities, we obtain the canonical measure by minimizing the Kullbuck-Leibler information criterion, subject to $\sum_{j=1}^N \pi_j^* = 1$ and the constraints in equation (6). It can be shown that the resulting canonical measure $\tilde{\pi}_j^*$, $j = 1, 2, \ldots, N$ is given by

$$\tilde{\pi}_j^* = \frac{\exp(\sum_{i=1}^m \lambda_i v(\omega_j))}{\sum_{j=1}^N \exp(\sum_{i=1}^m \lambda_i v(\omega_j))}, \quad j = 1, 2, \ldots, N,$$

where the Lagrangian multipliers $\vec{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_m)'$ can be expressed as

$$\vec{\lambda} = \underset{\gamma_1, \ldots, \gamma_m}{\arg\min} \sum_{j=1}^N \exp\left(\sum_{i=1}^m \gamma_i(v_i(\omega_j) - V_i)\right).$$

The intuition of the extension above can be demonstrated diagrammatically. The top panel in Figure 6 represents the case when there is only one primary security. As in Figure 1, the sheet is the set of all measures equivalent to P, while the line on the sheet is $\mathcal{Q}$, the set of all measures equivalent to P satisfying the constraint. On $\mathcal{Q}$, we can find the canonical measure Q_0, which is the point that is closest to P. The middle panel in Figure 6 represents the case when there are two primary securities. By requiring measures in $\mathcal{Q}$ to price both primary securities correctly, the locus for $\mathcal{Q}$ is effectively shortened. It is noteworthy that the introduction of an additional primary security may result in a different Q_0, since the previous Q_0 may no longer be encompassed by the locus for $\mathcal{Q}$. The bottom panel represents the extreme case when there are infinitely many primary securities, or equivalently speaking, a complete market. In this case, the locus for $\mathcal{Q}$ reduces to a single point, the only position that Q_0 can take, implying that Q_0 coincides with the unique equivalent martingale measure.

5 An illustration

5.1 Pricing Vanilla Survivor Swaps

We illustrate our pricing framework with vanilla survivor swaps, in which the parties involved agree to swap a series of payments, one of which depends on a longevity index, periodically until the swap matures.

Vanilla survivor swaps can be constructed in different ways. Following Dowd et al. (2006), we consider vanilla survivor swaps with a fixed proportional premium θ and a fixed time-to-maturity T. At $t = 1, 2, \ldots, T$, there is an exchange of, per \$1 notional principal, a preset amount $(1 + \theta)K(t)$ and a random amount $S(t)$ that is linked to the number of survivors in a certain reference population.

To keep mutual credit risks down, it makes sense for the agreement to specify that the two parties exchange only the net difference between the two payment amounts. Therefore, per \$1 notional principal, the fixed-payer pays the the fixed-receiver an amount of $(1 + \theta)K(t) - S(t)$ if $(1 + \theta)K(t) > S(t)$ and the fixed-receiver pays the fixed-payer an amount of $S(t) - (1 + \theta)K(t)$ otherwise. It is easy to see that the fixed-payer has a long exposure to longevity risk (i.e., the risk that $S(t)$ turns out to be low relative to $K(t)$), while the fixed-receiver has a short exposure.

In our illustration, the floating leg $S(t)$ is linked to the mortality of the same reference population as that for the BNP/EIB longevity bond. Specifically, we set $S(t)$ to the realized survival function for the reference population, that is,

$$S(t) = S(t-1)(1 - q_{64+t,2002+t}), \quad t = 1, 2, \ldots, T,$$

where $S(0) = 1$, and $q_{x,t}$ is the realized probability that an English/Welsh male aged x at the beginning of year t dies during year t.

A key difference between this and a vanilla interest-rate swap is that, rather than being constant, the fixed leg $K(t)$ for this swap declines over time in line with the values of $S(t)$, $t = 1, 2, ..., T$, anticipated at time zero. Here we set $K(t)$ to the projected survival function for the reference population, on the basis of the 2003-based principal mortality projection made by the UK Government Actuary's Department.[11] Values of $K(t)$ for $t = 1, 2, \ldots, 25$ are shown in Table 4.

In line with vanilla interest rate swaps, the premium θ is chosen so that the initial value of the swap is zero to each party. As such, we can calculate θ by using the following equation:

$$\sum_{t=1}^{T} B(0,t) \left(\mathbb{E}^{Q}[S(t)|\mathcal{F}_0] - (1+\theta)K(t) \right) = 0, \tag{7}$$

where $B(0,t)$ is the time-0 price of a fixed-principle zero-coupon bond that pays 1 at time t. Note that θ might be positive, zero, or negative.

All that then remains is to obtain $\mathbb{E}^{Q}[S(t)|\mathcal{F}_0]$ for $t = 1, 2, \ldots, T$. When we use canonical valuation, these expectations can be calculated as follows:

$$\sum_{j=1}^{N} S(t, \omega_j) \pi_j^*,$$

where $S(t, \omega_j)$ is the value of $S(t)$ in the jth mortality scenario, and π_j^* is the probability associated with the jth scenario under the canonical measure Q_0, which we identified in Section 4.

Assuming a risk-free rate of 4%, we calculate the swap premia for maturities ranging from 1 to 25 years. The solid line in Figure 7 shows the resulting values of θ based on the sample period of 1950–2002. To evaluate the robustness of canonical valuation

[11] The 2003-based principal mortality projection of age/sex specific mortality rates is available at http://www.gad.gov.uk/Demography_Data/Population/.

relative to the historical data used, we consider two additional sample periods: 1960–2002, and 1940–2002. The results, also shown in Figure 7, indicate that a change of the sample period by 10 years does not affect the swap premia significantly.

5.2 Comparing with Other Pricing Methods

We now compare our pricing framework with the Wang transform and the method that is based on the two-factor stochastic mortality model proposed by Cairns et al. (2006).

- *The Wang transform*

  Let $F^P(x)$ be the distribution function for a future lifetime random variable in the real-world probability measure (P measure). The Wang transform defines the 'distorted' distribution function $F^Q(x)$ for the random variable by

  $$F^Q(x) = \Phi(\Phi^{-1}(F^P(x)) + \lambda),$$

  where Φ is the distribution function for the standard normal random variable, and λ is the market price of risk, which reflects the level of longevity risk. Using the Wang transform, we can calculate the price of a mortality-linked security by discounting its expected payoff implied by $F^Q(x)$ at the risk-free interest rate.

  In pricing the vanilla survivor swaps we defined earlier, $F^P(x)$ represents the real-world survival distribution for the cohort of English and Welsh males who were aged 65 in year 2002. We calculate $F_P(x)$ from the 2003-based principal mortality projection made by the UK Government Actuaries Department.

  To obtain the market price of risk λ, we make use of the market price of the BNP/EIB longevity bond. Specifically, we find λ such that the price of the bond implied by the resulting $F^Q(x)$ is the same as the market price of the bond.

- *The two-factor model*

  The two-factor model is a discrete-time model which assumes that $q_{x,t}$, the single-year death probability at age x and time t, can be formulated as follows:

  $$q_{x,t} = \frac{e^{A_1(t)+A_2(t)x}}{1 + e^{A_1(t)+A_2(t)x}},$$

  where $\{A_1(t)\}$ is a stochastic factor that affects all the ages in an equal manner, and $\{A_2(t)\}$ is another stochastic factor that has a different effect for different ages.

In the real world probability measure (P measure), $\{A_1(t)\}$ and $\{A_2(t)\}$ follow a bivariate random walk with drift, that is,

$$A(t+1) = A(t) + \mu + CZ(t+1),$$

where $A(t) = (A_1(t), A_2(t))'$, μ is a constant 2×1 vector, C is a constant 2×2 upper triangular matrix, and $Z(t)$ is a bivariate standard normal random variable. We estimate μ and C from the historical death probabilities for the English and Welsh male population from 1950 to 2002.

In a risk-adjusted pricing measure (Q measure), the stochastic process for $A(t)$ has the following form:

$$A(t+1) = A(t) + \tilde{\mu} + C\tilde{Z}(t+1),$$

where $\tilde{\mu} = \mu - C\lambda$, $\tilde{Z}(t+1)$ is a bivariate standard normal random variable under the Q-measure, and $\lambda = (\lambda_1, \lambda_2)'$ is a vector of market prices of risk. Although λ might vary with time, it is assumed here that it is constant over time since it is difficult to assume anything more complicated in a lack of market price data.[12]

As before, we make use of the market price of the BNP/EIB longevity bond to find λ_1 and λ_2. In particular, we choose λ_1 and λ_2 that would result in an equality between the price implied by the model and the issue price quoted in the contract. Since there are two unknowns but only one equation, there are infinitely many pairs of λ_1 and λ_2 under which the price produced by the model would match market price. We consider the special case that $\lambda_1 = \lambda_2$.[13]

Given the market prices of risk, we calculate the price of a security that is linked to the mortality of the same reference population by discounting its expected payoff under Q at the risk-free interest rate.

In Figure 8 we show the swap premia θ on the basis of the three pricing methods. By requiring all three methods to price the 25-year BNP/EIB longevity bond correctly, they yield the same premium for the vanilla survivor swap with a maturity of 25 years. This is because, as Blake et al. (2006a) point out, the BNP/EIB longevity

[12] Cairns et al. (2006) also make this assumption.

[13] Cairns et al. (2006) consider three special cases: $\lambda_1 = 0$, $\lambda_2 = 0$, and $\lambda_1 = \lambda_2$. We find that these three cases yield similar premia for the vanilla survivor swap we defined earlier.

bond may be regarded as a combination of a survivor swap and some fixed cash flows.[14] In an incomplete market, the rest of the swap curve is a mere extrapolation. From Figure 8 we observe that the extrapolated swap curves take different shapes, depending on the pricing method used. The Wang transform renders a fairly linear extrapolation, while canonical valuation and the two-factor model give non-linear swap curves with different curvatures.

6 Discussion and Conclusion

This study develops an alternative framework for pricing mortality-linked securities, on the basis of the theory of canonical valuation. The framework is comprised of two components. The first component is a non-parametric method that allows us to generate scenarios of future mortality rates, while the second is a transformation of the real-world probability distribution for the mortality scenarios into its risk-neutral counterpart for pricing purposes. The empirical results indicate that this alternative pricing framework is reasonably robust relative to the amount of historical mortality data used.

Most other pricing methods are heavily based on an assumed stochastic process for the evolution of mortality. They are subject to model risk, because any stochastic process is only a simplified version of reality, and with any simplification there is the risk that something will fail to be accounted for. For example, a pricing method that is based on the Lee-Carter model might produce inaccurate prices if the temporal signal in the model has a non-constant volatility or significant structural changes.[15] Even if the pricing error is small, the problem might be made larger by several orders of magnitude if the same model is also used for designing the hedge portfolio. On the contrary, the framework we propose is largely non-parametric, effectively helping us avoid model risk, which might be significant in other pricing methods.

Although the longevity market is still very immature, there has been a raft of new entrants, such as Goldman Sachs and JP Morgan, competing for new business. It is therefore legitimate to expect more products coming to the market in the near future.

[14] There might be a small discrepancy, since the BNP/EIB longevity bond is based on central death rates ($m_{x,t}$) while our vanilla survivor swap is based on death probabilities ($q_{x,t}$).

[15] In the Lee-Carter model, it is assumed that the temporal signal of mortality (often denoted by k_t or κ_t) follows a simple linear time-series process with innovations that have a constant variance.

The prices of the new products, as we have demonstrated in Section 4, can be incorporated into the canonical measure easily by introducing additional constraints when we minimize the Kullbuck-Leibler information criterion. In the extreme case when there are infinitely many market prices available, the canonical measure converges to the unique equivalent martingale measure in a complete market. Nevertheless, in using the Wang transform, the incorporation of additional market prices is not that straightforward. Ideally, prices of securities linked to the mortality of the same cohort should yield the same market price of risk λ in the distortion operator, but in reality this may not be the case, since market prices are not necessarily consistent with the Wang transform. Should there exist multiple values of λ, a subjective decision on which to use will be needed. A similar problem may also occur when we base pricing on the risk-adjusted two-factor model in which there are only two market prices of risk.

Besides products like the BNP/EIB longevity bond, some insurance companies have entered into, on an over-the-counter (OTC) basis, financial contracts that are linked to their own mortality experience. For instance, in JP Morgan's 'q-forward,' the counterparty has the discretion to choose between a standardized index, which is linked to a larger population, and a customized index, which reflects the actual experience of individuals associated with a particular exposure, such as the policyholders of a life insurance portfolio or the members of a defined benefit pension plan. While customized deals involve less population basis risk, they are often difficult to price due to the paucity of data. In particular, maximum likelihood estimation might not work well when the data series is too short or when the number of exposures is too small. Our pricing framework, which is largely non-parametric, seems to be an attractive alternative way to value OTC deals that involve a smaller population with a thin volume of mortality data.

The pricing problem is sometimes complicated by cohort effects, which refer to situations when the mortality improvement for a group of birth years is systematically higher or lower than that of the neighboring cohorts. When a model-based method is used, we may factor cohort effects into security prices by considering a model that relates death rates to years of birth. The generalization of the two-factor model[16] is one example. However, it does not seem trivial to incorporate such effects into our pricing framework. An obvious avenue for future research is to investigate how we

[16]This model is labeled as Model M6 in Cairns et al. (2009).

may adapt the block bootstrap procedure so that cohort effects can be taken into account.

References

[1] Bauer D., Boerger M., Russ J. (2008). On the Pricing of Longevity-Linked Securities. Available at www.mortalityrisk.org.

[2] Bayraktar, E., Milevsky, M., Promislow, D., Young, V. (2009) Valuation of Mortality Risk via the Instantaneous Sharpe Ratio: Applications to Life Annuities. *Journal of Economic Dynamics and Control*, **33(3)**, 676-691.

[3] Blake, D., and Burrows, W. (2001). Survivor Bonds: Helping to Hedge Mortality Risk. *Journal of Risk and Insurance*, **68**, 339-348.

[4] Blake, D., Cairns, A.J.G., and Dowd, K. (2006a). Living with Mortality: Longevity Bonds and Other Mortality-Linked Securities. *British Actuarial Journal*, **12**, 153-197.

[5] Blake, D., Cairns, A.J.G.,, Dowd, K. and MacMinn, R. (2006b). Longevity Bonds: Financial Engineering, Valuation and Hedging. *Journal of Risk and Insurance*, **73**, 647-672.

[6] Cairns, A.J.G., Blake, D., and Dowd, K., (2006) A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance*, **73**, 687-718.

[7] 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.

[8] Carlstein, E. (1986). The Use of Subseries Methods for Estimating the Variance of a General Statistic from a Stationary Time Series. *Annals of Statistics*, **14**, 1171-1179.

[9] Continuous Mortality Investigation Bureau (1999). Standard Tables of Mortality Based on the 1991-94 Experiences. CMI Report no. 17. London: Institute of Actuaries and Faculty of Actuaries.

[10] Continuous Mortality Investigation Bureau (2002). An Interim Basis for Adjusting the "92" Series Mortality Projections for Cohort Effects. CMI Working Paper 1. London: Institute of Actuaries and Faculty of Actuaries.

[11] Cox, S.H. and Lin, Y. (2007). Natural Hedging of Life and Annuity Mortality Risks. *North American Actuarial Journal*, **11(3)**, 1-15.

[12] Currie, I. D., Durban, M. and Eilers, P. H. C. (2004). Smoothing and Forecasting Mortality Rates. *Statistical Modelling*, **4**, 279-298.

[13] Davison, A.C. and Hinkley, D.V. (1997). *Bootstrap Methods and Their Application.* New York: Cambridge University Press.

[14] Denuit, M., Devolder, P., and Goderniaux, A.-C. (2007). Securitization of Longevity Risk: Pricing Survivor Bonds with Wang Transform in the Lee-Carter Framework. *Journal of Risk and Insurance*, **74**, 87-113.

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

[16] Foster, F.D. and C.H. Whiteman. (1999). An Application of Bayesian Option Pricing to the Soybean Market. *American Journal of Agricultural Economics*, **81**, 722-727.

[17] Gerber, H.U. and Shiu, E.S.W. (1994). Option Pricing by Esscher Transforms. *Transactions of the Society of Actuaries*, **46**, 99-191.

[18] Gray, P. and Newman, S. (2005). Canonical Valuation of Options in the Presence of Stochastic Volatility. *Journal of Futures Markets*, **25**, 1-19.

[19] Hall, P., Horowitz, J.L. and Jing, B.-Y. (1995). On Blocking Rules for the Bootstrap with Dependent Data. *Biometrika*, **82**, 561-574.

[20] Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute of Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de (data downloaded on 1 Jan 2009).

[21] Jaynes, E.T. (1957). Information Theory and Statistical Mechanics. *Physical Review*, **106**, 620-630.

[22] Kapur, K. N. (1989). *Maximum-Entropy Models in Science and Engineering.* New York: Wiley Eastern, Ltd.

[23] Kullback, S. and Leibler, R. A. (1951) On Information and Sufficiency. *Annals of Mathematical Statistics*, **22**, 79-86.

[24] Künsch, H.R. (1989). The Jackknife and the Bootstrap for General Stationary Observations. *Annals of Statistics*, **17**, 1217-1261.

[25] Lahiri, S.N. (2003). *Resampling Methods for Dependent Data.* New York: Springer-Verlag.

[26] Lee, R. and Carter, L. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, **87**, 659-671.

[27] Lin, Y. and Cox, S.H. (2005). Securitization of Mortality Risks in Life Annuities. *Journal of Risk and Insurance*, **72**, 227-252.

[28] Lin, Y. and Cox, S.H. (2008). Securitization of Catastrophe Mortality Risks. *Insurance: Mathematics and Economics*, **42(2)**, 628-637.

[29] Milevsky, M.A., Promislow, S.D., Young, V.R. (2005). Financial Valuation of Mortality Risk via the Instantaneous Sharpe Ratio: Applications to Pricing Pure Endowments. Working paper, Department of Mathematics, University of Michigan. Available at: http://arxiv.org/abs/0705.1302.

[30] Pelsser, A. (2008). On the Applicability of the Wang Transform for Pricing Financial Risks. *ASTIN Bulletin*, **38**, 171-181.

[31] Rittelli, M. (2000). The Minimal Entropy Martingale Measure and the Valuation Problem in Incomplete Market. *Mathematical Finance*, **10**, 39-52.

[32] Ruhm, D. (2003). Distribution-Based Pricing Formulas Are Not Arbitrage-Free. Casualty Actuarial Society Proceedings, **XC**, 97-129.

[33] Stutzer, M. (1996). A Simple Nonparametric Approach to Derivative Security Valuation. *Journal of Finance*, **51**, 1633-1652.

[34] Stutzer, M. and M. Chowdhury. (1999). A Simple Nonparametric Approach to Bond Futures Option Pricing. *Journal of Fixed Income*, **8**, 67-76.

[35] Wang, S. (1996). Premium Calculation by Transforming the Layer Premium Density. *ASTIN Bulletin*, **26(1)**, 71-92.

[36] Wang, S. (2000). A Class of Distortion Operations for Pricing Financial and Insurance Risks. *Journal of Risk and Insurance*, **67**, 15-36.

[37] Wang, S. (2002). A Universal Framework For Pricing Financial and Insurance Risks. *ASTIN Bulletin*, **32(2)**, 213-234.

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

[39] Young, V.R. (2008). Pricing Life Insurance under Stochastic Mortality Via the Instantaneous Sharpe Ratio: Theorems and Proofs. *Insurance: Mathematics and Economics*, **42**, 691-703.

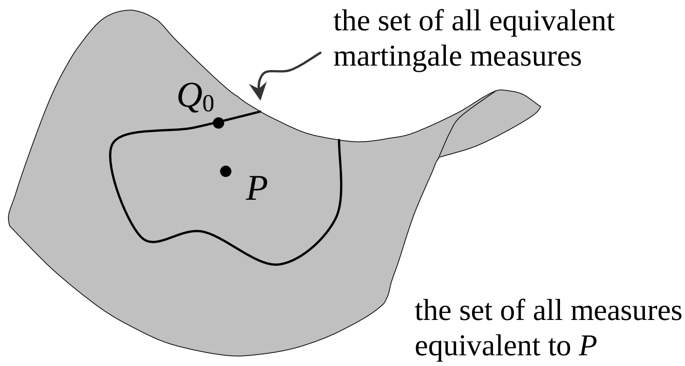

Figure 1: A geometric interpretation of the canonical valuation principle.

	CCM					Simplified CCM				
Lag 1	−0.18	0.04	0.08	−0.01	−0.14	·	·	·	·	·
	−0.12	−0.40	−0.15	−0.26	−0.40	·	−	·	·	−
	−0.17	−0.17	−0.41	−0.27	−0.34	·	·	−	·	−
	−0.09	−0.13	−0.06	−0.48	−0.47	·	·	·	−	−
	−0.09	−0.28	−0.22	−0.46	−0.62	·	−	·	−	−
Lag 2	0.12	−0.01	0.01	0.05	0.20	·	·	·	·	·
	0.02	0.07	0.00	0.14	0.29	·	·	·	·	+
	0.15	0.11	0.00	0.06	0.26	·	·	·	·	·
	−0.03	−0.03	−0.15	−0.01	0.27	·	·	·	·	·
	−0.11	0.15	0.02	0.18	0.33	·	·	·	·	+
Lag 3	0.20	0.14	0.04	−0.01	−0.12	·	·	·	·	·
	0.33	0.13	0.15	0.09	−0.08	+	·	·	·	·
	0.18	0.03	0.14	0.12	−0.02	·	·	·	·	·
	0.21	0.02	0.15	0.03	−0.09	·	·	·	·	·
	0.26	−0.04	0.18	−0.01	−0.21	·	·	·	·	·
Lag 4	−0.01	0.03	0.09	−0.03	−0.07	·	·	·	·	·
	−0.20	0.05	−0.06	−0.05	0.02	·	·	·	·	·
	−0.07	−0.18	−0.09	−0.04	−0.18	·	·	·	·	·
	0.00	0.17	−0.22	0.17	0.02	·	·	·	·	·
	−0.02	0.13	−0.18	0.11	0.05	·	·	·	·	·

Table 1: Sample CCM and simplified sample CCM constructed from the single-period mortality reduction factors at 5 representative ages: 70, 75, 80, 85, and 90.

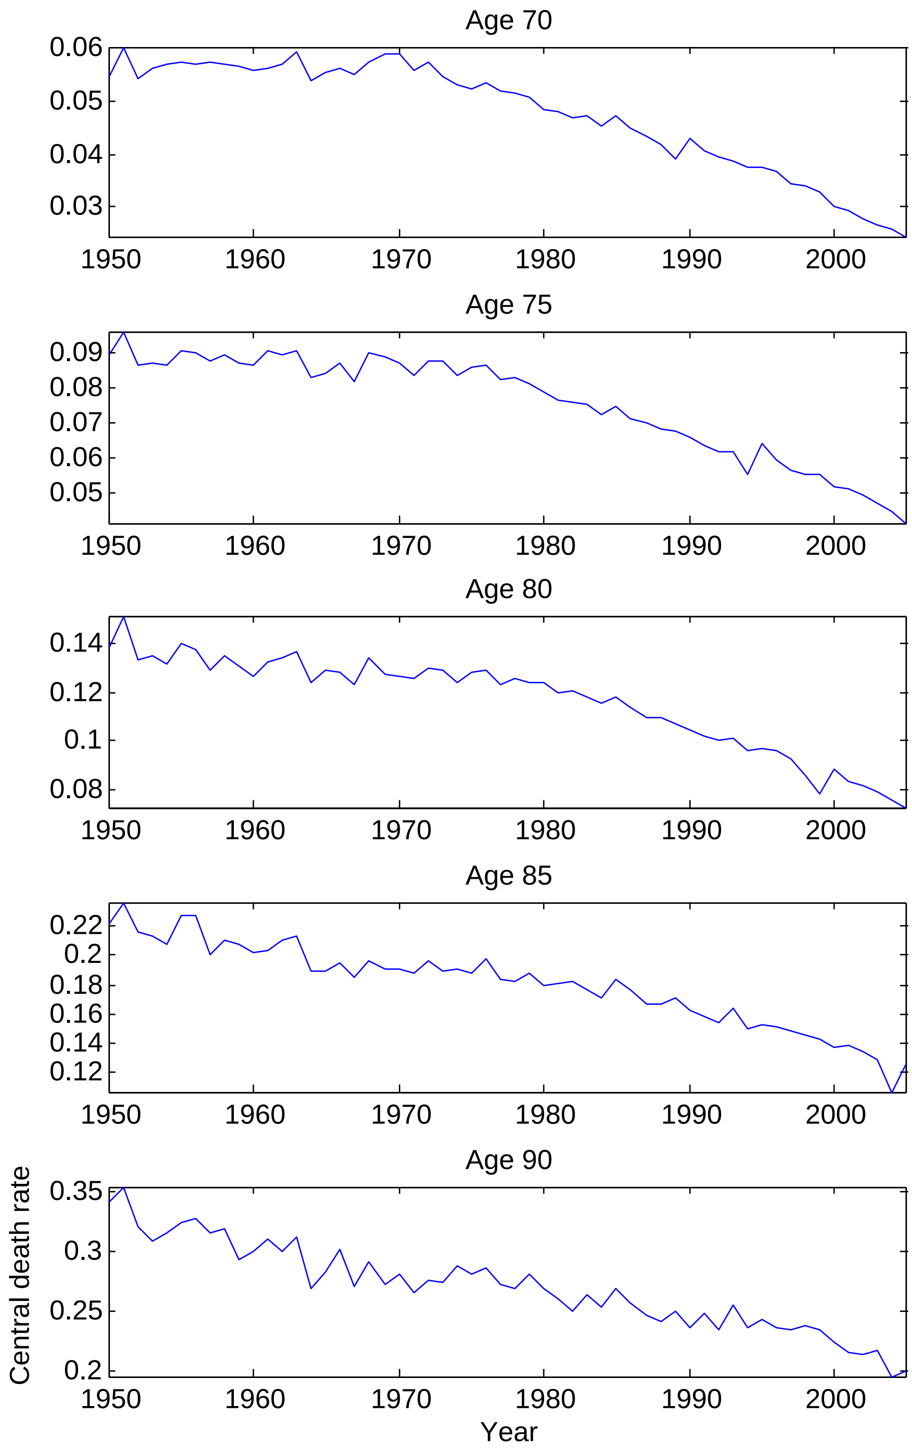

Figure 2: Central death rates at representative ages.

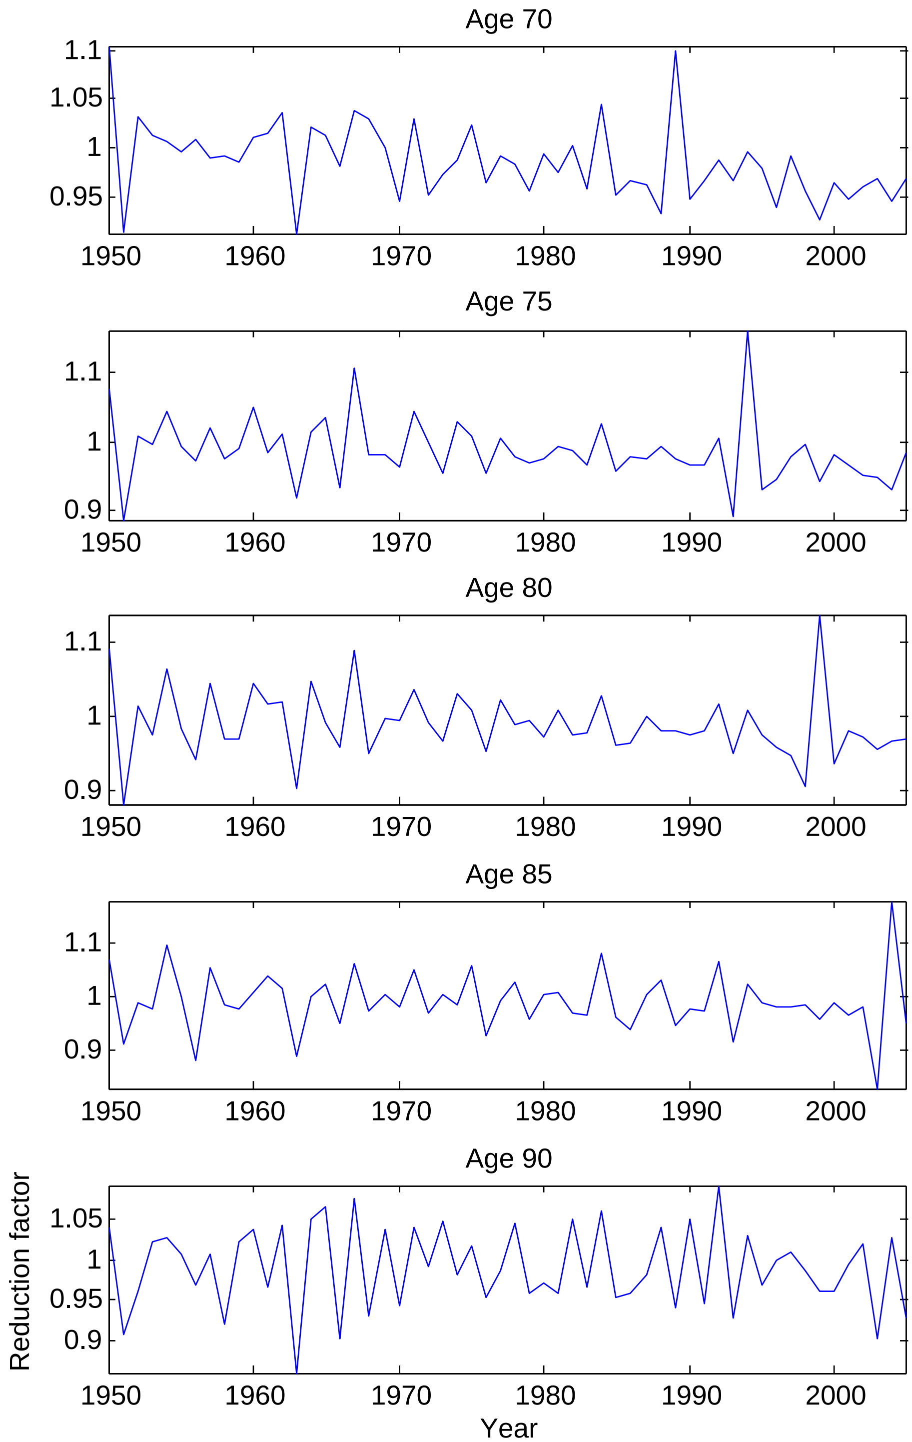

Figure 3: Mortality reduction factors at representative ages.

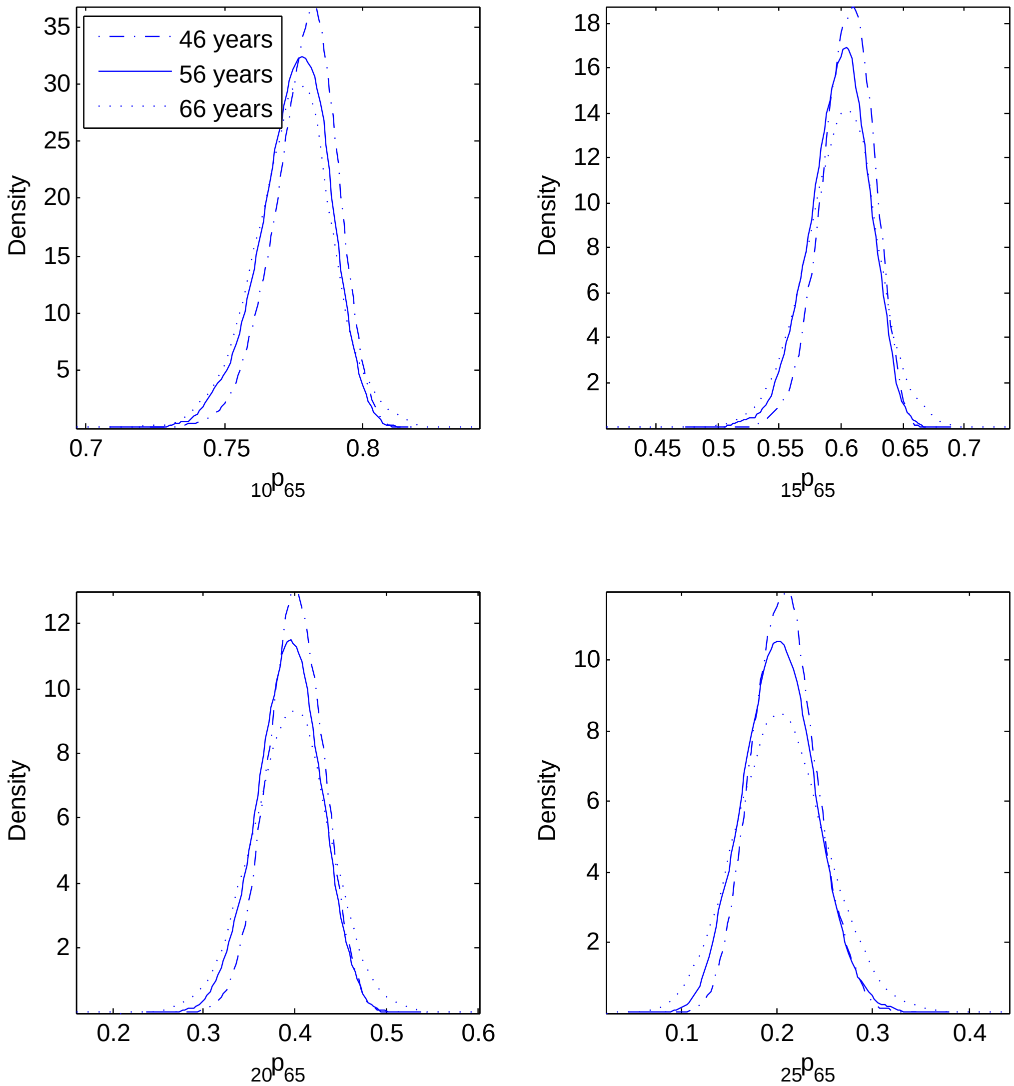

Figure 4: Empirical distributions of the survival probabilities for the cohort aged 65 in year 2005, on the basis of 46 years of data (dot-dash line), 56 years of data (solid line) and 66 years of data (dotted line).

	Sample period		
	1960–2005	1950–2005	1940–2005
${}_{10}p_{65}$	0.7790 (0.7541, 0.7987)	0.7753 (0.7475, 0.7971)	0.7749 (0.7449, 0.8029)
${}_{15}p_{65}$	0.6048 (0.5607, 0.6422)	0.5981 (0.5467, 0.6408)	0.5998 (0.5389, 0.6560)
${}_{20}p_{65}$	0.3999 (0.3385, 0.4584)	0.3928 (0.3226, 0.4579)	0.3954 (0.3084, 0.4798)
${}_{25}p_{65}$	0.2080 (0.1465, 0.2748)	0.2030 (0.1330, 0.2776)	0.2057 (0.1157, 0.3028)

Table 2: Estimates of the survival probabilities for the cohort aged 65 in year 2005. (The corresponding 95% confidence intervals are shown in parentheses.)

	Non-parametric	Lee-Carter	Two-factor
${}_{10}p_{65}$	0.7790	0.7755	0.7814
${}_{15}p_{65}$	0.6048	0.6011	0.6135
${}_{20}p_{65}$	0.3999	0.3995	0.4132
${}_{25}p_{65}$	0.2080	0.2039	0.2146

Table 3: Central estimates of the survival probabilities for the cohort aged 65 in year 2005, on the basis of the non-parametric bootstrap, the Lee-Carter model and the two-factor model.

t	$K(t)$	t	$K(t)$	t	$K(t)$	t	$K(t)$	t	$K(t)$
1	0.9800	6	0.8954	11	0.7813	16	0.6324	21	0.4512
2	0.9648	7	0.8754	12	0.7542	17	0.5991	22	0.4119
3	0.9488	8	0.8540	13	0.7257	18	0.5645	23	0.3727
4	0.9320	9	0.8312	14	0.6958	19	0.5280	24	0.3335
5	0.9143	10	0.8070	15	0.6646	20	0.4900	25	0.2951

Table 4: Values of the fixed leg, $K(t)$, for $t = 1, 2, \ldots, 25$.

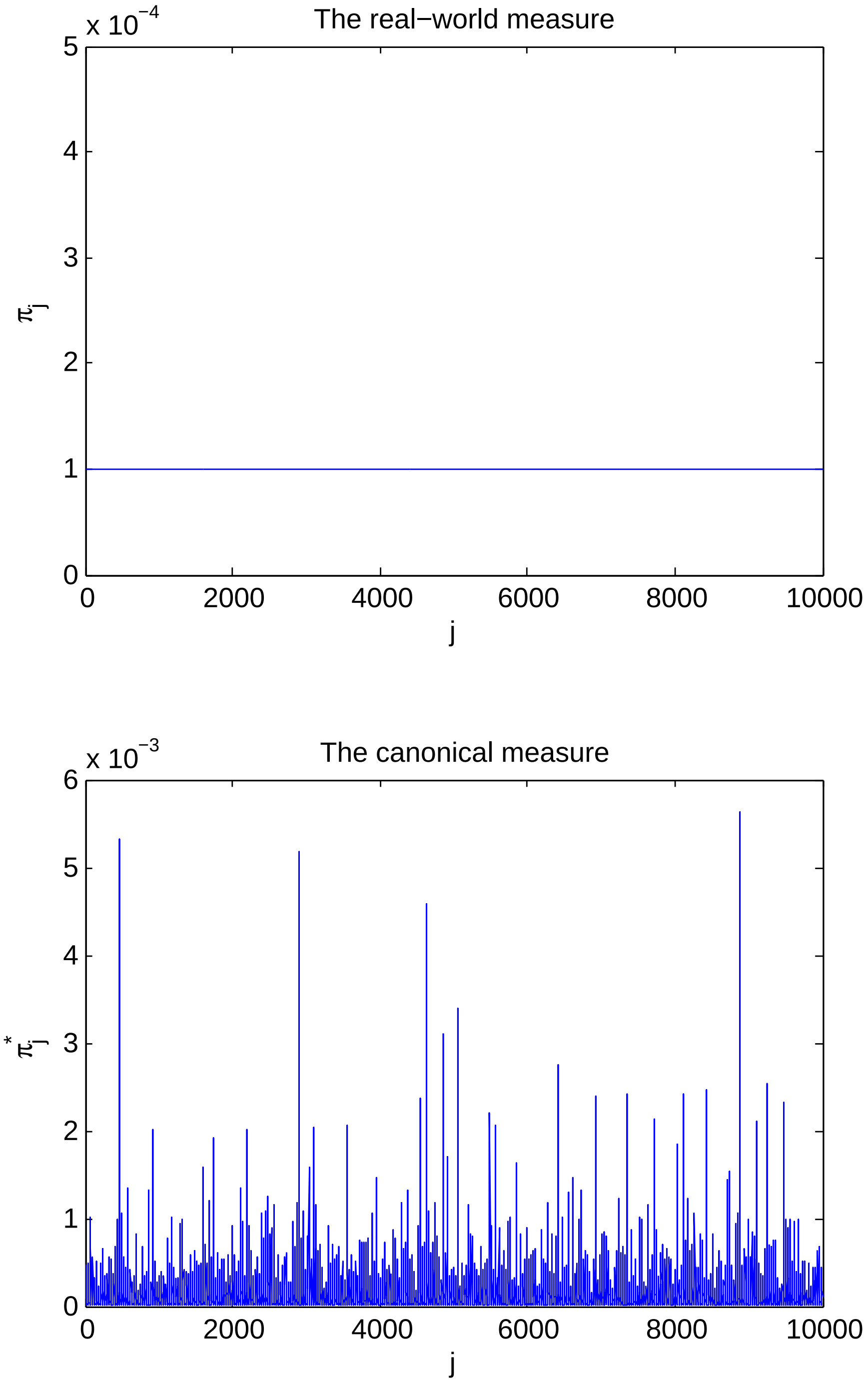

Figure 5: Probability distribution of $v(\omega)$ under the real-world measure P and the canonical measure Q_0.

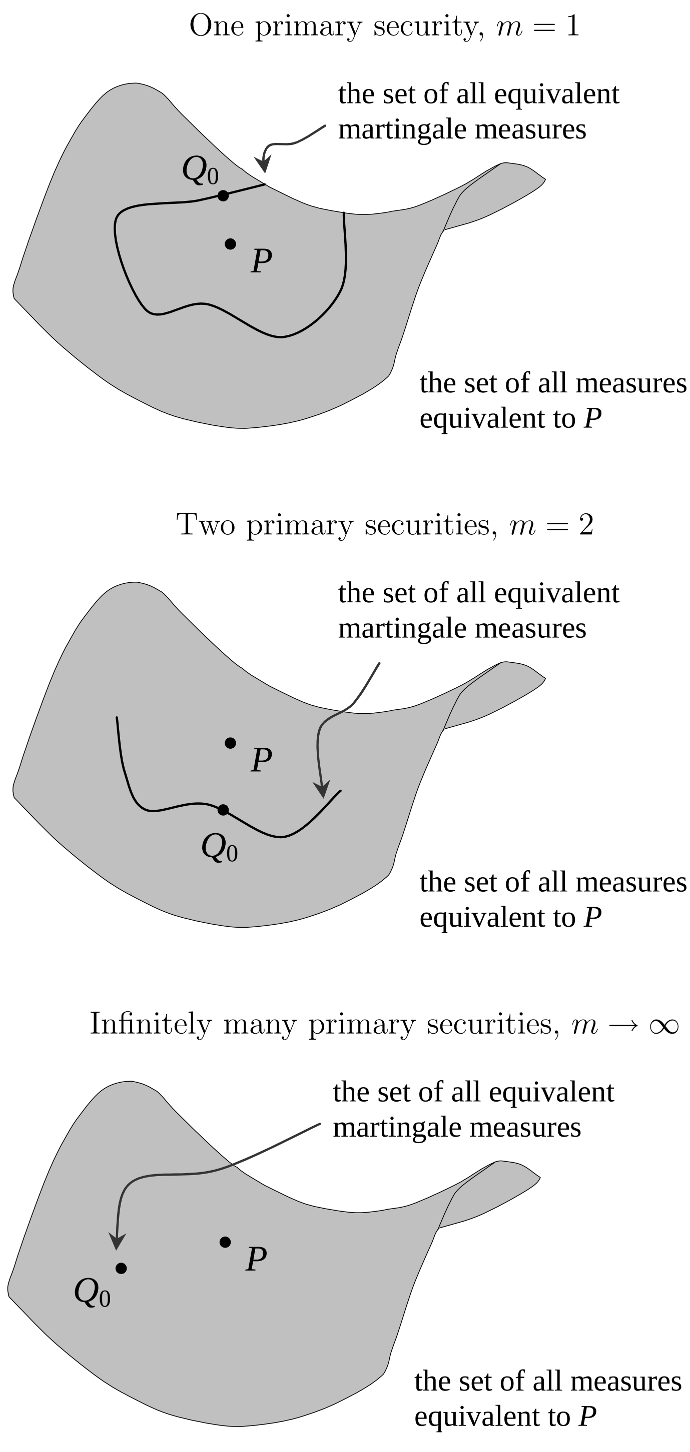

Figure 6: The canonical measure Q_0 derived from different numbers of primary security.

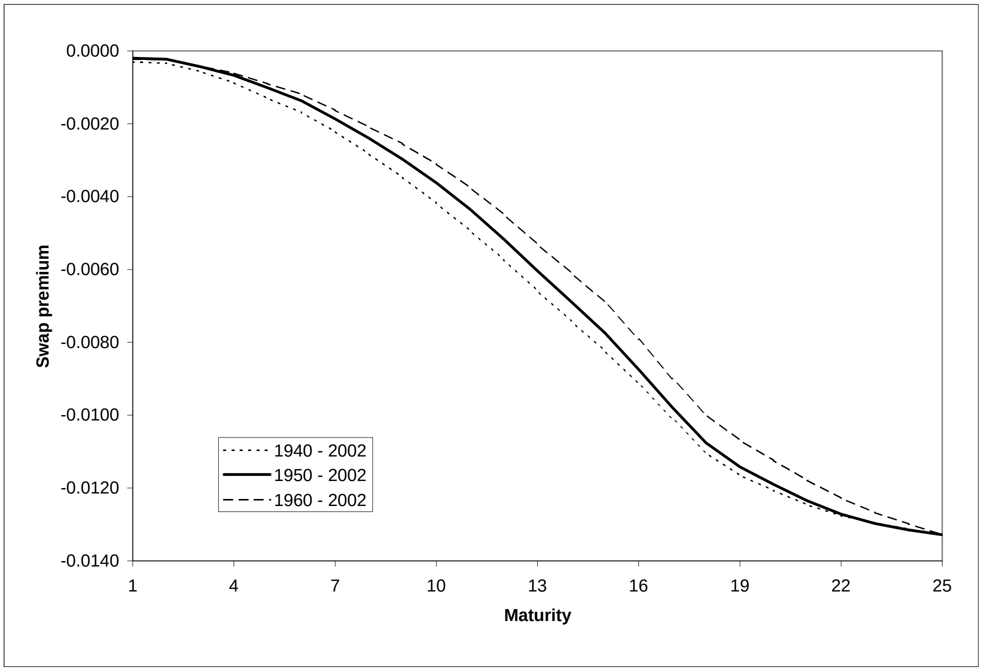

Figure 7: The swap premium θ calculated from the method of canonical valuation.

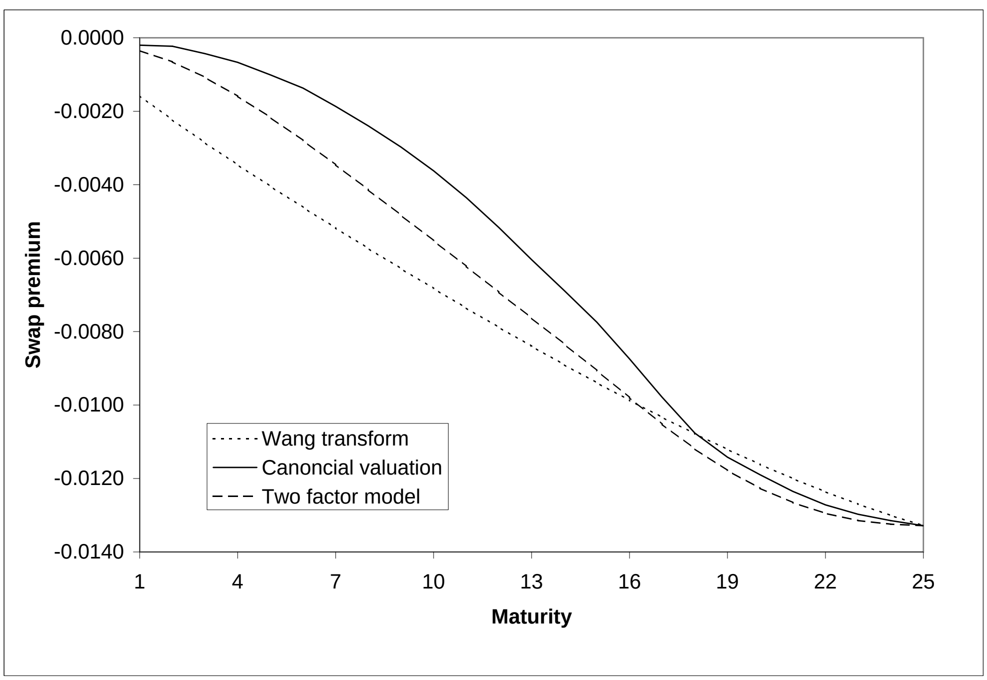

Figure 8: The swap premium θ calculated from different pricing methods.