




# Longevity Pricing Framework

A framework for pricing longevity exposures  
developed by the **LLMA**  
(Life & Longevity Markets Association)

29 October 2010



### **A framework for pricing longevity exposures developed by the LLMA**

This document describes a simple and transparent framework for pricing longevity exposures that has been developed by the Life & Longevity Markets Association. The purpose of the framework is to provide market participants with a means by which they can communicate pricing in a clear, transparent and unambiguous manner.


It is important to note that this framework is not intended to generate firm prices upon which longevity transactions can be executed. Rather it is a framework intended to facilitate a shared understanding of the nature of a particular longevity exposure and a 'benchmark price' against which an actual transaction price may be compared. It should also help longevity hedgers compare the pricing of different hedging solutions on a comparable basis.

A useful analogy may be drawn with the Black-Scholes price of an option. In the financial markets the prices at which options are traded are based on models that are generally proprietary and more complex than the Black-Scholes formula. Nevertheless, the Black-Scholes formula plays a valuable role in the market by providing a means of understanding and communicating actual market prices. In particular, it provides both a benchmark price against which actual option prices may be compared and the market's standard metric of implied volatility.

This document is the LLMA's second framework publication, the first being the *LLMA Longevity Index Framework* (LLMA (2010a)). Associated with this document are product technical notes on q-forwards and S-forwards and pricing spreadsheets which are available from the website.


The longevity pricing framework is the product of the LLMA's Technical Workstream and has benefited from the input of specialists from each of its current member organisations representing the insurance, re-insurance and investment banking industries. In developing this framework the LLMA has sought to capture the knowledge and experience of a broad set of individuals with different perspectives from across these firms.

## About the LLMA



The Life and Longevity Markets Association ('LLMA') is a non-profit organisation funded by its members which are drawn from across the insurance, reinsurance and banking industries. The LLMA aims to promote the development of a liquid traded market in longevity and mortality-related risk, through the development of consistent standards, methodologies, benchmarks and best practice.


## Disclaimer



This framework document is not intended to provide advice of any kind, including investment advice. No liability is accepted and no representation is made by the LLMA, its employees, agents or members in relation to the contents of this paper, including in relation to any errors or omissions. No independent verification of the information in this paper has been undertaken. No reliance should be placed on this paper.


Any listings of websites is provided purely for the convenience of readers. The LLMA has no control over the contents of such websites or the services that maybe be provided through them. Any reference, or hyper-link, to these sites does not constitute endorsement by the LLMA, and the LLMA makes no representation and accepts no liability, as to their contents. The Internet is not a failsafe means of communication and interruptions to service, through failure of software or hardware outside the control of the LLMA, may occur. The LLMA shall not have any liability for any failure in communication or service.





	Page
<b>Longevity pricing framework</b>	2
<b>About the LLMA</b>	3
<b>Disclaimer</b>	4
<b>1 Introduction</b>	6
<b>2 Longevity pricing framework: Overview</b>	7
<b>3 The key pricing inputs</b>	9
<b>3.1 Structure of the product</b>	9
<b>3.2 Data on the reference lives</b>	9
<b>3.3 Base mortality rates</b>	9
<b>3.4 Expected mortality improvements</b>	10
<b>3.5 Risk premium</b>	12
<b>4 Determining the price</b>	13
<b>5 Examples</b>	14
<b>5.1 Pricing a q-forward contract</b>	14
<b>5.2 Pricing an S-forward contract</b>	19
<b>References</b>	26

# 1. Introduction



The Life & Longevity Markets Association (LLMA) has developed a simple and transparent framework for pricing longevity exposures. The purpose of this framework is to provide a standard conceptual basis and a standard methodology for arriving at a benchmark price for any longevity exposure that is well understood, easily communicated and replicable by all market participants.


The LLMA believes that this kind of standardisation will help to promote its objective of encouraging the development of a liquid traded market in longevity risk transfer. Such standardisation should encourage much needed consistency in the complex area of longevity risk transfer, thereby providing greater comfort to first-time counterparties and reducing the time spent on due diligence for these transactions.

It is important to note that this framework is not intended to be used to produce firm prices upon which longevity transactions can be agreed and executed. Instead it is intended to facilitate a shared understanding of the nature of a particular longevity exposure and produce a 'benchmark price' against which an actual transaction price may be compared. It should also help longevity hedgers compare the pricing of different hedging solutions on a comparable basis.

The LLMA acknowledges that, in a sense, this framework is obvious and simplistic. Nevertheless, it is the experience of LLMA members from their involvement in longevity risk transfer transactions that pricing discussions and due diligence would be enhanced by a common understanding of pricing among all transaction stakeholders. This includes fiduciaries, advisors, asset managers and custodians, as well as counterparties. Furthermore, the market as a whole would benefit from a broader understanding of pricing principles across the industry at large, built on a shared conceptual basis. It is our hope that the longevity pricing framework presented in this document will help make this a reality.

Associated with this publication, product termsheets, technical notes and pricing spreadsheets are available from the LLMA website. The product technical notes are: *Technical note: The q-forward* (LLMA (2010b)) and *Technical note: The S-forward* (LLMA (2010c)).

This publication is organised in the following way. Section 2 gives an overview of the longevity pricing framework. Then Section 3 describes the key inputs required for pricing longevity exposures. This is followed by a brief description of how the inputs are combined to give a price in Section 4. Finally, Section 5 is devoted to two examples illustrating how the pricing framework can be applied.



## 2. Longevity pricing framework: Overview

Determining a price for any risk exposure, be it an exposure to financial risk or to insurance risk, requires a number of different inputs. Broadly speaking these may be summarised as follows:

- The definition of the risk exposure in question
- The current values of the input variables for the pricing algorithm or formula
- A method for determining future values of those variables

These inputs are then fed into an appropriate pricing algorithm or formula to determine the price.


As an example, let us consider the inputs required to price a floating-coupon bond. Pricing a floating-coupon bond requires the definition of the bond (e.g., maturity, floating index upon which the coupon is based, spread over the floating index, principal, credit rating, etc.), plus the current spot interest rate yield curve (e.g., spot swap curve and credit spread curve), plus the forward interest rate yield curve (e.g., the forward swap curve and forward credit spread curve). These inputs are then fed into the appropriate bond pricing formula (e.g., for a bullet bond the price is given by the sum of the present value of each payment cash flow discounted at the appropriate discount rate).

In the specific context of longevity risk, there are five key longevity-specific inputs required to price a longevity exposure. These are:

- A. Structure of the product
- B. Data on the reference lives
- C. Base mortality rates
- D. Expected mortality improvements
- E. Risk premium

Each of these inputs is described in detail below, and a brief summary of what is meant by each of these is given in Table 1. To determine the price of the exposure, these five inputs are fed into the relevant pricing formula, together with the appropriate discount rates.

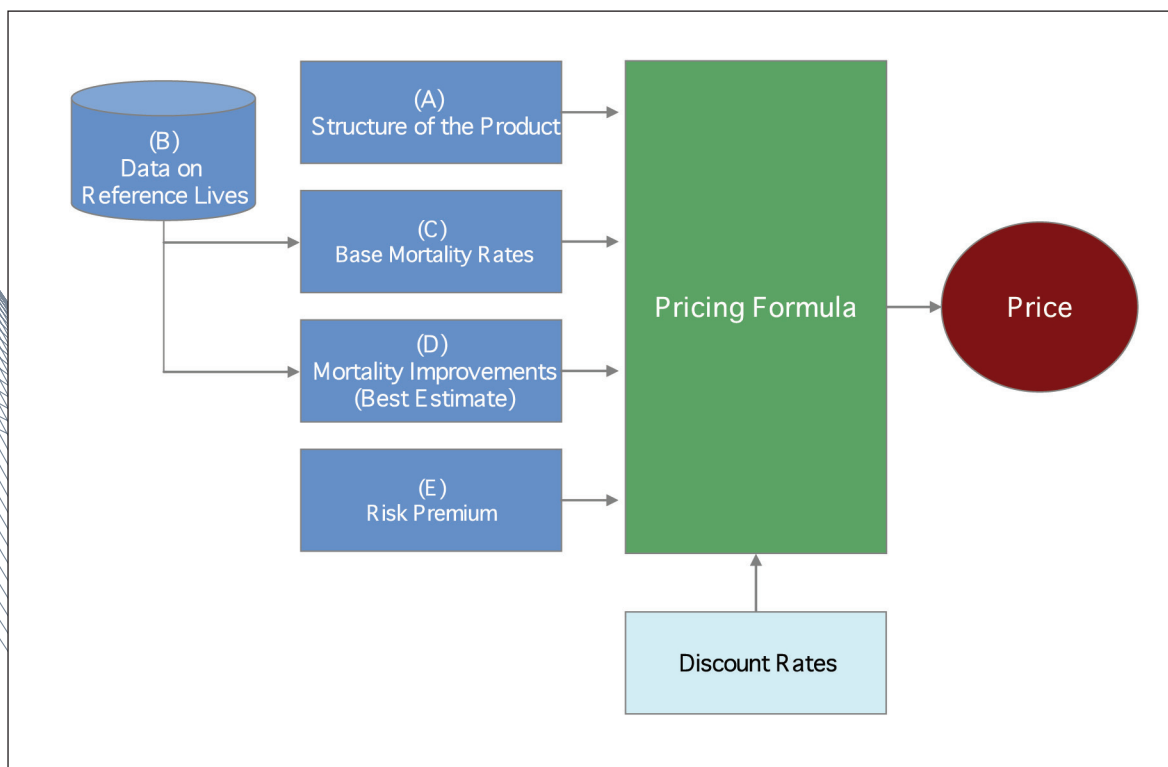
Figure 1 illustrates conceptually how the various inputs fit within the longevity pricing framework. The pricing formula requires information on the structure of the product, the base mortality rates for the reference population, the projected mortality improvements for the reference population, the longevity risk premium and the discount rates used to present value future cash flows. These are the direct inputs into pricing. The remaining input – the data on the reference lives – is an indirect input, which is used to determine base mortality rates and expected mortality improvements.



**Table 1.** Key inputs for the longevity pricing framework


Input	Description
A Structure of the product	The details of the product or instrument that is being priced, e.g., the structure of a pension benefit, an annuity or a longevity swap
B Data on the reference lives	Historical experience data on the mortality of the population of individuals on which the exposure is based
C Base mortality rates	The current or most recent set of mortality rates for the population in question. It forms the basis from which future mortality rates are projected
D Expected mortality improvements	The projected relative (percentage) changes in mortality rates with respect to an initial base mortality table. These essentially define the expected future levels of mortality rates – i.e., best estimate mortality rates – that are needed for pricing longevity
E Risk premium	A charge paid by longevity hedgers to compensate longevity investors for taking on longevity risk

**Figure 1.** Longevity pricing framework





## 3. The key pricing inputs



In this section we describe each of the five key inputs in the pricing framework in greater detail.

### **3.1 Structure of the product**

The 'structure of the product' refers to the detailed definition of the particular instrument being priced. For example, if we are pricing a pension, then it refers to the specific structure of the pension benefit, including the pension amount, payment frequency, the escalation rate, the pension amount payable to the spouse in the event of death, details of any option to retire early or any option to take a lump sum instead of the pension, etc. Similarly for an annuity, the structure refers to the annuity amount, escalation rates, benefits paid to dependents on the event of death, details of any guaranteed minimum payments, etc.

Other longevity-related products include capital markets instruments, such as derivatives used to transfer longevity risk. Examples of these include mortality forward rate contracts, known as 'q-forwards', survivorship forward rate contracts, known as 'S-forwards' and survivor swaps, which are commonly known simply as 'longevity swaps'.

In all these examples, the structure of the product simply defines the amount and timing of cash flows (including both deterministic and contingent/stochastic cash flows).

### **3.2 Data on the reference lives**


The pricing of longevity clearly requires information on the underlying pool of lives that are referenced in the transaction. This is equally true for specific pools of lives relating to the members of a particular pension plan or the beneficiaries of an annuity portfolio, as it is for the pool of lives comprising the national population.

This information typically includes:

- Historical experience data reflecting the size of the population at particular dates in the past and the deaths occurring in each historical period
- Demographic data on the lives relating, in particular, to age, gender, and spouse details
- Other data on the lives such as their addresses (in particular, postcode), pension amounts and occupations, all of which help provide a perspective on the socio-economic classification of each individual, which in turn influences projected mortality rates.

### **3.3 Base mortality rates**

Base mortality rates, or a base mortality 'table', refer to the current or most recent set of period mortality rates for the population of lives in question. Base mortality rates typically form the basis from which future mortality rates are projected by applying expected mortality 'improvements' to the base mortality rates.



There are different ways to arrive at a base mortality table. For example:

- *Standard table*: A number of pre-calculated standard mortality tables are available any of which may be used for the mortality base rates for a particular exposure.
- *Historical experience data*: Where good historical experience data exist for the population of lives in question, these data can be used to derive a bespoke base mortality table that is specific to these lives. This often involves a process called 'graduation' to produce 'graduated mortality rates'.
- *Standard table + Historical experience data*: This involves a combination of the first two methods above in which available data are fitted to standard tables based on goodness of fit criteria. This is necessary when there is insufficient data available to perform the analysis on experience data alone. This approach can be supplemented by postcode analysis to derive mortality rates for different socioeconomic groups.

### **3.3.1 Graduation and graduated mortality rates**


Typically, the base mortality rates used for pricing are what are called 'graduated' mortality rates. Graduated mortality rates are constructed from crude mortality rates using a smoothing, or 'graduation', procedure. The graduation of mortality rates is desirable to reduce noise and is justified because mortality rates for adjacent ages are similar and highly correlated. Graduation captures additional information about the mortality rate at a particular age from the rates at nearby ages. As a result, graduation does not destroy the integrity of the data, but instead brings the twin benefits of simplification and noise reduction. Graduated mortality rates are important in the valuation of longevity exposures and the calculation of life expectancy.

Graduation is a technical subject that we do not wish to go into detail in this document. But for any transaction, the parties involved will need to have clear agreement on an objective and well-defined methodology for graduation that will be applied consistently throughout its life. See *LLMA Longevity Index Framework* (LLMA (2010a)) for a more detailed discussion on graduation.

### **3.4 Expected mortality improvements**

The term 'mortality improvements' refers to relative changes in mortality rates with respect to an initial base mortality table. They are typically quoted as a percentage. For example a 1% mortality improvement means that next year's mortality rate is 99% of the current year's rate.

The pricing of longevity exposures requires an estimate of the expected mortality rates, so-called 'best estimate' mortality rates, at different times in the future. A convenient way of describing these best estimate mortality rates is in terms of expected, or best estimate, mortality improvements relative to a particular mortality base table. If, for example, expected mortality improvements are 1% per year, this means that every year mortality rates are 99% of the previous year's rates. So after 10 years, the mortality rate will have fallen to  $(100\% - 1\%)^{10} = 90.4\%$  of the base mortality rate.




We can use a number of different methods for determining the expected mortality improvements to be incorporated into a pricing process:

- *Standard mortality improvement tables:* A number of pre-calculated mortality improvement tables are available. Examples of these include the Scale-AA table in the US (see Society of Actuaries (2000)) and the tables published by the UK's Continuous Mortality Investigation (see for example CMI (2007) and CMI (2009)).
- *Mortality projection based on own experience data:* This involves developing a projection for future mortality improvements from the actual historical experience data for the specific pool of lives in question.
- *Mortality projection based on a different population:* This involves a projection for future mortality improvements from the historical experience data for a different pool of lives, such as the national population, or the population of assured lives.

Regardless of which population they are based on, projections based on experience data can be performed in different ways. For example:

- *Historical mortality improvements:* If sufficient historical data are available we can calculate the mortality improvements that have been realised in the past. These historical improvements can then be used as an estimate of expected future improvements. This approach has the advantage of simplicity, but may not be appropriate where large cohort effects are present.
- *Mortality projection models:* A large number of projection models are available that generate future mortality rates from historical death and population data. These models are much more complex than projections based directly on historical improvements, however they are more realistic, for example in the way they deal with cohort effects. Such models include the well-known Lee-Carter model (Lee & Carter (1992)), the Renshaw and Haberman model (Renshaw & Haberman (2006)), other cohort extensions to the Lee-Carter model such as Debonneuil (2010), so-called p-spline models (Durban et al. (2002)) and the Cairns-Blake-Dowd series of models (Cairns et al. (2006) and (2007)). Note that this list is for illustration only and is not exhaustive.

*Mortality projection models fall into two general classes:* deterministic and stochastic. It should be noted that stochastic models are often cumbersome to use for the purpose of pricing, but are clearly very useful for risk assessment.




### **3.5 Risk premium**

The risk premium is the final piece in the puzzle and represents the charge, or cost, that a hedge provider would levy to take on longevity risk from a party wishing to reduce or eliminate it. The risk premium may be incorporated into pricing in different ways, e.g., through adjustments to the base mortality rates, adjustments to improvement rates, upfront payments, or a combination of all of these.

There are a number of potential drivers that determine the size of the risk premium. These include:

- Investor related factors:
  - Investor demand
  - Return available on other products (opportunity cost)
  - Diversification benefit
  - Maturity of instrument that is transferring the risk
  - Ability of intermediary to provide liquidity
- Hedge provider/intermediary/reinsurer related factors:
  - Capital implications
  - Risk appetite
  - Return on capital
  - Diversifying positions
- Other factors:
  - Number of lives in the pool
  - Quality and amount of data available
  - Size of liability/underlying risk
  - Type of instrument that is transferring the risk
  - Amount of credit risk in the transaction





## 4. Determining the price


In addition to the five key inputs discussed above, the pricing framework also requires discount rates to calculate the present values of future cash flows. All these inputs are then fed into the appropriate pricing formula or algorithm to determine the price of the product or instrument.

The exact pricing formula depends on the nature of the product being priced, but it essentially involves a number of elements that are common to all products:


- Determination of the timing and size of cash flows generated by the product, including the risk premium
- Calculation of the present value of those cash flows using the appropriate discount rates
- Averaging over all potential scenarios, or outcomes, for mortality rates in an appropriate way.

For products that don't involve any option-like behaviour, there is a relatively simple way of addressing the third element, the averaging over mortality scenarios. This simplified approach involves using 'best estimate' mortality rates in addition to the risk premium to determine the size of the cash flows. Best estimate mortality rates reflect the unbiased expectation for the realised mortality rates at different times in the future. They essentially correspond to the average over different mortality scenarios.

Using this approach, the pricing process simplifies to (i) determining the product's cash flows based on best-estimate mortality rates and the risk premium, and then (ii) discounting those cash flows using the appropriate discount rates.



## 5. Examples



In this section we apply the pricing framework to two simple examples involving the pricing of a q-forward (a mortality forward-rate contract) and an S-forward (a survivorship forward rate contract). These are simple derivative contracts for transferring longevity risk that can be used as building-blocks to create practical longevity hedges for pension and annuity liabilities.

Both examples are based on the same population of males and use the same mortality data. For the purpose of calculating present values we use a flat discount rate of 5%.

### 5.1 Pricing a q-forward contract

The q-forward is in many ways the simplest instrument for transferring mortality and longevity risk. It involves a single exchange of cash flows between counterparties at the maturity date of the contract.

It is a cash-settled forward-rate contract linked to the mortality rate of a given population. The letter 'q' in the name of this contract corresponds to the symbol actuaries use to denote a mortality rate.


Suppose we wish to price a 10-year q-forward contract on a population of males currently aged between 55 and 59 years old, and who will be aged between 65 and 69 years old at maturity. Assume the pricing date is 1 January 2011 and the maturity date is 31 December 2020. For this particular example we assume that the q-forward settlement payment is based on the average mortality rate in 2020 across each of the five individual ages at the maturity date.

#### 5.1.1 Structure of the product

To be more specific, 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 (the floating leg) at that future date, in return for a fixed mortality rate (the fixed leg) agreed upon at the inception of the contract. The floating leg of the instrument references the uncertain future mortality rate of the reference population (i.e., males who will be between 65 and 69 in 10 years time). The fixed leg reflects the fixed mortality rate that a counterparty to the transaction would wish to receive for committing to pay the floating leg of the transaction. A counterparty hedging longevity risk will receive fixed and pay floating mortality. For more details on this instrument please see *q-Forwards: Derivatives for transferring longevity and mortality risk* (Coughlan et al. (2007)) and the LLMA publication *Technical note: The q-forward* (LLMA (2010b)).

#### 5.1.2 Data on reference lives

For this example, we assume we have good quality data on the mortality experience of the population in question which enables us to construct appropriate base mortality rates and mortality improvement projections.



### 5.1.3 Base mortality rates

Figure 2 shows the base mortality rates for the population for the year 2010, which reflects the inception of the transaction. The aggregate mortality rate for the population at this time is simply an average of the mortality rates for each of the individual ages (1.61%, 1.77%, 1.98%, 2.23%, 2.41%):


$$\begin{aligned} q_{\text{realised}}(2010) &= \text{Aggregate realised mortality rate in 2010} \\ &= 2.00\% \end{aligned}$$

**Figure 2.** Base mortality rates at the inception of the transaction.

AGE	q(2010)
65	1.61%
66	1.77%
67	1.98%
68	2.23%
69	2.41%
Average	2.00%

### 5.1.4 Mortality improvement rates

We assume that best estimate mortality improvements for this population of 65-69 year-old males are 2% per year. This means that the best estimate mortality rate in any year is 98% of that in the previous year. Moreover, the best estimate mortality rate at the maturity of the contract after 10 years is 81.7% of the base mortality rate at the inception of the contract. These mortality improvements are incorporated into the best estimate mortality table shown in Figure 3.



The aggregate best estimate mortality rate in 2020 for the group of 65-69 year-old males is a simple average of the best estimate mortality rates for the individual ages (1.32%, 1.45, 1.62%, 1.82%, 1.97%). The average of these rates yields the aggregate best estimate mortality rate for the q-forward as follows:

$$\begin{aligned}
 q_{BE}(2010:2020) &= \text{Aggregate best estimate mortality rate in 2020} \\
 &\quad \text{as determined at transaction inception} \\
 &= q_{\text{realised}}(2010) \times (100\% - 2\%)^{10} \\
 &= 1.63\%
 \end{aligned}$$


Here the notation  $q_{BE}(2010:2020)$  reflects the fact that the best estimate in 2020 is estimated based on data in 2010.

If there were no risk premium for transferring longevity risk, then this best estimate mortality rate would correspond to the fixed rate in the q-forward transaction. However, the need to pay a risk premium in order for an investor to take on the longevity risk means that the fixed rate will actually be lower than the best estimate, as we shall see in the next subsection.

**Figure 3.** Best estimate mortality table incorporating best estimate mortality improvements of 2% per year applied to the base mortality rates for 2010. (Note that ‘age’ refers to the age at the start of the year in question).

AGE / YEAR	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
65	1.61%	1.58%	1.55%	1.52%	1.49%	1.46%	1.43%	1.40%	1.37%	1.34%	1.32%
66	1.77%	1.73%	1.70%	1.67%	1.63%	1.60%	1.57%	1.54%	1.51%	1.48%	1.45%
67	1.98%	1.94%	1.90%	1.86%	1.83%	1.79%	1.75%	1.72%	1.68%	1.65%	1.62%
68	2.23%	2.19%	2.14%	2.10%	2.06%	2.02%	1.98%	1.94%	1.90%	1.86%	1.82%
69	2.41%	2.36%	2.31%	2.27%	2.22%	2.18%	2.13%	2.09%	2.05%	2.01%	1.97%
<b>Average</b>	2.00%	1.96%	1.92%	1.88%	1.84%	1.81%	1.77%	1.74%	1.70%	1.67%	1.63%





### 5.1.5 Risk premium

Assume that in this example the risk premium is quoted in terms of an increase in the level of mortality improvements of 1.00% per year on top of best-estimate improvements. In other words if best-estimate mortality improvements are 2% per year, then to take account of the risk premium we need to calculate cash flows based on mortality improvements of 3% per year. This is illustrated in Figure 4.

*NOTE: The risk premium used is an example taken purely for illustrative purposes. LLMA makes no statement as to whether or not this is an appropriate level for the risk premium in this or any other transaction.*

To calculate the impact of the risk premium we need to recalculate the aggregate mortality rate in 2020 taking account of this additional 1.00% per year mortality improvement. This aggregate mortality rate, reflecting best-estimate mortality improvements plus the risk premium corresponds to the fixed rate in the q-forward transaction. It is sometimes referred to as the 'mortality forward rate':


$$\begin{aligned}q_{\text{forward}}(2010:2020) &= \text{Fixed rate paid in 2020} \\ &\quad \text{as determined at transaction inception} \\ &= q_{\text{realised}}(2010) \times (100\% - 2\% - 1\%)^{10} \\ &= 1.47\%\end{aligned}$$

*Note that the forward rate differs from the best estimate rate by virtue of the risk premium. So the risk premium corresponds to a discount of 0.16 percentage points or 16 basis points relative to the best estimate mortality rate in 2020 (i.e., 1.63% – 1.47% = 0.16%).*

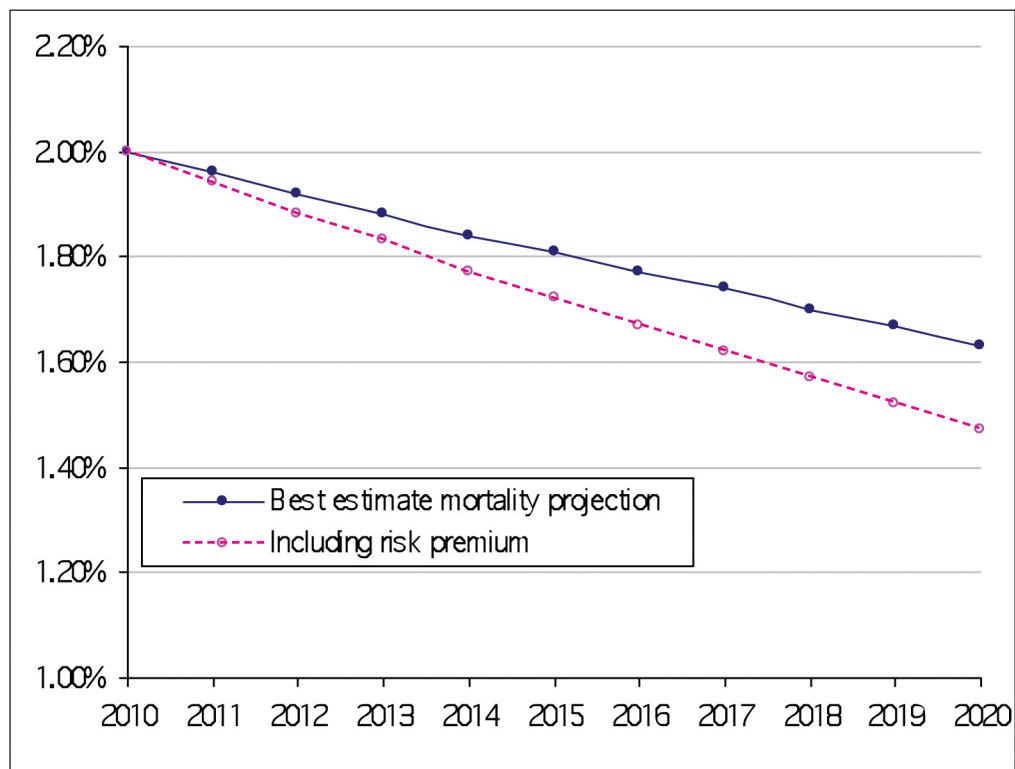
The size of the risk premium paid by the longevity hedger in monetary terms is given by the present value of the notional amount multiplied by difference between the best estimate mortality rate and the forward mortality rate. Assuming a discount rate of 5% then the present value of the risk premium is:

$$\begin{aligned}\text{PV of Risk Premium} &= \text{GBP}100\,000\,000 \\ &\quad \times (1.63\% - 1.47\%) / (1 + 5\%)^{10} \\ &= \text{GBP } 98\,226\end{aligned}$$

This reflects a reduction in the size of the payment for the fixed-rate side of the q-forward, since the fixed-rate receiver is hedging longevity risk and must pay the risk premium.



**Figure 4.** Best estimate mortality projection and the impact of the risk premium




### 5.1.6 Calculating the price

The price of this q-forward is simply the present value of the two legs of the contract: the fixed leg and the floating leg. The pricing of the fixed leg is straight forward – it is given by the present value of the notional amount multiplied by the forward rate. At inception of the transaction, if we ignore any bid-offer spread this is given by:

$$\begin{aligned}
 \text{PV of Fixed Leg} &= \text{GBP}100\,000\,000 \\
 &\quad \times 1.47\% / (1 + 5\%)^{10} \\
 &= \text{GBP } 902\,452
 \end{aligned}$$

Pricing of the floating leg involves calculating the floating cash flow taking account of both the best estimate mortality rate and the risk premium. At inception of the transaction, if we ignore any bid-offer spread, the price of the floating leg is exactly the same as that of the fixed leg. In other words the q-forward has a net value of zero at inception. (In practice the bid-offer spread leads to a non-zero value).



During the life of the transaction the value of the q-forward will deviate from zero as changes emerge in best estimate mortality projections and the size of the risk premium. Pricing the q-forward at any time during its life involves the same procedure described above. Specifically, the payoff from the fixed leg is present valued in the obvious way, and the payoff from the floating leg is calculated from the cash flows derived from the new best estimate mortality projection and the new risk premium. In practice this is done by determining the fixed leg available in the market at that time for the same floating leg. The value of the q-forward is given by the difference between the value of this new 'on-market' fixed leg and that of the actual fixed leg for the transaction.

## **5.2 Pricing an S-forward contract**


The S-forward is an intuitive instrument for transferring longevity risk since it directly relates to the survival rate for a given population over a particular period. The S-forward is the basic building-block from which the well-known longevity (survivor) swaps can be constructed. In particular, a longevity swap of the kind we have seen transacted in the market over the past two and a half years is essentially just a string of S-forwards maturing at different times in the future.

Like the q-forward, the S-forward also involves a single exchange of cash flows between counterparties at the maturity date of the contract. It is a cash-settled forward-rate contract linked to the survival rate of a given population.

Suppose we wish to price a 10-year S-forward contract on a population of males currently aged 65 years old, and who will be aged 75 years old at maturity. Assume the pricing date is 1 January 2011 and the maturity date is 31 December 2020.

### **5.2.1 Structure of the product**

To be more specific, an S-forward is an agreement between two counterparties to exchange at a future date (the maturity of the contract) an amount equal to the realized survival rate of a given population (the floating leg) at that future date, in return for a fixed survival rate (the fixed leg) agreed upon at the inception of the contract. The floating leg of the instrument references the uncertain future survival rate of the reference population (i.e., males who are currently 65). The fixed leg reflects the fixed survival rate that a counterparty to the transaction would wish to receive for committing to pay the floating leg of the transaction. A counterparty hedging longevity risk will receive floating and pay fixed survival. For more details on this instrument please see the LLMA publication. *Technical note: The S-forward* (LLMA (2010c)).



Survival rates are calculated from mortality rates as follows. The one-year survival rate is simply one minus the probability of death occurring over the following year, where the probability of death is the same as the initial mortality rate. We have the following relationship:

$$p_x = 1 - q_x$$

where:

$q_x$  is the 1-year initial mortality rate, for age group  $x$  at time  $t$

$p_x$  is the 1-year survival probability, for age group  $x$  at time  $t$

Ten-year survival rates, which are needed for this transaction, are simply a product of one-year survival rates for each of the ten years, taking account of the obvious fact that the individuals are one year older each year.


### **5.2.2 Data on reference lives**

For this example, we assume we have good quality data on the mortality experience of the population in question which enables us to construct appropriate base mortality rates and mortality improvement projections. From the realised and projected mortality rates we calculate realised and projected survival rates for this population, according to the method given above. It should be remembered that we start with a survival level of 100% at inception of the transaction.

### **5.2.3 Base mortality rates**

Figure 5 shows the base mortality rates for the population for the year 2010, which reflects the most recently available data at the inception of the transaction.





**Figure 5.** Base mortality rates at the inception of the transaction.

AGE	q(2010)
65	1.61%
66	1.77%
67	1.98%
68	2.23%
69	2.41%
70	2.61%
71	2.98%
72	3.30%
73	3.71%
74	4.18%

#### **5.2.4 Mortality improvement rates**


We assume that best estimate mortality improvements for this population are a flat 2% per year for each of the next 10 years. This means that the best estimate mortality rate in any year is 98% of that in the previous year. Moreover, the best estimate mortality rate at the maturity of the contract after 10 years is 81.7% of the base mortality rate at the inception of the contract. These mortality improvements are incorporated into the best estimate mortality table shown in Figure 6(a).

The mortality rates highlighted in the table correspond to those rates that are relevant for calculating the best estimate survival rate at the end of the year in question. Since the transaction start date is 1 January 2011, at which time the population is 65 years old, the relevant mortality rate for calculating the one-year survival rate is that for a 65 year old in 2011.

From Figure 6(a), we calculate the best estimate survival rate in 2020 as follows:

$$\begin{aligned} p_{BE}(2010:2020) &= \text{Best estimate survival rate between year-end 2010} \\ &\quad \text{and year-end 2020 as determined at inception} \\ &= 78.75\% \end{aligned}$$

This means that we can expect 78.75% of the initial population of males to be still alive at the end of 2020. The evolution of best estimate survival rates in different years is shown in Figure 6(b).




**Figure 6(a).** Best estimate mortality table incorporating best estimate mortality improvements of 2% per year applied to the base mortality rates for 2010. (Note that ‘age’ refers to the age at the start of the year in question).

AGE / YEAR	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
65	1.61%	1.58%	1.55%	1.52%	1.49%	1.46%	1.43%	1.40%	1.37%	1.34%	1.32%
66	1.77%	1.73%	1.70%	1.67%	1.63%	1.60%	1.57%	1.54%	1.51%	1.48%	1.45%
67	1.98%	1.94%	1.90%	1.86%	1.83%	1.79%	1.75%	1.72%	1.68%	1.65%	1.62%
68	2.23%	2.19%	2.14%	2.10%	2.06%	2.02%	1.98%	1.94%	1.90%	1.86%	1.82%
69	2.41%	2.36%	2.31%	2.27%	2.22%	2.18%	2.13%	2.09%	2.05%	2.01%	1.97%
70	2.61%	2.56%	2.51%	2.46%	2.41%	2.36%	2.31%	2.27%	2.22%	2.18%	2.13%
71	2.98%	2.92%	2.86%	2.80%	2.75%	2.69%	2.64%	2.59%	2.54%	2.48%	2.43%
72	3.30%	3.23%	3.17%	3.11%	3.04%	2.98%	2.92%	2.86%	2.81%	2.75%	2.70%
73	3.71%	3.64%	3.56%	3.49%	3.42%	3.35%	3.29%	3.22%	3.16%	3.09%	3.03%
74	4.18%	4.10%	4.01%	3.93%	3.86%	3.78%	3.70%	3.63%	3.56%	3.49%	3.42%

**Figure 6(b).** Best estimate survival rates from the best estimate mortality table. (Note that ‘age’ refers to the age at the start of the year in question, but the survival rates correspond to the end of the year in question).

AGE / YEAR	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
65		98.42%									
66			96.75%								
67				94.95%							
68					92.99%						
69						90.97%					
70							88.86%				
71								86.57%			
72									84.13%		
73										81.53%	
74											78.75%



### 5.2.5 Risk premium

Assume that in this example the risk premium is quoted in terms of an increase in the level of mortality improvements of 1.00% per year on top of best-estimate improvements. In other words if best-estimate mortality improvements are 2% per year, then to take account of the risk premium we need to calculate survival rate cash flows based on mortality improvements of 3% per year.

*NOTE: The risk premium used is an example taken purely for illustrative purposes. LLMA makes no statement as to whether or not this is an appropriate level for the risk premium in this or any other transaction.*

To calculate the impact of the risk premium we need to recalculate the survival rate at the end of 2020 taking account of this additional 1.00% per year mortality improvement. This survival rate, reflecting best-estimate mortality improvements plus the risk premium corresponds to the fixed rate in the S-forward transaction. It is sometimes referred to as the 'survival forward rate':


$$\begin{aligned} p_{\text{forward}}(2010:2020) &= \text{Forward survival rate between year-end 2010} \\ &\quad \text{and year-end 2020 as determined at inception} \\ &= 79.92\% \end{aligned}$$

Note that the forward rate differs from the best estimate rate by virtue of the risk premium. So the risk premium corresponds to an increase of 1.18 percentage points or 118 basis points relative to the best estimate survival rate at the end of 2020 (i.e., 79.92% – 78.75% = 1.18% with rounding). The evolution of the risk premium in terms of survival rates is shown in Figure 7.

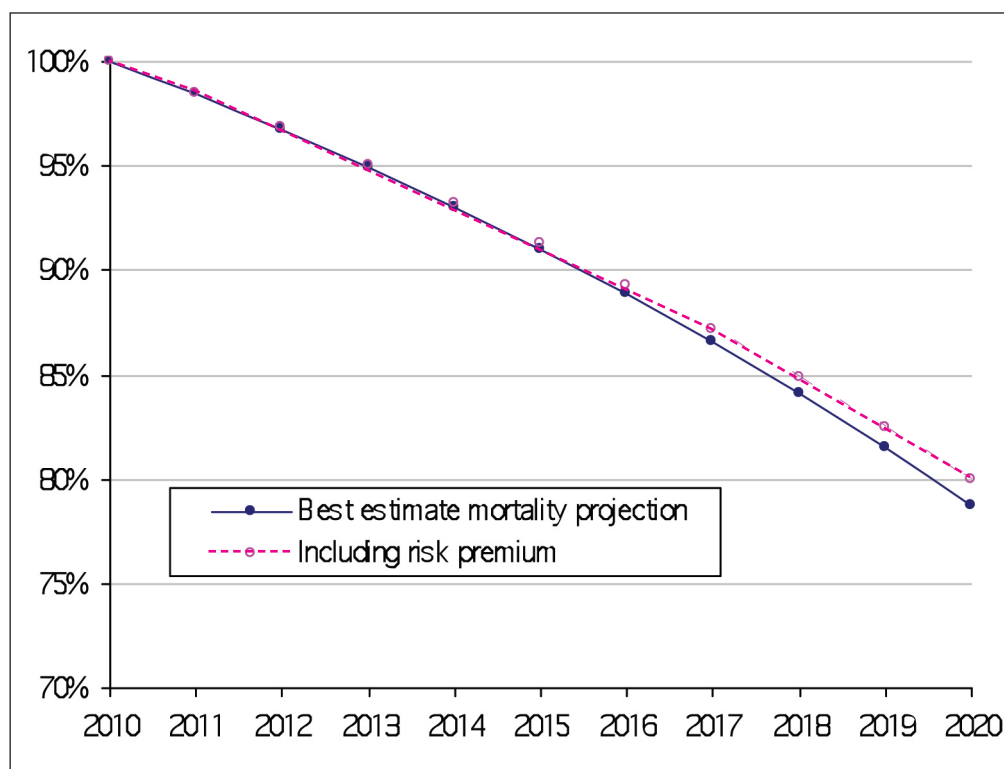
The size of the risk premium paid by the longevity hedger in monetary terms is given by the present value of the notional amount multiplied by difference between the best estimate survival rate and the forward survival rate. Assuming a discount rate of 5% then the present value of the risk premium is:

$$\begin{aligned} \text{PV of Risk Premium} &= \text{GBP100 000 000} \\ &\quad \times (79.92\% - 78.75\%) / (1 + 5\%)^{10} \\ &= \text{GBP 721 831} \end{aligned}$$


This reflects an increase in the size of the payment for the fixed-rate side of the S-forward, since the fixed-rate payer is hedging longevity risk and must pay the risk premium.



**Figure 7.** Best estimate survival rate projection and the impact of the risk premium. The dates shown in the chart correspond to 31 December in the given year.








### 5.1.6 Calculating the price

The price of this S-forward is simply the present value of the two legs of the contract: the fixed leg and the floating leg. The pricing of the fixed leg is straight forward – it is given by the present value of the notional amount multiplied by the forward survival rate. At inception of the transaction, if we ignore any bid-offer spread and assume a discount rate of 5% this is given by:

$$\begin{aligned} \text{PV of Fixed Leg} &= \text{GBP}100\,000\,000 \\ &\quad \times 79.92\% / (1 + 5\%)^{10} \\ &= \text{GBP } 49\,066\,165 \end{aligned}$$

Pricing of the floating leg involves calculating the floating cash flow taking account of both the best estimate mortality rate and the risk premium. At inception of the transaction, if we ignore any bid-offer spread, the price of the floating leg is exactly the same as that of the fixed leg. In other words the S-forward has a net value of zero at inception. (In practice the bid-offer spread leads to a non-zero value).

During the life of the transaction the value of the S-forward will deviate from zero as changes emerge in best estimate mortality projections and the size of the risk premium. Pricing the S-forward at any time during its life involves the same procedure described above. Specifically, the payoff from the fixed leg is present valued in the obvious way, and the payoff from the floating leg is calculated from the cash flows derived from the new best estimate mortality projection and the new risk premium. In practice this is done by determining the fixed leg available in the market at that time for the same floating leg. The value of the S-forward is given by the difference between the value of this new 'on-market' fixed leg and that of the actual fixed leg for the transaction.



### **From the LLMA:**

LLMA (2010a). *LLMA Longevity Index Framework*. [www.llma.org]

LLMA (2010b). *Technical note: The q-forward* [www.llma.org]

LLMA (2010c). *Technical note: The S-forward* [www.llma.org]

### **From other sources:**

Cairns, A.J.G., Blake, D., & Dowd, K. (2006). "A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration." *The Journal of Risk and Insurance*, 73 (4): 687-718.

Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., Epstein, D., Ong, A., and Balevich, I. (2007) "A quantitative comparison of stochastic mortality models using data from England & Wales and the United States." *North American Actuarial Journal*, 13 (1): 1-35.

CMI (2007). "The CMI library of mortality projections", Working Paper 30, Institute of Actuaries and Faculty of Actuaries: London, 23 November. [www.actuaries.org.uk/research-and-resources/pages/continuous-mortality-investigation-working-papers]

CMI (2009). "The graduations of the CMI Self-Administered Pension Schemes 2000-2006 mortality experience: Final 'S1' Series of Mortality Tables", Working Paper 35, Institute of Actuaries and Faculty of Actuaries: London, February. [www.actuaries.org.uk/research-and-resources/pages/continuous-mortality-investigation-working-papers]

Coughlan, G.D., Epstein, D., Sinha, A. and Honig, P. (2007). "q-Forwards: Derivatives for transferring longevity and mortality risk", J.P. Morgan: London, 2 July [www.lifemetrics.com].

Debonneuil, E. (2010) "A simple model of mortality trends aiming at universality: Lee Carter + Cohort". Arxiv, Populations and Evolution. March 9 [www.arxiv.org/pdf/1003.1802]

Durban, M., Currie, I. & Eilers, P. (2002). "Using P-splines to smooth two-dimensional Poisson data". Proceedings of 18th International Workshop on Statistical Modelling, Leuven, Belgium, 97-102. [www.ma.hw.ac.uk/~iain/research/Leuven.pdf]

Lee, R. D., and L. R. Carter (1992) "Modeling and forecasting U.S. mortality." *Journal of the American Statistical Association*, 87: 659-675.

Renshaw, A.E., and Haberman, S. (2006). "A cohort-based extension to the Lee-Carter model for mortality reduction factors." *Insurance: Mathematics and Economics*, 38: 556-570.

Society of Actuaries (2000). "The RP-2000 Mortality Tables". July [www.soa.org]

A decorative graphic consisting of numerous thin, parallel yellow lines that curve and flow across the right side of the page, set against a dark blue background.

[www.ILMA.org](http://www.ILMA.org)