Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness

Coughlan, Guy and Khalaf-Allah, Marwa and Ye, Yijing and Kumar, Sumit and Cairns, Andrew and Blake, David and Dowd, Kevin

2011

Pensions Institute

DISCUSSION PAPER PI-1013

Longevity Hedging 101: A Framework For Longevity Basis Risk Analysis And Hedge Effectiveness

Guy D. Coughlan, Marwa Khalaf-Allah, Yijing Ye, Sumit Kumar, Andrew J. G. Cairns, David Blake, and Kevin Dowd

Published in the North American Actuarial Journal, Volume 15, Number 2, 2011

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

LONGEVITY HEDGING 101: A FRAMEWORK FOR LONGEVITY BASIS RISK ANALYSIS AND HEDGE EFFECTIVENESS

Guy D. Coughlan,* Marwa Khalaf-Allah,† Yijing Ye,‡ Sumit Kumar,§ Andrew J. G. Cairns,‖ David Blake,# and Kevin Dowd¶

ABSTRACT

Basis risk is an important consideration when hedging longevity risk with instruments based on longevity indices, since the longevity experience of the hedged exposure may differ from that of the index. As a result, any decision to execute an index-based hedge requires a framework for (1) developing an informed understanding of the basis risk, (2) appropriately calibrating the hedging instrument, and (3) evaluating hedge effectiveness. We describe such a framework and apply it to a U.K. case study, which compares the population of assured lives from the Continuous Mortality Investigation with the England and Wales national population. The framework is founded on an analysis of historical experience data, together with an appreciation of the contextual relationship between the two related populations in social, economic, and demographic terms. Despite the different demographic profiles, the case study provides evidence of stable long-term relationships between the mortality experiences of the two populations. This suggests the important result that high levels of hedge effectiveness should be achievable with appropriately calibrated, static, index-based longevity hedges. Indeed, this is borne out in detailed calculations of hedge effectiveness for a hypothetical pension portfolio where the basis risk is based on the case study. A robustness check involving populations from the United States yields similar results.

1. INTRODUCTION

Longevity risk—the risk that life spans exceed expectation—is a significant concern for defined benefit pension plans and life insurers with large annuity portfolios. Until recently, the only way to mitigate longevity risk was via an insurance solution: Pension plans bought annuities from, or sold their liabilities to, insurers, and insurers bought reinsurance. Then 2008 saw the first capital markets solutions for longevity risk management executed by Lucida plc and Canada Life in the United Kingdom.[1] Both these

* Guy D. Coughlan, PhD, is Managing Director, Pension Advisory Group, J.P. Morgan, 125 London Wall, London EC2Y 5AJ, United Kingdom, guy.coughlan@jpmorgan.com.

† Marwa Khalaf-Allah, PhD, is Vice President, Pension Advisory Group, J.P. Morgan, 125 London Wall, London EC2Y 5AJ, United Kingdom, marwa.khalafallah@jpmorgan.com.

‡ Yijing Ye is Associate, Pension Advisory Group, J.P. Morgan, 125 London Wall, London EC2Y 5AJ, United Kingdom, yijing.x.ye@jpmorgan.com.

§ Sumit Kumar is Associate, Pension Advisory Group, J.P. Morgan, 125 London Wall, London EC2Y 5AJ, United Kingdom, sumit.z.kumar@jpmorgan.com.

‖ Andrew J. G. Cairns is Professor of Financial Mathematics, Maxwell Institute for Mathematical Sciences, and Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom, a.j.g.cairns@ma.hw.ac.uk.

David Blake is Professor of Pension Economics and Director of the Pensions Institute, Pensions Institute, Cass Business School, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom, D.Blake@city.ac.uk.

¶ Kevin Dowd is Visiting Professor, Pensions Institute, Cass Business School, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom, kevin.dowd@hotmail.co.uk.

[1] In January 2008, Lucida, a U.K. pension buyout insurer, executed a standardized index-based longevity hedge designed to hedge liability value (Lucida 2008; Symmons 2008). Then in July 2008, Canada Life executed a customized hedge of its U.K. liability cash flows (Trading Risk 2008; Life and Pensions 2008). See also Coughlan (2009b) for more details.

transactions were significant catalysts for the development of the longevity risk transfer market, bringing additional capacity, flexibility, and transparency to complement existing insurance solutions.

While customized (i.e., indemnity-based) capital markets longevity transactions have received more publicity following Canada Life's pioneering longevity swap, it is significant that the very first capital markets transaction to transfer the longevity risk associated with pension payments was a standardized, index-based hedge. In this transaction, Lucida plc executed a mortality forward rate contract called a "q-forward,"[2] the payoff of which was linked to the LifeMetrics longevity index[3] for England and Wales. Hedging longevity risk with index-based hedging instruments can be beneficial for several reasons. First, by standardizing the longevity exposure to reflect an index, there is the potential to create greater liquidity and lower the cost of hedging. Second, some pension plans are just too large to hedge the full extent of their exposure to longevity risk in other ways. Third, it is currently the most practical solution for hedging the longevity risk associated with deferred pensions and deferred annuities (Coughlan 2009a).[4] In fact, as this paper was going to press, the trustees of the Pall (U.K.) Pension Fund announced just such an index hedge of their deferred pensioner longevity risk (Davies 2011; Mercer 2011; Stapleton 2011). This transaction, executed with J.P. Morgan in January 2011, was also based on the LifeMetrics index and was calibrated according to the framework described in this paper. The hedge consists of a portfolio of q-forwards linked to male and female mortality rates in 10-year age buckets.

Against the benefits of index-based longevity hedges, one must weigh the disadvantages, the primary one being basis risk. Because the mortality experience of the index will differ from that of the pension plan or annuity portfolio, the hedge will be imperfect, leaving a residual amount of risk, known as basis risk. When contemplating whether and how to hedge, it is clearly essential to evaluate the size of this risk and weigh the degree of risk reduction against the cost of the hedge.

Unfortunately, until now, there has been little work published by academics or practitioners on longevity basis risk and its impact on the effectiveness of longevity hedges. Nor has hedge effectiveness as it relates to longevity risk been well understood by practitioners and consultants in the pension and insurance industries. Lacking a proper framework, it has been common practice in some quarters to "assess" hedge effectiveness by making qualitative value judgements based on the differences in observed mortality rates. This has led to a widely held misconception that index-based longevity hedges are ineffective.

This paper addresses these issues, first, by proposing a framework for assessing longevity basis risk and hedge effectiveness, and, second, by presenting a practical example—based on U.K. data—that illustrates this framework and demonstrates that index-based hedges can indeed be highly effective. As a robustness check, we have also evaluated a second example that gives similar results based on U.S. data, but in this paper we present only the results, not the details.

The framework we propose sets out the key principles and steps involved in a structured approach to determining the effectiveness of longevity hedges. The key initial step is a careful analysis of the basis risk between the population associated with the pension plan or annuity portfolio (the "exposed population") and the population associated with the hedging instrument (the "hedging population"). This is illustrated by the two examples mentioned above. Each example compares the experience of the national population with a particular affluent subpopulation, which, historically, has enjoyed, on average, lower rates of mortality and higher mortality improvements over time than the national pop-

[2] The q-forward instrument is described in detail in Coughlan et al. (2007c).

[3] The LifeMetrics longevity index is documented in Coughlan et al. (2007a, 2007b, 2008). See also www.lifemetrics.com.

[4] For deferred longevity risk, index-based hedges are a viable solution, whereas customized hedges are generally either not available or else very costly. Furthermore, the risk prior to retirement is a pure valuation risk, since no payments are made before retirement. This valuation risk is driven by uncertainty in mortality improvements, which are generally calibrated from improvement forecasts for large populations (e.g., national population indices). In addition, the underlying longevity exposure of deferred members of pension plans is not well defined owing to various member options, such as early retirement, lump sums, or spouse transfers, so exact, customized hedging is neither practicable nor desirable.

ulation. Using this basis risk analysis, we then conduct an evaluation of hedge effectiveness for a hypothetical pension plan with the same mortality characteristics as the affluent subpopulation for a static (i.e., not dynamically rebalanced) hedge, based on a longevity index linked to the national population.

This framework provides a practical approach that hedgers can use to calibrate index-based hedges, develop an informed understanding of basis risk, and evaluate hedge effectiveness. The framework does *not* assume any particular model for the future evolution of mortality rates and can be applied effectively with any modeling approach that the user might choose. Nor does it assume any particular model for the valuation (or pricing) of the longevity exposure or hedging instrument. In this sense, the framework can be considered to be model independent. However, the application of the framework requires modeling choices to be made. In the two examples presented in this paper, we have chosen a nonparametric approach to (or model for) basis risk that reflects the historical relationships observed in the data between the two populations concerned. We have also chosen a simple approach to the valuation of the pension liability and hedging instrument which is described later.

The examples demonstrate that, despite the different demographic profiles of the population pairs, there is evidence of stable long-term relationships between their mortality experiences. This has favorable implications for the effectiveness of appropriately calibrated, index-based longevity hedges. From this, we conclude that longevity basis risk between a pension plan, or annuity book, and a hedging instrument linked to a broad population-based longevity index can, in principle, be reduced very considerably.

Although our examples involve static longevity hedges, the framework could be extended to dynamic hedges with an appropriate hedge rebalancing criterion. In practice, dynamic hedging is currently untenable because of high transaction costs and lack of liquidity.

The paper is organized as follows. In the next section, we describe the relationship between longevity basis risk and hedge effectiveness, and review the existing literature on the subject. Section 3 then presents the framework for analyzing basis risk and hedge effectiveness. In Section 4, we apply the framework to the case study mentioned above, and Section 5 briefly considers a robustness check that we performed using U.S. data. Finally, Section 6 is devoted to conclusions.

2. BASIS RISK AND HEDGE EFFECTIVENESS

2.1 What Is Basis Risk?

Basis risk arises whenever there are differences, or mismatches, between the underlying hedged item and the hedging instrument. These differences can take many forms, ranging from differences in the timing of cash flows to differences in the underlying variables that determine the cash flows. The presence of basis risk means that hedge effectiveness will not be perfect and that, after implementation, the hedged position will still have some residual risk.

It is important to note that basis risk is present to some degree in most financial hedges, and it does not automatically invalidate the case for hedging. For example, the interest rate and inflation hedges used by pension plans and insurance companies almost always have some basis risk. Contrary to common practice, basis risk should always be quantified because, in many cases, it can be minimized through careful structuring and calibration of the hedging instrument to ensure high hedge effectiveness. A ‘‘good’’ hedge is therefore one in which the basis risk is small, relative to the risk of the initial unhedged position.

2.2 Longevity Basis Risk

In the context of longevity and mortality, basis risk often relates to mismatches in demographics between the ‘‘exposed population’’ (e.g., the population of members of a pension plan or the beneficiaries of an annuity portfolio) and the ‘‘hedging population’’ associated with the hedging instrument (i.e., the population that determines the payoff on the hedge).[5] These demographic mismatches can arise because the two populations are completely different, or because one population is a subpopulation of the other, or because just a few individuals are different. Regardless of how they arise, however, such mismatches can be classified according to a small number of demographic characteristics (Richards and Jones 2004), such as gender, age, socioeconomic class, or geographical location.

If two populations have similar profiles for these characteristics, then one would generally expect basis risk to be small. On the other hand, if the two populations have vastly different profiles, then basis risk could be large, but this is not necessarily the case, as we shall see below.

Examples of basis risk between populations include that originating from the mismatch in mortality rates between males and females (the ''gender basis''), the mismatch between mortality at different ages (the ''age basis''), the mismatch between national mortality and the mortality of a particular subpopulation (the ''subpopulation basis''), and the mismatch between mortality in different countries (the ''country basis'').

The basis risk associated with gender or age is generally not an issue for index-based hedges, because it can be minimized by appropriate structuring of the hedging instrument. Indeed, most broad-based population longevity indices are broken down into subindices by gender and age, thereby permitting the hedging instrument to be matched to the gender and age profile of the underlying pension plan or annuity portfolio through appropriate combinations of subindices. As a result, the most important determinant of basis risk in most index-based hedging situations is that associated with socioeconomic class.[6]

For the purposes of this paper, we distinguish between basis risk and sampling risk, the latter being the risk associated with small populations.[7] Our goal is to explore the elements of basis risk that cannot be explained by the size of the population, although in our examples, variability due to sampling risk will be a factor for higher ages, since the population at higher ages is relatively small and the numbers of deaths—and hence mortality rates—will be highly variable from one year to the next. Note that sampling risk could be taken into account by combining it with the population basis risk and including it in the simulation of mortality rates for the two populations.

2.3 Hedge Effectiveness

Although hedge effectiveness is an intuitive concept, it is not yet widely understood or applied in the context of longevity hedging.[8] Certainly, the presence of longevity basis risk reduces the effectiveness of longevity hedges, but the relationship between the two is not as straightforward as one might suppose. Whereas basis risk is typically measured in demographic terms, hedge effectiveness should be measured in economic terms, and demographic mismatches do not necessarily result in significant economic costs, as we shall illustrate.

The key to designing an appropriate method for assessing hedge effectiveness is to start from hedging objectives that reflect the nature of the risk being hedged and focus on the degree of risk reduction in economic terms. One way to do this is by using a Monte Carlo simulation to generate forward-looking scenarios for the evolution of mortality rates for both the exposed and hedging populations. Ideally, a two-population stochastic mortality model, calibrated from a suitable basis risk analysis, should be used to do this in a consistent fashion. However, we wish, in this paper, to avoid the additional

[5] In the context of longevity hedging, basis risk can also arise from the structure of the hedging instrument independently of any demographics, e.g., a mismatch in maturity between the underlying exposure and the hedging instrument. In this paper we ignore this kind of basis risk.

[6] Variations in mortality associated with different regions within a country can largely be explained by socioeconomic or lifestyle differences. See, for example, Richards and Jones (2004).

[7] Even if we had a sample from a larger population without basis risk, because it was free from gender, age, and subpopulation biases, the small population would still experience sampling—random variation—risk relative to the large population, which would cause its mortality experience to diverge from that of the large population over time.

[8] It is worth drawing the parallel that when the requirement to assess hedge effectiveness was first introduced in an accounting context under U.S. GAAP (SFAS 133) and IFRS (IAS 39), the subject was at that time poorly understood by both practitioners and accountants.

complexity of dealing with a formal model in order to concentrate on producing a straightforward exposition of longevity hedge effectiveness. Accordingly, the specific model we use involves a nonparametric approach, based largely on historical data on mortality rates and mortality rate improvements.

It is important to recognise that different approaches to hedge effectiveness will give different results. The most appropriate approach is that which is most closely aligned with the hedging objectives. In practice, this requires judgement and experience.

2.4 Existing Literature

Several authors have explored the basis risk between populations associated with annuity portfolios and life insurance portfolios. Cox and Lin (2007) found empirical evidence of a (partial) natural hedge operating between such portfolios, implying that the basis risk between them is relatively small.[9] Coughlan et al. (2007b, pp. 85–87) provided a calculation of the risk reduction between hypothetical annuity and life insurance portfolios using historical mortality experience data: The results suggest significant benefits in terms of reduction in risk and economic capital. Sweeting (2007) explored the basis risk associated with longevity swaps in a more qualitative fashion but draws similar conclusions.

Recently a number of researchers have developed mortality models for two or more related populations (Li and Lee 2005; Jarner 2008; Jarner and Kryger 2009; Plat 2009; Cairns et al. 2011; Dowd et al. 2011; Li and Hardy 2011). These models are all based on the principle that, on the grounds of biological reasonableness,[10] the mortality rates of related populations should not diverge over the long term. Whereas all these papers are motivated by a desire to develop coherent and consistent forecasts, only the last four express an explicit goal related to the measurement of basis risk and hedge effectiveness. Ideally, a coherent two-population model is needed to evaluate basis risk prospectively, but such a model also needs to be both intuitively appealing and appropriately calibrated, and these, in turn, depend on a sound understanding of historical mortality experience.

Hedge effectiveness testing in a general context has been addressed by Coughlan et al. (2004), and its application to longevity hedging has been briefly discussed in Coughlan et al. (2007b) and Coughlan (2009b). More recently, Ngai and Sherris (2010) have evaluated some hedge effectiveness metrics for hedges of various annuity products in the Australian market.

3. A FRAMEWORK FOR ANALYZING BASIS RISK AND HEDGE EFFECTIVENESS

Any decision to execute an index-based longevity hedge requires a framework for (1) developing a deep understanding of the basis risk involved, (2) calibrating the hedging instrument, and (3) evaluating hedge effectiveness.

In most situations involving real pension plans and annuity portfolios, the amount of historical data available will be too short to draw rigorous statistical conclusions about basis risk. Nevertheless, by examining available data carefully and trying to identify key demographic—especially socioeconomic—characteristics, one can usually develop an informed assessment of the nature and magnitude of this risk. As fits with risk management best practice in other areas, hedging decisions are ultimately based on professional judgment-supported data analysis and experience.

3.1 Basis Risk Analysis

Basis risk analysis should be appropriately aligned with the hedging objectives in terms of the metric, the time horizon, and the analytical method.

[9] Similar results have been found by Dahl and Møller (2006), Friedberg and Webb (2007), and Wang et al. (2010).

[10] A method of reasoning used to establish a causal association (or relationship) between two factors that is consistent with existing medical knowledge.

3.1.1 Metric

Many different metrics can be used to gain a perspective on the basis risk associated with longevity hedges. Because of the complex relationships between mortality experience across age, time (or period), and year of birth (or cohort), it is necessary to examine the historical performance of all key metrics:

- Mortality rates (either crude rates or graduated rates)
- Mortality improvements (i.e., percentage changes in mortality rates)
- Survival rates
- Life expectancies
- Liability cash flows
- Liability values.

Since mortality rates constitute the basic raw data associated with longevity, they have been the most commonly used metric for assessing basis risk. Unfortunately, a direct comparison between the mortality rates of two populations provides a naive and often misleading perspective on basis risk and the effectiveness of longevity hedges, for several reasons. First, mortality rates as metrics are not directly related to the effectiveness of longevity hedges. Therefore, one needs to be careful in drawing any conclusions regarding the impact of basis risk, as observed in mortality rate comparisons, on longevity hedge effectiveness.

Second, the data corresponding to annual mortality rates at particular ages contain a lot of noise or sampling variability. This noise is present both though time and across ages and is evidenced in the observed year-on-year fluctuations of mortality rates around their long-term trends.[11] The noise can be reduced by (1) graduating mortality rates across ages using a smoothing routine, (2) bucketing adjacent ages together when calculating mortality rates, and (3) evaluating the changes in mortality rates over the longer time horizons that are more typical of the timescales associated with longevity trends emerging. So, in using mortality-rate comparisons to evaluate the basis risk of longevity hedges, it is important to incorporate these three elements—graduation, age bucketing, and longer horizons—into the analysis.

Survival rates and life expectancy both address the above shortcomings as metrics for basis risk analysis. Because survival rates in a pensioner population correspond to the number of members who are still alive to receive a pension and life expectancy corresponds to the expected period over which a pension needs to be paid, these metrics are more closely related to the hedge effectiveness objective than mortality rates. Moreover, both survival rates and life expectancy are calculated from many different mortality rates for different ages and at different times, so there is natural smoothing out of the noise that is associated with individual mortality rates in individual years. For example, the 10-year survival rate for 65-year-old males depends on the mortality rates for males aged 65 through 74. Similarly, life expectancy for 65-year-olds depends on the mortality rates for every age above 65.

Although useful in developing an understanding of basis risk, none of the above metrics is ideal for quantifying hedge effectiveness. Since most hedging exercises are focused on mitigating the variability in the liability cash flows or the variability in the value of these cash flows, basis risk studies should ideally focus on the impact on cash flow and/or value. These metrics directly reflect the monetary impact of basis risk and are the appropriate metrics for evaluating the effectiveness of longevity hedges.

Although cash flow and value are more useful for quantifying basis risk than the other metrics discussed above, they suffer from one main disadvantage in that they depend on the specific details of the benefit structure of the particular pension plan or annuity portfolio and, as such, involve complex calculations, including discounting of future cash flows, that must be repeated in full for each situation.

---

[11] Indeed, mortality rates typically exhibit a negative autocorrelation, so that a high mortality rate in one year is followed by a low mortality rate the next year (Coughlan et al. 2007b).

By contrast, mortality rates, survival rates, and life expectancy are independent of the details of the specific benefit structure and can, with appropriate interpretation, give useful insights into basis risk.

3.1.2 Time Horizon

The choice of time horizon is important in assessing basis risk. Longevity risk, as it applies to large populations, is a slowly building, cumulative trend risk that should be evaluated over long time horizons. To be consistent with this, metrics should be evaluated over horizons of at least several years. For example, in comparing the evolution of mortality rates for two populations, it is desirable to evaluate changes in their mortality rates over multiyear horizons, rather than year on year.

Unfortunately, using long horizons means that there are fewer independent observations available from a given historical data set. So selecting the time horizon for analyzing basis risk involves making a trade-off between a horizon long enough to identify trends and short enough to provide enough independent data points to give a robust analysis.

3.1.3 Analytical Method

The analytical method for evaluating basis risk should, as for the metric and time horizon, also be appropriately aligned with the hedging objective. This means deciding on various details of the analysis, such as whether to compare the levels of a particular metric or changes in the metric for each population. If comparing changes, we need to specify how these changes should be defined. For example: What time period is optimal? Should we use overlapping periods or nonoverlapping periods? Should we use one-period changes or cumulative multiperiod changes?

Other methodological choices include, for example, whether and how to bucket age groups and whether and how to graduate mortality rates.[12] The use of both age-group bucketing and graduation are generally desirable to reduce noise and can be justified because mortality rates for adjacent ages are similar and highly correlated.[13] As a result, bucketing and graduation do not destroy the integrity of the data; rather, they bring the twin benefits of simplification and noise reduction, thereby rendering a clearer perspective on basis risk and hedge effectiveness. In practice, mortality curves—as represented in mortality tables—are graduated as a key part of the valuation process for liabilities whenever pension or annuity portfolios are transferred between different counterparties. Furthermore, hedges constructed using bucketed age groups have already been transacted in the capital markets, and their effectiveness in reducing risk has not been compromised by the age bucketing. In particular, the various mortality catastrophe bonds and swaps issued by Swiss Re, Scottish Re, Axa, Munich Re, and others, as well as the longevity index hedges completed by Lucida and Pall (U.K.), have all used bucketed age groups.

Once these aspects of the analysis have been decided, the next decision is how to compare the results across the two populations. This can be done qualitatively in a graphical format, or quantitatively using statistical analyses, such as correlation.

Correlation is a common way of evaluating basis risk in a general setting, but care should be taken in the context of longevity. Correlation in the annual improvements in mortality rates between two populations reflects the short-term relationship between these populations and, by virtue of the noise inherent in mortality data, can give a very misleading indication of the strength of the relationship between their long-term trends. By contrast, correlations between long-term mortality improvements are more relevant indicators, but, as mentioned, there will be far fewer independent data points for long-term improvements.

The essence of longevity basis risk analysis is the search for a stable long-term relationship between the two populations. If such a relationship can be identified, then an appropriate index-based longevity hedge can be calibrated by determining the optimal hedge ratios for the hedging instrument.

[12] The graduation and bucketing methods used in this paper are explained in the Appendix.

[13] The "age basis" is typically small for adjacent ages in a large population (Coughlan et al. 2007b).

3.2 Hedge Calibration

Hedge calibration refers to the process of designing the hedging instrument to maximize its effectiveness in reducing risk, relative to the hedging objectives. It involves two elements. The first is the determination of the appropriate structure and characteristics of the hedging instrument (e.g., type of instrument, maturity, or index to be used). The second is the determination of the optimal amount of the hedge required to maximize hedge effectiveness. This involves determining optimal "hedge ratios" for each of the subcomponents of the hedging instrument.

As a simple example, consider a hedging instrument with just one component designed to hedge the value of a pension liability at a future time, which we call the "hedge horizon." Suppose we have bought h units of the hedge for each unit of the liability: h is the hedge ratio. Then the total (net) value of the combined exposure is

$$V_{\text{Total}} = V_{\text{Liability}} + h \times V_{\text{Hedge}}. \tag{1}$$

The optimization element referred to above involves selecting h to maximize hedge effectiveness by minimizing the uncertainty in V_{Total}. It can be shown that, assuming the values are normally distributed and risk is measured by standard deviation, then the optimal hedge ratio is given by (Coughlan et al. 2004)

$$h_{\text{Optimal}} = -\rho \times (\sigma_{\text{Liability}}/\sigma_{\text{Hedge}}), \tag{2}$$

where $\sigma_{\text{Liability}}$ and σ_{Hedge} are the standard deviations of the values of the liability and hedging instrument, respectively, at the hedge horizon, and ρ is the correlation between them.

It is evident from this simple example that basis risk analysis is an essential prerequisite for optimal hedge calibration.

3.3 Hedge Effectiveness Methodology

Assessing hedge effectiveness requires taking account of the hedging objectives and the nature of the risk that is being hedged to develop a methodology that is appropriate. Table 1 summarizes the key steps involved in this process (derived from Coughlan et al. 2004). A key part of the methodological choice is whether hedge effectiveness is to be assessed *retrospectively* or *prospectively*.

Retrospective hedge effectiveness analysis involves using actual historical data to assess how well a hedging instrument would have performed in the past. In this kind of effectiveness test, basis risk is

Table 1
Framework for Assessing Hedge Effectiveness

Step 1	Define hedging objectives Metric Hedge horizon Risk to be hedged (full or partial)
Step 2	Select hedging instrument Structure hedge Calibrate hedge ratio
Step 3	Select method for hedge effectiveness assessment Retrospective vs. prospective test Basis for comparison (comparing hedged and unhedged performance, valuation model, etc.) Risk metric Simulation model to be used
Step 4	Calculate the effectiveness of hedge Simulation of mortality rates for both populations Evaluation of effectiveness based on the simulations
Step 5	Interpret the effectiveness results

taken account of by virtue of the historical relationships between the observed mortality outcomes for both the hedging and exposed populations.

By contrast, prospective hedge effectiveness analysis involves developing forward-looking scenarios to anticipate how well a hedging instrument might perform in the future. This involves a Monte Carlo simulation of potential future paths for mortality rates from which the performance of the hedging instrument can be assessed relative to the underlying longevity exposure. In this case, basis risk must be explicitly taken into account, with the simulation of scenarios for future mortality rates reflecting, in a consistent way, the observed relationship between the hedging and exposed populations. As mentioned above, this ideally requires a two-population stochastic mortality model to be used.

Let us discuss the hedge effectiveness framework presented in Table 1 in greater detail. Step 1 involves defining hedging objectives, in particular, designating the precise risk being hedged. This includes the risk class (i.e., longevity risk), as well as the precise nature of what is being hedged (e.g., longevity trend improvements above 2% per year over the next 10 years, or the total uncertainty in survivorship over 40 years). An essential part of this is defining the hedge horizon of the hedging relationship as well the performance metric (e.g., hedging liability cash flows or liability value). The choice of hedge horizon will be driven by the performance metric, the precise nature of the exposure being hedged (e.g., longevity risk up to the point of retirement or longevity exposure from retirement), and factors associated with the hedging instrument, such as cost and liquidity.

Step 2 in the process is to select the hedging instrument and calibrate the optimal hedge ratio. The latter should be chosen to maximize the degree of risk reduction and should be determined from an appropriate basis risk analysis, as suggested by the simple example in the previous section (see, for example, eq. (2)).

Step 3 in the process—which defines the hedge effectiveness methodology—is important, because an inappropriate choice can lead to spurious and misleading results with effective hedges being deemed ineffective, or vice versa. Defining the methodology involves several choices. The first choice involves selecting between a retrospective and prospective effectiveness test, which we have discussed above. The second choice to be made is the "basis for comparison," which involves specifying how the performance of the unhedged and hedged exposures are to be compared. A simple choice is in terms of the degree of risk reduction:

$$\text{Relative Risk Reduction} = RRR = 1 - \text{Risk}_{(\text{Liability + Hedge})}/\text{Risk}_{\text{Liability}}. \tag{3}$$

Clearly, a perfect hedge reduces the risk to zero, corresponding to 100% risk reduction.

If the hedging objectives are framed in terms of hedging the liability value (rather than the liability cash flows), another key element of the basis-for-comparison choice is the pricing (i.e., valuation) model used to determine the values of the liability and the hedging instrument under different scenarios at the hedge horizon.

The next choice to be made is the selection of the risk metric. If the hedging objectives are couched in terms of hedging liability value, then an example of an appropriate risk metric might be the value-at-risk (VaR) of the liabilities at the hedge horizon relative to an expected, or a median, outcome and calculated at a particular confidence level.

The final methodological choice relates to selecting the type of simulation model used to generate the scenarios needed for the test. Note that this framework can be used with any simulation model, including fully stochastic two-population mortality models, models with parametric trends, and non-parametric approaches that use historical mortality data directly.

Step 4 addresses the actual calculation of hedge effectiveness. This involves an implementation of the method defined in the previous step as a two-stage process: (1) simulation and (2) evaluation. Note that the simulation of mortality risk is a separate process from the evaluation of the impact of the scenarios on the liability and the hedging instrument, as illustrated in Figure 1. In particular, the evaluation process is the same for any set of mortality scenarios, regardless of how the set of scenarios is generated. The evaluation process involves calculating the cash flows and/or value of both the pension liability and the hedging instrument in each scenario. This requires a choice of valuation, or pricing,

Figure 1

Process of Hedge Effectiveness Assessment (Step 4 in the Hedge Effectiveness Framework) Involving Two Distinct Parts: Simulation and Evaluation, with Basis Risk Analysis a Key Input into the Former

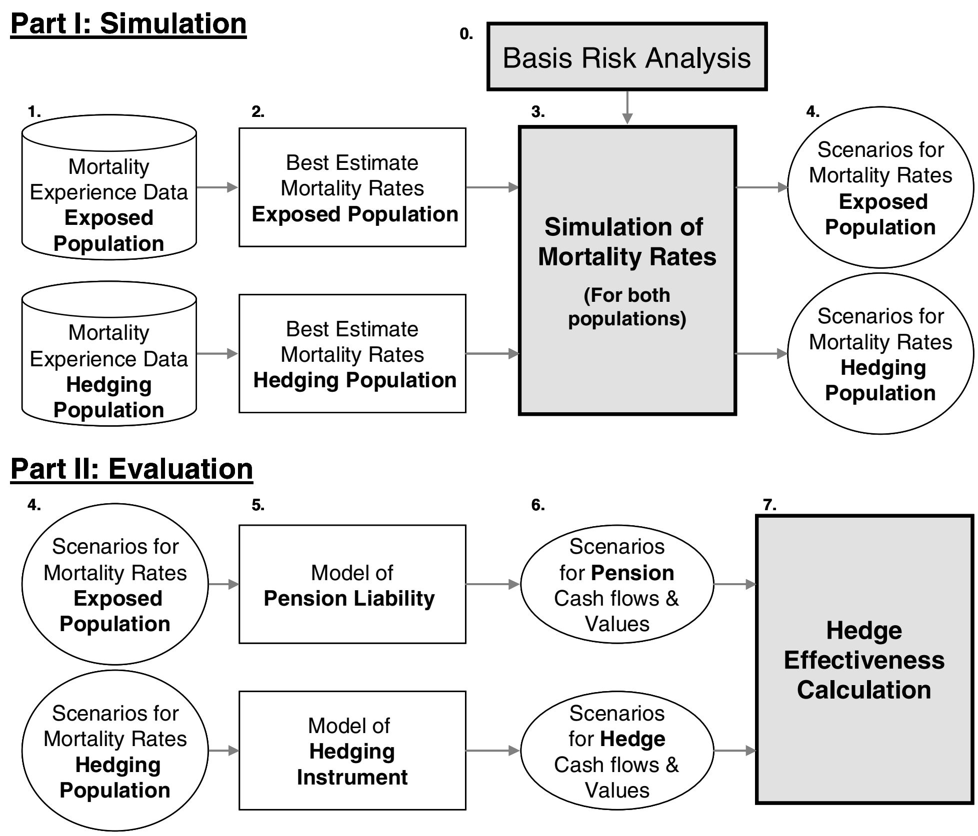

model to be made. Figure 1 also shows the important role of basis risk analysis in the calculation of hedge effectiveness.

Step 5 in the process—the final step in the framework—involves interpreting the hedge effectiveness results.

4. U.K. BASIS RISK CASE STUDY

In this section we present the results of an empirical analysis of the basis risk between the national population of England and Wales males (based on data from the Office for National Statistics [ONS]) and the population of U.K. males who own life assurance policies (based on data from the Continuous Mortality Investigation [CMI]). The ONS is the U.K. government agency that compiles official national mortality statistics. The CMI is a body, funded by the U.K. life insurance industry and run by the U.K. Actuarial Profession, which publishes mortality rates for assured lives, derived from data submitted by U.K. life insurers. The data used in this analysis cover the 45-year period 1961–2005.

The CMI data come from an affluent subset of the U.K. population, whose mortality rates have consistently been lower and mortality improvements higher than those of the national population.[14] It

[14] The CMI data are a subset of the U.K. population, so are not strictly a subset of the population of England and Wales. However, in the context of this analysis the difference is small.

is important to note that the population of assured lives behind the CMI data is a subset of the national population that changes from year to year, depending on which insurers choose to submit their data: The CMI data are therefore not only a subset of the national population, but also a (changing) subset of the population of assured lives. Furthermore, the number of lives in the CMI assured lives data set has fallen significantly over the past 20 years. Currie (2009) has a chart showing how the exposure by age has changed through time, from a peak of over 200,000 lives at an age of around 40 in 1985 to a peak of fewer than 50,000 lives for males in their late 50s in 2005. For higher ages, the exposures are even lower. Such changes have inevitably introduced additional noise into the CMI data and are likely to lead to a higher measured basis risk than genuinely exists between the two populations. As a consequence, the results of estimating basis risk that we present in this paper are likely to be conservative.

In this analysis, we used graduated initial mortality rates, whose calculation is described in the Appendix. Moreover, we consider only data up to age 89, because beyond this age the ONS publishes only aggregated mortality statistics, and the results may be affected by the modeling choice for higher ages used in the graduation of mortality rates. We begin by looking at mortality rates, before moving on to examine other metrics.

4.1 Mortality Rates and Mortality Improvements

Figure 2 shows a graphical comparison of graduated mortality rates for the assured population and the national population. The most obvious feature of the data, which is common to all ages over 35, is the significant difference in the level of mortality rates for the two populations: Assured mortality is much lower than national mortality. What is also evident is that the long-term downward trends are quite similar, suggesting that there might be a long-term relationship between the mortality rates of the two populations. Certainly from year to year, there is volatility around each trend, and the assured data set appears to be noisier (particularly at higher ages), but broadly speaking the observed improvements in mortality are moving together and not diverging.

Figure 2

Comparison of Male Mortality Rates for the U.K. Assured and England and Wales National Populations: (a) Spot Mortality Curves for 2005, (b) Historical Evolution of Graduated Mortality Rates for 65-Year-Old Males, 1961–2005

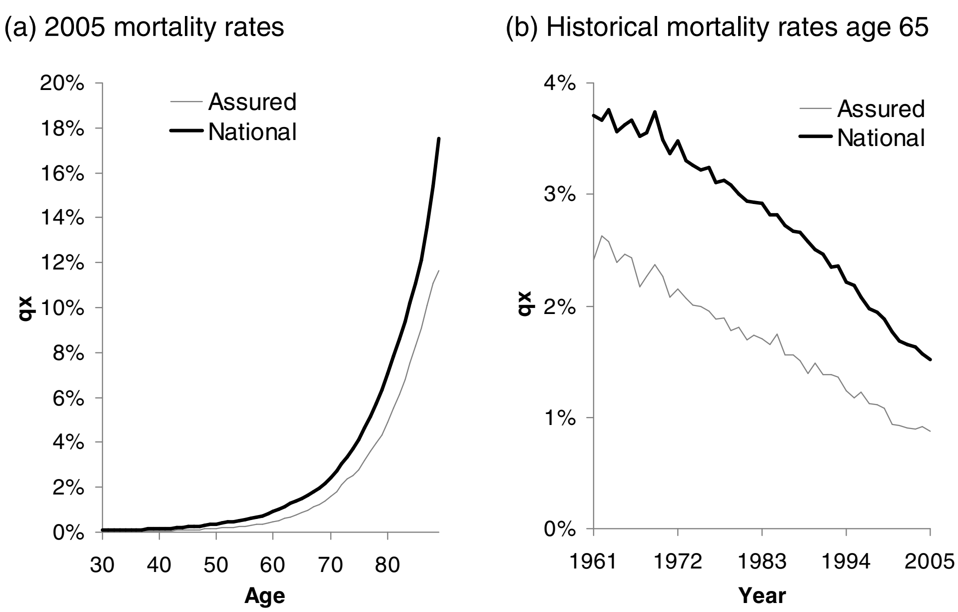

Table 2

U.K. Assured Male Mortality Rates as a Percentage of England and Wales National Male Mortality Rates, Averaged over Age

Ratio of Mortality Rates (Assured/National)	1961	2005
Overall: 40–89	68%	57%
Younger: 40–64	62%	46%
Older: 65–89	74%	68%

Comparing the average levels of mortality rates (Table 2) for the two populations, we see that assured mortality in 2005 was on average 57% of national mortality, having fallen from 68% in 1961. So there has been a pronounced decrease in relative assured mortality rates since 1961. Moreover, the relative rates vary by age, with assured mortality for the younger, pre-retirement ages of 40–64 currently averaging just 46% of national mortality and for older post-retirement ages of 65–89 averaging 68%.

Over the period 1961–2005, observed mortality improvements (Table 3) have averaged 2.04% p.a. for the assured population, compared with 1.62% p.a. for the national population. Furthermore, the younger pre-retirement ages have experienced much higher improvements of 2.32% p.a. and 1.67% p.a., respectively, for assured males and national males, while for the older post-retirement ages improvements have been lower, at 1.75% and 1.57%, respectively.

As we have already mentioned, the differences in both the levels of and the improvements in mortality rates between the two populations do *not* necessarily mean that the effectiveness of longevity hedges will be poor. Indeed, there appears to be a relatively stable long-term relationship between them that can be exploited to construct hedges that are highly effective. That this is the case is evidenced by evaluating aggregate correlations in the observed changes in mortality rates for the two populations (Tables 4 and 5).

Table 4 lists the ''aggregate correlations'' for changes in mortality rates over different horizons calculated from individual ages. The calculation of these correlations is described in Appendix A.3 and involves evaluating the correlation for changes in mortality rates over nonoverlapping periods jointly for each individual age. The correlations are calculated for both absolute changes in mortality rates and relative, percentage changes, that is, mortality improvements. Note that the aggregate correlations in year-on-year changes based on individual ages are quite small, just 36%, but they increase with the length of the time horizon. Correlations are around 97% or more for a 20-year horizon and around 80% for a 10-year horizon.

Using age buckets helps remove noise from the mortality data and leads to much higher correlations as shown in Table 5. With 10-year age buckets (50–59, 60–69, 70–79, and 80–89), the aggregate correlation for year-on-year changes is 54%, rising to 94% for a 10-year horizon and to 99% for a 20-year horizon.

Table 3

Annualized Male Mortality Improvements for the England and Wales National and U.K. Assured Populations, Averaged over Age Groups, 1961–2005

Mortality Improvements 1961–2005 (Annualized)	National (% p.a.)	Assured (% p.a.)	Difference (Percentage Points)
Overall: 40–89	1.62	2.04	0.42
Younger: 40–64	1.67	2.32	0.65
Older: 65–89	1.57	1.75	0.18

Table 4

Aggregate Correlations of Changes in Male Mortality Rates for Individual Ages between the U.K. Assured and England and Wales National Populations, 1961–2005

Individual Ages	Correlation between Absolute Changes in Mortality Rates: Individual Ages: 40–89	Correlation between Absolute Changes in Mortality Rates: Individual Ages: 50–89	Correlation between Improvement Rates (Relative Changes): Individual Ages: 40–89	Correlation between Improvement Rates (Relative Changes): Individual Ages: 50–89
20-year horizon	97%	97%	93%	96%
10-year horizon	80%	77%	81%	85%
5-year horizon	69%	66%	65%	70%
1-year horizon	36%	36%	29%	36%

Note: Correlations are calculated across time (using nonoverlapping periods) and across individual ages (without any age bucketing), using graduated mortality rates. See the Appendix.

Table 5

Aggregate Correlations of Changes in Male Mortality Rates for 10-Year Age Buckets between the U.K. Assured and England and Wales National Populations, 1961–2005

Age Buckets: 50–59, 60–69, 70–79, 80–89	Correlation between Absolute Changes in Mortality Rates	Correlation between Improvement Rates (Relative Changes)
20-year horizon	99%	98%
10-year horizon	94%	93%
5-year horizon	91%	85%
1-year horizon	54%	51%

Note: See note to Table 4.

It should be noted that for long horizons, however, that there is a limited number of data points. Small numbers of data points lead to an upward bias in the correlation results and increased sampling noise, so the results should be considered as indicative only. Despite this lack of formal statistical robustness, we should take comfort from the fact that the aggregate correlation results are collectively consistent and intuitive. Furthermore, the results of the other analyses in this section provide additional support for the existence a long-term relationship between the two populations.

4.2. Survival Rates

Comparing long-term survival rates provides a different perspective on basis risk and the relationship between the longevity experiences of the two populations. This is because long-term (multiyear) survival rates involve mortality rates for different ages across different years.[15] Figure 3(a) shows the evolution of 10-year survival rates for the two populations for 65-year-old males over the period 1970–2005. The 10-year survival rate at age 65 therefore shows the proportion of 65-year-olds surviving to age 75. Both survival rates have been increasing over time, but, more importantly, the ratio between them has been more or less constant over time, as shown in Figure 3(b). The latter chart suggests a relatively stable long-term relationship between the survival rates of the two populations.

Note that the survival ratio of assured to national survival rates is greater than one for all ages and increases with age. The average survival ratios over the period are listed in Table 6 along with some summary statistics on the variation in the survival ratio through time. For example, 45-year-old males have an average survival ratio of 1.03 compared with an average of 1.55 for 80-year-olds. This means that the assured population has 3% more 45-year-old males surviving to age 55 than the national population. Similarly the assured population has 55% more 80-year-olds surviving to age 90.

[15] Short-term survival rates, by contrast, provide the same perspective as mortality rates since a one-year survival rate is just one minus the corresponding mortality rate.

Figure 3

Historical 10-Year Male Survival Rates for the U.K. Assured and England and Wales National Populations Based on Data over the Period 1961–2005: (a) Historical Evolution of 10-Year Survival Rates for Males Reaching 65 in Different Years between 1970 and 2005, (b) Ratio of the 10-year Survival Rate for the Assured Population to the 10-Year Survival Rate for the National Population for Males Reaching Various Ages in Different Years between 1970 and 2005

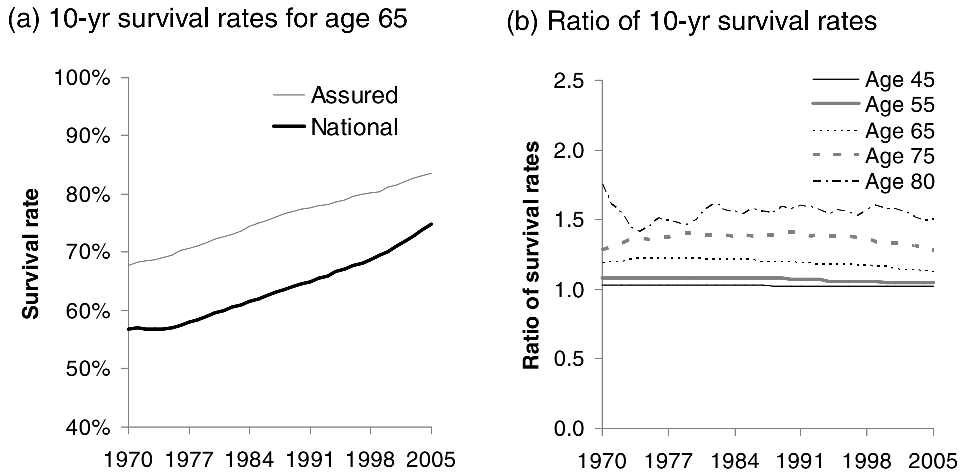

4.3 Life Expectancy

Another perspective on basis risk comes from *period* life expectancy.[16] This is a measure of how much longer on average individuals would be expected to live and, in a pension context, of how much longer one would expect that retirement income must be paid. Note that these results are dependent on the method used to estimate mortality rates at very high ages for which only limited mortality data are available.

Figure 4 shows the evolution of (curtate) period life expectancy for selected ages over 1961–2005. Note that despite the different levels of life expectancy, the ratio between the life expectancies of the two populations is relatively constant through time and increases with age. The ratio averages 1.14 at age 45, 1.22 at age 65, and 1.24 at age 80.

Table 6

Key Statistics on the Male Survival Ratio between the U.K. Assured and England and Wales National Populations

10-Year Survival Ratio (Assured/National)	Age 45	Age 55	Age 65	Age 75	Age 80
Average survival ratio	1.03	1.07	1.19	1.36	1.55
Standard deviation	0.004	0.011	0.029	0.037	0.063
Coefficient of variation (std dev/average)	0.4%	1.1%	2.4%	2.7%	4.1%
Worst case (max/average)	0.6%	2.0%	6.2%	6.1%	13.8%

Note: The survival ratio is defined as the 10-year survival rate for the assured population to the 10-year survival rate for the national population over the period 1970–2005. Survival rates are calculated for each age cohort using graduated mortality rates. The quoted age represents the age at the start of the 10-year period.

[16] Period life expectancy is calculated from the spot (i.e., current) mortality curve, assuming no further mortality improvements. It is generally acknowledged that period life expectancy underestimates actual (i.e., cohort) life expectancy because mortality rates are widely expected to continue to fall through time. However, period life expectancy has the advantage of being an objective metric.

Figure 4

Evolution of Male Period Life Expectancy for the U.K. Assured and England and Wales National Populations, 1961–2005: (a) Life Expectancy for 65-Year-Old Males Measured in Years, (b) Ratio of Life Expectancy for the Assured Population to the Life Expectancy for the National Population for Various Ages

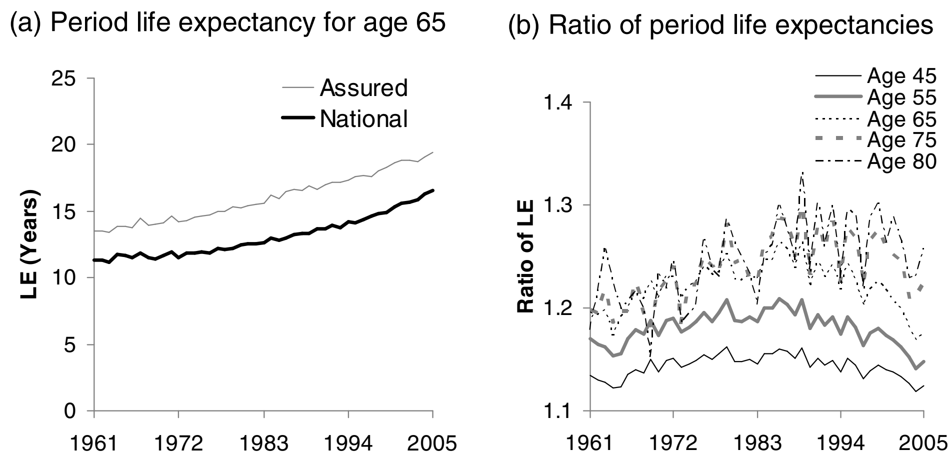

Over the entire 45-year period 1961–2005, period life expectancy has increased significantly for both populations at all ages. The assured data show greater increases in life expectancy than the national data. In particular, the highest increases have occurred for assured males in their 30s, with life expectancy for 33-year-olds increasing by 8.09 years, compared with 7.39 years for national population males of the same age. Figure 5 compares the change in life expectancy by age for the two populations. We stop calculating period life expectancy at age 80 to avoid the results being overly impacted by the method of graduating mortality rates for higher ages, that is, age 90 and above.

Between ages 30 and 80, Figure 5(a) shows that the difference in life expectancy between the two populations varies from 0.58 years for 50-year-olds to 0.98 years for 80-year-olds. It is evident that there

Figure 5

Increase in Male Period Life Expectancy for the U.K. Assured and England and Wales National Populations, 1961–2005: (a) Increase Measured in Years, (b) Increase in Percentage Terms

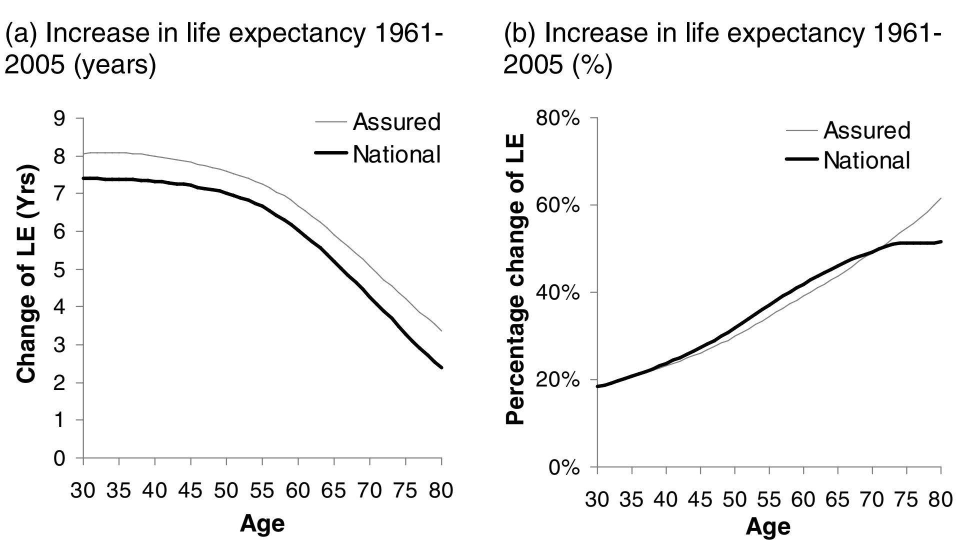

are clear differences between the populations in terms of the increase in period life expectancy measured in years. However, when we consider the relative percentage increase in life expectancy over the period, we find that the two populations have behaved in a very similar way, as illustrated in Figure 5(b). Both populations exhibit virtually the same percentage increases for ages 30 to 75. There is a divergence above 75, which may well be caused by differences in the graduation methodology used for higher ages.

The conclusion we draw from the analysis above is that the data for period life expectancy, like the other metrics we have examined, are indicative of a stable long-term relationship between the two populations, which is likely to have a favorable impact on hedge effectiveness.

4.4 Liability (Annuity) Cash Flows

Comparing the historical cash flows paid by annuities for different cohorts in each population provides yet another perspective on basis risk. To minimize the noise in comparing the two populations we focus on cumulative cash flows over periods of 10 years.

This particular metric is closely related to survivorship; in fact, it essentially corresponds to the average survival rate over the 10-year period. This should not be surprising. The annuity cash flow in any given year is proportional to the survival rate to the end of that year. So the 10-year cumulative cash flow is proportional to the sum of survival rates over periods ranging from one year to 10 years (Fig. 6(a)). The calculation assumes that the annuity pays £1 each year to each surviving member of each population.

Figure 6(b) shows the ratio of 10-year cumulative annuity cash flows for the assured population to those of the national population. Each line represents the ratio over time for the same initial age. The chart shows the evolution of this ratio between 1970 and 2005, where the age at the start of the 10-year period runs from 60 to 80 years.[17] Several points are worthy of note:

Figure 6

Cumulative Cash Flows over 10-Year Horizons for Liabilities (Annuities) Based on the U.K. Assured and England and Wales National Male Populations over the Period 1970–2005: (a) Cumulative Cash Flow for 65-Year-Old Males Measured in Pounds Sterling, (b) Ratio of Cumulative Cash Flows for the Assured Population to That for the National Population

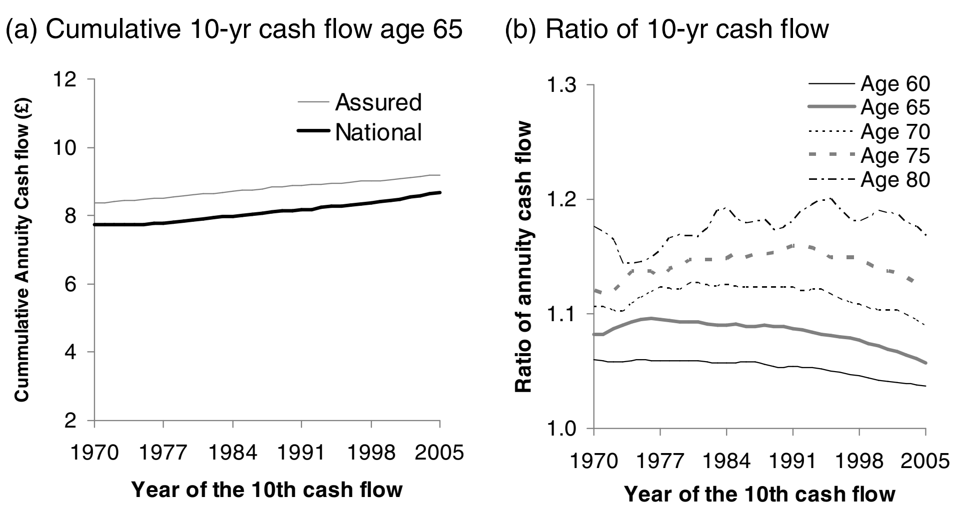

[17] Note that this metric cannot start before 1970, since it requires 10 years of historical cash-flow data to calculate the first cumulative cash-flow value.

- The ratios are all greater than one, again reflecting higher survival rates for the assured population. They vary between approximately 1.04 to 1.20, depending on the cohort and the year.
- The ratios are reasonably stable, and trend downward very slightly through time for cohorts with an initial age below 70. In particular, the ratio for the cohort with an initial age of 60 varies between 1.04 and 1.06 over the period and for an initial age of 70 between 1.09 and 1.13.
- Cohorts born in earlier years (and therefore of higher initial age) have consistently higher ratios than those born in later years.
- Volatility increases for older birth cohorts (i.e., as the initial age increases). This is consistent with higher levels of noise associated with the higher age data.

4.5 Liability (Annuity) Values

The prices of lifetime annuities give an alternative perspective on basis risk that relates directly to hedge effectiveness in monetary terms. A life annuity purchased from a life office by a pension plan for each of its members on their retirement provides a perfect hedge for the member's longevity risk. Since adequate data on actual annuity prices for the assured and national populations are not publicly available, we compare theoretical annuity prices based on the kind of pricing (or valuation) model that the life office might itself use to price annuities. The results are therefore more dependent on modeling considerations than the other analyses in the paper: In particular, a projection model needs to be used to forecast future mortality rates, and there also needs to be a model to calibrate mortality rates at ages higher than those available in published data (see the Appendix).[18]

The annuity pricing model takes the graduated mortality rates from the previous year as the mortality base table for each population in a particular year. Then future mortality rate forecasts relative to that base table are made by applying a projection model to historical data. We will suppose the life office uses a (single-population) Lee-Carter (1992) model with 30 years of historical data to generate these projections for each population for every year from 1991 to 2006.[19] The resulting set of mortality rates constitutes a complete cohort mortality table for each base year. From this cohort mortality table, the life office will calculate the stream of expected cash flows paid out for each annuity. Finally, the annuity price is just the present value of the expected cash flows, and we assume the life office uses a constant discount rate of 5%. One benefit of this assumption is that it allows us to focus on the change in annuity prices that are caused purely by changes in longevity.[20]

Figure 7 shows the results of comparing annuity prices based on assured and national mortality data for the same age. Figure 7(a) shows the stable, slightly upward-sloping relationship over time between the annuity prices for the assured and national populations for age 65. Figure 7(b) shows the ratio of the prices of assured population annuities to national population annuities through time for ages 60 to 80, with each line showing how the ratio evolves through time for a given age. The annuities are assumed to pay equal amounts as long as an individual is alive at the end of the relevant year. The annuities are first priced on January 1, 1991, using the 1990 graduated mortality rates as the initial mortality base table and mortality rate projections generated from this date using historical rates between 1961 and 1990. In the following year, 1991 graduated mortality rates are used as the base table and mortality rate projections are generated using historical rates between 1962 and 1991, and so on.

[18] The price of a lifetime annuity depends on mortality rates for all ages above the current age and, in particular, on the mortality rates at very high ages for which there are many fewer data points available and consequently higher levels of noise.

[19] Note that we cannot start before 1991, in order that we have 30 years of historical data to calibrate the projection model. This projection model is used only for the pricing (valuation) of the annuities, not for simulating scenarios.

[20] In practice, interest-rate risk does matter, but this could be hedged using an interest-rate swap.

Figure 7

Liability Values for Selected Ages through Time for the U.K. Assured and England and Wales National Male Populations over the Period 1991–2006: (a) Liability Values for 65-Year-Old Males Measured in Pounds Sterling, (b) Ratio of Liability Values for the Assured Population to That for the National Population

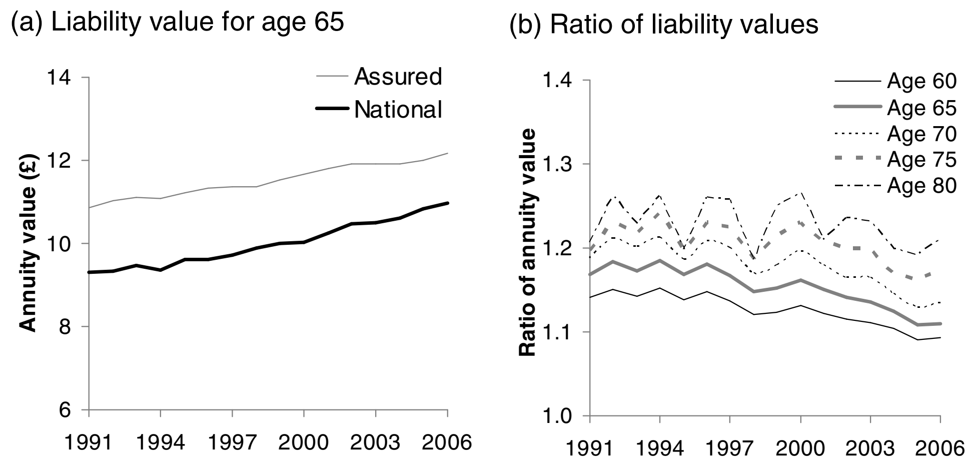

The main observations from Figure 7(b) are as follows:

- The ratios are all greater than one, again confirming that the average cost of providing pensions for the assured population is higher than that for the national population. The ratio varies between approximately 1.09 to 1.27, depending on the age and the year.
- Higher ages demonstrate consistently higher ratios.
- The ratios are reasonably stable, but trend downward slightly through time. In particular the ratio for age 60 varies between 1.09 and 1.15 over the period and for age 70 between 1.13 and 1.21.
- Volatility around the trend increases as the age increases.

Figure 8 provides a different presentation of the data in which the lines follow the same cohort through time (rather than the same age but at different times). Figure 8(a) shows that, as the cohort aged 65 in 1991 gets older and its remaining life expectancy falls, the liability value also falls in a stable way for both assured and national populations. Figure 8(b) shows the ratio of the prices of assured population annuities to national population annuities through time for cohorts aged between 60 and 80 years at the start of the analysis period. Each line shows how the ratio of annuity prices evolves through time for a single cohort. This chart exhibits features that are consistent with the previous chart, namely:

- The ratios are all greater than one.
- Older birth cohorts demonstrate consistently higher ratios than those born in later years.
- The ratios are reasonably stable, but trend upward slightly through time. In particular, the ratio for the cohort with an initial age of 60 varies between 1.14 and 1.19 over the period and for an initial age of 70 between 1.18 and 1.27.
- Volatility increases for older birth cohorts and as a cohort ages.

These results constitute a retrospective hedge effectiveness test and are consistent with those of Coughlan et al. (2007b, pp. 80–81).

Despite the modeling assumptions that are an inevitable part of the annuity pricing calculation, we again see evidence of a relatively stable relationship between the two populations from both a period and a cohort perspective.

Figure 8

Liability Values for Selected Cohorts through Time for the U.K. Assured and England and Wales National Male Populations over the Period 1991–2006: (a) Liability Values for Male Cohorts Aged 65 in 1991 Measured in Pounds Sterling, (b) Ratio of Liability Values for Selected Cohorts for the Assured Population to That for the National Population

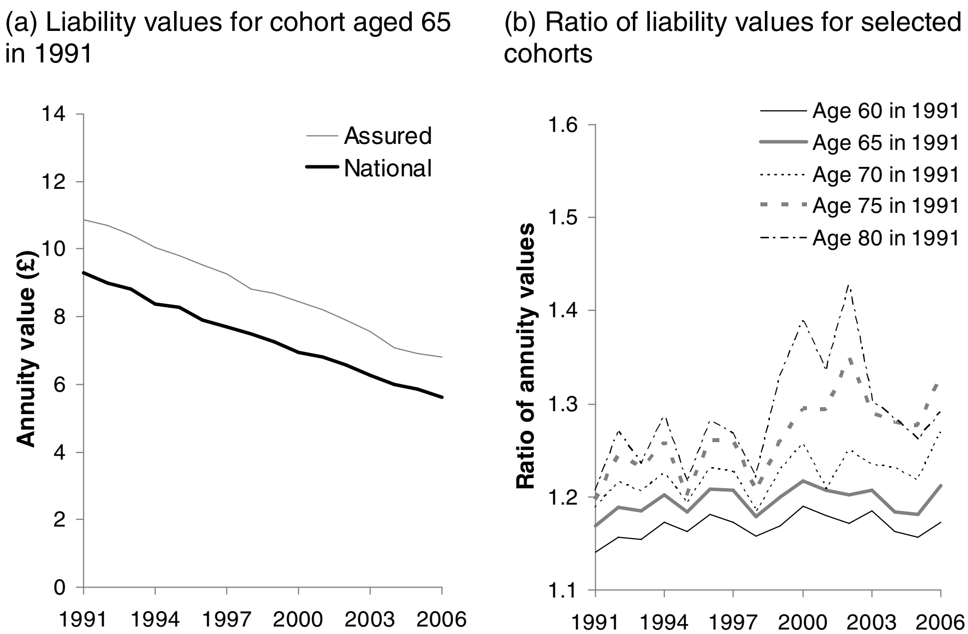

4.6 Hedge Effectiveness Example

Considering the above analyses along with the context of the CMI assured population as an affluent, but more volatile, subset of the national population, the evidence suggests that there has been a stable long-term relationship between their mortality experiences. This is particularly significant given their different mortality levels, their different mortality improvements, and the amount of noise in the CMI assured population data.

The implications of this for longevity hedges indexed to national population mortality data are clear. The long-term effectiveness of such hedges—providing they are appropriately calibrated—should be relatively high, leading to a significant reduction in longevity risk.

To illustrate this, we have evaluated the effectiveness of a static longevity hedge linked to the LifeMetrics Index for England and Wales in reducing the longevity risk of a hypothetical pension plan with the same mortality behavior as the CMI assured population. We focus on a particular retrospective hedge effectiveness test based on historical data.[21] The approach to calculating hedge effectiveness is the one outlined in Section 3.3 and summarized in Table 1 and Figure 1.

4.6.1 Step 1: Hedging Objectives

The pension plan is assumed to consist entirely of deferred male members currently aged 55, whose mortality characteristics are the same as the CMI assured population and who will receive a fixed pension of £1 for life, beginning at retirement in 10 years' time (the hedge horizon) at age 65. The hedging objective is to remove the uncertainty in the *value* of the pension at retirement due to longevity risk.

[21] This avoids the requirement, essential for a prospective test, of having to choose a two-population stochastic mortality model.

4.6.2 Step 2: Hedging Instrument

The hedging instrument is a 10-year deferred annuity swap that pays out on the basis of a survival index for the national population of England and Wales for 55-year-old males. As we are considering a hedge of value, we can assume (without impacting the economics) that the hedging instrument is cash-settled at the hedge horizon at the market value prevailing at that time. In other words, in 10 years' time, the pension plan receives a payment reflecting the market value of the hedging annuity at that time in return for making a fixed payment at that time. So the hedging instrument involves a net settlement that is the difference between the fixed payment and the market value of the hedging annuity in 10 years' time.

Because the hedging population is different from the exposed population, the hedge ratio needs to be calibrated to reflect the relationship between their mortality improvements. Applying equation (2) results in a hedge ratio of 0.987, meaning that to hedge a £1 liability requires £0.987 of the hedging instrument.

4.6.3 Step 3: Method for Hedge Effectiveness Assessment

As mentioned above, we are performing a retrospective effectiveness test, based on historical data. The basis for comparison that we use is twofold involving evaluation of (1) the correlations between the unhedged and hedged liability and (2) the degree of risk reduction. Since the objective is to hedge the value of the pension liability, we focus on a risk metric corresponding to the value-at-risk (VaR) in 10 years' time, where the VaR is measured at a 95% confidence level relative to the median. We measure hedge effectiveness by comparing the VaR of the pension before and after hedging. We use historical mortality data to directly evaluate historical scenarios for the evolution of mortality rates over a 10-year horizon, from which the VaR of the pension can be calculated. For asymmetric, nonnormal distributions a risk metric such as VaR is preferred over standard deviation, as it more accurately captures the downside risk and the associated implications for economic capital. However, other risk metrics generally give similar results.

The hedge effectiveness is calculated in terms of the relative risk reduction in equation (3) and implemented as follows:

$$RRR = 1 - \mathrm{VaR}_{(\mathrm{Liability+Hedge})}/\mathrm{VaR}_{\mathrm{Liability}}. \tag{4}$$

We construct scenarios for the hedge effectiveness analysis in a model-independent way directly from the historical mortality data. Because the amount of available historical data is limited to 45 years, we form historical scenarios using the approach summarized in Table 7. This approach involves combining the set of historically realized mortality improvements with the set of realized mortality base tables. In particular, we construct 1,575 scenarios for each population by applying realized mortality improvements coming from the full historical set of 35 overlapping 10-year periods or *windows* (1961–1971,

Table 7

The Approach Used to Generate 1,575 10-Year Scenarios for Mortality Rates Combining the 45 Historical Base Tables with the 35 (Overlapping) 10-Year Historical Mortality Improvements for Each Base Table

Component of the Approach	Applies at	Calculated from	Data Used for Pension Liability	Data Used for Hedging Instrument
1. Base mortality table	$t = 0$	Historical period mortality rates for each year 1961–2005	Member-specific mortality data	National mortality index data
2. Realized mortality rates for each scenario	$t = 1, 2, 3, \ldots, 10$	Historical 10-year mortality improvements applied to each base table in component (1)	Member-specific mortality data	National mortality index data

Note: Each component of the approach builds on the previous one. The inception of the hedge corresponds to time $t = 0$ and the hedge horizon to time $t = 10$.

1962–1972, . . . , 1995–2005) to each of the realized mortality *base tables* defined by the observed mortality rates in each of the 45 years (1961, 1962, . . . , 2005).

To make this more explicit, let us denote the inception date of the hedge $t = 0$ and the hedge horizon $t = 10$. Let q_x^b denote the realized historical mortality rate in year b for age x. For a given value of b, we will use q_x^b to provide us with a base table at time 0 in our generation of a set of scenarios. For each of the 45 base tables, we have $w = 1, \ldots, 35$ 10-year historical windows that we use to generate 35 scenarios. For each pair (b, w), we generate one mortality-rate scenario $q_x^{(b,w)}(t)$ for each time $t = 0, 1, \ldots, 10$. Specifically, the mortality rate scenarios for each population are given by

$$q_x^{(b,w)}(t) = q_x^b \times \frac{q_x^{w+t}}{q_x^w}, \tag{5}$$

where $b = 1961, 1962, \ldots, 2005$; $w = 1961, 1962, \ldots, 1995$; $t = 0, 1, \ldots, 10$. Equation (5) consists of the base table giving us the scenario's mortality rates for $t = 0$, and q_x^{w+t}/q_x^w is the t-year improvement factor at age x in window w.

This approach provides us with 45 sets of scenarios with each set comprising 35 different 10-year scenarios. We evaluate hedge effectiveness by examining (1) the 45 sets of scenarios separately and (2) all the sets in aggregate (1,575 scenarios in total).

The valuation, or pricing model, that we use to value both the liability and the hedging instrument at time $t = 10$ involves projecting the expected (i.e., best estimate) future stream of cash flows beyond time 10 for the liability and the hedging instrument, and then discounting those cash flows back to time 10, the time at which the valuation is to be made. In particular, the valuation under the different scenarios at time 10 requires us to come up with a best estimate projection of mortality rates for that scenario beginning at time $t = 10$.

Note that this pricing model is the only model-dependent element used in this example and models other than the one we describe below could also be used. (We reiterate that in this example the simulation of mortality rates and basis risk up to the hedge horizon is model independent, since it uses actual historically observed mortality data for each population.)

In order to value the pension liability and the hedging instrument at time 10, we use a very simple approach to project best estimate mortality rates into the future beginning from $t = 10$ in each scenario. This simple projection method averages the one-year percentage mortality improvements from $t = 0$ to 10 and applies this average improvement rate to the scenario mortality rate at $t = 10$ for all future years. In other words, in each scenario, starting at year 10, we develop a forecast for the future path of best estimate mortality based on the observed history up until that point in time.[22] We will assume for the purpose of valuation that the mortality improvement projections for the members of the pension plan are the same as those for national population. We justify this on the grounds that common industry practice is to use mortality projections derived from a large, broad-based population (e.g., the national population) to calibrate the mortality projections for a specific pension plan for the purposes of valuation. This is because, in practice, the vast majority of pension plans do not have sufficient data to develop their own mortality projections from their historical experience alone. It can also be justified intuitively on the grounds that the long-term relationship between the two populations that we have demonstrated above implies that the long-run projections for the two populations cannot diverge in a systematic way. This is the principle that lies behind the two-population mortality models discussed in Section 2.4. We have explored the implications of relaxing this assumption and find a modest fall in hedge effectiveness if separate, independent mortality improvement projections are used.

To value the liability in a particular scenario, we start from the best estimate mortality projection in year 10, calculate the stream of expected cash flows that will be paid out in that scenario, and then calculate the present value of the expected cash flows at $t = 10$. The value of the pension liability in that scenario is therefore given by

[22] So the value in each scenario at the horizon in 10 years' time depends on a projection based on the history of that particular scenario.

$$V^{(b,w)}_{\text{Liability}}(10) = {}_{10}p^{(b,w)}_{55}(0) \sum_{t=1}^{\infty} {}_{t}\hat{p}^{(b,w)}_{65}(10) \times DF(t). \tag{6}$$

Here ${}_{10}p^{(b,w)}_{55}$ is the (b,w) scenario's historical 10-year survival rate for 55-year-olds in scenario (b,w) from $t = 0$ until $t = 10$ and ${}_{t}\hat{p}^{(b,w)}_{65}(10)$ is the projected t-year survival rate at the hedge horizon for the same cohort (now aged 65) in that scenario, with the hat over the p to indicate its dependence on the choice of valuation model. The discount factor for calculating the present value is denoted by DF, and our calculations assume a flat discount rate of 5% for simplicity. The hedging instrument is valued in a similar manner.

4.6.4 Step 4: Calculation of Hedge Effectiveness

The first part of the calculation of hedge effectiveness requires the simulation of the 1,575 historical scenarios for the evolution of mortality rates using the method described above.

The second part of the calculation of hedge effectiveness involves determining the values of the pension and the hedging instrument in each scenario at the end of the 10-year hedge horizon. This is done by applying the valuation model described above in equation (6) to each of the scenarios at the hedge horizon.

This provides us with a distribution of 1,575 possible values for the pension and a distribution for 1,575 possible values for the hedging instrument in 10 years' time. From these, we calculate the correlation between the liability values and the hedging instrument values, as well as the relative risk reduction. The results of these calculations are summarized in Table 8. Implementing the scenarios as 45 different sets, each with a different base table, gives high correlations between the values in all scenario sets. However, the relative risk reduction is modest. One factor contributing to this modest risk reduction is the fact that we are dealing with a small number of scenarios in each set (i.e., only 35 scenarios), and these involve overlapping periods. As a result, there is a relatively small amount of volatility in each scenario set, leading to a modest risk reduction.

By contrast, aggregating the 1,575 scenarios introduces considerably more volatility in the scenarios set and leads to much higher hedge effectiveness results. For the aggregate scenario set, the hedge reduces the impact of longevity risk on the value of the pension liability by 82.4%, demonstrating a high level of hedge effectiveness. This is shown in the histogram in Figure 9. The light-colored bars represent the range of possible outcomes for the unhedged value of the pension liability in 10 years' time coming from the 1,575 scenarios. The dark-shaded bars reflect a much narrower range of possible outcomes for the value of the pension liability after hedging. The fact that the histogram for the hedged

Table 8

Results of the Hedge Effectiveness Analysis for the Hypothetical U.K. Pension Plan

Method of Using the Scenarios	Correlation between Value of Liability and Hedging Instrument	Hedge Effectiveness	Comments
45 sets of 10-year scenarios, each set with a different mortality base table and 35 scenarios	Range: 0.93–0.97 Average: 0.96	Range: 67–79% Average: 73%	Scenarios in each set involve overlapping 10-year improvements and are not fully independent 35 scenarios is a small number for a fully comprehensive analysis
Aggregate set of 1,575 scenarios	0.98	82.4%	Adequately large number of scenarios Scenarios include different base tables, which increases the range of scenario outcomes

Note: Hedge effectiveness is measured in terms of relative risk reduction with VaR as the risk metric, as in eq. (4).

Figure 9

Hedge Effectiveness Analysis Using an Index-Based Hedging Instrument Linked to England and Wales National Male Mortality to Hedge the Longevity Risk of a Pension Plan with the Same Mortality Characteristics as the U.K. Assured Male Population

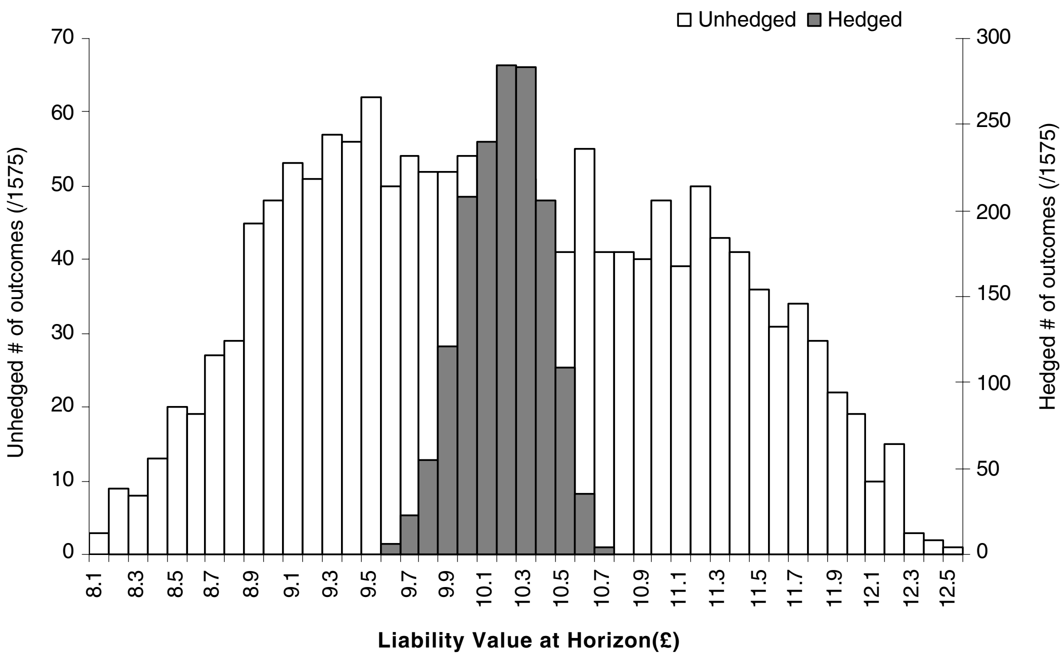

case is much narrower than that for the unhedged case provides a visual illustration of the high degree of hedge effectiveness.

5. A Robustness Check Using U.S. Data

The analysis of the previous section has been repeated in a U.S. context as a robustness check of the results. The populations used were the U.S. national population and the population of the state of California (both based on data sourced from the Centers for Disease Control and Prevention [CDC] and the National Census Bureau). The data in this analysis covered the 25-year period 1980–2004. California has a higher level of affluence than the nation as a whole with per capita GDP 11% above the national average.[23] This greater affluence is reflected in mortality rates that have been consistently lower and mortality improvements higher than those of the national population. The results are very similar to those obtained in the U.K. example above, leading to a similarly high degree of hedge effectiveness associated with an appropriately calibrated index-based hedge linked to U.S. national population longevity. Using an identical methodology to that described above, we evaluated the effectiveness of an index-based hedge linked to the U.S. national population in reducing the longevity risk associated with a pension plan whose mortality experience reflects that of the state of California. We obtained an aggregate hedge effectiveness result of 86.5%.

6. Conclusions

In this paper, we have developed a framework for analyzing longevity basis risk and its implications for the effectiveness of longevity hedges. Note that the framework does *not* assume any particular model for basis risk or for valuation. Such a framework is essential for

[23] U.S. Census Bureau and U.S. Bureau of Economic Analysis, 2008 figures.

- Understanding basis risk
- Calibrating two-population stochastic mortality models
- Calibrating index-based longevity hedges
- Measuring hedge effectiveness.

The framework is built on a quantitative analysis of data, together with a qualitative understanding of the contextual relationship between the populations involved. The quantitative analysis involves examining the historical experience of the populations in terms of different metrics. The nature of longevity risk dictates that the time horizon for the analysis should be long, but the framework acknowledges that there will not, in general, be enough data for a robust analysis over such horizons. As a result, an assessment of the nature and magnitude of basis risk must rest on professional judgment, informed by evidence coming from analysis, experience, and context.

A key element of this framework is a structured approach to hedge effectiveness assessment, based on a proven approach developed for derivatives accounting. To obtain a valid and meaningful hedge effectiveness test requires careful attention to (1) the design of the test, (2) the generation (simulation) of appropriate scenarios, and (3) the valuation methodology for the underlying exposure and the hedging instrument in each scenario. The latter two are the only places where model choice enters the application of the framework.

We have applied the framework to a detailed case study involving empirical analysis in a U.K. context of the relationship between the mortality experience of the national population and that of a more affluent subpopulation. Despite the different demographic profiles between these related populations, we demonstrate evidence of high correlations in mortality improvements between them and a stable long-term relationship across different metrics. Similar results were also obtained in a U.S. case study. These results have very favorable implications for the effectiveness of appropriately calibrated, index-based longevity hedges. From this, we conclude that longevity basis risk between a pension plan, or annuity portfolio, and a hedging instrument linked to a broad national population-based longevity index can in principle be reduced very considerably.

APPENDIX: TECHNICAL ISSUES

A.1 AGE BUCKETS

As discussed in the text, bucketing ages together is frequently beneficial in removing noise from mortality rates and is often implemented in actual transactions to hedge both mortality risk and longevity risk. Where age buckets are used in this paper, the buckets all involve simple averages over each age in the bucket. So, for example, in computing the mortality rate for the age bucket 60–69 years old, we take a simple unweighted average of the mortality rates for each age $q_{60}, q_{61}, q_{62}, \ldots, q_{69}$. Similarly, the survival rate for the age bucket 60–69 is an average of the survival rates for each individual age $p_{60}, p_{61}, p_{62}, \ldots, p_{69}$.

A.2 GRADUATION METHOD FOR MORTALITY RATES

The case study presented in this paper is based on mortality rates that have been calculated and graduated using a method very similar to that used in the LifeMetrics Longevity Index (see Coughlan et al. 2007b for details). The graduation method is an objective one that is applied consistently in each year to each data set. In particular, we have used the same cubic spline approach as used for the LifeMetrics Index to smooth the crude central mortality rates, but we have used a variant of the higher age methodology to come up with graduated mortality rates for ages above 80.

Graduated mortality rates for higher ages, which are needed to calculate period life expectancies and annuity prices, are obtained by the method described below:

1. Calculate crude central mortality rates m_x for each age x for the year in question.

2. Graduate the crude central mortality rates for all ages up to age 89 using the cubic splines approach described in Coughlan et al. (2007b). Denote these first-stage graduated rates by m'_x.
3. For ages below 80, set the graduated central mortality rate m^g_x to be m'_x.
4. For ages 80–89:
   a. Perform a regression of log m'_x against age, and back out the fitted central mortality rates m''_x for each age.
   b. Set the graduated central mortality rate m^g_x to be a blend of m'_x and m''_x, such that at age 80 the blend is 100% m'_x and at age 89 it is 100% m''_x. The following monotonic blending function is used:

$$f(z) = \begin{cases} \exp(1) \times \exp[-1/(1-z^2)] & \text{for } 0 \le z < 1 \\ 0 & \text{for } z = 1, \end{cases} \tag{A1}$$

   where $z = (x - 80)/9$ for ages x between 80 and 89.
5. For ages 90 and over, calculate the graduated central mortality rate m^g_x using the algorithm outlined in Coughlan et al. (2007b), which involves fitting a cubic polynomial.
6. Calculate graduated initial mortality rates q_x for each age using the transformation:

$$q_x = m^g_x/(1 + m^g_x/2). \tag{A2}$$

A.3 CALCULATIONAL METHOD FOR "AGGREGATE CORRELATIONS"

In calculating correlations in mortality rates between two populations, the correlations relate to both changes and relative changes in mortality rates δq for each population over a particular time horizon, H, that varies from one year to 20 years. Nonoverlapping periods are used, which means that, for long horizons, there are fewer historical data points available. Note that because of this lack of data when considering long horizons, we calculate what we call "aggregate correlations", which jointly compare changes for different periods and for different ages. In other words, we correlate matrices $[\delta q^1_{x,t}]$ and $[\delta q^2_{x,t}]$, where x labels the age and t labels the time period t to $t + H$. Consider the example of calculating aggregate correlations for the mortality rate changes for individual ages from 50 to 89 over a five-year horizon ($H = 5$ years). The calculation involves computing the correlation between two 40×9 matrices of the following form:

$$\begin{bmatrix}
\delta q(x=50, t=1961\text{–}1966) & \delta q(x=50, t=1966\text{–}1971) & \delta q(x=50, t=1971\text{–}1976) & \ldots & \delta q(x=50, t=2001\text{–}2005) \\
\delta q(x=51, t=1961\text{–}1966) & \delta q(x=51, t=1966\text{–}1971) & \delta q(x=51, t=1971\text{–}1976) & \ldots & \delta q(x=51, t=2001\text{–}2005) \\
\delta q(x=52, t=1961\text{–}1966) & \delta q(x=52, t=1966\text{–}1971) & \delta q(x=52, t=1971\text{–}1976) & \ldots & \delta q(x=52, t=2001\text{–}2005) \\
\ldots & \ldots & \ldots & \ldots & \ldots \\
\delta q(x=89, t=1961\text{–}1966) & \delta q(x=89, t=1966\text{–}1971) & \delta q(x=89, t=1971\text{–}1976) & \ldots & \delta q(x=89, t=2001\text{–}2005)
\end{bmatrix}$$

When aggregate correlations are calculated from age buckets rather than individual ages, the calculation is identical except that the rows in the $[\delta q^1_{x,t}]$ and $[\delta q^2_{x,t}]$ matrices are indexed by the age bucket rather than the individual single year of age.

Note that we use the same symbol above for absolute and relative changes in mortality rates, although they are defined differently. Absolute changes in mortality rates are defined by $q(t + 1) - q(t)$, and relative changes are defined by $[q(t + 1) - q(t)]/q(t)$.

DISCLAIMER

REFERENCES

CAIRNS, A. J. G., D. BLAKE, K. DOWD, G. D. COUGHLAN, AND M. KHALAF-ALLAH. 2011. Bayesian Stochastic Mortality Modelling for Two Populations. *ASTIN Bulletin* 41(1): 29–59.

COUGHLAN, G. D. 2009a. Hot Topics in Risk Management: Longevity Hedging. Paper presentated at the Actuarial Profession's Risk and Investment Conference, June 21–23, Brighton, UK.

COUGHLAN, G. D. 2009b. Longevity Risk Transfer: Indices and Capital Market Solutions. In *The Handbook of Insurance Linked Securities,* edited by P. M. Barrieu and L. Albertini, pp. 261–281. London: John Wiley.

COUGHLAN, G. D., S. EMERY, AND J. KOLB. 2004. HEAT (Hedge Effectiveness Analysis Toolkit): A Consistent Framework for Assessing Hedge Effectiveness under IAS 39 and FAS 133. *Journal of Derivatives Accounting* 1(2): 221–272.

COUGHLAN, G. D., D. EPSTEIN, J. HEVIA-PORTOCARRERA, M. KHALAF-ALLAH, C. S. WATTS, AND P. JOSEPH. 2007a. LifeMetrics: Netherlands Longevity Index. Technical Document Supplement. October 22. London: JPMorgan. www.lifemetrics.com.

COUGHLAN, G. D., D. EPSTEIN, A. ONG, A. SINHA, I. BALEVICH, J. HEVIA-PORTOCARRERA, E. GINGRICH, M. KHALAF-ALLAH, AND P. JOSEPH. 2007b. LifeMetrics: A Toolkit for Measuring and Managing Longevity and Mortality Risks. Technical Document. March 13. London: JPMorgan. www.lifemetrics.com.

COUGHLAN, G. D., D. EPSTEIN, A. SINHA, AND P. HONIG. 2007c. q-Forwards: Derivatives for Transferring Longevity and Mortality Risk. July 2. London: JPMorgan. www.lifemetrics.com.

COUGHLAN, G. D., D. EPSTEIN, C. S. WATTS, M. KHALAF-ALLAH, P. JOSEPH, AND Y. YE. 2008. LifeMetrics: Germany Longevity Index. Technical Document Supplement. April 7. London: JPMorgan. www.lifemetrics.com.

COX, S. H., AND Y. LIN. 2007. Natural Hedging of Life and Annuity Mortality Risks. *North American Actuarial Journal* 11(3): 1–15.

CURRIE, I. 2009. Basis Risk in Mortality Projections. Longevitas Limited. www.longevitas.co.uk/site/informationmatrix/basisriskinmortalityprojections.html.

DAHL, M., AND T. MØLLER. 2006. Valuation and Hedging of Life Insurance Risks with Systematic Mortality Risk. *Insurance: Mathematics & Economics* 39: 193–217.

DAVIES, P. J. 2011. JPMorgan Strikes an Important Deal for Longevity Risk Trading. *Financial Times,* February 1.

DOWD, K., A. J. G. CAIRNS, D. BLAKE, G. D. COUGHLAN, AND M. KHALAF-ALLAH. 2011. A Gravity Model of Mortality Rates for Two Related Populations. *North American Actuarial Journal* 15(2): 334–356.

FRIEDBERG, L., AND A. WEBB. 2007. Life Is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk. *B.E. Journal of Economic Analysis & Policy* 7(1): Article 31.

JARNER, S. F. 2008. Small-Region Mortality Modelling. Paper presented at the Fourth International Longevity Risk and Capital Markets Solutions Conference, Amsterdam, September 5, 2008.

JARNER, S. F., AND E. M. KRYGER. 2009. Modelling Adult Mortality in Small Populations: The SAINT Model. Pensions Institute Working Paper PI-0902, January.

LEE, R. D., AND L. R. CARTER. 1992. Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association* 87: 659–675.

LI, J. S. H., AND M. R. HARDY. 2011. Measuring Basis Risk in Longevity Hedges. *North American Actuarial Journal* 15(2): 177–200.

LI, N., AND R. LEE. 2005. Coherent Mortality Forecasts for a Group of Populations: An Extension to the Lee-Carter Method. *Demography* 42(3): 575–594.

LIFE AND PENSIONS. 2008. Canada Life Hedges Equitable Longevity with JPMorgan Swap. *Life & Pensions* (October): 6.

LUCIDA PLC. 2008. Lucida and JPMorgan First to Trade Longevity Derivative. February 15. www.lucidaplc.com/en/news.

MERCER 2011. World's First Longevity Hedge for Non-Retired Pension Plan Members Completed. Press Release. February 1. http://uk.mercer.com/press-releases/1406520.

NGAI, A., AND M. SHERRIS. 2010. Longevity Risk Management for Life and Variable Annuities: Effectiveness of Static Hedging Using Longevity Bonds and Derivatives. University of New South Wales, Australian School of Business Research Paper No. 2010ACTL03. April.

PLAT, R. 2009. Stochastic Portfolio Specific Mortality and the Quantification of Mortality Basis Risk. *Insurance: Mathematics & Economics* 45(1): 123–132.

RICHARDS, S. J., AND G. JONES. 2004. Financial Aspects of Longevity Risk. Paper presented at the Staple Inn Actuarial Society, London, October 26.

STAPLETON, J. 2011. Pall Scheme Completes World's First Longevity Hedge for Non-Retired Members. *Professional Pensions,* February 1. www.professionalpensions.com.

SWEETING, P. 2007. Pricing Basis Risk in Survivor Swaps. *Life & Pensions* (September): 44–48.

SYMMONS, J. 2008. Lucida Guards against Longevity. *Financial News*, February 19. www.efinancialnews.com.

TRADING RISK. 2008. JPMorgan Longevity Swap Unlocks UK Annuity Market. *Trading Risk* (September/October): 3. www.trading-risk.com.

WANG, J. L., H. C. HUANG, S. S. YANG, AND J. T. TSAI. 2010. An Optimal Product Mix for Hedging Longevity Risk in Life Insurance Companies: The Immunization Theory Approach. *Journal of Risk and Insurance* 77(2): 473–497.

Discussions on this paper can be submitted until October 1, 2011. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.