The Nature of Longevity Risk

Sacha Dhamani

Presented to Staple Inn Actuarial Society on the 09 June 2015

1. Introduction

There have been a number of papers produced discussing how to model longevity risk but it has been a number of years since a paper has specifically been produced to consider the fundamental nature of longevity risk[1].

The IAA breakdown of longevity risk [2] (Trend, Level, Catastrophe and Volatility) is still perceived to be the standard definition used despite many practitioners considering them insufficient as a foundation for modelling longevity risk.

Whilst they continue to provide a common language for professionals to discuss longevity risk it is time for a revised conceptual framework for longevity risk which can provide a stronger foundation to take into account how different types of provider[3] are exposed to longevity risk and the differences in their approach to longevity risk management. This paper is intended to trigger a debate in this area by proposing such a framework.

The framework will not attempt to define how longevity risk should be modelled but will rather outline the different behaviours of longevity risk:

- Trend Uncertainty
- Trend Volatility
- Catastrophe Risk
- Basis Risk
- Underwriting Risk
- Mis-estimation Risk
- Statistical Volatility

A key point to mention at the outset is that there is no right answer – what is important that all sources of and effects on longevity are accounted for in the risk model. The argument of this paper is that it is best to start from a conceptual understanding of the risk and then determine the appropriate modelling techniques. The risk is that if one focuses on application of a mathematical model in the first instance then one may end up with a theoretically robust and well documented risk model which doesn't necessarily measure the risk being taken.

The effect on liabilities and risk profiles for different longevity risk behaviours will be discussed. Two of the key variations are:

- Mortality rating approach
- Depth and breadth of experience

The intention is to outline qualitatively the expected variation in risk capital based upon the behaviours above. However, it is not the intention to describe the importance and ease of mitigation of the different behaviours in quantitative detail – the variation between providers is so great that an

[1] Richards et al (2014) (1) provides the most recent example of this kind of discussion but focuses more on the modeling rather than conceptual framework.

[2] A Global Framework for Insurer Solvency Assessment

[3] By "provider" it is intended to cover all types of longevity insurance providers – annuity companies, occupational pension schemes, reinsurers, etc. This framework as a result is intended to be applicable to all entities. The term provider shall be used throughout to indicate all potential longevity insurance providers.

exhaustive discussion would extend the paper significantly (and detract from the key message).

It is often required to model or measure the risk over a particular time frame or to a particular level of probability. For example, Solvency II requires an estimate of the capital required to support a longevity event over a one year time horizon that has a 0.5% probability of occurring.

This is a particular example of how an understanding of longevity risk can be employed but is not the purpose of this paper - rather it is to outline the nature of the risk as a practitioner is exposed to it. Therefore, longevity will be considered as a run off risk as this should be the more appropriate economic consideration for providers.

I would like to thank (in alphabetical order) Madhavi Bakekal, Andrew Chamberlain, Sam Gutterman, Robin Houghton, Stephen Richards and Richard Willets for their help in reviewing and challenging this paper.

Special thanks to Nicola Leckenby for helping to produce the analytics shown in this paper.

2. Importance of a Conceptual Model

Before attempting to mathematically model longevity risk (or indeed any risk) it is critical to understand the risk being modelled. It is all too easy to start with a mathematical model and convince oneself that it provides a measure of the full range of outcome. In truth, any mathematical model should be a simplified representation of the underlying reality, to ensure transparency and understanding; however there is a risk that it will fail to fully account for all the risk exposures and fail to provide the illumination needed to adequately manage the risk. Critical to this are the risks that are not naturally susceptible to mathematical modelling (or where current best thinking is yet to uncover an appropriate modelling approach). Uncertainty risks (described further below in section 3) fall potentially into this category more than volatility risks (also described below in section 3).

The diagrams below (Figures 1 & 2) attempt to explain the effect of the different approaches. The totality of longevity risk is shown as the larger circle within which the models used to measure risk depending on the approach taken. A risk model schema is attempting to cover the entirety of the risk whilst minimising any areas of double counting. The risk model schema are shown as yellow shapes with areas which are not covered shown in white and areas of double counting shown as a darker yellow.

What these diagrams are designed to highlight is that if a direct mathematical approach is used (without full consideration of the underlying reality) then there is a significant risk that some aspects of the risk are not accounted and others are double counted.

Figure 1 - Direct Mathematical Modelling Approach

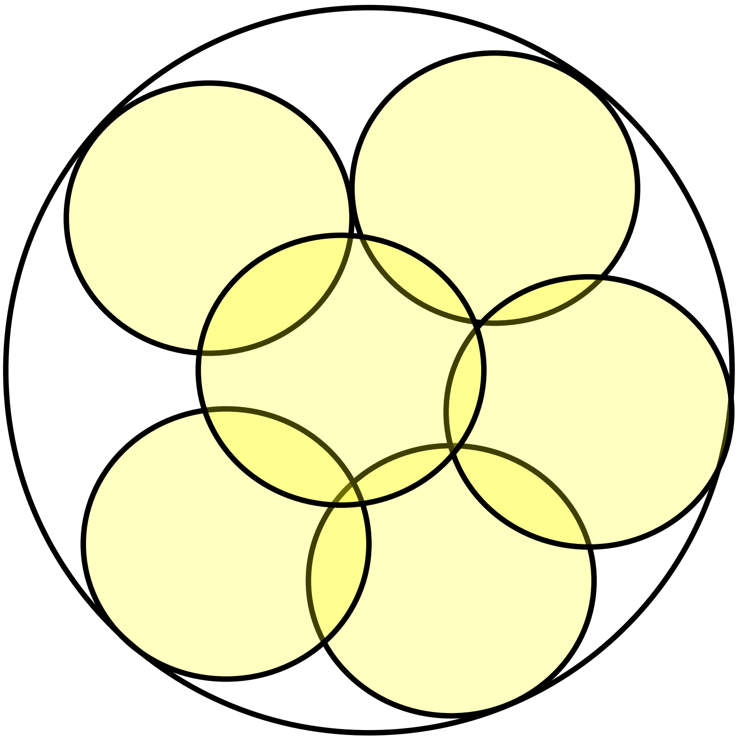

Figure 2 - Conceptual Modelling Approach

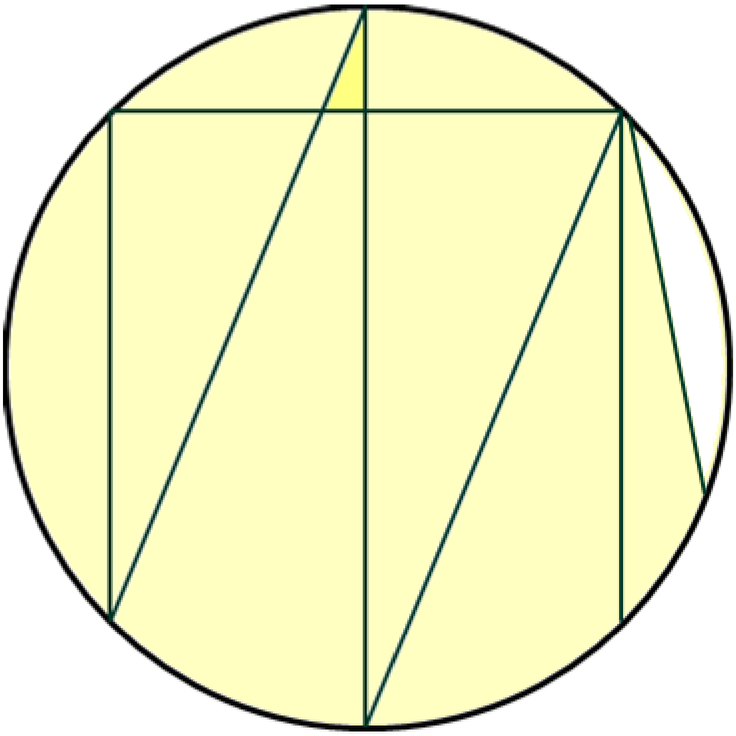

The mathematical modelling approach is shown as a set of circles to indicate that the models are likely to be robust and well developed; whereas the conceptual approach is shown as a jagged set of triangles to indicate that the modelling approach is likely to require more innovative approaches to cover the entirety of the risk. This is to indicate that whilst the conceptual approach may lead to a more holistic approach to longevity risk the modelling techniques may not be as robust.

With the conceptual approach the attempt is made to understand the fundamental nature of the risk and ensure that all aspects are accounted for and only once. Given the challenge of doing this it is unlikely to be a perfect process (as indicated by the stray area of brighter yellow in figure 2) but is likely to provide a better understanding of risk and lead to improved risk based decision making. The challenge then becomes constructing models that fit the understanding of the risk.

In summary, before one attempts to model a risk it is crucial to ensure that the model is being fitted to the risk rather than the risk being fitted to the model.

One final comment on this point – there is no one right answer. The nature of the specific portfolio may bias the exposure to longevity risk and lead to an appropriate conceptual foundation (and hence risk model) which differs significantly from other portfolios. As a result it should not be expected that different providers risk will develop a model for the risk in the same way.

Having said this, the framework described in the following section is intended to provide a suitably holistic basis for the individual provider to assess the longevity risk exposure and formulate an appropriate modelling approach. Therefore, it would be expected that any model of longevity would utilise many of the same concepts.

3. What is Risk?

Underpinning the conceptual framework is an understanding of risk that considers the differences in risk along the following lines:

- Uncertainty and volatility risks
- Uncertainty: parameter and model risks
- Systemic and specific risk

3.1. Uncertainty and Volatility Risks

As is described in a number of papers[4] longevity risk can be considered to be the result of two possibilities:

- the risk of getting the average wrong,
- the risk of getting the average right, but being unlucky.

In other words risk is a mixture of uncertainty and volatility. The latter might be considered as a stochastic process (and hence susceptible to models of that form) but the former is intrinsically harder to model.

The reason for this is that uncertainty to a great extent relates to the possibility of the future not merely being a continuation of the past – the potential for "black swan" events renders any assumption that this is the case suspect. So the form of the risk may change in the future and the relevance of past data to parameterize a model is questionable. In addition, past levels of risk variation may not be a relevant measure of future variation, and hence back testing provides limited validation of any particular modelling approach.

Over the short term, volatility tends to dominate a risk exposure but in the long term uncertainty is the key driver. As a result longevity is more of an uncertainty risk than a volatility risk.

3.2. Uncertainty: Parameter and Model Risks

Uncertainty can be considered to be the result of two elements: parameter and model risks. To explain the difference between the two the example of longevity trend uncertainty will be used. Parameter risk, in this example, is the uncertainty in the true level of mortality in the future and specifically relates to the uncertainty of the future. This risk is independent of the practitioner and their chosen approach. Model risk is the difficulty for the practitioner assessing the true level of mortality improvement and may be the result of the modelling approach or the information available for making an assessment. Model risk as a result is dependent on the practitioner and the model (and information) they utilise to assess mortality improvement.

Expanding from this example; uncertainty risks require an assessment of both the parameter and model risks for the risk to be captured correctly.

Typically when measuring risk the approach is to assess the parameter risk but many practitioners incorporate an allowance for model risk which depends upon the particular model the practitioner chooses to employ in estimating mortality rates, and the purpose for which the model is being

[4] Notably Sweeting (2007)

created. Model risk does not behave in a standard or definable way and for different practitioners may be more of a "trend" risk or more of a "level" risk, varying by the particular uncertainty risk. As such it would be difficult to outline an overarching approach to model risk which was independent of the model used or the nature of the longevity exposure. As a result it should be captured under each uncertainty risk specifically.

3.3.Systemic and Specific Risks

It might be ideal to have a conceptual framework for longevity risk that stands independent of exposure. However, nature of longevity risk is unique and specific to the provider and the approach taken to manage longevity risk influences the risk profile significantly such that some assumption has to be made about the approach.

A common approach to estimate mortality rates is to utilize mortality assumptions drawn from a reference population (such as England & Wales population) and apply modifications to account for the specific nature of the portfolio. This approach shall be assumed in designing this conceptual framework. As a result longevity risk will be considered to be exhibited from two sources:

- Reference population risk (or the systemic longevity risk)
- Portfolio risk (or the specific longevity risk)

This division is explained further in section 4.

4. Conceptual Framework

As outlined in the introduction the framework defines seven different behaviours of longevity risk:

- Trend Uncertainty
- Trend Volatility
- Catastrophe risk
- Basis risk
- Underwriting risk
- Mis-estimation risk
- Statistical Volatility

As described in section 3 there are considered to be two primary classifications of longevity risk: uncertainty/volatility risk and systemic/specific risk[5]. The behaviours above are shown against this split in Figure 3.

Figure 3 – Classification of Longevity Risk Behaviours

<table>
<tr><th></th><th>Uncertainty in setting the “right” assumptions</th><th>Volatility in experience relative to the “right” assumptions</th></tr>
<tr><td rowspan="2">Systemic (or population risks)</td><td>Trend Uncertainty</td><td>Trend Volatility</td></tr>
<tr><td colspan="2">Catastrophe Risk</td></tr>
<tr><td>Specific (or Portfolio risks)</td><td>Mis-Estimation Risk
Basis Risk
Underwriting Risk</td><td>Statistical Volatility</td></tr>
</table>

This division of risks should not be viewed as absolute; as indicated above catastrophe could be viewed as either an uncertainty risk or a volatility risk (discussed further in section 4.1.3). Furthermore, specific risks may arise ultimately from systemic risks, depending on how the longevity risk is managed by the provider. The important consideration is to ensure that the potential variation in longevity for the specific provider is captured conceptually for the risk model to be assessed against.

[5] As per the previous section uncertainty risks should be considered to incorporate both parameter and model risks.

4.1. Systemic[6]

As described above the systemic risk relates to potential variation in behavior of the reference population both in terms of the uncertainty in setting expectations about the future course of mortality improvement and the potential for actual future improvements to differ from expectations.

4.1.1. Trend Uncertainty Risk

Trend uncertainty is the risk relating to the ability to predict mortality rates in the future as mortality is influenced by a range of drivers, which are difficult to assess and appropriately allow for the complexity in the relationship between them, such as[7]:

- Development in medical treatments
- Lifestyle factors
- Economic condition
- Public policy

These issues lead to the challenges in constructing a projection model which appropriately accounts for the complexity in the drivers of mortality improvement and the difficulty in acquiring a sufficiently credible data set to calibrate a projection model.

The risk might manifest itself in a number of ways. For example, the long term level of improvement might be higher than estimated or the shape of improvement may be different than estimated.

This risk behavior is what is traditionally understood by longevity risk and whilst possibly the best understood is probably the hardest to model and calibrate relative to its materiality (for most providers).

As discussed in section 3.2 trend uncertainty incorporates elements of both parameter and model risk. Whilst the difficulty in assessing future improvement is well understood (or at least well discussed!) and is known to have significant ranges of outcome the model risk should not be underestimated. The projected level of improvement using different kinds of models can be significantly different. An example of this is shown in Appendix A.

4.1.2. Trend Volatility Risk

The risk considered is that the improvement rates experienced are different from those assumed for a period, after which improvement rates return to best estimate. The divergence in experience is caused by systemic variation, rather than random chance relating to individual lives. However, the underlying best estimate is correct and would be returned to as the cause of the systemic variation subsides. Examples of this are periods of harsh/mild winters or periods of economic recession or growth.

Two examples of scenarios that might occur are:

[6] It should be noted that it is assumed that the reference population is sufficiently large that there is negligible estimation error in the derived rates and hence the longevity risk solely pertains to future experience. If the reference population does not satisfy this condition then a further mis-estimation risk measure needs to be included as a systemic risk. For most general population level analyses the level of mis-estimation risk at material ages is arguably negligible. Moreover, the basis risk taken by applying assumptions derived from a reference populations to the specific portfolio are likely to dwarf any uncertainty in the reference population estimates.

[7] There have been a great number of papers written on this topic (One of the best sources for discussion on this topic remains Willets et al (2004)) and further description is not included in this paper.

1. Following the period of increased mortality improvement the improvement rate decreases to the best estimate and mortality rates are permanently lower relative to the best estimate mortality rates. The survival curve is shifted to the right
2. Following the period of increased mortality improvement the improvement rates fall below the best estimate until the mortality rates revert to the original projection. The survival curve is unchanged after the period of volatility.

Under the first scenario, the extra lives that have survived longer than expected over the period of volatility experience mortality rates under the best assumptions and an increased number of lives are alive for all future periods. They make a permanent gain in life expectancy.

Under the second scenario, the unexpectedly surviving lives following the period of lighter mortality have a higher force of mortality which suppresses the mortality improvement until the excess population reduces to the expected level. The increased life expectancy is temporary.

These two scenarios are shown in Figure 4.

Figure 4 – Variation from best estimate survival curve under the two scenarios

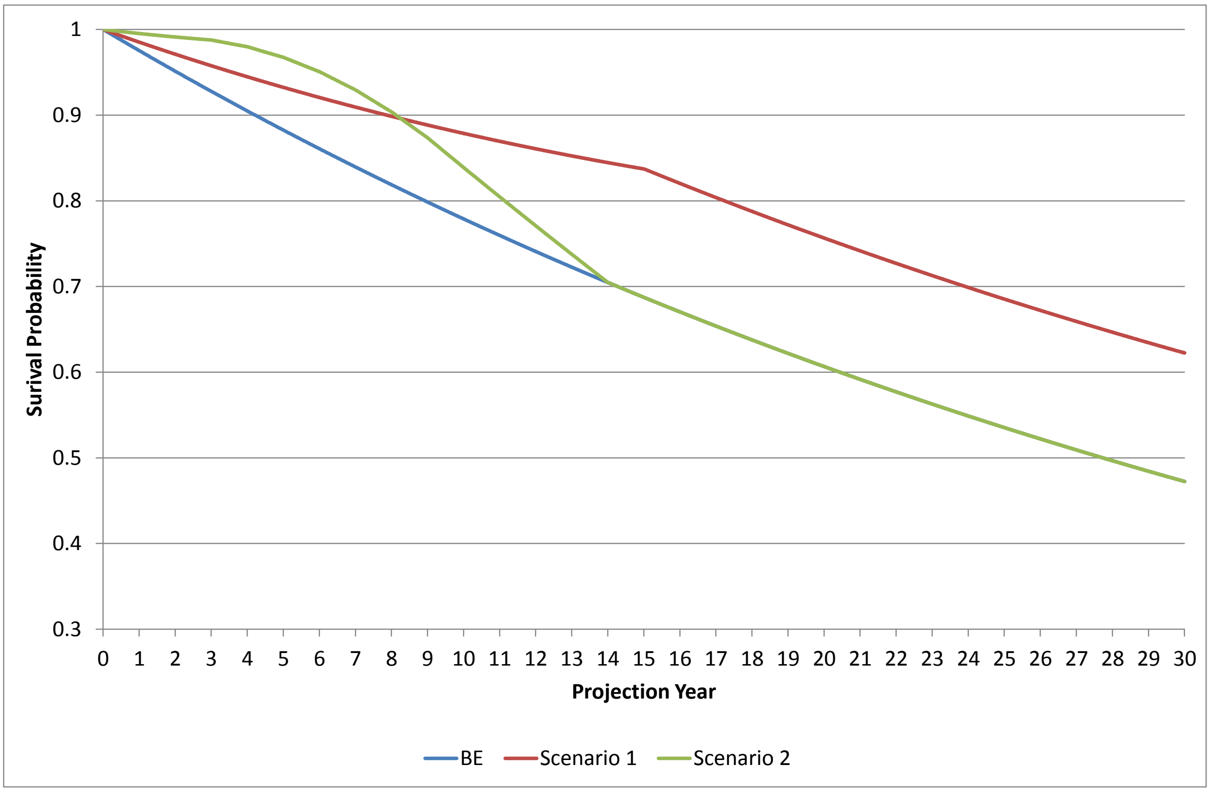

The key difference from trend uncertainty is that the underlying projection of improvement rates is "correct" but short term period effects which cause variation from the fundamental drivers in the trend in improvement rates.

4.1.3. Catastrophe Risk

Probably the most contentious behaviour proposed in this paper is that of "catastrophe" risk. If we consider the behaviour of "catastrophe" risk for mortality risk[8], a common view is that Longevity risk

[8] A sudden or widespread disaster leading to a significant number of non-independent deaths.

does not behave in this way, unlike mortality or equity risk where catastrophic events not only can be readily conceived but have happened in the recent past (such as flu pandemics).

However, it is proposed that there is the potential for "catastrophic shifts" in mortality rates – where mortality rates fall exceptionally rapidly owing to some singular change in a driver of mortality that takes effect far quicker than would be anticipated in a mortality trend projection model or as a consequence of systemic volatility.

One possible example of a catastrophic shift is the elimination (or substantive reduction) of mortality owing to a singular cause of death and a common scenario postulated is a cure for cancer[9]. The realism of this kind of behaviour is debatable but there is evidence that a "catastrophe" type event has happened in the past.

In the mid to late 80s there was considerable concern amongst insurance experts and medical professionals that an AIDS epidemic might occur in the western developed world (specifically UK and the US). The feared epidemic did not occur as a result of changes in sexual behaviour, education programs and the development of drugs which made the condition manageable if not quite a cure and the number of deaths has reduced dramatically.

Figure 5 shows the pattern of deaths in the 1990s. What can be seen is that the number of HIV and AIDs deaths was steadily increasing until 1995 and then over a 3/4 year period the number reduced by around 70%. Furthermore, the increasing number of HIV diagnoses indicates that the level of reduction may be higher than is first evident[10].

[9] It is noted that given the nature of cancer, a collection of many diseases, a singular cure is not considered realistic but a collection of strategies to address cancer mortality might be considered to have the same effect.

[10] However, it cannot be assumed that the number of deaths would have increased in line with the diagnosis as the availability of drugs to manage the disease is likely to have led to people to seek diagnosis and the true number of people with HIV prior to 1995 was higher than appeared, as there would be a perception that there is limited value in being diagnosed with an incurable disease. This would clearly change following the announcement that there are drugs which can make the condition manageable.

Figure 5 - HIV and AIDS diagnoses and deaths in HIV-infected people, UK, 1981 to 2005

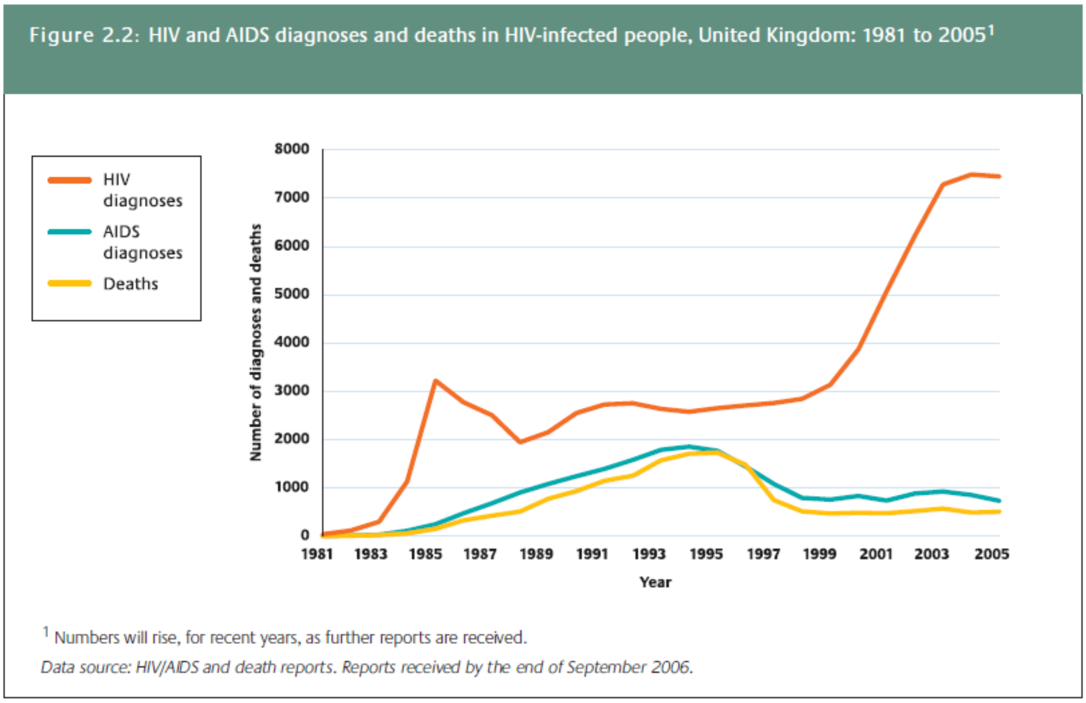

Source: Health Protection Agency

What can be drawn from this is that it is possible for a new regime of treatment to be implemented and take effect relatively quickly leading to significant reduction in mortality over a short period of time, and hence support the concept of a longevity catastrophe.

In addition to postulating the elimination of a major cause of death an alternative formulation can be considered regarding the elimination of a lifestyle factor that leads to multiple causes of deaths. For example, if a common behaviour that is currently considered relatively benign was found to have a serious impact upon mortality and was banned the effect on mortality might be in the form of a catastrophe event.

As suggested in section 4 catastrophe could either be considered an uncertainty risk (in as much as it is the difficulty in anticipating an event) or a volatility risk (in as much as the event leads to an affect which is a paradigm shift to the longer term progression).

4.1.4. Comparison of Systemic Shapes

Figure 6 provides examples of how the shapes of the three behaviours might compare[11].

Figure 6 - Comparison of Population Longevity Risk Behaviours

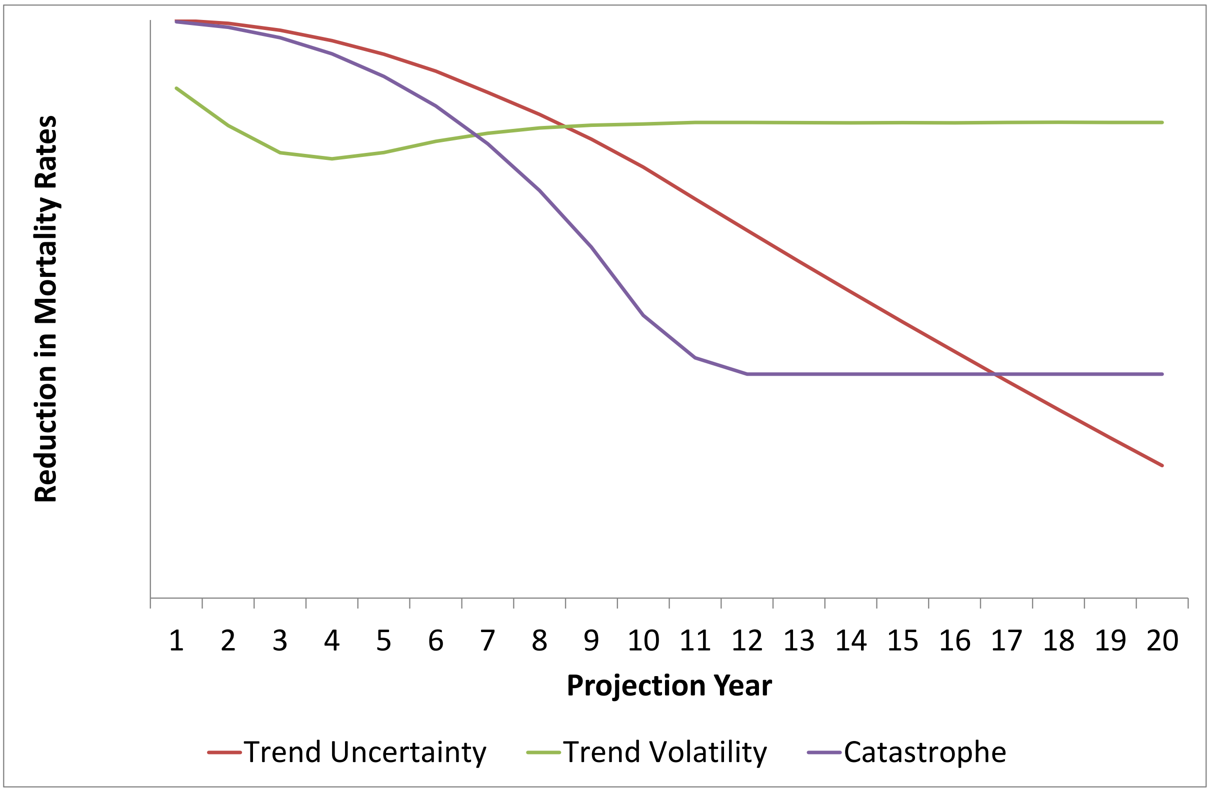

The chart indicates that the shape of these three behaviours might be quite different:

- Trend uncertainty – has relatively low impact initially but the cumulative effect is the significant risk in the long term. The reason being that understanding of the effect of the underlying drivers of mortality improvement is likely to be relatively measurable in the short term but over the long term the uncertainty increases
- Trend volatility - is significant at early durations but the effect is limited in the long term as the underlying best estimate is "correct"
- Catastrophe – has potentially the biggest impact over the medium term[12] but runs off as the catastrophe event concludes in its impact[13].

It is possible to view catastrophe as an extreme acceleration of expected (or unexpected) trend improvement. Moreover, these risks are not independent - if a catastrophe occurs the impact of significant unexpected mortality improvement is reduced (in absolute if not relative terms).

---

[11] These are representative of the differences in shape, not indications of the actual level of the risks which have been chosen to allow easy comparison of the shape. It should also be noted that no allowance for model risk has been included – so this chart solely accounts for parameter risk.

[12] Whilst a mortality catastrophe can occur with minimal warning a longevity catastrophe event would be detectable in advance to some extent. The difference between the two is that whilst the effect of a mortality catastrophe could be mitigated if known in advance (i.e. if an earthquake was known to occur at a particular time in the future the people potentially effected could be evacuated) the effect of the longevity catastrophe event could not be (unless the knowledge was in some way privileged).

[13] It is assumed that if a catastrophic event occurred it would be short but significant it terms of its impact. If it happened over a longer period it would be more appropriate to consider it as a trend uncertainty event.

However, it is important to capture the behaviours separately, rather than as one overall population risk model which might lead to a misstatement of the level of risk to providers with different exposures to longevity. For example, an impaired annuity provider may be exposed to a catastrophe more than a provider offering standard rates to healthy lives. If a treatment became available that eliminated the extra mortality risk for diabetics then an annuity provider with a concentration of diabetics would be more affected than a provider that wrote solely (or predominantly) to healthy lives.

4.2.Specific

The level of specific risk is a function of the approach to deriving mortality rates and the strengths or weaknesses of that approach in managing longevity risk. It is argued that the key to managing longevity risk is through the accumulation and use of pertinent information to determine the appropriate mortality rate.

There are conceptually three sources of information that can be used in determining the mortality assumptions for a particular life within a portfolio:

- Individual life information such as the lifestyle or medical conditions the life has
- Past mortality experience of the portfolio
- External mortality experience or related analysis from a reference population

Each of these approaches is attempting to set accurate mortality assumptions for the portfolio and manage the longevity risk by reducing the potential for lives to exhibit different mortality assumed by reducing the size of the group of lives covered by each assumption and limit heterogeneity. These approaches lead to longevity risks in the form of:

- Underwriting risk
- Mis-estimation risk
- Basis risk

However the approaches should lead to an overall reduction in the diversified longevity risk even if the undiversified risk could be higher[14]. These behaviours result from attempting to manage systemic risk or one of the other specific risks, examples of this latter scenario are:

- Underwriting risk results from using individual information to avoid the selection risk that occurs from purely using external mortality experience and so reduces the basis risk taken
- Mis-estimation risk results from using experience to override underwriting assumptions where the underwriting decisions appear to be out of line with the resulting experience and so reduces the underwriting risk taken
- Basis risk results from overriding the results of experience where external evidence suggests the past portfolio experience may no longer be as relevant in assessing the future and so reduces the mis-estimation risk taken

In some sense; these risks are the residual risks following the risk management of longevity.

[14] This is described further in section 5.1.

4.2.1. Basis Risk

Basis risk is the risk that mortality assumptions where derived from analysis of other lives (including past experience of the portfolio) do not reflect the future mortality experience of the particular life (or lives) under consideration. Examples of the causes of this risk:

- Assumptions based on a reference population (such as the general population), rather than relative to experience, do not reflect the mortality drivers of the portfolio population
- Within the current portfolio some lives may live longer and others shorter than assumed as a result of risk factors not allowed for in the assumptions. It can also be caused by over grouping of risk factors as the credibility would be insufficient to allow more granular analysis.
- New lives entering the portfolio may have materially different mortality resulting in selection against particular aspects of the basis. As in the previous bullet this risk may be caused by excluded risk factors and over grouping, but may additionally be caused by systemic bias exhibited in the past experience which becomes self-perpetuating.
- When past data is used to determine mortality rates there is the risk that prior experience is not relevant to current and future lives. This is described further in section 4.2.3.

How the risk behaves will depend on the particular mortality assumptions, the particular aspects which are derived from a reference population (and how these dove tail with any experience rating and underwriting), survivorship bias and selection.

The behaviour of this risk will be quite complex (not least because it will incorporate parameter and model risk elements) and will depend on the specific provider. As a result it could behave as either a "trend" or a "level" risk.

One specific example that is worth highlighting is the increasing usage of postcode rating[15]. Postcode rating reduces the longevity risk exposed to (if appropriately derived and applied) by setting assumptions that are more specific to the individual (as opposed to setting population level assumptions). There is a risk that the rating will not work appropriately and this should be captured.

Expanding upon this; the aim of postcode rating is twofold:

- To improve the accuracy of the mortality rate assumed
- To improve the certainty of the mortality rate assumed

This can be seen more clearly in Figures 7 & 8.[16] Figure 7 shows the spread in actual mortality rates by local authority (LA) and shows the potential uncertainty if population average was used (indicated by the arrows). If postcode rating is applied the accuracy and precision of the assumptions should improve (indicated by the arrows in Figure 8). Hence, overall the level of risk should be reduced.

[15] For further information see Richards (2008) and Madrigal et al (2011)

[16] The actual mortality rates are population mortality rates split by LA in England for males aged 70-74 over 2001-2010. The estimated mortality rates have been calculated using an LSOA level model of population mortality using the aggregate index of multiple deprivation for 2010. This has been chosen as it indicates the behavior one hopes to achieve from a postcode rating model.

Figure 7 – Spread of Actual Mortality Rates at the LA level

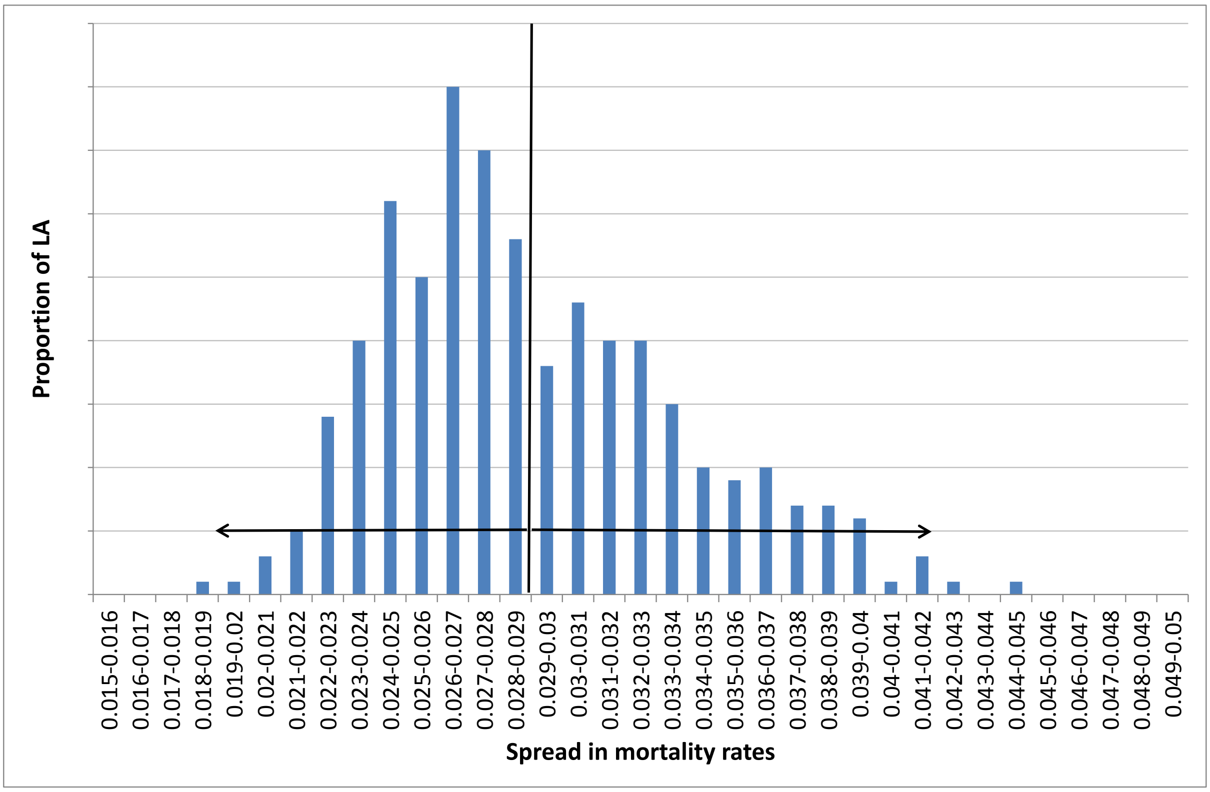

Figure 8 – Comparison of Actual and Estimated Mortality Rates at the LA level

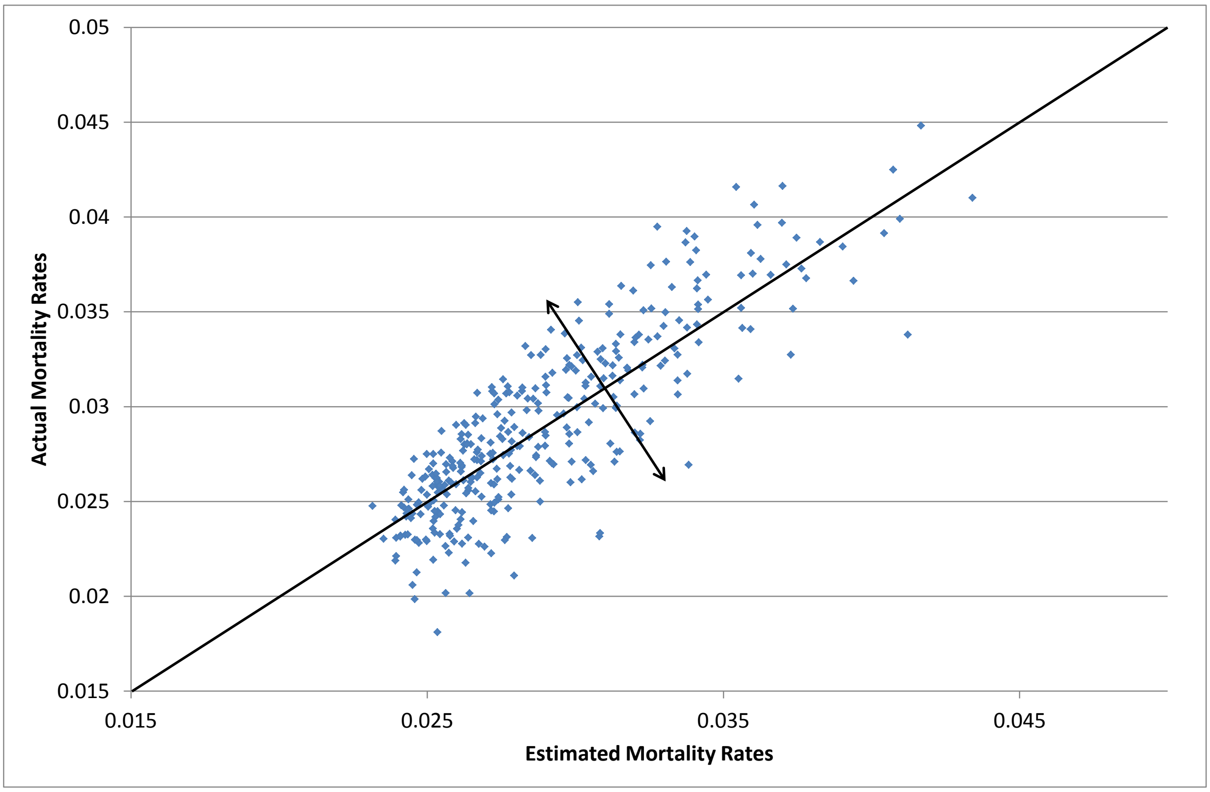

4.2.2. Underwriting Risk

The use of information sourced directly from individuals can be used to determine the appropriate mortality assumptions, the most well-known example is the use of medical information to price enhanced annuities[17]. There is a risk that the assessment of the information is an imperfect process leading to incorrect mortality assumptions.

The risk envisaged here is not the operational risk that individual underwriters allocate the wrong mortality adjustment to a specific life. Rather, the risk is that the guidance and/or underwriting tables provided to the underwriters are systemically in error.

The risk arises because of the difficulty in determining the impact that particular conditions and risk factors will have on mortality rates. This is an uncertainty risk and is not wholly susceptible to statistical analysis of past experience and needs to allow for trends in mortality and morbidity.

There are two primary causes of the risk:

- Adequate research is not carried out to justify the particular mortality assumption implied by a particular source of impairment
- The research is not interpreted correctly or fails to take into account other factors which effect the relevance or credibility of the research

To summarise, the nature of the risk is the uncertainty in the depth and accuracy in understanding when setting the underwriting assumptions.

The behaviour of the risk depends on the underwriting approach and derived underwriting measures that are used to produce mortality rate adjustments – hence is likely to be specific to the provider and the specific way in which they underwrite and determine the mortality adjustment to be applied.

This risk is only taken by providers who underwrite lives based on individual information – so non-underwriting providers are not exposed to this risk. However, they will be more reliant upon assumptions drawn from external population analysis (and their own experience) so would be expected to have a higher level of basis risk (and not have the advantage of any diversification benefit from utilising an increased number of approaches to manage longevity risk).

4.2.3. Mis-estimation Risk[18]

If portfolio experience is used to adjust the mortality basis then a risk results from the statistical error in the basis relating to the credibility of the data. Specifically, it looks at the risk that the estimated parameters of the basis are not correct.

Aside from the statistical uncertainty in the parameter estimates there is also the issue of relevance of the past experience in setting assumptions for future experience. Whilst the data used to determine the parameters should be restricted to only use the experience which is believed to be relevant there is a risk that there are unidentified biases in the experience which lead to the estimate being inappropriate for setting mortality assumptions for future experience. Hence, the use of past experience results in a basis risk as well as a mis-estimation risk.

---

[17] For further information see Ainslie (2000)

[18] Richards et al (2014) (2) provides an in depth study of estimating the effect of mis-estimation risk and contains further background on this risk so further description is not included in this paper.

Whilst there is likely to be some shaping to the assumptions derived by experience rating this risk is probably closest to behaving like a "level" risk[19].

4.2.4. Statistical Volatility Risk[20]

This is a relatively simple risk resulting from the random chance of different lives in the portfolio dying and the impact that would occur on reserves.

Two aspects occur under this risk:

- A lower than expected number of lives survive
- The deaths that occur have a lower average provision leaving lives with a higher average provision in the portfolio

The latter is typically understood as concentration risk.

4.3. Comparison

Table 1 shows a comparison with two other breakdowns of longevity risk; one published by the IAA[21] and one by Richards, Currie & Ritchie[22] (the latter has been deduced by the author).

Table 1 – Comparison of Different Risk Type Frameworks

<table>
<tr><th>IAA Risk Behaviours</th><th>Richards et al Risk Behaviours</th><th>Proposed Risk Behaviours</th></tr>
<tr><td>Volatility</td><td>Idiosyncratic</td><td>Statistical Volatility</td></tr>
<tr><td>Catastrophe</td><td rowspan="2">N/a</td><td rowspan="2">Catastrophe</td></tr>
<tr><td rowspan="6">Trend Uncertainty</td></tr>
<tr><td>Volatility</td><td rowspan="2">Trend Volatility</td></tr>
<tr><td rowspan="2">Trend</td></tr>
<tr><td rowspan="2">Trend Uncertainty</td></tr>
<tr><td>Model</td></tr>
<tr><td rowspan="2">Basis</td><td rowspan="2">Basis</td></tr>
<tr><td rowspan="3">Level Uncertainty</td></tr>
<tr><td>Mis-estimation</td><td>Mis-estimation</td></tr>
<tr><td>N/a</td><td>Underwriting</td></tr>
</table>

It can be seen that there are differences between the three models (notably the inclusion or exclusion of catastrophe and underwriting) – however the three different frameworks are largely consistent. Variation is mainly in the level of granularity of the conceptual breakdown.

As highlighted at the outset of the paper it is key to keep in mind that there is no right answer – what is important is that all sources and effects of longevity are accounted for in the risk model. The

[19] It should be noted that the choice of model the experience is assessed against will influence the results and hence there is additional model risk beyond the parameter risk determined by statistical analysis.
[20] Also referred to as "idiosyncratic" or "binomial" risk in other papers
[21] http://www.actuaries.org/LIBRARY/Papers/Global_Framework_Insurer_Solvency_Assessment-public.pdf
[22] Richards et al (2014) (1)

argument of this paper is that it is best to start from a conceptual understanding of the risk and then determine the appropriate modelling techniques.

4.4. Diversification

It is not the intention of this paper to outline the modelling method for these risks or consider the appropriate method of aggregation.

However, it should be noted that these risk behaviours are not 100% correlated and there is likely to be significant diversification benefit in ensuring that a provider's longevity risk exposure is not dominated by one particular risk behavior (and hence adopts a longevity risk management approach which utilizes multiple sources of information).

The precise level of diversification will depend on the approach to risk management and how interrelated the assumptions drawn in each behavior are. If the assumptions drawn from different sources interact (which might lead to an enhanced level of risk mitigation) then the behaviours are unlikely to be independent.

5. Variation in Exposure

The exposure to the different longevity risk behaviours will vary for different providers, depending on the business proposition and approach to managing the longevity risk. There are a number of reasons why the longevity risk will vary between providers but key variations are:

- Mortality rating approach
- Depth and breadth of experience

5.1. Mortality Rating Approach

The approach taken to rating lives can lead to different exposures to the different risk behaviours. The following are key differentiators:

- "Standard" provider; no account is taken of the mortality differences resulting health status or geographical location and the mortality basis is likely to be a modified base and trend.
- Postcode rating provider; account is taken of the geo-demographic variation as a proxy for health and socio-economic variation.
- Underwriting provider; account is taken of individual's specific circumstances.

Intuitively, the greater level of risk management (through postcode rating, underwriting or other risk factor rating approach) the lower the level of longevity risk the provider is taking (albeit at a cost which has to be balanced against the risk) assuming the approach is valid.

To provide an indication of the diversifying effect of employing multiple methods for managing longevity risk in pricing annuities the state of risk for an insurance company (Company A) at various points in its lifecycle are described below:

- Time 1 – Company A buys an annuity company (with existing liability and active in the market) which comes with no experience and limited policy holder information (DOB, gender, postcode, premium). As a result the assumptions are derived from external sources of information and so the specific risk is predominantly basis risk.
- Time 2 – Company A adopts an underwriting approach and for lives (past and new) medical information is available to base the mortality assumptions on. There is now the risk that the underwriting assessment is wrong but less reliance is placed upon the external assumptions and there is greater diversification benefit between longevity risks.
- Time 3 – The experience develops to the point where the underwriting assumptions can be credibility weighted against experience. This introduces a mis-estimation risk as a result of the finite credibility of the experience but reduces the reliance upon the underwriting assessment and the diversification benefit increases.
- Time 4 – Experience continues to grow and can now be used to credibility weight the externally sourced assumptions and so reduces the reliance on these but increases the mis-estimation risk further – and again the diversification benefit increases.

The potential effect on risk capital at each point in time is shown in Figure 9, both diversified and undiversified[23].

---

[23] The diversified risk capital estimates are calculated using a correlation matrix with the assumption that the risks are independent and the diversified risk capital is allocated according to the Euler method. This should not be taken as necessarily appropriate but has been used to allow assessment of the effect of diversification. For details on the Euler method see Tasche (2008).

Figure 9 – Comparison of Undiversified and Diversified Risk Capital

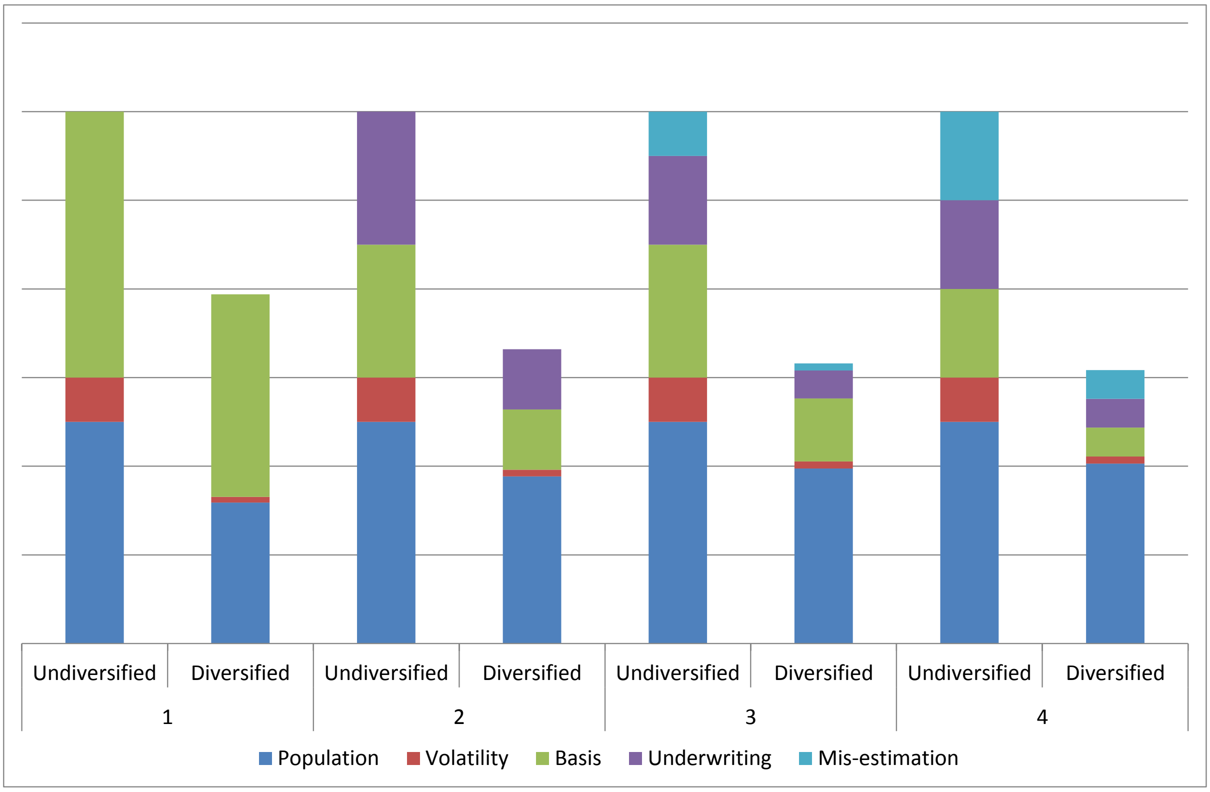

The undiversified risk capital is shown in Table 2 at each stage and is assumed not to vary over time so as to show the effect of diversifying the risk management approach more easily.

Table 2 – Undiversified Risk Capital

Risk	**Time 1**	**Time 2**	**Time 3**	**Time4**
Population	50	50	50	50
Volatility	10	10	10	10
Basis	60	30	30	20
Underwriting	0	30	20	20
Mis-estimation	0	0	10	20
Total	**120**	**120**	**120**	**120**

The movement in undiversified risk capital will vary by the effect of the risk management approaches on the best estimate reserves (and customer selection in response to the variation in risk management approach). If the best estimate is reduced then the undiversified risk capital might increase but the diversification effect would increase as population risks would make a smaller contribution to the undiversified risk total. Another example is shown in Appendix A where the undiversified risk capital increases.

The real impact of the diversification as indicated previously will depend on the assumed correlations and allocation method used.

5.2.Depth and Breadth of Experience

5.2.1. Deaths

Intuitively, the more experience (i.e. number of deaths) the greater reliance that can be placed upon portfolio experience in setting assumptions. This means less reliance is placed upon external information (reducing exposure to basis risk) and individual information (reducing exposure to underwriting risk).

However, it is likely that experience will be concentrated in specific areas of the basis rather than providing strong evidence for all assumptions. The "edges" of the basis are likely to remain heavily dependent on non-experience driven assumptions or based upon extrapolation from areas with greater credibility. This results in a basis risk as experience for one group of lives is being used to set assumptions for another group of lives.

As mentioned above, the relevance of past experience may be an issue as the experience of an older cohort may not be directly applicable to a younger cohort or there may have been changes in the portfolio mix (detectable or undetectable) that lead to the assumptions derived being inappropriate. As a result basis risk is incurred when experience rating a portfolio, and may be a material risk.

5.2.2. Lives

The size of the portfolio will affect the breakdown of exposure to the different types of longevity risk (all other things being equal). Figure 10 provides a picture of how the proportional exposure to different longevity risk behaviours might change with the size of the portfolio. The risks are split into portfolio volatility, population risks (trend uncertainty, trend volatility and catastrophe) and portfolio uncertainty risks (basis, medical underwriting and mis-estimation) for ease of understanding.

The proportions shown should not be taken as being accurate assessments of the relative mix of the risk types; they are not based upon analysis and are merely indicative of the behaviour to be expected. The significant points to be highlighted are:

- Portfolio volatility reduces exponentially as the number of lives increases and a relatively small annuity portfolio may be able to diversify the level of risk exposed to.
- Population risk increases as the portfolio becomes closer to a population level size and the chance of significantly different mortality experience reduces.
- Portfolio uncertainty risk is believed to increase as the volatility subsides but then reduces as the difficulty in identifying the mortality of the specific portfolio decreases (and population risk begins to dominate).

Figure 10 – Variation in Longevity Risk Exposure with the size of the portfolio

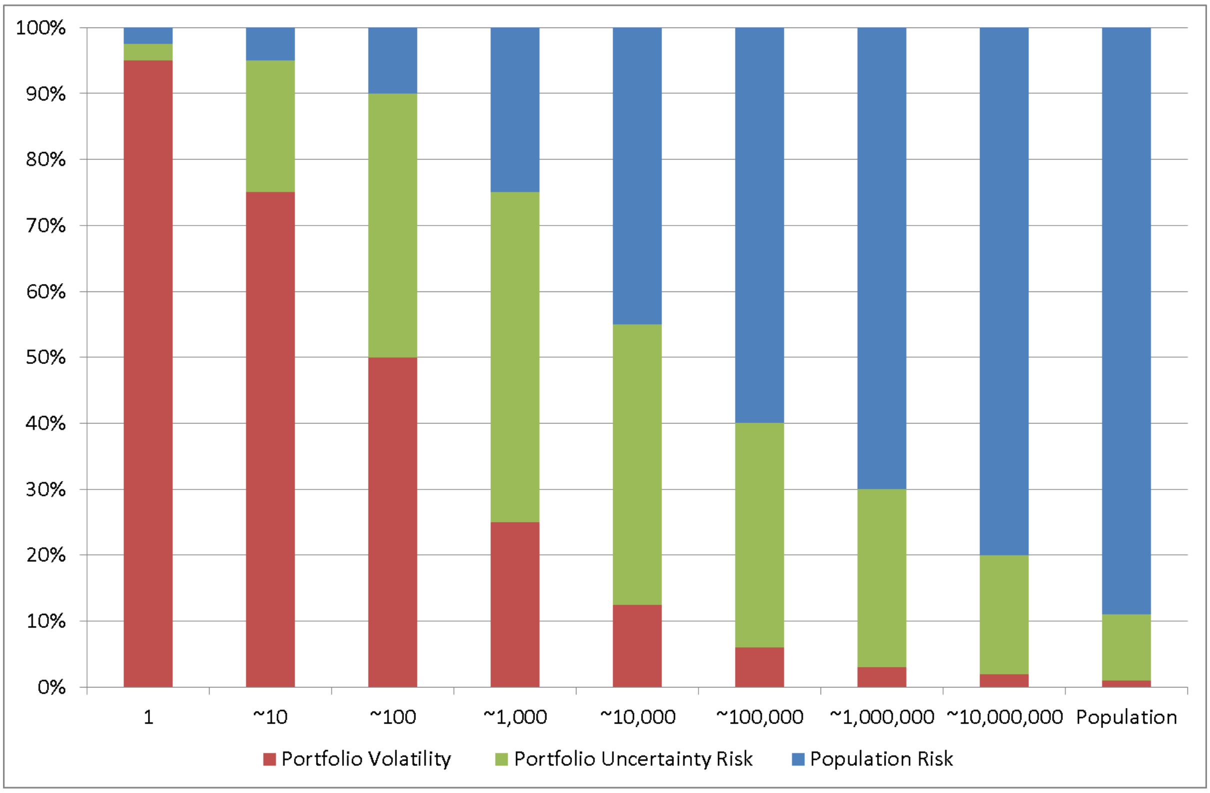

6. Final Thoughts

This paper has focused on the core elements that affect longevity risk and sought to outline how the risk behaves with respect to these elements. However, there are a number of additional aspects that need to be considered in assessing the risk exposure, notably the commercial environment.

Longevity exposure can be found in a number of different arrangements: occupational pension schemes, internal vesting life office pension schemes, IFA advised business, direct to consumer, etc. The differing sources of the business will influence the basis risk (as well as other behaviours to some extent); notably in the form of selection risk. Effectively the more choice an individual has exercised in the creation of a longevity risk exposure the greater the risk.

The changing nature of the annuity market, work patterns and pension saving mean that consideration needs to be given to how these effect the exposure to longevity risk, specifically the relevance of past experience and current methods of risk management for use in setting mortality assumptions for new lives exposed to.

However, it is intended that the framework outlined in this paper provides a basis for this consideration and, at the very least, stimulates a wider discussion as to the appropriate framework for considering longevity risk behavior.

7. Bibliography

Ainslie, R (2000) Annuity and Insurance Products for Impaired Lives. Staple Inn Actuarial Society, London.

IAA Insurer Solvency Assessment Working Party (2004) A Global Framework for Insurer Solvency Assessment.

Madrigal, AM, Mathews, FE, Patel, DD, Gaches, AT and Baxter, SD (2011) What longevity predictors should be allowed for when valuing pension scheme liabilities? British Actuarial Journal 16(1), 1-38.

Richards, SJ (2008) Applying survival model to pensioner mortality data. British actuarial Journal 14(2), 257-303.

Richards, SJ, Currie, ID and Ritchie, GP (2014) (1) A Value-at-risk framework for longevity trend risk. British Actuarial Journal 19(1), 116-139.

Richards, SJ, Currie, ID and Ritchie, GP (2014) (2) Mis-estimation risk: measurement and impact.

Sweeting, PJ (2007) Modelling and Managing Risk. British Actuarial Journal 13(3), 579-621.

Tasche, D (2008) Capital allocation to business units and sub-portfolios: the Euler principle. Pillar II in the New Basel Accord: The Challenge of Economic Capital, 423-453.

Willets et al (2004) Longevity in the 21st Century. British Actuarial Journal 10(4), 685-832.

8. Appendix A – Variations in projected mortality improvement

Figures 11 and 12 provide examples of the projected level of improvement for males and females aged 65 and 75 using:

- CMI_2012 (CMI)[24]
- A Mortality by Cause Model (MBC)[25]

What can be seen is that there are variations in the projected level of improvement which provide an indication of the potential level of variation that results from using different modelling approaches. The CMI and MBC models are quite similar in design so models that vary more significantly (such as a stochastic projection model) are likely to show greater variation – so the level of model risk is likely to be greater than these charts imply.

Figure 11 – Projected reduction factors for lives aged 65 in 2014

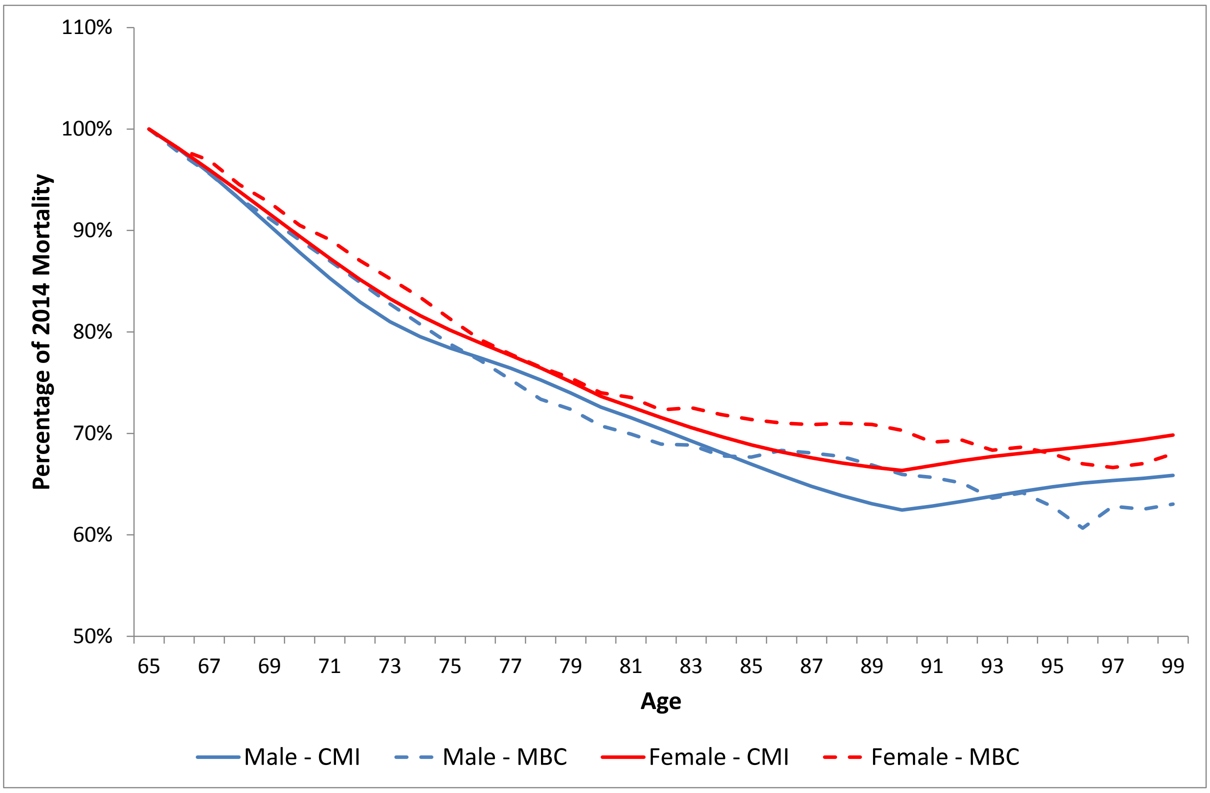

[24] Core parameters with long term rates of 1.75% and 1.5% for males and females respectively.

[25] These projections are taken from a mortality by cause model used by the author. No derivation is provided as these results are purely intended to be indicative rather than considered to be viewed as definitively showing the relative position of a mortality by cause approach relative to the CMI projections model.

Figure 12 – Projected reduction factors for lives aged 75 in 2014

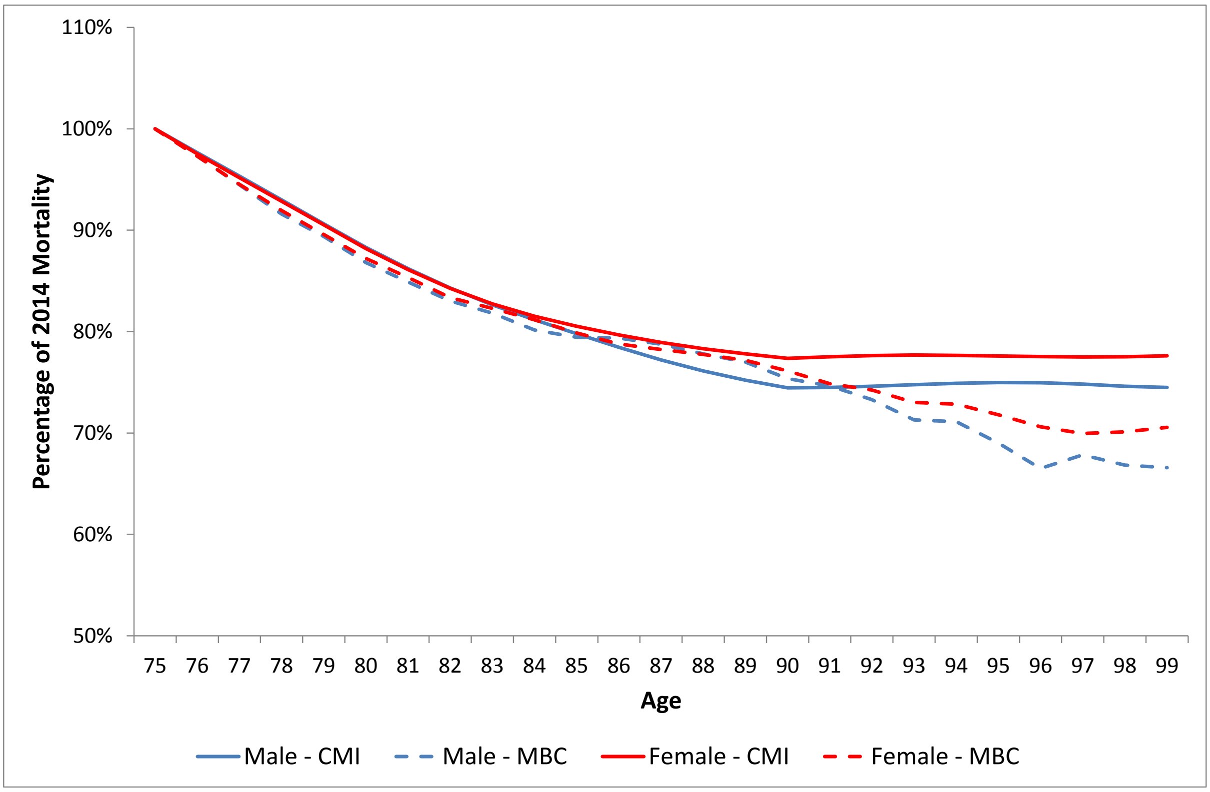

9. Appendix B – Second Risk Capital Variations Example

Table 3 shows the undiversified risk capital in this example.

Table 3 - Undiversified Risk Capital

Risk	Time 1	Time 2	Time 3	Time4
Population	50	50	50	50
Volatility	10	10	10	10
Basis	60	30	30	20
Underwriting	0	40	30	30
Mis-estimation	0	0	20	40
Total	**120**	**130**	**140**	**150**

The variation resulting from this set of undiversified values is shown in Figure 13.

Figure 13 - Comparison of Undiversified and Diversified Risk Capital