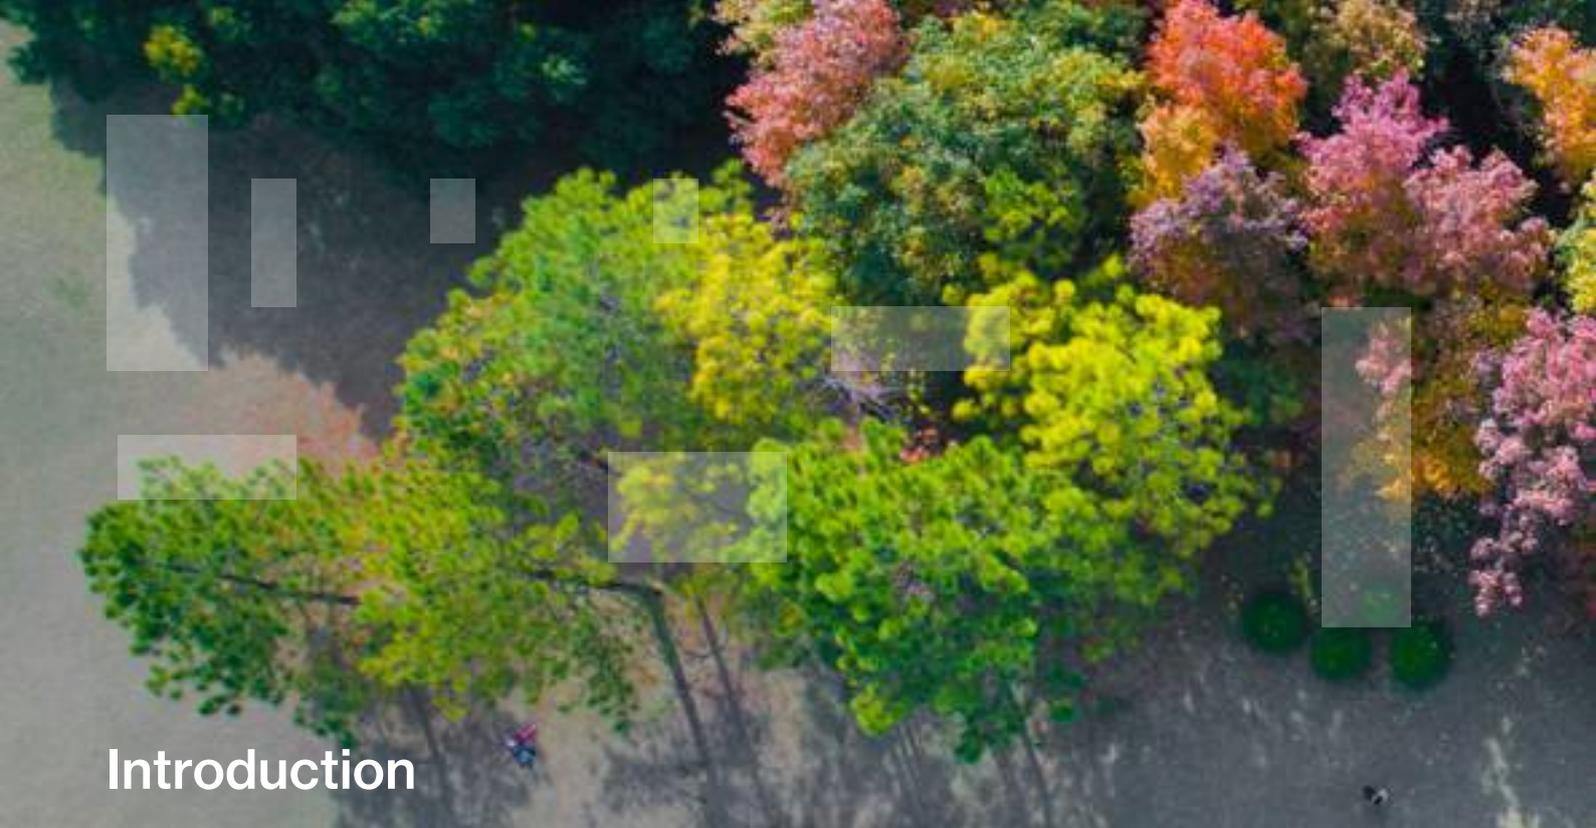


A new framework  
for future mortality



# Introduction

For insurers writing long-term protection products or annuities, the future evolution of mortality rates is of great importance, both for use in the pricing of new policies and reserving for existing policies.

However, compared to other important assumptions (such as yields on assets available to the insurer) which might be observable at the time of writing the business, changes in mortality over time are far more subjective. There are many and varied reasonable interpretations of even the effects of those events relevant to future mortality which we can observe in currently available data; how much wider then is the range of potential assumptions when considering those which we cannot observe?

In this article we examine some of the challenges arising in setting mortality improvement assumptions and propose a new framework. This is built around the different characteristics of short, medium and long-term outlooks, and puts a greater weight on forward-looking (rather than extrapolative) approaches.

The benefits of this new approach – other than what should be more accurate assumptions – include:

- an improved understanding of the main drivers of mortality variation at different stages of the future,
- a clearer process for quantifying the components of future mortality variation, facilitating discussion with other stakeholders inside and outside the organisation, and
- greater consistency with a firm's Internal Model (where applicable).

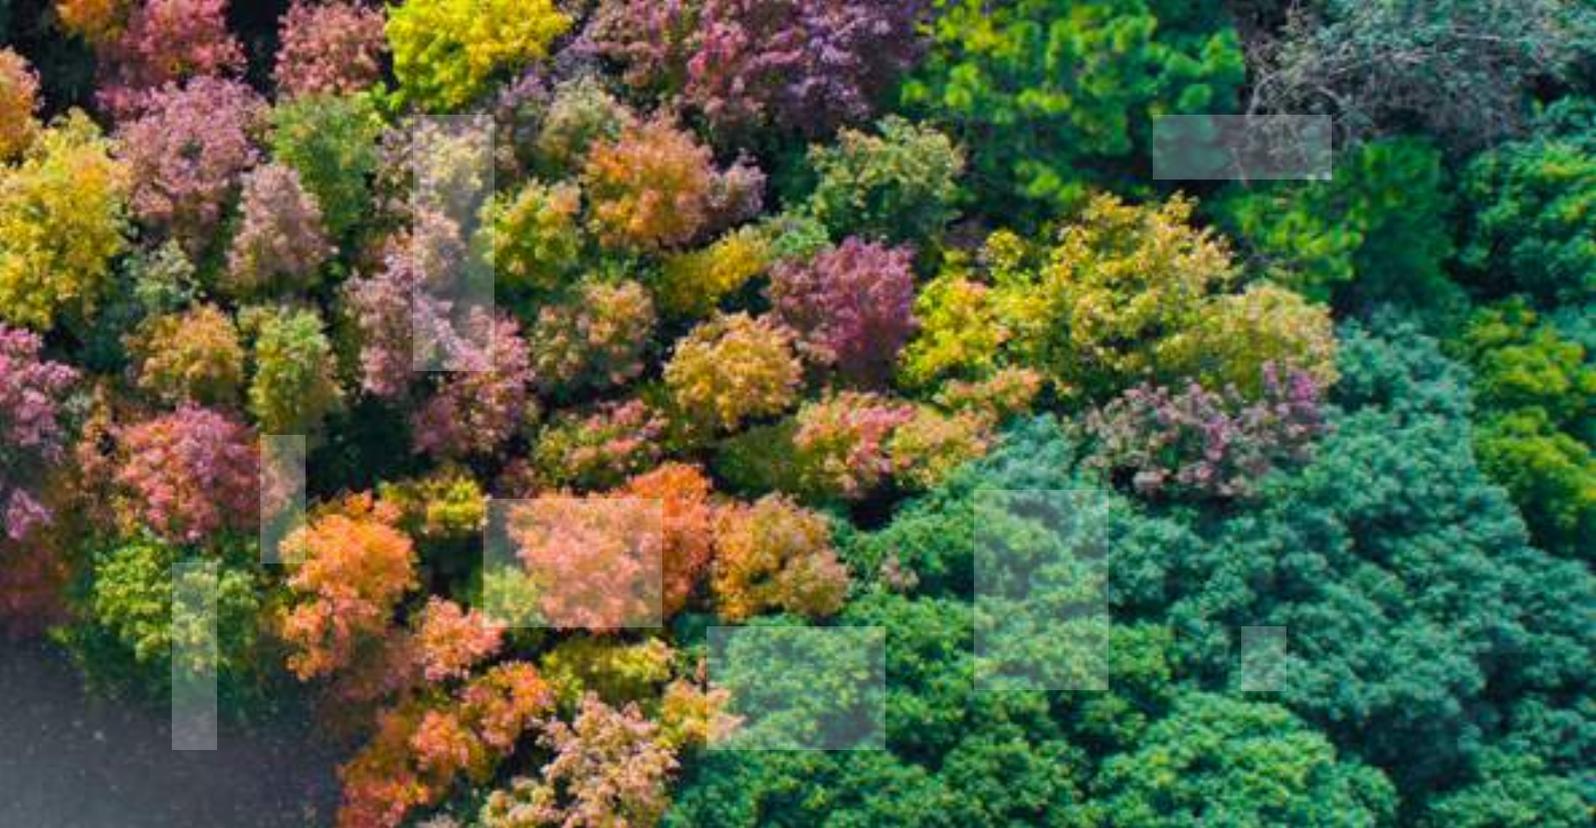
## Background

In an age where the UK's lingua franca for mortality improvements is the CMI's series of projection models, it's commonplace for the decision making around improvement assumptions to effectively be reduced to two or three questions.

- Which CMI model should be used?
- What is the appropriate long-term rate of improvement?
- What value should be used for the period smoothing parameter?

Changes in mortality improvement assumptions can be very material, particularly for insurers holding and writing significant volumes of (individual or bulk) annuities. In spite of this, to varying degrees depending on the insurer concerned, the assumption setting process may fail to take into account all available information.

Whilst there are clearly different levels of complexity in the modelling approaches adopted by different insurers, in some cases the decision-making process can be effectively decomposed into deciding whether to adopt the next in the series of CMI models each year and then deciding on one or two key parameters, perhaps by benchmarking against peers' assumptions from the previous year. It is fair to say that there is a significant amount of herding, with companies not wishing their assumptions to be out-of-line with those of their peers.



Moreover, in a case of “the tail wagging the dog”, these assumptions could be used by insurers to manage the annual changes in their balance sheets. In spite of the supposedly best-estimate nature of assumptions under the Solvency II regime, the decision as to how much to release from (or hold back within) technical provisions can influence the setting of these assumptions, rather than the required change in technical provisions being decided by the updated assumptions. This is particularly the case in light of the recent slowdown in mortality improvements.

There are a number of key questions which may not be appropriately addressed by the assumption setting framework within individual insurers.

- Ignoring the choice of model for improvements, what would we expect improvements to look like in the short- to medium-term?
- Does the current (or any) CMI model actually reflect this view of improvements?
- Assuming that the answer to this is yes, then upon the release of a new CMI projection model, should this model be adopted?
- Whether or not a new CMI projection model has been released/adopted, are the parameters which were used in a previous exercise still appropriate?
- How can the process of making these decisions be effectively governed, reducing the risk of (a) herd mentality, based on benchmarking of assumptions, and (b) choosing assumptions purely to control the impact on the balance sheet, contrary to the realism required of Solvency II bases?

These questions are, to a large extent, the same questions which insurers would typically set out to answer as part of their model and parameter risk assessments in an Internal Model of longevity trend risk. Interpreting these questions within the best-estimate basis setting process, what we find is that generating a realistic improvement assumption requires us to approach improvements from first principles, not allowing ourselves to be constrained by the popular models of the day.

Of course, there are some insurers and reinsurers with improvement assumptions set by more sophisticated methods; some or all of the questions above may not apply to these firms. However, even for those who don't follow the pack, there may be commercial and regulatory pressure to do so.

In this article, we discuss:

- what we mean by herding and “group think” and why these might be problematic;
- differences between current approaches and an ideal approach to mortality improvement assumption setting;
- considerations in developing such an approach; and
- the implications of implementing such an approach for insurers with significant longevity risk.

## Herding and “group think”

Whenever different people or institutions, be they from a scientific, political or industrial background, each come up with a point of view on a given topic, there may be commonality in the positions that they take.

Where this is as a result of each contributor going through independent thought processes, investigations and critical analysis of the issues from first principles, such consensus between the different parties can be a strong indicator of the reasonableness of the outcomes of their respective studies.

However, where the views of some of the contributors are formed on the basis of those views contributed by others, this can result in “group think”. This can be the case for insurers’ views of mortality improvement assumptions. For an insurer looking at benchmarking data, they may judge that most or all of the contributors have well thought-out rationales for their choice of model or choice of parameters and this may therefore anchor their own thinking about improvements for the coming year. In practice, each contributor may be giving too much weight to the prior year’s benchmarking when determining their own assumptions and then being comforted by the convergence of opinions across the industry. The effect of the challenging of outliers relative to any benchmarking information by either regulators or auditors can also have the effect of pressuring insurers to move their improvement assumptions towards the median industry view, at least for reserving purposes. Another less obvious form of group think is the passive acceptance of default assumptions within the CMI model, which we will discuss later.

There are plenty of good reasons why insurers might have similar improvement assumptions:

- Evaluation of the same set of systemic drivers of mortality resulting in similar views of the effects of those drivers;
- Reliance upon the same set of available scientific data and evidence of pan-industry experience (e.g. the CMI’s analysis of the self-administered pension scheme (SAPS) dataset) when determining assumptions;
- Similarities between the groups of policyholders insured by each insurer, meaning limited idiosyncratic variation to account for; and
- Technical support provided by a small number of reinsurers (who may themselves have similar views on improvements).

However, where convergence of best-estimate assumptions derives from group think and herding, there can be significant negative effects:

- There is inconsistency between insurers’ views of improvements (e.g. as espoused in their Internal Models of longevity trend risk, particularly in respect of new information) and their best-estimate models. More pointedly, there is a disconnect between how an insurer depicts (in their Internal Model) the potential changes to their basis in the following year and how they will actually set their basis. This may lead to decisions on reinsurance or other forms of risk transfer being made using unrealistic risk metrics.
- There is an effective outsourcing of responsibility for setting improvement assumptions, contrary to the requirements of Solvency II that insurers have a clear understanding of the risks that they run.
- There is the potential for an audit of the improvement assumptions to turn into a simple benchmarking exercise, rather than an assessment of whether the insurer has arrived at the assumption through appropriate analysis of those factors driving improvements. Such a tick-box approach does nothing to support investor confidence in the insurer’s assessment of their future liabilities, nor to add credibility to the audit process at a time when auditors are being criticised for their failure to identify problems in firms’ accounts prior to their collapse into administration.

Moreover, in the guidance accompanying the “Impartiality” principle of the revised Actuaries’ Code, published by the Institute and Faculty of Actuaries and entering into force on 18 May 2019, “group think” is specifically indicated as a threat to impartiality, meaning that such a decision-making process might potentially be in breach of this principle<sup>1</sup>.

It is worth noting, however, that the Solvency II Delegated Acts<sup>2</sup> (Article 10) require valuation methods to reflect market prices where these are available (implying a requirement for market-consistent assumptions). This would be an argument in favour of using benchmarking in the assumption setting process were there a deep liquid market in an instrument whose market value directly reflected mortality improvement assumptions.



Whilst the risk premiums offered by a panel of reinsurers during a treaty negotiation might give the insurer some information about the possible “market value” of improvements – the improvement rates at which reinsurers are willing to lock into an agreement to take on liabilities written by the direct insurer – this information is at least partially obscured (by expense loadings and profit margins) and the views are generally only obtainable from a small panel of reinsurers and at infrequent intervals, so that this does not represent a deep liquid market.

Benchmarking of assumptions across insurers also fails the test, since the assumptions are not tradeable – companies providing assumptions (for the calculation of their best-estimate liabilities) to a benchmarking provider are not offering to engage in market trades on the basis of those assumptions, so the benchmarking observations do not constitute a liquid market.

Finally, even if a deep liquid market were to develop, based upon which insurers could set their improvement assumptions for the purpose of determining technical provisions under Solvency II, there would still be a necessity to determine an improvements basis for pricing purposes.

To meet the additional requirements of Solvency II regarding an insurer’s understanding of its risks, an own view of improvements is essential.

## Is herding of improvement assumptions really occurring?

In 2018, we carried out a survey of mortality improvement assumptions. Survey participants contributed the assumptions that they had used for their end-2017 valuation and those which they intended to use at end-2018.

We saw a number of features suggestive of herding both for model selection and parameter setting:

- Almost all participants were planning to move to the next year’s CMI model instance for their end-2018 assumptions (e.g. CMI\_2017 if they had used CMI\_2016 at end-2017).
- The long-term rates of improvement were more tightly clustered for end-2018 than for end-2017 (50% of respondents had long-term assumptions between 1.5% and 1.75% for males at both year-ends, but for females this range narrowed from 1.25% to 1.75% at end-2017 to 1.5% to 1.7% at end-2018).
- For the period smoothing parameter (or  $S_k$ ), almost half of the participants had chosen to retain the default parameter specified by the CMI and most of the others had chosen to set  $S_k$  to between 7.5 and 8. It might be argued that this is to reflect socioeconomic differences between an annuitant population and the general population, but it also happens to have the (desirable) effect of reducing future volatility in reserves, as was noted in the Staple Inn Actuarial Society (SIAS) presentation, “Mortality improvements in the next decade”, April 2017.
- Other than tapering assumptions, very few firms were making any adjustments to the core model parameters in the CMI model.

Of course, whilst the closer clustering of assumptions from one year to the next may be suggestive of herding, it is entirely possible that some (or all) of the participants independently reached the same or similar conclusions based on their own assessments of the systemic underlying drivers of mortality.

<sup>1</sup>[www.actuaries.org.uk/upholding-standards/standards-and-guidance/actuarial-code/actuarial-code-principle-3-impartiality/section-5-principle-3-impartiality#Institutional\\_bias\\_or\\_Group\\_Think](http://www.actuaries.org.uk/upholding-standards/standards-and-guidance/actuarial-code/actuarial-code-principle-3-impartiality/section-5-principle-3-impartiality#Institutional_bias_or_Group_Think)

<sup>2</sup>Commission Delegated Regulation (EU) 2015/3, (<https://eur-lex.europa.eu/>)

# Frameworks for the setting of improvement assumptions

Figure 1 shows what a typical decision-making process for the setting of mortality improvement assumptions might involve at present (at its simplest).

## The focus of the process is largely on maintaining:

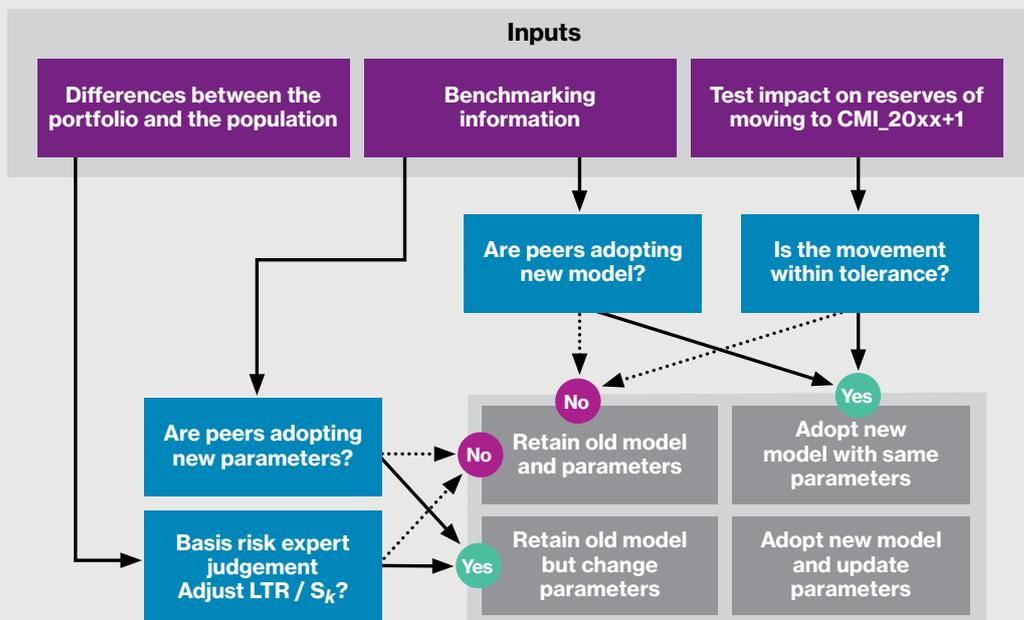
- Stability of liabilities over time
- Alignment with the consensus amongst other insurers (even if this is the result of herding)
- Avoiding being singled out by either the regulator or auditors as an outlier

It also makes an a priori assertion that the only reasonable options are future improvements from the CMI model currently being used or one of the updates of that model. The choice of parameters in the model is largely informed by benchmarking and a desire to avoid balance sheet volatility.

Some insurers may attempt to validate their choice of basis above by considering whether it is consistent with their views on a number of key drivers of mortality. For example, they may question whether the improvements are consistent with their expectations of future smoking cessation or increases in obesity. However, unless they believe that there is some clear inconsistency, it is likely that their basis will not change.

Such considerations as the stability of liabilities over time and whether there is an industry consensus are clearly useful. However, as we'll argue throughout this paper, these considerations should feature at a later point within the process of assumption setting.

Figure 1. Typical decision-making process



## Different time horizon, different approach

If we wish to evolve a more sophisticated approach to improvement modelling, taking into consideration a much greater proportion of the data available, we need to invert the thought process underlying the approach set out in Figure 1. Rather than consciously limiting any volatility in the balance sheet due to a change in improvements, and rather than selecting our model and parameters to be “amongst the pack”, we need to go back to first principles: focusing initially on the improvements which we might expect to see and only then working out which model and parameters we should use (and whether we should make any explicit adjustments to reduce volatility in the balance sheet).

This is no insignificant task. The drivers of improvements will almost certainly change over time. Also, our ability to judge how a driver will develop over time will vary between drivers. For some, (e.g. trends in the prevalence of smoking or alcohol consumption) a long-term view may be plausible, whereas for others (e.g. changes in climate) it may be difficult even to assert with any confidence a short-term best-estimate scenario.

It is convenient, therefore, to divide the future into three broad time periods and to approach the estimation of improvements over each of these differently when modelling an overall view of future improvements (see *Figure 2*).

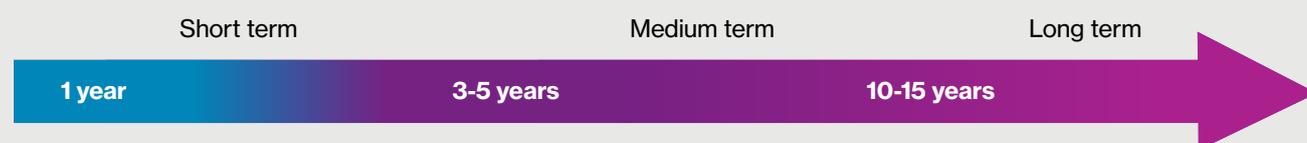
There is no precise rule defining where these three periods start or finish, and indeed some blending between them can help to reduce discontinuities. However, some guiding principles can help us to assess where it is reasonable to start and end any such blending of views:

- The short term should be the timeframe over which we have a good level of confidence in our assessments of how future mortality will vary (in the absence of

significant one-off events such as a very severe winter influenza season). It should be the time horizon over which we believe that there is continuity in the effects of key drivers of mortality and the trends in those drivers. It might be the case that this is only the next one to three years, or perhaps it could be up to five years into the future.

- The medium term is the period over which we can make reasonably accurate judgements in respect of a range of drivers of mortality (albeit there will be increasing uncertainty over time). This could be a period over which we might expect current information about the pipeline of pharmaceuticals or the licensing of (and coming off-patent of) drugs to have an impact on mortality experience. It could be the period over which current information about developments in technology to support lifestyle changes (for better or worse) might be expected to have an effect. It might also be the projection horizon for academic studies looking into, for instance, obesity rates in the population. This might feasibly be between 10 and 15 years into the future.
- The long term is, therefore, the period over which any assessment of the rates of improvement to be experienced in the population will necessarily be a matter of expert judgement, perhaps regarding more fundamental drivers such as overall economic growth, or per-capita health expenditure growth. No current information will likely be relevant to that period or have a predictable effect by that time. This approach may be similar to how some companies set their long-term rate of improvement for use in the CMI model (e.g. in the absence of benchmarking information). However, the interpretation of “long-term” here need not be consistent with that period in the future (in the CMI model with core parameters) in which the convergence of period improvements to the long-term rate has already taken place.

Figure 2. **Time horizons for improvement modelling**



It's worth considering what the relative materiality of improvement rates might be within the different time horizons. Taking as a base position the expectations of life for males aged 65 and 75 on 1 January 2019 with mortality as per 100% of PMA08 with CMI\_2017\_M[1.5%] (Core) improvements, we have calculated the increases in expectation of life due to a flat uplift of 0.5% or 1% across all ages (tapering the uplift to 0% linearly between ages 85 and 110) for each of three ranges of years of improvements (years 1 to 4, 5 to 14 and 15 to 30, where year 1 corresponds to those improvements required to derive 2019 mortality from that for 2018).

Perhaps unsurprisingly, the early years have the greatest impact per year of uplift (naively 0.31% per year for four years of uplifts of 1% for a male aged 65) – higher improvements in these years (unless offset by lower improvements later on) will affect all future years' mortality rates. However, in terms of the absolute impact of the short-, medium- and long-term improvements, the medium-term improvements are the most material, followed by the long-term improvements. Extending the long-term uplift to year 40 rather than year 30 only increases the 1.50% in the final row to 1.55%, so the cut-off did not materially affect the result. For deferred annuitants the weighting will shift from the medium term to the long term (increasingly so the younger they are).

With a proper understanding of the time horizons over which we're modelling, what characterises them and their relative materiality, it's possible to start constructing an overall view of mortality improvements.

## Moving towards an 'ideal' framework for the setting of improvement assumptions

As was intimated in the previous section, current processes for the setting of improvement assumptions may not be taking into account all available information on factors relevant to mortality improvements. This could mean that the link between the modelled improvements and the actuaries' and experts' understanding of why improvements are occurring is unclear; at worst, the views could be in conflict. Incorporating data on a sufficiently wide range of drivers of mortality mitigates against our choice of an improvements model being ill-informed, or a choice born out of convenience.

A more sophisticated alternative process, in which improvements are developed from the bottom up, starting with a detailed understanding of a wide range of drivers of mortality is shown in *Figure 4*.

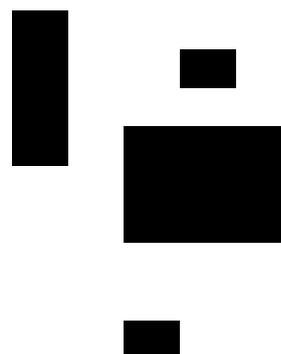


Figure 3. Relative importance of different periods of improvements

Increase years			Base EoL		% increase	
From	To	Increase.	EoL (M65)	EoL (M75)	EoL (M65)	EoL (M75)
N/A	N/A	0.0%	22.165	13.657	–	–
1	4	0.5%	22.303	13.764	0.62%	0.78%
1	4	1.0%	22.441	13.870	1.24%	1.56%
5	14	0.5%	22.435	13.810	1.22%	1.12%
5	14	1.0%	22.706	13.965	2.44%	2.25%
15	30	0.5%	22.330	13.686	0.74%	0.21%
15	30	1.0%	22.498	13.716	1.50%	0.43%

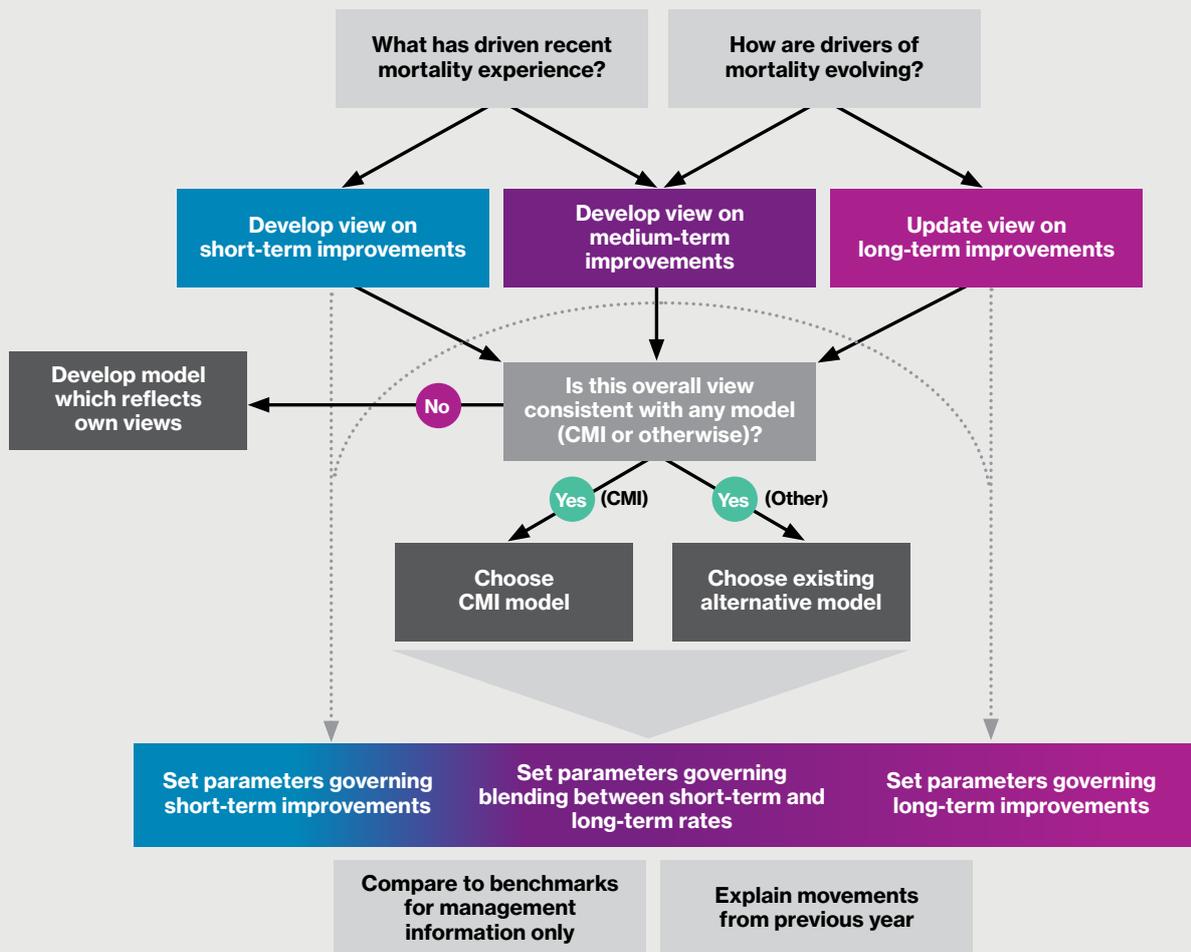
The key differences between this approach and what currently takes place in a typical UK insurer include:

- The focus of the exercise is on understanding why mortality is improving (or otherwise) and only then converting these views into a choice of model and a set of parameters which reflect this understanding.
- There is no a priori choice of model in the process – it is not assumed at the outset that a CMI model will be chosen and that the only questions are the long-term rate assumption and the short-term parameters.
- Even if the process alights upon the use of the CMI model, it does not assume that core assumptions are appropriate; rather, it requires the company to develop its own views for most of the material parameters, informed by its own understanding of the drivers of those improvements.

- The benchmarking information does not form part of the decision-making process, but instead is used to explain how internal views differ from those of other companies.
- The process attempts to determine a genuine view of realistic future rates of improvement; it doesn't consider balance sheet stability to be more important than realism.

It's worth noting that the processes described in both *Figure 1* and *Figure 4* are only the processes by which a recommendation for an assumption is arrived at. The final decision will generally be taken by the relevant governing body, which will be entitled to reach a different decision. However, one may expect that if detailed analysis underpins the assumption and movements year-on-year are appropriately explained, this will strongly support the recommendation.

Figure 4. **Ideal decision-making process**



# The role of the CMI projection models

It is important to make it clear that we do not suggest that the problem lies with the CMI projection models themselves. There are plenty of good reasons for choosing to use one of these models to project mortality improvements, not least to take advantage of the significant research and development investment made by the CMI, the combined wisdom of those actuaries with specialisms relating to mortality/longevity who fed into the design of the model, and of course the fact that it is a “common currency” by which actuaries can communicate improvement assumptions easily.

An under-appreciated aspect of the CMI projection models is their flexibility. There are “core” parameters, but the actuary is not obliged to use these parameters. Indeed, by changing the wide range of parameters available, an actuary can produce myriad different improvement projections with varying initial improvements (and “direction of travel” for CMI\_2016 onwards), different degrees of smoothing and persistence of cohort effects, and different shapes and speeds of convergence to a long-term improvement assumption which can vary by age.

It is regarded by many as a “default” position to use the core parameters. Indeed, it is too easy for actuaries to adopt core parameters (often simply because they are defaults) and then to rely upon benchmarking for the one parameter left unspecified by the CMI. That actuaries choose to do so could be a result of the many demands on their time, or perhaps due to senior level push-back against more sophisticated approaches (“if other firms are using core parameters, why aren’t they good enough for us?”).

Alternatively, it could be a lack of awareness of what the less-well-known parameters mean and what they do, how they were set, or why one might want to change them. Whatever the reasons, we need to recognise that choosing a “default” value is still making an assumption, an assumption for which we need to take responsibility, rather than delegating that responsibility to the CMI.

As shown in *Figure 3*, the process for determining an improvements basis may result in the use of a CMI model (and given the flexibility in those models, this might well be the case). The difference will be that the parameters are chosen to represent the views of the actuary, rather than the views of the actuary being determined by the choice of (default) parameters.



## Short-term trends

Over the next one to (up to) five years, it might be reasonable (in the absence of specific indications to the contrary) to base improvements on those improvements which have recently been experienced in the population.

The rationale behind an extrapolative approach is that we would expect continuity in terms of both:

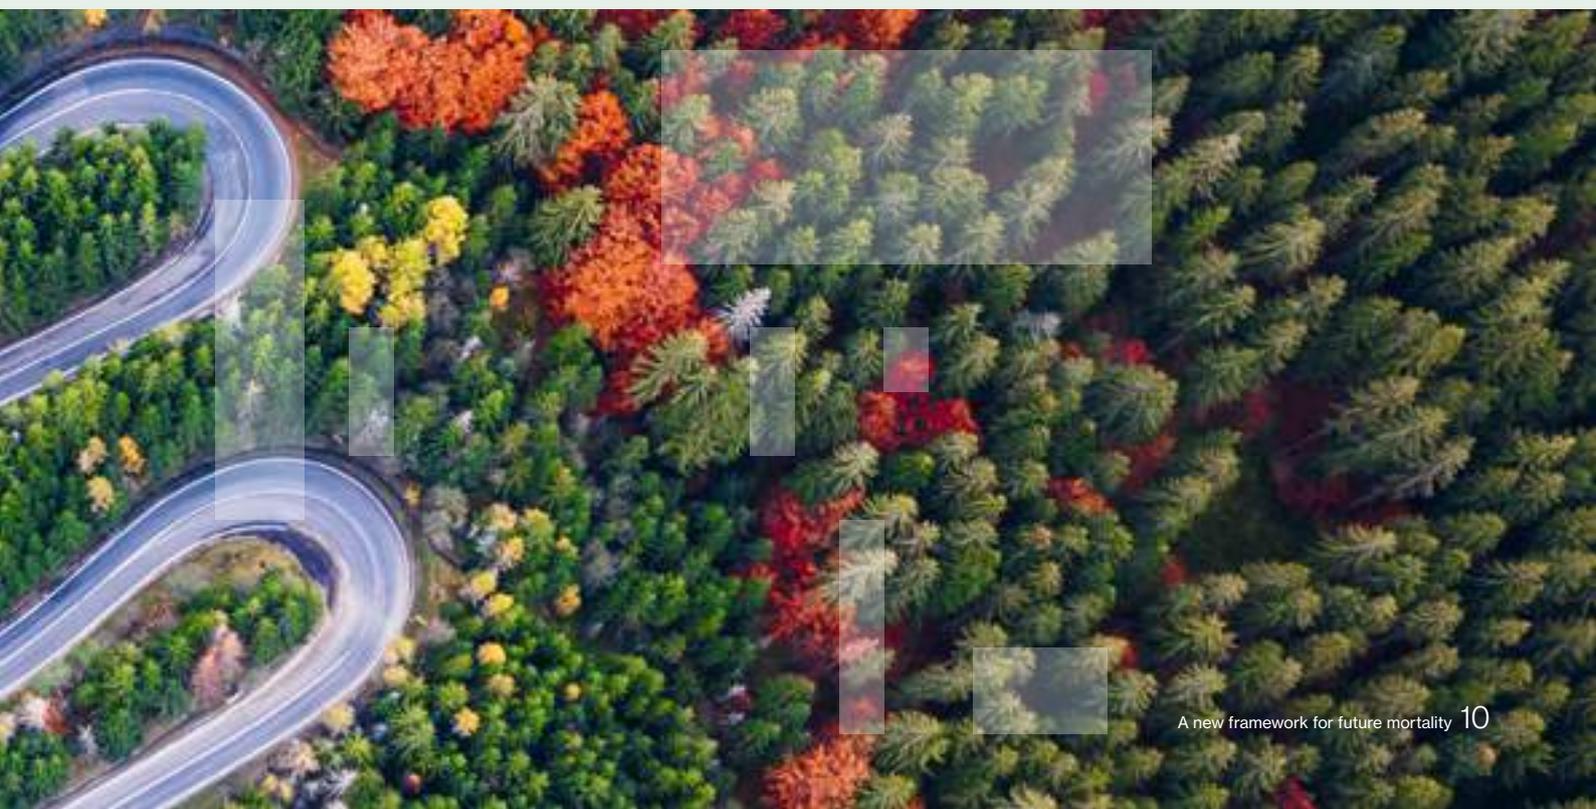
- the trends in drivers of mortality which are giving rise to improvements in the short term, and
- the effects of past changes in key drivers of mortality whose impacts are still being experienced.

It is important to note, however, that short-term improvements can deviate significantly from those extrapolated from short-term experience. One only has to look at the change in trend experienced around 2012, or perhaps most acutely as a result of 2015's heavy mortality experience, to see that a purely extrapolative approach will not guarantee that short-term improvements are predicted effectively.

Indeed, there are a number of good reasons for short-term volatility: extreme weather events, epidemics (e.g. influenza), sudden increases or decreases in public health and social care expenditure, spikes in utility prices (e.g. fuel poverty) and extreme economic conditions (recessions, unemployment spikes) to name a few. If such events occur (with positive or negative impact on mortality), they will not necessarily be present in the recent data,

so their effect cannot be simply extrapolated from that data; rather, their effects may need to be estimated based on earlier mortality data, or alternatively assessed from first principles. Short-term improvements could be assessed based on a range of datasets, so companies might consider evidence (for the effect of particular events) from primary care data and hospital episode statistics (through partnership with an academic institution) as well as using the usual ONS or Human Mortality Database sources.

If a recent CMI model is used to represent the views developed here, then the initial improvements (both age-period and cohort) and the direction of travel will be the key quantities to align to the short-term view of improvements. The choices of the smoothing parameters for the period and cohort components of the Age-Period-Cohort-Improvement (APCI) model would be improvement if initial improvements were being set by fitting to past data rather than by specifying initial improvements and direction of travel. For the CMI\_2018 model, an explicit adjustment to the initial rates of improvement can be made to avoid having to adjust initial rates via manipulation of the smoothing parameters.



## Medium-term improvements

Developing a view on medium-term mortality improvements will inevitably mean examining what has been driving recent past improvements. The premise behind this is that the default rate of improvement in the absence of any change would be 0%; every improvement or worsening in mortality must be due to some driver of change. Short-term drivers of improvement will include:

- Recent changes in health-related behaviours – e.g. smoking cessation
- Recent novel pharmaceuticals going on-licence or drugs coming off-patent and generic equivalents becoming available (increasing prescription rates)
- Changes in disease prevalence, e.g. diabetes or Alzheimer's disease
- Changes in health and social care funding and changes in healthcare policies
- Prevalence and severity of infectious disease
- Climate and short-term weather patterns, e.g. extreme/mild winters
- Potential new cohort effects, e.g. susceptibility to measles outbreaks due to low immunisation uptake in some birth cohorts.

A key step in converting this information on drivers into a view on improvements is calibrating the effect of a change in each driver on mortality from each of a range of causes of death. It is possible to glean information from academic studies, e.g. relative risk measures, which can be used to estimate the impact of a given unit of change in a driver on cause-specific mortality for each of a number of causes.

Armed with indicators of how the drivers have changed over the recent past and best-estimate projections (e.g. from academic studies or NHS and Public Health England publications) of how those drivers might continue to change in the short- to mid-term future, it is possible use the calibrated links between drivers and cause-specific mortality to estimate the overall evolution of mortality rates in the future. Understanding the interactions between drivers is important here, since a change in two drivers might mean a greater or lesser combined impact than the sum of the individual component changes.

Another highly useful source of information to feed into the decision-making process is the opinion of a range of medical experts, each able to provide insight into the likely developments in their own field of medicine.

Multi-state modelling is an excellent medium for the blending of such expert views. By approximating the prevalence of diseases in the population alongside combinations of medical and lifestyle factors, it is possible to use a state-based model to project future mortality rates for subsets of the population and then, by weighting survival curves appropriately, to determine changes in overall population mortality rates.

This has the advantage of capturing the changing importance of different conditions as their prevalence varies both by age and future timestep – improvements for those with individual conditions will contribute more or less to the overall mortality improvements as the prevalence of those conditions in the population increases or decreases.

The Willis Towers Watson PulseModel uses a multi-state approach to project mortality rates based on the medical and lifestyle information available for a given policyholder.

Medical expertise has been used by a number of companies in building the 'new information' components of their longevity trend risk internal models, often involving a process in which both 'best-estimate' views and 'stressed' views of improvements for a number of conditions or causes of death are developed. The equivalence between the best-estimate views and the actual best-estimate assumptions is often tested as part of the validation of the implied model stress. What would take place here, on the other hand, would be using the expert opinions (amongst other sources of information) to set those best-estimate views in the first place.

Where the CMI model is used to represent the views developed here, the key parameters will be the convergence rates (shape and speed of convergence) to the long-term rates for both age-period effects and cohort effects. These will need to be chosen with the initial improvements and direction of travel, as well as the long-term rate, already set, as the medium-term improvements will be determined by the shape of the blending from the short-term rates to the long-term rates.

# Long-term assumption setting

Establishing a view on what improvements will be in the long-term future is probably the hardest area of judgement in developing this type of framework. Even if a CMI model is assumed at the outset, this is a question which a company needs to grapple with, since the CMI (perhaps wisely) does not opine on what the appropriate level of this parameter should be.

There is a whole range of potential approaches to determining a long-term rate of improvements and a company may choose to use some combination of these when determining their choice for the long-term rate of improvement in their chosen model.

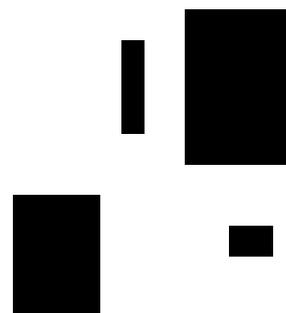
- One could consider a small number of fundamental drivers which are expected to be predictive of future mortality for mechanistic reasons. These could include economic growth (which might affect wealth in the population and tax receipts) or per-capita health expenditure growth, perhaps relative to expected demand for healthcare services.
- One could use 'horizon scanning' to determine what are likely to be the main drivers of mortality improvement over the long term. The potential effect of these drivers in the long term could be determined using some combination of information about similar strong drivers of mortality in the present and their impact on mortality, and an assessment of the likely distribution of conditions prevalent in the population (or major causes of death) in the future and the potential impact of these drivers on mortality in the presence of those conditions (or from those causes). The information gleaned in setting the company's views on short- to medium-term mortality evolution can be used to help to determine the likely future prevalence of such key conditions or future distribution of causes of death.
- One could consider international trends, identifying drivers of mortality improvements which have not had a significant impact to date in the UK and consider how these could affect the UK in the future. However, arguably these drivers could affect the UK over the short to medium term rather than the long term and their impact may not be sustained over a significant period.
- One could consider the annualised improvements implied over a sufficiently long period of history. The danger of this is that it may imply that particularly strong drivers of mortality improvements in that period in the past would be expected to be replicated over any similarly long period in the future (perhaps affecting different causes of death).

- One could also consider the 'cost of making future improvements' in terms of what has to be spent (in real terms) in order to achieve an increase in expectation of life of one year, based on research into the costs of achieving the same increases from earlier ages. The growth in this cost per additional year of age has, in the past, been touted by some academics as a way of showing that there is a practical upper limit to population-level expectations of life. This aside, it's sensible to consider whether the chosen long-term improvement assumption would imply a need for an unreasonable increase in spending on health and/or social care.
- Benchmarking against long-term rate assumptions from other insurers is common. The problem with this is that there are too few participants contributing to benchmarking to constitute a liquid market in mortality improvements and often the contributing companies desire to be close to the median from past benchmarking themselves, meaning that a herd mentality may develop.

Tapering of the long-term rate at older ages is also something that companies need to consider. The CMI's research indicates that there is some interdependence between the choice of ages over which to apply tapering and the choice of the long-term rate to apply pre-tapering. Tapering from an earlier age would, on the basis of past improvements, suggest that a higher long-term rate were required.

## Socioeconomic variations

Irrespective of the period during which improvement assumptions are being set, consideration may be given to whether an explicit adjustment is appropriate to reflect a higher socio-economic status on average amongst insured lives (relative to the population). However, there is an argument that one would expect, in the longer term, a degree of convergence between the improvements for higher and lower socio-economic groups.





## Aligning best-estimate and capital modelling assumptions

One advantage of using a process such as that in *Figure 4* to set the best-estimate is that the process would inform the internal model calibration, not just the best-estimate calibration.

- The effect of new information received over the course of the year would be interpreted in the context of a change in the evolution of the various drivers being considered. This would result in an update of the improvements over the medium term (and possibly long term).
- The effect of new data received over the course of the year would be to update the view on short-term improvements.
- Parameter uncertainty and model uncertainty would come through in the models used to combine the effects of the drivers – uncertainty around the future evolution of each driver, uncertainty around the link between those drivers and mortality from various causes, and uncertainty around the interaction between drivers. These would each result in changes in the answers to the question of consistency with a given model. Revising the internal views on improvements over each time horizon would potentially result in a different model being selected, or a different parameterisation of that model being selected.
- Basis risk would also come through in the parameters used to link drivers to mortality – having a different population could mean greater or lesser sensitivity to particular drivers (the best-estimate and stressed scenarios for each parameter may be out, e.g. due to population prevalence of smoking being unknowingly higher than that in the insurer’s portfolio, and predicted declines in each scenario being greater than those which would be expected in the portfolio were this known).

Insurers’ internal models of longevity trend risk will typically indicate “what would happen” if new information became available over the course of a year to indicate potential changes in future improvements.

However, in our experience, insurers rarely make any explicit adjustments to their best-estimate improvements basis in respect of information that they have received over the past year.

By switching from a near-default move from CMI\_20xx to CMI\_20xx+1 each year to a process in which improvement assumptions are derived from own views, an insurer can also switch from a position where the default is to disregard new information – behaving inconsistently with the principles underlying their new information risk models – to one where the new information risk model is genuinely consistent with the improvements model, not just because the expected improvement rates (in the statistical sense) under the new information risk distribution are the same as the best-estimate model, but also because the way that new information is treated in the capital model reflects exactly the process through which the best-estimate improvements will be set in the following year.



# Implications for financial reporting

Implementing this type of approach to setting mortality improvement assumptions has some potentially significant implications for insurers' financial disclosures.

The initial impacts of changing the approach to setting mortality improvements are potentially significant; if there is a significant difference between the resulting improvements and the prior CMI model, then this could mean a significant change in the best-estimate liabilities and capital requirements.

Arguably, if the insurer's informed and carefully reasoned views on the improvements basis were to suggest a significantly higher rate of improvement for annuitants than had been implied by their previous model, then we would expect insurers to want to reflect this view by holding higher technical provisions.

Alternatively, if the views represented a significant reduction in improvement rates, then a potentially sizeable release in technical provisions might be dictated by the Solvency II rules. Insurers would probably be concerned about whether taking credit for the whole of this release in a single step were appropriate.

Annuity writers with a matching adjustment will also need to be careful when making liability changes; changing the reserves may result in more surplus being locked into the matching adjustment fund, or alternatively more money needing to be added to this fund. In the former case, procedures need to be in place to demonstrate that a surplus has arisen and to allow this to be released. Furthermore, changes in the duration of the liabilities may mean that the matching adjustment fund needs rebalancing to meet the matching requirements.

If they wanted to unwind the technical provisions more gradually (for example, to allow for potential future volatility in short-term improvements, in particular, due to the new approach taking more account of new information received each year), this would not formally be possible within the best-estimate liabilities under the current Solvency II rules if the actuary held the belief that the new assumptions were a genuine best-estimate (based on properly-understood data meeting criteria set out by the insurer).

Moreover, whereas, under the old Solvency I, Pillar 1 Peak 1 reporting, an additional discretionary reserve could be set up to protect against volatility, the Solvency II rules do not allow for discretionary reserving.

One solution would be to make supplementary disclosures – an adjusted Solvency II balance sheet with additional reserves set up explicitly to smooth the impact arising from the initial change to the improvements basis (and any volatility in short term improvements in the future).

In practice, however, what actuaries are more likely to do is to rely on the clause in Article 77 of the Solvency II directive requiring the best-estimate to be “based upon up-to-date and credible information and realistic assumptions”, where they could effectively manipulate their best-estimate assumptions on the basis of needing to wait to confirm that their updated view was indeed credible and supported by further emerging data prior to accepting it in full. This approach would also allow them to deal with the potentially greater volatility in the short-term improvements basis. In any case, this might be expected to be less significant than the initial impact of the change in approach.

For those insurers using an internal model, a significant change in the improvements basis (and the approach to setting this) may give rise to a need for major or minor model changes with a knock-on impact on overall capital requirements.

For companies producing accounts in accordance with UK GAAP, the use of prudent margins in the reserving assumptions (per the old Solvency I, Pillar 1 Peak 1 valuation rules) means that there is far more discretion on the part of the Chief Actuary in respect of how to allow for the change in the best-estimate improvements. They can exercise their judgement as to how they allow for an initial change (due to the change in methodology) and the size of the margin that they apply in their reserving assumptions can be set so as to absorb any volatility in the short-term views of improvement rates over time.

Similarly, on a future IFRS 17 balance sheet, an increase in the volatility in the improvements basis may be less of a concern. Any change in assumptions resulting in a fall in liabilities will not flow through to the profit and loss (P&L) accounts, but instead accumulate within the contractual service margin ('CSM') to be released over the term of the contract. Assumption changes resulting in an increase in liabilities would reduce the CSM (to the extent that such a CSM was available) or otherwise be recognised as an immediate loss in the P&L accounts. This treatment leads to reduced balance sheet volatility, but the recognition of the benefit of a (positive) improvement assumption change will be (potentially greatly) deferred.

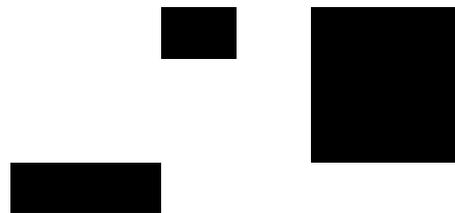
Arguably, if there were a change in the outlook for improvements with an expectation that (at least short term) improvements would become more volatile, then there could be an increase in the risk adjustment (the same confidence measure, or Value at Risk could imply a wider range of possible cash flow patterns – amounts and timings – and therefore a larger deviation from the best-estimate value of liabilities). The run-off of this risk adjustment would likely be quite different from that of the CSM.

### A change in the improvement assumptions could affect whether new contracts are considered “onerous”.

At present, guaranteed annuity options are generally heavily in-the-money for policyholders, so annuity contracts arising from such guarantees are likely to remain onerous unless there are quite extreme changes to improvement assumptions. For open market annuities (and perhaps to a lesser extent internally vesting non-guaranteed annuities) if improvement assumptions were to be strengthened

sufficiently, then some annuities could become onerous during the lifetime of the policies, although this would not affect whether the contracts were considered onerous at inception. So long as such annuities were repriced to reflect the updated improvements basis, the classification of the annuities under the onerous contracts test should not be affected.

Over time, if contracts in a cohort were to become onerous, the CSM for the cohort would be eroded or, if this reached zero, a loss would be recognised on the balance sheet.





## How Willis Towers Watson can help

Our mortality projection tools are ideally suited for use in implementing a bottom-up approach to projecting mortality improvements: Willis Towers Watson PulseModel is an ideal tool to support the development of rates of improvement in the medium term, whereas our Driver-based Model offers insight into improvements at all durations, but perhaps most importantly into long-term improvements in mortality.

### Willis Towers Watson PulseModel

Our multi-state model of morbidity and mortality risk tracks policyholders from health, through disease states and eventually through to their death. It allows for the effects of medical and lifestyle-related risk factors as well as the effect of duration-in-state.

Figure 5. PulseModel's multi-state structure

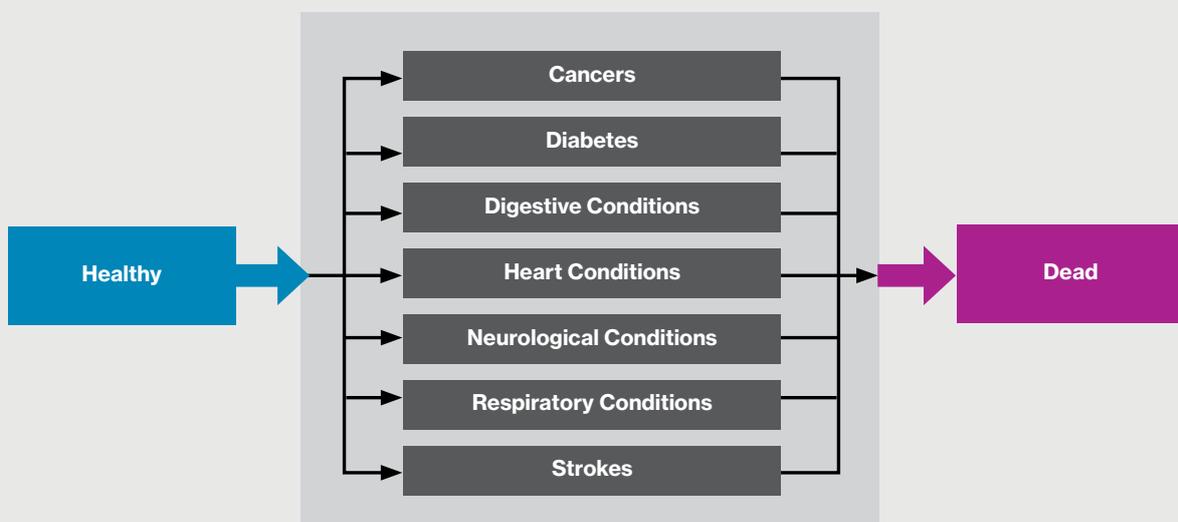
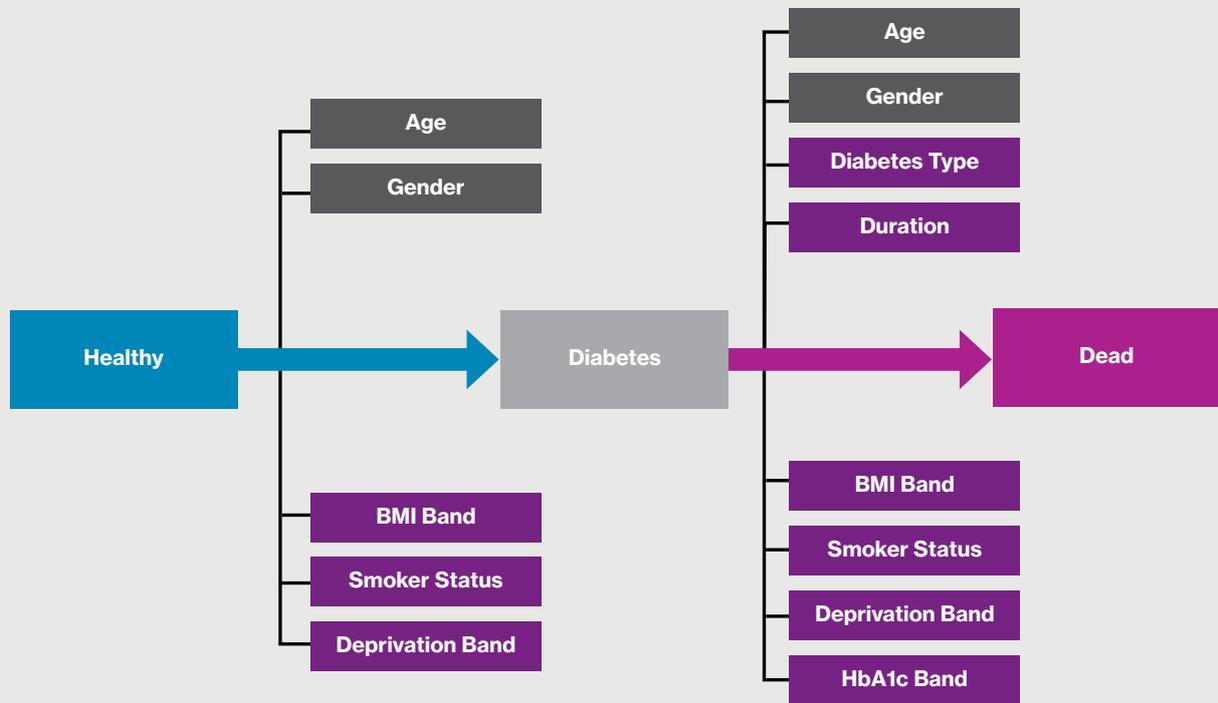


Figure 6. Factors affecting a typical transition model



■ Required input

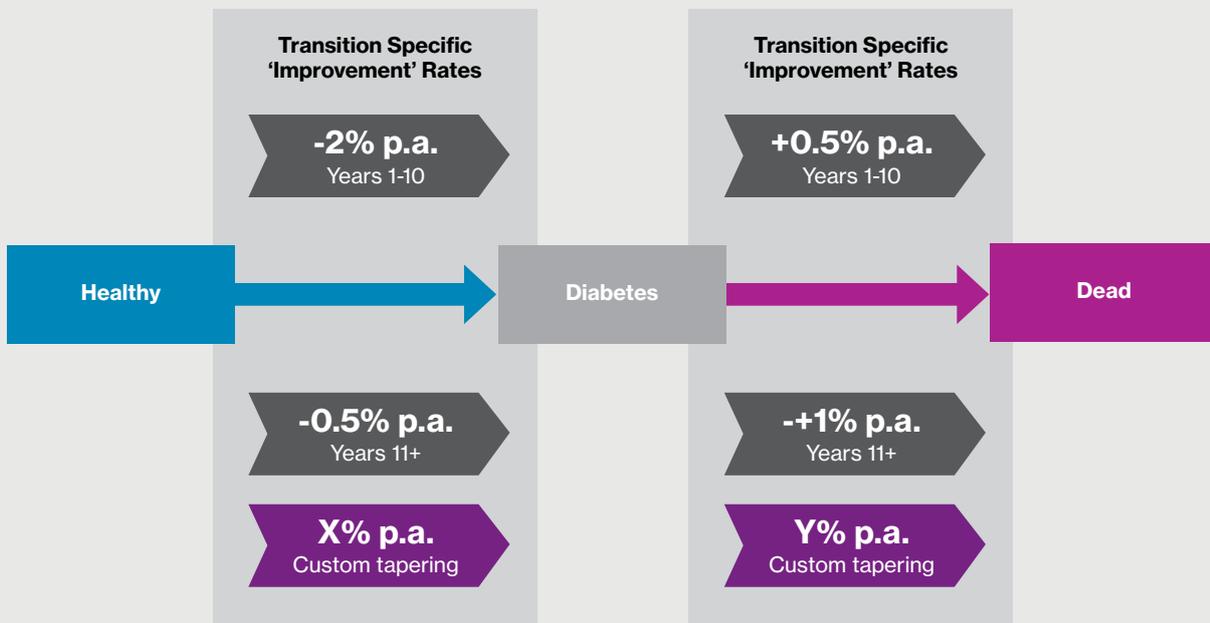
■ Can be "Unknown"

NB. Illustrative figures only

The evolution of morbidity and mortality rates in PulseModel derives from the opinions of a panel of medical experts, covering a range of specialisms (aligned with the disease groups in the model). These opinions, supported by detailed reports, differentiate between medium-term and long-term expectations for changes in morbidity and mortality rates and, in cases, between genders. Importantly, these views are forward looking, not bound to replicate the trends observed in past data – particularly appropriate when that past data includes the effects of strong drivers of improvements (e.g. smoking reduction and advances in cardiovascular medicine) which are not expected to be repeated in the future.

PulseModel's multi-state structure acts as a blending tool to combine the effects of these disease-specific improvements, taking into account the reweighting of the various diseases over time as both new diagnoses and deaths result in changing age-specific disease prevalence over the course of the projection.

Figure 7. Improvements per disease group – PulseModel acts as a blending tool



NB. Illustrative figures only

The opinions of medical advisors might be expected to be well-informed over the short- to medium-term future and then to become progressively less reliable beyond (say) 10 to 15 years. PulseModel is therefore an ideal tool to support the development of rates of improvement in the medium term. Information about projected population disease prevalence from PulseModel could also form one input into the decision for long-term improvements.



### Willis Towers Watson Driver-based Model

Our Driver-based Model (or “cause-of-death” model) considers the response of overall population mortality to changes in the levels of individual drivers of mortality.

For example, if we believe that there will be a continued trend towards decreasing smoking prevalence, then by understanding the link between smoking and mortality from a range of different causes, the model will predict a change in mortality which reflects this change in behaviour.

The drivers in the model can be varied, allowing for the fact that different actuaries (and those experts advising them) will likely have differing views on which are the key drivers underlying the mortality improvements of the future. Irrespective of the drivers chosen, the model is typically calibrated to a wide range of results from academic papers, overlain with the actuary’s own expert judgement.

Figure 8. Example of a best-estimate scenario in the Driver-based Model

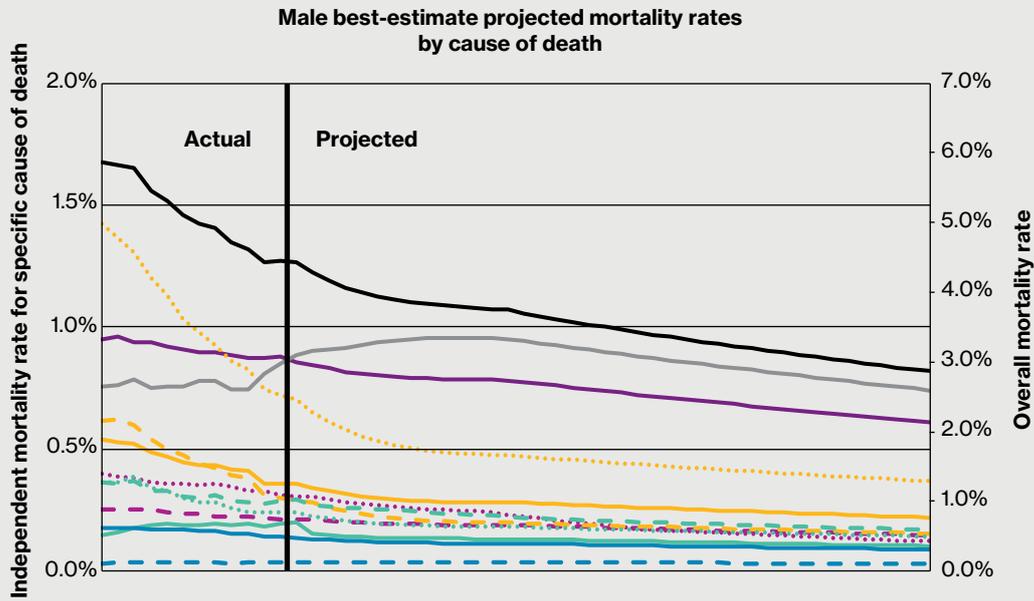
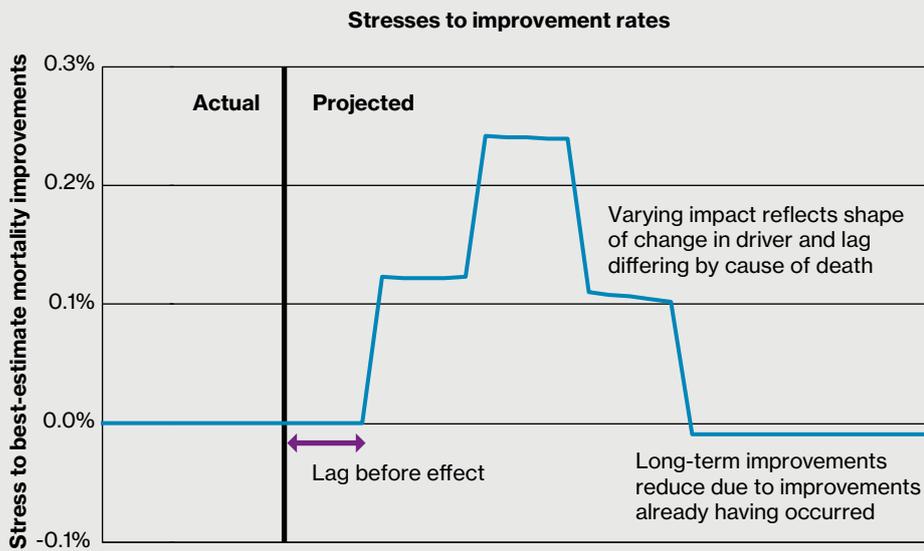


Figure 9. Example of the impact of setting an alternative scenario for one driver

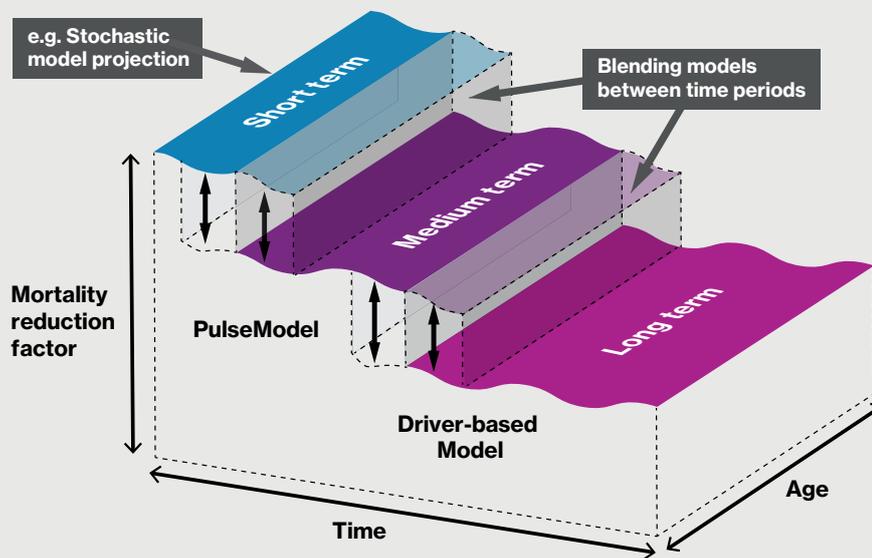


Depending on the drivers considered, it may be reasonable to consider long-term scenarios for their evolution. This means that it can be used to inform improvements in both the medium term and the long term.

The following diagram shows a configuration of three models of improvement rates which can be used to develop a reasoned view of improvements over the full spectrum of time horizons.

Taking this information as a guide and developing a combination of parameters in (for instance) one of the CMI series of projection models to reflect the views of the actuarial team should make it possible to then explain the overall improvements model to the board or relevant governance body in a more helpful and informative way, supporting a better understanding of mortality/longevity risk in general.

Figure 10. **Blending insights from different improvement models**



## Contact us:

### Further information

If you wish to discuss any of the issues raised in this publication, our mortality and longevity experts would be delighted to assist you – please get in touch using the contact details below.



**Matthew Edwards**

+44 1737 284 771

[matthew.edwards@willistowerswatson.com](mailto:matthew.edwards@willistowerswatson.com)



**Richard Marshall**

+44 161 833 6244

[richard.marshall@willistowerswatson.com](mailto:richard.marshall@willistowerswatson.com)



**Kenny McIvor**

+44 20 7170 2402

[kenneth.mcivor@willistowerswatson.com](mailto:kenneth.mcivor@willistowerswatson.com)



**Momtchil Iliev**

+44 20 7170 2577

[momtchil.iliev@willistowerswatson.com](mailto:momtchil.iliev@willistowerswatson.com)





## About Willis Towers Watson

Willis Towers Watson (NASDAQ: WLTW) is a leading global advisory, broking and solutions company that helps clients around the world turn risk into a path for growth. With roots dating to 1828, Willis Towers Watson has 45,000 employees serving more than 140 countries and markets. We design and deliver solutions that manage risk, optimise benefits, cultivate talent, and expand the power of capital to protect and strengthen institutions and individuals. Our unique perspective allows us to see the critical intersections between talent, assets and ideas – the dynamic formula that drives business performance. Together, we unlock potential. Learn more at [willistowerswatson.com](http://willistowerswatson.com).

Willis Towers Watson  
51 Lime Street  
London  
EC3M 7DQ

Towers Watson Limited (trading as Willis Towers Watson) is authorised and regulated by the Financial Conduct Authority in the UK.

The information in this publication is of general interest and guidance. Action should not be taken on the basis of any article without seeking specific advice.

To unsubscribe, email [eu.unsubscribe@willistowerswatson.com](mailto:eu.unsubscribe@willistowerswatson.com) with the publication name as the subject and include your name, title and company address.



[willistowerswatson.com/social-media](http://willistowerswatson.com/social-media)

Copyright © 2019 Willis Towers Watson. All rights reserved.  
WTW-HP-2019-0187

[willistowerswatson.com](http://willistowerswatson.com)

**Willis Towers Watson**