

Evaluating Biometric Trend Drivers

How to reflect medical breakthroughs and other drivers in forward-looking assumptions



Introduction


Forward-looking biometric assumptions, including mortality improvement assumptions, are among the most financially significant for life (re)insurers and pension providers. Actuaries must form a view on mortality rates many years into the future, knowing that new positive and negative drivers can emerge in the interim. Improvement bases should reference historical improvement trends over both the long term and short term but also reflect foreseeable changes to future trends. Similarly, bases should react to changes in forward-looking and historical information without overreacting and introducing unjustified basis instability.

A classical problem for actuaries is knowing how to interpret and respond to information about new or changing mortality drivers. It is easy to become overwhelmed by the sheer number of potential medical breakthroughs touted by media and academic literature at any given time, and it takes skill and experience to sift out the signal from the noise. A timely example is the rise in popularity of anti-obesity medications, but there are myriad other examples, including vaping, climate change, anti-aging medications, and multi-cancer early detection tests, to name but a few. In this paper, we lay out some of the considerations actuaries must grapple with when supplementing projections with new information – namely:

- Identifying drivers
- Quantifying their impact
- Determining relevance for insured lives or annuitants
- Forming a view on what is already allowed for in existing bases

The starting point for most actuaries will be an existing improvement basis and information regarding a potential new driver. As we will see, the nature of the existing basis is an important part of any decision for how to treat the new driver, and so we start there.

The paper focuses on mortality and longevity bases, but the principles are also applicable to morbidity.



A decade ago, immunotherapy was an emerging option for a few advanced cancers; today, it is a central pillar of cancer treatment. CAR T-cell therapy, depicted here, involves genetically engineering a patient's own T-cells to attack cancer cells.

The sheer number of potential medical breakthroughs touted by media and academic literature can be overwhelming, and it takes skill and experience to sift out the signal from the noise.

The importance of existing bases

All (re)insurers will include some “secret sauce” in their basis-setting process, but it is reasonable to assume that a large proportion of most future improvement bases can be reduced to the following steps:



STEP 1

Model historical mortality data to extract parameters or rates that can be forecast or extrapolated.

Example – Age-period and cohort improvement rates.



STEP 2

Determine a method for extrapolating those parameters, or modeled improvement rates.

Example – Convergence to a long-term rate or a statistical approach to time-series modeling.



STEP 3

Make adjustments for known model limitations, basis differences, or “one-offs.”

Examples – An overlay for COVID-19 deaths, or additional improvements to reflect faster improvements for insured lives vs. population improvements, or an adjustment to reflect the one-off nature of smoking cessation.

As a general rule, near-term improvements (initial or short rates) are assumed to resemble the recent past and over the longer term will revert to a more sustained average set using expert judgment and statistical analysis (long-term rates).

The approaches of many major actuarial bodies – for example, the CMI models in the UK and the MIM models produced by the SOA in the US – adhere to the same core principles.

Although the methodology seems quite formulaic and prescriptive, actuarial judgment is applied at multiple points. The long-term rates and the pace at which they are approached (the convergence period) are highly subjective, for example, and decisions about how to interpret and treat historical data, which to outside observers may seem esoteric and trivial, can be contentious and financially material. Actuaries will typically lean on mortality subject matter expertise when setting these parameters:

- Long-term rates may be set with reference to optimism or pessimism around sources of medical advancements.
- Factors that influence short-term rates may be set with reference to the current mortality environment; pessimism around healthcare provision may lead to a parameterization that gives lower short rates when the data and model could support higher rates, for example.

As a result, improvement bases are a blend of extrapolation and expert judgment. Reassuringly, this approach of taking “base rates” or “outside views” (looking at historical data) and adjusting for what is known about a specific situation has proven particularly effective in creating accurate forecasts.¹ However, it can mean that a substantial portion of future improvements are not attributed to sources or drivers of those improvements. There is nothing inherently wrong with this. Holding a firm view on exactly how the entirety of these improvements will arise is probably spuriously precise, but we shall see later that this can lead to difficulties when deciding how to incorporate forward-looking analyses into improvement bases.

¹ Superforecasting; The Art and Science of Prediction. Tetlock and Gardner, 2015.

Identifying drivers

For the most part, identifying drivers of future improvements is straightforward: scanning the literature, attending conferences, and consulting with internal and external experts will normally identify a wide range of topics that may have a material impact on future mortality rates. These will range from optimistic (anti-aging therapies) to pessimistic (extreme climate change scenarios), from high tech (personalized medicine) to low tech (smoking trends), and from near term (obesity) to long term (changes in population distributions). Actuaries initiate genuinely new avenues of research from time to time, but, in general, our role is to sift through potential drivers highlighted by research groups and industry, use our skills and knowledge to identify those that look the most impactful, and update our assumptions accordingly.

When triaging candidate drivers, there are three main factors to consider:



1. Magnitude of impact



A breakthrough with a sizable magnitude of impact on mortality for a large proportion of the population should be a higher priority for further analysis than something that applies only to a small subset of the population. It is important to consider both positive and negative drivers, and working with both life insurance and longevity teams is an effective way to ensure both types of drivers are considered without bias.



Proximity refers to when drivers are likely to have an impact. In general, impacts that will apply in the near future should take priority over those expected decades from now. Discounting means that near-term impacts are more financially significant than those that occur over the long term, and impacts too far in the future may come too late to impact today's insured lives and annuitants.



Finally, we should focus more on events with a high likelihood of occurring than those that are speculative.² For example, we are unlikely to overturn the laws of physics and see nano-robots repairing our cells by 2027. Effort is better spent understanding the current state of multi-cancer early detection trials.

Proximity and probability are interrelated. Events touted as being impactful many years from now are often less likely than those expected sooner. In simple terms, if drivers are close to or already having an impact, less can go wrong or change between now and their anticipated impact date. For drivers anticipated many years away, there is more scope for things to go wrong or stall. We should account for those hurdles when assessing long-term drivers.



...if drivers are close to or already having an impact, less can go wrong or change between now and their anticipated impact date. For drivers anticipated many years away, there is more scope for things to go wrong or stall.

2 This is true if your main concern is best-estimate assumptions. For stressed assumptions, unlikely but potentially impactful events are important.

It can be tempting to maintain a long list of drivers and incorporate all of them explicitly in basis-setting exercises, in the belief that this leads to a more-complete view of how mortality rates will evolve and therefore leads to higher-quality forecasts. This approach risks creating spuriously precise forecasts, a compounding propagation of uncertainty, and an unwieldy number of items to update between basis-setting cycles.

Instead, we advocate focusing on a smaller number of drivers that could plausibly lead to material deviations from existing mortality forecasts and making basis adjustments only for those that are compelling enough to justify doing so. This does not mean other drivers are ignored; we may wish to form a view on a driver for reasons other than incorporating into our bases, and concluding that bases should not be adjusted for a driver also requires research and quantification. It is useful to have a broad research base as a reference point when breakthrough research papers or clinical trials are published.

Drivers may be of interest to (re)insurers for wide-ranging reasons: some will be of interest because of the potential for tangible near-term impacts on mortality rates (anti-obesity medications and multi-cancer early detection tests, for example), but others will be part of a broader horizon-scanning effort that helps (re)insurers to react quickly to breakthrough papers or clinical trials regarding those drivers. Some drivers might be deemed low probability but potentially high impact, which means they are of more interest for setting capital requirements than for best-estimate assumption-setting.

At RGA, we devote significant resources to analyzing the drivers of future mortality and morbidity rates. This is an ongoing effort, monitoring academic literature and analyzing in-house and publicly available data, and incorporates a structured annual process of engaging with stakeholders across the business.

Given the long-term nature of mortality rates' evolution, we expect the relative importance of drivers to remain fairly stable year-on-year. Drivers would be prioritized for deep-dive research according to business needs and high-level estimates of potential impacts. Most drivers will not feature as explicit adjustments to bases; only the drivers with the largest and most tangible impacts are expected to be taken forward to basis-setting teams, and even these drivers require significant judgment to decide whether the impacts are already implicit in existing improvement assumptions. We discuss this in more detail later in this paper.



We advocate focusing on a smaller number of drivers that could plausibly lead to material deviations from existing mortality forecasts and making basis adjustments only for those that are compelling enough to justify doing so.

Quantifying impacts

Quantifying a driver's impact should involve two phases:

1. A preliminary calculation of an approximate value for triaging
2. A more-detailed calculation if the approximate impact warrants further investigation

The ability to quickly rule out immaterial drivers is key. Not every research project should be judged by whether it leads to a change in assumptions. Protecting the business from reactive assumption changes and expensive projects adds value. Communicating and documenting decisions to halt investigation early is important for consistency and transparency, especially in large organizations.

In many cases, the transition from phase 1 to phase 2 may be gradual. Initial impacts for a new drug may be calculated as the proportion who are eligible multiplied by mortality impacts from a clinical trial, with some reduction for anticipated uptake being less than 100%. Refinements may be made over time, according to business need or data availability. For example, refinements by age and/or sex may converge on a more-nuanced view of impacts, eligibility, and uptake.

A balance should be struck between complexity and robustness. It can be tempting to model every conceivable variable linked to a given driver, but parsimonious models with fewer parameters are preferable in most cases. Each additional parameter is another input to consider and another source of potential errors. If complex models are unavoidable, model owners should consider exposing only the most material parameters to end users. Changing advanced parameters should occur only when strictly necessary and under the supervision of model experts. Models should, of course, be built carefully, and limitations should be well-understood, but practitioners should be mindful of the broad uncertainty of predicting future mortality rates and not seek false comfort in model sophistication.

Insured lives and annuitant adjustments

Insured lives and annuitants are typically more affluent than the general population and benefit from lower mortality rates. Insured lives with standard or preferred underwriting ratings have demonstrated good health at the time of underwriting, and these groups experience lower mortality rates as a result, especially in the years immediately following underwriting. We might also expect insured lives to have better access to healthcare. If we consider anti-obesity medications briefly, we can see that better health status and improved access to healthcare can act in opposite directions: Insured lives may be more able to afford these treatments than the general population, but this may be counterbalanced by a lower average body mass index (BMI) among insured lives and, hence, less scope to offer mortality improvements.

At the aggregate level, the causes of death are likely to be quite different among an insured book vs. the general population. The nature of this will vary by factors such as country and underwriting philosophy. This will be a large factor when translating between population and insured impacts. In general, we know much more about causes of death (or claim) among insured lives than we do annuitants because cause of death information is not typically collected for annuity business. In these cases, wealth or other indicators of socioeconomic status can create a proxy for annuitants from population-level data where that exists.



Ruling out immaterial drivers is key. Not every research project should be judged by whether it leads to a change in assumptions.



Offsetting impacts for improvements already assumed

Given some modeled impacts of a driver on an insured population (or annuitants), what changes should be made to the existing basis? This is a fiendishly difficult question, given the previously discussed nature of existing bases: They generally assume future mortality improvements but do not ascribe precise sources. This means that when we identify a material driver of future mortality improvements, we cannot automatically assume any impacts beyond what we already predict in our existing basis. The same is true for drivers of negative mortality trends: Historical trends have included negative drivers (such as obesity) that have moderated the impact of positive drivers.

If we identify a material positive driver of future mortality improvements, we find ourselves in one of the following three scenarios:

1. The resulting improvements are more than what was previously anticipated (implicitly or explicitly), so assumed mortality improvements need to be adjusted upward. The adjustment should recognize that some of the improvements might already have been anticipated; thus, the basis adjustment will generally be smaller than the impact of the driver in isolation.
2. The improvements are consistent with what is required to realize what is already assumed in the existing basis.
3. The improvements fall short of what is required to sustain the assumed level of improvements, either because a larger impact was previously assumed for the driver or because not enough positive drivers have been identified to support an existing view of high future improvements, so assumed mortality improvements need to be adjusted downward.

In clear-cut cases where the magnitude of a driver impact is very large, it follows that a basis adjustment is justified. More often than not, it will be a matter of judgment to determine which of the three scenarios applies. Two actuaries could agree on the impact of a driver in isolation but take different actions based on the level of future improvements in their existing bases.

A practical approach to adjusting bases for new drivers

So, what should actuaries do with their improvement bases when faced with a new driver of mortality improvements? The answer will always involve significant judgment, but the following framework can help:

- Focus only on the top handful of drivers likely to drive basis changes. An extrapolative approach to setting future improvements is most likely to capture the net impact of the myriad positive and negative drivers that could ordinarily be expected to impact mortality rates.
- Understand trends in the causes of death thought to be impacted by the key drivers. Many cancers have experienced consistently high improvements historically, and these are implicitly extrapolated forward as part of methods such as the CMI models. Future cancer breakthroughs are required to sustain this rate, and, as such, we might view news of a new cancer therapy as “expected.” For more stubborn causes of death, such as dementia and Alzheimer’s disease, news of a proven therapeutic would be a clear departure from historical trends.

- Drivers with imminent and material impacts are more likely to deflect trends from their anticipated trajectory than those with more distant impacts. This is because, in general, the drivers of existing trends are less likely to be immediately exhausted and require replacement to sustain current trends.
- Allow for explicit modeling of the drivers in short-term projections, but taper this off over time and allow for more-distant impacts through altering the long-term rate (LTR).

By focusing on what we think are likely to be the most material drivers of future improvements, we can focus on the hard but tractable question of whether we think an extrapolative mortality projection is correct, given the magnitude of some of the most significant drivers. The alternative of forming a view on the entire composition of future mortality improvements is impractical and discards the important information we can glean from the properties of historical mortality improvements.

Sources of improvement many years from now are uncertain, and the composition of the LTR is highly uncertain. Making precise statements about the magnitude and timing of impacts many years in the future is therefore spurious, and even if an underlying model produces predictions with this kind of precision, it may not be desirable to reflect this precision in basis changes. Some countries are experiencing an environment of low initial improvements and an LTR that is higher than the initial improvements, reflecting the view that the low improvements will not continue indefinitely, and careful thought needs to be given to what is required to drive the transition to higher improvements. Drivers that seem like “game changers” may actually be what is required to realize the improvements already assumed; by considering the long- and short-term rates separately, we can be more explicit in our reasoning around how a new driver alters our existing view.

A convenient way to taper off the model impacts over time is to use the same methodology the CMI uses to transition from initial improvement rates to

the LTR. Very briefly, the contribution of the initial rate falls away as the contribution of the LTR is increased, and a parameter called the “convergence period” is used to control the rate at which this happens. For example, a short convergence period means the LTR is approached more quickly. By using this same mechanism and parameterization for drivers we want to reflect explicitly in our bases, we can neatly transition between precise near-term impacts and the longer-term impacts reflected in the LTR.

This is a pragmatic approach that seeks to avoid over-engineering the shape of the improvement surface, while still recognizing that the near- and long-term improvement outlooks have changed.

Most practitioners will use convergence periods that vary by age and may find it awkward that the model outputs are treated differently at different ages. We note that convergence periods are usually set to reflect the persistency of short-term mortality improvements, and that shorter convergence periods are generally associated with ages more prone to short-term mortality fluctuations, which are not conducive to precise prediction. The LTRs also typically vary with age and can be set to counteract any undesirable effects of the age shape of convergence periods.

A consistent and transparent mechanism to translate between research and modeling and changes to the basis ensures a pathway for translating R&D insights into business value. Moreover, being explicit about adjustments made in light of forward-looking research makes it easier to respond quickly to new information – publication of clinical trial results or changes in the price of a drug, for example. Having such a framework also encourages critical thinking and provides a common language for stakeholders and subject matter experts across a range of teams to discuss and sharpen their views. The future will always be uncertain, but having the right people, research, models, and processes in place ensures good decisions are made with the imperfect information inherent in future improvement assumptions.



A consistent and transparent mechanism... ensures a pathway for translating R&D insights into business value.

Conclusion

Qualitative research and discussion of potential future mortality drivers are important to a holistic approach to setting future improvements, but they add no value without a method for quantifying and adjusting bases accordingly.

Accounting for specific drivers in future improvement projections is fraught with difficulties and requires a careful blend of analysis and judgment. We recommend a framework that promotes consistent and transparent basis adjustments and errs toward robustness rather than spurious accuracy. This begins with researching potential future drivers of mortality rates and distilling these down to a small number of candidates for explicit basis adjustments. The modeled impacts of these drivers are translated into basis adjustments through an understanding of drivers of the initial rates and of the difficulty of making long-range predictions. All of this must be done with an appreciation that the existing mortality improvement basis already allows for future mortality improvements.

To further explore biometric trend drivers' potential insurance impacts, reach out to your local RGA representative.

Authors



Craig Armstrong, PhD, FIA
Vice President, Senior Actuary,
and Data Scientist, Global R&D

[View bio](#)



Chris Falkous, FIA
Vice President and Senior Biometric
Insights Actuary, Global R&D

[View bio](#)



Richard Russell, PhD
Vice President, Biometric
Research, Global R&D

[View bio](#)



Andrew Gaskell, FIA
Vice President and Senior Actuary,
Enterprise Pricing

[View bio](#)



Ben Johnson, PhD, FSA
Vice President and Managing
Actuary, Global Valuation

[View bio](#)



Elena Tonkovski, FSA, FCIA
Vice President and Senior Actuary,
Insurance Risk

[View bio](#)