2014 International Congress of Actuaries
Communication of mortality experience, by Sam Gutterman, FSA, CERA, FCAS, MAAA

Abstract

Although considerable time is spent in the actuarial literature on the projection and application of mortality rates, limited time and space is spent on how to communicate historical mortality experience and projections. The objective of this paper is to discuss some basic considerations and concerns regarding this communication, as well as selected observations regarding related mortality and longevity issues. Emphasis is given to communication of mortality and longevity in an insurance company environment, although many of the observations made may also be applicable to other circumstances.

In many cases the communication of mortality and longevity results can be enhanced, with the result being improved understanding of the financial effects, uncertainty, and causes of deviations from expectations can lead to enhanced decision-making of an insurer.

Table of Contents

1. Background

Although many actuaries spend considerable time analyzing mortality experience and developing mortality assumptions for a wide range of applications, they have not necessarily focused on the effective communication of their findings to users of their analyses. The objective of this paper is to provide perspective on how to calculate, group, assess and communicate historical and projected mortality results and uncertainties. Focus is given to experience of insurance entities.

Why focus on the communication of mortality and longevity results? Sometimes it is believed that current management has little control over mortality performance, so why bother. There are several reasons why this is important, including:

- Making changes in future underwriting policy;
- Revising prices for new sales;

- Updating dividends or other non-guaranteed elements;
- Understanding income trajectories and risks;
- Financial planning; and
- Development and keeping assumptions for liabilities and valuations current.

Best practice management reporting differs by the intended audience, incorporating the information and data consistent with the needs and preferences of its users that can differ by level of management or their functional needs or interests. For example, a CEO will usually be most interested in a highlight of key metrics or key performance indicators to highlight any significant deviations from expected that need addressing by management and to understand overall corporate performance. In contrast, staff responsible for analysis of performance, often actuaries, will need a data base with query/drill-down capabilities, in a more detailed actual-to-expected format in multiple dimensions to better understand the effects of the multiple set of factors that might influence experience trends that can hide problem areas, as well as serving as a basis for revisions of future expectations and benchmarks.

This paper does not include a recommendation for a specific format for reporting, as effective reporting is often based on the specific circumstances, concerns and personal preferences involved. Especially for management reporting purposes (in contrast to detailed actuarial experience analysis), a format consistent with management reporting of other performance factors is preferable. For example, the CEO might prefer graphical charts rather than tables filled with numbers, and a printed template rather than a drill down computer-based file or set of pivot tables. Depending on the volume and variety of data, different time periods or breakdowns of the data would be appropriate. The application of any of these reporting options can represent good practice; best practice represents getting the right information in areas of significant performance to relevant users in a timely manner so enable appropriate business decisions to be made.

It is important that different functional areas use a common data infrastructure and set of analytic definitions, suggesting a strong need for a corporate data dictionary. It has been found that if this is not available, different areas may report inconsistent trends, as metrics do not always follow the same pattern. In addition, the IT department (or other area responsible for data definition, aggregation and retrieval) may produce reports that are not consistent with those requested. Although this may not be as important as in other insurance lines, such as health insurance where the data items can be more complex, duelling internal information sources is never pleasant, leads to internal stress, and can contribute to inappropriate decisions.

The key to successful reporting is inclusion of appropriate objectives, benchmarks or trended information from which conclusions can be drawn or identification that more information is needed. Some users will want to focus on longer-term trends through charts or tables, with smoothed results, for example, showing rolling annual period results, while others are more concerned with early warning indicators that may trigger further study or immediate preventive action that would lead to emphasis on current individual period or recent results. The period emphasized will differ by company, coverage, and company situation (generally with timeframes ranging from three to five years).

In this regard, crude calendar period data (that may be updated as reporting lagged data becomes available) may suffice in many cases, although for the purpose of development or revision of prices or insurance liability assumptions, more detailed, fully developed, and highly scrubbed data are necessary, typically developed by actuarial experience analysis areas.

It is also important to note that external reporting should be consistent with what is reported and used internally.

2. Mortality and Longevity

Mortality and longevity risks are mirror images of each other, but apply to different insurance coverages. For example, mortality risk in life insurance is present when the rate of mortality is greater than expected, while longevity risk occurs in payout annuities, level premium long-term care, or disability income insurance prior to a claim (when the rate of mortality is less than expected).

It is often assumed that having both life insurance and payout annuity blocks will automatically result in a reduction in uncertainty, as there is bound to be a highly negative correlation between them. This is not, however, always the case. Although in some cases, the identical customers may have purchased both life insurance and payout annuities, in others different risk characteristics between the two groups may exist, with the result of limited correlation between them.

Studying or communicating aggregated mortality and longevity experience thus would make little sense, although comparisons between the level and trend in mortality rates in either case might provide some insight into the relative effect of selection, self-selection, and anti-selection, as well as mortality differentials in the population segments. Note that in terms of overall financial effect the two offset each other – the extent of offset is useful information.

The attention paid to mortality and longevity risks depends on how significant these risks are to the overall financial performance or competitive position of the entity, in part a function of their strategy and risk tolerance. The needs of financial report users, underwriters and pricing actuaries (through premiums, dividends, non-guaranteed elements, re-pricing) and appropriate feedback depend upon the level or type of mortality risks an entity is exposed to.

Mortality (life insurance) should always be separately analysed from longevity (annuity), although if the identical insured grouping is being provided the mirror risks, together they form a perfect hedge. However, to the extent that the population underlying the mortality risk is different than that underlying the longevity risk, this offset will be less than perfect, even though these populations are exposed to the same perils and conditions.

It is important to choose the group of those subject to mortality and longevity risks studied by means of meaningful metrics and corresponding benchmarks to assess mortality or longevity experience. Which one(s) to use depends on the insurance coverage, although usually both are measured in terms of mortality rates. There are several base measures that might be used – those denominated in numbers and those in currency. For example,

- Mortality is often measured in terms of number of policies (or certificates in the case of group insurance) or insured lives. The former is often more practical, unless the insurer can through their systems tie all policies on a life together (and even then tying these together when they may have been issued at different times is problematic), although the latter may more accurately affect the exposures. Since in some administrative systems, riders look like a separate policy, care may be needed not to treat those as separate policies. Measurement by numbers is often used in a situation in which demographics are being projected, rather than the financial effect of mortality.

- Mortality is also often expressed in terms of the amount of death benefits payable or the net amount at risk (death benefit less reserve held or the amount that would otherwise be paid out in case the insured voluntarily lapsed at that time). Either of these measures better relates to the amount of insurance payable and thus the financial effect of mortality or longevity than numbers of policies or lives, as amount is most likely the unit for which the mortality/longevity will be applied. In addition, it has to be kept in

mind that due to the effect of education and income on mortality, as well as more intensive underwriting, mortality results by amount tend to be less than those measured by number.

Actuarial exposure theory properly addresses the timing of increases and decreases in exposures. In most cases, (1) increases (or "ons") are only exposed to mortality (that is affect the denominator of a mortality rate) from the date that an increase occurs, (2) decreases not due to mortality (the competing risks, such as voluntary termination or disability, as applicable) are exposed for that part of the period those lives are exposed, and (3) decreases due to mortality are exposed for the entire period.

Assurance that the underlying data is accurate is always a priority in preparation for a mortality study, but is also relevant in a comparison with a benchmark, whether it be based on a company's history, industry aggregates, or an external source. Data coding can change, cause-of-death coding can differ over time depending on the ICD(X) coding structure, or the basic source of data (e.g., in the U.S., the National Bureau of Vital Statistics collects death information, but exposures from the Census Bureau are not as reliable at the very old ages; thus, most U.S. population studies look to general population exposure at older ages determined by the Medicare's administrative data base, as that is perceived as being the most reliable source. Another example is the use of death data from Social Security, to which the Social Security Administration for the last several years all states has not been able to disseminate its data to become public, due to privacy and accuracy concerns, with a result being that the death recording for a line of business such as long-term care insurance, may not be accurate.

Almost as important as utilizing sources that provide accurate data is the relevance of the data. If the company's market is not the same as that of the external benchmark source, its use should be considered with caution.

Practical internal data warehouses are ideal sources of experience, as they can gather comprehensive data input linking claim and exposure information. Unfortunately in some companies, effective and efficient data warehouses are more talked about than being available to user areas, such as actuarial departments.

Unless the mortality or longevity risks begin at birth, a life expectancy at birth metric will not be an appropriate surrogate for these risks. For example, changes in early age death seen throughout the twentieth century may be over-weighted in life expectancy calculations and thus may not relate well to the risks being measured. However, in some cases life expectancy at a certain age, e.g., at age 65 for certain retirement programs or at 85 for long-term care insurance, might consolidate relevant risk experience into a reasonable single value.

3. Data grouping

In the analysis and subsequent communication of any mortality result, grouping of insureds is necessary, as insuring a pool of insureds is by definition a fundamental element of many insurance and benefit programs, the means by which these programs can operate efficiently and effectively. The pools or group of pools cannot be either too large as to provide misleading trends or too small so as to provide actuarially credible information regarding the mortality experience. Ideally the reporting unit would group heterogeneous groups that have similar expected costs or mortality rates, but if too small they can provide too much statistical noise and detail to be able to interpret properly. To assist in determining a reasonable grouping, the following is a list of factors that can be considered in a particular situation:

- Intended purpose and audience. If, for example, the information is intended for use by pricing actuaries or underwriters, the detail needed would be substantial and would normally be available on a drill-down basis, to enable them to ascertain whether existing pricing or underwriting programs (and their risk classification systems) have been or are expected to be successful or needing to be modified. In contrast,

for investors or boards of directors, the detail should be highly aggregated, possibly with only one two categories, and possibly further based on specific presentation of say, the operation in a particular country or life insurance and annuities. This is because their decisions are usually made at the entity, rather than a product level.

- Benchmarks used. It is usually more meaningful if a common metric is used. If an actual-to-expected metric approach is used, then it is more appropriate for a larger set of aggregations to be applied. Another approach is to apply age-gender-other variable adjustments to overall results.

- Size of groupings. If a feel for overall performance is needed, an entity-wide A-to-E ratio is useful, as long the Expected rates reflect appropriate demographics and risk classification. To the extent that assumptions apply to a range of types of exposures, a change in the mix of those exposures can correspondingly influence results.

- Ease of calculation. It goes without saying that comparisons and calculations should be made in a practical manner.

- Period of issued policies. Issue age cohorts can be important to categorize because (1) its use can facilitate a comparison of underwriting effectiveness and (2) in certain countries a definite cohort effect can be seen (e.g., the U.K.).

- Size of policy. As underwriting tends to get stricter as size of a life insurance policy is greater, segmentation by size of face amount grouping can be revealing. Since different socio-economic groups tend to buy life insurance to meet their need, larger size life insurance policies are usually found to experience lower mortality rates than do smaller size policies.

- The largest risk classification category is often used as a base grouping, with relativities measured with respect to the mortality experience of that category (see Table 2).

- Concentration of older age exposures. Given their large mortality rate levels, the financial risks of a deviation in experience of these exposures can be significant. However, their effect may be manageable because (1) in most cases the number of exposures in this age category is relatively minor in comparison with the entire inforce, (2) in many cases these policies do not have a large liability then, that is, due to inflation in average size policy and small net amounts at risk, and (3) discounting from the time of valuation and reaching these attained ages reduces the effect of their current value.

- Grouping by occupation grouping, especially in employer group situations, in which concentration in a particular type of employee, e.g., miners in mining companies can be used as a differentiator.

If not adjusted for by relatively detailed category, for instance, the classic life insurance groupings would classify individuals by age group, gender and underwriting class (e.g., standard, substandard; smoker, non-smoker). If not done by this or something equivalent, trends will be affected as much by changes in the mix of insureds than by changes in mortality rate performance.

A study of smaller groupings is especially important if mortality is expected to vary in a continuous manner. For example, in the study of obesity, the rate of mortality can vary by body mass index (BMI) – some studies have indicated a lower mortality rate at moderate (class 1) obesity levels than those at normal BMI, while those at more extreme (class 2+) obesity experience higher than normal BMI; but overall, mortality is higher than normal for all

those obese combined. The greater than average mortality rates for all obese can hide the favourable mortality experience for those at moderate BMI levels.

4. Benchmarks

The choice of the mortality benchmark can enhance or mislead an assessment of mortality results. For example, a benchmark can be taken from historical mortality rates for similar policies for a certain period or mortality rates from a reference population, such as those gathered in an inter-company experience study or mortality rates from the general population. Some observations regarding the choice of a benchmark follow:

- One or multiple dimensional table (usually policy years since issue)
- Traditional actuarial benchmarks include factors including age and gender.
- An expected benchmark will usually represent an a priori base for the grouping being studied, either from a prior experience study or that used as a parameter for the company's pricing or valuation models. These expectations can be based on the assumption made in these models or from the company's internal business plan if different – or maybe more than one of these if alternative sensitivity tests are conducted.
- The benchmark should be based on a population as similar as possible to the population being evaluated. As a result, usually the most preferable population base is that of the same business unit, given that the business unit encompasses the same or similar market segment. In some cases, especially where underwriting or the market segment is similar, experience of mortality of other business units or other companies can be used. If those data sources aren't available or relevant, typically adjustments to industry experience, e.g., those of the Society of Actuaries in the U.S., are applied. Less relevant would be the country's total population, although for trends this experience might be used. Other countries' experience might also be used, if an appropriate adjustment could be made between the expected mortality of the two countries. In small countries, such as in the case of a small European country, one must almost always use experience from other countries in some way.
- Credibility approaches are often used, either on the basis of a prior distribution (a Bayesian approach) or based on subjective weights. If these are applied, the choice of the other case can be even more important than that of the actual experience studied – in some cases the reference population rates are the prior assumption; thus actual experience is really being used to indicate how the prior assumption should be modified or updated.
- Trend. In any comparison of more than a year or two, consideration of de-trending the benchmark may be appropriate if prior years are being studied (that is, if the benchmark is applied over a long period, the benchmark for previous years are adjusted for trend, if appropriate).
- An actual to Expected (A-to-E) study or a Generalized Linear Model (GLM) approach may be applied.
- A raw or smoothed benchmark. In some cases, the benchmark (without a smoothing adjustment) can be quite erratic. As a result, smoothed results, especially when the benchmark is based on a short period or limited volume, can be used. Before the smoothed benchmark is used, a comparison should be made between the raw and smoothed results to ensure that the smoothing process doesn't introduce an unanticipated bias.

- Historical comparison. If numerical performance is provided, it is usually useful to provide some historical comparison to identify if a deviation has arisen. The most common approach is to compare results on a year-to-date (or rolling year) basis compared with the prior year – that can also automatically correct for any seasonality in deaths (e.g., more deaths can occur in certain months of the year). In most cases, a trend analysis can serve as a useful comparison; however, in case of extreme volatility, it might be appropriate to at least include two prior years of experience.

If the benchmark is some subset of historical experience, then the time periods of study can be quite important. The longer the period, the clearer trends might become. If the period of actual experience is quite short – that is, actual monthly experience, is usually quite volatile. Sometimes experience for the prior quarter or yearly periods is used. Rolling results, say of actual mortality of a rolling twelve months of experience is often a period of time that is used (if rolling results are displayed, the monthly results are also often displayed in the same chart).

A typical Actual-to-Expected analysis looks like the following table, with the top block conducted in terms of number of policies (or lives) and the bottom block in terms of death benefit (or net amount at risk):

Table 1 – Template for Actual-to-Expected analysis

	Males				Females			
		Deaths		Actual /		Deaths		Actual /
Age group	Exposure (no.)	Actual	Expected	Expected	Exposure (no.)	Actual	Expected	Expected
< 40								
40 - 49								
50-59								
60-69								
70 +								
All								
		Deaths		Actual /		Deaths		Actual /
	Exposure (FA)	Actual	Expected	Expected	Exposure (FA)	Actual	Expected	Expected
< 40								
40 - 49								
50-59								
60-69								
70 +								
All								

Another approach would be to use the largest class as a benchmark, after age adjustment (which requires a standard assumption regarding the distribution of exposures – typically this would the distribution of exposures for the largest class). In the case of Table 2, the largest class is the non-smoker preferred class.

Table 2 – Relativities to current year base underwriting category experience

		Males					Females				
Risk classification		**Current year**					**Current year**				
Underwritten		**1 - 2**	**3 -5**	**6-10**	**11-20**	**20 +**	**1 - 2**	**3 -5**	**6-10**	**11-20**	**20 +**
Non-smoker	Preferred	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Smoker	Preferred										
Non-smoker	Standard										
Smoker	Standard										
Non-smoker	Sub-standard										
Smoker	Sub-standard										
Term conversion											
Group conversion											
Risk classification		**Prior year**					**Prior year**				
Underwritten		**1 - 2**	**3 -5**	**6-10**	**11-20**	**20 +**	**1 - 2**	**3 -5**	**6-10**	**11-20**	**20 +**
Non-smoker	Preferred	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Smoker	Preferred										
Non-smoker	Standard										
Smoker	Standard										
Non-smoker	Sub-standard										
Smoker	Sub-standard										
Term conversion											
Group conversion											
Non-smoker Preferred Current / Prior year											
		%	%	%	%	%	%	%	%	%	%

Often, depending on size, company studies are conducted every three or five years (or a study of the same length of time with the last year of the prior study replaced by the most recent year, so the study period becomes a rolling one), either to increase the credibility of the actual results or enable the actuary to examine risk classification and underwriting / marketing-related variables in detail with greater credibility, where a study of one month or one year in length would not provide credible data for this purpose.

5. Effects of selection and anti-selection

The effect of selection by an insurer by means of underwriting (either by the underwriter based on medical or non-medical information, or by means of a simple questionnaire, or by means of field underwriting (selection by an agent/broker who has experience with the company's underwriting philosophy)) is often considerable. In some cases there is asymmetric information, that is, the applicant knows more about his or her own health than the insurer could ever find out through underwriting means. This asymmetric knowledge has an offsetting effect.

After policy inception, insured anti-selection can also have an effect on mortality experience. That is, those healthy lives that could purchase a more competitive product several years afterwards tend to lapse their policies, with the remaining pool experiencing on average worse mortality experience.

For payout annuities, policyholder selection is usually present, in that only a good mortality risk would purchase such an annuity, especially those annuities with a short certain period. Self-selection can also occur up-front if a company is known for tight underwriting – only those who expect to pass the underwriting screen will bother to apply to that company.

Depending on the situation, the combined effect of selection/anti-selection can result in a select-and-ultimate period, with the select period being different depending on the situation. In some cases, this period has been shown to be up to twenty five or thirty years in length, although the period is usually shorter at older issue ages.

6. Other aspects of mortality analysis

Mortality study by cause-of-death and risk characteristics of population being studied at issue or observed (e.g., age, gender, marital status, past or current smoking habits, immediate cause-of-death, smoking habits, occupation, physical fitness, weight, and fatness) can provide information that affect pricing, underwriting, and dividend policy.

Cause of death can be assessed at at least two levels. First by proximate cause, that is those identified in the insured's death certificate, e.g., cardiovascular disease, malignant neoplasm, respiratory disease, diabetes, Alzheimer's disease, and accidents. Or second it can be looked at through basic drivers of the proximate cause, providing an early warning of deaths, including such factors as smoking, cholesterol, blood pressure, obesity, and fitness. Although the second type of causes may provide an earlier warning of premature death, death-attribution by these causes can be more difficult.

Actuarial analysis has often been conducted one or two variables at a time. For example, mortality has often been conducted by quinquennial or decennial issue (or attained) age group, gender and policy year. Depending on the size of the block of business, this multi-variable approach may lack statistical credibility in many categories. This credibility is then determined on the basis of the number of claims involved. Other statistical approaches, including multivariate regressions can provide supplementary analysis, especially when looking at detailed segmentation of the data. Thus, techniques that simultaneously solve for the effect of multiple demographic, underwriting and environmental variables, such as a Cox proportional hazards model, commonly used in epidemiological and general population studies, can be used.

Stochastic mortality modelling is becoming more common, in part to better reflect low frequency high severity pandemics. This more refined modelling points to the need for assessment of interactions and correlation among the contributing causes. For example, those who are obese also have a larger number of co-morbidities, so that being obese is inter-related with other factors, such as diabetes.

The lack of credibility (reliability) of the experience data can be an extremely important issue to address. Using apparent relativities or trends in experience (especially when the population is studied in a sliced and diced manner), unless careful and particularly in the application of risk classification relativities where a limited volume of data is available, can lead to incorrect decisions. Patience (or reliance on external data) and professional judgment is needed in many studies. Disclosure of data (and study) limitations can in some cases be just as important as the results themselves.

Outlier experience (usually small blocks with extremely high or low A-to-E ratios) are sometimes ignored because of small sample size, or if the sample size is so small that little credibility is given to the mortality experience of those blocks. In some cases, further research is warranted, as data problems may have arisen, and are generally worth a little more effort to investigate.

Average values are used in most analysis. Actuaries often rely on expected values (means), although depending on the circumstances, a median value may be more relevant, especially when there are a handful of large claims. In fact, skewness of results is a common reason why results expressed in terms of amount of insurance are different than those expressed in terms of numbers of policies. This skewness can be dealt with in some cases by showing experience with and without the outlier(s), although in some cases there seem to always be an outlier of some sort every year.

7. Communication of mortality experience

The most effective approach to communicate mortality/longevity experience depends on who the users are who receive this information and the expected use of the data. Most managers want to see what they have been used to seeing (inertia), but in most cases will be easily converted to more graphical presentations (however, some actuaries will always prefer tables with numbers).

It is important to determine the relative value of graphical vs tables to the users. Normally, graphical displays supplemented by highlights of them in accompanying text are more effective in communication to busy people. Nevertheless, it is also usually necessary to produce a comprehensive report documenting more fully the findings of an annual study, most importantly accompanied by (or made available upon request) documentation of the process followed, data issues identified and their resolution, and recommendations for future studies.

There are various ways to incorporate the study of the risk characteristics and possible other variables that may be available. One approach is to select the largest homogenous category(ies), such as preferred non-smoker underwritten male experience. And then study relativities (e.g., Actual-to-Expected analysis, age-adjusted mortality rates, or Cox proportional hazard ratio analysis) to this benchmark or assume the historical mix of insureds has been constant over time. Stratification by key variables needs careful planning, with the A-to-E analysis is particular useful in studying the effects of relatively small insured sub-populations.

The timeframe over which trends can be compared usually incorporate individual period (e.g., weekly or monthly) values that tend to gyrate up and down, together with the more smooth rolling average values. If the comparison isn't direct between a current period (less than a year) and a corresponding period in prior years, seasonal adjustments might be needed. Unlike morbidity, because of the relatively quick settlement of most death claims, claim backlog adjustments are unusual and if they do occur are reflected in the incurred but not reported (IBNR) liability, although explanation is needed if significant shifts in claim backlogs do arise.

A technical and infrequently significant issue in mortality reporting is whether to use information based on processing date or effective date of the transaction (e.g., claim incurral date). The accurate way is, of course, to use the date of death rather than when the death is first picked up by the valuation system. Actuarial experience analysis will almost always be based on effective dates. However, because reported financial reports have to use reported dates (with estimated adjustments approximate the effective dates), estimates of the effect of the reporting and processing lags on incurral date data and operational statistics are made using reporting dates, inconsistencies in the effect of these values regularly arise (although in most cases, they are not significant). A decision is needed as to how best to report on these.

A related issue is whether to study mortality on a calendar period or policy period basis. There are advantages to support the use of either – calendar period studies are more up-to-date (using the most recent bit of

information) and more consistent with financial results, while requiring consistent treatment of the beginning period and end-of-period results, while policy period (especially covering policy periods) tend to be able to overcome problems such as varying claim backlogs, seasonality and statistical fluctuations given the stub (partial) policy year periods.

Although graphs can be quite useful, care is needed in their construction, as it is quite easy to leave an incorrect impression. An example is shown in Figure 1, in which the vertical bars begin at different points.

Figure 1 – A different perspective to the same data

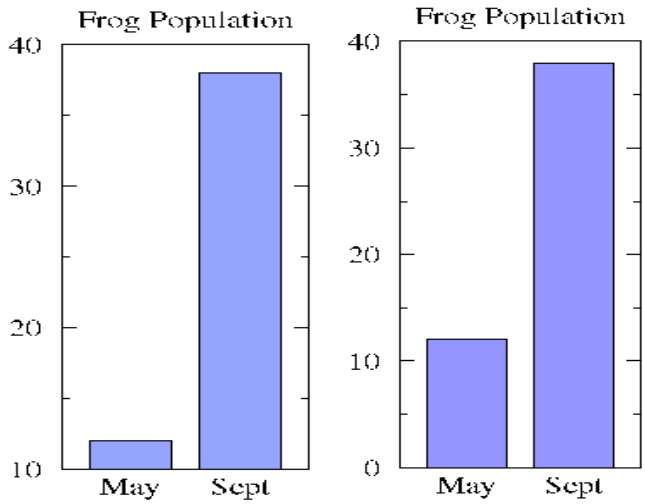

This is an example of how to "lie with statistics", unintentionally or intentionally, emphasizing stability or volatility of results. Other examples that need care, depending on the sophistication of the audience, include the use of logarithmic versus uniform scales.

The following charts (connected from left to right), based on graphics previously provided on the BBC website, demonstrate how volatility can be highlighted or diminished, depending on the scale used. The top charts are cosmic radiation rate in neutrons per hour, while the bottom charts show temperature and temperature change since 1975. The top charts demonstrate in a real life example, the same presentation effect as the above charts of the frogs, while the bottom charts compare the trend in temperature in contrast with the change in temperature on a very different scale. This suggests the importance of reporting on the financial effect of these trends to ensure that misleading results are not given and an appropriate level of attention is given.

Figure 2 – Two sets of two graphs providing different messages based on the same facts

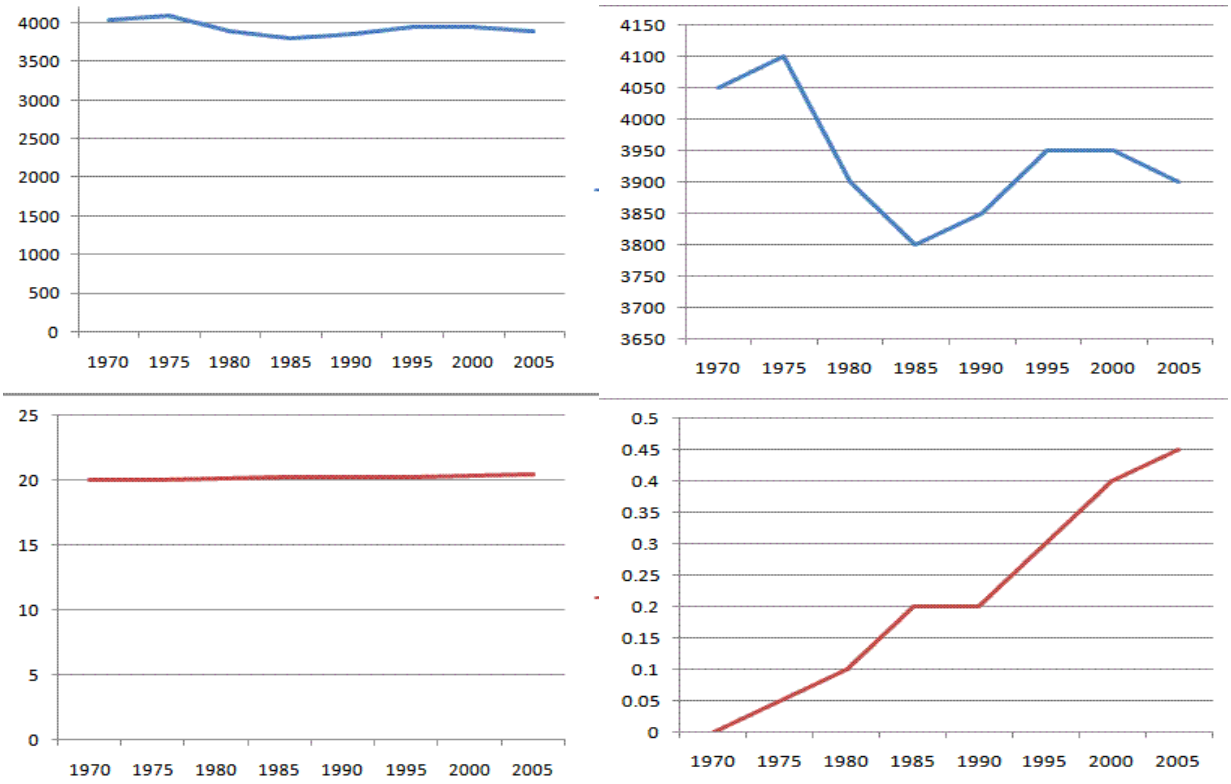

Heat maps are especially revealing if cohort effects have been observed to exist or other two-dimensional relationships are important. These can also illustrate the effect, if any, of changes in underwriting guidance.

8. A dashboard

A dashboard usually consists of a set of information presented in a manner that highlights key performance indicators (KPIs), primarily in terms of comparison with a set of benchmarks, whether historical trend, budget, pricing assumptions or liability assumptions. Although often one page in length, it doesn't have to be so restricted. However, since it is intended to overview of multiple performance factors, it is common that only one or two graphics (or a couple of lines of numbers) would be committed to an assessment of mortality performance. In addition, because dashboards are usually produced frequently, care is needed to over-reaction to volatile weekly or monthly swings in experience (or better yet, don't include mortality data in daily or weekly dashboards). On the other hand, if mortality is an important component in earnings or if prices are sensitive, a separate mortality performance dashboard page might be appropriate.

The objective of a dashboard is to serve as a source of both early warning and actionable information and an easy-to-access source of decision-useful reference information to its users. It isn't intended to provide detail, although drill-down capabilities might be provided digitally. Experience from recently issued business might be shown separate from experience from the entire inforce, as that is more likely to be indicative of recent management decisions.

The following are illustrative metrics in graphical or tabular format that might be included:

- Monthly number (and/or amounts) of death claims submitted or processed
- Year-to-date (or rolling) actual-to-expected mortality and longevity, and for major business units
- Source of earnings due to mortality/longevity

For some products where it is difficult to separately identify mortality from other policy terminations, such as is the case for long-term care (LTC) insurance, it may be more appropriate to monitor the combination of deaths and voluntary lapses. Or for long-term health insurance, show deaths separately for active lives and those in disabled lives (especially relevant since for LTC and disability income coverages, deaths more closely relate to a longevity rather than a mortality risk.

It is important to periodically request feedback relating to questions about these reports during the past month for which the dashboard did not provide desired perspective or insight, although in some cases a dashboard is intended to raise rather than answer questions. Nonetheless, it is rare that such comments/suggestions are provided without some prodding – you'll never find out if you don't ask. Similarly, in developing a dashboard, it is difficult for most people to dream up what they need without seeing live data in a report format – thus, pilot versions can be useful before new or revised reports "go live".

Consistent with good management information practise, it can be good practice to not send out a report for a while or send it out late, to see how many inquiries are received about its whereabouts – a lack of requests for the report may provide insight as to its actual value ("nice to have" in contrast to needed information).

9. Trends and projections

A projection of mortality and longevity experience can be approached in several ways. To develop proper input assumptions for such a projection, a common initial step is to estimate the current mortality or longevity level, as applicable, of the relevant population. This can be derived from observations of the current population (or at least a reasonable estimate of it based on an extrapolation from the recent past, for example, from a recent experience analysis. If this isn't possible or practical (e.g., this is a population that has not yet been observed or does not have a credible size), estimates based on experience of similar populations with suitable adjustments. A recent estimate could also serve a base, suitably adjusted to represent current conditions and populations. An applicable estimated trend to the current valuation date should also be applied, if appropriate. This estimate would be made based on a suitable number of risk classification variables (e.g., age, gender, cohort, time since policy inception, family history, marital status, and underwriting classification).

The projection would then be prepared using suitable trends. These trends could consist of an aggregate factor or be nuanced by category and time from the valuation date. Considerations regarding the development of the applicable trend include the reliability of any recent company experience mortality rates derived and the degree of confidence with the current assumptions. The projection could be derived from those currently used in pricing, valuation or another assumption can be used as a base. A common statistical approach used, with or with the major refinements that have been offered is the Lee-Carter approach, which assumes that recent trends will continue.

Other approaches are also possible – in the U.S., several annual trends developed by experts are available, especially as applied to annuities (alternatively sources such as the Office of the Chief Actuary of the Social

Security Administration publishes a set of mortality improvement rates it uses), but in any case the actuary should assure himself/herself that the trends are appropriate for the population being dealt with.

The trends could incorporate a cause-of-death approach, at least for a certain number of future years, say fifteen years. Adjustments may be appropriate for any given method, possibly affected by known influences, such as the expected effects of previous changes in smoking prevalence. Parametric approaches have also been applied, for instance, Gompertz and Konnisto. Hybrid approaches are also used, such as a by-cause projection over the first five or ten years, with an overall trend assumption thereafter, possibly varying by age and gender.

If improvement has been assumed, a key parameter is the time period over which it is applied. Practice has differed, with some projections incorporating an annual trend that continues for the length of the projection period, while some actuaries are reluctant to incorporate this type of trend for such a long period (at least for mortality in contrast with longevity projections). Some actuaries taper downwards their mortality trend projection over ten, fifteen or twenty years, or possibly at a very high attained age, such as 95.

An actuary is always challenged if the company does not have adequate and relevant experience base for direct use to develop a mortality / longevity assumption. Careful consideration of adjusted industry or population data (or the entity's experience in other lines) would normally be called for. In addition, in developing a trend assumption, an analysis of changes in conditions between the current date and the expected future should also supplement the development of these assumptions. Comparisons to trajectories of a relevant population of a recent period would also be conducted, if practical. Sensitivity testing relating to the desired application, e.g., pricing of an insurance product or valuation of a financial statement policy liability.

Whatever approach is taken to develop assumptions, it is good practice to establish a controlled governance process, incorporating a process to validate results, obtain internal approvals (hopefully independent of the person or area whose responsibility it is to develop assumption recommendations), and peer or other external review.

A report providing detailed documentation of the process used to develop experience studies or such projections need to be prepared. It should include a crisp and to the point Executive Summary – for at least this part of the report, it should not assume that the reader would be an actuary. Items included in the summary would usefully focus on the study's key findings, significant changes from the previous report, a summary of recommendations, and the estimated financial effect of new recommendations, if applicable.

10. Annual Report

It may be somewhat surprising how little information regarding mortality results is included in many life insurers' Annual Reports. Although qualitative discussion of mortality/longevity experience and risks are usually provided, the space allotted to a discussion of mortality performance is typically quite small in comparison with that of assets and investment risk. Some speculate that this is because there is significantly more volatility and risk associated with financial instruments than in pooled insurance risks such as mortality. In general, more can be done, depending on the relative sources of profit and risk attributable to mortality or longevity.

Aspects of mortality/longevity that might be covered in more depth in Annual Reports include:

- A source of earnings analysis. This would represent a decomposition of income, including returns from mortality, longevity, morbidity, policyholder behaviour, expense, and investment returns.

- Comparison of actual-to-expected performance. The usefulness may depend upon the basis for the expected mortality. For example, because of selection effects and mortality trends, comparison of reliance of the current experience to that of locked-in mortality assumptions of thirty years ago would produce such low ratios that the resulting A-to-E ratios would be misleading, at best. However, comparisons with current valuation, pricing or dividend assumptions may provide a useful performance.

- Facts and circumstances dictate the level of depth appropriate to cover. For example, in addition to separately identifying mortality from longevity results, experience be provided on the segment or major lines of business level, consistent with discussion or disclosure of overall results.

- Sensitivity of key financial indicators to alternative mortality and longevity assumptions which can provide users in focus.

Overall, external Annual Reports have provided limited quantitative information regarding the effect of the sensitivity of mortality performance on income or capital. A few examples of such information provided by life insurers include the following:

(1) A 10-K report (to the U.S. Securities and Exchange Commission) provided an illustration of the effect on the deferred acquisition cost (DAC) asset and the unearned revenue reserve (URR) of a one percentage change in aggregate mortality rates,

(2) An Annual Report showed the results of a sensitivity test on net income and shareholders' equity of a ten percent change in mortality rates, and

(3) An Annual Report showed the sensitivity to the insurer's economic capital of a lethal pandemic. It seems that further development of disclosure information regarding mortality and longevity is needed.

11. How to communicate uncertainty

Future mortality experience is certainly uncertain, although knowledge of a credible base of experience from a similar population can reduce this uncertainty significantly. A set of unbiased expectations (sometimes referred to as a best estimate, best estimate or expected value) is determined, derived by means of a set of assumptions. Uncertainty (that is, the lack of knowledge of the best estimate) can be contrasted with the risk of volatility, which represents the statistical fluctuations around the mean due to the limited size of the experience base) arises because future conditions are not known for certainty at the time the projection is made.

Volatility or differences from expected mortality or longevity can be caused by several types of factors, including: (1) inappropriate expectations, (2) being part of a longer-term trend, (3) a change in mix of risk characteristics of exposures subject to mortality or longevity risks as a result of changes in such factors as

underwriting, distribution, or lapses, (4) identifiable non-recurring causes, including a change in mix of a certain type(s) of claim, such as deaths of insureds with large benefits, or (5) random fluctuations.

Analysis of results (or just a passage of more than a single period) may be needed to ascertain the primary sources of the deviation. Where possible, management information should be designed to lead to action or a feedback loop—however, in some cases this information is useful just to deepen the understanding of the financial results of the company, while in others, once this understanding is obtained corrective actions (e.g., changes in premiums, dividends, or liability assumptions) can follow. Although it would be desirable for action trigger points (that is, a pre-determined trigger level such as a 10 percent change in Actual-to-Expected ratio) to be set, before it is understood what type of causes are involved, it may not be appropriate to take action until identification of the sources of change or differential is obtained.

It is generally accepted actuarial practice to provide the user(s) of a set of actuarial projections information regarding this uncertainty based on a range of key assumptions, both individually and collectively. These ranges are not intended to indicate the best or worst possible scenario; rather, they are developed to provide the user with an understanding of an alternative set of reasonably possible scenarios.

Sensitivity and scenario testing of some type are typically conducted. These are usually presented in any recommendation of assumptions intended for an internal audience (e.g., for pricing, valuation or financial reporting) to provide decision-makers a feel for the sensitivity to alternative set of assumptions. As indicated above, based on a limited informal review of public reporting, relatively limited sensitivity test results are published, although there is usually significant qualitative discussion of associated risks.

Given that many observers believe that the most significant areas of uncertainty are (1) uncertainty as to the expected level of mortality (e.g., as a result of being in a new market, having implemented a new risk classification system, or another situation with incomplete information about current mortality or longevity circumstances), (2) mortality trend (result of changes in conditions, such as a medical breakthrough), and (3) pandemic risks (often one-off fluctuations), consideration might be given to include one or more sensitivities using alternative scenarios for overall mortality (e.g., ±5%), an adverse trend (e.g., ±1% annually), and a pandemic (+20% in the next year; common practice is to use the 1918-19 Spanish flu epidemic as an example, although a different epidemic structure may be more appropriate depending on the distribution of exposures). In addition, it would be useful to identify the recent trend in mortality/longevity has been over the recent past to put the selected scenario in perspective, e.g. what has been the aggregate trend over the past five years.

To put currently provided information into context, the proposed disclosures are given in the Exposure Draft of the International Accounting Standards Board's project on *Insurance Contracts*[1] on the nature and extent of risks that arise from insurance contracts. "Information about the nature and extent of risks that arise from insurance contracts to enable users of financial statements to understand the nature, amount, timing and uncertainty of future cash flows that arise from insurance contracts. … and information about insurance risk on a gross basis and a net basis, before and after risk mitigation (for example, by reinsurance), including information about: (a) sensitivity to the insurance risk in relation to its effect on profit or loss and equity. This shall be disclosed by a sensitivity analysis that shows any material effect on profit or loss and equity that would have resulted from: (i) changes in the relevant risk variable that were reasonably possible at the end of

[1] International Account Standards Board's Exposure Draft on *Insurance Contracts* (June 2013)

the reporting period; and (ii) changes in the methods and inputs that are used in preparing the sensitivity analysis. However, if an entity uses an alternative method to manage sensitivity to market conditions, such as embedded value analysis or value at risk analysis, it can meet this requirement by disclosing that alternative sensitivity analysis and (b) concentrations of insurance risk, including a description of how management determines the concentrations and a description of the shared characteristic that identifies each concentration (for example, the type of *insured event*, geographical area or currency). Concentrations of insurance risk can arise if an entity has underwritten risks that: (i) are concentrated in one geographical area or one industry…."

This suggests that, at least for the communication of mortality results and risks, in the future we will see an expanded set of future disclosures.

In assessing uncertainty, a probability distribution for the mortality rates (or trends in these rates) needs to be developed. This may take the form of an asymmetric distribution, as an adverse scenario is the normal concern of management, regulators or investors. Although this distribution could be developed by means of observing past volatility, it could also be developed using one of several possible standard probability distributions (possibly augmented by means of a calamity / catastrophe distribution. Either a Monte Carlo simulation model or representative selected values (found from the distributions developed) might be selected to be applicable to the conditions being addressed.

Concentration risk of many insurers does not often represent a significant systemic risk to the entity, but periodic studies of concentration of mortality exposures may provide useful information of the total potential concern to management, including possible effect on such entity-specific factors as a decrease in the numbers of agents in the company's distribution system and a slowdown in the economy.

In a few cases, insurers have hedged their mortality risk, either through reinsurance or mortality derivatives. In neither case is the hedging perfect, the former because of counter-party credit risk, the latter because the risk covered, usually an index based on industry or population trends, may not match the experience of the company. Any remaining uncertainty still needs to be assessed.

Once the distribution or alternative scenarios are selected, a reasonable confidence range should be developed and properly communicated to convey to the user the level of credence that should be given to the estimates. In addition, depending upon its application, a provision for risk and uncertainty (sometimes referred to as a risk adjustment or provision for adverse deviation) should be calculated. This may be subject to external rules or principles (e.g., for regulatory reporting or pricing) or be determined based upon the risk aversion of the entity (or if being made in conjunction with pricing, together with considerations of desired pricing strategy that considers price sensitivity of the target market and level of competitors' rates). For example, a 99% Tail-VaR or a 99.5% VaR might be selected as a risk measure – it is likely that testing of alternatives would be conducted, with a comparison of its effect on prices (or liabilities if done in conjunction with financial reporting). In any event, approaches consistent with the treatment of uncertainty with respect to other assumptions and the overall risk preferences of the entity are useful considerations as well.

Practice differs regarding whether or how to incorporate any such margin for uncertainty on an implicit or explicit basis. An explicit margin could be developed either by means of an explicitly determined mortality margin (e.g., flat, percentage, or trend; via an overall factor or one that varies by risk classification, product or demographic characteristics) or as input to an explicit overall product margin, reflecting the statistical studies

mentioned above. Some believe that a separate provision for uncertainty is preferable to an implicit margin to remind people of the uncertainties of the projection or to facilitate a discussion of the risks and uncertainties the company has and is expecting to undertake of, say, five percent of mortality.

Communication of this uncertainty can be done in several ways. These include displaying the confidence intervals on both sides around the best estimate in terms of values or in graphical form. A more effective means is to illustrate the margins on the same graph as historical experience rates are shown.