WHITE PAPER

INSURANCE RISK MANAGEMENT

SIFTING MEANING FROM THE MIX – MANAGING RISK DATA FROM DIFFERENT SOURCES

BY JOHN WINTER,
DIRECTOR OF PRODUCT MANAGEMENT, RISK MANAGEMENT,
FIS' INSURANCE BUSINESS

Contents

Introduction

For financial risk in insurance, convergence is the new black. Many organisations have different departments looking at different aspects of risk – Finance, Credit, Cash Management, Asset Management and the actuaries all have their own view of a subset of the risk facing the organisation. In many organisations, this "natural separation" carries through into the risk mitigation strategies, with independent measures in place to contain each different risk class, often managed by separate teams. But the real world is often not as clear cut as this neat picture implies. After the financial crisis, there was intense debate about what TYPE of risk had led to the problem – was it concentration risk, or operational risk, or perhaps just straightforward market or credit risk? What these arguments served to illustrate was only that risk should not be approached as a series of silos – and that risk managers would be better off adopting a more holistic view. To quote Douglas Adams, risk managers should consider "the Fundamental Interconnectedness of All Things".

The regulators have certainly taken steps towards this, with regimes such as Solvency II or C-ROSS providing frameworks into which firms can insert the different risk classes in order to get an overall capital exposure. Solvency II makes provision for currency stresses, market stresses, and insurance-related stresses, such as longevity and mortality, as well as internal risks, such as operational risk. They have also driven the responsibility for risk assessment into the business units, looking for evidence that risk has been considered in the regular business decision making process, and demanding evidence that business units have identified (and mitigated) their own risks.

The applications that measure and manage risk need to reflect this trend, and solution providers need to ensure that the information held in their applications can be easily extracted, compared, and mixed together to give the overall picture required to guide a modern business.

JOHN WINTER
AUTHOR

About the Author

John Winter is the director of product management for the risk function with FIS' insurance business. In this role, he is responsible for developing the ongoing strategy for the Prophet suite of products, in particular the Prophet Data Management Platform, and also the wider integration of Prophet into FIS' risk management suite.

Prior to joining SunGard, John spent 6 years at Algorithmics (now part of IBM), heading up their operational risk practice. During that time, he developed his understanding of risk, and in particular how different types of risk interact.

Prior to Algorithmics, John worked at NatWest Bank, Siemens, and Logica. In total he has more than 30 years of experience in software developed for the finance industry.

Assets and liabilities

In order to ensure the long-term stability of the insurance industry, regulators need to satisfy themselves that insurers are solvent, and will remain so over the duration of their liabilities. Solvency II introduced the famous Asset/Liability picture shown in Figure 1. this shows how a firm's assets are matched against the "Best Estimate Liabilities" to give reassurance that the capital held is enough to ensure the solvency of the firm. This balance has to be maintained through a series of prescribed stresses designed to test how well a firm could respond to shocks like sudden currency shifts or asset value collapses.

Even if the firm should get into difficulties, the technical provisions include an additional sum to compensate a potential purchaser for the additional capital they would need to cover the liabilities taken on in the event they stepped in to rescue the firm – the Risk Margin.

This picture lends itself quite well to the idea that the right hand side could be calculated by the actuaries, and the left hand side managed by the fund managers, and all the compliance officer needs to do is subtract one number from the other. Simple!

In reality, there is an enormous amount of complexity in the calculation of the Best Estimate Liabilities (BEL) and the Risk Margin, not least caused by the interplay between assets and liabilities as the insurance liabilities age. This is particularly true of with-profits business, or business with guarantees or other non-linear cash flows. They must account for human behaviours – for example, management decisions will set different bonus levels depending on the prevailing economic conditions and policy holders will react to these bonus payments in different ways, normally reflected in changes in the lapse rates or the rate of acquisition of new business. There are choices to be made for each future period – is the surplus to be paid out as a dividend? Should any surplus be reinvested, and if so, what in? This means that a lot of assumptions about asset behaviour are built into the BEL calculation.

Of course, because the projection period for these liabilities can be anywhere from 10 to 100 years (particularly in the life insurance industry, less so in non-life), the asset behaviour is modelled at the level of generic assets, rather than specific asset values – "Corporate Bonds" rather than "Bond Finance 1988 6¼ Perpetual". But when calculating the actual value of the assets held by the insurer, the asset manager will use the real holdings. They will also select the real assets in order to maximise the return on investment for the same level of risk – optimising their portfolio to better play the market. This requires the ability to forecast probable market behaviour, and also stress various risk factors to ensure that the level of risk taken does not change significantly.

Figure 1: Solvency II balance sheet

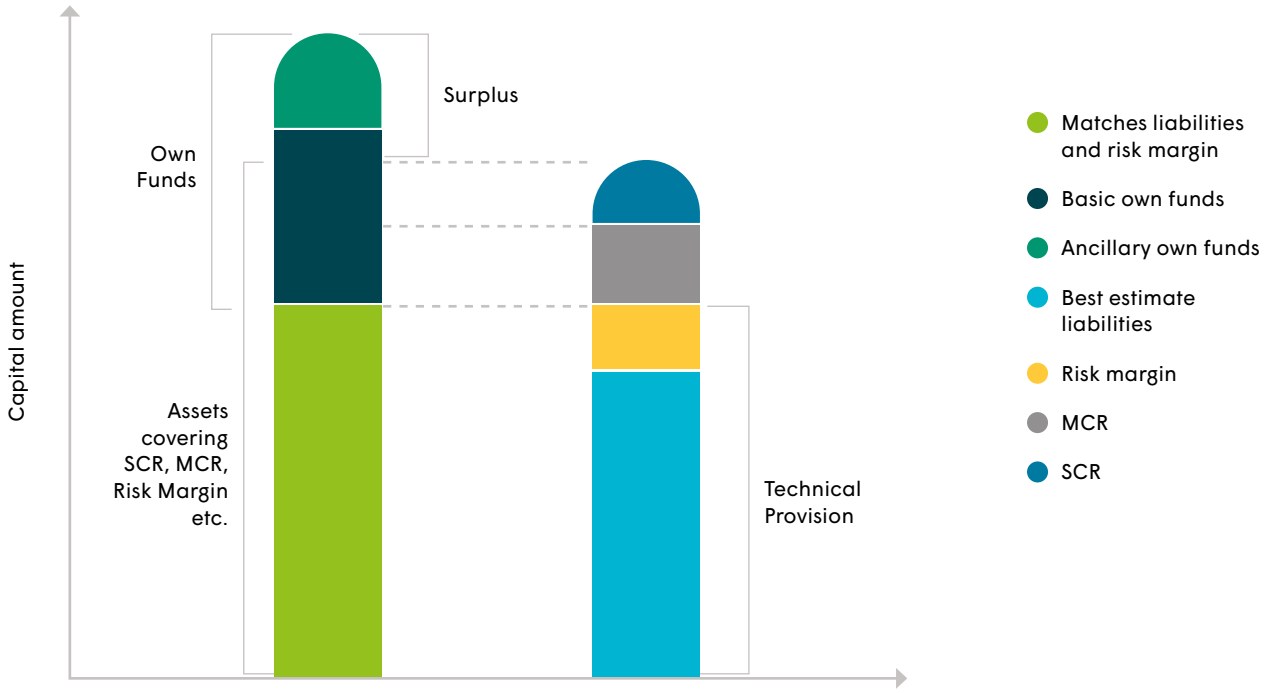

Applying stresses to the assets to see how they behave relative to the benchmark means that the boundaries of validity for the BEL data passed to the asset manager may be broken inadvertently. For example, if the actuary passed the 25th, 50th and 75th percentile cash flows to the asset manager, they might assume that a portfolio that matched well through a range of risk factor stresses would be good enough. However if the actuary passed them the 95th percentile values, these same stresses might highlight a massive potential loss. So there is an interaction between these two worlds which can be missed by the siloed approach.

Another important consideration for insurers is the capital efficiency of their chosen investments. In an attempt to immunise insurers against market risk, the regulators have defined capital levels deemed sufficient for each class of asset held. For example, if an insurer holds only government bonds, the level of capital they must hold to guard against a loss in value of those assets is significantly less that if the same value was held in equities. To complicate matters more for the insurers, the regulator insists that any fund for which the contents cannot be determined – true for many external funds – the worst case assumption must hold – in other words, the insurer must treat these funds as if they were completely made up of equities, and set aside the maximum capital reserve to cover them.

So decisions made by the asset manager to maximise returns can have a detrimental effect on the insurance company's capital requirements. It seems that there is little choice but for the asset managers and the actuaries to interact. Which means the systems on which they both rely should do so too.

Actuaries and asset managers – the great divide

Insurers structure themselves in different ways to deal with the split between assets and liabilities. Some, generally larger, insurers will have their own asset management function, taking data from the actuaries and performing the optimisations described above solely for their own liabilities. Others will pass the task to one or more external fund managers, describing their future cash flows to the external supplier, and looking for recommendations for investments which will provide the necessary cover, whilst hedging their risk as far as possible, and remaining efficient in their use of regulatory capital.

In many cases, the liabilities are passed to the asset managers in the form of projected future cash flows. This is the natural language of actuarial systems, which excel at processing millions of policies of many different types to obtain an overall picture of the financial health of the organisation as their book of business ages. One of the key inputs into these systems is the assumptions, which can take the form of single values (e.g. an assumed inflation rate), a range of time dependent values (e.g. a yield curve for bonds), or policy dependent data (e.g. specific mortality tables). The process may also contain stochastic elements, where the system simulates a range of possible outcomes, and produces values with an associated level of confidence (e.g. the 99.5% percentile).

Figure 2 shows a small set of insurance products projected forward over 20 years. These cash flows would be taken as the "benchmark" by the asset manager, who would attempt to find a series of assets which can

a. Produce an opposing series of cash flows, leaving the insurer with a balanced portfolio.

b. Maximise the available returns on the invested amount – increasing the overall profit of the firm.

c. Reduce as far as possible the regulatory capital needs required to cover the investment risk.

This fitting process is not a static one, as the asset manager has to assess how the portfolio behaves under a series of stresses, some of which will mirror those used by the actuaries, and some of which will be down to the asset manager's judgement. In most cases there will be a range of different potential investment portfolios to try to meet these often opposing goals.

For insurers with an internal asset management function, the asset management and actuarial departments will report their risk levels independently to the risk management group, who will merge the data as part of their reporting process. The merge process can often be unaware of differences in the underlying assumptions. For insurers using external fund managers, this is made more complex by the different recommendations and different data formats provided by each fund manager, and the different levels of transparency they provide about the contents of their selected portfolios.

The brief interaction described above is a very simplistic approach to modelling how an organisation responds to financial stress, and, given the sophistication of the models used today and the technology available, there are clearly ways in which it can be improved on. So what can be done to make this process better, and add value to the modelling process?

Building bridges

The following sections describe some ways in which the activities and models of the actuaries and the asset managers can be brought together, with a view to improving the ability of the business to meet its objectives, and also aiming for a level of consistency between the two that allows the risk manager to combine the data in a meaningful way on reports and dashboards. This is intended as a "discussion starter" and is by no means a comprehensive list. The order of the sections does not imply a relative importance – each organisation will have different priorities and different opinions, making a fixed sequence impossible. However there are enough areas to indicate that some substantial improvements can be achieved by convergence of models, data and assumptions.

Common economic scenarios and assumptions

When actuarial modelling systems project cash flows and asset values forward in time, they do so using a set of assumptions (for example, the inflation rate, or the target return on investment) and a series of economic scenarios. These scenarios are normally provided by an external vendor, who will undertake to provide scenarios which are properly calibrated – that is to say, they exhibit consistent, correlated behaviour between related indices (for example, equity values are more volatile than bonds, and when bonds and equities fall, commodity prices tend to appreciate, and so on). Actuaries will make use of sets of these scenarios, which are used particularly with stochastic models to try a range of possible outcomes.

Figure 2: Projected cash flows from insurance products

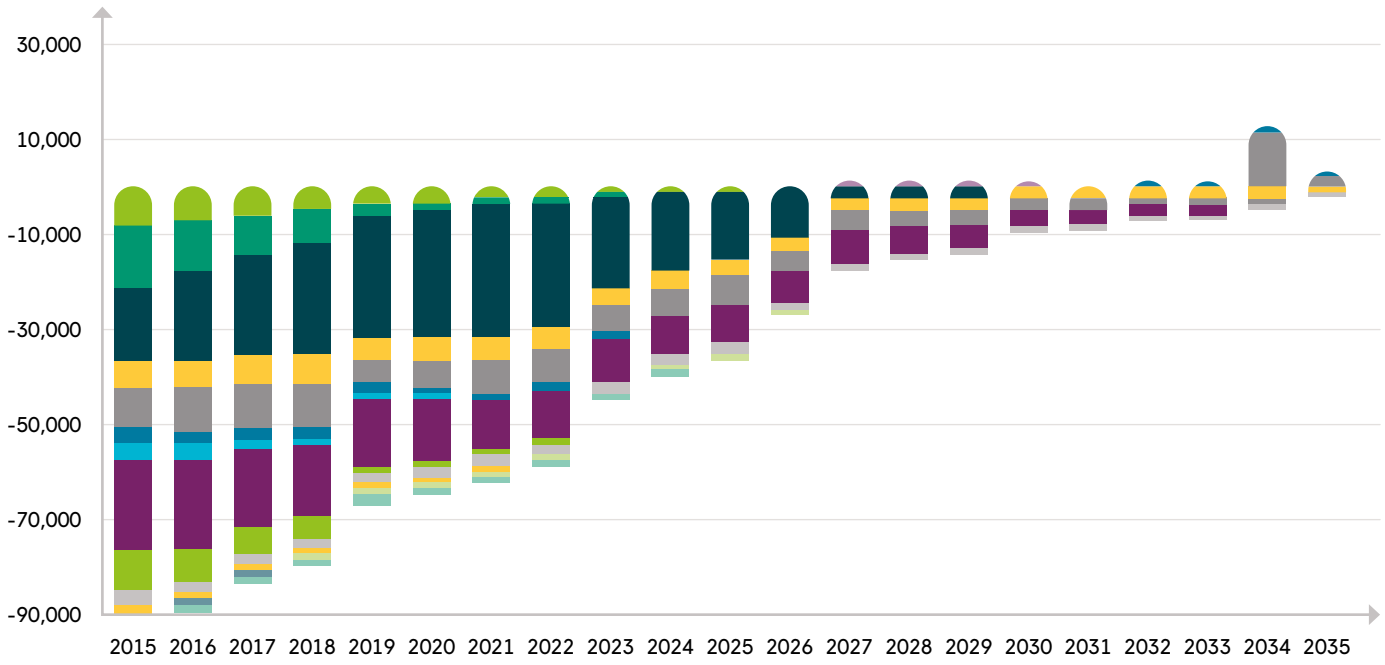

Asset Managers will also make use of assumptions and scenarios to test the behaviour of their investments under stress. The factors used will vary (actuaries will have scenarios including mortality – unlikely to be of interest to asset managers; asset managers will use factors relating to specific industries, which would be too detailed for the actuarial models), and the timescales used will also vary – for life actuaries, a 50 year scenario duration is not unusual, but asset managers (and General Insurance actuaries) tend to work over a 5 year timescale at most. However, where these scenarios coincide, it is important that they are consistent, particularly for assumptions – if the cash flows calculated by the actuaries are based on a 5% inflation rate, it makes little sense for the asset manager to assume 2%.

Common scenarios can be used to address one of the other problems facing the asset manager – that is the static nature of the cash flows. For example, if the actuaries provide one set of cash flows, and the asset managers use that as the benchmark, then a change in assumption within the asset management system will only affect the asset side of the equation. The liabilities remain constant. However, the actuaries can provide the results of multiple runs as multiple benchmarks – and the asset manager can select the appropriate benchmark to match the pre-agreed common scenario. In that way, the way the investment portfolio reacts to a particular stress can be reflected in a similar response in the benchmark, giving the asset manager a better understanding of how well the two interact. See "Transfer of Liability" below for more discussion on this topic.

Capital efficiency

The capital required by the regulator to offset the "riskiness" of the different asset classes is a key decision factor for an insurer selecting an investment portfolio. The calculation is pretty straightforward – with capital charges being assigned to each class. Equities are rated the highest, with a single stress, then Property. Bonds are more complex, with stresses varying depending on the maturity and type of bond (including a scaling factor for credit rating). The lowest capital cover is required for cash. The end result of this is that two investment portfolios of equivalent market value can have very different capital implications.

An important factor when using external asset managers, is that where the contents of a fund cannot be determined, the regulations insist on a worst case treatment. So care should be taken when selecting a fund to ensure that there is sufficient transparency to allow the calculation to accurately assess the capital requirement.

The actuaries need to do this calculation as part of the Solvency II pillar one process, and therefore are in a good position to repeat the asset part of the calculation for any proposed portfolio from the asset manager. So there is immediate benefit in passing these proposals back to the actuaries to assess the capital effect of each one, before the final selection is made. This is a good example of interaction between the various technology components within the organisation, as shown in Figure 3.

The Business Intelligence component is able to take data from either system, and in this case displays data which has already passed between the two. The comparison of the different proposed portfolios can then easily be presented to the decision makers for the final selection.

Figure 3: Data flows for capital requirements calculation

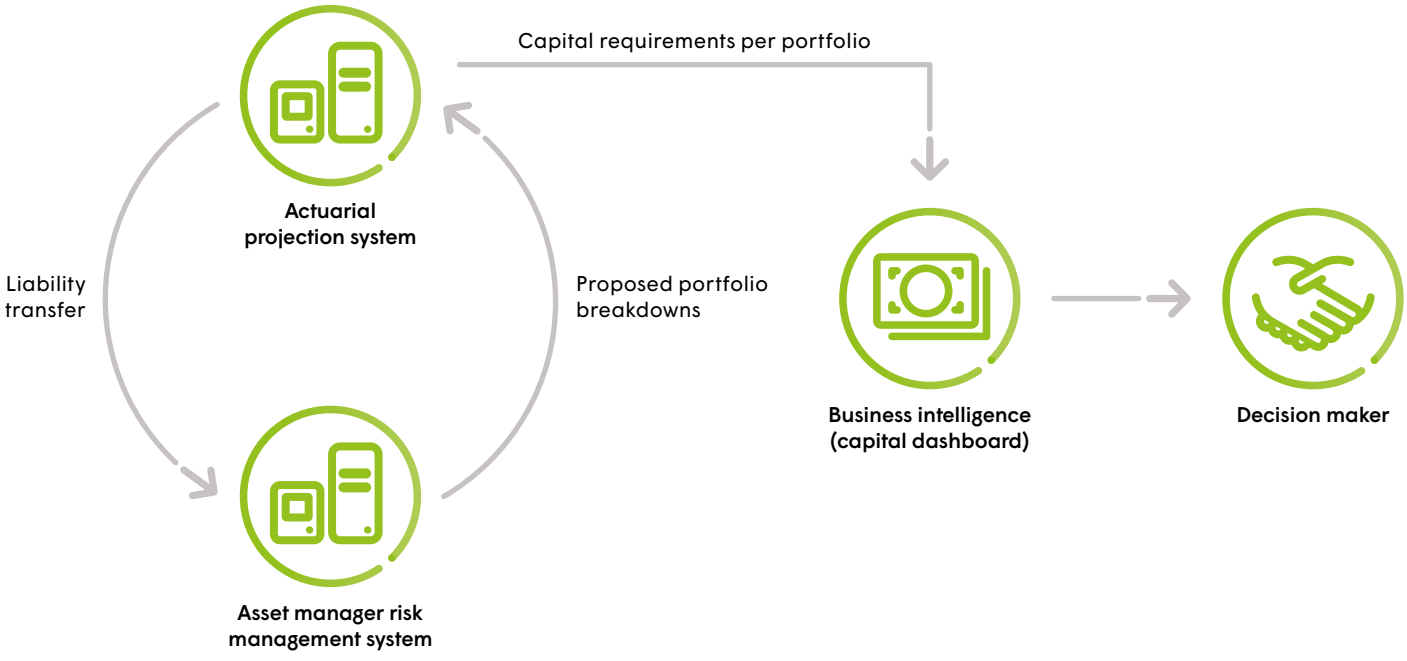

Transfer of liability

One of the key requirements for even a basic data exchange between these two worlds is for the actuaries to pass details about the liability cash movements to the asset manager. But there are far more possibilities that can be used to give the asset manager much better insight into their behaviour. The first of these was touched on earlier – passing individual cash flow collections corresponding to different scenarios, and allowing the asset manager to switch between them. But this requires large volumes of data, and pre-agreed scenarios, not particularly ideal, and not particularly flexible.

One relatively simple enhancement is to pass information about the modelled asset positions along with the cash flows. During an actuarial projection, the system will assume certain asset holdings, and simulate buying and selling those assets in line with the assumed investment strategy and the availability of cash surpluses or shortfalls. If these positions are passed to the asset managers, even in the form of generic assets, then as they flex the risk factors on their investment portfolios, the modelled asset positions in the benchmark portfolio will flex correspondingly – giving a more accurate picture of their behaviour.

Of course this process will at best give a partial picture, and a far more compelling mechanism is the use of proxies. Proxies are simple mechanisms that emulate the behaviour of the complete actuarial model by mimicking the cash flows it produces over a range of different risk factors. A commonly used form of proxy is a Replicating Portfolio – in which the system proposes a collection of assets whose cash flows match those of the liabilities. The portfolio is calibrated over a series of different scenarios, and the best fit can then be used to approximate the full set of liabilities. This type of proxy is the easiest for most asset management systems to deal with, as it changes the benchmark cash flows into a simple asset portfolio. The mechanism needed to stress this benchmark is exactly the same as that used to stress the proposed investment portfolio. All the asset manager needs to do is translate the assets used in the actuarial system (generic asset classes) into real assets – a more guided version of the process they already do with the cash flows.

A second mechanism is to replace the liabilities with a mathematical function which reproduces their behaviour. This kind of proxy is often referred to as curve fitting or Least Squares Monte Carlo, after the technique used to calibrate it. The end result of the process is a polynomial expression of the form

$$ar_1^{\alpha} + br_1^{\beta}r_2^{\gamma} + cr_1^{\delta}r_2^{\varepsilon}r_3^{\mu}$$

... where r_i are the risk factors, α, β, γ are the exponents, and a, b, c etc. are the coefficients. This potentially gives a much better fit than the Replicating Portfolio technique, and is very quick to execute – simply plug in the values for the risk factors, and out pops the answer. This does require the asset management system to be able to calculate polynomials, but this is a relatively simple extension.

Proxies represent a very powerful way for the asset managers to view the behaviour of the liabilities, and provides them with almost complete freedom to experiment with the effects of the different risk factors. It should not be forgotten however, that proxies ARE only approximations, and conclusions drawn using them should always be checked against the full model – at least on a periodic basis.

Common asset attributes

Many of the processes described above rely on the exchange of asset information between the two parts of the business, either in the form of benchmarks, or proposed investments. Making these assets common between the two, or at least providing a standard mechanism to translate between the specific and the generic makes this transfer simpler, and also eases the job of creating reports and dashboards which combine the data. If, for example, the actuarial projection system uses assets that are already defined on the asset management system, then exporting information about the assumed purchases and sales of these assets becomes a standard file transfer, with no need for intervention. Similarly, if an actuary is going to provide a proxy model based on a replicating portfolio, if the model is based on pre-defined assets, then using this proxy in the asset management system is trivial.

The master of this relationship is generally the asset management system, as this needs to hold significant amounts of data about the assets in order to model their behaviour accurately. The actuaries therefore need to ensure that their models and systems can consume data exported from the asset system, and export asset information using the same formats.

Investment strategy

When projecting liabilities forward in time, at each time period a number of cash movements will be assumed. Premiums will be paid in, claims will be paid out, policies will be surrendered and so on. These movements will be offset by asset sales or purchases (or by simply using cash reserves, in the simplest case). A good Asset Liability Model (ALM) will follow an investment strategy when modelling these sales and purchases, and "rebalance" the portfolio to keep to a particular mix. Figure 4 shows the modelled returns for a range of different investment strategies – evidently holding all of your assets in cash is the least risky strategy, but the best possible returns (and worst losses) are realised by holding everything in equities.

Asset management systems will often follow more detailed strategies than the selection shown here, and may also change strategy based on the values of economic indicators – for example, the relative yields of long-term to short-term bonds. These systems are designed to investigate the relative performance of the different strategies, at least in the short term.

Since the asset managers will be investigating the benefits of different strategies, and the liability models from the actuaries which they are taking as their benchmarks may embody the effects of a given strategy, clearly it makes sense for the two models to be aligned, at least over the near term where both models overlap. This is another area where it may be sensible to pass over a set of options for the asset manager to investigate, embodying a range of simple investment strategies, and then pass the results of their investigations back to the actuaries for incorporation into the projections.

For the actuarial team, it may be sensible to split the investment strategy into a near term and longer term section, to allow specific strategies to be used in the near term, and then switch to a simpler generic strategy in the longer term.

Figure 4: The spread of outcomes for different investment strategies

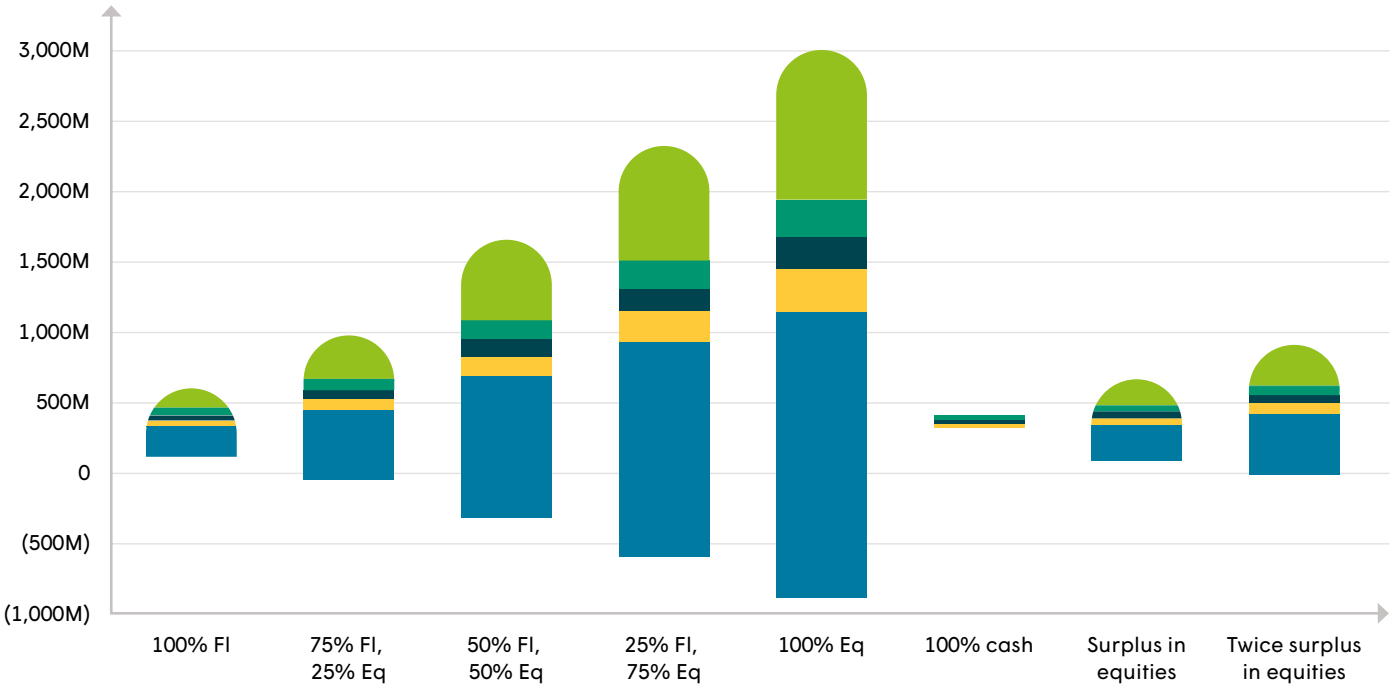

Common categorisation

Figure 3 shows a typical system architecture in an insurer, with separate systems for the actuaries and asset managers, and a business intelligence tool to host the reporting database or cubes. There are of course, many variants of this, and many systems have in-built BI capabilities, but when looking at combined data, it is more usual to consider a separate reporting engine. This allows easy addition of further systems, and also works better in the external asset manager model, where data from a number of different providers are brought together.

Whenever business intelligence is mentioned, and particularly dashboards, the expectation is always that some sort of dynamic drilldown will be supported – clicking through from one item to the next. This is supported by means of categories, sometimes referred to as "Reporting Dimensions", basically common attributes that apply to sets of data. Examples of these categories are Legal Entity, Asset Class, Product etc. These categories can be flat (Gender is a flat category – "Male", "Female", "Undeclared") or hierarchical (Legal Structures are often hierarchical, particularly in insurers which have grown by acquisition).

Placing the combined asset and actuarial information onto a common report or dashboard, and providing drill down, relies on a common categorisation scheme which can be used to group the data. All asset management systems have natural structures built into them to handle portfolio structures and asset classes at the very least. Because of the large data volumes, most actuarial systems perform their calculations in multiple runs, where a run is a subset of the data, and perhaps a subset of the overall model. They often provide some mechanism to identify sub-components – product codes, or variable groups within each run.
To assemble the complete picture, these individual runs need to be collected up into a reporting database, and mapped to the correct reporting categories. In the actuarial world, this aggregation process is often performed by a separate tool, and that tool can also be used to incorporate the asset management data, to provide a single reporting database spanning both areas.

Summary and conclusion

Both the regulations and, perhaps more importantly, good risk practice are forcing insurers to treat their risks holistically, but also to share risk information with a much wider audience. The private domains of the actuary, credit manager, auditor, asset manager or accountant are now part of the overall risk heat map that sits on the desk of every business decision maker. This drive towards convergence should be seen as a very positive influence – the more people that see the risks faced by the business, and the effects their decisions can have on those risks, the safer, and more profitable that business should be.

Yet increased transparency is only useful when the data which is being combined and displayed is accurate, based on common assumptions, includes a common understanding of how economic factors affect the business, and as far as possible includes the effect of each risk factor on any related risk factors. In the case of financial risk, as evidenced by the 2008 crisis, an incomplete picture is almost worse than no picture at all. As the economy recovers, and with IT systems better able than ever to exchange information between themselves, now is the right time to invest in getting quality and consistency in the data being shown to each business manager. As shown in the sections above, each part of the business has a unique perspective on how it operates, and only by combining those views can an organisation truly claim to have "The Truth, The Whole Truth, and Nothing But the Truth".

About FIS' APT

FIS' APT provides investment technology for a broad range of asset classes, countries and regions including data and software for understanding market risk, credit risk, liquidity risk and for portfolio construction and performance analysis. APT provides investors with statistical market risk models, performance and risk analytics and portfolio optimization and construction tools. APT's customers include institutional and retail asset managers, pension funds, private wealth managers, hedge funds, broker-dealers, prime brokers and proprietary traders. **www.fisglobal.com/apt/learnmore**

About FIS

FIS is a global leader in financial services technology, with a focus on retail and institutional banking, payments, asset and wealth management, risk and compliance, consulting and outsourcing solutions. Through the depth and breadth of our solutions portfolio, global capabilities and domain expertise, FIS serves more than 20,000 clients in over 130 countries. Headquartered in Jacksonville, Florida, FIS employs more than 55,000 people worldwide and holds leadership positions in payment processing, financial software and banking solutions. Providing software, services and outsourcing of the technology that empowers the financial world, FIS is a Fortune 500 company and is a member of Standard & Poor's 500® Index. For more information about FIS, visit **www.fisglobal.com**


www.fisglobal.com


twitter.com/fisglobal


getinfo@fisglobal.com


linkedin.com/company/fisglobal