

The Trustee toolkit downloadable

Investment in a DB scheme

Tutorial four: Stochastic modelling

By the end of this tutorial you will better understand:

- ▶ what stochastic modelling is
- ▶ how stochastic modelling can be used to understand the likelihood of various scenarios
- ▶ what percentiles are
- ▶ what 'Value at Risk' (VaR) analysis is

This tutorial is part of Scenario two.

Glossary

A detailed glossary of technical terms can be downloaded from the Resources tab when you log in at www.trusteetoolkit.com

Introduction

So far we have looked at how your scheme actuary and investment consultant can work together to build a simple model of your scheme, projecting the scheme's assets and liabilities into the future using different market and economic scenarios. This can be used to increase your understanding of the financial risks your scheme faces by stress testing or projecting a range of appropriate scenarios.

However, scenario testing has two main weaknesses.

1. It gives no indication of the likelihood of each scenario.
2. As only a limited number of scenarios are considered, important scenarios may be overlooked.

To address these weaknesses, tools have been developed to help understand the range of likely outcomes from particular investment and funding strategies. These tools are often referred to as 'stochastic models'.

Stochastic models can be very helpful when used in the right circumstances. However, they are more complex than the models covered earlier in this module and more costly to use, so it may not always be proportionate to make use of one.

As a broad principle, the bigger the risks being run by the scheme, especially relative to the sponsor's ability to support them, the more appropriate it might be to use stochastic modelling to understand how those risks might be mitigated.

What are stochastic models?

Stochastic models allow for randomness. In this case the randomness of future market conditions.

Other terms used in this context include 'asset liability modelling' and 'Value at Risk (VaR) analysis'.

Stochastic models are principally used to consider the risks involved in adopting different investment strategies, especially when the strategies are complex or the risks significant (eg because the scheme is significantly under-funded, mature, or perhaps both).

What are they used for?

Stochastic models are principally used to consider the risks involved in adopting different investment strategies, especially when the strategies are complex or the risks significant (eg because the scheme is significantly under-funded, mature, or perhaps both).

How are the risks and returns analysed?

Typically, the analysis process starts by modelling the scheme's current investment strategy, to aid understanding of the investment risks currently faced by the scheme and to provide a point of comparison. Other investment strategies are then modelled which are intended to reduce investment risk in the scheme. Each such strategy will do this to a different extent and will have a different impact on the expected return.

As we've covered in previous tutorials in this module, scheme assets and liabilities can be projected into the future using assumptions for key variables such as interest rates, inflation and investment returns. These assumptions will almost certainly not be borne out exactly in practice. Variables such as interest rates, inflation and investment returns are all random, and the assumptions are typically chosen to reflect their expected values. The projection then shows the expected or average outcome.

Why bother?

One way of allowing for random variation is carrying out scenario analysis: altering the assumptions to reflect a particular economic or market scenario. But this does not show how likely this scenario is to happen. Stochastic models take randomness directly into account, and produce a range of outcomes intended to show the extent of this randomness around the expected outcome.

Although the broad effect of each change could be worked out by general reasoning, this is not as good as stochastic modelling can be at predicting the extent of the risk reduction, or what happens when several changes are made at once and they all interact with each other. Stochastic modelling is also particularly helpful when considering complex investments such as some types of downside protection strategies, or when looking at the investment risks facing particularly mature schemes.

How is a stochastic model made?

In the previous tutorial, you learned how a simple model of your pension scheme can be used as part of an investment strategy review. There are three steps each time the model is run.

1. Create an economic scenario.
2. Project the scheme's finances using this scenario and record the outcome.
3. Analyse the result and draw conclusions from it.

Each run of the model creates just one projection of the scheme's finances.

A stochastic model is similar, but instead of using just one economic scenario, it generates a range of scenarios allowing for random variation in future outcomes. These are then used to project the scheme's finances. There are essentially the same three steps each time the model is run.

1. Create a range of economic scenarios allowing for random variation.
2. Project the scheme's finances using each scenario in turn and record the outcomes in a database.
3. Analyse the database of results and draw conclusions from it.

Each run of the model creates a range of projections of the scheme's finances, allowing for random variation.

Economic scenario generators

A stochastic model of a pension scheme is therefore made up of two components:

- ▶ a stochastic model of future market conditions which is called an economic scenario generator (ESG) and introduces the randomness into the modelling
- ▶ a model used for projecting the pension scheme's finances (the one that we have already described when discussing scenario analysis)

What does running the ESG involve?

The random scenarios produced by the ESG are input to the scheme projection model to produce output showing how the scheme's finances develop under each of these random scenarios. Running the economic scenario generator involves:

- ▶ inputting assumptions about certain investment variables' average values, the likelihood of high and low values, and the relationships between different variables
- ▶ generating random values for the variables based upon those assumptions
- ▶ building up a database of a large number of values for the investment variables, ie a large database of economic scenarios

How do you build up the database?

Each scenario is then used to run the scheme projection model, building up a database of a large number of scheme statistics such as assets, liabilities and contributions, ready for analysis.

Building up a database of economic scenarios saves effort compared with re-running the ESG each time a new set of scheme projections is required as part of the investment strategy review (such as projections using a different investment strategy, or a different pattern of contributions).

It is important for you to understand in broad terms the range of scenarios produced by the ESG. This will naturally reflect the assumptions made when programming the ESG. Your investment consultant should be able to explain in broad terms how the input assumptions have been arrived at, and why the resulting output scenarios are appropriate for modelling your scheme.

How do you analyse the database?

Analysing the database of returns can be a helpful way of understanding the ESG output in 'real world' terms. For example, you could analyse the database of equity returns to ask:

- ▶ what is the average return over the next 10 years expected to be?
- ▶ how likely is a return 2% pa below average over the next ten years?
- ▶ how likely is a return 2% pa above average over the next ten years?
- ▶ how likely is a return below -10% in any one year?

The expected equity return will almost certainly be an explicit assumption for the ESG.

Question one is really a way of checking that the ESG is working sensibly (the answer for a large database of runs should be very close to the input assumption). Questions two and three would give a feel for the spread of returns around the average. Question four would give a feel for the downside or 'tail' risks in the scenarios produced by the ESG.

Key modelling assumptions

The input assumptions (both for the ESG and scheme projection model) are very important and have a potentially material effect on the model output. Typically, the key assumptions for a pension scheme stochastic model are the:

- ▶ expected return from different asset classes
- ▶ likely variability of those returns (often expressed as 'volatility' of return)
- ▶ likelihood of extreme investment returns, especially extremely negative ones
- ▶ expected future path of interest rates and inflation, especially if these are assumed to be different from what is implied by the current market prices of bonds
- ▶ correlation in returns between different asset classes and with inflation, ie, the extent to which different asset classes and inflation are assumed to move together or separately
- ▶ how the discount rate is set for the valuation of the scheme's technical provisions

All but the last of these are, strictly speaking, assumptions for the ESG, the last is an assumption for the scheme projection model.

If your investment consultant is doing an asset liability study for your scheme using a stochastic model, it is important that you understand which assumptions are the most important and how they have been derived. The choice of assumptions, and especially the relationships assumed between different economic variables, can have material effects on the model output.

Case study: Stochastic model output for funding level

Let's follow a simplified example looking at funding level. Here we have input various appropriate assumptions into the ESG, and run it to produce a database of 5,000 scenarios.

Year	1	2	3	4	5	6	7	8	9	10
Run 1	89%	85%	93%	102%	93%	88%	100%	107%	90%	96%
Run 2	92%	88%	86%	100%	98%	90%	101%	105%	96%	108%
Runs 3-4,998										
Run 4,999	111%	108%	94%	97%	112%	103%	95%	104%	106%	109%
Run 5,000	109%	98%	106%	101%	118%	103%	119%	103%	103%	105%

Then, we have used each scenario as input for the scheme projection model and built up a database of 5,000 scheme projections.

For each projection, the model has produced output including assets, liabilities, contributions payable and funding levels. Of these, we have recorded the projected funding levels in to the table shown.

Case study: Analysing the resulting statistics for funding level

Just as with the ESG output you can then analyse the database of funding levels to ask:

- ▶ what is the expected funding level in 10 years' time?
- ▶ how likely is it that the funding level will be above 100% in 10 years' time?
- ▶ how likely is it that the funding level will be below 70% in 10 years' time?
- ▶ how are things expected to be on average?
- ▶ how bad or good might they be in extreme scenarios?

However, it can be difficult to understand the model output just by looking at the database. We will now look at how the data can be analysed and presented to help answer these questions.

Using percentiles

One way of indicating the range of model outcomes, such as the funding level after 10 years in the case of our example scheme, is to sort them in order from best (at the top of the list) to worst and see where values lie.

What does '5th' or '25th' percentile mean?

'Percentiles' are a way of saying what values lie at certain points along the list. The 5th percentile is the funding level that is 5% (ie 1/20th) of the way up the sorted list. The 25th percentile is the funding level that is 25% (ie a quarter) of the way up the sorted list.

In other words, 5% of the 5,000 model outcomes have less than 82% funding after 10 years. Similarly, in 25% of the 5,000 scenarios, the funding level after 10 years is below 90%.


Another way of saying this is that, according to the model and assumptions, there is a 5% chance of the funding level being below 82% after 10 years, and a 25% chance of it being below 90%.

How can I tell what the average is?

The 50th percentile, known as the 'median' is the figure half way down the sorted list. In our example, in half the scenarios the funding level after 10 years is above 97%, and in half of them, it is below 97%.

The median funding level is sometimes referred to as the average or expected funding level although mathematically the average funding level is a different, but generally similar figure.

Using a bar chart



A typical way of showing funding level output from a stochastic model is a bar chart like the one shown. Year one of the projection is the first bar, then to its right the bar for year two and then year three, all the way to year 10. On each bar, we show the 5th, 25th, 50th, 75th and 95th percentiles of the distribution of funding levels for that year.


In order to produce these bars, the investment consultant will have had to sort the database 10 times (first by the funding level in year one, then by the funding level in year two etc).

Each individual projection will appear at different points in the list, every time it is re-sorted. Effectively, each of these bars is giving a snapshot of the range of outcomes at each point in time, it does not show how any individual projection progresses over time.

Interpreting the charts

Charts like this need careful interpretation. For example, when considering downside risks, it is deceptively easy to read along the bottom of each bar and imagine the funding level gradually deteriorating each year towards being 95% percentile (or worse) by year 10. However, the scenarios which are in the bottom 5% of outcomes for the year 10 funding level will have taken a variety of paths to get there.


One such scenario (scenario one) might represent a run of bad market conditions, where the funding level has not really improved despite significant contributions being paid. This line stays within 80% and 90% funding over the full 10 year period.



Another (scenario two) might have been relatively favourable with limited contribution requirements before being impacted by very bad market conditions in years nine and 10. This line rises from just over 80% funded in year one to almost 110% funded in year eight, but falls to under 80% funded over years nine and 10. These projections have been added to the bar chart.

Funnel of doubt

To illustrate the need for careful interpretation, we chose to illustrate the funding levels as a series of separate bars. It is more usual to join the bars up with a series of smoothed lines, like the one shown.



When plotted these lines start in year one closer together than in year 10 which looks like a funnel.

This type of chart is sometimes referred to as a 'funnel of doubt' chart, because the range of (or uncertainty in the) outcomes increases as the projection period increases so that it looks like a funnel.

Other statistics

The database of stochastic model output will typically record many more statistics than just the funding level. These can also be summarised into tables and graphs, to assist with understanding the risks facing the scheme. Typical examples would be the deficit repair contributions paid and the surplus or deficit in £ terms.

Deficit repair contributions paid

Particular care needs to be taken when considering output concerning the contributions paid.

The stochastic model will almost certainly have the flexibility to assume that actuarial valuations take place during the projection period (say every three years) and that revised patterns of contributions are put into payment following each one. To do this, assumptions need to be made regarding matters such as the period over which any deficit repair contributions are to be paid, or any surplus is spread.

If the model is run with allowance for future valuations to alter the contributions, this will tend to narrow down the 'funnel of doubt' charts. This is because the purpose of adjusting the contribution rate is to steer the funding level towards 100%.

Example

A scenario that results in 90% funding after 10 years without any adjustment to contributions might result in 95% funding after 10 years when allowance is made for adjusted contributions. Likewise, a scenario that results in 85% funding after 10 years without altering the contributions might also result in 95% funding after 10 years, when the contributions are adjusted.

In the latter case, the contributions paid would be higher, a feature that would not be apparent just by looking the 'funnel of doubt' chart. However, the higher contribution amounts would affect any output tables or charts showing contributions paid.

When considering stochastic modelling output for your scheme, it is important to understand exactly what assumptions are being made about the future payment of contributions.

Surplus/deficit 'Value at Risk'

One way of looking at the £ surplus or deficit is known as a 'Value at Risk' or 'VaR' measure. Despite the different terminology, VaR is simply a type of statistic that can be calculated from the database of stochastic model output. Rather than answering 'how bad is the deficit likely to be in X years' time' it seeks to answer 'how much worse than expected could it be?'

Deficit example

Suppose that you have the following output from a stochastic model:

Initial deficit (calculated as initial liabilities less initial assets, both of which are actually model inputs rather than outputs) = £10.5 million

Percentile	Deficit
5th	£14.5m
25th	£11.5m
50th	£8.5m
75th	£6.0m
95th	£4.5m

Distribution of the deficit in three years' time (percentiles from the database of model output) is shown in the table.

In this case the expected position in three years is to have a deficit of £8.5 million, according to the model assumptions. (This improvement would typically be on account of assumed deficit recovery contributions as well as expected investment returns.) However, there is a 5% chance that the deficit could be more than £14.5 million according to the model and assumptions used.

Deficit example VaR

The VaR way of looking at this would be to say that there is a 5% chance of the deficit being more than £6.0 million worse than expected or the 'Value at Risk' is £6.0 million.

However, just to say that the VaR is £6.0 million is incomplete since you need to state the probability level and the period over which the calculation is being done. In this instance it would be expressed by saying that the three-year VaR_{95} is £6.0 million.

(The subscript 95 indicates the percentile level the convention being that 95 indicates 5% probability, 99 indicates 1% probability, etc.) To be properly complete you would also indicate the model and key assumptions being used.

Separating VaR into its component risks

A further use of stochastic models is to illustrate the relative size of the various types of investment risk that the pension scheme faces.

VaR is a measure of how much worse than expected a pension scheme deficit can become. Above we gave the example of a scheme whose three-year VaR_{95} statistic was £6.0 million. This took into account the combined effect of variability in inflation, interest rates and investment returns.

It is common for this statistic to be broken down into components, by placing values separately on the inflation, interest rate and investment return parts.

Component risks

One way of doing this is to run the model four times showing:

- ▶ the combined effect (giving the VaR_{95} figure of £6.0 million)
- ▶ inflation only (ie turn off the variability in interest rates and investment returns) eg £4.5 million
- ▶ interest rates only (ie turn off the variability in inflation and investment returns) eg £1.5 million
- ▶ investment returns only (ie turn off the variability in interest rates and inflation) eg £2.0 million

The components invariably add up to more than the total VaR here they add up to £8.0 million, whereas the total VaR is £6.0 million. The reason for this is that the risks are diversified, ie to a degree (dependent on the model assumptions), poor outcomes for inflation, say, tend not to be associated with poor outcomes for interest rates. The benefit of this diversification is a reduction in the VaR of £2.0 million compared with the sum of the components.

One of the facts brought out by this type of VaR analysis is that, for many pension schemes, the interest rate and inflation risks can be as big as, if not bigger than, the investment return risks (such as equity market risks). This is especially the case if the scheme is under-funded.

Under-funded schemes

The underlying reason is that changes in interest rates and inflation can have big effects on the value placed on the scheme liabilities, which are not offset by the much smaller changes in the value of the scheme's bond holdings.

That is one of the reasons why interest rate and inflation hedging are important matters to consider as part of an investment strategy review.