portfolio construction forum

Building a robust portfolio for future unknowns

Mark Foster & Chris Nichols | Standard Life Investments | 08 August 2014

While historical stress scenario analysis is now a well-established approach in tail-risk management, there is less consensus on how best to perform forward-looking scenario analysis. This paper examines how expert opinions about future world states can be incorporated into projections of portfolio outcomes. It also illustrates that using a diversified multi-asset approach can effectively help mitigate losses in such extreme market scenarios.

INTRODUCTION AND OVERVIEW

Recognising the inadequacies of VaR (value at risk) metrics in measuring portfolio risk, investment practitioners and regulators have increasingly focused on stress testing and scenario analysis as a means of exploring so-called tail risks. However, while measuring portfolio behaviour under historical stress scenarios is now established practice in risk management, there is much less agreement on how to perform forward looking scenario analysis.

Major regulators noted two common weaknesses in the stress testing methodologies used during the financial crisis, namely:

- lack of qualitative expert judgment; and/or
- over-reliance on historical statistical relationships such as correlations.

Ideally, a forward-looking scenario methodology should be able to integrate a limited set of expert opinions with all relevant financial market data. These inputs could then be translated into portfolio gains and losses.

This paper aims to show:

- how to quantitatively construct a population of data points for tail-risk analysis;
- how to blend expert opinion to produce estimates of tail behaviour for different assets; and,
- how this methodology can be applied in practice, using one particular tail-risk scenario.

Analysis of the asset allocation of a typical Australian Super Fund illustrates the robustness of an absolute return investment strategy. This shows the benefits of spreading return-seeking investment risk within a single portfolio, by investing across asset classes and geographies, and using long and relative value positions.

CONSTRUCTING A POPULATION OF DATA POINTS FOR TAIL-RISK ANALYSIS

Fat tails exist if you believe the market is single-state distribution. Financial market data is often described as 'fat-tailed'. The concept of fat tails implies that, under a given scenario, a normal distribution of outcomes is not normal but indeed skewed and therefore has fatter tails. In practice, more extreme variations of returns occur.

Figure 1 highlights the fat tails commonly observed in financial data. It shows a histogram of historical daily returns for the Australian equity ASX200 Total Return index over the last 13 years, together with a normal approximation.

Figure 1: Actual Australian equity return distribution versus a fitted normal curve

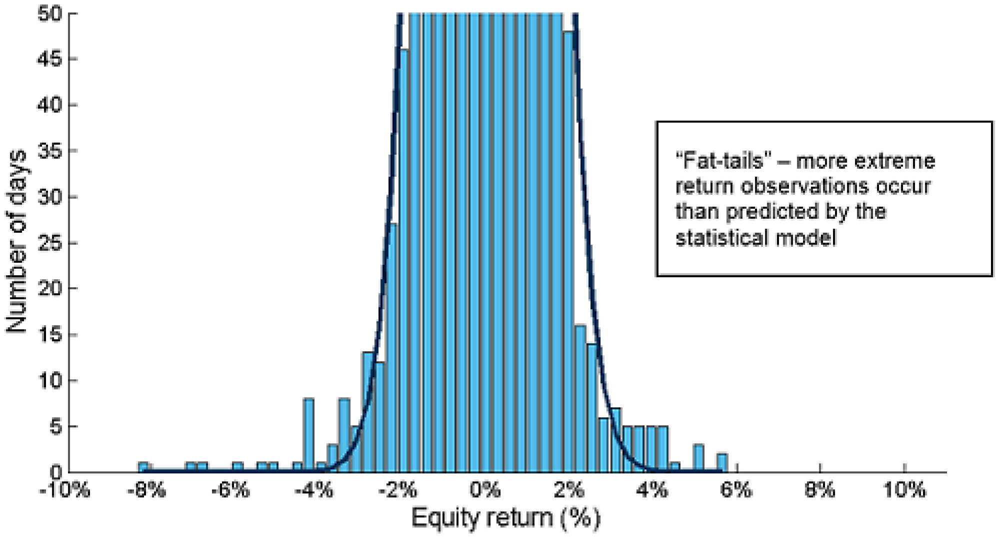

Sources: Bloomberg, Standard Life Investments (as of March 2014)

Traditional risk systems assume that financial market returns are independent and identically distributed (i.i.d.) and follow a normal distribution. (An example of i.i.d. is the repeated tossing of a coin – we expect as many heads as tails and the outcome of any one toss of the coin should be unaffected by any other.)

It is well known that these models significantly underestimate the probability of so-called tail events. Although more complex models attempt to describe and simulate fat tails, these approaches are still essentially built on the belief that one single-state distribution exists.

INTRODUCING MULTI-STATE DISTRIBUTIONS

A more intuitive approach is to consider market behaviour as a function of many different states of behaviour. This better reflects the observations of actual correlations and asset return variations throughout a market cycle.

For instance, investors have long recognised that economies tend to oscillate between steady, low-volatility states (or 'regimes') characterised by economic growth, and nervous, high-volatility states characterised by economic contraction.
Consider the same return data of the ASX200 Total Return Index of Figure 1 but this time, dividing the daily returns into a 'high-volatility' regime (in red) and a 'low-volatility' regime (in blue). For each of these two regimes, a normal distribution (Figure 2) is fitted. Even with this simple approach, it is visible that the fat tails observed earlier are less apparent.

Figure 2: Actual Australian equity return distributions for low and high volatility regimes versus fitted values of normal distributions

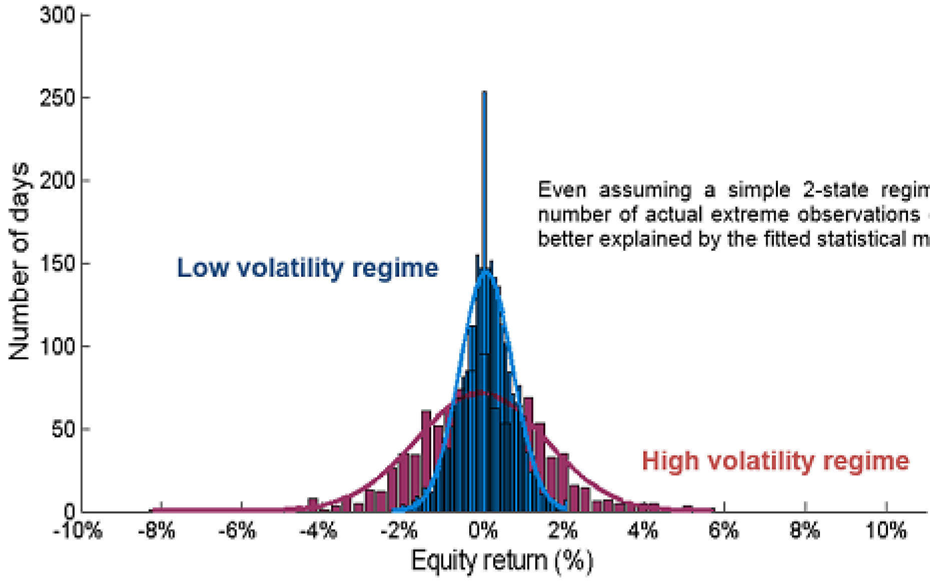

Sources: Bloomberg, Standard Life Investments (as at March 2014)

For the purpose of scenario analysis, this approach can be very useful as a regime-switching model. Such a model is able to produce fat tails as well as 'tail dependence' without the need for exotic parameter estimates. Therefore, the assumption that market returns are identically distributed is relaxed, and instead market behaviour is modeled by using a large number of overlapping regimes. These are encoded as a library of Monte Carlo simulations (or computational algorithms that rely on repeated random sampling to obtain numerical results; typically simulations are run many times over, in order to obtain the distribution of an unknown probability).

BLENDING EXPERT OPINION WITH QUANTITATIVE ANALYSIS

In the absence of any other information, each simulation has to be treated as equally likely. In blending expert opinion, a forward-looking distribution that is consistent with the expert views but remains as close to the prior distribution as possible is created. It is through the process of *maximum entropy* that it is possible to mathematically derive these sought after estimates.

For decades, the principle of maximum entropy has been used in the world of physics for a wide range of applications, from statistical mechanics to information theory and logical inference. The introduction of entropy maximisation in the context of stress testing and scenario analysis is relatively recent, with the global financial crisis providing the spur for better understanding of financial risk.

The use of *Entropy Pooling* is probably best explained by analogy. Imagine looking at a fuzzy picture – expert opinion states that the underlying image is a monkey – and the mathematics carried out produces a 'best' picture of that monkey, given the pixels (data points) that produce the image. However, if expert opinion instead states that the underlying image is that of a cat, the output of this mathematical output will be an image of a cat.

In the world of investment, the underlying image is the expert opinion of the 'picture' of a stressed market condition. The pixels are all the possible extreme outcomes, based on different regimes observed in the past. Entropy pooling aims to give the best representation of that stressed market with the data points we have available. So the 'picture' of the future stressed market condition still contains a footprint of the past, this being the interdependencies that have existed during specific regimes.

CASE STUDY – THE CHANGING RELATIONSHIP OF US EQUITY AND US CREDIT RETURNS

Figure 3 studies the relationship between equity and credit returns through the market cycle. It maps a scatter-plot of weekly returns of the S&P 500 total return index on the x-axis against weekly returns of the Merrill Lynch Corporate Master index on the y-axis over a 13-year period. In addition, this also shows the outline of a single multivariate normal distribution that has been calibrated using the entire data set.

Figure 3: Scatter-plot of weekly returns of the S&P 500 total return index against weekly returns of the Merrill Lynch Corporate Master index

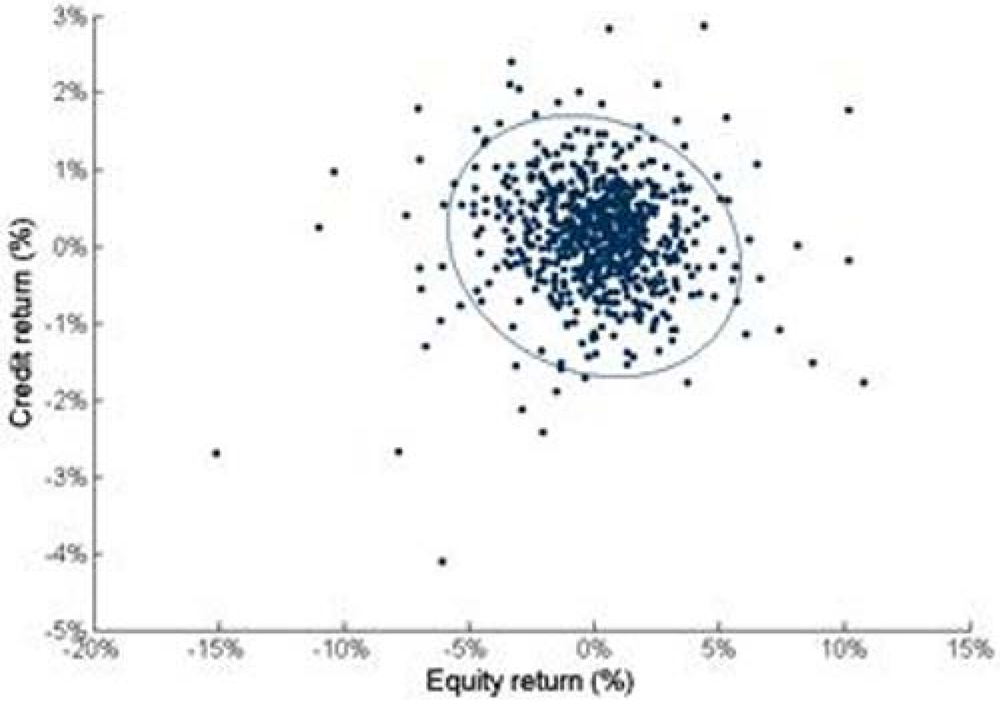

Sources: Bloomberg, Standard Life Investments. Data set March 2000 to July 2013

Figure 3 clearly shows that use of a single distribution does not capture the tail dependence of these two return series. So, while the average correlation of just -0.04 suggests bonds and equities do not move together, experience shows that during extreme falls (such as happened in 2008), they in fact tend to behave in unison.

As with the Australian equity example in Figure 3, this obtains a much better fit to the actual data when using a multi-regime model – that is to say, recognising that the relationship between equity and credit returns changes throughout the market cycle.

Figure 4 shows the library of simulated annualised returns (using Monte Carlo techniques) that incorporates the various correlation structures observed during the last 13 years, based on a large number of different regimes. This distribution retains the tail-dependence observed in the historical data and cannot be parameterised by a single distribution.

Figure 4: Scatter-plot of simulated annual returns of the S&P 500 total return index against annual returns of the Merrill Lynch Corporate Master index

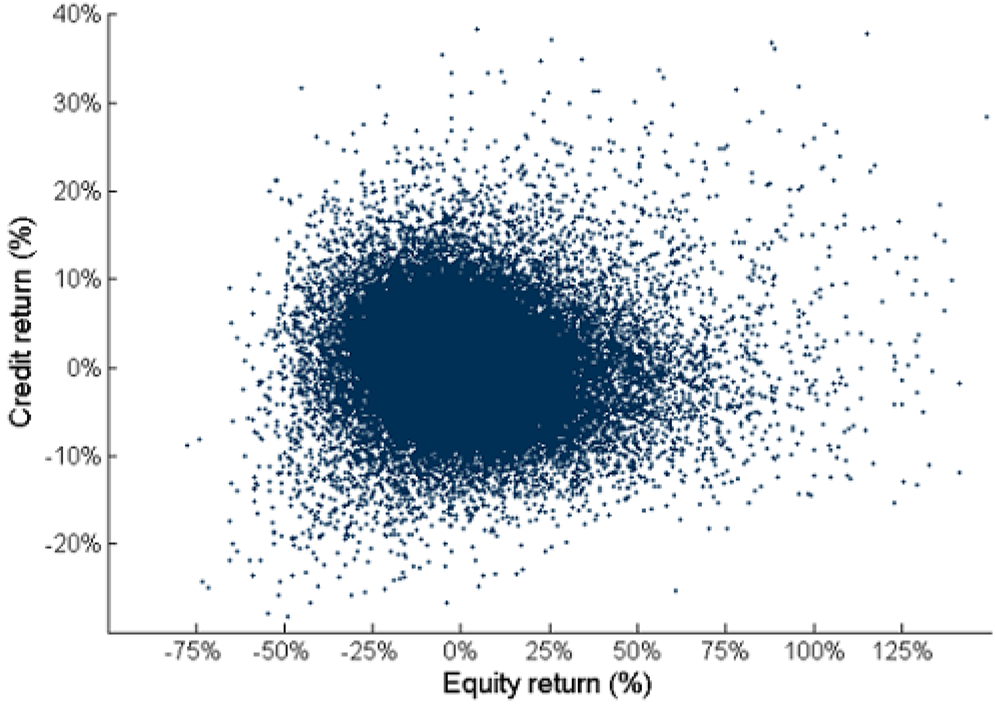

Sources: Bloomberg, Standard Life Investments (as at March 2014)

Now consider a 'credit collapse' scenario –which expert opinion has defined as a -20% shock to the credit index. The prior belief (fuzzy picture) has to be updated in light of this new information using the Entropy Pooling methodology (Figure 5).

Figure 5: Reweighting of data points consistent with 20% credit collapse scenario

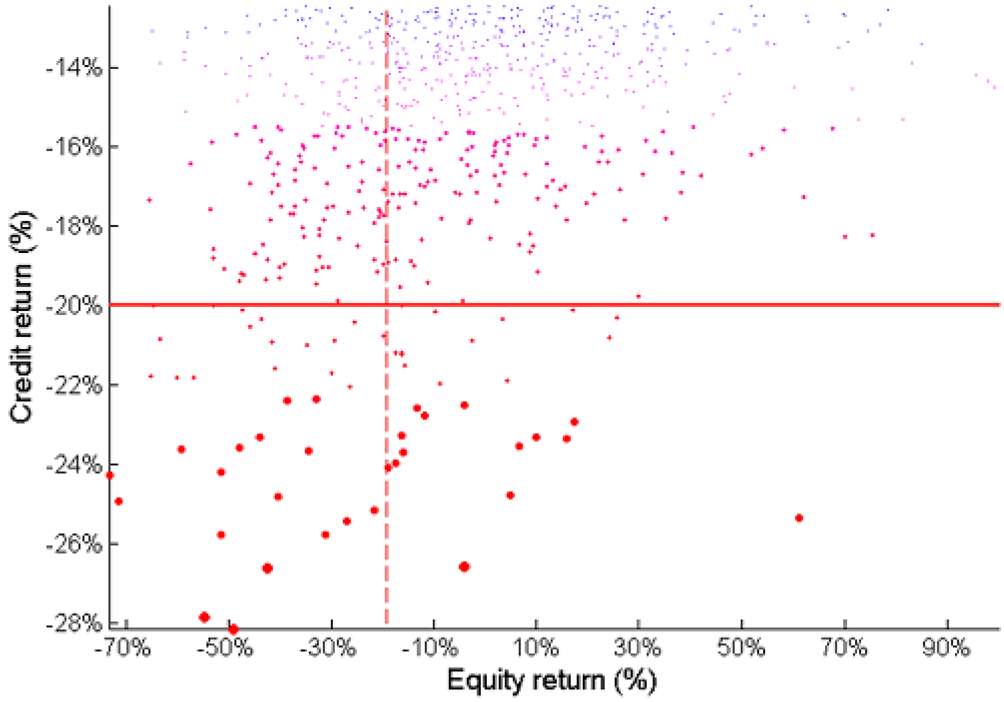

Sources: Standard Life Investments (as at March 2014)

Practically, this objective can be achieved by re-weighting the simulations in the library, such that the average loss for the credit position is -20%. In Figure 5, a -20% shock to the credit index is used to infer a -20% drop in the equity index.

The methodology correctly infers that, for extreme shocks such as a -20% drawdown in the credit index, the co-dependence between equity and credit is much stronger than for 'normal' periods; the expected loss for the equity position under this methodology is also about -20%. This contrasts markedly with a traditional 'single state' approach where a -20% shock in the credit return would imply a +3% rise in the equity index.

TAKING AN AUSTRALIAN PERSPECTIVE – CHINA CRISIS SCENARIO

The methodology demonstrated above can be applied to a wide range of asset classes whose returns can be modelled, and to a specific set of expert opinions about unlikely events. For the purpose of the next set of analyses, the asset weightings of a typical Australian super fund are used to test its robustness against a scenario defined as 'China crisis'

China crisis – key assumptions:

- Economic rebalancing causes China's growth to slow significantly;

- The slowdown is compounded by rising inflation/wages;
- Demographics limit future development;
- Debt-fuelled investment boom / misallocation of resources / debt quality; and,
- Productivity growth remains low or falls further.

China Crisis – key expert inputs:

- US$ vs AUS$ +15.0%
- Copper price -60.0%
- MSCI EM Index -75.0%

Treatment of liquid and illiquid asset groups

By using a typical Super Fund's asset allocation, a sense of a portfolio's likely behaviour, can be derived, given the above information. Although modelling liquid assets (i.e. where daily data is available) is straightforward, more thought is required for how to treat less liquid investments such as private equity and real estate. While price direction will ultimately be the same as liquid equivalents, the less frequent pricing of these assets will typically create time delays in their behaviour in tail-risk events.

The most optimistic approach is taken here, treating less liquid investments as static in nature. That is, prices are assumed to remain the same and that in such a crisis, these assets become effectively untradeable. Modelling these assets as their liquid equivalents (such as public equity for private equity and REITS for real estate) would cause materially worse outcomes in tail-risk modelling. As it is, memories of the 2008 global financial crisis and the resultant distortions on the super fund's asset allocation weightings should provide a recent reminder of the downside potential of the illiquidity premium.

Output Analysis

Figure 6 shows the likely negative impacts of a China crisis on various asset classes. Unsurprisingly, equities suffer significantly, with fixed income providing some protection. For 'Alternatives', consider only the liquid portion of this allocation, i.e. hedge funds. It can be seen that, as a universe, these funds have provided significant downside protection. However, it can be appreciated that in practice there will have been a wide dispersion of returns around this central figure, dependent upon manager and strategy selection.

Figure 6: The impact of 'China crisis' scenario on asset prices

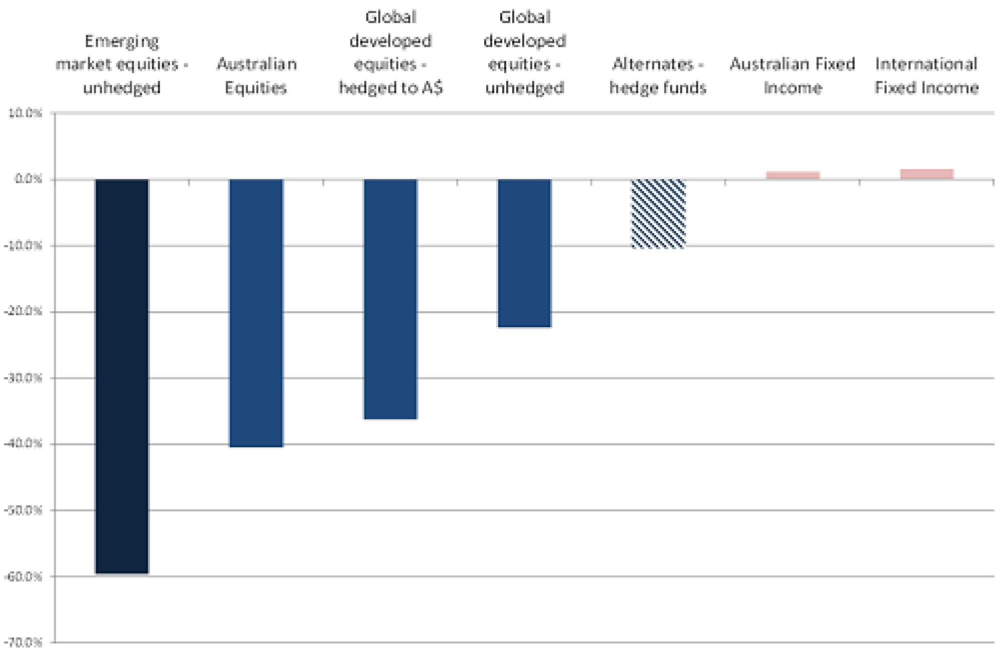

Sources: Standard Life Investments (as at 31 March 2014)

To calculate the expected losses for the superfund the following weighting assumptions[1] are used:

Australian equities	26%
Global developed equities – hedged to A$	11%
Global developed equities – unhedged	11%
Australian real estate	9%
Australian fixed	9%
Global fixed	5%
Alternates – hedge funds	7%
Alternates – infrastructure	5%
Alternates – private equity	5%
Cash	12%
Total	**100%**

For modelling purposes, given the earlier assumptions on illiquid assets, this is effectively looking at asset movements on 69% of the portfolio (31% of the portfolio being illiquid assets and cash).

Combining the return forecasts (Figure 6) with the asset weightings produces an expected negative total return of -16.7% for the entire portfolio (Figure 7).

Figure 7: China crisis forecast loss of -16.7% for a typical Super Fund broken down by asset contribution

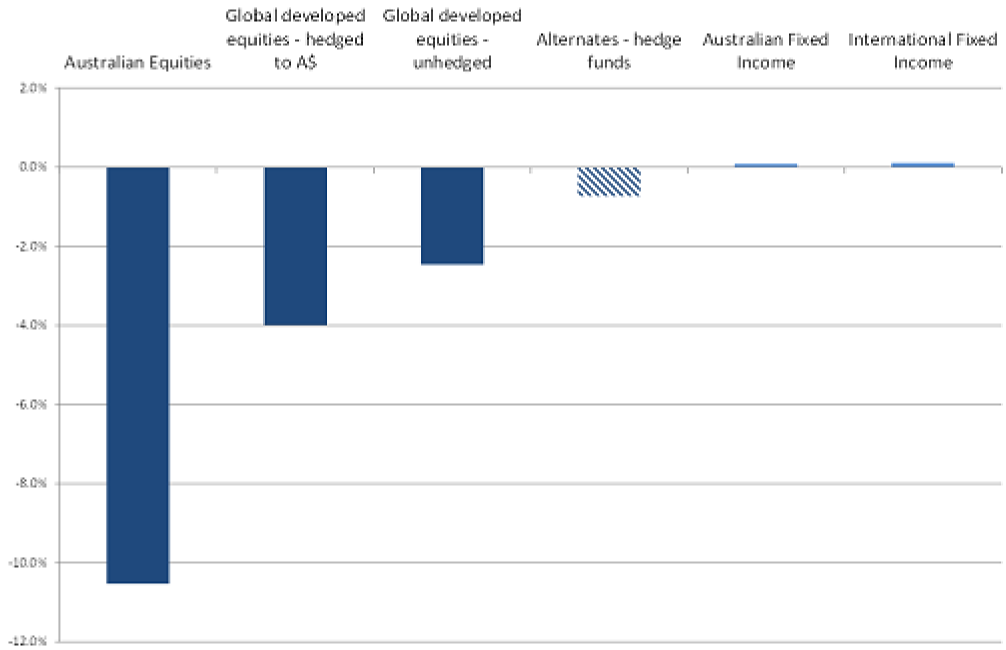

Source: Standard Life Investments

Unsurprisingly, given the large bias toward the domestic Australian equity market, the losses are greatest here. As pointed out earlier, using liquid market assumptions for some of the illiquid portfolio components would lead to significantly worse outcomes for the portfolio.

While fixed income provides good protection in this extreme scenario, it offers limited long-term returns, especially with prevailing fixed-income yields near 30-year lows in many developed markets. As discussed earlier, while hedge funds as an overall investment group offer good protection their well-documented shortcomings (high fees, selective access, disclosure opacity, indifferent performance) may preclude greater allocations to this asset class group in its own right.

ABSOLUTE RETURN FUNDS

The emergence of absolute return funds as a mainstream global investment class has been partly driven by a desire for the strong diversification benefits hedge funds can offer but without the traditional drawbacks.

The next part of the paper considers a representative absolute return fund, which has the transparency of underlying investment positions. This allows observing how future scenario analysis might help fund managers understand the likely behaviour of their holdings in extreme market conditions.

By way of brief introduction, the multi-asset absolute return fund under analysis is a multi-asset portfolio consisting of a diverse range of investment positions (long and relative value in nature) in traditional markets and in selective currency, interest rate and volatility markets. All portfolio positions are extremely liquid, hence ideal for this type of analysis.

The investment objective of this multi-asset fund is to deliver an absolute return of cash + 5% per year, gross of fees (i.e. consistent with the long-term return on equities), irrespective of market conditions, over a rolling 3-year period.

Because of the 'all-weather' nature of this portfolio, studying its potential behaviour in extreme scenarios is a core element of the investment approach.

Performing the above China crisis scenario analysis, the absolute return fund in this case would deliver a return of -6.2. This compares favourably with the 'alternatives - hedge funds' group (-10.6% from Figure 6).

The multi-asset absolute return strategy in this analysis has at any time around 30 return seeking positions, both long and relative value, invested across geographies and asset classes. Thus, it seeks to hold a very diverse set of investment risks. Figure 8 shows the contributions to return of the constituents of the multi-asset absolute return fund in the China crisis scenario.

Figure 8: absolute return portfolio constituent breakdown in China crisis totalling -6.2%

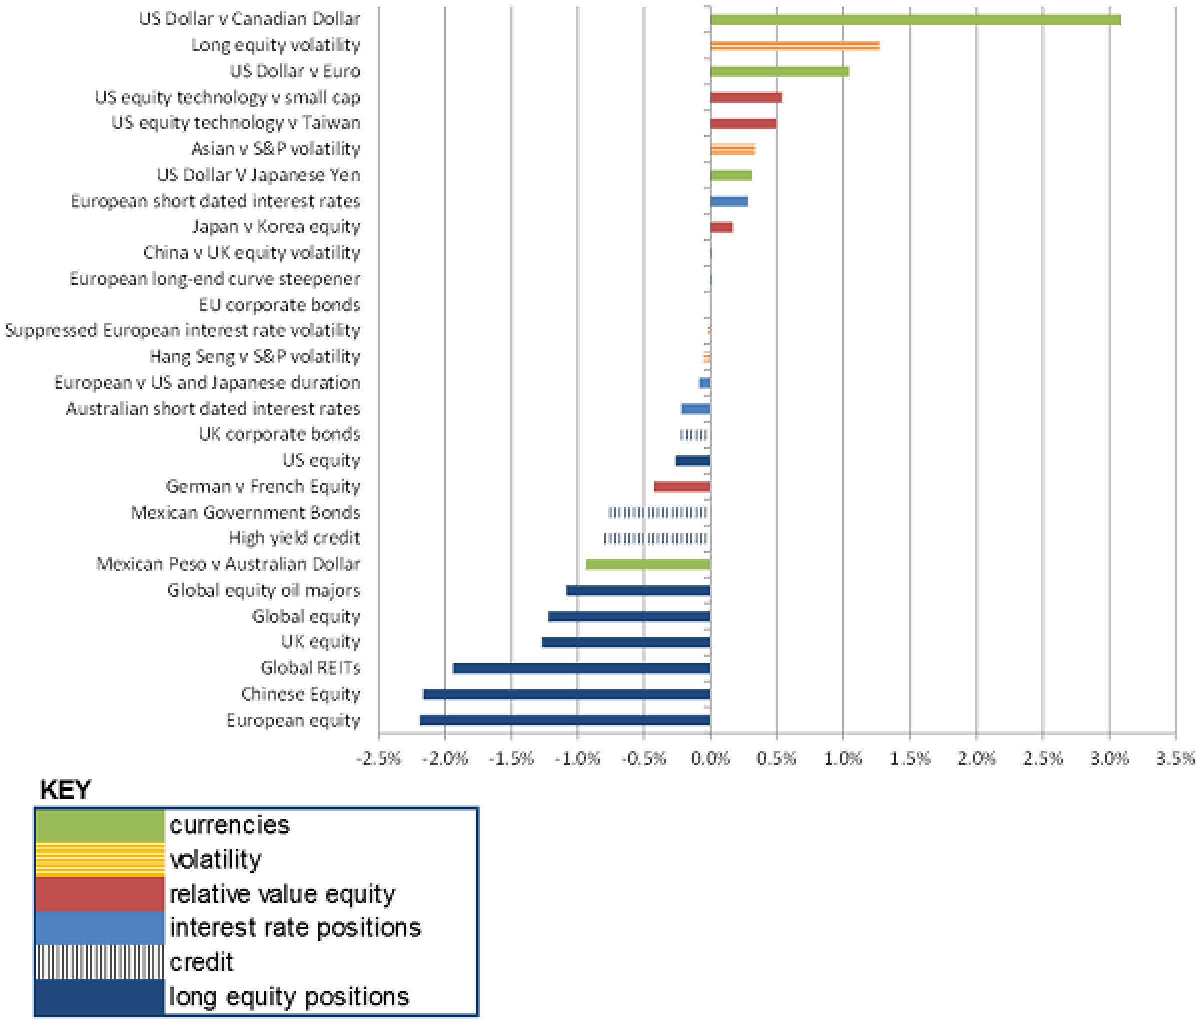

Sources: Standard Life Investments (as at March 2014)

As can be observed, carefully chosen investment strategies in currencies, volatility and relative value equity, which are chosen in the belief they can make money in normal market conditions, can also provide good returns in the China crisis scenario. Such strategies can thereby offset losses in more traditional return seeking strategies.

There been maintained a particular focus on the likely response of an Australian Super Fund to a China crisis. The elegance of forward looking scenario analysis allows for the specification of many additional unlikely but plausible 'bad outcomes' for markets. This involves little extra work except for the actual specification, which is thereby the most important input of any given scenario.

Organisation challenges for constructing extreme scenarios

However mathematically pleasing the above analysis, it constitutes only a small (albeit crucial) aspect of scenario analysis. It must be remembered that the output is just an accurate mathematical articulation of expert opinion and that the expert opinion in itself may be deeply flawed. Put more simply, garbage in, beautifully-modelled garbage out (or, in the context of Entropy Pooling, the underlying picture ends up being a cat, rather than the monkey we were given to expect!).

Two dimensions to scenario analysis that not yet discussed, but that are essential for its usefulness are:

- the process by which expert opinions are formed; and,
- the organisational framework around scenario analysis that ensures the process itself and its outcomes can help fund managers in their decision-making.

It is important to recognise that the value of successful scenario analysis should not lie in the exact specification of portfolio gains and losses in a specific scenario – the results should be intuitively correct but not spuriously accurate. Rather, its benefits lie in the process and interaction it enforces. Ultimately, forward-looking scenario analysis is a tool to explore potential portfolio weaknesses, through the interactions of:

- experts, who specify scenario shocks;
- risk managers, who model the inferred losses; and,
- fund managers, who can use the results to challenge and enhance their intuitive understanding of the behaviour of their portfolios in extreme scenarios.

CONCLUSIONS

In this paper, a methodology to quantitatively translate expert opinion about future world states into portfolio outcomes has been demonstrated. Furthermore, it has also been shown that using an absolute return investment approach can be beneficial in helping mitigate losses in such extreme markets. Finally, it was also highlighted that the real value of such analysis is in the engagement of fund managers in the scoping and output of the work, allowing them the opportunity to improve current portfolio construction techniques.

portfolio construction forum

ENDNOTES

1. Rainmaker average weighted asset allocation from corps, industry and government funds, September 2013. Source: Standard Life Investments.

REFERENCES

1. Fully Flexible Views: Theory and Practice, A Meucci, Risk 21.10, 2008.

2. Principles for sound stress testing practices and supervision, Basel Committee on Banking Supervision, 2009.

3. Asset allocation: combining investor views with market equilibrium, F Black and R Litterman, *The Journal of Fixed Income* 1.2, 1991.

4. Why banks failed the stress test, A Haldane, *BIS Review* 18, 2009.

5. Macrofinancial Stress Testing – Principles and Practices, J Vinals, International Monetary Fund (2012)

6. Consistent Information Multivariate Density Optimizing Methodology, M Basurto, Financial Markets Group, London School of Economics 2006.

7. Technical Note on Stress Testing, *IMF Country Report* No. 10/244, 2010.

8. Stress and Scenario Testing – Feedback on CP08/24 and final rules, Financial Services Authority, 2009.

9. Information Theory and Statistical Mechanics, E Jaynes, *Physical Review* 106 (4), 1957.

10. Maximum Entropy: The Universal Method for Inference, Adom Giffin, PhD thesis at the State University of New York.

ACKNOWLEDGEMENTS

Many thanks to the multi-asset risk team for the both the intellectual and quantitative content of this paper. In particular, Dr Jens Kroeske has been a key figure in helping understand the mathematical methodologies being employed in this process.

Mark Foster is Investment Director Absolute Return & Multi Asset Investing for Standard Life Investments. Chis Nichols is an Investment Director for Standard Life Investments.
