Forward-Looking Scenario Analysis

Preparing your portfolio for the unknown unknowns

Standard Life Investments

May 2016

This document is intended for institutional investors and investment professionals only and should not be distributed to or relied upon by retail clients.

Modelling future scenarios – blending belief systems and quantitative discipline

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. In this paper, we examine how expert opinions about future world states can be incorporated into projections of portfolio outcomes. We also illustrate that using a diversified multi-asset approach can effectively help mitigate losses in such extreme market scenarios.

Recognising the inadequacies of value at risk (VaR) 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 far 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 judgement
- 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
- to blend expert opinion to produce estimates of tail behaviour for different assets
- we can apply this methodology in practice, using one particular tail-risk scenario.

We then analyse the asset allocation of a typical US Global Tactical Asset Allocation (GTAA) portfolio to illustrate the robustness of an absolute return investment strategy. We show 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 a single-state distribution. Financial market data is often described as 'fat-tailed'.

This concept implies that, under a given scenario, a normal distribution of outcomes is not normal but indeed skewed and therefore has fatter tails. In other words, more extreme variations of return occur.

Chart 1 highlights the fat tails commonly observed in financial data. It shows a histogram of historical daily returns for the US equity S&P 500 Total Return Index over the 26-year period to July 2013, together with a normal approximation.

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.

Chart 1: Actual US equity return distribution versus a fitted normal curve

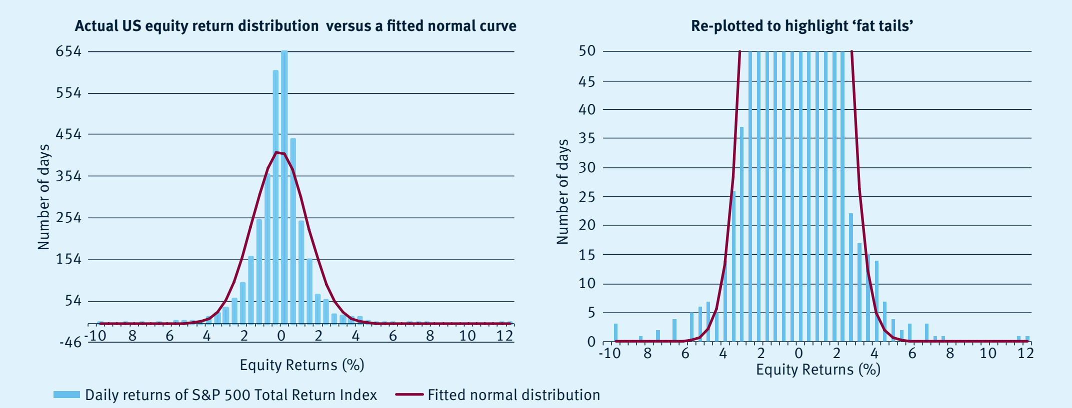

Source: Bloomberg; authors' computation as at July 2013

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 our 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. Let us consider the same return data of the S&P 500 Total Return Index but this time, we divide the daily returns into a 'high-volatility' regime (in red) and a 'low-volatility' regime (in blue). For each of these two regimes we fit a normal distribution (see Chart 2). Even with this simple approach, you can visibly note that the fat tails observed earlier are less apparent.

For the purpose of scenario analysis, this approach can be very useful as a 'regime-switching' model. A regime-switching model is able to produce fat tails as well as 'tail dependence' without the need for exotic parameter estimates.

We are therefore able to relax the assumption that market returns are identically distributed, instead modelling market behaviour by using a large number of overlapping regimes. These are encoded as a library of Monte Carlo simulations (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).

Chart 2: Actual US equity return distributions for low and high-volatility regimes versus fitted values of normal distributions

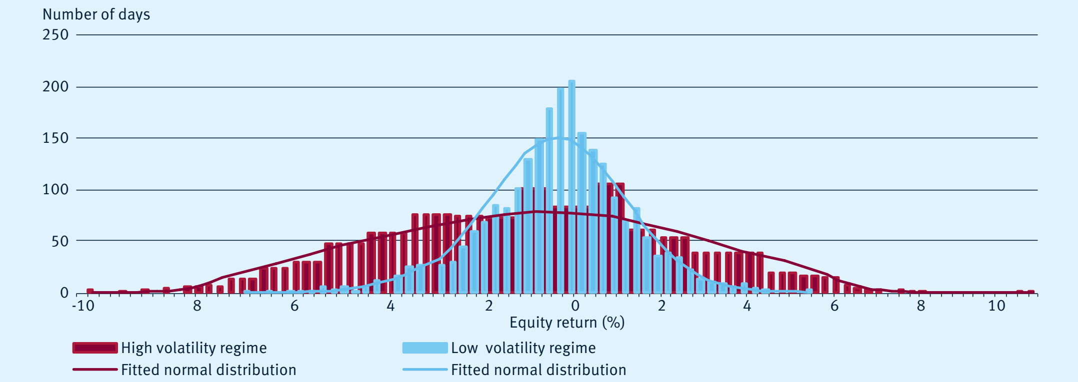

Source: Bloomberg, authors' computation as at July 2013

Blending expert opinion with quantitative analysis

In the absence of any other information, we have to treat each simulation as being equally likely. In blending expert opinion, we should create a forward-looking distribution that is consistent with the expert views but remains as close to the prior distribution as possible. It is through the process of maximum entropy that we can mathematically derive the estimates we are seeking.

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 we are looking at a fuzzy picture – expert opinion tells us the underlying image is a monkey – the mathematics we carry out produces a 'best' picture of that monkey, given the pixels (data points) that produce the image. However, if expert opinion instead tells us the underlying image is that of a cat, the output of our mathematical calculations 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 us 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

Chart 3 shows the relationship between US 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, we show the outline of a single multivariate normal distribution that has been calibrated using the entire data set.

The chart 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 tells us that during extreme falls (such as happened in 2008), they in fact tend to behave in unison. As with the US equity example in Chart 2, we obtain a much better fit to the actual data when we employ a multi-regime model – that is to say, we recognise that the relationship between equity and credit returns changes throughout the market cycle.

Chart 3: Equity returns versus credit returns using a single-state distribution

Source: Bloomberg, authors' computation; dataset March 2000 to July 2013

Chart 4 shows the library of simulated annualised returns (using Monte Carlo techniques) that incorporates the various correlation structures observed during the same 13-year period, 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.

We now consider a 'credit collapse' scenario – our expert opinion has defined this 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 (see Chart 5).

Chart 4: Equity returns versus credit returns using a multi-state distribution

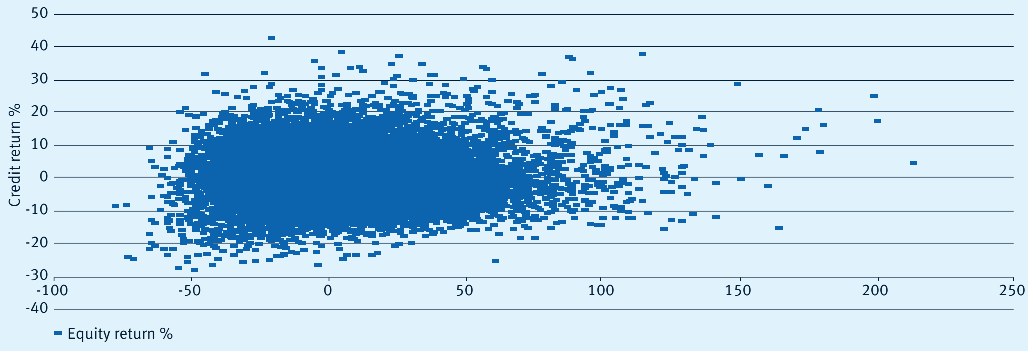

Source: Bloomberg, authors' computation, dataset March 2000 to July 2013

Chart 5: Equity returns versus credit returns under a 'credit collapse' scenario

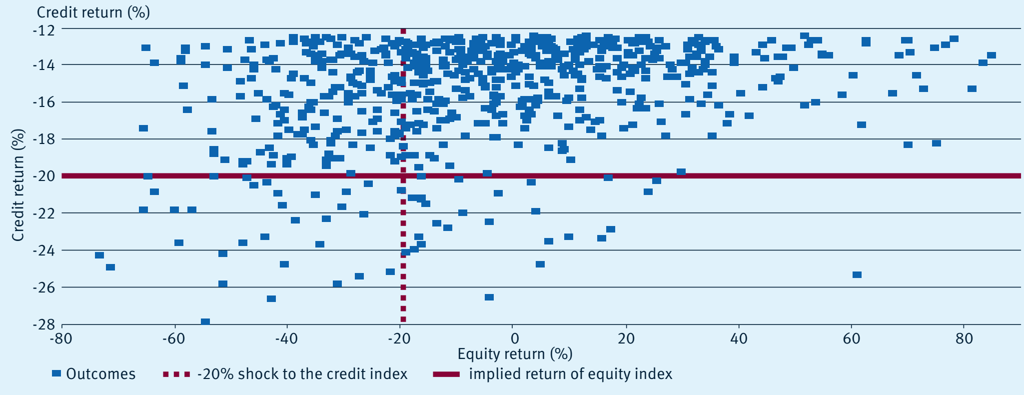

Source: Bloomberg, authors' computation as at June 2014

Practically, we can achieve our objective by re-weighting the simulations in our library, such that the average loss for the credit position is -20%. In Chart 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 a US perspective: China crisis scenario

The methodology we demonstrated can be applied to a wide range of asset classes whose returns we can model, and to a specific set of expert opinions about unlikely events.

For the purpose of the next set of analyses, we used the asset weightings of a typical US GTAA portfolio to test its robustness against a scenario we define 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.
- The economy is further hampered by the debt-fuelled investment boom/misallocation of resources/debt quality.
- 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 GTAA portfolio's asset allocation, we can derive a sense of its likely behaviour given this information. Although modelling liquid assets (i.e. where we have daily data) 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 for liquid equivalents, the less frequent pricing of these assets will typically create time delays in their behaviour in tail-risk events.

We have taken the most optimistic approach here and treated less liquid investments as static in nature. That is, we assume their prices 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 GTAA portfolio's asset allocation weightings should provide a recent reminder of the downside potential of the illiquidity premium.

Output analysis

Chart 6 shows the likely negative impacts of a China crisis on various asset classes. Unsurprisingly, equities suffer significantly, while fixed income provides some protection, albeit returns are still negative. For 'alternatives', we have considered only the liquid portion of this allocation, i.e. hedge portfolios. We can see that, as a universe, these portfolios have provided limited downside protection. However, we appreciate that in practice there will have been a wide dispersion of returns around this central figure, dependent upon manager and strategy selection.

Chart 6: The impact of a China crisis scenario

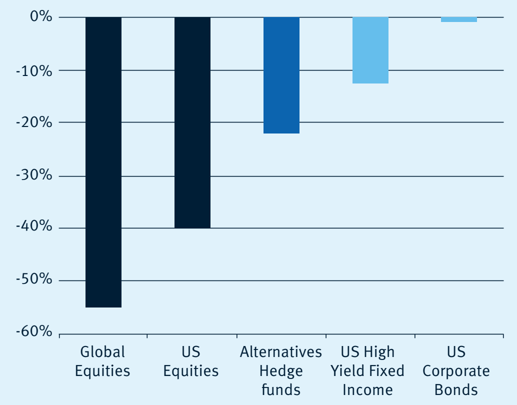

Source: Bloomberg, authors' computation

To calculate the expected losses for the GTAA portfolio we use the following weighting assumptions[1].

Weighting assumptions	
US Equities	30%
US Corporate Bonds	30%
Global Equities	20%
Global Real Estate	5%
Alternatives - Hedge portfolios	5%
US High Yield Fixed Income	5%
Alternatives – Private Equity	5%
Total	**100%**

[1] Asset allocation based on internal and Greenwich Consulting analysis

For modelling purposes, given the earlier assumptions on illiquid assets, we are effectively looking at asset movements on 90% of the portfolio (10% of the portfolio being illiquid assets).

Combining the return forecasts (see Chart 6) with the asset weightings produces an expected negative total return of -25.0% for the entire portfolio (see Chart 7).

Chart 7: China crisis forecast

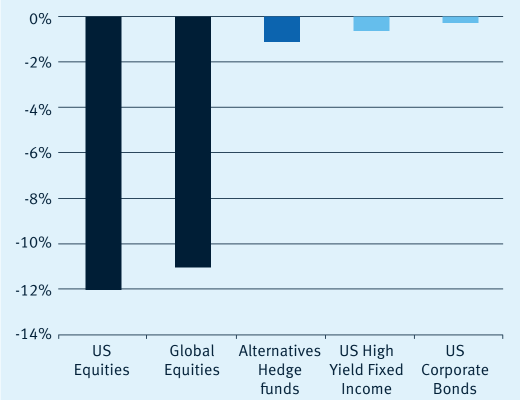

Source: Bloomberg, authors' computation

Unsurprisingly, given the large bias toward the domestic 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 suffers considerably less than equities in this extreme scenario it is still negatively impacted. Moreover, 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 portfolios as an overall investment group can give 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 portfolios

The emergence of absolute return portfolios as a mainstream global investment has been partly driven by a desire for the strong diversification benefits that hedge portfolios can offer but without the traditional drawbacks. For the next part of the paper, we consider a representative absolute return portfolio, where we have the transparency of underlying investment positions. This lets us see how future scenario analysis might help portfolio managers understand the likely behaviour of their holdings in extreme market conditions.

By way of a brief introduction, the multi-asset absolute return portfolio under analysis consists 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 portfolio 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 three-year period.

Because of the absolute return nature of this portfolio, studying its potential behaviour in extreme scenarios is a core element of the investment approach.

Performing the China crisis scenario analysis, the absolute return portfolio in this case would deliver a return of -7.5%. This compares favourably with the 'Alternatives – Hedge funds' group (-22.0%, Chart 6).

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

As can be observed, carefully selected 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.

We have focused in particular on the likely response of a US GTAA portfolio 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.

Chart 8: Absolute return portfolio constituent breakdown in China crisis totalling -7.5%

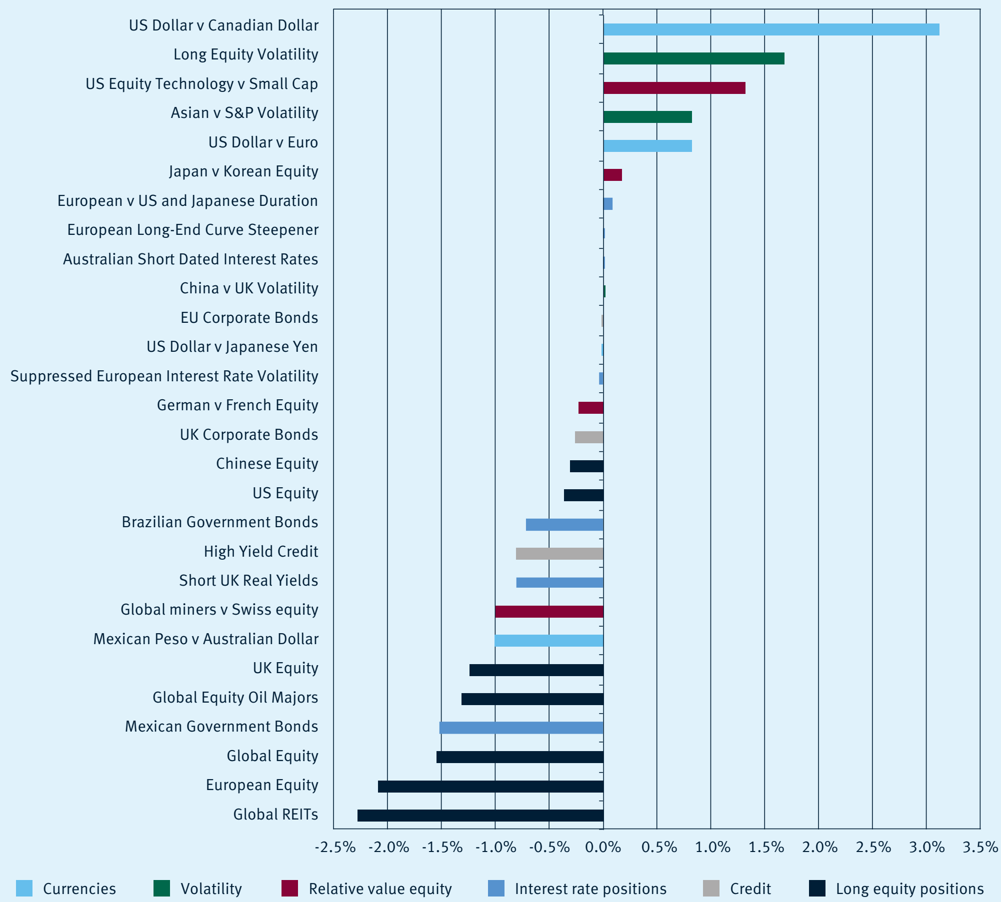

Source: Bloomberg, authors' computation as at October 2014

Organisation challenges for constructing extreme scenarios

However mathematically pleasing this analysis is, 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 may itself 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 we have not discussed, but that are essential for its usefulness are:

- the process by which expert opinions are formed
- the organisational framework around scenario analysis that ensures the process itself and its outcomes can help portfolio 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
- portfolio 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, we have demonstrated a methodology to quantitatively translate expert opinion about future world states into portfolio outcomes. Furthermore, we have shown how using an absolute return investment approach can be beneficial in helping mitigate losses in such extreme markets.

Finally, we highlight that the real value of such analysis is in the engagement of portfolio managers in the scoping and output of the work, allowing them the opportunity to improve current portfolio construction techniques.

Acknowledgements

Many thanks to our 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 us understand the mathematical methodologies employed in this process.

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 Portfolio (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.

Important Information

This material is for informational purposes only. This should not be relied upon as a forecast. research or investment advice. It does not constitute an offer, or solicitation of an offer, to sell or buy any securities or an endorsement with respect to any investment vehicle. Any offer of securities may be made only by means of a formal confidential private offering memorandum. This material serves to provide general information and is not meant to be legal or tax advice for any particular investor, which can only be provided by qualified tax and legal counsel.

Due to among other things, the volatile nature of the markets and the investment strategies discussed herein, they may only be suitable for certain investors. No investment strategy or risk management technique can guarantee return or eliminate risk in any market environment.

Third party data services disclaimer

Any data contained herein which is attributed to a third party ("Third Party Data") is the property of (a) third party supplier(s) (the "Owner") and is licensed for use by Standard Life**. Third Party Data may not be copied or distributed. Third Party Data is provided "as is" and is not warranted to be accurate, complete or timely. To the extent permitted by applicable law, none of the Owner, Standard Life** or any other third party (including any third party involved in providing and/or compiling Third Party Data) shall have any liability for Third Party Data or for any use made of Third Party Data. Past performance is no guarantee of future results. Neither the Owner nor any other third party sponsors, endorses or promotes the fund or product to which Third Party Data relates.

**Standard Life means the relevant member of the Standard Life group, being Standard Life plc together with its subsidiaries, subsidiary undertakings and associated companies (whether direct or indirect) from time to time."

If you would like to find out more about our strategies, please visit www.standardlifeinvestments.com where you will find contact details for your location.

Visit us online

standardlifeinvestments.com

INVBGEN_14_1221_Forward_Looking_Scenario_Analysis_White_Paper_TCM 0516