


Diversified Thinking.

Comparing risk parity and risk-based models in asset allocation

For investment professionals only. Not for distribution to individual investors.



Martin Dietz – Fund Manager



John Southall – Senior Investment Strategist

A growing number of risk-based concepts, and risk parity models in particular, promise an intuitive solution to the strategic asset allocation problem. In this edition of *Diversified Thinking*, we look at a range of these new approaches and explore their merits and shortfalls.

Our analysis is divided into parts. First, we discuss traditional asset allocation approaches and how risk parity attempts to address some of their potential pitfalls, focusing on naïve risk parity. Second, we explain how a risk parity approach has a bias towards low-risk asset classes and consider the likely characteristics of portfolios with large allocations to bonds. Next, we consider the wider investment beliefs and implicit assumptions that underlie risk parity and risk-based models, before discussing and comparing some more refined approaches. Lastly, we summarise what can be learned from risk parity and suggest how this could be integrated into a robust strategic asset allocation process.

Asset allocation approaches

Traditional asset allocation approaches

The key portfolio management problem is centred on maximising expected return for a given level of risk or, equivalently, minimising risk for a given expected return. As Nobel Prize winner Markowitz famously outlined, the classic quantitative approach is ‘mean-variance optimisation’. The approach characterises asset classes purely by their expected returns, and the historical volatility and correlation between them.

Although it is mathematically elegant and tractable, this theory relies upon a number of assumptions that are unrealistic in practice¹. A second criticism is that typical optimisation approaches are particularly sensitive to the expected return assumptions, which is exactly the information that is least reliable. In practical applications, small differences in expected returns, e.g. just a few basis points higher or lower, often result in large swings in the ‘optimal’ portfolio allocation. To help, constraints such as minimum and maximum weightings can be imposed. However, the portfolio asset allocation can then end up being defined more by the constraints than the optimisation.

Risk parity

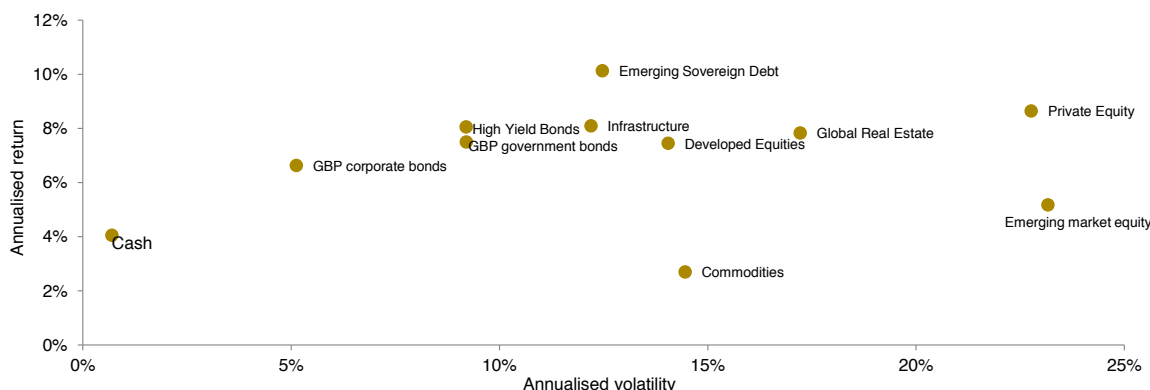
Risk parity is an approach that links asset allocation exclusively to the efficient allocation of portfolio risk. At its core, it argues that portfolios need to balance their risk exposure between underlying building blocks or risk drivers and maximise diversification.

Naïve risk parity is the most basic application of the approach, targeting the same stand-alone risk for every asset class. As individual volatilities differ, capital weights are driven by the inverse of an asset class’ volatility. In other words, volatile assets such as equities are given a low capital weighting in the portfolio, while low-volatility assets such as bonds are given a relatively high, or even leveraged, weighting.

While risk parity does not have any explicit return estimates, portfolios following this approach will still aim to have high risk-adjusted returns. Proponents of risk parity argue that efficient risk management is the key step towards generating efficient risk-adjusted returns, pointing out that historic data indicates a link between risk and return in most standard asset classes, as shown in **Figure 1**.

¹These include the assumptions that asset returns follow a normal distribution and that volatility (or standard deviation) of returns is a suitable measure of risk. In practice, many asset classes do not have this distribution and when assessing risk there is often a desire to focus on specific asset class features when markets are falling, something that standard deviation does not capture.

Figure 1. Return and historic volatility for different asset classes



Source: Bloomberg, LGIM. Data from 31 Dec 1993 to 31 Jul 2015. For illustrative purposes only.

There are two structural challenges to the risk parity approach. First, risk parity assumes a strong link between risk and return. In particular, the approach suggests that low-risk asset classes such as sovereign bonds achieve the same risk-adjusted returns as standard return-seeking assets. Risk parity portfolios thus tend to use leverage and have a significant overweight to bonds compared with more conventional portfolios. Second, the stand-alone risk focus of a naïve risk parity approach ignores the complicated dependencies between asset classes. More advanced risk parity approaches aim to incorporate these dependencies into the risk allocation.

The bond debate

Bonds are a focal point for the debate around risk-based approaches. Most risk parity approaches allocate a significant weighting to bonds, given that their low volatility and low correlation to other assets seems to justify relatively large capital allocations in order to achieve balanced risk exposure.

Historic back-testing

Proponents of risk parity approaches point out that the risk-adjusted returns on bonds have been attractive historically. The fact that the correlation between equities and bonds has been low seems to provide additional

support in a portfolio context. For example, a basic risk parity portfolio of UK equities and gilts would have outperformed UK equities over the last 20 years, and would have done this with a considerably lower level of volatility.

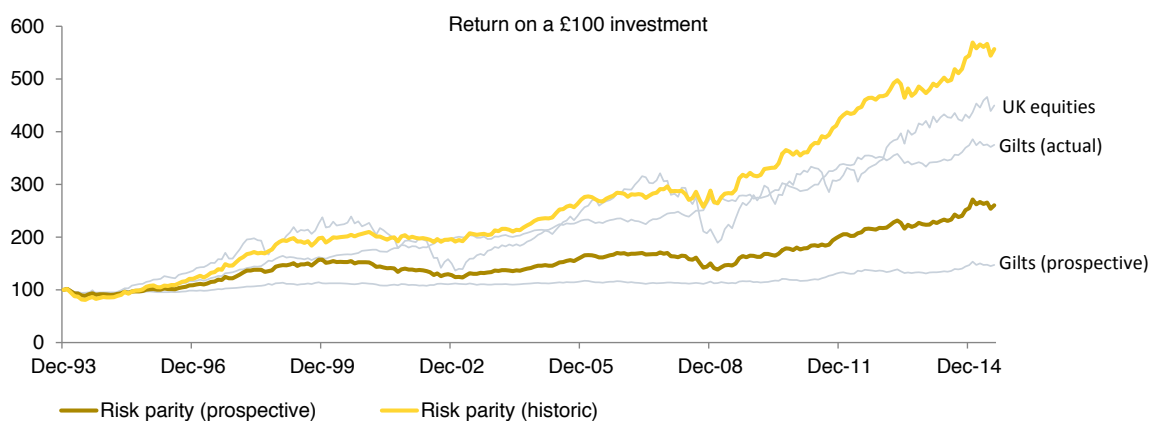
However, bond returns over the past few decades have benefited structurally by an incredible fall in yields. With both real and nominal yields at record-low levels, this bond rally would appear to be impossible to repeat over the coming years. To shed some light on the possible impact of today's depressed level of bond yields, **Figure 2** shows a back-test of a basic risk parity portfolio² from 1993 to 2015.

Table 1: Back-test of a basic risk parity portfolio

Asset class	Historic return			Bonds recalibrated to current yields		
	Annualised return	Volatility	Sharpe ratio	Annualised return	Volatility	Sharpe ratio
UK equities	7.2%	13.8%	23%	7.4%	13.8%	44%
Gilts (5-15 years)	6.3%	5.4%	42%	1.8%	3.8%	16%
Cash	4.1%	0.7%	0%	1.2%	0.0%	0%
Risk parity ²	8.3%	8.0%	51%	4.5%	7.1%	47%

Source: Bloomberg, LGIM. Data from 31 Dec 1993 to 31 Jul 2015. For illustrative purposes only.

Figure 2: Back-test of a basic risk parity portfolio



Source: Bloomberg, LGIM. Data from 31 Dec 1993 to 31 Jul 2015. For illustrative purposes only.

²For illustrative purposes, we use a portfolio of 40% UK equities and 120% gilts (implying a leverage ratio of 1.6).

The annualised returns, volatilities and Sharpe ratios (the returns achieved above the risk-free rate, per unit of volatility) are all shown in **Table 1**.

The 10-year gilt yield declined from more than 6% to under 2% over the time period and this allowed bonds to generate much higher risk-adjusted returns than equities. A risk parity approach outperformed both equities and bonds and generated a much higher Sharpe ratio.

Given current bond pricing, the decline in gilt yields of the past 30 years cannot be repeated. To get an idea of what future returns on gilts could be possible, we have recalibrated historical changes in yields and applied them to the current yield curve.³ If we start history again at the end of 1993, but base it on the current yield curve, our projection of bond returns still outperforms cash, but to a much lesser degree. The recalibrated returns of the risk parity strategy trail that of UK equities but achieve a similar risk-adjusted return (Sharpe ratio). Overall, risk parity no longer looks superior to holding plain equities.

Term premium

The risk parity approach assumes that risk taking will be rewarded. In bond markets, this implies that the duration risk inherent in long-dated bonds should result in them providing higher returns than that available on cash. However, our previous analysis suggests that the strong performance of bonds in recent decades has been largely driven by a fall in yields that cannot be repeated. When

it comes to prospective returns and the attractiveness of bonds, the important question is whether there is a term premium and, if so, how large it is.

The basic theory of the term structure of interest rates is the expectations hypothesis. According to this, the expected return of holding a long bond until maturity is equal to the expected return of rolling over a series of shorter bonds with a total maturity equal to that of the long bond. The expectations hypothesis makes clear sense for investors who focus on the certain cash flows of a bond. Other investors who do not aim to hold their bonds to maturity are exposed to interest rate risk and could require compensation for this risk. Any compensation for this risk is called a term premium.

Assessing the current size of the term premium mostly comes down to making an empirical estimate. Although this is notoriously difficult, a range of recent analysis⁴ indicates that the term premium has structurally declined over the past few decades, and that the expected returns on bonds are now close to the expected returns on cash. This may be related to increased demand for long-dated bonds from pension schemes and insurance companies (who are primarily concerned with matching liability cash flows and therefore do not demand a term premium) and, more recently, the effects of quantitative easing.

A possible conclusion from this analysis is that having high or even leveraged

exposure to bonds may not increase expected returns relative to cash and that duration risk may therefore be unrewarded in the current market environment.

Assessing the risk-reducing benefits of bonds

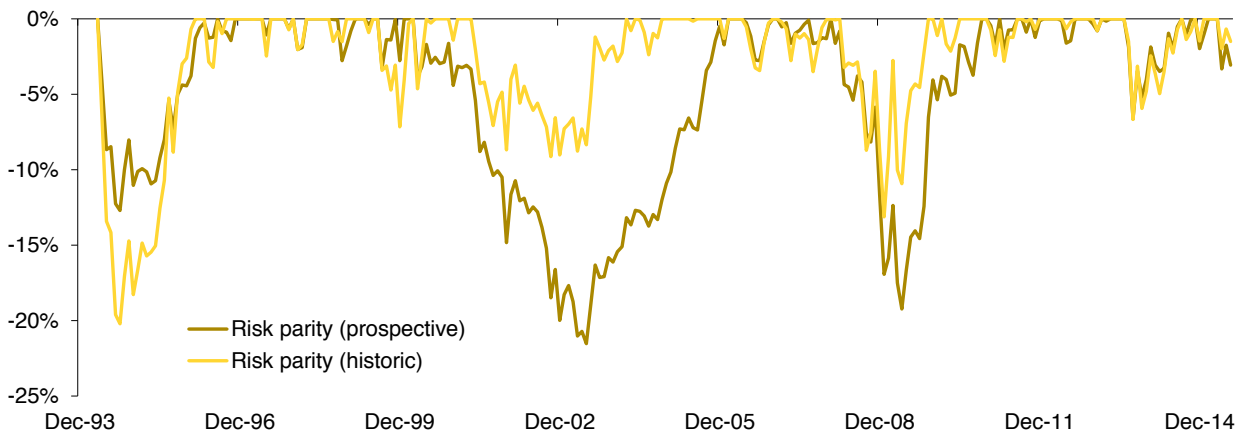
Today's lower yields are likely to moderate the risk-reducing benefits of bonds, which rely heavily on bonds' historic tendency to appreciate in a market crisis (as yields tend to fall strongly). With lower yields, the ability of bonds to generate strong returns to help offset falls in equity prices is much more limited. **Figure 3** shows how drawdowns in the risk parity strategy would generally have been more severe, had returns been recalibrated to allow for the current depressed level of yields.

Leverage

Another common concern regarding leveraged bond exposure focuses on the use of financial leverage itself. Historically, the use of financial leverage was relatively difficult to achieve. This has changed and the use of leverage is now well established as a tool for risk management, for example in liability-driven investment strategies.

Leverage removes portfolio constraints which can help to increase investment efficiency. However, the use of leverage in return-seeking portfolios creates additional complexity and governance issues for investors. At the most basic level, an unlevered multi-asset portfolio

Figure 3: Historic and prospective drawdown for a risk parity approach



Source: Bloomberg, LGIM. Data from 31 Dec 1993 to 31 Jul 2015. For illustrative purposes only.

³Our quantitative approach is to extract the historic and dynamic behaviour of the bond market and to impose the same curve shifts, twists and other adjustments, onto the current curve. We assume that short-term interest rates cannot fall below zero, and that any declines in forward rates towards zero therefore face increasing resistance.

⁴Sources: US Federal Reserve analysis in Kim/Wright (2003), Bank of England estimates in Guimaraes (2012).

spreads exposure across a range of asset classes and the limited liability of each exposure can be an important risk management feature that protects investors from the potential collapse in any individual asset class. With inappropriate use of leverage, shocks in a particular asset may spill over to the overall portfolio level. This may be of particular concern for bonds in the current low-yield environment, given that they are likely to exhibit negative skew, i.e. have a high probability of making a small gain and a small probability of making a large loss.

Overall, investors should check if they are comfortable with leverage and understand its use. This is particularly important if leverage is not offset by movements in liabilities and its usage is not capped to limit the risk of nasty surprises.

Beyond bonds: investment beliefs underlying risk parity approaches

We will now investigate the investment beliefs that underpin risk parity approaches. There are two key assumptions made by risk parity.

Assumption I: all risks are equally rewarded

The risk parity concept is based on the assumption that all risks are rewarded and, more strongly, that there is a linear relationship between underlying risk and return. A risk parity investor therefore aims to maximise exposure to underlying risks (assuming that

these drive returns) while limiting total portfolio risk. This can be achieved by efficient risk management and by diversifying risk sources (for example, if we combine uncorrelated risks, diversification means that the total risk is less than the simple sum of the parts). Risk management and diversification are key for any asset allocation approach, but the assumed relationship between underlying risks and returns allows a risk parity investor to ignore return estimation and focus on the risk component alone.

However, a ‘textbook’ assumption would be that only certain risks (systematic risk) should be rewarded, and that other risky bets (single stock risks, country or regional concentrations or buying lottery tickets as an extreme) shouldn’t yield excess returns. Effectively, this assumption puts a lot of emphasis on identifying the range of rewarded risks, and diffusing the idiosyncratic risks. The concern is therefore that a simplistic application of risk parity may incorporate inefficient allocations to unrewarded risks.

Assumption II: risk can be accurately measured

Naïve risk parity only considers the volatilities of asset classes and doesn’t consider the correlations between them, even though they impact the overall level of risk. In reality, volatility can be a poor measure of risk and ignoring interactions between asset classes is clearly an inappropriate

assumption. Other risk parity models allow for a more realistic risk structure. We discuss these in the next section.

Fundamentally, risk parity relies heavily on the measuring and managing of risk. While the estimation of expected asset class returns is perhaps more challenging than estimating risk, this is not to say that measuring risk is without its problems. Historical correlations and volatility values are dependent on the time periods chosen and tend to neglect events that have not occurred often or recently. It is easy to underestimate the risk of an asset that suffers infrequent collapses.

Different types of risk parity

Within the broader risk parity concept, there are a number of competing approaches that apply different methodologies to identify and combine the underlying risk factors.

Naïve risk parity assigns equal risk budgets to each underlying asset class, implicitly assuming that asset classes are the main risk drivers. In real-world applications, it is easy to construct cases where naïve risk parity results in counterintuitive allocations. For example, the total allocation to equities differs depending on the use of a single global equity index or a set of regional equity market indices. These issues mostly arise as a result of ignoring the correlation between asset classes. Other flavours of risk parity aim to improve on this simplistic approach.

Table 2: Comparing different risk parity model approaches

Assumption / approach	Naïve risk parity	Maximum diversification	Quantitative factors	Qualitative factors
Process	Adjust all asset classes to a common risk level (e.g. by adding leverage). Equally weight all of the re-risked asset classes	Maximise the sum of individual asset class risks (volatilities) subject to a certain portfolio volatility (calculated from the asset class weights and a covariance matrix)	Decompose the correlation matrix into independent factors. Allocate equal risk budgets to each independent risk factor	Specify / quantify the return drivers. Adjust the return drivers to make them uncorrelated. Allocate equal risk budgets to each risk factor
Allows for interactions and correlation structure between different asset classes or drivers	No*	Yes	Yes	Yes
Looks at underlying risk factors / drivers	No**	No**	Yes	Yes
Allows for unrewarded risks	No	No	No***	Yes
Allows for different return expectations per unit of risk	No	No	No	No
Allows for fact that volatility may be a poor measure of risk	No	No	No	No

* Implicitly assumes that correlations between asset classes are zero (or that they are all the same number)

** Assumes that asset classes are the most granular building blocks

*** However, there is sometimes a focus on the most dominant factors which may coincide with rewarded risks

Maximum diversification considers the correlation structure of the underlying assets in order to calculate the volatility of any given asset allocation. Within an overall target risk for a portfolio, it seeks to maximise the sum of individual risks taken. This means finding an allocation which maximises diversification between underlying asset classes.

Both naïve risk parity and maximum diversification focus their core building blocks on asset class exposure. A more flexible view is that there is an underlying structure to all individual asset classes, with a number of ‘return drivers’ or ‘risk factors’ that span across them. Each asset class can then be thought of as the sum of a number of factors. For example, small cap equities can be expressed as a combination of an equity risk factor, plus a small cap factor.

Quantitative factor models aim to isolate these factors solely by using statistics⁵. While the analysis throws out a long list of factors, only the

most significant ones are usually easy to understand and rationalise. The prime factor typically refers to equity risk. The many additional minor factors typically identified are harder to interpret. A pure application of the asset allocation process assigns an equal risk weight to all these factors and assumes that they are equally rewarded. In reality it seems more appropriate that only the first few factors reflect rewarded risks, while the remaining ones are statistical noise. There appears to be a need for a qualitative overlay that would attempt to weight the factors more appropriately.

Qualitative risk factor models are similar in concept but try to identify the rewarded risk and return drivers in a more qualitative way. As with the quantitative approach, the factors can stretch across the underlying asset classes. The aim of the qualitative approach is to have a clear economic rationale as to why exposure to any given factor should be rewarded, albeit with this rationale backed up by quantitative research.

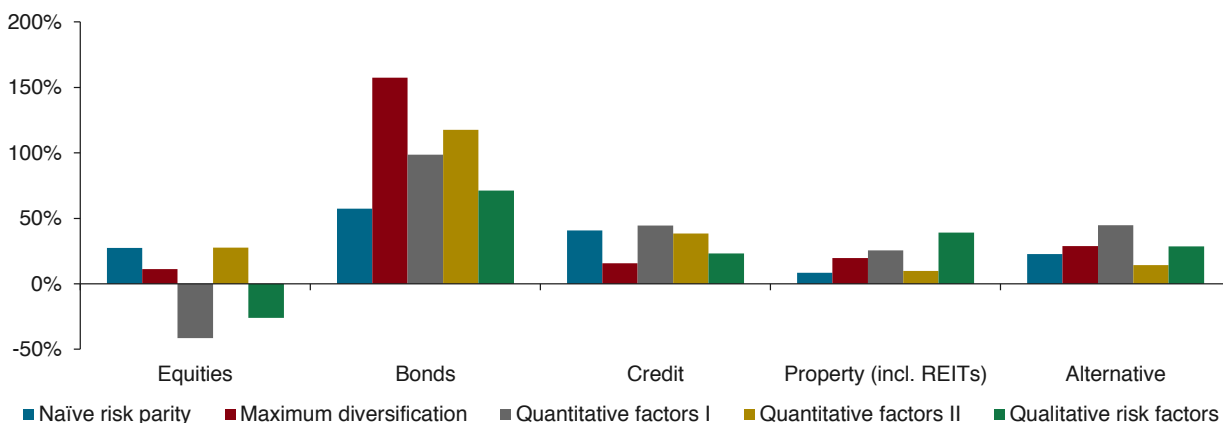
Comparing the different approaches

Table 2 summarises the implicit assumptions of each risk parity approach. All four approaches typically require the use of leverage. The quantitative and qualitative factor approaches may also require the use of short positions.

Figure 4 illustrates the range of different asset allocations that result from applying the different flavours in a multi-asset setup for unconstrained portfolios.⁶

As expected, high bond exposure is a feature of all risk parity approaches. Beyond this, there is a large variation in allocations, with the equity allocation being positive in some and negative in others. When we calculate historic asset allocation, we find considerable variability both over time and across approaches.

Figure 4: Sample asset allocation for different risk parity approaches



Source : LGIM. For illustrative purposes only.

Maximum diversification:	To improve stability, asset weights are required to be non-negative Choueifaty/Coignard (2008)
Quantitative factors I	Eigenvalue decomposition of the covariance matrix (see Meucci (2009)). Equal risk allocations to the five largest eigenvectors (to improve stability)
Quantitative factors II	Minimum torsion approach (see Meucci/Santangelo/Deguest (2014))
Qualitative risk factors	Contains equity, bond term, credit, small cap and value premia; factors are de-correlated using the minimum torsion approach

⁵Factor identification is achieved using a Principle Component Analysis (PCA) of the covariance matrix.

⁶Calculations are based on a broad selection of 24 asset classes, using a historic 10-year covariance matrix. All strategies are allowed to use leverage and are calibrated to a common volatility level of 10% to ease comparisons.

Lessons

There is no single risk parity approach and, as **Figure 4** shows, there is little consistency across the methods. While we have not discussed them all in detail, in our view none of the models are convincing on a stand-alone basis. The relationship between risk and return is more complicated than that implicitly assumed by risk parity models. Our analysis also indicates that risk parity can suffer from some of the same sensitivity issues as the methods it intends to replace.

Nevertheless we believe that there are a number of important lessons. Risk parity approaches underline the need for investors to focus on risk-adjusted returns and for risks to be managed across all asset classes. They also emphasise the need for genuine diversification and provide a framework for how to think about measuring this. Traditional approaches that don't allow for the uncertainty and high sensitivity of asset allocation to return estimates may miss this need. Finally, some of the approaches look for underlying risk drivers. This may have advantages over traditional approaches that don't take the underlying economic structure of asset returns into account.

To complement traditional approaches, there are a number of steps that could be taken. Taking account of the various 'diversification scores' implied by the different approaches is one possibility. For example, the maximum diversification approach involves maximising the 'diversification ratio' (defined as the sum of the stand-alone risks of the asset classes divided by the risk of the overall portfolio).

Given the flaws mentioned, we would be reluctant to base portfolio construction entirely on this ratio. However, it can form a useful metric to measure diversification alongside other risk/return measures. Given the apparent weaknesses of individual approaches, we prefer to look at a variety of indicators to help support robust portfolio construction.

Conclusion

Risk-based models are appealing in their simplicity and highlight some important considerations relevant to the interaction between asset classes. In particular, the intuitive focus on diversification is attractive and offers food for thought in terms of ways in which traditional approaches could be improved. But risk parity offers no panacea for the shortcomings of other asset allocation approaches and suffers from a number of shortcomings that make it hard to recommend as an approach in isolation. No matter what quantitative or qualitative process is adopted in a strategic asset allocation exercise, practical oversight and a clear focus on the objectives of the investor remain paramount throughout.

Literature

Choueifaty/Coignard (2008): Toward Maximum Diversification, *Journal of Portfolio Management*, Fall 2008, pp.40-51.

Guimaraes (2012): What Accounts for the Fall in UK Ten-Year Government Bond Yields?, *Bank of England Quarterly Bulletin* 2012 Q3, pp. 213-223.

Kim/Wright (2005): An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates, *Federal Reserve Working Paper* 2005-33.

Meucci (2009): *Managing Diversification, Risk*, pp. 74-79.

Meucci/Santangelo/Deguest (2014): *Measuring Portfolio Diversification Based on Optimized Uncorrelated Factors*, mimeo.

Important Notice

Views and opinions expressed herein are as at September 2015 and may change based on market and other conditions. This document is designed for our corporate clients and for the use of professional advisers and agents of Legal & General. No responsibility can be accepted by Legal & General Investment Management or contributors as a result of articles contained in this publication. Specific advice should be taken when dealing with specific situations; investment decisions should be based on a person's own goals, time horizon and tolerance for risk. The information contained in this document is not intended to be, nor should be, construed as investment advice, nor deemed suitable to meet the needs of the investor. All investments are subject to risk.

© 2015 Legal & General Investment Management Limited. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, without the written permission of the publishers.

Legal & General Investment Management Ltd, One Coleman Street, London, EC2R 5AA

www.lgim.com

Authorised and regulated by the Financial Conduct Authority.

M0515