Portfolio optimisation and bootstrapping

Rahul Srivatsa, Andrew Smith and Jon Lekander
Aberdeen Property Investors, London, UK

Received August 2009
Accepted October 2009

Abstract

Purpose – The purpose of this paper is to develop a more robust methodology for asset allocation for the property investment market which takes into account inherent valuation and data issues.

Design/methodology/approach – The methodology applied is that of a bootstrap, borrowed from Carlstein, and is applied to an investment universe consisting of UK equities, gilts and property. The bootstrap selectively re-samples the return time series by maintaining the economic cycle. The resulting return series is then used in the standard mean-variance optimisation (MVO) on an unconstrained basis. Finally, a "sanity" test is applied on the correlation matrix to ensure that spurious instances do not skew the results.

Findings – The bootstrapped optimisation provides a range within which the portfolio weights can be manoeuvred instead of a static point under the standard MVO. It provides a more robust methodology for asset allocation and without giving any undue significance to one year of extreme result.

Research limitations/implications – The current analysis is based on unconstrained portfolio optimisation, with a very limited investment universe. Additionally, by conforming with the MVO methodology, normality of asset returns is implicitly assumed, which is clearly not the case in the data used. Future work will also focus on an all-property portfolio.

Practical implications – The proposed methodology will prove to be useful for making asset allocation decisions, particularly in turbulent financial markets.

Originality/value – The paper focuses solely on bootstrapping with the IPD UK annual index and is particularly significant after one year of extremely poor performance of UK property. The results will be of use to fund managers and portfolio analysts.

Keywords Property, Financial markets, Market value, Assets management, Computer bootstrapping, United Kingdom

Paper type Research paper

1. Introduction

One of the long-standing issues with the property investment market has been the lack of data due to infrequent property valuations and transparency concerns surrounding the direct property market. As property valuations occur monthly (at the most) in the UK, its relevant returns series will be restricted to a monthly timeline. This is not helped by the fact that direct property transactions do not happen over an exchange but rather by a general agreement between a willing buyer and seller. This is unlike equities or bonds where the price is decided and available for all participants to view on the exchange or other such market places. Such issues make many statistical techniques widely applied in equities and fixed income incompatible with property. One of the main tools in the investment arsenal of a property fund manager is portfolio optimisation, which can be used to conduct their hold/sell analysis on their portfolio,

The authors would like to acknowledge Dr Andrew Patton, at Oxford University, for his work in bootstrapping, among others.

Journal of Property Investment & Finance
Vol. 28 No. 1, 2010
pp. 24-33

1463-578X
DOI 10.1108/14635781011020029

and verify if the model portfolio structure is compatible with the fund's investment objective. However, the key assumptions of Markowitz's Modern Portfolio Theory (MPT) (Markowitz, 1952) theory are that asset returns are normally distributed and that investors face a risk-return trade-off. It is widely accepted that most asset returns are non-normally distributed and this can be seen in the extreme tail risks in the current crisis and the long term capital management crisis in 1998. Such events are not covered adequately by a normal distribution function. In the property industry, most portfolio optimisation practices ignore the normality assumption of asset returns. To complicate matters further, the short time series of property returns data further compromises the stability of the estimated returns and covariance matrix. In portfolio literature such issues are referred to as estimation errors. Such deficiencies in the optimisation methodology could provide statistically incorrect outputs, i.e. portfolio weights. The appeal of this paper is that it works around these shortcomings rather than ignoring them altogether.

2. Literature review

There are a number of methodologies mentioned in the literature to overcome statistical issues with asset returns and MPT. Avouyi-Dovi *et al.* (2004) use an Omega function to accurately capture asset distribution, instead of the normality assumption, and asset allocation is done using a meta-heuristic optimisation method. Davies *et al.* (2003) included higher moments such as skewness and kurtosis in the portfolio construction process. By doing so, the analysis went beyond the standard risk return trade-off theory in MPT. They concluded that investors prefer increasing skewness and reducing kurtosis. Under those assumptions, they found asset allocation decision to be greatly influenced by co-skewness and co-kurtosis matrices. However, such research may be well beyond direct property analysis for the foreseeable future as historic returns are quite restrictive and extension to higher moments would require significantly more extensive historic data. There are other streams of research focusing on minimising the estimation error. One such technique is the model developed by Black and Litterman (1991) in an internal Goldman Sachs document. This model provides a quantitative framework to blend the prior view, the capital asset pricing model market portfolio, with the investor's view to arrive at a combined new distribution. The asset allocation from the Black-Litterman model is claimed to be more diversified than a plain mean-variance optimisation (MVO).

Initial studies on asset allocation decisions, in the USA, by Webb and Rubens (1987) have indicated a property weighting of at least 43 per cent in a multi-asset portfolio. More conservative estimates of property weight in a US multi-asset portfolio have been around 19-28 per cent by Gilberto (1993). Globally, Hoesli *et al.* (2004) recommend property allocation between 15-25 per cent in a multi-asset portfolio. However, this has rarely been replicated in reality with average UK pension funds' property allocation at about 4.3 per cent, according to Mercer (2008). This figure would have a downward bias as it includes fund managers with no property mandates. The cause of the unusually high weighting from academic asset allocation research has been attributed by Newell *et al.* (1996) to appraisal-smoothing and valuation lags in the UK direct property market. They also suggest increasing the appraisal-based risk estimates by a factor of 3.5 to reflect actual property risk. By adjusting returns for property smoothing, property returns can then be made comparable to other asset classes.

However, there will still be distributional issues with the adjusted returns data and its implications for parameter estimation.

The approach used in this paper to address this issue is that of bootstrapping. This paper focuses on a UK domestic investor with only UK-based assets as the investment universe. It is intended to expand this study to include more markets as previous studies, such as Hoesli *et al.* (2004), have found that based on the size of the domestic property market international investment will improve diversification of the portfolio. This approach has been used by Gold (1995) in real estate portfolios. However, their implementation of bootstrapping was that of random selection and applied to the Russell/NCREIF index in the USA.

The concept of bootstrapping data was initially implemented by Carlstein (1986). This was built upon by Politis and Romano (1995) and Politis and White (2004). Essentially, the idea is to randomise the returns series within itself based on certain rules, which can be customised to the user's choice. The choice of bootstrap selection will have bearing on the re-sampled data's relevance to the market cycle and thus on the output. Upon using the bootstrap with the right rules to match the economic cycle, we intend to iterate the optimisation across the bootstraps to obtain a (possibly) more stable output, i.e. asset weighting. This is done by bootstrapping the existing dataset, say N times, and maintaining the economic cycle within them. The bootstrapped data is then used for N portfolio optimisation scenarios. Thus, the idea would be to use these N resulting outputs (asset weightings) to indicate a range in which the fund manager's portfolio should lie. It is essential to maintain the structure within the time series of the asset classes considered otherwise the correlations generated by the bootstrap would be infeasible and hence results would be irrelevant.

The added advantage of the bootstrap is that the resulting output would be that of a range of portfolio weights. This could be more useful to a fund manager as building a property portfolio takes time, and this methodology could possibly take into account portfolio drifts.

The following sections will first introduce the dataset used in this analysis and briefly evaluate the issues surrounding them. Next, the standard MVO framework is used to provide asset allocation decision on a UK multi-asset investment universe. The shortcomings of this allocation are discussed and the need for an alternate framework, i.e. bootstrapped simulation is established. The concept of bootstrap is introduced along with its application to portfolio optimisation. Finally, the results are analysed with implications for a fund manager.

3. Data

One of the best markets, in terms of data availability, to work on a property optimisation is the UK market. The UK IPD annual index has one of the longest available historic series in global property returns. So this would be one of the best breeding grounds for any such researching techniques, with a property perspective. This paper only considers UK equities, gilts and property in the portfolio. The data used in this analysis goes back to 1948, which means it covers enough economic cycles to provide the bootstrap food for thought. The UK IPD data have been historically re-created by Aberdeen Property Investors using a cashflow methodology with results comparable to Scott (1996) (Figure 1).

From Table I, it can be seen that none of the three asset classes fits the criteria for normal distribution, i.e. zero skewness and kurtosis of three.

It should be noted that the property returns in this analysis are de-smoothed IPD returns with 1 per cent transaction cost on top, as this creates a level playing field between the asset classes (Figure 2). The 1 per cent transaction cost is approximately based on an assumption of 6 per cent purchase cost, 1 per cent selling cost and a 20 per cent portfolio turnover rate. The UK IPD return series are de-smoothed using Geltner's adjusted (Geltner, 1993) filter, which is given below:

$$R_u(t) = [R^*(t) - (1-a)^* R^*(t-1)]/a$$

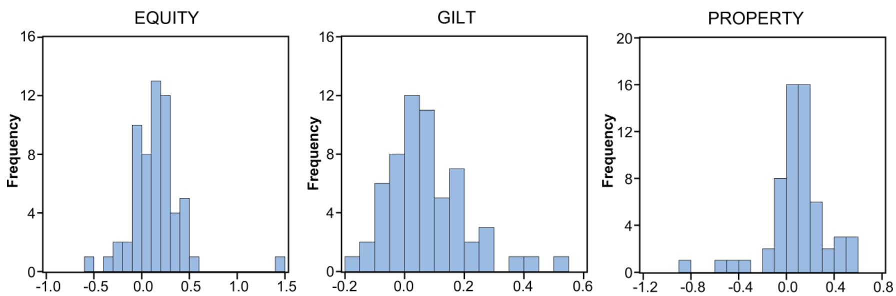

Figure 1. Histograms indicating the distribution of the asset returns

(%)	Equities	Gilt	Property
Mean	14.71	7.59	9.69
SD	27.30	13.29	23.95
Skewness	1.75	1.02	−1.10
Kurtosis	11.40	4.36	6.26

Table I. Key statistics of asset classes (from 1948-2008)

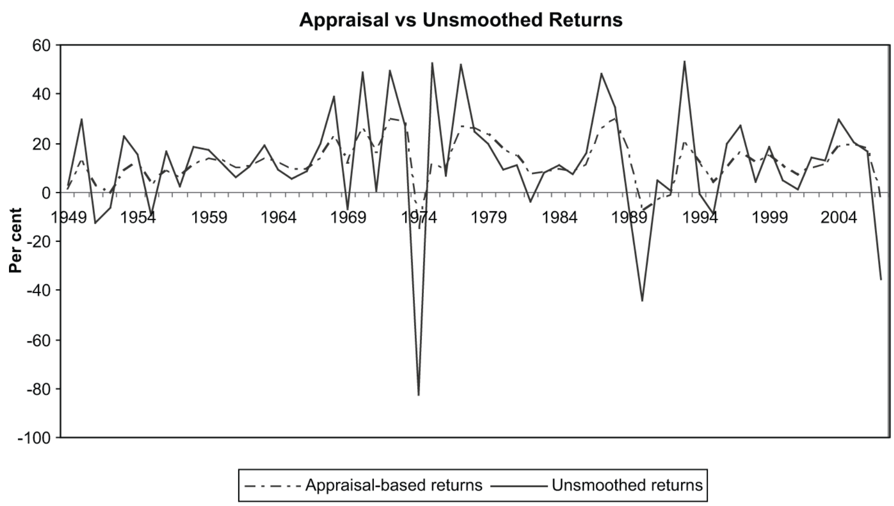

Figure 2. Appraisal-based returns against unsmoothed property returns (1948-2008)

where:

$R_u(t) =$ the unsmoothed property returns at time t.

$R^*(t) =$ is the observed appraisal-based return at time t, a is the smoothing parameter.

4. Methodology

4.1 Standard mean-variance optimisation

The typical asset allocation process would involve the MVO, whereby the portfolio is optimised on a risk-return basis with the underlying set of assets. The optimal portfolio is computed by picking the portfolio with the maximum Sharpe ratio and the underlying weighting is used for asset allocation decisions.

For this analysis, a risk-free rate of 5 per cent is used. Upon implementing a traditional MVO on an unconstrained basis between 1977 and 2008, we get the following results (Figures 3 and 4).

The result indicates that gilt and property weighting comes down as we move higher on the risk spectrum, which makes intuitive sense as equities are typically on the upper end of the risk scale. For a risk-free rate of 5 per cent, the optimal portfolio has an equity weight of 42 per cent, gilt weighting of 58 per cent and no property weighting.

It should be noted that when this exercise was repeated for data from 1947 onwards, an unusually high risk spectrum was generated for a similar range of returns compared with the result from using 1977-2008 (Figure 4). This is mainly due to the early 1950s that saw extreme volatility in equity returns. Such extremities can be unduly accounted for in the MVO. This instance of optimisation resulted in an equity weight as high as 85 per cent with property and gilts at about 7 per cent. Such a portfolio would be ill-equipped to weather a downturn in the equity markets (see Figures 5 and 6).

Upon running an optimisation excluding this period, the risk fell back to more conventional levels (below 15 per cent) with the portfolio better diversified.

4.2 Bootstrap portfolio optimisation

In order to take into account the lack of diversification in the result above and the short sample space used for a stable portfolio output, further analysis was conducted by

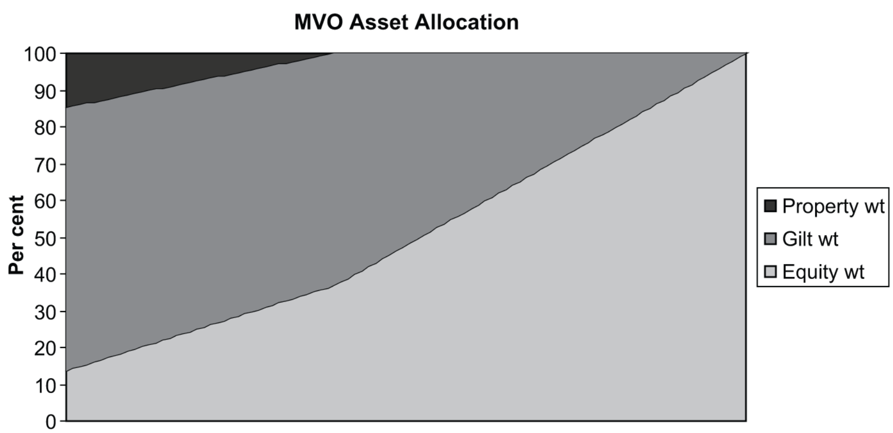

Figure 3.
MVO asset allocation strategy using data from 1977-2008

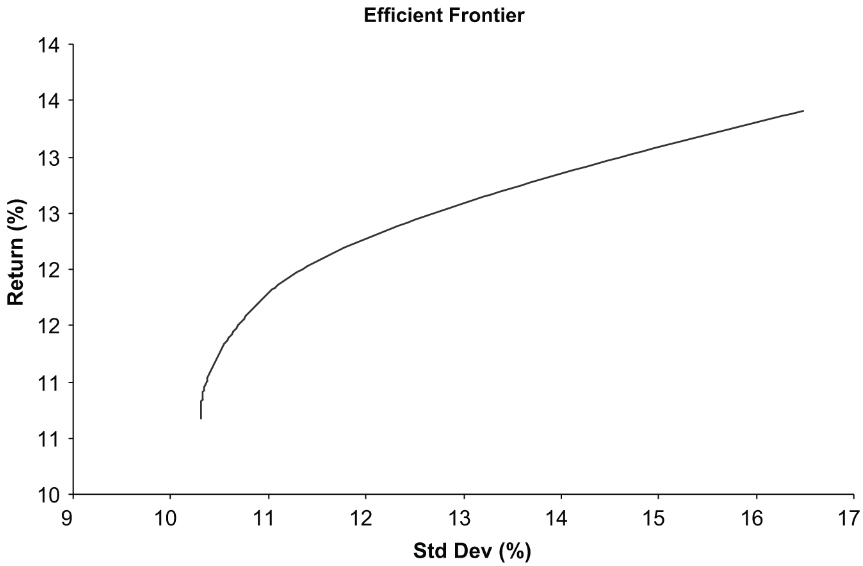

Figure 4.
MVO efficient frontier for a standard MVO (data from 1977-2008)

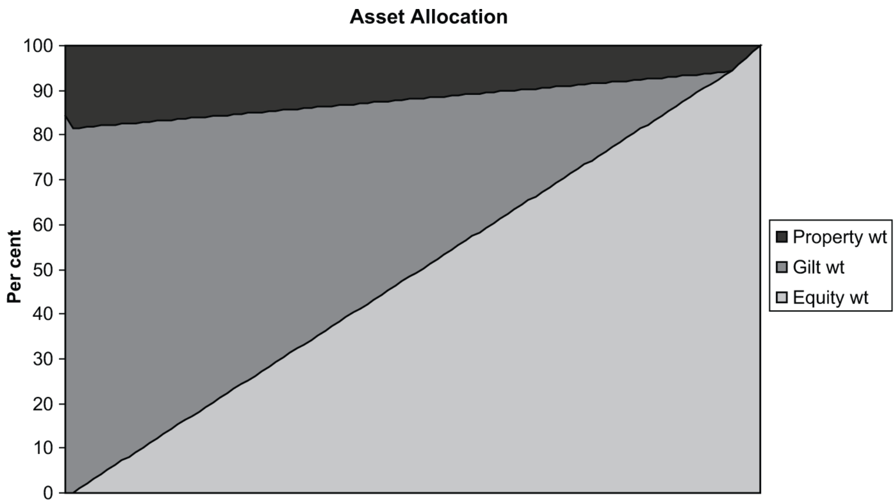

Figure 5.
MVO asset allocation strategy using data from 1948-2008

running a bootstrap on the underlying data. This re-creates a return history for each of the three asset classes based on certain rules on autocorrelation. Upon creating the bootstrapped data, portfolio optimisation was run for each of these instances. At this stage, optimisation was run on an unconstrained basis. From the resulting efficient frontier, the optimal "market" portfolio is picked by calculating the Sharpe ratio with a risk free of 5 per cent. Thus, if 100 bootstrap simulations are run the final output would be 100 efficient frontiers each representing the optimisation output of a possible return history for the three asset classes. Finally, from each of the 100 efficient frontiers we

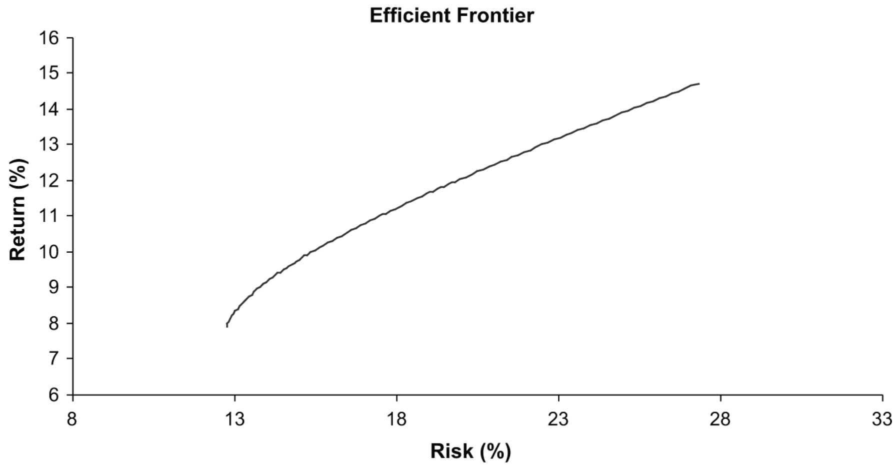

Figure 6.
MVO efficient frontier for a standard MVO (data from 1948-2008)

pick the optimal portfolio based on the Sharpe ratio. The portfolio weights output for such an analysis, with data from 1977-2008, is presented in Figure 7.

The most interesting output from this analysis is the distribution of weighting received for each asset class. Each vertical intercept would represent the optimal portfolio for a particular bootstrapped simulation. For example, for an optimal portfolio risk of 8 per cent the suggested asset weighting is:

- *Equity* – 29.5 per cent;
- *Gilt* – 63.8 per cent; and
- *Property* – 6.7 per cent.

By analysing the entire distribution spectrum, it was found that the average property weighting was at 8.44 per cent, with a standard deviation of 7.1 per cent. Under assumption of normality distribution of asset weighting, this would suggest the standard MVO output, of optimal property weighting being 0 per cent, is an event over

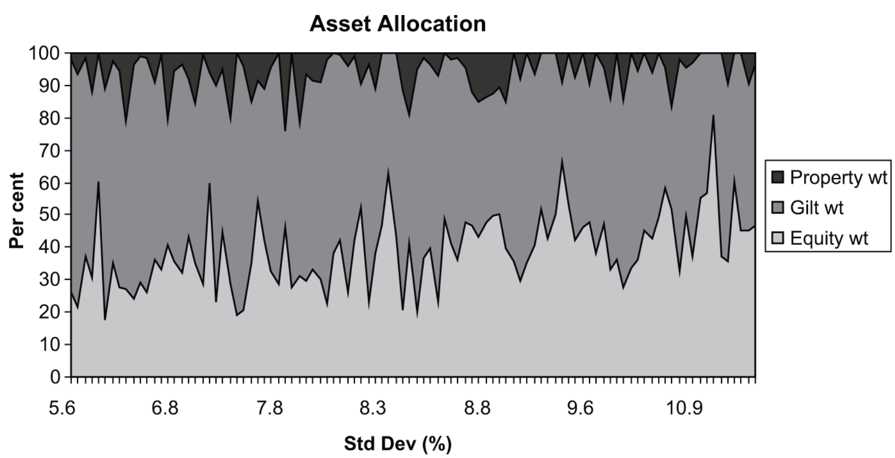

Figure 7.
UK multi-asset allocation using a bootstrap technique (data from 1977-2008)

one standard deviation from the mean of the bootstrapped distribution. Gold (1995) referred to this bootstrap distribution as "fuzziness".

Among the three asset classes, property is perceived to be more of a long-term investment. However, any changes in the underlying asset performance and valuations will create a drift in the portfolio weights and causing it to be "sub-optimal". The important question to be asked is: will the fund manager need to re-align the portfolio with the output from the portfolio optimisation? The implication of Table II is that it provides a fund manager with an acceptable region within which to operate the fund, as re-balancing and changing property weights on an annual basis might prove expensive. From Table II, we can also observe that property has the least variance in its weighting. This could indicate that it is the least sensitive (among the three asset classes) to changes in asset allocation decision.

To further test the bootstrap, a "sanity" filter was added to ensure that spurious correlations between asset classes do not skew the final results. As an initial screen, a filter is applied to ignore bootstrap simulations that have negative correlation between the asset classes. The motivation for this rule was that the horizon used for calculating the correlations is 30 years, and over such a long time period a hypothesis could be to expect asset classes to have positive correlation with one another (Figure 8).

Upon adding the sanity filter, it was found that the risk spectrum had been translated to a narrower band and property weighting had come down further. However, the general points remain the same. This chart provides an indication to a fund manager on where the fund's asset allocation could be. This still signals that

Table II. Asset weighting statistics for the bootstrapped portfolio optimisation

(%)	Equity	Gilt	Property
Mean	38	53.6	8.4
SD	11.9	13.4	7.1

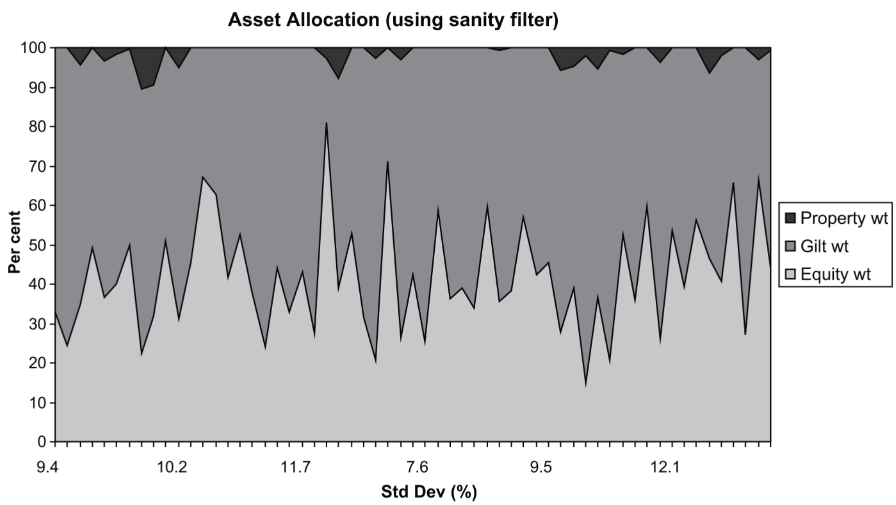

Figure 8. UK multi-asset allocation using a bootstrap technique (data from 1977-2007) and a sanity filter on the correlation matrix

property is the least volatile in terms of asset re-allocation decisions and that property weightings of up to 6 per cent can be expected in multi-asset portfolio with 95 per cent confidence, under the normality assumption (Table III).

5. Conclusion

This paper contrasts the performance of the standard MVO and a bootstrapped version, in the context of their application to the property market. The instant advantage of using the bootstrap in the property market is the option to ignore the limited historic returns data as a base case scenario, and instead focuses on a range of possible outcomes. Traditional MVO solutions are static in nature and rely extensively on the stability of the input data. Instead, we propose a bootstrapped portfolio optimisation that provides an acceptable region within which the fund manager can allow their portfolio weighting to vary. By doing so, the methodology accepts the inherent flaws in the property returns data and works with it. It should be noted that the investor in this analysis is UK-domiciled with a domestic portfolio alone. It is widely accepted in the industry that property weighting in a multi-asset portfolio should be around 5-15 per cent to be meaningful. The bootstrap simulation with a sanity filter on the correlation matrix, on the other hand, suggests a range of between 0-6 per cent with 95 per cent confidence. The bootstrapped simulation also suggests that property would be the least sensitive in terms of asset re-allocation. Similar results are observed in the simulation without the sanity filter with the bootstrap suggesting range up to over 15 per cent weighting for property.

In this paper, the optimisation is done on an unconstrained basis and is not realistic. Upon constraining the asset allocation, one would expect the equity weighting to come down to the benefit of gilts and property. However, the bootstrap provides increased confidence in results as it covers realistic bounds in terms of asset weighting.

The final conclusion of this paper is the robustness of the bootstrap methodology. 2008 was a turbulent year in the global financial markets with one year of data completely changing portfolio parameters. The MVO output, for data up to 2007, indicates a property weighting of 3 per cent. This drops all the way to zero when 2008 is included. Such an abrupt reaction to one year of data may be unwelcome. The bootstrap optimisation, on the other hand, does not indicate so. For the data up to 2007, the average property weighting was 6 per cent with the standard deviation of 6 per cent. When 2008 is included in the parameter estimation, the average property weighting was 1.5 per cent. Although the trend of asset allocation to property is similar to the MVO output, that of a property underweight rating, the allocation is not zero. Similar conclusions can be drawn from equity asset allocation in 2007 and 2008. The MVO methodology suggests a 16 per cent decrease in equity weight whereas the bootstrap methodology suggests a 12 per cent decrease. This is of crucial importance as it limits any knee-jerk reactions in asset allocation.

Table III.
Asset weighting statistics for the bootstrapped portfolio optimisation with sanity filter

(%)	Equity	Gilt	Property
Mean	41.9	56.6	1.5
SD	14.1	13.7	2.6

This analysis could be of significant use to a fund manager with a direct property target as building a property weighting takes time, unlike equities and bonds. With the recent financial markets volatility, a previously optimal portfolio could soon become sub-optimal and require an asset re-allocation, which can be expensive in property markets. It is here where the bootstrap analysis could help by providing the manager with a range within which to operate the portfolio and to limit any knee-jerk reaction. The more crucial application of this methodology could be in an all-property portfolio where data issues are more prevalent.

Future work will focus on combining non-normal distribution functions with the bootstrap simulation. This will combine a statistically correct view on property as an asset class with the bootstrap to provide a range within which the portfolio weighting can be permitted.

References

Avouyi-Dovi, S. and Morin, S. (2004), "Optimal asset allocation with Omega function", working paper, Banque de France, Paris.

Black, F. and Litterman, R. (1991), "Asset allocation: combining investor views with market equilibrium", *Journal of Fixed Income*, Vol. 1 No. 1, pp. 7-18.

Carlstein, E. (1986), "The use of subseries values for estimating the variance of a general statistic from a stationary sequence", *The Annals of Statistics*, Vol. 14 No. 3, pp. 1171-9.

Geltner, D. (1993), "Estimating market values from appraised values without assuming an efficient market", *Journal of Real Estate Research*, Vol. 8 No. 3, pp. 325-46.

Gilberto, S. (1993), "Measuring real estate returns: the Hedged REIT Index", *Journal of Portfolio Management*, Vol. 9 No. 3, pp. 94-9.

Gold, B. (1995), "Why the efficient frontier for real estate is fuzzy", *Journal of Real Estate Portfolio Management*, Vol. 1 No. 1, pp. 59-66.

Hoesli, M., Lekander, J. and Witkiewicz, W. (2004), "International evidence on real estate as a portfolio diversifier", *Journal of Real Estate Research*, Vol. 26 No. 2, pp. 161-206.

Markowitz, H. (1952), "Portfolio selection", *The Journal of Finance*, Vol. 7 No. 1, pp. 77-91.

Mercer (2008), *Asset Allocation Surveys and Market Profiles, 2008*, Mercer, London.

Newell, G. and McFarlane, J. (1996), "Risk estimation and appraisal-smoothing in UK property returns", *Journal of Property Research*, Vol. 13 No. 1, pp. 1-12.

Politis, D.N. and Romano, J.P. (1995), "Bias-corrected nonparametric spectral estimation", *Journal of Time Series Analysis*, Vol. 16, pp. 67-103.

Politis, D.N. and White, H. (2004), "Automatic block-length selection for the dependent bootstrap", *Econometric Reviews*, Vol. 23 No. 1, pp. 53-70.

Scott, P. (1996), *The Property Masters – A History of the British Commercial Property Sector*, E & FN Spon, London.

Webb, J. and Rubens, J.H. (1987), "How much in real estate? A surprising answer", *Journal of Portfolio Management*, Vol. 13, pp. 10-14.

Corresponding author
Rahul Srivatsa can be contacted at: rahul.srivatsa@aberdeen-asset.com