

Andrew Baum, Tony Key, George Matysiak, Joakim Franson

Attribution Analysis of Property Portfolios

ERES Conference, Athens, June 1999

1. Introduction

Property investors, increasingly, use performance measurement - or 'benchmarking' - services. They exist, first and foremost, to show whether a portfolio has achieved a rate of return better or worse than the 'market' average, or met investment objectives specified in a more sophisticated fashion. After benchmarking has answered the question *by how much* did we out- (under-) perform the benchmark?, there is an inevitable demand for 'portfolio analysis' which addresses the question *why* did we out- (under-) perform the benchmark?

An ideal system of portfolio analysis would identify the contribution of all aspects of portfolio strategy and management to relative returns. It would separate, for example, profits earned on investments from returns on held properties. Those are two distinctly separate activities with different return and risk characteristics, and reflect different features of management 'skill'. Among held properties, relative return may be influenced by anything and everything from the broadest allocation of investment between sectors to skill in selecting tenants, negotiating rent reviews, and controlling operating expenses.

In practice, the heterogeneity of individual properties and complexity of property management mean that the contributions of different functions and skills to portfolio performance are hard to disentangle. This paper is concerned with the one tool - 'attribution analysis' - which is found in all performance measurement systems in a precisely quantified form. Attribution analysis seeks to separate (at least) two components of a portfolio's relative return. The first is relative return which is due to 'structure' - the allocation of investment to 'segments' of the market with different average rates of return. The second is 'stock selection' - the choice of individual assets within each market segment which have returns above or below the averages for that market segment.

Attribution analysis is, therefore, of growing importance in property fund management - not just in terms of analysis, but also in the specification of investment objectives, the selection of managers, and setting performance-related rewards. Yet the academic and professional literature which deals with attribution of relative returns in property fund management is very thin. The literature on portfolio analysis for equities - the original source of the attribution technique - is not only surprisingly scanty, but sets out several apparently different methods of defining and calculating attribution components. Following that literature, suppliers of property performance measurement services are also adopting conflicting conventions.

This paper aims to clarify the potential confusion about the application of attribution analysis to real estate portfolios. Its three primary objectives are:

- To review, and as far as possible reconcile, the varying approaches to attribution analysis evident in the literature.
- To give a clear statement of the purposes of attribution analysis, and its meaning for real-world property managers.
- To show, using real portfolio data from IPD's UK performance measurement service, the practical implications of applying different attribution methods.

2. Attribution analysis - definitions and methods

The standard approach to the analysis of equity portfolios (see, for example, Hamilton and Heinkel, 1995) starts from three primary contributors to portfolio return: policy, structure and stock. (Unfortunately, the terminology for the last two contributors varies between sources. 'Structure' may alternatively be described as 'timing' or 'asset allocation'; 'stock' as 'selection' or 'property score'.)

Policy is the fundamental selection of the benchmark against which the portfolio's performance is to be measured. The simplest and most common benchmark would be the 'universe' average of all investors. In all measurement systems, however, many portfolios will be measured against sub-sets of that universe consisting of portfolios comparable in size, or structure, or investment objectives, or sets of properties (for example shopping centres, or high-yielding industrials) in which the portfolio has chosen to specialise. The policy selection of the appropriate benchmark is clearly critical to all subsequent conclusions about portfolio performance - and a question neglected by the literature - but is not dealt with in this paper.

Here we are concerned only with structure and stock selection. By structure we mean: the allocation of portfolio weights to 'segments' of the market - typically but not necessarily defined by a mixture of property types and geographical locations. By stock we mean the selection of individual investments within each segment which deliver returns above or below the average for that segment.

This is also a commonly accepted classification in equities analysis: Burnie, Knowles and Teder, for example, confirm that

Attribution models typically differentiate the overall effect of active management into two sub-effects: asset allocation and security selection.

2.1 The choice of segmentation

The simple statement conceals an initial choice in any attribution system which is critical to all that follows: what segments of the investible universe should be used to define 'structure'? Burnie, Knowles and Teder state that:

To be useful as a tool for evaluating portfolio management, performance attribution analysis should be carried out within a framework that mirrors the investment policy and the decision-making process particular to the fund under examination. A comprehensive attribution methodology will account explicitly for each key component of the portfolio management process.

In that view, the segment structure should reflect the way in which the managers of each individual portfolio choose to regard the 'structure' of their investible universe - how that universe is broken down for the purposes of analysis, forecasting and setting target weights. In practice, it would be extremely difficult for performance measurement services to operate such a finely differentiated analysis; and the value of the analysis would be much diminished, because it would not be possible to compare allocation and selection skills across portfolios.

For practical purposes, there has to be a standardised segmentation applied to the attribution analysis of all investors, at least as a first step. Several considerations bear upon the choice of segmentation: statistical, practical and convention.

- Statistically each segment should contain a sufficient number of properties for the average return to be reasonably robust: that is, each segment should ideally only reflect *systematic risk*.
- Following on from the previous point, the optimum segmentation of the market is that which statistically explains the most variance in individual property returns.
- Practically, segments most usefully cover property categories or areas for which property market information, with supporting information on (say) demographic and economic factors, are readily available to support analysis and forecasting.
- And, by convention, segments will be most acceptable to investors where they follow the generally accepted ways of dividing and analysing the market: it would be difficult to offer an analysis service in the UK, for example, which did not show City of London offices as a 'segment'.

Empirical work on the first two aspects indicates that: a large number of properties are required in order to get down to systematic risk levels and, on average, some 10 per cent of an individual property's return is accounted for by a broad market factor (Brown and Matysiak, 1999).

The considerations listed are, moreover, not always compatible. It may well be that statistical analysis would identify an optimum segmentation based on (say) quality of building, or yield band which is more powerful than one which lends itself to rigorous analysis, or conforms with the normal way of viewing the market. This point is discussed below. In real-world performance analysis services, the search for an appropriate segmentation will tend to resolve quite rapidly to a mixture of the dominant property types (shops, shopping centres, offices, industrials) and the geographical areas (either towns or regions) linked either to well-recognised property 'markets', or the city/regional boundaries used in the production of official statistics. For illustration, the market segmentations used in IPD's UK and Swedish performance analysis services are shown in the Technical Appendix.

2.2 Style

Property fund managers may adopt asset allocation positions which are different from the segment weighting of the benchmark for a variety of reasons. This may be the result of tactical asset allocation, so that views of likely market returns influence a manager to adopt an underweight or overweight position relative to the benchmark in an attempt to produce out-performance. It may be the result of strategic asset allocation or policy, where issues other than pricing – for example, liability matching – influence the asset allocation mix. It may also be the conscious or unconscious result of the style of the fund manager.

To some extent, this term has been appropriated (or misappropriated) by followers of Sharpe's so-called 'style analysis' (Sharpe, 1988). This purely statistical analysis method attempts to measure investment policy retrospectively by estimating the goodness of fit of returns with benchmark returns on investible asset types. This is not

what we mean by style. We use the term in an attempt to reflect more commonly-used judgements of investment style in fund management. Is the manager top-down or bottom-up? Is the manager a value manager or a growth manager? Such style judgements are, we argue, very applicable in property fund management, and yet are rarely used.

Our definition of style is some persistent bias in the property portfolio structure which is the result of preference or of habit. It may lead to long-term out-performance, or it may not. In equities fund management, value managers have produced very poor returns in recent years. Yet it is not to be expected that value managers cease to be value managers. In UK property fund management, large prime stock has out-performed secondary stock in recent years; yet it is not to be expected that smaller active managers (many property companies, for example) will change their style.

Style may be associated with investment houses, with individuals or with funds. Arguably, there is far too little explicit differentiation between house styles in property fund management. This may change.

2.3 Themes

As noted above, segment structure will typically be defined by reference to property use type and broad geographical region. Property fund management houses such as Henderson Investors invest in forecasting systems which enable managers to take a tactical view on prospective returns in the market ‘segments’ which are determined by this classification. It can be seen, then, that definitions of fund structure are of necessity rather stable. This is rational: data shows that use types in the UK have shown persistent cycles of under- and out-performance over a long period, a strong example being the continued out-performance of retail in the early 1980s and the early 1990s.

Table 1: Mean average deviations, IPD Irish Funds, 1986-1995

Segment	Mean average deviation
Sector	2.70
Locations within retail	4.40
Locations within offices	1.80
Sub-sectors within retail	4.60
Age within retail	4.30
Age within offices	1.30
Age within industrials	5.70
Size within retail	2.40
Size within offices	1.20
Size within industrials	2.40

Source: IPD

However, sector (type)/region segments are not necessarily optimal in permitting out-performance by asset allocation. Table 1 shows the mean average deviation between the mean return on the IPD index and the returns across different segment classifications for the Irish market over the period 1986 – 1995. The table suggests that the mean difference between the return on the individual sectors and the market as a whole in each individual year (‘the window of opportunity’) is less than the mean difference between the return on different age groups within the industrial market.

Hence an industrial fund benchmarked against the industrial market would be expected to have a higher tracking error than a diversified fund against the market – and potentially greater out-performance, if the correct allocations were made.

While sector allocations may not, in Ireland over the period 1986-1995, present the maximum potential for out-performance, there is no reason why this might not be the case over some future period. An excellent manager may be expected to anticipate when this might be. Equally, he would be expected to anticipate at what point size becomes important – or age. This is what we mean by themes.

The asset allocation process ideally takes accounts of themes as well as of standard segmentation. These may be new themes – sensitivity to changes in electronic shopping, for example – or they may be standard, such a high yield/low yield. Themes differ from styles, because there is no necessary persistence in the manager's preference for segments; and themes differ from structure, because there is no persistence in the segment classification nor reliance on external performance measurement standards.

2.4 City selection

An attribution system will preferably be stable and holistic. One major attraction of a regional classification in the UK, for example, is its completeness of coverage of UK property. However, this does not mean that fund managers will more effectively control risk and seek out-performance by categorising their holdings in this way.

A regional forecasting system may or may not be effective in identifying regional markets which will out-perform a national benchmark. Even if it is, this may not be of much use to the fund manager, because he/she may not recognise the region as a useful way to think about the market. A more technical challenge to the usefulness of the region is the possibility that there may be greater windows of opportunity within a region than between regions.

The north west of England, for example, has been suffering net out-migration and forecasts for population growth have been negative throughout the 1990s. Within the region are two major cities, Manchester and Liverpool, one of which has strengthened its position in the national hierarchy and one of which has suffered extreme decline. A fund manager is likely to have considerably different views of Liverpool and Manchester; a property forecast for the north west region is of very limited value to him.

Work at Henderson has suggested that for US and UK cities the windows of opportunity (mean average deviations from the mean) have been considerably greater at the city level than the regional level. In addition, work at Henderson has suggested that considerably greater forecasting accuracy is associated with town or city level work than with regional forecasts. We believe that city selection is a vital input into fund management strategy.

However, portfolio structure is difficult to categorise by city. This is not a holistic system, because even if every city and town in the UK were covered by the benchmark's databank there would still be outliers that fall outside defined city boundaries. This presents an attribution problem.

3. Calculating attribution scores

From the above, we may conclude that no attribution system should be regarded as absolute. The choice of segmentation that 'structures' the whole of the analysis will often represent a compromise between conflicting objectives. The use of a standardised segmentation across all investors will inevitably tend to be a baseline for the analysis of portfolio performance. But the existence of a standardised segmentation does not preclude the possibilities that alternative ways of dividing the market may offer a more powerful explanation of variation in portfolio returns, or that segmentations customised to the objectives or decision-making processes of individual investors may represent the ultimate ideal.

After the choice of segmentation, a second critical choice is the precise method of calculating the attribution scores. Here the literature not only offers a morass of varying terminology, calculation methods and mathematical notations, but also a disagrees on how many attribution components there are, and how they should be interpreted.

The pioneers in the field are Brinson, Hood and Beebower. They identify three attribution components: timing (which analogous to structure in our terminology), stock selection, and an 'other' or 'cross-product' term. Indeed, in their formulation of the attribution components, the cross-product term is effectively a residual component that, mathematically, reflects an additional *combined* contribution of timing and selection. Their interpretation of what they term timing and selection components broadly coincides with structure and stock selection components as defined in this paper. However, they do not offer an explanation of how the 'other' term relates to the objectives or management of the portfolio. Subsequent authors, and suppliers of performance measurement services, divide into two camps.

Liang, Hess, Bradford & McIntosh, Burnie, Knowles & Teder, and the main European performance measurement suppliers WM (all assets) and IPD (property) either follow a decomposition method which calculates structure and selection scores that account for the whole of relative returns without a cross-product component, or prefer to incorporate the cross-product term in either the structure or selection component.

According to Burnie, Knowles and Teder, the cross-product term:

...represents the interaction of two other attribution effects but which is not itself directly attributable to any one source of active management. It is therefore usually reallocated to another attribution effect or, if it remains isolated, is an ambiguous term whose value may exceed the measured effects of active management, thus rendering analysis results inconclusive.

While Liang et al state that the use of a two-component method is recommended:

... on the basis of simplicity and ease of interpretation. Little is lost in terms of usable information, and much 'noise' is avoided in efforts to explain the results to persons unfamiliar with the nuances of the calculation.

Hamilton & Hienkel, and the Property Council of Australia follow the three-component route, and go beyond Brinson et al in suggesting how the cross-product term may be related to management decisions. So, as put by Hamilton & Hienkel:

...Cross Product credits a manager for overweighting an asset class in which he or she outperforms the properties in that asset class in the RCPI (Russell Canadian Property Index).

Similarly, the PCA view the cross-product as showing the potential gain from choosing to allocate to a segment on a prior view that manager selection skills in that segment are strong.

Burnie et al are alone in offering a both clear statement of the reasons why they prefer a two component approach, and a suggestion that different ways of constructing the two components may be appropriate for portfolios constructed in different ways. On the first point:

To permit an unequivocal evaluation of each type of fund management decision, an attribution analysis should avoid error terms or unattributable, ambiguous components of management value.

On the second point, they propose that if the portfolio is structured by 'top-down' decisions (ie target weights by segment set on the basis of expected average segment returns), a calculation method which in effect combines the stock selection and cross-product terms is appropriate. If the portfolio is structured 'bottom up' (ie selecting assets on the basis of their expected individual return, and letting those choices determine the segment weighting), a calculation method which combines the structure and cross-product terms is appropriate.

The attribution schemas from different sources differ in some other respects, mostly concerning definitions of relative return and portfolio weights which avoid meaningless but small residual terms in annual and compounded analysis. But the central questions for suppliers and users of portfolio analysis services flow out of the central choice between two or three attribution components, or the flexible combination of both approaches. The remainder of this paper will address, using theoretical and empirical approaches, the following questions:

- Is there any way of deciding from the underlying mathematics which of these choices is right or wrong?
- How different do the results from different methods look when applied to illustrative examples or a large number of real portfolios?
- Are there features of property as an asset which imply a choice of attribution method different from other assets?

4. Results from different attribution methods

Case 1, derived from one of Henderson Investors' consulting projects, stands as an example of the differences in the message delivered to a fund manager by different choice of attribution methods.

Case 1

A fund achieved the following result (using simplified arithmetic for demonstration purposes) in 1994. Taking the three component attribution method:

Out-performance (1.0) = structure (0.1) + stock (-0.4) + cross-product (1.3)

What do these results signify concerning the relative importance of structure and stock?

If the cross-product is treated as part of stock selection, as in the most common two component system used by IPD and others:

Out-performance (1.0) = structure (0.1) + stock (0.9)

If the cross-product is allocated to structure, as proposed by Burnie et al for a portfolio constructed by bottom-up selection of individual assets with passive structure:

Out-performance (1.0) = structure (1.4) + stock (-0.4)

The choice of method is clearly non-trivial in this example. Different methods show results which differ in direction as well as scale.

Case 2

The performance of a European property share vehicle which was managed by a UK fund management house and which Henderson Investors analysed over 1998 was as follows, net of the effect of cash:

Out-performance (-2.9) = structure (-0.1) + stock (-2.0) + cross-product (-0.8)

The fund was overweight in countries where stock selection was poor and underweight in countries where stock selection was good, especially the UK. It would not be a surprise to the UK manager to learn that the stock selection score was better in the UK, but it may be distressing for him to realise that the stock selection under-performance was exaggerated by nearly a full point because of fund structure. Did he take account of expected superior UK stock selection in his asset allocation?

Case 3

'Specialist IPD funds come up trumps again' (Property Week, April 1 1999, p15). This report shows that specialist managers achieved 14% in 1999 against a benchmark of 12.2%.


Specialist managers include several large property owners with non-diversified portfolios, sector or location-biased. This may offer evidence that sector specialisation increases the likelihood of out-performance. There is also some evidence of this in the literature regarding REIT performance, and a 1992 University of Reading report (Ackrill, Barkham and Baum, 1992) found clear evidence of better share price performance among specialist UK property companies.

However, subsequent econometric work by Matysiak and Brown (1997) did not find any statistical evidence in favour of consistent out-performance, i.e. stock selection ability, for a mixed sample of specialist and non-specialist UK property companies. It appears that the jury is still out on the question of value-added by UK property companies.

IPD records for a large number of real portfolios over a long run of years can give a fuller picture of the results for real portfolios produced by different attribution methods. Figure 1, as an introduction, how much of the total variation in relative returns across portfolios is explained by each element of a three component attribution, and how that level of explanation changes over lengthening analysis periods. In the chart, the stock selection score from the two component method is the sum of the two upper elements.

(1997 is the last year of the analysis. The number of portfolios covered runs from 102 over 17 years to 235 over the final year.)

Figure 1 - % of Variation in Relative Returns Explained by Attribution Components



Source: IPD UK Portfolio Analysis Service

The contribution of structure to variation in returns obviously depends on the scale of differences in return across market segments. It reached a maximum in the boom and slump of the late 1980s and early 1990s, when there were spreads of up to 30 points between the strongest and weakest markets. Structure therefore accounts for 18% of the variance in relative returns annualised over 17 years, but 42% of the variance in returns over nine years when the influence of the boom and slump is at its greatest.

Tables 2 to 4 add more descriptive statistics on relative returns and the attribution components from all methods for three analysis periods. These results are of general interest in showing the ranges in performance and manager skills across property portfolios, and the extent to which skills are inter-related. High specific risks mean that property portfolio returns are more widely spread than equity and bond portfolios: 10% of managers underperformed the IPD benchmark by more than 6 points through 1997; over longer periods, good and bad years naturally tend to even out, but 10% of managers have still underperformed the benchmark by at least 2 points per year when measured over ten years.

Over a run of years, there is a weak correlation between structure and stock selection scores across portfolios - so managers who are good at allocation are a little more likely to be good at selection, or generally good at management - but not a strong enough correlation to suggest the skills are closely integrated.

Table 2 - Attribution Analysis Summary Statistics: over 10 years

	Relative Return	Structure Component	2 Component Stock	3 Component Stock	Cross Product
Mean	0.4	0.5	-0.2	0.0	-0.2
90 th percentile	2.6	1.8	1.4	1.3	0.7
75	1.4	1.2	0.6	0.7	0.2
50	0.4	0.5	-0.2	0.0	-0.1
25	-0.7	-0.2	-1.1	-0.5	-0.7
10 th percentile	-2.0	-0.7	-1.8	-1.2	-1.1
Minimum	-4.5	-2.3	-3.7	-6.0	-3.2
Maximum	6.8	5.7	5.1	3.0	5.7
Std Deviation	1.9	1.0	1.4	1.2	0.9
Explanation	100	41	59	35	23

Correlation matrix

Relative Return	1.00	0.70	0.85	0.64	0.44
Structure Component		1.00	0.22	0.17	0.10
2 Component Stock			1.00	0.75	0.53
3 Component Stock				1.00	-0.17
Cross Product					1.00

Table 3 - Attribution Analysis Summary Statistics: over 5 years

	Relative Return	Structure Component	2 Component Stock	3 Component Stock	Cross Product
Mean	-0.6	-0.1	-0.5	-0.4	-0.1
90 th percentile	1.8	0.8	1.5	1.1	1.0
75	0.7	0.3	0.6	0.4	0.4
50	-0.6	-0.2	-0.4	-0.2	-0.1
25	-1.6	-0.6	-1.4	-1.0	-0.7
10 th percentile	-3.3	-0.9	-2.7	-2.1	-1.2
Minimum	-6.9	-2.6	-6.7	-10.4	-5.7
Maximum	6.2	2.0	6.6	3.5	10.8
Std Deviation	2.0	0.7	1.8	1.5	1.3
Explanation	100	13	87	51	36

Correlation matrix

Relative Return	1.00	0.46	0.94	0.69	0.46
Structure Component		1.00	0.11	0.14	-0.01
2 Component Stock			1.00	0.71	0.51
3 Component Stock				1.00	-0.23
Cross Product					1.00

Table 4 - Attribution Analysis Summary Statistics: over 1 year

	Relative Return	Structure Component	2 Component Stock	3 Component Stock	Cross Product
Mean	-1.4	-0.3	-1.1	-0.8	-0.3
90 th percentile	3.0	1.0	2.8	1.5	2.0
75	0.8	0.4	0.8	0.5	0.7
50	-1.3	-0.2	-0.8	-0.8	-0.3
25	-3.8	-1.0	-3.1	-2.1	-1.7
10 th percentile	-6.2	-1.7	-5.8	-3.5	-3.0
Minimum	-12.2	-7.0	-11.1	-11.5	-7.1
Maximum	12.4	8.0	11.6	7.4	11.3
Std Deviation	3.9	1.3	3.6	2.3	2.5
Explanation	100	11	89	43	46
Correlation matrix					
Relative Return	1.00	0.42	0.94	0.68	0.73
Structure Component		1.00	0.09	0.08	0.06
2 Component Stock			1.00	0.71	0.78
3 Component Stock				1.00	0.11
Cross Product					1.00

In this form, however, the statistics offer no assistance in deciding between the merits of different attribution methodologies. The ranges and standard deviations in the component values are broadly similar, so it is not the case that using one attribution method is more likely to generate wild and destabilising results than another. And all correlate strongly enough with relative total returns to suggest they have a role to play in the explanation of performance (though this is more or less guaranteed by the way they are constructed).

Different attribution methods do, however, would offer managers different messages in a fairly large number of cases. The stock selection scores from the two methods carry a different sign in 15% of portfolios measured over 1 year, rising to 20% of portfolios measured over 17 years. This raises the question of whether the differences in results from the attribution methods are systematically related to differences in the characteristics of portfolios.

Table 5 shows the correlations between performance analysis and two portfolio characteristics: the Gini Coefficient is a measure of specialisation, with a high Gini score indicating a larger departure from benchmark weights; size is simply the average capital employed in the portfolio.

As before, results are shown for different periods of analysis up to the end of 1997. They vary with the period of analysis employed. In 1997, more specialised portfolios showed a weak tendency to under-perform, due to adverse selection, and large portfolios a weak tendency to out-perform, due to positive structure and selection.

Table 5 Portfolio Performance & Characteristics: Correlation Coefficients

	Relative Return	Structure Component	2 Component Selection	3 Component Selection	Cross Product
Portfolio Specialisation (Gini coefficient)					
1 year	-0.29	-0.04	-0.29	-0.19	-0.24
5 years	-0.04	-0.03	-0.03	0.03	-0.08
10 years	0.00	0.03	-0.02	0.04	-0.09
Portfolio Size (average capital employed £m)					
1 year	0.20	0.13	0.17	0.17	0.10
5 years	0.15	0.11	0.12	0.09	0.06
10 years	-0.12	-0.28	0.04	0.02	0.04

Source: IPD UK Portfolio Analysis Service

Over longer periods, there has been no relationship at all between any feature of portfolio performance and degree of specialisation. Over five years, larger portfolios have maintained a mild advantage in both structure and selection, but over ten years they have had poor structure (due to a high London office weight in larger portfolios), and no advantage in selection.

For the purposes of this discussion, the main point to note from Table 5 is that neither of the attribution methods appears to be superior in picking up links between portfolio characteristics and achieved returns - in particular a relationship between specialisation and return that might be assumed to appear in the cross-product. A closer inspection of the relationships, moreover, does not indicate that the overall lack of correlation does not mask clusters of portfolios for which that relationship is significant.

Further statistical work along these lines remains to be done. We propose to examine the stability of the stock selection scores within each portfolio over runs of years, with the hypothesis that a method which shows more consistency in this aspect of management is more likely to be conveying useful information to the manager. And we will examine whether there is any association within individual segments between high weights and strong selection scores to point to the merits of specialisation.

5. Attribution and portfolio management

From the above, we conclude that there is nothing in the mathematical construction of different attribution methods, nor (pending further tests) in the real-world portfolio results they produce, that dictates a preference from one attribution method over another. We therefore move on to consider whether the choice of attribution methods might depend upon the way in which a portfolio is managed, taking up Burnie et al's suggestion that different attribution formulae might be applied to portfolios structured top-down from those assembled from the bottom up. We would add the observation that the nature of that process is in some key respects different in property portfolios from equity and bond portfolios.

5.1 Top down portfolio building

The two component method embodies the classic top-down model of portfolio construction. Policy dictates a benchmark against which the portfolio is to be measured, specified in terms of a portfolio weighting by segment. An 'allocator', working with market analysis and forecasts, decides which segments are likely to out-perform or underperform the overall benchmark return, and (perhaps taking into account relative risks) determines a target weighting for the portfolio. Other things being equal, segments expected to out-perform will be over-weighted, taking 'bets' against the market. The scale of the bet will depend on confidence in forecasts, perhaps on permitted deviations from the benchmark specific in policy, otherwise the manager's willingness to accept tracking error against the benchmark.

Once the target weights have been set, the management task passes to 'selector'. Selector chooses the specific assets to be held in each segment, with the target of choosing assets which are expected to out-perform the benchmark average for that segment. In equities, the assets will (most likely) be shares in individual companies. In property, they will (most likely) be individual buildings.

Evaluating the contributions of allocator and selector to the relative performance of a portfolio constructed in this fashion is given a strong focus if it is assumed that an overall performance bonus will be split between them in proportion to that contribution.

The contribution of allocator is unambiguous. In both attribution methods, it is calculated as the product of segment weights and the benchmark performance in each market segment. The structure components within each segment and overall will show whether or not allocator's forecasts of benchmark performance have turned out to be correct, and the extent to which allocator was prepared to depart from benchmark weightings.

Under the two component attribution methodology, the task of the selector is to hold in each segment of the portfolio assets which will on average produce returns above the segment average for the benchmark. If all selector's choices match the segment averages, the selection score will be zero, under both methods of attribution. In the two component attribution, out-performance in any segment will be weighted by the portfolio weight; in the three component method, it will be weighted by the benchmark weight. Under the two component method, therefore, selector's contribution to return is 'how much *did* you add to the performance of this portfolio *given* the portfolio weights set by allocator?'. Under the three component method, selector's contribution to return is 'how much *could* you have added to the performance of *average* portfolio?'.

There is a strong difference between property and other assets in the underlying nature of selector's task, which bears upon the choice between these two specifications of selector's contribution to performance. It is set by the fact that individual property stocks are indivisible. In an equity portfolio, it can be assumed that the relative return within each segment achieved by selector is independent of the portfolio weighting set by allocator. The weight segment weight could have been varied by buying or selling the same company stocks originally selected, maintaining the same segment return. In property, that is not the case. Varying the segment weight would involve buying or selling properties with varying rates of return, and thus changing the rate of return.

The stock selection score in a three component attribution might, accordingly, be seen as a 'pure' reflection of selector's skills. It shows the contribution those skills could have made to the performance of *any* portfolio. It is therefore perhaps the most appropriate method where - as in equities - selector's skills *could* have been applied in any portfolio. It is less clearly appropriate, and arguably inappropriate, if - as in property - selector's skills may have turned out differently is a different weighting had been set.

The point may be sharpened by considering the extreme cases of rewards and incentives offered to selector awarded a bonus proportionate to the stock selection score. Given high specific property risks, segment portfolio return is most likely to show large differences from the benchmark segment average where portfolio weights are very low. For a single building, with would not be uncommon to see deviations from a benchmark of 50 points or more in any one year. Under two component attribution, stock selection scores from these extreme results are given a low weight in the overall analysis because the portfolio weight is low. Under three component attribution, these extreme stock selection scores for individual segments may be magnified by grossing up with the benchmark weight.

In property, there is a strong possibility that rewarding a selector using stock scores from the three component attribution method would produce a perverse incentive structure. Adding to segment return in segments where the portfolio is over-weight - and that task is more difficult - would be diminished in the overall stock selection score for the portfolio. Adding to segment return in segments where the portfolio is underweight - which will happen from time to time purely by chance - will be magnified in the overall stock selection score for the portfolio. And, since selection scores in segments where very few buildings are held are likely to be more erratic, the overall measure of selector's skills provided by the three component stock selection score will be more likely to vary wildly from on year to the next.

And, finally, the three component method leaves the cross-product term which cannot be awarded to either allocator or selector - an indeterminate element of performance which cannot be related to the manager inputs to relative return in the way we have defined them.

We would therefore see a strong case in favour of the two component attribution method in property portfolios under two limiting conditions. First, that the portfolio is structured 'top-down', setting the tasks for allocator and selector in the way we have defined them. Second, that these definitions of tasks are clearly understood by the users of the results.

Some property portfolios are, in our experience, structured by the top-down process we have outlined. Among large investors, most make extensive use of forecasting and quantitative portfolio analysis; many will have an explicit division of responsibilities between 'strategic' allocators and 'tactical' selectors - at least in principle. In practice, property illiquidity, heterogeneity and indivisibility all, of course, make any highly systematic approach to portfolio management difficult to stick to. For smaller investors, the structure of the process may be less formal and more opportunistic. And there are investors in all size ranges who choose to structure their portfolio by a completely different process. Where the portfolio is not structured in line with our 'standard' model,

the two component attribution method may not provide managers with the most appropriate indicators of performance.

Three alternative methods of portfolio construction will be tested against the choice of attribution methods: passive structure, backing selection skills, specialist portfolios, and alternative segmentations.

5.2 Stock selection with passive structure

Passive structure, in its purest form, could be defined as the selection of properties expected to out-perform the all-property benchmark average irrespective of their type or location. It would be followed by an investor who either does not believe that there are any segmentations of the market which contain meaningful information about individual property performance, or who has no ability to predict segment performance. Under this regime, all the relative performance of the portfolio is attributable to stock selection skills. Under either attribution method, the structure score could be interpreted as the 'cost' of eschewing a segment allocation strategy. If the structure score varies randomly over a run of years, the investor may be correct in sticking to a passive strategy. If it proves systematically negative, the investor might conclude that others are making better use of forecasting and allocation skills.

Since, by definition, no allocation skills are being applied to the construction of the passive portfolio, there are splits of rewards or incentives issues to be considered. Logically, the structure component would be taken as random with respect to manager decisions, and management reward would be based wholly on the stock selection score from the two component attribution method.

5.3 Backing selection skills

A portfolio constructed by backing selection skills offers a more interesting, and probably more common, case. Here managers choose to hold high weights in segments where selection skills are believed to be strong (perhaps on the evidence of track record). Here the task of the allocator is redefined to take account of *both* the overall performance of market segments and the skills of the selector when setting portfolio weights. In this case, the three component method of attribution offers a useful distinction between the relative inputs to portfolio performance. As before, structure score measures allocator's forecasting ability. The stock selection score measures selector's skills in the purest form. And the cross-product measures how far pre-judgements of selection skills have proved to be correct. For the reasons given in our discussion of top-down portfolio construction, it would still be problematical to use the three component stock selection score as basis for selector's rewards. It may be possible with further consideration to devise a quantitative method for splitting the cross-product contribution to overall portfolio performance fairly between allocator and selector.

5.4 Specialist portfolios

A specialist portfolio could be taken as an extreme case of backing selection skills. Here the portfolio is narrowly structured on segments where selection skills are believed to be exceptionally strong, possibly in the belief that such as concentration will in itself improve selection skills. Attribution analysis as it has been defined above may no longer apply to these portfolios, because portfolio structure is defined by 'policy'

rather than manager discretion. Under these conditions, an attribution analysis using a standard segmentation would show the benefits or otherwise of the overall policy choice. The performance of the manager is most appropriately measured against a benchmark limited to the segments pre-determined by policy. Within those segments, special attribution analysis by sub-segment could provide information on the skills applied within that specialist area.

6. Conclusions and further work

6.1 Conclusions

Attribution analysis of property portfolios is a tool, not a theorem. We do not believe there are grounds for a definitive choice between attribution methods which lie in either the mathematics of their construction, or the character of the results they will produce. Like any tool, attribution analysis may need to be adapted to different tasks and circumstances, and should be employed only with clear understanding of its function.

In the ideal, an attribution analysis which was 'carried out within a framework that mirrors the investment policy and decision making process particular to the fund under examination' might be flexible in choice of benchmark, the segmentation of the portfolio and the benchmark for analysis purposes, and the attribution method used in that analysis.

This is certainly the direction which mature property performance services such as IPD's in the UK have taken. There a standardised two-component attribution analysis is merely one element of performance measurement which separates out all the underlying inputs to the attribution, and allows the fund manager to see directly, segment by segment, where relative performance is being determined. For many managers, the standard analysis will be no more than a starting point for more specific research using custom benchmarks, and alternative segmentations. The next step along that path – currently under development – is a system which will allow managers to construct benchmarks and segmentations based on any of the array of attributes which may be associated with variation in property return.

Offering informed users a choice of attribution method would be a desirable extension of that flexibility. In passing, we would add that the challenge of quantifying manager contributions to return in property is far greater than in equity markets, because property 'stocks' once 'selected' are also directly managed. For property investors, a positive stock score may reflect the quality of on-going management rather than the quality of the original choice of investment. A fully developed property attribution structure would also be capable of making that distinction – though the quality of information on the costs and effectiveness of property management makes that a more distant objective.

Overall, we suggest that there are three potential sources of performance. Out-performance can be derived from allocating by sector according to forecasts of sector performance, from stock selection, or from allocating by sector according to superior stock selection skills.

While this may be true, it is not necessarily of value in deciding remuneration. It is difficult to imagine how the sector allocation decision, whether dependent upon expected superior stock selection within fund segments or on forecasts of superior returns in certain segments, could be separated. One allocation decision will be made, whether determined primarily by forecast available returns in the market or by forecast available returns from the current portfolio. Given this, the two term attribution makes most sense for establishing rewards.

In the standard attribution method, it may be over-simplistic to attribute returns to stock and structure, and to allocate the cross-product to stock. The cross-product is a function of the stock selection contribution in each sector and the bet made in that sector. Hence, if a manager is highly competent in one sector and not in any other, the rational course of action would be to increase the bet in that sector (subject, of course, to the manager not holding a very negative view about that sector's performance relative to the market). Is this, then, a measure of the effectiveness of a fund's style?

The cross-product may be an economically powerful explanation of investment strategy. It has been suggested that it measures the contribution of sector specialisation. Property companies have been encouraged to specialise (see literature): the cross-product may be the contribution to performance they are encouraged to access. This specialisation is a clear example of style. A successful forecast-driven allocator should earn high structure scores. A bottom-up manager who is good at stock selection should earn high 'pure stock' scores. A successful specialist should earn high cross-product scores.

However, interest in the value of the cross-product per se may be misguided if it is 'not itself directly attributable to any one source of active management'. This will be true if the same person/process allocates money to segments for whatever reason: either because average property in that segment is expected to out-perform, or because the managers in that segment are expected to out-perform, or because of a combination of both. Subject to risk, and within the confines of benchmarks, allocations will be made because existing stock plus new cash in segment A is expected to do better than existing stock plus new cash in segment B. This can be due to the people involved, the potential of the existing stock, or because of the market. Whatever the rationale, it is likely that one person or process can make that decision, hence identification of the cross-product is not essential.

It may therefore be that the two term system should continue to set the standard for popular consumption, but that the professional manager should begin to understand the relative contributions of forecast-driven structure decisions and the sector-specific stock selection skills of his teams(s). This may require a high level of competence in understanding the situations in which meaningless results may be produced.

6.2 Further work

Within the agenda that includes this working paper, there remain some empirical areas to explore, digging further into the IPD data reported in Section 4. Both two the three component attribution results will be tested to find which are associated with specific portfolio characteristics such as specialisation in individual segments, and which may offer more powerful and consistent explanations of the relative returns on individual portfolios.

From that, we will offer final conclusions on the most appropriate applications of attribution for portfolios with varying characteristics or methods of construction.

More widely, this work highlights how little is known, in a formal rather than anecdotal sense, about the decision-making processes in property portfolios, or the ways (outside of performance measurement services) managers attempt to evaluate the costs and efficiency of their decisions, structures and staff. There is a fruitful ground for behavioural and business management research in these areas.

References and further reading

- Ackrill, A, Barkham, R and Baum, A (1992), *The Performance of UK Property Companies*, University of Reading
- Brinson, G, Hood L and Beebower, G (1986) Determinants of Portfolio Performance, *Financial Analysts Journal*, 42:4, 39-44
- Brown, G and Matysiak, G (1999) *Real Estate Investment: A Capital Market Approach*, Prentice-Hall (forthcoming).
- Burnie, S, Knowles, J and Teder, T, (1998), Arithmetic and Geometric Attribution, *Journal of Performance Measurement*, Fall, p59-68
- Hamilton, S and Heinkel, R (1995), Sources of Value-added in Canadian Real Estate Investment Management, *Real Estate Finance*, Summer, pp 57-70
- Liang, Y, Hess, R, Bradford, D and McIntosh, W (1999), Return Attribution for Commercial Real Estate Investment Management, *Journal of Real Estate Portfolio Management*, Vol 5, 1, pp 23-30
- Matysiak, G and Brown, R (1997) A Time-Varying Analysis of Abnormal Performance of UK Property Companies, *Applied Financial Economics*, 7, pp 367-377
- Sharpe, W (1988), Determining a Fund's Effective Asset Mix, *Investment Management Review*, November/December, pp 59-69

Technical appendix

The dominant method of performance measurement expresses the performance of the portfolio against a benchmark as a relative return, based on the ratio of the two rates rather than the simple difference:

$$\text{Relative Return} = ((1 + \text{Portfolio Return}) / (1 + \text{Benchmark Return}) - 1)$$

So a portfolio return of 10% against a benchmark return of 5% gives a relative return of 4.8%:

$$\text{Relative Return} = 1.10 / 1.05 - 1 = 4.8\%$$

This formula ensures that components of return and returns annualised over a run of years maintain consistent relative results, which is not possible if simple differences are used to compare returns. Attribution scores are built up from comparisons of weights and returns in each segment of the market. Separate structure and selection scores in each segment are summed across the portfolio, to produce the portfolio level structure and selection scores which account for relative return.

The two and three component methods of attribution calculate structure scores in exactly the same way. In each segment:

Segment Structure Score

$$= (\text{Portfolio Weight} - \text{Benchmark Weight}) \times \text{Benchmark Return}$$

The alternative ways of calculating stock selection scores are:

Two component attribution method segment selection score

$$= \text{Portfolio Weight} \times ((1 + \text{Portfolio Segment Return}) / (1 + \text{Benchmark Segment Return}) - 1)$$

Three component attribution calculates the segment selection score as:

$$= \text{Benchmark Weight} \times ((1 + \text{Portfolio Segment Return}) / (1 + \text{Benchmark Segment Return}) - 1)$$

The difference lies in a single term. The three component multiplies segment relative returns by the benchmark weight, while the IPD method multiplies by the portfolio weight. When calculated on the IPD methods, the structure score and IPD selection score in each segment add up to the weighted contribution to relative return. Summed across segments, the structure score and IPD selection score add up to the portfolio's relative return.

In the Three Term Method, the structure and selection scores do not add up in this way, leaving a 'residual' term, the cross product, which is calculated as:

$$\text{Cross Product} = \text{Relative Return} - ((1 + \text{Structure Score}) \times (1 + \text{Selection Score}) - 1) \times 100$$