

What has the fund manager done for me?
A value-based approach to investment performance measurement

Gordon Bagot and Seth Armitage

January 2003

Gordon Bagot is an actuary specialising in investment consultancy.

Seth Armitage is a Reader in Finance, Heriot-Watt University, Edinburgh, UK.

Address for correspondence:

Dr Seth Armitage
School of Management
Heriot-Watt University
Edinburgh EH14 4AS
UK

Tel: 0131 451 3297

Email: S.E.Armitage@hw.ac.uk

What has the fund manager done for me?

A value-based approach to investment performance measurement

Abstract

In many circumstances it is unsatisfactory to measure portfolio performance using time-weighted rates of return. There are well known problems with performance attribution analysis using time-weighted returns in a multi-interval context, and it is impossible to measure the contribution of a portfolio's manager to the wealth of an individual investor. This paper shows that measurement of outcomes in terms of value solves the problems of multi-interval attribution analysis and enables precise customised assessment to be made of the manager's contribution for each investor. Working with values is also simple and transparent.

Introduction

The use of a benchmark portfolio is ubiquitous in portfolio performance appraisal. The benchmark enables the rate of return achieved by a fund over a particular period to be broken down into several components, consisting of the rate of return on the benchmark and of incremental rates which capture the effects of active management. This form of analysis was first presented by Fama (1972) in the context of the capital asset pricing model. He uses a portfolio's beta as an explicit measurement of its risk; the benchmark portfolio is defined as a passive portfolio with a beta chosen by the investor. The managed fund's rate of return is then explained as the return on the benchmark, plus the difference for taking more or less risk, plus the difference due to stock selection. Brinson, Hood and Beebower (1986) re-cast the analysis to reflect industry practice. No explicit risk measure is employed; instead, the benchmark is defined as a passive portfolio with an asset allocation chosen by the investor, and the rate of return on the managed fund is explained as the return on the benchmark, plus the difference due to deviation of the fund's asset allocation from that of the benchmark, plus the difference due to stock selection. This attribution framework has since been extended to identify the incremental rates of return from risk positioning within an asset class (Brinson, Singer and Beebower, 1991) and from currency selection (Allen, 1991; Ankrum and Hensel, 1994; Singer and Karnosky, 1995).

In all the above work, and in current industry practice, it is taken for granted that outcomes are measured by rates of return. If performance is assessed over more than one measurement interval, a rate of return is a rate compounded over the intervals, or a geometric mean rate per interval. Such a 'time-weighted' rate of return (TWR) for a fund is unaffected by the amounts of money an investor injects or withdraws at various dates. The main argument for using TWRs rests on the assumptions that the purpose of the analysis is to assess the manager's skill, and that the manager does not influence the investor's decisions regarding cash inflows and outflows (for example, Sharpe, Alexander and Bailey, 1999, pp. 827-8). Given these assumptions, it is appropriate to assign the fund's rate of return for each interval an equal weight in calculating a multi-interval TWR. However, investors are not only interested in the verdict on the manager's skill. A natural question for an investor to ask is 'What has the fund manager done for me, given my initial investment and the cash inflows and outflows by me along the way?' The investor wants to know the manager's contribution to the value of, or rate of return on, his or her investment, which is not the same thing as a

judgement regarding the manager's skill. Attribution analysis provides a better understanding of the manager's contribution.

If the assessment period consists of one interval, attribution analysis can equally well be conducted in terms of rates of return or values, because they are arithmetically equivalent, and the results for a pooled portfolio are applicable to each investor. But in a multi-interval context with intervening cash flows between the start and finish, we can not measure the manager's contribution to the value of the portfolio using TWRs, and attribution analysis can not be done correctly. Furthermore, the results of the analysis for the whole portfolio are no longer applicable to different investors in the portfolio. Even if there are no intervening cash flows, attribution using TWRs will be inexact unless the constituent asset classes held by the portfolio and its benchmark are identical in every interval. These problems are recognised by performance measurers but, to our knowledge, they have yet to be resolved satisfactorily.

In this paper, we present a value-based method of analysis which gives correct answers to the question of the manager's contribution to an investor's holding in a multi-interval context, and which enables precise attribution analysis to be conducted. The value-based method measures the manager's contribution by the difference between the final market value of the investor's holding in the fund, with its associated cash flows over time, and the final value of the equivalent holding in the fund's benchmark. In comparison, the industry standard method measures the manager's contribution by the difference between the TWR on the fund and the TWR on its benchmark. The value-based method is a form of money weighting, in that the impact of the manager's interval-specific decisions on the final value of the holding is positively related to the preceding cash flows made by the investor. But the method we present does not use money-weighted rates of return.

Application is straightforward, as we shall demonstrate. The data requirements are no more onerous than they are if TWRs are used. Attribution analysis in terms of contributions to market value is clear and exact, if unfamiliar. It can accommodate multi-currency portfolios and cash flows associated with expenses and taxes. Attribution can be taken down to the level of individual assets, if desired. The contribution can be calculated from not holding assets which are in the benchmark, and from holding assets which are not in the benchmark.

In summary, performance measurement in terms of values is advantageous if there is more than one interval. It enables attribution analysis to be conducted precisely and transparently, and it enables each investor to be provided with a customised report of the

manager's contribution to his or her holding. This information can not be provided satisfactorily using TWRs to measure performance.

The value-based method

Let the start of an assessment period be date $t = 0$, and the end be $t = T$. The assessment period consists of more than one measurement interval, so $T > 1$. At the end of each interval, the market value of the managed portfolio and of its benchmark is measured, and investors can inject or withdraw cash. At date 0, the value of investor i 's holding in managed portfolio P is $\$V_{P0,i}$. This could be the value of an existing holding in P , or an initial cash investment, or a combination of the two. Portfolio P is defined at any date t by the proportions held in the constituent assets $a = 1, 2 \dots n$. The proportion held in, or the weight of, asset a is $w_{Pat} = l_{Pat}\$p_{at}/\P_t , where l_{Pat} is the number of units of a held at date t , $\$p_{at}$ is the price per unit of a and $\$P_t$ is the market value of the portfolio. Thus, $P_t = w_{P1t}, w_{P2t}, \dots w_{Pnt}$, and $\sum_{a=1}^n w_{Pat} = 1$. If performance is being analysed at the level of asset allocation, an asset class A can be substituted for individual assets of the same class. The composition of the portfolio will change over time due to changes in asset prices, and due to changes in the units held which result from active investment decisions, from cash inflows and outflows between the portfolio and investors, and from investment of income received from assets held.

To establish the outcome from a passive policy of investing in a benchmark, we proceed as follows. A notional investment $\$V_{B0,i}$, is made in P 's benchmark portfolio, B , of the same amount as was invested in P : $\$V_{P0,i} = \$V_{B0,i}$. The benchmark is defined as $B_t = w_{B1t}, w_{B2t}, \dots w_{Bnt}$, and $\sum_{a=1}^n w_{Bat} = 1$. At any given date, including the start date, an actively managed portfolio P may hold assets in proportions which differ from those in B , and the constituents may differ; P may have zero holdings of some assets in B , and positive holdings of other assets not in B . If an investor injects, or withdraws, cash in portfolio P at some date t , assets of the same value as the cash are bought, or sold, in the benchmark in line with their weights in the benchmark at date t . The composition of the benchmark will change over time due to changes in asset prices and to re-investment of income,¹ but it is not affected by active management decisions, nor by the timing of cash inflows and outflows by investors. The value of an investor's benchmark is specific to the investor, but the benchmark's composition is common to all.

An alternative conception of the benchmark is that the weights are *fixed* over time. This implies that the benchmark is rebalanced at each date to ensure that weights w_{Bat} remain

constant, and that cash inflows or outflows are made by buying or selling assets in proportion to their constant weights. The question of whether to assume that the benchmark is rebalanced arises in the multi-interval context whether one uses TWRs or our proposed value-based method. An assumption of rebalancing raises the problem of the associated extra transactions costs. We assume in what follows that the asset weights in the benchmark are allowed to drift from their start-date weights during an assessment period. This is not restrictive; a value-based analysis similar to that presented could be done under an assumption of rebalancing each date.

Measurement with no intervening cash flows

The manager's contribution to the wealth of investor i is measured by comparing the end-date market value of the investor's holding, $\$V_{PT,i}$ with the end-date value of the investor's notional holding in i 's benchmark, $\$V_{BT,i}$. The difference can be expressed, if desired, as a percentage of the benchmark value: $(\$V_{PT,i} - \$V_{BT,i})/\$V_{BT,i}$. This tells the investor how much value the manager has gained or lost during the assessment period, per dollar of the investor's notional end-date wealth, had he or she invested in the benchmark. Suppose first that investor i makes no injections or withdrawals during the assessment period. In this case, the percentage gain or loss for i 's holding in relation to the value of i 's notional holding in the benchmark is the same as the percentage gain or loss for the whole portfolio P in relation to its benchmark B , $(\$V_{PT} - \$V_{BT})/\$V_{BT}$. Since $\$V_T = \$V_0(1 + R)$, where $1 + R = \prod_{t=1}^T(1 + R_t)$, the same figure for i could obviously be calculated from the TWRs on P and B , given the starting values $\$V_{P0} = \V_{B0} . However, attribution analysis can not be carried out exactly using rates of return if there are intervals in which the portfolio has zero holdings of some assets in the benchmark, or in which the portfolio has positive holdings of assets not in the benchmark. A zero weighting for an asset in one or more intervals means that either portfolio or benchmark, or perhaps both, will lack a compounded rate of return for the asset in question for the assessment period. In addition, when a fund goes short in an asset, or when it uses derivatives, it may be impossible to calculate a sensible compounded rate of return. These contingencies are not a problem for the value-based method, as will be shown in Example 1 for the case of a zero weighting in an asset.

Measurement with intervening cash flows

The value-based method really comes into its own if there are cash injections or withdrawals by investors during the assessment period. In this case, exact attribution analysis for the whole portfolio, or for any of the investors' holdings, is never possible using rates of return, and the relative performance figure for the portfolio based on values, $(\$V_{PT} - \$V_{BT})/\$V_{BT}$, is not the same conceptually or numerically as the figure based on TWRs, $(R_P - R_B)/(1 + R_B)$. A further consequence of intervening cash flows is that the TWRs on each investor's holding will, in general, differ from each other, and from the return on the whole portfolio. This is ignored in standard attribution analysis, which relates to the whole portfolio, not to the holdings of specific investors. But each investor's holding, together with the whole portfolio, should in principle be given its own attribution analysis. This can be done using values.

The procedure under the value-based method is that, each date on which an investor injects or withdraws cash, there is a matching notional cash flow for the investor's benchmark and for the whole portfolio's benchmark. It is only necessary to know the composition of the benchmark when there is a cash flow, which need not be every interval. Because the benchmark is precisely adjusted for the same cash flows as the portfolio, the effect on the final portfolio value of the timing of the cash flows is removed. This is appropriate because such cash flows are the responsibility of the investor, not the manager. Having outlined value-based performance measurement in general terms, we now present an example to illustrate measurement and attribution analysis using values.

Example 1A: no intervening cash; one investor

Table 1 shows the composition and performance of a managed portfolio and of its benchmark over an assessment period of two intervals. We assume temporarily that there is no new money invested at the end of interval 1 (= date 1). Performance is considered here at the level of asset allocation, and there are three asset classes available; equity, bonds and cash. The actual portfolio has an initial value of \$1,000 invested 55% in equity with 45% in bonds. The benchmark is different; the same \$1,000 is invested 50% in equity, 40% in bonds and 10% in cash. The returns for the portfolio are 4.48% in interval 1 and -2.80% in interval 2, giving a two-interval TWR for the assessment period of 1.55%. The TWR for the benchmark is 2.12%, so the portfolio has underperformed by -0.56% [= $(0.0155 - 0.0212)/1.0212$]. In terms of values, the portfolio has a final value of \$1,015.46, compared

with \$1,021.20 for the benchmark. The shortfall of $-\$5.74$, expressed as a percentage of the benchmark value is, of course, -0.56% , exactly the same as the underperformance derived from the two TWRs.

Table 1 around here

Example 1B: intervening cash; one investor

We now assume that a further \$200 is invested at date 1. The portfolio manager decides to hold this \$200 in cash for the second interval, but the notional \$200 invested in the benchmark is allocated amongst the asset classes in line with the weighting of each class in the benchmark at date 1. Table 2 shows the benchmark for the new money, and the performance of the new money compared with its benchmark. The composition of the benchmark in interval 2 is the same for the new money as for the existing investment, and has the same return over interval 2 as above in Table 1. The portfolio's return on the new money is quite different, since all \$200 was retained in cash and earned 1.25% rather than the -1.95% which the benchmark earned. The 'portfolio' consisting of this new cash alone outperformed the benchmark by $+3.26\%$ in interval 2, whether calculated from the rates of return or from the portfolio's value added of $\$6.40$ divided by its benchmark value of $\$196.10$.

Tables 2 and 3 around here

The next step is to add together the values relating to the initial investment and to the new cash. This sum is shown in Table 3. The new cash changes the value of the benchmark but not its composition, so the benchmark including the new cash has the same rate of return in interval 2 as in Tables 1 and 2, and the same two-interval return of 2.12%. But the return for the portfolio including the new cash is now -2.15% in interval 2, rather than -2.80% . This is the result of combining the interval 2 return of -2.80% on the initial investment of \$1,000 with the interval 2 return of 1.25% on the new investment of \$200 at date 1. The two-interval TWR for the total portfolio is now 2.23%, rather than 1.55%, and the relative performance is $(0.0223 - 0.0212)/0.0212 = +0.10\%$. In values, the portfolio is worth \$1,217.96 compared with \$1,217.30 for the benchmark, a difference of $+\$0.66$. This difference expressed as a percentage of the benchmark value gives the relative performance of $+0.05\%$. Now that we

have an intervening cash flow, the performance measure in relation to the benchmark derived from TWRs does not match the measure derived from values. This means that attribution analysis which seeks to account for the generation of excess value of +\$0.66 can not be done exactly using TWRs. The reason for the mismatch is that TWRs give equal weighting to the rates of return in each interval, whereas in reality the rate in interval 2 has more impact on the final value, because there was extra cash invested during interval 2.

It is true that, in this example, both value-based and standard methods result in a similar inference, which is that the managed portfolio did slightly better than its benchmark during the assessment period. The discrepancy between the relative performance measures would widen were the intervening cash flows larger. For example, if \$10,000 were invested at date 1 instead of \$200, relative performance would be +3.19% using TWRs and +2.90% using values. More dramatically, the inferences from the methods can diverge, because they can differ in sign. That is, the relative performance from TWRs can be positive, yet value has been lost in relation to the benchmark, and vice versa.

Example 1C: intervening cash; two investors

To see how the value-based method enables investor-specific attribution to be carried out, and to see how a difference in sign can arise compared with the TWR measure, suppose now that there were two investors in the portfolio, as shown in Table 4. Investor X placed \$900 with the manager at date 0, with no further payments, while Investor Y placed \$100 at date 0 and was responsible for the entire cash inflow of \$200 at date 1. The composition of the whole portfolio, and the rates of return on each asset class, are exactly the same as before.

Table 4 around here

The situation in interval 1 is that X has 90% and Y has 10% of the managed portfolio, so both earn the rate of return of the portfolio in interval 1, 4.48%, and the benchmark return for both is the benchmark return for the portfolio, 4.15%. The situation in interval 2 needs a little more explanation. After Y's investment at date 1, the composition of *both* portfolios changes. After interval 1 but before Y's investment, the managed portfolio consists of 56% equity, 44% bonds and 0% cash; X owns 90% and Y owns 10%. After Y's investment, the portfolio consists of 47% equity, 37% bond and 16% cash; X owns 76% and Y owns 24%. Thus, even though X does nothing, the composition of his portfolio changes, as does Y's,

because of the manager's decision to keep the new \$200 in cash.² For example, X now has \$151.08 in cash. The composition of the benchmark portfolio does not change as a result of Y's investment, because the new cash is allocated in line with the existing proportions, which are 51% equity, 39% bonds and 10% cash at date 1.

X's initial \$900 in the managed portfolio grew to \$920.04 by date 2. If the \$900 had been invested in the benchmark, then the final value would have been \$919.08. Thus the manager added value for X in relation to the benchmark of \$0.96, giving a relative performance of +0.10%. The outcome for Y is different. The final value for Y is \$297.92, compared with \$298.22 had Y's money been invested in the benchmark, so the manager lost \$0.30, giving a relative performance of -0.10%. Also, Y suffered an absolute capital loss; the \$100 invested at date 0 plus the \$200 at date 1 is only worth \$297.92 at date 2.

The results using TWRs are those for the whole portfolio in Table 3, which means X and Y both have a two-interval TWR of +2.23%, and both have a relative performance of +0.10%. X's portfolio had no intervening cash flow, so the absolute value did indeed increase by 2.23%, and the value-based and standard methods produce the same relative performance figure of +0.10%. However, the results from the standard method for Y's portfolio are misleading. The relative performance for Y from TWRs is +0.10%, but in fact Y lost value in relation to the benchmark. The absolute TWR for Y is 2.23%, but in fact Y suffered an absolute capital loss. The explanation is that the managed portfolio did badly in interval 2 both absolutely and relatively, and TWRs do not capture the fact that Y invested more money in it at the start of interval 2.

Attribution analysis in Example 1C

The aim of attribution analysis using values is to explain the difference between the value of an investor's portfolio and that of its benchmark at the end of the assessment period. This difference is the manager's contribution; it is additional to the impact on the benchmark of the timing of the cash flows. For example, had Y's end-date benchmark value been calculated on the assumption that the \$200 was invested at date 0, rather than date 1, then the benchmark value would have been \$306.36 rather than \$298.92. The difference of -\$7.44 in the benchmark represents a loss attributable to the later actual payment of \$200 at date 1. Since the portfolio's performance is judged against the lower benchmark value of \$298.92, the loss is accountable to the investor rather than the manager.

Tables 5, 6, 7 and 8 around here

The manager's contribution to value can be analysed in terms of asset allocation policy, stock selection and a cross-product term, as usual. Table 5 shows the analysis for the total portfolio, for each interval and for the full assessment period. The analyses for each investor are given in Tables 6 and 7, and the underlying formulas are in Table 8. We present Tables 5, 6 and 7 so that readers can see how attribution using values works out numerically, and can satisfy themselves that it delivers the precision and transparency which we claim for it.

The analysis proceeds in three stages. The first stage is to carry out the attribution for interval 1 in same way as in Brinson et al (1986, 1991), but expressed in terms of values rather than rates of return. The second is to carry out the attribution for both intervals combined, which explains how the gain or loss in relation to the benchmark has arisen during the assessment period as a whole. The third is to deduct the attribution values for interval 1 from the values for the assessment period, to obtain the values for interval 2.

This procedure is necessary because the relative value added or lost which we are seeking to explain is the relative value which has arisen during the whole assessment period. The attribution values for the whole period can not be arrived at by adding the values for each interval analysed separately. A correct attribution analysis for interval 2 would involve re-setting the benchmark at the start of interval 2 so that it had the same value as the managed portfolio. The final values of the benchmark and of its constituent assets would also be re-set, so they would no longer match the true final benchmark values. This would be an alternative solution to the problem of multi-interval attribution. The method we present shows a correct attribution analysis for interval 1 and for the assessment period, but not for interval 2; the values for interval 2 are calculated as residuals. The alternative method would provide correct attribution values for both intervals, but the sum of these would not provide a correct analysis for the whole period.³

If the assessment period consists of one interval, attribution analysis in terms of rates of return is consistent with analysis in terms of values. But if the assessment period consists of more than one interval, attribution analysis using rates of return can be problematic. In our example, the portfolio has no cash in interval 1 but \$200 cash in interval 2, so a rate of return on cash in the portfolio can not be calculated for the assessment period. This means that performance attribution can not be done for the assessment period using TWRs, at least not

without fudging the issue somehow. In contrast, Table 5 shows the contribution from the absence of cash in the portfolio at the start, in relation to the benchmark which has 10% in cash.

The attribution calculation in Table 5 is for the total portfolio. Unless investors inject or withdraw amounts of cash which are the same as a proportion of their initial investments, the attribution will be different for each of them and for the total portfolio. The value-based method allows for customised reports to be produced, one for each investor, as Tables 6 and 7 show for X and Y respectively. The values in these tables sum to give the values for the total portfolio in Table 5. The analysis explains the manager's contribution of \$0.96 for X and –\$0.30 for Y, relative to the benchmark. It is precise, accurate to the cent; there are no 'buckets' of unexplained contributions to performance for any investor. Such investor-specific attribution can not be done using TWRs.

Exactly the same value-based analysis can be done at the level of individual assets, instead of asset classes, although naturally the data requirements are greater. It is also straightforward to accommodate expenses and tax payments. These payments are cash outflows from a portfolio which can be identified and dated. However, we do not pursue these aspects of appraisal here, but turn instead to the problem of a multi-currency portfolio.

Currency analysis

The currency dimension adds complexity to the task of performance measurement, and has yet to be accommodated satisfactorily if the assessment period has more than one interval. Singer and Karnosky (1995, p. 89), for example, note that 'multiperiod attributions require an accounting for changes in active weights and returns over time. While it may be convenient to add monthly attribution results or to use longer-term average weights and returns, the results can be misleading at best and often wrong. Multiperiod attributions are best accomplished by creating weighted return indexes to compute each attribution component.' But use of weighted return indexes will not result in a clear and exact analysis. We now develop a second example which shows that the value-based method can accommodate changes in active weights and returns over time clearly and exactly.

Example 2: two-currency portfolio with cross-currency switch

Table 9 shows the composition and performance of an international portfolio and of its benchmark over an assessment period of two intervals. There are only two asset classes; holdings in the USA and in the UK. The portfolio starts with \$1,000 and, for simplicity, there are no intervening cash flows in this example. The portfolio does well in relation to its benchmark in the first interval. It is overweight in the US, which is advantageous because the manager's stock selection is better in the US than in the UK, and the dollar appreciates against sterling. The manager increases the weighting further towards the US at the end of the first interval, but this turns out to be a bad move. The manager's stock selection in the US deteriorates, and the dollar depreciates against sterling in the second interval. As in Example 1, the performance in relation to the benchmark is the difference between the final value of the portfolio and of its benchmark, which in this case is $-\$2.07$, or -0.18% of the final value of the notional holding in the benchmark. Since there is no intervening cash flow, the relative performance is exactly the same using TWRs.

Tables 9, 10 and 11 around here

An attribution analysis is shown in Table 10 for the full assessment period, with the formulas given in Table 11. The numbers for the two intervals are not shown. Value added is broken down into three primary components and four cross-product terms, the rationale for this being as follows.

Assume for simplicity that there is only one interval. An asset class is now a holding in a foreign market M , the portfolio return on M is R_{PM} , the benchmark return on M is R_{BM} , and the return on the whole benchmark is R_B . Let E be the proportionate change in the exchange rate during the interval. We wish to account for $V_{PM1} - V_{BM1}$, the difference between the final holdings in M of portfolio P and in its benchmark B .

$$\begin{aligned} V_{PM1} - V_{BM1} + (V_{PM0} - V_{BM0}) &= [V_{PM0}(1 + R_{PM}) - V_{BM0}(1 + R_{BM})](1 + E) \\ V_{PM1} - V_{BM1} &= [V_{PM0}(R_{PM}) - V_{BM0}(R_{BM})](1 + E) + (V_{PM0} - V_{BM0})E \\ &= [(V_{PM0} - V_{BM0})(R_{BM1} - R_{B1}) + V_{BM0}(R_{PM1} - R_{BM1}) \\ &\quad + (V_{PM0} - V_{BM0})(R_{PM1} - R_{BM1})](1 + E) \\ &\quad + (V_{PM0} - V_{BM0})E \\ &= [\text{market selection (MS) + stock selection (SS)} \\ &\quad + \text{market selection} \times \text{stock selection (MS.SS)}](1 + E) \\ &\quad + \text{currency selection [CS} = (V_{PM0} - V_{BM0})E] \end{aligned}$$

$$= MS + SS + CS + MS.SS + MS.E \\ + SS.E + MS.SS.E,$$

as in Tables 10, 11 and 12. Hence, the three primary components are market selection and stock selection, both ignoring currency, and currency selection, which captures the exchange rate gain or loss from market selection. The first cross term, MS.SS, is the value attributable to market selection times stock selection, ignoring currency, as in the single-currency analysis. The other three cross terms capture the interaction between the change in the exchange rate and the local-currency value gained or lost through, respectively, market selection, stock selection and the cross product of these. If M is the market of the portfolio's master currency, $E = 0$, and the analysis reduces to that for the single-currency case.⁴

The clear and unequivocal calculation of the cross-terms is a particular benefit of the value approach. There is no reason to factor these numbers 'back into' the three principal numbers of markets, currency and stock selection. The values of the cross terms are added up over time, and kept separate from the principal components. Managers usually decide on markets, currency and stocks and are not thinking about cross terms. Correct attribution of the principal components of performance should not be muddled by the second-order cross terms which, if factored back in, can seriously change the principal components, even turning a positive contribution negative or vice versa.

A problem arises for TWRs when there is a negative market value, for example from going short, using derivatives or from selling assets. In our example there is a switch from the UK to the US at date 1. Using the value-based method, it is possible to quantify and analyse the gain, or in this case the loss, arising from the switch, something which is not possible using TWRs. Table 12 shows the attribution analysis for the switch in value terms. \$5.90 was lost, mostly due to the rise of sterling against the dollar.

Table 12 around here

A further problem in currency attribution is created by 'currency tainting'. For example, suppose that the benchmark of a Far Eastern equity portfolio, reporting in US\$, is the relevant component of the FTSE All-World Index. Some of the Far Eastern portfolio is held in assets denominated in HK\$, which is tied to the US\$. The currency-related performance of the benchmark is determined primarily by the Yen/\$ exchange rate. If the Yen falls against the US\$, the HK\$ holding will show a positive currency contribution, when in

fact there is no currency benefit from being in the HK\$. This problem is avoided using values, since the value of the portfolio's holding in a currency is assessed against the value of the benchmark's holding, both expressed in the master currency. Thus, there would be currency contribution of zero for the HK\$ holding, which is correct.

Conclusion

The value-based method presented here measures the contribution of a fund manager to an investor's holding in a managed portfolio, in terms of dollars gained or lost in relation to a benchmark. The method is equivalent to the standard method using TWRs if there is one interval in the assessment period. If there is more than one interval, the value-based method has two advantages: it solves the problems which can beset attribution analysis when TWRs are used, and it enables customised performance reports to be provided for each investor, as well as for the whole portfolio. The value added or lost by the manager can be expressed as a percentage of the fund's final benchmark value, which will often be a more convenient way of expressing the manager's contribution than is the value added itself. The drawback with our method, other than its unfamiliarity, is that it does not provide a rate of return which can be compared with the rates achieved by other funds. The results of a value-based analysis, whether expressed as values or as percentages, may not be appropriate for use in assessment of the manager's skill.⁵ Rather, value-based analysis enables the manager's contribution to a fund, and to the holdings of the individual investors, to be analysed correctly.

Notes

1. A fuller account would include rules regarding treatment of cash receipts from individual assets, of cash flows arising from use of derivatives, and of expenses and taxes.
2. We are assuming that X and Y do not have segregated portfolios, so their portfolios always have the same composition as that of the managed portfolio.
3. The underlying problem here is that $R_P - R_B \neq \sum_{t=1}^T (R_{Pt} - R_{Bt})$, where $R_Q = \prod_{t=1}^T (1 + R_{Qt}) - 1$, for $Q = P, B$. In view of this, the assessment period should be defined in advance, because a manager's behaviour may differ if he knows he will be assessed over one interval at a time rather than over several intervals combined. Recent papers by Carino (1999) and Menchero (2000) present adjustments to single-interval returns which reduce the difference between a compounded return and the sum of its constituent single-interval returns. The same adjustments could be made to the single-interval returns expressed in values. However, Carino and Menchero do not consider the problems caused either by intervening cash flows or by mismatch between the constituents of a portfolio and its benchmark.
4. Ankrim and Hensel (1994) and Singer and Karnosky (1995) present analyses which isolate the effect of the interest rate parity relation between exchange and interest rates. Ankrim and Hensel do this by splitting the currency selection component into a 'forward premium effect' and a term for the difference between the actual change in the exchange rate and the relevant forward premium. Singer and Karnosky deduct the local-currency eurodeposit rate from the market rate of return in the 'market selection' and 'stock selection' components. 'Currency selection' is then defined as the sum of the local-currency euro-rate and the change in the exchange rate. Both papers also allow for currency hedging. We have not isolated the effect of the interest rate parity relation, nor included a hedged position, in order to keep the example as simple as possible. There is no problem in conducting either an Ankrim-Hansel or a Singer-Karnosky analysis in value terms.
5. The value added relative to the benchmark is a measure of the manager's skill. The issue is that the effect of manager skill per interval is weighted by the value of the fund at the start of the interval. This need not always rule out cross-fund comparison. Manager skill can be compared across funds using value added expressed as a percentage of the final benchmark, if the funds concerned have an initial value and intervening cash flows which are the same or which differ by a scaling factor. For example, one may wish to compare a group of mutual funds on the assumption that each fund starts with \$Y and receives a constant \$X per month. Such a comparison would be fair; the effect of manager skill in each interval

would receive the same weighting across the funds. If the comparison is fair, there is a case for value-weighting the effect of skill. The measure of skill would be the same as the correct measure of the manager's contribution to the fund's final value. There are also arguments that it is easier to add value to a small fund than to a large one. The use of value-based analysis in assessing skill merits further study.

References

- Allen, GC, 1991, 'Performance attribution for global equity portfolios', *Journal of Portfolio Management*, Fall, 59-65.
- Ankrim, EM and CR Hensel, 1994, 'Multicurrency Performance Attribution', *Financial Analysts Journal*, March-April, 29-35.
- Brinson, GP, LR Hood and GL Beebower, 1986, 'Determinants of portfolio performance', *Financial Analysts Journal*, July-August, 39-44.
- Brinson, GP, BD Singer and GL Beebower, 1991, 'Determinants of portfolio performance II: an update', *Financial Analysts Journal*, May-June, 40-48.
- Carino, D, 1999, 'Combining attribution effects over time', *Journal of Portfolio Measurement*, Summer, 5-14.
- Fama, EF, 1972, 'Components of investment performance', *Journal of Finance* 27, 551-67.
- Menchero, JG, 2000, 'An optimized approach to linking attribution effects over time', *Journal of Portfolio Measurement*, Fall, 36-42.
- Sharpe, WF, GJ Alexander and JV Bailey, 1999, *Investments*, 6th ed, Prentice Hall.
- Singer, BD and DS Karnosky, 1995, 'The general framework for global investment management and performance attribution', *Journal of Portfolio Management*, Winter, 84-92.

Table 1. Example 1A: A managed portfolio and its benchmark over two intervals

	<i>Date 0</i> Compos- ition	<i>Interval 1</i> Rate of return	<i>Date 1</i> Compos- ition	<i>Date 1</i> Allocation of new cash	<i>Interval 2</i> Rate of return	<i>Date 2</i> Compos- ition	Two- interval TWR ¹	Value added	Difference in value added
Managed portfolio									
Equity	\$550.00	6.50%	\$585.75	\$0.00	-5.00%	\$556.46			
Bonds	\$450.00	2.00%	\$459.00	\$0.00	0.00%	\$459.00			
Cash	\$0.00	n/a	\$0.00	\$0.00	1.25%	\$0.00			
Total	\$1,000.00	4.48%	\$1,044.75	\$0.00	-2.80%	\$1,015.46	1.55%	\$15.46	
Benchmark portfolio									
Equity	\$500.00	6.00%	\$530.00	\$0.00	-4.60%	\$505.62			
Bonds	\$400.00	2.50%	\$410.00	\$0.00	0.50%	\$412.05			
Cash	\$100.00	1.50%	\$101.50	\$0.00	2.00%	\$103.53			
Total	\$1,000.00	4.15%	\$1,041.50	\$0.00	-1.95%	\$1,021.20	2.12%	\$21.20	-\$5.74
Relative performance ²							-0.56%		-0.56%

Notes

1. The two-interval TWR, R_{1+2} , is $(1 + R_1)(1 + R_2) - 1$, where R_t is the rate of return over interval t for the managed portfolio P or its benchmark B .
2. Relative performance is given by $(R_{P1+2} - R_{B1+2}) / (1 + R_{B1+2}) - 1$ for TWRs, and by $(V_{P2} - V_{B2}) / V_{B2}$ for values, where V_2 is the value of portfolio P or benchmark B at date 2.

Table 2. Example 1B: New cash for interval 2

	<i>Date 1</i> Benchmark proportions ¹	<i>Date 1</i> Allocation of new cash	<i>Interval 2</i> Rate of Return	<i>Date 2</i> Compos- ition	Rate of return	Value added	Difference in value added
Managed portfolio							
Equity		\$0.00	-5.00%	\$0.00			
Bonds		\$0.00	0.00%	\$0.00			
Cash		\$200.00	1.25%	\$202.50			
Total		\$200.00	1.25%	\$202.50	1.25%	\$2.50	
Benchmark portfolio							
Equity	50.89%	\$101.78	-4.60%	\$97.09			
Bonds	39.37%	\$78.73	0.50%	\$79.13			
Cash	9.75%	\$19.49	2.00%	\$19.88			
Total	100.0%	\$200.00	-1.95%	\$196.10	-1.95%	-\$3.90	\$6.40
Relative performance					3.26%		3.26%

Note

1. The benchmark proportions are the value of each asset class in the benchmark at date 1 divided by the total value of the benchmark at date 1, before the new cash arrives. These numbers are from Table 1.

