THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Value-based performance measurement: further explanation

Citation for published version:
Armitage, S & Bagot, G 2009, 'Value-based performance measurement: further explanation' Journal of Performance Measurement, vol 13, no. 2, pp. 58-74.

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Author final version (often known as postprint)

Published In:
Journal of Performance Measurement

Publisher Rights Statement:
© Value-based performance measurement: further explanation. / Armitage, S.; Bagot, G.
In: Journal of Performance Measurement, Vol. 13, 2009, p. 58-74.

OPEN ACCESS

Value-based performance measurement: further explanation
Seth Armitage and Gordon Bagot

Seth Armitage is a professor of finance at the University of Edinburgh, UK. His teaching and research interests are mainly in corporate finance and capital markets. He became interested in performance measurement after hearing Gordon Bagot give a talk on the subject. He holds a PhD in finance from the University of Edinburgh. He was previously at Heriot-Watt University, and was a merchant banker in London before becoming an academic.

Gordon Bagot specializes in investment analysis and consultancy. He is a former director of The WM Company in Edinburgh where his responsibilities lay principally in performance measurement and analysis, plus investment research and consultancy. He is an actuary and formerly a vice president of The Faculty of Actuaries in Scotland and chairman of The Actuarial Profession's Finance, Investment & Risk Management Board.

Value-based performance measurement: further explanation

Abstract

A value-based method of performance measurement was presented in our 2004 paper in this *Journal.* Here we extend our earlier paper in several ways. First, we compare the value method with the internal rate of return. Both incorporate money weighting, but values are easier to use and understand. We therefore recommend performance measurement in values as the preferred money-weighted complement to measurement using time-weighted returns. Second, we highlight the fact that the contribution to terminal value of every external or internal cash flow can be measured separately. It is this feature that gives the value method its great transparency and flexibility when analysing performance over multiple time periods. Third, we discuss several aspects of implementation that are neglected in our earlier paper.

INTRODUCTION

A few years ago the authors presented a method of performance measurement that is based on values, rather than returns (Armitage & Bagot, 2004). The method shows how much money the manager of a given fund has gained or lost during an assessment interval, compared with a benchmark fund with exactly the same cash flows on the same dates as the flows in and out of the portfolio. That is, performance is measured by $V_{P,T} - V_{B,T}$, where $V_{P,T}$ is the value of the portfolio at a terminal date T, and $V_{B,T}$ is the value of the benchmark.

We argued that there are three main advantages of the value-based method. First, $V_{P,T} - V_{B,T}$ is a direct measure of the manager's contribution to the fund, in relation to the benchmark. It is always a true picture. Returns can give 'funny' answers, for example a superior rate of return for the fund compared with the benchmark, when the client would have been better off from investing in the benchmark. We believe that investors are interested in the manager's contribution to their wealth, measured in terms of value. Second, attribution analysis is easier and more transparent than it is using returns, when there is more than one time period and particularly when there are transactions and/or cash flows. An exact analysis of the difference $V_{P,T} - V_{B,T}$ can be carried out, without any prior adjustment or averaging of the cash flows, values or returns for each period. This is not possible when we are trying to explain a difference in single-period returns that have been added or compounded. Third, in the case of a pooled fund with several clients who have made differing cash flows, it is straightforward to calculate the value added or lost for each client separately, and to provide each client with a separate attribution analysis. The attributes measured in values will sum exactly across the clients to give the attributes for the pooled fund as a whole.

The purpose of the current paper is to develop further the case for the value method, and to explain how it works in more detail. After outlining the value method, we discuss the fundamental question of the role of performance measurement using values. The value method incorporates money weighting. The standard measure to date that incorporates money weighting is the internal rate of return (IRR). Measurement with money weighting, via IRR, has been strongly recommended as a complement or alternative to measurement with time-weighted returns (see, for example, Campisi, 2004; Spaulding, 2005). The GIPS standards require presentation of the IRR for private equity funds and recommend it for real estate assets (Maginn et al, 2007, pp. 832-40). Our argument is that values are easier to work with than IRRs, and easier to understand. In particular, exact attribution and sub-fund analyses are straightforward using values, but not using IRRs; and IRRs can be misleading. We therefore believe that there is a strong case for using values instead of IRRs in the many instances in which money-weighted performance measurement is called for.

The transparency and flexibility of the value method was demonstrated in our earlier paper by means of examples. Here we aim to show *why* performance analysis using values is so transparent and flexible, and why it always explains exactly the difference $V_{P,T} - V_{B,T}$. The reason is that the contribution to terminal value of each cash flow, whether made by the investor or by the manager, can be measured separately. This means that, in attribution analysis, the cross-product terms that cause the problems using returns do not arise using values. The separate contribution of each cash flow also explains the value method's ease of implementation in other respects. This is apparent when we discuss various potentially awkward aspects of implementation which our previous paper says little about. The aspects we discuss are negative holdings (short positions) in the portfolio or benchmark, maintenance of fixed weights for asset classes in the benchmark, voids in asset classes, income from securities held, and expenses and taxes. All can be accommodated without difficulty.

OVERVIEW AND THE CASE FOR VALUES

The value-based method

Performance over a pre-agreed assessment interval is measured in values by $V_{P,T} - V_{B,T}$. $V_{P,T}$ is whatever the terminal value of the portfolio turns out to be. The question is how $V_{B,T}$, the terminal value of the benchmark, is arrived at. The *modus operandi* of the value method is that, to calculate $V_{B,T}$, every external cash flow, ie a cash flow between a client and the fund, is matched by a cash flow of the same amount and at the same time for the benchmark. Cash invested in the fund or withdrawn from it is allocated or withdrawn across asset classes as the manager sees fit. Cash is allocated or withdrawn in the benchmark across asset classes in accordance with the benchmark weightings for each asset class at the time. If the fund manager makes internal cash flows across asset classes, and there is no external flow, then there are internal cash flows within the fund but not within the benchmark.

Suppose that there are two asset classes, a and b. Then the entries for an external and internal cash flow are:

	External cash flow, between investor and portfolio	Internal cash flow within portfolio, across asset classes
Portfolio	$Y_{Pa,t} + Y_{Pb,t} = Y_t$	$Y_{Pa,t} = -Y_{Pb,t}$ or $-Y_{Pa,t} = Y_{Pb,t}$
Benchmark	$w_{Ba,t}Y_t + w_{Bb,t}Y_t = Y_t$	no changes

where the subscripts P, B, a, and b denote the portfolio, its benchmark, and asset classes a and b, respectively; Y_t is an external cash flow; and $w_{Ba,t}$ is the weight for asset class a in the benchmark at date t. The benchmark for each asset class would normally be an index, so cash flows for the benchmark entail flows for the notional holdings in the benchmark's asset-class indices.

The benchmark weights can be allowed to drift from period to period, due to differences in the returns on the asset classes in the benchmark. Or the weights can be held constant, which means there will be internal cash flows in the benchmark each period to re-set the values of the holdings per asset class in accordance with the fixed weights. This is explained later in the paper.

Example 1. Cash flows and returns for a portfolio and its benchmark

---

Date	0	1	2	3
External cash flows	Y_0	$-Y_1$	Y_2	0
Portfolio returns		$R_{P,1}$	$R_{P,2}$	$R_{P,3}$
Benchmark returns		$R_{B,1}$	$R_{B,2}$	$R_{B,3}$

Each cash flow, whether external or internal, positive or negative, has an impact on $V_{P,T} - V_{B,T}$. To see how this works, consider the structure of cash flows and returns for a portfolio and its benchmark in Example 1. The terminal value arising in the portfolio from an external cash flow at date t, $V_{P,T}(Y_t)$, is given by

$$V_{P,T}(Y_t) = Y_t \prod_{\tau=t+1}^{T} (1 + R_{P,\tau})$$

where $R_{P,\tau}$ is the return on the portfolio in period τ. So the terminal value of the portfolio, $V_{P,3}$, is given by

$$V_{P,3} = Y_0(1 + R_{P,1})(1 + R_{P,2})(1 + R_{P,3}) - Y_1(1 + R_{P,2})(1 + R_{P,3}) + Y_2(1 + R_{P,3})$$

Notice that the cash outflow at date 1 reduces $V_{P,3}$ by the amount of the outflow multiplied by the returns that would have be earned on the money were it to have remained in the portfolio. This might seem odd at first, but it makes sense: $-Y_1(1 + R_{P,2})(1 + R_{P,3})$ is simply the impact, in value terms, that the withdrawal of the amount Y_1 has on the terminal value, given the subsequent returns for the portfolio. The terminal value of the benchmark, $V_{B,3}$, is given by

$$V_{B,3} = Y_0(1 + R_{B,1})(1 + R_{B,2})(1 + R_{B,3}) - Y_1(1 + R_{B,2})(1 + R_{B,3}) + Y_2(1 + R_{B,3})$$

Therefore, the terminal value added or lost in relation to the benchmark is

$$V_{P,3} - V_{B,3} = + Y_0[(1 + R_{P,1})(1 + R_{P,2})(1 + R_{P,3}) - (1 + R_{B,1})(1 + R_{B,2})(1 + R_{B,3})]$$
$$- Y_1[(1 + R_{P,2})(1 + R_{P,3}) - (1 + R_{B,2})(1 + R_{B,3})]$$
$$+ Y_2[(1 + R_{P,3}) - (1 + R_{B,3})] \quad (1)$$

This shows how the gain or loss arising from each external cash flow sums to give the total gain or loss for the whole assessment interval. All the returns used in the analysis are single-period returns. The value method does incorporate money weighting, but not by using IRRs.

Example 2. External and internal cash flows

Date	0	1	2
External cash flows	Y_0	0	$-Y_2$
Allocation by asset class			
Portfolio	$Y_{Pa,0}$	$-Y_{Pa,1}$	$\Delta Y_{Pa,2}$
	$Y_{Pb,0}$	$+Y_{Pb,1}$	$\Delta Y_{Pb,2}$
Benchmark	$w_{Ba,0}Y_0$	0	$-w_{Ba,2}Y_2$
	$w_{Bb,0}Y_0$	0	$-w_{Bb,2}Y_2$

Example 2 illustrates the treatment of internal cash flows, between asset classes or sub-funds. At date 1 the fund manager switches an amount $Y_{Pa,1}$ out of asset class a and into b. There is no external cash flow at that date, and no cash flows for the asset classes in the benchmark (assuming drifting weights). At date 2 the investor withdraws Y_2. The manager must withdraw Y_2 from the asset classes in total, so $\Delta Y_{Pa,2} + \Delta Y_{Pb,2} = -Y_2$. But he might also carry out a further re-allocation. So the cash flow for one of the asset classes need not be negative; hence the Δ prefix rather than a minus sign. Cash is withdrawn from the benchmark in line with the benchmark's asset-class weightings. The impact of each internal flow on the terminal value of the portfolio or benchmark is calculated in exactly the same way as the impact of an external flow. For example, the impact of the manager's switch at date 1 on $V_{P,2}$ is $-Y_{a,1}(1 + R_{Pa,2}) + Y_{b,1}(1 + R_{Pb,2})$.

To use the value method accurately, values need to be measured when an external or internal cash flow occurs. The periods between valuations need not be of equal length. This is also a requirement for accurate calculation of time-weighted returns. The GIPS standards are moving in the direction of more frequent valuations. From 2010, GIPS will require portfolio valuations on the dates of all 'large' external cash flows.

What do we want to measure? TWRs, IRRs and the value method

The standard measure in assessing fund performance is the time-weighted rate of return (TWR). For T periods, the TWR for a portfolio between dates 0 and T, $R_P(0,T)$, is the compound return over the T periods from date 0 to date T.

$$R_P(0,T) \; = \; \prod_{t=1}^{T}(1 + R_{P,t}) - 1$$

The TWR will not explain the difference between a starting value at date 0, $V_{P,0}$, and the terminal value, $V_{P,T}$, if there are external cash flows between dates 0 and T. This is because external cash flows will affect $V_{P,T}$ but not the single-period returns and therefore not the TWR. They will also affect the difference $V_{P,T} - V_{B,T}$, although the impact of market timing or mis-timing by the investor will normally be much less for the difference than it will for the fund or benchmark terminal value considered on its own.

Example 1 above shows exactly how the manager's contribution to the fund, $V_{P,T} - V_{B,T}$, arises from the interaction between his decisions or skill and the investor's cash flows. For the initial investment, Y_0, the relevant skill is that displayed over all three periods, measured by $R_P(0,3) - R_B(0,3)$. For the outflow $-Y_1$, the relevant skill is $R_P(1,3) - R_B(1,3)$, and for the inflow Y_2, it is $R_P(2,3) - R_B(2,3)$ only. If Y_2 is much larger than Y_0, for example, then the manager's skill will matter much more in period 3 than in periods 1 or 2, in terms of the manager's impact on the terminal value gained or lost.

The apparent problem for TWRs caused by external cash flows is the reason why TWRs are the industry norm. Assuming that external cash flows are controlled by the investor rather than the manager, it is deemed desirable to measure performance by $R_P(0,T) - R_B(0,T)$, which is unaffected by the cash flows, even though the difference between the TWRs can not be reconciled exactly with $V_{P,T} - V_{B,T}$, if at all.

The main complement or alternative to the TWR is the IRR, also known as the money-weighted rate of return. An IRR for a portfolio is a value for IRR that solves

$$Y_{P,0} \; = \; \sum_{t=1}^{T} Y_{P,t} / (1 + IRR_{P,T})^T \qquad (2)$$

for dates of equal length apart. $Y_{P,T} = V_{P,T}$ and has a negative sign, ie the terminal value is assumed to be paid out. The IRR *is* affected by intervening cash flows. Its use in performance measurement has probably been becoming more widespread, and several authors have recommended that IRR should feature much more prominently. Campisi (2004) makes the point that attribution using money weighting, though approximate, provides a more correct explanation of how the manager's contribution to the fund's TWR was made than does attribution using TWRs. Spaulding (2005) argues not only that IRRs should be reported alongside TWRs, but that IRRs should be the primary performance measure in many situations: 'in those cases where a money manager isn't controlling the

flows, TWRR is the best way to operate. But that's <u>it</u>! For just about every other circumstance, we should use... the internal rate of return' (p. 20).[1]

In addition, the standard performance measures used in academic research on private equity funds are values and IRRs (see, for example, Phalippou & Gottschalg, 2008). Values are used in the guise of the profitability index, which is the present value of the cash flows paid out by the fund divided by the present value of the cash flows paid in. The profitability index enables values to be used to compare the performance of funds of different sizes (ie different cash flows). The difference in the profitability index for fund and benchmark is equal to $V_{P,T} - V_{B,T}$ divided by the future value of all the inflows as at date T, using the benchmark's IRR as the rate for compounding in the denominator.[2]

We now discuss how measurement using values compares with IRRs. The first point to make is that, if there is only one change of sign in the cash flow series for both fund and benchmark, and if the terminal value is positive, both methods will provide the same answer to the question of whether the fund outperformed its benchmark. That is, $IRR_{P,T} - IRR_{B,T} > 0$ when $V_{P,T} - V_{B,T} > 0$. This can be seen by re-arranging equation (2). Assuming that there is one change of sign, after date x, the portfolio must have inflows up to date x followed by outflows:

$$V_{P,T} = Y_{P,T} = \sum_{t=0}^{x} Y_{P,t}(1+IRR_{P,T})^{T-t} - \sum_{t=x+1}^{T-1} Y_{P,t}(1+IRR_{P,T})^{T-t}$$

Since all the inflows are earlier than any of the outflows, and $V_{P,T} > 0$, an increase in $IRR_{P,T}$ must increase the value of the first summation more than the value of the second summation, holding constant all the cash flows before date T. Therefore, since the cash flows before T are the same for portfolio and benchmark, $V_{P,T} - V_{B,T} > 0$ requires that $IRR_{P,T} - IRR_{B,T} > 0$. This result is not guaranteed if there is more than one change of sign. There are series of cash flows for which a higher terminal value implies a *lower* IRR,[3] and others for which there is no solution in real numbers to equation (2). But for most portfolios, $IRR_{P,T} - IRR_{B,T}$ will be consistent with $V_{P,T} - V_{B,T}$, to the extent that both will have the same sign.

However, there are various problems with IRRs, of which we highlight two. One is that simple and exact attribution analysis of $IRR_P - IRR_B$ is impossible, as it is with TWRs. This is because the IRR for each asset class or attribute can not be added to arrive at the IRR for the whole fund. Campisi (2004) offers an approximate attribution analysis with money weighting. But we do not have to live with approximation: precision is possible using values.

A second problem with IRRs is their interpretation. If the investor withdraws money before the terminal date, the IRR is not the rate of return that he receives on his money, unless the rate of return on the money withdrawn is the same as the fund's IRR. Rather, the IRR is the rate of return on the cash invested that is just sufficient for the fund to provide the cash outflows that were made, including the final withdrawal of all the remaining value at the terminal date. The difference, between the rate of return on the investor's money, including that withdrawn before the terminal date, and the rate that is sufficient to provide the cash outflows, is important, but it is not easy to understand. IRR

has also been described as the average growth rate of all money invested in a fund (Maginn et al, 2007, p. 727). This is a helpful description, but IRR is not an average in the normal sense.

Consider Example 3, in which most of the money in the fund is withdrawn at date 1. Note that the $140 withdrawal from the benchmark at date 1 exceeds the $104 in the benchmark at that date. In effect the investor has to short-sell the benchmark by $36 at date 1 in order to obtain the full $140. He must close out the short position at date 4. So the $40.5 cash flow for the benchmark at date 4 is a payment *from* the investor.

Example 3. Data for comparison of IRR and the value method

Year	0	1	2	3	4	IRR
Fund returns		50%	3%	3%	3%	
Benchmark returns		4%	4%	4%	4%	
Withdrawal at date 1						
Cash flows for fund	$100.0	–$140.0	$0.0	$0.0	–$10.9	43.7%
Contribution to terminal value	$100.0	× 1.50	× 1.03	× 1.03	× 1.03 = $163.9	
		–$140.0	× 1.03	× 1.03	× 1.03 = –$153.0	
Cash flows for benchmark	$100.0	–$140.0	$0.0	$0.0	$40.5	4.0%
Contribution to terminal value	$100.0	× 1.04	× 1.04	× 1.04	× 1.04 = $117.0	
		–$140.0	× 1.04	× 1.04	× 1.04 = –$157.5	
No withdrawal at date 1						
Cash flows for fund	$100.0	$0.0	$0.0	$0.0	–$163.9	13.1%
Cash flows for benchmark	$100.0	$0.0	$0.0	$0.0	–$117.0	4.0%

The IRR for the fund is 44% pa, but this is nothing like what the investor obtained most years. In fact, the main determinant of what was obtained is whatever rate was available outside the fund on the $140 withdrawn at date 1. Had the money been left in the fund and then withdrawn at date 4, the IRR would only have been 13% pa (assuming the fund's returns would have been the same had the fund had more money in it). The benchmark IRR is 4% pa.[4] So the difference in IRRs is +40% if the $140 is withdrawn at date 1, and +9% if it is not withdrawn. This is hard to understand. The returns on the *fund* are not affected by whether the money is withdrawn. The returns on the investor's *money* are affected. But the IRR is not the rate of return on the investor's money. The investor will not really receive an extra 40% pa by withdrawing the money, if the interest rate outside the fund is three or four per cent. The problem is that IRR is neither the time-weighted rate of return on the fund nor the rate of return on the investor's money. It is a sort-of average rate on the money in the fund. It is not a rate that

would have been available had the investor left more money in the fund. Chasing 44% pa would have been like trying to find the end of a rainbow; the more money left in the fund, the lower its IRR would have been.[5]

Using values, $V_{P,4} - V_{B,4}$ is +\$51.4 if the \$140 is withdrawn, and +\$46.9 if it is not withdrawn. The withdrawal has some impact on the performance result, as we would expect with money weighting. But it is not the enormous and misleading impact that arises using IRRs. The difference of \$4.5 is the contribution to value from the withdrawal, which is $-\$140[(1.03)^3 - (1.04)^3] = \$4.5$. Under the value method, the reason why the investor is in fact \$4.5 better off for having withdrawn \$140 at date 1 is because he would have been \$4.5 worse off had he kept the money in the fund, earning 3% pa rather than 4% pa as in the benchmark.

In general, if money is put into the fund, the manager's decisions over time add to or detract from its value compared with the value of the same amount put into the benchmark. If money is taken out, the investor forgoes the subsequent gain or loss in value that would have arisen had the amount been kept in. In other words, after money is withdrawn, some of the subsequent impact of the manager on value is subtracted. That is how money weighting works using values. This assumption has to be made – cash flows entering a fund achieve the same returns as assets already there, cash flows leaving a fund lose the same returns as those remaining. The difference $V_{P,T} - V_{B,T}$ is the sum of the gains or losses on all the cash amounts invested in the fund at various times, *less* the gains or losses that would have arisen on cash amounts which were withdrawn, over the periods after the withdrawals, had these amounts been kept in the fund (see equation 1). The difference measures how much value the manager has added or lost, given the funds made available to him by the investor. This concept is readily understood, unlike the concept of IRR.

In the light of the above discussion, the question we ask is this. If money-weighted performance measurement is accepted as appropriate for many funds, why not do it using values instead of IRRs? The arguments for – and against – money weighting apply equally to the value-based method but the latter has the additional and sometimes significant benefits that attribution analysis is easy and that what is being measured is easier to understand.

Comparing and remunerating fund managers

How can fund managers be compared using values? We can adjust for funds of different sizes by using the profitability index, or by $(V_{P,T} - V_{B,T})/V_{B,T}$, the value gain or loss relative to the terminal benchmark value of the fund. If there are no cash flows between dates 0 and T (intervening cash flows), then TWRs and the value method are equivalent in their measure of overall performance (though not in their attribution analyses):

$$(V_{P,T} - V_{B,T})/V_{B,T} = (R_{P,T} - R_{B,T})/(1 + R_{B,T})$$

But if there are intervening cash flows, performance according to a value-based measure will always be the outcome of the interaction between the cash flows and the decisions of the manager. Thus, to

compare the performance of two or more managers using values, ideally the cash flows for the managers should be identical. If the cash flows are not identical, a measure using values can be calculated that is independent of the intervening cash flows, and that gives the same number as $(R_{P,T} - R_{B,T})/(1 + R_{B,T})$, This calculation is done by excluding the value gained or lost that results from the intervening cash flows, to give the gain or loss arising from the initial investment only. The calculation is explained in the next section.

A related question is the basis on which the manager should be remunerated. The measurement method used is, or should be, determined by the nature of the performance according to which the client wishes to pay the manager.[6] If the client wishes the manager to be paid for his work over a given T periods so that the impact of his actions on his remuneration is the same for each period, then TWR should be chosen. If the client wishes the manager to be paid so that the impact of his actions varies over time with the external cash flows by the client, then IRR or, better, the value-based method should be used. Part of the manager's fee could be a percentage per dollar gained in relation to the benchmark, or in relation to any 'hurdle value' specified by the client. There are arguments for both approaches, but it is not clear a priori that TWR is the better measure for the purpose of performance-based remuneration.

Our case is that the value method is a complement to TWRs rather than a substitute. The value method provides a bulletproof measure of the wealth actually gained or lost by a manager for a client during an assessment interval, given the intervening cash flows made by the client. TWRs do not reflect the gain or loss in wealth correctly. They do provide a convenient measure of the manager's performance or skill, although it is easy enough to adapt the value method so that it gives the same result as TWRs. The result using TWRs ignores the interaction between the intervening cash flows and the manager's decisions. It is up to the client whether the manager should be rewarded for skill measured in this way, or rewarded for wealth creation.

Why attribution for multiple periods works using values

In this section we explain why attribution analysis works exactly using values but not using returns. For a single period t, conventional attribution analysis is a way of explaining $R_{P,t} - R_{B,t}$. With two asset classes, the difference can be analysed as

$$\underline{R}_{Pa,t} - \underline{R}_{Ba,t} = w_{Pa,t-1}R_{Pa,t} - w_{Ba,t-1}R_{Ba,t}$$
$$+ \ \underline{R}_{Pb,t} - \underline{R}_{Bb,t} = w_{Pb,t-1}R_{Pb,t} - w_{Bb,t-1}R_{Bb,t}$$

$$R_{P,t} - R_{B,t}$$

where $R_{Pa,t}$ is the portfolio's return on a in period t, and $\underline{R}_{Pa,t}$ is the contribution to the portfolio's total return, $R_{P,t}$, from its investment in a: $\underline{R}_{Pa,t} + \underline{R}_{Pb,t} = R_{P,t}$. The difference $\underline{R}_{Pa,t} - \underline{R}_{Ba,t}$ can then be broken down into asset allocation, stock selection, and interaction, à la Brinson, Hood and Beebower (BHB, 1986), or in other ways.

If we are measuring returns over more than one period, then the problem of attribution for multiple periods arises. Suppose there are two periods. The difference we want to explain is that between the compound return on the portfolio and on its benchmark:

$$R_P(0, 2) - R_B(0, 2)$$
$$= (1 + R_{P,1})(1 + R_{P,2}) - (1 + R_{B,1})(1 + R_{B,2})$$

We know that

$$R_P(0, 2) = (1 + w_{Pa,0}R_{Pa,1} + w_{Pb,0}R_{b,1})(1 + w_{Pa,1}R_{Pa,2} + w_{Pb,1}R_{Pb,2}) - 1$$
$$= w_{Pa,0}R_{Pa,1} + w_{Pb,0}R_{b,1} + w_{Pa,1}R_{Pa,2} + w_{Pb,1}R_{Pb,2}$$
$$+ w_{Pa,0}R_{Pa,1}w_{Pa,1}R_{Pa,2} + w_{Pa,0}R_{Pa,1}w_{Pb,1}R_{Pb,2}$$
$$+ w_{Pb,0}R_{b,1}w_{Pa,1}R_{Pa,2} + w_{Pb,0}R_{b,1}w_{Pb,1}R_{Pb,2} \quad (4)$$

The above expansion shows why neither adding nor compounding the single-period contributions for each asset class will provide an exact breakdown of $R_{P,T}$. Adding the contributions ('arithmetic linking') gives

$$\underline{R}_{Pa,1} + \underline{R}_{Pa,2} + \underline{R}_{Pb,1} + \underline{R}_{Pb,2}$$
$$= w_{Pa,0}R_{Pa,1} + w_{Pa,1}R_{Pa,2} + w_{Pb,0}R_{Pb,1} + w_{Pb,1}R_{Pb,2} \quad (5)$$

Compounding the contributions ('geometric linking') and summing gives

$$(1 + \underline{R}_{Pa,1})(1 + \underline{R}_{Pa,2}) - 1 + (1 + \underline{R}_{Pb,1})(1 + \underline{R}_{Pb,2}) - 1$$
$$= w_{Pa,0}R_{Pa,1} + w_{Pa,1}R_{Pa,2} + w_{Pa,0}R_{Pa,1}w_{Pa,1}R_{Pa,2}$$
$$+ w_{Pb,0}R_{Pb,1} + w_{Pb,1}R_{Pb,2} + w_{Pb,0}R_{Pb,1}w_{Pb,1}R_{Pb,2} \quad (6)$$

Neither (5) nor (6) is equal to (4). In fact, the multiperiod contribution for each asset class cannot be defined unambiguously using returns. The cross-product terms in (4), such as $w_{Pa,0}R_{Pa,1}w_{Pb,1}R_{Pb,2}$, tell us that the contribution of a given asset class to the compound return on the portfolio is a function of the period-by-period weights and returns on *other* asset classes. That is why one cannot arrive at the compound return on the portfolio by adding the arithmetic or geometric returns for each asset class considered in isolation, without cross-terms: they do not measure the contribution to the compound return correctly. This was first noted by Burnie, Knowles and Teder (BKT, 1998, p. 66). Although there are solutions whereby adding or compounding the single-period contributions can be made to provide an exact breakdown of $R_P(0,T)$, they all involve making adjustments to the single-period contributions, and a consequent loss of transparency.

The value-based method enables the analyst to provide an exact attribution analysis without making any adjustments to the raw ingredients, which are the cash flows in and out of the portfolio and its asset classes, and the single-period returns on the asset classes. The reason is that *there are no cross-terms using values*. The contribution of any cash flow at a given date t to the final value of the portfolio at date T is the value of the cash flow multiplied by the relevant compound return. For an investment in an asset class, the relevant compound return is the compound return between t and T for the asset class in question. In the simple two-asset class, two-period case, the difference in values we are trying to explain, $V_{P,T} - V_{B,T}$, is analysed as follows:

$$V_{Pa,T} - V_{Ba,T} = Y_{Pa,0}(1 + R_{Pa,1})(1 + R_{Pa,2}) - Y_{Ba,0}(1 + R_{Ba,1})(1 + R_{Ba,2})$$
$$+ \ V_{Pb,T} - V_{Bb,T} = Y_{Pb,0}(1 + R_{Pb,1})(1 + R_{Pb,2}) - Y_{Bb,0}(1 + R_{Bb,1})(1 + R_{Bb,2})$$
$$\overline{\phantom{+ \ V_{Pb,T} - V_{Bb,T}}}$$
$$V_{P,T} - V_{B,T}$$

The contribution of the investment in asset class a to the portfolio's terminal value, $Y_{Pa,0}(1 + R_{Pa,1})(1 + R_{Pa,2})$, is not affected by the returns or weights for the other asset classes, unlike the contribution of the investment in a to the portfolio's compound return. This is why the contributions in values from the investments in each asset class considered in isolation add up to give the terminal value.

If the manager sells an amount, $Y_{Pa,1}$, of the portfolio's holding in class a at date 1 and invests this amount in class b, then two new, internal cash flows appear for the portfolio:

$$V_{Pa,T} - V_{Ba,T} = Y_{Pa,0}(1 + R_{Pa,1})(1 + R_{Pa,2}) - Y_{Ba,0}(1 + R_{Ba,1})(1 + R_{Ba,2}) - Y_{Pa,1}(1 + R_{Pa,2})$$
$$+ \ V_{Pb,T} - V_{Bb,T} = Y_{Pb,0}(1 + R_{Pb,1})(1 + R_{Pb,2}) - Y_{Bb,0}(1 + R_{Bb,1})(1 + R_{Bb,2}) + Y_{Pb,1}(1 + R_{Pb,2})$$
$$\overline{\phantom{+ \ V_{Pb,T} - V_{Bb,T}}}$$
$$V_{P,T} - V_{B,T}$$

There are no internal cash flows for the benchmark, assuming drifting weights.

It can now be seen why attribution analysis for the value added or lost over multiple periods is exact using values. The BHB attributes, for example, for a given asset class, a, measured in values, are derived from the contributions to $V_{P,T} - V_{B,T}$ of all the cash flows for a in the portfolio and its benchmark. The formulae parallel the BHB formulae for returns, but with the cash flows instead of the weights at each date (Armitage & Bagot, 2004, p. 23):

$$\text{asset allocation} = \sum_{t=0}^{T}(Y_{Pa,t} - Y_{Ba,t})[R_{Ba}(t,T) - R_B(t,T)]$$

$$\text{stock selection} = \sum_{t=0}^{T}Y_{Ba,t}[R_{Pa}(t,T) - R_{Ba}(t,T)]$$

$$\text{interaction} = \sum_{t=0}^{T}(Y_{Pa,t} - Y_{Ba,t})[R_{Pa}(t,T) - R_{Ba}(t,T)] \tag{7}$$

Because the above formulae are summations of contributions to terminal value from cash flows, there are no cross-terms in arriving at the contributions of each attribute for each asset class. So the analysis for investments in asset classes given above applies to their attributes. The grand summation across the attributes and asset classes reconciles with $V_{P,T} - V_{B,T}$ exactly.

For example, consider an investment in the portfolio by the client at date 1, Y_1. Let there be two further dates in the assessment interval. The fund manager invests $Y_{Pa,1}$ out of the inflow Y_1 into asset class a. The corresponding investment for the benchmark is given by $Y_{Ba,1} = w_{Ba,1}Y_1$. The contribution to $V_{P,T} - V_{B,T}$ from the asset allocation decision at date 1 relating to the inflow Y_1 and asset class a is

$$(Y_{Pa,1} - Y_{Ba,1})[(1 + R_{Ba,2})(1 + R_{Ba,3}) - (1 + R_{B,2})(1 + R_{B,3})]$$

This expression correcly measures the gain or loss in terminal value arising from the difference $Y_{Pa,1} - Y_{Ba,1}$, given the subsequent benchmark returns on a and on the benchmark as a whole.

The absence of cross-terms not only makes for an exact attribution analysis. It also enables the impact of each of the manager's decisions to be measured unambiguously. A cross-term normally does not correspond to a type of decision; it is a by-product of the analysis. Fund managers try to make money through asset allocation, currency selection, stock selection, and so on. These types of decision are normally made independently of each other, so it is helpful to be able to provide fully independent measures of the contributions from each type. The interaction term in (7) captures the contribution that cannot be assigned unambiguously to a single type. This is a different matter from the cross-terms in (4). They arise because of ambiguity in measuring the contribution of any particular investment to the compound return of a portfolio, not from the effect of interaction between two or more types of investment decision.

There can be a separate attribution analysis for *every* external cash flow, if desired. This makes the potential nature of a performance report very flexible. For example, the analyst could calculate the impact on $V_{P,T} - V_{B,T}$ of all the intervening cash flows made by the client. The report could then show (i) the contribution of the manager had the client invested the initial amount in the fund, and done nothing subsequently, and (ii) the contribution resulting from all the intervening cash flows. (i) is a pure measure of the manager's skill during the assessment interval, stripped of the impact of money weighting. (ii) is the outcome from the interaction of the client's cash flows and the manager's subsequent investment decisions. Exact attribution analyses could be done for (i) and (ii) separately, which would sum to give the analysis for the fund.

One problem that the value method does not solve is that the correct attribution analyses for each period treated separately will not give $V_{P,T} - V_{B,T}$ when added. This is because the fund and its benchmark will have different values after date 0, due to differences in their returns. A correct analysis using values for a given period t must start with equal values for the fund and its benchmark. Such an analysis can be done without difficulty, though the results for each period will not add to give the correct results for the full assessment interval. However, the value method assumes that an assessment interval has been agreed in advance between the client and the fund manager. In this case the client's focus will normally be on the interval as a whole. A pre-agreed interval with, say, two periods can be compared with an 800m race, two laps of the track. What matters is the runners' performance over the two laps together. We cannot properly assess performance over 800m by adding together the times for two separate 400m races, with two quite distinct starts.

MORE ON USING VALUES
Benchmark with weights fixed over time

In the rest of the paper, we examine several specific issues that can arise when using the value method. We start with the weights in the benchmark. Many analysts will prefer to work with weights that are fixed, rather than weights that drift due to differences between asset classes in their returns. Fixed weights are perfectly feasible using values, but they complicate the analysis somewhat. To

maintain fixed weights, there will need to be internal cash flows for the benchmark from time to time. This will add to the number of cash flow entries in the calculation of the benchmark's terminal value, and will make the results of an attribution analysis harder to interpret. But otherwise the value-based methodology is unchanged. The terminal values gained or lost from the internal flows necessary to maintain the fixed weights are calculated in the same way as the terminal value gained or lost from any other cash flow. Because the decision regarding the benchmark weights will make a difference to $V_{B,T}$, it is one of the items that should be agreed in advance when specifying how the manager will be evaluated.

Fixed weights are illustrated in Part A of Example 4. To keep the weights at 40% for asset class a and 60% for b, the benchmark requires a switch of $9.1 from b to a at date 1. This results in a terminal value for the benchmark that is $2.0 higher than if the weights are allowed to drift, because the return on a is higher than the return on b in period 2. The manager adds value of $1.7 with drifting weights in the benchmark, and loses $0.3 with fixed weights. In the attribution analysis that follows in the table, we have extra entries of $Y_{Ba,1} = +\$9.1$ and $Y_{Bb,1} = -\$9.1$. Since they relate to cash flows within the benchmark only, the analyst might wish to exclude the values of the 'attributes' for these flows from the main analysis – or to avoid them altogether by accepting a benchmark with drifting weights. However, the internal flows to maintain the fixed weights can be incorporated in a standard attribution analysis, as the example shows.

If the weights are allowed to drift, then no internal flows are needed in the benchmark. But the weight for a, for example, will drift from 40.0% at date 0 to 38.2% at date 2:

$$
\begin{aligned}
w_{Ba,2} &= w_{Ba,0}(1 + R_{Ba,1})(1 + R_{Ba,2})/(1 + R_{B,1})(1 + R_{B,2}) \\
&= 0.4(1.020)(1.020)/(1.248)(0.872) \\
&= 38.2\%
\end{aligned}
$$

Voids in asset classes

Voids cause no particular problem for the value method. A void arises because of a cash outflow from an asset class, and this negative flow affects terminal value in the same way as any other. The fact that there is no money left in the asset, in either portfolio or benchmark, is irrelevant. This can be seen in Part B of Example 4. The manager withdraws all $65 from asset class b at date 1, and transfers it to a, so that the portfolio is void in b in period 2. The transfer gives rise to cash flow entries in the attribution analysis, as expected, but does not require any other change. The transfer improves fund performance, because the return on a is higher than on b in period 2.[7]

Example 4. A portfolio and its benchmark with two asset classes
The attribution analyses use the formulae in (7).

Year	0	1	1	2	*Compound return*
Asset class *a* rtns: portfolio		2%		2%	4.0%
Asset class *a* rtns: benchmark		2%		2%	4.0%
Asset class *b* rtns: portfolio		30%		–10%	17.0%
Asset class *b* rtns: benchmark		40%		–20%	12.0%
Portfolio	Cash flow	Value	Cash flow	Value	
Asset class *a*	$50.0	$51.0	$0.0	$52.0	4.0%
Asset class *b*	$50.0	$65.0	$0.0	$58.5	17.0%
Total	$100.0	$116.0	$0.0	**$110.5**	
Returns		16.0%		–4.7%	10.5%
Benchmark with drifting weights					
Asset class *a*	$40.0	$40.8	$0.0	$41.6	4.0%
Asset class *b*	$60.0	$84.0	$0.0	$67.2	17.0%
Total	$100.0	$124.8	$0.0	**$108.8**	
Returns		24.8%		–12.8%	8.8%

Example 4, Part A. Benchmark with fixed weights

	0	1	1	2	*Compound return*
Asset class *a*	$40.0	$40.8	$0.0	$41.6	4.0%
			$9.1	$9.3	2.0%
Asset class *b*	$60.0	$84.0	$0.0	$67.2	17.0%
			–$9.1	–7.3	–20.0%
Total	$100.0	$124.8	$0.0	**$110.8**	
Returns		24.8%		–11.2%	10.8%

Attribution analysis of $0.3 loss in relation to benchmark with fixed weights

Asset class *a*

Asset allocation	($50.0 – $40.0)(4.0% – 10.8%) – $9.1(2.0% + 11.2%)	=	–$1.9
Stock selection	$40.0(4.0% – 4.0%) + $9.1(2.0% – 2.0%)	=	$0.0
Interaction	($50.0 – $40.0)(4.0% – 4.0%) – $9.1(2.0% – 2.0%)	=	$0.0

Asset class *b*

Asset allocation	($50.0 – $60.0)(12.0% – 10.8%) + $9.1(–20.0% + 11.2%)	=	–$0.9
Stock selection	$60.0(17.0% – 12.0%) – $9.1(–10.0% + 20.0%)	=	$2.1
Interaction	($50.0 – $60.0)(17.0% – 12.0%) + $9.1(–10.0% + 20.0%)	=	$0.4

Total			–$0.3

Example 4, Part B. Void in portfolio in asset class *b* in period 2

Year	0	1	1	2	Compound return
Portfolio	Cash flow	Value	Cash flow	Value	
Asset class *a*	$50.0	$51.0	$0.0	$52.0	4.0%
			$65.0	$66.3	2.0%
Asset class *b*	$50.0	$65.0	$0.0	$58.5	17.0%
			−$65.0	−$58.5	−10.0%
Total	$100.0	$116.0	$0.0	**$118.3**	18.3%

Attribution analysis of $9.5 gain in relation to benchmark with drifting weights

Asset class *a*

Asset allocation	($50.0 − $40.0)(4.0% − 8.8%) + $65.0(2.0% + 12.8%)	=	$9.1
Stock selection	$40.0(4.0% − 4.0%)	=	$0.0
Interaction	($50.0 − $40.0)(4.0% − 4.0%) + $65.0(2.0% − 2.0%)	=	$0.0

Asset class *b*

Asset allocation	($50.0 − $60.0)(12.0% − 8.8%) − $65.0(−20.0% + 12.8%)	=	$4.4
Stock selection	$60.0(17.0% − 12.0%)	=	$3.0
Interaction	($50.0 − $60.0)(17.0% − 12.0%) − $65.0(−10.0% + 20.0%)	=	−$7.0

Total	$9.5

Example 4, Part C. Short position in asset class *b* in period 2

Year	0	1	1	2	Compound return
Portfolio	Cash flow	Value	Cash flow	Value	
Asset class *a*	$50.0	$51.0	$0.0	$52.0	4.0%
			$165.0	$168.3	2.0%
Asset class *b*	$50.0	$65.0	$0.0	$58.5	17.0%
			−$165.0	−$148.5	−10.0%
Total	$100.0	$116.0	$0.0	**$130.3**	30.3%

Attribution analysis of $21.5 gain in relation to benchmark with drifting weights

Asset class *a*

Asset allocation	($50.0 − $40.0)(4.0% − 8.8%) + $165.0(2.0% + 12.8%)	=	$23.9
Stock selection	$40.0(4.0% − 4.0%)	=	$0.0
Interaction	($50.0 − $40.0)(4.0% − 4.0%) + $165.0(2.0% − 2.0%)	=	$0.0

Asset class *b*

Asset allocation	($50.0 − $60.0)(12.0% − 8.8%) − $165.0(−20.0% + 12.8%)	=	$11.6
Stock selection	$60.0(17.0% − 12.0%)	=	$3.0
Interaction	($50.0 − $60.0)(17.0% − 12.0%) − $165.0(−10.0% + 20.0%)	=	−$17.0

Total	$21.5

Short positions in asset classes

A short position results in a negative holding in the relevant asset class, and a cash flow received by the portfolio. In Part C of Example 4, the manager not only sells all of his holding in *b* at date 1, but goes short in *b* by $100. The resulting $165 is invested in *a* in period 2. Going short in *b* further improves fund performance, because *b* has a negative return in period 2. It can be seen that there are no additional complications from a short position.

Income from securities held

Dividends and other payments from securities held are cash inflows. Implementation of the value method should include an agreed procedure for such income. The simplest approach is to assume that income will be re-invested immediately in the asset class of the security from which it is received. Then no separate cash flow entries are needed for the relevant inflows, either for the portfolio or the benchmark, assuming that attribution analysis is taken to the level of asset classes rather than individual securities. Only if income is transferred to another asset class would an internal cash flow need to be recorded. Example 5 illustrates re-investment of a dividend.

_____

Example 5. Re-investment of a dividend
A portfolio's holding in a given asset class consists of two shares, *j* and *k*. The manager receives dividends of $400 at date 1a from the holding in *j*. He uses it to buy two more units of *j* for $160 and 12 more of *k* for $240. There is no other cash flow at date 1a, and the portfolio need not be valued then. Date 2 is the next review date. The total portfolio value is then $10,910. If the benchmark has the same two shares, and the rule for the benchmark is that a dividend is to be invested in the share generating it, then all $400 would be invested in share *j*, and the overall benchmark value at date 2 would be $10,925. Thus the portfolio would have underperformed by $15, the result of investing $240 of the dividend from *j* in 12 shares of *k*, rather than 3 further shares of *j* whose price then changed from $80 to $85.

_____

Date	1	1a *Before re- investment*	1a *After re- investment*	2
Share *j*				
Units	100	100	102	102
Price	$78	$80	$80	$85
Dividend received	$0	$400	$0	$0
Share *k*				
Units	100	100	112	112
Price	$20	$20	$20	$20

_____

It is a different matter if the manager pays out income to the client. Then the portfolio and benchmark will need to be valued each time income is paid out, and each outflow $-Y_t$ will be matched by the same outflow for the benchmark, drawn from the benchmark's asset classes in the usual way, ie $-Y_t(w_{Ba,t} + w_{Bb,t})$.

It is convenient for the benchmark to be a total-return index, or indices if there is more than one asset class. Then the analyst need not worry about the benchmark's income. Otherwise, the benchmark's income will need to be re-invested according to an agreed rule. For example, if a capital-only index is used, the rule can be that income is used to buy additional units in the index.

Expenses and taxes

The starting point for the value method is the actual terminal value of the fund, which will be net of the various expense and tax payments it has made. The precise amounts by which these negative cash flows reduce the terminal values can be calculated. The amounts can be added back as desired to obtain a terminal value that is, say, gross of tax but net of all expenses, or gross of tax and management fees but net of trading costs. To estimate $V_{B,T}$ on a net basis, the benchmark's cash outflows, from which $V_{B,T}$ is calculated, must subtract the expense and tax payments that have *not* been added back in calculating the fund's terminal value. Presentation of the amounts by which expenses and taxes reduce $V_{P,T}$ and $V_{B,T}$ will add to the number of headings in a value-based report, but, once again, the methodology is unaffected. The grand summation of the components of terminal value, gross or net, will continue to be exact.

Example 6 includes expense rates, applied to cash flows, and a tax payment. The purpose is to show how their impact on terminal value is arrived at, and also how the single-period returns are calculated net of expenses, before and after tax. The fund's actual terminal value is $5,744.68, which has arisen after expenses and tax have been paid along the way. It would be usual to measure performance gross of payments for tax. There is a single payment of $300 at date 2, and adding this back increases the fund's terminal value by $300 × 1.025 = $307.5, to $6,052.18. The benchmark's comparable terminal value, at date 3, is $6,025.75, calculated net of typical expense rates but gross of the $300 payout at date 2.

Whether the fund beats the benchmark depends, in the example, on the assumption made about the benchmark's expense rates. If we assume that the rates are typical for the type of fund in question, the fund's terminal value net of expenses is $26.43 greater than the benchmark's (= $6,052.18 – $6,025.75). If we assume that the rates are the same as the fund's rates, the fund's terminal value is $36.11 less than the benchmark's (=$6,052.18 – $6,088.29). So the fund's lower expense rates result in added value of $62.54 (check: $6,088.29 – $6,025.75), offsetting the relatively poor returns of the manager versus the benchmark.

To calculate the benchmark's terminal value net of tax, the same outflows for tax that the fund paid are applied to the benchmark. In the example, this causes the benchmark's terminal value to fall by $300 × 1.03 = $309.0, whichever assumption is made about its expense rates.

The example shows that the expenses and tax payments are incorporated clearly and precisely. In fact, it is necessary to work with values to arrive at fully accurate single-period returns net of expenses or tax, even if the final figure one wants is a TWR.

Example 6. Values and returns after expenses and taxes

In this example the expense rates are applied to the cash flows. The value at the end of a period t is calculated by applying the gross return to the money in the fund at date $t–1$, ie the fund's value plus any cash flow net of expenses at $t–1$, excluding any tax payment but including the expense associated with the payment. To calculate the value after tax at t, any tax payment at $t–1$ is also subtracted from the value for $t–1$. The net return for period t is the return on the money in the fund at $t–1$ with the expense charge at $t–1$ added back. For example, the value before tax of the portfolio at date 3 is ($5,907.57 – $3.00)(1.025) = $6,052.18, and the net return before tax in period 3 is $6,052.18/$5,907.57 – 1 = 2.45%. The value after tax at date 3 is ($5,907.57 – $303.00)(1.025) = $5,744.68, and the net return after tax is $5,744.68/$5,907.57 – 1 = –2.76%.

	Date 0 Cash flow	Period 1 Return	Date 1 Value	Date 1 Cash flow	Period 2 Return	Date 2 Value	Date 2 Cash flow	Period 3 Return	Date 3 Value
Portfolio									
Cash inflows and gross returns	$5,000.00	7.00%		+$500.00	3.00%			2.50%	
Cash outflow for tax payment							–$300.00		
Expense rates	2.00%			1.50%			1.00%		
Net cash flows	$4,900.00			+$492.50			–$303.00		
Values before tax			$5,243.00			$5,907.57			$6,052.18
Net returns before tax		4.86%			2.87%			2.45%	
Values after tax			$5,243.00			$5,907.57			$5,744.68
Net returns after tax		4.86%			2.87%			–2.76%	
Benchmark									
Cash inflows and gross returns	$5,000.00	6.00%		+$500.00	4.00%			3.00%	
Cash outflow for tax payment							–$300.00		
Typical expense rates	3.00%			2.00%			2.00%		
Values and net rtns before tax		2.82%	$5,141.00		3.82%	$5,856.24		2.89%	$6,025.75
Values and net rtns after tax		2.82%	$5,141.00		3.82%	$5,856.24		–2.38%	$5,716.75
Portfolio's expense rates	2.00%			1.50%			1.00%		
Values and net rtns before tax		3.88%	$5,194.00		3.86%	$5,913.96		2.95%	$6,088.29
Values and net rtns after tax		3.88%	$5,194.00		3.86%	$5,913.96		–2.28%	$5,779.29

CONCLUSION

In this paper we have explored further how the value method works and what its role is in performance measurement. We have suggested that values be used instead of IRRs for money-weighted measurement. The key to the value method is that every external and internal cash flow results in a separate, measurable change in the fund's terminal value, given the relevant returns following the date of the cash flow. There is nothing complicated in the calculations or the attribution analysis. We encourage readers to try the value method for themselves.

An aspect we have not discussed is the measurement of risk using values. This is a large topic that merits attention in the future.

NOTES

1. Other authors who recommend money weighting, especially for attribution analysis, include Darling & MacDougall (2002) and Illmer & Marty (2003).

2. The benchmark's profitability index (*PI*) is defined as

$$PI_{B,T} = [\text{PV}(V_{B,T}) + \text{PV}(\textit{Outflows})] \div \text{PV}(\textit{Inflows})$$

where PV(.) denotes the present value of the term in brackets, the *Outflows* are all the cash withdrawals from the portfolio and its benchmark between dates 0 and *T*–1, and the *Inflows* are all the cash inflows (if $V_{B,T}$ is negative, it is an inflow). The benchmark must have a *PI* of one. For this to be the case, the discount rate for the benchmark's outflows and inflows must be the benchmark's IRR. It is also the discount rate to use in calculating the *PI* for the portfolio, because to achieve a *PI* of more than one, the present value of the portfolio's outflows and terminal value must exceed the present value of the inflows had they been invested in the benchmark, earning the benchmark's IRR. So we can now write

$$PI_{P,T} - PI_{B,T} = [\text{PV}(V_{P,T}) + \text{PV}(\textit{Outflows}) - \text{PV}(V_{B,T}) - \text{PV}(\textit{Outflows})] \div \text{PV}(\textit{Inflows})$$
$$= (V_{P,T} - V_{B,T})/\text{FV}(\textit{Inflows})$$

where FV(.) denotes the future value. For a given cash flow Y_t, $\text{FV}(Y_t) = \text{PV}(Y_t)(1 + R)^T$, where *R* is the discount rate (in this case, $R = IRR_B$).

3. An example is the following series of cash flows:

| $200 | –$1,000 | $800 | –$150 | $150 | $150 | $150 | –$150 |

One of the IRRs is 15.9%. If the terminal payout is increased to, say, $160, the IRR falls to 15.3%.

4. But because there is more than one change of sign in the cash flow series for the benchmark, there is more than one IRR. The second is 6.0% pa.

5. Phalippou (2008) discusses problems with IRR in the context of private equity funds. He describes a modified IRR that alleviates some of the problems. He does not consider attribution analysis.

6. See Schliemann & Stanzel, 2008, for a recent discussion of performance-based compensation.

7. Footnote 6 in Armitage & Bagot (2004) discusses voids, and suggests a special procedure to deal with them. But no special procedure is needed (and the one we originally proposed is incorrect).

REFERENCES

Armitage, Seth, and Gordon Bagot, 'What Has the Manager Done for Me? A Value-based Method of Measuring Fund Performance in Relation to a Benchmark', *Journal of Performance Measurement* 9, Fall 2004, pp. 19-34.

Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower, 'Determinants of Portfolio Performance', *Financial Analysts Journal*, July-August 1986, pp. 39-44.

Burnie, J. Stephen, James A. Knowles, and Toomas J. Teder, 'Arithmetic and Geometric Attribution', *Journal of Performance Measurement* 3, Fall 1998, pp. 59-68.

Campisi, Stephen, 'The Case for Money-weighted Performance Attribution', *Journal of Performance Measurement* 8, Spring 2004, pp. 31-41.

Darling, Robert, and Alastair MacDougall, 'Using Performance Statistics: Have Measurers Lost the Plot', *Journal of Performance Measurement* 7, Winter 2002, 22-32.

Illmer, Stefan, and Wolfgang Marty, 'Decomposing the Money-weighted Rate of Return', *Journal of Performance Measurement* 7, Summer 2003, 42-50.

Maginn, John L., Donald L. Tuttle, Jerald E. Pinto, and Dennis W. McLeavey, *Managing Investment Portfolios*, 3rd ed., CFA Institute Investment Series, Wiley, 2007.

Phalippou, Ludovic, 'The Hazards of Using IRR to Measure Performance: The Case of Private Equity', *Journal of Performance Measurement* 12, Summer 2008, pp. 55-67.

Phalippou, Ludovic and Oliver Gottschalg, 'The Performance of Private Equity Funds', *Review of Financial Studies*, 2008, forthcoming.

Schliemann, Martin and Matthias Stanzel, 2008, 'Performance-based Compensation Contracts in the Asset Management Industry', *Journal of Performance Measurement* 12, Spring, pp. 61-70.

Spaulding, David, 2005, 'Contrasting Time- and Money-weighted Returns: When Each Should be Used', *Journal of Performance Measurement* 10, Fall, pp. 20-30.