
Why maximising information ratios is incorrect

Arun S. Muralidhar

M^{cube} Investment Technologies, 80 Saratoga Drive, West Windsor, NJ 08550, USA.
Tel. +1 22 554 6831; E-mail: asmuralidhar@mcubeit.com

Received: 10th February, 2005

Arun S. Muralidhar is Chairman and Founder of M^{cube} Investment Technologies LLC (www.mcubeit.com), a company that assists investors in making high value-added decisions effectively through innovative web-based technologies. He also serves as Managing Director at FX Concepts Inc. Prior to that, he was Head of Research and a member of the Investment Management Committee at the World Bank, where he helped manage approximately \$40 billion in assets. He is the author of *Innovations in Pension Fund Management* (Stanford University Press, 2001) and co-author (with the late Professor Franco Modigliani) of *Rethinking Pension Reform* (Cambridge University Press, 2004). He currently serves on the Advisory Board for the *Journal of Performance Measurement* and the *Journal of Asset Management*.

Practical applications

Many practitioners focus on trying to maximise the information ratio of active assets for a given tracking error, as there is a perception that higher information ratios imply higher skill. This paper first demonstrates that higher information ratios may not transfer into higher confidence in skill and, more importantly, that a new class of risk-adjusted performance measures are providing invaluable advice on optimal portfolio construction. In short, for a given tracking error budget, the paper demonstrates that the correct approach is to maximise risk-adjusted return. Using performance measures such as the M-square or the M-cube to evaluate risk-adjusted performance also provides clients with invaluable advice on optimal portfolio construction; specifically, allocations to cash, the passive benchmark ('the beta') and the active strategy (which may include multi-manager portfolios). In short, a more effective way of creating optimal portfolios is to integrate the active-passive and leverage-deleverage decisions to achieve a desired risk budget (with certain constraints on volatility of the optimal portfolio) and not just maximise the information ratio on the active component.

Abstract

Many papers on active management argue for maximising information ratios using a risk-budgeting framework. Recent innovations in risk-adjusted performance measures show why maximising information ratios on active portfolios could be the wrong policy and also provide a different twist to the discussion on separating

alphas from betas. The literature on maximising information ratios focuses only on the active management process and ignores two actions used by clients or managers to improve risk-adjusted performance: passive management and leverage/deleverage using cash. This paper demonstrates the impact of maximising the wrong objective function and shows the benefit of

maximising risk-adjusted returns for the entire fund, rather than the information ratio on the active component.

INTRODUCTION

Since the prospective returns for passive benchmarks are moderate (having declined over the last three years), clients are searching for higher returns. Many are exploring different ways to generate excess returns over the benchmark (alpha) through a mix of strategies. The challenge for clients, is to construct robust portfolios using the many possible investment ideas that they have researched or have access to. All this, however, needs to be done within existing risk budgets.

Portfolio managers are trying to create portfolios of the best investment rules for their clients to either hold onto or gain new assets. Pension funds are trying to implement the best active management decisions (tactical asset allocation, sector allocation etc) and hire the best external managers and allocate appropriate amounts to them (where portfolios are managed by external managers). In essence, risk management today is about making the best investment decisions to generate the highest possible returns per unit of risk.

There has been a flood of papers on the fundamental law of active management and how clients should select optimal portfolios. In their widely read book, Grinold and Kahn¹ make three distinct statements: (1) Larger information ratios are better than smaller (p. 6); (2) The notion of success is captured and quantified by the information ratio (p. 110); and (3) every investor seeks

the strategy or manager with the highest information ratio (p. 125). These statements (and related research) will be contrasted in this paper. Grinold and Kahn make an interesting point that the excess returns are uncorrelated with the benchmark returns, but this is rarely examined in practice and finding true uncorrelated alpha from investment managers or strategies is often very difficult, if not impossible.² Many academics and practitioners make the assumption that investors wish to maximise the information ratio (defined as the excess return over the benchmark divided by the standard deviation of excess returns) or, alternatively, to maximise the return over benchmark for a target risk budget (called the tracking error budget) by optimally allocating funds across many competing active alternatives. The idea behind this approach, not always explicit, is that investors make a choice between how much of the portfolio will be active rather than passive (benchmark performance), and then try to maximise the returns on the active portion for the given risk budget. By doing so, however, they are creating inefficient portfolios, because they have either created the portfolio in two stages or excluded useful investment options which can allow investors to maximise risk-adjusted returns. The idea of using investments in risk-free assets or passive benchmarks to increase risk-adjusted returns has been examined in the context of risk-adjusted performance measures and, increasingly, such performance measures are providing guidance on portfolio construction. In short, investors should try to maximise risk-adjusted returns on the

whole portfolio, not just information ratios on the active portfolio.

WHY INFORMATION RATIOS ARE INCOMPLETE STATISTICS

For a long time, investors satisfied themselves with using the Sharpe ratio (ratio of excess return over the risk-free rate to the standard deviation of that excess³) and the information ratio (ratio of excess returns *vis-à-vis* a benchmark to standard deviation of excess returns) to distinguish whether risk-adjusted performance was good or poor. The standard deviation of excess returns is referred to as tracking error in the asset management industry. The formulae for tracking error and information ratio are provided in the Appendix. The simplest reason for disregarding the strategy with the highest information ratio is that it may not produce adequate returns if the risk taken is too low.⁴ For example, a strategy with an information ratio of 1, but which takes only 0.1 per cent risk generates only 0.1 per cent of returns. This paper moves beyond this simple shortcoming to show that even in the context of full utilisation of risk budgets, it could potentially be an inadequate measure.

A major contribution from Modigliani and Modigliani⁵ highlighted that such measures were inadequate for investors, as the ratios said nothing about performance and provided no guidance as to how portfolios should be constructed. After all, one cannot use a ratio to pay pensioners or any other liability; risk-adjusted performance is a more direct measure of success.

By adapting the Sharpe ratio, they were able to provide a true performance measure (ie measured in basis points of performance) which adjusted for differences in volatility between a portfolio and its benchmark. This measure also provided some guidance as to how much should be invested in the risk-free rate and the portfolio being reviewed. This measure came to be called the M^2 measure of risk-adjusted performance, and the authors were able to demonstrate that rankings of portfolios would be changed when this measure was used as opposed to raw performance. They were also able to show the flaws in using the information ratio — a measure that potentially rejected portfolios with high risk-adjusted performance because they had negative information ratios. Research by Muralidhar⁶ extended this approach to include the possibility of adding passive portfolios to the optimal mix (ie in addition to the risk-free asset and the active opportunities) to demonstrate that this allowed for an optimal normalisation of risk-adjusted excess returns subject to a tracking error budget. More recent work⁷ shows that information ratios may not say enough about confidence in skill, as high information ratios could lead to low confidence in skill, depending on how the performance has been generated (ie an examination of the excess returns, the variances of the benchmark and active strategy and the correlation between the two is required). In addition, all these measures suffer from the problem of not adjusting for time, as a high information ratio over a short period of time may convey the same confidence in skill as a much lower information ratio over longer periods of

time. This paper ignores the time dimension, which will be addressed in future research.⁸

Therefore, investors who seek to maximise the information ratio over a number of different active strategies or managers run a big risk — that of excluding the possibility of two basic alternatives, namely using the risk-free asset (or cash) or passive alternatives to create portfolios with the optimal risk-adjusted excess returns. None of these possibilities is out of the realm of use either by an active manager or within an active strategy and, hence, it shows how a simple, naïve assumption of excluding investment possibilities can lead to the wrong result.⁹

M^2 measure

Modigliani and Modigliani⁵ make an important contribution by showing that the portfolio and the benchmark must have the same risk in order to compare them in terms of basis points of risk-adjusted performance. They define risk as the volatility or standard deviation of a portfolio's returns. They propose that the portfolio be leveraged or deleveraged using the risk-free asset. In effect, combining cash with the active management strategy creates a more effective portfolio. If B is the benchmark being compared with portfolio 1, r denotes the return and σ denotes the standard deviation, then the leverage factor d is defined as follows

$$d = \sigma_B / \sigma_i \quad (1)$$

The M^2 adjustment creates a new portfolio, called the risk-adjusted portfolio (RAP), which return $r(\text{RAP})$ is equal to the 'leverage' factor multiplied by the original

return plus one minus the leverage factor, multiplied by the risk-free rate. Thus, if portfolio F is the riskless asset with zero standard deviation and is uncorrelated with other portfolios, the risk-adjusted return

$$r(\text{RAP}) = d \times r(\text{active portfolio}) + (1 - d)r(F) \quad (2)$$

where

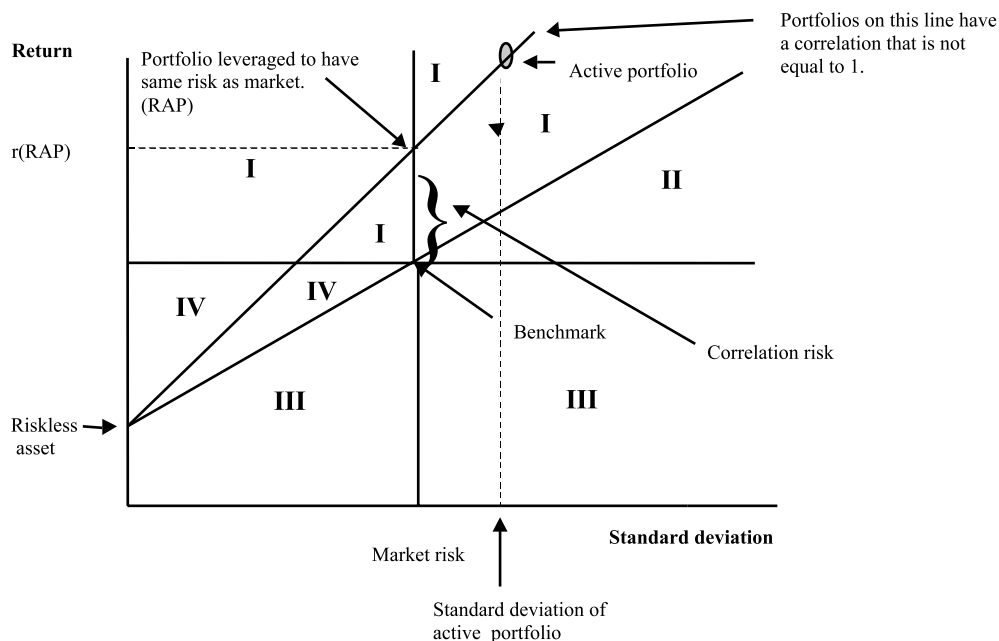
$$\sigma_{\text{RAP}} = \sigma_B \quad (3)$$

Figure 1 demonstrates this transformation by carving up the return-standard deviation space into four regions and showing how a strategy may be levered or delevered to have the same standard deviation as that of the benchmark. The correlation of the original portfolio to the benchmark is identical to the correlation of the RAP to the benchmark, as 'leverage or deleverage' using the risk-free rate does not change the correlation characteristics (because of the zero correlation). The correlation of the active portfolio with the benchmark is normally less than unity. If the correlation = 1, it could lead to a riskless arbitrage. Notice, this discussion is not dependent on the correlation of the excess return versus the benchmark.

Specifically, the four regions in Figure 1 using the M^2 measure are where an investor would experience:

- (1) portfolio outperformance on an absolute and risk-adjusted basis (positive information ratio);
- (2) portfolio outperformance on an absolute basis and *underperformance on a*

Figure 1: Evaluation of the M^2 measure



risk-adjusted basis (positive information ratio);

- (3) portfolio underperformance on an absolute and risk-adjusted basis (negative information ratio); and
- (4) portfolio underperformance on an absolute basis and *outperformance on a risk-adjusted basis* (negative information ratio).

This paper suggests that this M^2 adjustment allows for a comparison of ‘apples to apples’, ie returns from the benchmark and the RAP have the same volatility. It shows that making this adjustment can reverse peer rankings of mutual funds, managers or active strategies, especially those based on the information ratio. As a result, the paper discards the use of the information ratio for evaluating active managers or strategies, as

it could lead to incorrect decisions. For example, portfolios in region IV would have a negative information ratio, but would have a risk-adjusted performance greater than the benchmark, or one would accept a manager in region II who has a negative risk-adjusted performance. This first step in adjusting for risk is adequate to provide a contrast to the three major claims of Grinold and Kahn¹ which were highlighted in the introduction. The M^2 measure is also preferred to the Sharpe Ratio as it expresses risk-adjusted performance in terms of basis points of outperformance and provides guidance on assets allocated to the active asset allocation strategy or external manager (allocation of d) and the risk-free asset (allocation of $1 - d$). Graham and Harvey¹⁰ propose a variation of this method, assuming the

riskless asset need not be an asset uncorrelated with other assets. This only leads to different allocations across funds rather than suggesting a new approach.

The M^2 adjustment made the comparison in terms of basis points of outperformance by ensuring all portfolios had the same variance as the benchmark. It ignored the possibility of investing in the passive benchmark, however. As a result, the one major shortcoming was that two investment strategies, normalised for the benchmark volatility, could have different correlations with the benchmarks and hence different tracking errors (see equation A1). Tracking error is important to investors, especially institutional investors or those to whom responsibility has been delegated (ie in pure economic parlance, principals who have hired agents to manage their portfolios), because it provides a measure of the variability of a strategy/manager's returns around the benchmark. Investors would prefer, all else being equal, funds/agents with lower tracking error (and hence greater predictability in returns and higher confidence in skill). Hence, the M^2 rankings could provide investors with incorrect information about the relative risk-adjusted performance of funds for certain types of investors.¹¹

M^3 methodology: Adjusting for differences in correlations

An investor has to rely on the available data to make projections for the future. Assuming historical distributions are preserved in the future, the three-dimensional problem of comparisons of return, standard deviations and

correlations (ρ) has to be synthesised into a simple two-dimensional space of return and risk.¹² In mean-variance space, the riskless asset is portfolio F (with returns $r(F)$) and it can be used to leverage or deleverage the desired mutual fund/manager or active strategy. In tracking error space, the only riskless portfolio is the one with zero tracking error (TE) or the benchmark portfolio, as it is perfectly correlated with itself (where $\rho = 1$, TE = 0, as $\sigma_B = \sigma_1$). Therefore, combining active strategies/managers with passive benchmarks and the riskless asset can be used to alter the overall portfolio's standard deviation and its correlation with the benchmark.

To create measures of correlation-adjusted performance, the investor needs to invest in the active strategy (or strategies), the riskless asset and benchmark to ensure: (a) the volatility of this composite is equal to that of the benchmark;¹ and (b) the tracking error of this composite is equal to the target tracking error.⁶ This new measure, termed the M^3 (M -cube) measure, recognises that the investor has to consider basis points of risk-adjusted performance after ensuring that correlations of various funds versus the benchmark are also equal. In effect, now all tracking errors are normalised to the same value.

M^3 model as applied to a retail investor

While the principle is identical for an investor hiring an external manager or implementing an active management strategy (ie an investment manager), this section assumes that an investor is invested in a defined contribution pension plan and

must evaluate several mutual funds. Hammond¹³ states that, to establish performance-related thresholds for managers, the investor must set a target tracking error and compare funds with the target. A similar approach is proposed by Litterman *et al.*¹⁴

Assume that the investor is willing to tolerate a certain target annualised tracking error around the benchmark, say 300 basic points (bps) (TE(target)).¹⁵ The investor essentially wants to earn the highest risk-adjusted alpha for a given tracking error and variance of the portfolio. Now define a , b and $(1 - a - b)$ as the proportions invested in the mutual fund, the benchmark and the riskless asset. Let CAP be the correlation adjusted portfolio and therefore the returns of a CAP

$$r(\text{CAP}) = ar(\text{mutual fund}) + br(\text{B}) + (1 - a - b)r(\text{F}) \quad (4)$$

As one can see, this is an extension of the M^2 measure. Further, the investor must hold appropriate proportions of each to ensure the final portfolio has the target tracking error and the standard deviation of the benchmark. For a specific mutual fund, say mutual fund 1 with a risk-adjusted return $r(\text{CAP}-1)$, equation (4) can be rewritten as

$$r(\text{CAP} - 1) = ar(1) + (1 - a - b)r(\text{F}) + br(\text{B}) \quad (4)$$

where the coefficients of each portfolio represent the optimal weight of that specific portfolio to ensure complete risk adjustment. In addition, from the constraint on tracking error, there is a unique target

correlation between the CAP and benchmark B. As demonstrated by Muralidhar,⁶ this target correlation of the portfolio with that of the benchmark ($\rho_{T,B}$) is given by the equation for tracking error (equation A1) when $\sigma_B = \sigma_{\text{CAP}-1}$, namely

$$\rho_{T,B} = 1 - \frac{\text{TE}(\text{target})^2}{2\sigma_B^2} \quad (5)$$

By maximising the $r(\text{CAP})$ subject to the condition that the variance be identical to the benchmark, and its correlation to the benchmark equal to the target correlation, one finds that for mutual fund 1

$$a = + \sqrt{\frac{\sigma_B^2(1 - \rho_{T,B}^2)}{\sigma_1^2(1 - \rho_{1,B}^2)}} = \frac{\sigma_B}{\sigma_1} \sqrt{\frac{(1 - \rho_{T,B}^2)}{(1 - \rho_{1,B}^2)}} \quad (6)$$

$$b = \rho_{T,B} - a \frac{\sigma_1}{\sigma_B} \rho_{1,B} \quad (7)$$

The details of these calculations are provided in Muralidhar.⁶ If one substitutes for a in equation (7), the allocation to the benchmark is independent of variances and is only a function of the correlation terms. While b and $(1 - a - b)$ may be greater than or less than zero (negative coefficients being equivalent to shorting the futures contract relating to the benchmark and borrowing at the risk-free rate), a is constrained to being positive, as it is not currently possible to short mutual funds.¹⁶ With active management strategies, however, a can be positive and negative.

This method is preferred to the M^2 , as it: (1) expresses risk-adjusted performance in basis points; (2) gives advice on portfolio construction — specifically between the risk-free asset, the benchmark (passive investing) and the active portfolio (active

Table 1: Information on mutual fund managers

	Return (%)	Risk (%)	Correlation					Excess ret (%)	Tracking err (%)	IR	rRAP (%)
			Rf	Bench	Mgr A	Mgr B	Mgr C				
Risk free	5.00	0.00	1	0.00	0.00	0.00	0.00	-18.01	13.39	-1.35	
Benchmark	23.01	13.39	0	1.00	-0.18	0.36	0.14	0.00	0.00		
Manager A	22.07	10.88	0	-0.18	1.00	0.29	0.78	-0.94	18.71	-0.05	26.0
Manager B	25.02	15.71	0	0.36	0.29	1.00	0.82	2.01	16.58	0.12	22.1
Manager C	34.44	15.58	0	0.14	0.78	0.82	1.00	11.43	19.02	0.62	30.3

management); and (3) provides rankings that are identical with rankings based on skill for *equal time* horizons.

Muralidhar^{17,18} extends this approach to multiple manager portfolio or multiple investment strategies which is briefly summarised here.¹⁹ Define K as the portfolio of active management strategies, each with weight w_i , and now the objective function is to select optimal weights for each active investment opportunity (or aw_i), passive benchmark b and risk-free asset $(1 - a - b)$, so as to maximise the risk-adjusted return subject to a tracking error target or budget.

Mathematically, this is shown as choosing w_i to

$$\text{maximise } r(\text{CAP} - \text{K}) = \text{maximise}[ar(\text{K}) + br(\text{B}) + (1 - a - b)r(\text{F})] \quad (8)$$

subject to

$$\sigma_{\text{CAP-K}} = \sigma_{\text{B}} \quad (9)$$

$$\rho_{\text{CAP-K}} = \rho_{\text{T,B}} \quad (10)$$

$$\sum w_i = 1 \quad (11)$$

Muralidhar⁶ demonstrates that maximising the $r(\text{CAP})$ as opposed to maximising the

information ratio provides the best outcome as well as outcomes that provide the highest confidence in skill in the strategy. Even where clients are not allowed to ‘leverage’ as Figure 1 demonstrates, some strategies are better delevered to make for optimal portfolio construction. As a result, this technique can be used with any restrictions that a particular client may have.

A quick numerical simulation

While this discussion, especially Figure 1, provides adequate reason to disregard the maximisation of information ratios as an objective, the same is demonstrated with some numbers. Compare the strategy of maximising the risk-adjusted return versus the process of either maximising the information ratio on the active component or the classic two-stage process of differentiating between an active and passive component and then maximising information ratios within the active component. Assume that there are three active strategies to choose from, and their performance relative to the benchmark is provided in Table 1. Also, assume a risk-free asset (F) with an annualised return

Table 2: Allocations for maximising risk-adjusted performance versus maximising information ratios

	<i>Max $r(\text{CAP})$ (%)</i>	<i>Max IR (Active only) (%)</i>
Risk free	-73.2	0.0
Benchmark	80.1	0.0
Active (of which)	93.1	100.0
Manager A	92	120
Manager B	14	156
Manager C	-6	-175
Total	100.00	100.00
Return	36.48	4.99
Risk	13.39	6.97
Tracking error	10.00	12.70
IR	1.35	-1.42
Correlation with BM	0.72	0.36

of 5 per cent with no variance. Table 1 provides information on the correlation across investment strategies and also information ratios and risk-adjusted performance (using the single active strategy) assuming a target tracking error of 10 per cent. Notice also that all managers have a tracking error far in excess of the target. Table 1 demonstrates how Manager A, with a negative information ratio, has a positive risk-adjusted performance (higher than that of Manager B who has a positive information ratio). While not reported in the table, Manager A has the highest $r(\text{CAP})$ of 35.1 per cent, followed by Manager C with an $r(\text{CAP})$ of 34.1 per cent. This shows that, if an investor normalises for standard deviation or tracking error in the correct manner to determine the highest risk-adjusted return, the

information ratio conveys little information about the value or skill of the manager. For simplicity, no constraints are imposed on the sign or size of any allocation, as this would provide a very specific example of the more general proof attempted here.

A more interesting aspect is when multiple investment opportunities are included in a portfolio. Two strategies are constructed in Table 2: (1) maximising the risk-adjusted performance versus the target tracking error using M^3 ; and (2) maximising the information of the active component in isolation and then combining with cash and passive benchmarks in the second stage. The second simulation is conducted in two ways: (1) try to maximise the information ratio on the active portion only assuming the tracking error budget of 10 per cent (which is not achievable in this case for the

Table 3: Allocations for maximising risk-adjusted performance versus maximising information ratios through a two-stage process (active-only for a 20 per cent budget and then active versus passive)

	<i>Max r(CAP)</i> (%)	<i>Max IR (active only)</i> (%)	<i>Max portfolio return (based on active portfolio)</i> (%)
Risk free	-73.2	0.0	-8.8
Benchmark	80.1	0.0	117.1
Active (of which)	93.1	100.0	-8.3
Manager A	92	-33	3
Manager B	14	6	0
Manager C	-6	127	-11
Total	100.00	100.00	100.00
Return	36.48	37.95	23.36
Risk	13.39	17.98	14.66
Tracking error	10.00	20.00	10.00
IR	1.35	0.75	0.04
Correlation with BM	0.72	0.21	0.75

investment options chosen); and (2) maximise the information ratio on double the budget (20 per cent tracking error) and then decide between this active portfolio and passive management and cash. As Table 2 shows, if one tries to maximise the information ratio on the active-only portion, the lowest possible tracking error achievable is 12.7 per cent, which gives a negative information ratio and a poor overall return. Compare this with the situation where the risk-adjusted performance is maximised. This is not a surprise, as all the active managers are high tracking error managers. More important, the allocation to Manager C is negative, and this was the manager with the highest information ratio. In a perverse way, the

high information ratio manager is desired to help improve the portfolio characteristics through a *short* position.

Table 3 demonstrates the case where the tracking error on the active-only portion is increased to 20 per cent and the active portfolio allocations are locked in. In this case, the total return is more than that of the ‘maximise risk-adjusted performance’ simulation, but with higher risk (both absolute and relative). When this active allocation is combined with the risk-free asset and the benchmark, the overall return collapses in order to achieve the tracking error target. In short, maximising information ratios on active components, either in isolation or as a precursor to active-passive allocation decisions for a

tracking error budget, cannot do as well as strategies that maximise risk-adjusted performance.

CONCLUSION

Research on risk-adjusted performance measures is providing guidance on optimal portfolio construction for given risk budgets. It has been demonstrated that, with respect to information ratios on active portfolios, either in choosing a single strategy/manager or in constructing portfolios based on multiple active strategies, bigger is not necessarily better. A naïve approach of maximising the information ratio across many active alternatives or even maximising the excess return subject to a tracking error budget suffers from either excluding certain investment alternatives (leveraging or deleveraging using cash and the passive benchmark) or using risk budgets inefficiently, so that the chosen alternative may not have the highest return or confidence in skill level. Instead, clients with a target tracking error relative to a benchmark should maximise the risk-adjusted performance of the portfolio through optimal allocations to active, passive and risk-free alternatives, while ensuring that the standard deviation of returns is equal to that of the benchmark and the correlation of the returns is equal to the target correlation.²⁰ While this in effect leads to the highest possible information ratio on the overall portfolio, it does not require the highest information ratio on the active component as suggested by previous research. This calls for a

paradigm shift and a more open approach to optimal portfolio construction.

Acknowledgments

My debt of gratitude to the late Prof. Franco Modigliani is immeasurable for the hours spent on discussing this topic. My thanks also to Sanjay Muralidhar, Rahul Rauniyar, Srinivas Bette and Arjan Berkelaar for helpful discussions and finding errors in my work. Thanks also to Leah Modigliani, David Buckle, Ronald van der Wouden, Kenneth Miranda, Barton Waring and Larry Seigel. All errors are mine.

References and Notes

- 1 Grinold, R. C. and Kahn, R. N. (1999) 'Active Portfolio Management: A Quantitative Approach for Providing Superior Returns and Controlling Risk', 2nd edn, McGraw-Hill, New York, NY.
- 2 The path-breaking work referred to in all the research is Grinold and Kahn (ref. 1). The current papers use this book as the starting point for their discussions. See, for example, Blitz, D. C. and Hottinga, J. (2001) 'Tracking Error Allocation', *Journal of Portfolio Management*, Vol. 27, pp. 19–33. Mashayekhi-Beschloss and Muralidhar (Mashayekhi-Beschloss, A. and Muralidhar, A. (1996) 'Managing the Implementation Risks of a Currency Overlay Program', *Journal of Pension Plan Investing*, Vol. 1, No. 3, pp. 79–93) make this assumption in the context of implementing currency management strategies. I should like to thank David Buckle for highlighting this distinction between the theoretical assumption of Grinold and Kahn (ref. 1) and the practical application of the concepts, though it does not change the results of this paper.
- 3 Where the risk-free rate does not have zero variance, the denominator is the standard deviation of the portfolio.
- 4 Thanks to Ronald van der Wouden for this point.
- 5 Modigliani, F. and Modigliani, L. (1997) 'Risk-Adjusted Performance', *Journal of Portfolio Management*, Vol. 23, No. 2, pp. 45–54.
- 6 Muralidhar, A. (2000) 'Risk-Adjusted Performance — The Correlation Correction', *Financial Analysts Journal*, Vol. 56, No. 5, pp. 63–71.
- 7 Muralidhar, A. (2002) 'Skill, Horizon and Risk-Adjusted Performance', *Journal of Performance Measurement*, Vol. 48, pp. 28–43.

- 8 Many papers adjust for this by allowing for a subjective allocation of confidence in skill (Black, F and Litterman, R. (1992) ‘Global Portfolio Optimization’, *Financial Analysts Journal*, September/October, pp. 28–43) or by using other measures to convey such confidence (Lam, K. and Lee, W. (2001) ‘Implementing Optimal Risk Budgets’, *Journal of Portfolio Management*, Vol. 28, No. 1, pp. 73–80). Muralidhar⁷ uses an explicit measure of skill to create a time-adjusted ‘expected’ risk-adjusted performance measure.
- 9 The author’s experience as a plan sponsor and an asset manager has shown that no guideline ever prevented an active strategy from using cash or the passive benchmark to improve risk-adjusted performance. It is just a bad assumption made by active managers who focused on the wrong objective function, namely, to maximise the information ratio. In fact, often there is no constraint on the volatility of the manager and it is the case that the manager is not penalised for ‘leverage.’
- 10 Graham, J. R. and Harvey, C. R. (1997) ‘Grading the Performance of Market Timing Newsletters’, *Financial Analysts Journal*, Vol. 53, No. 6, pp. 54–66.
- 11 In addition, when benchmark returns are negative (eg in currency mandates), the M^2 measure will incorrectly rank underperforming funds with high volatility as preferred to funds with low volatility. This is a quirk of the method, as it is generally assumed that benchmarks have positive returns.
- 12 These are heroic assumptions to say the least. Some forecast needs to be made on expected outperformance, variability of performance to achieve this outperformance and correlations between portfolio and benchmark returns. Historical performance is one way of making forecasts, but the M^3 measure is independent of the forecasting technique. In addition, one must believe that markets are inefficient at conducting such analyses.
- 13 Hammond, D. (1997) ‘Establishing Performance-Related Termination Thresholds for Investment Management’, in Fabozzi, F. (ed.), ‘Pension Fund Investment Management’, Frank Fabozzi Associates, New Hope, PA.
- 14 Litterman, R., Longestaey, J., Rosengarten, J. and Winkelmann, K. (2001) ‘The Green Zone ... assessing the Quality of Returns’, *Journal of Performance Measurement*, Vol. 5, No. 3, pp. 29–50.
- 15 This measure is independent of the level of tracking error and hence is applicable across all tracking error targets.
- 16 This may change with the development of exchange-traded funds (ETFs) on active portfolios. In some cases, it may be difficult to short the benchmark as well, and then b will need to be constrained to being greater than, or equal to, zero. This would not change the analysis of the measure. In general though, most benchmarks can be shorted either through their futures contract or through a swap.
- 17 Muralidhar, A. (2001) ‘Optimal Risk-Adjusted Portfolios with Multiple Managers’, *Journal of Portfolio Management*, Vol. 27, No. 3, pp. 97–104.
- 18 Muralidhar, A. (2001) ‘Innovations in Pension Fund Management’, Stanford University Press, Stanford, CA.
- 19 See also Waring, B. and Castile, C. (1998) ‘A Framework for Optimal Manager Structure’, *Barclays Global Insight*, Vol. 1, No. 2, for the simple approach of maximising multi-manager returns subject to a risk budget.
- 20 There are issues about the stability of correlations and standard deviations that have been addressed in previous research and, since the goal of this paper is to change the paradigm first, these issues are not addressed again in this research.

APPENDIX: DEFINING TRACKING ERROR AND INFORMATION RATIO

Define 1 as the active manager or asset allocation strategy, B as the benchmark, r denotes the return, σ denotes the standard deviation and ρ is the correlation of returns between the manager and the benchmark.

If the tracking error of portfolio 1 versus the benchmark is defined as the standard deviation of excess returns, it is trivial to define TE(1) as follows

$$TE(1) = \sqrt{(\sigma_1^2 - 2\rho\sigma_1\sigma_B + \sigma_B^2)} \quad (A1)$$

The annualised information ratio is equal to the annualised excess return divided by the annualised tracking error, or alternatively,

$$IR(1) = [r(1) - r(B)]/TE(1) \quad (A2)$$