3 Revealing the information content of investment decisions

Noriyuki Okuyama and Gavin Francis

INTRODUCTION

Conventional performance measurement methods concentrate on investment outcomes rather than the underlying investment process. This chapter examines the effectiveness of the investment process by considering the essential part of any investment strategy: the investment decision. As recognized by Kahneman and Tversky (1979), the essential first step is to decompose returns into gains and losses.

We introduce tools to indentify the asymmetries in investment returns available to passive and active investors. The approach fits naturally into an enhanced risk budgeting framework for more effective portfolio construction.

INVESTMENT DECISIONS

An investment decision results in a change of exposure in an underlying portfolio, exchanging one stream of returns for another. Since this must be based on an assessment of relative performance, we define an investment exposure as a zero size long/short portfolio, holding a position in a risky asset against an equal and opposing position in a risk-free asset. A risk-free asset earns the risk-free rate, but cannot suffer a loss. This is the numeraire against which all other returns are measured. Each currency's cash rate is considered as the risk-free rate for the assets denominated in that currency. We can broaden the scope for a long/short position created by two assets denominated in different currencies, each considered as a risk exposure against the local cash rate. This approach incorporates the difference of the two currencies' cash rates: the interest rate differential.

It is widely recognized in the field of behavioral finance that the basis of an investment decision is an assessment of the trade-off between gain and loss. Ultimately, all investment decisions are binary: to act or not to act, that is the question. Having implemented the decision, the outcome can be success or error. These outcomes are summarized in Figure 3.1.

The four quadrants of the matrix represent the principal outcomes of an investment decision. Investment success can be either the capture of gains, a true positive, or the avoidance of loss, a true negative. Investment error occurs either by incurring a loss, a false positive or 'type I error' or as a result of missing out on a gain, a false negative or 'type II error.' See Luce and Raiffa (1989) for a general discussion of decision theory.

This simple matrix demonstrates that there are two sides to investment error: holding

		Actual outcome (*ex post*)	
		Gain	Loss
Investment decision (*ex ante*)	Take exposure (visible)	Capture gains (true positive) success	Incur losses (false positive) type I error
	Avoid exposure (invisible)	Give up gains (false negative) type II error	Avoid losses (true negative) success

Figure 3.1 Fundamental investment decision matrix

an asset when it falls in value or holding cash when the asset is rising. Investment error is defined as any uncertainty that contributes negatively to an investor's potential wealth. This decomposition of risk demonstrates that a full evaluation of investment error must include both false positives and false negatives. Following this line of reasoning, an investment success is any uncertainty that contributes positively to returns. There are two ways to improve returns: by capturing gains through true positives or by avoiding losses through true negatives. Making a 5 percent gain and avoiding a 5 percent loss both have an equal effect on an investor's wealth.

This chapter proposes that active investment skill lies in the ability to maximize investment success relative to investment error. When making an assessment of the historic performance of an investment decision maker, this is superior to measuring return versus volatility.

However, one should ask why traditional approaches to risk management, such as mean-variance analysis and value at risk, focus on only one of the two sources of investment error: the risk of incurring losses. The answer lies in the fact that captured gains and incurred losses are both 'visible': they are the returns observed in the portfolio, whereas the losses one avoids and the gains given up are both 'invisible'. This means that it is normally hard to evaluate the costs of false negatives and true negatives. We can define this visibility as the degree to which the investor sees *ex post* gains and losses in the portfolio. Any investment decision contains a trade-off between visibility and invisibility of its investment outcome, which is determined by the size of the risk exposure.

Investment Error versus Volatility

Several consequences ensue from equating investment risk with investment errors (both false positives and false negatives). It becomes highly questionable to consider volatility as a proper measure of investment risk, since it simply looks at a degree of deviation from an achieved return on both negative and positive sides. In other words, the traditional mean-variance framework assumes that any uncertainty in possible outcomes should be considered as risk. This is a typical heuristic approach in many of the performance metrics that have emerged from the framework of modern portfolio theory that fails to

		Actual outcome (*ex post*)	
		Gain	Loss
Investment decision (*ex ante*)	Take exposure (visible)	Higher volatility (true positive) success	Higher volatility (false positive) type I error
	Avoid exposure (invisible)	Lower volatility (false negative) type II error	Lower volatility (true negative) success

Figure 3.2 Shortcomings in the use of volatility

distinguish between loss and gain. It ranks different returns according to their degrees of uncertainty without disentangling constituent factors behind: losses and gains are mixed together.

Consider the Sharpe ratio, described by Sharpe (1994). Losses and gains are mixed to obtain an average return. This is divided by the volatility of returns, which would only be a useful measure of loss if the returns were described by a normal distribution with a mean of zero. The Sharpe ratio, like the information ratio, measures a signal-to-noise ratio, but it does not measure the balance between investment success and investment error. Noise contains both gains and losses. A signal-to-noise ratio can be meaningful only if perpetually matching a target outcome is the top priority. But in investment, exceeding the target is a benefit. A positive outcome (even if due to luck) is always welcome and should not be penalized. This asymmetry in investment utility cannot be captured by return and volatility alone.

From the perspective of the decision matrix, the uncertainty of an achieved return should be decomposed into the uncertainty of positive returns and the uncertainty of negative returns. This suggests that a larger negative return and a smaller positive return both represent real investment errors to the same extent. However, while larger uncertainties in negative returns will increase overall volatility, smaller uncertainties in positive returns reduce overall volatility. Even though both cases are investment errors, the impact on volatility is opposite.

In contrast, the chances of smaller negative returns or larger positive returns both constitute investment success. Yet these also have opposite effects on volatility. This means that volatility can properly represent neither investment errors nor investment success.

Figure 3.2 shows that volatility makes an incorrect assessment of realized gains. In order to analyze investment error/success appropriately, one needs to go to the trouble of decomposing returns into negatives (losses relative to the risk-free portfolio) and positives (or gains). Unlike the mean-variance 'two-tails-combined' approach, this risk decomposition approach enables us to reveal to what extent either skill in reducing losses or skill in capturing gains has contributed to a total return.

Aware of the shortcomings in the use of volatility, Markowitz (1959) proposed the use of a downside risk measure. The Sortino ratio (Sortino and Price Lee, 1994) looks at negative observations rather than volatility, but the numerator of the ratio remains a mixture of losses and gains. Although this is an improvement on volatility, because it does not penalize gains, it is not an effective measure of investment error because it still ignores false negatives. The weakness in the conventional approach is its focus on the final outcome, not on the decision-making process that generates it.

Passive and Active Management

A passive investment management strategy results from a single investment decision to match a pre-defined benchmark payoff. For example, a static buy-and-hold approach makes no attempt to change exposures over time. Since no further investment decisions are required, this is a matter of implementation rather than management. It is almost a tautology to say that a pure passive strategy has no expected return relative to its benchmark (before costs).

An active manager, however, makes decisions to change the risk exposure through time, so exposures are not held statically. The key distinction is that a successful active manager uses information to enhance the balance between losses and gains, whereas a passive manager does not enhance this balance. Biglova et al. (2004) describe the Rachev ratio, which compares the probabilities of extreme gains against extreme losses by defining an upper and a lower threshold. When both of these thresholds are aligned to zero return, we move closer to a useful measure of loss versus gain.

In the long term, however, the return on an investment is the difference, not the ratio, between gain and loss. Utility theory has been used to attempt to disentangle the impacts of losses and gains from inputs that mix them, such as return, volatility and correlation. Superior portfolios would be available to investors who could directly measure the loss and gain of investment processes. Fishburn (1977) defined the lower partial moments of a distribution relative to a threshold. This chapter extends that approach by measuring the upper partial moments of a distribution and comparing these against the lower partial moments. This is called conditional risk attribution (CRA).

FORMAL DEFINITIONS OF PARTIAL MOMENTS

The starting point is to measure the log returns[1] of the outcomes of the investment process relative to the appropriate risk-free rate.[2] This determines a set of n outcomes, x_i. CRA segments the set of observations into losses (the negative returns) and gains (the positive returns). As a result, rather than analyzing a single distribution of outcomes, CRA analyses it in two parts: the distribution of losses and the distribution of gains (see Figure 3.3).

Clearly neither losses (L) nor gains (G) are distributed according to a Gaussian curve. Both distributions are bounded at zero and their modes would normally lie close to the origin. The frequency distribution of losses may be characterized by the lower partial moments: the averages of L, L^2, L^3, L^4 etc., and similarly for the distribution of gains. Recognizing that these distributions were originally derived from a single distribution of n outcomes, the moments are scaled according to the total number of original observations, n.

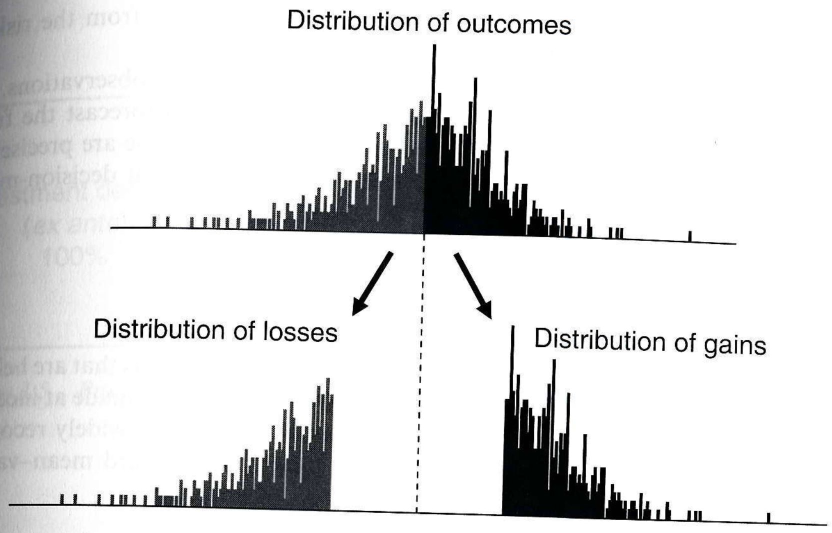

Figure 3.3 *Disentangling gains from losses*

	Lower partial moment	Upper partial moment
First	$\overline{L} = \frac{1}{n}\sum_{x_i<0} x_i$	$\overline{G} = \frac{1}{n}\sum_{x_i\geq 0} x_i$
Second	$\overline{L^2} = \frac{1}{n}\sum_{x_i<0} x_i^2$	$\overline{G^2} = \frac{1}{n}\sum_{x_i\geq 0} x_i^2$
Kth	$\overline{L^k} = \frac{1}{n}\sum_{x_i<0} x_i^k$	$\overline{G^k} = \frac{1}{n}\sum_{x_i\geq 0} x_i^k$

Figure 3.4 *Definitions of partial moments*

This leads to the formal definitions of partial moments[3] (see Figure 3.4). The general definition allows any (integer or non-integer) partial moment to be calculated. The 'zeroeth' moments define the fractions of observations that were losses or gains: the proportions of so-called 'down months' and 'up months', for a monthly track record. The first partial moments represents average loss and gain, while the second partial moments represents dispersion of loss and gain, with the third partial moments being more strongly influenced by the tails. If the distribution is asymmetric, the segregation of losses and gains extracts more information from the frequency distribution than simply taking the average.

The second moment, $\overline{L^2}$, is a measure of the dispersion of losses. Here we start to diverge more significantly from the traditional approach of calculating the variance of the returns. The variance is a measure of the average squared deviation from the mean observation. But since CRA is concerned with losses and gains, the second partial moment is based on the average squared deviation from zero (i.e. from the risk-free return).

Higher moments are more strongly influenced by the extreme observations. This higher moment effect can be a disadvantage when attempting to forecast the future. However, from the perspective of historic performance analysis, these are precisely the observations that are the most important in assessing the investment decision-making process.

CONDITIONAL RISK ATTRIBUTION

Conditional risk attribution (CRA) is a tool for comparing investments that are held passively. A passive strategy is the result of a single investment decision, made at inception. No attempt is made to change the exposure through time. It is not widely recognized that this assumption of passivity is implicit in the output of standard mean–variance optimizers.

Taking Risk Passively in a Symmetric World

We begin by applying the fundamental investment decision matrix (Figure 3.1) to a simplified, constrained scenario. If we relax the constraints, the scenario becomes more representative of the real world. This leads to the concept of CRA.

Consider a world in which there is a risk-free asset and a single risky asset that generates symmetric returns with 50 percent gains and 50 percent losses. The passive investor has the choice of statically holding one asset or the other.

We would normally expect the investor to begin with the risk-free asset. In this case, the decision to switch to the risky asset would be a risk-taking decision. This approach successfully captures all gains, but the portfolio incurs the full impact of losses.

Figure 3.5 shows that false positives are expected to occur 50 percent of the time, i.e. type I errors are maximized. Note that when the risky asset is held, losses and gains are fully visible as portfolio returns.

Compare this with holding the risk-free asset, shown in Figure 3.6, where the decision threshold has shifted. This successfully avoids all losses, but the portfolio is deprived of ever making a gain. As a result, false negatives, i.e. type II errors, are maximized. The effect is to incur type II error rather than type I error.

This fact is not obvious to investors who consider only the volatility of observed outcomes, because both the missed gains and losses are invisible, in the sense that they do not show up as portfolio returns. Behavioral bias tends to prefer visibility.

We now relax one of the constraints and allow the investor to be partially invested in the risky asset. Clearly, this is simply an intermediate between the two extremes. For example, with a 50 percent exposure position to the risky asset, one would expect to see a quarter of observations in each quadrant, resulting in 25 percent of type I error and 25 percent of type II error. Note from Figure 3.7 that the passive approach trades off type I errors for type II errors on a one-for-one basis, so total investment error remains at 50 percent.

The level of a passive exposure determines the horizontal decision threshold between

		Actual outcome (*ex post*)	
		Gain	Loss
Investment decision (*ex ante*) 100%	Take exposure 100% visible	Capture gains (true positive) 50% success	Incur losses (false positive) 50% type I error
	Avoid exposure 0% invisible		

Figure 3.5 *Passively holding a risk exposure*

		Actual outcome (*ex post*)	
		Gain	Loss
Investment decision (*ex ante*) 100%	Take exposure 0% visible		
	Avoid exposure 100% invisible	Give up gains (false negative) 50% type II error	Avoid losses (true negative) 50% success

Figure 3.6 *Passively holding a risk-free exposure*

		Actual outcome (*ex post*)	
		Gain	Loss
Investment decision (*ex ante*) 100%	Take exposure 50% visible	Capture gains (true positive) 25% success	Incur losses (false positive) 25% type I error
	Avoid exposure 50% invisible	Give up gains (false negative) 25% type II error	Avoid losses (true negative) 25% success

Figure 3.7 *Passively holding a partial risk exposure*

visibility and invisibility. The threshold in the box simply slides up or down in parallel, depending on the choice of benchmark static exposure.

Any reduction of incurred losses is accompanied by a reduction of captured gains. Incurring a loss (type I error) and missing a gain (type II error) have exactly the same negative impact on expe return. In the simplified case, where risk exposures that have an equal chance of risin falling, we conclude that:

(a) all passive exposur ave the same expected return;
(b) different passive b mark selections represent different utilities on visibility of outcomes (differen ference/aversion to regret);
(c) all passive exposu ncur exactly the same amount of investment error!

Reducing Risk Passively Symmetric World

In the above, it was assu that the investor started with the risk-free asset. If, however, an investor is already percent invested in the risky asset, it is possible to take a risk-reducing decision. can be achieved most simply by cutting back exposure in exchange for the risk-fr set. Alternatively, risk can be reduced by hedging. This is done by adding a risk sure that acts as an offset, by reliably generating returns opposite to the risky as eld in the portfolio. In order to be effective, a hedge must have a correlation relia lose to −1, in all market environments. The magnitude of the hedge will determine ther it partially or fully offsets the other risk exposure. The degree of offset is called hedge ratio, which, by definition, can lie only between zero and one. Hedges can be naged either passively or actively. A hedge has the special property of reducing the of the maximum potential loss of the portfolio. So hedging is a risk-reducing activit

As an example, we as e that an investor has currency exposure in a portfolio of international assets. A p e currency hedging strategy consistently maintains a predetermined benchmark he ratio throughout an investment horizon, regardless of any subsequent subjective vi There is no attempt to distinguish dynamically between currency gain environme nd currency loss environments over time. As we concluded earlier, for a risky asset symmetric returns, any level of passive hedge results in 50 percent investment error

Identifying Asymmetry in sive Investments

Most market participan cognize that financial assets exhibit asymmetric returns. An investment return is hematically equivalent to the excess of gains over losses. The entire rationale for ive investment therefore relies on finding risky assets with asymmetric returns. In t ay, investors can benefit from investment success of more than 50 percent (and inve nt error of less than 50 percent) without making any active investment decisions. In the active decisions are delegated. For example, a passive investment in the S&P 50 lex delegates active decisions to the managers of the largest 500 US companies and t set of rules used to construct the index.

CRA is designed to n re the asymmetries or 'market bias' between the partial moments of a distributio returns. Often a positive asymmetry in the first moment (a positive return) is associ with negative asymmetries in the higher moments (a large negative tail).

In order to create the damental investment decision matrix (Figure 3.1), it was necessary to decompose ns into losses and gains. This matrix can be used to analyze

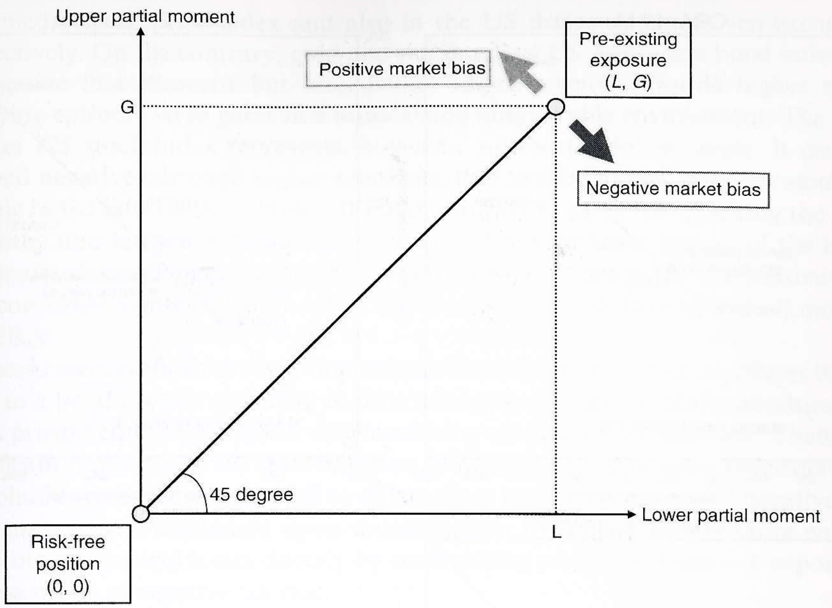

Figure 3.8 Conditional risk attribution

binary decisions, such as the success or failure of a medical test. However, investment outcomes can be quantified in numerical terms. So we must move from a discrete framework to a continuum. The underlying principle of CRA is to compare the result of passive management against a set of perfectly symmetric outcomes. Having calculated the upper and lower partial moments, conditional upon the time period of the observations, we plot each pair in a CRA diagram. It is convenient to rescale each moment to the same dimension, i.e. to plot $(\overline{G^k})^{1/k}$ versus $(\overline{L^k})^{1/k}$, for $k > 0$. The vertical axis represents the upper partial moment, while horizontal axis represents the corresponding lower partial moment (see Figure 3.8).

Now, if the evolution of the underlying risk exposure were determined by the toss of a coin, it would follow a random walk. In this case, we would expect that in the short term the upper partial moment of a pre-existing exposure G would deviate from the lower partial moment L. However, the long-run expectation would be for G to match L, so that the CRA diagram becomes square.

Holding a partial exposure statically (or passively hedging) reduces risk by avoiding both losses and gains, without employing any active information to discriminate between them. Therefore the impact on the upper and lower partial moments is identical. This one-for-one ratio gives the slope of a diagonal line between the pre-existing exposure (L, G) and the risk-free position (0, 0). Under the random walk assumption, the slope would be 45 degrees. Any deviations from the 45 degree passive line indicate either positive or negative market bias.

When investing in financial assets, we hope to identify clues from the higher moments of the distribution to give us confidence that a positive deviation in the first moment (i.e. positive return) is sustainable and not just a short-term trend in a random walk.

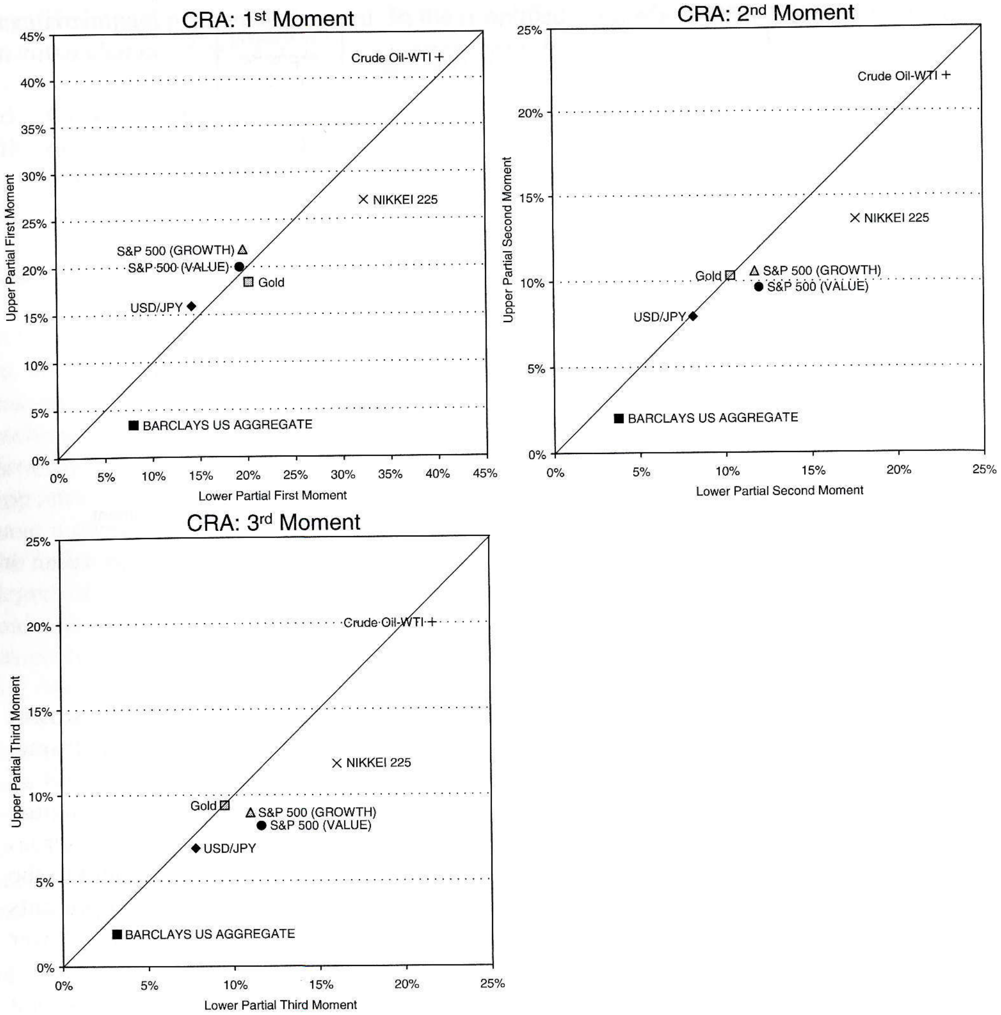

Source: DataStream.

Figure 3.9 CRA examples of different moments

Empirical Examples

With historical price data from January 1988 to July 2009, we calculated both upper and lower partial moments (first, second and third moments) for various investments (see Figure 3.9).[4] Looking at the S&P 500 composite data, both growth and value sectors showed a positive first moment but an incremental deterioration towards higher moments, implying episodic large losses over the long-term investment horizon. This is consistent with an overall distribution of outcomes possessing a positive mean, but a pronounced negative tail. The same observation was identified in the WTI (West Texas Intermediate) oil price index and also in the US dollar–Japanese yen exchange rate respectively. On the contrary, gold and the Barclays US Aggregate bond index showed a negative first moment but incremental improvements towards higher moments, implying episodic large gains in a long-lasting unfavorable environment. The Japanese Nikkei 225 stock index represents, however, an unorthodox example. It consistently showed negatives through higher moments, due to the fall-out from the stock market bubble in the late 1980s in Japan. It burst abruptly in early 1990, sinking the Japanese economy into ten years of asset depreciation. The significant impact of the initial fall dominates through higher moments. Nevertheless, it is rare to see in a passive exposure any consistent significant bias, either positive or negative, throughout all moments in the CRA.

These considerations are very important when searching for risk exposures to buy and hold in a portfolio, for example, in determining a strategic portfolio allocation, selecting a private equity investment or when hiring an investment manager. There are also significant implications for diversification in portfolio construction. Investors who rely on volatility as a risk measure fail to differentiate between positive and negative tail risk. This makes them dependent upon diversification to control losses. More enlightened investors can control losses directly by constructing portfolios from risk exposures that are less prone to negative tail risk.

GENERALIZED CONDITIONAL RISK ATTRIBUTION

Generalized conditional risk attribution (GCRA) is a tool to indentify active investment decision-making skill. We have argued that successful investment decisions achieve greater investment success than investment error. It is possible to achieve this passively if the underlying investments consistently exhibit asymmetry between gain and loss. Active management involves taking decisions through time. A skillful active manager should be able to introduce additional positive asymmetry into a distribution of returns, over and above that available to a passive investment.

Maximum Permitted Active Risk Exposure

We now turn our attention to the assessment of active investment processes. We show that a skillful active manager is able to vary the size (visibility) of a risk exposure through time in order to improve the trade-off of gain versus loss. While the purpose of CRA was to compare passive investments against a symmetric random walk, the objective of GCRA is to compare the result of active management against passive management. This requires a definition of the range of outcomes that could have been available from a static exposure, determined by the maximum permitted active risk exposure. In the context of unleveraged, long-only investment, this is well defined as the size of the initial amount of capital.

The concept of a maximum permitted active risk exposure is crucial to the analysis. This is what defines the potential passive returns, against which the active returns should be compared. The maximum permitted exposure determines the 'size of the box' in the CRA diagram (Figure 3.8). In the case of a hedging mandate, the maximum exposure equals the underlying exposure. Many active mandates fail to specify maximum permitted exposures, relying on less precise concepts, such as *ex ante* volatility and notional portfolio sizes. This omission makes it very difficult to assess the decision-making skill of the manager.

Normalization of Market Contribution to Asymmetry

The first step is to eliminate any deviations from the 45 degree line that were available to passive management. These deviations, which make G no longer equal to L, do not contain any active information; they simply represent market bias over the time period of observation that creates episodic asymmetry in gains and losses. Clearly, in assessing the skill of an active manager, this market bias needs to be removed. This process is called 'normalization'.

We first calculate the degree of market bias by taking the ratio of the upper and lower partial moments.

$$\text{Market bias} = \frac{G}{L}$$

This ratio represents the market-biased slope of the passive line, which can be replicated by holding a static exposure. Normalization is a process to reset to zero any incremental deviations caused by market bias. This adjustment is a perpendicular shift from the market-biased passive line to the 45 degree market-neutral passive line. Although the 'market bias' ratio is indifferent to any exposure level, the degree of asymmetry is different at each exposure level. The perpendicular distance from the 45-degree line is a quantity to be normalized. It is called 'market contribution to asymmetry' (MCA) (see Figure 3.10).

The MCA is determined for all passive benchmark exposure levels, but not for the outcomes of active management. The appropriate shift for active outcomes should always refer to the MCA normalization corresponding to the benchmark exposure against which the active manager is measured. This preserves the displacement of the active manager's outcome relative to the benchmark.[5] Normalization removes the market bias, revealing a genuine active skill that is the remaining perpendicular distance from the 45-degree normalized passive line.

After normalization, the CRA diagram becomes a square and therefore symmetric. However, the side of the square will be equal to the average of the upper and lower partial moments of the maximum permitted risk exposure. The outcome of the actively managed portfolio must lie within this square. We take one further step to generalize the approach by rescaling the axes to one unit. This is achieved by taking the coordinates of every point on the diagram and dividing them by the length of the side of the square. Finally we arrive at a unit square that permits a comparison between different managers managing different investment exposures over different time periods. We call this active information diagram the generalized conditional risk attribution (GCRA) (see Figure 3.11).

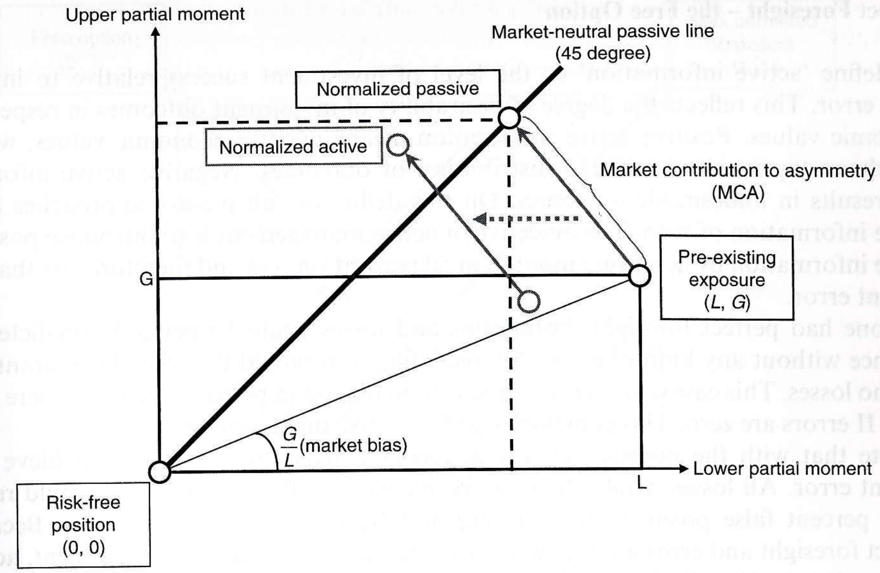

Figure 3.10 Normalizing market contribution to asymmetry

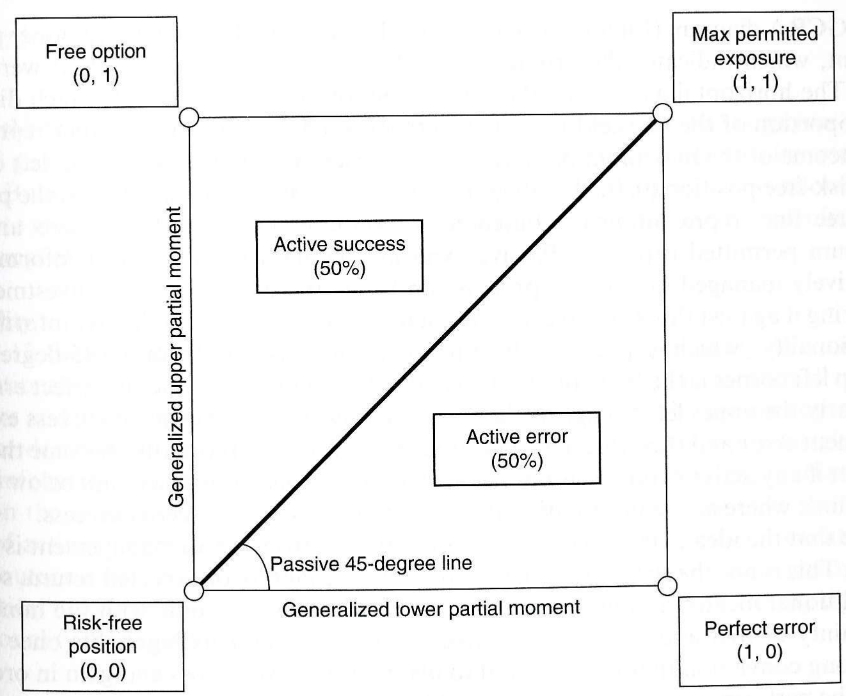

Figure 3.11 Generalized conditional risk attribution

Perfect Foresight – the Free Option

We define 'active information' as the level of investment success relative to investment error. This reflects the degree of desirability of investment outcomes in respect of economic values. Positive active information means positive economic values, which introduces 'optionality' into the distribution of outcomes. Negative active information results in undesirable outcomes. On this definition, all passive approaches have active information of zero. The objective of active management is to introduce positive active information by achieving more than 50 percent success and therefore less than 50 percent error.

If one had perfect foresight, both gains and losses could be perfectly predicted in advance without any kind of error. All available gains would therefore be guaranteed, with no losses. This case is known as a free-option pay-off (a perfect outcome) where type I and II errors are zero. This constitutes perfect active management.

Note that with the poorest active management, it is also possible to achieve 100 percent error. All losses would therefore be incurred, with no gains. This would result in 50 percent false positives and 50 percent false negatives (perfect error). Because perfect foresight and error are the two most extreme cases of active management, active information ('optionality') should be constrained by these two cases.

Optionality and Visibility

In the GCRA diagram (Figure 3.11), the vertical axis shows the generalized upper partial moment, which indicates the proportion of the available market gains that were captured. The horizontal axis shows the generalized lower partial moment, which displays the proportion of the market losses that were incurred. The top right corner represents the outcome of the maximum permitted risk exposure (1, 1) and the bottom left corner is the risk-free position (0, 0). The diagonal line connecting (0, 0) and (1, 1) is the passive 45-degree line, representing the outcomes of static exposures between zero and the maximum permitted exposure. Passive exposures contain no investment information. An actively managed investment process can be compared with passive investment by measuring it against this 45 degree line. We hereby define the degree of active information as 'optionality', which is quantified by a perpendicular distance from the 45-degree line. The top left corner is the free option (0, 1) and the bottom right corner is perfect error (1, 0). Clearly the upper left triangle is the desirable area, where investment success exceeds investment error and therefore active information is positive. It can thus become the area of regret if any active exposure is not taken. The area of disappointment lies below the 45 degree line, where active information is negative, because error exceeds success.

Note that the ideal, 100 percent successful outcome from active management is a free option. This is not the outcome with the minimum variability of expected return, sought by traditional mean–variance approaches. In fact it is the outcome with the minimum uncertainty of loss and the *maximum* uncertainty of gain! This highlights once again our strong conviction that it is essential to distinguish between loss and gain in order to assess the performance of an active manager.

It is no coincidence that Figure 3.11 resembles the receiver operating characteristic (ROC) diagrams used in decision theory (see Marques de Sa, 2001). An ROC diagram

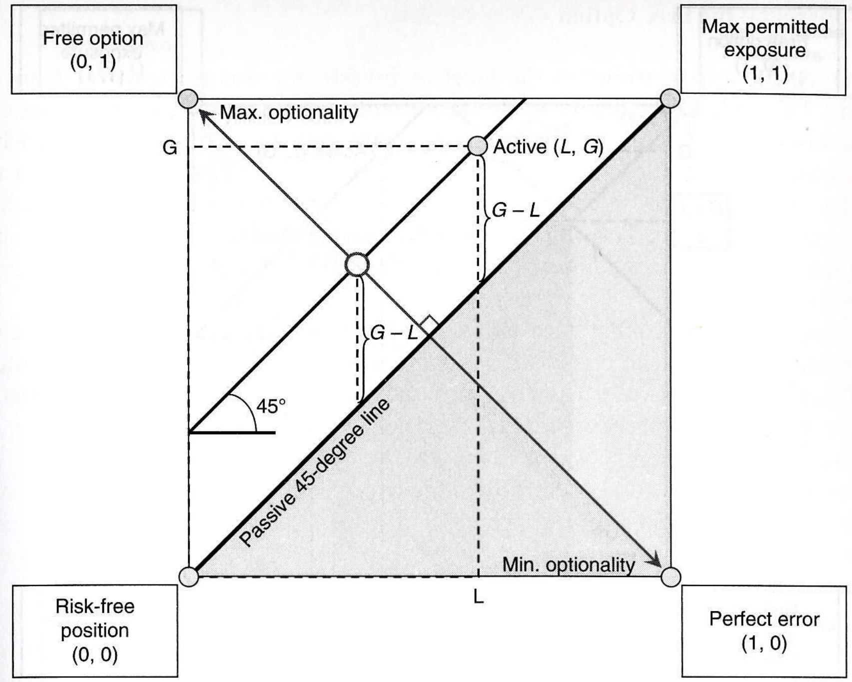

Figure 3.12 Definition of optionality

displays the trade-off between the proportions of false positives and false negatives. The 'zeroeth' generalized partial moments show exactly this information for an investment process. But GCRA goes beyond this, by plotting the higher moments. This allows the investor to measure not just the percentage of investment errors, but also the significance of their impact.

Optionality

Figure 3.12 shows that the free option represents the maximum optionality. We use the term 'optionality' because the ideal objective of an active manager is to capture all available gains and avoid all potential losses. Such outcomes are offered by option contracts, at a price determined by the market. Genuine investment skill lies in using insights or private information to create option-like outcomes more cheaply than the market. All points on the same line running parallel with the passive 45-degree line have the same degree of active information, i.e. given that the generalized upper partial moment of a specific outcome is *G* and the generalized lower partial moment of the same outcome is *L*, a vertical distance (G – L) is constant for all points on the line. This quantity defines the optionality demonstrated by active management.

Visibility

The degree to which the investor sees the impact of the actively managed exposure in the partial moments of the portfolio returns is called 'visibility', defined as the sum of

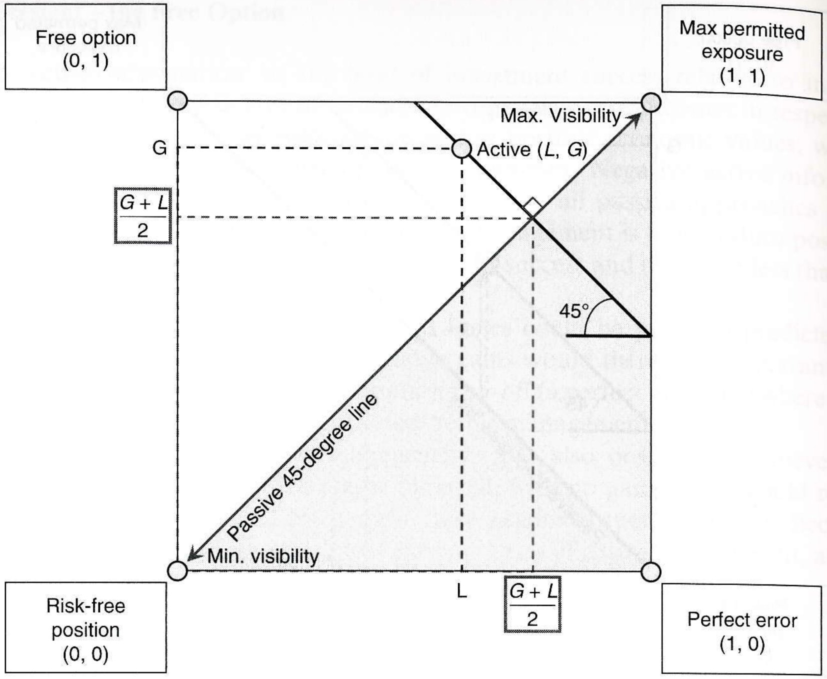

Figure 3.13 Definition of visibility

exposures to both the upper and lower generalized partial moments ($G + L$). Visibility is zero at the risk-free position and the distance from the risk-free position to each line running perpendicular to the passive 45 degree line represents incremental visibility. Figure 3.13 shows that the generalized maximum permitted exposure in the top right corner represents the maximum visibility. Visibility itself does not represent active information, as it is always measured on the passive 45-degree line, where no active information is deployed. All decision making on visibility is therefore a passive judgment.

The investor's utility on visibility pre-determined by an initial benchmark selection should distinguish investment outcomes which contain the same economic value (see Figure 3.14).

Outcome A would therefore be appropriate for an investor who is cautious about incurring loss and therefore willing to be exposed to potentially large regret of missing out on gains. This type of regret-tolerant investor is an invisibility seeker, trying to reduce regret through active manager's skill rather than a passive judgment.

Outcome B would suit an investor who is keen to capture the full extent of potential gains and minimize regret at the expense of being exposed to potentially large losses. This type of regret-averse investor is a visibility seeker, limiting regret by a passive judgment while trying to reduce losses through active manager's skill.

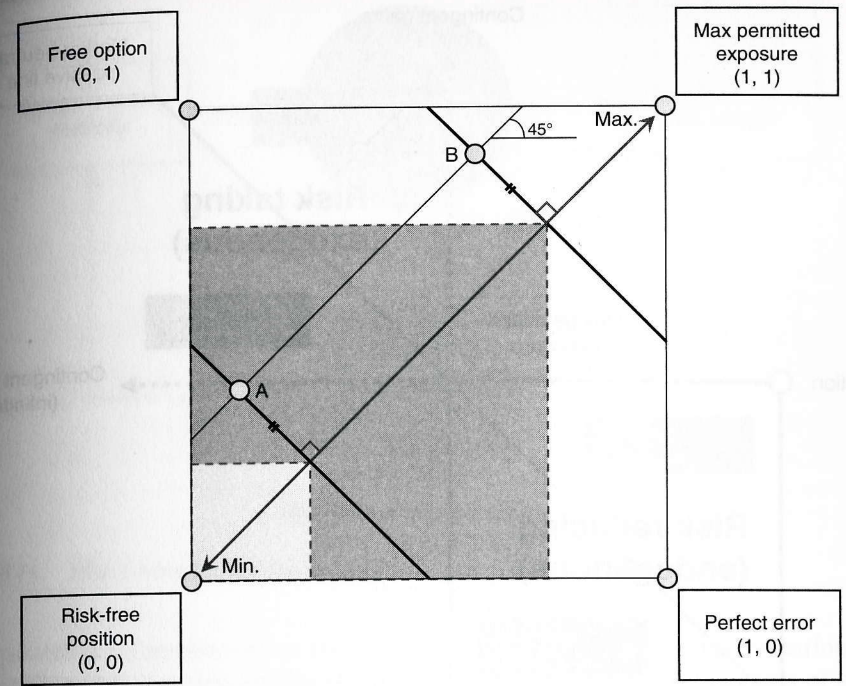

Figure 3.14 Invisibility seeker versus visibility seeker

RISK BUDGETING

Risk-reducing decisions and risk-taking decisions are different choices about different outcomes. Risk reduction considers only the exposures already held in the portfolio, whereas risk taking considers only exposures that are not held in the portfolio (see Figure 3.15).

The initial portfolio construction phase is a risk-taking activity, investing risk-free assets, such as cash, in simple long-only strategies in conventional asset classes, such as equities and bonds. This can be achieved either passively or actively within a pre-determined risk budget. Since a risk budget is always finite, the only way to achieve better risk-adjusted returns is to reduce risk where it is poorly rewarded. The role of risk-reducing decisions is therefore of fundamental importance to successful investment strategies. This risk-reducing activity can also be achieved either passively or actively.

The effect of reducing risk is to free up part of the risk budget so that an investor can reallocate it to better rewarded risk-taking activities. An investment process should always follow a route from risk taking through risk reducing to further risk taking. In other words, a risk-reducing activity is considered as a bridge from an initial long-only investment portfolio to a more efficient portfolio where long-only strategies, hedging and long-short strategies coexist.

The framework developed here proposes enhancing the wealth of the investor by

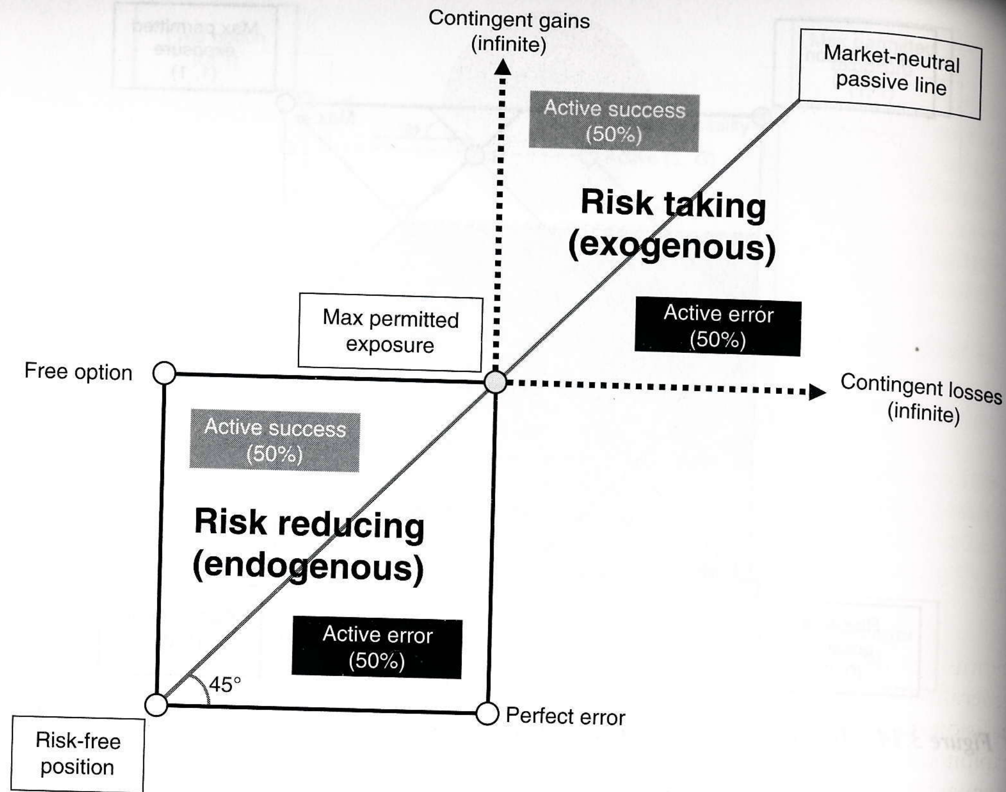

Figure 3.15 Geometric representation of risk reducing and risk taking

improving the balance of investment success to investment error. Whether an investor is reducing risk or increasing risk, the objective is the same: to exploit active information (optionality) in order to increase investment success. See Figure 3.16 for a clear depiction of the repeating cycle of risk re-budgeting.

The risk re-budgeting approach described above highlights the significant importance of distinguishing risk reducing from risk taking, so that one can achieve a better risk-adjusted return within a pre-determined risk budget. Without making a clear strategic distinction, neither the first process (a creation of budget) nor the second process (an enhancement of returns) will be implemented efficiently. Both will be achieved through increasing investment successes.

CONCLUSION

This chapter examined investment decisions using the fundamental investment decision matrix to define investment success and investment error. The framework was effective in describing passive investment. CRA is a tool for identifying the asymmetric

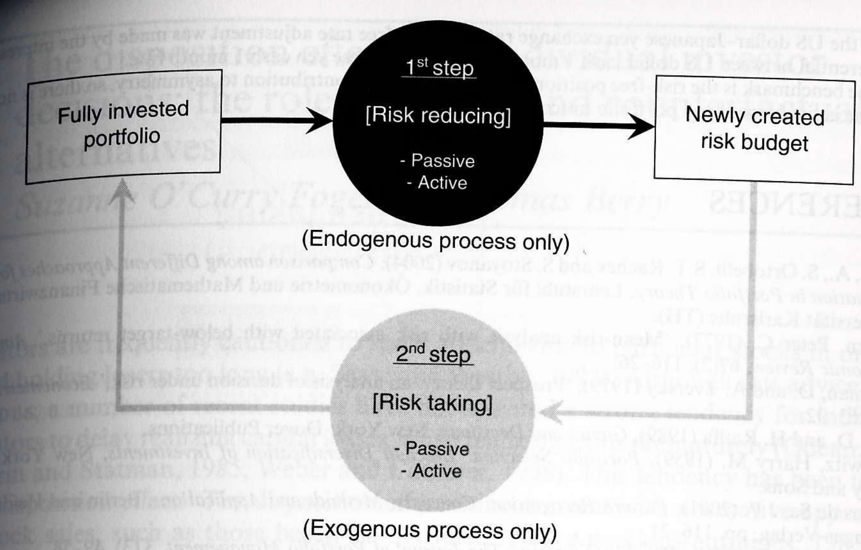

Figure 3.16 Risk re-budgeting framework

returns available to passive investors. We expanded this geometric interpretation into generalized conditional risk attribution (GCRA) to measure the effectiveness of active decision-making processes. Manager skill is associated with active information that is exploited by the investment process. GCRA avoids cognitive confusion between impacts from active and passive decision making, and conducts a heuristic approach to identify genuine investment contributions from active skill.

Modern portfolio theory assumes, as an axiom, that financial returns are symmetric. The tools presented in this chapter are designed to measure the degree of asymmetry (market bias and optionality) in investment returns. The skill of passive management is to find naturally occurring forms of market bias. The role of active managers is to create optionality.

This led to a clear distinction between risk-reducing and risk-taking decisions. By disentangling the cognitive bias inherent in the conventional mean–variance approach, risk budgeting in a GCRA framework allows the portfolio construction process to reflect a better assessment of investment decisions.

NOTES

1. Since our objective is to distinguish asymmetry from symmetry, it is important to avoid performance metrics that contain asymmetric properties, such as geometric returns, where a 25 percent gain is offset by a 20 percent loss.
2. Risk-free rate in the general sense defined by Sharpe (1994) is the risk-free benchmark portfolio against which the investment is evaluated.
3. For convenience, we adopt the convention that a zero return is considered a gain.
4. The risk-free rate used for the US dollar denominated investments (S&P 500 Value & Growth, Barclays US Aggregate bond index, WTI, gold) was JPMorgan US dollar cash 1 month rate, while the rate for the Japanese yen denominated investments (Nikkei 225) was the JPMorgan Japanese yen cash 1 month rate.

For the US dollar–Japanese yen exchange rate, the risk-free rate adjustment was made by the interest rate differential between US dollar cash 1 month rate and Japanese yen cash 1 month rate.

5. If the benchmark is the risk-free position, there is no market contribution to asymmetry, so there is no shift in the actively managed portfolio outcome.

REFERENCES

Biglova, A., S. Ortobelli, S.T. Rachev and S. Stoyanov (2004), *Comparison among Different Approaches for Risk Estimation in Portfolio Theory*, Lehrstuhl für Statistik, Ökonometrie und Mathematische Finanzwirtschaft, Universität Karlsruhe (TH).

Fishburn, Peter C. (1977), 'Mean-risk analysis with risk associated with below-target returns,' *American Economic Review*, **67**(2), 116–26.

Kahneman, D. and A. Tversky (1979), 'Prospect theory: an analysis of decision under risk,' *Econometrica*, **47** (2), 263–92.

Luce R.D. and H. Raiffa (1989), *Games and Decisions*, New York: Dover Publications.

Markowitz, Harry M. (1959), *Portfolio Selection: Efficient Diversification of Investments*, New York: John Wiley and Sons.

Marques de Sa, J.P. (2001), *Pattern Recognition: Concepts, Methods and Applications*, Berlin and Heidelberg: Springer-Verlag, pp. 116–21.

Sharpe, William F. (1994), 'The Sharpe ratio,' *The Journal of Portfolio Management*, **3** (3), 49–58.

Sortino, Frank A. and Lee N. Price (1994), 'Performance measurement in a downside risk framework,' *The Journal of Investing*, **3** (3), 59–64.