The Journal of Behavioral Finance
2006, Vol. 7, No. 2, 75–87

Disentangling Cognitive Bias in the Assessment of Investment Decisions: Derivation of Generalized Conditional Risk Attribution

Noriyuki Okuyama and Gavin Francis

Conventional performance measurement methods concentrate on investment outcomes rather than the underlying investment process. This paper 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. Using a fundamental investment decision matrix broadens the analysis beyond realized gains and losses to include an assessment of missed opportunities and losses avoided. This leads to an examination of investment strategies in terms of investment success versus error. We contrast these terms with traditional return and risk measures.

The second step is to decompose investment decisions into the mutually exclusive categories of risk-taking and risk-reducing. Uncertainty observed in risk-taking activities is newly created at each manager's discretion, while the uncertainty in risk-reducing activities is already defined by an investor's portfolio. The conventional performance measurement approach fails to differentiate between these two types of decisions.

The contrast between passive and active management styles of risk reduction becomes apparent from the perspective of investment success and error, particularly when using currency risk management as an illustration. Active information diagrams describe a geometric approach called conditional risk attribution. An active manager's skill lies in the ability to use information to maximize investment success and therefore minimize investment error. The choice of benchmark determines the degree of visibility of an investment exposure.

We also address the problem of distinguishing between an active manager's skill and the uncontrollable impact of market movements. Having adjusted for the bias in an investment environment, it becomes possible, as a final step, to measure the information content of an active manager's strategy as the balance between investment success and error within a generalized conditional risk attribution framework.

The unlimited nature of risk-taking decisions is evident when they are included in the framework. However, having defined an ex ante risk-taking limit, we can assess whether the choice of active manager has resulted in an enhancement of investment success over error. This allows investors to compare risk-taking managers with risk-reducing managers in the context of overall portfolio construction. The entire analysis is framed in terms of risk budgeting to demonstrate its general application to all asset classes and consequently to enhance total portfolio returns within the predetermined loss budget.

Investment Decisions

This paper analyzes the fundamental part of any investment strategy: the investment decision. An investment decision results in a change of exposure in an underlying portfolio, exchanging one stream of returns for another. Since it must be based on an assessment of relative performance, we define an investment exposure as a long position in one stream of returns and a short position in another set.

An investment decision is any deliberate act to change an investment exposure. It could be the choice to switch from cash to equities, to change exposures by shifting capital between assets, or the choice to move from an unhedged to a fully hedged stance on foreign investments, where an investment exposure is changed

Noriyuki Okuyama is a Senior Vice President of Pareto Investment Management Limited in London.

Gavin Francis is a Managing Director of Pareto Investment Management Limited in London.

Requests for reprints should be sent to: Noriyuki Okuyama, Senior Vice President, Pareto Investment Management Limited, 160 Queen Victoria Street, London EC4V 4LA, UK. Email: noriyuki.okuyama@paretopartners.com

by reallocating contingent capital. For example, a board resolution to replace a passive bond manager with an active bond manager would illustrate an investment decision requiring an evaluation of the relative performance of two alternatives.

The field of behavioral finance recognizes 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. Having implemented the decision, the outcome can be success or error. These outcomes are summarized in Table 1, the Fundamental Investment Decision Matrix. When the outcomes are represented by frequencies, this framework is consistent with the receiver operating characteristic approach.[1] Marques de Sa [2001] describes its use in signal detection theory and medical research.

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 as a result of missing a gain, a false negative or "type I error," or by incurring a loss, a false positive 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: incurring losses and failing to make gains. Investment error is defined as any uncertainty that contributes negatively to returns. 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. And there are two ways to improve returns: by capturing gains through true positives, or by avoiding losses through true negatives.

This paper proposes that active investment skill lies in the ability to maximize investment success relative to investment error. When assessing the historic performance of an investment decision maker, this is superior to measuring return versus volatility. However, why is it that traditional approaches to risk management like mean-variance analysis and value at risk focus only on 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. The losses one avoids and the gains one gives up are both "invisible."

This means that it is usually 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, and the invisibility as the degree to which the investor does not see ex post gains and losses resulting from an exposure in the portfolio. Any investment decision contains a trade-off between the visibility and invisibility of its investment outcome.

Risk-Taking and Risk-Reducing Decisions

However, there is another important difference between the visible and invisible outcomes. Acquiring an investment exposure is a *risk-taking* decision, in that it increases uncertainty by adding both gains and losses. The primary concern of investors is often to seek returns, or, more accurately, to seek gains while maintaining a rough idea of the size of a potential loss.

Figure 1 shows how overall uncertainty is increased by the investment decision to take an unhedged USD/JPY position.[2] In this example, as the investor decides to take on full foreign exchange rate risk exposure, both potential gains and losses are maximized. This results in perfect visibility, where the investor is exposed to the maximum true positive and the maximum false positive.

On the other hand, the decision to hedge an existing exposure is a *risk-reducing* decision that narrows the range of outcomes leading to perfect invisibility at the extreme. The critical distinction between risk-taking and risk-reducing is that a risk-reducing decision can only affect exposures that are already present in the portfolio. Therefore, risk-reducing decisions concern all exposures in a portfolio, while risk-taking decisions concern all possible incremental investment exposures that are *not* held in the portfolio.

The motivation for making a risk-reducing investment decision is that the potential gains are insufficient to offset potential losses. As shown in Figure 2, this

Table 1. *Fundamental Investment Decision Matrix*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante)	Take exposure (visible)	Capture Gains (True Positive) Success	Incur Losses (False Positive) Type II Error
	Avoid exposure (invisible)	Give Up Gains (False Negative) Type I Error	Avoid Losses (True Negative) Success

FIGURE 1
Risk-Taking Investment Decision

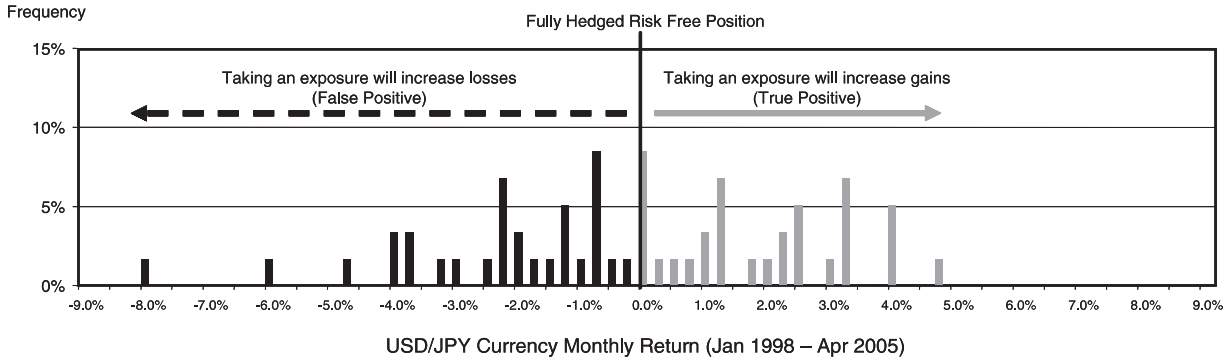

kind of decision reduces both true positives and false positives. A risk-reducing investment decision increases certainty and reduces the visibility of investment outcomes (invisibility-seeking); a risk-taking investment decision increases both the uncertainty and visibility of investment outcomes (visibility-seeking). Due to the cognitive bias of overconfidence, anticipation of potential gains from risk-taking tends to divert attention from the significant benefits of risk-reducing activities.

Risk-reducing decisions are fundamentally different from risk-taking ones. Risk reduction is constrained in the sense that it is only possible to reduce an existing exposure to zero. Furthermore, the activity is asymmetric, because an exposure can only be reduced, while uncertainties associated with an existing exposure vary asymmetrically over time. This means that a risk-reducing strategy can always be compared ex post with market outcomes.

Risk-taking, on the other hand, is a speculative, alpha-oriented activity that always adds new exogenous uncertainty to an existing portfolio. Since we have defined an investment exposure as a long/short position, no actual funding is required for an initial investment. The area is theoretically infinite, unconstrained by a specific boundary set by the existing underlying exposures. This means that risk-taking decisions must be standardized by their ex ante investment constraints. Moreover, risk-taking is symmetric: One can choose either a long or a short position, which can go up or down. All risk-taking activities add extra uncertainties, making the distribution of outcomes wider.

Investment Error versus Volatility

There are several consequences from equating investment risk with investment error (both false positives and false negatives). It becomes highly questionable to consider volatility as a proper measure of investment risk, since volatility only measures the 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 is considered as risk.

But this typical heuristic approach fails to distinguish between the different natures of uncertainty. It makes it possible to realign different returns according to their degrees of uncertainty, but it does not disentangle the constituent factors behind the uncertainty; losses and gains are mixed together.

FIGURE 2
Risk-Reducing Investment Decision

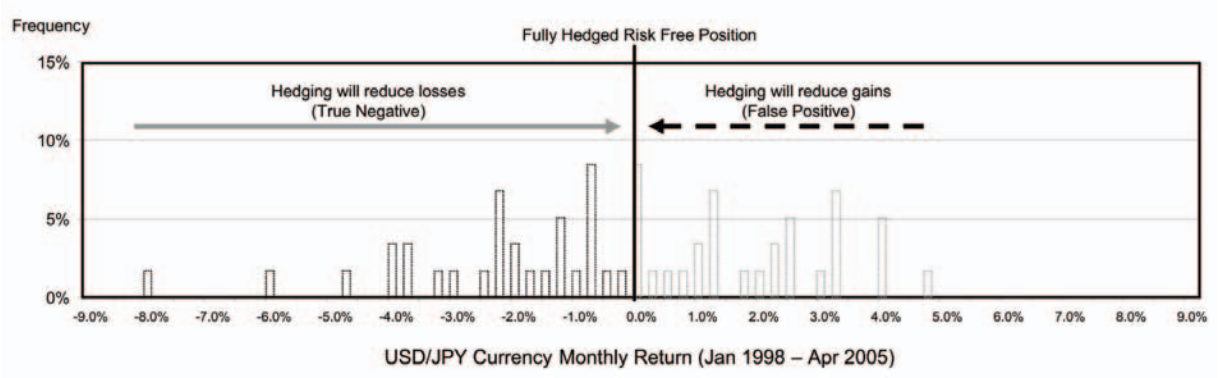

From the perspective of the decision matrix, the uncertainty of an achieved return should be decomposed into the uncertainty of positive and negative returns. This suggests that either a larger negative return or a smaller positive return will represent real investment errors to the same extent. However,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

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.

Table 2 illustrates how volatility incorrectly assesses realized gains. In order to analyze investment error/success properly, returns should be decomposed into negative returns (losses relative to the risk-free portfolio), and positive returns (gains). Unlike the mean-variance "two-tails-combined" approach, this risk decomposition approach enables us to determine the extent to which either skills to reduce losses or skills to retain gains have contributed to the total return.

Markowitz [1959], aware of the shortcomings of volatility, proposed using a downside risk measure. This was an improvement because it does not penalize gains, but it is not an effective measure of investment error because it still ignores false negatives.

Passive and Active Management

A passive investment management strategy is a single investment decision to match a benchmark portfolio consisting of a predefined set of investment exposures. This static approach makes no attempt to change exposures over time. Since no further investment decisions are required, a passive manager does not need to exploit any source of information. It is almost a tautology to say that a pure passive strategy has no expected return relative to its benchmark (before costs).

Active management, however, does not hold exposures statically. Investment decisions are made over time, using the manager's skill to generate a positive expected return relative to the passive benchmark alternative. The key distinction is that a successful active manager uses information to enhance returns, while a passive manager uses no information and therefore does not enhance returns.

The fundamental decision matrix has applications in many areas of investment. The field of currency risk management provides a very good example of the difference between active and passive risk-reducing strategies. This would be of interest to investors who are exposed to foreign exchange rate risk via investments in overseas assets.

Risk-Reducing Decisions

Passive Management—Controlling Currency Exposure

To illustrate further, we use as an example an investor with currency exposure in a portfolio of international assets. A passive currency hedging strategy consistently maintains a predetermined benchmark hedge ratio throughout an investment horizon, regardless of any subsequent subjective views. All active decision making on currency gains and losses is already completed once the specific benchmark hedge ratio is selected. There is thus no need to distinguish dynamically between currency gain environments and currency loss environments over time. Passive currency hedging is indifferent to both currency gains and losses.

Initially, we assume that the returns of a currency exposure are symmetric (consistent with a random walk), with 50% gains and 50% losses. A passive unhedged benchmark position takes on all the currency exposure. This approach successfully captures all the gains, but the portfolio also incurs the full impact of any losses.

As Table 3 shows, in this situation, false positives would be expected to occur 50% of the time, i.e., type II errors are maximized. Note that with an unhedged benchmark, losses and gains are fully visible as portfolio returns.

Table 2. *Failure of Volatility to Measure Investment Error*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante)	Take exposure	***Higher*** volatility (True Positive) **Success**	Higher volatility (False Positive) Type II Error
	Avoid exposure	***Lower*** volatility (False Negative) **Type I Error**	Lower volatility (True Negative) Success

Table 3. *Passive Risk-Taking—No Hedge on an Existing Exposure*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante) **100%**	Take exposure **100% visible**	Capture Gains (True Positive) Success **50%**	Incur Losses (False Positive) Type II Error **50%**
	Avoid exposure **0% invisible**		

Compare this with the passive fully hedged position, shown in Table 4, where the decision threshold has shifted. This choice of benchmark is consistent with eliminating all currency exposure. The passive hedge successfully protects against all currency losses, but the portfolio is deprived of making any currency gains. As a result, false negatives, or type I errors, are maximized. The effect is to exchange the type II errors for type I errors.

This fact is not obvious to investors who consider only the volatility of observed outcomes, because both the missed currency gains and the losses avoided are invisible. They do not show up as portfolio returns, and behavioral bias always tends toward visibility.

Clearly, the passive implementation of a partial hedge is simply an intermediate between the two extremes of fully hedged and unhedged. For example, with a 50% passive hedged position, we would expect to see 25% of observations in each quadrant, resulting in 25% of type I errors, and 25% of type II errors. Note from Table 5 that once again the passive approach trades type II errors for type I errors on a one-for-one basis, so total investment error remains at 50%.

The level of a passive hedge determines the horizontal decision threshold between visibility and invisibility. It is as if the horizontal decision threshold in the box simply slides up or down in parallel, depending on the choice of benchmark. Any reduction in incurred losses is accompanied by a reduction in captured gains. Incurring a loss (type II error) and missing a gain (type I error) have exactly the same negative impact on expected return. For risk exposures that have an equal chance of rising or falling, we conclude that

- all passive benchmarks have the same expected return,

Table 4. *Passive Risk Reduction—100% Hedge on an Existing 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) Type I Error **50%**	Avoid Losses (True Negative) Success **50%**

Table 5. *Passive Partial Risk Reduction—50% Hedge on Existing Exposure*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante) **100%**	Take exposure **50% visible**	Capture Gains (True Positive) Success **25%**	**Incur Losses (False Positive) Type II Error 25%**
	Avoid exposure **50% invisible**	**Give Up Gains (False Negative) Type I Error 25%**	Avoid Losses (True Negative) Success **25%**

- different passive benchmark selections represent different utilities on outcome visibility, and
- all passive benchmarks incur exactly the same amount of investment error!

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 with respect to economic values. Positive active information means positive economic values, which introduces optionality into the outcome distribution. Negative active information results in undesirable outcomes. Under this definition, all passive approaches have zero active information. The objective of active management is to introduce positive active information by achieving more than 50% success and therefore less than 50% errors.

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 is known as a free option payoff (a perfect outcome), where types I and II errors are zero. It constitutes perfect active management.

Table 6 shows the type of result an active manager aspires to obtain. Note that with the poorest active management, it is also possible to achieve 100% error. The result would be all losses and no gains, or 50% false positives and 50% false negatives. We refer to this as a perfect error, which an active manager must avoid (see Table 7). Because perfect foresight and perfect error are the two most extreme cases of active management, active information should be constrained by these two cases.

The Geometry of Risk-Reducing Decisions

To create the fundamental investment decision matrix, we decomposed returns into gains and losses. This matrix can be used to analyze a single decision. However, an active investment process consists of many ongoing investment decisions, so we must move from a discrete framework to a continuum. In Figure 3, we plot the uncertainty of gains against the uncertainty of losses for a given portfolio with a defined risk-free position.

Risk-reducing decisions are constrained, because they are defined by a preexisting (endogenous) exposure, which can only be reduced until the risk-free position is attained. The boundary of the square in Figure 3 represents the constraints on the possible outcomes of endogenous risk reduction.

Continuing with the example of currency hedging with an even balance of outcomes, the unhedged position stands at the top right corner, fully exposed to both losses and gains. It experiences all true positives and all false positives visibly in the stream of outcomes.

In contrast, the fully hedged position is, by definition, insulated from currency movements. It therefore represents the risk-free position, which lies at the lower left corner. Although the fully hedged position benefits from avoiding all losses, it is deprived of capturing any gains. All passive outcomes are represented by the

Table 6. *Perfect Foresight*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante) **100%**	Take exposure **50%**	Capture Gains (True Positive) Success **50%**	↓
	Avoid exposure **50%**	↑	Avoid Losses (True Negative) Success **50%**

Table 7. *Perfect Error*

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante) **100%**	Take exposure **50%**	↓	Incur Losses (False Positive) Type II Error **50%**
	Avoid exposure **50%**	Give Up Gains (False Negative) Type I Error **50%**	↑

FIGURE 3
Active Information Diagram (conditional risk attribution)

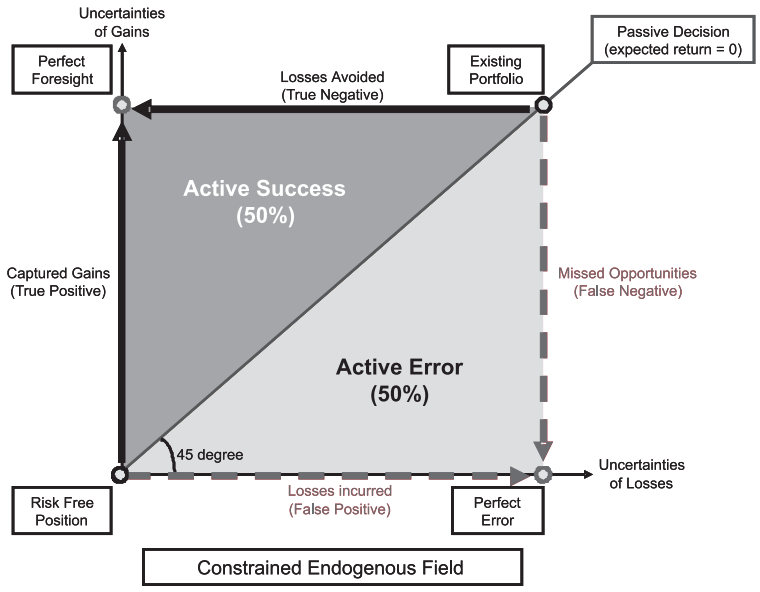

45-degree line segment from fully hedged to unhedged. These are the outcomes of strategies that contain no information and therefore have no expected return. Since investment success is offset by equal investment error, all passive outcomes contain zero active information.

The maximum information for active management would be perfect foresight, an outcome that appears at the top left corner: 100% investment success. Obviously, making exactly opposite investment decisions would generate 100% investment error, shown in the lower right corner.

The upper left triangle is clearly the desirable area, where investment success exceeds investment error, and therefore active information is positive. The area of regret lies below the 45 degree line, where active information is negative because error exceeds success. Under the random walk assumption, active managers have a 50% chance of falling into either triangle. Skillful active managers make use of information that is not used by a passive manager. The deviation of an active manager's track record away from the passive line toward the point of perfect foresight is a measure of the positive information contained in the investment strategy.

Note that Figure 3 presents an entirely different concept from the traditional risk/return charts used to display efficient frontiers. Here we decompose returns into gains and losses in order to measure the manager's skill in reducing preexisting risk. In this chart, the vertical axis shows gains, not returns, which are a mixture of gains and losses. The horizontal axis shows losses, not volatility, which is a mixture of the squares of the gains and losses. We call this active information diagram the conditional risk attribution (CRA). This logical departure from the conventional approach breaks the spell of the mean-variance paradigm.

The diagonals of the CRA square are significant. Moving away from the lower left corner toward the top right corner increases passive exposure to uncontrolled losses and gains, resulting in higher visibility. As we move from the lower right corner to the upper left corner, successive points represent progressively improving active information, passing through zero from negative to positive as one crosses the 45-degree passive line.

Active Management—Controlling Currency Exposure

Suppose an investor is exposed to currency movements as a result of investing in overseas assets. Considering the currency risk in isolation, the investor stands at the upper right corner of Figure 3, where, for the sake of clarity, we assume a fully hedged position is risk-free. A passive hedge would clearly reduce the uncertainty of both gains and losses. However, assuming no implementation costs, the outcome would lie on the 45-degree line, offering no expected return.

It would clearly be desirable to find an active manager with skill in adjusting the level of the hedge to arrive at an outcome in the upper left triangle. Figure 4 shows examples of such outcomes. All points on the same line running parallel with the passive 45-degree line have the same degree of active information. The distance from the 45-degree line represents incremental active information to the passive benchmarks. Managers A, B, and C demonstrate positive active information by achieving greater success than error. All three

FIGURE 4
Indifferent Active Information Line (45 degree)

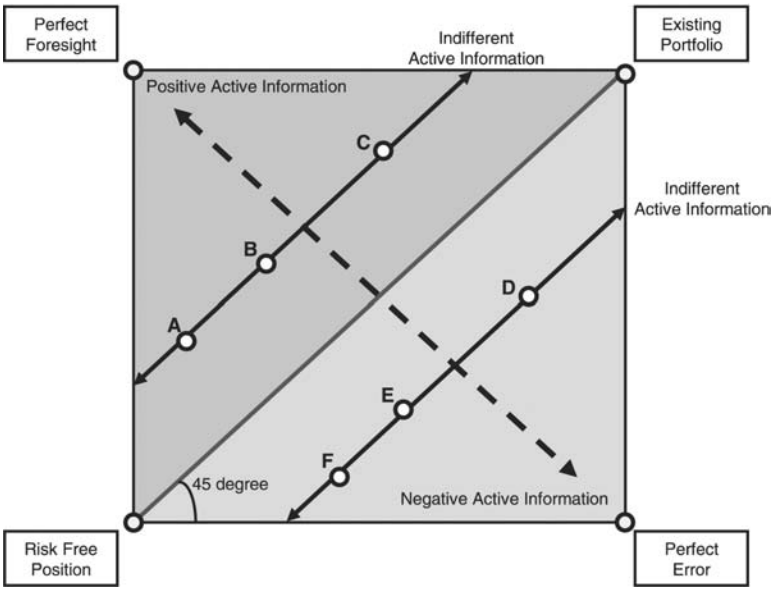

dominate all possible passive outcomes. Managers D, E, and F have introduced more error than success, to the detriment of the portfolio.

The key observation is that the only way for an active manager to attain positive active information is to treat the uncertainties of losses and the uncertainties of gains asymmetrically. This is dramatically different from a manager who uses volatility as a risk measure and considers the risk uncertainties of gains to be the same as the risk uncertainties of losses. However, it does make volatility an appropriate risk measure for a passive manager, because the objective of a passive management is no incremental change in economic value to the benchmark. Any deviations, either positive or negative, could be errors for a passive manager. This is why the tracking error has been a conventional risk measure for passive management.

Figure 5 shows a different interpretation of the conditional risk attribution. All points on the same line running perpendicular to the 45-degree line have the same level of visibility of investment outcomes. The distance from the risk-free position to each line repre-

FIGURE 5
Indifferent Visibility Line (45 degree)

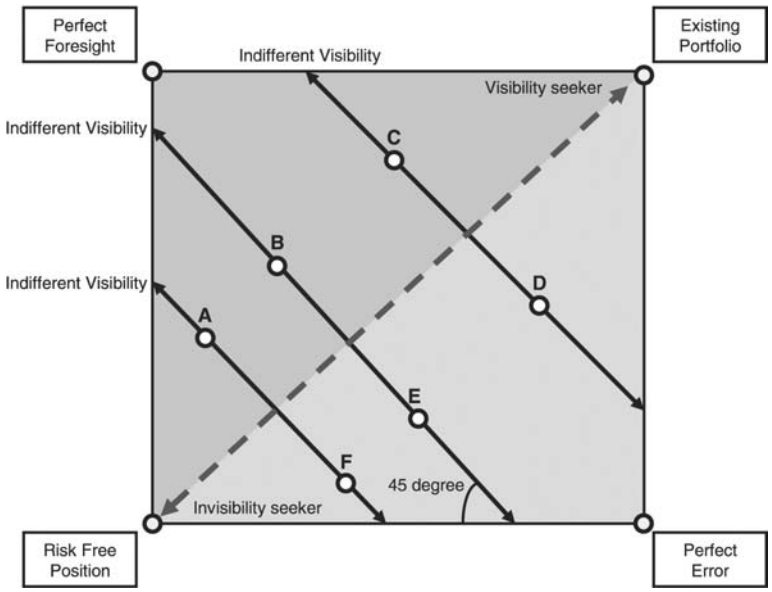

sents incremental visibility, i.e., the degree to which the investor sees the impact of the risk exposure in the portfolio returns.

Although manager A has positive active information and manager F has negative active information, both face the same level of visibility. This is true for managers B and E, and for managers C and D. Visibility itself does not explain economic value; it simply reveals whether the degree to which each manager has taken exposures is consistent with the predefined benchmark.

Although managers A and C achieve the same level of active information and the same expected return, they represent very different streams of outcomes and therefore very different visibility. Investor utility on visibility predetermined by an initial benchmark selection should distinguish investment outcomes that contain the same economic value.

Outcome A would therefore be appropriate for investors who prefer to be fully hedged because they are cautious about incurring loss (low visibility). Yet these investors may be willing to risk some limited losses to benefit from the active manager's skill at capturing gains by reducing the hedge when foreign currencies rise. This type of investor is an invisibility seeker.

Outcome C would suit investors who are keen to capture the full extent of unhedged gains (high visibility), but would also like to benefit from the active manager's skill in avoiding the worst losses by increasing hedges when foreign currencies fall. This type of investor is a visibility seeker.

Manager A successfully captured some gains while limiting losses (false positives); manager C successfully avoided some losses while limiting missed opportunities (false negatives). Remember that both types of error have the same negative impact on expected return, so value can be added by controlling either false positives or false negatives.

The Effect of the Environment

At this point we relax the assumption that 50% of observations are losses and 50% are gains. Even if this assumption were valid over the long term (as it may be for developed market currency exposures), an exactly even balance is very unlikely over any particular sample period. It is more likely that the exposure will experience episodic upward or downward trends over the period, which would change the shape and size of active success and error (see Figure 6). Note that the outcome for a passive unhedged manager, with no information, has diverged from the 45-degree line purely as a result of an uncontrolled market move and *not* as a result of any manager's active skill.

Clearly, it is highly desirable to remove this market effect when assessing manager skill. This is particularly important for an active manager, who might be tempted to claim responsibility for outperformance, but blame the market for underperformance. This normalization is achieved by translating all points perpendicular to the 45-degree line by an amount appropriate for the relevant benchmark[3] (see Figure 7). As a result, the line joining all passive portfolios becomes aligned with the 45-degree line. But since the active portfolios are shifted along the same vector as their corresponding benchmarks, relative performance is preserved.

For example, suppose that portfolio G was managed against an unhedged benchmark. This point would be shifted by the same amount as the unhedged benchmark to G′, which is in the upper left triangle of investment success. After normalization, environmental bias

FIGURE 6
Effect of the Environment

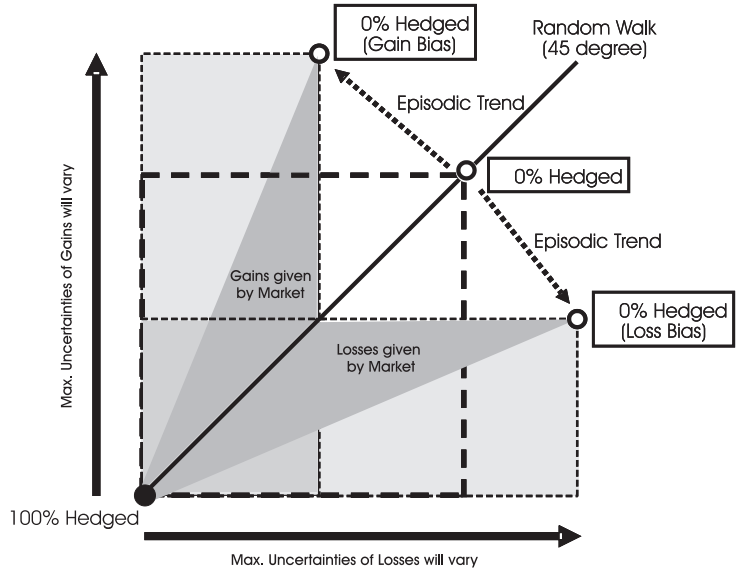

FIGURE 7
Normalization to Remove Environmental Bias

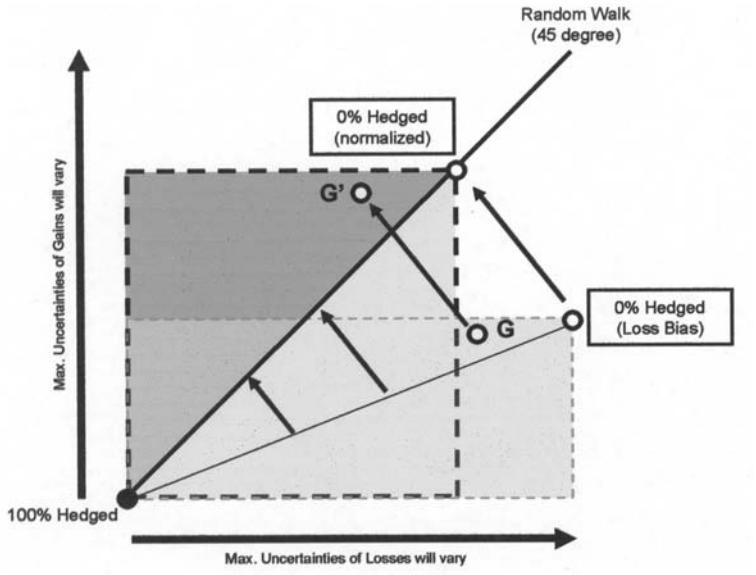

is removed and the active information diagram becomes a square and therefore symmetric. However, the side of the square will be equal to the average of the uncertainties of losses and gains of each individual data series.

One further step generalizes the approach: to rescale the axes to one unit. This is achieved by taking the coordinates of every point on the diagram, dividing each gain by the maximum normalized gain, and each loss by the maximum normalized loss. We arrive at a unit square that permits a comparison between different managers managing different investment exposures over different time periods. We call this generalized conditional risk attribution (GCRA) (see Figure 8).

Risk-Taking Decisions

We now shift our framework to the field of risk-taking, swiftly running through a similar logical sequence. However, one glance at the risk-taking decision table in Figure 9 highlights an important point: It is only possible to take an exposure, because it makes no sense to avoid an exposure you do not have. Thus it

FIGURE 8
Generalized Conditional Risk Attribution

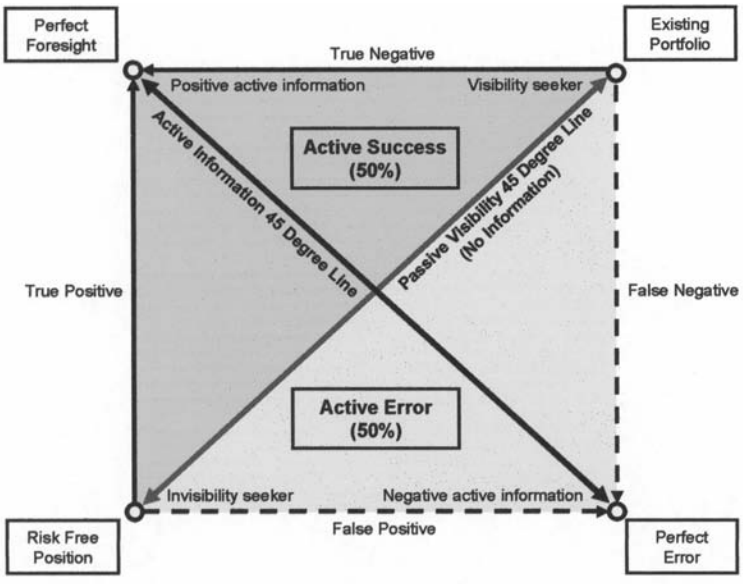

FIGURE 9
Fundamental Investment Decision Matrix—Risk-Taking

		Actual Outcome (ex post)	
		Gain	Loss
Investment Decision (ex ante) **100%**	Take exposure **100%**	Capture Gains (True Positive) Success **50%**	Incur Losses (False Positive) Type II Error **50%**

is clear that risk-taking decisions can only add gains and add losses.

Since an investment exposure has been defined as a combination of long and short positions, no investment of capital is required. In risk-taking decisions, sponsors are investing contingent capital, i.e., the amount of money needed to cover the worst losses (see Layard-Liesching [2004]). As the decision table shows, all risk-taking decisions should be assessed in terms of the level of success relative to the amount of error.

In principle, it is possible to add an infinite amount of any risk to a portfolio by taking more long/short positions. Risk-taking is unbounded. Therefore, unlike risk reduction, the active risk-taking information diagram has no limits (see Figure 10). Since there is no benchmark level of risk, it is not possible to normalize or generalize performance in the same way as for risk reduction.

The only way to define a bounded risk-taking area is in terms of an ex ante defined limit. This often takes the form of a maximum exposure or a volatility-based constraint. This type of limit defines a symmetric range of outcomes on a scale that can be compared with the risks inherent in the existing portfolio.

Figure 11 gives the geometric interpretation of risk-reducing and risk-taking decisions. Any risk-taking investment strategy that contains no information (i.e., passive risk-taking) lies on a projection of the 45-degree line beyond the existing portfolio. However, the area of investment success or error is potentially infinite. Thus, it is important to impose strict constraints on the size of losses incurred by risk-taking managers.

The uncertainty observed in risk-reducing activities is defined by an investor's existing portfolio, so visibility and invisibility are both measurable. However, uncertainty in risk-taking activities is newly created at each manager's discretion. In this case, only visibility can be measured. Since many conventional performance measurements such as the Sharpe ratio, the information ratio, and the Sortino ratio suffer from cognitive "visibility" bias, they fail to measure the impact of the invisible risk that we successfully capture within the GCRA framework.

FIGURE 10
Geometric Representation of Risk-Taking

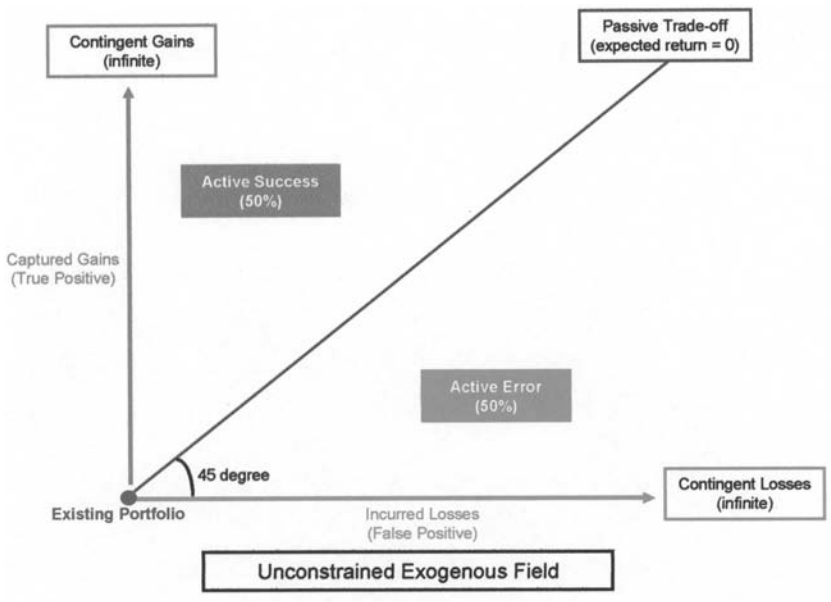

FIGURE 11
Geometric Representation of Risk-Reducing and Risk-Taking

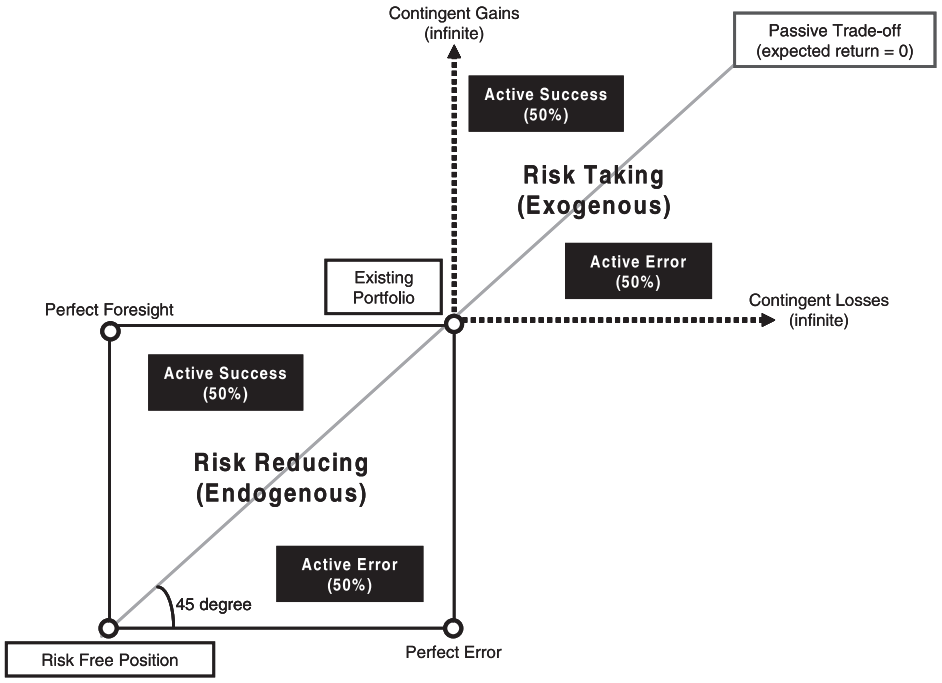

Risk Budgeting

Risk-reducing decisions and risk-taking decisions are different choices about different outcomes. Risk reduction considers only the exposures already held; risk-taking considers only exposures that are not held.

But the role of risk-reducing decisions is fundamentally important to successful investment strategies. The initial portfolio construction phase is a risk-taking activity: investing risk-free assets such as cash in simple long-only strategies within conventional asset classes such as equities and bonds. This can be achieved either passively or actively within a predetermined risk of loss 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. This can also be achieved either passively or actively.

The effect of reducing risk is to free up part of the risk budget so it can be reallocated to better rewarded risk-taking activities. The 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 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 we develop here proposes to enhance returns by improving the ratio of investment success to investment error. Whether investors are reducing or increasing risk, the objective is the same: to exploit active information to increase investment success. Before considering risk-taking activities, however, investors should always consider other ways to reduce risk, preferably by using a skillful risk-reducing active manager. See Figure 12 for a clear depiction of the repeating cycle of risk of loss rebudgeting.

The risk of loss rebudgeting approach we describe highlights the significance of distinguishing risk-reducing from risk-taking to achieve a better risk-adjusted return within a predetermined loss budget. Without making a clear strategic distinction, neither the first process (creation of loss budget) nor the second (return enhancement) will be implemented efficiently.

Conclusion

We examined investment decisions using our fundamental investment decision matrix to define investment success and error. This led to a clear distinction between risk-reducing and risk-taking decisions. The framework was effective in describing risk-reducing decisions, as illustrated by the examples of passive and active currency hedging. A geometric interpretation gave insights into the distinction between environmental effects and manager skill.

We expanded the geometric interpretation to the generalized conditional risk attribution framework (GCRA), where different managers' results and different period results can be compared through normalizing environmental effects and generalizing different uncertainties into one unit. This allowed a simple extension of the structure to incorporate risk-taking deci-

FIGURE 12
Risk of Loss Rebudgeting Framework

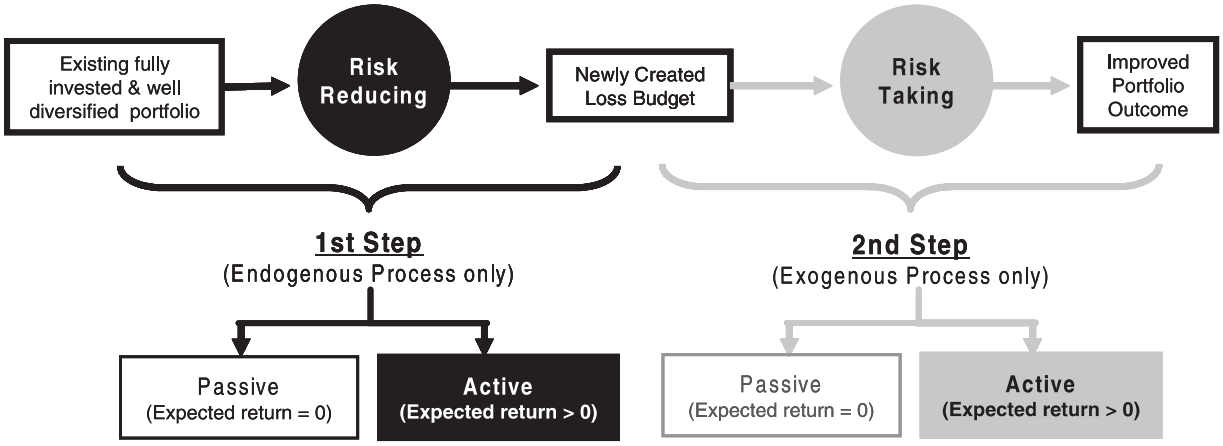

sions. Manager skill is associated with active information that is exploited by the investment process.

Numerical examples used simple frequencies of success and failure. Clearly, investment outcomes depend on the magnitude of outcomes. Future research should consider the precise formulas of the performance metrics for investment success and error.

Our purpose in setting the analysis against the background of risk of loss budgeting was that we believe this approach has broad application to investment management, far beyond the confines of currency hedging. By disentangling the cognitive bias inherent in the conventional mean-variance approach, the GCRA framework allows the portfolio construction process to better reflect the assessment of investment decisions.

Notes

1. This approach is "model-free" in that it is not necessary to model the underlying process.
2. All monthly returns are adjusted for interest rate differentials.
3. Since the size of the translation depends on the distance from the origin, the mathematical transformation of the passive outcomes is called a shear. In order to preserve benchmark relative performance, the coordinates of an active manager are translated by the same amount as the benchmark against which the portfolio is managed.

Reference

Kahneman, D. and A. Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47, (1979), pp. 263–292.

Layard-Liesching, R. "Investing Contingent Capital." *Journal of Portfolio Management*, 30th Anniversary Issue, 2004.

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

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

Marques de Sa, J.P. *Pattern Recognition: Concepts, Methods and Applications*. Berlin and Heidelberg: Springer-Verlag, August 2001.