The Journal of Behavioral Finance
2007, Vol. 8, No. 3, 121–137

Quantifying the Information Content of Investment Decisions in a Multiple Partial Moment Framework: Formal Definition and Applications of Generalized Conditional Risk Attribution

Noriyuki Okuyama and Gavin Francis

Investment decisions are based on a trade-off between profit and loss. This paper aims to measure the effectiveness of active investment decision-making processes by comparing the distributions of positive and negative outcomes against those available to a passive investor. A genuinely skillful active manager should generate outcomes with more attractive loss/gain balances than a passive buy-and-hold strategy. Generalized conditional risk attribution is a method of assessing whether a decision-making process has created this benefit.

We initially describe this methodology in the context of passive investing, where no ongoing investment decisions are made. We then compare the returns resulting from a static risk exposure against a symmetric random walk. The partial moments of a distribution describe the balance between losses and gains relative to the risk-free position. Conditional risk attribution can quantify the asymmetry between the upper and lower partial moments. This technique is useful for characterizing the risk of a set of passively managed outcomes measured over a given time period.

In assessing active decision-making processes, we illustrate how to normalize partial moments to account for any bias in the historic environment. By rebasing the results to a unit scale, we obtain a general result that we can compare against other time periods for different risk exposures. We use the term "optionality" to define the degree to which active decisions have skillfully improved the balance between gain and loss. A second term, "visibility," is used to describe the extent to which available returns were captured; it is not related to manager skill.

We then move on to generalized conditional risk attribution (GCRA), which offers several advantages over traditional performance measurement methods. Approaches that summarize outcomes using only two parameters, such as the mean and the volatility, implicitly assume returns are symmetric. However, the entire purpose of active management is to create asymmetric distributions to capture gains and avoid losses. Two-parameter models cannot effectively capture this asymmetry.

The results of our analysis allow investors to characterize different investing styles. Some managers focus primarily on forecasting the first moment of the distribution, relying on diversification to control risk. Others concentrate on improving the second moment of the distribution, with styles based on volatility. A third style of decision-making explicitly controls the probability of loss (higher moments).

keywords: *Upper and Lower Partial Moment, Generalized Conditional Risk Attribution, Maximum Permitted Risk Exposure, Optionality, Visibility*

Introduction

Many of the performance metrics that have emerged from the framework of modern portfolio theory fail to distinguish between loss and gain. Consider the Sharpe ratio [1994]. One of its advantages is that, by subtracting the risk-free rate from investment returns, outcomes can be clearly separated into losses and gains. However, this distinction is not used in the calculation. In fact, losses and gains are mixed to obtain an average return, which is then divided by the volatility of

Noriyuki Okuyama Senior Vice President, Pareto Investment Management Limited in London, 160 Queen Victoria Street, London EC4V 4LA, UK. Tel: +44 20 7163 1090; Fax: +44 20 7163 1066; Email: noriyuki.okuyama@paretopartners.com

Gavin Francis Managing Director, Pareto Investment Management Limited in London, 160 Queen Victoria Street, London EC4V 4LA, UK. Tel: +44 20 7163 1023; Fax: +44 20 7163 1066; Email: gavin.francis@paretopartners.com

returns. This measure of loss would only be useful if the returns were described by a normal distribution with a mean of zero.

Thus the Sharpe ratio, like the information ratio, measures a signal-to-noise ratio, but it does not measure the balance between loss and gain. Noise contains both gains and losses, and a signal-to-noise ratio is meaningful only if perpetually matching a target outcome is the top priority. But in investment, exceeding the target (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.

The Sortino ratio [1994] comes closer to distinguishing between loss and gain by defining a threshold return that segregates the data into two parts. If the threshold is set to zero, negative observations can be used to obtain a better estimate of expected loss than volatility would provide. However, the numerator of the ratio remains a mixture of losses and gains. The weakness of this conventional approach is its focus on the final outcome rather than on the process that generates it. These approaches fail to identify investment skill, instead mixing it with pure luck.

Biglova et al. [2004] describe the less well-known Rachev ratio, which compares the probabilities of extreme gains with extreme losses by defining an upper and a lower threshold. When both align to zero returns, we move closer to a useful measure of loss versus gain.

All of these ratio measures can be misleading, however, because they do not distinguish between gains obtained by holding a passive leveraged position, and gains obtained through skillful active management. This distinction is essential to properly evaluate the economic value of an active manager versus a purely mechanical process that contains no insight.

Furthermore, each of these ratios describes the outcomes with only two parameters. They fail to identify the skill of active managers to modify the skew of the returns distribution. In the long term, investment return is the difference, not the ratio, between gain and loss.

Utility theory has been used to separate losses and gains from inputs that combine them such as return, volatility, and correlation. This approach is used to create portfolios in standard mean-variance optimization under the assumption that holdings remain static. But superior portfolios would be available to investors who can directly measure the loss and gain of investment processes.

Fishburn [1977] defined the lower partial moments of a distribution relative to a threshold. This paper extends that approach by measuring the upper partial moments of a distribution and comparing them against the lower partial moments. This is called conditional risk attribution. In developing this approach to assess the active skill of an investment manager, it is necessary to compare the range of outcomes available to managers who use a mechanical process such as passive indexation, which requires no information input. This leads to a framework called generalized conditional risk attribution.

This paper also builds on earlier work by Okuyama and Francis [2006] to create a suite of loss and gain measures that can easily be calculated from a set of monthly or daily returns. This provides a rich framework in which to characterize and identify the effectiveness of active investment decision-making processes. We explain the use of these measures and the interpretation of graphical presentations of the results below.

A skillful active manager should be able to modify the distribution of returns available to a passive investor so as to generate a more attractive distribution of outcomes. The generalized conditional risk attribution (GCRA) method is designed to measure this skill.

Definition of Terms

Some of the terms we use here are somewhat open to interpretation. Since our purpose is to give a formal definition of GCRA, it is helpful to define certain terminology. A risk exposure is defined 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. To clarify, consider an active manager who varies a portfolio allocation between the risk-free asset and a risky asset. Any portfolio mix can be considered as a 100% allocation to the risk-free asset, plus a risk exposure that is long the risky asset and short the risk-free asset. Similarly, an active equity manager holds the (risk-free) benchmark index, plus a set of risk exposures. The concept of a risk exposure comfortably accommodates leveraged positions and derivative exposures.

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. If a long/short position is created by using two different assets that are both denominated in U.S. dollars, it consists of two risk exposures against the U.S. dollar cash rate.

We 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, or the interest rate differential.

Passive management is characterized by a static risk exposure; active management is characterized by changing the risk exposure through time to enhance the balance between losses and gains. To analyze the skill of an active manager, it is necessary to know which risk exposures were managed, the size of the maximum permitted position in each risk exposure, and what performance was attributable to each return source. Third parties rarely have access to all of this information. Active managers are the most qualified to perform a comprehensive analysis.

Some investors hold a particular type of risk exposure, which plays a very specific role within a portfolio. If a risk exposure acts as an offset by reliably generating returns opposite to another risk exposure in the portfolio, it is called a hedge. In order to be effective, a hedge must have a correlation reliably close to -1 in all market environments. The magnitude of the hedge determines whether it partially or fully offsets the other risk exposure.

The degree of offset is called the hedge ratio, which, by definition, can only lie between 0 and 1. Hedges can be managed either passively or actively. A hedge has the special property of reducing the size of the maximum potential loss of the portfolio. Hedging is therefore a risk-reducing activity. Adding any other risk exposure to a portfolio is a risk-taking activity, because it increases the maximum potential loss of the portfolio.

Although active hedging and active risk-taking are two mutually exclusive activities, they are connected as follows: The investment guidelines of an active hedging manager define the set of existing risk exposures to be hedged. In contrast, the investment guidelines of a risk-taking manager would normally list a set of permitted risk exposures, possibly with some kind of limit on the size of the maximum position. So an active hedging manager continually adjusts the visibility of well-defined existing risk exposures, while an active risk-taking manager continually adjusts the visibility of a set of "potential" risk exposures. In both cases, active skill lies in improving the balance of loss and gain available to a passive manager.

Formal Definitions of Partial Moments

Our analysis starts by measuring 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. The GCRA approach segments the set of observations into losses (negative returns), and gains (positive returns). As a result, rather than analyzing a single distribution of outcomes, GCRA analyzes it in two parts: the distribution of losses, and the distribution of gains (see Figure 1).

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, and so on; the distribution of gains would be characterized in

FIGURE 1
Disentangling Gains from Losses

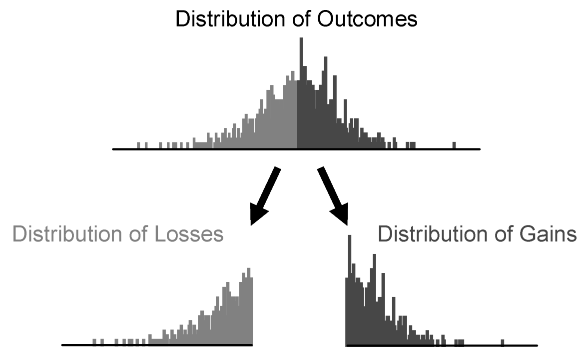

a similar way. Because these distributions were originally derived from a single distribution of n outcomes, the moments are scaled according to the total number of n original observations.

Table 1 summarizes the formal definitions of partial moments.[3] The general definition allows any (integer or non-integer) partial moment to be calculated. The "0-th" moments define the fractions of observations that were losses or gains: the proportions of so-called down and up months for a monthly track record. The first partial moments represent average loss and gain, while the second partial moments represent uncertainty of loss and gain. The third partial moments are more strongly influenced by the tails.

By defining the partial moments this way, we observe that the corresponding moment of the original distribution can be recovered. For example, the first moment or average return is the sum of the average loss and the average gain:

$$\bar{X} = \frac{1}{n}\sum_{1}^{n} x_i = \frac{1}{n}\sum_{<0} x_i + \frac{1}{n}\sum_{>=0} x_i = \bar{L} + \bar{G}$$

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 uncertainty of loss. 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 GCRA is concerned with losses and gains, the

Table 1. *Definitions of Partial Moments*

	Lower Partial Moment	**Upper Partial Moment**
First	$\bar{L} = \frac{1}{n}\sum_{x_i<0} x_i$	$\bar{G} = \frac{1}{n}\sum_{x_i\geq 0} x_i$
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	$\overline{L^k} = \frac{1}{n}\sum_{x_i<0} x_i^k$	$\overline{G^k} = \frac{1}{n}\sum_{x_i\geq 0} x_i^k$

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 extreme observations, which can be a disadvantage when attempting to forecast the future. But from a historic performance analysis perspective, those observations are the most important in assessing risk management.

Conditional Risk Attribution

The underlying principle of conditional risk attribution (CRA) is to compare the results of passive management against a set of perfectly symmetric outcomes. We use the outcomes of passive management to illustrate the graphic representations and calculations used in CRA.

Having calculated the upper and lower partial moments, we plot each pair in a CRA diagram conditional on the time period of the observations. For convenience, we 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, and the horizontal axis represents the corresponding lower partial moment (see Figure 2).

Now, if the evolution of the underlying risk exposure were determined by a coin toss, it would follow a random walk. In this case, we would expect that the upper partial moment of a preexisting exposure G could deviate from the lower partial moment, L, in the short term. However, the long-run expectation would be for G to match L so that the CRA diagram becomes square. The top right corner is the unhedged preexisting exposure (L, G). The bottom left corner is the risk-free position (0, 0).

Passive hedging reduces risk by avoiding both losses and gains without using 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 preexisting exposure (L, G) and the risk-free position (0, 0). Under the random walk assumption, the slope becomes forty-five degrees. In the long term, losses can only be avoided to the same extent that upside opportunities are missed. With respect to active information content, any partially hedged passive outcome should be symmetric on upper and lower partial moments. Any deviations from the forty-five-degree passive line should account for either positive or negative market bias.

However, no one would consider betting on a coin toss to be an investment. Real financial returns need properties that make them sufficiently attractive for investors to take a passive exposure. The most obvious is

FIGURE 2
Conditional Risk Attribution

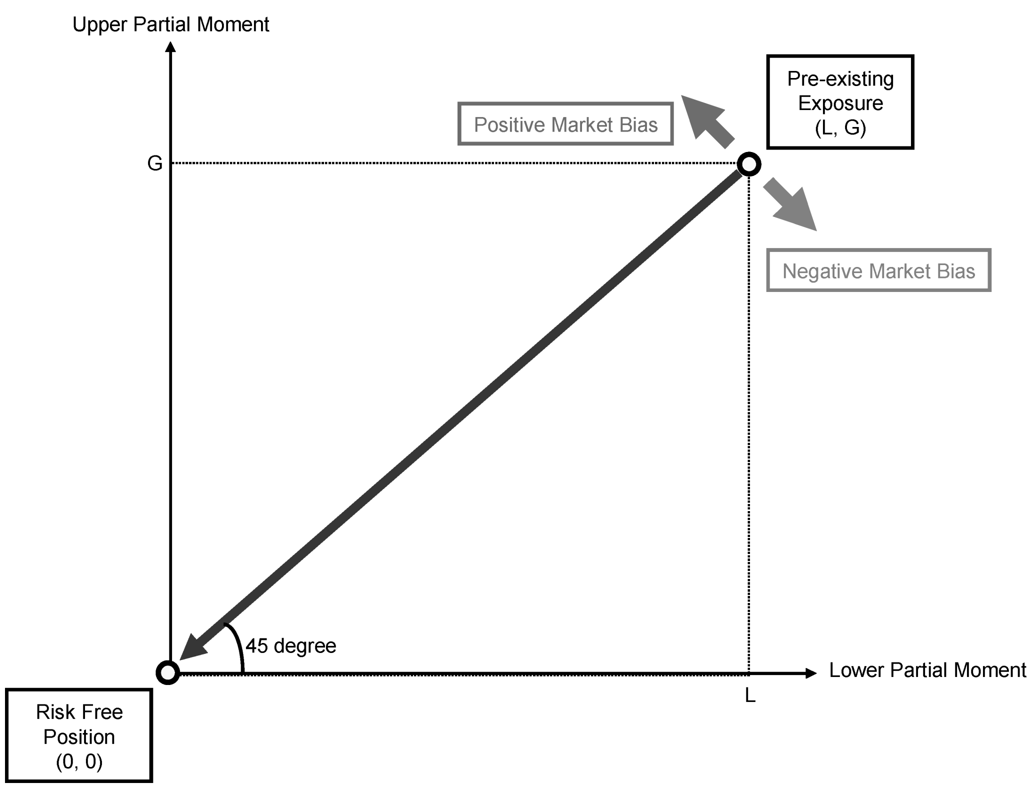

FIGURE 3
CRA of Different Moments

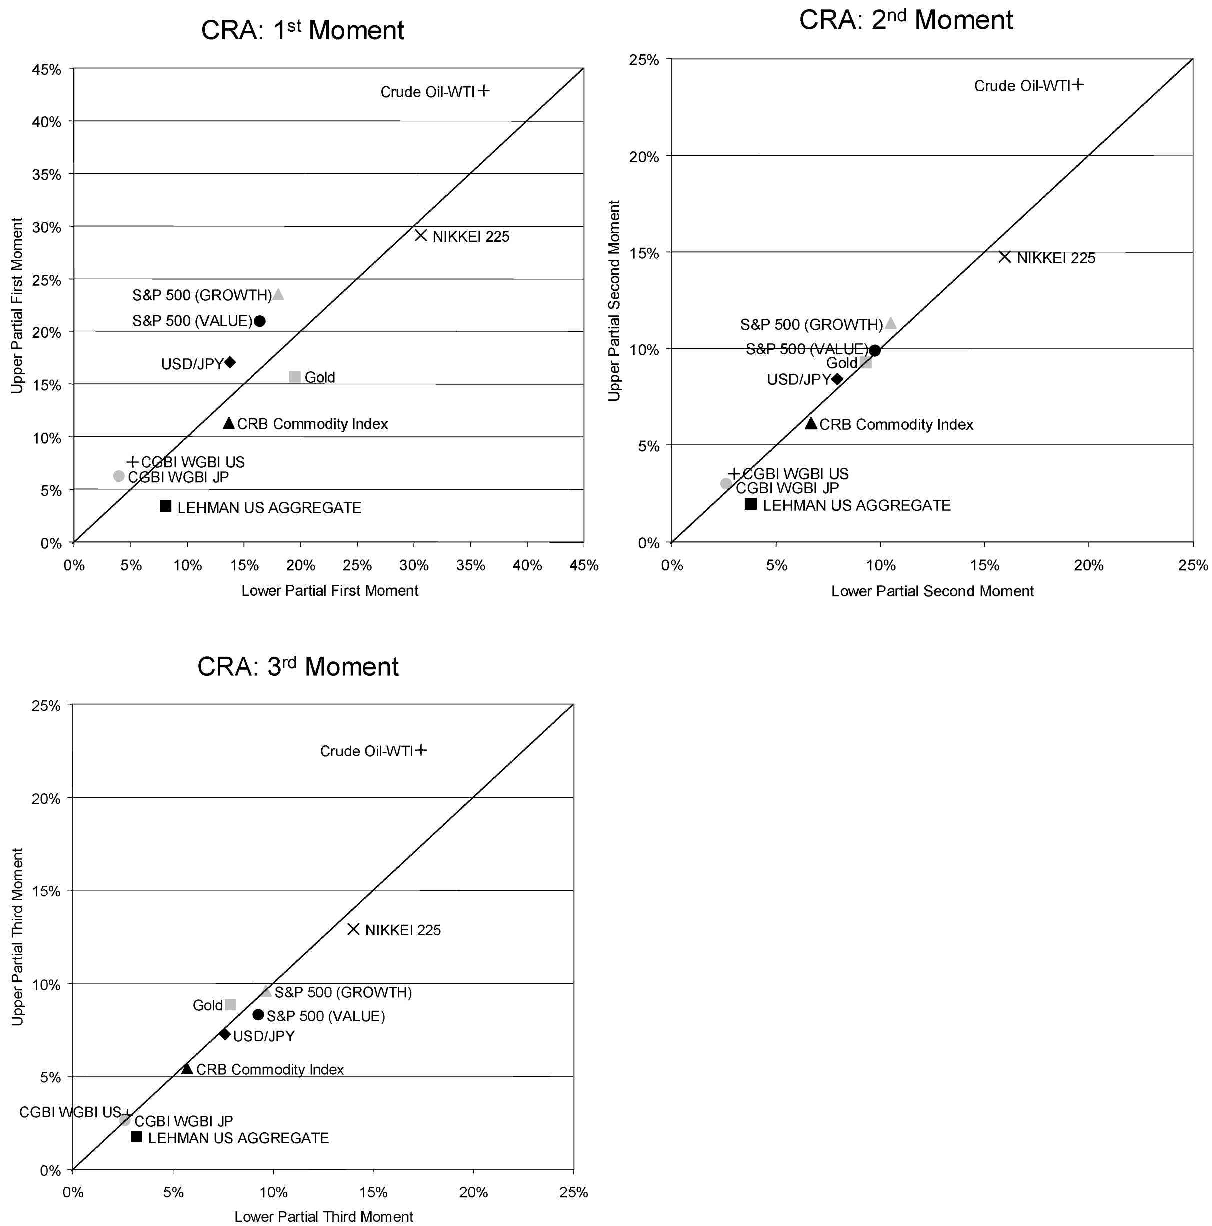

Source: DataStream.

that average gains should be larger than losses relative to the risk-free rate. In other words, the coordinate of the first moment (L, G) should lie above the forty-five-degree line.

However, most investments that offer a return premium, such as equities, expose the investor to a low probability of a potentially large loss. This characteristic is represented on the CRA diagram by points below the forty-five-degree line for the higher moments. Some risk exposures, like lottery tickets or insurance, have a negative expected return but offer potentially large gains. On the CRA diagram, these are represented by a first moment that appears below the forty-five-degree line, with higher moments above it. The rules of market efficiency ensure that investors are unlikely to find a *passive* exposure for which all moments lie above the forty-five-degree line.

Empirical findings support this argument (see Figure 3). Using historical price data from January 1988 through December 2005, we calculate both upper and lower partial moments (the first, second, and third moment) for various investments.[4] The S&P 500 composite data show that the growth and value sectors both show a positive first moment, but an incremental deterioration toward higher moments, implying episodic large losses over the long-term investment horizon. This is consistent with an overall distribution of outcomes that has a positive mean but a pronounced negative tail.

The same observation was seen for the U.S. dollar-Japanese yen exchange rate, as well as for the U.S. and Japanese bond indices. However, gold, the CRB Commodity index, and the Lehman U.S. Aggregate Corporate Bond index all showed a negative first moment but incremental improvements toward higher moments, implying episodic large gains under a long-lasting unfavorable environment.

There are, however, some exceptions, such as the WTI crude oil price index and the Nikkei 225 Stock index. The WTI index shows a neutral first moment with improving higher moments from January 1988 through December 2001. But since January 2002, geopolitical risks and oil supply/demand imbalances have caused all CRA moments to show an extraordinary and attractive bias: a positive mean with a long positive tail.

The Japanese Nikkei 225 stock index consistently showed negatives through higher moments because of the fallout from the late 1980s stock market bubble, which sank the Japanese economy into ten years of asset depreciation. The significant impact of the initial fall dominated through higher moments. Nevertheless, it is rare in a passive exposure to see any consistent significant positive or negative bias throughout all moments in the CRA.

These considerations are very important when selecting risk exposures to buy and hold in a portfolio, for example, to determine a strategic portfolio allocation, to select a private equity investment, or to hire 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.

Maximum Permitted Active Risk Exposure

We now turn to assessing active investment processes. The *raison d'être* of an active manager is to vary the size of a risk exposure through time in order to generate a better trade-off of gain versus loss than can be obtained through passive management.

While the purpose of CRA was to compare passive investments against a symmetric random walk, the objective of our next step is to compare the result of active management against passive management. This requires a definition of the range of outcomes available from a static exposure, which is determined by the maximum permitted active risk exposure. In the context of active hedging, this is well defined as the preexisting unhedged position. To fully analyze other decision-making processes, we need to isolate each individual risk exposure under active management and to define the maximum exposure permitted by the mandate.

In the case of portable alpha strategies, where the benchmark is zero, active management consists only of long and short positions (which we could also refer to as unconstrained investment). However, no active manager is allowed to take unlimited risk exposure. Since portable alpha strategies are not physically constrained by any well-defined existing risk exposures, the investment guidelines must include virtual constraints that either directly or indirectly determine the maximum sizes of a set of "potential" risk exposures. The most direct way to do this is to fix maximum permitted static risk exposures as completely customized exposure constraints, based on a notional amount.

For example, suppose an active alpha manager claims to generate absolute returns within clearly defined limits. Assume that maximum permitted exposure is $+/-100\%$ of the notional amount. The returns generated by long notional exposures can now be compared against the outcomes of holding a passive $+100\%$ notional exposure over the entire period. Separate analyses should be performed for the returns generated by short notional exposures. But the advantage of such a constraint is that it enables us to analyze active skills (success and error) in long and short positions separately.

Given the very complicated long/short setups commonly used by alpha-seeking managers, it is generally considered even simpler to use volatility as a proxy constraint to define the largest permitted static risk exposure for such activities. A volatility constraint is certainly not ideal, however, compared with the notional exposure constraint.

First, the volatility constraint is not an absolute limit, and thus provides less strict target protection under extreme events. This means both uncontrollable positive and negative surprises can occur. Second, as long as volatility stays within the 1-standard deviation limit, there is no constraint on maximum permitted active exposure. Active positions can go infinitely beyond the notional limit provided the balance of gains and losses lies within the predetermined volatility range. This means final outcomes can be heavily biased by the uncontrollable sequence of gains and losses, which will lead to confusion of active skills and pure luck.

The concept of a maximum permitted active risk exposure is crucial to our 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. 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 like 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

Having defined the maximum active risk exposure, we can plot the upper and lower partial moments of both active and passive returns on a CRA diagram. The next step is to eliminate any deviations from the forty-five-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 sample period that creates episodic asymmetry in gains and losses. Clearly, in assessing active manager skill, this market bias needs to be removed. We call this process "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 passively hedging the preexisting exposure at a fixed hedge ratio. The process of normalization resets to zero any incremental deviations caused by market bias. This normalization adjustment is a perpendicular shift from the market-biased passive line to the forty-five-degree market-neutral passive line. Although the "market bias" ratio is indifferent to any hedge level, the degree of asymmetry is different at each hedge level. The perpendicular distance from the forty-five-degree line is the "market contribution to asymmetry (MCA)," a quantity to be normalized. It is defined as follows:

$$\text{Market Contribution to Asymmetry (MCA)} = \frac{\sqrt{2}}{2} \times (G - L) \times \text{Exposure}$$

where $0 \leq \text{exposure} \leq 1$.

Normalizing the MCA quantity rebases the market-biased preexisting exposure back to the forty-five-degree passive line where G and L, respectively, converge into the average of the two.[5] Note that this adjusts G and L by the same amount, which is appropriate for rebasing the outcome of passive management.

Figure 4 shows that the top right corner of the CRA diagram therefore moves to the market-neutral existing exposure $(\frac{G+L}{2}, \frac{G+L}{2})$. MCA for the unhedged exposure is $\frac{\sqrt{2}}{2} \times (G - L)$, and MCA for the fully hedged risk-free positions is zero. The MCA quantity is determined by the passive exposure; its value differs for each benchmark hedge level. The MCA's value is

FIGURE 4
Normalizing Market Contribution to Asymmetry

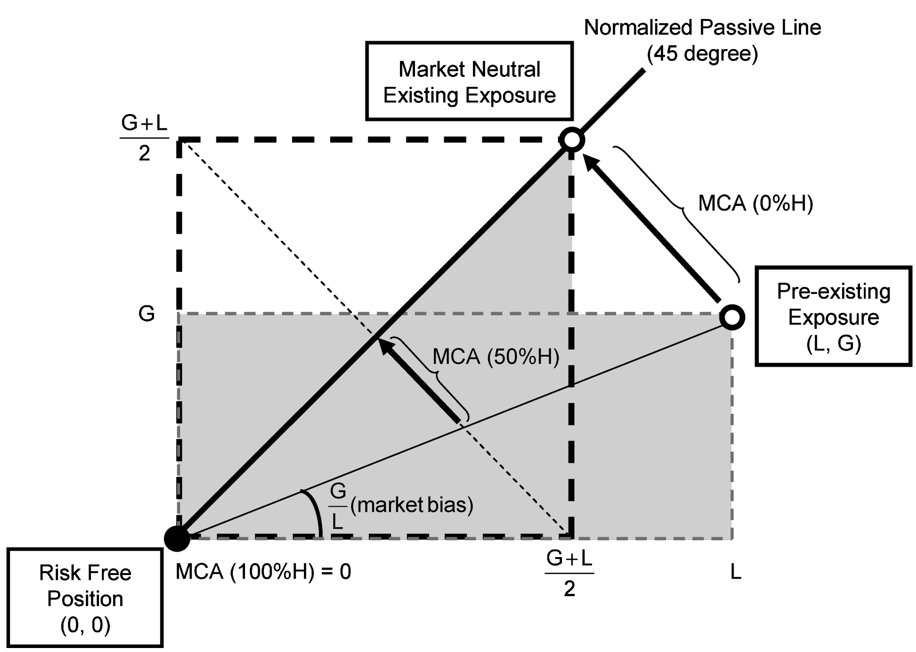

FIGURE 5
Generalized Conditional Risk Attribution Diagram

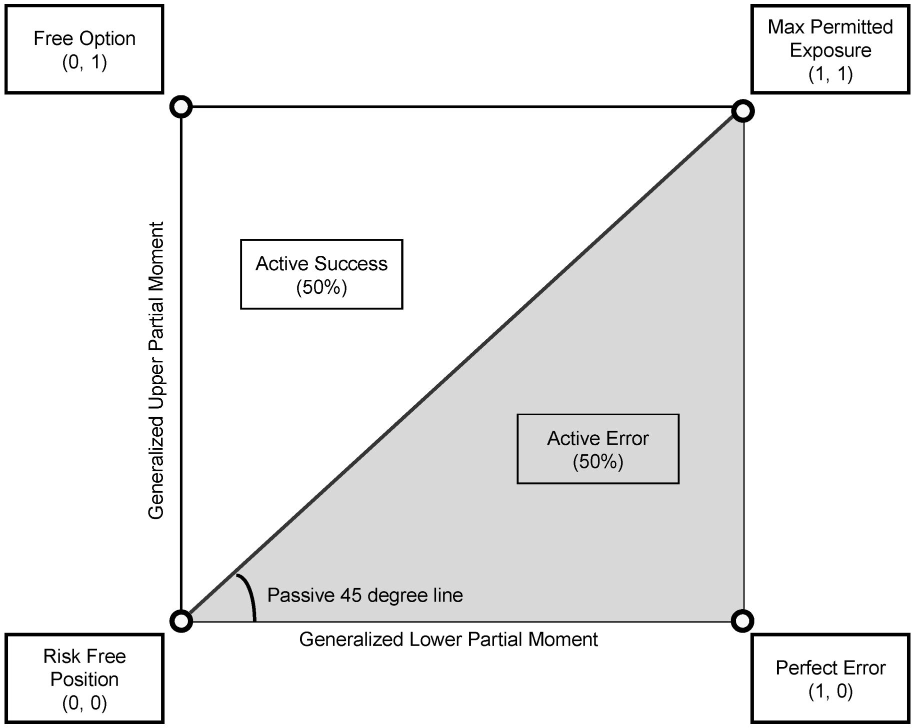

proportional to the retained exposure (= 1 – benchmark hedge ratio).

Note that the MCA is determined for all passive benchmark hedge ratios, but not for the outcomes of active management. The appropriate shift for active outcomes should always refer to the MCA normalization that corresponds to the benchmark against which the active manager is measured. This preserves the displacement of the active manager's outcome relative to the benchmark.[6] Normalization removes the market bias, revealing a genuine active skill that is the remaining perpendicular distance from the forty-five-degree normalized passive line.

GCRA – Optionality and Visibility

After normalization, market bias is removed and the CRA diagram again becomes a symmetric square. However, the side of the square is 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 a 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 generalized conditional risk attribution (GCRA).

In the GCRA diagram (see Figure 5), the vertical axis shows the generalized upper partial moment, which indicates the proportion of available market gains that were captured. The horizontal axis shows the generalized lower partial moment, which displays the proportion of market losses 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 forty-five-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 decision making by measuring it against this forty-five-degree line. We hereby define the degree of active information as "optionality," which is quantified by a perpendicular distance from the forty-five-degree line. The top left corner is the free option (0, 1), and the bottom right corner is perfect error (1, 0).

Note that the ideal 100% successful outcome from active management is a free option. This is not the minimum variability of expected return, the outcome sought by traditional mean/variance approaches. In fact, it is the outcome with the minimum uncertainty

FIGURE 6
Maximum and Minimum Optionality

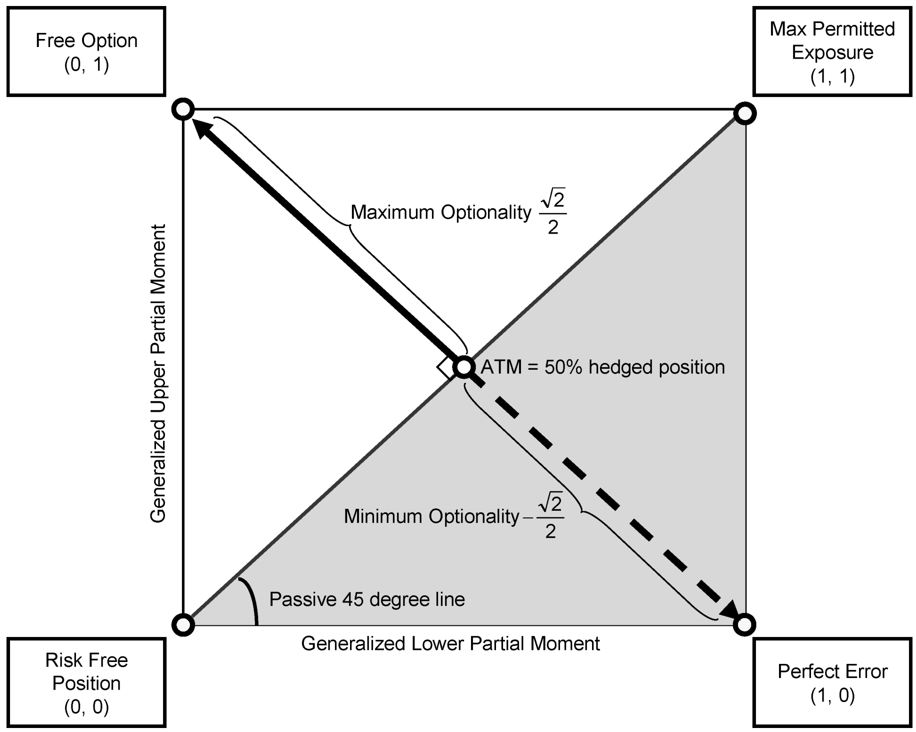

of loss and the *maximum* uncertainty of gain! This highlights once again how essential it is to distinguish between loss and gain in order to assess active manager performance.

It is no coincidence that Figure 5 resembles the receiver operating characteristic (ROC) diagrams used in decision theory (see Marques de Sa [2001]). An ROC diagram displays the trade-off between the proportions of false positives and false negatives. The 0-th generalized partial moments show exactly this information for an investment process. But GCRA goes beyond this by plotting the higher moments, which allows investors to measure not just the percentage of investment errors, but the significance of their impact.

Figure 6 shows that the free option represents the maximum optionality ($= \frac{\sqrt{2}}{2}$). Note that, in the first moment, the outcome of a market-traded at-the-money option lies somewhere on the diagonal line connecting the free option with the perfect error, depending on its price. Since the expected return of a fairly priced at-the-money option equals the risk-free return, its outcome lies on the other diagonal. Therefore, the fairly priced option lies at the center of the diagram. Since "at-the-money" means a 50% uncertainty of gains and a 50% uncertainty of losses, a passive option-rolling strategy would appear at the center of the figure (ignoring costs), and would coincide with the position of the passive 50% exposed strategy. GCRA can detect the skill of managers who buy cheap options.

All points on the same line running parallel with the passive forty-five-degree line have the same degree of active information. In other words, 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, as Figure 7 shows, a vertical distance (G – L) is constant for all points on the line. We therefore define the optionality demonstrated by active management as:

$$\text{Optionality (degree of active information)}$$
$$= \frac{\sqrt{2}}{2} \times (G - L)$$

Note that all passive outcomes have an optionality of zero, due to the normalization process used to calculate the generalized partial moments. Optionality is a true measure of an active manager's skill: the ability to modify a distribution of passive outcomes in an asymmetric fashion that captures available gains while avoiding losses. If this is achieved for each moment of the distribution, the positive tail is retained and the negative tail is diminished. This would be the active management equivalent to the CRA outcome observed earlier for a passive crude oil exposure.

Optionality is attributed to each of the moments of the distribution of outcomes by taking the differences between upper and lower generalized partial moments.

FIGURE 7
Definition of Optionality

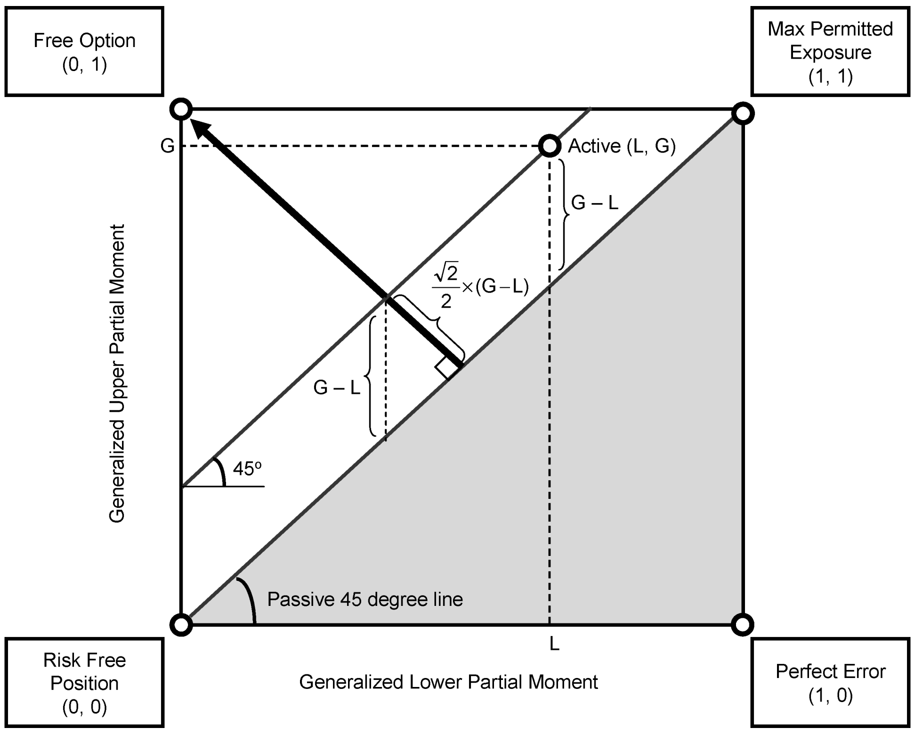

Note that for this to show up in the GCRA diagram, the same total number of sample observations must be used in the original definitions of both lower and upper partial moments. Therefore, a subtraction of the two represents the degree of skew in the entire distribution of all outcomes relative to the risk-free position at each moment.

The degree to which investors notice the impact of the actively managed exposure in the partial moments of the portfolio returns is called "visibility." We define visibility as the sum of exposures to both the upper and lower generalized partial moments. It is zero at the risk-free position, and the distance from the risk-free position to each line running perpendicular to the passive forty-five-degree line represents incremental visibility. As Figure 8 shows, the generalized maximum permitted exposure in the top right corner is the maximum visibility ($= \sqrt{2}$).

All points on the same line running perpendicular to the passive forty-five-degree line have the same degree of visibility. In other words, as we have already observed during the normalization process, the average exposure to both the generalized upper and lower partial moments ($\frac{G+L}{2}$) is constant for all points on that line (see Figure 9). Visibility is therefore defined as follows:

$$\text{Visibility} = \sqrt{2} \times \frac{(G+L)}{2}$$

Visibility has nothing to do with active information. Its value is always defined by the distance from the risk-free position to the point of intersection with the line that runs through the point perpendicular to the passive forty-five-degree line. Thus it is always measured on the passive forty-five-degree line, where no active information is deployed. All decision making on visibility is a passive judgment. This can be due to the average position size an active manager uses, or, as with active hedging, to benchmark selection.

Because the optionality of an active manager is measured as the perpendicular distance from the passive forty-five-degree line, the manager is compared against the passive outcome with equal visibility. But since *all* passive outcomes have zero optionality regardless of their visibility, the optionality of an active manager measures the performance against *all* active outcomes.

Table 2 summarizes the key measurements of optionality and visibility in the GCRA diagram. The calculation of optionality and visibility contains constant parameters $\frac{\sqrt{2}}{2}$ and $\sqrt{2}$. Constant parameters have no impact on comparative analyses, so we eliminate $\sqrt{2}$ from the calculation.

We now rescale the GCRA square diagram by multiplying all points in the diagram by $\sqrt{2}$. By rotating the GCRA diagram forty-five degrees, it becomes a diamond shape (see Figure 10).

FIGURE 8
Maximum Incremental Visibility

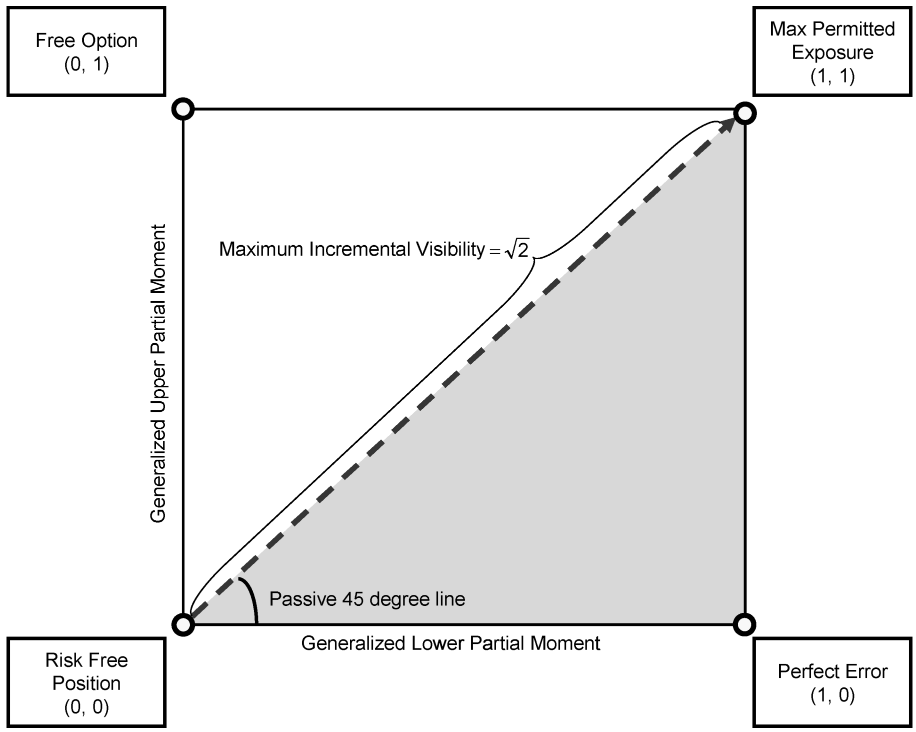

FIGURE 9
Definition of Visibility

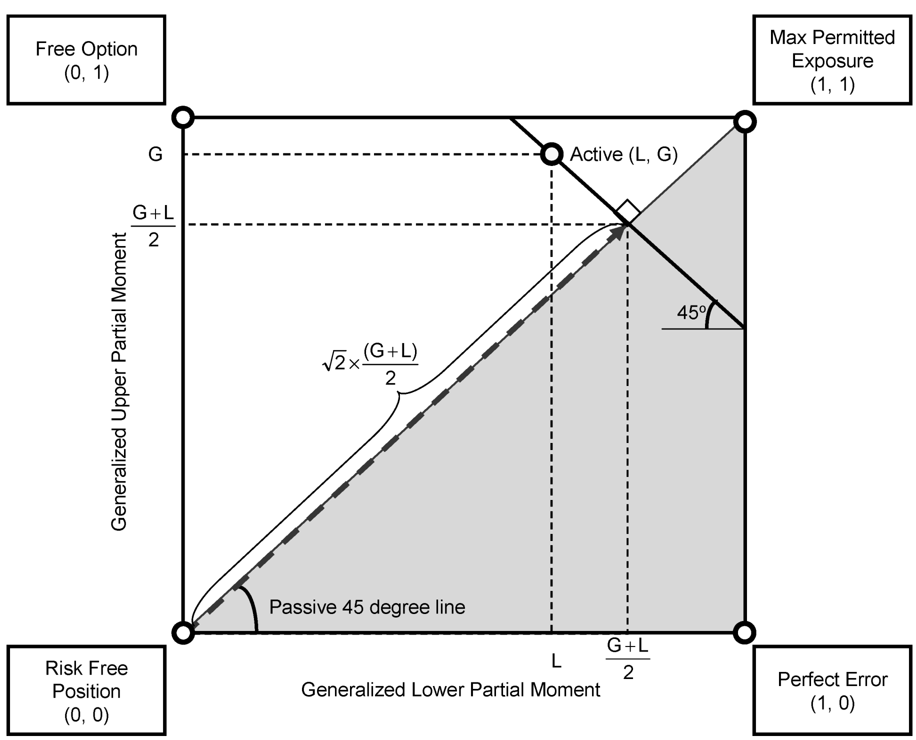

Table 2. *Key Measurements of Optionality and Visibility*

	Optionality	Visibility
Passive 0% Hedged (max permitted exposure)	0	$\sqrt{2}$
Passive 50% Hedged (center point)	0	$\frac{\sqrt{2}}{2}$
Passive 100% Hedged (risk-free position)	0	0
Perfect Foresight (free option)	$\frac{\sqrt{2}}{2}$	$\frac{\sqrt{2}}{2}$
Perfect Error	$-\frac{\sqrt{2}}{2}$	$\frac{\sqrt{2}}{2}$

Table 3. *Redefinition of Optionality and Visibility*

	Optionality	Visibility
Passive 0% Hedged (max permitted exposure)	0	2
Passive 50% Hedged (center point)	0	1
Passive 100% Hedged (risk-free position)	0	0
Perfect Foresight (free option)	1	1
Perfect Error	-1	1

The vertical axis represents optionality, and the horizontal axis represents visibility. We then redefine optionality and visibility in Table 3.

Identifying Active Skill

Note that active information is defined as the difference between the upper and lower partial moments. It is not the ratio of the two, which is the slope of the line drawn from the risk-free position to the active result (L, G). Although this slope may be considered an intuitive measure of active skill, it is a mixture of optionality and visibility.

This is clear in the rotated GCRA diagram, and is a flaw of the ratio-based approach. Modern portfolio theory assumes that both gains and losses are symmetrically distributed, and thus returns per unit of volatility relative to the benchmark represent a proxy of active information. Symmetric returns represent a special case in the GCRA. But, as we discussed above, active information is embedded in the asymmetry of gains and losses, so results derived simply from the returns and volatility ratio are contaminated. This contamination is filtered out of the GCRA diagram, so that optionality changes with the difference between

FIGURE 10
Rescaled and Rotated GCRA

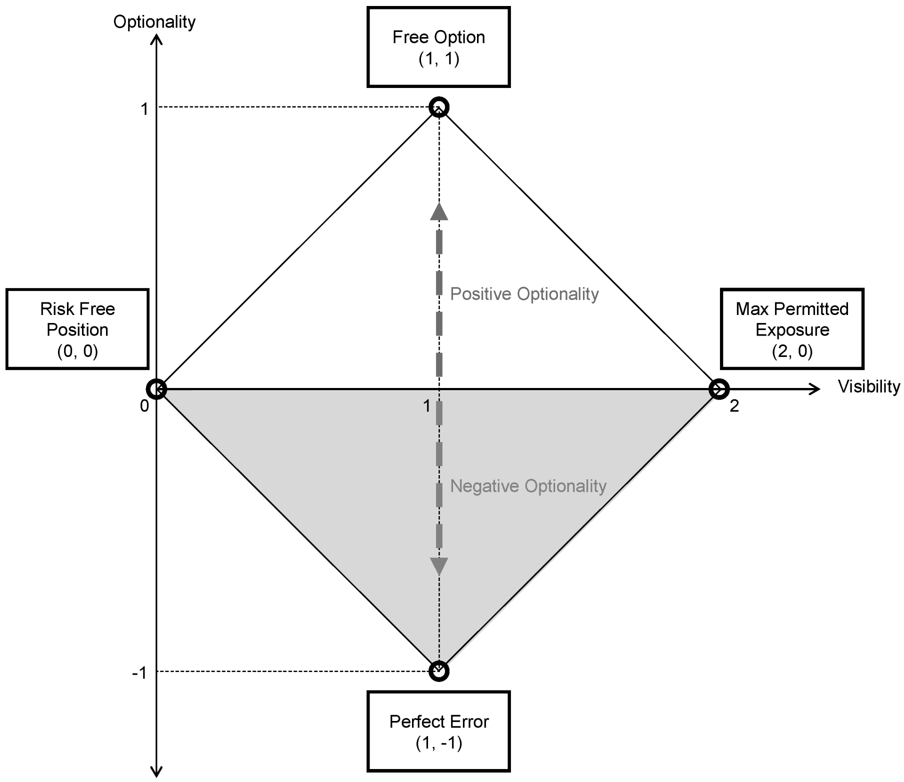

FIGURE 11
Definition of Incremental Optionality and Visibility

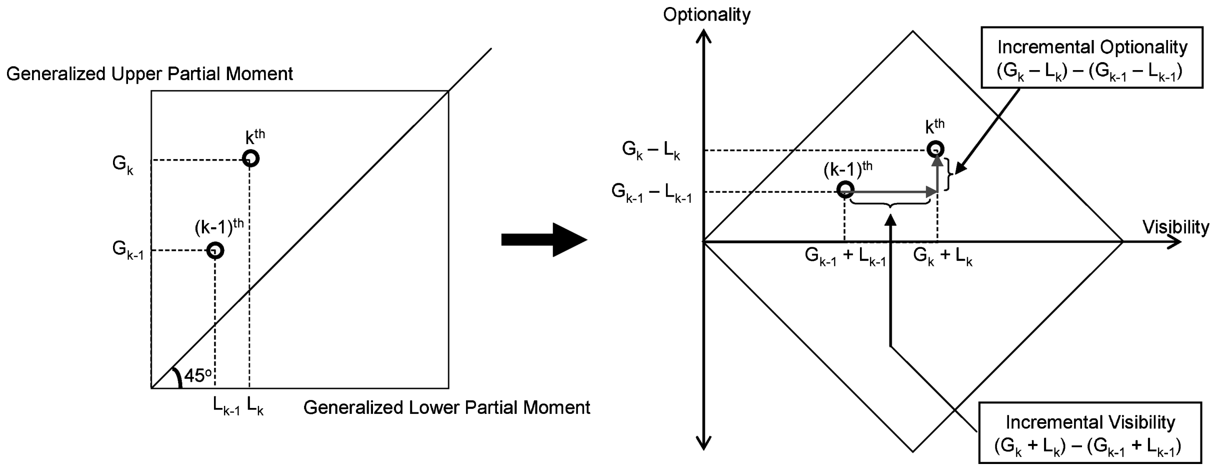

the upper and lower partial moments as a result of separating gains and losses. Because there is no overlap between optionality and visibility, the optionality represents the degree of asymmetry created by active information.

Optionality can be calculated through higher moments. Thus far, we have described partial moments in general terms, without describing how each represents a different characteristic of asymmetry. The first moment represents asymmetric returns, while the second moment contains asymmetric uncertainties of returns. The third moment is impacted more by asymmetric fat tails of returns. We consider only up to the third moment here, as the fourth moment and beyond amplify the impact from fat-tailed returns, and there is little extra optionality information to be gained.

Figure 11 plots different moments on the same GCRA diagram. The vertical displacement from each previous moment to the next represents the "incremental optionality" at each moment. If there is no active information shown, the displacement will be a horizontal shift parallel to the passive visibility axis. A displacement other than a horizontal shift indicates the existence of incremental optionality, either positive or negative in the higher moment.

We define incremental optionality as the change in optionality calculated using the kth partial moments versus the optionality calculated using the $(k-1)$th partial moment.

1st moment optionality (skew in returns) =

$G_1 - L_1$

2nd moment *incremental* optionality

(volatility control) $= (G_2 - L_2) - (G_1 - L_1)$

3rd moment *incremental* optionality

(surprise control) $= (G_3 - L_3) - (G_2 - L_2)$

kth moment *incremental* optionality =

$(G_k - L_k) - (G_{k-1} - L_{k-1})$

Incremental visibility at the kth moment can be similarly defined. A horizontal transition from each previous moment to the next represents incremental visibility at each moment. The degree of incremental visibility at the kth moment is defined by the horizontal distance between the kth moment and the $(k-1)$th moment.

1st moment visibility (returns) =

$G_1 + L_1$

2nd moment *incremental* visibility (volatility) =

$(G_2 + L_2) - (G_1 + L_1)$

3rd moment *incremental* visibility (tail sensitivity) =

$(G_3 + L_3) - (G_2 + L_2)$

kth moment *incremental* visibility =

$(G_k + L_k) - (G_{k-1} + L_{k-1})$

Style Characterization by Optionality Sequences

GCRA measures the extent to which active returns exhibit option-like distributions. Thus it is no surprise that sets of points representing the moments of the active returns in the GCRA diagram correspond to a representative payoff profile in a conventional payoff chart. Visibility is connected to the slope of the

Table 4. *Sequence of Incremental Optionality*

	Investment Success				Investment Error			
	Case 1	**Case 2**	**Case 3**	**Case 4**	**Case 5**	**Case 6**	**Case 7**	**Case 8**
1st Moment	+	+	+	+	−	−	−	−
2nd Incremental	+	+	−	−	−	−	+	+
3rd Incremental	+	−	+	−	−	+	−	+

payoff profile, while optionality measures its curvature. The incremental visibility and incremental optionality describe subtleties in the payoff profile by providing more detail about the relative shape of the tails of the return distributions.

Positive optionality indicates the existence of positive active information. The question, however, is how stable and reliable is the optionality? Even after removing market bias, some pure luck resulting from "market timing" might show up as optionality. Market timing can create positive asymmetry in the first moment, but it comes at the expense of negative asymmetry in the higher moments because the first moment is the most susceptible to pure luck. If incremental optionality turns negative toward higher moments, it will tend to be less meaningful even if the first moment is positive, because of the existence of negative tails. On the other hand, if incremental optionality consistently stays positive toward higher moments, this suggests it is stable and reliable.

Visibility is a barometer to assess how far active exposures deviate from the risk-free position. The first moment represents only the average risky exposure taken to deliver the return. Over the course of delivering such returns, further exposures must be taken in higher moments. The second moment explains the incremental exposure observed in volatility; the third moment shows a further incremental exposure in the tails. Increased exposure to higher moments increases the chance of a negative surprise.

But there is a relationship between optionality and visibility, particularly at higher moments. An investment process with a high level of optionality can capture available gains while avoiding losses, which effectively controls the strategy's visibility. Strategies that do not exhibit optionality in the higher moments are vulnerable to highly visible losses. This means that the transition of incremental visibility in the higher moments can be a useful indicator of the reliability of the optionality.

It is possible to observe the incremental optionality of an investment process by plotting the higher moments on the same GCRA diagram. This allows us to identify which moments are managed by an active strategy by observing the sequence of optionality over higher moments. Because we are only considering the first through third moments, there are eight possible sequences of optionality (see Table 4).

We can define four different active management styles based on the characteristics of their moments. A sequence characterized by the first moment is a "return-driven" style. Characterization by the second moment indicates a "volatility-driven" style, while characterization by the third moment suggests a "tail-driven" style. If the sequence is defined by both the second and third moments, the active style is an integrated case of a "higher moment-driven" process.

In the short run, "return-driven" styles are characterized by either investment success (Case 1), or investment error (Case 5). However, only the first moment is directly managed. No active skill is deployed at higher moments, which means episodic optionality at higher moments results only from luck. Since extreme events are not managed, it is only a matter of time before optionality in the higher moments converges into the negative.

The "return-driven" style aims to achieve positive optionality in the first moment to compensate for negative optionality at higher moments. Over the long run, the "return-driven" approach is characterized by Case 4 (success at the first moment only), or by Case 5 (errors at all moments). The "return-driven" strategy depends on accurate forecasts of market timing. Because market timing is dominant, optionality is less stable and reliable.

The "volatility-driven" approach is characterized by investment success (Case 2), and investment error (Case 7). "Tail-driven" strategies are characterized by investment success (Case 3), and investment error (Case 6). In both cases, no active skill is deployed at the first moment, resulting in only episodic optionality at the first moment. However, higher moments are directly managed.

If a strategy can produce an integrated "higher moment-driven" result, all optionality beyond the first moment will tend to be positive in the long run. In this case, the "higher moment-driven" skill aims to achieve consistently positive optionality at higher moments to compensate for episodic optionality at the first moment, which depends on the market bias. Therefore, in the long run, "higher moment-driven" strategies are characterized by Case 1 (success at all moments), and

Table 5. *Characterization of Incremental Optionality by Sequence*

	Return-Driven		Volatility-Driven		Tail-Driven		Higher Moment-Driven	
1st Moment	+	−	+	−	+	−	+	−
2nd Incremental	−	−	+	+	−	−	+	+
3rd Incremental	−	−	−	−	+	+	+	+

Case 8 (errors at the first moment only). This implies that "higher moment-driven" strategies depend on effective management of the probability of loss. Because market timing is less significant at higher moments, optionality is stable and reliable. Table 5 summarizes the style characterization.

The set of opportunities available for active management is affected by the degree of short-term market bias. Any 0-th moment optionality represents skew in available opportunities. In a strongly biased market environment, the opportunity for successful active management can temporarily diminish. During such periods, even with the "higher moment-driven" style, negative optionality in the first moment may coincide with negative optionality at higher moments. Such inconsistency in style and optionality sequence in a particular moment will be resolved in the long run, however, indicating that a short-term analysis of optionality may be unreliable. Long-term analyses are critical for obtaining reliable and robust results.

GCRA Performance Assessment

Table 6 summarizes the GCRA performance assessment. In the GCRA diagram, optionality explains the degree of active skill, while visibility explains how much the active result is due to its benchmark exposure level. The GCRA analysis is not biased by any specific data because it is generalized by adjusting for market environment. Therefore, these statistics are independent of the underlying exposure, the benchmark, the time period, and so on. The generalized unit square permits us to compare different managers managing different investment exposures over different time periods by specifically focusing on optionality, i.e., the ability of an active investment decision maker to enhance the balance between loss and gain.

If two different managers achieve identical optionality, investors might prefer the one who achieved visibility closer to its benchmark, because this manager's investment style is more relevant to its benchmark. Alternatively, if investors are more concerned with the exposure attained after active management, visibility might determine investor preference in manager selection. However, if superior manager skill is the top priority, stable and reliable optionality, rather than visibility, should be the primary concern, because visibility does not contain any active information.

The sequence of optionality in the management of higher moments indicates how dependent it is on market timing, and consequently how stable and reliable it is. Based on the sequence of optionality up to the third moment, Figure 12 characterizes our four different active "driven" styles.

The style characterization is useful in identifying the need for diversification. Many investors blindly believe that diversification is always a useful and effective tool, but that is not the case. Looking at Table 5, the "return-driven" style relies primarily on market timing. If a manager has skill at the first moment only, the rest of the higher moments are exposed to market systematic risk. This negative tail risk can be reduced to some extent by diversification.

This is in contrast to investment styles with expertise in directly managing the higher moments, which demonstrate stable and reliable skill. Like any strategy, these approaches may also be susceptible to unexpected failure of active skill. However, the likelihood of negative tails is directly managed under the higher moment control, almost like insurance. The GCRA approach implies that investors must diversify the embedded systematic risk in the "return-driven" strategies by continuous selection of a spread of managers. But if investors can identify a skillful manager with a "higher moment-driven" strategy, there is no loss-controlling benefit from diversification. This renders the need for

Table 6. *GCRA Performance Assessment*

	Active Result	Benchmark	Assessment
Optionality	G – L	0	• Active Skill
Visibility	G + L	• Max Permitted Exposure = 2 • Risk-Free Position = 0	• Use of Available Range • Benchmark Relevance

FIGURE 12
Active Style Characterization

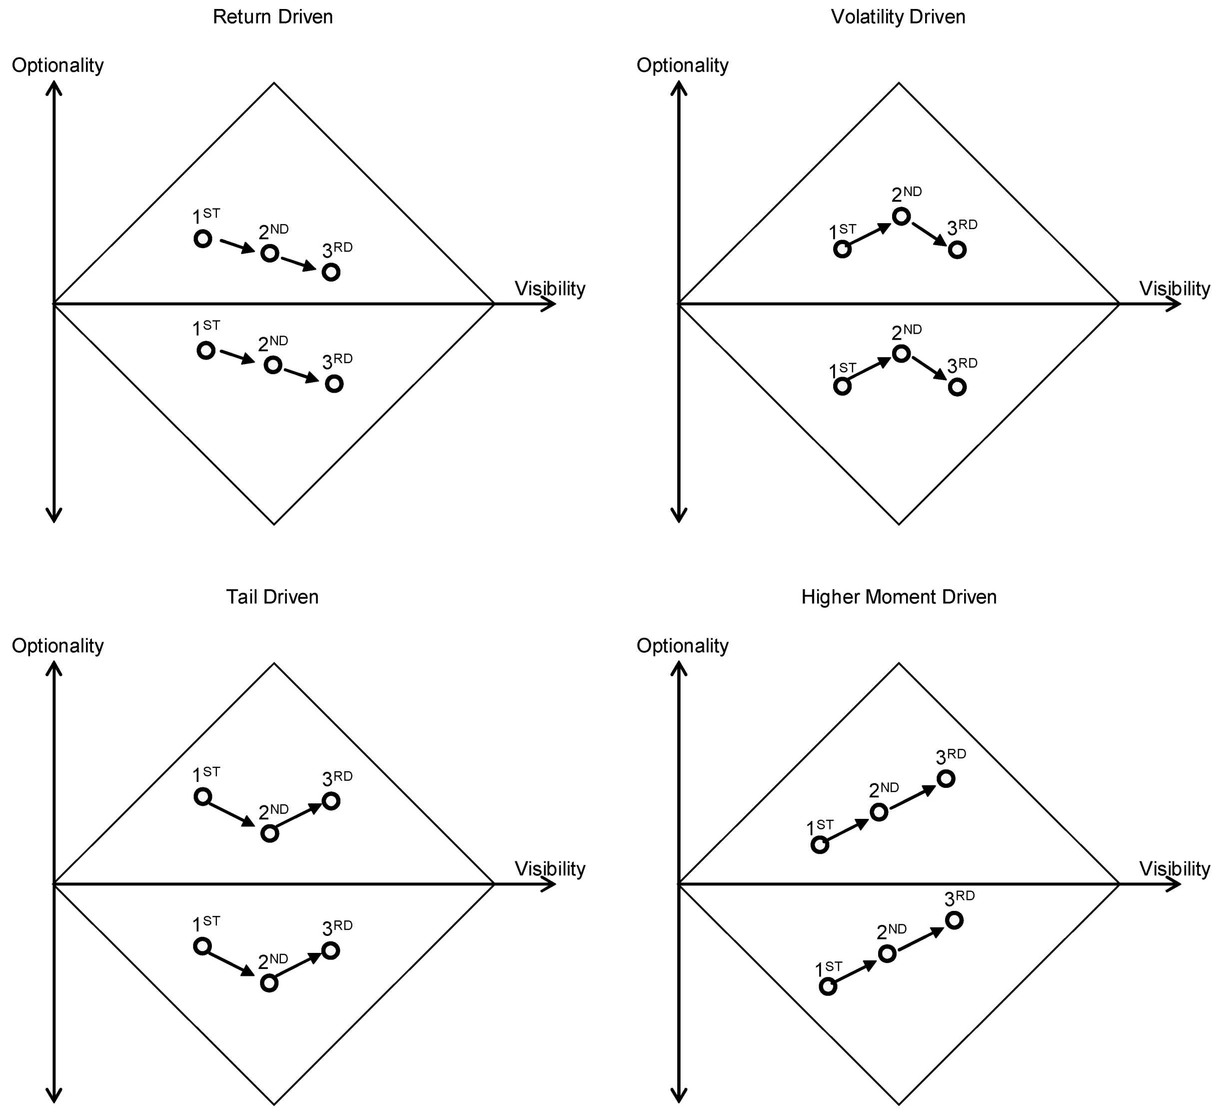

continuous selection minimal because performance is more stable and reliable.

Conclusion

This paper describes the use of 1) conditional risk attribution to analyze passive risk exposures, and 2) generalized conditional risk attribution to analyze active investment outcomes.

The formal definitions of partial moments are fundamental to the application of this approach to investments that are defined by maximum permitted active risk exposure. Maximum permitted exposure is a more certain and transparent boundary constraint than tracking error, which can only provide an uncertain loss limit. Nevertheless, maximum permitted exposure can be biased by market environment.

The elimination of the market bias inherent in a passive investment allows active outcomes to be displayed in a unit square. GCRA clearly separates active investment outcomes into two parts: active skill, defined as *optionality*, and the degree of passive impact from a retained exposure, defined as *visibility*. 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.

Positive optionality means positive active skill. Optionality varies through higher moments. Incremental optionality between moments represents active skill at different moments. Active styles are characterized by the sequence of incremental optionality. This style characterization highlights an embedded reliance of active managers on either market timing, i.e. pure luck, or on an ability to transform uncertain and unbiased outcomes into certain and biased outcomes.

The approach we describe here seeks out the ability to create certainty out of uncertainty in an investment decision-making process. Future research could address the far broader applications we believe are apparent, such as policy asset mix and subsequent tactical portfolio rebalancing.

Notes

1. Our objective is to distinguish asymmetry from symmetry, so it is important to avoid performance metrics that contain asymmetric properties such as geometric returns, where a 25% gain is offset by a 20% loss.
2. The risk-free rate in the general sense was defined by Sharpe [1994] as the risk-free benchmark portfolio against which the investment is evaluated.
3. For convenience, we assume that a zero return is considered a gain.
4. The risk-free rate used for U.S. dollar-denominated investments (the S&P 500 Value & Growth, the Citigroup U.S. World Government Bond Index, Lehman U.S. Aggregate, the CRB Commodity Index, the WTI, gold) was the JPMorgan U.S. dollar cash one-month rate. The rate for the Japanese yen-denominated investments (the Nikkei 225, Citigroup Japanese World Government Bond Index) was the JPMorgan Japanese yen cash one-month rate. To determine the risk-free rate adjustment for the U.S. dollar-Japanese yen exchange rate, we use the interest rate differential between the dollar cash one-month rate and the yen cash one-month rate.
5. Because the degree of normalization is proportional to the exposure, the required geometrical transformation is a shear, not a rotation or a translation.
6. 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 Stoyanov, "Comparison Among Different Approaches for Risk Estimation in Portfolio Theory." Lehrstuhl für Statistik, Ökonometrie und Mathematische Finanzwirtschaft, Universität Karlsruhe (TH) 2004.

Fishburn, Peter C., "Mean-Risk Analysis with Risk Associated with Below-Target Returns." *American Economic Review*, Vol. 67, 2, 1977, pp. 116-126.

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

Okuyama, N., and G. Francis, “Disentangling Cognitive Bias in the Assessment of Investment Decisions.” *Journal of Behavioral Finance* 2006.

Sharpe, William F., “The Sharpe Ratio.” *Journal of Portfolio Management*, Fall, 1994.

Sortino, Frank A., and Lee N. Price, “Performance Measurement in a Downside Risk Framework.” *Journal of Investing*, Fall, 1994.