Informationless Investing and Hedge Fund Performance Measurement Bias

Dangerous attractions.

Andrew B. Weisman

Hedge funds have emerged as an important investment category for institutional investors. They are attractive because they promise superior, non-correlated, rates of return compared to traditional industry benchmarks. This popular perception of hedge funds, in conjunction with the institutional desire to improve risk-adjusted performance, typically characterized by some form of mean-variance efficiency, has led to dramatic capital flows into this investment class.

Institutional investors have not, however, always had pleasant experiences with their hedge fund investments. For example, Brown, Goetzmann, and Ibbotson [1999] observe in their examination of off-shore hedge funds that the hedge fund industry is characterized by very high rates of attrition, estimated to be about 20% per year, as compared to the approximate 5% rate for mutual funds.[1]

The failure to meet investor expectations is typically the result of two related factors. First, many institutional investors, investment consultants, and academicians lack a basic understanding of the return-generating processes of many hedge funds. Second, because of this basic lack of understanding, there tends to be an over-reliance on conceptual frameworks and technologies that are appropriate to the traditional investment world, but highly inappropriate for hedge funds.

To avoid such pitfalls, it is useful to consider some of the basic investment techniques that are widely employed in the hedge fund industry, in particular, a

Andrew B. Weisman is a managing partner at Strativarius Capital Management in New York (NY 10152).
weisman@strativarius.com

widely used class of investment techniques here referred to as "informationless investment strategies." Informationless strategies tend to produce return enhancements over relatively long periods even though they frequently provide no *theoretical* long-term benefit.

The most important consequence of such strategies is that they tend to systematically bias statistically derived performance measures, such as mean, variance, and measures of association. Subsequently, quantitative optimization will tend to systematically worsen overall portfolio performance in the context of hedge fund investing.[1] Indeed, as demonstrated below, when they select managers to maximize an ex post measure of risk-adjusted return, portfolio managers may be virtually guaranteeing a bad outcome.

My purpose, therefore, is to present three specific informationless investment strategies peculiar to the asset management industry in general, and the hedge fund industry in particular, and their consequences with respect to performance measurement and asset allocation.

SHORT-VOLATILITY INVESTING

The first informationless strategy relates to the reasonably common practice of structuring investments that are essentially equivalent to writing insurance policies against low-probability events, i.e., *short-volatility investing*. Short-volatility investing is typically operationalized by using derivative securities that possess optionality. Options (or certain active management strategies that mimic them) permit a trader to collect a premium for assuming the risks associated with low-probability events.

A wide variety of hedge fund investment strategies derive their returns from short-volatility investing. These investment strategies typically involve the purchase of one or more securities and simultaneous short sale of one or more securities, where the long security is viewed to be undervalued relative to some perceived equilibrium relationship with respect to the short security. Positive payouts accrue to the investor as the relative valuations of the securities converge to the perceived equilibrium, while losses accrue as the relationship becomes increasingly strained; thus the term short-volatility investing.

Strategies such as merger arbitrage, various forms of fixed-income arbitrage, and statistical arbitrage (pairs trading) can all be classified as short-volatility investing programs.

Short-volatility investments are typically initiated when the relationship between the long and short securities is estimated to be at an extreme valuation, so a continuation or further straining of the relative valuation is determined to be a low-probability outcome. In fact, the tendency to structure individual investments with a high probability of a successful outcome is a hallmark of such strategies.[2]

Such investment strategies are usefully thought of as a process of selling insurance policies written against perceived low-probability events. Viewed in such terms, the general performance characteristics of short-volatility investing become analytically tractable, and, most important, it can be demonstrated that short-volatility investment strategies can be easily constructed that appear to provide performance enhancement for reasonably long periods, without in fact doing so. In so doing, such strategies can systematically bias statistically derived estimates of risk, return, and association.

Sample Short-Volatility Investment Program

To clarify this point, consider an investment strategy I discuss in Weisman [1998]. Assume the current risk free-rate is 5%. A hypothetical manager invests all of his or her capital at the risk-free rate. At the beginning of every month, the manager writes (sells) a series of fairly valued calls and puts that expire at the end of the month. The strike prices are, respectively, 2.5 standard deviations (with respect to the prevailing market volatility) above and below the current market price of some unspecified financial instrument. The manager writes (sells) a sufficient number of these strangles so that in the event the market remains within the 2.5 standard deviation collar, the manager will take in enough premium to double the risk-free rate.

Using Monte Carlo simulation, we can define the probabilities associated with various related outcomes. Exhibits 1 and 2 depict two randomly generated five-year outcomes for this investment strategy.

The performance of this investment strategy can be summarized as follows. The manager has (approximately) an 88% chance of outperforming the risk-free rate in any year, and almost an 86% chance of doubling it. The manager has an almost 50% chance of doubling the risk-free rate over any five-year period. The expected time to a "volatility event" (when the underlying security trades outside the collar by any month-end, resulting in a loss of capital) is almost seven years.

As we note, the options are assumed to be fairly valued, so the "informationless" process of selling options is

EXHIBIT 1
Randomly Generated Five-Year Performance—T-Bill versus Short-Volatility Strategy

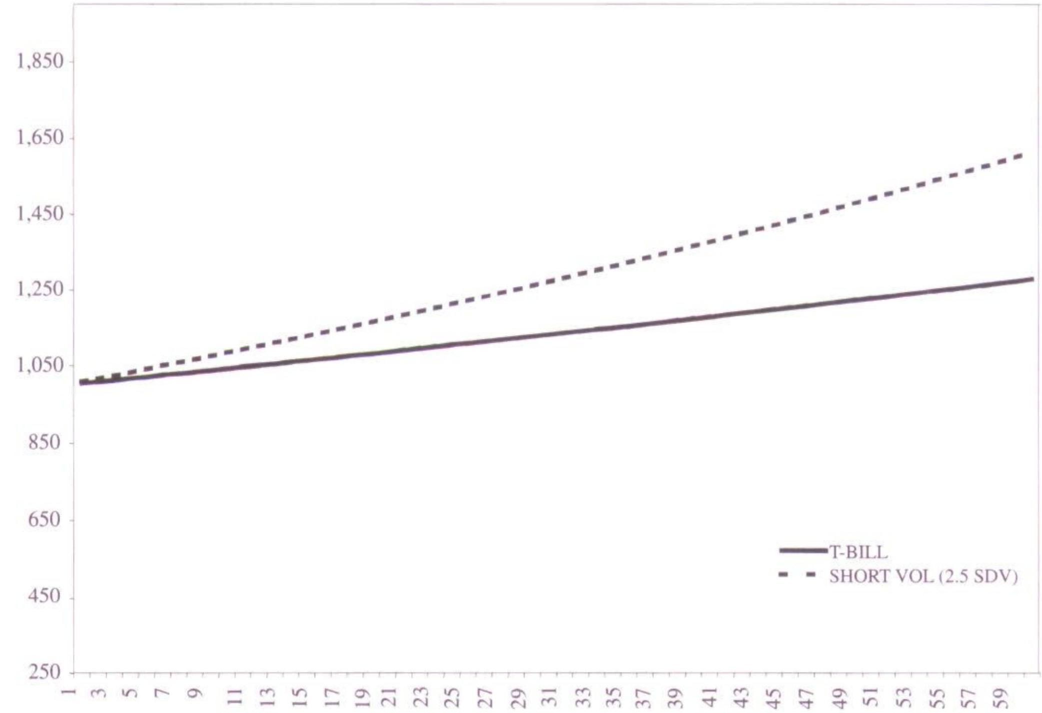

EXHIBIT 2
Randomly Generated Five-Year Performance—T-Bill versus Short-Volatility Strategy

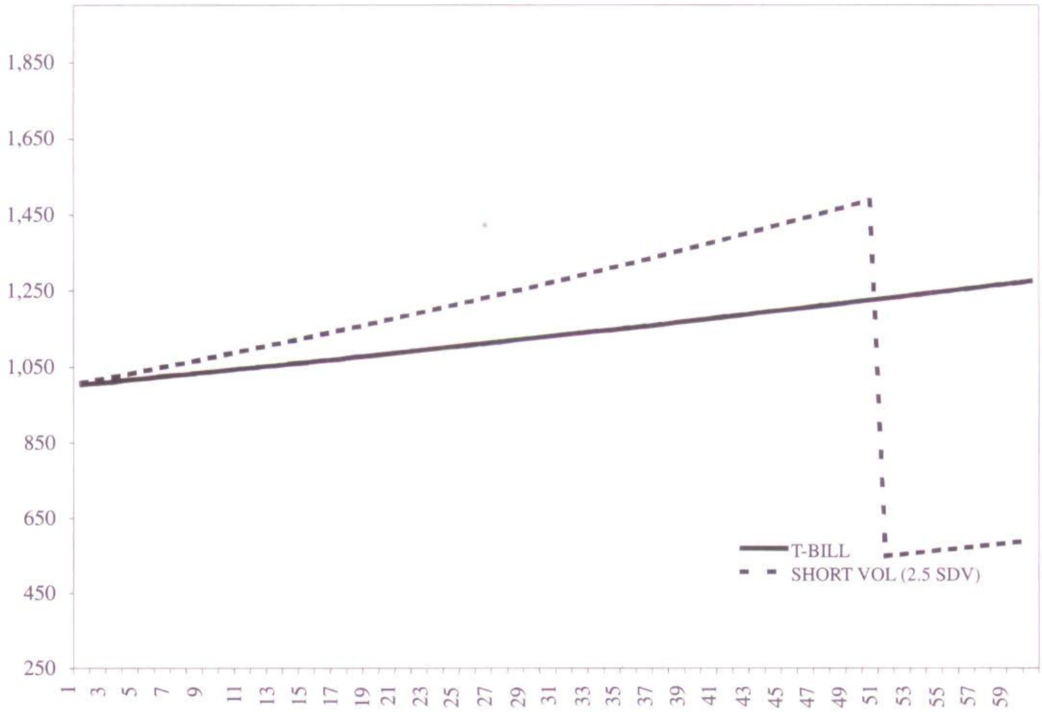

assumed to have a zero expected value, and an equalizing event is therefore necessitated. The equalizing event is that when a volatility event occurs the expected loss of capital is approximately 32%.

This example can be extended by including additional options with different strikes in order to clarify the relationship between the probability of outperforming the risk-free rate in any year and the extent of the expected loss of capital. Exhibits 3 through 6 present the results of this analysis, while Exhibit 7 summarizes the relationship between the probability of outperforming the risk-free rate in any year and the magnitude of expected future periodic loss of capital.

Exhibit 7 illustrates that as the probability of outperforming the risk-free rate increases, the extent of the anticipated loss of capital grows at an increasing rate. This graph illustrates one of the most serious issues associated with the interpretation of hedge fund performance data. For managers who use short-volatility strategies, a stellar performance history, characterized by (for example) a high Sharpe ratio, may be an indication of a very high degree of assumed risk. Most important, statistically derived estimates of the manager's risk-return characteristics will be diametrically incorrect, i.e., high in-sample risk-adjusted returns may imply poor out-of-sample performance.

As per Exhibit 7, in the context of short-volatility investments, it is rational to argue that the strategy of selecting managers by maximizing an ex post measure of risk-adjusted return is, in fact, a *negative selection* process. Such a conclusion is especially likely when the track records are brief enough to exclude a major volatility event.

This reality is clearly demonstrated in the example, where it is shown that an investment strategy could be devised that is simultaneously constrained to provide no long-term performance enhancement and a high likelihood of generating high risk-adjusted rates of return for a fairly significant period of time.

Perhaps the most startling conclusion with respect to short-volatility investments is that very high, statistically derived, estimates of risk-adjusted return can be directly linked with an increasing probability of an unacceptably large loss of capital.

EXHIBIT 3
Probability of Outperforming Risk-Free Rate of Return

	Time Period		
Distance to Strike	1 Year	3 Year	5 Year
Standard Deviations			
1.50	0.57	0.50	0.47
2.00	0.70	0.58	0.52
2.50	0.88	0.73	0.65

EXHIBIT 4
Probability of Doubling Risk-Free Rate of Return

	Time Period		
Distance to Strike	1 Year	3 Year	5 Year
Standard Deviations			
1.50	0.18	0.01	0.00
2.00	0.57	0.19	0.06
2.50	0.86	0.63	0.46

EXHIBIT 5
Expected Time to Draw-Down (Capital Loss)

Distance to Strike	Length of Time
Standard Deviations	Years
1.50	0.80
2.00	1.97
2.50	6.83

EXHIBIT 6
Expected Draw-Down (Capital Loss)

Distance to Strike	Percent Loss of Capital
Standard Deviations	
1.50	-2.96
2.00	-9.00
2.50	-31.92

Short-Volatility Regression Bias

Short-volatility investing also severely complicates the process of determining a measure of association between a manager's returns and likely return-generating factors. I show below that regression analysis is unlikely to reveal the importance of the association between a short-volatility manager's performance and movements in the price and volatility of the underlying asset class traded. As long as the analyzed period includes no major volatility events, i.e., as long as the market remains within the collar, the manager's outcomes will be positive regardless of market direction.

EXHIBIT 7
Probabilities of Outperforming Risk-Free Rate of Return versus Expected Draw-Down

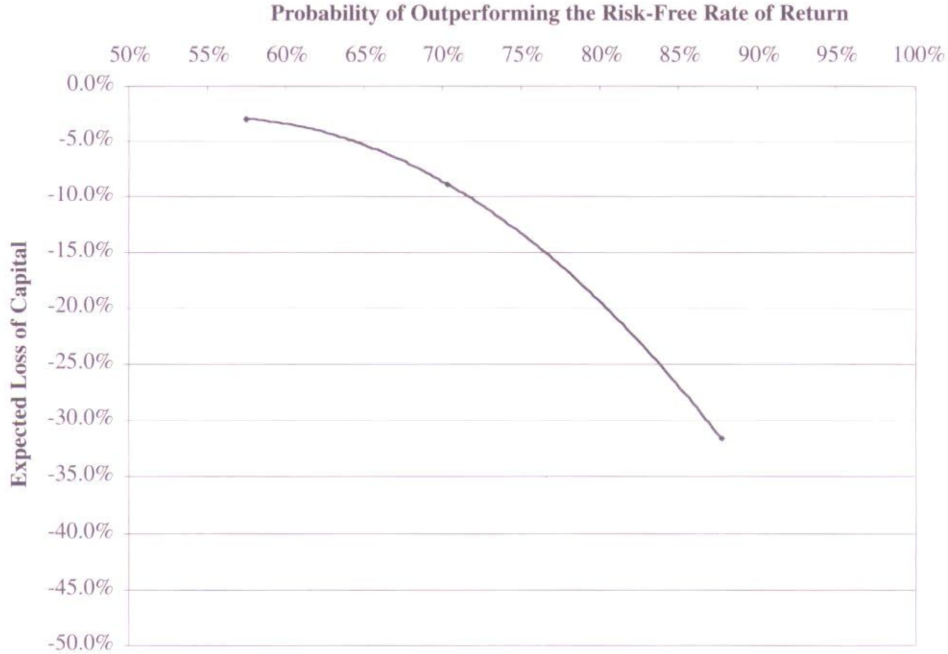

Exhibit 8 should serve to clarify this issue. The first half of the randomly generated market performance presented in Exhibit 8 depicts a positive trending market without any major volatility events, while the second half depicts a negatively trending market without any major volatility events.

With no major volatility events, the manager's returns will be statistically positively associated with the market for the first half, negatively associated for the second half, and unrelated for the entire period.

Most important, however, a statistically derived measure of association is unlikely to adequately describe the highly elastic response the manager's returns will exhibit during a sharply down-trending market; that is, derived regression coefficients will underestimate the tendency for the manager to become highly correlated during such turbulent conditions.

This point has very serious implications for both hedge fund investors and academicians who are attempting to analyze hedge fund performance, and probably necessitates a reexamination of much of the research.[3]

Short-Volatility Summary

The net result is that portfolio managers who naively make use of certain standard optimization strategies, in conjunction with statistically derived inputs, will tend to systematically overallocate to managers who have a short-volatility profile and systematically maximize a future period loss. The tendency for portfolio managers to overallocate to such investment strategies is here referred to as short-volatility bias.

Short-volatility bias is a direct result of an overestimate of the manager's risk-adjusted returns and an underestimate of the manager's correlation during volatile market conditions.

ILLIQUID SECURITY INVESTING

The second informationless investment technique simply involves expressing basic market exposures using illiquid securities. To better understand the consequences of this simple informationless strategy, consider a simple two-manager world. Managers 1 and 2 operate investment programs that have precisely the same performance characteristics, except that Manager 2, due to the illiquidity of the securities in her portfolio, is unable or unwilling to accurately value the portfolio on a periodic basis. Manager 2 therefore employs the simple informationless strategy of systematically understating both the periodic increases and decreases in value of the portfolio and subsequently generates the appearance of performance enhancement.

EXHIBIT 8
Short-Volatility Regression Bias

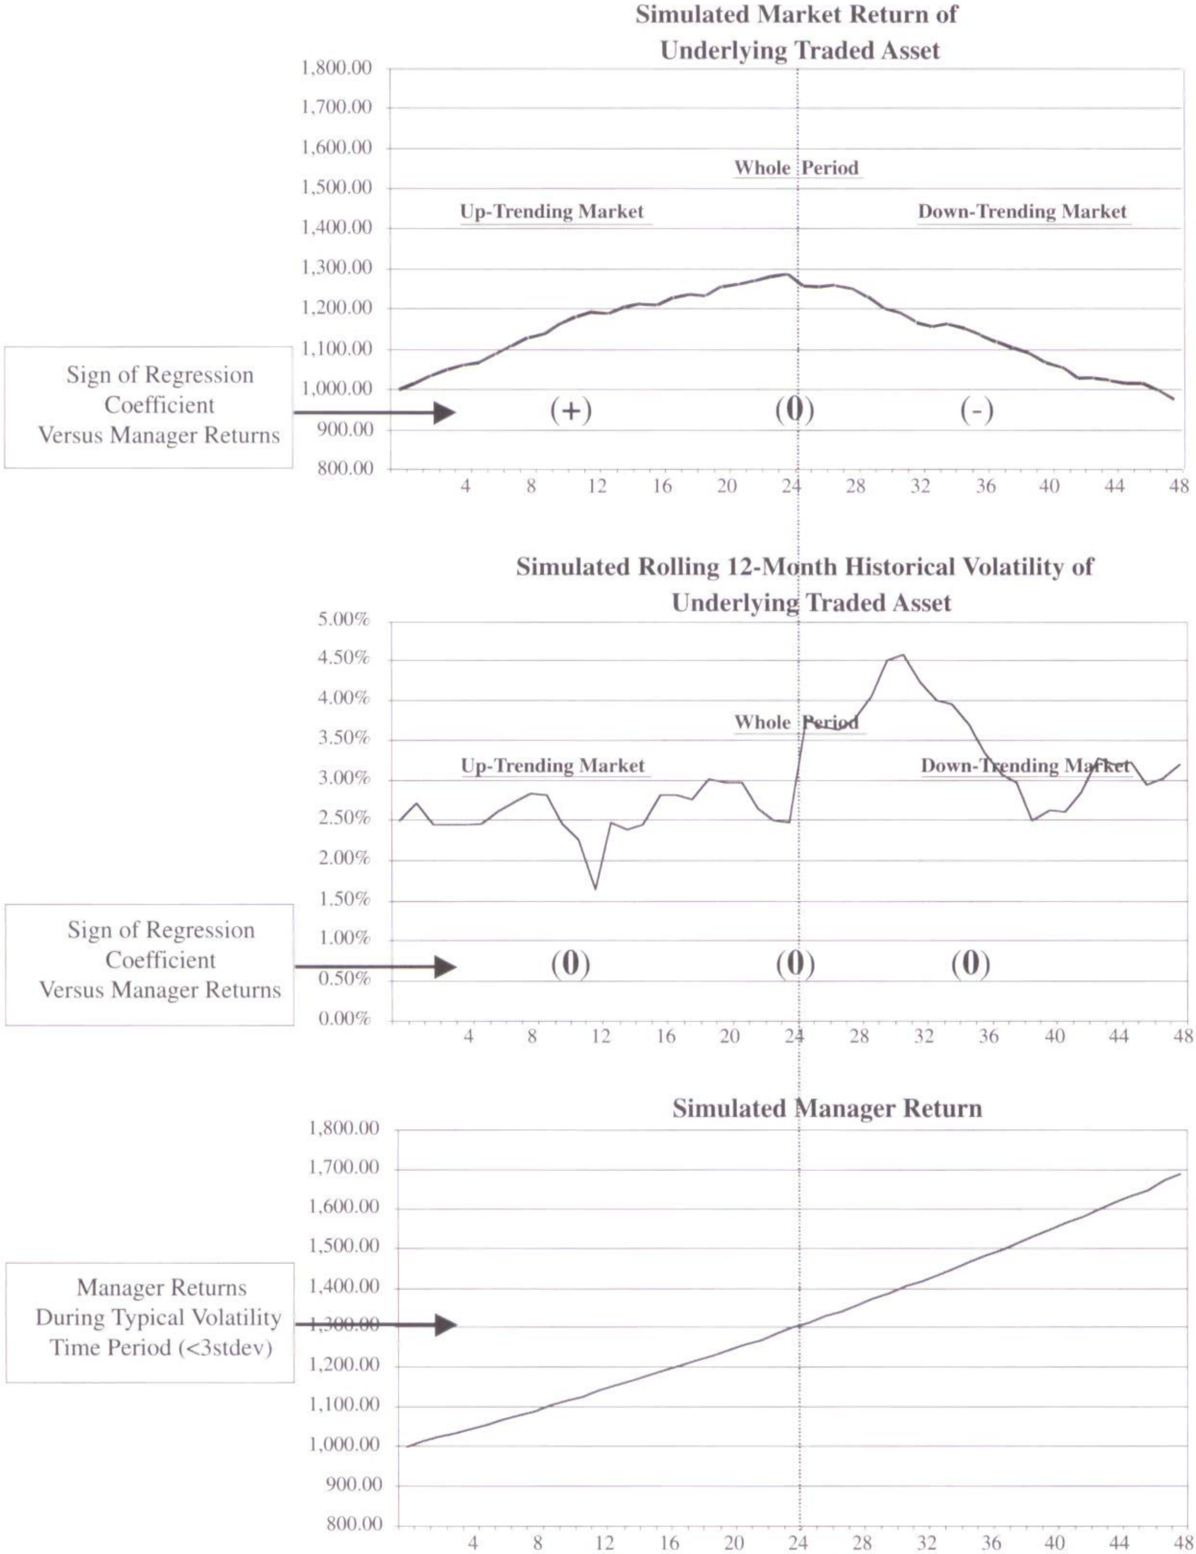

EXHIBIT 9
Illiquidity Bias

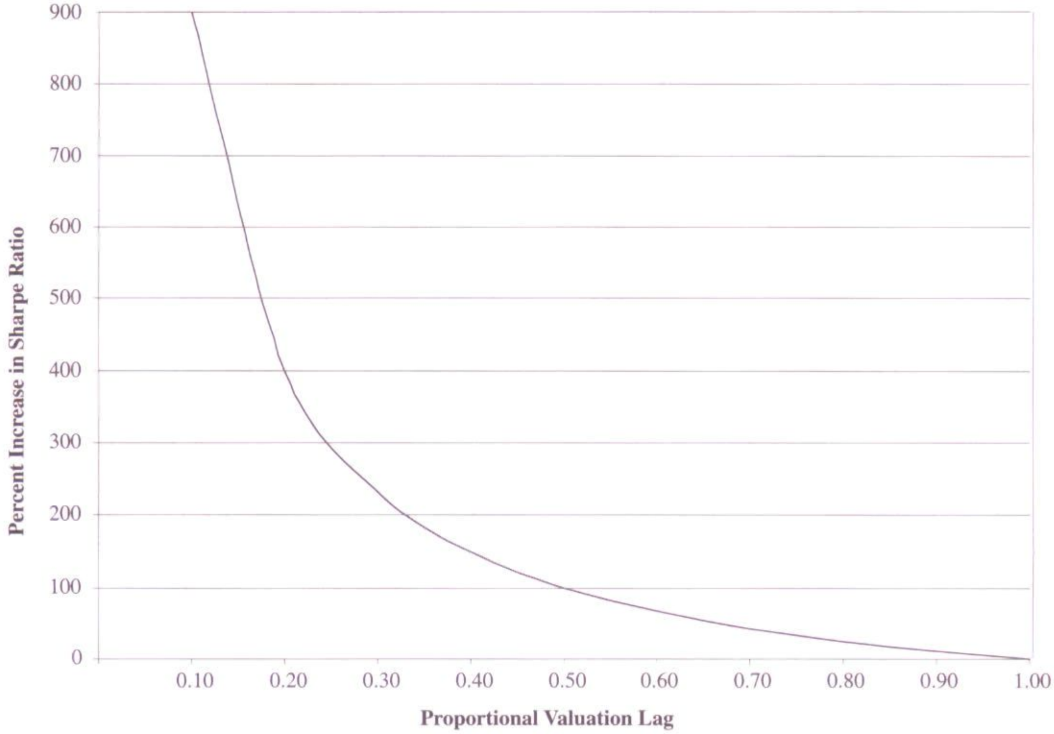

Assume notation as follows:

$\bar{x}_i$ = trend (average) return for Manager i;
σ_i = reported standard deviation of returns for Manager i;
$\bar{x}_i = \bar{X}$;
$\sigma_i = \sigma$; and
δ = proportion of the standard deviation of return that is reported by the manager, referred to as proportional valuation lag, where $0 \leq \delta \leq 1$.

Therefore:

$$[(\sigma - \delta\sigma)/\sigma]100 = (1-\delta)100$$

= percent reduction in reported volatility

Similarly, where r_f = the risk-free rate:

$$\left[\left(\left(\bar{X} - r_f\right)/\sigma\right) - \left(\left(\bar{X} - r_f\right)/\delta\sigma\right)\right] / \left(\left(\bar{X} - r_f\right)/\sigma\right)100$$
$$= [(1/\delta) - 1]100$$

= percent improvement in reported Sharpe ratio.

It is worth noting that the performance of Manager 2 is in no way superior to Manager 1; Manager 2 merely represents herself as being superior due to the stability of her returns—which is in fact merely a consequence of inability or unwillingness to accurately value the portfolio.

Exhibit 9 illustrates the relationship between the proportional valuation lag and the improvement in reported risk-adjusted returns as represented by the Sharpe ratio. Note that Sharpe ratios can be highly sensitive to proportional valuation lags. For example, a lag factor of 0.5 will result in a 100% overstatement of risk-adjusted returns, while a lag factor of 0.15 will result in a 567% overstatement.

Therefore, when managers face difficulties in performing accurate periodic valuations of their portfolios, there is substantial opportunity for overstating risk-adjusted performance. When performance is systematically overstated by individual managers, or systematically overstated for certain classes of managers, statistically derived performance measures will by definition mischaracterize performance. Subsequently, there will be a tendency to overweight such individuals or investment classes.

Once again, one could rationally argue that ex post Sharpe ratio optimization is a negative selection process when applied to unadjusted manager or index performance data. If manager data are not appropriately normalized, quantitative optimization strategies will systematically bias a portfolio toward illiquidity and de facto reduce risk-adjusted returns.[4] This tendency is here referred to as *illiquidity bias.*

Illiquid securities create problems for an investor beyond suboptimal allocation. Most important, the tendency to under-report volatility implies the occurrence of predictable financial calamities. To better understand this issue, consider the anatomy of a typical illiquidity-based financial crisis.

At the commencement of trading, a manager invests in a selection of illiquid securities. The securities are valued at the purchase price. Therefore, initially the reported net asset value (NAV) of the portfolio is approximately equal to the true or liquidation value of the portfolio. At the end of a period, the securities will have changed in value. Given the illiquidity of the securities, the manager cannot determine their precise values, nor can any objective third party. Consequently, the manager will tend to systematically understate the periodic change in the NAV of the portfolio.

The manager produces a periodic NAV by augmenting the prior period's NAV by some proportion of the difference between where the portfolio was previously valued and its current true value. This strategy results in periodic over- and undervaluations of the reported NAV compared to the true NAV. The extent of these misstatements is a function of three key variables: the true mean ($\bar{x}$) and standard deviation (σ) of the underlying portfolio, and the extent to which the manager captures the periodic difference in value between the prior period's reported NAV and the current true NAV, i.e., the proportional valuation lag (δ).

A crisis typically results when a manager's prime broker or investment partners become concerned about the possible difference between the reported and the actual NAVs and force a liquidation of all or a portion of the investment portfolio. A difference that exceeds some crisis threshold value (L) typically evokes such concern.

Exhibit 10 presents a randomly generated example of just such a manager. It graphs actual and reported NAVs as well as a histogram of the periodic valuation differences. In this example σ is assumed to be 30%, $\bar{x}$ is assumed to be 15%, δ is assumed to be 0.15, and L is assumed to be 20%. Using these assumed parameter values in conjunction with a simple Monte Carlo simulation, we can determine an estimated time to financial crisis (T). Simply put: $T = f(\bar{x}, \sigma, \delta, L)$. In this example the expected time to crisis is 49 months.[5]

EXHIBIT 10
Reported NAV versus Actual NAV

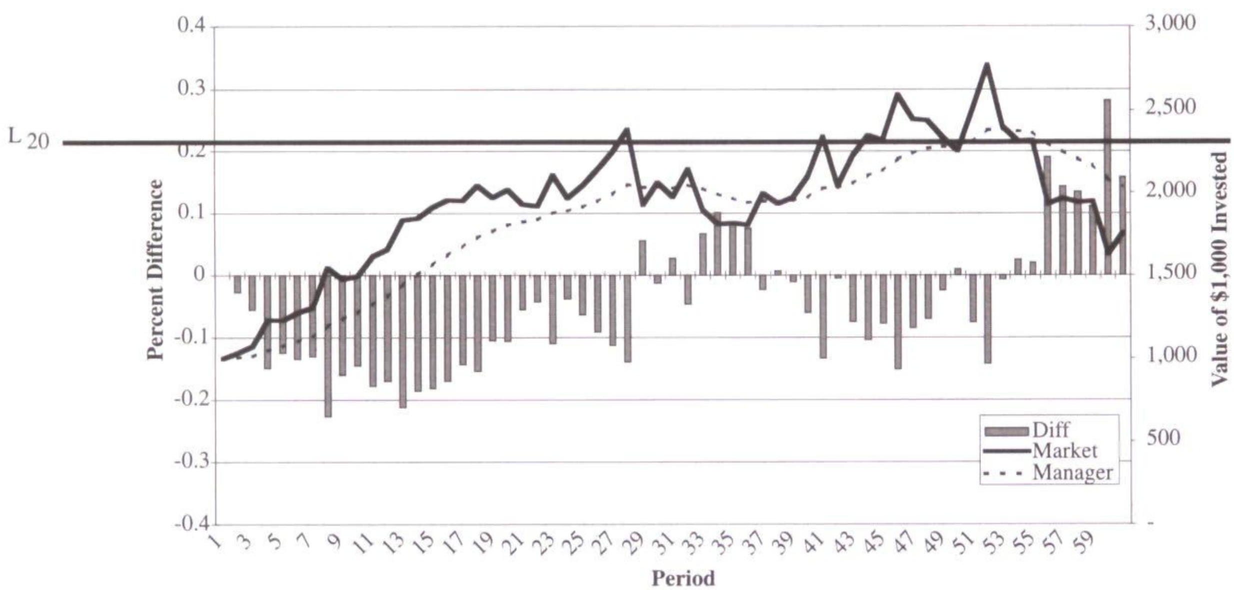

Mean = 0.15, sigma = 0.30, lag valuation = 0.15.

EXHIBIT 11
Proportional Valuation Lag versus Estimated Months to Disaster

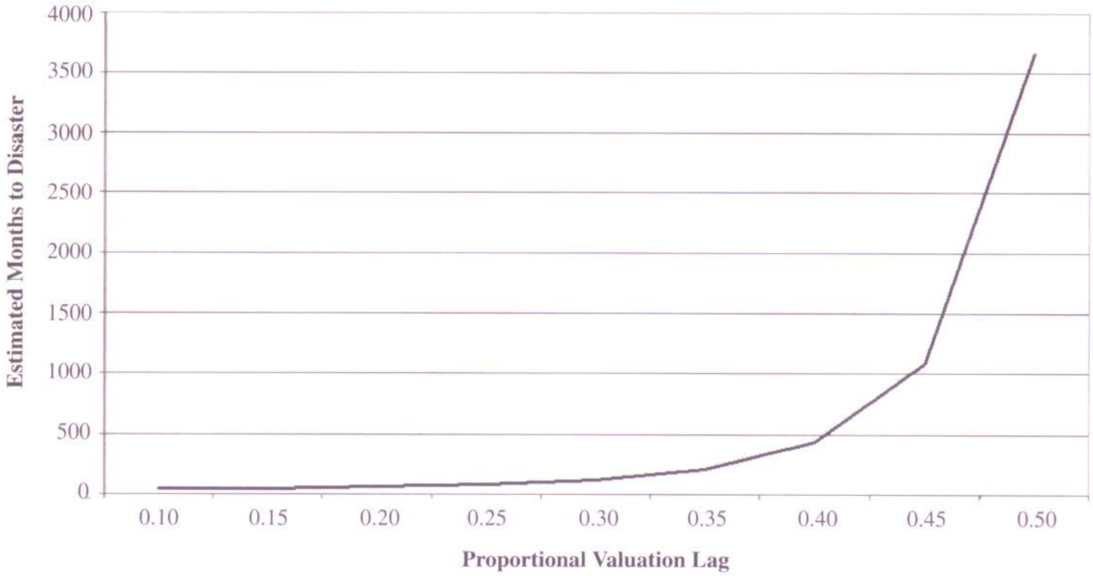

Mean = 0.15, sigma = 0.30, threshold = 0.2.

Exhibit 11 is a graphical representation of the relationship between the extent to which a manager understates volatility and the expected time to crisis expressed in months. Exhibit 12 presents the discrete data points that are used to generate Exhibit 11.

Finally, it is worth noting that the use of illiquid or over-the-counter securities is particularly endemic to the hedge fund industry, and is rarely satisfactorily addressed in academic studies of hedge fund performance. When you are researching the return-generating factors of hedge funds, or evaluating such research, it is a good idea to give serious consideration to the impact of illiquidity.[6]

EXHIBIT 12
Values for Generating Exhibit 11

Lag Valuation Parameter (δ)	Expected Time to Crisis (Percentile$_{10}$)	Expected Time to Crisis (Percentile$_{50}$)	Expected Time to Crisis (Percentile$_{90}$)
0.10	5	44	161
0.15	7	49	167
0.20	9	59	192
0.25	12	77	253
0.30	18	113	390
0.35	32	204	691
0.40	67	437	1411
0.45	165	1088	3702
0.50	560	3668	6000+

$\bar{x} = 0.15;\ \sigma = 0.30;\ L = 0.20.$

ST. PETERSBURG INVESTING

To understand this informationless investment technique, it is first necessary to consider a concept known as the St. Petersburg Paradox. This concept refers to the seemingly paradoxical expectations associated with a simple betting strategy.

This informationless strategy involves making a single unit bet on the outcome of a binomial process such as a coin toss. If you win, you bet again with the same unit size. If you're wrong, you "double up" by betting two units on the subsequent trial. If you're wrong again, you "double up" once more by betting four units. You continue doubling up until you eventually win, at which point you return to betting the starting unit amount.

This betting strategy has some unique properties. First, even though the coin is assumed to be fair, this strategy has an infinite expected value. Second (and here's the paradox), you will, with a probability of one, eventually become bankrupt. With absolute certainty, you will eventually encounter a long enough series of losing bets so that, for any finite amount of capital, you will lose everything. Clearly, if a manager increases leverage as he goes into a draw-down (as he loses capital as result of investment losses), he is subjecting his investors to a substantial amount of risk.

The solution to this problem is to avoid managers who engage in this sort of behavior. Yet, due to the return patterns generated by managers who employ some form of informationless, "double-up," betting strategy, it is frequently the case that inexperienced asset allocators actively select for such managers.

Sample St. Petersburg Investment Strategy

To better understand such a money management strategy, it is worth considering the returns associated with a simple St. Petersburg-like investment strategy. At the start of the first week of trading, a hypothetical manager makes an investment (bet) risking 50 basis points (0.5%) of capital. If the bet fails, the manager makes a second bet at the beginning of week two, risking 100 basis points. The manager continues to double up, as a percentage of the remaining equity, at the beginning of every week until successful, at which point the manager reverts to making the initial unit bet of 50 basis points of capital. Finally, to introduce an "opportunity" component, the manager reports returns only at month-end.[7]

With this limited amount of information, we can use Monte Carlo simulation to characterize the manager's likely future performance. For the purpose of this simulation we assume that the manager has no systematic skill, or lack of skill, i.e., there is a 50% chance of being right in any given week.[8]

Exhibit 13 presents a randomly generated sample monthly performance history for our fictitious manager. It depicts precisely what we would expect from a firm employing inappropriate (St. Petersburg style) money management. There is a fairly prolonged period of consistent profitability. The manager appears to recover brilliantly, and rapidly, from any loss of capital. Finally, the manager goes out of business in spectacular fashion. It is worth noting that this is a very common life cycle for hedge funds.

Given our precise specification of this strategy, we can once again use Monte Carlo simulation to accurately describe the associated expectations. These expectations are summarized in Exhibit 14.[9]

EXHIBIT 13
ACME Hedge Fund Performance

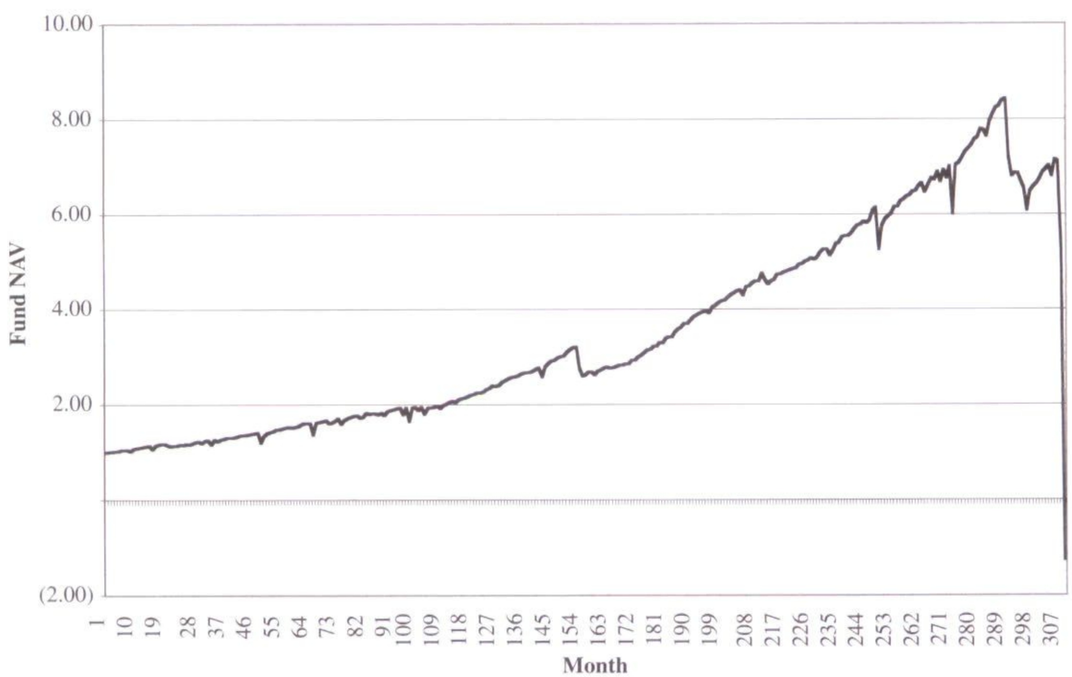

EXHIBIT 14
Comparison of Percentage Draw-Down versus Mean Time to Occurrence (Derived via Monte Carlo Simulation)

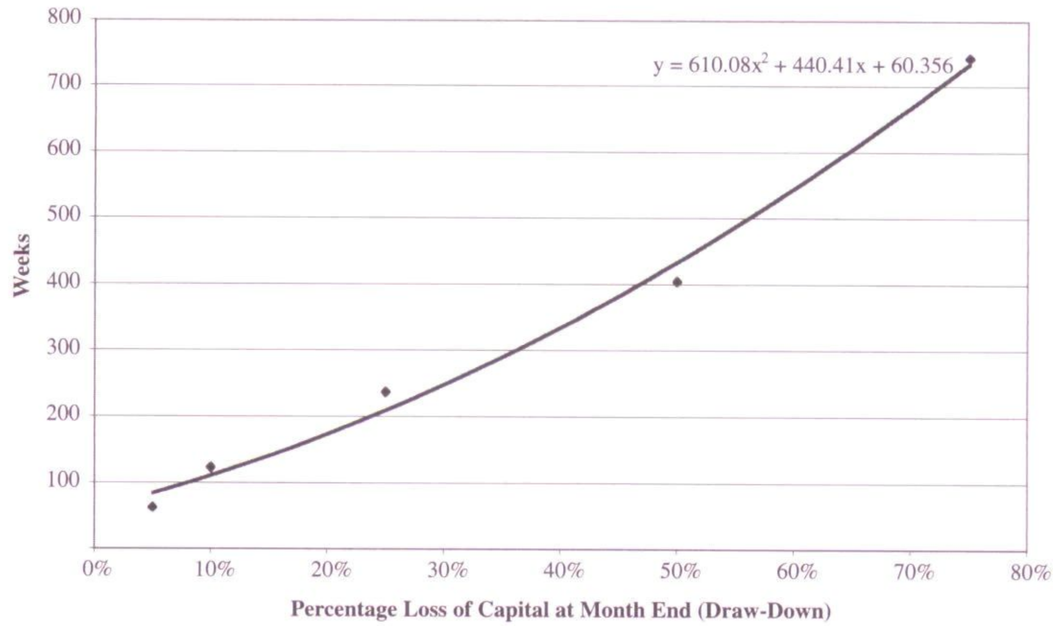

Exhibit 14 describes the relationship between a specific percentage loss of capital and its expected time to occurrence. You can see from the graph that the expected time until a month-end loss of 50% of capital is approximately 400 weeks, or seven and three-quarter years.

It is worth considering the significance of a seven-and-three-quarter-year expected time to a 50% loss of capital. First, the sales cycle for a hedge fund is far shorter than seven and three-quarter years. Typically, hedge funds come up for serious consideration after three to five years. Additionally, as our sample performance history in Exhibit 13 indicates, prior to experiencing a large loss of capital, the manager's performance is quite compelling.

Monthly reporting tends to obscure much of the fund's volatility; the draw-downs (losses of capital) have a very limited duration, and the returns are consistently positive. In fact, right up until its fiery death, such a fund would generate approximately a 15% annualized rate of return with about a 12% annualized standard deviation, and would be profitable approximately 78% of all months.

One further interesting trait associated with this strategy is, given that the manager is simply making a series of unrelated weekly wagers, it is unlikely that the overall return series will have any long-term systematic correlation with any particular index. In short, we have defined a seemingly high risk-adjusted return product with a low correlation with other managers and indexes.

Bear in mind that all of this wonderful performance is consistent with the a priori structure of this experiment, i.e., that the manager is employing an informationless strategy and is assumed to have no systematic skill. As a consequence, statistically derived in-sample performance measures will by definition significantly mischaracterize potential out-of-sample results.

St. Petersburg Summary

The most frightening result of this experiment is just how easy it is to create a St. Petersburg-type investment program that will probably generate a long period of superior performance and very low correlation relative to many traditional benchmarks. Subsequently, when structuring a portfolio by naively maximizing an ex post measure of risk-adjusted return, one may actually be selecting for managers who employ money management strategies that imply a catastrophic loss of capital.[10]

I call the tendency to allocate to such managers "St. Petersburg bias." Anecdotal evidence indicates that the St. Petersburg bias, allocating to managers who increase leverage as they go into a draw-down, is quite prevalent in the alternative investment industry.

CONCLUSION

Short-volatility bias, illiquidity bias, and the St. Petersburg bias are important considerations to bear in mind in attempting to apply established investment concepts and technology to the world of hedge funds. As I demonstrate, hedge fund managers have the ability to engage in essentially "informationless" strategies that can produce the appearance of return enhancement without necessarily providing any value to an investor. Consequently, statistically derived estimates concerning risk, return, and association frequently mischaracterize hedge fund returns.

These mischaracterizations have significant negative implications for both the asset allocation process and the validity of considerable academic research.

ENDNOTES

The author thanks Jerome Abernathy, Mark Anson, and Masao Matsuda for their thoughtful comments; Richard Michaud for his guidance and wisdom; Tim Birney for excellent quantitative research assistance; Adam Albin for editorial assistance; and finally the Institute for Quantitative Research in Finance and its April 2001 conference participants for their thoughtful questions and comments. The views expressed in this article are the opinion of the author, and should not be taken to represent those of his employer.

[1]As Michaud [1998] notes, mean-variance optimization is highly prone to "error maximization" because such procedures tend to overuse statistically estimated information, and thereby magnify the impact of estimation errors.

[2]Such strategies tend to produce very compelling "stick-like" performance histories (i.e., rates of return, when graphed, that appear smooth and upward-sloping over fairly long periods). A high probability of success on a given trial, however, should not be confused with the notion of a positive expected value that takes into account the payoffs associated with various outcomes.

[3]See, for example, Fung and Hsieh [1997], McCarthy and Spurgin [1998], Schneeweis and Spurgin [1998], and Liang [1999].

[4]Normalized to compensate for the impact of estimated lagged valuation.

[5]In my opinion, these parameter values represent an eminently realistic example. Furthermore, this framework for analyzing illiquidity goes a long way toward explaining the highly deterministic and cyclical nature of such events.

[6]The pervasive tendency for certain managers, or classes of managers, to include illiquid securities in their portfolios calls for a reexamination of much of the published research in this area.

[7]Providing an opportunity to engage in unreported, intramonth, overly aggressive trading.

[8]In my experience, the example is not an extraordinarily unrealistic example of the money management practices employed by certain hedge fund managers, especially with respect to the directional investment strategies typically employed by commodity trading advisors.

[9]We use Monte Carlo simulation for solving this problem primarily because we are interested in incorporating the effect of periodic month-end reporting. If this were not the case, this problem could be solved deterministically.

[10]Interestingly, Edwards and Ma [1988] note that there is actually a significant empirical negative relationship between in-sample pro forma track records provided in commodity fund offering memoranda and subsequent out-of-sample performance.

REFERENCES

Brown, S., W. Goetzmann, and R. Ibbotson. "Offshore Hedge Funds: Survival and Performance 1989-95." *Journal of Business*, January 1999, pp. 91-117.

Edwards, F., and C. Ma. "Commodity Fund Performance: Is the Information Contained in Fund Prospectuses Useful?" *Journal of Futures Markets*, Vol. 8, No. 5 (1988).

Fung, W., and D. Hsieh. "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds." *The Review of Financial Studies*, Vol. 10, No. 2 (Summer 1997).

Liang, B. "On the Performance of Hedge Funds." Association for Investment Management and Research, July/August 1999.

McCarthy, D., and R. Spurgin. "A Review of Hedge Fund Performance Benchmarks." *The Journal of Alternative Investments*, Summer 1998.

Michaud, Richard O. *Efficient Asset Management*. Boston: Harvard Business School Press, 1998.

Schneeweis, T., and R. Spurgin. "Multifactor Analysis of Hedge Fund, Managed Futures, and Mutual Fund Return and Risk Characteristics." *The Journal of Alternative Investments*, Fall 1998.

Weisman, A. "Conservation of Volatility and the Interpretation of Hedge Fund Data." Alternative Investment Management Association, June/July 1998.