Quantitative Topics in Hedge Fund Investing

To inspire healthy discourse.

Jimmy Liew and Craig French

The fund of hedge funds industry is experiencing unprecedented growth as institutional investors have become aware of the benefits of combining alternative and traditional investments. Alternative investments have historically benefited traditional portfolios in two ways: 1) reduction of risk through diversification, and 2) consistent historical returns. Yet these benefits are not without their costs, including the time and resources required to understand the intricacies of trading strategies, develop a network of hedge fund contacts, flesh out and truly understand the hidden risks, train and retain staff for proper ongoing monitoring of managers, and construct proper institutional-grade risk management platforms.

Funds of hedge funds have become the institutional vehicle of choice to gain exposure to hedge funds. With institutional aversion to headline risk and the need for experienced and properly trained staff to properly execute a sound investment process, the pendulum has swung away from direct investment toward institutional-quality funds of hedge funds. The top funds of funds are upgrading their investment process by hiring institutional-quality people, including those trained in mathematical finance and financial engineering programs. The best funds of hedge funds will typically provide knowledge transfer as part of their services for their clients.

We introduce some topics we feel are important for clients to consider in making investments in funds of hedge funds. We also want to provide an introduction to industry realities for those now sharpening their quantitative skills.

Jimmy Liew
is a lecturer at Columbia University's Financial Engineering Program in New York City.
jimmyliew100@gmail.com

Craig French
is a partner and director of Risk Management and Quantitative Research at Corbin Capital Partners in New York City.
cfrench@corbincapital.com

As the industry has evolved, we have seen an increased demand for quantitative analysts who can help untangle the challenging world of hedge fund investing. In this work, we highlight a few areas of current interest on hedge fund investing from the perspective of two quantitatively trained investors:

1. Separating alpha and beta.
2. Accounting for serial correlation.
3. Portfolio construction realities.
4. Manager selection versus strategy allocation.
5. Mispriced negative skew.

We hope our work will contribute to a broader understanding of practical topics and spark debate about these important issues in the fund of hedge funds industry.

SEPARATING ALPHA AND BETA

The critical questions of separating alpha from beta and how much compensation each should command were posed years ago in the long-only active management world; they have since then been revisited in the hedge fund industry by Anson [2000]. Schneeweis and Spurgin [1996], Fung and Hsieh [1997], Weisman [2002], Jensen and Rotenberg [2003], Asness [2004a], Siegel [2004], and others have proposed innovative approaches that provide a powerful framework for answering these questions.

Traditionally, alpha has been thought of as outperformance over a given benchmark, but hedge funds have historically been considered *absolute return* investment vehicles; many have no explicit benchmark. For hedge fund managers who can buy, sell, and lever positions, alpha can be achieved from both sides—long and short. Alpha can also be magnified through the use of leverage, or carved out appropriately through the use of derivatives.

In the real world of hedge fund investing, distinguishing between alpha and beta becomes a difficult endeavor. From a practical perspective, a model that captures the systematic component of returns is necessary. One equilibrium-type approach uses the Treynor [1962]–Sharpe [1964] capital asset pricing model. This is a common approach in the long-only world of traditional investments, and it can also be employed, with varying degrees of success, for hedge funds.[1]

Given the broad opportunity set available to hedge fund managers, however, the linear CAPM with alpha defined rigorously as Jensen's alpha can become less tractable. If hedge fund managers trade their portfolios dynamically or hold non-linear payoff securities such as derivatives, the linear CAPM may not be the correct framework. Other typical concerns are a lack of precision in estimation of betas and the lack of a long history of data.

Multifactor models based on Ross's [1976] arbitrage pricing theory can be a useful starting place. Some of these models include the Fama-French [1993] three-factor model, the Carhart [1997] four-factor model, and the Edwards and Caglayan [2001] six-factor model. Some researchers use hedge fund indexes as factors; others use statistical factors such as principal components analysis.

We have recently seen some practitioners construct hedge fund trading strategies by mechanically replicating known hedge fund trading processes. Two recent developments have occurred for merger arbitrage and convertible arbitrage.

Asness [2004b] suggests a pay schedule; namely, hedge fund alpha should be more expensive than hedge fund beta, and both should be more expensive than regular beta. Using security-level indexes as factors may be the purest way to distinguish alpha from beta (or, more precisely, alpha from hedge fund beta), but this is not generally practical (because security-level indexes do not exist for all hedge fund strategies).

We believe the best factors are market observable factors. With market observable factors, we can gain substantial insight about the hedge fund manager's return-generating process.

It is our view that, in general, hedge fund return streams are influenced primarily by four main market observable factors: 1) equity markets, 2) bond markets, 3) volatility, and 4) credit. These four general factors have been shown to explain some but not all of the variation in hedge fund returns. We suggest that prudent institutional investors should be able to fine-tune these factors.

For example, if the hedge fund under consideration is an equity long-short fund, the factors used should be equity markets, small-minus-big, high-minus-low, previous one-year returns, equity volatility, and so on (see Edwards and Liew [1999] and Edwards and Caglayan [2001]). Similarly, for a fixed-income arbitrage fund, one should use strategy-appropriate factors including proxies for bond markets, steepeners/flatteners, fixed-income volatility, swaption volatility, and so on (see Darnell and Billings [2004]). Ideally the factors should be simple tradable strategies, thus allowing for the most intuitive interpretation of the hedge fund returns.

Framing the separation of alpha from beta in such a multifactor context appears, in our view, to be the best

EXHIBIT 1

Hedge Fund Beta Distribution—139 Long-Short Equity Managers

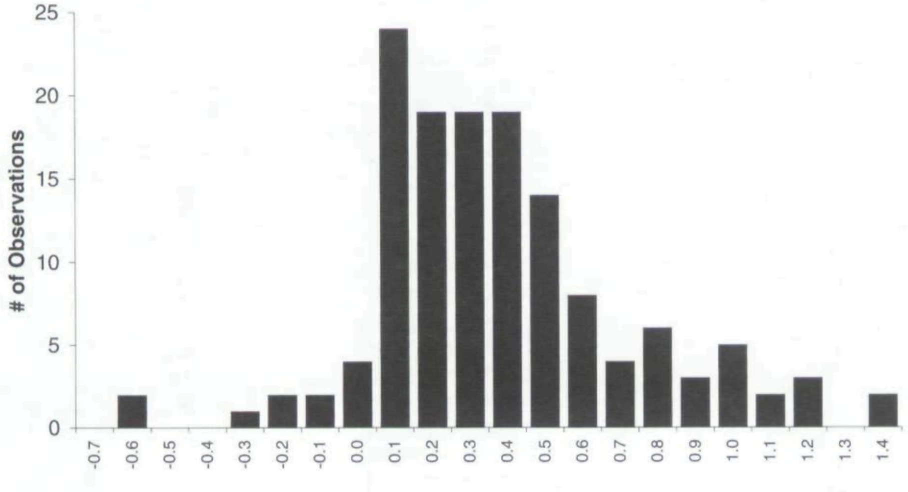

Source: CSFB/Tremont data May 1984–October 2004.

approach, as it is simple, powerful, flexible, and intuitive. While no model can capture all the intricacies of hedge fund return streams, we have found it successful to favor simple market observable factors with tractable intuition as opposed to unidentified statistical factors.

To move from the theoretical multifactor world to a simple real-world single-factor example, in practice we are limited by hedge fund return streams over a relatively short time, as well as by biases in the databases. A cross-sectional distribution of betas for equity long-short hedge fund managers demonstrates the heterogeneous risk profiles of hedge fund managers.

Exhibit 1 shows that, on average, the hedge fund beta is 0.31 over our sample of 139 equity long-short managers. We use an ordinary least squares (OLS) regression of excess monthly hedge fund returns against excess monthly S&P 500 returns, for the moment ignoring any serial correlation problems.

Some hedge funds have betas to the U.S. equity market above 1.0, while others have betas below –0.5. Given the wide dispersion of betas, we can see it is inadequate to use broad generalizations to describe hedge funds. Hedge funds continue to be a motley bunch.

A single beta estimate, however, does not convey the whole story for hedge funds. Given the dynamic trading strategies of many hedge funds, betas or factor sensitivities generally vary over time. Time variation of both alpha and beta should thus be estimated.

Exhibit 2 displays a sample analysis of time variation in betas. The top panel depicts time series of betas representing three investment vehicles: 1) the fund of hedge funds portfolio, *sample portfolio*; 2) the candidate fund, *sample global trading manager*; and 3) the original portfolio with the candidate fund included at a hypothetical 5% allocation, *combined portfolio*. This time we estimate OLS betas iteratively, using a rolling 12-month window. The S&P 500 proxies for the equity market factor.

The solid line represents the sample portfolio's beta over time; the dotted line just underneath that solid line represents the combined portfolio's beta over time had we included the candidate fund. The candidate fund's beta over time is represented as a dashed line.

The middle panel of Exhibit 2 shows the difference between the combined portfolio and the original portfolio. It is apparent that inclusion of the candidate fund would have reduced the portfolio's exposure to the S&P 500 viewed on a historical basis.

In the bottom panel of Exhibit 2 we compute the regression t-statistics to see if the exposures are statistically significant over time. Failure to perform such statistical significance testing would provide a less meaningful framework for quantitative analysis. It is much less informative to say that a hedge fund has a beta of 0.2 than it is to say that a hedge fund has a beta of 0.2 that is statistically significant with a t-statistic of over 2.0.

The stability of beta can be questionable in so-called tail events. Liew [2003] provides evidence that hedge fund indexes have unstable betas in the tails. A good quantitative process should estimate exposures during tail events for hedge fund returns; one commercial application of this concept is described in Wilmott [2001]. Recent academic evidence seems to suggest that, even at the individual equity level, stocks that covary highly with market declines have a higher risk premium; see, e.g., Ang, Chen, and Xing [2004]. We believe it is reasonable to apply similar intuition when we analyze hedge funds.

Exhibit 3 displays one type of tail risk analysis for a hedge fund with a significant positive beta to the S&P 500 index of 0.34 over the entire sample distribution. Notice that the beta in this example inverts in times of market stress. During the tail event, beta is estimated at negative 0.18. This is a valuable piece of evidence on a character-

EXHIBIT 2
Rolling Factor Exposures

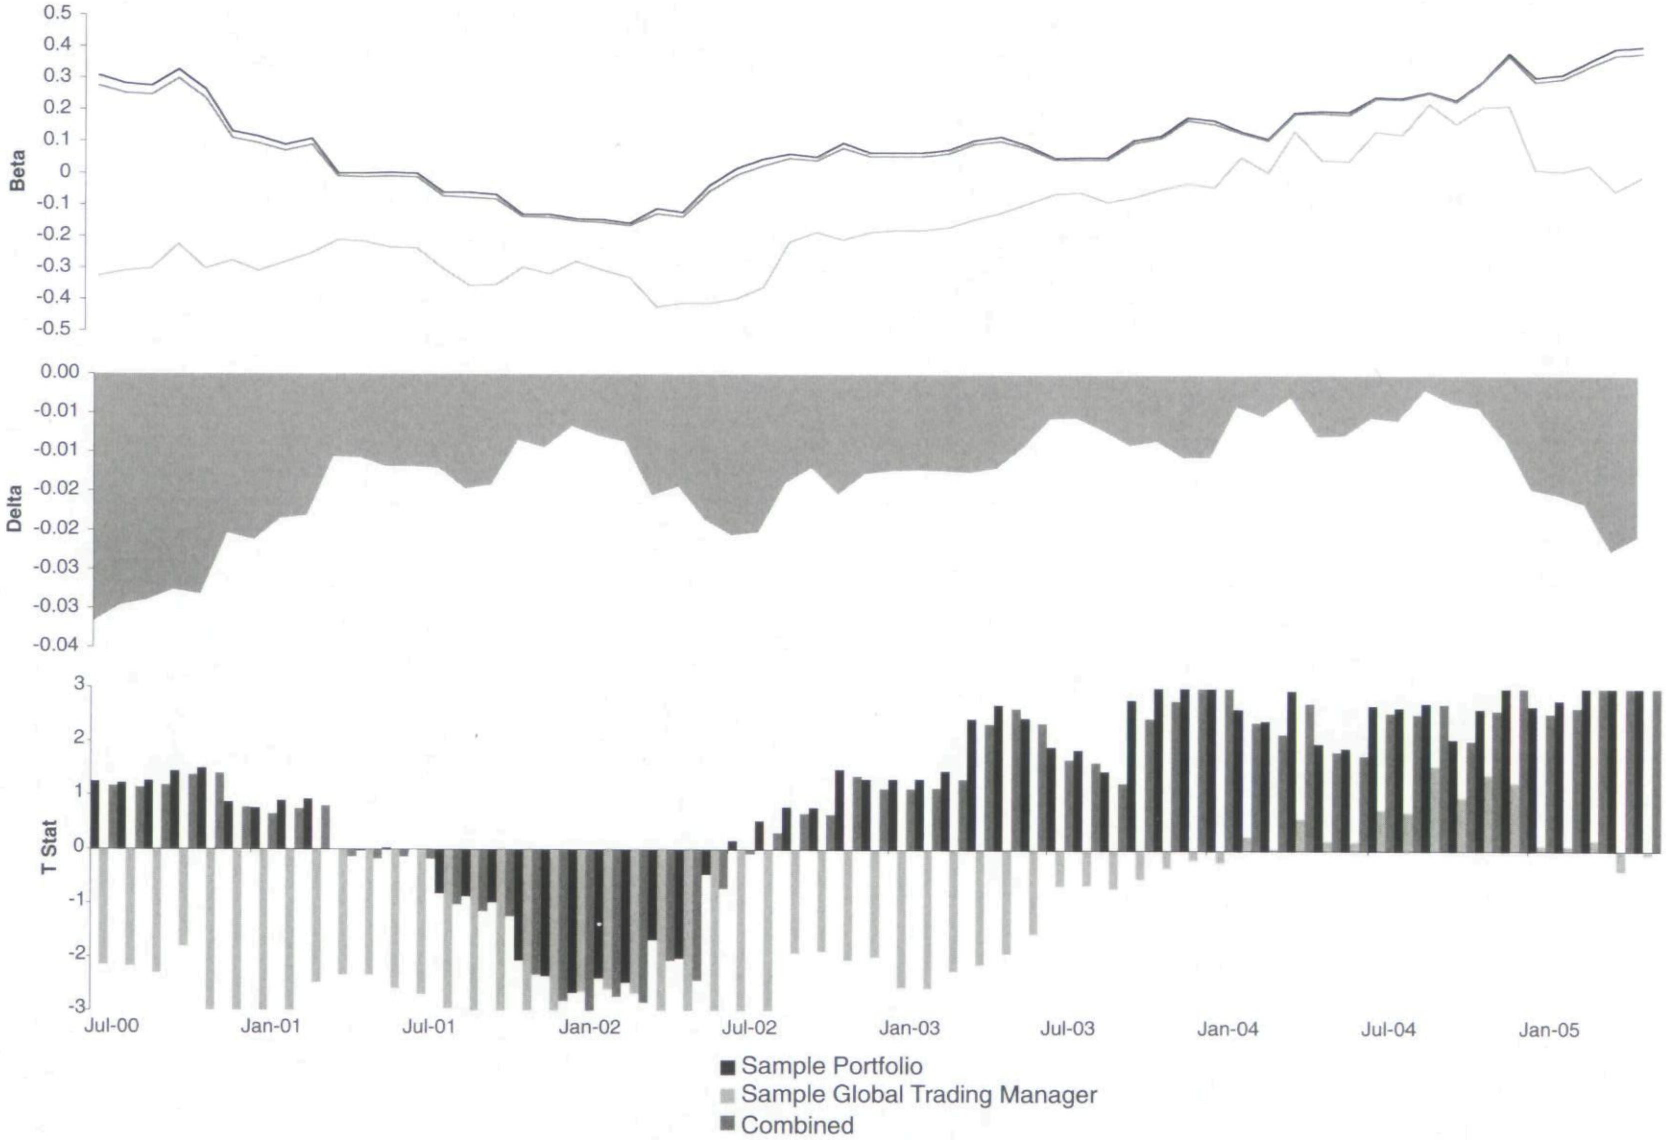

Exhibit 3
Tail Risk Analysis

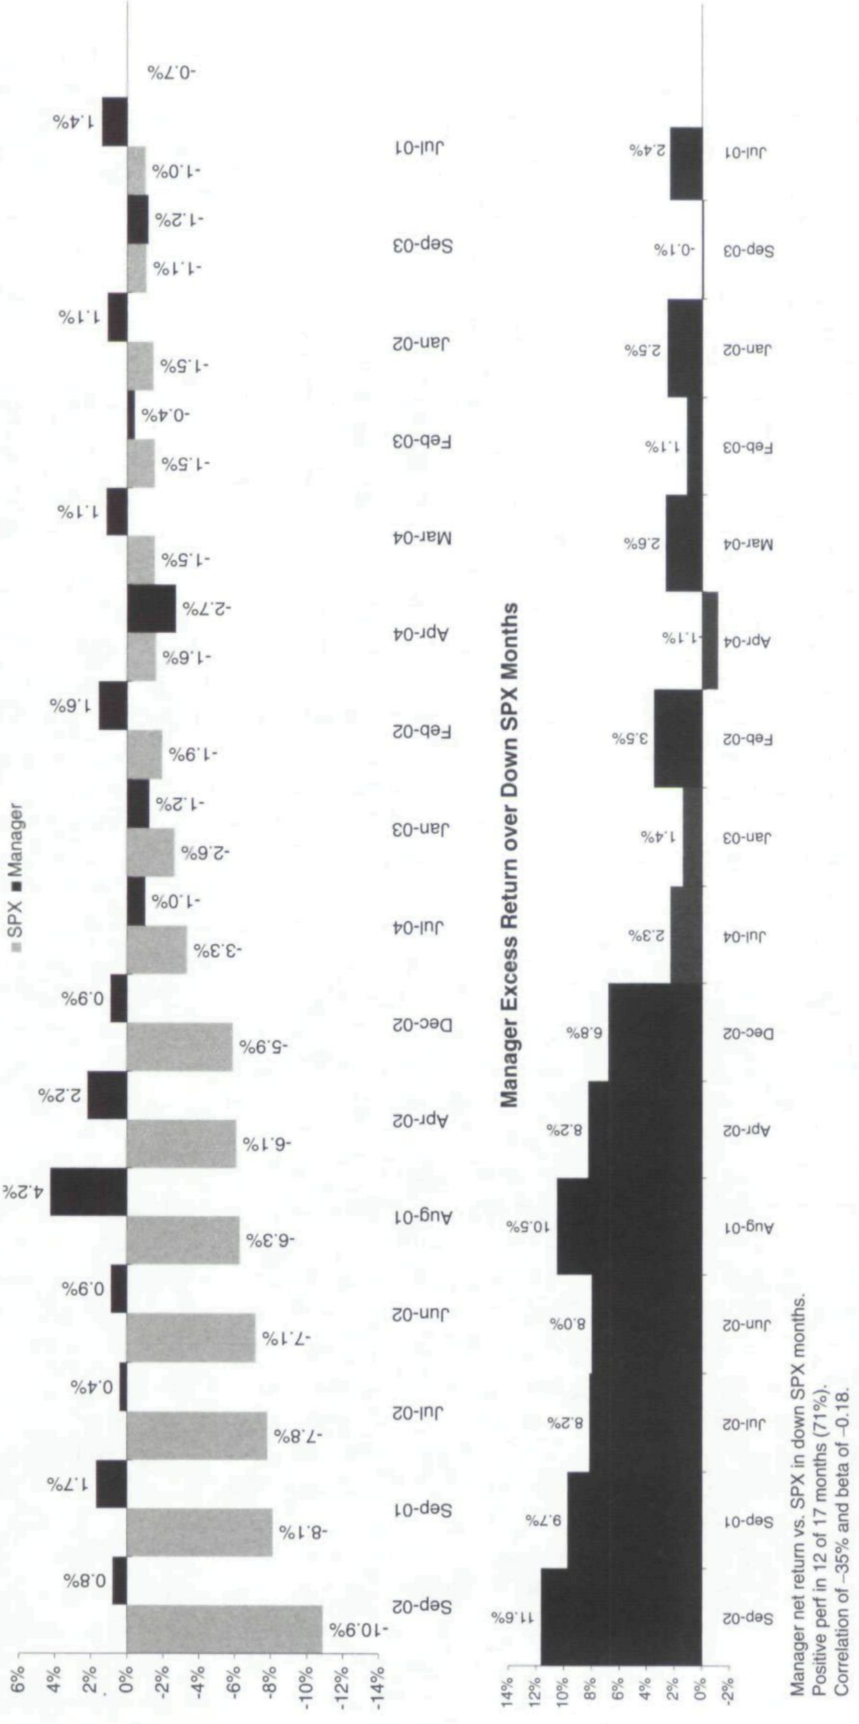

istic that could be overlooked without tail risk analysis.

Another important concern is the non-linear nature of hedge fund returns. The long volatility characteristics of some strategies (e.g., commodity trading advisors, CTAs), and the short volatility profile of others (e.g., merger arbitrage) may not be uncovered through traditional OLS regression techniques. Mitchell and Pulvino [2000] and Anson and Ho [2003] note the non-linear characteristics of some hedge fund strategy returns, and use a piecewise linear regression model to replicate the option-like payoffs in such strategies. Such non-linear models can be extremely useful in the analysis of hedge funds.

ACCOUNTING FOR SERIAL CORRELATION

One problem we believe has not been sufficiently fleshed out in the industry is positive serial correlation in hedge fund returns. Positive serial correlation biases traditional measures of volatility, beta and correlation. Asness, Krail and Liew [2001] find that certain hedge fund indexes exhibit positive serial correlation, and once standard measures of risk and reward are corrected, returns are less attractive. If the hedge fund monthly return series suffers from positive serial correlation, and the quantitative analyst fails to make the appropriate adjustments, true annualized volatilities will be underestimated and true diversification benefits will be overestimated.

To analyze hedge fund returns, one must first determine whether positive serial correlation is present in the returns through a statistical test such as Durbin-Watson [1950] or Ljung-Box [1978]. Then, if it is warranted, one makes the proper adjustment to the Sharpe ratio.

Lo [2002] proposes an innovative way to account for positive serial correlation in non-independent and identically distributed (iid) but stationary returns. We use this approach in the example shown in Exhibit 4. We examine a sample of CSFB/Tremont's convertible arbitrage universe and sort the top ten funds by the Sharpe ratio using data from November 1988 through August 2004. The question is whether we would invest in the manager with the highest Sharpe ratio.

Exhibit 4 illustrates that, if we use Sharpe ratios to evaluate hedge funds, we would conclude that convertible arbitrage fund 1 has the best risk-adjusted return. Notice, however, that once positive serial correlation is accounted for using the Lo ratio, convertible arbitrage fund 3 becomes the most attractive.

Asness, Krail, and Liew [2001] examine CSFB/Tremont Hedge Fund Index data over the 1994-2000 period and document evidence of stale pricing; stale pricing may be incidental (due to non-synchronous pricing quotations) or intentional (due to managed pricing). While the latter reason is significant, in either case, such non-synchronous data can lead to underestimation of true market exposures (relative risk). Asness, Krail, and Liew [2001] correct for this effect by applying the Scholes and Williams [1977] lagged beta approach to estimate equity market betas of the hedge fund indexes.

Simple measures of volatility (absolute risk) also seem to understate actual hedge fund risk. This effect is documented directly by Conner [2003], who finds variance smoothing at the aggregate hedge fund index level, using CSFB/Tremont Hedge Fund Index data over 1994-2002. He uses a model proposed in Geltner [1991] to de-lag the return series and construct an economic, or unsmoothed, return series.

While Asness, Krail, and Liew [2001] suggest modifying Sharpe ratio estimates (absolute risk-adjusted return) using their lagged beta methodology to generate a hedged Sharpe ratio, Lo [2002] approaches the Sharpe ratio question from a different angle. Given that serial correlation tends to dampen volatility, and thus simple volatility in the denominator of the Sharpe ratio will be underestimated when annualized in the traditional square root of N way, Lo [2002] proposes a generalized method of moments approach to correct for this problem when annualizing the monthly Sharpe ratio estimates.

Malkiel and Saha [2004] use a three-lag beta estimation approach as proposed by Asness, Krail, and Liew [2001] on underlying hedge fund data from the Tremont TASS database over 1996-2003, and confirm the results previously obtained (although only at the aggregate hedge fund index level).

We construct a meta-dataset using a proprietary database, which includes a significant number of hedge funds that do not report to any of the publicly available data vendors, and for which we have collected data over the past 20 years. The dataset also incorporates all underlying hedge fund data from the publicly available sources: 1) Altvest, 2) Hedgefund.net, 3) HFR, 4) MSCI, and 5) TASS. We thus mitigate some of the problems associated with vendor-supplied hedge fund data, such as survivorship, backfill, end-of-life reporting, and self-selection biases.

We examine all funds operating at least over the five-year period January 2000 through December 2004 in four categories: convertible arbitrage, event-driven, long/short equity, and global macro. We are interested in two statistics: the proportion of the identifiable hedge fund universe that exhibits positive serial correlation in net returns, and the amount of overestimation in standard risk-adjusted returns due to the autocorrelation.

For example, in long-short equity we find 3,925 funds in the meta-data; 1,259 have return histories spanning at least the five-year 2000–2004 period. Using a multistage procedure to eliminate the duplicate funds, we select a sample of 615 independent long-short hedge funds with full track records over the five-year period.[2]

Of these 615 funds, 197 funds, or 32%, exhibit significant positive serial correlation (momentum) in their

EXHIBIT 4
Rank-Order—Sharpe Ratio versus Lo Ratio

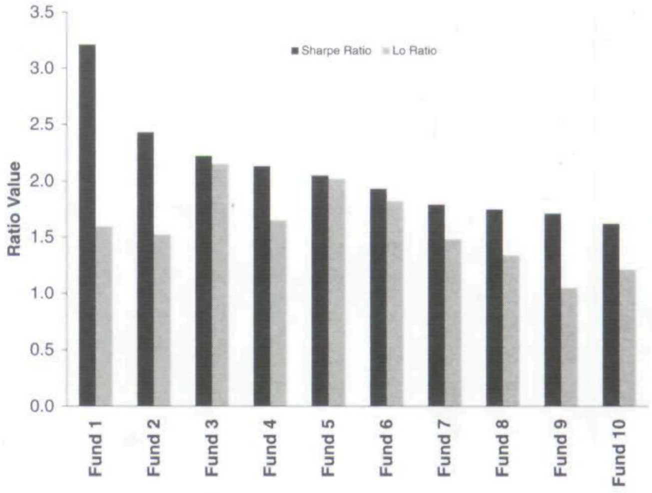

EXHIBIT 5

Positive Serial Correlation in Hedge Fund Returns—Jan 2000–Dec 2004

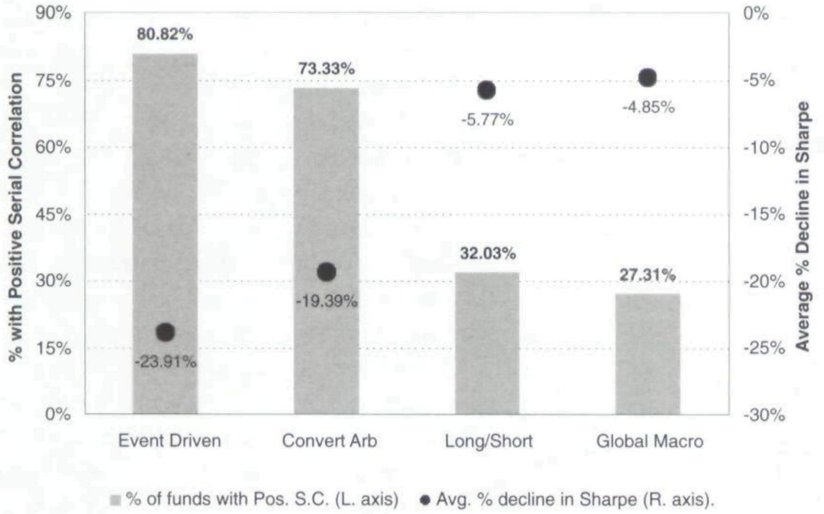

	Event-Driven	Convertible Arbitrage	Long-Short	Global Macro
HFs Reported (all sources)	335	169	1,259	457
Unique HFs (post-deduplication)	146	75	615	271
HFs w/ Signif. Positive Serial Correlation	118	55	197	74
% of HFs w/ Signif. Positive Serial Correlation	81%	73%	32%	27%
Average Percent Decline in Sharpe	23.9%	19.4%	5.8%	4.9%

returns as measured at the 95th level of confidence by the Durbin-Watson test statistic. For the remaining 197 funds, we then estimate the Sharpe [1966] ratios in the standard fashion and then in the method proposed by Lo [2002]. We find a mean Sharpe ratio of 1.04 and a mean Lo ratio of 0.98, for a difference of 0.06, a 5.8% reduction in risk-adjusted return due to autocorrelation.

The results of all four strategies are summarized in Exhibit 5. We see that the vast majority of event-driven hedge funds in this sample (81%) exhibit significant positive serial correlation. While such a high level of smoothing is not pervasive across all strategies (only 27% of our sample of global macro managers suffer from positive serial correlation), it can have a strong effect on the risk estimation of funds to which it applies. We also see that accounting for smoothing reduces Sharpe ratio estimates by 5% to 24%.

Some funds of hedge funds may allocate to managers using risk budgeting. Frequently this approach is accompanied by marginal value at risk analysis or maximum risk limits across strategies and managers. To the extent that the underlying hedge fund return streams suffer from positive serial correlation, allocating to risk buckets without any adjustments to estimated risks will bias the overall portfolio and unduly overweight those managers.

PORTFOLIO CONSTRUCTION REALITIES

Markowitz's mean-variance optimization technique is commonly used in portfolio construction. Unfortunately, this technique can introduce errors, both because the future tends not to resemble the past and because the technique generally takes historical point estimates for the vector of mean returns and the variance-covariance matrix as inputs. To the extent that high past returns and low correlations do not persist over time, mean-variance optimization may produce misleading results.

There is also often considerable uncertainty about the estimates input into the optimizer. We have found the optimizer can be most useful as a tool in portfolio construction, when used as a guide to the potential efficiency of a portfolio, rather than when it is blindly followed for allocation purposes.

To avoid some of these problems, the diligent analyst should use some sort of mechanism that allows for the input of realistic and forward-looking estimates of return and risk, such as a Bayesian approach, and should evaluate a variety of potential outcomes based on multiple scenarios.

Exhibit 6 displays output in the mean-variance space for a portfolio and its constituents using Bayesian priors blended with historical data. The Bayesian approach can be considered a generalization of the Black-Litterman [1990] model, which allows investors to incorporate current views by tilting around equilibrium allocations. The equilibrium benchmark allocations are usually determined by the capitalization weights for long-only portfolios. In the world of hedge fund investing, the Black-Litterman equilibrium approach becomes self-contradictory, because in an equilibrium pricing context hedge funds should not exist.

One way to address this contradiction is to formu-

EXHIBIT 6
Mean-Variance Analysis of Hedge Funds

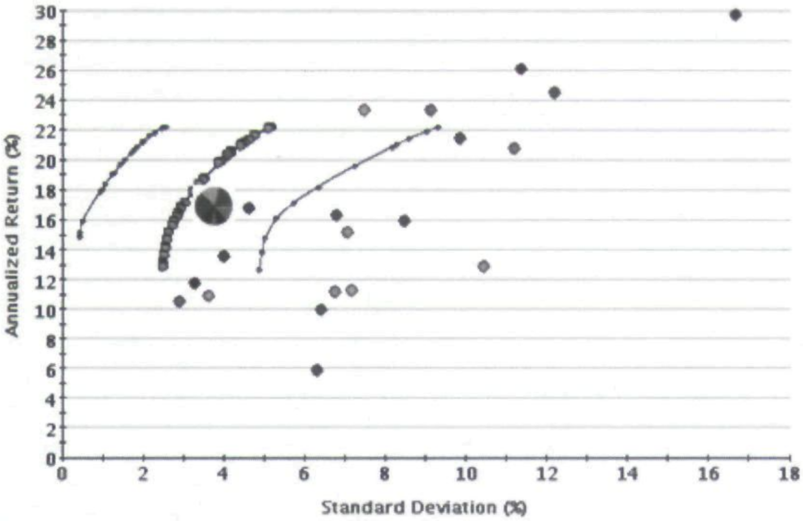

Efficient frontier for sample portfolio.
Data for 05/1992–09/2004.
Includes max constraints and Bayesian priors.

late a pseudo-equilibrium benchmark allocation following the long-only world, and choose an observable cap-weighted index of hedge funds such as the CSFB/Tremont Hedge Fund Index as the starting point. Given this pseudo-equilibrium benchmark, we could then apply our views by tilting around the allocations of the benchmark.

In practice, hedge funds have varying liquidity terms such as monthly, quarterly, and annual (or longer) redemption periods; they may also have multiyear lock-up periods, and require various notice periods for redemption. Given these real-world liquidity frictions, typical hedge fund investors have minimal flexibility when it comes to reallocating across strategies. Some practitioners, though, do attempt to overcome these obstacles and reallocate capital opportunistically through a Bayesian technique.

Michaud [1998] proposes one averaging method that uses Monte Carlo simulations. A similar iterative technique involves perturbing the variance-covariance matrix to stress the portfolio in a comparative statics framework. Output from this approach, which we call the Frontier Rangefinder, is illustrated in the far left and far right efficient frontiers in Exhibit 6.

Other advanced approaches have been employed with varying success. Harvey et al. [2003], for example, have developed a three-moment optimizer that incorporates skewness into the objective function.

Covariance matrix estimation can be particularly troubling, since the optimization algorithm tends to overweight outlying observations. In a typical hedge fund of hedge funds cross-correlation matrix, most optimization algorithms will overemphasize the highest and the lowest values.

Ledoit and Wolf [2001] propose a technique to diminish the effect of this mathematical perversion. We find a similar technique useful; shrinking the correlation estimates toward a unit matrix results in conservative estimates in the sense of generating higher portfolio variances.

These can be interpreted heuristically as generic stress tests—we assume some future unknown economic conditions will cause future cross-correlations to be higher than they were in the past, and they can occur for any unspecified reasons. We are interested in the potential result of reduced diversification, given that we cannot predict what its cause will be, and shrinkage estimators are a simple yet effective way to accomplish that goal.

The literature provides mixed evidence about performance persistence in the hedge fund industry. Capocci [2001] replicates Carhart's [1997] mutual fund persistence study, and finds some evidence of time-varying persistence in hedge fund returns. Malkiel and Saha [2004] find that a small majority of winners repeat annually. Agarwal and Naik [2000] find evidence of persistence at the quarterly horizon, and Edwards and Caglayan [2001] using a six-factor model find evidence of persistence at one-year and two-year horizons.

We find that risk appears to be more persistent than returns. For example, a global macro manager with past high volatility typically will continue to employ a similar trading style in the future, and thus exhibit high future volatility. Similarly, a statistical arbitrage manager with low past volatility tends to have low future volatility, as tight risk controls incorporated into the process continue over time. In practice, there is empirical evidence of significant persistence in the short term for risk exposures. Risk exposure persistence monitoring is a critical component of portfolio construction in hedge fund portfolios.

We believe that multiple risk statistics, and even multiple asset pricing models, should be used to develop a better understanding of hedge fund investments. Liew, Mainolfi, and Rubino [2002] propose a bifurcated fund

EXHIBIT 7

Selection versus Timing

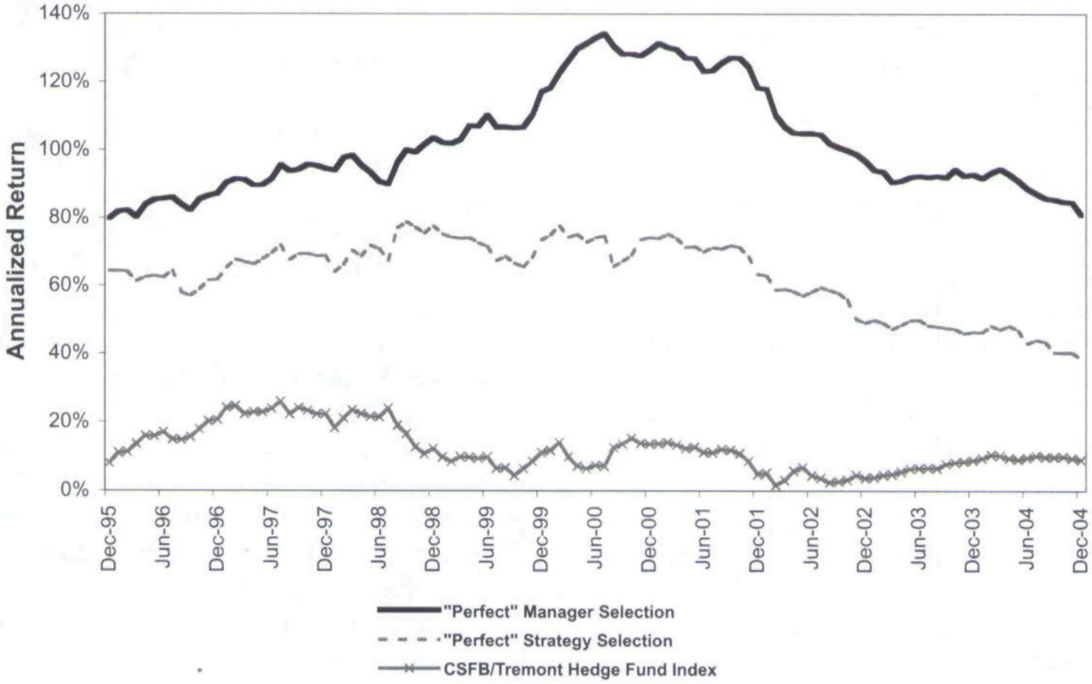

analysis model (BFAM), which accommodates multiple models and acts as an efficient filtering device. Some institutional-quality firms use more advanced filter models in practice to narrow the scope of attractive alternative investments.

Of course, quantitative work should always be coupled with thorough qualitative work such as understanding the manager's approach and the mid- and back-office procedures. Some of these qualitative assessments can be incorporated into such a model.

Finally, achieving true diversification can be challenging in the hedge fund universe. French [2003a] shows that, while diversification can reduce volatility, the marginal benefits of adding more holdings to a portfolio can decline rapidly. Careful manager selection and rigorous portfolio construction are required.

MANAGER SELECTION VERSUS STRATEGY ALLOCATION

One interesting question is the relative importance of selection versus timing—the ability to pick the best managers versus the ability to move across hedge fund strategies. The answer could inform fund of funds managers about where they should focus their efforts and spend most of their time. Some funds of funds claim to be skillful at taking tactical bets across strategies, while others claim expertise in manager selection, and still others claim that they excel at both. We can provide some empirical evidence that manager selection may be more important than strategy allocation over time.

Several attribution models could prove useful in the hedge fund world: the Treynor [1965] model, the Fama [1972] selectivity model, the regression method of Brinson, Hood, and Beebower [1986], and the Kritzman and Page [2002] opportunity set approach. We provide an empirical example that lends some support to the Kritzman and Page thesis.

We use data from the CSFB/ Tremont hedge fund universe, which has ten strategy classifications. As of October 31, 2004, long-short represented the most common strategy allocation, at 26%, while dedicated short was the smallest, at only 1%. We could identify 372 of the 382 managers in the CSFB/Tremont index, and we use their performance data from January 1994 through August 2004 in the analysis.

We construct three different hypothetical trading strategies. The first is a passive buy-and-hold strategy in the CSFB/Tremont Hedge Fund Index. CSFB/Tremont's index is asset-weighted—that is, constructed by assets under management for each manager.

Second, we construct a hypothetical *perfect strategy allocation* trading strategy. This strategy assumes perfect foresight and rotates to the best strategy each month. Perfect strategy allocation invests at asset-weight to each manager, thus incorporating no manager selection component.

Finally, we construct a hypothetical *perfect manager selection* trading strategy, which assumes perfect foresight and chooses the best manager in each strategy monthly. This trading strategy has no tactical rotation ability as strategy allocations are equal weighted over time.

Clearly the last two of these three strategies would be impossible to implement in practice, because no one has perfect foresight, and, practically speaking, rebalancing monthly is difficult, if not impossible. But let us continue in this vein for a moment just for educational purposes.

A $1 investment in January 1994 in the passive CSFB/Tremont index would have grown to approximately $3 in August 2004. A $1 investment at the same time in the active perfect *strategy allocation* would have yielded $585. But the same $1 invested in perfect *manager selection* would have yielded a whopping $22,520. Manager selection has by far the most potential.

Exhibit 7, depicting a rolling 24-month analysis, shows that perfect manager selection outperforms perfect strategy selection (which outperforms index investing) over time.

This analysis demonstrates that the wide cross-sectional distribution of hedge fund returns within strategies dominates the constraint of equal weights across strategies. Thus, theoretically, manager selection has more potential opportunities than tactical rotation across strategies. In essence, a fund of funds may be better off pursuing great managers than in trying to build tactical asset allocation models.

It is easy to say our comparison of perfect strategy allocation and perfect manager selection is an unrealistic test that compares apples to oranges. We intend it merely to shed light on an interesting debate in the industry that might motivate other quantitatively oriented researchers to investigate this question in more detail.

MISPRICED NEGATIVE SKEW

Ask most investors about negatively skewed returns, and you may be surprised at their aversion to such distribution characteristics. Negatively skewed returns are typically generated from strategies that sell volatility or have positive carry profiles, including merger arbitrage, negative gamma trades, credit carry trades, mortgage carry trades, or emerging market carry trades. Typically, historical returns are quite consistent with low positive monthly returns in most months and then a terrible double-digit down month, such as that observed in August 1998.

Cremers, Kritzman, and Page [2004] note that hedge fund returns are frequently found with significant non-normal higher moments. They also show that investors with bilinear utility functions, or S-shaped preferences, should be attracted to both excess kurtosis and negative skewness.

To explore this question, we introduce a fund of funds called the Negative Skew Fund. The fund invests only in hedge fund return streams characterized by negatively skewed returns. The point of the exercise is to highlight the importance of portfolio construction. The critical consideration is that they are all uncorrelated with each other. That is, not all the managers in the portfolio will experience the double-digit down month simultaneously.

This hypothetical portfolio consists of 30 managers, all with a negatively skewed return stream. In most months, each manager generates returns of +1% per month; when the tail event occurs, however, the manager loses 10% in that month.

If the correlation across all funds is 1.0, then the portfolio of negatively skewed managers would be exactly the same as the individual manager distribution and would represent a very dangerous portfolio. When the funds have cross-correlations that are anything less than perfectly positive, though (i.e., they are not all perfect substitutes for each other), the portfolio skewness benefits from diversification across managers.

We construct a correlation matrix of all the managers of -0.017 using 60 months of data. The equally weighted portfolio's performance with a tail event is $(1/30)(-10\%) + (29/30)(1\%) = + 0.63\%$ in months one of the managers has a tail event episode and +1% when no manager has a tail event. Interestingly enough, the portfolio generates returns that have a skew of essentially zero (–2.1E-14).

This Negative Skew Fund has been constructed to give all positive monthly returns. Just as we learned in the long-only investing world, it's more important to know how things fit together into the overall portfolio (coskew) rather than focus on each individual's investment risk.

CONCLUSION

Institutions must consider many important issues regarding hedge funds. Just knowing the right questions to ask is a good starting point for investment success in hedge fund investing. Such knowledge can help institutional investors avoid many of the potential pitfalls.

In our view, the primary issues to address include alpha-beta separation, positive serial correlation, portfolio construction realities, manager selection versus strategy allocation, and negative skewed returns. We hope for a healthy discourse about these issues to inspire future quantitative researchers and institutional investors to begin their journey into hedge fund investing.

ENDNOTES

This article builds on a presentation given to the Financial Engineering Seminar at Columbia University's School of Engineering and Applied Science, Department of Industrial Engineering and Operations Research on November 29, 2004. It can be accessed at: http://www.ieor.columbia.edu/fe_seminars_French.html. The authors thank the seminar participants for their thoughtful comments and suggestions.

[1]For a summary of development of the Treynor-Sharpe CAPM, see French [2003b].

[2]The dataset includes substantial overlap because many data vendors report the same fund's returns, so we apply a six-stage deduplication procedure: 1) eliminating funds with the same exact name and net return series; 2) eliminating 3c1 replicas of 3c7 funds; 3) eliminating euro-denominated replicas of U.S. dollar-denominated funds; 4) eliminating offshore replicas of domestically domiciled funds; 5) eliminating hot-issues-eligible replicas of otherwise non-hot-issues-eligible funds; and 6) eliminating replicas due to redundant series of a fund. In stages 2-6 of the procedure, replica is defined as a fund with a very similar name and a very similar net return series.

REFERENCES

Agarwal, Vikas, and Narayan Y. Naik. "Multi-Period Performance Persistence Analysis of Hedge Funds." *Journal of Financial and Quantitative Analysis*, Vol. 35, No. 3 (September 2000).

Ang, Andrew, Joseph Chen, and Yuhang Xing. "Downside Risk." Working paper, Columbia Business School, July 2004.

Anson, Mark. "Hedge Fund Risk Management for Institutions." In V.R. Parker, ed., *Hedge Fund Risk Management for Institutions—Managing Hedge Fund Risk*. New York: Risk Books, 2000.

Anson, Mark, and Ho Ho. "Short Volatility Strategies: Identification, Measurement, and Risk Measurement." *Journal of Investment Management*, Vol. 1, No. 2 (2003), pp. 30-43.

Asness, Clifford. "An Alternative Future." *The Journal of Portfolio Management*, 30th Anniversary Issue, 2004a.

——. "An Alternative Future: Part II." *The Journal of Portfolio Management*, Fall 2004b.

Asness, Clifford S., Robert Krail, and John M. Liew. "Alternative Investments: Do Hedge Funds Hedge?" *The Journal of Portfolio Management*, Vol. 28, No. 1 (2001), pp. 6-19.

Black, Fischer, and Robert Litterman. "Asset Allocation: Combining Investor Views With Market Equilibrium." Goldman Sachs Fixed Income Research, September 1990.

Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. "Determinants of Portfolio Performance." *Financial Analysts Journal*, July/August 1986.

Capocci, Daniel P.J. "An Analysis of Hedge Fund Performance 1984-2000." Working Paper No. 0109, Université de Liège, November 2001.

Carhart, Mark. "On Persistence in Mutual Fund Performance." *Journal of Finance*, Vol. LII, No. 1 (March 1997), pp. 57-82.

Conner, Andrew. "The Asset Allocation Effects of Adjusting Alternative Assets for Stale Pricing." SEI Investments Research Paper, January 2003.

Cremers, Jan-Hein, Mark Kritzman, and Sébastien Page. "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" Revere Street Working Paper 272-13, September 2004.

Darnell, Max, and Mark Billings. "Partner's Message: Alpha & Beta In Market Neutral." First Quadrant, 2004.

Durbin, J., and G.S. Watson. "Testing for Serial Correlation in Least Squares Regression I." *Biometrika*, Vol. 37 (1950), pp. 409-428.

Edwards, Franklin R., and Mustafa O. Caglayan. "Hedge Fund Performance and Manager Skill." *The Journal of Futures Markets*, Vo1. 21, No. 11 (2001).

Edwards, Franklin R., and Jimmy Liew. "Hedge Funds versus Managed Futures as Asset Classes." *The Journal of Derivatives*, Vol. 6, No. 4 (1999), pp. 45–64.

Fama, Eugene F. "Components of Investment Performance." *Journal of Finance*, Vol. 27, No. 3 (June 1972), pp. 551-567.

Fama, Eugene F., and Kenneth R. French. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, 33 (1993).

French, Craig. "Diversification in Multi-Manager Hedge Funds: A Flexible Approach." Corbin Capital Partners, January 2003a.

——. "The Treynor Capital Asset Pricing Model." *Journal of Investment Management*, Vol. 1, No. 2 (2003b), pp. 60-72.

Fung, William, and David A. Hsieh. "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds." *The Review of Financial Studies*, Vol. 10, No. 2 (1997), pp. 275-302.

Geltner, David. "Smoothing in Appraisal-Based Returns." *Journal of Real Estate Finance and Economics*, Vol. 4 (1991), pp. 327-345.

Harvey, Campbell, John Liechty, Merrill Liechty, and Peter Muller. "Portfolio Selection with Higher Moments." Working paper, Duke University and NBER, October 2003.

Jensen, Greg, and Jason Rotenberg. "The Flood Into Hedge Funds." *Bridgewater Daily Observations*, May 9, 2003.

Kritzman, Mark, and Sébastien Page. "Asset Allocation versus Security Selection: Evidence from Global Markets." Revere Street Working Paper 272-9, February 2002.

Ledoit, Olivier, and Michael Wolf. "Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection." Working paper, Credit Suisse First Boston, November 2001.

Liew, Jimmy. "Hedge Fund Index Investing Examined." *The Journal of Portfolio Management*, Vol. 29, No. 2 (Winter 2003), pp. 113-123.

Liew, Jimmy, Francesco Mainolfi, and Dean Rubino. "Bifurcated Fund Analysis Model." *MFA Reporter*, November 2002, pp. 12-15.

Ljung, G., and G. Box. "On a Measure of Lack of Fit in Time-Series Models." *Biometrika*, Vol. 65, No. 1 (1978), pp. 297–303.

Lo, Andrew. "The Statistics of Sharpe Ratios." *Financial Analysts Journal*, July/August 2002, pp. 36-52.

Malkiel, Burton G., and Atanu Saha. "Hedge Funds: Risk and Return." Working paper, Princeton University, 2004.

Michaud, Richard O. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Cambridge: HBS Press, 1998.

Mitchell, Mark, and Todd Pulvino. "Characteristics of Risk and Return in Risk Arbitrage." *Journal of Finance*, Vol. 56, No. 6 (2000), pp. 2135-2175.

Ross, Stephen A. "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory*, Vol. 13, No. 3 (December 1976), pp. 341-360.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, Vol. 19, No. 3 (1964), pp. 425-442.

——. "Mutual Fund Performance." *Journal of Business*, January 1966, pp. 119-138.

Schneeweis, Thomas, and Richard Spurgin. "Multi-Factor Models in Managed Futures, Hedge Fund and Mutual Fund Return Estimation." Working paper, University of Massachusetts, December 1996.

Scholes, Myron, and Joseph Williams. "Estimating Betas from Nonsynchronous Data." *Journal of Financial Economics*, Vol. 5 (1997), pp. 309-327.

Siegel, Laurence B. "The Risk of Paying Alpha Fees for Beta Performance." Presentation to *PRMIA*, August 2004.

Treynor, Jack. "How to Rate Management of Investment Funds." *Harvard Business Review*, January/February, 1965, pp. 63-75.

——. "Toward a Theory of Market Value of Risky Assets." Working paper, Arthur D. Little, 1962. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=628187

Weisman, Andrew. "Dangerous Attractions: Informationless Investing and Hedge Fund Performance Measurement Bias." *The Journal of Portfolio Management*, Summer 2002.

Wilmott, Paul. *Paul Wilmott Introduces Quantitative Finance*. West Sussex: John Wiley & Sons, 2001.