

Canadian Investment Review

WINTER 2007 | VOLUME 20, NUMBER 4



OVER THE TOP

**CONFIDENT
INVESTORS PUSH
RISING MARKETS
TOO HIGH.**

PLUS:

**REPLICATION NATION:
WINNERS OF THE 2007
AIMA CANADA AWARD**

THE DUAL-CLASS DIVIDE

WWW.INVESTMENTREVIEW.COM

GETTING Real

**A look at the limits of
hedge fund replication.**

BY NEIL SIMONS AND ADRIAN HUSSEY

Within the past year, replication techniques have captured a great deal of media attention. In particular, the linear factor replication approach considered by Hasanhodzic and Lo and distributional replication pioneered by Kat and Palaro are frequently highlighted. At the same time, several investment banks as well as smaller investment management firms have launched products based on replication processes. All the attention is certainly not surprising given the fact that hedge fund replication has been proposed as a method for accessing hedge fund performance without the associated illiquidity, fees, necessity of performing extensive due-diligence, and exposure to single-manager risk. But can hedge fund replication really deliver? To find out, we focused our research on understanding the underlying fundamentals associated with the various replication processes. Contrary to what many say, we have determined that the underlying approaches are limited in their ability to access the performance of hedge funds. In the paper that follows, we look at both linear factor replication and distributional replication to determine the best approach to portfolio optimization.

LINEAR FACTOR REPLICATION

We started with linear factor analysis, which

has been extensively applied to investigate the market factors underlying hedge fund returns. The application of factor models to perform out-of-sample hedge fund replication is studied in Hasanhodzic and Lo. Two criteria are required for linear factor models to successfully replicate a return series. The return series of the hedge fund or index needs to have a large component of its behaviour explained by linear dependence on market factors and the hedge fund exposure to the market factors varies slowly over time. In addition to the above requirements for successful replication, profitable replication requires the trades derived from the factor model to provide good risk-adjusted returns in the future.

Individual hedge funds possess a larger amount of idiosyncratic risk than hedge fund indices. The idiosyncratic component of hedge fund returns tends to diversify away as many funds are combined into an index. Therefore, the out-of-sample linear factor replication of individual funds is less successful than the replication of hedge fund indices. We consider the replication of hedge fund indices as opposed to individual hedge funds. The linear factor-based replication products that have been released thus far target well-diversified, broad-based hedge fund indices such as the HFRI Composite. In the examples that follow, we consider the linear factor replication of the CS Tremont, CS Long Short Equity, CS Market

Neil Simons is vice-president, Northwater Capital Management Inc.; Adrian Hussey is vice-president, Northwater Capital Management Inc.

Neutral Equity, and HFRI Composite index using: S&P500, Russell 2000, USD Index, MSCI EAFE, MSCI Emerging Markets, and 30-Day US T-Bills. The R^2 from performing a 24-month rolling regression is provided in Figure 1.

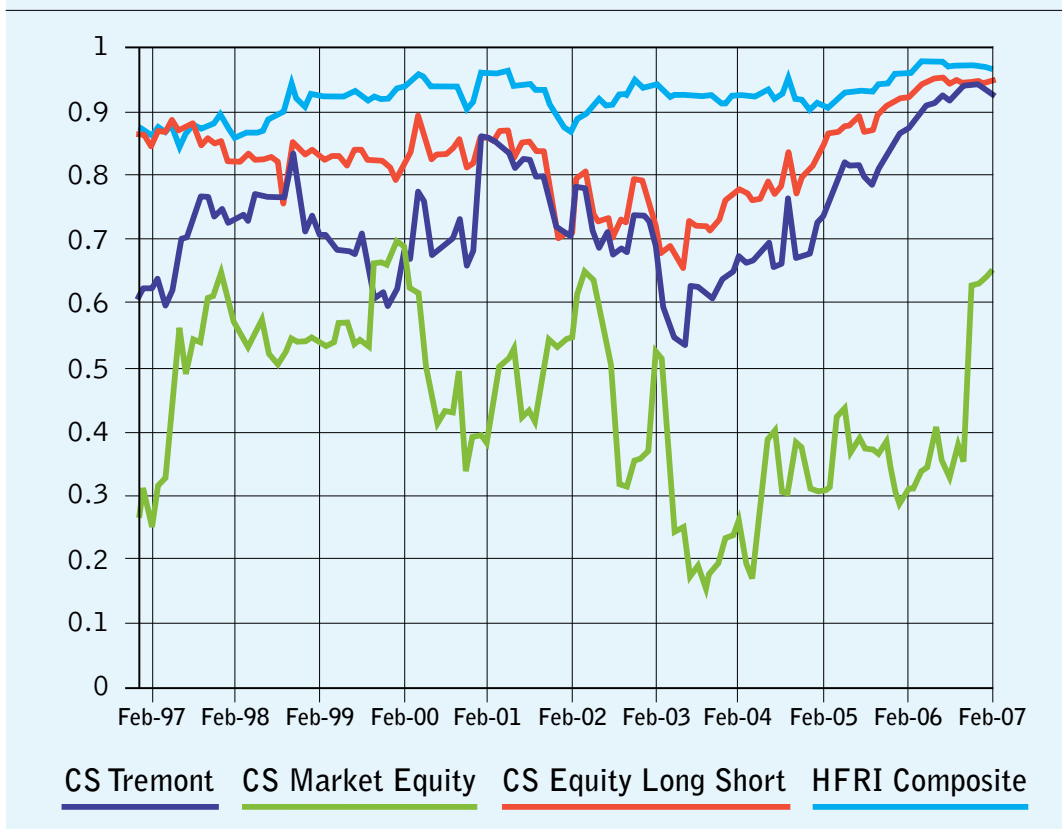
The R^2 is a measure of the success of the linear regression, through which one determines how well the known variables (in this case the market factors such as the S&P500, Russell 2000, USD Index, MSCI EAFE, MSCI Emerging Markets, and 30-Day US T-Bills) are able to explain the variable under investigation (in this case a relevant hedge fund index such as the CS Tremont, CS Long Short Equity, CS Market Neutral Equity, and HFRI Composite). The R^2 ranges between 0 and 1. A value close to zero implies almost no explanatory

fund return series. For a market-neutral hedge fund, one would expect to find that simple market factors are unable to explain the variation in hedge fund returns, and therefore a relatively low value for R^2 .

We find that the best fit occurs for the HFRI Composite (rarely lower than 0.9); and the worst fit for the CS Market Neutral Equity index (rarely greater than 0.6). We have investigated other hedge fund indices in addition to the four indices analyzed here. There are a few strategy indices that possess R^2 values similar to that of the equity market-neutral index. The remainder possesses R^2 values that are distributed between the market-neutral and composite indices.

The presence of a large rolling R^2 indicates

FIGURE 1: R^2 ACHIEVED FROM A ROLLING 24-MONTH REGRESSION OF 4 HEDGE FUND INDICES



Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

power (i.e., the market factors are unable to explain any variation in the hedge fund return series). A value close to 1 implies that the market factors are able to explain most of the variation in the hedge

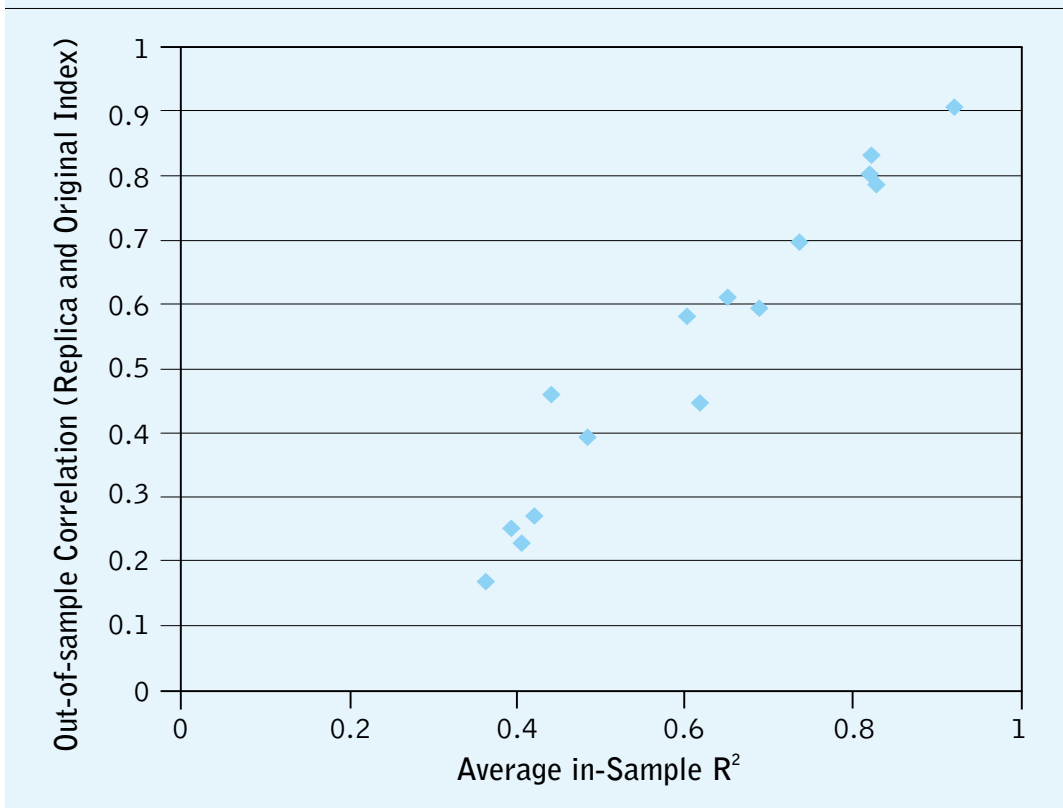
the returns can be modelled accurately looking backward. We have examined the out-of-sample replication of the four indices considered in Figure 1. For each index, a regression is performed using

Volatility and correlation targeting can be achieved using a monthly trading process derived through mean-variance portfolio optimization with additional constraints.

to December 2006), we have compared the Sharpe ratios for the original four indices and their linear replicas.

The comparison indicates that linear replication is capable of providing Sharpe ratios similar to the original indices for all cases, with the exception of the CS Equity Market Neutral index. This index has a low rolling R^2 value and, therefore, the specific factors applied in the model are unable to accurately explain the behaviour of the index. We

FIGURE 2: RELATIONSHIP BETWEEN THE OUT-OF-SAMPLE CORRELATION AND THE AVERAGE IN-SAMPLE R^2 FOR 15 HEDGE FUND INDICES



Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

the previous 24 months of index returns as the dependent variable, with the six market factors listed in the previous section as the independent variables. The factor loadings determined from the linear regression are used as position weights. A portfolio of market factors built using the position weights is held out-of-sample over the following month. Replication is simulated by rolling the process forward in time. Using a recent 3.5-year time horizon to perform the replication (July 2003

have also determined that out-of-sample linear factor replication is less successful over longer time horizons. The improved replication success for the broad-based indices over more recent time horizons is partially attributed to the increased rolling R^2 for the long/short equity strategy indices (see Figure 1). Long/short equity strategies comprise a larger weight within the composite indices (such as the CS Tremont and the HFRI Composite) in recent years.

The out-of-sample correlation of the replica and the underlying index provides an indication of replication success. If the replication process has successfully reproduced the time series behaviour of the original index, a correlation close to 1 is expected. The relationship between the out-of-sample correlation achieved by the replication process and the average R^2 observed over the replication time period is provided in Figure 2. We have expanded the set of indices considered and each point within the figure corresponds to a particular index. We find a positive dependence between the achieved out-of-sample correlation and the average R^2 with a y-intercept less than 0. The positive-sloped relationship is intuitively obvious: replication success depends on the extent of the factor contribution to the strategy returns. However, the negative y-intercept implies that a threshold exists for the average R^2 before the replica resembles the original index, and therefore acts as a guideline for selecting indices or funds suitable for linear replication.

It is possible to use more complex factors in order to improve the in-sample R^2 , and therefore potentially the out-of-sample correlation. For example, Fung and Hsieh, and Foerster have investigated factors that are useful in the explanation of CTA/Managed-Futures and Equity Market-Neutral strategies, respectively. While the refinement of the factor basket will help to improve replication success, the behaviour exhibited within Figure 2 remains. A continuum of strategies exists where some are more market-neutral than others. We conclude that successful linear factor replication requires the underlying funds or indices to possess significant exposures to market factors. The replication process provides access to the beta component of hedge fund returns. Due to its inherent idiosyncratic nature, the alpha component of hedge fund returns, which investors generally seek, is not accessible through linear factor replication.

DISTRIBUTIONAL REPLICATION FINDINGS

We now turn to distributional replication, which attempts to reproduce a desired return distribution (as described by volatility, skewness, kurtosis,

and correlation to a target portfolio) through a daily trading process using a set of liquid futures contracts. Distributional replication has been pioneered by Kat and Palaro and is sometimes reported to provide returns that are sometimes greater than the original hedge funds or hedge fund indices while matching their distributional characteristics. We have found that the performance attributed to distributional replication is due to the selection of market factors applied within the process.

Distributional replication can be separated into two subcomponents—distributional adjustments and correlation targeting. Distributional adjustments transform the distribution of a market factor or portfolio of market factors (referred to as the reserve portfolio), into a more desirable distribution. The process can scale volatility, implement a performance floor, or increase the skewness of an underlying market factor. Implementation of a distributional adjustment is therefore similar to dynamic option replication using delta hedging.

The idea of correlation targeting is an appealing one. It creates a replica with a desired correlation with respect to one specific target portfolio, and one desirable attribute of hedge funds is a low correlation to traditional market factors. One expects that a replica with a low correlation to an investor's portfolio will improve portfolio performance when combined with the original investor's portfolio. However, correlation targeting represents a synthetic approach to achieving the correlation target. Lower desired levels of correlation are achieved through a larger short position in the investor's portfolio within the replication process. This synthetic approach to achieving correlation becomes redundant when considered within a mean-variance portfolio optimization.

Volatility and correlation targeting can be achieved using a monthly trading process derived through mean-variance portfolio optimization with additional constraints. We have compared the results obtained from our monthly volatility and correlation targeting process to those achieved from the process applied by Kat and Palaro. The comparison utilizes the same

While hedge fund replication may be limited in its potential, such analysis leads to an improved understanding of hedge fund returns.

underlying hedge fund indices (10 Edhec Hedge Fund indices), the same time period (March 1999 to September 2006), and the same reserve and target portfolios as those considered by Kat and Palaro. The monthly process is capable of providing correlation, volatility, and mean excess returns very similar to those produced by the Kat and Palaro process. For each index, the correlation and volatility of the replicas are similar to the targeted hedge fund indices, indicating the success of the replication processes.

With respect to risk-adjusted return, our analysis indicates that the Sharpe ratios of the replicas are sometimes greater than the original Edhec indices and therefore one might conclude that the performance is due to the replication process. However, we demonstrate that the performance is due to the market factors selected for use within the replication process. The results provided in Figure 3 demonstrate the impact of using a different mix of market factors within the reserve portfolio. The replication of 10 Edhec Indices is performed using three different reserve portfolios (see x-axis). For each reserve portfolio, the correlation and volatility targets are successfully met, but mean returns are significantly different. As demonstrated by the minimum, maximum and average replica Sharpe ratio, the replicas have performance similar to that of the reserve portfolio in each case. Reserve Portfolio 1 consists of the same set of market factors as Kat and Palaro (S&P500, GSCI, Russell 2000, Long US Treasury Bonds, Medium Term US Treasury Bonds, and Eurodollar Futures). Reserve Portfolio 2 contains a modification to the original reserve portfolio. The S&P500 Information Technology and Telecommunication Services sectors (with poor performance over the time period considered) are included within

the reserve portfolio, and the Russell 2000 and GSCI (good performance over the time period) are excluded. The replica Sharpe ratios achieved from using Reserve Portfolio 2 are poor in comparison to those achieved from using Reserve Portfolio 1. Reserve Portfolio 3 has been modified in an attempt to achieve superior performance to the original reserve portfolio. In this case, S&P500 Aerospace & Defense, and the S&P500 Homebuilding sub-industry groups; and short positions in the S&P500 Information Technology and S&P500 Telecommunication Services sectors have been added.

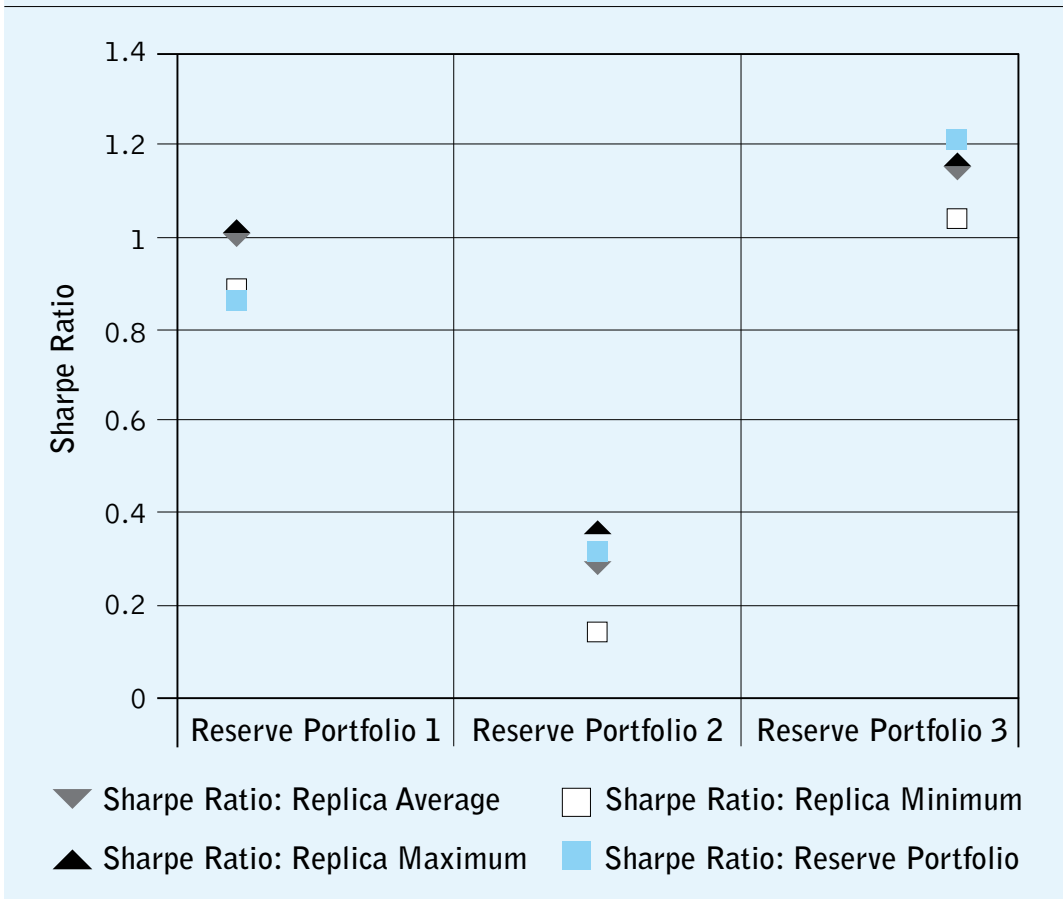
To a large extent, the excess return earned from distributional replication is obtained from the contents of the reserve portfolio. The Sharpe ratio of the reserve portfolio serves as an estimate for the Sharpe ratio of the replicas. Since the contents of the reserve portfolio consists of a long position in traditional market factors (to satisfy the constraint of using liquid instruments), expectations for the long-term performance of distributional replication will be equivalent to that of a portfolio of traditional market factors. Volatility and correlation targeting is robust with respect to changes in the reserve portfolio. However, mean returns are not robust to changes in the reserve portfolio.

OPTIMAL PORTFOLIO CONSTRUCTION

Let's recall that replication is an attempt to replace hedge funds (i.e., an investor allocates a portion of their portfolio to replicas rather than hedge funds). In order to understand the implications of this, we have investigated optimal portfolio construction using distributional replicas. The target portfolio is a 50/50 mix of S&P500 and Long Bonds; and the reserve portfolio is the original reserve portfolio used by Kat and Palaro. A Tobin frontier is generated using the optimal combination of the target and reserve portfolios along with a cash/funding allocation (in order to leverage this optimal combination and achieve any desired level of volatility).

In addition, we have created optimal portfolios combining the monthly replicas with the original target portfolio. For each of the 10 Edhec replicas,

FIGURE 3: SHARPE RATIOS OF DISTRIBUTIONAL REPLICAS FOR THREE DIFFERENT RESERVE PORTFOLIOS.



Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Edhec=Risk

without exception, the optimal portfolios lie along the Tobin frontier. Therefore, for each of the 10 replicas, the optimal combination of replica and original target is equivalent to the optimal portfolio created by combining the reserve asset and the target asset. In each of the 10 cases, it was unnecessary to follow the monthly volatility and correlation targeting process, and then subsequently use the replica to build an optimal portfolio. The end result is the same as that which could have been achieved through the creation of an optimal combination of the reserve and target portfolios.

The process of replication essentially adds new assets (in the form of the reserve portfolio) into an investor’s existing portfolio. In the examples provided, this entails adding commodities, small capitalization stocks, and bonds with different maturities through the distributional replication

process. The addition of these new asset classes represents a change to the original 50/50 S&P500 and Long Bond policy mix.

VERDICT ON DISTRIBUTIONAL REPLICATION

In terms of understanding the fundamentals of distributional replication and the sources of returns, our analysis leads to several conclusions. First, a comparative framework can be used to benchmark the performance of distributional replication and investigate the relative cost of different distributional parameters. Second, the performance achieved via distributional replication is dependent upon the market factors contained in the reserve portfolio. Parameters such as volatility and correlation to a target portfolio are robust with respect to the selection of reserve portfolio; however, the mean return is dependent upon the contents of the

reserve portfolio.

Next, correlation targeting represents a partial match of the correlation structure of the original hedge fund within the capital markets. A replica can be constructed to achieve a desired correlation to a specific target portfolio; however, a factor analysis would indicate a potentially large R^2 with correlations to market factors within the reserve/target portfolios. These factors and relative value trades may represent entirely new exposures for an investor, and are not likely to be present within the original targeted hedge fund. Finally, we have investigated the construction of optimal portfolios using the distributional replicas and an investor's original portfolio. For each distributional replica considered, the result is equivalent to a portfolio created using traditional mean-variance portfolio construction, without using replication. This demonstrates a redundancy associated with correlation targeting.

CONCLUSIONS

In this paper, we have investigated two common approaches to hedge fund replication and have concluded that both are limited in their ability to access the performance of hedge funds. Linear factor replication is limited to broad-based indices containing significant and slowly time-varying exposures to simple market factors. This does not represent the ideal hedge fund sought by investors. Distributional replication has been demonstrated to be redundant when placed within a mean-variance portfolio optimization. This is due to the synthetic nature of correlation targeting. In addition, for an investor with an existing well-diversified portfolio, one runs out of market factors to include within the reserve portfolio.

While hedge fund replication may be limited in its potential, such analysis leads to an improved understanding of hedge fund returns (although not necessarily useful as a replication exercise). Some of the more complex factors contained in hedge fund returns are sometimes referred to as Alternative Betas. These factors may represent trading strategies that have been profitable in the past (but perhaps won't be profitable in the future), and may also represent alternative risk

premia that are expected to be profitable over longer time horizons. This has two implications for building efficient portfolios. First, these alternative risk premia are useful for inclusion within a portfolio although, not necessarily through a replication process. Second, as these risk premia and/or trading strategies become widely recognized, they are no longer classified as alpha. The formerly alpha-rich strategies begin to resemble complex betas, implying the search for alpha is a continuous process. ■

REFERENCES:

- J. Hasanhodzic and A. W. Lo, "Can Hedge-Fund Returns Be Replicated?: The Linear Case" (August 16, 2006). Available at SSRN: <http://ssrn.com/abstract=924565>
- G. Amin and H. Kat, "Hedge Fund Performance 1990-2000: Do the Money Machines Really Add Value?" *Journal of Financial and Quantitative Analysis*, vol. 38, no. 2, June 2003, pp. 1-24.
- H. M. Kat and H. P. Palaro, "Hedge Fund Indexation the Fund Creator Way: Efficient Hedge Fund Indexation without Hedge Funds," Working Paper #0038, Alternative Investment Research Centre Working Paper Series, Cass Business School, City University, London, UK, December 7, 2006.
- Northwater Capital Management, "Northwater Capital Management's Thoughts on Hedge Fund Replication", May 2007.
- T. Schneeweis and R. Spurgin, "Multi-Factor Models of Hedge Fund, Managed Futures, and Mutual Fund Return and Risk Characteristics," *Journal of Alternative Investments*, Fall 1998.
- W. Fung and D. A. Hsieh, "Asset-Based Style Factors for Hedge Funds," *Financial Analysts Journal*, vol. 58, no. 5, September/October 2002, pp. 16-27.