NORTHWATER

NORTHWATER CAPITAL MANAGEMENT INC.

Northwater Capital Management's Thoughts on Hedge Fund Replication

May 2007

Neil Simons and Adrian Hussey

Table of Contents

List of Figures:

List of Tables:

Summary

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, etc. Linear factor replication [1] and distributional replication [2-3] have attracted the most attention. Products have been recently launched based on these replication processes. These products succeed in providing liquid investments, with fees that are typically quoted at less than 100bps per year. Our research has focused on understanding the underlying fundamentals associated with the various replication processes. We have determined that the underlying approaches are limited in their ability to access the performance of hedge funds.

Linear factor replication is the out-of-sample application of linear factor decomposition of hedge fund returns. This style of replication has reasonable success when applied to broad-based indices that possess significant linear dependence on simple market factors. Linear replication is not successful where market factors are unable to explain a large proportion of the return series, and therefore not applicable to market neutral funds or indices, funds that exhibit idiosyncratic behavior, or funds with significant alpha content. A relationship exists between the in-sample R^2 determined within the model fitting and the out-of-sample correlation, and provides a guideline for the selection of indices or funds for linear factor replication. The relationship does not imply the out-of-sample performance of the replica will match that of the original index; the relationship indicates that a reasonable out-of-sample correlation will be achieved [4].

Distributional replication attempts to re-create the distributional characteristics of an individual hedge fund, or create an investor's desired distributional characteristics (as described by parameters such as volatility, skewness, kurtosis, and correlation to an existing portfolio). Four conclusions are provided based on our research on distributional replication [5-6]. (1) A comparative framework can be used to benchmark the risk-adjusted returns obtained from distributional replication. (2) The risk-adjusted returns achieved from distributional replication are dependent upon the selection of market factors utilized within the replication process. Alteration of the market factors results in large differences in risk-adjusted performance. (3) The return series of the replica may possess a large correlation to market factors included within the replication process but excluded from the correlation specifications, and a factor analysis would indicate a large R^2. (4) Construction of optimal portfolios using replicas leads to the same solution that can be achieved using traditional mean-variance portfolio design, indicating a redundancy associated with correlation targeting.

1: Linear Factor Replication

Linear factor analysis has been extensively applied to investigate the beta component of hedge fund returns. The application of factor models to perform out-of-sample hedge fund replication is studied in [1]. Two criteria are required for linear factor models to successfully replicate a return series. (1) The return series of the hedge fund or index needs to have a large component of its behavior explained by linear dependence on market factors. (2) The hedge fund exposure to the market factors varies slowly over time. This is necessary in order for the positions suggested by the linear factor modeling to resemble the current positions within the hedge fund as opposed to the positions held on average over recent history. If net hedge fund exposures are time varying (a reasonable assumption), then there will always be a lag between the positions held by the funds and those held within the replica.

In addition to the above requirements for successful replication, the performance of replication as a trading strategy requires the trades derived from the factor model to provide good risk adjusted returns in the future.

In the following examples, we consider the linear factor replication of the CSFB Tremont, CSFB Long Short Equity, CSFB Market Neutral Equity, and HFRI Composite index using: S&P500, Russell 2000, USD Index, MSCI EAFE, MSCI Emerging Markets, and 30-Day US TBills. This factor set corresponds to that applied within the Merrill Lynch factor index [7]. Other hedge fund indices and market factors are considered in [4].

The R-squared from performing a 24-month rolling regression is provided in Figure 1. The best fit occurs for the HFRI Composite (rarely lower than 0.9); and the worst fit for the CSFB Market Neutral Equity index (rarely greater than 0.6).

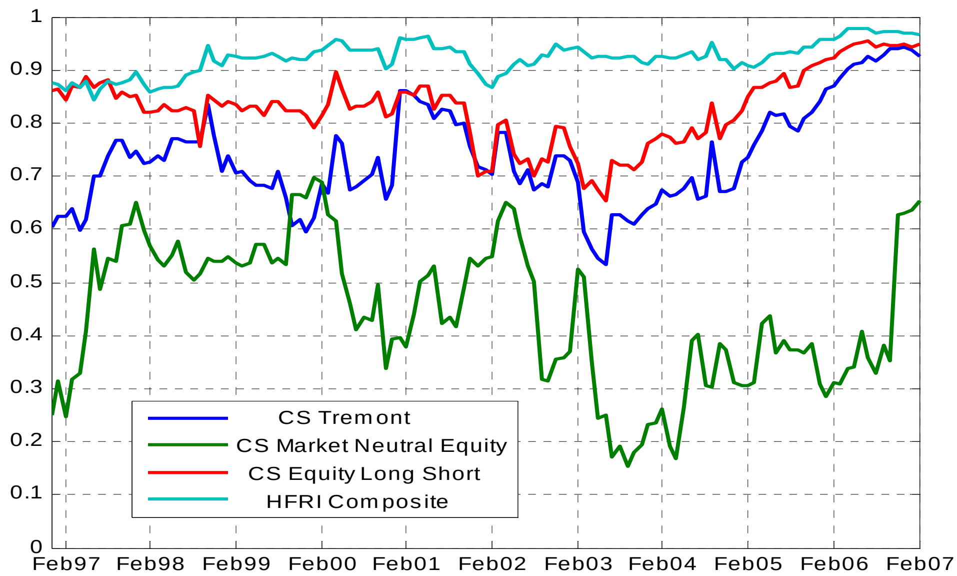

Figure 1: R^2 achieved from a rolling 24-month regression of 4 hedge fund indices.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

Performance of Linear Factor Replication – Recent 3 Year Time Horizon

The presence of a large rolling R-squared indicates the returns can be modeled accurately looking backward. In the following figures we report on the performance of rolling the factor model forward in time to perform the replication. Using the time period examined within [7], the Sharpe ratios for the four indices and their linear replicas are provided in Figure 2.

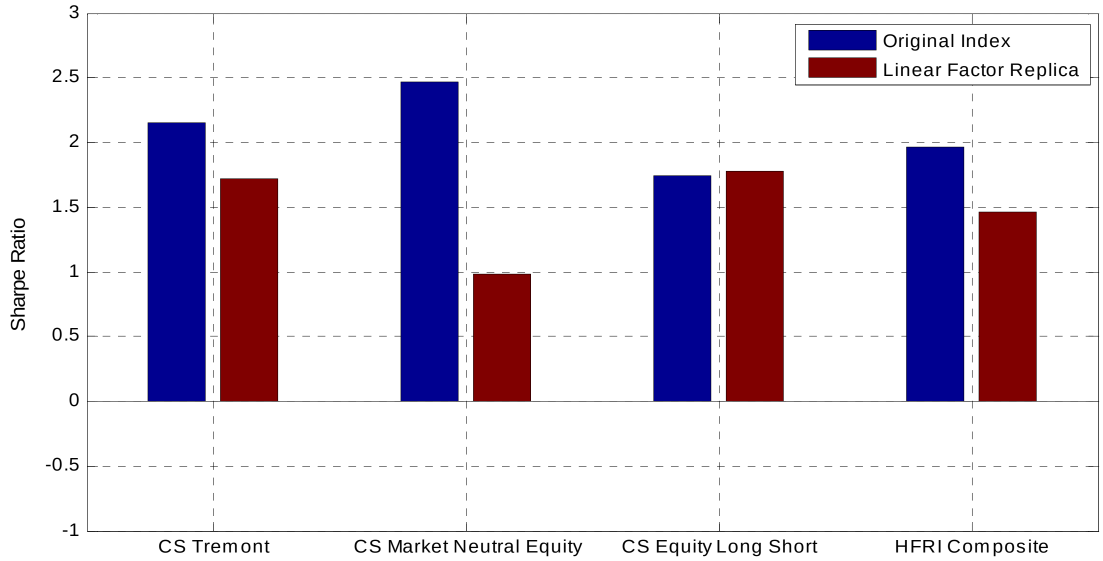

Figure 2: Comparison of the Sharpe Ratio of the Linear Factor Replica to the Original Index over the time period of July 2003 until December 2006.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

The above examples indicate that linear replication is capable of providing Sharpe ratios similar to the original indices for all cases with the exception of the CSFB 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 behavior of the index. As expected, replication process is unable to match the performance of the original index. We have found the same behavior when attempting to replicate individual funds that possess a large idiosyncratic component, indices consisting of a small number of funds, or market neutral indices with low correlation to the specified market factors [4].

Performance of Linear Factor Replication – Longer 10 Year Time Horizon

Considering a much longer time horizon (10 years), the performance is somewhat different. The results are provided in Figure 3 and indicate that the CSFB Tremont index replica provides a Sharpe ratio close to 1, capable of matching the performance of the original index. The linear replica of the HFRI index performs poorly over the 10 year time horizon (although it did provide good performance over the shorter time horizon considered above). The replica of the CSFB Equity Market Neutral index is poor on an absolute basis, and much worse in comparison to the original index.

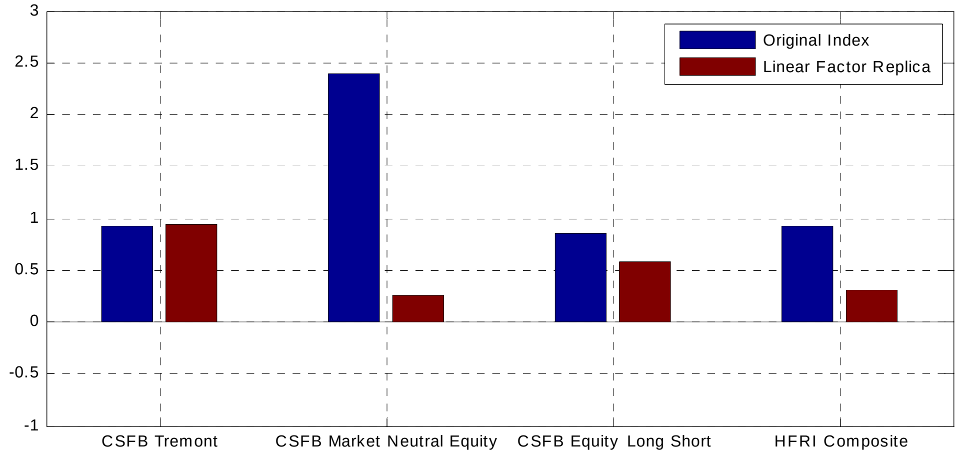

Figure 3: Comparison of the Sharpe Ratio of the Linear Factor Replica to the Original Index over the time period of January 1997 until December 2006.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

A closer examination of the return and volatility indicates that the replication has reliably reproduced the volatility of the underlying indices, however in all cases except for the CSFB Tremont index, the returns are lower (Figure 4). The replication process has delivered all of the risk of the underlying indices, but in some cases not all of the return.

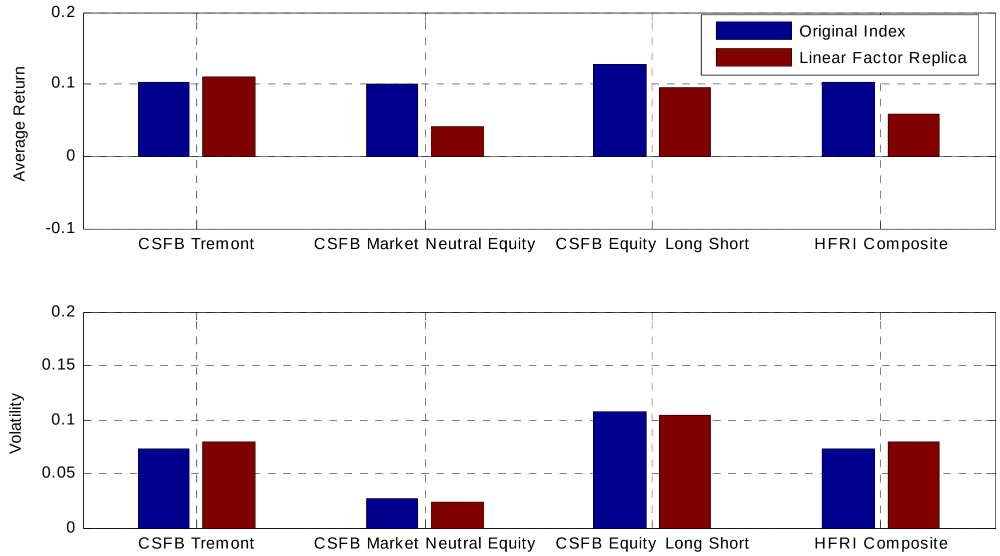

Figure 4: Comparison of the average return and the volatility of the Linear Factor Replica to the Original Index over the time period of January 1997 until December 2006.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

Comparison of In-Sample R^2 and Out-of-Sample Correlation

In Figure 5, the out-of-sample correlation of the replica and the underlying index is provided for the four indices. If the replication process has successfully reproduced the behavior of the original index, we expect a correlation close to 1. The results indicate that the CSFB Tremont, CSFB Equity Long/Short, and HFRI

Composite replicas possess a high correlation to the original index. As expected, the correlation is much lower for the CSFB Equity Market Neutral index.

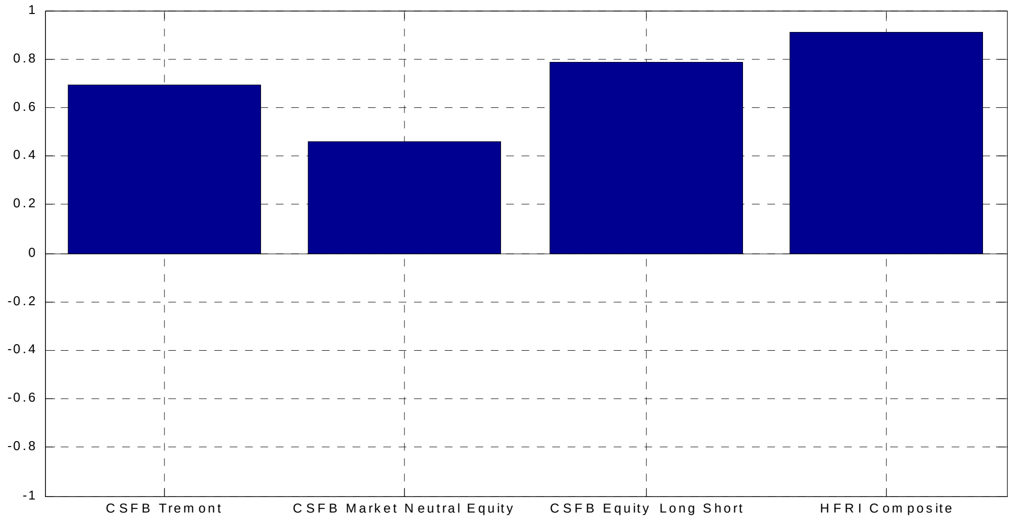

Figure 5: Out-of-sample correlation of the Linear Factor Replica and the Original Index over the time period of January 1997 until December 2006.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

An inconsistency arises when comparing the achieved replica correlation with the performance. Although replication of the CSFB Tremont index results in a lower correlation to the original index than the HFRI composite index, the CSFB Tremont replica provides better performance. This indicates the selection of an index to replicate and the expected out-of-sample performance is not straightforward. A significant historical R^2 is a requirement of linear factor replication, but it is not an indication that out-of-sample replica performance will match the original index performance in the future.

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 6. Each point within the figure corresponds to a particular index[1]. The five indices within the top right hand corner of the chart (successful replication) are the HFRI Composite, CSFB Equity Long/Short, CSFB Emerging Markets, CSFB Dedicated Short, and CSFB Tremont; and the four indices in the bottom left hand corner (unsuccessful replication) correspond to the CSFB market Neutral Equity, CSFB Managed Futures, CSFB Convertible Arbitrage, and CSFB Fixed Income Arbitrage.

A positive approximately-linear dependence is found 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: if there are no factors appearing within the historical time period, there is little chance for replication success. The relationship is not expected to be linear for R^2 values close to zero. The negative y-intercept implies that a threshold exists for the average R^2 before the replica resembles the original index, and therefore a guideline for selecting indices or funds for replication. The relationship does not imply the out-of-sample performance of the replica will match that of the original index; the relationship indicates that a reasonable out-of-sample correlation will be achieved.

[1] The set of indices considered are: the CSFB Tremont, CSFB Convertible Arbitrage, CSFB Event Driven, CSFB Market Neutral Equity, CSFB Dedicated Short Bias, CSFB Emerging Markets, CSFB Distressed Security, CSFB Risk Arbitrage, CSFB Fixed Income Arbitrage, CSFB Global Macro, CSFB Long/Short Equity, CSFB Managed Futures, CSFB Event Driven Multi-Strategy, CSFB Multi-Strategy, and the HFRI Composite.

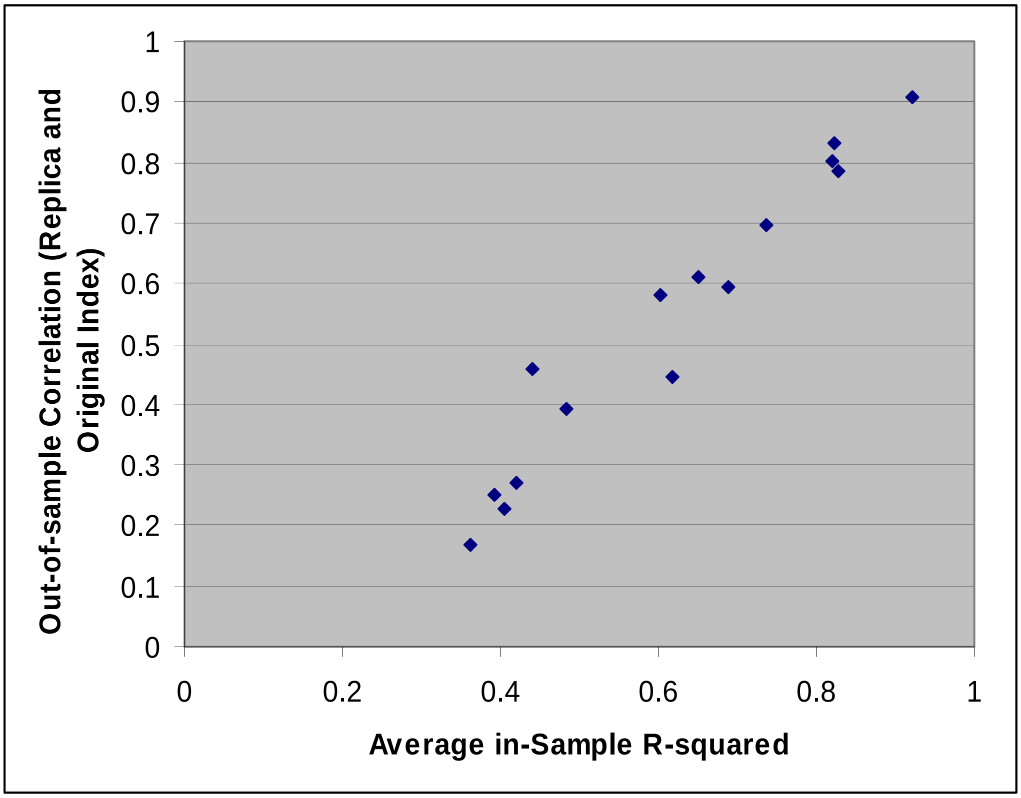

Figure 6: Relationship between the out-of-sample correlation and the average in-sample R^2 for 15 hedge fund indices.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Credit Suisse Tremont Index LLC; Hedge Fund Research, Inc.

Conclusions: Linear Factor Replication

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. The alpha component of hedge fund returns is not accessible through linear factor replication. If an investor is attempting to utilize hedge funds to provide returns that contain significant alpha and a small beta to market factors, then linear replication is not of much use.

Linear hedge fund replication is equivalent to implementing a series of time varying trades in underlying market factors. The most recent positions resulting from the linear replication of the HFRI Composite index implement the view of a strengthening US Dollar (a small long USD Index position); and the out performance of emerging market, international, and US small-cap equities relative to large-cap US equities (long MSCI EM, long MSCI EAFE, long Russell 2000 vs. short S&P500). What is the basis for implementing a short-term position in these trades? Is it rational to implement these trades simply because that is what the aggregate of hedge funds have done over the past two years?

Proponents of replication suggest that some market neutral strategies will eventually be replicated through the use of more complex market factors. The search for more complex factors is a worthwhile exercise because it helps the understanding of hedge fund returns. These factors may represent trading strategies that have been profitable in the past (but perhaps not 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. (1) These alternative risk premia are useful for inclusion within a portfolio although not necessarily through a replication process. (2) 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.

2: Distributional Replication

Distributional replication attempts to reproduce a desired return distribution through a daily trading process using a set of liquid futures contracts [2-3], [8-9]. The distribution may be that of a hedge fund or any other distribution with desirable characteristics (as described by volatility, skewness, kurtosis, and correlation to a target portfolio). The return series generated by distributional replication will not match the exact time series of the original hedge fund (i.e., it does not match the complete correlation profile). Over time the distribution of the series will converge to the desired distribution. Distributional replication has been reported to provide returns that are sometimes greater than the original hedge funds or hedge fund indices while matching their distributional characteristics [2-3].

The specific implementation of distributional replication developed by Amin, Kat, and Palaro described in [2-3], [8-9] is a complex process involving bivariate distribution fitting using copulas to specify dependence, and the derivation of a transaction-cost-efficient delta hedging schemes. Despite the complexity, the behavior of the process can be captured using models with minimal complexity. For example, in Appendix A the characteristics of the time series behavior of two of Kat and Palaro's replicas [9] are matched using static positions in the reserve and target portfolios.

We have found it useful to separate distributional replication into two subcomponents – distributional adjustments and correlation targeting. This separation has allowed our research to focus on understanding the underlying fundamentals of the approach and the sources of returns. Our conclusions are:

1) A comparative framework can be used to benchmark the performance of distributional replication and investigate the relative cost of different distributional parameters.
2) 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. A reserve portfolio with a high Sharpe ratio will tend to produce a replica with a high Sharpe ratio, and a reserve portfolio with a low Sharpe ratio will tend to produce a replica with a low Sharpe ratio.
3) 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 both individual market factors and relative value trades representing long/short combinations of 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.
4) 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.

In the following sections we provide an overview of our research and demonstrate the above points. Each specific example we provide illustrates one or more of the above conclusions.

Separation into 'Distributional Adjustments' and 'Correlation Targeting'

In order to understand the fundamentals underlying distributional replication, it is useful to separate the approach into two components. The first component (referred to as 'Distributional Adjustments') transforms the distribution of a market factor or portfolio of market factors (referred to as the reserve portfolio), into a desired univariate return distribution. The process can scale volatility, implement a performance floor, or increase the skewness of an underlying market factor or portfolio of market factors. The transformation is accomplished through a daily trading process that is similar to dynamic option replication using delta hedging. We utilize the percentile by percentile approach introduced by Amin and Kat [2] to perform distributional adjustments.

The second component to distributional replication is correlation targeting. One desirable attribute of hedge funds is a low correlation to traditional market factors. Correlation targeting creates a replica with a desired correlation with respect to one specific target portfolio. The idea of correlation targeting is an appealing one. 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. We demonstrate that volatility and correlation targeting can also be achieved through a monthly trading process where the reserve and target portfolio allocations are determined through mean-variance portfolio optimization with additional constraints.

Comparative Framework for Benchmarking the Performance of Distributional Replication

Separation of the two components of distributional replication and various numerical experiments has led us to develop a comparative framework for benchmarking the performance of distributional replication. The results of the following sections demonstrate that the performance of replication is dependent upon the selection of the market factors utilized within the reserve portfolio and therefore dependent on the excess return of the reserve portfolio (R_{RES} - r_f). Distributional adjustments alter the distributional characteristics of an underlying portfolio of market factors. A daily trading process equivalent to delta hedging is required to achieve the desired distributional characteristics. An implicit net option premium and payoff, P_{dist} will result from distributional adjustments. Targeting a negative correlation to an investor's portfolio will have a cost associated with paying away risk premium, referred to as a correlation premium, P_{corr}.

The following relationship can be used to quantify the excess return obtained from distributional replication (R_R - r_f):

$$R_R - r_f = \frac{\sigma_R}{\sigma_{RES}}\left(R_{RES} - r_f\right) - P_{dist} - P_{corr} - C_{trans}$$

The above relationship allows insight into comparative replica performance for different correlation targets and distributional adjustments; and allows for the benchmarking and comparison of replication processes. The framework can provide answers to the questions:

1) How much of the replica performance is attributable to the market factors within the reserve portfolio?
2) What is the average leverage required to achieve the volatility target (σ_R/σ_{RES})?
3) What is the impact of the distributional adjustments, P_{dist}?
4) What is the impact of the correlation targeting, P_{corr}?
5) What are the transaction costs of running the replication process, C_{trans}?

For small distributional adjustments and moderate levels of target correlation (i.e., small values of P_{dist} and P_{corr}), replica performance will be equivalent to the performance of the reserve portfolio. The reserve portfolio typically consists of market factors for which liquid futures contracts exist [3]. Over long time horizons, the expected performance of distributional replication will be very similar to a portfolio of market factors.

The use of the above framework is demonstrated through examination of the four replicas created by Kat and Palaro within [9]. The results are provided in Table 1. The cost of increasing skewness and kurtosis is indicated by the performance difference between Replica Fund 1 and Replica Fund 2 over the time horizon considered (a distributional adjustment). Similarly, the cost of decreasing correlation to a 50/50 S&P500 and Treasury Bond portfolio from 0% to -50% (representing a correlation premium) is indicated by the performance difference between Replica Fund 1 and Replica Fund 3; and the cost of implementing a floor of -5% (a distributional adjustment) is indicated by the performance difference between Replica Fund 1 and Replica Fund 4. These results represent the empirical distributional and correlation premia paid over the backtesting period considered within [9]. In this example the costs are positive. In general, the cost will depends on the path dependency of markets during the replication.

Table 1: Demonstration of the comparative framework: Benchmarking distributional replication.

	Fund 1 [9]	Fund 2 [9] +ve S, +ve K	Fund 3 [9] Negative Correlation	Fund 4 [9] (Performance Floor)
Mean Return (%)	11.42	9.52	6.81	10.41
Volatility (%)	12.35	12.80	12.21	11.79
Sharpe Ratio	0.59	0.43	0.22	0.54
‘Cost’ relative to Fund 1 (%)	0	1.90	4.61	1.01

Source: Northwater Capital Management Inc; Kat and Palaro [9].

Distributional Adjustments

We have investigated the first component of distributional replication within [5] and present a summary of our analysis in this section. Researchers from Edhec have also investigated distributional adjustments in isolation [10]. The use of distributional adjustments was introduced by Amin and Kat [2] where a percentile by percentile mapping is used in the derivation of the payoff function. This original paper did not include correlation targeting.

A variety of target distributions were considered in our research with different parameters such as positive skewness, negative skewness, low volatility, and excess kurtosis. Various reserve portfolios were considered including market factors such as the S&P500, Lehman Aggregate Index, USD Index, GSCI, and Russell 2000. A long time horizon was investigated (January 1990 to October 2006) as well as shorter time periods selected to correspond to both bull and bear markets. The conclusions from our investigation are:

1) Mean returns from the replication process are dependent on the contents of the reserve portfolio. Unless extreme path dependency of the reserve portfolio occurs within a given month, the returns from the replication process will be negative when the reserve asset returns are negative; and positive when the reserve asset returns are positive.
2) The delta functions derived were intuitively easy to understand given the desired distributions. Positive (negative) skewness resulted in increasing delta as the reserve portfolio increased (decreased) in value. Excess kurtosis required larger deltas in both tails of the distribution. Leverage is used to adjust volatility. An increase (decrease) in volatility relative to the reserve asset required a delta greater than (less than) 1.
3) Volatility is the easiest moment to modify or target.

4) Skewness can be moved in a desired direction as opposed to exactly targeted. If positive skewness is desired, it is best to start with a reserve portfolio that possesses positive or zero skewness. If a positive skewness is desired, and the reserve portfolio exhibits negative skewness over the backtesting period, the replication process may only achieve skewness close to zero. Kurtosis is the hardest moment to target, requires a long time horizon to converge, is impacted by path dependency of the reserve portfolio. Kurtosis is sensitive to the relatively small number of monthly observations.

The last conclusion does not represent a severe limitation of the replication process considering that an investor is not likely to have strong views on skewness and kurtosis (as compared to returns and volatility).

Insight into performing distributional adjustments can be gained by considering the case of dynamic option replication using delta hedging. Dynamically replicating a call option is equivalent to purchasing a call option[2]. A call option is one of many potential payoff functions. It is possible to synthesize an investor's ideal payoff function using a delta hedging process. There is an implicit price (the option premium) for the payoff that is obtained within the replication process; we refer to this net cost as P_{dist}. Receiving only the upside from a market above a certain threshold requires an investor to pay a call option premium. Requesting extreme positive skewness or a lower performance bound will require an investor to pay a premium for these features. Returns are dependent on the underlying market factor (or basket of factors). The ideal payoff function provides no return if the underlying market variables do not terminate in the profitable area of the payoff function. Replicating a call option is a profitable strategy only if the underlying market factor increases above the strike price and option premium.

Correlation Targeting

The second component to distributional replication is correlation targeting. One desirable attribute of hedge funds is a low correlation to traditional market factors. Correlation targeting creates a replica with a desired correlation with respect to one specific target portfolio. The idea of correlation targeting is an appealing one. 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. While intra-month dynamic trading is required to perform distributional adjustments, it is not necessary to achieve correlation and volatility targeting.

Volatility and correlation targeting can be achieved using a monthly trading process derived using mean-variance portfolio optimization with additional constraints. The replica is constructed using a position in the reserve portfolio, w_{RES} and the target portfolio, w_T using recent historical returns and then applied out-of-sample over the following month. The replica's returns can be represented as:

$$r_R = w_{RES} r_{RES} + w_T r_T + (1 - w_{RES} - w_T) r_{TBILL}$$

where r_{RES}, r_R, and r_{TBILL} represent the returns of the reserve portfolio, target portfolio, and treasury bills, respectively. An optimization is used to select the weights, w_{RES} and w_T such that the return stream of the replica r_R would match the volatility target, and the correlation target. The MATLAB function fmincon (constrained minimization) is utilized where the objective function is defined as

$$f_{obj} = -(E(r_R) - \sigma_R \sigma_R)$$

with the constraints: σ_R equal to the target volatility, and $\rho(r_R, r_T)$ equal to the target correlation. The optimization is free to select long or short positions in the reserve or target portfolio. Leverage is implicitly allowed within the replica definition given above and is financed at the Treasury-bill rate.

[2] Using the usual assumptions within a Black-Scholes world with a constant realized volatility equal to the implied volatility, no jumps in the underlying market variable, etc.

A historical window of 24 months (using monthly data) is used to perform the optimization. The process is rolled forward in time to determine the new monthly weights for the reserve and target portfolios. The above process will only achieve a match of the correlation and volatility, the skewness and kurtosis will be determined by the market behavior of the two portfolios.

The reserve and target portfolios specified within [3] are utilized: the target portfolio is a 50/50 mix of the S&P500 and 30-year Treasury Bonds; and the reserve portfolio consists of the GSCI, Russell 2000, S&P500, Eurodollar Futures, 2-year Treasury Bonds and 10-year Treasury Bonds. The Eurodollar and the 2-year Bonds are given a weight of 5 and 4 respectively in the reserve portfolio, with all other assets given a weight of 1. A small adjustment to the reserve and target portfolios has been made in order to utilize monthly return data provided by Ibbotson and Associates. We substitute the Ibbotson U.S. Long Bond return series for the 30-year Treasury, and the Ibbotson U.S. Intermediate Bond return series as a replacement for the 2-year and 10-year bonds. The intermediate bond series is leveraged by a factor of 4 within the reserve portfolio.

Correlation Targeting and Portfolio Diversification

The reserve and target portfolios described above have a positive correlation of 70% over the time period considered. Specification of a correlation less than 70% to the target portfolio will require establishing a short position in the investor's portfolio within the replication process. For example, Figures 1 and 9 within reference [9] indicate a short position in the target portfolio based on the slope of the payoff function. A short position in an asset exhibiting a positive risk premium will result in paying away risk premium. As discussed previously, we refer to this implicit cost as a correlation premium, P_{corr}.

Correlation targeting can be considered as an indirect method of achieving portfolio diversification. A change in the asset mix occurs when a portion of the investor's existing portfolio is sold (short position in the target) and new asset classes are added (long position in the reserve portfolio, which contains market factors such as the Russell 2000, GSCI, and collateralized Eurodollar futures that are absent from the investor's original portfolio).

The short target portfolio position within the replica can be avoided if the reserve portfolio is constructed using assets that exhibit a low correlation to the target portfolio. This is difficult to achieve, since there are not many asset classes that exhibit this type of behavior that would not already be contained in a well diversified investor's portfolio.

Performance of the Monthly Correlation Targeting Process

We have applied the above process to replicate 11 Edhec Hedge Fund indices over the time period of March 1999 to September 2006. This is the same set of indices replicated and time period used by Kat and Palaro in [3]. In Figure 7, the monthly replicas are compared to the replicas produced by Kat and Palaro [3]. The numbers on the x-axis refer to the specific Edhec Indices[3]. The results indicate that the monthly process has provided correlation, volatility, and mean excess returns very similar to those produced by Kat and Palaro in [3].

We have not applied distributional adjustments, and therefore the skewness and kurtosis of the monthly replicas will not match that of the original Edhec indices. Recall our research focuses on understanding the fundamentals of the approach, and not building an exact 'replica' of a replication process.

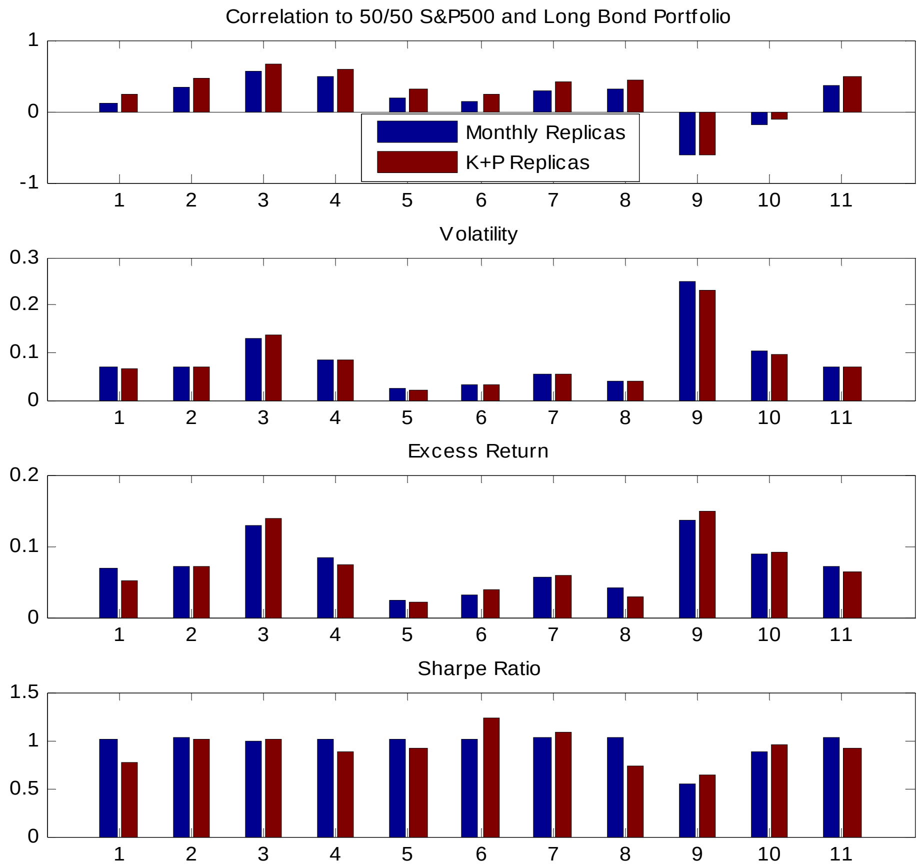

Figure 7: (Replica Set 1) Comparison of the performance of the Monthly Replicas to the Kat and Palaro Replicas from [3]: Correlation to Target Portfolio, Volatility, Excess Return, and Sharpe Ratio.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Edhec-Risk; Kat and Palaro [3].

The results of Figure 7 demonstrate that the monthly replication process is sufficient to achieve a wide variety of behavior exhibited by the Edhec indices, and matches the risk-adjusted performance of Kat and Palaro

[3] Edhec Indices Examined:
1 – Convertible Arbitrage; 2 – Distressed Securities; 3 – Emerging Markets; 4 – Long/Short Equity
5 – Equity Market Neutral; 6 – Fixed Income Arbitrage; 7 – Global Macro; 8 – Merger Arbitrage
9 – Short Selling; 10 – CTA Global; 11 – Fund of Funds

process. Given that we have constructed a suitable replication process, we demonstrate the performance attributes of distributional replication i.e., the impact of the reserve portfolio on performance, and the factor exposure of the replicas.

The results (Replica Set 2) provided in Figure 8 were generated using a different mix of market factors within the reserve portfolio. The performance of this set of replicas is poor in comparison to the replicas presented in Figure 7. The S&P500 Information Technology and Telecommunication Services sectors are included within the reserve portfolio, and the Russell 2000 and GSCI are excluded. The correlation and volatility targets are successfully met; however the mean excess returns are now very poor when compared to the previous results.

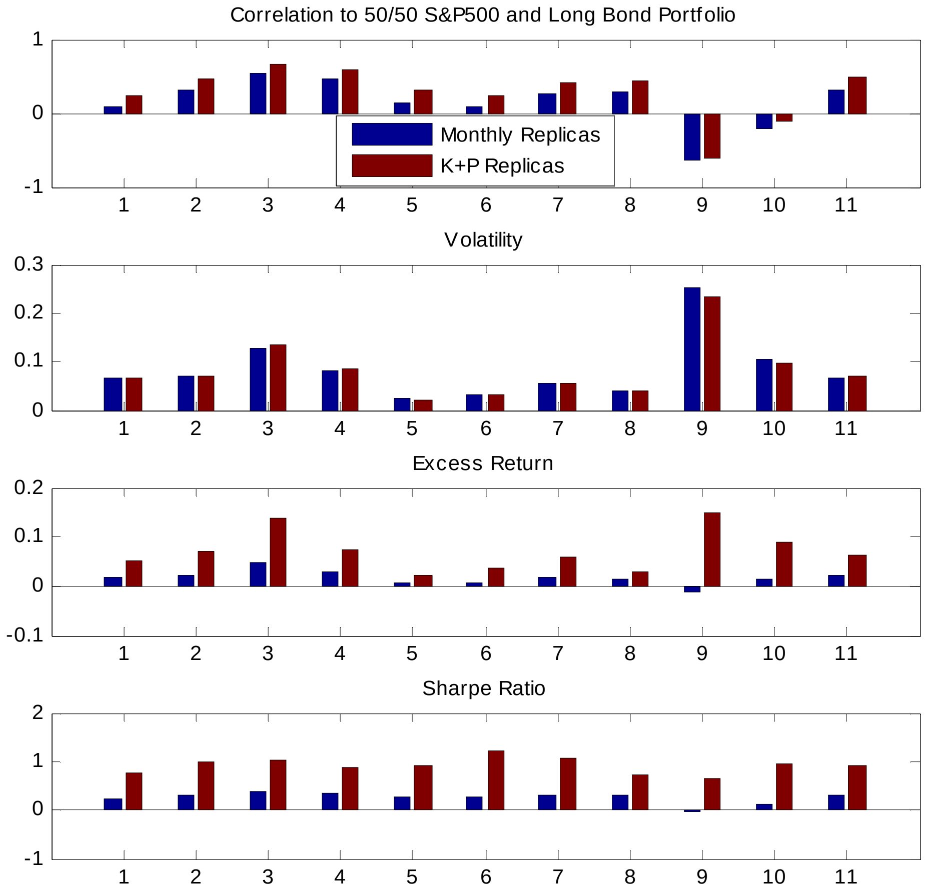

Figure 8: (Replica Set 2) Impact of a change in Reserve Portfolio composition (GSCI and Russell 2000 excluded; and S&P500 Information Technology and Telecom Services sectors included): Correlation to Target Portfolio, Volatility, Excess Return, and Sharpe Ratio.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Edhec-Risk; Kat and Palaro [3].

Replica Set 3 utilizes a different mix of market factors again, but this time provide mean excess returns higher than the previous replicas in every case except one (index 6 which corresponds to the Fixed Income Arbitrage index; and is likely due to a positive payoff from distributional adjustments within [3]). The results are provided in Figure 9. To achieve this performance, the reserve portfolio is altered to contain long positions in the 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.

Changes in the reserve portfolio did not affect the ability of the process to match the volatility or correlation targets. Therefore, volatility and correlation targeting is robust with respect to changes in the reserve portfolio, but mean returns are not robust to changes in the reserve portfolio.

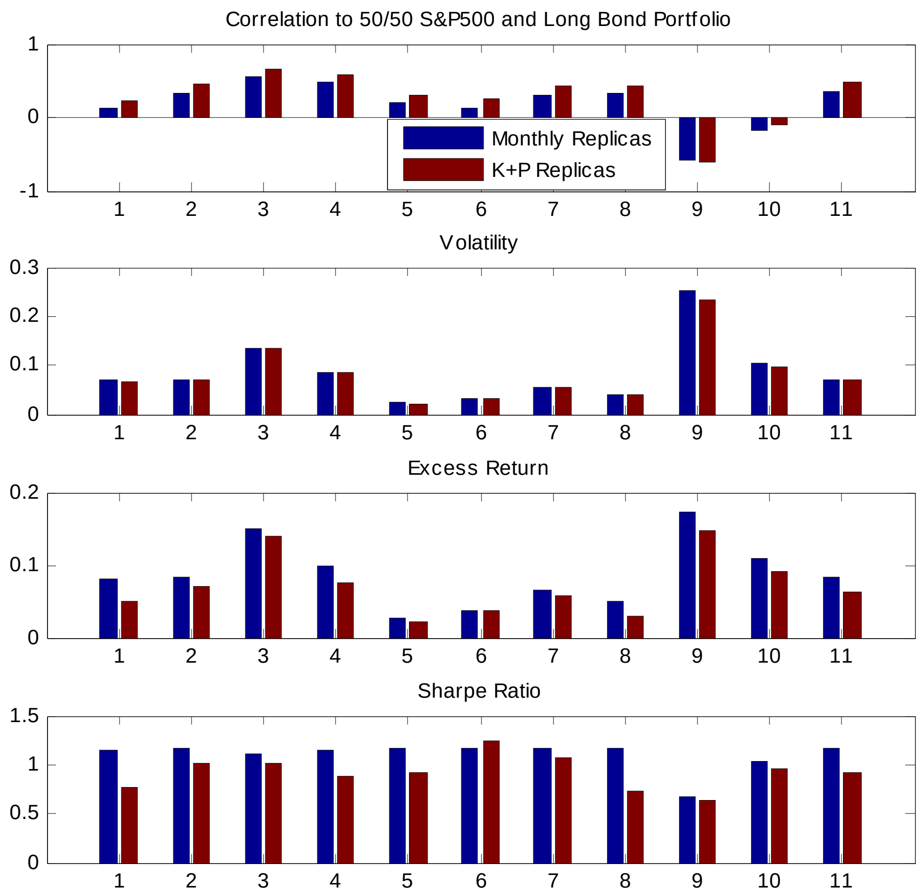

Figure 9: (Replica Set 3) Impact of a change in Reserve Portfolio composition (short positions in the S&P Information Technology and Telecom Services sectors and long positions in the Homebuilding and Aerospace & Defense sub-industry groups): Correlation to Target Portfolio, Volatility, Excess Return, and Sharpe Ratio.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Edhec-Risk; Kat and Palaro [3].

Implications of a Partial Match to the Correlation Profile: GSCI and Hurricane Katrina

In this section, we examine the time series behavior of the replicas. The results in Appendix A indicate the performance of replicas can resemble that of a static linear combination of the reserve and target portfolio. The typical correlation constraint references a 50/50 S&P500 and Long Bond portfolio, without matching the correlation profile of the original fund to other market factors. This is expected since the goal of distributional replication is to only replicate the distribution of the underlying hedge fund or index and not the exact time series of the returns of the original fund or index. A complete match of the correlation profile would result in an exact match of the time series. While this is not the goal of distributional replication, it is informative to examine the time series behavior of the replica. In this section, we use the original reserve portfolio considered in [3].

The behavioral match provided in Appendix A is not surprising given the structural replica composition of a long position in the reserve portfolio and a position in the target portfolio which depends on the specific correlation target. If there is a positive correlation between the reserve and target portfolios, the position in the target is typically short in order to reduce the replica's correlation to the target portfolio. Given the structural replica position, one would expect the replicas to possess correlations to the components of the reserve portfolio as well as correlations to relative value trades consisting of long a market factor in the reserve portfolio and short a market factor in the target portfolio. Essentially, the replication process is adding new market exposures not present within the investor's existing portfolio.

As indicated in Table 2, the replicas have a large positive correlation to the GSCI, where the original index does not. The replicas examined in Table 2, represent a sub-set of those presented in Figure 7.

Table 2: Large correlation of replicas to market factors contained within the reserve portfolio.

Correlation to GSCI	Original Edhec Index	Monthly Replica
Edhec Convertible Arbitrage	0.07	0.79
Edhec Distressed Securities	0.00	0.74
Edhec Equity Market Neutral	0.27	0.78

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

The large GSCI correlation is due to the inclusion within the reserve portfolio. We expect the correlation of the Kat and Palaro replicas to the GSCI [3] to be less than the values presented above due to the daily trading used to perform the distributional adjustments.

To illustrate the impact of the GSCI correlation, we compare the returns of the replica and the original Edhec index in a few notable months in 2005 and 2006. In Table 3, the monthly returns of the GSCI, the original Edhec Distressed Securities index, and the monthly Distressed Securities replica are provided. The replica mirrors the moves in the GSCI, while (as expected) the original Edhec index does not.

Table 3: Comparison of the returns obtained from the replica to the returns of the original Edhec index for months with large movements in the GSCI.

Monthly Return	GSCI	Monthly Replica	Original Edhec Distressed Securities Index
August 2005	15.1%	4.9%	1.2%
October 2005	-10.3%	-4.4%	-0.3%
February 2006	-10.1%	-3.4%	0.6%

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

August and October represent the largest monthly gain and loss, respectively for the GSCI in 2005 (recall Hurricane Katrina's impact on energy markets). These months also represent the largest monthly gain and loss for the replica of the Edhec Distressed Securities index for 2005. The largest monthly gains and losses for the actual Edhec Distressed Securities index occurred in different months (July and April of 2005, respectively). The above example illustrates the implications of matching only a portion of the correlation profile of the original index, and also illustrates the dependence of replica performance on the contents of the reserve portfolio.

Implications of a Partial Match to the Correlation Profile: Dedicated Short Replicas

Kat and Palaro's replicas significantly outperform a variety of short-selling hedge fund indices. The performance of the Edhec Short Selling index is compared to the replica in Table 4.

Table 4: Comparison of the un-smoothed Edhec Short-Selling index to the Kat and Palaro replica.

	Un-Smoothed Edhec Index[4]	Kat+Palaro Replica [3]
Excess Return	-3.58%	14.80
Volatility	21.57%	23.24%
Correlation to 50/50 Portfolio	-62%	-61%

Source: Kat and Palaro [3].

The replica provides an excess return 18% greater than the un-smoothed index while matching the volatility as well as the correlation to a 50/50 S&P500 and Long Bond portfolio.

The monthly process presented in the previous section is used to produce a short-selling replica. The results for the replica generated using our monthly process are compared to those of Kat and Palaro's replica in Table 5.

[4] Kat and Palaro use an un-smoothing process [3] to remove autocorrelation from the original Edhec index data [10]. The un-smoothing will increase the volatility, and therefore decrease the Sharpe ratio of the index.

Table 5: Comparison of the Monthly Replica to the Kat and Palaro short-selling replica.

	Monthly Replica	Kat + Palaro Replica [3]
Excess Return	13.50%	14.80%
Volatility	24.89%	23.24%
Correlation to 50/50 Portfolio	-61%	-61%

Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Edhec-Risk; Kat and Palaro [3].

The results contained in Table 5 indicates that a replica has been produced that is similar to that published in [3] (also previously demonstrated in Figure 7, where the short selling index corresponds to index 9). This monthly replica also outperforms the targeted Edhec index by a wide margin.

In Table 6, the monthly replica is compared to the original Edhec index and a few more correlation relationships are included.

Table 6: A more complete correlation profile of the Short-Selling replica. Comparison of the replica and original index for correlation to GSCI, S&P500, and Russell 2000.

	Original Edhec Index	Monthly Replica
Excess Return	1.75%	13.50%
Volatility	18.63%	24.89%
Correlation to 50/50 Portfolio	-64%	-61%
Correlation to GSCI	***-7%***	***65%***
Correlation to S&P 500	***-78%***	***-50%***
Correlation to Russell 2000	***-85%***	***-3%***
Correlation to R2000-SPX	***-40%***	***45%***

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

The last four rows of Table 6 explain the out-performance of the replica in comparison to the actual Edhec Short-Selling index. The Edhec index has a large negative correlation to the S&P500 and Russell 2000 (expected behavior for a short selling index). The replica has much lower correlations to the two indices, especially the Russell 2000. The replica pays away much less of the equity risk premium (especially to the Russell 2000) which is a large performance drag within the actual short-selling index. The Edhec index has a low correlation to the GSCI (again expected behavior) while the replica has a large positive correlation to the GSCI. As well, the Edhec index has a negative correlation to the 'long Russell2000 and short S&P500' relative value trade, while the replica has a positive correlation (essentially a co-linearity effect given the relationship with the S&P500 and Russell2000). Again, the positive correlation to the GSCI and R2000/SPX relative value trade are sources of positive return for the replica due to their strong performance over the time period considered. All of the above correlations are due to the selection of market factors within the reserve portfolios and specification of the target portfolio.

The replica has reproduced the distributional parameters as specified in terms of volatility and correlation to the 50/50 portfolio. However, the replica is not a 'clone' of a short selling hedge fund index. It has a much smaller correlation to the S&P500 and has an almost zero correlation to the Russell 2000. The replica would no longer be a useful small cap hedge. However, it was not specified to be a small cap hedge, it was only specified to possess a -60% correlation to a 50/50 portfolio of the S&P500 and Long Bonds. The replication processes provide a portfolio to match this one specific correlation relationship.

In the example given above, the significant out-performance of the hedge fund replica over the original index is due to a partial match of the correlation structure of the original index. A more rigorous test of the replica performance would be to complete the correlation specifications to include the large negative correlation to the S&P500 and Russell 2000 and the low (slightly negative) correlation to the GSCI. This is a more complete description of the behavior of the Edhec index, but a contradiction to distributional replication. Recall that the goal of distributional replication is to replicate a return distribution, but not the exact time series of returns. A complete match of the correlation profile would result in replication of the time series of returns (this is the goal of other replication processes, such as linear factor replication).

Given the above behavior, analysis of the time series produced by distributional replication would indicate a significant R^2. Factor loadings would correspond to the contents of the reserve portfolio and also relative value trades representing long/short combinations of market factors within the reserve/target portfolios. These may represent entirely new exposures for an investor.

Optimal Portfolio Construction

Recall that replication is an attempt to replace hedge funds i.e., an investor would allocate a portion of their portfolio to replicas rather than the actual hedge funds. Correlation targeting is intended to provide replica with desirable characteristics (low correlation to the target portfolio). In this section, optimal portfolio construction is investigated using the correlation targeted replicas of the 11 Edhec hedge fund indices presented in the previous sections (see Figure 7). The analysis contained in this section indicates:

1) The reserve portfolio is superior to the target portfolio in terms of risk and return over the time period considered.
2) As a result of this superior performance, the optimal combination of the reserve and target portfolios corresponds to an approximately 120% long position in the reserve and -20% short position in the target. Therefore, in order to create this optimal combination, an investor is required to establish a short position in their existing portfolio and a leveraged long position in new assets.
3) As a result of this unusual situation, the replicas represent near-optimal combinations of the reserve and target portfolios. This near-optimal replica behavior is not due to the replication process. The near-optimal behavior is a result of the selection of the specific reserve and target portfolios over the time period considered, and the structural replica position of being long the reserve portfolio and typically short the target portfolio (in order to reduce correlation to the target).
4) Alteration of the problem specification (reversing the role of the reserve and target portfolios) leads to sub-optimal replicas.
5) Optimal combination of the replicas with the investor's original portfolio leads to the same solution as that provide in 2) above. Therefore, we demonstrate a redundancy associated with correlation targeting.

On a conceptual level, distributional replication adds new assets (the contents of the reserve portfolio) into an investor's existing portfolio (typically specified as the 50/50 target portfolio). To examine this situation in more detail, consider the construction of an optimal portfolio using two assets. The reserve and target portfolios represent the two assets. In Figure 10, the return and risk characteristics are provided for: the target portfolio (50/50 mix of S&P500 and Long Bonds), the reserve portfolio (the original reserve portfolio used by Kat and Palaro from [3] defined in a previous section), and various combinations of the reserve and target portfolios (the yellow line). The reserve portfolio is a much better portfolio than the 50/50 portfolio over the time horizon considered. The optimal combination consists of a 120% allocation to the reserve portfolio and a -20% allocation to the target portfolio. Therefore, if we consider the reserve portfolio as a new asset available to an investor, in order to create an optimal portfolio, an investor is required to establish a short position in their existing portfolio and a leveraged long position in the new asset. The blue line represents the Tobin-frontier constructed using leveraged positions in this optimal combination [11].

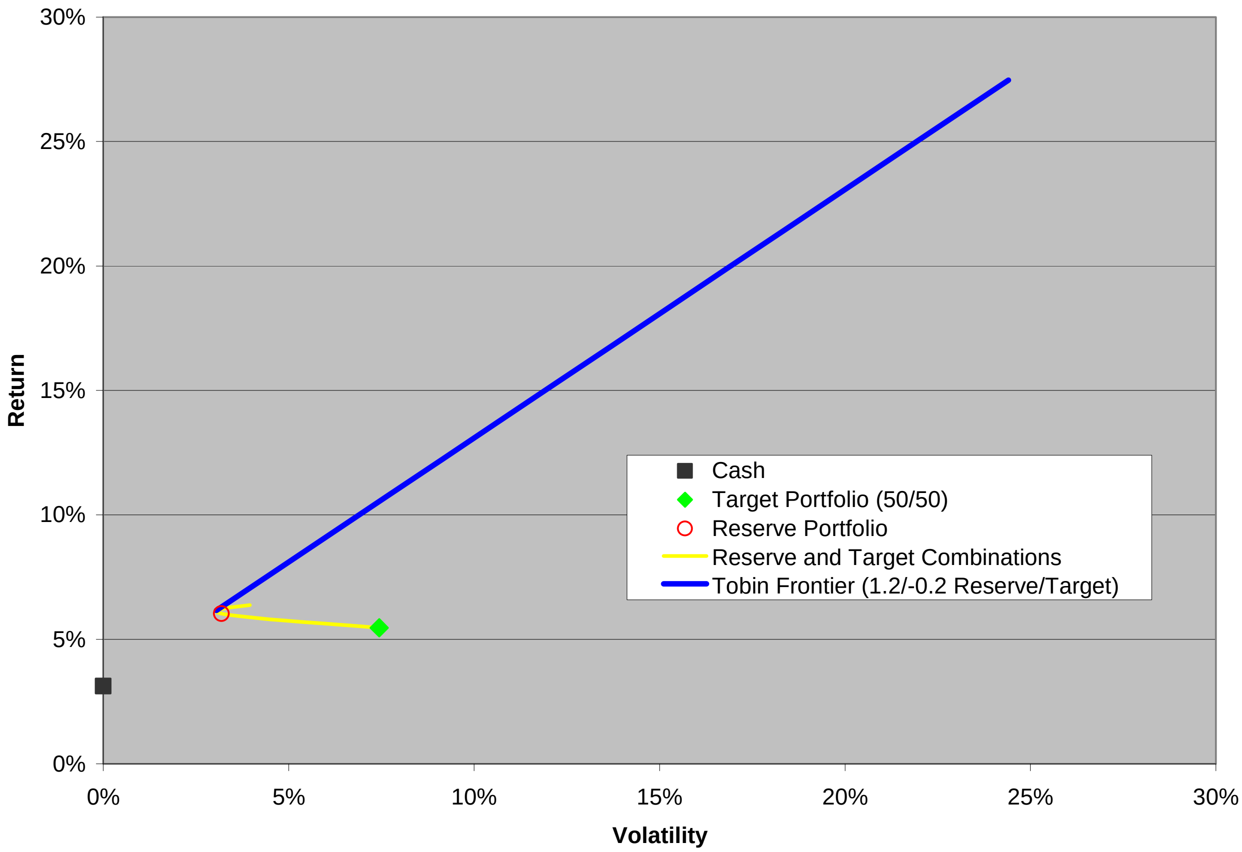

Figure 10: Risk and Return Characteristics: Target Portfolio, Reserve Portfolio, Combinations of the Reserve and Target Portfolios, and the Tobin Frontier.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg.

Next, we compare the risk and return characteristics of the replicas to the optimal combination. In Figure 11, the risk and return characteristics of the Kat and Palaro replicas [3] and those generated using the monthly process described in this document are provided, representing replicas of the 11 Edhec hedge fund indices. These are the same risk and return results provided in Figure 7. Many of the Kat and Palaro replicas and the monthly replicas exist along or close to the Tobin-frontier. One might conclude that the near-optimal behavior of the replicas is a result of the replication process. However, it is actually due to the behavior of the reserve and target portfolios over the time horizon considered and the basic structural position embedded within the replicas. Distributional replication establishes a long position in the reserve portfolio, and then typically a short position in the target portfolio in order to reduce correlation to the target. We refer to this as the 'structural position' of distributional replication. We will demonstrate that reversing the roles of the reserve and target portfolios results in replicas well below the Tobin-frontier.

The amount of leverage applied within each of the replicas can be extracted from the information within Figure 11, given the reserve portfolio volatility of approximately 3%. For example, the replica with a volatility of approximately 14% (corresponding to the Edhec Emerging Markets index) requires a 4 to 5 times leveraged position in the reserve portfolio to achieve this volatility target.

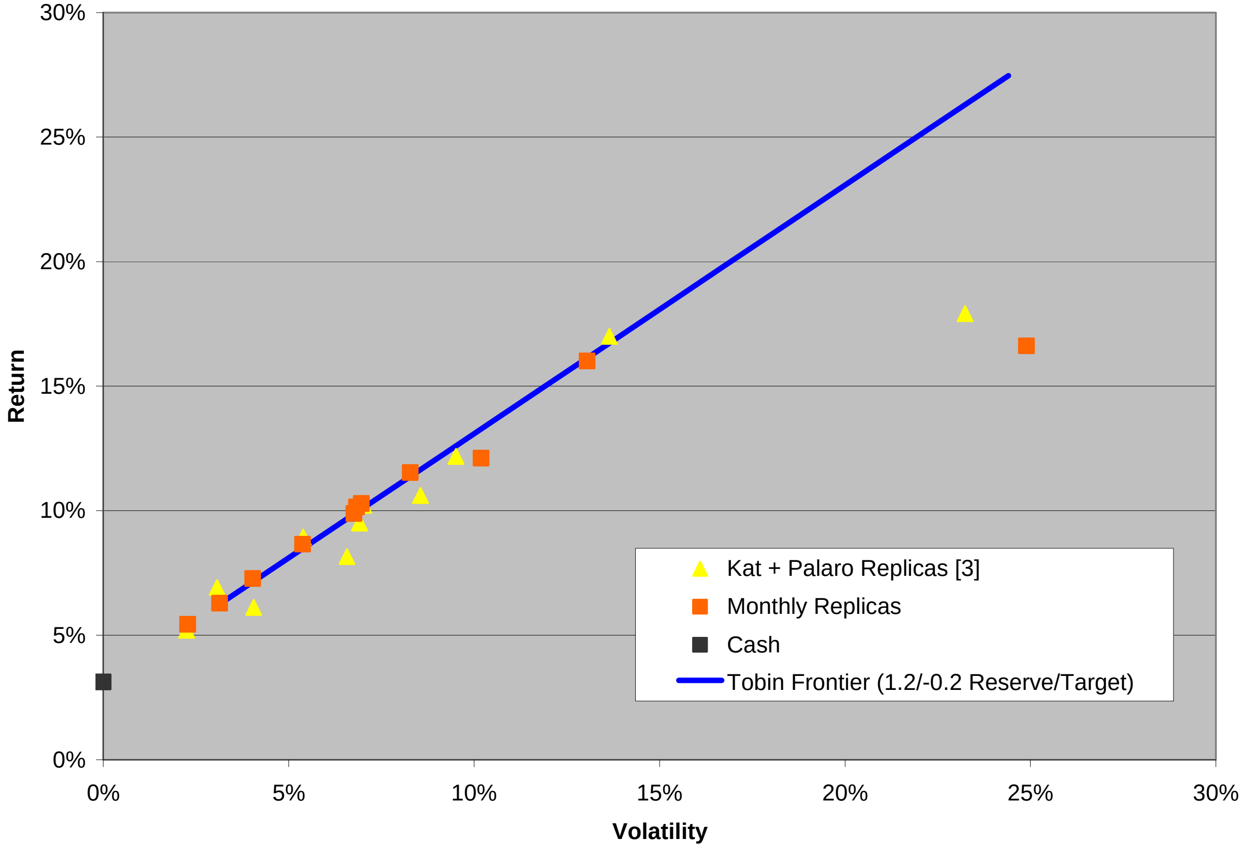

Figure 11: Risk and Return Characteristics: Comparison of the Monthly Replicas and the Kat and Palaro replicas from [3], and the Tobin Frontier.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Kat and Palaro [3].

We have created optimal portfolios combining the monthly replicas with the original target portfolio. For each of the 11 Edhec replicas, we determine the optimal combination of the replica and the original target portfolio. The risk and return characteristics of the optimal portfolios are presented in Figure 12. The optimal portfolios lie along the Tobin-frontier. Therefore, for each of the 11 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 11 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.

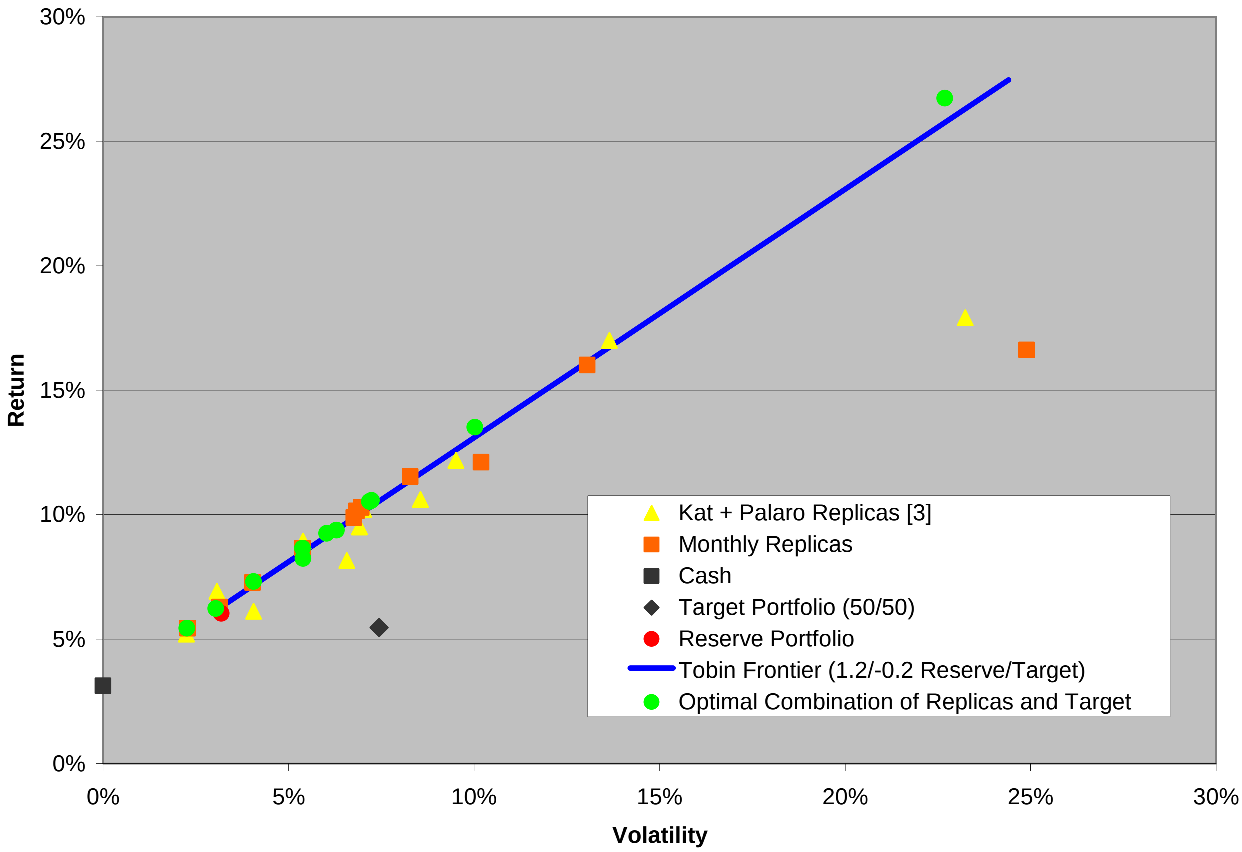

Figure 12: Risk and Return Characteristics: Optimal Portfolio constructed using the monthly replicas and original target portfolio, and comparison to the Tobin Frontier.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg; Kat and Palaro [3].

We have created an efficient frontier freely allocating to all of the assets within the reserve and target portfolios (S&P500, Russell 2000, Intermediate Bonds, Long Bonds, collateralized Eurodollar Futures, and GSCI) rather than allocating to the target and reserve portfolios as two assets. This frontier is superior to the Tobin-frontier related to the optimal mix of reserve and target portfolios (see Figure 13). The performance gain is due to the complete flexibility in selecting the weights of the market factors, where the weights are fixed within the reserve and target portfolios.

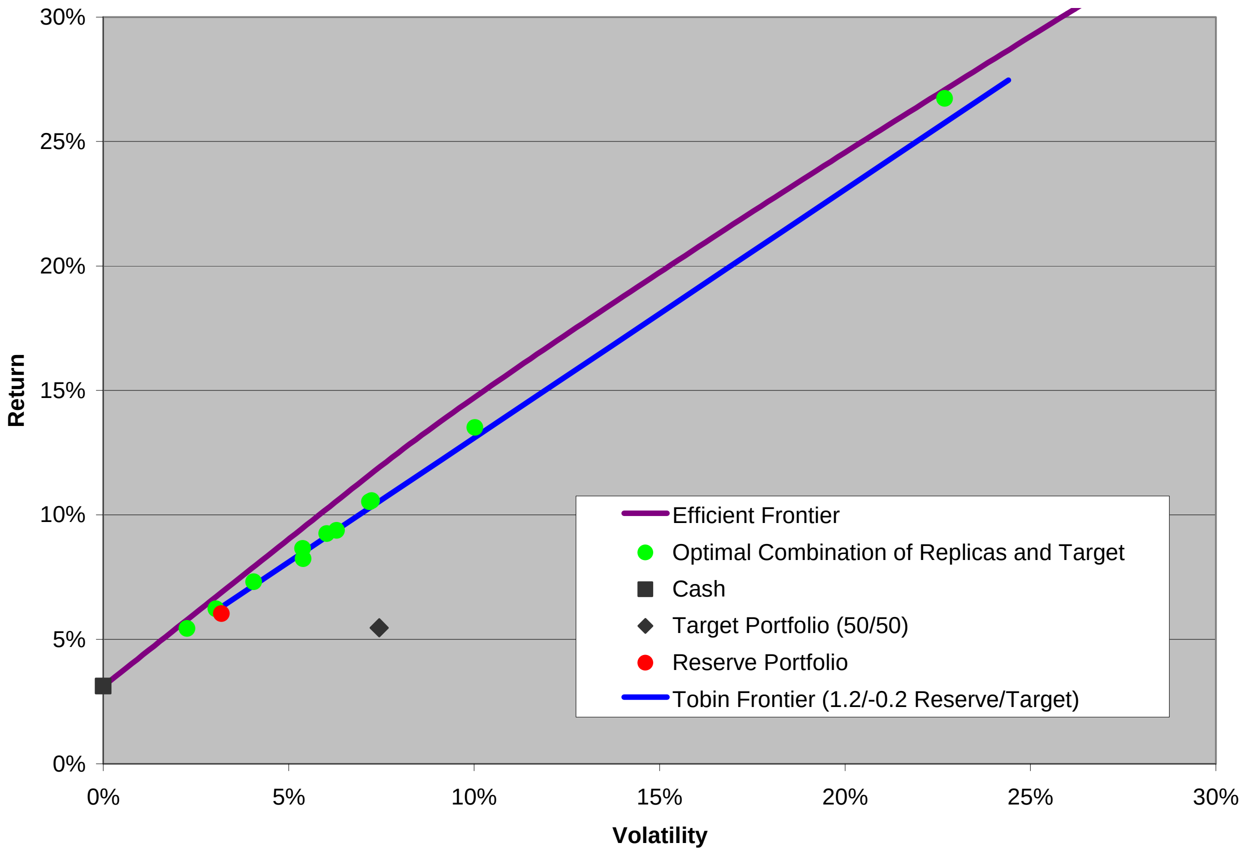

Figure 13: Risk and Return Characteristics: Comparison of the completely flexible efficient frontier and the Optimal Portfolios constructed using the monthly replicas.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg.

Structural Position of Distributional Replicas

The typical example considered in [3] and within this document consists of an investor's portfolio represented by a 50/50 mix of the S&P500 and Long Bonds; and the reserve portfolio represented as a mix of commodities, large and small caps, bonds, and collateralized Eurodollar futures. The structural position of the replica (long the reserve and short the target to reduce correlation) is beneficial over the time period considered since we have demonstrated that the optimal portfolio mix for the reserve and target portfolios is a position of 1.2 in the reserve and -0.2 in the target. The result of the structural position and behavior of the reserve and target portfolios over the time period considered is that most replicas represent near-optimal combinations of these two assets.

Consider an example where the roles of the reserve and target portfolios are reversed. An investor's portfolio could contain exposure to large cap stocks (S&P500), small cap stocks (Russell 2000), commodities (GSCI), Intermediate Bonds, and collateralized Eurodollar futures. Therefore in this case, the target portfolio is the previously applied reserve portfolio. We can select the 50/50 portfolio of S&P500 and Long Bonds as the reserve portfolio. A long position in the 50/50 portfolio and a short position in the diversified portfolio will be required. In this case, the replicas will possess poor performance due to the structural position of distributional replication. The results in Figure 14 indicate that each replica is a sub-optimal combination of the reserve and target portfolios, despite achieving the desired volatility and correlation targets.

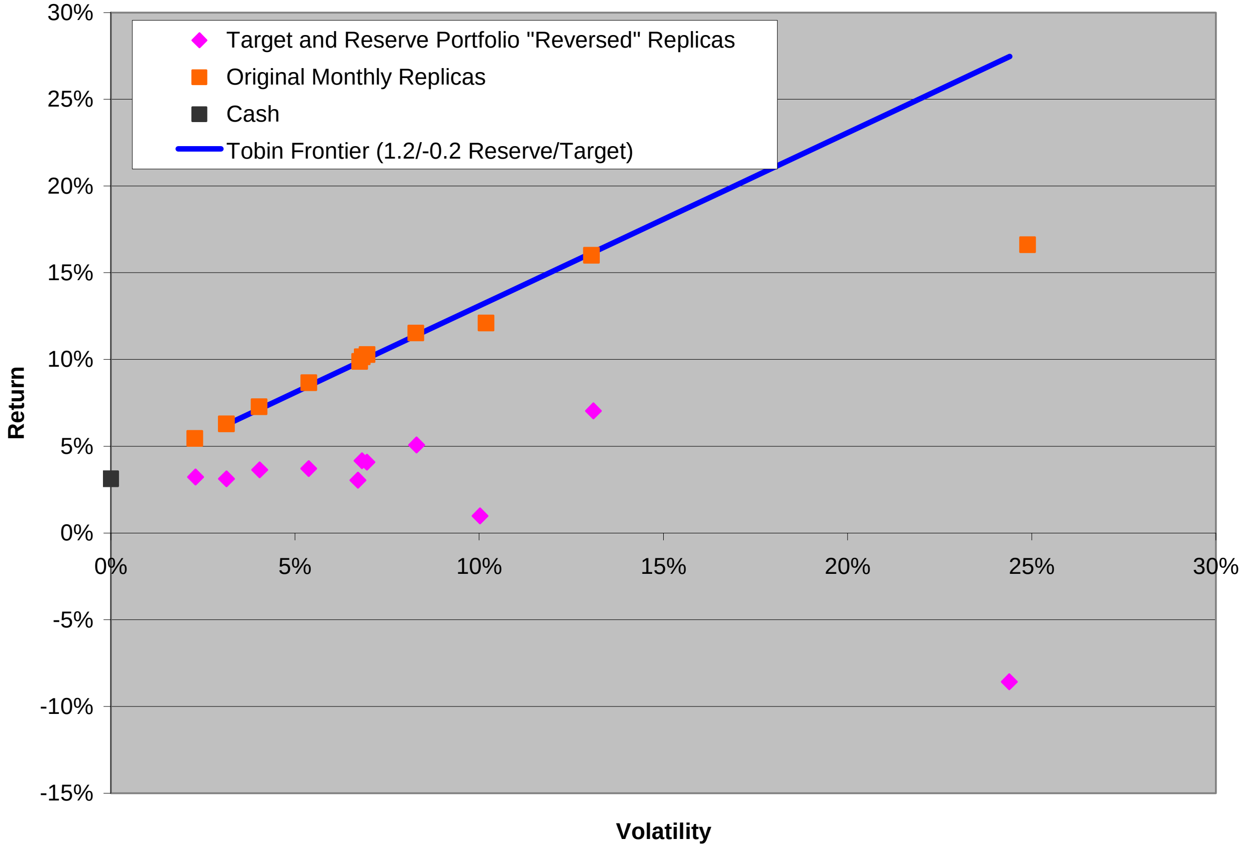

Figure 14: Risk and Return Characteristics: Replica performance when the role of the original Reserve and Target Portfolios is reversed.
Source: Northwater Capital Management Inc; Ibbotson Associates, Inc.; Bloomberg.

Another way of looking at the original set of replicas is that they represent a switch out of the S&P500 and Long Bonds (the typical short position in the target portfolio) and into new asset classes: Commodities, small caps, and bonds across the maturity spectrum (due to the long position in the reserve portfolio). Investors might believe this is a useful exercise, as part of their overall portfolio design, but perhaps not embedded within a replication process.

The replication process utilizes market factors as return drivers. If these market factors are not already part of an investor's portfolio, it is worth considering why the factors would they be added via a replication process. These market factors may have been overlooked within the portfolio construction, and therefore worth adding into the portfolio (although not necessarily through a replication process). Alternatively, they may have been explicitly excluded, and therefore should not be added back through a replication process. Regarding the construction of a well-diversified reserve portfolio, if an investor's portfolio contains a wide variety of market factors, what is left over to use within the replication process?

Static vs. Dynamic Efficient Frontiers

The efficient frontier provided in Figure 13 represents the optimal selection of static positions within the market factors contained within the reserve and target portfolios. This frontier is superior to the Tobin-frontier representing the optimal mix of reserve and target portfolios (the individual market factors have fixed weights within the reserve and target portfolios). The replication process of Kat and Palaro and the monthly replication process presented in this document perform a time-varying allocation to the reserve and target portfolios. Therefore, it is possible for these replicas to outperform the Tobin-frontier due to the dynamic as opposed to static allocation. For example, as shown in Figure 12, two of the optimal portfolios constructed using the monthly replicas (around volatility levels of 10% and 22%), and one of the original Kat and Palaro replicas (volatility level of 3%) exist slightly above the Tobin-frontier.

The above observation leads to the idea of portfolio construction using a completely dynamic asset allocation model. It is not obvious if this allocation would provide systematically better returns than the static allocation.

Conclusions: Distributional Replication

Our research has focused on understanding the fundamentals of distributional replication and the sources of returns. The analysis contained in this document has led to the following conclusions:

1) A comparative framework can be used to benchmark the performance of distributional replication and investigate the relative cost of different distributional parameters.
2) 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. A reserve portfolio with a high Sharpe ratio will tend to produce a replica with a high Sharpe ratio, and a reserve portfolio with a low Sharpe ratio will tend to produce a replica with a low Sharpe ratio.
3) 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 both individual market factors and relative value trades representing long/short combinations of 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.
4) 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.

Appendix A: Time Series Behavior of a "Static" Distributional Replica

The results of this section demonstrate that although the Kat and Palaro process is complex, the behavior of the process can be understood through the use of simple models. We recreate the characteristics of the time series behavior of the replicas referred to as Fund 1 and 3 in [9]. The process used in this Appendix is different from that applied in other parts of this document. In this section we utilize a static position in the reserve and target portfolios throughout the entire 9-year replication time period. The performance of these static positions is dependent on the performance of the reserve and target portfolios. While this static model will not match the higher order moments of skewness and kurtosis, the volatility and correlation targets are matched. The draw-downs and periods of strong performance correspond well with those of the replicas published in [9]. These examples further demonstrate that the fundamental behavior of the replication process is driven by the performance of the reserve and target portfolios.

The evolution of the mean excess return of Fund 1 is provided in *Figure 5 from reference [9]*. For the first month of the replication process applied within [9], the payoff profile is provided in *Figure 1 from reference [9]*. This payoff profile provides an intuitive guide to the positions that will be run within the replica for the first month. Using the slope of the payoff profile from *Figure 1 of reference [9]*, the trading process begins with a 2.5 leveraged (long) position in the reserve portfolio and a -1.5 leveraged (short) position in the target portfolio. For the moment, consider an investor requires no distributional adjustments, and as well as we assume no change in the correlation relationship of the target and reserve throughout the remainder of the backtesting period. This would imply that the replica's positions in the reserve and target portfolio do not change throughout the entire 9-year period.

The resulting evolution of the mean excess return is provided in Figure 15 on the following page. The chart is very similar to the results contained in *Figure 5 of reference [9]* for the replica referred to as Fund 1. The small bump in performance around January 1998 is matched, as well as the drop in performance at the beginning of 1999 and the subsequent rise in performance until mid-2000 is clearly matched. The correlation of this static portfolio is 1% to the 50/50 S&P500 and Long Bond target portfolio. The desired target is 0%. Within the Kat and Palaro process there is a considerable amount of dynamic trading to perform the distributional adjustments as well as monthly updates to re-calibrate the entire process. Despite this complexity, the behavior over a long time period is modeled reasonably well using the static model.

Consider another example, where a static model is constructed for the replica from reference [9] referred to as Fund 3 using a 2.5 leveraged (long) position in the reserve portfolio and a -2.2 leveraged (short) position in the target portfolio. A larger short position as compared to Fund 1 is necessary given the negative correlation specification, and is evident from the slope of the payoff function (*compare Figure 9 to Figure 1 within reference [9]*). Again, the positions are held statically over the 9-year time horizon. The resulting evolution of the mean excess return is provided at the bottom of the following page. The behavior illustrated in the chart is very similar to that of *Figure 11 of reference [9]* for the replica referred to as Fund 3. The correlation of the static portfolio is -45% to the 50/50 S&P500 and Long Bond target portfolio, close to the desired target of -50%.

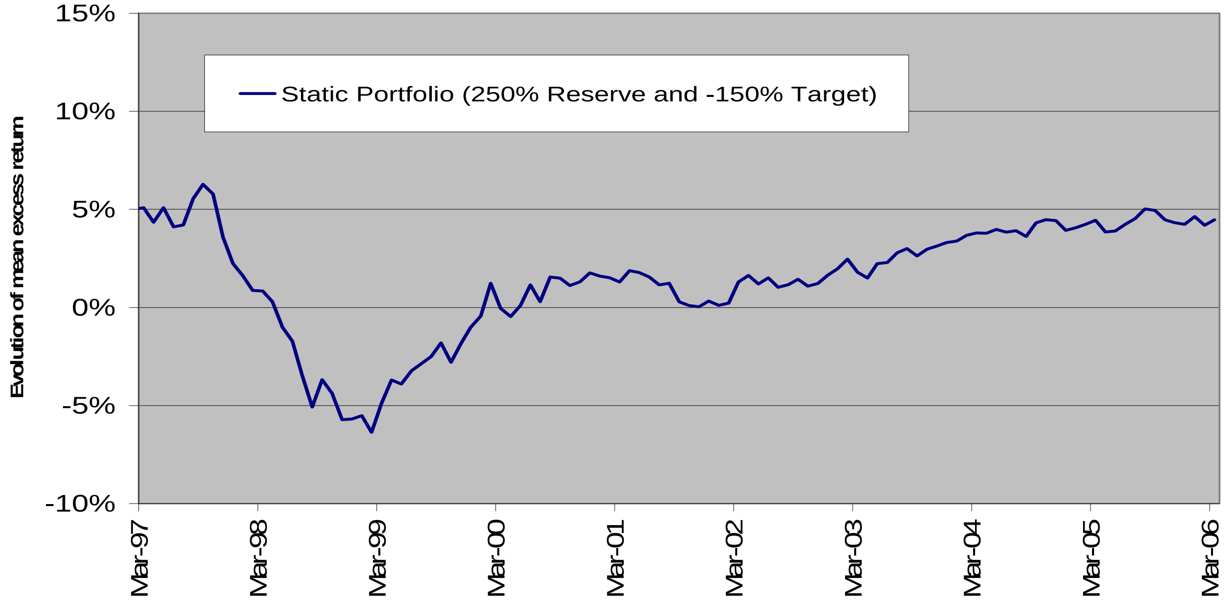

Figure 15: Evolution of Mean Excess Return: this static portfolio has a 1% correlation to the target portfolio (compare this time series to *Figure 5 from reference [9]*).

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

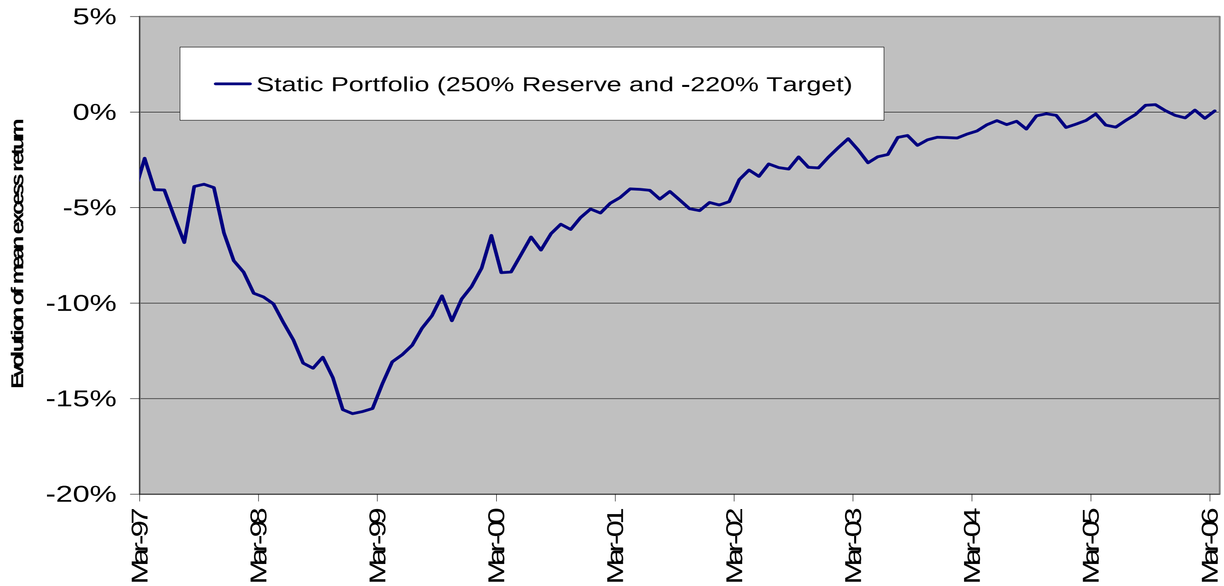

Figure 16: Evolution of Mean Excess Return: this static portfolio has a -45% correlation to the target portfolio (compare this time series to *Figure 11 from reference [9]*).

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

References:

[1] J. Hasanhodzic and Andrew W. Lo, "Can Hedge-Fund Returns Be Replicated?: The Linear Case" (August 16, 2006). Available at SSRN: http://ssrn.com/abstract=924565

[2] 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.

[3] Harry M. Kat and Helder 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.

[4] Northwater Capital Management, "Hedge Fund Replication Part 1: Linear Factor Models", March 2007.

[5] Northwater Capital Management, "Hedge Fund Replication Part 2: Distributional Replication", March 2007.

[6] Northwater Capital Management, "Hedge Fund Replication Part 3: Distributional Replication with Correlation Targeting", March 2007.

[7] Steven Umlauf and Ben Bowler, "Merrill Lynch Factor Index: A New Approach to Diversified Hedge Fund investing", presented at the Hedge Fund Replication and Alternative Beta Conference, London UK, February 2007.

[8] Harry M. Kat and Helder P. Palaro, "Who Needs Hedge Funds? A Copula-Based Approach to Hedge Fund Return Replication", Working Paper #0027, Alternative Investment Research Centre Working Paper Series, Cass Business School, City University, London, UK, November 23, 2005.

[9] Harry M. Kat and Helder P. Palaro, "Tell me what you want, what you really, really want! An Exercise in Tailor-Made Synthetic Fund Creation", Working Paper #0036, Alternative Investment Research Centre Working Paper Series, Cass Business School, City University, London, UK, October 9, 2006.

[10] The Edhec study on Distributional Adjustments is referenced at: http://www.allaboutalpha.com/

[11] B. Dumas and B. Allaz, Financial Securities: market equilibrium and pricing methods, Chapman and Hall, 1996.

Disclaimers

The opinions, estimates and projections contained in this research report are those of Northwater Capital Management Inc. as of May 15, 2007, unless indicated otherwise, and are subject to change without notice. **Past performance is not indicative of future results, which may vary.**

Northwater makes effort to ensure that the contents herein have been compiled or derived from sources believed to be reliable and that any assumptions used in its modeling process are appropriate. However, Northwater makes no representation or warranty, express or implied, in respect thereof, and accepts no liability arising from its use. All formulae, models and backtested data contained in this research report are for illustrative purposes only, and Northwater makes no guarantee of any kind of the investment returns or any other outcomes that may result from the application of such formulae or models. Unlike an actual performance record, models have certain inherent limitations, as modeled or hypothetical performance results involve certain assumptions that may or may not be correct. In addition, modeled or hypothetical performance results do not represent the results of actual trading and may over- or under-compensate for the impact, if any, of certain market factors.

The investment strategies discussed herein are speculative and involve substantial investment, liquidity, derivative, and other risks. This brief statement does not disclose all of the risks and other significant aspects of investing in products such as those described herein and you are advised to speak to your investment advisors before making any such investment. Northwater expressly disclaims any liability thereto.

This research report is not and should not be construed as an offer, recommendation or solicitation to buy or sell, nor is it a confirmation of any terms of, any interests or shares in any of the investment funds managed by Northwater.