Fund of Hedge Fund vs Multi-Strategy Providers:

Implications for cost–effectiveness and portfolio risk

By Igor Lomtev PhD Quantitative Research Analyst, State Street Global Advisors, **Chris Woods** Senior Managing Director, State Street Global Advisors and **Vladimir Zdorovtsov** PhD Quantitative Research Analyst, State Street Global Advisors

ABSTRACT

This study investigates the relative merits of fund of hedge funds and multi-strategy approaches and their implications for cost-effectiveness and portfolio optimality. We present evidence documenting unambiguous fee savings generated by a multi-strategy solution. Furthermore, we show that access to more detailed performance information and less constrained liquidity allow a multi-strategy provider use of a richer suite of risk management tools. The additional performance information and closer access to the underlying strategy managers afforded in a multi-strategy context also allows for better identification of the risk factors driving the underlying strategy returns. This in turn can assist in the production of more robust covariance matrices and offers an opportunity to optimize the strategy allocation by trading factor exposures rather than disrupting the positions of the underlying funds. Taken together, these cost, information and construction advantages can potentially go a long way to offsetting the greater range of managers accessible within a fund of fund structure.

INTRODUCTION

In a world of pressing pension fund deficits, institutional investors have started to look in earnest at hedge funds as a means of augmenting the return delivered by traditional long-only strategies. Although not showing quite the same stellar returns as they once did during the equity bull market, hedge funds still promise equity-like returns with bond-like risks whilst at the same time possessing low correlations with the major equity and bond markets.

Once an institutional investor has decided to commit to hedge funds, the question becomes how best to get hedge fund exposure. An institution with a large and sophisticated investment staff may have the confidence to evaluate strategies, conduct due diligence on managers, and optimise the allocation to funds using its in-house resources. But for investors without the luxury of a large investment staff, the task is daunting. For such investors, a more pragmatic point of entry is through funds of hedge funds. A fund of funds effectively shares its own expertise to carry out the evaluation of strategies and managers, and to allocate investment capital efficiently among them. Funds of funds have become increasingly popular, and now control approximately one-half of hedge fund assets and two-thirds of current flow. Funds of funds were a particularly attractive proposition when hedge funds regularly delivered double-digit returns. However, in an era when hedge fund returns are harder to come by, investors find that the fee take from the fund of funds represents an increasing proportion of their net returns. As a result, investors have started to consider multi-strategy funds as an alternative to investing in funds of hedge funds.

Multi-strategy funds consist of several hedge funds all managed by the same provider. The provider may be a large institutional money manager that has developed a hedge fund capability, or a successful boutique hedge fund that over time has gathered sufficient assets to enable diversification into additional strategies.

Multi-strategy funds have a number of pros and cons compared with funds of hedge funds. The obvious disadvantage is that no multi-strategy provider can claim to have the best manager within each hedge fund category. A fund of funds, by contrast, may interview hundreds of managers and use this experience to identify strategies with a discernible edge and talented managers able to exploit the opportunity. Set against this is the fact that the multi-strategy manager is not handicapped by the additional layer of fees. Furthermore, while a fund of funds charges performance fees by strategy, a multi-strategy fund determines them as a function of aggregate portfolio returns. Thus, on a net of fees basis the latter may prove very competitive.

While the fee differential is a clear and well-known benefit of multi-strategy solutions, a less discussed but no less significant advantage is the greater transparency and flexibility the approach affords for risk management and portfolio construction.

This study investigates the relative merits of fund of hedge funds and multi-strategy approaches in terms of their cost-effectiveness and portfolio optimality. We present simulation evidence documenting unambiguous fee savings generated by a multi-strategy solution. Furthermore, we show that access to more detailed performance information and less constrained liquidity allow a multi-strategy provider use of a richer suite of risk management tools.

The rest of the paper is organised as follows: Section 2 analyses the differential in fees as a function of performance; Section 3 presents an overview of the relevant literature, methodology, and the empirical results on risk management implications of different levels of transparency; Section 4 discusses factor risk models; Section 5 examines implementation of allocation changes; Section 6 concludes.

2. PORTFOLIO PERFORMANCE AND THE LEVEL OF FEES

The savings offered by the multi-strategy funds can be decomposed into two elements: (1) while a fund of funds is charged performance fees by each constituent strategy, a multi-strategy provider nets out excess performance to determine portfolio-level incentive fees; (2) a fund of funds provider carries an additional layer of management and (optionally) incentive fees for their due diligence in selecting hedge fund managers as well as for dynamic portfolio re-allocations.

The effects of the first component can be succinctly demonstrated by the following proof. Let X_1, X_2 be random variables distributed bivariate normal. Assume also that their means and standard deviations are zero and one, respectively, and that the correlation between them is ρ. Let $F_1 = \max(X_1+X_2,0)$ and $F_2 = \max(X_1,0) + \max(X_2,0)$. Random variable $Y = X_1+X_2$ will have a normal distribution with $\mu=0$, $\sigma^2=2+2\cdot\rho$. Consider the expectation of F_1:

$$E[F_1] = E[\max(Y,0)] = \int_0^\infty \frac{1}{\sqrt{2\cdot\pi\cdot\sigma}} \cdot \exp(-\frac{x^2}{2\cdot\sigma^2})\cdot x\cdot dx =$$

$$\frac{1}{\sqrt{2\cdot\pi\cdot\sigma}}\cdot\int_0^\infty \exp(-\frac{x^2}{2\cdot\sigma^2})\cdot d(\frac{x^2}{2}) =$$

$$= \frac{\sigma}{\sqrt{2\cdot\pi}} = \frac{\sqrt{1+\rho}}{\sqrt{\pi}}$$

Similarly, the expectation of F_2:

$$E[F_2] = E[\max(X_1,0)] + E[\max(X_2,0)] = \frac{2}{\sqrt{2\cdot\pi}} = \sqrt{\frac{2}{\pi}}$$

Hence, $E[F_1] \leq E[F_2]$ and $E[F_1]=E[F_2]$ if and only if $\rho=1$. Since the expected values of performance fees are proportional to the expectations of F_1 and F_2 above, we can see that the multi-strategy approach will dominate as long as the strategies comprising the portfolio are not perfectly correlated.

To confirm this intuitive result empirically, we carry out the following Monte Carlo exercise. Using a five-year monthly history of ten absolute return series, we generate 50,000 sets of annualised returns for two scenarios: a multi-strategy manager with portfolio-level fees and a fund of funds manager with constituent-fund level fees. We presume 1 per cent management and 20 per cent performance fees in both cases. Table 1 and Figure 1 below summarise the results.

Clearly, a multi-strategy solution outperforms on a net-of-fees basis. Average and median annual savings are 23 and 19 basis points, respectively. The savings are bounded from below by zero and are a decreasing function of portfolio performance. Intuitively, if each strategy in the portfolio yields non-negative returns after management fees, the two simulated cases will produce identical results in this simple set-up. As performance deteriorates and insofar as the constituent returns are not perfectly correlated, the fund of funds approach begins to underperform on a net-of-fees basis. While the average savings across all 50,000 simulated years may not be extreme, they tend to rapidly grow in magnitude when it matters most, reaching as high at 1.54 per cent. This is certainly a non-trivial amount, particularly given the portfolio performance level at which these savings are attained Note, for example, that the fund of funds is still charging performance fees even when the portfolio return is negative. Furthermore, this presumes the fund of funds manager does not charge an additional layer of fees. Incorporating this extra level would significantly amplify the difference.

TABLE 1. SUMMARY STATISTICS: MULTI-STRATEGY VS FOF

		Return After Fees		
	Return Before Fees	Multi-Strategy	Fund of Funds	Fee Difference
Mean	8.43	5.94	5.71	0.23
Median	8.33	5.87	5.68	0.19
StdDev	3.07	2.46	2.59	0.21
Min	-2.73	-3.73	-4.25	0
Max	22.66	17.33	17.22	1.54

3. OPTIMAL CAPITAL ALLOCATION

Another advantage that a multi-strategy manager has over a fund of funds is the additional transparency afforded by having individual strategy managers in-house. This proximity to strategy managers should allow the construction of a better optimised portfolio, increased flexibility in allocating between strategies, and a greater understanding of – and potentially an ability to hedge – the risks embodied in the portfolio. We develop these themes in the following sections.

A. LITERATURE OVERVIEW

A fund of funds manager wishing to allocate capital to managers may feel constrained by the lack of available information. Many hedge fund managers, particularly those near capacity and with enviable track records, may prove unwilling to disclose any position information. Furthermore, the return information may be limited to the data provided in monthly fund updates.

In light of this, fund of funds managers may choose to allocate capital equally across the selected funds and instead devote energies to devising the appropriate allocation to the various categories of strategies. Some fund of funds managers may choose not even to do this, taking the view that the identification of good managers is both more objective and more important than subjectively trying to assess the likely performance of strategy categories in the anticipated market environment.

How optimal is a naïve equal weighting of managers? That depends on the expected returns and volatilities of the managers' funds. If all funds have the same expected returns and risks, then equal weighting makes sense. But if they have the same expected return with different expected volatilities, then the fund of funds manager should allocate a lower weight to those strategies with higher volatilities in order to optimise the Sharpe ratio (expected excess return divided by volatility) of the portfolio.

Theoretically, we can show that the appropriate weights in this case are proportional to the inverse of the individual funds' expected variances. This raises the question as to how we will forecast these variances. One solution is simply to use the variance of the funds' historical returns. However, if the only returns we have are monthly and the track record is short, this may be insufficient. We examine alternative forecast procedures later in this paper.

One can continue to add degrees of complexity to the allocation procedure. So far we have ignored the fact that the returns of individual funds may be correlated with each other. But if we select several managers implementing similar strategies, this is assumption is likely invalid. To overcome this, the aggregate weight to these managers should be reduced. A quantitative mechanism for achieving this is to submit the forecasts for expected return, volatilities, and correlations to a classical mean-variance optimisation procedure.[1]

At this point, purists may object that hedge funds typically do not exhibit normal (Gaussian) return distributions. It is well documented hedge fund returns exhibit undesirable higher moments (Cremers, Kritzman, and Page, 2005; Alexiev 2005). The sources of this non-normality can be traced to the use of dynamic option-like strategies, derivative instruments, and non-linear fee structures embedded in reported returns.[2] To address this, a number of alternative measures of risk have been proposed in the literature, such as Value at Risk (VaR), Conditional Value at Risk (CVaR), Conditional Drawdown at Risk (CDaR), Tail Risk etc.[3]

FIGURE 1. FEE DIFFERENCE AS A FUNCTION OF RETURN

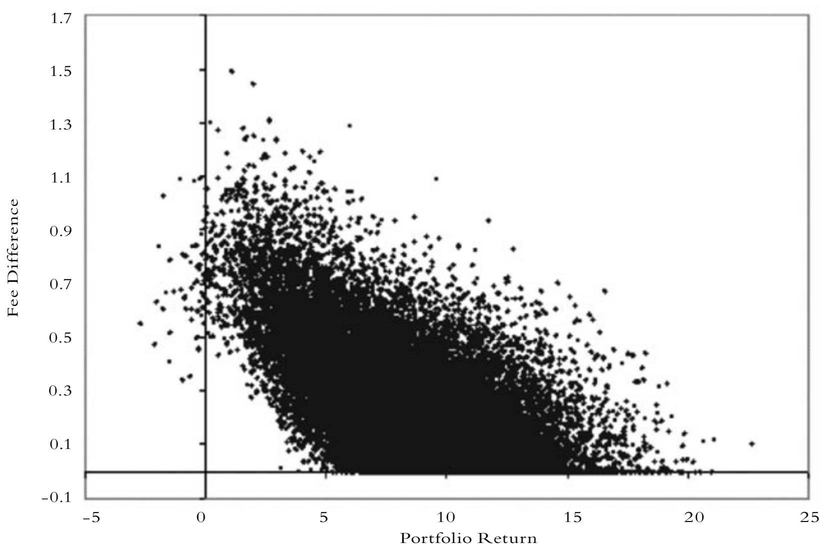

1. The subject of forecasting strategy returns is beyond the scope of this analysis.

2. E.g., see Goetzmann, Ingersoll, Spiegel and Welch (2002); Spurgin (2001); Mitchell and Pulvino (2001); Goetzmann, Ingersoll and Ross (2003); Taleb (2004); and Chan, Getmansky, Haas and Lo (2005).

3. Several recent studies in this strand of research are Estrada (2001); Jorion (2000a); Jorion (2000b); Gupta and Liang (2005); Bali and Gokcan (2004); Agarwal and Naik (2004); Alexander and Baptista (2004); and Chekhlov, Uryasev, Zabarankin (2003).

While the pitfalls of using conventional mean-variance to construct portfolios of hedge funds are well established, survey evidence suggests that the technique is still widely used in practice (e.g., see Amenc, Giraud, Martellini, and Vaissie, 2004). The reluctance to adopt alternative portfolio construction techniques likely stems from the need to estimate a number of additional parameters. Given the growing recent evidence on the ruinous effects of estimation noise even in the context of simple mean-variance optimisations, hesitation to rely on progressively more noisy estimates of higher moments is understandable. Furthermore, while a multi-strategy fund may have access to the large samples of in-depth daily performance records necessary to estimate the requisite parameters, a fund of funds manager is forced generally to rely on low frequency (monthly at best) histories plagued with various reporting biases.

To help establish the impact of data granularity on portfolio construction, we compare the efficacy of several alternative approaches as a function of the length of parameter estimation horizon and the frequency with which observations are sampled. Specifically, we compare (1) conventional mean-variance (CMV); (2) diagonal sample covariance matrix; (3) Random Matrix Theory filtering (RMT); (4) Ledoit shrinkage estimators; (5) CVaR; and (6) CDaR.

The first two are plain-vanilla mean-variance optimisations based on the sample covariance matrix and the diagonal thereof, respectively. The RMT technique draws upon the insights from theoretical physics to noise-undress the sample covariance matrix.[4] Specifically, by comparing the eigenvalues of each sample matrix to their known distributions for random matrices of the same size, we can retain only the significant eigenvalues and repopulate the matrix to be used in portfolio optimisation.

Ledoit (1994) develops an alternative covariance matrix estimator designed to reduce the effect of estimation uncertainty.[5] Intuitively, the technique creates a weighted average of the sample covariance matrix and a "shrinkage target".

$$\Sigma = w \cdot F + (1-w) \cdot S \qquad (1)$$

Where:

F - shrinkage target

S - sample covariance matrix

w - weight, or "shrinkage intensity"

We consider three different shrinkage targets: a one-factor covariance matrix (Ledoit1); a diagonal covariance matrix (Ledoit2); and a two-parameter covariance matrix (Ledoit3).

The one-factor covariance matrix estimator assumes that strategy returns are generated by a one-factor (market) model.

$$x_{it} = \alpha_i + \beta_i \cdot x_{0t} + \varepsilon_{it} \qquad (2)$$

Where x_{0t} is the market return and ε_{it} is the residual. Then,

$$F = \sigma_{00}^2 \cdot \beta \cdot \beta^T + \Delta$$

Where σ_{00}^2 is the variance of market return, β is the vector of slopes, Δ is the diagonal matrix containing residual variances δ_{ii}. The diagonal elements of F are the same as diagonal elements of S.

In the second approach, the shrinkage target is the diagonal of the covariance matrix, while in the third, it is a two-parameter covariance matrix wherein all variances are equal to the first parameter and all covariances to the second.

Lastly, CVaR and CDaR optimisations rely on sample conditional value at risk and conditional drawdown at

FIGURE 2. SIMULATED SURROGATE RETURN SERIES EXAMPLE

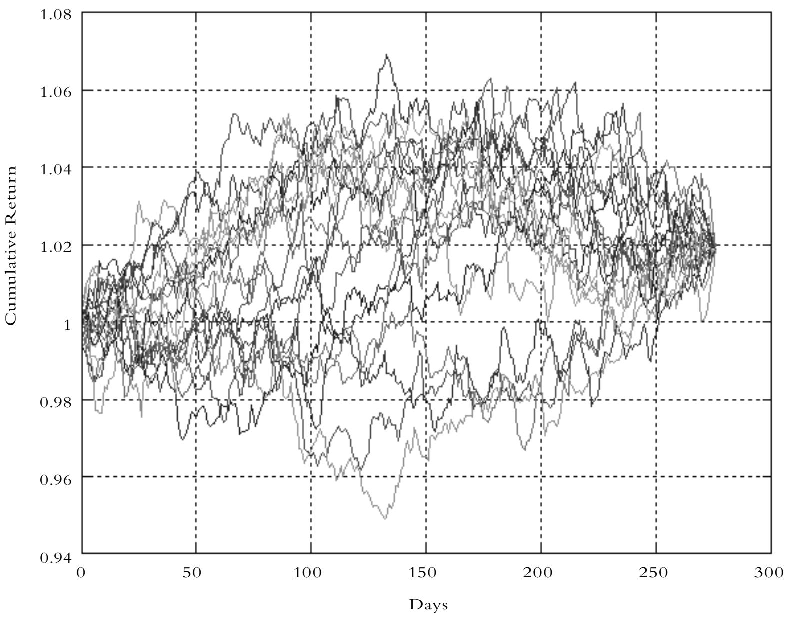

4. E.g., see Plerou et al (2002).

5. See also Ledoit and Wolf (2003)

FIGURE 3. PORTFOLIO PERFORMANCE VS ESTIMATION WINDOW SIZE: MONTHLY RETURNS

PANEL 1

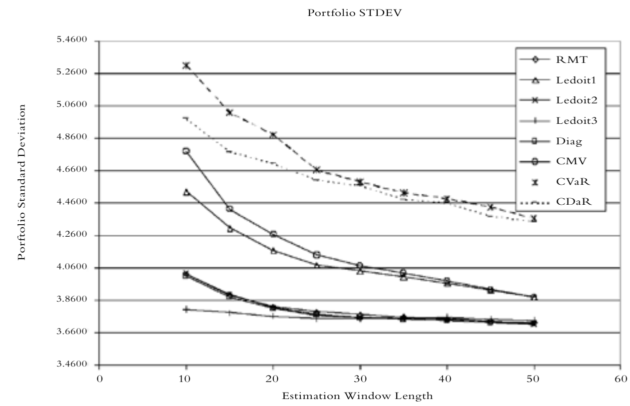

PANEL 2

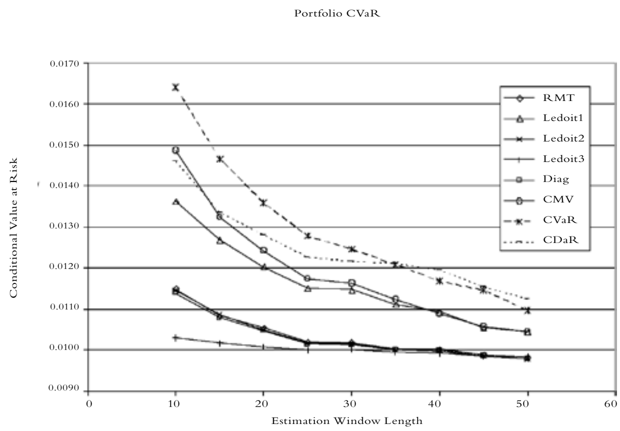

PANEL 3

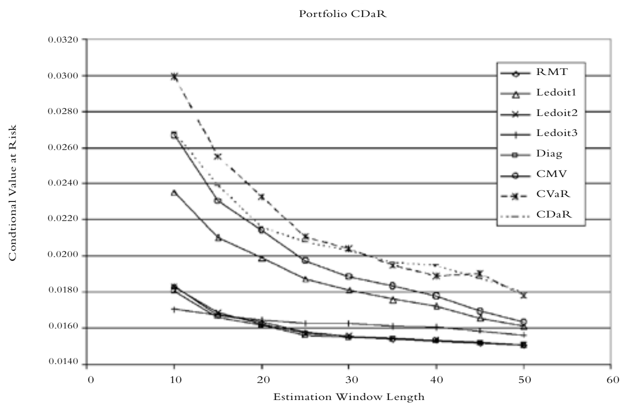

PANEL 4

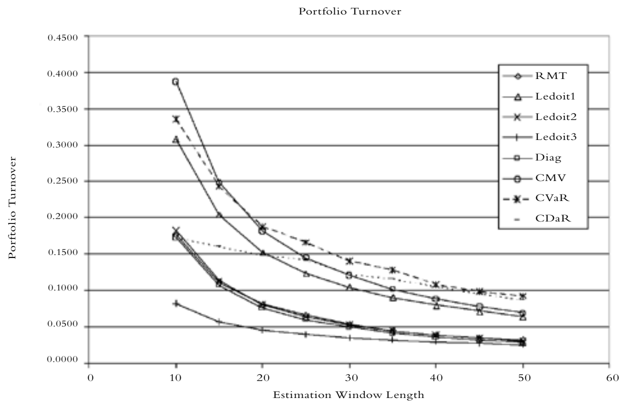

risk, respectively. We start with the usual VaR measure ($VaR(\alpha)$): an α – quantile of the loss distribution, i.e., the loss that is not exceeded in $\alpha \cdot 100\%$ cases. We define conditional VaR($CVaR(\alpha)$) as the conditional expectation of losses exceeding the $VaR(\alpha)$ level: $CVaR(\alpha)= E[L|L>Var(\alpha)]$. CVaR defined this way is a convex function of portfolio positions, allowing use of standard optimisation methods.

CDaR is an extension of CVaR. By definition, a portfolio's drawdown is the drop in portfolio value as compared with the maximum level attained earlier in the sample path. The drawdown accounts not only for the amount of losses, but also for their sequence. In this sense it is a loss measure "with memory". We define CDaR as the following conditional expectation:

$CDaR(\alpha)=E[DD|DD>DD(\alpha)]$, where $DD(\alpha)$ is the drawdown that is not exceeded in $\alpha \cdot 100\%$ cases.

Our monthly return sample is comprised of 7 absolute return strategies from April 1999 to July 2006. The daily returns for the same strategies cover the period from August 1, 2005, to August 31, 2006. Using these initial samples, we employ the surrogate data algorithm developed in Schreiber and Schmitz (2000) to generate multivariate time series with properties similar to the original samples. Specifically, the process aims to preserve the means, standard deviations, autocorrelations, and higher moments of each time series.[6]

Figure 2 presents sample surrogate paths for one of the constituent strategies.

A hundred pseudo samples are generated for each dataset. We then take a specified length of estimation window, optimise the portfolio according to the chosen objective function, measure portfolio return during the next time period, roll the estimation block one period forward and repeat the process for the balance of the sample. This is repeated for each of the simulated surrogate sets and the resulting portfolio risk metric is averaged across them.

C. RESULTS

Figure 3 presents the results of various allocation procedures based on monthly surrogate returns to the seven absolute return strategies. Optimal minimum risk portfolio allocations are created using risk estimates calculated from different sizes of rolling estimation window. Panels 1, 2, 3, and 4 show the realised portfolio standard deviation, CVaR, CDaR, and turnover levels, respectively, as a function of the estimation horizon length and the risk metric utilised in the optimisation objective function. The dynamic rebalancing process described in the methodology section is performed for each simulated surrogate sample to arrive at the time series of portfolio returns and the respective metric of interest. The plotted points represent averages over the 100 simulated surrogate samples.

The results point to the dangers of pushing the degree of sophistication in determining the optimum allocation to managers – particularly if the only data available is a history of monthly returns. As the historical sample window over which a given risk measure is computed grows, the performance of different techniques in terms of the realised risk tends to converge. Intuitively, regardless of the specific objective function employed, as the estimation sample becomes progressively more reliable and yields more robust estimates of risk, minimising either risk measure indirectly helps minimise the others.

For example, if one were primarily concerned with a portfolio's CDaR or CVaR profile, optimising the allocation via any of the chosen methods will produce generally similar results once the estimation sample contains close to five years of monthly observations. Clearly, some techniques are better at capturing a given risk dimension than others. For instance, as expected, if reducing annualised portfolio standard deviation is the main goal, the objective functions that rely on the entire return distribution fare better.

Not surprisingly, portfolio turnover is decreasing in estimation accuracy: as the estimation window expands and/or noise reduction techniques are employed, the average monthly change in the portfolio weights declines.

Most importantly, there appears to be little value in resorting to sophisticated portfolio construction methods when the available return data is of monthly frequency. While indeed Ledoit shrinkage and RMT noise-reduction dominate conventional mean-variance, a simple diagonal covariance matrix optimisation does nearly as well or better, regardless of the risk dimension one investigates.

Figure 4 summarises similar analysis performed on the daily return series.[7] While the general conclusions are similar and a simple diagonal covariance approach works well for portfolio standard deviation and CVaR, the results are notably different for CDaR. Clearly, unlike the monthly observations, the much richer dynamics of daily returns allow CDaR optimisation to select a more optimal portfolio with respect to expected future drawdown exposure. Once the estimation horizon exceeds approximately three months of daily returns, CDaR-optimised portfolios exhibit much lower realised CDaR's than those optimised with respect to other risk measures.

Figures 3 and 4 show that the allocations obtained from using a diagonal covariance matrix (equivalent to weighting in proportion to the inverse of the funds' variances) performs surprisingly well when compared against optimisations using full covariance matrices as well as covariance matrices transformed to eliminate estimation noise. The simple diagonal matrix approach also performs well compared with optimisations that use downside risk measures. Given these results, we can see that it is entirely sensible for fund of funds

7. The reported portfolio standard deviations in Figures 3 and 4 are annualized; CVaR's are over 1 month and 1 day, respectively; CDaR's are measured over the corresponding monthly or daily paths; turnover is the average monthly or daily change in strategy weights.

6. Note that the version of the algorithm we use does not preserve cross-correlations. An alternative specification aims to reproduce the means, standard deviations, autocorrelation and cross-correlations, although it does not preserve skewness and kurtosis. In unreported results, we show that the conclusions are not sensitive to the choice of the algorithm.

FIGURE 4. PORTFOLIO PERFORMANCE VS ESTIMATION WINDOW SIZE: DAILY RETURNS

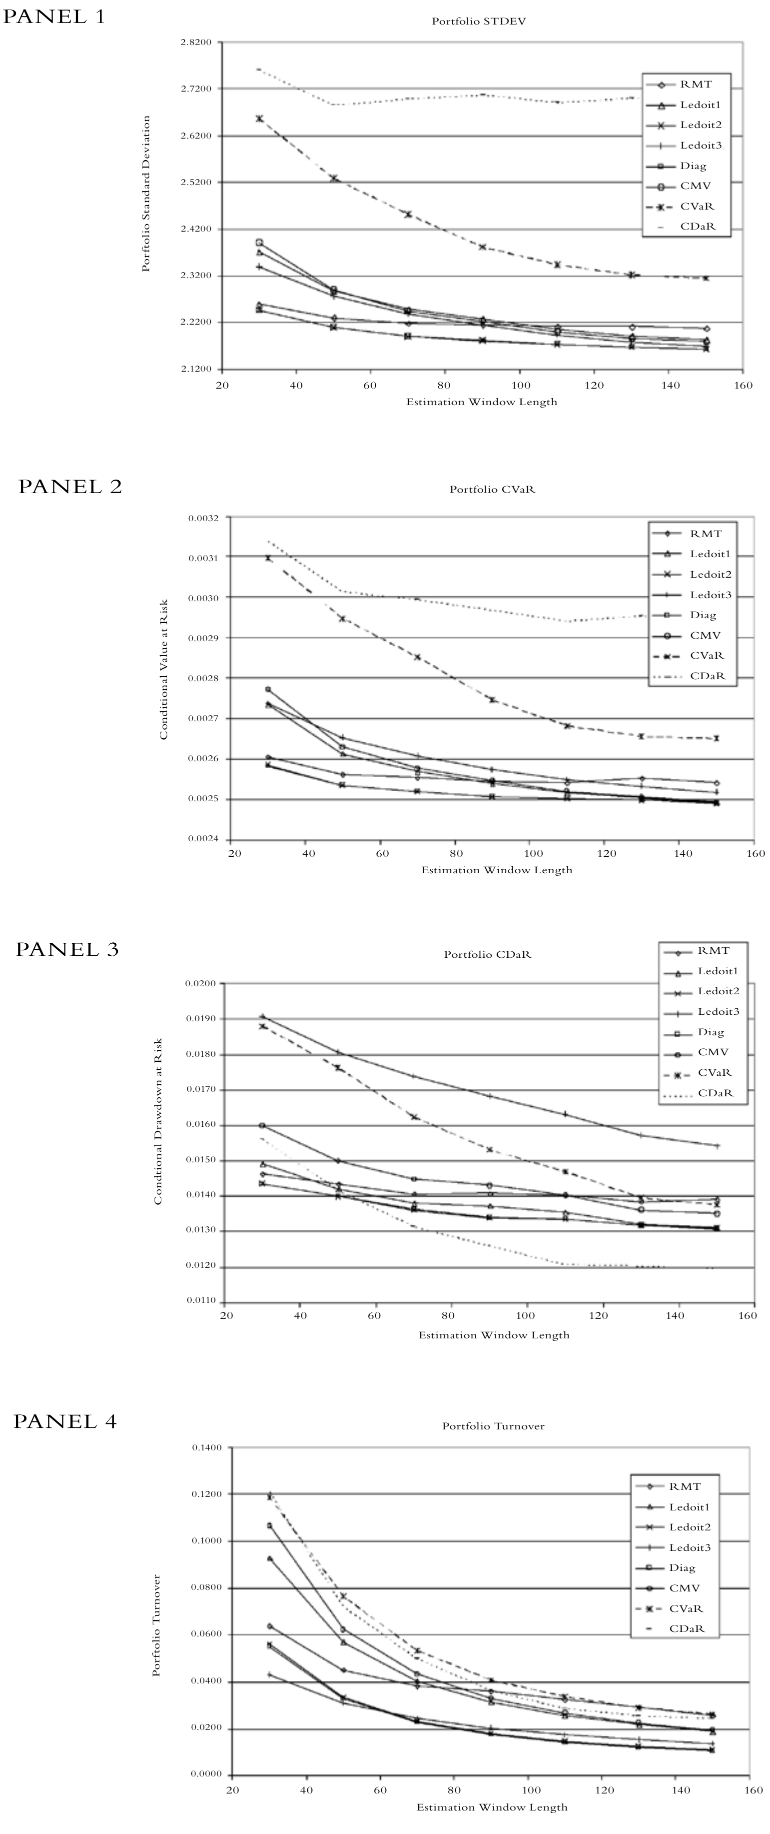

managers to take the pragmatic approach of either equally weighting funds, or weighting them in accordance with the inverse of their variances. However, if daily returns are available, the additional information makes the effort of extra sophistication worthwhile.

4. FACTOR MODELS

The previous section has shown that access to daily return information allows a more sophisticated approach to optimising the allocation of capital to various hedge fund strategies and that, in the absence of daily returns, a simple approach using a diagonal covariance matrix can perform better than more complex alternatives. In the world of equity risk models, firms such as Barra have demonstrated the benefit of employing factor risk models over sample matrices calculated from stock returns and it is natural to ask whether similar models might be usefully applied to the hedge fund allocation problem. A number of recent academic papers have sought to show that common hedge fund risk factors do exist, and that to a non-negligible extent, some of these risk factors are rewarded and may even contribute a large fraction of hedge fund returns.[8]

There are, however, some important differences when it comes to comparing common factors for hedge funds with those identified for equities. Individual companies are relatively stable and their risk factor exposures change only slowly through time. In contrast, hedge funds by their nature seek to adapt to their environment and exposure to their factors can be expected to vary more rapidly. For example, convertible arbitrage managers may, depending on the market conditions prevailing at the time, seek to capture option mispricings or try to exploit changes in credit spreads and yields. Another observation that we might usefully attempt to capture is that hedge fund return distributions can change with the market environment. Some environments are conducive to good returns with low volatility and others to poor returns with high volatility. Identification of such regimes can help in forecasting the expected returns and risk of these strategies.

These risk factors can include well known generic hedge fund risk premia ("alternative betas" in the jargon) such as the small cap minus large cap premium, the value minus growth premium, the credit premium and so on. They can also include factors specific to individual strategies such as the prices of crude oil and other commodities. The relevant factors can be allowed to vary over time and their volatility regimes identified. Estimates of the expected returns and variances of each factor in the current regime can be used to construct estimates of the expected returns and variances of individual strategies. Because several strategies may be exposed to the same factor, the strategy covariance matrix built in this fashion will have non-zero correlations between the various strategies.

By removing spurious correlations present in the historical return streams and replacing them with those driven by exposure to common factors, the return and risk forecasts for strategies obtained in this way are likely to be superior to those constructed solely from the returns of the strategies themselves. In determining the appropriate factors to use in the analysis, the multi-strategy manager is likely to have an information advantage over his funds-of-funds counterpart. Not only will the multi-strategy manager have the benefit of more frequent return information, he will also have the ability to quiz the underlying strategy managers on the factors they believe are relevant to their returns. This additional insight will help limit the possibility of wrongly identifying factors through simplistic regression techniques. However, we should not ignore the lessons of the simulation study above that showed how shrinking correlation estimates towards zero can improve the robustness of the optimisation procedure. Accordingly, one can combine the factor derived covariance matrix with a diagonal covariance matrix in a Bayesian fashion as advocated by Ledoit (1994).

5. IMPLEMENTATION

So far we have seen that a multi-strategy manager with access to more frequent return information can improve on the naïve equal-weighted (or reciprocal-variance weighted) allocation that a fund of funds manager would do well to rely on. However, a better allocation is useless if it cannot be acted upon for reasons of fund liquidity or because of the associated transaction costs.

A multi-strategy manager may well be in a better position to change allocations than his fund of funds counterpart due to greater access to the underlying strategies. For example, a fund of funds manager may be subject to lock-ups and gates imposed by external managers. In contrast, a multi-strategy manager may be able to take advantage of countervailing flows to cross allocations, or the presence of sufficient cash in a strategy due to have a reduced allocation to make changes.

However, the presence of such opportunities cannot always be relied on and there will be times when the multi-strategy manager cannot change allocation to a strategy because such a change would unduly disrupt the underlying strategy's ability to keep its trades in place. It is here that the factor analysis used to estimate strategy risks and returns can come in useful. If an allocation change is in part being prompted by views on the return of a specific factor, it may be possible to trade that factor directly outside of the constituent strategies rather than to rebalance the strategies themselves to obtain the desired factor exposure.

For example, suppose the current mix of strategies is identified as being susceptible to a rise in the price of oil. The multi-strategy manager has a choice of either reducing the allocation to the underlying strategy most susceptible to the price rise, or alternatively the multi-manager can hedge this exposure by buying oil futures. A similar strategy could be used if the multi-strategy fund suffers in the event of a fall in equity market volatility; in this case the multi-strategy manager can sell volatility either by trading VIX futures or variance swaps.

In some cases, the offending risk factor may not be directly traded. For instance, the optimal allocation may overweight small cap stocks and underweight large cap stocks. Such a position may be desirable in the long run but if for any reason the multi-strategy manager wishes to cut exposure in the short

8. E.g. see Fung and Hsieh (2002), Fung, Hsieh, Naik and Ramadorai (2006).

term, his options may be limited. A number of managers are looking to provide passive exposure to such risk factors through the creation of so-called synthetic hedge funds. These will provide an alternative hedging mechanism going forward.

6. CONCLUSION

As more institutional investors seek to use hedge funds to boost returns whilst reducing risk, the number of ways to obtain hedge fund exposure continues to rise. The traditional routes of DIY manager selection and entry via a fund of funds are increasingly being supplemented by new methods which include multi-strategy funds. Although the latter are unable to claim that they invest in the best managers in each strategy, this deficiency may be compensated by lower fees and, by virtue of access to better and more frequent information, a greater ability to optimise the allocation to underlying strategies. This ability will be further enhanced with the development of synthetic hedge funds which will allow the multi-strategy manager to optimise exposure to hedge fund risk factors without causing unnecessary disruption to the underlying managers.

REFERENCES

Agarwal, V, and Naik, N 2004, Risks and Portfolio Decisions Involving Hedge Funds, *Review of Financial Studies* 17, 63-98.

Alexander, G, and Baptista, A 2004, A Comparison of VaR and CVaR Constraints on Portfolio Selection with the Mean-Variance Model, *Management Science*, vol. 50, no. 9, pp.1261-1273.

Alexiev, J 2005, The Impact of Higher Moments on Hedge Fund Risk Exposure, *Journal of Alternative Investments*, vol. 7, no 4, pp. 50-65.

Amenc, N, Giraud, J, Martellini, L and Vaissie, M 2004, Taking a Close Look at the European Fund of Hedge Funds Industry, *Journal of Alternative Investment*, Winter 2004, pp. 59 – 69.

Bali, T G, and Gokcan, S 2004, *Alternative Approaches to Estimating VaR for Hedge Fund Portfolios*, Edited by Barry Schachter, Intelligent Hedge Fund Investing, pp. 253-277, Risk Books, Publisher: Incisive Media PLC.

Chan, N, Getmansky, M, Haas, S and Lo, A, 2005, *Systemic Risk and Hedge Funds*, Working Paper, NBER Conference on the Risks of Financial Institutions.

Chekhlov, A, Uryasev, S, and Zabarankin, M, 2003, *Portfolio Optimization with Drawdown Constraints*. B. Scherer (Ed.) Asset and Liability Management Tools, Risk Books, London, 2003

Chekhlov, A, Uryasev, S, and Zabarankin, M 2005, Drawdown Measure in Portfolio Optimization, *International Journal of Theoretical and Applied Finance*, 8, pp. 13-58

Cremers, Kritzman and Page 2005, Optimal Hedge Fund Allocations: Do Higher Moments Matter?, *Journal of Portfolio Management*, vol 31, i3, pp70(12).

Estrada, J. 2001, *The Cost of Equity of Internet Stocks: A Downside Risk Approach*, Working Paper, IESE Business School, Spain.

Fung, W, Hsieh, D, 2002, Asset-Based Style Factors for Hedge Funds, *Financial Analyst Journal*, 58, pp. 16-27.

Fung, W, Hsieh, D, Naik, N, Ramadorai, T 2006, *Hedge Funds: Performance, Risk and Capital Formation*, Working Paper, Duke University.

Goetzmann, W, Ingersoll, J and Ross, S, 2003, High-Water Marks and Hedge Fund Management Contracts, *Journal of Finance*, vol. 58, no. 4, pp.1685-1717.

Goetzmann, W, Ingersoll, J, Spiegel, M, and Welch, I, 2002, *Sharpening Sharpe Ratios*, Working Paper, No. 02-08, Yale International Center for Finance.

Gupta, Anurag, and Bing Liang 2005, Do Hedge Funds Have Enough Capital? A Value at Risk Approach, *Journal of Financial Economics,* 77, pp. 219-253.

Jorion, P, 2000a, Value at Risk: The New Benchmark for Managing Financial Risk, McGraw-Hill, 2nd ed.

Jorion, P, 2000b, *Risk Management Lessons from Long Term Capital Management*, European Financial Management, 6, pp. 277-300. 33

Ledoit, O 1994, *Portfolio Selection: Improved Covariance Matrix Estimation*, Working Paper, MIT.

Ledoit, O, and Wolf, M 2003, Improved Estimation of the Covariance Matrix of Stock Returns with Application to Portfolio Selection, *The Journal of Empirical Finance*, vol. 10, Issue 5, Dec. 2003, pp.603-621.

Mitchell, M and Pulvino, T 2001, Characteristics of Risk and Return in Risk Arbitrage, *Journal of Finance*, vol. LVI, no. 6, pp.2135-2176.

Plerou, V and Gopikrishnan, P 2002, A Random Matrix Approach to Cross Correlations in Financial Data, *Physical Review E*, 65, 066126

Spurgin, R B 2001, How to Game your Sharpe Ratio, *The Journal of Alternative Investment*, winter 2001, pp. 38-46. 34

Schreiber, T and Schmitz, A 2000, Surrogate Time Series, Physica D, 142, pp. 346-382

Taleb, N N 2004, Blowup versus Bleed: What Does Empirical Psychology Say About the Preference for Negative Skewness?, *Journal of Behavioral Finance*, vol 5, 1, pp. 2-7.

Note

Simulations were conducted to show that a Fund of Hedge Funds is an inferior alternative to Multi-Strategy Funds due to (1) higher expected fees, and (2) less optimal portfolio construction. Fee comparison analysis relied on a Monte Carlo simulation presuming a multivariate normal distribution for portfolio constituents. Comparison of portfolio construction techniques involved the use of surrogate data simulations as per Schreiber, T. and Schmitz, A., 2000, "Surrogate Time Series", Physica D, 142 (2000) 346-382.

All simulations were conducted by State Street Global Advisors. There will be differences between the simulated returns and what would actually happen managing money in our investment strategies based upon each client's risk and return objectives. The criteria used to create the simulated performance is objective and the results are verifiable.

Monte-Carlo "Fee Analysis" simulations assumed a 1 per cent management fee and a 20 per cent performance fee as per the discussion in the paper. Surrogate data series used in portfolio risk analysis are all before fees.The simulated performance results shown do not represent the results of actual trading using client assets, but were achieved by means of the retroactive application of a model that was designed with the benefit of hindsight. The simulated performance was compiled after the end of the period depicted and does not represent the actual investment decisions of the advisor. These results do not reflect the effect of material economic and market factors on decision-making.

The simulated returns are not necessarily indicative of future performance, which could differ substantially.

NOTES