Optimal allocation to hedge funds

Lionel Martellini and Mathieu Vaissié argue that it is only by taking into account the exact nature and composition of their existing portfolio that institutional investors can maximise the benefits they can expect from investing in hedge funds. To this end, they introduce suitably designed measures of the contribution of various hedge fund strategies to the risk in an existing stock/bond portfolio, and use them in the context of optimal selection of hedge fund strategies from an investor's standpoint

One of the by-products of the bull market of the 1990s has been the consolidation of hedge funds as an important segment of the financial markets. It was recently announced that the value of the hedge fund industry globally accounts for nearly $1.35 trillion, with more than 10,000 hedge funds and funds of hedge funds in the world.[1]

There seem to be two main reasons behind the success of hedge funds in institutional money management. On the one hand, hedge funds may provide abnormal risk-adjusted returns, due to the superior skills of hedge fund managers and flexibility in trading strategies. These benefits have been labelled 'return enhancement benefits', and have been the focus of much literature on the presence (or absence) of skills in hedge fund management (see, for example, Amenc, El Bied & Martellini, 2003, as well as references therein). A careful quantitative and qualitative due diligence process is needed to ensure that the selection process will increase the likelihood of including high risk-adjusted performance funds in a portfolio. On the other hand, hedge funds seem to provide diversification benefits with respect to other existing investment possibilities. These have been labelled 'risk reduction benefits'. In a nutshell, the diversification argument states that investors can take advantage of hedge funds' linear and non-linear exposure to a great variety of risk factors, including volatility, credit and liquidity risk, to reduce the risk of their global portfolio (see, for example, Amenc & Martellini, 2002, Georgiev, Karavas & Schneeweis, 2002, Terhaar, Staub & Singer, 2003, and Alexander and Dimitriu, 2004, among many others).

Although the existing literature seems to agree about hedge funds as valuable investment alternatives, there seem to be several shortcomings in current industry practice when it comes to fully capitalising on the advantages of including hedge funds in an investor's asset allocation. So far, the only solution to gaining exposure to alternative investment strategies for most investors was to invest in funds of hedge funds. Fund-of-hedge-funds solutions, however, are off-the-shelf products that are not designed to meet the specific needs of a specific investor. The key to implementing a structured top-down approach to investing in hedge funds is to properly define a target asset allocation across various hedge fund styles. A variety of investment constraints and objectives need to be taken into account in the design of the target allocation. In particular, the target allocation should be designed to allow for an optimal mix with the client's existing stock and bond portfolio.

In this article, we argue that it is only by taking into account the exact nature and composition of an investor's existing portfolio – as opposed to regarding hedge fund investing from a stand-alone approach – that institutional investors can truly customise and maximise the benefits they can expect from investing in these modern forms of alternative investment strategies. To achieve this objective, we first introduce suitably designed quantitative measures of the contribution of various hedge fund strategies to the risk in an existing stock/bond portfolio, and use them in the context of optimal selection of hedge fund strategies from an investor's standpoint. We then present an example of the design of a hedge fund benchmark in line with an investor's initial asset allocation, with a focus on extreme risk management.

Selecting the right hedge fund strategies

Amenc, El Bied & Martellini (2003) showed that different alternative investment strategies present dramatically different exposures to various risk factors. It is thus important to identify the strategies that are likely to offer the best diversification benefits with respect to traditional asset classes. To do so, we have to assess the extent to which they could lead, in a portfolio context, to a marginal reduction in average risk (measured by volatility) as well as in downside and extreme risks (measured in terms of skewness and kurtosis, respectively).

Since seminal work by Markowitz (1952), it is well known that this trade-off can be expressed in terms of mean-variance analysis

[1] *These numbers have been extracted from the 2005 SFS Hedge Fund Database Study, as reported for instance in The Hedge Fund Blog (www.cgml.co.uk/blog/archives/2006/01/hedge_fund_tran.html)*

under the restrictive assumption that asset returns are normally distributed. Recent research has shown that the returns of alternative funds are clearly not Gaussian (see, for example, Brooks & Kat (2002)). In the case of portfolios that include derivatives, the assumption of Gaussian returns is not tenable. Even if the return of the traditional asset classes were Gaussian, the return of funds using derivatives or dynamic strategies relating to those traditional classes would not be. Derivatives generally generate cashflows that are non-linear functions of the underlying asset value, and it is well known that a non-linear function of a Gaussian variable is not distributed in a Gaussian manner.

In line with the latest academic and practitioner research on the subject, we therefore recognise in our methodology the need to look beyond the first- and second-order moments of alternative fund return distributions, that is, their mean and variance. Because investors generally display a non-trivial preference for the third- and fourth-order moments of return distribution (see Scott & Horvath, 1980), these quantities must be taken into account when searching for strategies allowing good diversification properties.

To achieve this objective, we first introduce the following co-moments (co-variance, co-skewness and co-kurtosis):

$$CoV\left(R_i,R_j\right)=E\left[\left(R_i-E\left(R_i\right)\right)\left(R_j-E\left(R_j\right)\right)\right]$$
$$CoS\left(R_i,R_j\right)=E\left[\left(R_i-E\left(R_i\right)\right)\left(R_j-E\left(R_j\right)\right)^2\right]$$
$$CoK\left(R_i,R_j\right)=E\left[\left(R_i-E\left(R_i\right)\right)\left(R_j-E\left(R_j\right)\right)^3\right]$$

While the importance of higher order-co-moments in asset pricing theory has been already recognised (see, for example, Harvey & Siddique, 2000), existing measures of co-variation lack either a theoretical background or a statistical interpretability. In what follows, we will show that a simple and natural normalisation of higher-order co-moments exists, which generalises the concept of standard beta measures, and which can be used in the context of optimal selection of hedge fund strategies from an investor's standpoint.[2]

If we denote the initial portfolio by P, and the new portfolio as $P' = (1 - \varepsilon)P + \varepsilon A$, the marginal impact of the introduction of some small amount ε invested in a new asset A (for example, a hedge fund or hedge fund portfolio) on the second moment of portfolio distribution is, as a first-order approximation:

$$Var\left(R_{P'}\right)-Var\left(R_P\right)\approx -2\varepsilon Var\left(R_P\right)+2\varepsilon CoV\left(R_A,R_P\right)$$

From this, we obtain the following condition, which states that the introduction of the new asset A has led to a decrease in portfolio variance if the beta of this asset with respect to the initial portfolio is less than one:

$$Var\left(R_{P'}\right)\le Var\left(R_P\right)\Leftrightarrow \beta^{(2)}_{A/P}=\frac{CoV\left(R_A,R_P\right)}{Var\left(R_P\right)}\le 1$$

Similarly, we obtain that the marginal impact of the introduction of some small amount ε invested in a new asset A on the third moment of portfolio distribution is, as a first-order approximation:

$$\mu^{(3)}\left(R_{P'}\right)-\mu^{(3)}\left(R_P\right)=\mu^{(3)}\left(\left(1-\varepsilon\right)R_P+\varepsilon R_A\right)-\mu^{(3)}\left(R_P\right) \underset{\varepsilon\to 0}{\approx} -3\varepsilon\mu^{(3)}\left(R_P\right)+3\varepsilon CoS\left(R_A,R_P\right)$$

which leads to the following conditions:

$$\text{If } \mu^{(3)}\left(R_P\right)>0,\ \ \mu^{(3)}\left(R_{P'}\right)\ge\mu^{(3)}\left(R_P\right)\Leftrightarrow \beta^{(3)}_{A/P}\equiv\frac{CoS\left(R_A,R_P\right)}{\mu^{(3)}\left(R_P\right)}\ge 1$$

$$\text{If } \mu^{(3)}\left(R_P\right)<0,\ \ \mu^{(3)}\left(R_{P'}\right)\ge\mu^{(3)}\left(R_P\right)\Leftrightarrow \beta^{(3)}_{A/P}\equiv\frac{CoS\left(R_A,R_P\right)}{\mu^{(3)}\left(R_P\right)}\le 1$$

Finally, we obtain that the marginal impact of the introduction of some small amount ε invested in a new asset A on the fourth moment of portfolio distribution is, as a first-order approximation:

$$\mu^{(4)}\left(R_{P'}\right)-\mu^{(4)}\left(R_P\right)=\mu^{(4)}\left(\left(1-\varepsilon\right)R_P+\varepsilon R_A\right)-\mu^{(4)}\left(R_P\right) \underset{\varepsilon\to 0}{\approx} -4\varepsilon\mu^{(4)}\left(R_P\right)+4\varepsilon CoK\left(R_A,R_P\right)$$

which leads to the following condition:

$$\mu^{(4)}\left(R_{P'}\right)\le\mu^{(4)}\left(R_P\right)\Leftrightarrow \beta^{(4)}_{A/P}\equiv\frac{CoK\left(R_A,R_P\right)}{\mu^{(4)}\left(R_P\right)}\le 1$$

Hence, a low or negative second-, third- and fourth-moment beta indicates good diversification potential, in the sense of potential for a decrease in the overall portfolio average risk (that is, volatility), in the bias toward lower than average returns (that is, skewness), and in fat-tails (that is, kurtosis), respectively.[3] Standard beta coefficients indicate the extent to which an asset may generate second-order moment (that is, volatility) diversification effects when introduced in an existing portfolio. In the same spirit, third- and fourth-order betas indicate the extent to which an asset may produce third- and fourth-order moment (that is, skewness and kurtosis) diversification effects when mixed.

For the sake of illustration, we focus on a specific investor with a given initial asset allocation, that is, 20% in stocks and 80% in bonds.[4] This is consistent with a typical institutional investor in the presence of liability constraints. While no existing commercial index is perfectly suited to representing a specific investor's liability structure, bond indexes are typically the asset class most correlated with institutional investors' liabilities, and as a result should be predominant in institutional investors' allocation.

As can be seen in table A, all hedge fund strategies do not present the same diversification potential when considering the initial allocation of our traditional investor. Given the investor's original 80%/20% bond/stock allocation, Convertible arbitrage, CTA global, Equity market neutral, Fixed-income arbitrage and Short selling strategies present appealing diversification properties, while Emerging markets, Global macro or Long/short equity strategies show limited diversification potential.

An alternative to eliminating some strategies before launching the optimisation phase would be to include them all in an optimal allocation exercise such as the one presented in the next section, and let the optimiser eliminate the strategies, if any, that would be found to be unfit for the task at hand, that is, reducing the risk of an investor's existing portfolio. This alternative approach, however, leads to a potentially significant lack of robustness in the analysis because of greater sample-dependence. From a theoretical standpoint, it is well known (see Jagannathan & Ma, 2003) that the presence of portfolio constraints allows one to achieve a better trade-off between specifi-

[2] See Martellini & Ziemann (2005) for greater details on the estimation of higher order betas

[3] This is under the assumption of a portfolio with negative skewness, which is the case in our sample for the stock and bond proxies (see table A)

[4] We use here the S&P 500 to proxy equity market returns, and a portfolio made up of 20% US government bonds, 45% mortgage bonds and 35% corporate bonds to proxy bond market returns

A. Higher moment betas: January 1997–December 2004

	Convertible arbitrage	CTA global	Distressed securities	Emerging markets	Equity market neutral	Fixed-income arbitrage	Global macro	Long/short equity	Merger arbitrage	Short selling
2nd moment beta bond	0.20	0.24	0.50	1.44	0.25	0.04	0.80	0.99	0.39	–2.79
3rd moment beta bond	0.14	–0.51	1.61	3.54	0.26	0.38	1.09	1.28	0.72	–3.33
4th moment beta bond	0.24	0.55	0.46	1.41	0.32	0.01	0.87	0.88	0.39	–2.35

No diversification benefits · Limited diversification benefits · Significant diversification benefits · Very significant diversification benefits

Note: these results were obtained using the series of indexes published by Edhec. It is worth stressing that single hedge funds following these strategies might present different diversification profiles. However, as highlighted in Learned & Lhabitant (2002), portfolios comprising five to 10 hedge funds appear to be fairly representative of their investment universe. In an attempt to diversify idiosyncratic risk away, most multi-managers generally construct funds of hedge funds composed of more than 15 single hedge funds. The results presented above are therefore relevant for most indexes and funds of hedge funds purely by strategy

cation error and sampling error, similar to what can be achieved by statistical shrinkage techniques (see Jorion, 1986, and Ledoit, 1999). In other words, we advocate first selecting suitable hedge fund strategies presenting the most favorable higher-order betas in an attempt to 'help' the optimisation model, as opposed to letting the model decide to allocate freely across all strategies.

Optimising the hedge fund portfolio

Many studies have shown that in a mean-variance framework, mixing hedge funds with traditional assets leads to an enhancement of the return of the traditional portfolio with a constant risk level, or conversely, to a reduction of the risk level of the traditional portfolio with a constant return. This might easily be explained by the fact that hedge funds typically have low volatility, together with a low correlation with traditional asset classes, and relatively high returns. However, low volatility is not a free lunch. It is possible to show, through a statistical model integrating fatter tails than those of the normal distribution, that minimising the second-order moment (the volatility) is often accompanied by a significant increase in extreme risks (Sornette, Andersen & Simonetti, 2000). This is confirmed in Amin & Kat (2003), where the authors find empirical evidence that low volatility is generally obtained at the cost of lower skewness and higher kurtosis.

As stressed in Cremers, Kritzman & Page (2004), in the presence of asymmetric and/or fat-tailed return distribution functions, the use of mean-variance analysis can lead to a significant loss of utility for investors. Many attempts have therefore been made to better account for hedge fund-specific risk features. Building on Markowitz's initial proposition, some of them used downside risk measures, such as semi-deviation or lower partial moments (see McFall Lamm, 2003) or robust estimations of extreme risks, to define the risk dimension. In the latter case, some authors opted for the modified value-at-risk (see Favre & Galeano, 2002, and McFall Lamm, 2003) based on the expansion introduced in Cornish & Fisher (1937). Some others have preferred conditional VAR based on empirical distribution of hedge fund returns (see De Souza & Gokcan, 2004, or Agarwal & Naik, 2004) or using extreme value theory (see Bacmann & Gawron, 2004).

Another stream of research consists in maximising alternative ratios such as the omega ratio (Keating & Shadwick, 2002) or the alternative Sharpe ratio (Lee & Lee, 2004). The principal interest of such ratios is to take into account the whole return distribution function, or some approximation of the distribution function. While in Favre-Bulle & Pache (2003) the authors use the empirical returns of hedge funds, Passow (2005) suggests modelling higher-order moments applying the Johnson cumulated densities to calculate a so-called Johnson-omega. Johnson-omega's advantage is to extract persistent information from track record and exploit persistence of higher-order moments (that is, up to four). It therefore overcomes the instabilities within non-parametric approaches to omega, and allows for more reliable portfolios.

In what follows, we approach the question of optimal strategic asset allocation in a pragmatic manner. Because of the presence of large estimation risk in the estimated expected returns, we focus on selecting the one portfolio on the efficient frontier for which no information on expected returns is required, that is, the portfolio with the minimum amount of risk. This is similar to Amenc & Martellini (2002) or Alexander & Dimitriu (2004), except that we choose VAR, as opposed to portfolio volatility, as our risk measure. We use a pragmatic application of the VAR calculation in a fat-tail distribution environment based on a Cornish-Fisher expansion, allowing the skewness and kurtosis present in hedge fund returns to be taken into account in a semi-parametric way (see Favre & Galeano, 2002).

The Cornish-Fisher expansion is derived from the general Gram-Charlier expansion, using the standard normal distribution as the reference function. For a four-moment approximation of α-percentiles, the following formula is given:

$$\tilde{z}_a = z_a + \frac{1}{6}\left(z_\alpha^2 - 1\right)S + \frac{1}{24}\left(z_\alpha^3 - 3z_\alpha\right)K - \frac{1}{36}\left(2z_\alpha^3 - 5z_\alpha\right)S^2$$

where S denotes the sample skewness, K the sample's excess kurtosis and z_α the α-percentile of the standard normal distribution. $\tilde{z}_\alpha$ denotes the modified α-percentile.

This approximation is built on the assumption that the underlying distribution is close to a normal distribution. We obtain the modified VAR measure with confidence $(1 - \alpha)$[5]:

$$VAR_{\text{mod}}(1-\alpha) = -\left(\mu + \tilde{z}_\alpha \sigma\right)$$

where S denotes the standard deviation and μ the mean of the sample. In what follows, we will use a confidence level of 95%.

The reason why we focus on risk minimisation, as opposed to considering a more general risk-return trade-off, is because there is a general consensus that expected returns are difficult to obtain with a reasonable estimation error. Hedge fund performance measurement biases have largely been documented in the literature (see Fung & Hsieh, 2000 and 2002, among other examples). What makes the problem worse is that optimisation techniques are very sensitive to differences in expected returns, so portfolio optimisers typically allocate the largest fraction of capital to the asset class for which estimation error in the expected returns is the largest (for example, Britten-Jones, 1999, or Michaud, 1998).

On the other hand, it should be noted that this minimum-risk allocation can subsequently be used to infer neutral views on hedge fund strategies through a reverse-engineering process, and these neutral views may eventually be mixed with an investor's active views to implement a more standard return maximisation problem subject to

[5] *It should be noted that if the distribution is normal, S and K (representing the excess kurtosis in the formula) are equal to zero and consequently, z = Zc, and we come back to the Gaussian VAR*

a certain level of tail risk (see Martellini, Vaissié & Ziemann, 2006, for a suitable extension of the Black-Litterman Bayesian approach to a four-moment framework, with an application to hedge fund allocation style decisions).

We now provide a numerical illustration of the results.[6]

Measuring the benefits in terms of risk management

For illustration purposes, we will assume that the investor is willing to allocate as much as 15% to hedge funds. The resulting allocation is therefore 17% to stocks, 68% to bonds and 15% to hedge funds.[7] The control variables are how much to allocate to each selected hedge fund strategy so as to minimise the VAR of the whole portfolio.

We rebalanced the optimally designed portfolio every three months, using the 36 trailing months as the calibration period in the optimisation procedure. Since Edhec Alternative Indices' history starts in January 1997, the first calibration period was January 1997–December 1999. We therefore disposed of out-of-sample returns for the optimally designed portfolios from January 2000 onward. In figure 1, we show the dynamic evolution of this customised hedge fund portfolio.

To assess the benefits provided by the hedge fund portfolio in terms of both average and extreme risk diversification, we calculated the decrease in volatility and in modified VAR it entailed over the out-of-sample sample period. For comparison, we conducted the same experiment with a typical off-the-shelf fund of hedge funds.[8] We expect the diversification benefits allowed by the introduction of hedge funds into the portfolio to be much weaker in the latter case. As can be seen from figure 2, normal and extreme risks are reduced by 19.8% and 21.4% respectively when traditional asset classes are mixed with our optimally designed portfolio. When traditional asset classes are mixed with a standard fund of hedge funds, on the other hand, the benefits are significantly lower. Hence, using optimally designed portfolios generates a diversification effect as measured by a decrease in normal and extreme risks that is on average 52.2% and 58.7% higher than that allowed by the use of off-the-shelf products (see figure 3).

An interesting question relates to whether the diversification benefits stem from the selection stage or from the optimisation stage. To try to estimate the relative importance of these two steps, we proceed as follows. We use in the base case an equally weighted portfolio that comprised all the hedge fund strategies presented in table A (which we denote as Portfolio 1). We then estimate the risk level of a portfolio comprising 15% of this Portfolio 1 and 85% of our traditional portfolio. We subsequently do the same experiment with an equally weighted portfolio solely comprising the strategies chosen in the selection stage (which we denote as Portfolio 2), and an optimised portfolio solely comprising the strategies chosen in the selection stage (which we denote as Portfolio 3). Let us finally denote by $\delta_{2/1}$ the risk reduction (in absolute terms) provided by Portfolio 2 over Portfolio 1 at the global portfolio level, $\delta_{3/2}$ the risk reduction (in absolute terms) provided by Portfolio 3 over Portfolio 2 at the global portfolio level, and $\delta_{3/1}$ the risk reduction (in absolute terms) provided by Portfolio 3 over Portfolio 1 at the global portfolio level. The portion of the risk reduction attributed to the selection stage is then estimated as $\delta_{selection} = (\delta_{2/1})/(\delta_{3/1})$, while the one attributed to the optimisation stage is finally calculated as follows: $\delta_{optimisation} = (\delta_{3/2})/(\delta_{3/1})$. We therefore have: $\delta_{3/1} = \delta_{selection} + \delta_{optimisation}$.

Interestingly, the application of this methodology leads our sample to the following result: we find that 80% of the risk reduction benefits can be attributed to the choice of the strategies and 20% to the optimisation of the weights. We take this as a further indication of the relevance of the selection stage.

1 Time-evolution of strategy allocations within the customised hedge fund portfolio

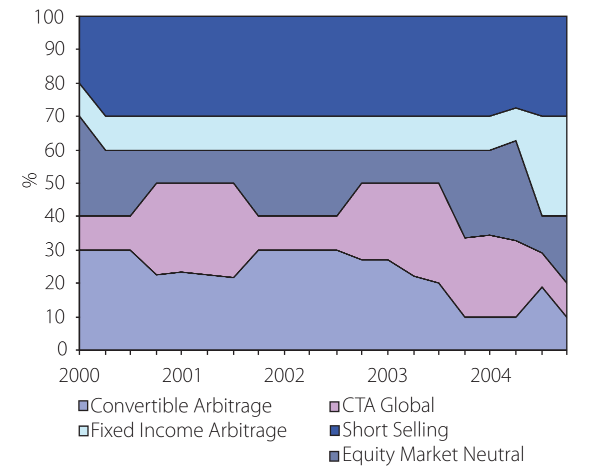

2 Diversification benefits

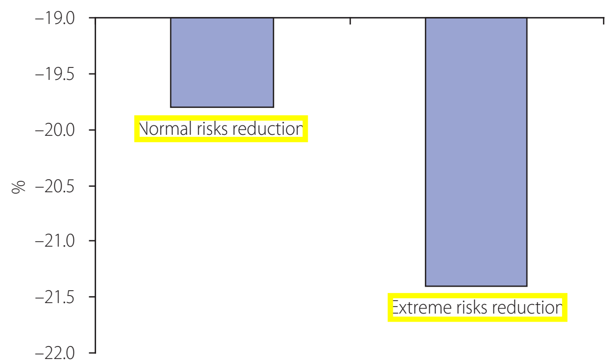

3 Diversification benefits' improvement with respect to fund-of-hedge-funds index

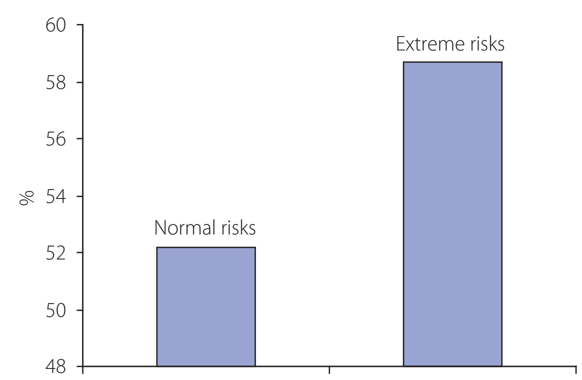

[6] We also impose portfolio constraints, ensuring that not less than 5% and not more than 40% is allocated to a single hedge fund strategy

[7] We have performed a number of robustness analyses that show that the added benefits of using strategy indexes, as opposed to composite indexes, are robust with respect to the original composition of the investor's traditional portfolio. Of course, the benefits of diversification increase with the fraction allocated to hedge funds

[8] We have used the Edhec funds of hedge funds index to proxy the return of a typical fund of hedge funds

Overall, these results confirm that hedge fund strategies can provide institutional investors with appealing diversification properties. However, most of them are still reluctant to invest a significant proportion of their assets in hedge funds. They must therefore strive to make the best out of their limited exposure to hedge fund strategies. The construction methodology we have presented is designed to maximise the benefits obtained by investors in terms of diversification. By doing so, investors reduce by 52.2% and 58.7% the normal and extreme risks they would have obtained with off-the-shelf products.

Conclusion

Investors, especially institutional investors, have long disposed of one single option to get exposed to hedge fund strategies, namely funds of hedge funds. However, funds of hedge funds generally do not provide them with sufficient information on underlying risk factor exposures, and have only limited liquidity. It was therefore difficult for investors to integrate hedge fund strategies into their global asset allocation process properly and, in turn, to fully benefit from their diversification properties.

Fortunately, in the meantime several index providers have launched an alternative to actively managed funds of hedge funds, namely investable hedge fund indexes. The analysis we have conducted in this article strongly suggests that these indexes appear to be natural tools for implementing strategic style allocation decisions in the hedge fund universe. As outlined in Amenc, El Bied & Martellini (2003), such indexes can also be used to implement active style timing decisions due to their full transparency and increased liquidity.[9] ■

Lionel Martellini is a professor of finance at Edhec Graduate School of Business and the scientific director of Edhec Risk and Asset Management Research Center. Mathieu Vaissié is a research engineer at Edhec Risk and Asset Management Research Center. Email: lionel.martellini@edhec.edu

[9] *Investable hedge fund indexes allegedly provide investors with a lower-cost solution to get exposed to hedge fund strategies. Investors should however be cautious. As stressed in Jacobson Fund Management (2004): "The subscription and redemption costs, notice periods and annual fees make the actual performance that investors can expect to realise from buy and hold investing substantially less than that reported for the underlying indexes"*

References

Agarwal V and N Naik, 2004
Risks and portfolio decisions involving hedge funds
Review of Financial Studies 17(1), pages 63–98

Alexander C and A Dimitriu, 2004
The art of investing in hedge funds: fund selection and optimal allocations
In Intelligent Hedge Funds Investing, edited by Barry Schachter, Risk Books

Amenc N, S El Bied and L Martellini, 2003
Evidence of predictability in hedge funds returns
Financial Analysts Journal 5(59), pages 32–46

Amenc N and L Martellini, 2002
Portfolio optimization and hedge fund style allocation decisions
Journal of Alternative Investments 5(2), pages 7–20

Amin G and H Kat, 2003
Stocks, bonds and hedge funds: not a free lunch!
Journal of Portfolio Management 29(4), pages 113–120

Bacmann J and G Gawron, 2004
Fat tail risk in portfolios of hedge funds and traditional investments
Working paper, RMF

Britten-Jones M, 1999
The sampling error in estimates of mean-variance efficient portfolio weights
Journal of Finance 54(2), pages 655–671

Brooks C and H Kat, 2002
The statistical properties of hedge fund returns and their implications for investors
Journal of Alternative Investments 5(2), pages 26–44

Cornish E and R Fisher, 1937
Moments and cumulants in the specification of distributions
Revue de l'Institut International de Statistique 4, pages 1–14

Cremers J, M Kritzman and S Page, 2004
Optimal hedge fund allocations: do higher moments matter?
Journal of Portfolio Management 31(3), pages 70–81

DeSouza C and S Gokcan, 2004
Allocation methodologies and customizing hedge fund multi-manager multi-strategy products
Journal of Alternative Investments 6(4), pages 7–21

Favre L and J Galeano, 2002
Mean modified value-at-risk optimization with hedge funds
Journal of Alternative Investments 5(2), pages 21–25

Favre-Bull A and S Pache, 2003
The omega measure: hedge fund portfolio optimization
MBF Master's Thesis, University of Lausanne

Fung W and D Hsieh, 2000
Performance characteristics of hedge funds and commodity funds: natural versus spurious biases
Journal of Financial and Quantitative Analysis 35(3), pages 291–307

Fung W and D Hsieh, 2002
Benchmark of hedge fund performance, information content and measurement biases
Financial Analysts Journal 58(1), pages 22–34

Georgiev G, V Karavas and T Schneeweis, 2002
Alternative investments in the institutional portfolio
Working paper, CISDM

Harvey C and A Siddique, 2000
Conditional skewness in asset pricing tests
Journal of Finance 55(3), pages 1,263–1,295

Jacobson Fund Management, 2004
Investable hedge fund indices: an assessment and a review
White paper

Jagannathan R and T Ma, 2006
Risk reduction in large portfolios: why imposing the wrong constraints helps
Journal of Finance 58, pages 1,651–1,683

Jorion P, 1986
Bayes-Stein estimation for portfolio analysis
Journal of Financial and Quantitative Analysis 21(3), pages 279–292

Keating C and W Shadwick, 2002
A universal performance measure
Journal of Performance Measurement 6(3), spring, pages 59–84

Learned M and F Lhabitant, 2002
Diversification: how much is enough?
Journal of Alternative Investments 5(3), pages 23–49

Ledoit O, 1999
Improved estimation of the covariance matrix of stock returns with an application to portfolio selection
Unpublished, UCLA

Lee B and Y Lee, 2004
The alternative Sharpe ratio
In Intelligent Hedge Fund Investing, edited by Barry Schachter, Risk Books

McFall Lamm R, 2003
Asymmetric returns and optimal hedge fund portfolios
Journal of Alternative Investments 6(2), pages 9–21

Markowitz H, 1952
Portfolio selection
Journal of Finance 7, pages 77–91

Martellini L, M Vaissié and V Ziemann, 2006
Investing in hedge funds: adding value through active style allocation decisions
Forthcoming in Journal of Portfolio Management

Martellini L and V Ziemann, 2005
Marginal impacts on portfolio distributions
Working paper, Edhec Risk and Asset Management Research Centre

Michaud R, 1998
Efficient asset management: a practical guide to stock portfolio optimization and asset allocation
Harvard Business School Press

Passow A, 2005
Omega portfolio construction with Johnson distributions
Risk April, pages 85–90

Scott R and P Horvath, 1980
On the direction of preference for moments of higher orders than the variance
Journal of Finance 35(4), pages 915–919

Sornette D, J Andersen and P Simonetti, 2000
Portfolio theory for 'fat tails'
International Journal of Theoretical and Applied Finance 3(3), pages 523–535

Terhaar K, R Staub and B Singer, 2003
An appropriate policy allocation for alternative investments
Journal of Portfolio Management 29(3), pages 101–110