IPE
INVESTMENT & PENSIONS EUROPE
SUMMER 2011

EDHEC-Risk Institute
Research Insights

EDHEC-RISK Institute

EDHEC-Risk Institute
Research Insights

Introduction Noël Amenc

It is my privilege to introduce the third EDHEC-Risk Institute Research Insights supplement to Investment & Pensions Europe. We are again delighted to be teaming up with IPE to provide information on research-based solutions to some of the key challenges facing institutional investors today.

In the first article in the present supplement, 'Is there a risk/return trade-off across stocks? An answer from a long-horizon perspective', Felix Goltz and Dev Sahoo address one of the key questions in modern finance from both an academic and practitioner perspective: are investors rewarded for investing in high-risk stocks by enjoying higher expected returns? By adopting a long-horizon standpoint, our research provides an unambiguously affirmative answer to this question: the trade-off between risk and return does indeed exist over long horizons as postulated by financial theory.

The second article discusses what is probably the single most important subject to come out of the financial crisis: risk management. As Stoyan Stoyanov explains, risk management is about maximising the probability of achieving certain objectives at the investment horizon while staying within a risk budget. Diversification, hedging, and insurance can be relied on to make optimal use of risk budgets. These three techniques involve different aspects of risk management, but they are complementary techniques rather than competing ones.

In the article on 'Efficient indices and benchmarks', we look at the difference between a reference index and a custom benchmark and explain how this distinction leads to different approaches to passive investment. We also introduce EDHEC-Risk Indices & Benchmarks' efficient relative return benchmark approach, which allows investors to benefit from the performance of efficient diversification while continuing to rely on the popularity and simplicity of traditional cap-weighted indices.

Felix Goltz also looks in a subsequent article at the advantages and shortcomings of minimum variance portfolios, and how the shortcomings can be addressed. Since minimum variance portfolios have only the objective of lowering risk, rather than aiming to optimise the risk/reward ratio, minimum variance portfolio optimisation leads to a pronounced concentration in low volatility stocks at the expense of exploiting correlation properties. While minimum variance portfolios have been shown to lead to relatively poor performance, they may be suitable for investors who wish to load up on low-risk or 'defensive' stocks.

In an article drawn from the Advanced Modelling for Alternative Investments research chair at EDHEC-Risk Institute, supported by the Prime Brokerage Group at Newedge, Lionel Martellini examines the question of optimal hedge fund allocation. This article discusses an application of improved estimators for higher-order comoment parameters in the context of hedge fund portfolio optimisation. In recent research we have found that the use of these enhanced estimates generates considerable benefits for investors in hedge funds.

In the final article in the supplement we analyse dynamic core-satellite strategies with exposure to value and momentum strategies. In this research, produced as part of the Core-Satellite and ETF Investment research chair in partnership with Amundi ETF, we find that these investment strategies alone could achieve higher returns but are exposed to high extreme risk because they consist of equity portfolios that are concentrated in the sectors with the highest value or momentum exposure. Combining these strategies with the dynamic core-satellite approach, however, improves portfolio returns while also keeping downside risk in check.

We wish you an informative and enjoyable read of the supplement and look forward to continuing this editorial partnership with IPE. Our mutual objective with this supplement is to provide academic insights that will genuinely contribute to improving institutional investment practices.

Noël Amenc, Professor of Finance, EDHEC Business School, and Director, EDHEC-Risk Institute

Contents

EDHEC-Risk Institute Research Insights is published as a supplement to Investment & Pensions Europe
IPE International Publishers Ltd, 320 Great Guildford House, 30 Great Guildford Street, London SE1 0HS
Tel: +44(0)20 7261 0666, Fax: +44(0)20 7928 3332, Web site: www.ipe.com, ISSN 1369-3727

Is there a risk/return trade-off across stocks?

An answer from a long-horizon perspective

Felix Goltz, Head of Applied Research, EDHEC-Risk Institute
Dev Sahoo, Research Engineer, EDHEC-Risk Institute

Theory and – perhaps more importantly – financial common sense suggest that there should be a trade-off between a stock's riskiness and its expected returns. On the one hand, standard asset pricing models suggest that systematic risk should be positively rewarded – ie, stocks with higher betas should earn a higher expected return (see Ross's Arbitrage Pricing Theory, 1976). Subsequently, research has underlined the explanatory power of stock-specific or so-called idiosyncratic risk for expected returns (Merton, 1987). Taken together, these results suggest that total volatility, which is the model-free sum of systematic volatility explained by a factor model, and idiosyncratic volatility, should also be positively rewarded (Martellini, 2008).

In contrast to this consensus regarding the existence of an unambiguously positive risk-return relationship from a theoretical perspective, a number of older as well as more recent papers have reported a number of puzzling or, at least, contrasted findings from an empirical perspective. First, the 'low beta anomaly' stipulates that the relationship between systematic risk as measured by a stock beta and return is much flatter than predicted by the Capital Asset Pricing Model (see early papers by Black, 1972; Black, Jensen and Scholes, 1972) and sometimes even inverted (paper by Haugen and Heins, 1975). More recently, Ang, Hodrick, Xing and Zhang (2006, 2009) have drawn new attention to these results with a focus on the specific risk component, finding that high idiosyncratic volatility stocks have had "abysmally low returns" in longer US samples and in international markets. This result is now widely known as the 'idiosyncratic volatility puzzle'. Yet other papers have documented a rather flat or even negative relationship between total (as opposed to specific) volatility and expected return, an anomaly that some call the 'total volatility puzzle' (Haugen and Baker, 2008; Blitz and Van Vliet, 2007; Baker, Bradley and Wurgler, 2011).

Several attempts have been made to explain these puzzling empirical results. A number of recent papers have questioned the robustness of Ang *et al*'s (2006, 2009) results. Among other concerns, the findings are not robust to changes in data frequency, portfolio formation period, to the screening out of illiquid stocks (Bali and Cakici, 2008), or to adjusting for short-term return reversals (Huang *et al*, 2010). Several other authors have changed the short-term measure of volatility in Ang *et al* (2006) with conditional measures estimated from returns over longer calibration periods and found a positive relationship (Fu, 2009; Brockmann and Schutte, 2007).

In order to understand the risk-return relationship further, recent research at EDHEC-Risk Institute adopts a long-horizon perspective. Rather than looking at return realisations over a monthly horizon to proxy for expected returns, we look at longer investment horizons beyond one year, which are arguably closer to horizons relevant for a typical institutional equity investor. Taking a long-horizon perspective seems to be the natural approach since the theoretical predictions of standard asset pricing models relate to the relationship between a stock's risk estimate and expected return on that stock across many varying market conditions. This can only be assessed by looking at a stock performance over long horizons. A similar approach is taken by Bandi and Perron (2008) in recent work on the long-term risk/return relationship, with a focus on the time-series perspective. In a related effort, from a cross-sectional perspective, Bandi *et al* (2010) find empirical support for an approximate long-run version of the CAPM, where betas and returns are both measured over long horizons.

1. Risk/return profile for portfolios using long and short horizon

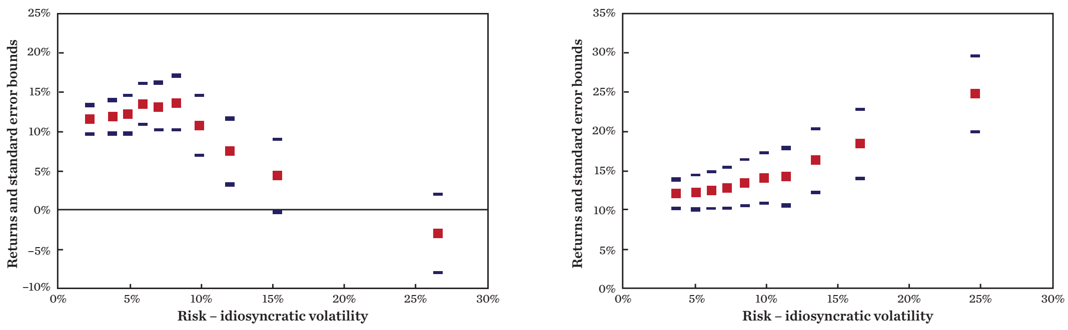

The graphs show the arithmetic average return and average risk of decile portfolios built on idiosyncratic volatility (relative to Fama-French Factor exposures). The values plotted are the average values over all cross sections and are annualised. The bounds of the one-standard-deviation error are shown along the return axes, along with the average values. The left-hand graph replicates the short horizon results of Ang *et al* (2006) while the right-hand graph uses a longer horizon of 24 months. The period of analysis runs from July 1963 to December 2009.

Below we report results obtained using a broad cross-section of US stocks over the period from July 1963 to December 2009. First, we replicate the short-horizon findings in the early literature on the idiosyncratic volatility puzzle. To study the effect of idiosyncratic volatility on long-horizon expected returns, we use a simple trading strategy similar to that used by Jegadeesh and Titman (1993). This trading strategy sorts stocks into portfolios every month by their volatility and holds the portfolios over longer horizons of up to three years to allow expected return differences to materialise[1].

From the short-horizon results, we can see that the high volatility portfolio has negative returns over the following month. As Figure 1 shows, the finding of a negative relation between risk and return can be largely attributed to the portfolios of high-volatility stocks. The finding of a negative relation does not hold when comparing returns across the first six portfolios. This is similar to the findings of Ang *et al*. It is also well known (Huang *et al* 2010) that the short-horizon underperformance of high-volatility stocks can be attributed to the short-term return reversal effect. The cap-weighting scheme employed to weight stocks within the portfolio further increases the exposure to short-term return reversals as it overweights the past winners. Using equally-weighted portfolios (which are not sensitive to past returns) and holding the portfolios for longer horizons avoids short-term reversal effects and provides a clear positive risk/return relationship across all the portfolios.

In order to show the horizon effect more clearly we show the geometric mean returns of the high volatility and low volatility portfolios over different horizons in figure 2. The low-volatility portfolio yields higher return over a one-month horizon (ie, the risk/return relationship appears to be negative). But holding the stocks for longer horizons shows the opposite relationship (high-volatility portfolios have higher returns and lower-volatility portfolios have lower returns).

Taking the risk/return relationship beyond the mean-variance setting, theoretical models have also shown that investors are willing to accept lower expected returns and higher volatility compared to the mean-variance benchmark in exchange for higher skewness and lower kurtosis of returns. High skewness and low kurtosis have been shown to be associated with lower expected returns in theory (see, for example, Barberis and Huang, 2004, Mitton and Vorkink, 2007, among many others). The intuition behind this result is that investors like to hold portfolios with positive skew and low kurtosis. In terms of idiosyncratic and total risk measures, Boyer, Mitton and Vorkink (2010) and Conrad, Dittmar and Ghysels (2008) also provide converging empirical evidence that individual stocks' skewness and kurtosis are indeed

1 Multiple portfolios are held at the same time and only those that arrive at their horizon are rebalanced in a given month.

2. Geometric mean returns for multi-portfolio analysis

Geometric average returns		1 month	3 months	12 months	24 months	36 months
Idiosyncratic volatility:	Low	14.21%	14.10%	13.34%	11.92%	12.07%
	High	12.63%	15.89%	21.89%	21.63%	24.09%

The table shows the results using geometric averaging of portfolio returns over the whole analysis period. The annualised returns of the high and low portfolios are provided for both risk measures. The period of analysis runs from July 1963 to December 2009.

3. Multivariate regression results using three risk measures

	Idiosyncratic volatility	Idiosyncratic skewness	Idiosyncratic kurtosis	R-squared
Slope coefficient	0.1428	–0.0065	0.0002	6.45%
t-stat	2.00	–3.98	1.55	

Each month we run stock-level regressions (similar to a Fama-MacBeth stock regression) using the stock returns as the dependent variable and the idiosyncratic/total risks (volatility, skewness, and kurtosis) as the independent variables. The stock returns used are the average monthly returns for the next 24 months and the risk measures use the historical daily data for the previous 12 months. This table shows the average values of the coefficients of regression and R-squared values over all of the cross-sections. Newey-West is used to correct the t-stats. The period of analysis runs from July 1963 to December 2009.

positively related to future returns.

To assess the impact of higher-order risk measures along with volatility we performed a multivariate regression analysis. This assessment is potentially important to take into account any link between volatility, skewness and kurtosis. Also, using such analysis we can test the joint impact of all of these risk measures. We run a monthly regression by using volatility, skewness and kurtosis as the independent variables. Figure 3 shows the average slope coefficient estimate and R-squared, over all the cross sections, as well as the autocorrelation-adjusted t-statistics.

The regression results confirm the strong positive risk/return relationship for volatility and skewness[2]. The effect of kurtosis is insignificant when used along with volatility and skewness. On the whole, these multivariate results suggest that the greatest effects of risk on expected stock returns stem from volatility and skewness.

Although the results of past research into the cross-sectional relationship between idiosyncratic/total risk and expected stock returns are puzzling, the results are not universal and the puzzle exists only as a short-term effect that depends on how we go about measuring volatility and its effects. Various authors have shown that the risk-return relation is positive, for example using value-at-risk instead of volatility (Bali *et al*, 2004) and after adjusting for reversal effects (Huang *et al*, 2010). Overall, the case for a negative relationship is not only contrary to common sense and theory but also weak empirically given the opposing evidence in empirical papers. Our findings suggest that the ambiguous nature of research into the risk/return trade-off may be accounted for in part by the horizon used in that research. Our results provide evidence that, although there may well be short-term anomalies of higher risk not leading to higher expected returns, the trade-off between risk and return plays out over longer horizons much as is posited by financial theory.

2 For skewness, we expect a negative relationship since negatively skewed stocks are riskier.

References

Amenc, N., F. Goltz, L. Martellini and D. Sahoo. 2011. *Is There a Risk/Return Tradeoff across Stocks? An Answer from a Long-Horizon Perspective,* EDHEC-Risk Institute Publication

Ang, A., R. J. Hodrick, Y. Xing and X. Zhang. 2006. The cross-section of volatility and expected returns. *Journal of Finance* 61 (1): 259–99.

Ang, A., R. J. Hodrick, Y. Xing and X. Zhang. 2009. High idiosyncratic volatility and low returns: International and further US evidence. *Journal of Financial Economics* Elsevier 91(1): 1–23.

Baker, M. P., B. Bradley and J. Wurgler. 2011. Benchmarks as limits to arbitrage: Understanding the low volatility anomaly. *Financial Analysts Journal* 67 (1): 1–15.

Bali, T. G., and N. Cakici. 2004. Value at Risk and expected stock returns. *Financial Analysts Journal* 60 (2): 57–73.

Bali, T., and N. Cakici. 2008. Idiosyncratic volatility and the cross-section of expected returns? *Journal of Financial and Quantitative Analysis* 43: 29–58.

Bandi, F. M., and B. Perron. 2008. Long-run risk-return trade-offs. *Journal of Econometrics* 143 (2): 349–74.

Bandi, F. M., R. Garcia, A. Lioui and B. Perron. 2010, A long-horizon perspective on the cross-section of expected returns, working paper.

Barberis, N., and M. Huang. 2007. Stocks as lotteries: The implications of probability weighting for security prices. NBER Working Papers 12936, National Bureau of Economic Research.

Black, F. 1972. Capital Market Equilibrium with Restricted Borrowing. *Journal of Business* 45 (3): 444–454.

Black, F., M. C. Jensen and M. Scholes. 1972. The Capital Asset Pricing Model: Some Empirical Tests. *Studies in the Theory of Capital Markets*. Michael C. Jensen, ed. New York: Praeger: 79–121.

Blitz, D. C., and P. Van Vliet. 2007. The volatility effect: Lower risk without lower return. *Journal of Portfolio Management* 34 (1): 102–13.

Boyer, B., T. Mitton, and K. Vorkink. 2010. Expected idiosyncratic skewness. *Review of Financial Studies* 23 (1): 169–202.

Brockman, P., and M. Schutte. 2007. Is idiosyncratic volatility priced? The international evidence. Working paper, University of Missouri-Columbia.

Conrad, J., R. F. Dittmar and E. Ghysels. 2008. Ex ante skewness and expected stock returns. Working paper.

Fu, F. 2009. Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics* 91 (1): 24–37.

Haugen, R. A., and A. J. Heins. 1975. Risk and the Rate of Return on Financial Assets: Some Old Wine in New Bottles. *Journal of Financial and Quantitative Analysis* 10 (5): 775–784.

Haugen, R. A., and N. L. Baker. 2008. Case Closed. *The handbook of portfolio construction: Contemporary applications of Markowitz techniques*, John B. Guerard Jr ed. Forthcoming.

Huang, W., Q. Liu, S. G. Rhee and L. Zhang. 2010. Return reversals, idiosyncratic risk, and expected returns. Review of Financial Studies 23 (1): 147–68.

Jegadeesh, N., and S. Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48 (1): 65–91.

Martellini, L. 2008. Toward the Design of Better Equity Benchmarks: Rehabilitating the Tangency Portfolio from Modern Portfolio Theory. *Journal of Portfolio Management* 34 (4): 34–41.

Merton, R. C. 1987. A simple model of capital market equilibrium with incomplete information. *Journal of Finance* 42 (3): 483–510.

Mitton, T., and K. Vorkink. 2007. Equilibrium Underdiversification and the Preference for Skewness. *Review of Financial Studies* 20: 1255–1288.

Ross, S. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13 (3): 341–60.

A post-crisis perspective on diversification for risk management

Noël Amenc, Director, EDHEC-Risk Institute
Felix Goltz, Head of Applied Research, EDHEC-Risk Institute
Stoyan Stoyanov, Head of Research, EDHEC-Risk Institute – Asia, Singapore

The global financial crisis of 2008 has shifted the attention of investors to risk management. From the standpoint of long-term investors, the goal of risk management is to maximise the probability of achieving long-term objectives while satisfying short-term constraints. Although different institutions have different short-term constraints, a unifying feature is that they usually concern the downside – eg, a minimum funding ratio or maximum drawdown constraint. Since the global crisis, there has been a special focus on management of extreme risks motivated by a genuine interest from industry.

Critical reviews of some industry practices blamed the negative impact of the crisis on inadequate diversification largely caused by the shortcomings of modern portfolio theory (MPT). Although technically justified, the criticisms reveal common misconceptions not only about the benefits and limitations of diversification, but also about its relationship with hedging and insurance. In this article, our goal is to review diversification as a general method and also hedging and insurance, sometimes incorrectly regarded as competing techniques.

Diversification: advantages and disadvantages

Diversification is one of the most widely used concepts in modern finance. Although the idea behind it has long existed, a scientifically consistent framework for diversification, MPT, was first posited by Markowitz (1952). He introduced variance as a proxy for risk and formulated a portfolio construction technique ►

◄ by means of an optimisation problem – combine risky assets in such a way as to minimise variance at each level of expected return. The resulting set of portfolios describes the efficient frontier. Diversification – international diversification, sector and style diversification, and so on – has since become the pillar of many investment philosophies. It has also become a very important risk management technique, so much so that it is often considered, erroneously, as being synonymous with risk management.

It is a common misconception to regard diversification as a method of risk reduction, especially in the context of extreme risk management. In fact, according to MPT diversification is related to risk reduction as much as it is to improving performance and, therefore, it is most effective when it is used to extract risk premia. A key insight of MPT is that the notion of efficient portfolios depends critically on the joint behaviour of asset returns.

This insight identifies one limitation of the method – diversification is less effective when asset returns are more highly correlated and, empirically, correlations increase in times of market crashes. In the extreme case of 100% correlation between stocks, holding a diversified portfolio is as efficient as holding only a few stocks. In such an environment, if we were to rebalance the global minimum variance portfolio, which is the lowest risk portfolio of the efficient frontier, it would be composed of only one stock.

Amenc, Goltz and Stoyanov (2011) emphasise that the importance of joint behaviour for the effectiveness of diversification extends beyond mean-variance analysis. Diversification opportunities are in fact determined by the multi-dimensional dependence structure of asset returns. The choice of risk measure influences how these opportunities are translated into actual allocations but if no opportunities are present, then the functional form of the risk measure is irrelevant. As a consequence, switching to a downside risk measure may improve portfolio construction in general but will not provide better loss protection in big market crashes.

A central point regarding choosing MPT versus alternative methods as a framework for implementing diversification is the classic trade-off between model sophistication and estimation complexity. As a consequence, theoretical benefits may not materialise because of difficulties with empirical estimation. In applied research, model enhancement should balance both theoretical and empirical concerns – see Martellini and Ziemann (2010) for a discussion.

Finally, being a method for efficient extraction of risk premia, diversification is designed to work over long horizons rather than in specific market conditions. This insight alone renders diversification useless in times of market downturns. Furthermore, perhaps more importantly, it underlines the significance of a robust quantitative methodology for proper implementation of diversification – see Amenc *et al* (2010b).

Misunderstanding the limitations of diversification can mislead investors into concluding that, since it did not protect them from big losses in 2008, diversification is a useless concept in general when it was never meant to provide loss protection in the first place.

Hedging: fund separation and risk reduction

Diversification is unreliable in highly correlated markets and it is not an efficient technique for loss control in the short term. To limit losses in market downturns, investors should identify techniques that can complement it and offset its shortcomings. One potential technique is hedging.

1. A comparison of simulated average annual return distribution and their PSP components with a 10% maximum drawdown floor

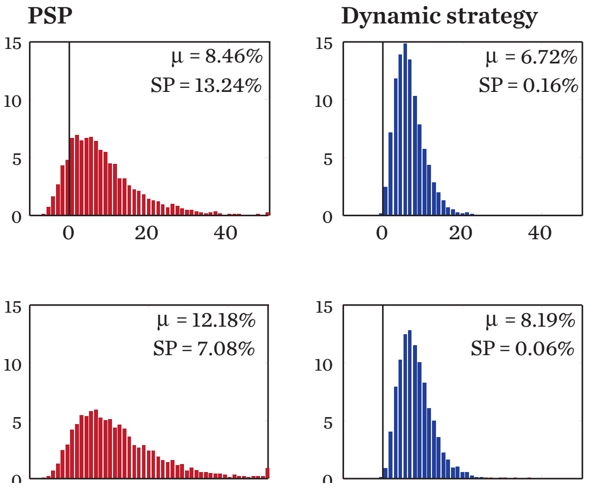

The bottom set is obtained after a 50% improvement of the Sharpe ratio of the PSP component. The investment horizon is 10 years.

In the presence of a risk-free asset, Tobin (1958) argued that risk-averse investors should hold portfolios of only two funds – the risk-free asset and a fund of risky assets characterised by the property of having the highest Sharpe ratio. The degree of risk aversion influences only the relative weight of the two funds, not their composition. The implication for portfolio construction is that lower-risk portfolios are best obtained by increasing the weight of the risk-free asset rather than by re-optimising the fund of risky assets and, as a consequence, global minimum variance portfolios are inefficient in the presence of a risk-free asset.

This arrangement is the result of a so-called two-fund separation theorem and can be regarded as a simple form of hedging. Amenc *et al* (2010a) point out that the two-fund separation implies that any risk-averse investor can construct portfolios in two steps: (a) build the maximum Sharpe ratio (MSR) portfolio designed to extract risk premia from risky assets and (b) depending on the degree of risk-aversion, partially hedge the risk present in the MSR portfolio by allocating a fraction of the capital to the risk-free asset. In the second step, the risks in the MSR portfolio are partially hedged, which limits the downside of the return distribution. There is, however, a caveat. Along with the left tail, this technique scales down the right tail as well. As a consequence, limited drawdown comes at the cost of lower upside potential.

Because of the static nature of hedging, the allocation to the risk-free asset does not depend on the state of the market and a specific drawdown limit can be achieved only ex-post. Although in the presence of liabilities the two-fund separation implies that the risk-free asset is not cash but a portfolio designed to hedge the liabilities, the same disadvantage holds – see Amenc *et al* (2010a) for the asset-liability management (ALM) context).

Insurance: dynamic risk management

Drawdown, being a dynamic characteristic, is best managed through dynamic portfolio theory. Posited by Merton (1969), the theory presents the most natural form of risk management, generalising substantially the static portfolio selection model of MPT. The theory has been extended with absolute or relative constraints on asset value that can accommodate, for example, maximum drawdown and rolling performance floors. Simple insurance strategies such as CPPI and OBPI arise as dynamic optimal strategies for investors subject to particular explicit or implicit floor constraints respectively.

Dynamic fund separation in an ALM context suggests constructing a performance-seeking portfolio (PSP) and a building block of safe assets dedicated to hedging the stream of liabilities.[1] The allocation to the two building blocks is generally not constant and depends on the market state. Thus, in a market downturn, the closer the portfolio value gets to the floor constraint, the higher the allocation to the safe building block becomes, limiting the potential for further losses. See Martellini and Milhau (2010) for a more general discussion.

Conceptually, the functional separation of the building blocks provides a key insight: diversification, hedging, and insurance are complementary techniques. The first two are responsible for the optimal construction of the PSP and the safe building blocks and the method of insurance guarantees that the floor constraint is satisfied through the dynamic and state-dependent weights of the building blocks.

Unlike diversification, however, insurance always comes at a cost. The cost can materialise in different ways depending on implementation: as an implicit opportunity cost, if the implementation is through dynamic trading (eg, CPPI), or explicitly as the price of a derivative overlay (eg, OBPI). The top two plots in figure 1 provide an illustration with a dynamic strategy with a 10% maximum drawdown floor constraint. The benefit of the maximum drawdown constraint results in a much lower shortfall probability (SP) of getting a negative return and the implicit opportunity cost of insurance materialises as a lower expected return of the dynamic strategy compared to that of the PSP component.

Concerning the PSP component itself, the common approach is to adopt a stock market index: a cap-weighted portfolio that is concentrated and highly inefficient.[2] Therefore, we can improve the performance of the dynamic strategy by improving the efficiency of the PSP component. Amenc *et al* (2010b) demonstrate that optimal diversification based on a robust methodology can consistently improve the Sharpe ratio of S&P 500 by more than 50%.

The bottom set of plots in figure 1 is produced after a realistic improvement of 50% of the Sharpe ratio of the PSP component. The expected return of the dynamic strategy increases and almost matches the expected return of the PSP component before the improvement, nearly compensating for the cost of insurance. This illustration underlines an important practical benefit of carefully implemented optimal diversification: it can partially offset the implicit cost of insurance.

Conclusion

In a broad context, risk management is about maximising the probability of achieving certain objectives at the investment horizon while staying within a risk budget. Diversification, hedging and insurance can be relied on to make optimal use of risk budgets. These three techniques involve different aspects of risk management, but they are complementary techniques rather than competing ones.

Diversification provides investors with the best reward per unit of risk through a smart combination of individual assets. It is designed to

1 In the presence of a stochastic opportunity set, there are also hedging demands that represent another component, see Martellini and Milhau (2010).

2 See, for example, Amenc *et al* (2006) for an extensive empirical study.

work in the long run across different market conditions and is, therefore, helpless in such specific conditions as severe market downturns. Since the purpose of diversification is efficient extraction of risk premia, it is most effective in the construction of performance-seeking portfolios.

Hedging can be combined with diversification to reduce risks that cannot be diversified away. Hedging is achieved through a portfolio of safe assets, or simply through cash, which is another dedicated building block. A non-diversifiable risk that can be handled in this way is the risk of a large drawdown.

Insurance, unlike diversification and hedging, combines the building blocks optimally to comply with the corresponding risk budgets. So downside risk control is best achieved through dynamic asset allocation. This technique makes it possible to control the downside of the return distribution while preserving access to the upside through the performance-seeking portfolio. In this context, a well-diversified portfolio is a building block of crucial importance. A carefully designed performance-seeking portfolio with an improved Sharpe ratio resulting from good diversification can reduce the implicit cost of insurance.

References

Amenc, N., F. Goltz and V. Le Sourd. 2006. Assessing the quality of stock market indices. EDHEC-Risk Institute Publication (September).

Amenc, N., F. Goltz, L. Martellini and P. Retkowsky. 2010b. Efficient indexation: An alternative to cap-weighted indices. EDHEC-Risk Institute Publication (January). Forthcoming in *Journal of Investment Management*.

Amenc, N., F. Goltz and S. Stoyanov. 2011. A post-crisis perspective on diversification for risk management. EDHEC Risk-Institute publication (May).

Amenc, N., L. Martellini, F. Goltz and V. Milhau. 2010a. New frontiers in benchmarking and liability-driven investing. EDHEC Risk-Institute publication (September).

Markowitz, H. M. 1952. Portfolio selection. *Journal of Finance* 7 (1): 77–91.

Martellini, L., and V. Milhau. 2010. Hedging versus insurance: Long-term investing with short-term constraints. Working paper, EDHEC-Risk Institute.

Martellini, L., and V. Ziemann. 2010. Improved estimates of higher-order comoments and implications for portfolio selection. *Review of Financial Studies* 23 (4): 1467–1502.

Merton, R. C. 1969. Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics* 51: 247–57.

Tobin, J. 1958. Liquidity preference as behavior towards risk. *Review of Economic Studies* 25: 65–86.

Efficient indices and efficient relative return benchmarks

Noël Amenc, Professor of Finance, EDHEC Business School, Director, EDHEC-Risk Institute and Chairman, EDHEC Risk Indices & Benchmarks

The words 'index' and 'benchmark' are often used indiscriminately in practice even though they are two *a priori* very different concepts:

- A reference index is a portfolio that should represent the performance of a given segment of the market, so the focus is on representativeness;
- A custom benchmark is a portfolio that should represent the fair reward expected in exchange for risk exposures that an investor is willing to accept, so the focus is on efficiency.

For most investors this distinction may be semantic, but it leads clearly to different approaches to passive investment.

For example, an index that is constructed differently to a cap-weighted index will always be considered a substitute for the latter, so it seems normal that investors would expect this new reference index to have the same level of transparency, and perhaps the same level of popularity, as the previous one. In the end, what determines the success of a new reference index will be as much its financial characteristics as its 'popularity', not only with investors but also with consultants.

Naturally, implementation of a new form of reference index is not risk-free. All rebalancing schemes, with the notable exception of cap-weighting or equal-weighting, assume a certain level of out-of-sample stability in the structures that led to the in-sample estimation of the parameters. Whether one tries to reduce the dimensions of a variance-covariance matrix with a factor model, or uses accounting attributes to define the size of a company and *de facto* its position in an index, or creates a link between the risk and return of a stock, all of these methodological choices are more or less relevant depending on the period chosen. That is why we have always considered that the evaluation of an alternative weighting scheme for an index can only be carried out over a long period; that globally this evaluation could not concern the ability of diversification to reduce portfolios' risks (see Amenc, Goltz and Stoyanov, 2011) but instead involves obtaining greater efficiency in the investment over long periods – ie, a better return for each unit of risk. This serious approach to the performance of alternative forms of indices will probably lead investors to diversify the alternative forms of investment. As an example, it is interesting to observe that minimum volatility and efficient indices do not have the same outperformance in relation to cap-weighted indices in different market conditions.

A custom benchmark does not necessarily aim to replace an index because the objective in using it relates to the implementation of a passive investment strategy. The goal of the custom benchmark is not to serve as an external reference for the investment but to be a genuine representation of the investor's inter- or intra-class allocation choices. Ultimately, it is not so much the 'popularity' of a benchmark that will lead to its success but its customisation and appropriateness to reflect the investor's strategic choices in terms of risk and reward. Since it is not being used as an external reference, a custom benchmark will be judged less on its transparency or its relative simplicity than on its capacity to enable investors to achieve their diversification objectives, notably with regard to an external reference represented by a cap-weighted index.

The Efficient Relative Return Benchmark methodology enables institutional investors to benefit from the latest progress in the area of diversification in order to avail themselves of customised benchmarks that are representative not only of their choice of absolute risks (geographic, sector, style, etc) but also of their relative risks by implementing a particularly innovative and efficient process for managing tracking error with respect to a market index.

This relative return approach allows investors to limit the risk of eventual underperformance when market conditions do not allow efficient indices to outperform (which is the case in speculative bubbles when diversification is not useful and momentum is the best investment) and obviously given the fact that, as with the vast majority of alternative forms of indices, there can be moments when the in-sample estimation, through significant deformation of the structures (eg, correlation), loses its out-of-sample robustness. The relative return benchmark approach represents a choice of implementation of the efficiency concept that is more modest, and less high-performance, but also less risky.

Ultimately, this efficient relative return benchmark offering allows investors to benefit from the performance of efficient diversification while continuing to rely on the popularity and simplicity of traditional cap-weighted indices for their global asset allocation and also for their communication to all their stakeholders.

References

Amenc, N., F. Goltz, L. Martellini and P. Retkowsky. 2010. Efficient Indexation: An Alternative to Cap-Weighted Indices. EDHEC Risk-Institute publication (January).

Amenc, N., F. Goltz, and S. Stoyanov. 2011. A post-crisis perspective on diversification for risk management. EDHEC Risk-Institute publication (May).

1. Risk-adjusted performance and relative risk of relative return benchmarks and indices: US large-cap universe (500 stocks)

	Return	Volatility	Sharpe ratio	Excess return	Tracking error	Information ratio	Worst tracking error (95% confidence)	Worst relative return (95% confidence)
Relative return efficient benchmark 3% TE	11.0%	15.4%	0.37	1.4%	2.3%	0.63	3.8%	–3.0%
Efficient index	12.2%	14.8%	0.46	2.6%	4.2%	0.63	7.6%	–6.1%
Equal-weighted index	11.8%	16.3%	0.39	2.2%	4.8%	0.46	9.3%	–7.7%
S&P 500 cap-weighted (CRSP)	9.6%	15.5%	0.27	0.0%	0.0%	0.00	0.0%	0.)%

This table shows performance statistics computed based on weekly total returns data from January 1959 to December 2010. Worst tracking error and relative returns refer to the 5th percentile of most extreme values observed over the entire period for a rolling one-year window when assessing this at the end of each quarter. The relative return benchmarks use a target tracking error level of 3% and aims at reliably controlling the extreme tracking error. Data for the EDHEC-Risk Relative Return Efficient Benchmark Series is used. The efficient index is based on EDHEC-Risk's Efficient Index long-term US data. The equal-weighted index is rebalanced quarterly using the same constituents. The cap-weighted index for the S&P 500 universe is computed by CRSP.

Advantages and shortcomings of minimum variance portfolios

Felix Goltz, Head of Applied Research, EDHEC-Risk Institute

Scientific diversification is based on reaching a high risk/return objective through portfolio construction techniques. It corresponds to the desire of investors to position their portfolios on the efficient frontier of Modern Portfolio Theory. Approaches based on scientific diversification are increasingly being used to construct equity core portfolios in institutional investment management, with the aim of gaining an optimal risk-reward profile from exposure to equities.

In practice, to obtain a decent proxy for efficient portfolios, one needs to use careful risk and return parameter estimates. In essence, practical approaches to equity portfolio construction based on scientific diversification make different choices regarding the challenge of risk and return estimation.

The minimum variance portfolio is a remarkable portfolio that provides the lowest possible portfolio volatility. This means that the only optimisation inputs required are correlations and volatilities. Since the estimation risk inherent in expected returns is well known, the fact that the minimum volatility portfolio relies only on risk parameters is an appealing feature (see for example Amenc and Martellini, 2002, among many others who have made this argument). To be sure, estimating risk parameters is also a serious challenge, with issues such as the 'curse of dimensionality' when dealing with a large number of assets and the 'curse of non-stationarity' of risk in the stock market. However, ever since Markowitz published his theory of 'portfolio selection' in the 1950s, constant progress has been made on dealing with these challenges and today we dispose of a rich set of tools that allows risk estimates to be improved, including sophisticated factor models as well as dynamic risk models.

Despite the reasonable idea of avoiding expected returns estimation, minimum variance portfolios have been shown to lead to relatively poor performance. DeMiguel, Garlappi and Uppal (2009) for example evaluate a range of minimum variance portfolios across seven empirical datasets, and they find that none is consistently better than a simple equal-weighting rule in terms of Sharpe ratio.

In addition to such empirical findings, there are two main interrogations with minimum variance portfolios. From an ex-ante perspective, minimum variance portfolios are not optimal portfolios. They will be dominated by a combination of the risk/reward optimal portfolio (tangency portfolio) with cash. In principle, investors should only care about designing this tangency portfolio, using cash holdings, if they wish to reduce their portfolio's volatility to a lower level. This tangency portfolio will only coincide with the minimum variance portfolios if one is ready to assume that expected returns of all assets are identical, clearly a strong and rather unrealistic assumption.

From an ex-post perspective, minimum variance portfolios are typically heavily concentrated in the assets with the lowest volatility. The high concentration in GMV portfolios is a widely recognised issue. Clarke, De Silva and Thorley (2011) note that their long-only minimum variance "portfolio averages about 120 long securities, ie, about 12% of the 1,000-security investable set". Likewise, DeMiguel, Garlappi, Nogales and Uppal (2009) note that "shortsale-constrained minimum-variance portfolios [...] tend to assign a weight different from zero to only a few of the assets". This concentration can be seen in figure 1, where we sort categories of stocks according to their volatility and analyse the weights allocated to each category in a minimum volatility portfolio and in a cap-weighted portfolio.

In equity portfolio construction, such concentration in low-volatility stocks leads to a pronounced sector bias towards utility stocks. Chan, Karceski and Lakonishok (1999, Table 4) report that a typical minimum variance portfolio invests 47% in the utility sector while the corresponding weight in the market cap-weighted portfolio is 9% and the corresponding weight in the equally-weighted portfolio is 15%.

Such concentration may correspond to the requirements of investors who – for whatever reason – want to bet on low volatility stocks. Whether or not such a bet is promising depends on the properties of such stocks. For example, it is sometimes argued that the highest volatility stocks come with lower returns in a short-term perspective (Ang *et al*, 2006), though this finding has been shown to lack robustness and does not hold for long holding periods where holding high volatility stocks is actually rewarded with higher returns (see Fu, 2009, and Huang *et al*, 2010, and the article, "Is there a risk/return trade-off across stocks?" on page 2 of this supplement). Researchers have also studied the general risk properties of low volatility stocks and have found that low volatility stocks – while they have low risk in terms of volatility – may display high extreme risks (Boyer, Mitton and Vorking, 2010; Chen, Hong and Stein, 2001) or unfavourable exposures to shocks in aggregate market volatility (Barinov, 2010).

Irrespective of whether low-volatility stocks are attractive or unattractive, it is clear that a minimum variance strategy leads to poorly diversified portfolios and does not fully exploit correlations. In the end, for investors that have made the decision to move away from market cap-weighted portfolios, one can reasonably question whether replacing the concentration in the largest capitalisation stocks inherent in cap-weighting with concentration in the lowest volatility stocks addresses their concerns.

Researchers have recognised this limitation of minimum volatility portfolio construction, and have proposed various ways to remedy the concentration of optimised portfolios in low-volatility stocks.

The most straightforward solution to any concentration problem is to impose weight constraints. Imposing lower and/or upper bounds on weights provides quite rigid constraints which leave reduced room for optimisation, but can help to obtain more reasonable portfolios. Recently, more flexible weight constraints have been proposed by DeMiguel, Garlappi, Nogales, and Uppal (2009), who use so-called 'norm constraints'. Such constraints put limits on the overall amount of concentration in the portfolio (eg, on the sum of squares of portfolio weights) rather than limiting the weight of each stock in the portfolio, thus leaving more room for the optimiser while avoiding concentration overall. An alternative to the use of weight constraints is to avoid making a difference between stocks based on their volatilities. Christofferson *et al* (2010) minimise volatility with the assumption that volatilities are identical across stocks. Hence the only difference across stocks that the optimiser then takes into account are differences in correlations. The minimum volatility portfolio under this assumption will thus be unaffected by a concentration in low volatility stocks, at the cost of an implicit assumption that all volatilities are equal.

A different way of avoiding concentration in

1. Concentration of minimum volatility portfolio in low-volatility stocks

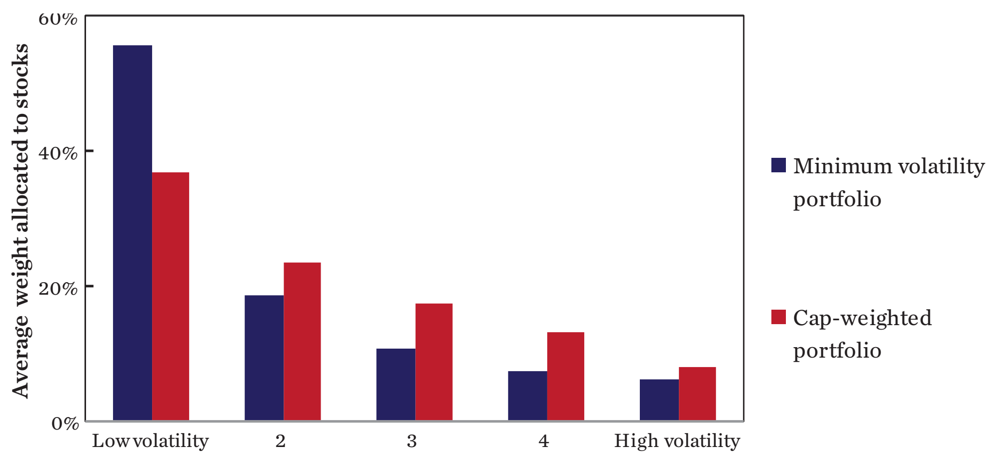

Each month, stocks in the S&P 500 universe are sorted into quintiles by their past volatility (12 months). The graph shows the average weight allocated to the stocks belonging to each of these quintiles for two different portfolios: minimum volatility portfolio and cap-weighted portfolio. The analysis period runs from January 1959 until December 2008 (weekly returns).

low volatility stocks is to penalise these stocks in the portfolio optimisation. Amenc *et al* (2010) construct efficient indices and benchmarks by maximising the Sharpe ratio, rather than minimising volatility. The approach effectively penalises low-risk stocks through assuming a low expected return. This penalty on the expected returns side counterbalances the attractiveness of low-risk stocks from a risk perspective. To ensure parsimony and robustness, they group stocks by their total downside risk (in particular a stock's semi-deviation, which incorporates higher moments) and distinguish stocks based on their riskiness only across groups rather than on a stock-by-stock basis.

What the abovementioned approaches have in common and what distinguishes them from a pure minimum variance approach is that they avoid concentration in low-risk stocks and try to exploit more fully the information available on correlations in the relevant equity universe. Using the covariance matrix solely to minimise volatility tends to result in concentrated portfolios, dominated by the low-volatility entries on the diagonal of the covariance matrix. Penalising the low-volatility stocks that add no diversification benefit either directly or indirectly through constraints and implicit assumptions is meant to exploit better the covariance (off-diagonal) entries. In fact the important insight that carries through from Markowitz's early work on efficient diversification to recent work on portfolio construction is that, while portfolio returns are a simple weighted average of component returns, the portfolio volatility is not just the weighted average of the stock's volatilities. For a given level of return, one can lower the risk by intelligently combining stocks according to their correlations. This diversification principle is not what drives minimum variance portfolio construction. Since minimum variance portfolios have only the objective of lowering risk, rather than aiming to optimise the risk/reward ratio, minimum variance portfolio optimisation leads to a pronounced concentration in low-volatility stocks at the expense of exploiting correlation properties. Such portfolios are therefore suitable for investors who wish to load on low-risk or 'defensive' stocks, while alternative approaches may be relevant for investors who want to manage risk and reward properties through combining both low-risk and high-risk stocks in a broadly diversified portfolio.

References

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

Amenc, N., F. Goltz, L. Martellini and P. Retkowsky. 2010. Efficient Indexation: An Alternative to Cap-Weighted Indices, EDHEC-Risk Institute Publication.

Ang, A., R. J. Hodrick, Y. Xing and X. Zhang. 2006. The cross-section of volatility and expected returns. *Journal of Finance* 61 (1): 259–99.

Boyer, B., T. Mitton and K. Vorkink. 2010. Expected idiosyncratic skewness. *Review of Financial Studies* 23 (1): 169–202.

Chan, L. K. C., Jason Karceski and Josef Lakonishok. 1999. On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model, *Review of Financial Studies*, 12 (5): 937–974.

Chen, J., H. Hong and J. C. Stein. 2001. Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices. *Journal of Financial Economics* 61 (3): 345–81.

Christoffersen, Peter F., Vihang Errunza, Kris Jacobs and Xisong Jin. 2010. Is the Potential for International Diversification Disappearing?, working paper.

Clarke, R., H. de Silva and S. Thorley. 2011. Minimum-Variance Portfolio Composition. *Journal of Portfolio Management*, 37 (2): 31–45.

DeMiguel, Victor, Lorenzo Garlappi, Javier Nogales and Raman Uppal. 2009. Optimal versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy? *Review of Financial Studies* 22 (5): 1915–1953.

DeMiguel, Victor, Lorenzo Garlappi, Javier Nogales and Raman Uppal. 2009. A Generalized Approach to Portfolio Optimization: Improving Performance By Constraining Portfolio Norms. *Management Science*, forthcoming.

Fu, F. 2009. Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics* 91 (1): 24–37.

Huang, W., Q. Liu, S. G. Rhee and L. Zhang. 2010a. Return reversals, idiosyncratic risk, and expected returns. *Review of Financial Studies* 23 (1): 147–68.

Optimal hedge fund allocation with improved estimates for coskewness and cokurtosis parameters

Lionel Martellini, Professor of Finance, EDHEC Business School and Scientific Director, EDHEC-Risk Institute

Given that hedge fund returns are not distributed in a Gaussian manner, in the classic bell curve distribution around the mean, mean-variance optimisation techniques, which would be sub-optimal and impact negatively on the investor's welfare, need to be replaced by optimisation procedures that incorporate higher-order moments and comoments. As such, optimal portfolio decisions relating to hedge fund style allocation require estimates not only for covariance parameters, but also for coskewness and cokurtosis parameters. This is a considerable challenge that significantly augments the dimensionality issue that already exists with mean-variance analysis.

In a recent research paper that is part of the Advanced Modelling for Alternative Investments research chair at EDHEC-Risk Institute, supported by the Prime Brokerage Group at Newedge, we present an application of enhanced estimators for higher-order comoment parameters, introduced by Martellini and Ziemann (2010), in the context of hedge fund portfolio optimisation. We find that using these improved estimates leads to a considerable improvement for investors in hedge funds. We also find that it is only when enhanced estimators are used that portfolio selection with higher-order moments is consistently superior to mean-variance analysis from an out-of-sample perspective.

Diversification benefits

One of the principle reasons why asset owners generally are willing to include hedge funds in their portfolios is that they expect to achieve diversification benefits with respect to other existing investment possibilities.

Many academics (see for example Terhaar *et al*, 2002) have stressed that mixing hedge funds with traditional assets leads to a reduction in the volatility of the traditional portfolio. If they wish to capitalise fully on the benefits of diversification in a top-down approach, investors or (funds of hedge funds) managers must be able to rely on robust techniques for optimising portfolios that include hedge funds. Standard mean-variance portfolio selection techniques are known to suffer from a number of shortcomings, and the problems are exacerbated in the presence of hedge funds. First, because hedge fund returns are not normally distributed (see for example Brooks and Kat, 2002), a mean-variance optimisation would be severely ill-adapted, except in the case of an investor who possesses quadratic preferences. For example, it can be shown, through a statistical model integrating fatter tails than those of the normal distribution, that minimising the second order moment (volatility) can be accompanied by a significant increase in extreme risks (Sornette *et al*, 2000).

This finding is confirmed in Amin and Kat (2003), where the authors present empirical evidence that low volatility is generally obtained at the cost of lower skewness and higher kurtosis. As a result, as stressed in Cremers *et al* (2005), in the presence of asymmetric and/or fat-tailed return distribution functions, the use of mean-variance analysis can potentially lead to a significant loss of utility for investors.

Extending portfolio optimisation techniques

As a consequence of the shortcomings of mean-variance optimisation, many attempts have been made to account for the specific risk features of hedge funds in a better way and to extend portfolio optimisation techniques in order to account for the presence of fat-tailed distributions, mostly by introducing some risk objective (eg, value at risk as in Favre and Galeano, 2002, or conditional value at risk as in De Souza and Gokcan, 2004, and Agarwal and Naik, 2004), that is more general than volatility, integrating the presence of non-trivial higher moments in asset returns. ►

In the presence of non-normally distributed asset returns, optimal portfolio selection techniques require not only estimates for variance-covariance parameters, but also estimates for higher-order moments and comoments of the return distribution. However, the need to estimate coskewness and cokurtosis parameters considerably exacerbates the dimensionality problem, which is already a serious concern in the context of covariance matrix estimation. This concern is particularly acute in the hedge fund universe, where data is scarce, with a short history and low frequency, and where a number of performance biases are present (see for example Fung and Hsieh, 1997, 2000, 2002). In this context, given the dramatic increase in dimensionality involved, one might wonder whether portfolio selection techniques that rely on higher-order moments can efficiently be implemented at all in realistic situations.

In a recent paper, Martellini and Ziemann (2010) shed some light on this question by introducing improved estimators for the coskewness and cokurtosis parameters. They extend to the skewness and kurtosis dimensions several improved estimates that had been proposed for the covariance matrix, including most notably the factor-based approach (Sharpe, 1963), the constant correlation approach (Elton and Gruber, 1973) and the statistical shrinkage approach (Ledoit and Wolf, 2004). In an empirical analysis based on US large-cap stock returns, Martellini and Ziemann (2010) subsequently find that the use of these enhanced estimates generates a significant improvement in investors' welfare.

In our new research, we complement these results by providing the first application of improved estimators for higher-order comoment parameters in the context of hedge fund portfolio optimisation. We find that the use of these enhanced estimates generates a significant improvement for investors in hedge funds. We also find that it is only when improved estimators are used that portfolio selection with higher-order moments dominates mean-variance analysis from an out-of-sample perspective. More specifically, we construct portfolios based on various hedge fund style indices using a 4th order approximation of expected CARA utility, using shrinkage estimators so as to alleviate the concern over robustness of purely sample-based estimates. We find that the use of improved estimators leads to substantial increases in the investor's utility as compared to using sample estimators. On the other hand, we find that extending the objective function to encompass higher-order moments of hedge fund return distribution can lead to value being destroyed, as opposed to added, when sample-based estimators are used. In our research paper, we introduce the improved estimators for hedge fund return covariance, coskewness and cokurtosis parameters and then present our empirical analysis.

Conclusion

This article discusses an application of the improved estimators for higher-order comoment parameters, introduced by Martellini and Ziemann (2010), in the context of hedge fund portfolio optimisation. We have found in recent research that the use of these enhanced estimates generates considerable improvements for investors in hedge funds. We also find that it is only when improved estimators are used that portfolio selection with higher-order moments is consistently superior to mean-variance analysis from an out-of-sample perspective. Our results have important potential implications for hedge fund investors and fund of hedge funds managers who routinely use portfolio optimisation that incorporates higher moments without a formal analysis of the induced increase in parameter uncertainty and related lack of robustness of the results.

The research from which this article was drawn was produced as part of the research chair on Advanced Modelling for Alternative Investments at EDHEC-Risk Institute sponsored by the Prime Brokerage Group at Newedge.

References

Agarwal, V., and N. Naik. 2004. Risks and Portfolio Decisions Involving Hedge Funds. *Review of Financial Studies*, 17 (1): 63–98.
Amin, G., and H. Kat. 2003. Stocks, bonds and hedge funds: not a free lunch! *Journal of Portfolio Management*, 29 (4): 113–120.
Brooks, C., and H. Kat. 2002. The statistical properties of hedge fund returns and their implications for investors. *Journal of Alternative Investments*, 5, (2): 26–44.
Cremers, J. H., M. Kritzman and S. Page. 2005. Optimal hedge fund allocations: do higher moments matter? *Journal of Portfolio Management*, 31 (3): 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): 7–21.
Elton, E., and M. Gruber. 1973. Estimating the dependence structure of share prices – Implications for portfolio selection. *Journal of Finance*, 28 (5): 1203–1232.
Favre, L., and J.-A. Galeano. 2002. Mean-modified Value-at-Risk optimization with hedge funds. *Journal of Alternative Investments*, 5 (2): 2–21.
Fung, W., and D. A. Hsieh. 1997. Empirical characteristics of dynamic trading strategies: the case of hedge funds. *Review of Financial Studies*, 10 (2): 275–302.
Fung, W., and D. A. Hsieh. 2000. Performance characteristics of hedge funds and commodity funds: natural versus spurious biases. *Journal of Financial and Quantitative Analysis*, 35 (3): 291–307.
Fung, W., and D. A. Hsieh. 2002. Benchmark of hedge fund performance, information content and measurement biases. *Financial Analysts Journal*, 58 (1): 22–34.
Ledoit, O., and M. Wolf. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10 (5): 603–621.
Martellini, L., and V. Ziemann. 2010. Improved estimates of higher-order comoments and implications for portfolio selection. *Review of Financial Studies*, 23 (4): 1467–1502.
Sharpe, W. 1963. A simplified model for portfolio analysis. *Management Science*, 9 (2): 277–293.
Sornette, D., J.V. Andersen and P. Simonetti. 2000. Portfolio theory for "fat tails". *International Journal of Theoretical and Applied Finance*, 3 (3): 523–535.
Terhaar, K., R. Staub and B. Singer. 2003. An appropriate policy allocation for alternative investments. *Journal of Portfolio Management*, 29 (3): 101–110.

Value and momentum effects across exchange-traded funds

Combining trading strategies with risk management

Elie Charbit, EDHEC Business School; Jean-René Giraud, CEO, Koris International and Research Associate, EDHEC-Risk Institute; Felix Goltz, Head of Applied Research, EDHEC-Risk Institute; and Lin Tang, Research Engineer, EDHEC-Risk Institute

Investors are willing to take on risk only if they are compensated with greater expected reward. There are many types of equity exposure that can lead to risk premia. Value and momentum are among the most robust return drivers in the cross-section of expected equity returns. Value effects often refer to the fact that stocks with low price/earnings ratios or high dividend yields tend to outperform stocks with high price/earnings ratios or low dividend yields, while momentum effects usually refer to the fact that stocks with high returns in the past yield high returns in the future. There is ample evidence in the academic literature on these two strategies (eg, Jegadeesh and Timan, 1993; Graham and Dodd, 1934; and Fama and French, 1992).

Although exposure to value and momentum effects is expected to yield attractive performance over the long run, in moving away from the market factor and trying to exploit value or momentum effects, investors' portfolios tend to become more concentrated, increasing drawdown risk. In this article, we use exchange-traded funds (ETFs), which are a liquid investment medium, to apply value or momentum trading strategies across sectors in a dynamic core-satellite (DCS) portfolio to assess the risk-control benefits of the DCS portfolios.

The next section describes our method and data choices. Section three discusses our findings. A final section summarises our conclusions.

Data and methodology

In this section, we describe both the data we use to build our portfolios and the dynamic core-satellite strategy.

All of our data is on a monthly basis and covers the period from 31 January 1989 to 31 December 2009. Our investment universe is limited to Europe equities, in particular the STOXX Europe 600 and its sector sub-indices.

Methodology
The value factor is computed from the 15 STOXX Europe 600 sector sub-indices we initially select. We compute the aggregate book-to-market (BM) ratio of each index and then rank these BM ratios from highest to lowest. Every month, we go long the five sectors with the highest BM ratios in the previous month. We can therefore create a long-only equally-weighted value portfolio.

We then build a momentum portfolio. To do so, we calculate the cumulative returns over the previous 12 months up to two months earlier (as in much of the literature, observations of the most recent returns are discarded to prevent short-term reversal effects). Once we have all the sectors' 12-2 month cumulative returns, we rank them from highest to lowest and go long the five sectors exhibiting the highest cumulative returns in every month.

We apply the dynamic core-satellite (DCS) framework developed by Amenc *et al* (2004) to build the risk-control strategy with either the long-only value strategy or the long-only momentum strategy as a performance-seeking satellite portfolio [1] (see Amenc *et al*, 2010). The core portfolio in each case is a euro cash investment comparable to a money market ETF.

The DCS approach will combine the core and satellite portfolios such that we generate a participation in the upside potential of the satellite (ie, the momentum or value strategy's returns) while ensuring that the investment value respects a floor level which in particular limits the downside risk at a maximum level of 10%. Our examples show that dynamic asset allocation techniques make it possible to better address investor concerns over drawdown risk.

The DCS portfolio is constructed by first specifying a maximum drawdown floor equal to 10% and a performance cap (investment goal) set at the wealth achieved by compounding twice the cash rate over the 20-year period[2]. The maximum allocation to the satellite is set at 50%. Now we turn to the analysis of the performance of these portfolios

Access to value and momentum premia with downside risk control

Amenc *et al* (2010) conclude, with a different dataset and over a different time period, that the DCS can offer better returns and at the same time limit downside risk. We now look into whether the DCS could be used to gain access to value and momentum premia. Keeping downside risk under control is, of course, even more important in strategies in which the investment universe is reduced to concentrate the portfolio in stocks with high exposure to value or momentum.

Figure 1 summarises the performances of value, momentum and the market index portfolio. We find higher returns and also higher volatilities for both value and momentum portfolios compared to the market index STOXX 600. Value and momentum portfolios however generated higher Sharpe ratios in our sample.

1. Risk and return statistics

	Return	Volatility	Sharpe ratio
STOXX 600	8.37%	16.43%	0.39
Value portfolio	10.62%	20.24%	0.43
Momentum portfolio	10.16%	17.38%	0.47

From 31 January 1989 to 31 December 2009

Though historically there is a higher risk/return ratio for value and momentum portfolios, by design, such strategies are more highly concentrated, as they are built on a handful of sectors (five, in our case). The result is greater exposure to downside risk (see figure 2). The maximum drawdown of the value strategy is nearly two-thirds of peak portfolio wealth. The maximum drawdown of the momentum strategy is also greater than that of the STOXX 600. The shortfall probability is indicated as the probability of experiencing a maximum drawdown that breaches the 10% limit. Although it has the largest maximum drawdown, the value strategy, interestingly, has the lowest shortfall probability, whereas the momentum strategy, despite its smaller maximum drawdown, has the highest shortfall probability. This finding implies that the extreme losses of the value strategy are larger but less frequent than those of the momentum strategy. The higher downside risk of the strategies (value and momentum) are also shown by the one-month 99% value at risk, which is 15% for value and 14% for momentum but 12.6% for the STOXX 600. As for the trailing returns, the value strategy has consistently higher extreme losses in both the short term and the long term than the broad market index. The momentum strategy, conversely, has higher short-term but lower long-term extreme drawdown than the broad market index. The Calmar ratios[3] for these three portfolios are roughly the same, indicating that the premia for bearing one additional percentage point of drawdown risk are similar for all three portfolios. On the whole, the value strategy is exposed to greater downside risk as a result of its more highly concentrated portfolios. The momentum strategy is also riskier in the short term, but, for the sample we study here, it is less prone to drawdown over the longer term.

"On the whole, the value strategy is exposed to greater downside risk as a result of its more highly concentrated portfolios. The momentum strategy is also riskier in the short term, but, for the sample we study here, it is less prone to drawdown over the longer term"

Our finding implies that though value and momentum offer attractive returns, portfolio construction should also take into account the downside risks they are exposed to. Now we explore how the DCS approach could reduce the downside risk exposure.

Figure 3 shows the cumulative returns of the value DCS (bold blue line). To better understand the results, we also show the core, the satellite, the floor and the goal.

It is clear that the DCS value strategy performs smoothly throughout the entire period. It reduces the fluctuation of the value satellite and limits downside risk. This relatively high return compared to the core portfolio suggests

2. Summary of downside risk exposure for satellite and DCS portfolios

	Maximum drawdown	Shortfall probability	99% VAR over a month	3-month trailing return 1st percentile	12-month trailing return 1st percentile	Calmar ratio
STOXX 600	54.34%	43.75%	12.57%	–21.91%	–42.06%	0.15
Value portfolio	65.13%	36.25%	15.07%	–25.88%	–49.99%	0.16
Momentum portfolio	54.98%	50.00%	14.08%	–24.12%	–39.76%	0.18

From 31 January 1989 to 31 December 2009

3. Evolution of the value DCS and its parameters

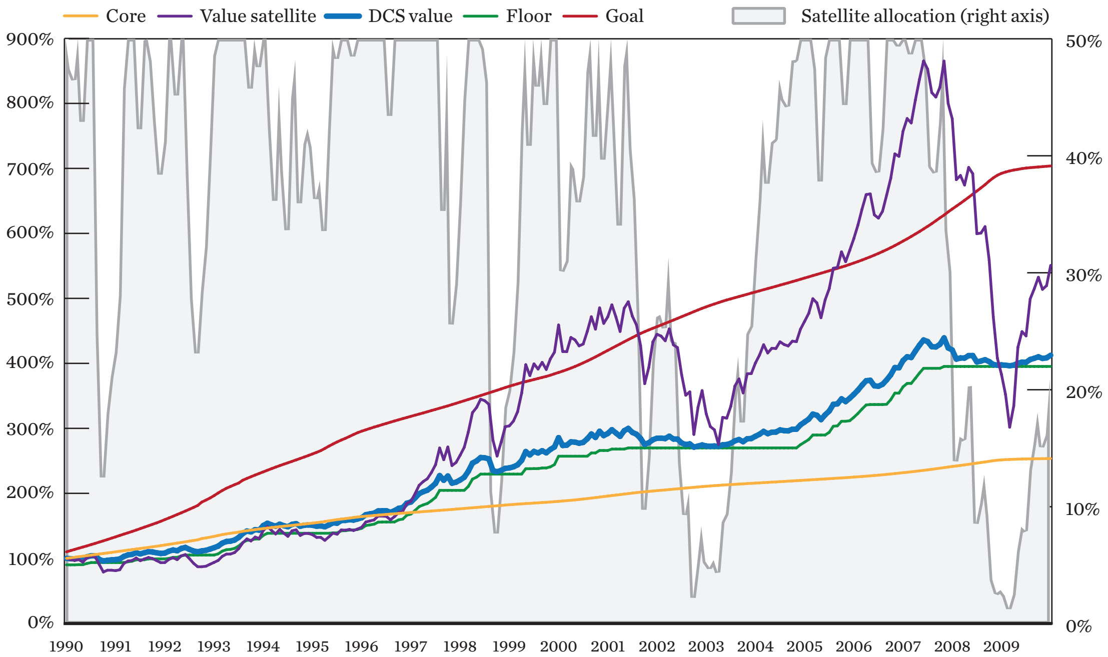

1 For comparison purposes, the core is the bond ETF on the PIBOR rate and the satellite is the ETF on the STOXX Europe 600 index.
2 This cash rate used is the France PIBOR one-month interest rate from 1989 to 1999 and the EONIA from 2000 to 2009. In this case, the goal is 700% of the initial wealth at the end of the 20-year horizon.
2 The Calmar ratio, which is calculated by dividing the annual return by the maximum drawdown (Young, 1991), indicates the premium for bearing one additional percentage point of drawdown risk. A higher Calmar ratio implies better downside risk-adjusted performance.

4. Evolution of the momentum DCS and its parameters

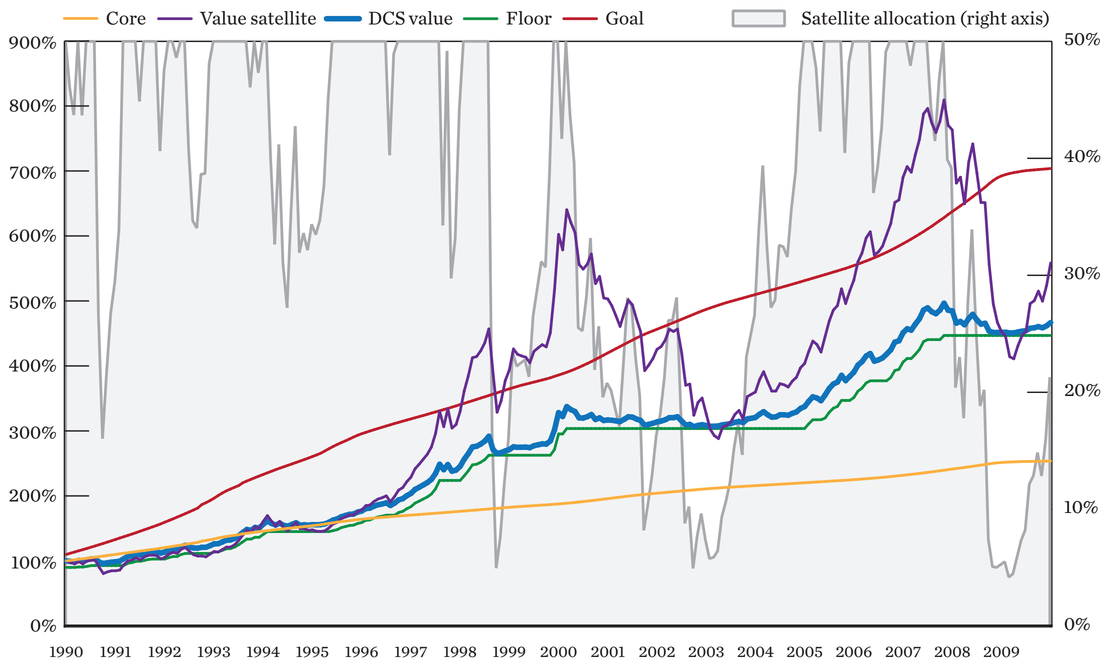

however that the value DCS may help investors gain access to the value premium and, at the same time, limit the huge drawdown that may otherwise afflict the value portfolio.

Similarly, figure 4 shows the cumulative returns of the DCS momentum strategy. As usual, the momentum DCS (bold blue line) is unaffected by market downturns but profits from upturns.

Now we show the risk/return measures of the DCS portfolios in one table (figure 5).

Compared to the results shown in figure 2, we find that DCS approaches significantly reduce the extreme risk exposures while maintaining relatively high returns. The volatilities drop to about 6% from 20%. Both DCS portfolios respect the 10% limit on maximum drawdown and have a one-month 99% VaR of less than 5%. The first percentiles of trailing returns are also significantly reduced. In addition, both the short-term and the long-term trailing returns are reduced to around –7 to –6%. Higher Calmar ratios mean that, for about the same drawdown risk, portfolios achieve higher returns.

Dynamic versus static

We also build the fixed-mix portfolio with cash and the satellite portfolios (ie, value and momentum portfolios). From a risk management point of view, we set the constant weight of the fixed-mix strategy so that the historical maximum drawdown (MDD) is equal to the preservation objective of the simple DCS, which results in 10% maximum peak-to-valley drawdown. Thus, by authorising the fixed-mix strategy with value and momentum satellites to lose up to 10% historically, the fixed-mix portfolios have to allocate more than 90% to cash. Figure 6 summarises the comparison of the risk/return profiles between static fixed-mix strategies and DCS approaches. Though respecting the 10% MMD limit, DCS portfolios in general have higher volatilities because the allocation to the satellite could go up to 50% of the entire portfolio as long as the total wealth is above the floor. Hence, in the short term, DCS strategies are more risky. However, DCS strategies could also deliver higher returns compared to the fixed-mix strategies. Higher Calmar ratios mean that for bearing an additional percentage of drawdown risk, DCS strategies could achieve higher returns. In addition, DCS portfolios are less risky in the long term; in particular, the 1st percentile of 12-month trailing returns is lower.

Conclusion

In this article, we analyse dynamic core-satellite strategies with exposure to the value and momentum strategies. We find that these investment strategies alone could achieve higher returns but are exposed to high extreme risk because they consist of equity portfolios that are concentrated in the sectors with the highest value or momentum exposure. Combining these strategies with the DCS approach, however, dopes portfolio returns and, at the same time, keeps downside risk in check. In addition, by comparing DCS approach with the ex-post fixed-mix strategies, we find that DCS portfolios outperform the static strategies as the DCS approach allows access to the upside potential of the satellite portfolios while fixed-mix strategies forfeit the upside by limiting the majority of the allocation to the cash. Exchange-traded funds on sectors rather than on stocks can be used to put these strategies into effect; ETFs would also greatly facilitate the shifts – required by dynamic strategies – from core to satellite.

This research was conducted as part of the second year of research for the Core-Satellite and ETF Investment research chair at EDHEC-Risk Institute, in partnership with Amundi ETF

References

Amenc, N., F. Goltz and A. Grigoriu. 2010. Risk control through dynamic core-satellite portfolios of ETFs: Applications to absolute return funds and tactical asset allocation. *Journal of Alternative Investments* 13 (2): 47–57.

Amenc, N., P. Malaise and L. Martellini. 2004. Revisiting core-satellite investing – A dynamic model of relative risk management. *Journal of Portfolio Management* 31 (1): 64–75.

Fama, E., and K. French. 1992. The cross-section of expected stock returns. *Journal of Finance* 47 (2): 427–65.

Graham, B., and D. Dodd. 1934. *Security analysis.* New York: McGraw Hill.

Jegadeesh, N., and S. Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48: 65–91.

Young, T. 1991. Calmar ratio: A smoother tool. *Futures Magazine* (1 October).

5. Summary of risk/return measures of DCS portfolios

	Return	Volatility	Sharpe ratio*	Maximum drawdown	Shortfall probability	99% VAR over a month	3-month trailing return 1st percentile	12-month trailing return 1st percentile	Calmar ratio
DCS with value satellite	7.05%	6.41%	0.79	9.76%	0	4.69%	–7.42%	–6.24%	0.72
DCS with momentum satellite	7.70%	6.24%	0.91	9.25%	0	4.09%	–5.97%	–6.81%	0.83

* The annual risk-free rate is assumed to be 2%. From 31 January 1989 to 31 December 2009

6. Comparison between static and dynamic allocations

	Return	Volatility	Sharpe ratio*	Maximum drawdown	99% VAR over a month	3-month trailing return 1st percentile	12-month trailing return 1st percentile	Calmar ratio
DCS allocation with value satellite	7.05%	6.41%	0.79	9.76%	4.69%	–7.42%	–6.24%	0.72
Static allocation with value satellite	5.09%	2.47%	1.25	10.00%	1.95%	–3.32%	–6.46%	0.51
DCS allocation with momentum satellite	7.70%	6.41%	0.91	9.25%	4.09%	–5.97%	–6.81%	0.83
Static allocation with momentum satellite	5.31%	3.39%	0.98	10.00%	2.82%	–4.73%	–7.68%	0.53

* The annual risk-free rate is assumed to be 2%. From 31 January 1989 to 31 December 2009