Intelligent Hedge Fund Investing: An Introduction

Barry Schachter

By just about any metric, the assets allocated to hedge funds continue to grow rapidly. At the time this volume is being published, data for the calendar year 2003 is just becoming available. The approximate number of extant hedge funds is 6,700 (including 1,700 funds of funds), and the estimated total US dollar value of the assets of these funds is US$725 billion (or more). The estimates provided here are according to TASS as reported by Reuters (2004). More interesting, perhaps, is the (changing) composition of the investor base. Traditionally, the majority of invested capital came from high net worth individuals (with greater than US$1 million in financial assets). Merrill Lynch Cap Gemini Ernst & Young (2003) estimated that 60% of hedge fund investments came from such individuals in 2002. That percentage has been declining as institutional investments have been growing rapidly. Endowments have become a major source of funding. The National Association of College & University Business Officers (2004) estimated that, in 2003, about 13% of investment assets were allocated to hedge funds. A breakdown by size of endowment is shown in Table 1. Most recently, pension funds have begun to devote more attention to hedge funds. Surveying more than 1,000 pension funds and endowments in 2003, Greenwich Associates (2004) found a continued increase in planned allocations to hedge funds, with pension funds significantly increasing their targeted allocations from under 1% of assets

Table 1 Fiscal year 2003 average allocations (in pct)

Investment pool assets	Equity	Hedge funds	All other
Greater than US$1 billion	44.8	19.9	35.3
US$501 million–US$1 billion	54.4	13.4	32.2
US$101 million–US$500 million	56.5	8.3	35.2
US$51 million–US$100 million	58.7	4.3	37
US$26 million–US$50 million	60.2	4.2	35.6
Less than US$25 million	57.0	1.6	41.4

Source: National Association of College & University Business Officers 2003 Endowment Study of 705 institutions. Used with permission

to about 5% of assets. Greenwich Associates estimated that, in US dollar terms, this represents over US$250 billion in allocations.

It's a risk manager's job to worry about potential danger signs. If a sudden, massive increase in demand for maple syrup were to materialise, the response in supply in the short term would be limited by production capabilities, and only by cutting corners or diluting quality, could significant new quantities be produced. Although there may be a better analogy, this one serves well enough. The production of "alpha" from hedge funds in the short run is not highly elastic. Many strategies don't scale well and there is no latent army of new, quality hedge fund managers. Furthermore, as the universe of hedge fund investors expands (the latest wave due to pension fund managers), lack of familiarity with the asset class becomes more prevalent.

For both these reasons, investors face obstacles to intelligent hedge fund investing which can only be surmounted with deeper understanding. One such obstacle is a lack of understanding of the characteristics of hedge funds that are different from other asset classes. The difference between the focus on hedge funds on absolute return, and the relative return focus that is ubiquitous in the mutual fund industry, is significant. Trading on the short side in addition to the long side is another important attribute of hedge funds. The nature of manager compensation and the fiduciary relationship is another. The regulatory environment is also special to hedge funds. The diversity of strategies also belongs on this list. All these differences have implications for properly evaluating the asset class for asset allocation.

The feasibility of properly evaluating the asset class is another obstacle. Hedge fund data, even those provided privately to prospective investors, have the potential to lead investors to incorrect decisions. For example, data can suffer from biases from return smoothing attributable to positions in illiquid securities. Data can embed biases from the process used to accumulate the data history. Even clean data can exhibit difficult properties. For example, departures from normality have been widely documented in hedge fund returns, meaning that measures such as Sharpe Ratios and tools such as Markowitz style portfolio allocation models may not work well.

The chapters in this volume focus on the questions that investors should be asking about hedge funds. The first goal of this volume is to illuminate the hedge fund landscape (so to speak), and thereby to identify the pitfalls to sound investment decisions. The second goal is to suggest a path (or paths) to help investors reach those decisions (such as performance evaluation and asset allocation). It is important to note that the research presented here is part of an ongoing process of discovery. Future research, building on the insights offered here, will continue to deepen our understanding of hedge funds.

REFERENCES

Greenwich Associates, 2004, "Pension Windfall Transforms Hedge Fund Industry", Press Release, Greenwich, CT: 22 January.

Merrill Lynch Cap Gemini Ernst & Young, "World Wealth Report 2003", URL: http://www.GloriaMundi.org/picsresources/rb-wwr2003.pdf.

National Association of College & University Business Officers, 2004, "2003 Endowment Study", Mimeo.

Reuters, 2004, "Hedge Fund Assets Up Estimated US$60 Billion in 2003", 26 January.

Introduction

Barry Schachter

In the movie classic *The Wizard of Oz*, Dorothy's first epiphany after touching down in Munchkin land led her to aver, "I've a feeling we're not in Kansas anymore". For hedge fund investors (and potential investors), the equivalent realisation is, we're not in mutual fund land anymore. The chapters in this section discuss some of the issues unique to the hedge fund landscape and the current state of thinking about those issues.

Kat, in his contribution, gives us a piquant illustration of the "GIGO Principle", or garbage in – garbage out, applied to hedge fund investments, where inappropriate use of the tools of portfolio analysis, such as the Sharpe Ratio, can lead to misleading and costly decisions. He first explains why hedge fund returns data are fundamentally different from data on stocks and bonds, embedding various biases in their reported returns and volatility. A significant research effort over the last several years has focused on identification and quantification of these data biases found in individual hedge fund returns and in hedge fund index returns.

Looking at the properties of hedge fund return data, Kat points out that departures from normality in hedge fund returns are the norm (so to speak), not the exception. Hedge fund returns exhibit significant negative skewness and excess kurtosis, which are both "bads" to a risk-averse investor. He goes on to show that these "bads" can be correlated with the mean and standard deviation of returns, such that apparently more attractive returns can also carry penalties in terms of skewness and kurtosis.

Kat goes on to argue (compellingly) that, in light of data biases and departures from normality, the usual application of the standard tools of portfolio analysis, such as the Sharpe Ratio and mean-variance portfolio optimisation, does not properly account for these "bads". As a result, the mechanical application of these tools can yield incorrect conclusions about the attractiveness of hedge fund investments. Some portion of the apparent "excess" returns earned by hedge fund managers according to standard measurements is not "excess" at all, but in reality, is just compensation for taking on the risks that result in the observed negative skewness and kurtosis. He states that new tools need to be developed to account for these complexities in evaluating hedge funds.

In the end, Kat issues a challenge to hedge fund investors to use common sense and develop a deep, if qualitative, understanding of hedge funds as part of their asset allocation regimen. Sometimes, the best advice is the simplest advice.

Arabadjiev does us all a great favour by choosing the word "understanding" for the title of his contribution. He explains the importance to investors of appreciating the qualitative factors that circumscribe the quantitative task of evaluating hedge fund risks. It doesn't matter if there is a well-established theory for the optimisation of asset allocation, when an investor does not have access to the essential asset return information necessary to apply the theory. Arabadjiev takes us through the frequently discussed risks of hedge fund investing, investment risk, operational risk, liquidity risk, and business risk, looking behind the standard methodologies. We find that understanding is the result of a process that takes as inputs the parameter estimates from models, not a process that ends with the production of those estimates.

From this perspective, Arabadjiev provides his thoughts on achieving understanding (if, alas, not inner peace). For example, he notes that even if an investor does not have the systems and resources necessary to effectively use position level information from a hedge fund to measure risk and performance, the position information is still valuable in service of the investor's need to monitor strategy drift.

When I meet with hedge fund investors, I emphasise the importance of understanding both the risk management culture and the extent of risk management processes over the importance of the level of fund VAR or the geographic distribution of the fund's risk at a particular date. The result of this understanding is a vastly improved employment of the tools of quantitative risk analysis and the improved use of the results of those quantitative tools.

Recently, we have observed the creation and marketing of a proliferation of hedge fund indices. This has paralleled the increased interest in hedge fund investment. Amenc, Martellini, and Vaissié (AMV) examine these indices and show us why all hedge fund indices are not created equal. AMV identify large differences among indices in reported performance for the same time period. They also find that correlations with various benchmarks vary wildly among indices. As a result, identifying the "best" index presents a significant problem for assessing individual hedge fund performance and for asset allocation decisions. AMV document how the differences among indices can be traced to differences in their construction methods, in their constituent funds, and in the impact of well-known biases (discussed in several chapters in this volume) that affect measured returns in hedge fund indices.

AMV suggest that the best index may in fact be a combination of several individual indices. That is, it may be best to employ multiple indices to construct a super strategy-based index. They implement this approach using Principal Component Analysis. In this way, they extract the primary driver of variability across competing indices for a given strategy. In the process, they argue, the resulting super index is likely to have neutralised the impact of biases in measured returns, at least to the extent that those biases do not affect various indices in a highly correlated fashion.

For a trader, or an investor, understanding the potential effects of illiquidity on an investment is not an afterthought but rather an integral part of the trading decision. Illiquidity is a bit like the weather. We complain about it loudly, because we can't really do anything about it. Till, in her chapter, brings us up to date on what the talk is all about in a comprehensive tour of the subject. Illiquidity is a complicated and changing thing. Despite significant advances in our thinking about illiquidity in the last few years, this measurement problem remains a significant challenge in risk management.

Till tells us that behavioral effects related to illiquidity play a greater part than we might have thought. Traders commonly struggle with challenges such as, "I don't want to have to move the price by a point if I want to get out", or "everyone is a buyer out there, but me". Trying to measure liquidity risk is not just about measuring the volatility of the bid–ask spread. Till explains that in one important case, at least an argument can be made that the behavioural implications may cause an investor to favour an illiquid investment (other things being equal, of course).

Kat has asked (and he is not alone in this) whether financial markets are really so far short of strong-form market efficiency that hedge fund managers can generate enough "alpha" to be able to charge fees of 2-and-20 (or more) and still have investors conclude that investment is expected utility enhancing. He suspects that, in fact, what is happening is that hedge fund returns in part reflect compensation for bearing illiquidity (or the obverse, supplying liquidity). This line of thinking suggests that investors need to ask, not only what are the costs and benefits of illiquidity associated with hedge fund investing, but also, is the expected return to that investing sufficient compensation for both the price risk being taken and the illiquidity risks being borne? Till's review will help clarify current thinking on these issues.

Ross focuses on a problem unique to a large (institutional) investor who wishes to allocate risk to hedge funds. For such an investor, the absolute size of the investment is large, even if it represents a small fraction of the investor's portfolio. Yet, the act of allocating large sums poses risks if hedge funds have limited capacity in generating alpha. Groucho Marx is supposed to have said, "I don't care to belong to any club that will have me as a member". Ross asks, should an institutional investor be leery of the fund that is willing to accept the large sums the investor needs to put to work? To address this question, she examines the relationship between assets under management (AUM) and hedge fund returns.

She finds that small funds seem to perform better, supporting the hypothesis that there are limits to scale. However, she also finds that in the short-term, increases in AUM improve performance.

Additionally, she finds evidence that there is more at work here than just size, because the very largest funds do not seem to suffer from the same capacity problems associated with other funds.

To the extent that her results suggest that a large investor does not shower its largesse on a single fund, she asks, how many funds make a diversified portfolio of hedge fund investments? She finds that diversification can be attained with relatively few funds, perhaps 25 or fewer. Additionally, she finds that the benefit of diversification does not dissipate as the number of funds held in the portfolio increases. This is an interesting result to reconcile with another observation made by other authors, namely that portfolios of hedge funds (or more specifically hedge fund indices) display low diversification benefit to tail events.

Introduction

Barry Schachter

Acar and Middleton remind us of an important lesson, namely that a statistic is not necessarily meaningful just because you can calculate it. Evaluating performance should be easy. It isn't. For example, much recent research (including some of the chapters in this section), and not a little common sense, has robbed the Sharpe Ratio of some of its lustre as a performance measure for hedge funds. Non-normal returns, illiquid or smoothed pricing, and the effects of estimation errors all negatively affect its usefulness. In the hedge fund universe, the use of maximum drawdown is as ubiquitous as the use of the Sharpe Ratio in traditional portfolio analysis. If I did not ask a prospective manager "what has been your largest drawdown?" I would almost surely be considered by my peers to be derelict in my assessment of the manager's potential. Yet, we see in this chapter that the popular maximum drawdown statistic has its problems too.

The problem with maximum drawdown begins when you realise that it is not sufficient to simply examine observed drawdowns in assessing performance. Observed or experienced maximum drawdowns will, other things being equal, be larger the longer is the time interval over which you observe returns for a manager. The longer the observation interval, the greater is the opportunity to experience a drawdown worse than has previously been known. For this reason, to use maximum drawdown as a performance measure, one needs to evaluate drawdowns in relation to some absolute or idealised benchmark, such as the maximum drawdown expected from a passive buy-and-hold strategy (appropriately defined). This is what Acar and Middleton do. They find that unless markets are inefficient, it can be difficult for maximum drawdown as a statistic to successfully detect managers with superior skill at generating returns while controlling drawdowns.

In the best tradition of the "belt-and-suspenders" approach to risk management, it is common to use multiple performance statistics, in part to overcome the lack of power or potential biases in individual statistics. Therefore, statistics based on maximum drawdown, such as the Calmar Ratio, are commonly looked at along with others, such as the Sharpe Ratio, Information Ratio, Sortino Ratio, or M-Squared. The danger here, of course, is that these statistics are not independent. That is, they do not provide independent information about performance. For example, the ordering of portfolio managers that would be obtained from the Sharpe Ratio is the same as that obtained from M-Squared. Several authors have also shown, as discussed by Acar and Middleton, that maximum drawdown can be stated as a function of the portfolio standard deviation (in a nonlinear way), meaning that it may contain very similar information. It is possible then, that using multiple measures of hedge fund performance may not be much more useful than using a single measure.

The "take-away", as in a lot of the research presented in this volume, is not "don't go there", but rather "use with care". Maximum drawdown can be a useful diagnostic, used intelligently, by understanding and being aware of its limitations.

Lee and Lee in their chapter apply the tools of portfolio optimisation to the problem of selecting a portfolio of hedge funds, ie, building a fund of funds (FoF). Their approach provides an important supplement to the qualitative work of FoF managers, in the assessment of individual portfolio managers, for possible inclusion in their portfolio of funds. Their quantitative approach is guided by very practical issues. First, they recognise that elegant mathematical models, while interesting in the abstract, need to be balanced against the need for ease of implementation and interpretation. Second, they recognise that actual return characteristics of hedge funds, more specifically the negative skewness and excess kurtosis observed empirically, raise questions about the applicability of standard Markowitz mean-variance portfolio optimisation (just as they raise questions about the applicability of the Sharpe Ratio).

In this chapter, they develop an alternative performance ratio that is robust to the departures from normality that lessen the usefulness of the Sharpe Ratio. Sharpe has argued for some time that the appropriate way to measure the desirability of an asset is by it's impact on the portfolio's Sharpe Ratio, not by its stand alone Sharpe Ratio. Lee and Lee extend this idea to show that the optimal FoF portfolio is found by choosing the particular allocations to various funds, such that each individual candidate fund's Alternative Sharpe Ratio (ASR), measured by its contribution to the FoF's risk, is equal to the FoF's ASR (in some cases with an added adjustment factor). Lee and Lee describe a very practical approach to implementing this optimisation in a way that recognises that there is little value to being overly precise, if in practice investors are unlikely to allocate their assets in very tiny bits, maybe no smaller than one percent or one-half of one percent. They are able to illustrate with hedge fund index data that their method can yield better cumulative returns and be more sensitive to exposure to drawdowns accompanying extreme market events, such as the Liquidity Crisis of 1998.

The recognition by researchers and investors that the standard tools of portfolio theory and performance measurement are not generally applicable to hedge funds has resulted in a significant shift in the focus of research. The good news is that investors armed with this knowledge are less likely to make incorrect inferences from the standard metrics. For example, Sharma notes that optimisation of FoF portfolios on the basis of Sharpe Ratios will result in the over-weighting of non-directional strategies, which tend to have higher Sharpe Ratios (arising from low market betas), which in turn may result in unintended illiquidity and short volatility bias in the portfolio (as these are prominent characteristics of some non-directional strategies). The bad news is that it takes a while before new metrics are developed, comparative studies are done, and conclusions are drawn about what's best. In this chapter, Sharma uses a new performance measure (one of several new approaches presented in this volume) that addresses some of the known issues, namely, negative skewness and excess kurtosis in returns, the effect of leverage, and applicability when measured mean returns are negative. Sharma's method accomplishes this by assuming that an investor's utility for risk and return can be expressed with a particular, fairly general, class of utility functions. From this assumption, it is possible to go in several directions. For example, it becomes simple to infer the risk premium that investors would demand from a hedge fund, given its actual observed returns. Among other things, this allows investors to compare performance in much the same way that would be done using the Sharpe Ratio.

Sharma shows that traditional measures such as the Sharpe Ratio, Jensen's Alpha and the Treynor ratio yield qualitatively very different conclusions about hedge fund performance than does the new measure, alternative investments risk adjusted performance (AIRAP). Still, AIRAP measurements of various styles reveal that hedge funds do indeed provide risk adjusted added value for investors. Additional support for this conclusion is found in Sharma's analysis of persistence in hedge fund returns. Sharma finds strong persistence in returns on all hedge fund styles with the exception of two notoriously volatile (and typically highly directional) styles, Global Macro and commodity trading advisor (CTA). The expanded use of alternative metrics like AIRAP for alternative investments, both in performance evaluation and in portfolio construction, can only benefit investors' decision making.

In contrast to Sharma's development of an alternative performance metric, Lee and Lee in this (their second) contribution, point out that the popularity of metrics like the Sharpe Ratio is such that they are likely to continue to be used, despite the existence of known problems and the existence of alternatives. Most people hate change, and having gotten accustomed to using the Sharpe Ratio in one context, investors are loathe to stop using it in other contexts. We could therefore try to identify those cases in which the familiar metric can still be used effectively, and decry any application outside these cases. For the remaining cases in which the standard measure does not apply well, we may prefer the approach advocated by Sharma, ie, to create new metrics which address some of the known problems, but that have the look and feel of the familiar metrics in need of repair.

Focusing on the use of Sharpe Ratios, Lee and Lee argue that intuition suggests that measured returns to portfolios of hedge funds or, FoFs should exhibit more normally distributed returns than in the individual funds (a kind of Central Limit Theorem). If that were so, then some of the basic issues with using the Sharpe Ratio would be resolved. However, given the problems with such return series documented elsewhere, in large part based on publicly available data on returns to hedge fund indices, we have good reason to be skeptical. Lee and Lee offer up some hope by arguing that examining FoFs directly, rather than indirectly (by looking at hedge fund indices), will more than likely better reflect actual experiences of investors, and be free of much of what is known to be problematic about published returns to hedge fund indices.

Lee and Lee find that under pretty general conditions, the Sharpe Ratio of a portfolio of hedge funds is asymptotically normally distributed, irrespective of the shape of the return distributions of the component funds. They then find empirical support for their results by conducting Monte Carlo simulations using a proprietary database of hedge funds. In particular, they find that FoFs that contain 40 or more (randomly selected) hedge funds have normally distributed Sharpe Ratios. Lee and Lee conclude that there is still room for using the Sharpe Ratio in evaluating investment performance for the FoF alternative investment vehicle.

Introduction

Barry Schachter

When we think about risk, we think about normal distributions, not so much because we think everything is distributed normally, but rather because we're comfortable with normality as a starting point. There's a lot of history behind the analytics of risk measurement with normal distributions. Unfortunately, the empirical truths of hedge fund returns don't fit nicely into this framework (in general, of course). The most widely documented departure from normality is "too many large returns", both positive and negative (and high autocorrelation in large returns). In hedge fund returns, as several authors in this volume note, we find too many extreme negative returns (negative skewness) relative to the normal distribution. If returns were normally distributed, then a 2.65 daily standard deviation loss event would happen about one day a year. In practice, a rule of thumb in many markets is to expect a 4 standard deviation loss event about once a year. This non-normal reality is a problem for traditional, normality-based, parametric risk measurement. Researchers have been working for some time now on the development of useful risk measures that are not beholden to the normality assumption.

Bali and Gokcan put forward an alternative parametric approach to VAR estimation that is robust to the types of deviations from the normal probability distribution seen in hedge fund returns. In broad brush terms, the usual parametric VAR calculation proceeds first by calculation of the standard deviation of portfolio returns, then by multiplication by a constant, eg, 2.33, based on the known quantiles of the normal distribution. Their alternative VAR method, is represented in very similar manner. The standard deviation is first calculated and a "modified" multiplicative constant, one that incorporates information about the skewness and excess kurtosis in the return distribution, is applied to that portfolio standard deviation.

Bali and Gokcan examine the returns to 17 hedge fund strategy indices, all of which exhibit statistically significant departures from normality, including the negative skewness we have come to expect in hedge fund returns. They show that their alternative to the standard parametric VAR gives a very close correspondence to the empirically observed tail percentiles, much closer than the usual parametric approach, which significantly underestimates the empirical tail percentile. Their estimates are also better than when they estimate VAR using an extreme value approach, a result that they argue is further enhanced by the greater ease of applying their approach.

It should be obvious why hedge funds are attractive to investors. For example, they have, typically, low correlation to the equity market index and they offer returns comparable to the index. In the asset allocation world, the hot topic of whether hedge funds were a separate asset class has given way to the hot topic of how much to allocate to hedge funds. Still, despite all the enthusiasm for the asset class, there remains a very lively debate, noted elsewhere in this volume, about how it is possible for hedge funds to offer these attractive characteristics. Christiansen, Madsen, and Christensen (CMC) extend one line of recent research in this area. They attempt to identify risk factors that explain hedge fund returns. The first benefit of this analysis is to gain a better understanding of return drivers, and the second is to get a proper handle on the extent of "alpha" that is contributed beyond that which can be explained by the factors. That is, CMC want to find the incremental return not explained by compensation for risks that a buy-and-hold investor could obtain, in principle, without the aid of a hedge fund manager.

Key to this analysis is the recognition that the dynamic nature of hedge fund trading has some strong similarities to the return profile of options. Failure to account for the option-like characteristics of returns can result in failing to capture the essence of the returns of hedge funds. It's important to note that it is not necessary for hedge funds to be heavy traders of options for their returns to be option-like. The analogy can be made to the old strategy of equity portfolio insurance, through dynamic trading of equity indices and cash in such a way as would replicate the payoff of a put option. More directly, traders who are quick to cut their positions as the market price moves against them, and who add to their positions as the market moves in their favour, are essentially dynamically replicating an option position on their underlying position. Further, it is very likely that they themselves do not think of their trading in these option terms at all.

Rather than using self-styled strategy classifications, they use principal components analysis (PCA) to permit the hedge funds' return data to tell the story of the number of different styles that effectively exist. Then, they qualitatively assign those styles names, eg, Global Macro, based on other information about the funds that seem to align themselves most closely with the PCA-identified style factors. They settle on five distinct hedge fund styles. Then they try to explain the returns to portfolios of hedge funds corresponding to each style using various indices for equities, fixed income and commodities, and importantly, options on those indices.

Finally, CMC find that option-like return dynamics do play a role in explaining return. One notable exception is ambiguous evidence for option-like dynamics in Global Macro strategies, which characteristically are trend-following, directional bet-makers, and should not be expected to evince significant option-like returns. They find only mixed evidence of a significant "alpha" to be attributed to the skill of the hedge fund managers taken as a group.

A repeating theme in this volume is that what you get with hedge fund investments may be more than what you see, and that may not be a good thing. The extensive literature on hedge fund styles, to which Okunev and White add their voice here, is an attempt to reveal more about hedge fund risk and return than is found in standard measures used in the mutual fund industry. Okunev and White provide a method for correcting reported hedge fund returns for "smoothing", in part an artifact of infrequent trading and illiquid pricing of positions taken by the fund managers. Their methodology for removing first and second order autocorrelations in effect shows that "de-smoothed" hedge fund index returns may have as much as two times the volatility as the originally reported returns for some categories of hedge funds. They then use their de-smoothed returns to identify linear and non-linear (ie, option-like) risk factors that best explain the returns. By associating hedge fund index returns with factors, investors can understand the actual risk drivers behind those returns. It may be the case that hedge fund categories have returns explained by similar risk drivers, even when their measured correlations are low. Okunev and White find, just as CMC and others have recently found, that non-linear factors are important to understanding the distribution of hedge fund index returns. In particular, short option structures consistently appear in the explanation of returns.

The finding of important non-linear effects motivates the use of VAR historical simulations to estimate the risk in hedge fund indices. The historical simulation VAR approach lets the data speak for itself in measuring risk. Okunev and White use the estimated risk factor mappings from their style analysis in risk estimation. The de-smoothed returns are likely a better historical record than the reported returns. However, for risk estimation, a forward-looking estimation is required. The risk factor mappings indicate how a hedge fund return will respond to specific future realisations of possible market movements. Therefore a simulation using the factor mappings will simulate the effects of the trading dynamics that take place, and provide a more truthful estimate of likely future extreme returns. The VAR estimates of the various hedge fund styles differ dramatically in both absolute and relative terms when de-smoothed returns are used compared to the use of naïve estimates.

The expression "knowledge is power" was coined by Francis Bacon, arguably the leading Renaissance thinker in England and proponent of modern scientific methods of research, in reference to the power over the forces of nature that is granted by knowledge of the workings of nature. As James Madison, 4th president of the United States, said, "knowledge will forever govern ignorance". For hedge fund investors, knowledge of the manner in which publicly available hedge fund data is gathered, explains the biases in observed hedge fund returns. This knowledge thus empowers investors to assess the risks associated with investment decisions through new tools that account for these biases. Posthuma and van der Sluis implement a new approach for cleaning hedge fund index data affected by the well-known backfill bias in hedge fund returns data. This bias exists because hedge fund managers choose the most favourable timing for adding their funds' returns to common databases. The effects of self-selection and liquidation results in hedge fund index returns appearing better than they actually are. Posthuma and van der Sluis show, using their methodology, that backfill bias may be much greater than has previously been thought. They find that it is pervasive across hedge fund styles and time periods.

They then use their cleaned data to examine whether they can find persistent patterns in hedge fund returns. The persistence of abnormally high returns provides support for the argument that hedge funds (at least some at any rate) do provide excess returns. They divide hedge funds into deciles and examine the decile migration probabilities in consecutive periods. They find that a fund in the highest (lowest) decile in one month has the greatest chance of remaining in that decile the following month. This is evidence that supports the theory of return persistence. Interestingly, and somewhat contradictorily, they find that the second most likely return decile for the top (bottom) decile funds in the following month is the lowest (highest) decile. When they assign funds to deciles according to their information ratios (to account for volatility), they obtain similar results. This picture, while a bit cloudy, is actually more informative than the apparently clearer results obtained when naively using raw data on hedge fund index returns.

Introduction

Barry Schachter

Bacmann and Pache tackle the problem of optimisation of asset allocation to hedge funds, taking into account non-normality in hedge fund returns. They implement an approach that is robust to the empirically observed negative skewness and excess kurtosis of hedge fund index returns, two characteristics of return distributions which are discussed elsewhere in this volume and which are, in general, disliked by risk averse investors. Bacmann and Pache accomplish this by making somewhat different behavioral assumptions about investors. In the first case, they assume that investors choose a portfolio to minimise the chance that returns will fall below a threshold level at some future date. The optimum portfolio is characterised by a statistic called the Stutzer index, a higher index value being more desirable. In the second case, they assume that investors choose a portfolio to maximise the ratio of the expected gains above a threshold level to the expected losses below that threshold level. The optimum portfolio is characterised by the Omega Ratio, a higher value corresponding to a more desirable portfolio.

To evaluate the effects of non-normality on portfolio selection, they construct an empirical portfolio optimisation exercise. They assume that an investor is allocating among 10 strategy-specific hedge fund indices and an index intended to proxy for a commodity futures fund. Bacmann and Pache then compare optimal portfolios using the Stutzer index and the Omega Ratio, and use as benchmarks the mean/standard deviation portfolio allocation framework and a skewness-and-kurtosis-adjusted parametric VAR model that replaces the standard deviation in the mean/standard deviation framework.

Bacmann and Pache find that mean-variance portfolios usually overweight indices with negative skewness and high kurtosis relative to the new measures, a point in the favour of the new measures. As a corollary, they find that the new measures tend to produce optimal portfolios with less negative skewness and less excess kurtosis. These results are broadly the same for both in-sample and out-of-sample experiments. Their findings suggest that their proposed measures may have something useful to say to investors, when returns exhibit departures from normality.

For several years pundits have been proclaiming that alternative investments (and hedge funds in particular) deserve a place in investors' portfolios. Notwithstanding the concerns of supervisory authorities on the level of sophistication required of investors to properly evaluate hedge funds, the weight of empirical evidence to date seems to support this salutary view. Hagelin and Pramborg contribute further to this evidence by evaluating returns to portfolios that include hedge fund investments, where the investor is assumed to allocate optimally among asset classes, and to rebalance those allocations regularly.

Like Bacmann and Pache, Hagelin and Pramborg do not assume that investors only care about expected return and standard deviation of return. Instead they allow investors to care about all the properties of asset return distributions. In their general approach, investors explicitly maximise their expected utility from investing in risky assets. Hagelin and Pramborg show how this approach can be employed relatively simply in practice.

When the expected utility maximisation apparatus is set into motion, the chapter shows us, at a fundamental level, that investors will optimally choose to allocate wealth to hedge funds, even when they take into account deviations from normality in hedge fund returns. Hagelin and Pramborg also show that those optimal allocations result in portfolio returns that have, in some cases, statistically significantly greater returns compared to the case in which hedge fund investing is not permitted. These results appear to be robust to a wide range of assumptions about the extent to which the investor is averse to risk.

As with any empirical study, the quality of the inferences that can be made are limited by the quality of the data employed. Hagelin and Pramborg remind us not to take the results of their chapter completely at face value because hedge fund return data is known to have a variety of deficiencies, discussed elsewhere in this volume. They look at the possible impact of some of those deficiencies on their results, and find that their results remain qualitatively intact.

Lhabitant joins the quest for a sound method to use in evaluating risk in portfolios of funds. He finds attractive the use of a small number of risk drivers (factors) to reduce the complexity of the task. Like other authors, he proposes a linear factor model (a linear regression) to identify the risk drivers. Since there is ample research demonstrating that hedge fund returns have option-like, ie, nonlinear, features, he proposes using as candidate factors, variables that are known to be nonlinear functions of more primitive risks. The natural set of candidates, then, are hedge fund style indices.

Lhabitant then illustrates how this estimation approach can be used to compare a hedge fund's stated approach to the styles that actually explain the fund's returns. This technique might also be useful for evaluating style drift, a risk that is particularly difficult to monitor for hedge fund investors. This approach also suggests which index benchmarks might be most useful for evaluating performance. Finally, this approach provides a coarse means for an investor to establish (or verify) that their hedge fund portfolio effectively contains a diversified set of styles.

Lhabitant suggests that VAR could be calculated for any fund or portfolio of hedge funds, even without the benefit of obtaining position information from the managers. First, the VAR of each of the style factor indices is calculated. The fund VAR is composed of two (appropriately combined) elements. The first element is, in effect, the VAR of a portfolio of style indices, with weights provided by the funds' betas with respect to each style. The second element is the VAR of the fund variance that is unexplained by the style factors.