HEDGE FUND PERFORMANCE 1990-2000

DO THE 'MONEY MACHINES' REALLY ADD VALUE?

Gaurav S. Amin*

Harry M. Kat#

Working Paper

First version: January 4, 2001

This version: December 6, 2001

Please address all correspondence to:

Harry M. Kat
ISMA Centre
The University of Reading
Whiteknights Park
Reading RG6 6BA
United Kingdom
Tel. +44-118-9316428
E-mail: h.kat@ismacentre.rdg.ac.uk

---

*Ph.D student, ISMA Centre, University of Reading, #Associate Professor of Finance, ISMA Centre, University of Reading. We would like to thank Stephen Brown (the editor) for many helpful comments and suggestions that have greatly improved the paper. In addition, we would like to thank Philip Dybvig, Stewart Hodges, David Hsieh, Alexander Ineichen, Mark Rubinstein and discussants and participants of the MFS 2001 Conference in Garda, the EFMA 2001 Conference in Lugano, the SIRIF Conference on the Performance of Managed Funds in Edinburgh, the IAFE Risk Conference in Budapest, the APFA 2001 Conference in Bangkok, the EIR 2001 Conference in Paris, the OVFA 2001 Annual Meeting and the DGF 2001 Conference, both in Vienna, for helpful comments and LaPorte for providing the hedge fund and hedge fund index data. Of course, we are fully responsible for any errors or omissions.

HEDGE FUND PERFORMANCE 1990-2000

DO THE 'MONEY MACHINES' REALLY ADD VALUE?

ABSTRACT

We investigate the claim that hedge funds offer investors a superior risk-return trade-off. We do so using a continuous time version of Dybvig's (1988a, 1988b) payoff distribution pricing model. The evaluation model, which does not require any assumptions with regard to the return distribution of the funds to be evaluated, is applied to the monthly returns of 77 hedge funds and 13 hedge fund indices over the period May 1990 – April 2000. The results show that as a stand-alone investment hedge funds do not offer a superior risk-return profile. We find 12 indices and 72 individual funds to be inefficient, with the average efficiency loss amounting to 2.76% per annum for indices and 6.42% for individual funds. Part of the inefficiency cost of individual funds can be diversified away. Funds of funds, however, are not the preferred vehicle for this as their performance appears to suffer badly from their double fee structure. Looking at hedge funds in a portfolio context results in a marked improvement in the evaluation outcomes. Seven of the 12 hedge fund indices and 58 of the 72 individual funds classified as inefficient on a stand-alone basis are capable of producing an efficient payoff profile when mixed with the S&P 500. The best results are obtained when 10-20% of the portfolio value is invested in hedge funds.

.

I. INTRODUCTION

A hedge fund is typically defined as a pooled investment vehicle that is privately organised, administered by professional investment managers, and not widely available to the public[1]. Due to their private nature, hedge funds have less restrictions on the use of leverage, short-selling, and derivatives than more regulated vehicles such as mutual funds. This allows for investment strategies that differ significantly from traditional non-leveraged, long-only strategies.

Due to its private nature, it is difficult to estimate the current size of the hedge fund industry. Van Hedge Fund Advisors[2] estimates that by the end of 1998 there were 5,830 hedge funds managing $311 billion in capital, with between $800 billion and $1 trillion in total assets. So far, hedge funds have primarily been an American phenomenon. About 90% of hedge fund managers are based in the US, 9% in Europe and 1% in Asia and elsewhere. In the last five years the number of hedge funds has increased by at least 150%. Around 80% of hedge funds are smaller than $100 million and around 50% is smaller than $25 million, which reflects the high number of recent new entries.

With management fees based on assets under management, marketing is extremely important in asset management. Most hedge fund managers have substantial experience in capital markets, either as an investment manager, investment analyst or as a proprietary trader. This expertise is often presented to investors as a virtual guarantee for superior performance. A recent report by KPMG Consulting (1998, p. 3) for example boldly states that '…. the long-term average performance of hedge funds as a group can be estimated to be in the range of 17-20%, several percentage points higher than traditional equity returns'. Especially in today's market environment with low interest rates and falling equity markets, many private as well as institutional investors are very sensitive to such arguments.

In this paper we investigate whether hedge funds indeed offer investors a superior risk-return profile. We are not the first to do so. Several authors have used traditional performance measures such as Jensen's alpha, the Sharpe ratio, and asset class factor models to analyse hedge fund returns.[3] In general, the conclusion from this type of research is that hedge funds indeed generate superior results. There is a problem, however. All these methods assume hedge fund returns to be normally distributed and to be linearly related with other asset classes. Recent research, however, has shown that neither of these assumptions is correct.[4] Due to the special nature of the investment strategies adopted, hedge fund returns may exhibit a high degree of non-normality as well as a non-linear relationship with the stock market. Since this makes the use of traditional performance measures questionable, in this paper we use a new dynamic trading based performance measure that does not require any assumptions about the distribution of fund returns.

The analysis proceeds in two steps. First, we investigate whether hedge funds offer superior performance on a stand-alone basis. In other words, we study whether hedge funds offer good value for money to investors that invest in nothing else but hedge funds. Hedge fund managers may all be experts in their field, but the presence of certain special skills does not guarantee superior performance. The opportunity costs of potentially poor diversification across assets as well as through time, the transaction costs incurred and the management and incentive fees charged, all have to be borne by the investor. The question therefore is not whether hedge fund managers have special skills per se, but whether they have enough skill to compensate for all these costs, which can be very substantial. Only in that case can we speak of truly superior skill and performance. The next step is to look at hedge funds in a portfolio context. We do so by mixing hedge funds with equity to see whether this creates portfolios that offer a superior risk-return profile. These results are relevant for investors that, apart from hedge funds, also invest in equity.

Throughout we concentrate on the question whether in terms of risk and return hedge funds offer investors value for money. It is important to note from the outset, however, that strictly speaking this is a different question than whether hedge funds should be included in an investment portfolio. The fact that an investment offers a superior risk-return profile does not automatically mean that investors should buy into it as it may not fit their preferences and/or fit in with other available alternatives.

The remainder of the paper is organized as follows. In the next section we briefly discuss traditional evaluation methods and their limitations. An alternative performance measure is presented in section III and IV. In section V we use the latter to evaluate 13 hedge fund indices and 77 individual funds over the period May 1990 – April 2000. In section VI we study the impact of transaction costs, in section VII we study hedge funds in a portfolio context, while in section VIII we look at the sensitivity of the analysis for outliers and the reference index used. Section IX concludes.

II. TRADITIONAL PERFORMANCE MEASURES

Although over the years much academic work has been done in this area,[5] practitioners typically use either one of two performance measures: the Sharpe ratio and Jensen's alpha. The first measure was introduced in Sharpe (1966) and is calculated as the ratio of the average excess return and the return standard deviation of the fund that is being evaluated. As such it measures the excess return per unit of risk. Assuming all asset returns to be normally distributed, the CAPM tells us that in equilibrium the highest attainable Sharpe ratio is that of the market index. A ratio higher than that therefore indicates superior performance. In more general terms, the market index's Sharpe ratio represents the set of return distributions that is obtained when statically combining the market index with cash. With the market index being highly diversified, these distributions offer the highest achievable expected return for every possible standard deviation. If a fund produces a return distribution with an expected return lower than that of a similar (in terms of standard deviation) distribution from this set, the fund is deemed to be inefficient and the other way around. The test that we discuss in section III follows exactly the same line of reasoning, except that we explicitly allow for the possibility of dynamic trading in the index and cash. This widens the span of these two primitive assets, which allows us to fully incorporate any non-normality found in the fund return into the performance evaluation without having to make far reaching assumptions about investors' preferences.

Jensen's alpha was introduced in Jensen (1968) and equals the intercept of the regression

$$\left(R_h - R_f\right) = \alpha + \beta\left(R_i - R_f\right) + e_h\,, \tag{1}$$

where R_h is the fund return, R_f is the risk free rate and R_i is the total return on the market index. Like the Sharpe ratio, Jensen's alpha is rooted in the CAPM. According to the CAPM, in equilibrium all (portfolios of) assets with the same beta will offer the same expected return. Any positive deviation therefore indicates superior performance. Jensen's alpha also has a more practical interpretation as we can see (1) as comparing the fund return, through the fund beta, with the return on an investible benchmark portfolio. If the average fund return is significantly higher than one would expect given the fund beta and the average benchmark return we speak of superior performance. The above reasoning is easily extended to the case of multiple benchmarks, which can be justified in an APT context. The benchmark portfolios can be fairly standard, such as in Sharpe (1992) for example, but they can also be explicitly constructed to capture observed non-linearities. One way to extend the span of one or more benchmarks from purely linear to non-linear is to include a number of ordinary put or call options on those benchmarks into the return generating process. This approach was first introduced by Glosten and Jagannathan (1994) and recently applied in a hedge fund context by Agarwal and Naik (2001). Apart from the selection of the relevant benchmarks, there are two problems with this approach though. First, a priori it is not clear how many options and which strike prices should be included. Ideally, one would want to infer this from the available data. In practice, however, the data sample will be too small to avoid ad hoc specification. Second, since only a small number of ordinary puts and calls can be included, there is a definite limit to the range and type of non-linearities that can be captured. The method discussed in the next section solves both problems by, instead of relying on a limited number of standard options, constructing a completely new type of option contract for every fund that is being evaluated, with a payoff function that is fully determined by the empirically observed return distribution.[6] This allows us to fully capture any observed deviation from normality, irrespective of its causes.

To confirm that non-normality and non-linearity are indeed problems to be reckoned with when evaluating hedge fund performance, we studied data from the Zurich Capital Markets (formerly MAR/Hedge) database, concentrating on the 77 funds for which at least ten years of consecutive monthly return data (net of fees) was available, from May 1990 until April 2000. Obviously, this introduces survivorship bias. There is no simple way to mitigate this problem though.[7] Note that the sample includes the period 1997-1998. With crises in Asia and Russia and the subsequent collapse of LTCM, this was a difficult time for hedge funds. A further classification of the 77 funds used can be found in table 1. Apart from individual funds, we also studied the monthly returns on 13 different hedge fund indices calculated by Zurich over the same period.[8]

<< Insert Table 1 >>

Looking at the monthly returns of the 13 hedge fund indices and 77 individual hedge funds in more detail we found the return distributions of most hedge fund indices as well as individual funds to be highly skewed. The Jarque-Bera (1987) test for normality confirmed that at a 5% significance level for none of the index distributions and for only 14.1% of the individual fund distributions normality is a satisfactory approximation. We also found strong non-linearities in the relationship between hedge fund returns and equity market returns. For brevity we do not report these results here but they are available on request.

III. AN ALTERNATIVE PERFORMANCE MEASURE

To properly evaluate the performance of portfolios with a non-normal return distribution, the entire distribution has to be considered. Ideally, this should be done without having to make any prior assumptions regarding the type of distribution. The performance measure that we will use in this study does exactly that. It is based on the following reasoning. When buying a fund participation, an investor acquires a claim to a certain payoff distribution. If we wanted to investigate whether a fund manager had any superior investment skills the most direct line of attack would therefore be to re-create the payoff distribution that he offers his investors by means of a dynamic trading strategy of some sort and compare the cost of that strategy with the price of a fund participation. If the manager in question indeed had superior skills, the strategy should be more expensive than the fund participation. Of course, the same payoff distribution can be generated in many different ways. The critical issue is therefore to find the strategy that does so most efficiently, i.e. at the lowest cost.

Assuming we live in the world of Black and Scholes (1973), we determine the cost of the cheapest dynamic trading strategy, trading some reference index and cash, which generates the same payoff distribution as the hedge fund in question in two steps.[9] First, we construct a payoff function that in combination with the distribution of the reference index yields the same end-of-month payoff distribution as the fund. Second, we price this payoff function using standard option pricing technology. The motivation and exact procedure, which we will refer to as 'the efficiency test', are explained in more detail below.

<< Insert Figure 1 >>

Having collected monthly return data on the fund to be evaluated, the first step is to use these returns to create an end-of-month payoff distribution, assuming we invest $100 at the beginning of the month. The same is done for the reference index. We do not make any assumption about the distribution of fund returns. Living in a Black-Scholes world, however, we explicitly assume the reference index to be normally distributed. An example of the resulting cumulative distributions can be found in figure 1. Of course, the index distribution is smooth due to its assumed normality.

Next, we construct a payoff function that, in combination with the index distribution, yields exactly the same payoff distribution as produced by the fund. Since there are many functions that will map one into the other, we make the additional assumption that the payoff must be a path independent non-decreasing function of the index value at the end of the month. The latter assumption derives from the theoretical work of Cox and Leland (2000) and the payoff distribution pricing model of Dybvig (1988a, 1988b). Cox and Leland (2000) showed that in a Black-Scholes world all path dependent strategies are inefficient because they generate payoff distributions that can also be obtained by a path independent strategy, but at lower costs. In addition, from

Dybvig (1988a, 1988b) we know that any investor who maximizes expected utility and prefers more money to less will want his wealth at the end of his investment horizon to be a monotonic non-increasing function of the state-price density. Investment strategies that do not have this feature will again be inefficient. In a Black-Scholes world with a positive risk premium the state-price density is a decreasing function of the terminal value of the index. This means that for a strategy to be efficient, final wealth must be a monotonic non-decreasing function of the terminal index value. Intuitively, this is a plausible result. A non-decreasing payoff will be positively correlated with the index. As a result, the rebalancing trades required by the strategy generating that payoff will tend to be relatively modest, which keeps trading costs down.

<< Insert Figure 2 >>

Given the above assumption, construction of the required payoff function is straightforward. Suppose that the fund distribution told us there was a 10% probability of receiving a payoff lower than 100. We would then look up in the index distribution at which index value X there was a 90% probability of finding an index value higher than X. If we found X=95, the payoff function would be constructed such that when the index ended at 95 the payoff would be 100. Next, we would do the same for a probability of 20%. If the fund distribution told us there was a 20% probability of receiving a payoff lower than 110 and the index distribution said there was a 80% probability of finding an index value higher than 100, the payoff function would be constructed such that when the index ended at 100 the payoff would be 110. This procedure is repeated until we get to 100%. An example of a typical payoff function (using 0.2% instead of 10% steps as in the above example) can be found in figure 2.

The third step is to find the initial investment required by the dynamic trading strategy, trading the reference index and cash, which generates the derived payoff function. Following Harrison and Kreps (1979), the latter can be calculated as the discounted risk neutral expected payoff. Since the payoff is not a neat function, however, we have to do so using Monte Carlo simulation. Having constructed a payoff function, we therefore generate 20,000 end-of-month index values using a discretized and risk neutralized geometric Brownian motion, calculate the corresponding payoffs, average them and discount the resulting average back to the present at the risk free rate to give us the price of the payoff function in question. If the price thus obtained is higher than \$100, we take this as evidence of superior performance and the other way around.

Since the efficiency test does not rely on a very specific asset pricing model, the choice of reference index is less critical than in many other performance evaluation models. The most important requirement is that, to avoid the problem of inefficient diversification across assets, the index is well diversified. If not, the benchmark itself will not be completely free of inefficiencies. Since we assume the index to follow a geometric Brownian motion, we also need index returns to be (approximately) normally distributed. In addition, and somewhat more practical, to keep transaction costs to a minimum we would prefer the index to trade not only in the cash market but also in a liquid futures market. The above requirements make the S&P 500 index, which is also often used in traditional performance evaluation, an obvious choice.

Choosing the reference index without reference to the market in which a fund actually invests may seem inadequate to those used to working with multi-factor models. This is, however, purely a matter of approach. Like the Sharpe ratio, the efficiency test evaluates the bottom line risk-return profile offered by a fund. The test concentrates solely on the fund return and not on the underlying return generating factors. It therefore requires neither information nor assumptions about the actual sources of risk and return.[10] The above can also be cast in terms of the general stochastic discount factor model. Working from the assumption that the S&P 500 contains little or no idiosyncratic volatility, the efficiency test simply extracts a stochastic discount factor from the available S&P 500 returns and uses that to price the fund payoff.[11] In this more general context, the finding of Chernov (2000) that, after comparing a number of different estimates of the pricing kernel, the best estimate involves only information on S&P 500 returns, is quite reassuring.

By using the cost of the most efficient dynamic trading strategy that generates the same payoff distribution as the fund as a benchmark we test whether a hedge fund manager has sufficient skill to compensate not only for transaction costs and management fees (which simply do not exist in a Black-Scholes world), but also for the inefficiency costs of potentially poor diversification across assets as well as through time. Implicit in this is the assumption that investors are only interested in the end-of-month payoff of a strategy and not in its intermediate values. This is an adequate characterization of hedge fund investors as most hedge funds have strict lock-up and exit rules that effectively force investors to take a longer-term view.

As mentioned before, the efficiency test is the dynamic trading equivalent of the Sharpe ratio. Assuming normal distributions, the Sharpe ratio tests whether a fund's return distribution offers an expected return that is higher than achievable on a static portfolio of the index and cash with the same standard deviation. The efficiency test follows the same reasoning, only this time in a dynamic trading framework. This eliminates the normality restriction since by dynamically trading the index and cash we are not only able to replicate the fund return's standard deviation but also its higher moments. Subsequently, we compare the expected fund return with that of the replicating distribution to decide whether the fund is efficient. If fund returns are normally distributed the efficiency test collapses into the Sharpe ratio as in that case, given a normally distributed index, no dynamic trading is required. To test this explicitly we calculated the Sharpe ratios and applied the efficiency test on 33 UK based equity mutual funds for which over the period May 1990 – April 2000 the hypothesis of a normal return distribution could not be rejected. We used the 1-month T-bill rate as a proxy for the risk free rate and the FTSE All Share index as the reference index. The efficiency test results showed a correlation of 0.99 with the Sharpe ratio results.

We are not the first to approach performance evaluation from a contingent claims perspective. Merton (1981) showed that the returns from successful market timing will resemble the returns from certain option strategies. Glosten and Jagannathan (1994) extended the traditional model with a number of additional factors corresponding with the payoffs of ordinary puts and calls. Agarwal and Naik (2001) revisited this approach an applied it to a number of hedge fund indices. Fung and Hsieh (2001) show that the returns from trend following strategies are similar to those from lookback straddles. Finally, Mitchell and Pulvino (2001) find that the returns from risk arbitrage strategies are very similar to the results from writing ordinary put options. All these authors link fund payoffs with the payoffs of specific option contracts. The efficiency test, however, is much more general as we do not replicate fund returns by trading specific types of options. Instead, for every fund we construct and price a completely new type of option with a payoff function that is fully determined by the empirically observed fund return distribution. As such, the above studies can be interpreted as special cases of the efficiency test.

IV. SAMPLING ERROR

We have 10 years of monthly data available to estimate hedge funds' payoff distributions. As a result, the efficiency test will be confronted with sampling error. Since we do not make any assumptions about the nature of the distributions involved, a formal study is problematic. We can, however, obtain an indication of the possible extend of the error by studying the efficiency test's application on a payoff function that we know to be efficient. We therefore applied the efficiency test to a so-called buy-write, i.e. an index plus a short call on that same index. Since the payoff of such a package is neither hampered by transaction costs or management fees nor by inefficient diversification, the test should produce a value of exactly 100. With only 120 observations available, however, this need not always be the case. Assuming monthly S&P 500 price returns to be normally distributed with a mean of 1.24% (14.88% per annum) and a standard deviation of 3.59% (12.43% per annum),[12] we generated 120 end-of-month index values and calculated the corresponding payoffs. Next, we applied the efficiency test to the 120 payoff values thus obtained, assuming the S&P 500 to follow a geometric Brownian motion with a volatility equal to the above standard deviation and a drift equal to the difference between the risk-free rate and the S&P 500 dividend yield. The former was set equal to the 10-year historic mean of 3-month USD LIBOR (5.35%) and the latter to its 10-year historic mean of 0.22% (2.65% per annum). We repeated the above procedure 20,000 times. A frequency distribution of the resulting annualised error, i.e. the difference between the actual test result and 100, can be found in figure 3

<< Insert Figure 3 >>

Figure 3 shows that with only 120 observations the efficiency test may produce an error that significantly differs from zero. The error distribution, however, has a high peak around zero, meaning that compared to a normal distribution there is a relatively high probability of a small error. In addition, figure 3 shows that the efficiency test is unbiased. The average error is –0.05% per annum. With more observations the sampling error will drop quickly, as was confirmed by further simulations. Given the currently available data on hedge funds, however, this is the best one can do. We repeated the above analysis for a number of other efficient payoff profiles, which yielded similar results.

V. APPLICATION TO HEDGE FUNDS

Using monthly total return data from May 1990 to April 2000, we calculated the alphas and Sharpe ratios of the 13 hedge fund indices and the 77 individual funds mentioned earlier in section II. We used 3-month USD LIBOR as a proxy for the risk-free rate and the S&P 500 as the relevant market index. The results can be found in the second and third column of table 2 and 3. Eleven indices show significant positive alphas. Twelve indices generate a Sharpe ratio higher than that of the S&P 500 (0.28). Of the individual funds, 71 show positive alphas while 28 funds produce a Sharpe ratio higher than that of the S&P 500. Based on this, it is very tempting to conclude that hedge funds have indeed shown superior performance over the last decade; a fact often quoted by hedge fund managers and marketers.

<< Insert Table 2 and 3 >>

Next, we applied the efficiency test, using the same S&P 500 parameter values as in the previous section. Note that over the period studied the ex-post risk premium has been relatively high, which potentially allows for significant inefficiency costs. The annualised evaluation results on the 13 indices can be found in the fourth column of table 2. Twelve of the 13 indices show signs of inefficiency with the average efficiency loss on these 12 indices amounting to 3.00% per annum. With an average efficiency loss of 4.15%, the three fund of funds indices make an important contribution to this figure. Excluding the latter, the average efficiency loss drops to 2.61%.

<< Insert Figure 4 >>

In practical terms the above means that by dynamically trading the S&P 500 and cash we would have been able to generate return distributions very similar to those generated by the hedge fund indices studied, except that in 12 out of 13 cases our engineered distributions would have had a substantially higher mean return. This is graphically illustrated in figure 4, which shows the mean, standard deviation, skewness and kurtosis of the distributions of the actual hedge fund index return as well as the return calculated from the payoff function resulting from the efficiency test.

<< Insert Figure 5 >>

The evaluation results on the 77 individual funds are reported in the fourth column of table 3 and are sorted and graphically reproduced in figure 5. We decided it would be inappropriate to name the individual funds by name. Instead, we number the funds in each class of funds arbitrarily and refer to them only by number. We do, however, provide information on the type of fund. The capitals denote whether a fund is classified as global (G), event driven (E), market neutral (M) or as a fund of funds (F). The lowercase letters denote whether the fund is US based (u) or not (n). Of the 77 funds studied, 72 show signs of inefficiency. For these 72 funds the average efficiency loss amounts to 6.97% per annum. Five funds offer superior performance with an average efficiency gain of 1.49% per annum. Even without taking survivorship bias into account, these results clearly contradict the claim that hedge funds generate superior investment results on a stand-alone basis.

With an average efficiency gain over individual hedge funds of 3.66%, the hedge fund indices perform significantly better than the individual funds, indicating that inefficiency costs can be reduced by investing in a portfolio of hedge funds instead of a single individual fund. Combining different types of funds seems to offer additional gains. If we were to combine all 13 indices into an equally-weighted portfolio, the efficiency loss would drop to 1.82%. Excluding the three fund of funds indices it would come down to 1.40%.

<< Insert Table 4 >>

To see whether there are any significant differences between different types of hedge funds, we divided the 77 individual funds in the sample in seven groups: fund of funds, non fund of funds, event driven, global, market neutral, offshore and US based. The results are summarized in table 4. All of the 15 event driven funds show signs of inefficiency with an average efficiency loss of 3.76% per annum. Of the 28 global hedge funds, 24 show some level of inefficiency with an average cost of 8.51% per annum. Of the 11 market neutral funds studied, 10 show some level of inefficiency. For this group the average efficiency loss is 6.80% per annum. US based funds appear to be more efficient than offshore funds. Four out of 49 US based funds classify as efficient with an average efficiency gain of 1.54% per annum. This is in line with the results of Ackermann, McEnally and Ravenscraft (1999) and Liang (1999), who found offshore funds to be substantially more volatile than US funds but without a corresponding higher mean return.

It is worrisome to see funds of funds perform so badly. Despite the extra layer of diversification, the three fund of funds indices are almost 2% behind on the other 10 hedge fund indices. The efficiency loss of the average individual fund of funds is 7.51% per annum, meaning that the average fund of funds is 1.56% less efficient than the average non-fund of funds and 5.17% less efficient than the average non-fund of funds index. A priori one would expect funds of funds to show results similar to those obtained for the non-fund of funds indices. The above result therefore strongly suggests that the fees charged by fund of funds managers outweigh the efficiency gains of additional diversification and potentially superior fund selection. Although smaller investors have little choice, larger investors should therefore think twice before externalising their hedge fund portfolio management.

VI. TRANSACTION COSTS

So far we have assumed transaction costs to be zero. Since our benchmark value equals the costs of a specific dynamic investment strategy, however, one might argue that this is not correct.[13] As shown by Leland (1985), transaction costs can be incorporated in the form of a volatility adjustment.[14] For strategies generating a payoff that is a convex function of the index at maturity, transaction costs can be modelled as an increase in volatility. The reverse is true for strategies generating a payoff that is a concave function of the index. If we denote the volatility of the index as σ, round-trip transaction costs (as a percentage of the index value) as c, and the time between subsequent rebalancings as Δt, then Leland's transaction costs adjusted volatility σ_L is given by

$$\sigma_L = \sqrt{\sigma^2 \left(1 + D \sqrt{\frac{2}{\Pi}} \frac{c}{\sigma \sqrt{\Delta t}} \right)} \qquad (2)$$

where D is +1 when dealing with a convex and –1 when dealing with a concave payoff. If we assume daily rebalancing and cash market execution at round trip transaction costs of 0.5%, the volatility to be used for convex strategies would be 15.34% and 8.74% for concave strategies. If we assumed futures market execution, transaction costs would of course be much lower, say 0.05% round trip. In that case convex strategies should be priced at 12.76% and concave strategies at 12.13%.

With futures market execution the adjusted volatilities are hardly different from the 12.43% that we have used so far, implying that in that case transaction costs will not change the previous conclusions. This is not the case with cash market execution. To see how cash market transaction costs change our results we investigated how high or low volatility needs to be to take the value of the hedge fund payoff distributions to 100. This information can be found in the fifth column of table 2 and 3. The numbers indicate the volatility level at which the value of the payoff distribution reaches 100. 'DNT 100' stands for 'does not touch 100', which means that there is no volatility for which the value of the payoff distribution equals 100. Concentrating on the 72 funds that have a value below 100 at 12.43% volatility, there are three interesting observations to be made.

1. 42 out of 72 reach 100 at a volatility level higher than 12.43% (18.79% on average). The fact that the value of these funds' payoff distributions increases with volatility implies a convex payoff. The benchmark volatility is therefore 15.43%, meaning that 8 of these funds can no longer be classified as inefficient.

2. 21 out of 72 reach 100 at a volatility level lower than 12.43% (8.81% on average), implying a concave payoff. The benchmark volatility is therefore 8.74%, which means that 4 of these funds can no longer be classified as inefficient.

3. 9 out of 72 fail to reach 100, even at 0% volatility. Apart from a concave payoff, this also implies that these funds are definitely inefficient.

In sum, with round-trip transaction costs at 0.5%, 12 of the 72 funds classified earlier as inefficient can no longer be classified as such. This still leaves us with 60 inefficient funds though. Looking at the 5 funds that show a value higher than 100 at 12.43% volatility, we see that 3 out of 5 show a decrease in value as volatility rises. Again, this implies a concave payoff.

The fact that the value of the payoff distribution offered by 33 out of 77 funds is a decreasing function of volatility is in line with the observation that many hedge fund return distributions offer relatively limited upside potential. Since a concave payoff is the most efficient way to generate such a distribution, this is what the mapping procedure produces. This does not mean that the hedge funds studied really offer investors a payoff that is a concave function of the index. All we are saying is that the payoff distributions offered by many hedge funds can be re-created cheapest by a dynamic trading strategy that aims to generate a payoff that is a concave function of the index. The last columns in table 2 and 3 show the correlation between the return calculated from the payoff function resulting from the mapping procedure and the actual hedge fund (index) return. These low correlation coefficients underline that although our payoffs replicate hedge funds' payoff distributions, they do not replicate their time series behaviour very well. In essence, this is one of the main reasons why we find hedge funds to be inefficient.

VII. Hedge Funds in a Portfolio Context

By construction, the mapped payoffs generate returns that are heavily correlated with the reference index. In reality, however, the relationship between hedge fund and index returns is unusually weak. For example, over the period studied the average correlation between the S&P 500 and the 77 individual hedge funds in the sample was only 0.29. The efficiency test used so far does not take this into account as it only aims to replicate hedge funds' unconditional payoff distribution and not their correlation profile. If we were to introduce an explicit correlation restriction into the efficiency test, i.e. require our trading strategies not only to replicate the funds' payoff distributions but also their correlation with the index, hedge funds would come out better as the restriction would make our replication strategies more expensive. There are two problems with this, however. First, incorporating an explicit correlation restriction into the procedure is technically complicated. Second, concentrating on correlation implicitly assumes the existence of a linear relationship between hedge fund returns and index returns. As discussed before, however, this is not the case.

To solve these problems we used the following procedure. First, we constructed portfolios of the S&P 500 and hedge funds, with the fraction invested in hedge funds ranging from 0% to 100% in 1% steps. Subsequently, we ran the efficiency test on these portfolios' 120 monthly returns to check for the existence of a combination that offers a risk-return profile that cannot be obtained with a mechanical trading strategy at a lower price. With the S&P 500 being efficient by definition, this procedure tests whether hedge funds' relationship with the S&P 500 is sufficiently weak to make up for the efficiency loss observed on a stand-alone basis. Table 5 and 6 show for every hedge fund index and individual hedge fund the correlation with the S&P 500, the highest efficiency value achieved and the portfolio mix at which this occurs.

<< Insert Table 5 and 6 >>

From table 5 we see that when mixed with the S&P 500, 7 of the 12 indices that were found to be inefficient on a stand-alone basis are able to produce an efficient payoff profile. In all 7 cases the most efficient mix consists of around 20% hedge fund index and 80% S&P 500. Given that the correlation coefficients of these indices with the S&P 500 (column 2) vary between 0.09 and 0.71, this indicates that there is more between hedge fund index returns and S&P 500 returns than a simple linear relationship. The results on individual funds can be found in table 6. Of the 72 previously inefficient funds, 58 are capable of producing an efficient payoff profile when mixed with the S&P 500. As with the hedge fund indices, the most efficient mix varies much less than expected given the funds' varying correlations with the S&P 500. For all 58 funds the most efficient mix consists of 10-20% hedge fund and 80-90% S&P.

Despite their marked inefficiency on a stand-alone basis, 13 of the 23 funds of funds are capable of producing an efficient payoff profile when mixed with the S&P 500. Although the achieved maximum efficiency level tends to be lower than for non-fund of funds, this clearly underlines the power of the portfolio effect. Another interesting result concerns the fact that with a 20/80 mix there is relatively little difference between the results for hedge fund indices and individual hedge funds. One might be inclined to conclude from this that in a portfolio context hedge fund diversification is much less of an issue than on a stand-alone basis. Such a conclusion would be premature, however, as it ignores the survivorship bias introduced by only including funds with at least 10 years of history.

VIII. SENSITIVITY ANALYSIS

With regard to the methodology used, a few questions remain. First, would the results have been different if instead of assuming S&P 500 returns to be normally distributed we used the S&P 500's empirical return distribution? To answer this question we repeated our study using the observed S&P 500 return distribution. The results obtained were hardly different from the ones reported previously. The correlation between both sets of results was 0.99. Second, how sensitive is the method to outliers? To gain insight in this matter we removed the top and bottom 2.5% of the return observations, leaving 114 instead of 120 monthly returns. The results again did not change much. The correlation between the results from 120 observations and 114 observations was 0.94. Finally, would the results have been different if we had used another reference index? Given the nature of the efficiency test, there is no reason to expect another, maybe even somewhat less diversified, index to yield completely different results. To test this, we repeated our tests using the Dow Jones Industrial (DJI) index as the reference index. The DJI is substantially different from the S&P 500. The former is made up of only 30 stocks, while the latter contains 500. In addition, instead of being value-weighted like the S&P 500 and most other major stock market indices, the DJI is price-weighted. Despite these differences, the test results did not change much when we used the DJI as the reference index instead of the S&P 500. The correlation between the S&P 500 results and the DJI results was 0.99. As an extreme case, we repeated the tests using the German DAX index (converted into US dollars). Although using the DAX instead of the S&P 500 yielded somewhat different results, the correlation with the S&P 500 results was still 0.81.

IX. Conclusion

In this paper we have used a continuous time version of the payoff distribution pricing model introduced by Dybvig (1988a, 1988b) to evaluate hedge fund performance over the period 1990 - 2000. The main results can be summarized as follows:

1. **Hedge fund returns and performance evaluation**. Because hedge fund returns distributions tend to be non-normal and non-linearly related to equity returns, traditional performance measures such as the Sharpe ratio and Jensen's alpha are not suitable for the evaluation of hedge fund performance

2. **The proposed performance measure**. The proposed efficiency test can be interpreted as a highly generalized Sharpe ratio. Thanks to its dynamic nature it is able to deal with any type of return distribution. The test appears to be unbiased, while with 120 observations sampling error risk does not seem to be prohibitively high.

3. **Hedge funds as a stand-alone investment.** With an efficiency loss of 6.42%, the average hedge fund makes for quite an inefficient investment. The 3.66% lower average efficiency loss observed on hedge fund indices, however, suggests that a major part of the inefficiency costs of individual funds can be diversified away by investing in a portfolio of hedge funds instead of a single hedge fund.

4. **Hedge funds in a portfolio context.** Hedge funds score much better when seen as part of an investment portfolio. Due to their weak relationship with the index, 7 of the 12 hedge fund indices and 58 of the 72 individual funds classified as inefficient on a stand-alone basis are capable of producing an efficient payoff profile when mixed with the S&P 500. The best results are obtained when 10-20% of the portfolio value is invested in hedge funds.

5. **Funds of funds.** The average stand-alone fund of funds' efficiency loss exceeds that of the average non-fund-of-funds hedge fund index by 5.17%. On a stand-alone basis, the average fund of funds therefore makes for quite a wasteful investment. Despite this, 13 of the 23 funds classified as inefficient on a stand-alone basis are able to produce an efficient payoff profile when mixed with the S&P 500. Again, the best result is obtained with 10-20% invested in the fund in question.

Our results make it clear that the main attraction of hedge funds lies in the weak relationship between hedge fund returns and the returns on other asset classes. It is interesting to note, however, that this is primarily the result of the general type of strategy followed by many hedge funds and not special manager skills. Any fund manager following a typical long/short type strategy can be expected to show low systematic exposure, whether he has special skills or not. This leads us to the question why investors should pay those high fees if the main attraction of hedge funds is not a manager specific feature? The answer of course is that investors have little choice. So far, only hedge funds provide hedge fund type returns.[15]

FOOTNOTES

1. See for example President's Working Group on Financial Markets (1999, p. 1).

2. See Van Hedge Fund Advisors International, Inc. at www.vanhedge.com.

3. See for example Fung and Hsieh (1997), Ackermann, McEnally and Ravenscraft (1999), Liang (1999) or Agarwal and Naik (2000).

4. See for example, Agarwal and Naik (2001), Mitchell and Pulvino (2001), Fung and Hsieh (2001), Lo (2001), or Brooks and Kat (2001).

5. An extensive bibliography on performance evaluation can be found on www.stern.nyu.edu/~sbrown/performance/bibliography.html.

6. Glosten and Jagannathan (1994) mention an approach similar to ours but they reject it on the ground that it may yield highly unusual contingent claims that are difficult to understand and value by practitioners.

7. Brown, Goetzmann and Ibbotson (1999), Fung and Hsieh (2000), Liang (2001) and Amin and Kat (2001) estimate the survivorship bias in hedge fund data at between 1.5% and 3% per annum.

8. The 13 indices considered are listed below with the number of funds included as of April 2000 between brackets: Event Driven (106), Event Driven: Distressed (45), Event Driven: Risk Arbitrage (61), Fund Of Funds (265), Fund Of Funds: Niche (232), Fund Of Funds: Diversified (33), Global: Emerging (85), Global: Established (245), Global: International (34), Global: Macro (58), Market Neutral (231), Market Neutral: Long/Short (109), and Market Neutral: Arbitrage (122). Note that the Event Driven, Fund of Funds and Market Neutral indices are simply baskets of the relevant sub-indices.

9. This is not unlike the approach nowadays taken to solve for optimal consumption and portfolio policies under uncertainty. See for example Cox and Huang (1989).

10. Factor model based evaluation of hedge fund performance is problematic as hedge funds follow complex and highly opportunistic strategies, often combining several strategies at the same time. As shown in Fung and Hsieh (1997), standard multi-factor models therefore tend to exhibit low to insignificant factor loadings as well as determination coefficients.

11. More details on performance evaluation with stochastic discount factors can be found in Chen and Knez (1996), He, Ng and Zhang (1999) or Dahlquist and Soderlind (1999).

12. These estimates were obtained from monthly S&P 500 return data over the period May 1990 - April 2000.

13. The effects of transaction costs in the binomial Dybvig (1988b) model are analysed in Pelsser and Vorst (1996).

14. Strictly speaking, Leland's volatility adjustment is only valid when gamma does not change sign, which need not be true in our case. If so, one should use the more general result of Hoggard, Whalley and Wilmott (1994), which requires the solution of a non-linear parabolic PDE with the payoff function to be priced as a boundary condition. Although this may cause underestimation of the actual expected transaction costs, to avoid additional complication we decided for the Leland approximation.

15. Inspired by the work of Mitchell and Pulvino (2001), a synthetic merger arbitrage fund was recently launched in the US. Although generated in a completely different way, the fund's returns appear to be highly correlated with the main merger arbitrage indices.

REFERENCES

Ackermann, C., R. McEnally and D. Ravenscraft (1999), The Performance of Hedge Funds: Risk, Return and Incentives, Journal of Finance, Vol. 54, pp. 833-874.

Agarwal, V. and N. Naik (2000), On Taking the Alternative Route: Risks, Rewards, Style and Performance Persistence of Hedge Funds, Journal of Alternative Investments, Vol. 2, pp. 6-23.

Agarwal, V. and N. Naik (2001), Characterizing Systematic Risk of Hedge Funds with Buy-and-Hold and Option-Based Strategies, Working Paper London Business School.

Amin, G. and H. Kat (2001), Welcome to the Dark Side: Hedge Fund Attrition and Survivorship Bias 1994-2001, Working Paper ISMA Centre, University of Reading.

Black, F. and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, Journal of Political Economy, Vol. 3, pp. 637-654.

Brooks, C. and H. Kat (2001), The Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors, Working Paper ISMA Centre, University of Reading.

Brown, S., W. Goetzmann and R. Ibbotson (1999), Offshore Hedge Funds: Survival & Performance, Journal of Business, Vol. 72, pp. 91-117.

Chen, Z. and P. Knez (1996), Portfolio Performance Measurement: Theory and Applications, Review of Financial Studies, Vol. 9, pp. 511-555.

Chernov, M. (2000), A Case of Empirical Reverse Engineering: Estimation of the Pricing Kernel, Working Paper Columbia Business School.

Cox, J. and C. Huang (1989), Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process, Journal of Economic Theory, Vol. 49, pp. 33-83.

Cox, J. and H. Leland (2000), On Dynamic Investment Strategies, Journal of Economic Dynamics and Control, Vol. 24, pp. 1859-1880.

Dahlquist, M. and P. Soderlind (1999), Evaluating Portfolio Performance with Stochastic Discount Factors, Journal of Business, Vol. 72, pp. 347-383.

Dybvig, P. (1988a), Distributional Analysis of Portfolio Choice, Journal of Business, Vol. 61, pp. 369-393.

Dybvig, P. (1988b), Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market, The Review of Financial Studies, Vol. 1, pp. 67-88.

Fung, W. and D. Hsieh (1997), Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds, Review of Financial Studies, Vol. 10, pp. 275-302.

Fung, W. and D. Hsieh (2000), Performance Characteristics of Hedge Funds and Commodity Funds: Natural vs. Spurious Biases, Journal of Financial and Quantitative Analysis, Vol. 35, pp. 291-307.

Fung, W. and D. Hsieh (2001), The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers, Review of Financial Studies, Vol. 14, pp. 313-341.

Glosten, L. and R. Jagannathan (1994), A Contingent Claim Approach to Performance Evaluation, Journal of Empirical Finance, Vol. 1, pp. 133-160.

Harrison, J. and D. Kreps (1979), Martingales and Arbitrage in Multi-Period Securities Markets, Journal of Economic Theory, Vol. 20, pp. 381-408.

He, J., L. Ng and C. Zhang (1999), Asset Pricing Specification Errors and Performance Evaluation, European Finance Review, Vol. 3, pp. 205-232.

Hoggard, T., E. Whalley and P. Wilmott (1994), Hedging Option Portfolios in the Presence of Transaction Costs, in: D. Chance and R. Trippi (eds), Advances in Futures and Options Research, Vol. 7, pp. 21-35.

Jarque, C. and A. Bera (1987), A Test for Normality of Observations and Regression Residuals, International Statistical Review, Vol. 55, pp. 163-172.

Jensen, M. (1968), The Performance of Mutual Funds in the Period 1945-1964, Journal of Finance, Vol. 23, pp. 389-416.

KPMG Consulting (1998), The Coming Revolution of the Hedge Fund Industry: A Case for Growth and Restructuring, KPMG Peat Marwick LLP (downloadable from www.rrcm.com).

Leland, H. (1985), Option Pricing and Replication with Transaction Costs, Journal of Finance, Vol. 40, pp. 1283-1301.

Liang, B. (1999), On the Performance of Hedge Funds, Financial Analysts Journal, July/August, pp. 72-85.

Liang, B. (2001), Hedge Fund Performance: 1990-1999, Financial Analysts Journal, January/February, pp. 11-18.

Lo, A. (2001), Risk Management for Hedge Funds: Introduction and Overview, forthcoming Financial Analysts Journal.

Merton, R. (1981), On Market Timing and Investment Performance. I. An Equilibrium Theory of Value for Market Forecasts, Journal of Business, Vol. 54, pp. 363-406.

Mitchell, M and T. Pulvino (2001), Characteristics of Risk and Return in Risk Arbitrage, Journal of Finance, Vol. 56, pp. 2135-2175.

Pelsser, A. and T. Vorst (1996), Transaction Costs and Efficiency of Portfolio Strategies, European Journal of Operational Research, Vol. 91, pp. 250-263.

President's Working Group on Financial Markets (1999). Hedge Funds, Leverage and the Lessons of Long-Term Capital Management, Report of The President's Working Group on Financial Markets, April 28, 1999 (downloadable from www.cftc.gov).

Sharpe, W. (1966), Mutual Fund Performance, Journal of Business, Vol. 39, pp. 119-138.

Sharpe W. (1992), Asset Allocation: Management Style and Performance Measurement, Journal of Portfolio Management, Vol. 18, No. 2, pp. 7-19.

Table 1: Hedge Fund Classification

This table classifies the 77 individual hedge funds in our data set into different categories.

	Number of Hedge Funds		
Category	Total	Onshore	Offshore
Market neutral	11	11	0
Global	28	16	12
Event driven	15	11	4
Fund of funds	23	11	12

Table 2: Hedge Fund Index Efficiency

In this table the second and third column show the alpha and Sharpe ratio of the 13 hedge fund indices in our sample based on monthly return data from May 1990 to April 2000. The relevant benchmark values are 0 and 0.28 (S&P 500) respectively. The fourth column shows the annual, i.e. monthly times 12, efficiency loss (-) or gain (+) according to the efficiency test. The fifth column shows how high/low implied volatility needs to be for the value of the indices' payoff distribution to reach 100. DNT 100 stands for 'does not touch 100'. This means that there is no volatility for which the price of the hedge fund payoff distribution is equal to 100. The last column shows the correlation between the actual hedge fund index return and the return calculated from the payoff function resulting from the efficiency test.

	Alpha (%)	Sharpe Ratio	Efficiency Yearly	Volatility	Correlation
EVENT DIST	0.4591	0.3495	-2.8685	DNT 100*	0.4920
EVENT RISK	0.5357	0.4751	-1.1157	0.19	0.3777
EVENT-DRIV	0.4713	0.5429	-1.6807	DNT 100*	0.3862
FUND DIV	0.3283	0.3471	-3.6614	DNT 100*	0.4515
FUND NICHE	0.2464	0.2865	-4.8756	DNT 100*	0.5011
FUNDOFFUND	0.2971	0.3425	-3.9225	0.05*	0.4784
GL EMER	0.1965	0.1714	-7.2672	0.09*	0.4893
GL EST	0.5445	0.4018	-0.7388	0.19	0.6844
GL INTL	0.4501	0.3526	-2.4846	0.16	0.4669
GL MACRO	0.4998	0.3510	-3.0333	0.18	0.4308
MKT ARB	0.7443	0.3720	0.1380	0.29*	-0.0777
MKT LONG	0.3901	0.8629	-2.4836	DNT 100*	0.3045
MKTNEUTRAL	0.4319	1.0865	-1.8591	DNT 100*	0.3574

* Inverse relationship between implied volatility and value payoff distribution

Table 3: Hedge fund Efficiency

In this table the second and third column show the alpha and Sharpe ratio of the 77 individual hedge funds in our sample based on monthly return data from May 1990 to April 2000. The relevant benchmark values are 0 and 0.28 (S&P 500) respectively. The fourth column shows the annual, i.e. monthly times 12, efficiency loss (-) or gain (+) according to the efficiency test. The fifth column shows how high/low implied volatility needs to be for the value of the funds' payoff distribution to reach 100. DNT 100 stands for 'does not touch 100'. This means that there is no volatility for which the price of the hedge fund payoff distribution is equal to 100. The last column shows the correlation between the actual hedge fund index return and the return calculated from the payoff function resulting from the efficiency test.

	Alpha (%)	Sharpe Ratio	Efficiency Yearly	Volatility	Correlation
Fn1	-0.0396	0.0568	-11.8074	DNT 100*	0.4435
Fn2	0.1670	0.1594	-8.4584	0.16	0.4153
Fn3	-0.2256	-0.2349	-33.5125	DNT 100*	-0.5951
Fn4	0.8592	0.3843	-1.5647	0.25	0.1724
Fn5	0.4866	0.2819	-4.7847	0.21	0.4908
Fn6	0.7297	0.3472	-3.6023	0.18	0.0151
Fn7	0.2255	0.1419	-9.7881	DNT 100*	0.3042
Fn8	0.5151	0.3164	-3.5152	0.04*	0.5820
Fn9	0.2830	0.2166	-6.7306	DNT 100*	0.2205
Fn10	0.3262	0.2740	-4.7176	0.10*	0.6916
Fn11	0.3557	0.2721	-5.0544	0.08*	0.4894
Fn12	0.3331	0.0721	-13.9695	0.20	0.0248
Fu1	0.1056	0.1786	-6.9299	0.06*	0.4234
Fu2	0.5449	0.3910	-3.0182	0.08*	0.2115
Fu3	0.1015	0.2136	-5.9387	0.11*	0.5405
Fu4	0.1979	0.2382	-6.0076	0.20	0.2435
Fu5	-0.4072	-0.0660	-15.5024	DNT 100*	0.2624
Fu6	0.3806	0.2465	-4.9686	0.21	0.0907
Fu7	0.5822	0.3881	-2.7776	0.15	0.2026
Fu8	0.3183	0.2764	-4.8574	0.17	0.6398
Fu9	0.1335	0.2213	-5.6178	0.08*	0.5076
Fu10	0.3684	0.2674	-4.9446	0.17	0.4251
Fu11	0.2101	0.2647	-4.7463	0.04*	0.3947
En1	0.4362	0.2629	-5.0871	0.15	0.2456
En2	0.7832	0.1706	-11.3449	0.14	0.3573
En3	0.4043	0.4998	-2.8266	0.17	0.2631
En4	0.5296	0.4657	-2.8010	0.19	0.1958
Eu1	0.5547	0.4465	-2.7851	0.03*	0.1767
Eu2	0.9556	0.3474	-1.4041	0.20	0.2006
Eu3	0.6881	0.4288	-1.7720	0.18	0.2336
Eu4	0.6628	0.3410	-3.1087	0.17	0.1186
Eu5	0.5122	0.3974	-2.9212	0.08*	0.2688
Eu6	0.4647	0.2552	-5.6319	0.19	0.0338
Eu7	0.4978	0.2973	-3.5697	0.16	0.2753
Eu8	0.4985	0.2531	-6.0954	0.19	0.2638

Eu9	0.6329	0.4490	-1.2385	0.11*	0.3701
Eu10	0.7341	0.2715	-5.3747	0.18	0.2578
Eu11	1.0324	0.3274	-0.4999	0.25	0.2943
Gn1	1.3852	0.3213	-2.0539	0.15	0.0498
Gn2	0.7954	0.1602	-10.1198	0.05*	0.2038
Gn3	-0.0551	0.0730	-15.6298	0.21	0.2625
Gn4	0.2613	0.2711	-3.6516	0.09*	0.6148
Gn5	0.7725	0.3123	-1.5353	0.15	0.3052
Gn6	-0.0610	0.1870	-9.4891	0.18	0.6693
Gn7	0.4158	0.3460	-3.3880	0.19	0.2816
Gn8	1.0613	0.3272	1.2936	0.29*	0.2938
Gn9	0.4588	0.1989	-8.6432	0.04*	-0.0491
Gn10	0.3519	0.0671	-15.8663	DNT 100*	-0.0544
Gn11	0.7148	0.1362	-8.9879	DNT 100*	0.1530
Gn12	0.1059	0.1791	-9.4294	DNT 100*	0.4642
Gu1	0.4859	0.1269	-13.7170	0.17	0.1235
Gu2	1.0360	0.2486	-5.5885	0.13	0.3828
Gu3	-0.0874	0.0893	-22.3700	0.23	0.3737
Gu4	0.7289	0.2739	-4.2534	0.15	0.4390
Gu5	0.8257	0.2856	-3.6816	0.16	0.4092
Gu6	0.0384	0.1765	-9.9442	0.21	0.6025
Gu7	0.3643	0.2566	-5.3228	0.05*	0.5553
Gu8	0.9358	0.5159	2.1847	DNT 100	0.6250
Gu9	0.5504	0.2690	-4.7188	0.18	0.4788
Gu10	1.6489	0.3069	1.6055	DNT 100	0.2257
Gu11	0.6242	0.4381	-1.8923	0.06*	0.2122
Gu12	0.2268	0.2154	-8.8535	0.18	0.6012
Gu13	0.2433	0.2395	-6.8602	0.03*	0.6670
Gu14	0.0972	0.1331	-17.2618	0.25	0.4052
Gu15	1.0023	0.3024	0.7961	0.28*	0.2756
Gu16	0.0456	0.1611	-10.9957	DNT 100*	0.6500
Mu1	0.1013	0.0476	-13.0194	0.24	0.1472
Mu2	0.0865	0.2003	-7.6970	0.16	0.5719
Mu3	0.2848	0.1431	-6.0134	0.11*	0.2817
Mu4	0.2634	0.1661	-9.9971	0.05*	0.3104
Mu5	0.3881	0.3740	-3.8130	0.25	0.1512
Mu6	0.0464	0.1933	-6.6852	0.19	0.4362
Mu7	0.9054	0.6846	1.5556	0.45*	0.1458
Mu8	0.1907	0.3526	-5.3320	0.15	0.2888
Mu9	0.0333	0.1323	-6.8739	0.07*	0.2009
Mu10	1.6204	0.2643	-2.2178	0.20	-0.0469
Mu11	0.4312	0.2326	-6.3766	0.02*	0.4886

* Inverse relationship between implied volatility and value payoff distribution

Table 4: Hedge Fund Efficiency: Summary

This table summarizes the efficiency test results on hedge fund indices and individual hedge funds. With regard to the latter we distinguish between the following categories: Fund of Funds (FOF), Non-Fund of Funds (Non FOF), Global, Market Neutral, Event Driven, Offshore and US Based.

	Overall		Efficient		Inefficient	
	No.	Avg. Yearly	No.	Avg. Yearly	No.	Avg. Yearly
Indices	13	-2.7579	1	0.1380	12	-2.9992
FOF Indices	3	-4.1532	0	-	3	-4.1532
Non FOF Indices	10	-2.3393	1	0.1380	9	-2.6146
Individual Funds	77	-6.4171	5	1.4871	72	-6.9660
FOF	23	-7.5137	0	-	23	-7.5137
Non FOF	54	-5.9501	5	1.48708	49	-6.7088
Global	28	-7.0848	4	1.4700	24	-8.5106
Market Neutral	11	-6.0427	1	1.5556	10	-6.8025
Event Driven	15	-3.7641	0	-	15	-3.7641
Offshore	28	-7.7523	1	1.2936	27	-8.0874
US Based	49	-5.6542	4	1.5354	45	-6.2932

Table 5: Hedge Funds Index Portfolio Efficiency

This table shows the efficiency test results on portfolios constructed by investing in varying proportions in the hedge fund indices and the S&P 500. The second column shows the correlation between the hedge fund index returns and S&P 500 returns. Column 3 shows the maximum efficiency level achieved. Column 4 gives the percentage to be invested in the S&P 500 to obtain that maximum efficiency level.

	Correlation with S&P 500	Maximum Efficiency	Investment in S&P 500
EVENT DIST	0.4918	0.2790	79%
EVENT RISK	0.3795	0.3142	81%
EVENT-DRIV	0.4200	0.1936	78%
FUND DIV	0.4358	0.0000	100%
FUND NICHE	0.4398	0.0000	100%
FUNDOFFUND	0.4381	0.0000	100%
GL EMER	0.4663	0.3844	83%
GL EST	0.7056	0.7331	82%
GL INTL	0.4699	0.3475	81%
GL MACRO	0.4017	0.1363	78%
MKT ARB	0.0850	0.6032	77%
MKT LONG	0.2342	0.0000	100%
MKTNEUTRAL	0.2777	0.0000	100%

Table 6: Hedge Fund Portfolio Efficiency

This table shows the efficiency test results on portfolios constructed by investing in varying proportions in the hedge funds and the S&P 500. The second column shows the correlation between hedge fund returns and S&P 500 returns. Column 3 shows the maximum efficiency level achieved. Column 4 gives the percentage to be invested in the S&P 500 to obtain that maximum efficiency level.

	Correlation with S&P 500	Maximum Efficiency	Investment in S&P 500
Fn1	0.2449	0.0000	100%
Fn2	0.3460	0.1234	96%
Fn3	-0.6787	0.0000	100%
Fn4	0.1224	1.0053	80%
Fn5	0.4297	0.3420	82%
Fn6	-0.0621	0.3866	78%
Fn7	0.2444	0.0000	100%
Fn8	0.5136	0.4878	78%
Fn9	0.2152	0.0000	100%
Fn10	0.6534	0.3004	78%
Fn11	0.5013	0.4012	96%
Fn12	-0.0519	0.0000	100%
Fu1	0.3888	0.3024	96%
Fu2	0.1502	0.2912	81%
Fu3	0.5090	0.0000	100%
Fu4	0.1609	0.0000	100%
Fu5	0.2523	0.0000	100%
Fu6	0.1807	0.0000	100%
Fu7	0.2356	0.3282	78%
Fu8	0.5666	0.2191	96%
Fu9	0.3230	0.1163	98%
Fu10	0.4592	0.0953	89%
Fu11	0.1666	0.0000	100%
En1	0.2906	0.0555	94%
En2	0.3387	0.7931	89%
En3	0.3370	0.2285	97%
En4	0.1198	0.0938	82%
Eu1	0.0630	0.2149	76%
Eu2	0.2164	1.3597	78%
Eu3	0.2268	0.6177	80%
Eu4	0.1483	0.4611	78%
Eu5	0.2075	0.1318	93%
Eu6	0.0877	0.0516	91%
Eu7	0.3253	0.1987	80%
Eu8	0.3012	0.2290	79%
Eu9	0.4049	0.6641	80%
Eu10	0.2838	0.9104	78%
Eu11	0.1916	1.8296	78%
Gn1	-0.0067	1.9820	77%
Gn2	0.1805	0.5607	90%
Gn3	0.2953	0.0000	100%
Gn4	0.7097	0.3487	89%

Gn5	0.3869	1.0281	78%
Gn6	0.7269	0.1217	96%
Gn7	0.3139	0.1911	96%
Gn8	0.3722	2.2509	79%
Gn9	-0.1058	0.2729	96%
Gn10	-0.0555	0.1805	96%
Gn11	-0.0015	0.3810	96%
Gn12	0.4905	0.0000	100%
Gu1	0.1717	0.4147	96%
Gu2	0.2503	1.2430	76%
Gu3	0.3500	0.0000	100%
Gu4	0.4583	1.1370	82%
Gu5	0.4418	1.3494	81%
Gu6	0.6044	0.0000	100%
Gu7	0.5818	0.4580	76%
Gu8	0.6184	2.1847	0%
Gu9	0.4896	0.7908	79%
Gu10	0.1796	2.8748	78%
Gu11	0.2506	0.4995	74%
Gu12	0.6345	0.3215	90%
Gu13	0.6861	0.4677	89%
Gu14	0.4294	0.5209	96%
Gu15	0.3585	2.1840	77%
Gu16	0.6667	0.1512	96%
Mu1	0.0358	0.1888	97%
Mu2	0.5819	0.2436	96%
Mu3	-0.0451	0.3093	96%
Mu4	0.3703	0.3604	96%
Mu5	0.2606	0.1798	96%
Mu6	0.3609	0.3956	96%
Mu7	0.2654	1.8706	15%
Mu8	0.1795	0.5204	96%
Mu9	0.2150	0.2867	96%
Mu10	-0.0755	2.1405	78%
Mu11	0.5565	0.8625	76%

Figure 1: Cumulative Probability Distribution

This figure shows the cumulative probability distribution of the end-of-month payoff of a hedge fund and the S&P 500. Both graphs are based on monthly return data from May 1990 to April 2000. S&P 500 returns are assumed to be normally distributed.

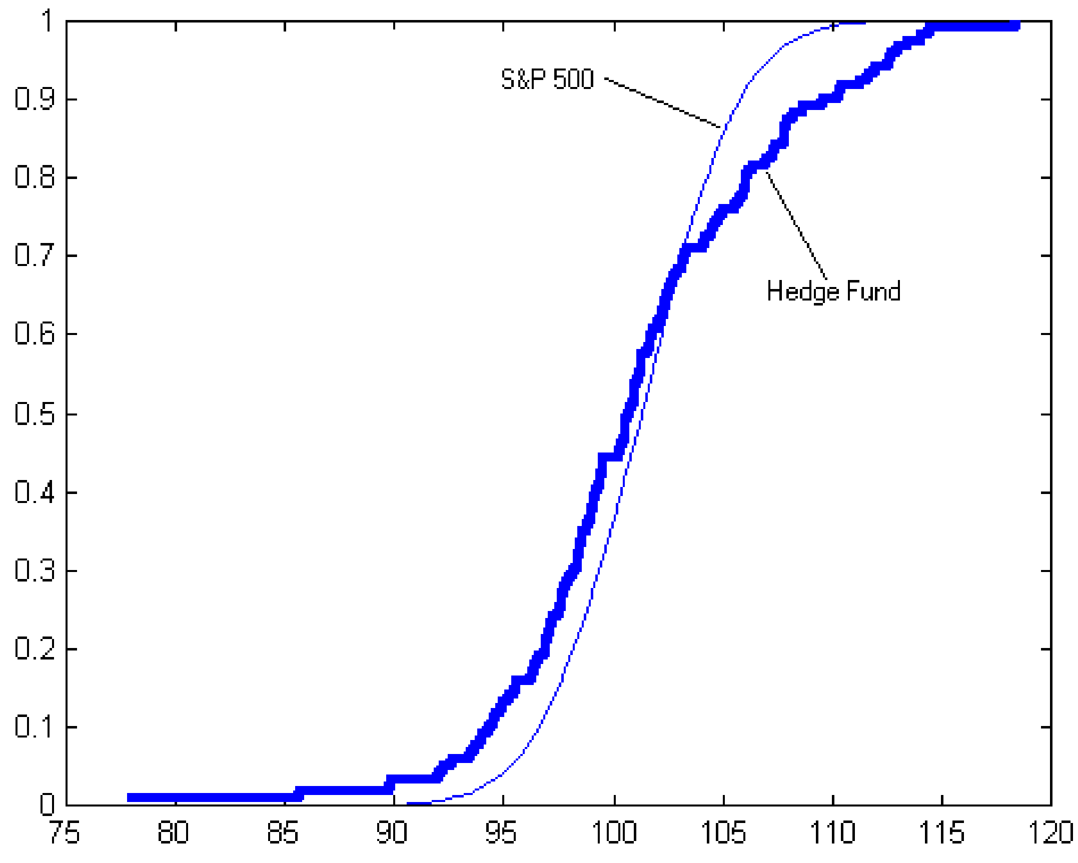

Figure 2: Payoff Function

This figure shows an example of a payoff function resulting from the mapping procedure discussed in Section III. Given the S&P 500 distribution, this payoff function implies the same payoff distribution as offered by the hedge fund (index) in question.

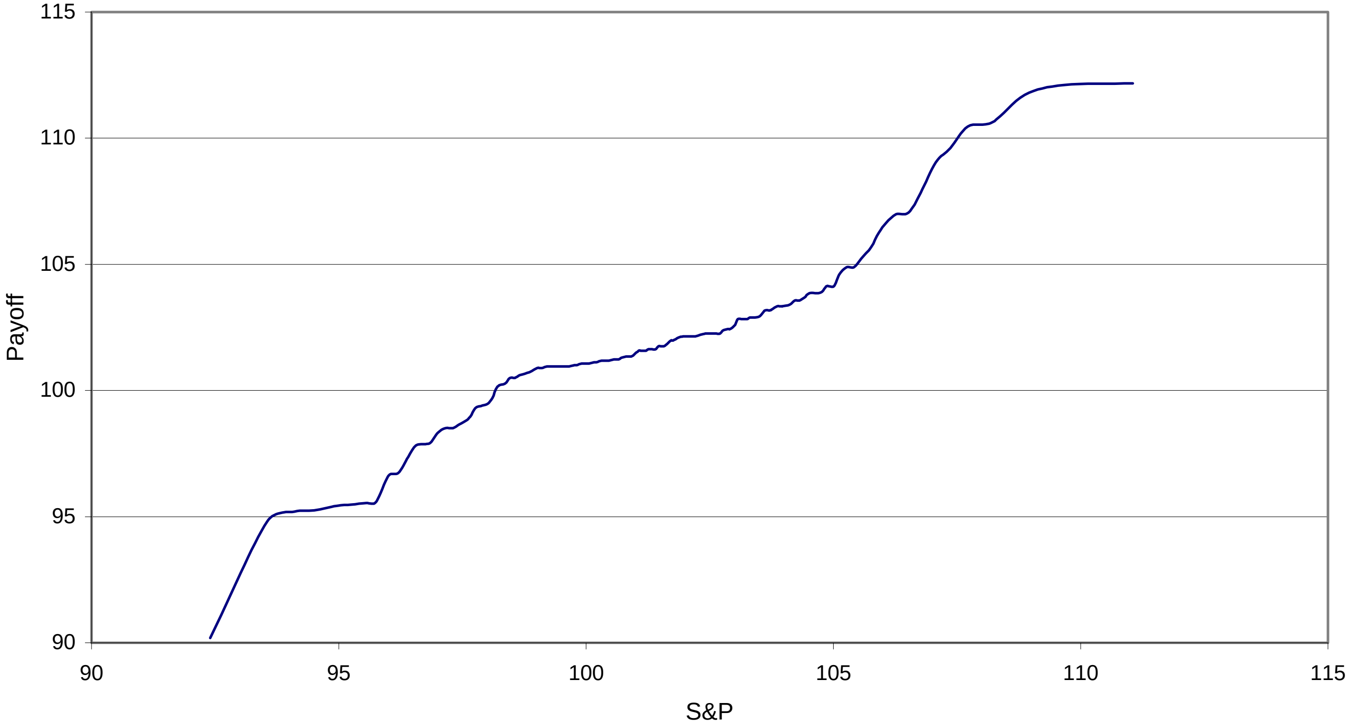

Figure 3: Sampling Error

This figure shows the frequency distribution of the annualised errors from performing the efficiency test on a combination of the index and a short at-the-money call. To calculate the errors we first sampled 120 end-of-month index values, assuming a monthly mean return of 1.24% and a standard deviation of 3.59%, and calculated the corresponding payoffs from the combination (including 0.22% dividend yield). Subsequently, we applied the efficiency test to these data and calculated the sampling error as the difference between the actual test result and 100. To obtain the frequency distribution shown, this procedure was repeated 20,000 times. The normal distribution shown has the same mean (-0.05) and standard deviation (2.14) as the sampling error distribution.

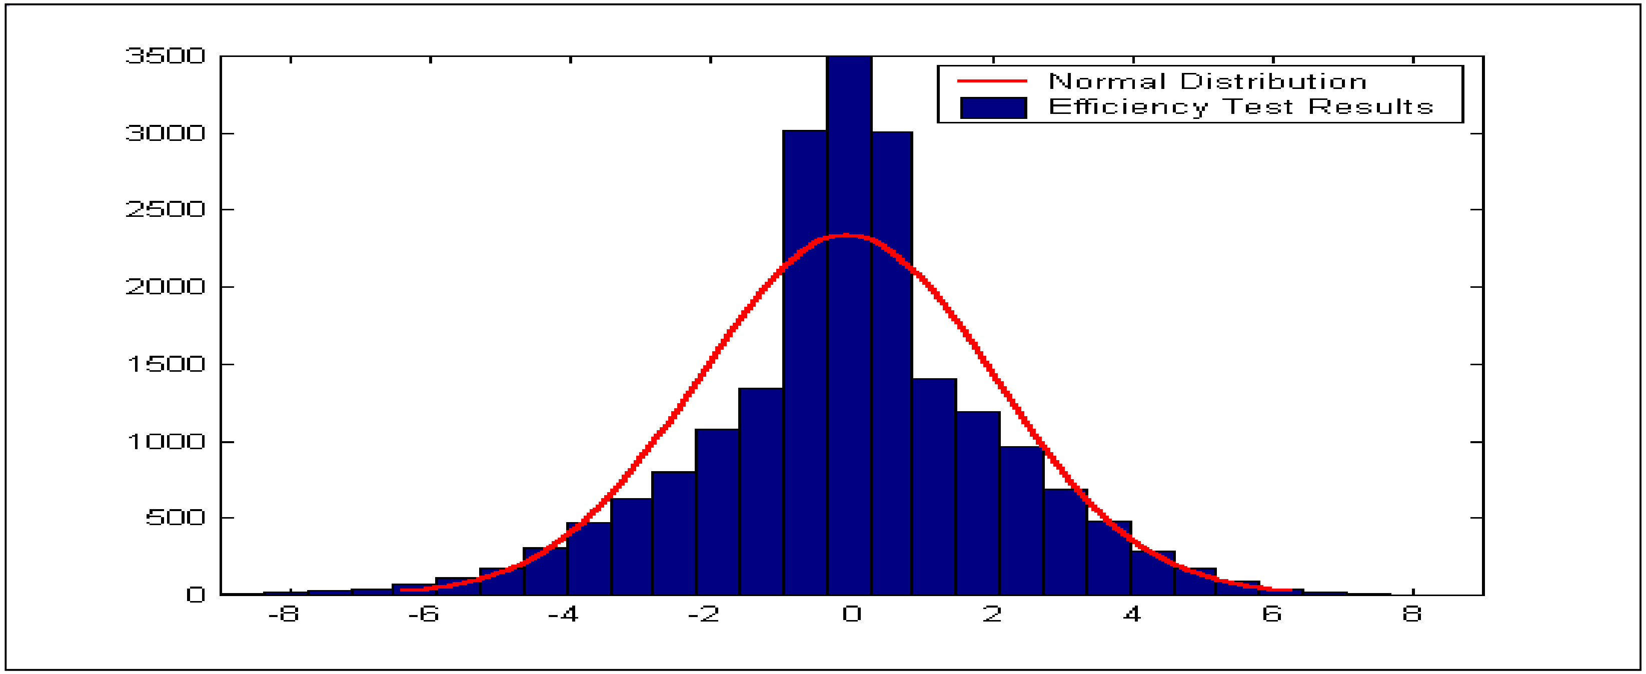

Figure 4: Replication Efficiency

This figure shows for 13 hedge fund indices the mean, standard deviation, skewness and kurtosis of the distribution of the actual hedge fund index return over the period from May 1990 to April 2000 as well as the return calculated from the payoff function resulting from the efficiency test.

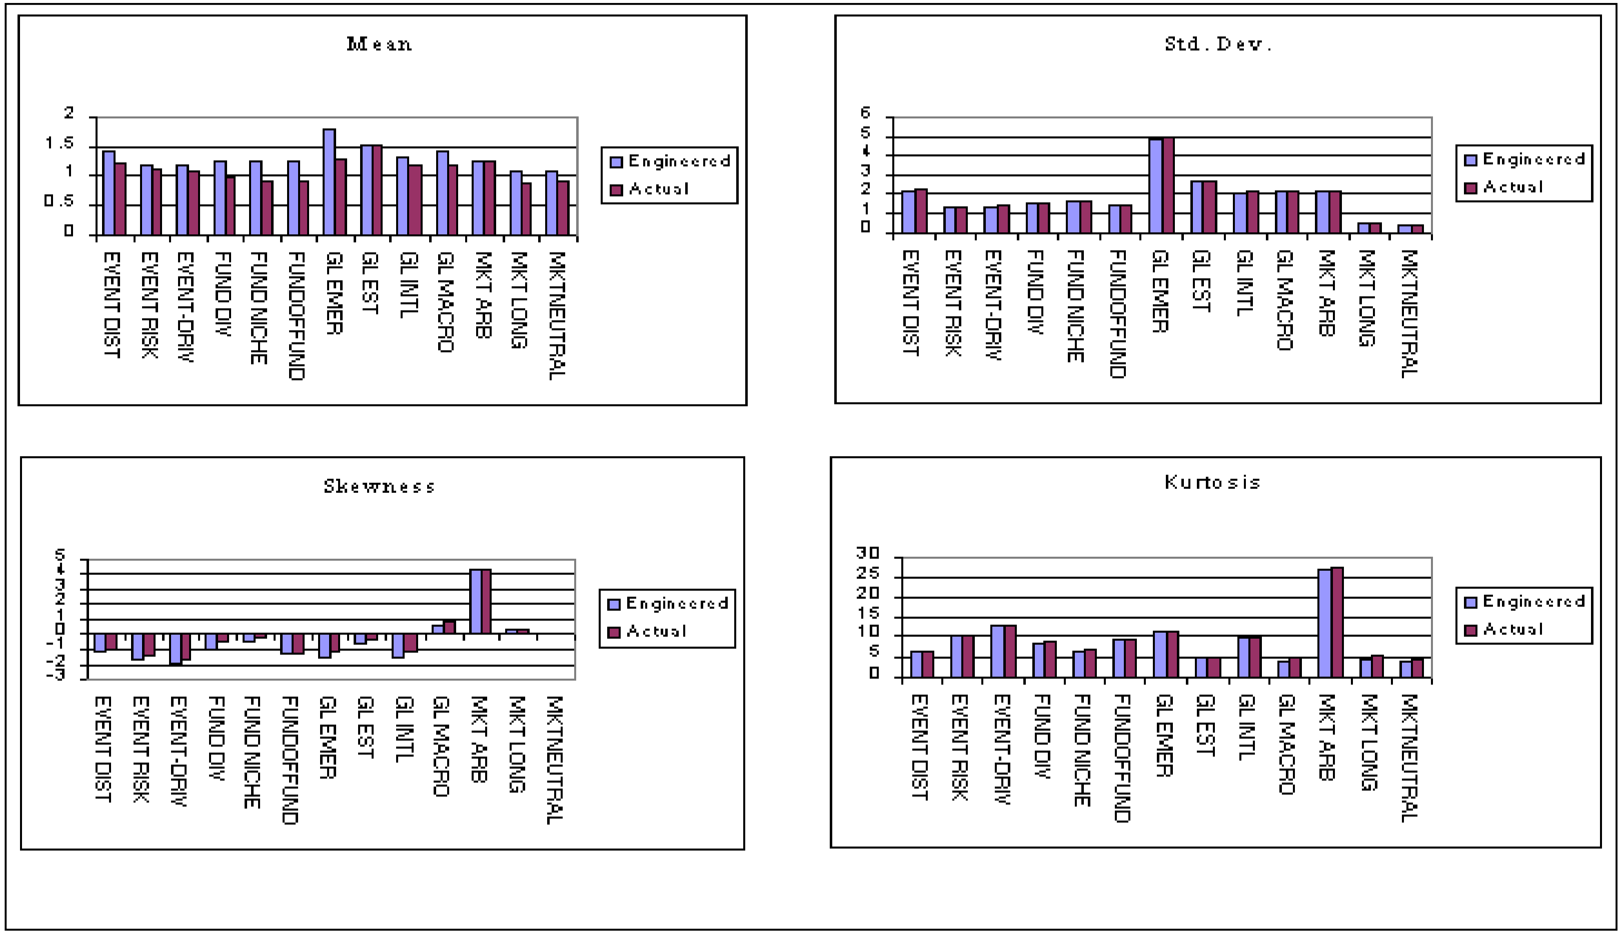

Figure 5: Hedge Fund Efficiency

This figure shows the annual efficiency loss (-) or gain (+) of 77 individual hedge funds based on monthly return data from May 1990 to April 2000. The entries are taken from the fourth column of table 3.