Replicating Hedge Fund Returns Using Futures

A European Perspective

Harry M. Kat
Professor of Risk Management and Director Alternative Investment Research Centre,
Cass Business School, City University, London.

Helder P. Palaro
PhD Student, Cass Business School, City University, London.

March 2006

Abstract
Hedge funds tend to put a lot of effort into generating their returns and charge substantial fees to do so. However, with the latest performance evaluation studies indicating that hedge fund performance is not truly superior (anymore), the question arises whether it is possible to generate similar, hedge fund-like, returns in a more mechanical way and with less effort. In other words, is it possible to design dynamic trading strategies, mechanically trading stocks, bonds, etc., that generate hedge fund-like returns? If such strategies indeed exist, then this would solve a respectable number of problems surrounding hedge fund investments and alternative investments in general, including the need for extensive due diligence, liquidity, capacity, transparency, style drift and regulatory problems, as well as excessive management fees. In this paper we develop and demonstrate the workings of a technique that allows the derivation of dynamic futures trading strategies, which generate returns with statistical properties similar to hedge funds. Trading nothing else than Eurex DAX® 30 and Euro-Bund Futures, we show that this technique is not only capable of replicating fund of funds returns, but is equally capable of replicating individual hedge fund returns. Accurately replicating the risk-return profile, but sharing none of the drawbacks of real hedge funds, our synthetic hedge fund returns are a worthwhile alternative to direct hedge fund investment.

Please address all correspondence to: Harry M. Kat, Professor of Risk Management and Director Alternative Investment Research Centre, Cass Business School, City University, 106 Bunhill Row, London, EC2Y 8TZ, United Kingdom, Tel. +44.(0)20.70408677, E-mail: Harry@airc.info or Investors@eurexchange.com

The authors like to thank CSI and Rudi Cabral for allowing us to access the CSI futures database. Financial support from Eurex is gratefully acknowledged.

1. Introduction

Rising from relative obscurity, over the last 15 years hedge funds have become increasingly popular with high net worth individuals as well as institutional investors. As a result, the number of hedge funds has risen from around 500 in 1990 to an estimated 9,000 in 2006. Over the same period, assets under management are estimated to have increased from USD 50 billion to USD 1 trillion. Apart from the success of the prime brokerage concept, one of the crucial factors behind the spectacular growth of the hedge fund industry has been the rise of the fund of funds structure as the preferred way of investing in hedge funds. Currently, most money invested in hedge funds flows through funds of funds, with the total number of such funds and structured products linked to them being estimated at around 4,000.

Initially, hedge funds were sold on the promise of superior performance, the story being that hedge fund managers' long experience and proven investment skills were a virtual guarantee for superior returns. Especially high net worth investors proved sensitive to these arguments and fuelled much of the early growth of the industry. Towards the end of the 1990s the story began to change, however. No longer were hedge funds sold on the promise of superior performance, but more and more on the basis of a diversification argument, pointing at hedge funds' relatively low correlation with stocks and bonds and the beneficial effects on risk and return from including hedge funds in a traditional investment portfolio. The reason for this rather remarkable change in sales tactics was twofold. Firstly, starting in the late 1990s, hedge fund performance took a turn for the worst, with every next year being worse than the year before. Secondly, driven by historically low interest rates and substantial losses in the equity markets, institutional investors, keen to be seen taking action, started to consider hedge funds more seriously as well. Given institutions' emphasis on risk management, the hedge fund story changed to accommodate this new clientele[1].

Hedge fund managers typically put a lot of effort into generating their returns. However, with the industry implicitly admitting, and recent performance evaluation studies confirming[2], that hedge fund performance is not truly superior (anymore), the question arises whether it is possible to generate similar returns in a more mechanical way and with less effort. More precisely, is it possible to design dynamic trading strategies, mechanically trading stocks, bonds, etc., that generate hedge fund-like returns? If indeed we could design such strategies, this would solve a respectable number of problems surrounding hedge funds as well as many other "alternative" investments, including:

The Need for Extensive Due Diligence

Without any publicly available information and research, investors are forced to invest substantial amounts of time and energy in visiting hedge fund managers, asking questions, interpreting the answers and doing background checks. A number of third parties do provide these services on a stand-alone basis, but at significant costs.

[1] The current trend is to introduce retail investors to hedge funds as well. Since the typical retail investor is unlikely to appreciate the special nature of hedge fund investment, this will intensify the call for more profound regulation, which in turn will force the industry to reshape itself once again.

[2] See for example Amin and Kat (2003), Bailey et al. (2004), Fung et al. (2005) or Kat and Palaro (2006a, 2006b).

Lack of Liquidity

Most hedge funds use lock-up structures to tie in new investors for periods ranging from 6 months up to 5 years[3]. After the lock-up has expired, investors typically need to give one or three months notice if they want to disinvest. In addition, some funds charge departing investors an additional fee of up to 5% to compensate remaining investors for the costs of having to rebalance the fund portfolio. However, it is hard to see why a fund would require an exit fee if there is already a notice period in place. Given proper notice, freeing up money should not cost an arm and a leg. Imposing an exit fee therefore seems nothing more than a subtle way of extending the lock-up period.

Lack of Transparency

All hedge funds claim to do something highly exclusive and proprietary and anxiously guard their trading secrets. Although transparency has improved with the arrival of institutional investors, hedge fund investors are seldom told exactly what goes on inside the black box. As a result, it can sometimes be very hard to properly assess the risk-return characteristics of a fund[4].

Lack of Capacity

In an attempt to preserve the level of returns, successful hedge funds may close for new investors or close for new money altogether[5]. However, this does not prevent money from flowing to other managers in the same category. As a result, when opportunities are in limited supply, performance may come under pressure. This is especially true for arbitrage-type strategies, where the arrival of more money and managers will significantly increase market efficiency. Recently, convertible bond arbitrage has suffered quite badly from this form of over-investment, reporting an average return of –1.92% for 2005 (HFRI Convertible Arbitrage Index).

Excessive Management Fees

The average hedge fund charges its investors “2 plus 20”, i.e. a flat management fee of 2% plus an additional incentive fee equal to 20% of any profits over a hurdle rate. Funds of funds tend to charge an additional “1 plus 10” on top of this. With interest rates and hedge fund performance at historically low levels, this means that nowadays pre-fee hedge fund returns are split more or less equally between investors and fund (of funds) managers.

Style Drift

Hedge fund managers may sometimes change their style or strategy, which in turn may cause a significant change in a fund’s risk-return profile. When not explicitly notified of this change and without sufficient transparency, investors can only find out about this from the returns that the fund generates. With returns reported on a monthly basis, however, it could take a long time before it becomes clear that something has changed.

[3] Investors are becoming increasingly resistant to lock-up periods. According to Dyment et al. (2005), in 2004 68% of investors would only invest with managers with lock-ups of one year of less. In 2005, this rose to 77%.

[4] Dyment at al. (2005) report that only 14% of investors requires full transparency. Even more surprisingly, 19% of investors do not require any transparency at all.

[5] Anticipating closure, according to Dyment et al. (2005), in 2005 45% of investors required future capacity rights when investing in hedge funds.

Regulation
External and/or internal regulation may make it difficult, if not impossible, for (certain types of) investors to invest in hedge funds directly. In Germany for example, before the introduction of the new German Investment Act and Investment Tax Act in January 2004, hedge fund investment has been quite problematic for a long time.

Of course, we are not the first to attempt to replicate hedge fund returns. Following the work of Sharpe (1992) on equity mutual funds, previous authors have primarily relied on the use of factor models to replicate month-to-month returns[6]. In theory, the factor model approach should work well. Once the relevant risk factors have been identified and the fund's sensitivity to these factors has been determined, one can construct a portfolio of stocks, bonds, and other securities with the same factor sensitivities as the fund in question. Since it has the same factor sensitivities, the resulting portfolio will generate returns that are similar to those of the fund.

The problem when applying the above approach in a hedge fund context is that in practice we often have little idea how hedge fund returns are actually generated, i.e. which risk factors to use. As a result, factor models typically explain only 25–30% of the variation in individual hedge fund returns, which compares quite unfavourably with the 90–95% that is typical for mutual funds. Although the procedure works better for portfolios of hedge funds, funds of funds and hedge fund indices, where much of the idiosyncratic risk is diversified away, factor models do not appear to offer a particularly fruitful alternative when looking to replicate hedge fund returns accurately[7].

Given the failure of the factor model approach, we take a step back and reconsider the problem at hand. When an investor likes a hedge fund, it is (or should be at least) because of the statistical properties of the fund's returns, i.e. their mean, standard deviation, etc. and their relationship with the returns on the existing portfolio. This implies that we do not necessarily have to replicate a fund's month-to-month returns. For most applications it will be enough if we can generate returns with the same statistical properties as the returns generated by the fund.

So far, there has only been one study, which followed this route. Based on the early theoretical work of Glosten and Jagannathan (1994) and Dybvig (1988a, 1988b), and primarily aiming to evaluate hedge fund performance, Amin and Kat (2003) developed mechanical trading strategies, trading the S&P 500 and cash, which aimed to generate returns with the same marginal distribution as the returns of a given hedge fund. Although interesting from a theoretical perspective, from a practical perspective only replicating the marginal distribution is not enough though. Most of today's investors are attracted to hedge funds because of their relatively weak relationship with traditional asset classes and their own

[6] See for example Schneeweis et al. (2003) or Agarwal and Naik (2004).

[7] Despite the lack of explanatory value, several parties have recently announced the intention to launch factor model based products that aim to provide investors with hedge fund-like returns at lower costs and in a more convenient format.

portfolio in particular. To properly replicate hedge fund returns we therefore not only have to replicate the marginal distribution, but also the relationship between a fund and the investor's existing portfolio. In this paper we develop a procedure, which does exactly that.

The basic idea behind the proposed procedure is straightforward. From the theory of dynamic trading it is well known that with complete markets one can perfectly hedge any payoff function. Therefore, if we can find a payoff function which, given the probability distribution of the underlying index or indices, implies the desired return distribution, we will also have found the dynamic trading strategy which generates returns that are drawings from that distribution.

Of course, there are a number of hurdles to take. First, we are not interested in just any strategy. To maximize expected return, we want the cheapest strategy possible. Second, since we are aiming to replicate not only a fund's marginal return distribution but also its relationship with the investor's existing portfolio to which the fund is going to be added, we are confronted with bivariate distributions, which can take on a large variety of shapes and forms. Third, real markets are less well behaved than typically assumed in theoretical models. As a result, an inconsistency may arise between the determination of the desired payoff function, which is a purely empirical matter, and the subsequent derivation of the dynamic trading strategy generating that payoff. A second consequence of relying on an abstract model is that in practice our dynamic trading strategies may not be able to exactly generate the desired payoff. We therefore perform extensive out-of-sample tests of our strategies, using daily data over the period 1990–2004.

The replication procedure concentrates on replicating a fund's risk profile without explicitly considering the fund's expected return. The underlying assumption is that, in an efficient market, in the longer run investors will receive a return in line with the risk that they have taken. This is why the empirical finding that hedge fund returns are not truly superior is fairly crucial. If they were superior, we would still be able to replicate their risk profile, but we could not expect to replicate their average as well. If it is superior, it can't be replicated and vice versa. The latter observation points at another application of the replication technique developed in this paper: the evaluation of hedge fund returns. Explicitly constructed to offer the same risk profile, when the average replicated return is significantly higher than the average fund return, the fund is the inefficient alternative. We investigate this line of thought further in Kat and Palaro (2006a, 2006b).

The paper proceeds as follows. In the next section we discuss the determination of the desired payoff function, i.e. the payoff function, which, given the distribution of the assets to be traded, implies the desired return distribution. In section 3 we discuss the practical implementation of the procedure and the results of some out-of-sample tests, replicating the returns on three well-known hedge funds (of funds). Section 4 contains our conclusions.

2. Determination of the Replication Strategy

Apart from cash, which is purely used to move money through time, in our replication strategies we will trade two different risky assets. To replicate the relationship between the investor's existing portfolio and the fund to be replicated, we will need to trade (a good approximation of) the investor's portfolio. Second, as our main source of uncertainty, we trade what we will refer to as a "reserve asset". In theory this reserve asset could be anything. In practice, however, since the reserve asset acts as the main risk factor underlying the replication strategy, some choices make more sense than others. We need the reserve asset to be liquid so we can move in and out without too much hassle and costs. We also need daily volatility. A reserve asset that doesn't move is unable to fulfil its most important task: the supply of uncertainty. In addition, to help the statistical modelling behind the replication as well as the execution of the replication strategy, we need the returns on the reserve asset to be statistically well behaved.

Given the fund to be replicated, the investor's existing portfolio and the reserve asset, the replication procedure consists of a number of distinct steps. First, we collect return data on the fund, the investor's portfolio, and the reserve asset. Second, we analyse the data to infer the joint distribution of the fund return and the investor's portfolio return. We refer to this as the "desired distribution". We do the same for the joint distribution of the investor's portfolio return and the return on the reserve asset, which we refer to as the "building block distribution". Third, we determine the cheapest payoff function, which turns the building block distribution into the desired distribution. We call this the "desired payoff function". Fourth, we price the latter payoff function and derive the required allocations to the investor's portfolio and the reserve asset and their evolution over time. In this section we discuss the above steps in more detail.

Estimation of the Desired and Building Block Distributions

Recent research in finance has uncovered various deviations from not only univariate, but also multivariate normality[8]. One powerful, and at the same time convenient, way to model this is by the use of so-called "copulas", as it allows the decomposition of any n-dimensional joint distribution into n marginal distributions and a single copula function[9]. In the replication procedure we allow three different marginal distributions (Normal, Student-t and Johnson SU)[10] and six different bivariate copulas. The first two copulas are part of the class of elliptical copulas, since they are derived from elliptical distributions. The normal copula is extracted from the bivariate normal distribution. If we combine the bivariate normal copula with two normal marginal distributions, we end up with the bivariate normal distribution. However, if either one or both marginal distributions are non-normal, then the joint distribution produced will be a completely different distribution. The Student-t copula, which is extracted from the bivariate Student-t distribution, is also an elliptical copula, but it differs from the normal copula in that it allows for some extreme dependence in the lower and upper tails. However, since the Student-t copula is symmetric, this dependence must be the same for both tails.

8 Longin and Solnik (2001) for example find clear evidence of asymmetric dependence in international equity markets. A similar conclusion can be found in Ang and Chen (2002) with respect to US stocks.

9 Copulas have been widely used in the statistical literature. Joe (1997) and Nelsen (1999) provide a good introduction. Cherubini et al. (2004) discuss copulas in a finance context.

10 See Johnson (1949, 1965) for details on the Su distribution.

The next three families of copulas, Gumbel, Cook-Johnson and Frank, are part of the Archimedean copulas class, a rich class of copulas that allow for very different types of dependence. The "Gumbel copula" is asymmetric. It has more dependence in the upper tail than in the lower tail. The "Cook-Johnson copula", also known as the Clayton copula, is also asymmetric, but with more dependence in the lower tail than in the upper tail. As shown by Longin and Solnik (2001) and Ang and Chen (2002), this is quite common behaviour in equity market returns. The "Frank copula" implies the same dependence between positive returns as between negative returns. Like the Normal and Student-t copulas, it allows for positive and negative dependence. The sixth and final copula is the "symmetrised Joe-Clayton (SJC) copula", proposed by Patton (2005a). It is the most flexible of the copulas discussed here. It has two parameters, which separately control the dependence in the lower and upper tail. As a result, this copula can fit data with very different patterns of dependence in the tails.

Figure 1. Random Drawings from Various Copulas, Assuming Standard Normal Marginals and a Linear Correlation Coefficient of 0.7.

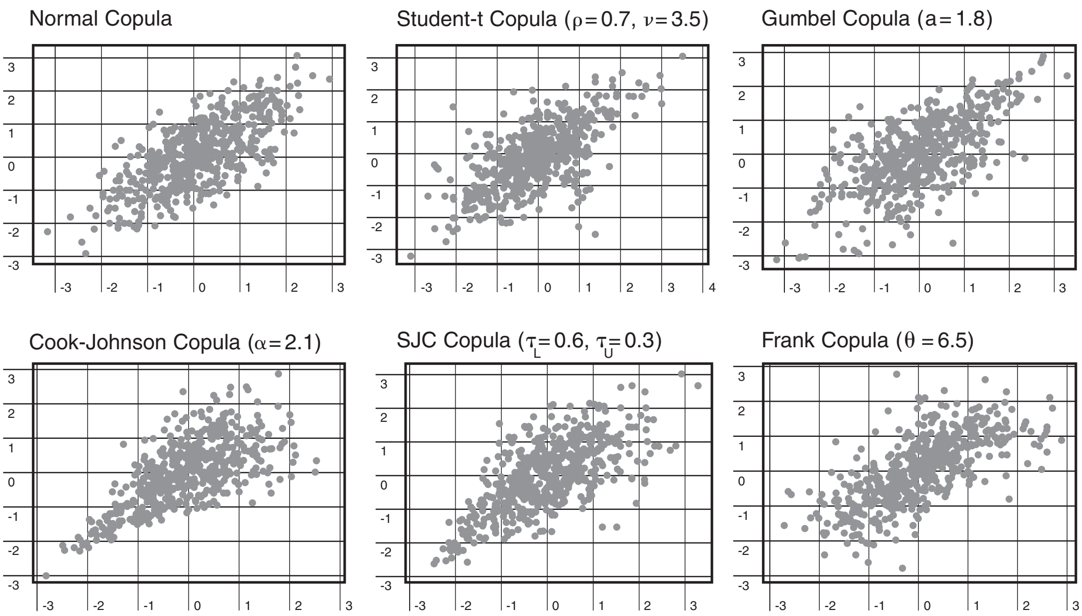

Figure 1 shows 500 simulated drawings from six bivariate joint distributions. In all cases, the marginal distributions are standard normal and the linear correlation is 0.7. Despite this, the plots show six different patterns of dependence, underlining the impact and different characteristics of each of the six copula families. Note that only in the bivariate normal case is the linear correlation coefficient sufficient to fully describe the observed dependence structure.

With three possible candidates for the marginal distribution and six for the copula, we have 54 possible joint distributions to choose from. To estimate these models we use the so-called Inference Functions for Margins (IFM) method[11]. Subsequently, we select the final model using the Akaike information criterion (AIC)[12]. We considered other selection criteria as well, including the quadratic distance between the estimated copula and the empirical copula for example. The advantage of the AIC, however, is that it penalises models with a large number of parameters.

Determination of the Desired Payoff Function

Having selected the desired and building block distributions, the next step is to determine the cheapest payoff function, which turns one into the other, i.e. the cheapest function g* such that:

$$P\,(S_P \leq \chi,\ g^*\,(S_P,\ S_R) \leq y) = P\,(S_P \leq \chi,\ S_I \leq y),\ \forall \chi,\ y \qquad (2)$$

with S_I denoting the end-of-month payoff of the fund, S_P the end-of-month payoff of the investor's portfolio, and S_R the end-of-month payoff of the reserve asset.

We start by assuming the current value of all assets is equal to 100. Rescaling to log-returns, this means looking for the cheapest function

$$g\,(\chi,\ y) = \log\left(\frac{g^*\,(100\exp(\chi),\ 100\exp(\gamma))}{100}\right)$$

such that:

$$P\,(X_P \leq \chi,\ g\,(X_P\ ,X_R) \leq y\,) = P\,(X_P \leq \chi,\ X_I \leq y\,) = F_{P,\,I}(\chi,\ y\,),\ \forall \chi,\ y \qquad (3)$$

with

$$X_I = \log\left(\frac{S_I}{100}\right),\ X_P = \log\left(\frac{S_P}{100}\right),\ \text{and}\ X_P = \log\left(\frac{S_P}{100}\right)$$

or equivalently, the cheapest function g such that:

$$P(g(X_P,\ X_R) \leq y \,|\, X_P = \chi) = P(X_I \leq y \,|\, X_P = \chi\,) = F_{I\,|\,P}\,(y\,|\,\chi),\ \forall \chi,\ y \qquad (4)$$

[11] Details on this method can be found in Xu (1996) and Patton (2005b).
[12] See Akaike (1973) for details.

Following a similar reasoning as in Dybvig (1988b) and with the help of some plausible assumptions with respect to the Sharpe ratio of the reserve asset and its correlation with the investor's portfolio, it can be shown that the cheapest payoff function will be a non-decreasing function of the reserve asset. The function g in expression (4) is therefore given by:

$$g(\chi, y), = F^{-1}_{I|P}\ (F_{R|P}(y|\chi)|\chi), \forall y \in \Re \quad (5)$$

where $F^{-1}_{I|P}(y|\chi)$ denotes the pseudo-inverse of $F_{I|P}(y|\chi)$.

This is a composed function, with two non-decreasing components. The composition is therefore also non-decreasing, as required.

Next, we have to prove that (4) holds:

$$P(g(X_P, X_R) \leq y | X_P = \chi) = P(g(\chi, X_R) \leq y | X_P = \chi) =$$
$$P(F^{-1}_{I|P}(F_{R|P}(X_R|\chi)|\chi) \leq y | X_P = \chi) = P(F^{-1}_{I|P}(U|\chi) \leq y | X_P = \chi), \quad (6)$$

where U ~ Uniform [0,1] by the probability integral transformation. Then, by the same reasoning, $F^{-1}_{I|P}(U|\chi)$ has the same distribution as X_I given $X_P = \chi$, so we finally have:

$$P(F^{-1}_{I|P}(U|\chi) \leq y | X_P = \chi) = P(X_I \leq y | X_P = \chi) = (F_{I|P}(y|\chi), \quad (7)$$

and (4) holds as required.

In order to obtain the function g, we need to model the conditional distributions $F_{I|P}$ and $F_{R|P}$. Let $C_{P,I}$ denote the copula between X_P and X_I and let $C_{P,R}$ denote the copula between X_P and X_R. Then from (1) we have:

$$F_{P,I}(\chi, y) = C_{P,I}(F_P(\chi), (F_I(y)), \chi \in \Re, y \in \Re. \quad (8)$$

$$F_{P,R}(\chi, y) = C_{P,R}(F_P(\chi), (F_R(y)), \chi \in \Re, y \in \Re. \quad (9)$$

We can write the conditional distributions $F_{I|P}$ and $F_{R|P}$ as:

$$F_{P|I}(y|\chi) = \kappa^{P,I}_{\chi}(y), \chi \in \Re, y \in \Re, \text{ where } \kappa^{P,I}_{\chi}(y) = \frac{\partial C_{P,I}(u,v)}{\partial u}\bigg|_{u = F_P(\chi),\, v = F_I(y)}$$

$$F_{R|P}(y|\chi) = \kappa^{P,R}_{\chi}(y), \chi \in \Re, y \in \Re, \text{ where } \kappa^{P,R}_{\chi}(y) = \frac{\partial C_{P,R}(u,v)}{\partial u}\bigg|_{u = F_P(\chi),\, v = F_R(y)}$$

So the cheapest function g in expression (5) can be rewritten as:

$$g(\chi, y) = \kappa_{\chi}^{(-1)P,I}\left(\kappa_{\chi}^{P,R}(y)\right), \chi \in \Re,\ y \in \Re \qquad (10)$$

We can now rewrite everything in terms of the end-of-month payoff to obtain the desired payoff function. The end-of-month replicated values from a monthly initial investment of 100 will be equal to:

$$S_g = g^*(S_P, S_R) = 100 \exp g\left(\log\left(\frac{S_P}{100}\right), \log\left(\frac{S_P}{100}\right)\right) \qquad (11)$$

Theoretically, the vector (S_P, S_g) will have the same joint distribution as the vector (S_P, S_I), meaning that, as intended, we are not only replicating the end-of-month payoff of the fund, but also its dependence with the investor's existing portfolio.

Pricing and Generating the Desired Payoff Function

Having determined the desired payoff function, the next step is to price it. This is of course not a new problem. It is what option pricing theory has concentrated on for the last 35 years. To price the desired payoff function we use the multivariate option pricing model of Boyle and Lin (1997). We could have used a standard bivariate Black-Scholes (1973) type model for this, but the Boyle and Lin (1997) model explicitly allows for transaction costs, which makes it ideally suited for the current application. Once the desired payoff function is priced, we can subsequently work out the controls of the dynamic trading strategy generating it.

Two points are worth noting at this stage. First, only after pricing the payoff function do we know what the expected return on the replicating strategy will be. The desired payoff function explicitly aims to replicate all aspects of the desired distribution, except the fund's expected return. The latter follows from the expected return on the investor's portfolio and the reserve asset, the desired payoff function, and the pricing environment for the latter, i.e. interest rates, expected dividends, volatilities, etc. In other words, it is the capital market that sets the expected return on the replicating strategy. Second, although determined in a much more flexible setting, the desired payoff function is priced in the standard model where asset returns are normally distributed. As long as we don't have access to a more sophisticated theoretical pricing model, we cannot escape this inconsistency.

3. Three Out-of-Sample Tests

Having worked out the replication procedure in theory, the next step is to check whether it also works in practice. We therefore proceed with some out-of-sample tests. Taking the replication procedure into the real world introduces a new set of problems. Where the model assumes all relevant parameters to be known, we are confronted with a significant degree of uncertainty about future parameter values for example. Fortunately, these problems are not new. They are characteristic to all model-based dynamic trading strategies. A number of authors have studied and suggested solutions to the above problems[13]. None of these, however, appears to be able to improve the efficiency of dynamic trading strategies to a very large extent. We therefore assume the simplest possible set-up. If our replication strategies do not work under these conditions, it is unlikely they will work in a more elaborate set-up.

The out-of-sample tests that follow are all structured in the same way. Given a fund to be replicated, we take the first 24 months of its track record as given, assuming we do not know anything about what is to come. If a fund's track record starts in January 1995 for example, we assume to be living on January 1, 1997. Subsequently, we determine the desired payoff function from the available 24 monthly returns, calculate the accompanying strategy controls and set up the required positions. During the month, we adjust our portfolio on a daily basis, driven by the daily changes in the underlying index values. At the beginning of the next month, we include the hedge fund return over the previous month in our dataset and repeat the whole procedure, now using 25 monthly returns instead of 24[14]. The above is repeated until we arrive at the end of October 2004 (where our hedge fund database ends).

Based on the idea that hedge funds provide investors with equity-like returns and bond-like risk, it has been suggested[15] that investors should replace their bond holdings by hedge funds. In line with this recommendation, we assume that the investor's existing portfolio consists of 100% equity, in the form of the DAX® 30 tracking portfolio. We use German Bunds as the reserve asset. To minimize transaction costs, all exposure management is done in the futures markets, trading DAX® 30[16] and Euro-Bund Futures on Eurex[17]. Put another way, we use fully collateralised DAX® 30 and Euro-Bund Futures as the portfolio and reserve asset respectively. All collateral is invested at the going one-month Euribor rate. We hold nearby futures contracts, rolling over into the next nearby contract on the first trading day of the nearby contract's expiry month to avoid liquidity drying up. Assuming sufficient liquidity, transaction costs for both futures contracts are set equal to BP 1 one-way. The necessary volatility and correlation inputs are obtained from historical estimates, using all available data at the time of determining the desired payoff function.

13 See for example Figlewski (1989), Kat (1996) or Clewlow and Hodges (1997).

14 In practice hedge funds typically take one or two weeks to report their end-of-month net asset value. For simplicity, we refrain from this complication here.

15 See for example McFall Lamm (1999).

16 We use DAX® instead of Dow Jones EURO STOXX 50® Index Futures as our data on DAX® Futures go back to 1990, as opposed to 1998 for the Dow Jones EURO STOXX 50® Index Futures.

17 See **www.eurexchange.com** for contract details.

In what follows we discuss the out-of-sample replication results for three different hedge funds (of funds). We selected these funds because they are well known within the industry and because they have relatively long track records[18]. The latter requirement stems from the fact that when comparing the statistical properties of the fund and the replicated returns we are basically comparing two bivariate distributions, which is best done using as many data points as possible. All fund returns are net of fees and were taken from the TASS database, with all data series ending per October 2004. We do not charge any management fees in the replication strategy.

Several studies have shown that reported monthly hedge fund returns may exhibit highly significant levels of autocorrelation[19]. This primarily results from the fact that many hedge funds invest in illiquid securities, which are often hard to mark to market. When confronted with this problem, hedge fund administrators will either use the last reported transaction price or a conservative estimate of the current price, which creates artificial lags in the evolution of hedge funds' net asset values, resulting in artificial smoothing of the reported monthly returns. As a result, estimates of hedge fund volatility for example, can be biased downwards by 30–40 % in some cases.

One possible way to correct for the above autocorrelation is found in the real estate finance literature. Due to smoothing in appraisals and infrequent valuations of properties, the returns on direct property investment indices suffer from similar problems as hedge fund returns. The approach employed in the literature has been to "unsmooth" the observed returns to create a new set of returns which are more volatile and whose characteristics are believed to more accurately capture the characteristics of the underlying property values. Nowadays, there are several unsmoothing methodologies available. We use the method originally proposed by Geltner (1991).

A Well-Known Fund of Hedge Funds

Our first example concerns a fund of hedge funds, which is one of the top funds in the fund of funds ranking performed recently by French financial newspaper La Tribune in cooperation with the French business school Edhec. We will refer to this fund as "Fund1".

[18] More in particular, we did not select these funds because the replication procedure works especially well for them, nor do we mean to boost or damage these managers' business in any way.

[19] See for example Brooks and Kat (2002) or Lo et al. (2004).

Figure 2. 3D Plot Payoff Function Fund1

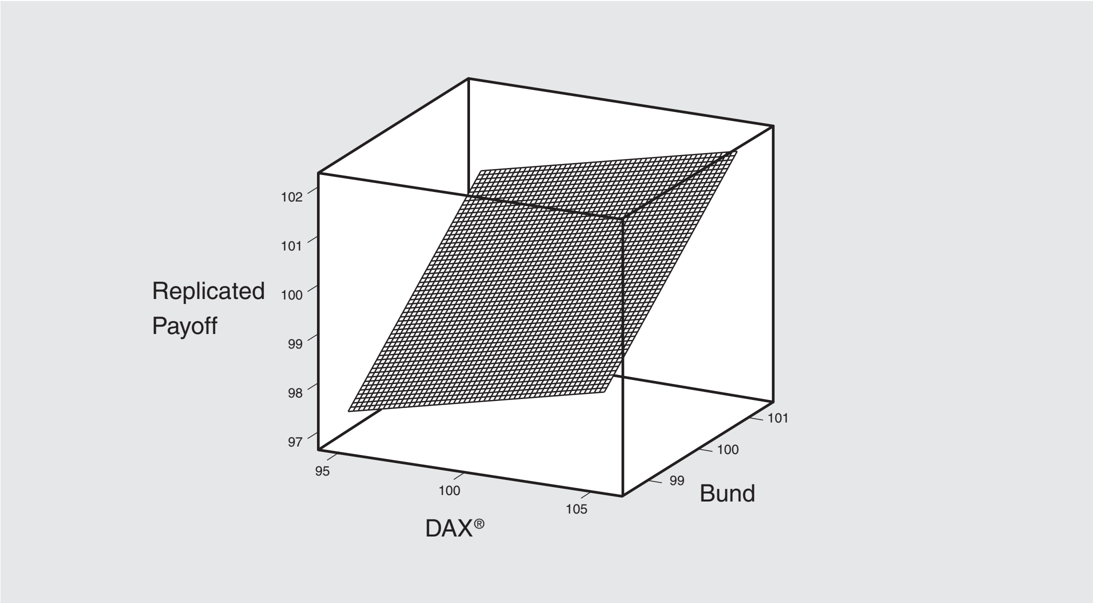

Figure 2 shows the payoff function used for the replication of the Fund1 return per October 1, 2004 (the last month for which we had fund return data available). From the graphs we see that the desired payoff is a positive function of the DAX® 30 as well as the reserve asset, implying that the replication strategy will be long in both DAX® 30 and Euro-Bund Futures. We also see some slight variation in the slope of the payoff surface, with the payoff being a slightly steeper function of the DAX®, the lower the Bund. This signals that generating the desired payoff is likely to require slightly more rebalancing when Euro-Bunds come down during the month.

Figure 3. Scatter Plot DAX® 30 Futures Returns Versus Fund1 Returns (Left) and Replicated Returns (Right), 1997–2004.

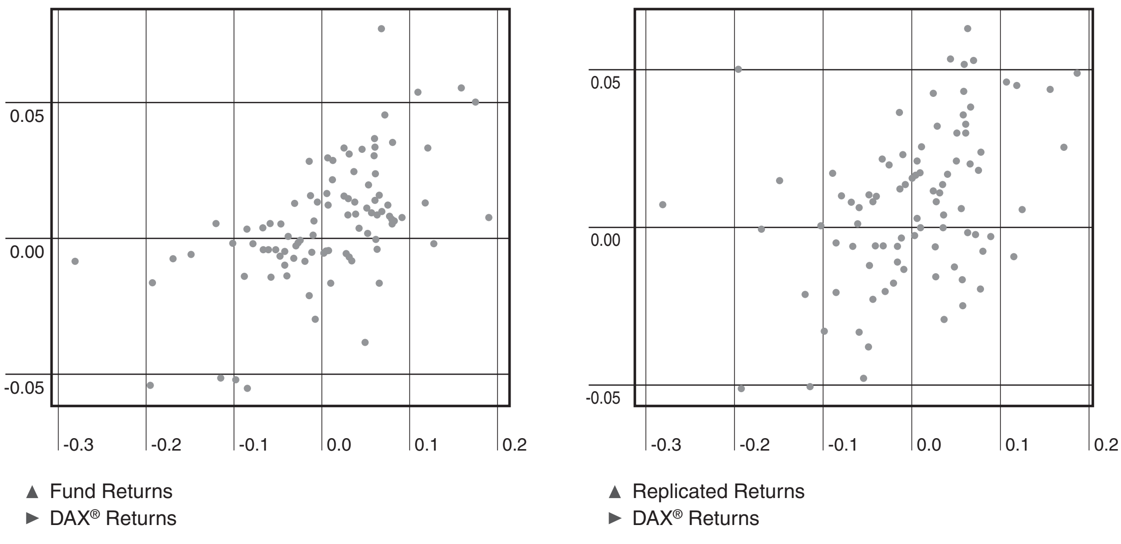

Next, we turn to the results of the replication exercise. The left hand side of figure 3 shows a scatter plot of the monthly returns on DAX® 30 Futures versus the Fund1 returns. The right hand side of figure 3 shows a scatter plot of the monthly returns on DAX® 30 Futures versus the replicated returns. Comparing both plots, we see that they are quite similar, which is a first indication that the replication strategy is indeed able to successfully replicate the statistical properties of Fund1's returns.

Table 1: Monthly Return Statistics Fund1 and Replication Strategy, 1997–2004.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt (robust)	Corr. with DAX	Kendall's Tau
Fund1	0.0059	0.0229	-0.0874	-0.0843	1.6046	2.0857	0.575	0.409
Replica	0.0080	0.0248	-0.1004	-0.1108	-0.1745	0.5420	0.408	0.308
Univariate K-S Statistic = 0.110, (approximated) p-value = 0.617								
Bivariate K-S Statistic = 0.121, (approximated) p-value = 0.665								

A better indication of the accuracy of the replication strategy comes from comparing the actual mean, standard deviation, skewness and kurtosis of Fund1's returns with those of the replicated returns. The latter statistics can be found in table 1, together with the correlation and Kendall's Tau[20] with the DAX® 30. Since the conventional skewness and kurtosis measures are very sensitive to extreme observations, we also report two more robust skewness and kurtosis measures[21]. To test whether the marginal distribution of the replicated returns and the joint distribution of the replicated returns and the investor's portfolio are significantly different from the original distributions, we use the univariate and bivariate Kolmogorov-Smirnov (K-S) tests[22].

Comparing the entries in table 1, it is clear that, despite the obvious limitations, the statistical properties of Fund1's returns have been successfully replicated. The replicated returns' standard deviation and skewness are very similar, while the kurtosis in the replicated returns is only slightly lower than in the Fund1 returns. The replication strategy has not only replicated the marginal distribution of Fund1's returns but also its relationship with the DAX® 30. This is also the conclusion from both the K-S tests. The Fund1 returns have a correlation of 0.575 with the DAX® 30, while the replicated returns have a correlation of 0.408. Kendall's Tau points in the exact same direction.

It is interesting to note that the average monthly replicated return exceeds the average Fund1 return by no less than BP 21, or 2.52 % per annum. This indicates that, measured over (most of) the fund's life, Fund1's returns have not been as good as many investors may have thought. Trading nothing else than DAX® 30 and Euro-Bund Futures, investors could have generated the same risk profile as Fund1 and obtained a higher average return at the same time.

[20] Kendall's Tau is a more robust measure of correlation than the conventional (Pearson) correlation coefficient, as it is rank-based instead of parametric. Like the conventional correlation coefficient, Kendall's Tau takes on values between –1 and +1.

[21] See Hinkley (1975) and Crow and Siddiqui (1967). These measures are also discussed in Kim and White (2004).

[22] See Fasano and Franceschini (1987) for details. Since the mean is not explicitly replicated, we subtract the mean from both the fund and the replicated returns before performing these tests.

Figure 4. Scatter Plot Euro-Bund Futures Returns Versus Replicated Returns (Left) and Fund1 Returns Versus Replicated Returns (Right), 1997–2004.

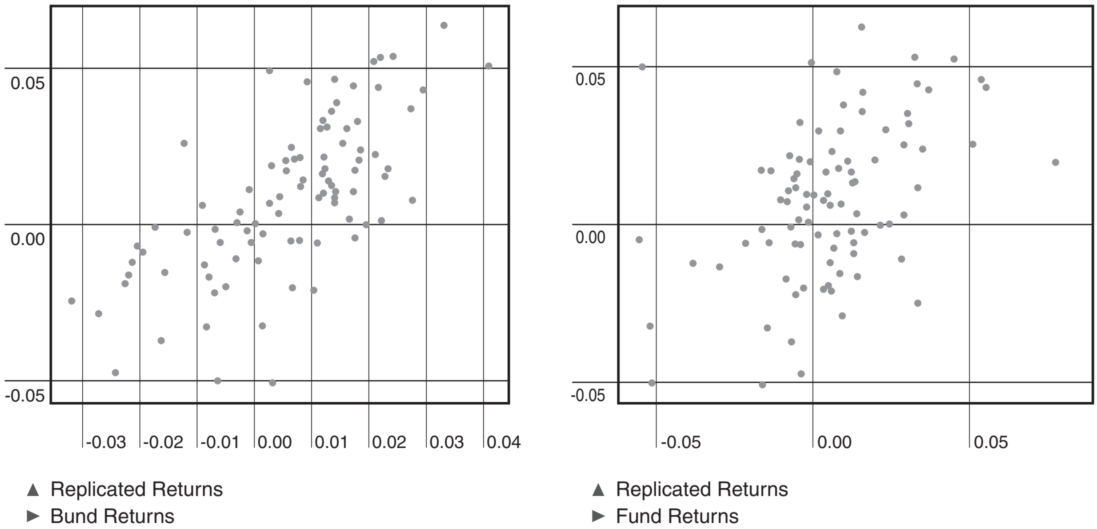

It is interesting to delve a bit further into the workings of the replication strategy. The left hand side of figure 4 shows a scatter plot of the Euro Bund futures returns versus the replicated returns. The positive relationship confirms the efficiency of the replication strategy (see section 2). The right hand side of figure 4 shows a scatter plot of the Fund1 returns versus the replicated returns. The plot makes it clear that although the replicated returns have statistical properties, which are very similar to those of Fund1, they come to the investor in a completely different order. It is exactly this feature of the replication process, i.e. giving up the requirement that returns need to be similar on a month-to-month basis as well, which allows us to do so much better than the standard factor model approach.

Figure 5. Evolution of Controls Fund1 Return Replication Strategy, March 1997–October 2004.

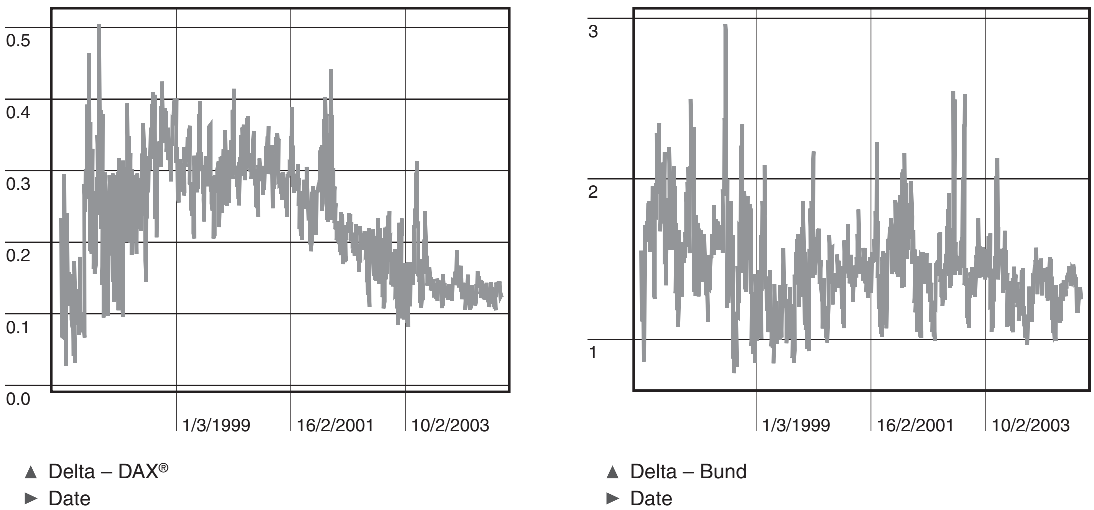

Figure 5 shows the evolution of the replication strategy's controls over the period March 1997– October 2004. The graph confirms that the replication strategy holds long positions in both DAX® 30 and Euro-Bund Futures. It also shows that the number of units of the Euro-Bund Future held is higher than for the DAX® 30 Future. This is because the volatility of the Euro-Bund is relatively low compared to that of Fund1 and the DAX® 30. It therefore requires some additional leveraging. The strategy is quite dynamic, with the strategy's controls exhibiting a number of peaks and troughs. The latter are the result of a combination of strong inter-month index movement, a steep payoff function and monthly strategy resetting.

A Well-Known Managed Futures Fund

Our second example is a well-known managed futures fund. We will refer to this fund as "Fund2". Studying the returns of Fund2, we noticed that its returns exhibit negative correlation with the DAX® 30. This makes replicating this fund's returns quite a challenge. The desired payoff function for Fund2 as per October 1, 2004 is shown in figure 6. The payoff function for Fund2 is quite different from what we found for Fund1. By construction, the payoff is a positive function of the reserve asset. The Fund2 payoff, however, is a negative function of the DAX® 30. The replication strategy will therefore go long in Euro-Bund Futures, but short in DAX® 30 Futures. This is of course what one would expect for a fund generating returns, which are negatively correlated with the stock market.

Figure 6. 3D Plot Payoff Function Fund2.

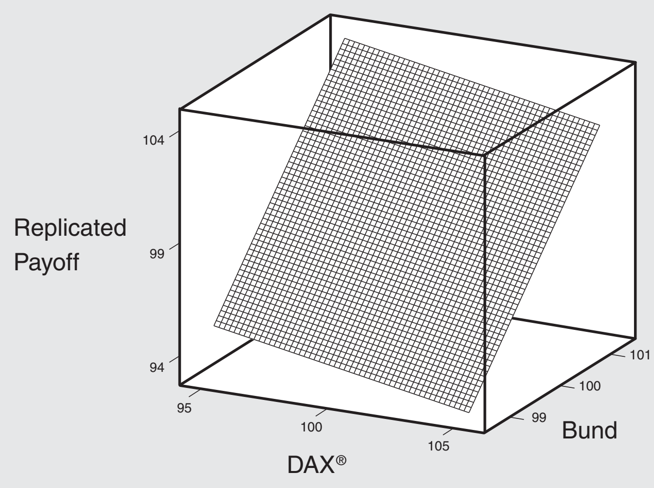

Figure 7. Scatter Plot DAX® 30 Futures Returns Versus Fund2 Returns (Left) and Replicated Returns (Right), 1998–2004.

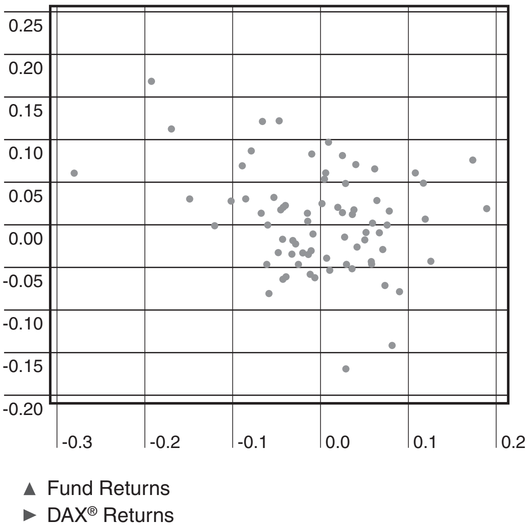

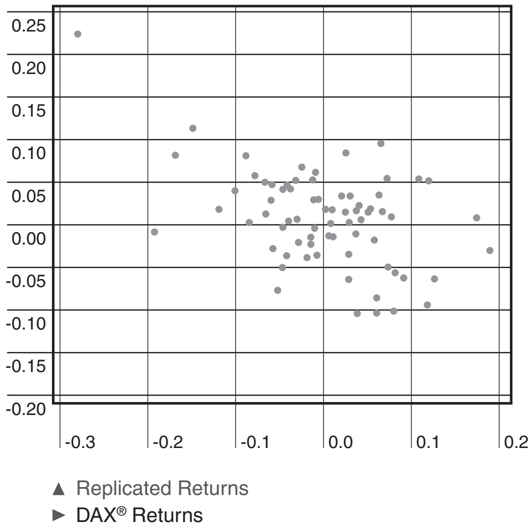

Figure 7 shows the same scatter plots as in figure 3: DAX® 30 Futures returns versus the Fund2 returns on the left and DAX® 30 Futures returns versus the replicated returns on the right. Comparing both plots, we again see that they are very similar, indicating the replication strategy performed quite well.

Table 2: Monthly Return Statistics Fund2 and Replication Strategy, 1998–2004.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt (robust)	Corr. with DAX® 30	Kendall's Tau
Fund2	0.0064	0.0596	0.0610	-0.0833	0.6954	1.0462	-0.271	-0.121
Replica	0.0089	0.0550	0.4867	-0.0820	2.3844	0.3912	-0.472	-0.221
Univariate K-S Statistic = 0.096, (approximated) p-value = 0.874								
Bivariate K-S Statistic = 0.116, (approximated) p-value = 0.822								

Table 2, which shows the same statistics as table 1, confirms the latter observation. As with Fund1, all statistics are very similar, including the negative correlation with the DAX® 30. In fact, the replicated returns exhibit more negative correlation than the Fund2 returns. Both the univariate and bivariate K-S test confirm that, aside from the means, there is no significant difference between the original and replicated risk profile. The average monthly Fund2 return, however, is BP 25 lower than the average replicated return. This means that, as with Fund1, trading DAX® 30 and Euro-Bund Futures over the period studied, investors could have obtained the same risk profile as with Fund2 but with a 3% higher average return.

Figure 8. Scatter Plot Euro-Bund Futures Returns Versus Replicated Returns (Left) and Fund2 Returns Versus Replicated Returns (right), 1998–2004.

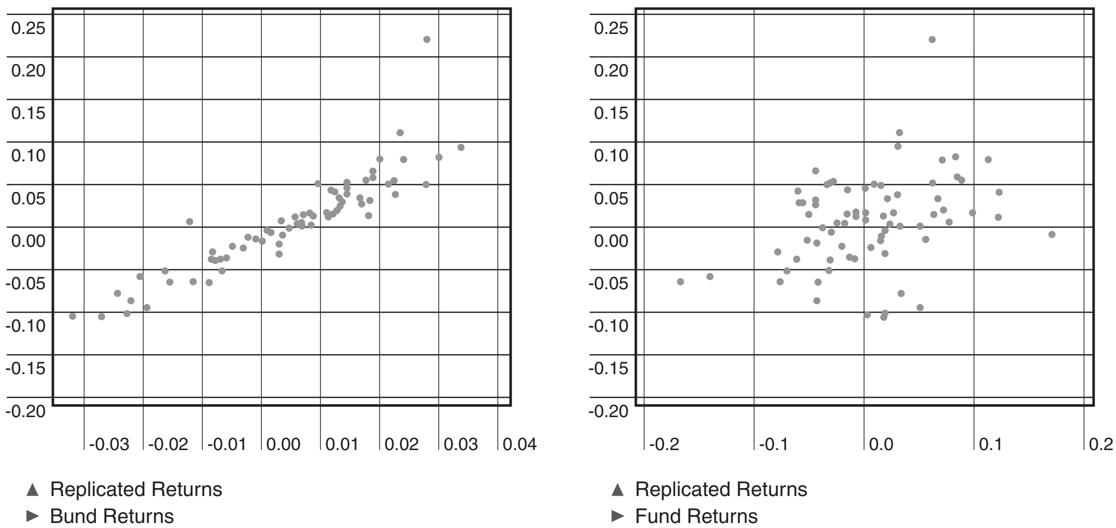

Figure 8 shows the same scatter plots as in figure 4. We again see a positive relationship between the Euro-Bund Futures returns and the replicated returns, confirming the efficiency of the replication strategy. In fact, the relationship appears to be quite strong. This suggests that the Bund element in Fund2's returns is quite substantial and may reflect a substantial holding of fixed income instruments as collateral for futures traded. The plot of the Fund2 returns versus the replicated returns shows a random scatter, underlining that although the replicated returns have similar statistical properties as Fund2, they come in a completely different order.

A Well-Known Event Driven Fund

Both funds we have looked at so far are denominated in Euros. There is, however, no reason why we could not incorporate exchange rate risk into the replication procedure as well. All this takes is to first convert the fund's returns into the desired currency (Euros in our case) before feeding them into the procedure. To illustrate this, our third example concerns a well-known event driven fund denominated in USD, which we will refer to as "Fund3".

Figure 9: 3D Plot Payoff Function Fund3.

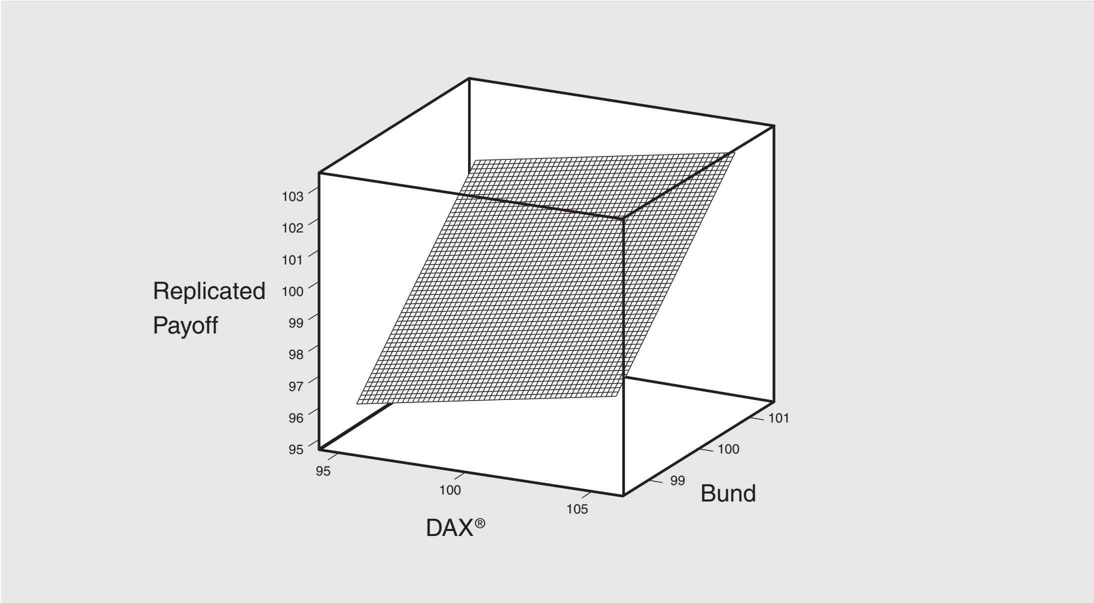

The desired payoff function for Fund3 as per October 1, 2004 can be found in figure 9. From this graph we see that the payoff function for Fund3 is similar to that for Fund1. Again, the relationship between the payoff and both underlying assets is positive, implying that the replication strategy will take a long position in both DAX® 30 and Euro-Bund Futures.

Figure 10. Scatter Plot DAX® 30 Futures Returns Versus Fund3 Euro Returns (Left) and Replicated Returns (Right), 1992–2004.

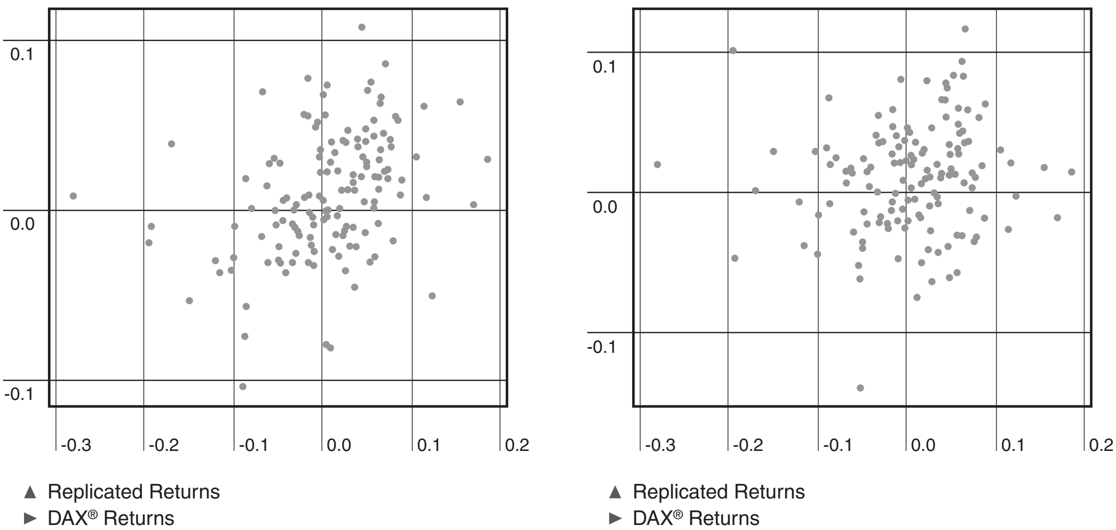

Table 3: Monthly Return Statistics Fund3 and Replication Strategy, 1992–2004.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt (robust)	Corr. with DAX® 30	Kendall's Tau
Fund3	0.0090	0.0362	-0.1382	-0.0905	0.3191	0.1662	0.355	0.278
Replica	0.0107	0.0391	-0.2471	-0.0481	0.9733	-0.1265	0.119	0.113
Univariate K-S Statistic = 0.091, (approximated) p-value = 0.576								
Bivariate K-S Statistic = 0.105, (approximated) p-value = 0.598								

From figure 10 and table 3 we see that, despite the fact that we are dealing with a completely different type of strategy, denominated in USD instead of Euros, the replication strategy performs in the same way as before. The replicated returns' standard deviation, skewness and kurtosis are very similar to the Euro returns of Fund3. In addition, the replicated returns exhibit the same low correlation with the DAX® 30 as the Fund3 Euro returns. As before, the mean of the replicated returns is higher than the average Fund3 Euro return. As for Fund1 and Fund2, the difference is again in the order of 2 % per annum. In other words, by simply trading DAX® 30 and Euro-Bund Futures investors could have generated the same risk profile as Fund3, but with a 2 % higher average return.

Figure 11. Scatter Plot Euro-Bund Futures Returns Versus Replicated Returns (Left) and Fund3 Euro Returns Versus Replicated Returns (Right), 1992–2004.

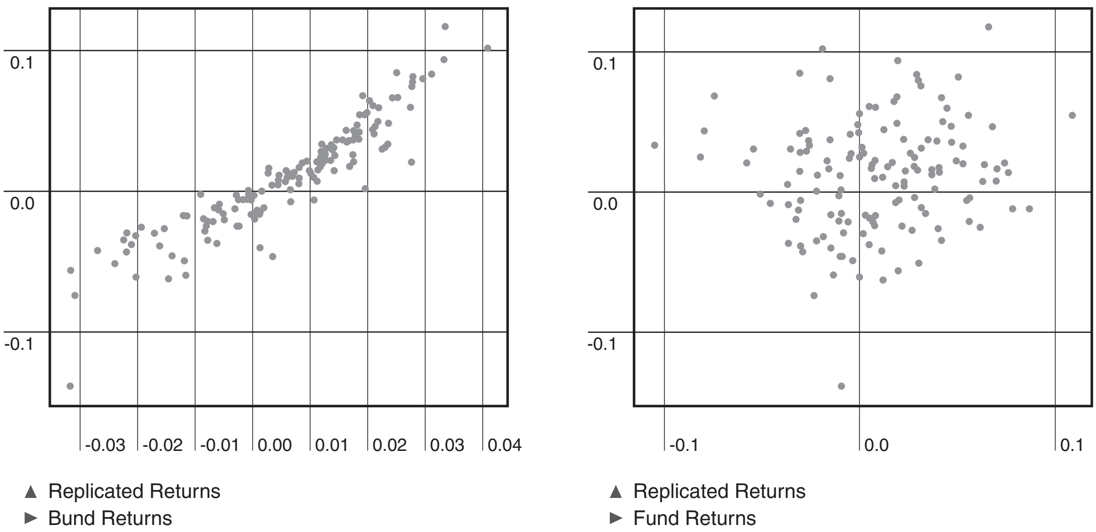

Finally, figure 11 shows a scatter plot of the Euro-Bund Futures returns and Fund3 Euro returns versus the replicated returns. As before, we see a clearly positive relationship between the Bund returns and the replicated returns, as well as a more or less random pattern relative to the Fund3 returns. Although the replicated returns have the same statistical properties, they arrive in a completely different order than the Fund3 returns.

4. Conclusion

Much of investors' current interest in hedge funds derives from the fact that traditional asset classes seem to lack opportunity these days. Stock markets have suffered considerably over 2000–2002 and are still somewhat hesitant, bond prices will come down when interest rates go up again, and the yield curve is flattening. With fresh memories of double-digit returns, this has driven investors towards commodities, emerging markets, credit-based structures, and of course hedge funds. Having generated high returns in the early years, the average return on hedge funds over the last 10–15 years has been quite im-pressive and many investors seem more than happy to use this as a guide for future returns. However, given today's low interest rates, low risk premiums across the board, as well as the current size of the hedge fund industry itself, a repeat of the last 10–15 years is extremely unlikely.

Investing in alternatives comes with many drawbacks, including the need for extensive due diligence, liquidity, capacity, transparency, style drift and regulatory problems, and excessive management and incentive fees. As long as investors believe they will be rewarded with (close to) double-digit returns, they will take these problems for granted. However, if reality kicks in and investors realize that hedge funds are no longer the money machines they once were (thought to be), their attitude may change. The above drawbacks will become more and more important and may ultimately become a reason to say farewell to hedge funds altogether and migrate to other alternative asset classes like emerging markets for example, which have shown stellar performance over the last 3 years. In fact, according to HFR, during the third quarter of 2005, funds of hedge funds were confronted with their first net outflow of funds in the amount of USD 1.2 billion. In the fourth quarter of 2005 funds of funds saw additional net outflows of USD 2.1 billion[23].

In this paper we have shown that apart from investing in real hedge funds, there is a viable alternative. Using a copula-based procedure, it is possible to design dynamic futures trading strategies, which generate returns with the same statistical properties as the returns on individual hedge fund or fund of funds investments. These strategies are not only capable of replicating a fund's risk profile, but, in the majority of cases, will also offer investors a higher expected return. Not because we are able to create something out of nothing, but simply thanks to a lack of fees and optimised, low cost trading[24]. Since this is accomplished by trading futures on traditional assets only, these strategies avoid the many drawbacks surrounding hedge funds and other alternative investments. As such, our synthetic hedge fund returns are clearly to be preferred over real hedge fund returns.

Finally, it should be noted that the applications of the technique introduced here are not limited to replication only. The same technique can also be used for performance evaluation for example. When the average replicated return is significantly higher than the average fund return, that fund cannot claim superiority. After all, superior returns can't be replicated.

[23] For the year 2005, Funds of Funds brought in only USD 9.5 billion in net new assets compared to USD 33 billion in 2004 and USD 59.4 billion in 2003.

[24] See Kat and Palaro (2006a, 2006b) for a more extensive, replication-based evaluation of hedge fund performance over the last 20 years.

As it essentially allows one to design trading strategies that generate returns with pre-defined statistical properties, the technique can also be used for the creation of completely new risk-return profiles. This means that investors no longer have to go through the usual process of finding and combining assets and funds in a costly and often unsuccessful attempt to construct a portfolio with the risk-return characteristics they require. Given the proper trading strategy, investors can now generate directly whatever risk-return profile they are after. Based on this technique, a whole new industry could develop!

5. References

Akaike, H., Information Theory and an Extension of the Maximum Likelihood Principle, in: B. Petrov and F. Csaki (eds.), Second International Symposium on Information Theory, Academiae Kiado, Budapest, 1973, pp. 267-281.

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

Amin, G. and H. Kat, Hedge Fund Performance 1990-2000: Do the Money Machines Really Add Value?, Journal of Financial and Quantitative Analysis, Vol. 38, June 2003, pp. 1-24.

Ang, A. and J. Chen, Asymmetric Correlations of Equity Portfolios, Journal of Financial Economics, Vol. 63, 2002, pp. 443-494.

Bailey, W., H. Li and X. Zhang, Hedge Fund Performance Evaluation: A Stochastic Discount Factor Approach, Working Paper Johnson School of Management, Cornell University, 2004.

Black, F. and M. Scholes, The Pricing of Options and Corporate Liabilities, Journal of Political Economy, Vol. 81, 1973, pp. 637-654.

Boyle, P. and X. Lin, Valuation of Options on Several Risky Assets When There are Transaction Costs, in: P. Boyle, G. Pennacchi and P. Ritchken (eds.), Advances in Futures and Options Research, Vol. 9, Jai Press, 1997, pp. 111-127.

Brooks, C. and H. Kat, The Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors, Journal of Alternative Investments, Fall 2002, pp. 26-44.

Calamos, J., Convertible Securities: The Latest Instruments, Portfolio Strategies, and Valuation Analysis, McGraw-Hill, 1998.

Calamos, N., Convertible Arbitrage: Insights and Techniques for Successful Hedging, John Wiley & Sons, 2003.

Cherubini, U., E. Luciano, and W. Vecchiato, Copula Methods in Finance, John Wiley & Sons, 2004.

Clewlow, L. and S. Hodges, Optimal Delta-Hedging Under Transaction Costs, Journal of Economic Dynamics and Control, Vol. 21, 1997, pp. 1353 – 1376.

Crow, E. and M. Siddiqui, Robust Estimation of Location, Journal of the American Statistical Association, Vol. 62, 1967, pp. 353-389.

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

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

Dyment, J., J. Olstein and A. Jones, 2005 Alternative Investment Survey, Deutsche Bank Hedge Fund Capital Group, 2005.

Fasano, G. and A. Franceschini, Monthly Notices of the Royal Astronomical Society, Vol. 225, 1987, pp. 155-170.

Figlewski, S., Options Arbitrage in Imperfect Markets, Journal of Finance, Vol. 44, 1989, pp. 1289-1311.

Fung, W., D. Hsieh, N. Naik and T. Ramadorai, Hedge Funds: Performance, Risk and Capital Formation, Working Paper London Business School, 2005.

Geltner, D., Smoothing in Appraisal-Based Returns, Journal of Real Estate Finance and Economics, Vol. 4, 1991, pp. 327-345.

Glasserman, P., Monte Carlo Methods in Financial Engineering, Springer Verlag, 2003.

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

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

Hinkley, D., On Power Transformations to Symmetry, Biometrika, Vol. 62, 1975, pp. 101-111.

Joe, H., Multivariate Models and Dependence Concepts, Chapman and Hall, 1997.

Johnson, N., Systems of Frequency Curves Generated by Methods of Translation, Biometrika, Vol. 36, 1949, pp. 149-176.

Johnson, N., Tables to Facilitate Su Frequency Curves, Biometrika, Vol. 52, 1965, pp. 547-558.

Kat, H., Delta Hedging of S&P 500 Options: Cash versus Futures Market Execution, Journal of Derivatives, Spring 1996, pp. 6-25.

Kat, H. and H. Palaro, Replication and Evaluation of Fund of Hedge Funds Returns, Working Paper No. 28, Alternative Investment Research Centre, Cass Business School, City University London, 2006a (downloadable from www.cass.city.ac.uk/airc).

Kat, H. and H. Palaro, Superstars or Average Joes? A Replication-Based Performance Evaluation of 1917 Hedge Funds and Funds of Funds, Working Paper No. 30, Alternative Investment Research Centre, Cass Business School, City University London, 2006b (downloadable from www.cass.city.ac.uk/airc).

Kim, T.H, and H. White, On More Robust Estimation of Skewness and Kurtosis: Simulation and Application to the S&P500 Index, Finance Research Letters, Vol. 1, 2004, pp 56-73.

Lo, A., M. Getmansky and I. Makarov, An Econometric Analysis of Serial Correlation and Illiquidity in Hedge Fund Returns, Journal of Financial Economics, Vol. 74, 2004, pp. 529-609.

Longin, F. and B. Solnik, Extreme Correlation of International Equity Markets, Journal of Finance, Vol. 56, 2001, pp. 649-676.

McFall Lamm Jr., R., Portfolios of Alternative Assets: Why Not 100% Hedge Funds?, Journal of Alternative Investments, Winter, 1999, pp. 87-97.

Nelsen, R., Introduction to Copulas, Springer Verlag, 1999.

Patton, A., Modelling Asymmetric Exchange Rate Dependence, Working Paper London School of Economics, forthcoming in the International Economic Review, 2005a.

Patton, A., Estimation of Multivariate Models for Time Series of Possibly Different Lengths, Working Paper London School of Economics, forthcoming in the Journal of Applied Econometrics, 2005b.

Schneeweis, T., H. Kazemi and V. Karavas, Eurex Derivative Products in Alternative Investments: The Case of Hedge Funds, Working Paper CISDM, University of Massachusetts, 2003.

Sharpe, W., Asset Allocation: Management Style and Performance Measurement, Journal of Portfolio Management, Winter 1992, pp. 7-19.

Xu, J., Statistical Modelling and Inference for Multivariate and Longitudinal Discrete Response Data, PhD Thesis, University of British Columbia, 1996.

Published
March 2006

Order Number
E5E-164-0306

ARBN Number
Eurex Frankfurt AG ARBN 100 999 764