ESSAYS IN HEDGE FUND REPLICATION, EVALUATION AND SYNTHETIC FUNDS

by

Helder Parra Palaro

A Thesis Submitted in Partial Fulfilment of
the Requirements for the Degree of

Doctor of Philosophy in Finance

CITY UNIVERSITY

Faculty of Finance, Cass Business School

May/2007

Supervisor : Prof Harry M Kat

Table of Contents

List of Tables

List of Illustrations

Acknowledgements

First I would like to thank my supervisor, Professor Harry M Kat. He introduced me to the fantastic world of the alternative investments, and he was always available to discuss any question which I had. His experience and objectivity were very important in this project. We worked 24/7 in the last 2.5 years.

The development of this project would not be possible without the financial support from my Brazilian sponsoring agency, CAPES. I would also like to thank my friend Benilton Carvalho de Sa and his family, who provided me accommodation for the interview which guaranteed this grant.

I also thank the Cass Business School staff for the research resources.

I thank my father, Geraldo Palaro (in memoriam) and my mother, Miriam Parra Palaro, for the education, love and encouragement in all these 23 years as a student. My father always said to me that the most valuable thing he would pass to me was my education, and I am very grateful for all his efforts and his support when I decided to pursue a PhD degree abroad.

I am very grateful to Chrystine, now my wife, for her love and patience during the whole PhD period.

I also thank my friend Alexandre Rubesam, who has been in this whole higher education journey with me, since the beginning of our undergraduate studies in 1998.

Finally, I am grateful to the United Kingdom, a great nation which welcomed me very warmly.

Copyright waiver

Abstract

In this thesis it is developed and demonstrated the workings of a copula-based technique that allows the derivation of dynamic trading strategies, which generate returns with statistical properties similar to hedge funds. It is shown that this technique is not only capable of replicating fund of funds returns, but is equally well suited for the replication of individual hedge fund returns. Since replication is accomplished by trading futures on traditional assets only, it avoids the usual drawbacks surrounding hedge fund investments, including the need for extensive due diligence, liquidity, capacity, transparency and style drift problems, as well as excessive management fees. This replication technique is also used to evaluate the net-of-fee performance of 875 funds of hedge funds and 2073 individual hedge funds, up to an including November 2006. Comparing fund returns with the returns on dynamic futures trading strategies with the same risk and dependence characteristics, no more than 18.6% of the funds of funds and 22.5% of the individual hedge funds in the data sample convincingly beat the benchmark. Besides the replication and evaluation of funds which already exist in the market, this technology can also be used to create new funds with previously unavailable return characteristics, the so-called 'synthetic funds'. In a set of four out-of-sample tests over the period January 1998 – February 2007, it is shown that the replication-based strategies are indeed capable of accurately generating returns with a variety of properties, including negative correlation with stocks and bonds and high positive skewness. The synthetic funds also produce impressive average excess returns. Disappointing performance is leading hedge fund investors to look for cheaper alternatives to invest, such as indices of hedge funds. Unfortunately, investable hedge fund indices are nothing more than funds of funds in disguise, with performance similar or even worse than real funds of funds. The replication technology generates returns with statistical properties very similar to those of hedge fund indices, and a higher average return for most hedge fund categories, but without actually investing in hedge funds.

Introduction

A hedge fund is a private fund which can take both long and short positions, leverage itself using derivatives and invest in any local or global market, with the goal of generating high returns, in absolute terms or compared with a benchmark. Unlike other investments, hedge funds are usually not regulated by any government commission.

Over the last 10-15 years hedge funds have become very popular with high net worth investors and are currently well on their way to acquire a significant allocation from many institutional investors as well. During the period 1990-2006, the number of hedge funds has risen from around 500 to more than 9000, and the assets under management are estimated to have increased from 50 billion to more than 1.5 trillion of US dollars[1].

Initially, hedge funds were sold on the story of superior performance, where high skilled managers could 'beat' the market. Seduced by this argument, private wealthy investors were responsible for the initial growth of the industry. Toward the end of the 1990s the performance of hedge funds started to deteriorate. Also, institutional investors with better risk management practices got more interested in these funds. Given the two reasons above, hedge funds began to be sold not on the promise of superior performance, but on the basis of the diversification argument, due to their low correlation with stocks and bonds[2].

Despite its popularity, investments in hedge funds come with some serious drawbacks. The first one is the lack of liquidity. Most hedge funds use lock-up clauses to tie in new investors for some period of time, up to five years[3]. After that, investors need to give a prior notice if they want to withdrawal their money, and in some cases they even need to pay an additional fee of up to 5%. The second major drawback is the lack of transparency. Hedge funds are usually 'black box' investments, which turns more

[1] HFI Press Release.

[2] 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.

[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%.

difficult the assessment of their risk profile, and any change of style or strategy overtime. Moreover, some strategies can suffer from capacity problems, specially the arbitrage-related ones. This has already happened for the convertible arbitrage strategy, for example[4]. Another major drawback is the level of fees imposed by most hedge funds. The average fund charges a flat management fee of 2% plus an additional performance fee of 20% of gross returns. This is usually known as "2 plus 20". For example, considering an annual return of 10%, this would mean 40% of the gross returns being paid as fees. Funds of funds tend to charge an additional "1 plus 10" on top of the "2 plus 20" structure.

If returns generated by hedge funds were far superior, the drawbacks discussed above would at least be compensated in some way. However, recent performance studies[5] show that the return-risk characteristics of hedge funds are no longer superior, but just different[6]. Then the question which naturally arises is whether it is possible to generate hedge fund-like returns mechanically trading cash, stocks, bonds and other asset classes. This so-called 'replication' would eliminate the liquidity, transparency, capacity and excessive fees problems.

Given these considerations, a new technique is introduced in Chapter 2. This technique extends the work of Amin and Kat (2003b), to replicate not only the marginal distribution of a given hedge fund, but also its dependence with the investor's existing portfolio distribution.

The execution of the strategy is based on dynamic trading. The basic idea of dynamic trading was put forward a long time ago by Arrow (1964), who pointed out that, instead of following a buy-and-hold strategy, by trading more often investors can exert greater control over the evolution of the value of their investment portfolio. This is an extremely important observation as it implies that when a given payoff profile is not directly available in the market, either as an individual asset or as a combination of different assets, investors may still be able to create it themselves by trading the available primitive assets in a specific way. The idea of complementing a market by

[4] The annual return of the HFRI Convertible Arbitrage Index was 1.18% in 2004 and -1.86% in 2005.

[5] These studies are briefly reviewed in Chapter 1

[6] See for example Amin and Kat (2003b), Bailey et al. (2004) or Fung et al. (2007).

dynamic trading was taken to the extreme in modern option pricing theory, which is rooted in the fact that, under certain simplifying assumptions, when investors can trade continuously, they will be able to generate any payoff profile imaginable. Black, Scholes and Merton used this observation to develop their famous option pricing formula. Over the 30 years that followed, others have used the same argument to price a large variety of other, more exotic options. The reasoning is always the same though. If it is possible to find a payoff function which, given the probability distribution of the underlying index or indices, implies the desired distribution, then the dynamic trading strategy which generates (returns that are drawings from) that distribution will also be found.

Of course, there are a number of serious hurdles to take. First, the interest does not lie in any strategy. To maximize expected return, the cheapest possible strategy is wanted. Second, since the aim is to replicate not only a fund's marginal return distribution but also its relationship with the investor's existing portfolio, it will be necessary the use of bivariate distributions, which can take on a large variety of shapes and forms. Third, real markets are a lot less well behaved than assumed in the standard theoretical model. 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 these dynamic trading strategies may not be able to exactly generate the desired payoff. Therefore extensive out-of-sample tests of these strategies were performed, using daily data over the period 1985-2006.

In Section 2.1 it is briefly discussed the theoretical setting in the form of Dybvig's (1988a) Payoff Distribution Pricing Model (PDPM), which is extended to a bivariate setting. In Section 2.2 the univariate replication is briefly reviewed. In Section 2.3, the determination of the desired payoff function is discussed, i.e. the payoff function, which, given the distribution of the assets to be traded, implies the desired return distribution. In Section 2.4 a number of simulation-based analyses is carried out, investigating how the size of the available data sample influences the accuracy of the procedure. In Section 2.5 the practical implementation of the procedure and the results of some out-of-sample tests are discussed. Three well-known hedge funds (of funds) are replicated. Proofs relating to the univariate and bivariate PDPM can be found in Appendix A.1 and A.2.

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 long 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, their risk profile could still be replicated, but it would not be possible to replicate their average as well. If it is superior, it cannot be replicated and vice versa. The latter observation points at another application of the replication technique discussed above: 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.

Using this idea, the performance of 2073 individual hedge funds and 875 funds of hedge funds is evaluated in Chapter 3, up to and including November 2006. Section 3.1 introduces the approach. In Section 3.2 one illustrative example is provided. In subsections 3.3.1, 3.3.2 and 3.3.3 it is discussed the data description, distributional analysis and evaluation results for funds of hedge funds. In Subsections 3.4.1, 3.4.2 and 3.4.3 the same is discussed for individual hedge funds.

In Chapters 2 and 3 the replication technique is used, respectively, to replicate and evaluate the returns of existing hedge funds. *There is no reason, however, why the same technique could not be used to create completely new funds, providing investors with previously unavailable return characteristics.* Finding and selecting new diversifiers is a very laborious and costly process. Typically, a fund's risk-return profile is not immediately obvious and investors may have to dig long and hard to gather sufficient information. This is where being able to create any type of risk-return profile pays off huge dividends, as it allows us to structure exactly what investors are looking for. No longer do investors have to work with what happens be available and guess what a fund's true risk-return profile is. Given an investor's existing portfolio, a special tailor-made strategy (or 'synthetic fund' as these strategies will be called) can be structured, that produces returns, which fit in optimally with what is already there. Clearly, this is a much more natural approach than the usual beauty parades held by investors.

The above idea could be even taken one step further and, instead of creating a synthetic fund as an addition to an investor's existing portfolio, the investor's entire portfolio can be replaced by a synthetic fund. This means that investors no longer would have to go through the usual process of finding and combining individual assets and funds into portfolios in an, often only partially successful, attempt to construct an overall portfolio with the characteristics they require. Using dynamic trading technology, a synthetic fund that produced returns with exactly the characteristics they were after could be designed.

In Chapter 4 it is tested whether, in practice, it is really possible to create synthetic funds, which generate returns with predefined statistical properties. In Section 4.1, four different synthetic funds with a variety of return characteristics are created and their out-of-sample performance are studied. In Section 4.2 some sensitivity analyses are performed to check the influence of the transaction costs, and the underlying chosen. Section 4.3 contains a brief comment on synthetic funds' alpha.

Chapter 5 will deal with the so-called hedge funds indexation. Over the past 20 years, indexation of equity portfolios has become very popular with institutional as well as private investors. Strongly advocated by big names from academia as well as the industry, such as Princeton's Burton Malkiel and Vanguard's John Bogle, assets under management by index funds have grown from almost negligible in 1976, when the first index fund was introduced, to several trillions of dollars now. It is estimated that (excluding closet-indexers) currently 40% of institutionally managed assets are indexed, with US institutions quite far ahead of the rest of the world.

The idea behind indexation is simple. Casual observation as well as large-scale academic research shows that the added value of traditional alpha chasing tends to be negative. Put simply, most traditional active managers are unable to earn back the fee that they charge their investors. This means that hiring these managers is counter-productive. On balance they will reduce the after-fee return. The above observation has lead many investors to abandon every form of security analysis and active trading altogether. They just buy and hold the market in the form of some index-replicating portfolio.

The rise of indexation over the past 20 years resulted from investors becoming aware of and admitting to the fact that the costs of traditional active management often exceed the value added. *A similar process is currently underway in the alternative investment, and especially the hedge fund industry*, where, as already discussed in this Introduction, fees tend to be a multiple of what they are in traditional investment management. The HFRI Fund of Funds Composite Index for example, returned 4.07% in 2000, 2.8% in 2001, 1.02% in 2002, 11.61% in 2003, 6.86% in 2004, 7.49% in 2005 and so far until October 2006 no more than 6.52%. Other well-known indices show a similar picture.

Given the above, and very similar to what has happened in traditional investment management, more and more investors are currently looking to improve their after-fee return by cutting costs. This brings us to the concept of hedge fund indexation. There are many different hedge fund indices around these days, so why not simply buy one of those, just as one does for stocks and bonds? Unfortunately, there are several reasons why this is going to be problematic:

1. As hedge fund managers seldom report to more than two databases, different indices cover different subsets of the hedge fund universe. This makes choosing an index to invest in far from trivial.

2. Some of the better-known hedge fund indices contain over 1000 individual hedge funds. Some contain even more than 2000 funds. In addition, most index providers do not explicitly list their index's components, which means it is impossible to find out what funds to actually invest in.

3. Most hedge fund indices contain a large number of funds that are closed for new, and sometimes even existing, investors. Also, most funds are highly illiquid due to lock-ups and notice periods. This makes any periodic rebalancing of an index-replicating portfolio highly problematic.

4. Most hedge fund administrators take a couple of weeks to work out the end-of-month Net Asset Value (NAV). This means that an index-replicating portfolio can only be rebalanced with a very significant delay.

5. Buying a hedge fund index will only allow one to eliminate the fund of funds manager's fee, but not the fees of the underlying hedge fund managers, which tend to be twice as high.

As a result of the above, the majority of existing hedge fund indices are simply not investable. Recently, several high profile index providers have attempted to solve this problem by introducing what they refer to as 'investable' hedge fund indices.[7] Unfortunately, the latter are nothing more than (more or less mechanically managed) funds of funds in disguise. Roughly speaking, investable hedge fund indices result from a joint venture between an index provider and a fund of funds manager. First, the manager puts together a portfolio of funds, in line with a number of criteria with respect to track record,[8] liquidity, valuation, transparency, capacity, etc. Subsequently, the index provider declares that portfolio to be the index. Apart from the obvious commercial incentive, the problem with investable indices is that there are not many hedge funds that fit all applicable criteria. As a result, investable indices are even less representative of the hedge fund universe than their non-investable counterparts.

The main problem with all existing hedge fund indexation schemes is that they still require investors to invest in hedge funds, which leaves the main cost factor, i.e. managers' fees, unaffected. Obviously, this is not a satisfactory solution for investors looking for a cheaper alternative. *To eliminate hedge fund managers' fees we need to find a way to obtain hedge fund index-like returns without actually investing in hedge funds.* This is possible by using the same replication machinery developed in Chapter 2. The returns generated by this approach will be statistically very similar to the actual hedge fund index returns. There is one important difference though: the synthetic fund returns are likely to arrive in a different sequence than the actual hedge fund index

[7] It is estimated that currently around $15 billion of the total $1200 billion invested in hedge funds is index-linked in some way.

[8] Since indices are big business these days, there appears to be a tendency for providers of investable indices to initially only include funds with good track records. Since there is little or no persistence in hedge fund returns, this explains why the actual performance of some investable indices has been so much worse than their pro-forma historical performance. In addition, investable indices have not outperformed the average fund of funds, which is not surprising as they are funds of funds themselves.

returns. For investment purposes this should not make a difference, however, as all that matters in an investment context are the statistical properties of the returns on the various assets and asset classes, and not the exact sequence in which those returns come in.

In Chapter 5 this idea is put in practice. Section 5.1 shows the main results and in Section 5.2, some sensitivity analyses are again performed to check the influence of the transaction costs, and the underlying indices chosen.

Chapter 6 concludes based on the results obtained in the replication (Chapters 2 and 5) and evaluation (Chapter 3) of hedge funds and funds of funds and also on the creation (Chapter 4) of synthetic funds.

1. Literature Review

As mentioned in the Introduction, two important applications of the copula-based method introduced in this thesis are hedge fund return replication and hedge fund performance evaluation. In this chapter we review previous research in both these areas.

Early studies of hedge fund performance used traditional performance measures such the Sharpe ratio or Jensen's alpha. Ackermann et. al (1999) investigate the effect of incentive fees, location, and age on hedge fund performance. They find a positive relationship between incentive fees and performance, but no strong location or age effect. Brown et. al (1999) examine the performance of offshore funds over the period 1989-1995, finding positive risk-adjusted performance for offshore funds as a group. They find below average returns for funds of funds though. Edwards and Caglayan (2001) analyse hedge fund performance over the period 1990-2001. They find that about 25 percent of hedge funds earn positive excess return and that the frequency and magnitude of funds' excess returns differ markedly by investment style.

Because of the non-linearities and non-normalities present in hedge fund returns, the traditional performance measures used in the above studies are not appropriate. For example, most hedge fund managers use dynamic strategies, which are quite different from the buy-and-hold strategy used in traditional investments. Also, these managers often use derivatives, especially options, which also create a non-linear relationship between fund returns and market returns. Even if the underlying assets were normally distributed, these non-linearities would result in non-normally distributed fund returns. The basic problem is that CAPM and the Sharpe ratio are derived assuming that investors only care about the first two moments of the returns distribution, i.e. the mean and the variance. As observed in Brooks and Kat (2002), often the attractive mean-variance characteristics of hedge fund returns is accompanied by negative skewness and positive excess kurtosis, indicating an increased probability of a relatively large losses. As negative skewness and excess kurtosis are both non-desirable features in a return distribution, this implies that traditional performance measures will tend to overestimate the true risk-return performance of hedge funds. In addition, as noted by Amenc, Martellini and Vaissi (2003), this could lead hedge fund managers to implement 'short volatilities' strategies based on the sale of out-of-the-money put and/or call options. Using these strategies, the premium received by writing the options will increase the mean return while the risk (measured only by the second moment of the returns distribution) would be limited. Some authors have used modifications of the Sharpe Ratio, which take higher distributional moments into account. This includes Kouwenberg (2003), Gregoriou and Gueyie (2003) and Eling (2006). However, these measures do not incorporate the non-linear relationship with other asset classes in a portfolio context, they only look at the marginal return distribution.

The style analysis technique first introduced by Sharpe (1992) for equity mutual funds, has also been used as a tool to replicate and evaluate hedge funds returns. Sharpe showed that only a limited number of asset classes are necessary to replicate the performance of a large sample of U.S. mutual funds. Once the relevant risk factors of a fund are selected and the sensitivities to these factors are estimated, one can construct a portfolio of stocks, bonds, cash and other securities, with the same sensitivities, which will then generate returns similar to the fund.

A number of studies have used the factor model approach to evaluate hedge funds, including Fung and Hsieh (1997, 2001), Schneeweis and Spurgin(1998), Liang (1999), Lhabitant (2001), Schneeweis et al. (2003), Capocci and Hubner(2004), Agarwal and Naik (2000a, 2000b, 2004), Hill et al (2004), Jaeger and Wagner (2005), Ibbotson and Chen (2006) and Hasanhodzic and Lo (2006).

Fung and Hsieh (1997) incorporate factors which reflect strategy-specific return components and a quantity component (leverage). They apply their 12-factor model to 3,327 mutual funds and 409 hedge funds/CTAs, finding that unlike mutual funds, hedge funds returns have low correlation with standard asset classes. Schneeweis and Spurgin (1998) propose factors that allow for the possibility of trending prices, short sales and intramonth volatility. Doing so, they find that hedge funds and CTAs provide beneficial diversification to stocks and bonds. Liang (1999) finds low correlation between different hedge fund strategies, and concludes that hedge funds offer a more attarctive risk-return trade-off than mutual funds. Capocci and Hubner (2004) use a large hedge fund database, with 2,796 individual funds. They consider the models used previously by Carhart (1997) and Agarwal and Naik (2000a), concluding that around one fourth of hedge funds delivers significant positive excess returns. This over-performance is, in most cases, constant over time, except during the 1997 Asian crisis. Agarwal and Naik (2004) extend the span of one or more benchmarks from purely linear to non-linear by including a number of ordinary put or call options on those benchmarks into the return generating process. They show the extent of the understimation in the tail risk in the mean-variance framework. They also show that a wide range of hedge fund strategies exhibit returns similar to those from writing a put option on the equity index. Finally, Ibbotson and Chen (2006) use a traditional factor model to analyse about 3,500 hedge funds and find that the pre-fee average return of 12.72% cna be split into a fee (3.74%), an alpha (3.04%) and a beta return (5.94%) component.

There are some serious problems with the factor-model approach, though. First, we often have little or no idea how a hedge fund's returns are actually generated. As a result, it is not clear which factors to use. One factor usually left out is liquidity. Many hedge funds act as market makers, buying illiquid assets and hedging the position with liquid ones. Omitting liquidity would therefore mean a loss of return potential in the replicating portfolio. In addition, some factors, such as volatility for example, can be difficult to trade in practice, with relatively high associated costs. Another problem is the assumption of a linear relationship between hedge funds and other asset classes. Although some authors have attempted to incorporate non-linearity introducing options in the replicating portfolio (as, for example, Agarwal and Naik (2004)), this also makes the execution of the replication strategies more complicated and expensive.

As a result of the above, the explanation power of factor models is quite disappointing, especially for individual hedge funds. In Hasanhodzic and Lo (2006) for example, the variation in hedge fund returns explained by their six-factor model is very low for most categories: convertible arbitrage (17.3%), emerging markets (19.4%), equity market neutral (10.4%), event driven (19.5%), fixed income arbitrage (14.9%), global macro (14.8%) and long/short equity (21.6%). Although the procedure works better for funds of funds and hedge fund indices, a model that leaves 80-85% of a fund's return variability explained is unlikely to be a good starting point when its aim is to replicate month-to-month returns.

Table 1.1 summarizes the out-of-sample replication results of Schneeweis et al. (2003) for five equally-weighted indices of European hedge funds over the period Jan 2001 – March 2003. From Table 1.1 it is clear that, despite the much higher systematic component in the index returns, the factor model used is unable to accurately replicate the returns on the above indices. As to judge from the correlation between the index return and the replicating return, the best results are obtained for long/short equity and event driven. Given the straightforward nature of these strategies, this is not really surprising. More complex strategies, like fixed income and convertible arbitrage, do a lot worse, however.

	Index		**Replica**		
Strategy	**Mean**	**StDev**	**Mean**	**StDev**	**Corr. index and replica**
Composite	-2.97%	3.35%	-7.07%	7.92%	43%
Fixed Income	7.87%	2.96%	2.89%	2.58%	16%
Long/Short	-0.98%	3.83%	-9.99%	7.13%	46%
Event Driven	-2.67%	4.79%	-6.34%	6.97%	90%
Convertible Arb	8.28%	1.82%	1.88%	1.54%	17%

Table 1.1: European hedge fund index return replication. Source: Schneeweis et al. (2003, Exhibit 2a-2f).

It is important to note that a high correlation between index return and replicating return does not guarantee that the replicating and index returns exhibit similar statistical properties. Looking at the standard deviations for event driven in Table 1.1, we see that, despite the 90% correlation, the standard deviation of the replicating returns is 46% higher than that of the index return. Obviously, this could cause major problems in portfolio risk management where the replica will typically be assumed to have properties similar to the target.

The Schneeweis et al. (2003) results are not unique. Table 1.2 shows to what extent the factor model used in Jaeger and Wagner (2005) was able to explain the variation in the well-known HFRI indices over the period Jan 1994 – Dec 2004. The message we get from table I.2 is not different from what we saw before in table 1.1. Relatively straightforward strategies, like long/short equity, score quite well, but more complex strategies, like managed futures and equity market neutral, come out a lot worse.

HFRI Index	Variation Explained
Managed Futures	34.3%
Equity Market Neutral	35.3%
Fixed Income Arbitrage	40.5%
Global Macro	49.7%
Merger Arbitrage	52.9%
Convertible Arbitrage	54.0%
Distressed	68.4%
Long/Short Equity	88.5%

Table 1.2: Percentage of HFRI return variation explained. Source: Jaeger and Wagner (2005, Table 1).

Given the failure of factor models, one could say that by trying to replicate hedge funds' month-to-month returns we are aiming too high. When an investor likes a hedge fund's returns, it is because of its statistical properties, such as mean, standard deviation, skewness and correlation with other asset classes. 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 the above route. Based on the early theoretical work of Glosten and Jagannathan (1994) and Dybvig (1988a, 1988b), and primarily aimed at evaluating hedge fund performance, Amin and Kat (2003b) 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. This innovative idea was an important step forwards from the factor model approach. A briefly review of such an univariate replication is provided in Section 2.2.

A new replication technique is introduced in Chapter 2. This technique extends the work of Amin and Kat (2003b) and aims to replicate not only the marginal return distribution of a given hedge fund of hedge fund index, but also its dependence with the investor's existing portfolio. The evaluation procedure explicitly takes transaction costs into account by, instead of a Black-Scholes type option pricing model, using the Boyle and Lin (1997) model. In factor model based evaluations, transaction costs are typically ignored, despite the fact that maintaining the replicating portfolio's factor loadings at their desired levels is likely to require significant periodic rebalancing. In addition, when dealing with hedge funds the risk factors used may be quite unusual and may therefore be accompanied by significant levels of transaction costs.

The replication-based evaluation also takes account of the fund's dependence pattern with other asset classes. According to theory as well as casual empirical observation, expected return and systematic co-variance, co-skewness and co-kurtosis are directly related. In other words, it is not just the marginal distribution, but its dependence structure with other assets that determines an asset's expected return. An asset, which is highly correlated with stocks and bonds, offers investors very little in terms of diversification potential. As a consequence, there will be little demand for this asset. Its price will be low and its expected return therefore relatively high. On the other hand, an asset that offers substantial diversification potential will be in high demand. Its price will be high and its expected return relatively low. Although hedge funds are not priced by market forces in the same way as primitive assets are, they do operate in the latter markets. It therefore seems plausible that a similar phenomenon is present in hedge fund returns as well. This is confirmed by the results in Kat and Miffre (2006).

In Chapters 2, 3 and 5, the return series used in the replication procedure are not the hedge funds' raw returns. The reason is that, as shown in Brooks and Kat (2002) and Lo et al. (2004) for example, monthly hedge fund returns may exhibit high levels of autocorrelation. This primarily results from the fact that many hedge funds invest in illiquid securities, which are 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 market price. This creates artificial lags in the evolution of hedge funds' net asset values, i.e. artificial smoothing of the reported returns. As a result, estimates of volatility, for example, will be biased downwards.

One possible method to correct for the above bias is found in the real estate finance literature. Due to smoothing in appraisals and infrequent valuations of properties, the returns of direct property investment indices suffer from similar problems as hedge fund returns. The approach employed in this 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. Throughout this thesis the method originally proposed by Geltner is used (1991).

In the evaluation procedure in Chapter 3 the sample contains live and dead funds. The reason for that is to minimize survivorship bias. Survivorship bias occurs if the database only contains information on "surviving funds". Ibbotson and Chen (2006) estimate survivorship bias in average returns to be 2.75% per year. In addition, Amin and Kat (2003a) mention that survivorship bias generates a downward bias in the standard deviation, and upward bias in the skewness and a downward bias in the kurtosis. Other studies attempting to quantify the degree and impact of survivorship bias are Brown et al (1992), Schneeweis and Spurgin (1996), Goetzmann and Park (1997), Fung and Hsieh (1997, 2000), Brown et al (1997), Hendrick et al (1997), Brown et al (1999), Carpenter and Lynch (1999), Liang (2000), Horst et al (2001) and Baquero et al (2004). Although estimates vary due to the use of different databases and different definitions of survivorship bias, these studies suggest that when left uncorrected survivorship bias inflates average hedge fund returns by 2-3%.

2. Replication of Hedge Funds[9]

2.1 Theoretical Setting

In principle, a given payoff distribution can be generated by many different payoff functions. Different payoff functions come with different price tags, however. We therefore need to know more about the general characteristics of the cheapest alternative. This is where Dybvig's (1988a) Payoff Distribution Pricing Model (PDPM) comes in. The PDPM can be derived from a simple set of primitive assumptions:

1) Investors' preferences depend only on the probability distribution of terminal, i.e. end-of-horizon, wealth.
2) Perfect capital markets (no taxes, transaction costs or information asymmetries).
3) Investors prefer more to less.

This set of assumptions allows investors' preferences to depend on all moments of the distribution of terminal wealth.

Suppose there are n possible states of the world. The *state price* of state i is the price of an elementary security which pays $1 if state i occurs and 0 otherwise. The *state-price density* is defined as the price per unit of probability of terminal wealth in a particular state, and is given by the ratio of the state price and the probability of occurrence of that state.

Consider, for example, a simple example with three equally probable states, with probabilities denoted as $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$. Suppose that an individual has to choose one of the following payoff vectors for consumption at time 1 : $c_1 = (1, 2, 2)$, $c_2 = (2, 1, 2)$ and $c_3 = (2, 2, 1)$. These three consumption patterns have the same distribution of consumption, giving consumption of 1 with probability 1/3 and consumption of 2 with

[9] The material in this chapter appeared as Alternative Investment Research Center WP 27, Cass Business School.

probability 2/3. Therefore, an agent with von Neumann-Morgenstern preferences would find all these consumption vectors equally attractive. However they do not cost all the same unless the state price density (and in this example, the state price) is the same in all states[10].

Dybvig (1988a) shows that the cheapest way to obtain a given payoff distribution is to allocate terminal wealth as a decreasing function of the state-price density. The cheapest price to obtain this payoff is called *payoff distributional price.* In Dybvig (1988b), the author applies this result, assuming a binomial tree model for the underlying index, and shows that for a payoff function to be efficient it should *allocate terminal wealth as a non-decreasing function of the final value of the underlying index.*[11] Intuitively, this is a plausible result as it implies that payoff and index will be positively correlated, which, when it comes to actually generating the payoff, will serve to keep the required rebalancing trades down.

We now propose a more general set of assumptions. Suppose that apart from being concerned about the terminal wealth obtained from some new investment opportunity, investors are also concerned about the dependence between this investment and their existing portfolio. This means replacing assumption 1 by:

1) Investors' preferences depend only on the joint probability distribution of terminal wealth derived from the investment and their existing portfolio.

Equivalently, since the distribution of the investor's existing portfolio will be given, we can say that:

1) Investors' preferences depend only on the probability distribution of terminal wealth derived from the investment conditional on the distribution of terminal wealth derived from their existing portfolio.

The non-satiation and perfect capital markets assumptions remain unchanged. Given this new set of assumptions, it is possible to derive an allocation rule for the cheapest

---

[10] Extracted from Dybvig and Ross (2003).

[11] Proof can also be found in Appendix A.1.

payoff function similar to the univariate case[12]. This time, however, the rule depends on the value of the investor's existing portfolio, which makes it a little more awkward to incorporate in the replication procedure. We will return to this issue in Section 2.3.

Another important paper in this context is Cox and Leland (2000). The latter show that in a Black-Scholes (1973) world all path-dependent payoff functions are inefficient because they generate payoff distributions that can also be obtained with a path-independent payoff function, but at lower costs. Our replicating payoffs will therefore not only have to allocate terminal wealth in a specific manner, but always be path-independent as well.

2.2 Univariate Replication

In this section the univariate replication proposed by Amin and Kat (2003b) is briefly reviewed.

When buying a fund participation, an investor acquires a claim to a particular payoff distribution. Therefore to evaluate the performance of a fund, one can compare the cost of the direct investment in this fund with the costing of replicating it.

The first step of their approach is finding the payoff function which maps a benchmark portfolio return distribution into the fund distribution. Not any payoff function should be considered, but the cheapest one.

In other words, what is wanted is the cheapest (path-independent) payoff function f^* that replicates the end-of-month payoff S_I of the investment or strategy, using a benchmark index or portfolio, which has end-of-month payoff S_P, i.e.,

$$P(f^*(S_P) \leq x) = P(S_I \leq x), \forall x \in \Re \tag{1}$$

Instead of working with the end-of-month payoff, the function can be written in terms of log-returns (from a monthly initial investment of 100). This is interesting because it makes the statistical modelling procedure easier. The log-returns are denoted by $X_P = \log\left(\frac{S_P}{100}\right)$, $X_I = \log\left(\frac{S_I}{100}\right)$ and the rescaled function by

$$f(x) = \log\left(\frac{f^*(100\exp(x))}{100}\right).$$

[12] Proof is provided in Appendix A.2.

The problem can be then restated as:

$$P(f(X_P) \le x) = P(X_I \le x) = F_I(x), \quad \forall x \in \Re \qquad (2)$$

The result in Dyvbig(1988a) implies that the cheapest function which satisfies (2) should be a non-decreasing function. It should be then given by:

$$f(x) = F_I^{-1}(F_P(x)), \forall x \in \Re \qquad (3)$$

It is straightforward to prove that (2) holds:

$$P(f(X_p) \le x) = P(F_I^{-1}(F_P(X_P)) \le x) = P(F_I^{-1}(U) \le x),$$

We know that U ~ Uniform[0,1], since F_P is continuous. This is the well-known probability integral transformation. We also know by the same reason that $W = F_I^{-1}(U)$ has probability distribution F_I. So $P(F_I^{-1}(U) \le x) = P(X_I \le x)$ and the result holds.

Everything can be rewritten in terms of end-of-month payoff now. The end-of-month replicated values from a monthly initial investment of *100* will be $S_f = 100\exp f(X_p) = f^*(S_p)$, where $f^*(s) = 100\exp f\left[\log\left(\frac{s}{100}\right)\right]$. S_f will theoretically have the same distribution as the original investment end-of-month payoff S_I.

For example, simulating 500 observations from a Normal(0.01 ; 0.015²) distribution for X_P and 500 observations from a Johnson-SU($\xi = 0$, $\lambda = 0.015$, $\gamma = 1.0$, $\delta = 1.658$) for X_I, we can see in Figure 2 the function f^* and the original and replicated end-of-month payoff histograms. As it was expect, f^* is a non-decreasing function.

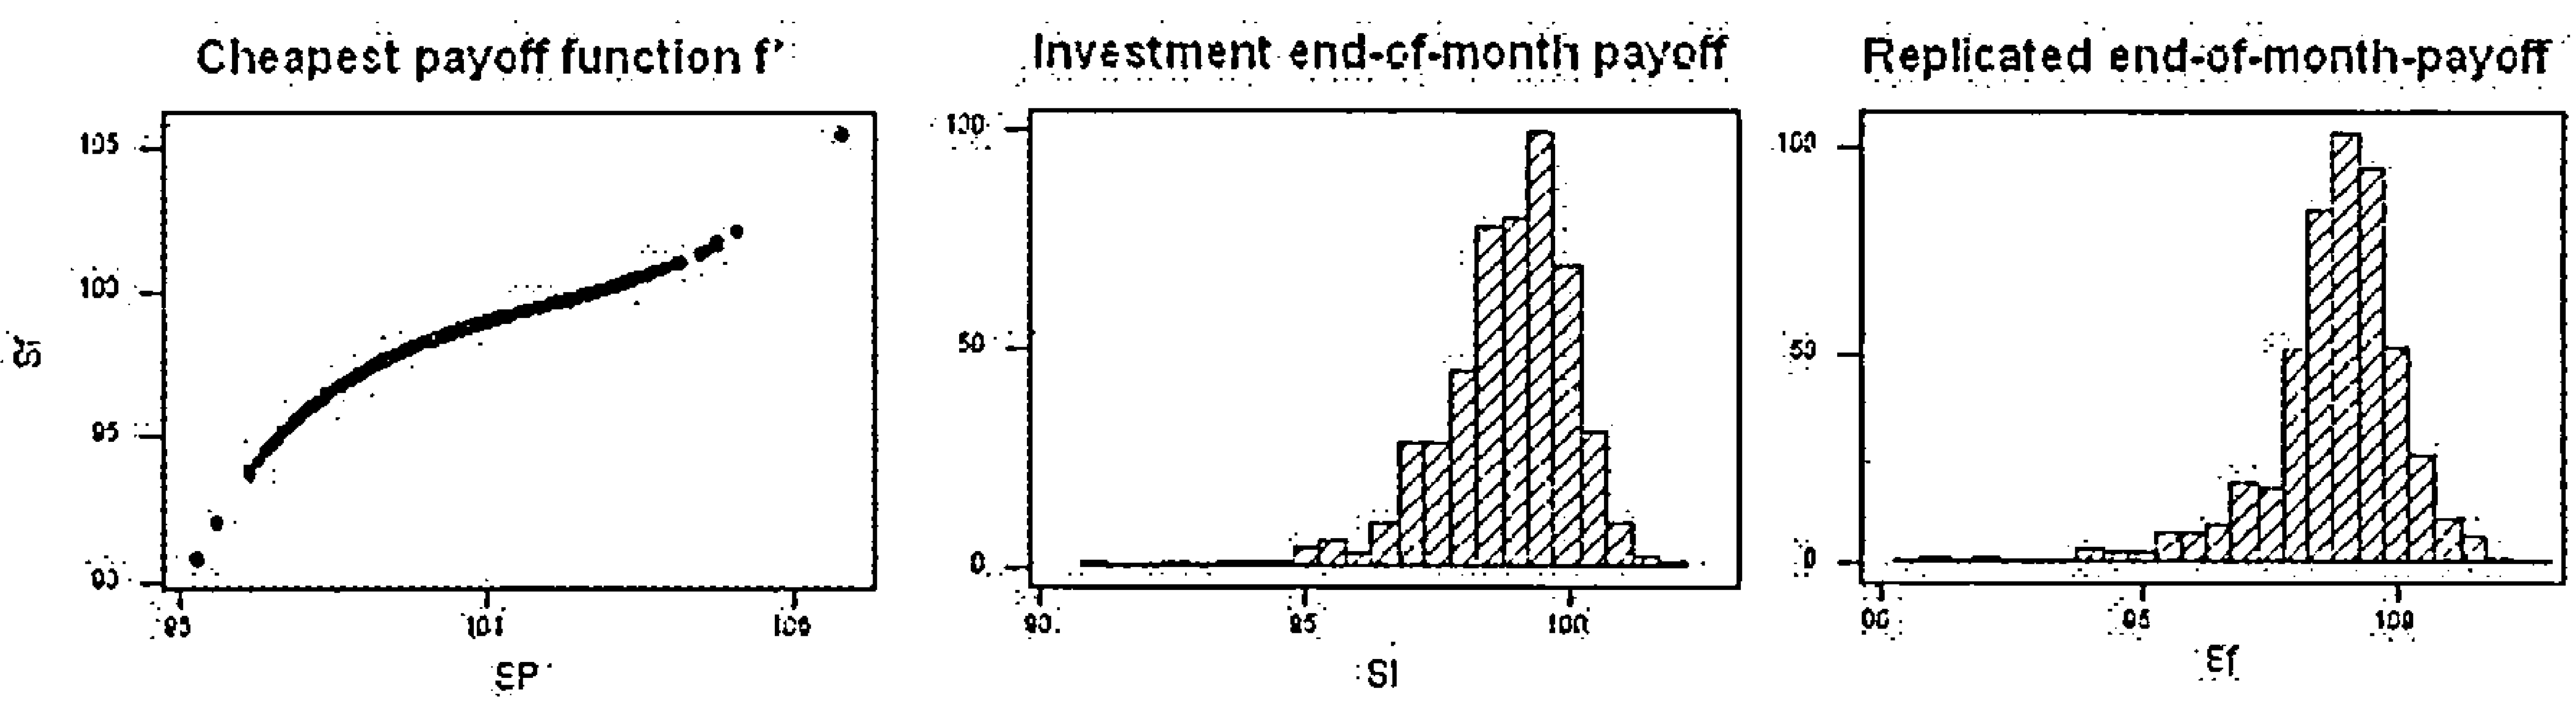

Figure 2.1. Cheapest payoff function and histogram of original and replicated payoffs for simulated data.

After finding the function f^*, it can be priced, following Harrison and Kreps (1979), as the discounted risk neutral expected payoff. This can be done by two well-known approaches in derivatives pricing: Monte Carlo simulation or Binomial tree. This technology is widely used to price options in real world applications.

Amin and Kat (2003b) considered as their benchmark portfolio the S&P500 index, assuming a normal distribution for its returns. They used the empirical distribution to model the funds returns. They applied the replication method to evaluate 77 individual hedge funds and 13 hedge fund indices over the period May-1990 to April-2000. They found 12 hedge fund indices and 72 individual funds to be inefficient, in the sense that they could be replicated at a cheaper price.

One critique of the Amin and Kat (2003b) approach is that due to their use of the S&P 500 index as the only risky asset, their replicas are all highly correlated with the US stock market. It can be clearly seen in Figure 2.1, for simulated data, that the replica is highly correlated with the benchmark portfolio. Therefore when comparing the original fund with the replica, the level of diversification provided by both series is very different. Given that hedge funds are nowadays primarily sold on the basis of their diversification benefits, replicating only their marginal distribution is not sufficient.

Also, since their work focuses only in evaluation of hedge funds, they only applied their method in-sample. The out-of-sample performance of their technique was not examined.

These two issues are addressed in this thesis. The next section of this chapter will extend their work for the bivariate case. Therefore the investor can compare the original fund and the replica not only on the basis of their marginal distribution, but also on the basis of the diversification benefit provided by both. Throughout the thesis the out-of-sample performance of this bivariate replication technique will be studied for individual funds and hedge fund indices.

2.3 Determination of the Replication Strategy

The replication procedure consists of a number of distinct steps. First, we collect return data on the fund to be replicated, the investor's portfolio, and the reserve asset (see Appendix A.2). 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. Fourth, we price the latter payoff function. Fifth, we derive the required allocations to the investor's portfolio and the reserve asset from the resulting value function.

In this section we discuss the above steps in more detail. Before we do so, however, we provide a brief introduction to copulas and their use in multivariate dependence modelling. As will become clear, copulas are a crucial ingredient in the replication procedure as they allow us to easily capture a large variety of non-normal dependence structures.

2.3.1 Copulas

Recent research in finance has uncovered various deviations from not only univariate, but also multivariate normality[13]. One powerful and at the same time convenient way to model this is by the use of copulas, as it allows the decomposition of any n-dimensional joint distribution into n marginal distributions and a single copula function[14]. Assume a random vector of two random variables. A bivariate copula can then be defined as follows.

Definition 1: *The copula C of the random vector (X,Y) is the joint distribution of the random vector (U, V), where $U = F_X(X)$ and $V = F_Y(Y)$, and where F_X, F_Y, are the distribution functions of X and Y respectively.*

The above definition implies that:

$$F_{X,Y}(x,y) = C(F_X(x), F_Y(y)), \forall x \in \Re, y \in \Re, \tag{4}$$

where F_{XY} is the joint distribution of the random vector (X,Y). Intuitively, the copula function divides the characteristics of the joint distribution between the marginal distributions, which contain the univariate characteristics of each random variable, and

[13] 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.

[14] 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.

the copula, which contains all information concerning the dependence between these random variables.

Next, we present a key result in copula theory[15]. Let $\overline{\Re} = \Re \cup \{-\infty, \infty\}$ denote the extended real line.

Sklar's Theorem: *Let $F_{X,Y}$ be a 2-dimensional joint distribution function with marginal distributions F_X, and F_Y. Then there exists a copula C such that for all (x,y) in $\overline{\Re}^2$ $F_{X,Y}(x,y) = C(F_X(x), F_Y(y))$. If F_X and F_Y are continuous then C is unique; otherwise, C is uniquely determined on Ran(F_X) $\times$ Ran(F_Y). Conversely, if C is a copula and F_X, F_Y are distribution functions, then the function F_{XY} defined by (4) is a joint distribution with margins F_X, F_Y.*

From a multivariate financial modelling perspective, it is the converse of Sklar's Theorem that is most interesting, as it implies that any combination of two univariate distributions and a copula defines a valid bivariate distribution. This solves the problem that in statistics, although we do have a large set of flexible parametric univariate distributions available, the set of parametric multivariate distributions is quite limited.

2.3.2 Estimation of the Desired and Building Block Distributions

In the replication procedure we allow three different marginal distributions (Normal, Student-t and Johnson SU)[16] 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. Since the

[15] Proof of this theorem can be found in Nelsen (1999, p. 18).

[16] See Johnson (1949, 1965) for details on the Su distribution.

Student-t copula is symmetric, however, this dependence must be the same for both tails.

The next three families of copulas, Gumbel, Cook-Johnson and Frank, are part of the *Archimedean* copulas class, a rich class of copulas that allows 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 (2006a). 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.

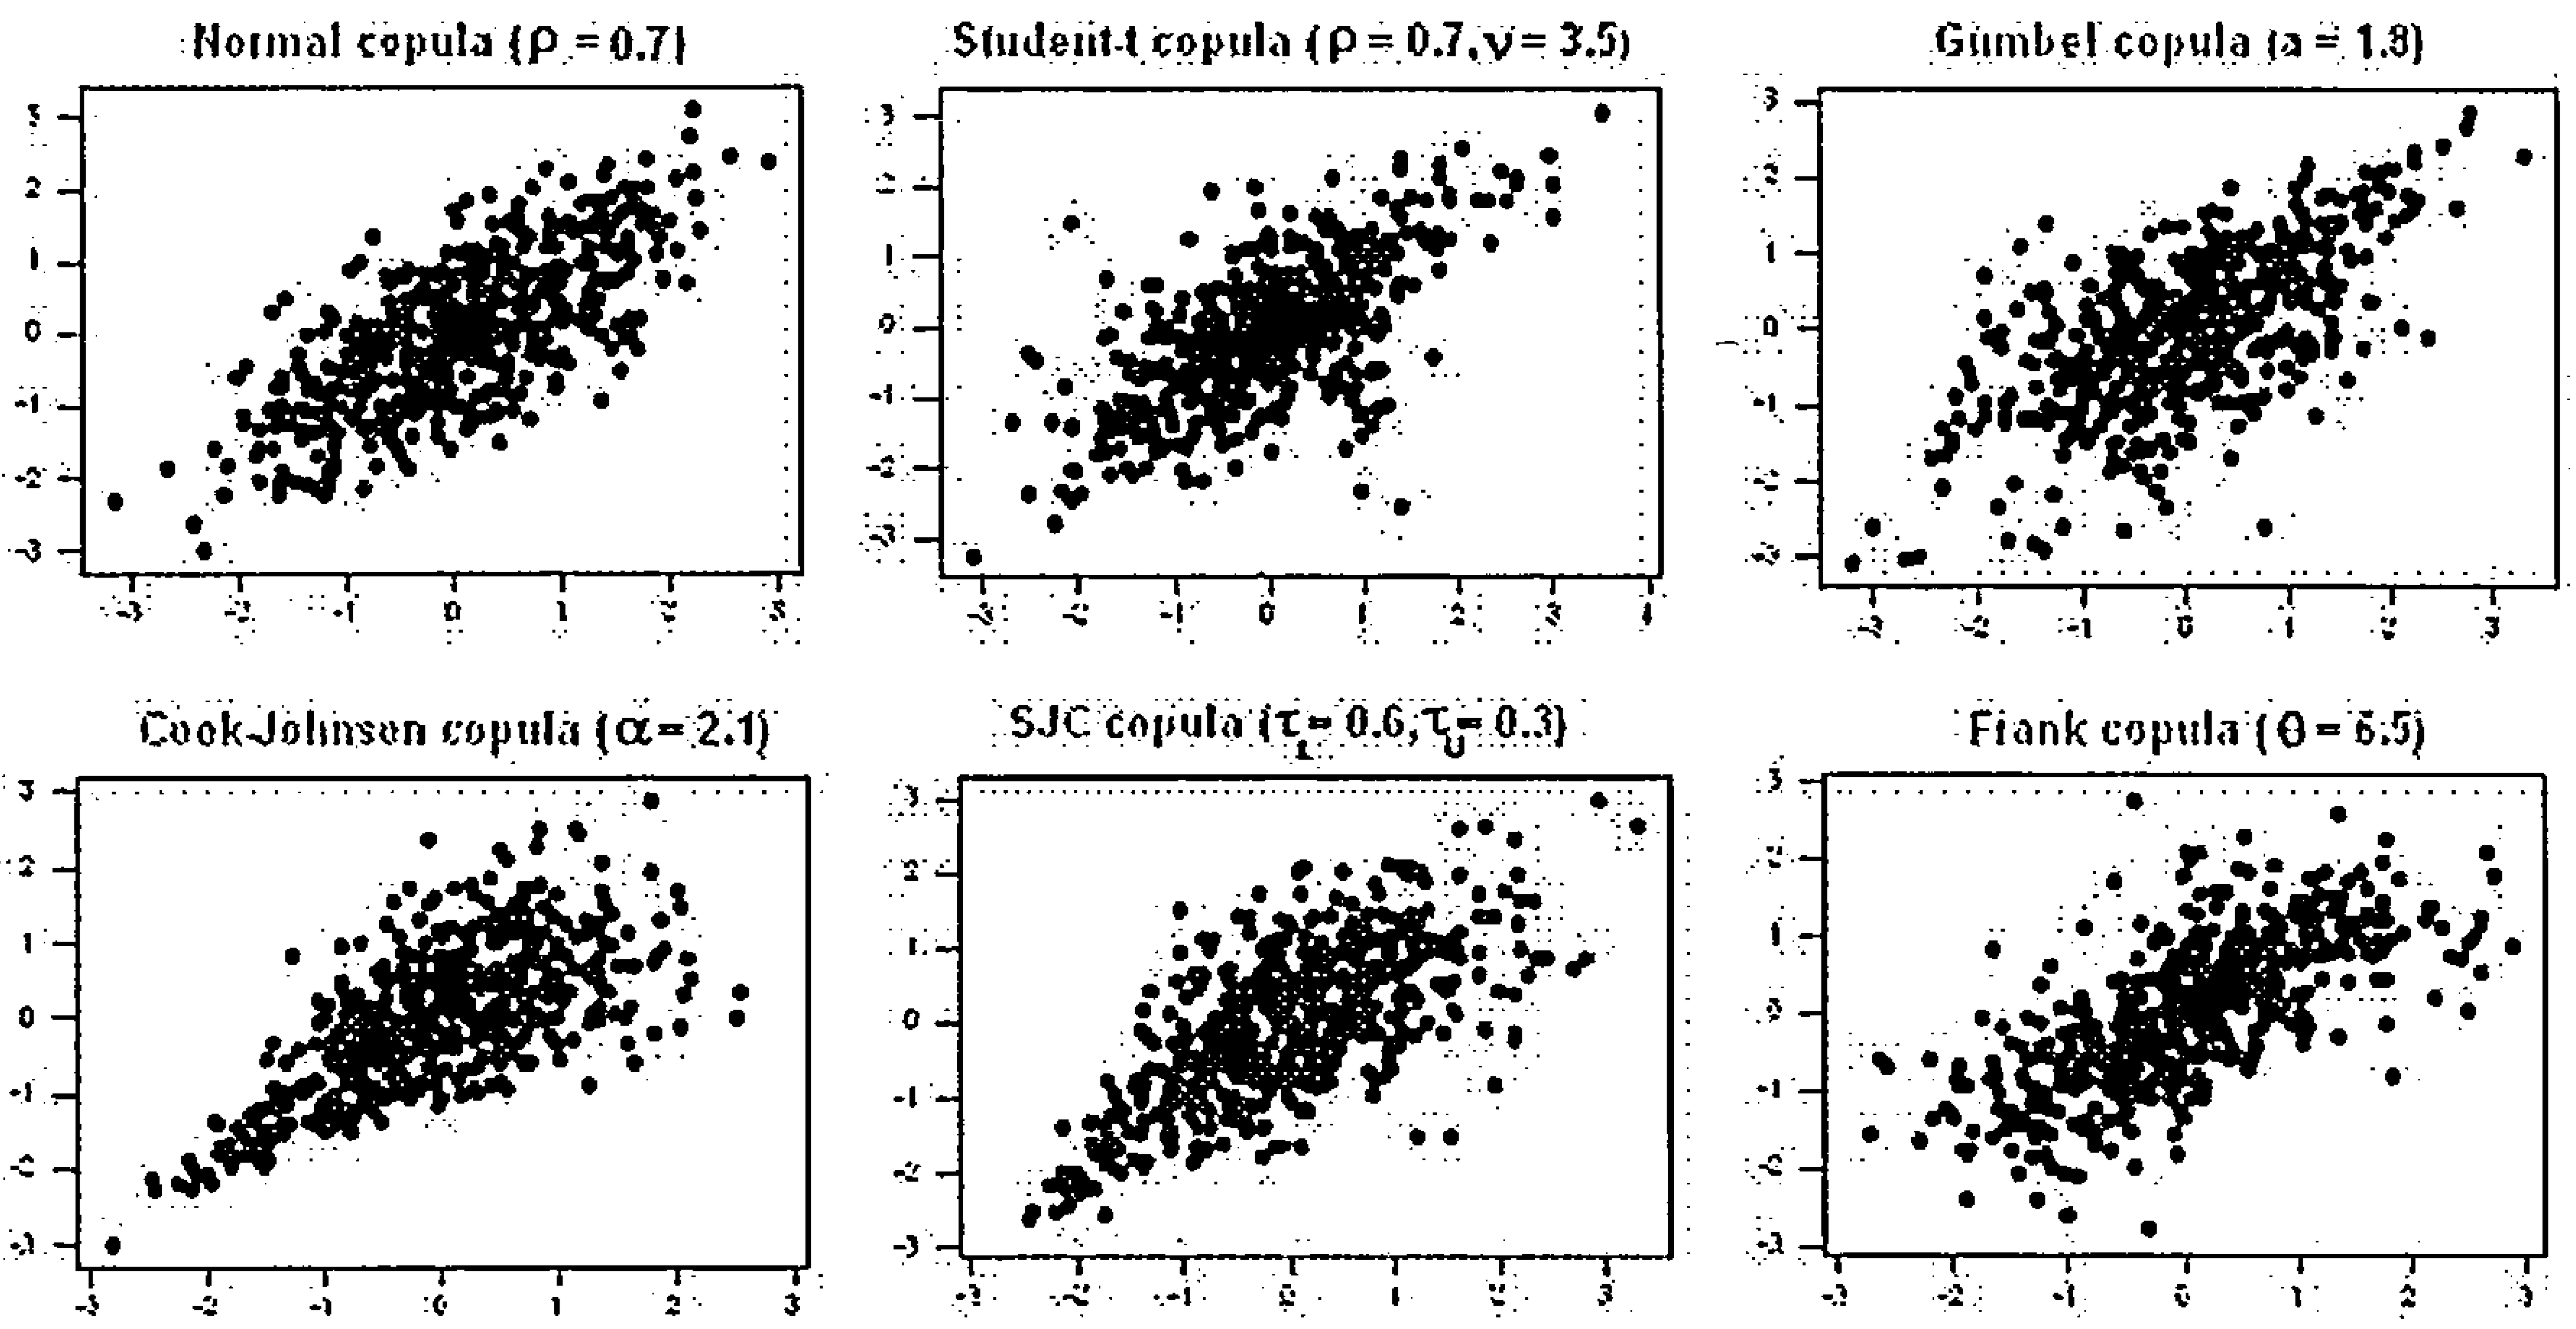

Figure 2.2. Random drawings from various copulas, assuming standard normal marginals and a linear correlation coefficient of 0.7.

Figure 2.2 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. Only in the bivariate normal case is the linear correlation coefficient sufficient to fully describe the observed dependence structure[17].

The estimation method that we use is known as the Inference Functions for Margins (IFM) method[18]. It is a two-step maximum likelihood method. Let (X,Y) be a vector of two random variables with joint distribution function F_{XY} and marginal distribution functions F_X and F_Y respectively. The marginal distribution F_X depends only on the set of parameters Θ_X and the same for F_Y and Θ_Y. Let Θ_C be the vector of parameters of the bi-dimensional copula C. So the unknown vector of parameters is given by $\Theta = (\Theta_X, \Theta_Y, \Theta_C)$. We know from Definition 1 that $F_{XY}(x,y;\Theta) = C(F_X(x;\Theta_1), F_Y(y;\Theta_2);\Theta_C)$. So the joint distribution F_{XY} is completely specified by the vector of parameters Θ. Differentiating with respect to both variables, we have $f_{XY}(x,y) = c(F_X(x), F_Y(x)) f_X(x) f_Y(y)$, where $c(u,v) = \dfrac{\partial C(u,v)}{\partial u \partial v}$ is the copula density.

For a bivariate random sample of size T $\{(x_i, y_i)\}_{i=1}^{T}$, the log-likelihood function is therefore given by:

$$l(\Theta) = \sum_{t=1}^{T} \ln c(F_X(x_t;\Theta_X); F_Y(y_t;\Theta_Y);\Theta_C) + \sum_{t=1}^{T} \ln f_X(x_t;\Theta_X) + \sum_{t=1}^{T} \ln f_Y(y_t;\Theta_Y).$$

Estimating all parameters at the same time would be very cumbersome and time-consuming. We therefore do so in two consecutive steps. First, we estimate the marginal set of parameters Θ_X and Θ_Y (separately) by maximum likelihood. Subsequently, we create the series $\hat{u}_t = F_X(x_t;\hat{\Theta}_X)$ and $\hat{v}_t = F_Y(y_t;\hat{\Theta}_Y)$ and estimate Θ_C by maximum likelihood using the likelihood function $l(\Theta_C) = \sum_{t=1}^{T} \ln c(\hat{u}_t;\hat{v}_t;\Theta_C)$.

With three possible candidates for the marginal distribution and six for the copula, we have 54 possible joint distributions to choose from. To select the final model, we use the

[17] Kat (2003) discussed this point in a hedge fund context.

[18] See Xu (1996) and Patton (2006b) for details on the statistical properties of this method.

Akaike information criterion (AIC)[19]. We considered some 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.

The copulas functions and marginal distributions used are listed in Appendix A.3.

2.3.3 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. 'Cheapest' means that we want the payoff function of the lowest possible price which generates the desired conditional distribution . In probabilistic terms, we want the cheapest function g^* such that:

$$P(S_P \leq x, g^*(S_P, S_R) \leq y) = P(S_P \leq x, S_I \leq y), \forall x, y, \tag{5}$$

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 $g(x,y) = \log\left(\frac{g^*(100\exp(x), 100\exp(y))}{100}\right)$ such that:

$$P(X_P \leq x, g(X_P, X_R) \leq y) = P(X_P \leq x, X_I \leq y) = F_{P,I}(x,y), \forall x, y, \tag{6}$$

with $X_I = \log\left(\frac{S_I}{100}\right)$, $X_P = \log\left(\frac{S_P}{100}\right)$, and $X_R = \log\left(\frac{S_R}{100}\right)$. Or equivalently, the cheapest function g such that:

[19] See Akaike (1973) for details.

$$P(g(X_P, X_R) \leq y \mid X_P = x) = P(X_I \leq y \mid X_P = x) = F_{I|P}(y \mid x), \forall x, y \tag{7}$$

Now we use the extension of Dybvig model provided in Appendix A.2, considering S_P and S_R as the underlying indices.

From Appendix A.2., we know that the cheapest payoff function depends on the conditioning value x. As a result, the bivariate function g may not be a 'smooth' function, i.e. the derivatives of this function will 'jump' around the line $x = x_{min}$, making the execution of the replication strategy derived from the payoff function quite awkward. From Appendix A.2. we find that the desired payoff function should only be a non-decreasing function of the reserve asset if

$$\frac{\mu_R - r}{\sigma_R} > \rho\left(\frac{\mu_P - x}{\sigma_P}\right), \tag{8}$$

for $\forall \rho \in [-1,1]$ and where r denotes the risk-free interest rate. The expression on the left is nothing more than the Sharpe ratio of the reserve asset. From (8) it therefore follows that as long as the Sharpe ratio of the reserve asset is high enough and the correlation with the investor's portfolio low enough, the desired payoff function should be a non-decreasing function of the reserve asset.

Assuming the reserve asset satisfies the above condition[20], the function g in expression (7) is given by:

$$g(x, y) = F_{I|P}^{-1}(F_{R|P}(y \mid x) \mid x), \forall y \in \Re \tag{9}$$

where $F_{I|P}^{-1}(y \mid x)$ denotes the pseudo-inverse of $F_{I|P}(y \mid x)$. 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 (7) holds:

[20] Extensive simulations showed that, under reasonable assumptions, this does not introduce any significant error if not true for some values of x.

$$P(g(X_P,X_R)\le y\,|\,X_P=x)=P(g(x,X_R)\le y\,|\,X_P=x)=$$
$$P(F_{I|P}^{-1}(F_{R|P}(X_R\,|\,x)\,|\,x)\le y\,|\,X_P=x)=P(F_{I|P}^{-1}(U\,|\,x)\le y\,|\,X_P=x), \tag{10}$$

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

$$P(F_{I|P}^{-1}(U\,|\,x)\le y\,|\,X_P=x)=P(X_I\le y\,|\,X_P=x)=F_{I|P}(y\,|\,x), \tag{11}$$

and (7) 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 (4) we have:

$$F_{P,I}(x,y)=C_{P,I}(F_P(x),F_I(y)),x\in\Re,y\in\Re. \tag{12}$$

$$F_{P,R}(x,y)=C_{P,R}(F_P(x),F_R(y)),x\in\Re,y\in\Re. \tag{13}$$

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

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

$$F_{R|P}(y\,|\,x)=\kappa_x^{P,R}(y),x\in\Re,y\in\Re\text{, where }\kappa_x^{P,R}(y)=\frac{\partial C_{P,R}(u,v)}{\partial u}\Big|_{u=F_P(x),v=F_R(y)}.$$

So the cheapest function g in expression (9) can be rewritten as:

$$g(x,y)=\kappa_x^{(-1)P,I}\left(\kappa_x^{P,R}(y)\right),x\in\Re,y\in\Re. \tag{14}$$

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:

$$S_g = g^*(s_P, s_R) = 100 \exp g\left(\log\left(\frac{s_P}{100}\right), \log\left(\frac{s_R}{100}\right)\right). \qquad (15)$$

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.

2.3.4 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 arbitrage-based option pricing theory has concentrated on for the last 35 years. Following Harrison and Kreps (1979), the desired payoff function can be priced by calculating the discounted risk neutral expected payoff. Using the notation in the previous subsections, we would have the price ψ of function g^* given by:

$$\psi = \frac{1}{r} E_Q[g^*(S_P, S_R)],$$

where Q denotes de risk-neutral expectation, hence supposing that

$$\begin{pmatrix} X_P \\ X_R \end{pmatrix} \sim N\left(\begin{pmatrix} r - \frac{\sigma_P^2}{2} \\ r - \frac{\sigma_R^2}{2} \end{pmatrix}, \begin{pmatrix} \sigma_P^2 & \sigma_{PR} \\ \sigma_{PR} & \sigma_R^2 \end{pmatrix}\right).$$

In the absence of transaction costs, the two most obvious methods to do so would be either bivariate Monte Carlo simulation or a trinomial tree[21].

But transaction costs need to be considered, otherwise it would be an unfair comparison between the replica and the original fund. Therefore we use the multivariate option pricing model of Boyle and Lin (1997). Their model examines the incorporation of transaction costs when there is more than one risk asset and it is costly to trade in each risk asset. A discrete time framework is used and the problem of option valuation if reformulated as a linear programming problem. The price of the payoff function will be called 'KP measure'. In Subsection 2.3.3 we assumed an initial investment of 100, so a

[21] See Jaeckel (2002) or Glasserman (2003) for an introduction to Monte Carlo methods. Details on the trinomial tree approach can be found in He (1990).

KP measure less than 100 would mean that the replication strategy is cheaper than the payoff fund, and the other way around if KP is greater than 100.

Once we are able to price the desired payoff function, we can work out the controls of the dynamic trading strategy generating it by straightforward partial differentiation of the value function. In computational terms, we work backwards on the trinomial tree and estimate the delta using the values of the payoff function g^* and the vector (S_P, S_R) in the three states of the first step of the tree.

Only after pricing the payoff function 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.

In Appendix A.4 the influence of the moments of the reserve asset distribution and its correlation with the reference portfolio on the KP measure is examined. The conclusion is that we want the highest possible risk premium and the smallest possible volatility for the reserve asset, as one would expect. The influence of the correlation between the reserve asset and the reference portfolio is not so simple.

2.4 Simulation Analysis

Given the desired and building block distributions, the above results allow us to derive, price and generate the cheapest payoff function that turns one into the other. The procedure is exact, so by itself it does not require any testing. Taking this procedure into the real world and using it to replicate fund returns, however, we are confronted with a number of problems. First, we do not know the true population distribution. The best we can do is estimating it from a small data sample. Second, the latter distribution may not be stationary over time. Third, due to market imperfections and insufficient information on the underlying price processes, we may not always be able to exactly generate the desired payoff function.

In this section, we use simulation methods to study the error resulting from determining the desired payoff function from a relatively small sample, instead of from the population distribution. In these simulations, we assume that the population distribution is stationary and that it is possible to generate the desired payoff function without any error. In the next section we perform a number of out-of-sample tests on real-life data to also include the error contributions of non-stationarity and sub-optimal dynamic trading. In the simulations, we study two different cases, selected to capture different distributional conditions. Throughout, we assume that the returns on the investor's portfolio and the reserve asset are both normally distributed with the parameter values given below. In addition, we assume they are related through a Gaussian copula with a correlation coefficient of 0.3.

Investor's portfolio

Log-returns - $X_P \sim N(0.01, 0.043301^2)$

Mean = 12% p.a.

Volatility = 15% p.a.

Reserve asset

Log-returns – $X_R \sim N(0.00833, 0.028867^2)$

Mean = 10% p.a.

Volatility = 10% p.a.

2.4.1 Gaussian Fund Marginal, Higher Dependence in the Lower Tail.

Our first case assumes that the fund return is normally distributed, but that the relationship with the investor's portfolio is such that there is more dependence in the lower than in the upper tail. This could be the risk profile of a fund of funds with a bias towards risk arbitrage for example. The marginal distribution of the fund return and the relevant copula are specified as follows:

Fund

Log-returns – $X_I \sim N(0.015, 0.057735^2)$

Mean = 18% p.a.

Volatility = 20% p.a.

Copula (investor's portfolio, fund) = SJC (0.75, 0.10)

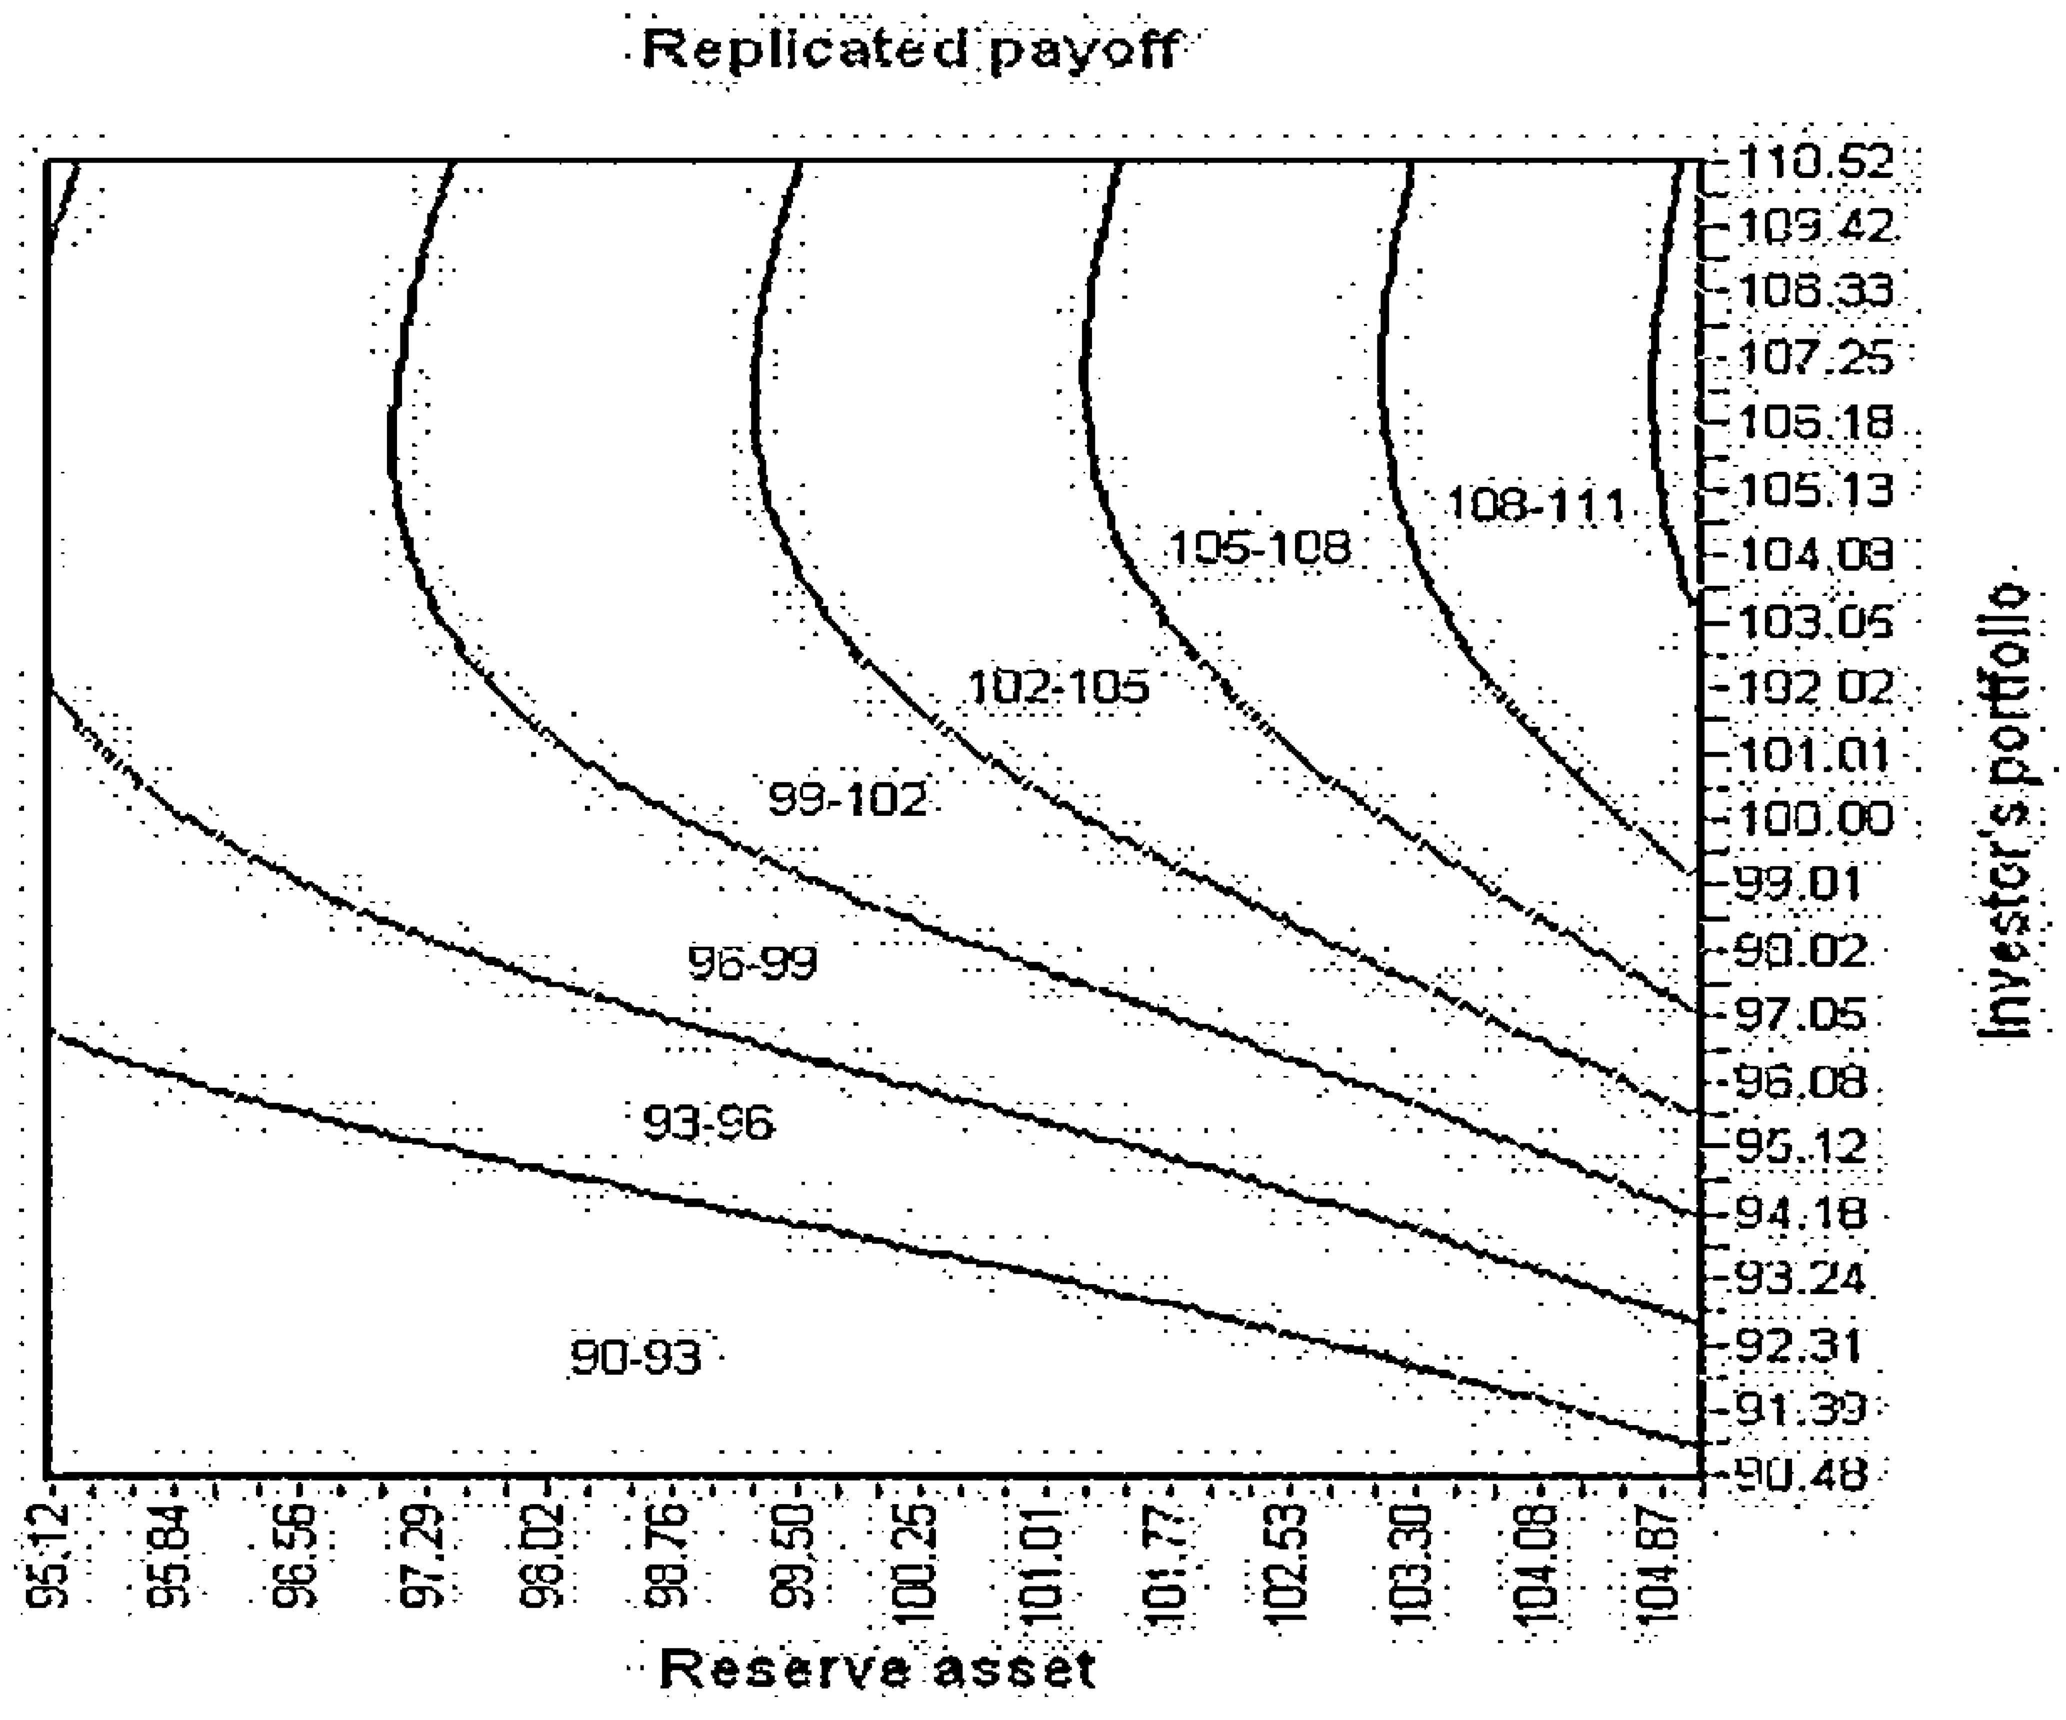

Figure 2.3. Contour plot payoff function from population distribution case 1.

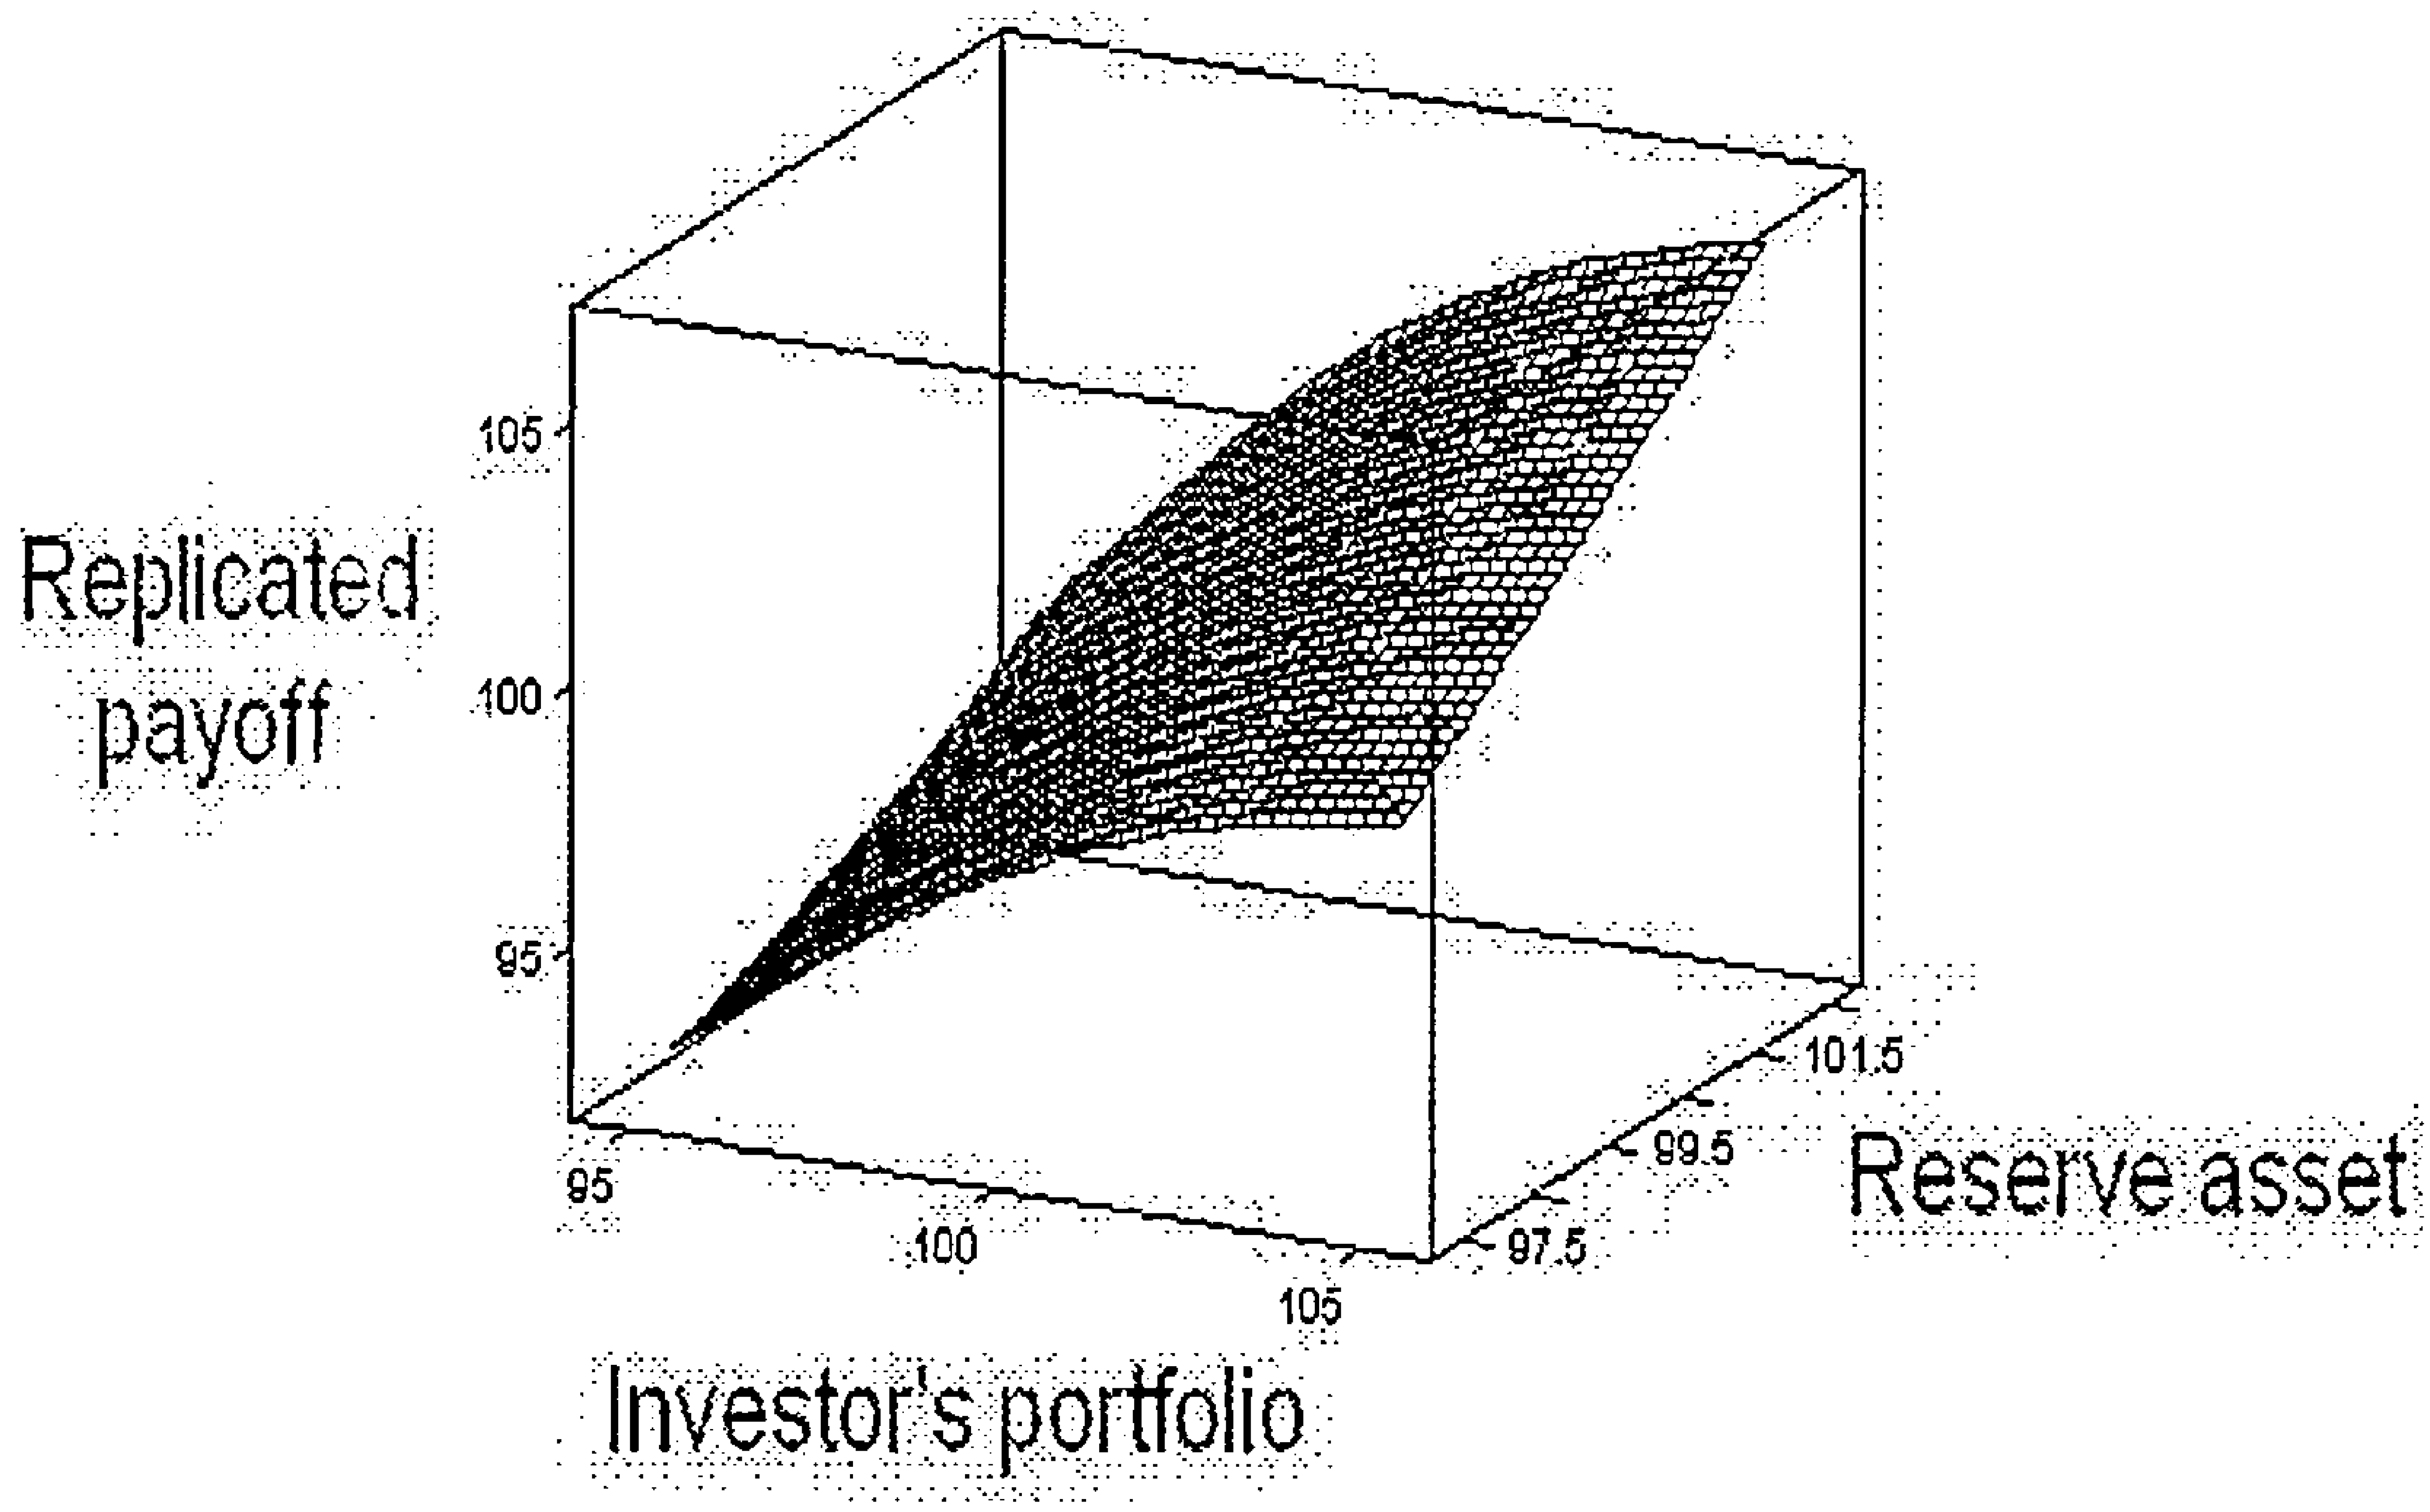

Figure 2.4. 3D plot payoff function from population distribution case 1.

Give the desired and building block distributions, we derived the desired payoff function using the results of Section 2.3. Figure 2.3 and 2.4 depict the latter graphically, as a contour plot as well as a 3D graph. From the graphs we see that the desired payoff is an increasing function of the reserve asset (by construction) as well as the investor's portfolio. The strategy's controls will therefore tell us to hold long positions in both assets. As is especially clear from the contour plot, the payoff function is quite curved. This of course serves to generate the required difference in dependence between the upper and lower tail.

To gain insight into the potential error when deriving the desired payoff function from only a small sample, instead of the population distribution, we took a sample of size N and derived a payoff function from it. Subsequently, we took 2000 observations from the building block distribution, and fed these observations through the latter payoff function to produce a joint distribution of replicated payoffs and the investor's portfolio. From the latter distribution we calculated the mean, standard deviation, skewness, and kurtosis of the replicated payoff as well as its correlation with the investor's portfolio. The above procedure was repeated 100 times, for different values of N (= 24, 48, 72, 96, 120, 240). Across each set of 100 runs, we subsequently calculated the mean, standard deviation and skewness of the replication errors, i.e. the differences between the above sample statistics and the true fund parameters. The results can be found in Table 2.1.

	Mean	St. Dev	Skewness	Excess Kurtosis	Corr. with Portfolio
Fund	101.6805	5.8267	0.0000	-0.0447	0.7400
Differences (replicated - fund)					
Avg 24	-0.0315	0.0416	0.0174	0.0952	-0.0171
Avg 48	0.1010	-0.0090	0.0156	0.0835	-0.0071
Avg 72	0.1390	0.0689	0.0052	0.0843	-0.0020
Avg 96	0.0151	0.0114	0.0187	0.0781	-0.0148
Avg 120	0.0142	-0.0079	0.0038	0.0935	-0.0011
Avg 240	-0.0209	-0.0118	0.0053	0.0595	0.0015
SD 24	1.0386	1.0253	0.2052	0.4181	0.1234
SD 48	0.8287	0.6701	0.1206	0.1641	0.0933
SD 72	0.5975	0.5496	0.1099	0.1613	0.0726
SD 96	0.5023	0.5620	0.1299	0.1712	0.0656
SD 120	0.4733	0.4794	0.1026	0.1437	0.0576
SD 240	0.3352	0.2918	0.0884	0.1465	0.0402
SK 24	0.3765	-0.0241	0.9582	4.4101	-1.1399
SK 48	0.3071	0.6894	0.3059	0.1667	-0.8320
SK 72	-0.2267	-0.3145	0.0991	0.1401	-0.5869
SK 96	-0.0114	0.2892	-0.1812	0.4813	-0.7360
SK 120	-0.0748	0.4211	-0.1757	0.1810	-0.3530
SK 240	-0.0991	0.3519	-0.1098	-0.0260	-0.4008

Table 2.1. Variation due to payoff construction from small sample case 1.

To be able to properly interpret the entries in the table, the first row in Table 2.1 shows the mean, standard deviation, skewness, kurtosis and correlation of the fund payoff, as implied by the assumed fund return distribution. The rows that follow show, for various sample sizes N (= 24, .., 240) and each over 100 runs, the average (Avg), standard deviation (SD) and skewness (SK) of the replication errors. Table 2.1 confirms that the larger the sample, the more accurate the desired payoff function will be. It also shows that even with a relatively small sample the procedure still works quite well and is unbiased. For all parameters and sample sizes, the average error is statistically insignificant at 5% (-1.96 SD, + 1.96 SD).

2.4.2 Negatively Skewed Fund Marginal, Gaussian Copula

Our second case is somewhat more extreme. It assumes that the fund return exhibits a high degree of negative skewness. To make up for that, however, it also has a relatively high mean and low correlation with the investor's portfolio. With a little imagination, this could be the risk profile of a fixed income arbitrage fund for example. The marginal distribution of the fund return and the relevant copula are specified as follows:

Fund

Log-returns – X_I ~ Johnson-SU (0.058604, 0.046978, 0.926426, 1.390468)

Mean = 18% p.a.

Volatility = 20% p.a.

Skewness = -2.0

Excess kurtosis = 10

Copula (investor's portfolio, fund) = Gaussian (0.2).

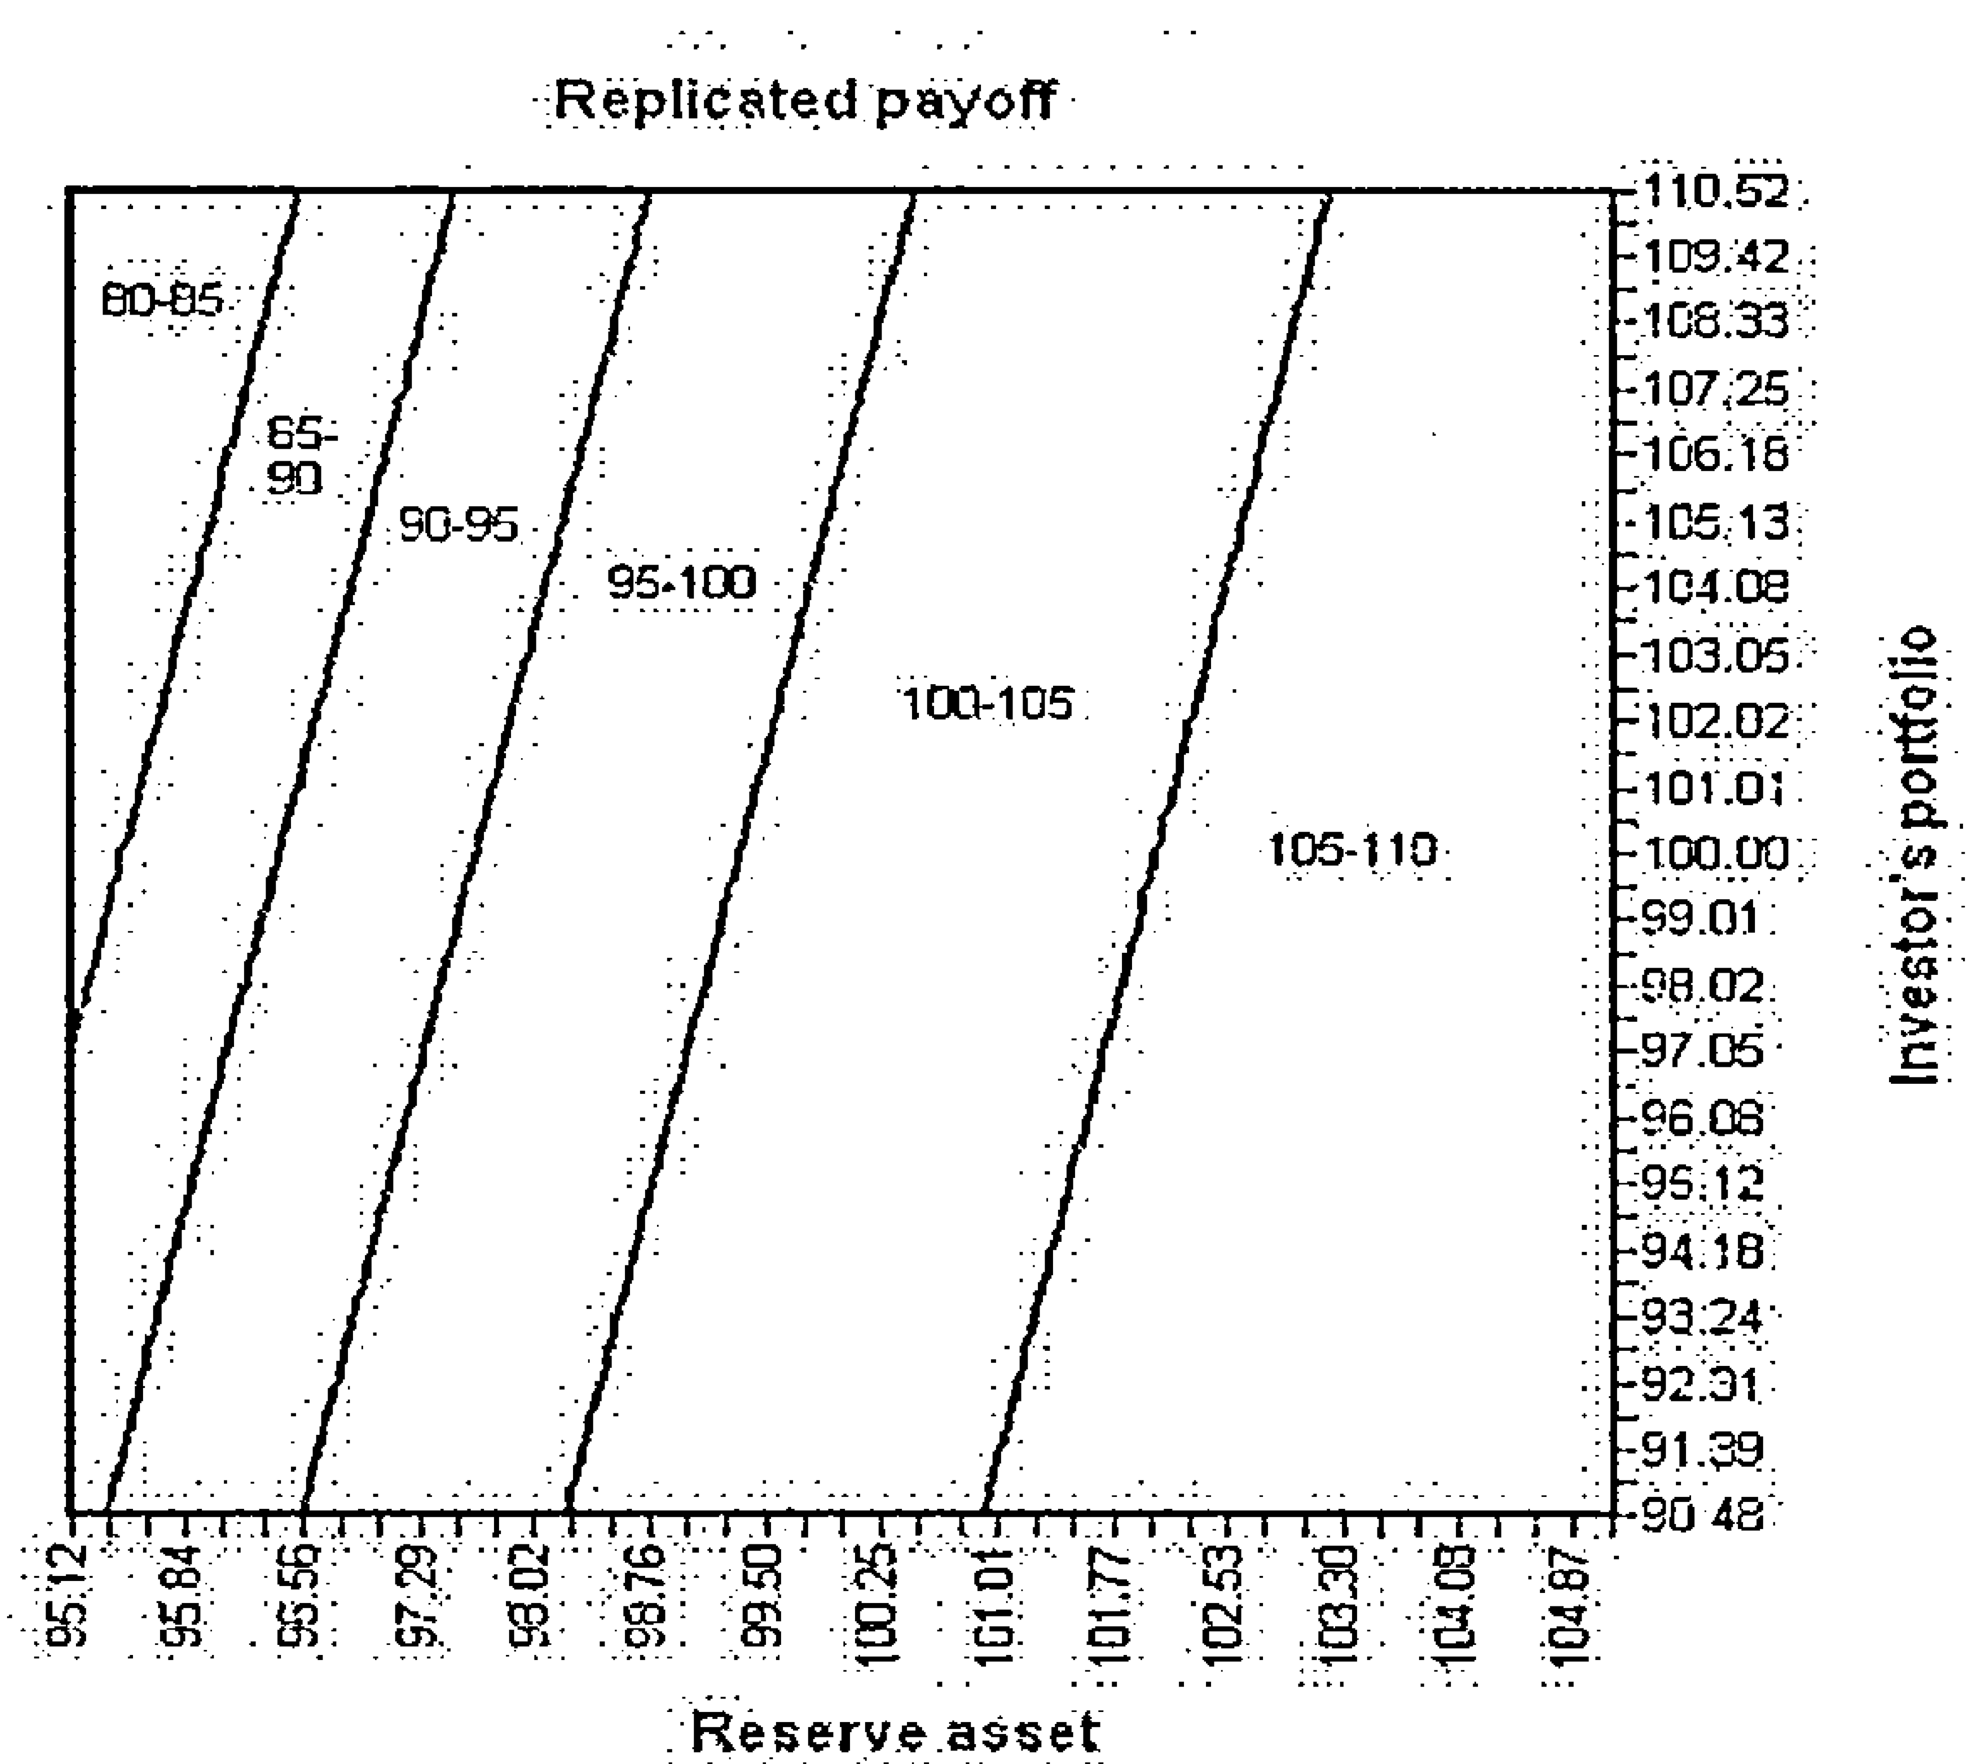

Figure 2.5. Contour plot payoff function from population distribution case 2.

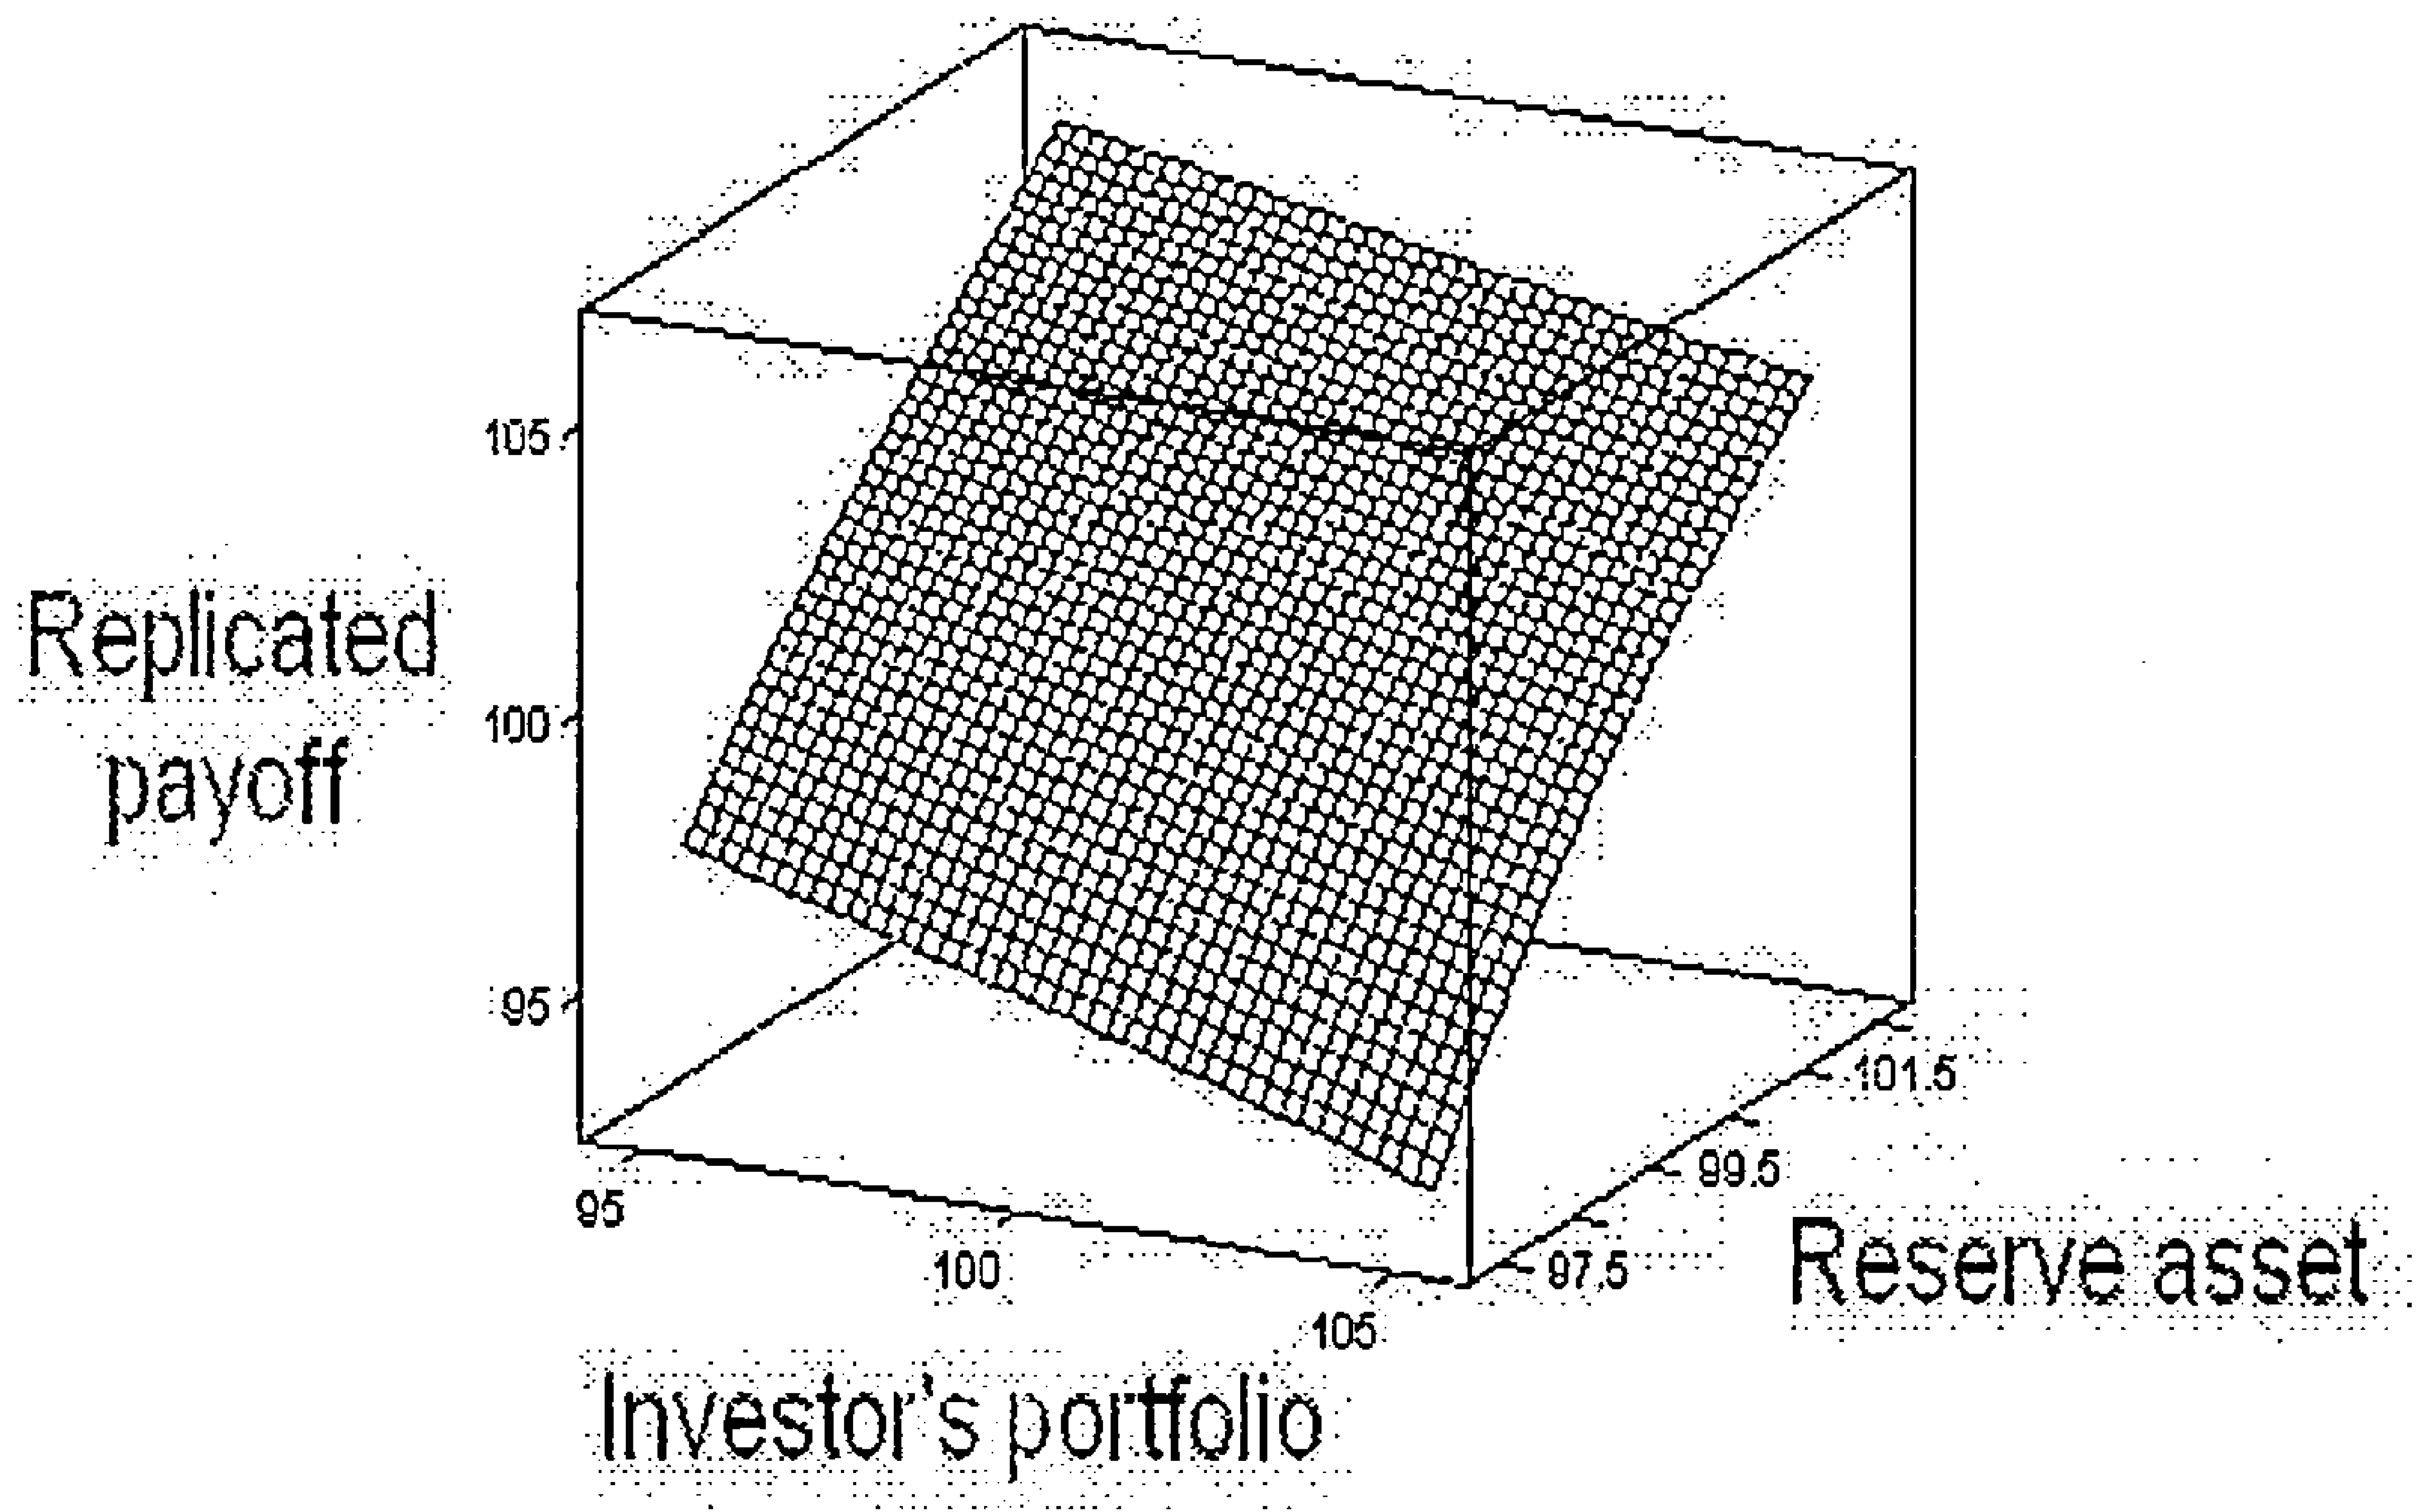

Figure 2.6. 3D plot payoff function from population distribution case 2.

From the above population distributions we again derived the desired payoff function, which is graphically depicted in Figure 2.5 and 2.6. As required, the payoff is a positive function of the reserve asset. However, since the assumed correlation between the fund and the investor's portfolio is lower than the assumed correlation between the investor's portfolio and the reserve asset, the payoff is a negative function of the investor's portfolio. The strategy's controls will therefore want us to go long in the reserve asset, but short in the investor's portfolio. The slope of the payoff function increases as the investor's portfolio rises and the reserve asset drops, which serves to generate the required negative skewness.

	Mean	St. Dev	Skewness	Excess Kurtosis	Corr. with Portfolio
Fund	101.6805	5.8604	-1.9851	9.8511	0.2000
Differences (replicated - fund)					
Avg 24	-0.0798	0.3568	1.1919	-0.0162	-0.0454
Avg 48	0.0221	0.5960	1.2455	4.5453	-0.0341
Avg 72	-0.1577	0.3589	0.7876	-2.4691	0.0450
Avg 96	-0.0201	0.1341	0.7929	-1.9950	-0.0152
Avg 120	0.0300	0.0985	0.6538	-1.6140	0.0090
Avg 240	-0.0762	0.1846	0.5238	-1.0256	0.0086
SD 24	1.6094	2.0242	1.1313	19.8187	0.2882
SD 48	1.3526	1.5985	1.7192	41.6658	0.2205
SD 72	0.8821	1.2115	0.6218	2.9127	0.1374
SD 96	0.7701	1.1864	0.6591	3.5340	0.1385
SD 120	0.7622	0.9344	0.5907	2.7828	0.1136
SD 240	0.3818	0.7366	0.6266	4.4289	0.0812
SK 24	-0.2246	0.9298	3.9961	8.2333	-0.1064
SK 48	0.0573	0.5893	5.6316	8.5348	0.1809
SK 72	0.0405	0.6589	0.1550	-0.2631	-0.0549
SK 96	0.0785	0.7394	0.9729	4.0508	-0.4042
SK 120	-0.1274	0.4989	0.0159	0.2908	-0.1725
SK 240	0.0587	0.6556	1.0541	5.0328	0.0376

Table 2.2. Variation due to payoff construction from small sample case 2.

To gain insight into the potential error from deriving the desired payoff function from a small sample in this particular case, we repeated the procedure used earlier in case 1. The results can be found in Table 2.2. Not unexpectedly, given the much more extreme distributional assumptions, we see quite some variation for small sample sizes. Especially the replication of the assumed –2.0 skewness meets with some difficulty. Since, by definition, tail events only occur infrequently, many smaller samples will not contain enough information to estimate skewness accurately. This is reflected by the strong positive skew of the error distribution for small *N*.

The above two case studies suggest that, depending on the distributions involved, the error from working with a small sample may sometimes be quite substantial. It is important to note though, that when applying the procedure in practice, one will typically re-estimate the payoff function periodically as new fund return data becomes available. Through time therefore, these errors may diversify away to some extent.

2.5 Out-of-Sample Tests

We proceed with some out-of-sample tests. The out-of-sample tests that follow are all structured in the same way. Given a fund, 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 1985 for example, we assume to be living on January 1st, 1987. 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[22]. The above is repeated until we arrive at the end of the series.

Throughout we assume the investor's portfolio consists of 50% US equity, in the form of the S&P 500 tracking portfolio, and 50% long-dated US Treasury bonds. We use nearby Eurodollar futures as the reserve asset[23]. To minimize transaction costs, all trading is done in the futures markets[24]. Throughout, we trade the nearby futures contract, rolling into the next nearby contract on the first day of the expiry month, assuming transaction costs of 1bp one-way[25]. The necessary volatility and correlation inputs are obtained from historical estimates, using all available data at the time of determining the desired payoff function.

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 among investors and because they have relatively long track records. 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

[22] 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.

[23] The decision to use Eurodollar futures is primarily based on liquidity and simplicity, since these are illustrative examples. In Chapters 3, 4 and 5 a more diversified reserve asset will be used.

[24] S&P 500 (SP) and Eurodollar futures (ED) are traded on the CME, while T-bond futures (US) are traded on the CBOT.

[25] Commission costs in futures tend to be extremely low, while quoted bid-ask spreads in the most liquid contracts are typically not high as well.

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. We do not charge any management fees in the replication strategy.

2.5.1 Leveraged Capital Holdings N.V.

Our first example concerns one of the first funds of hedge funds. Leveraged Capital Holdings (LCH) was started in 1969 (our return data, however, only start in 1985) by Georges Karlweis of Banque Privee Edmond de Rothchild in Geneva. Over the years, LCH has (been) invested in all well-known hedge fund managers, such as George Soros, Martin Zweig and Joseph DiMenna, and Michael Steinhardt. LCH is publicly listed on the Amsterdam Stock Exchange and in October 2005 had $1.32 billion under management (TASS database).

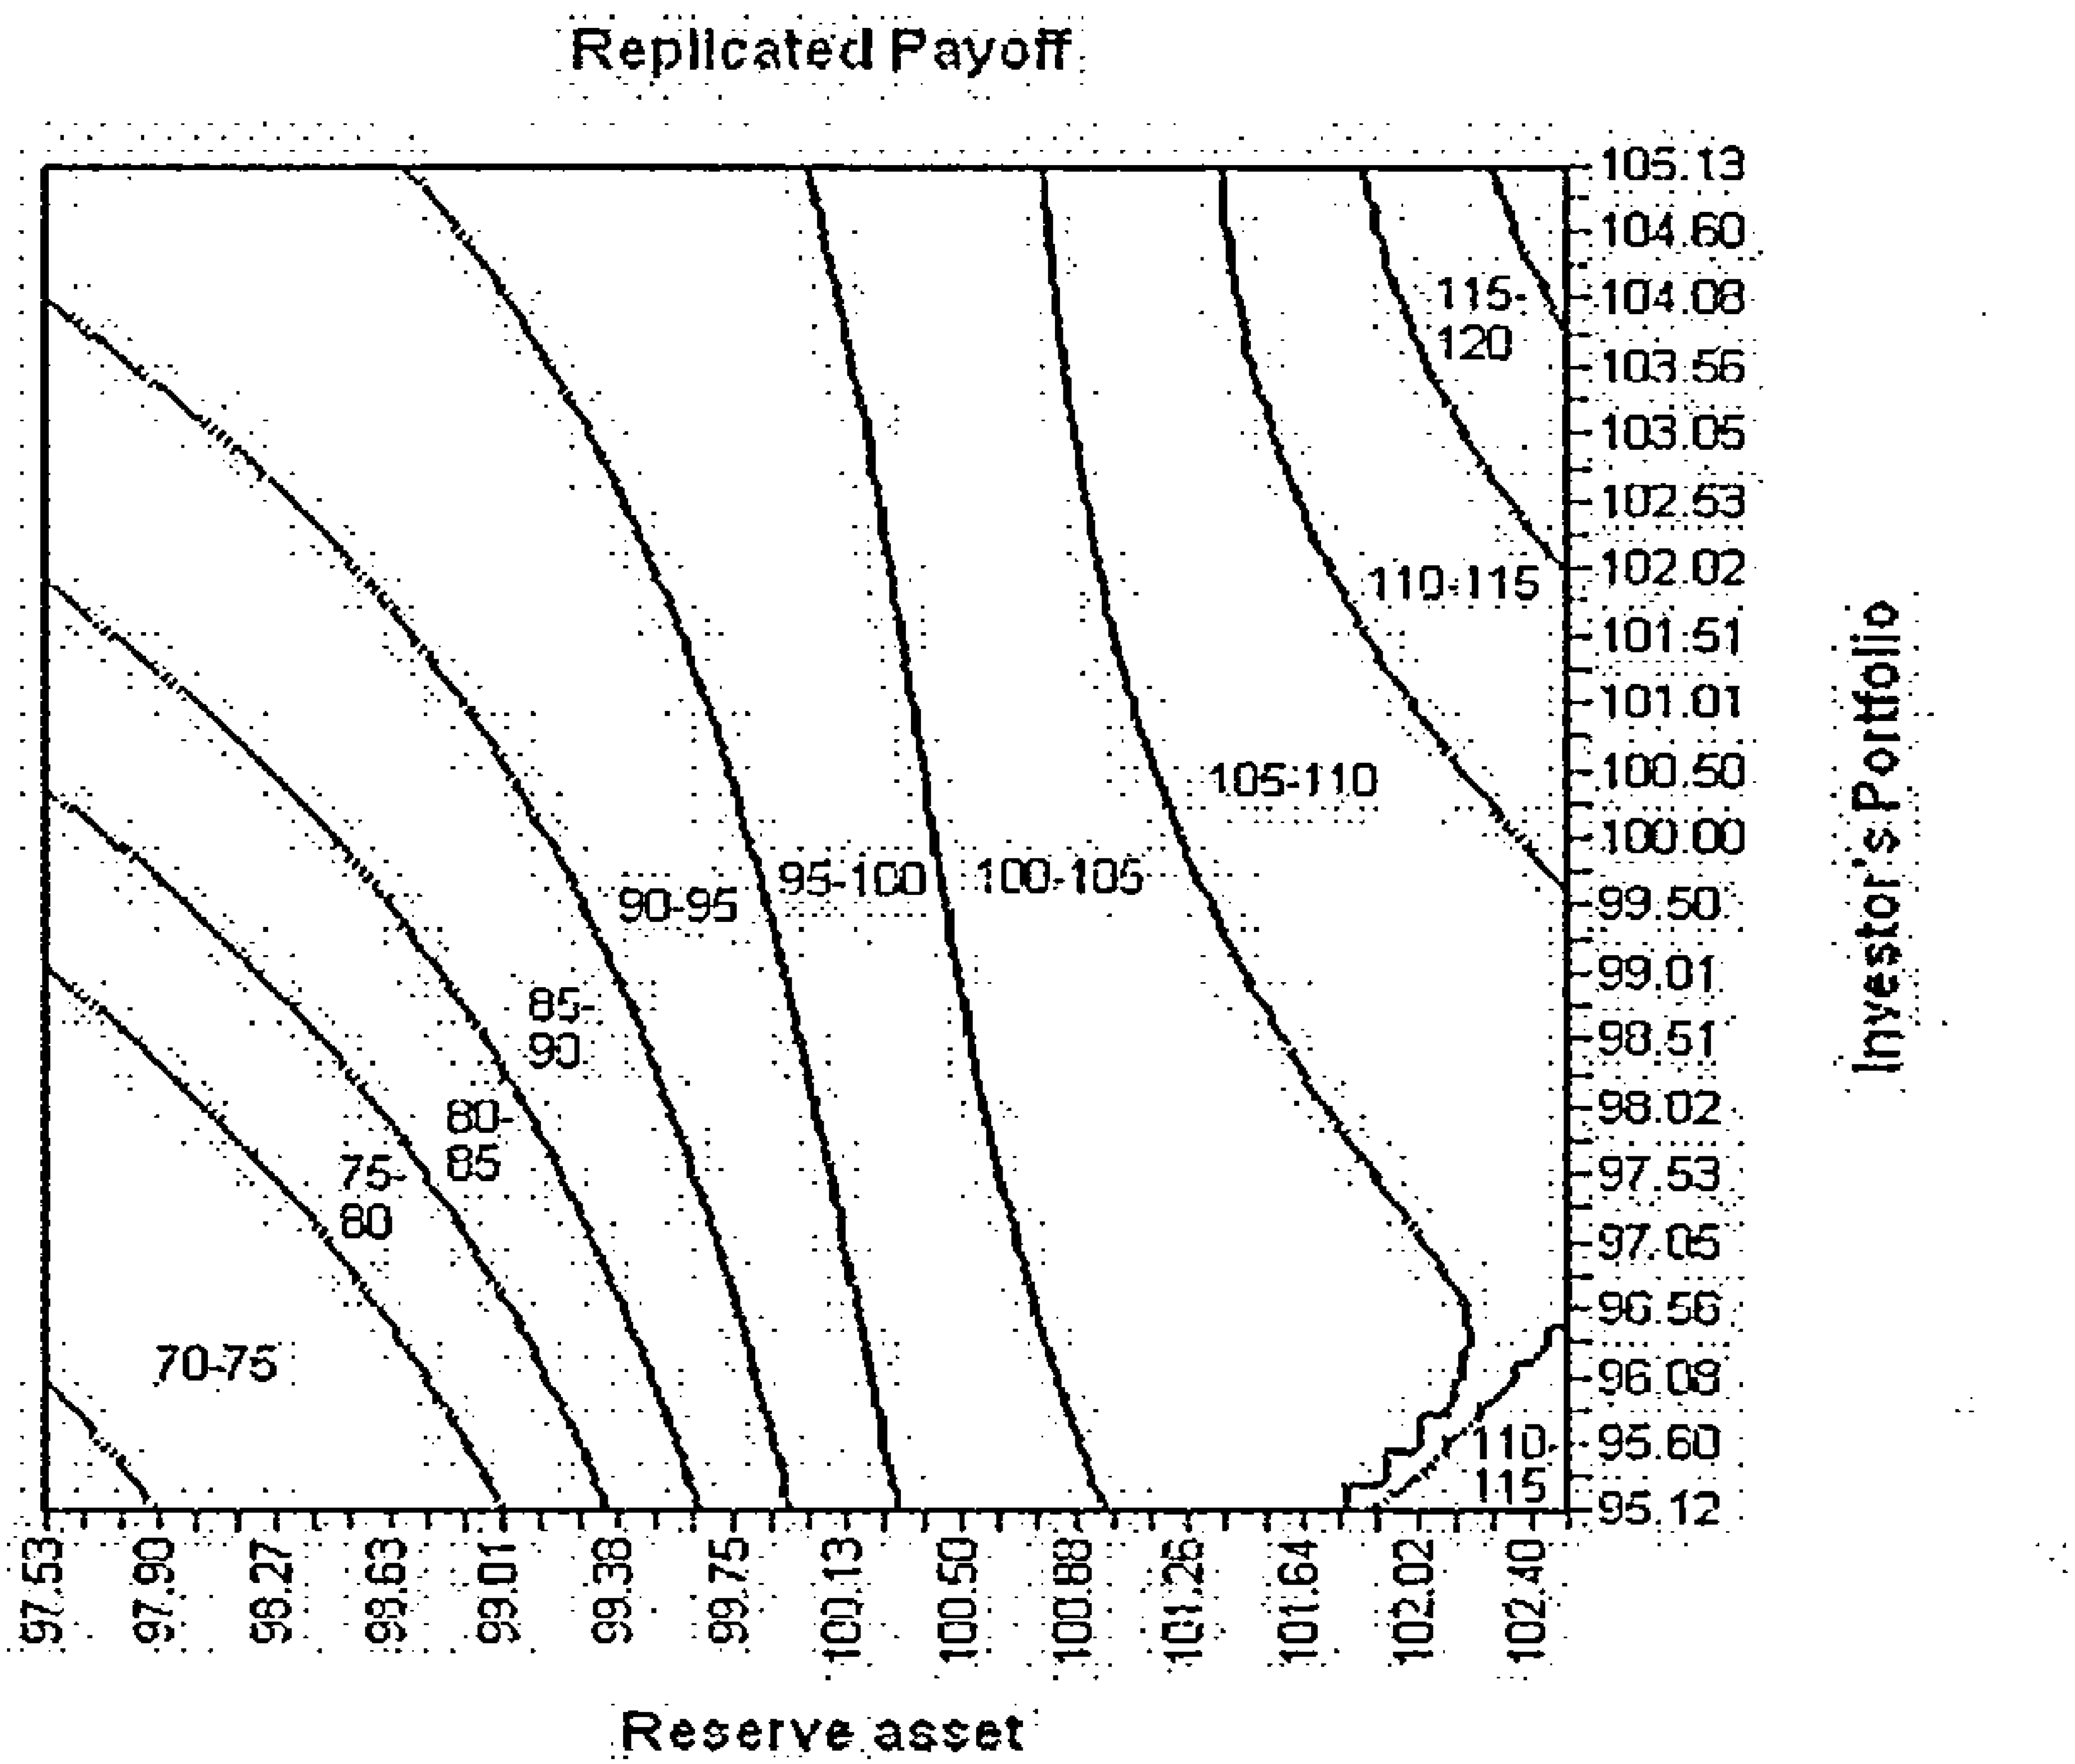

Figure 2.7. Contour plot payoff function Leveraged Capital Holdings.

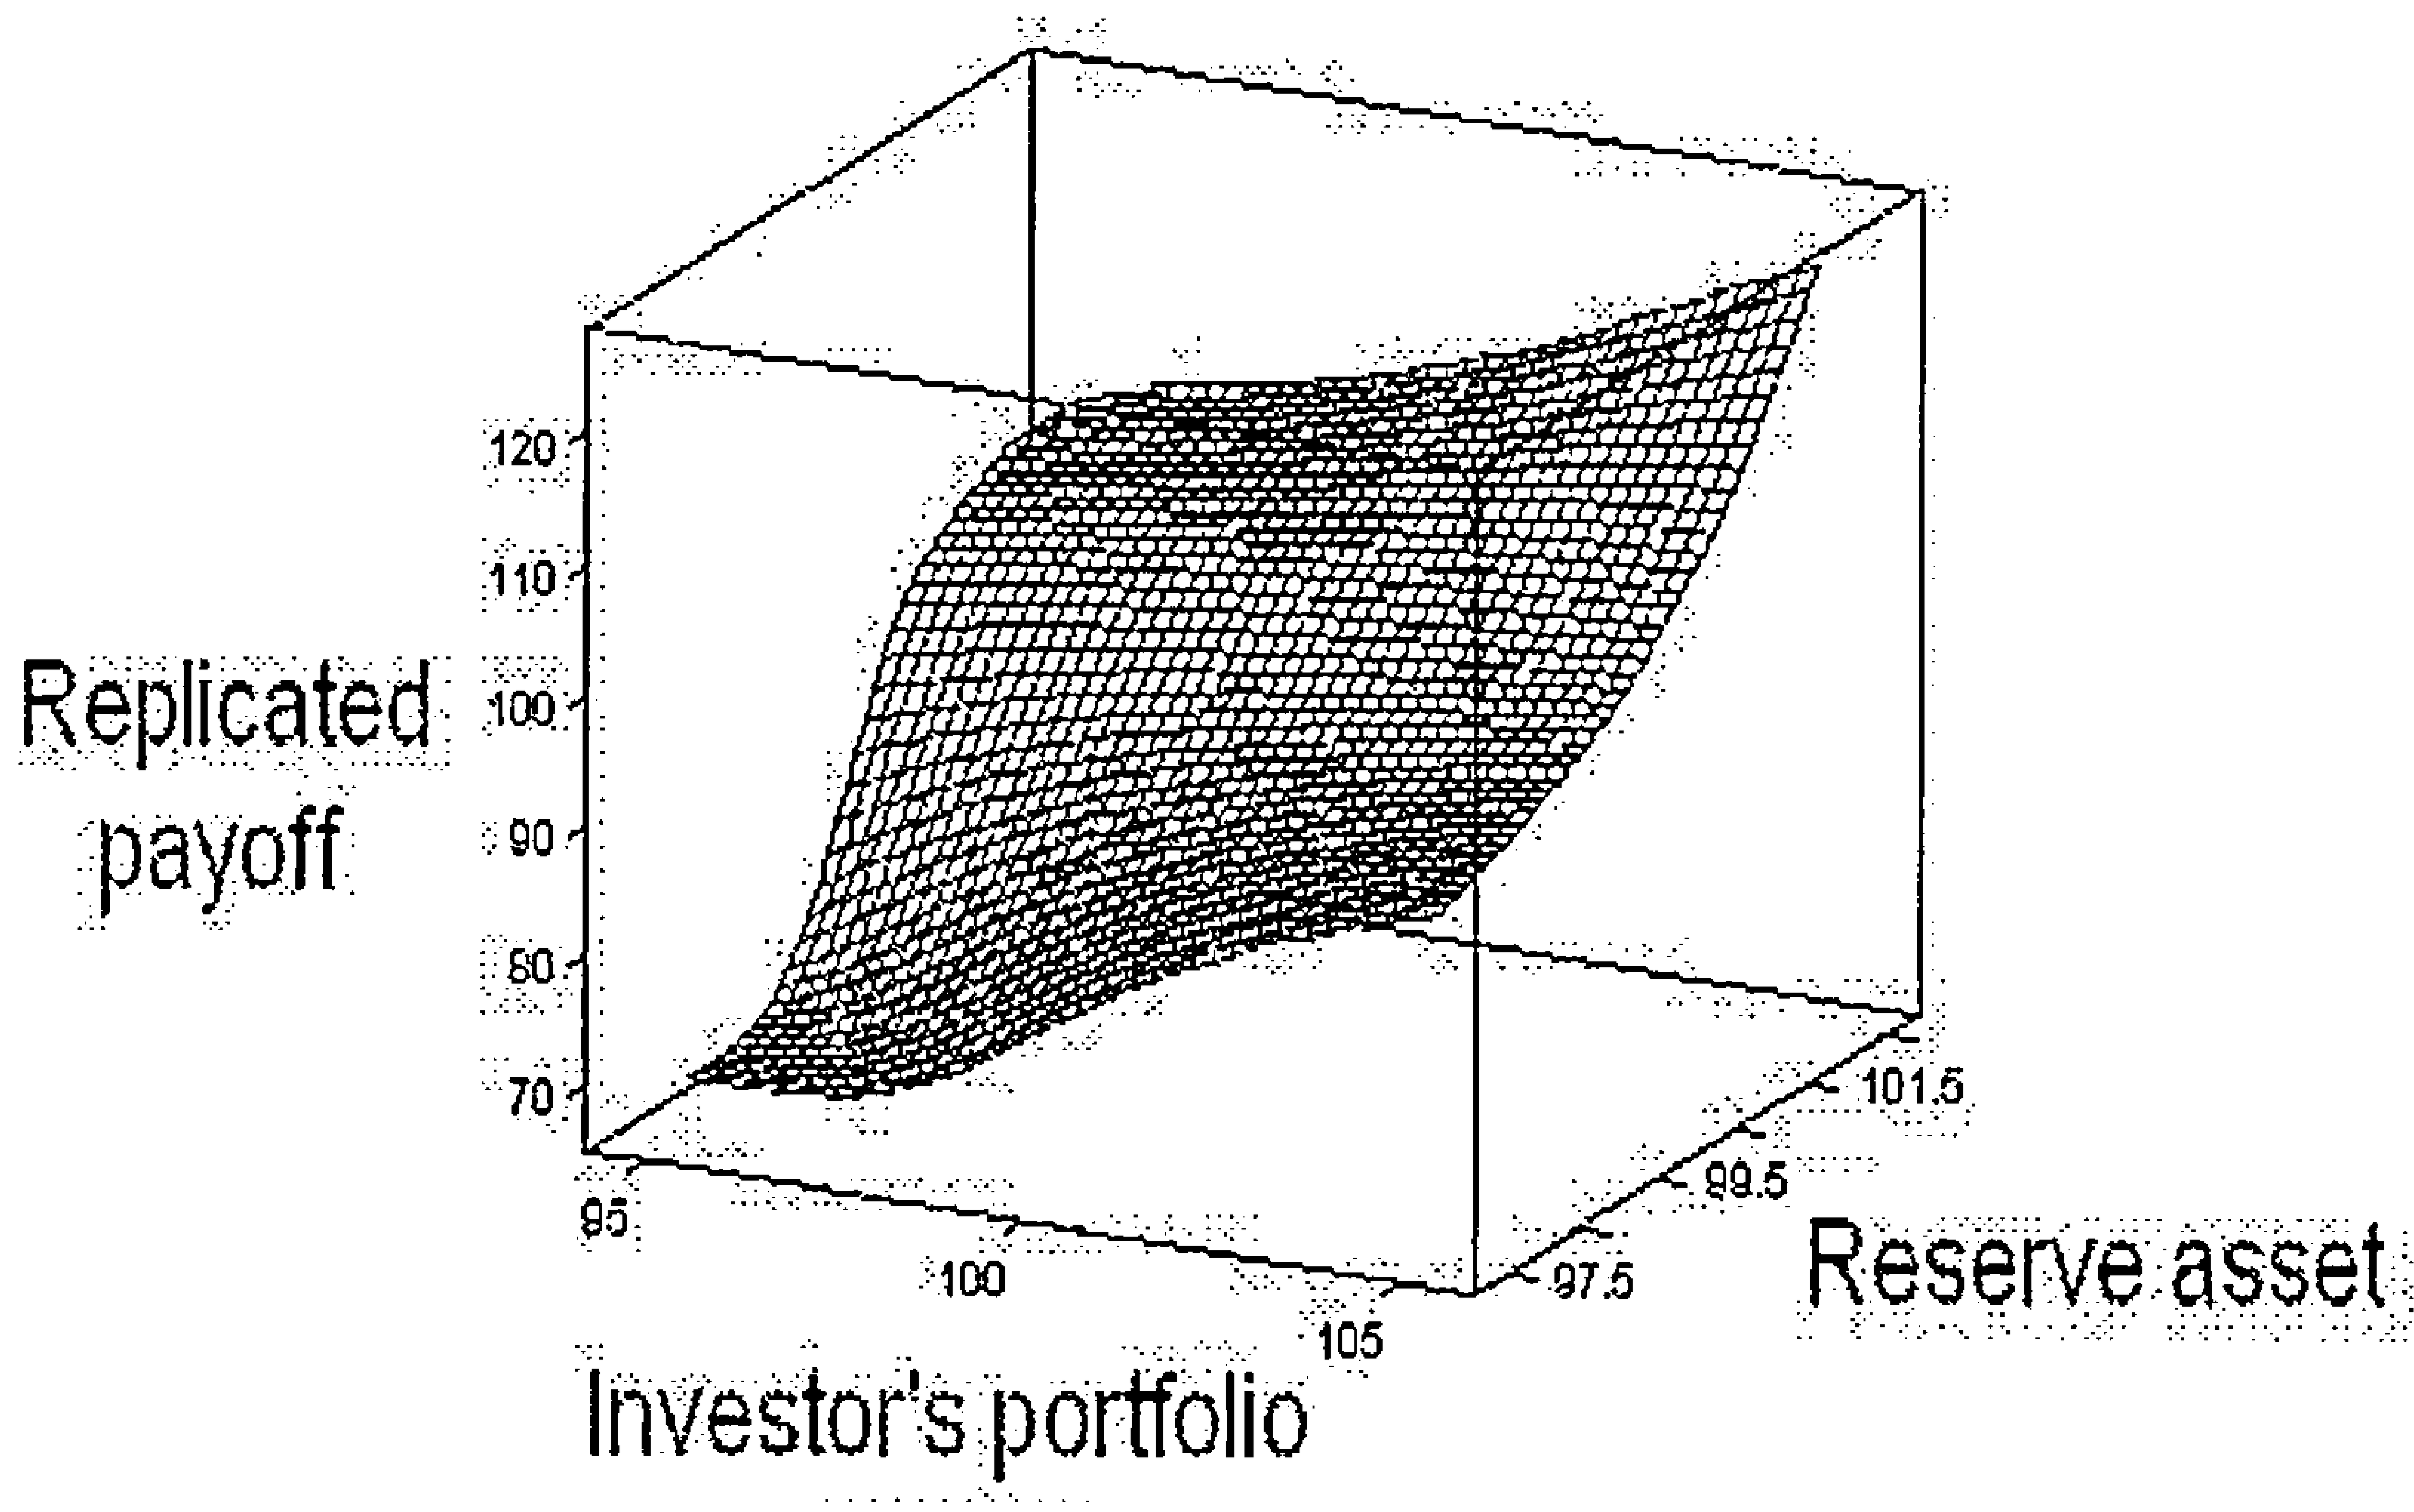

Figure 2.8. 3D plot payoff function Leveraged Capital Holdings.

Figure 2.7 and 2.8 show the payoff function used for the replication of the LCH return per October 2004 (the last month for which we have fund return data available)[26]. Notice that the LCH payoff function is a lot more 'lively' than the payoff functions encountered in the previous section. This underlines the complexity of real-life hedge fund returns. The graphs show that the desired payoff is a positive function of the investor's portfolio as well as the reserve asset, implying that the replication strategy will be long in both assets. We also see quite some variation in the slope of the payoff surface. Since the controls of the replication strategy are nothing more than the slope coefficients of the payoff value function, this signals the presence of 'hot spots', where relatively small changes in the investor's portfolio and/or reserve asset will generate relatively large changes in the strategy's controls.

[26] The jagged profile in the bottom right-hand corner of the contour plot is due to some numerical instability in the extremes. A similar phenomenon is observed in the two cases that follow.

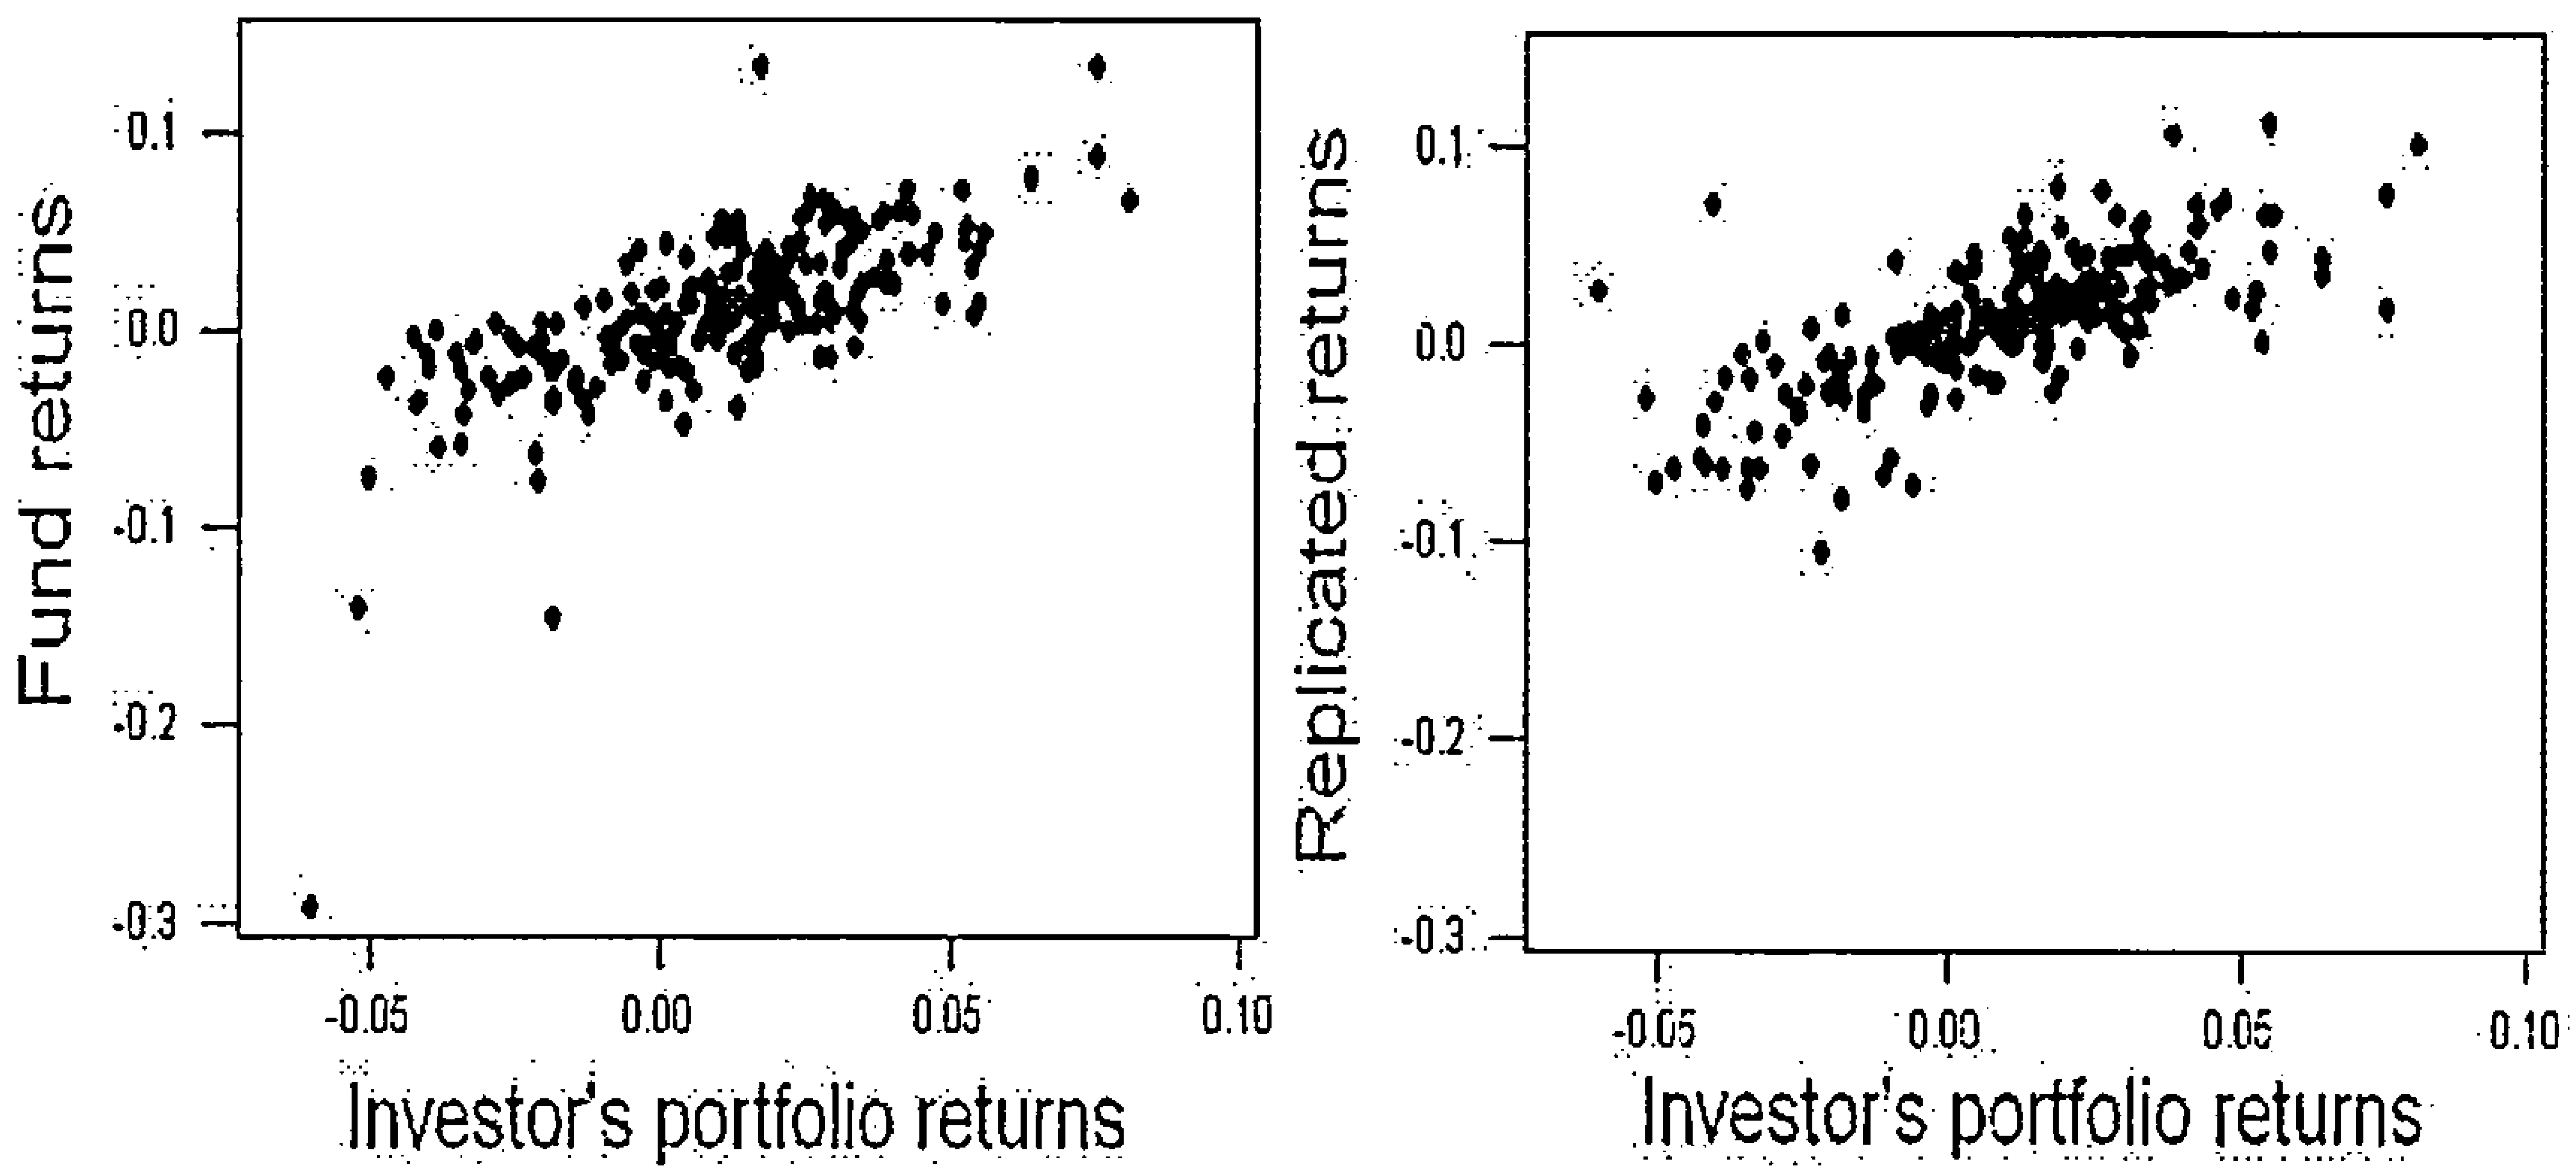

Figure 2.9. Scatter plot investor's portfolio returns versus Leveraged Capital Holdings returns (left) and replicated returns (right), 1987-2004.

The left hand side of Figure 2.9 shows a scatter plot of the monthly returns on the investor's portfolio versus the LCH returns. The right hand side of Figure 2.9 shows a scatter plot of the monthly returns on the investor's portfolio versus the replicated returns. Comparing both plots, we see that they are very similar, which indicates that the replication strategy is indeed able to successfully replicate LCH's returns' statistical properties. We also see that the replication strategy is unable to replicate the three large losses that LCH reported in October 1987 (-22.52%), August 1998 (-11.45%) and April 2000 (-10.83%). Since these are clearly outliers, it is not surprising that the replication procedure was unable to capture them out-of-sample. Given the size of these losses, it is unlikely investors will consider this a real shortcoming though.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt. (robust)	Corr. with Portfolio	Kendall's Tau
LCH	0.0095	0.0419	-1.9675	-0.1641	13.4015	0.3156	0.704	0.536
Replica	0.0125	0.0355	-0.3541	-0.1681	0.7021	0.5736	0.728	0.571
Univariate K-S Statistic = 0.056, (approximated) p-value = 0.884								
Bivariate K-S Statistic = 0.053, (approximated) p-value = 0.968								

Table 2.3. Monthly return statistics Leveraged Capital Holdings and replication strategy, 1987 - 2004.

Another indication of the accuracy of the replication strategy comes from comparing the actual mean, standard deviation, skewness and kurtosis of LCH's returns with those of the replicated returns. The latter statistics can be found in Table 2.3, together with the correlation and Kendall's Tau with the investor's portfolio. Since LCH's returns exhibit some clear outliers, apart from the standard skewness and kurtosis measures we also report more robust skewness and kurtosis measures[27]. 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[28].

Comparing the entries in Table 2.3, it is clear that, despite the obvious limitations, the statistical properties of LCH's returns have been quite successfully replicated. The replication strategy has not only replicated the marginal distribution of LCH's returns but also its relationship with the investor's portfolio. This is also the conclusion from both the K-S tests. Although slightly higher (14.76% pa versus 12.48% pa), the mean of the replicated returns is similar to that of the LCH returns as well. This confirms the assumption underlying the replication procedure that in the longer run investors receive a return which is in line with the risk profile they take on, irrespective of how that risk profile is acquired.

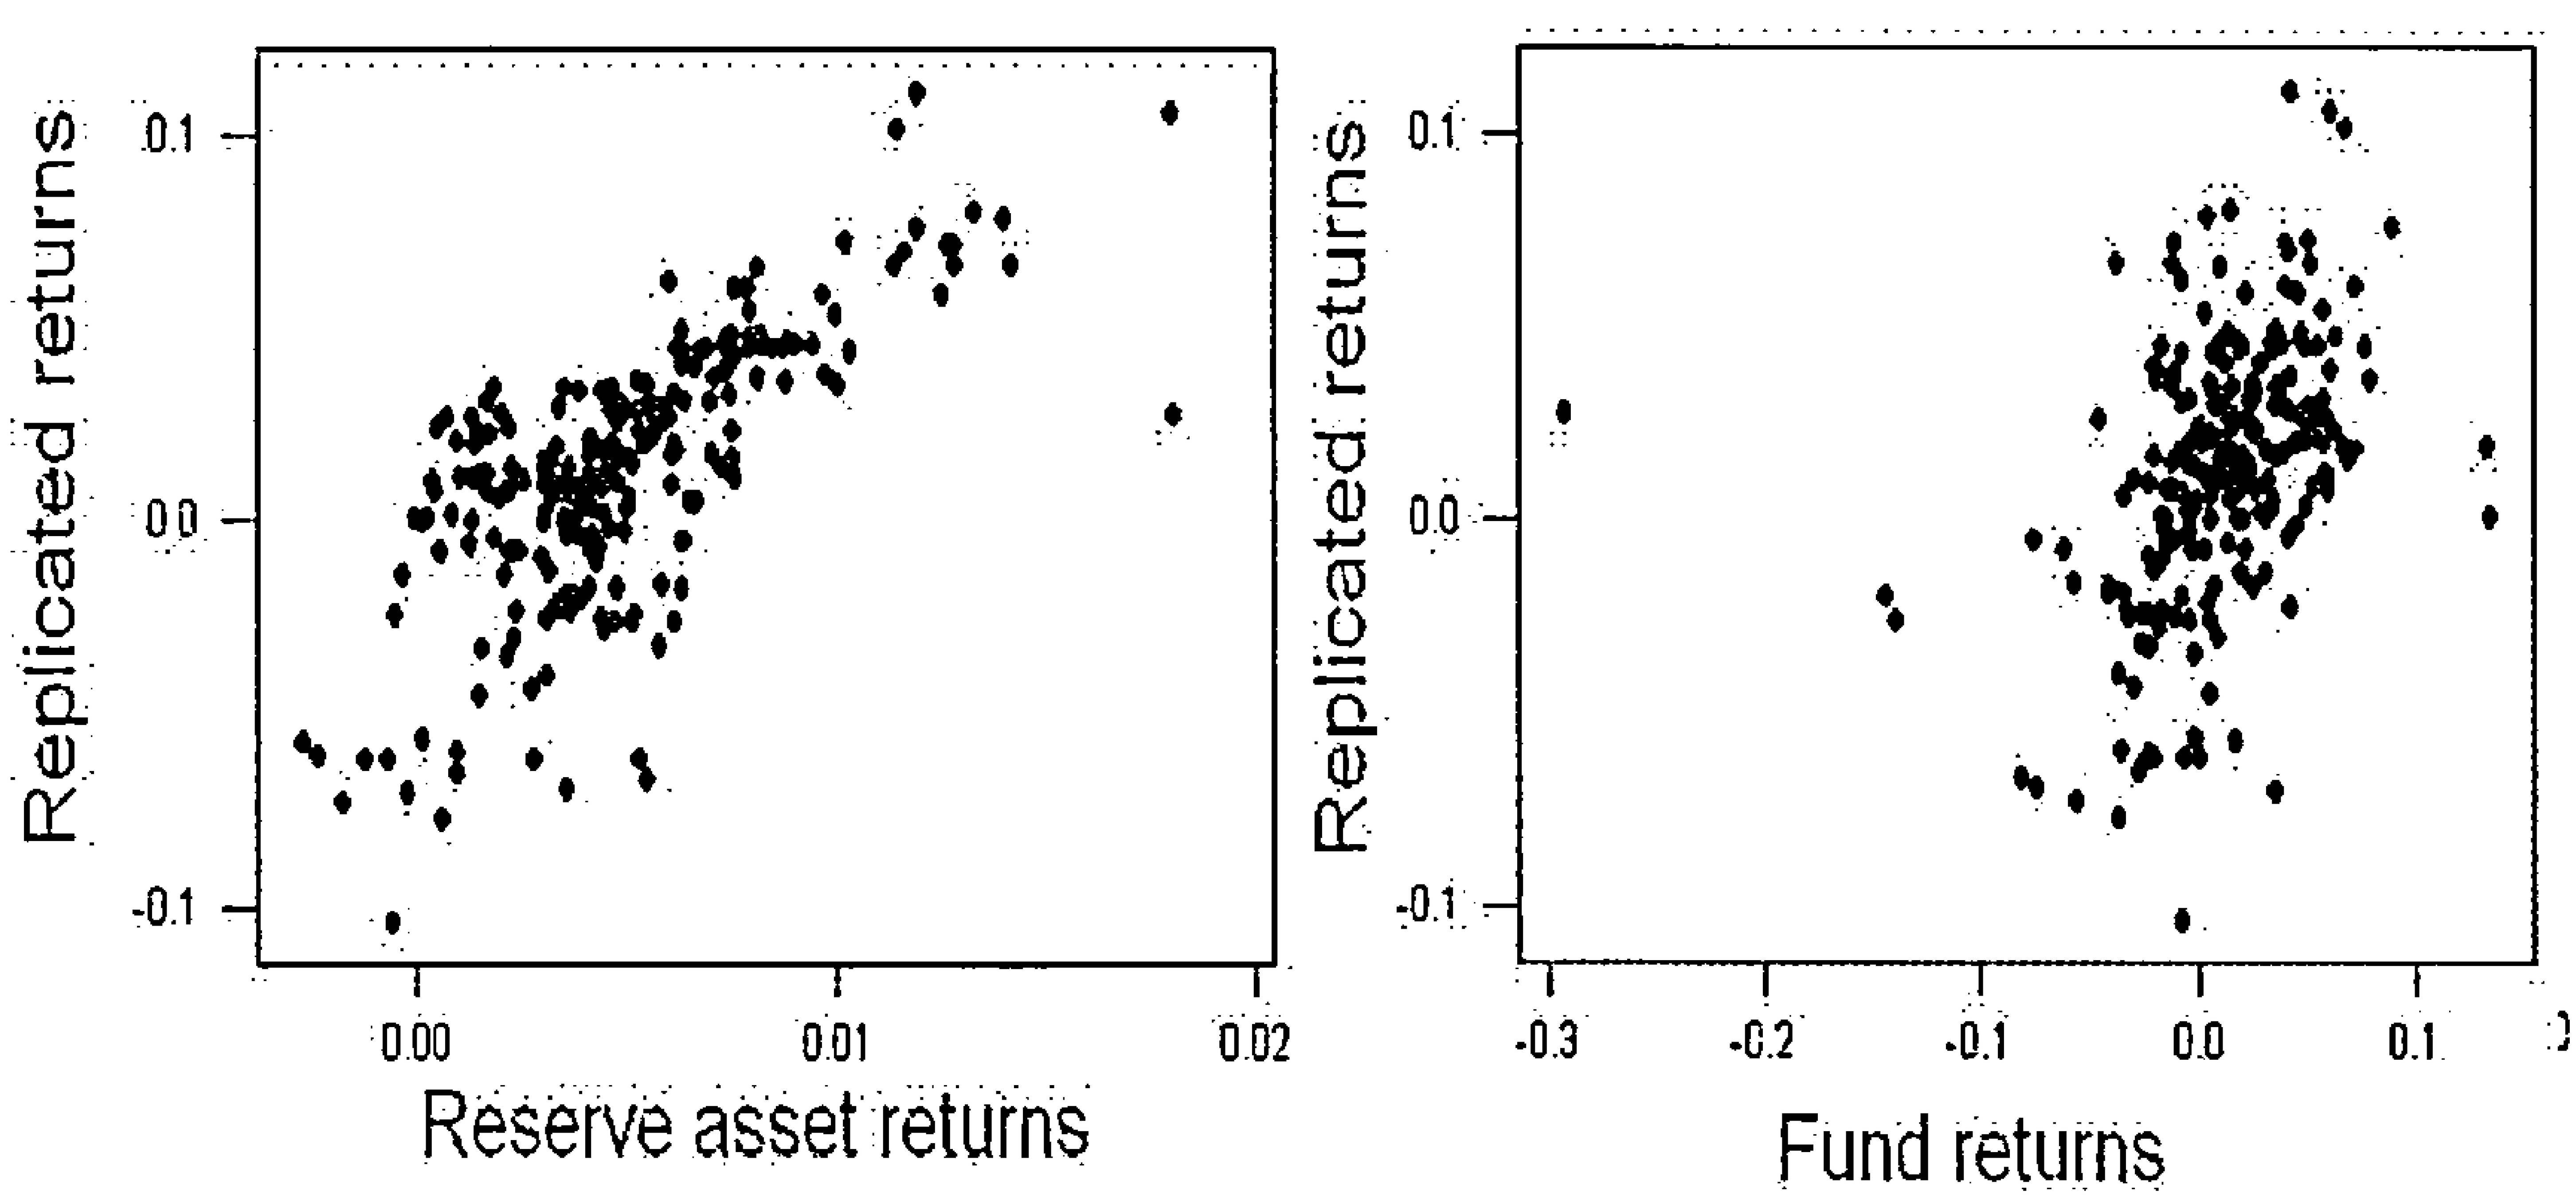

Figure 2.10. Scatter plot reserve asset returns versus replicated returns (left) and Leveraged Capital Holdings returns versus replicated returns (right), 1987 - 2004.

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

[28] 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.

It is interesting to delve a bit further into the workings of the replication strategy. The left hand side of Figure 2.10 shows a scatter plot of the reserve asset returns versus the replicated returns. The positive relationship confirms the efficiency of the replication strategy (see Section 2.3). The right hand side of Figure 2.10 shows a scatter plot of the fund 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 LCH, 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.

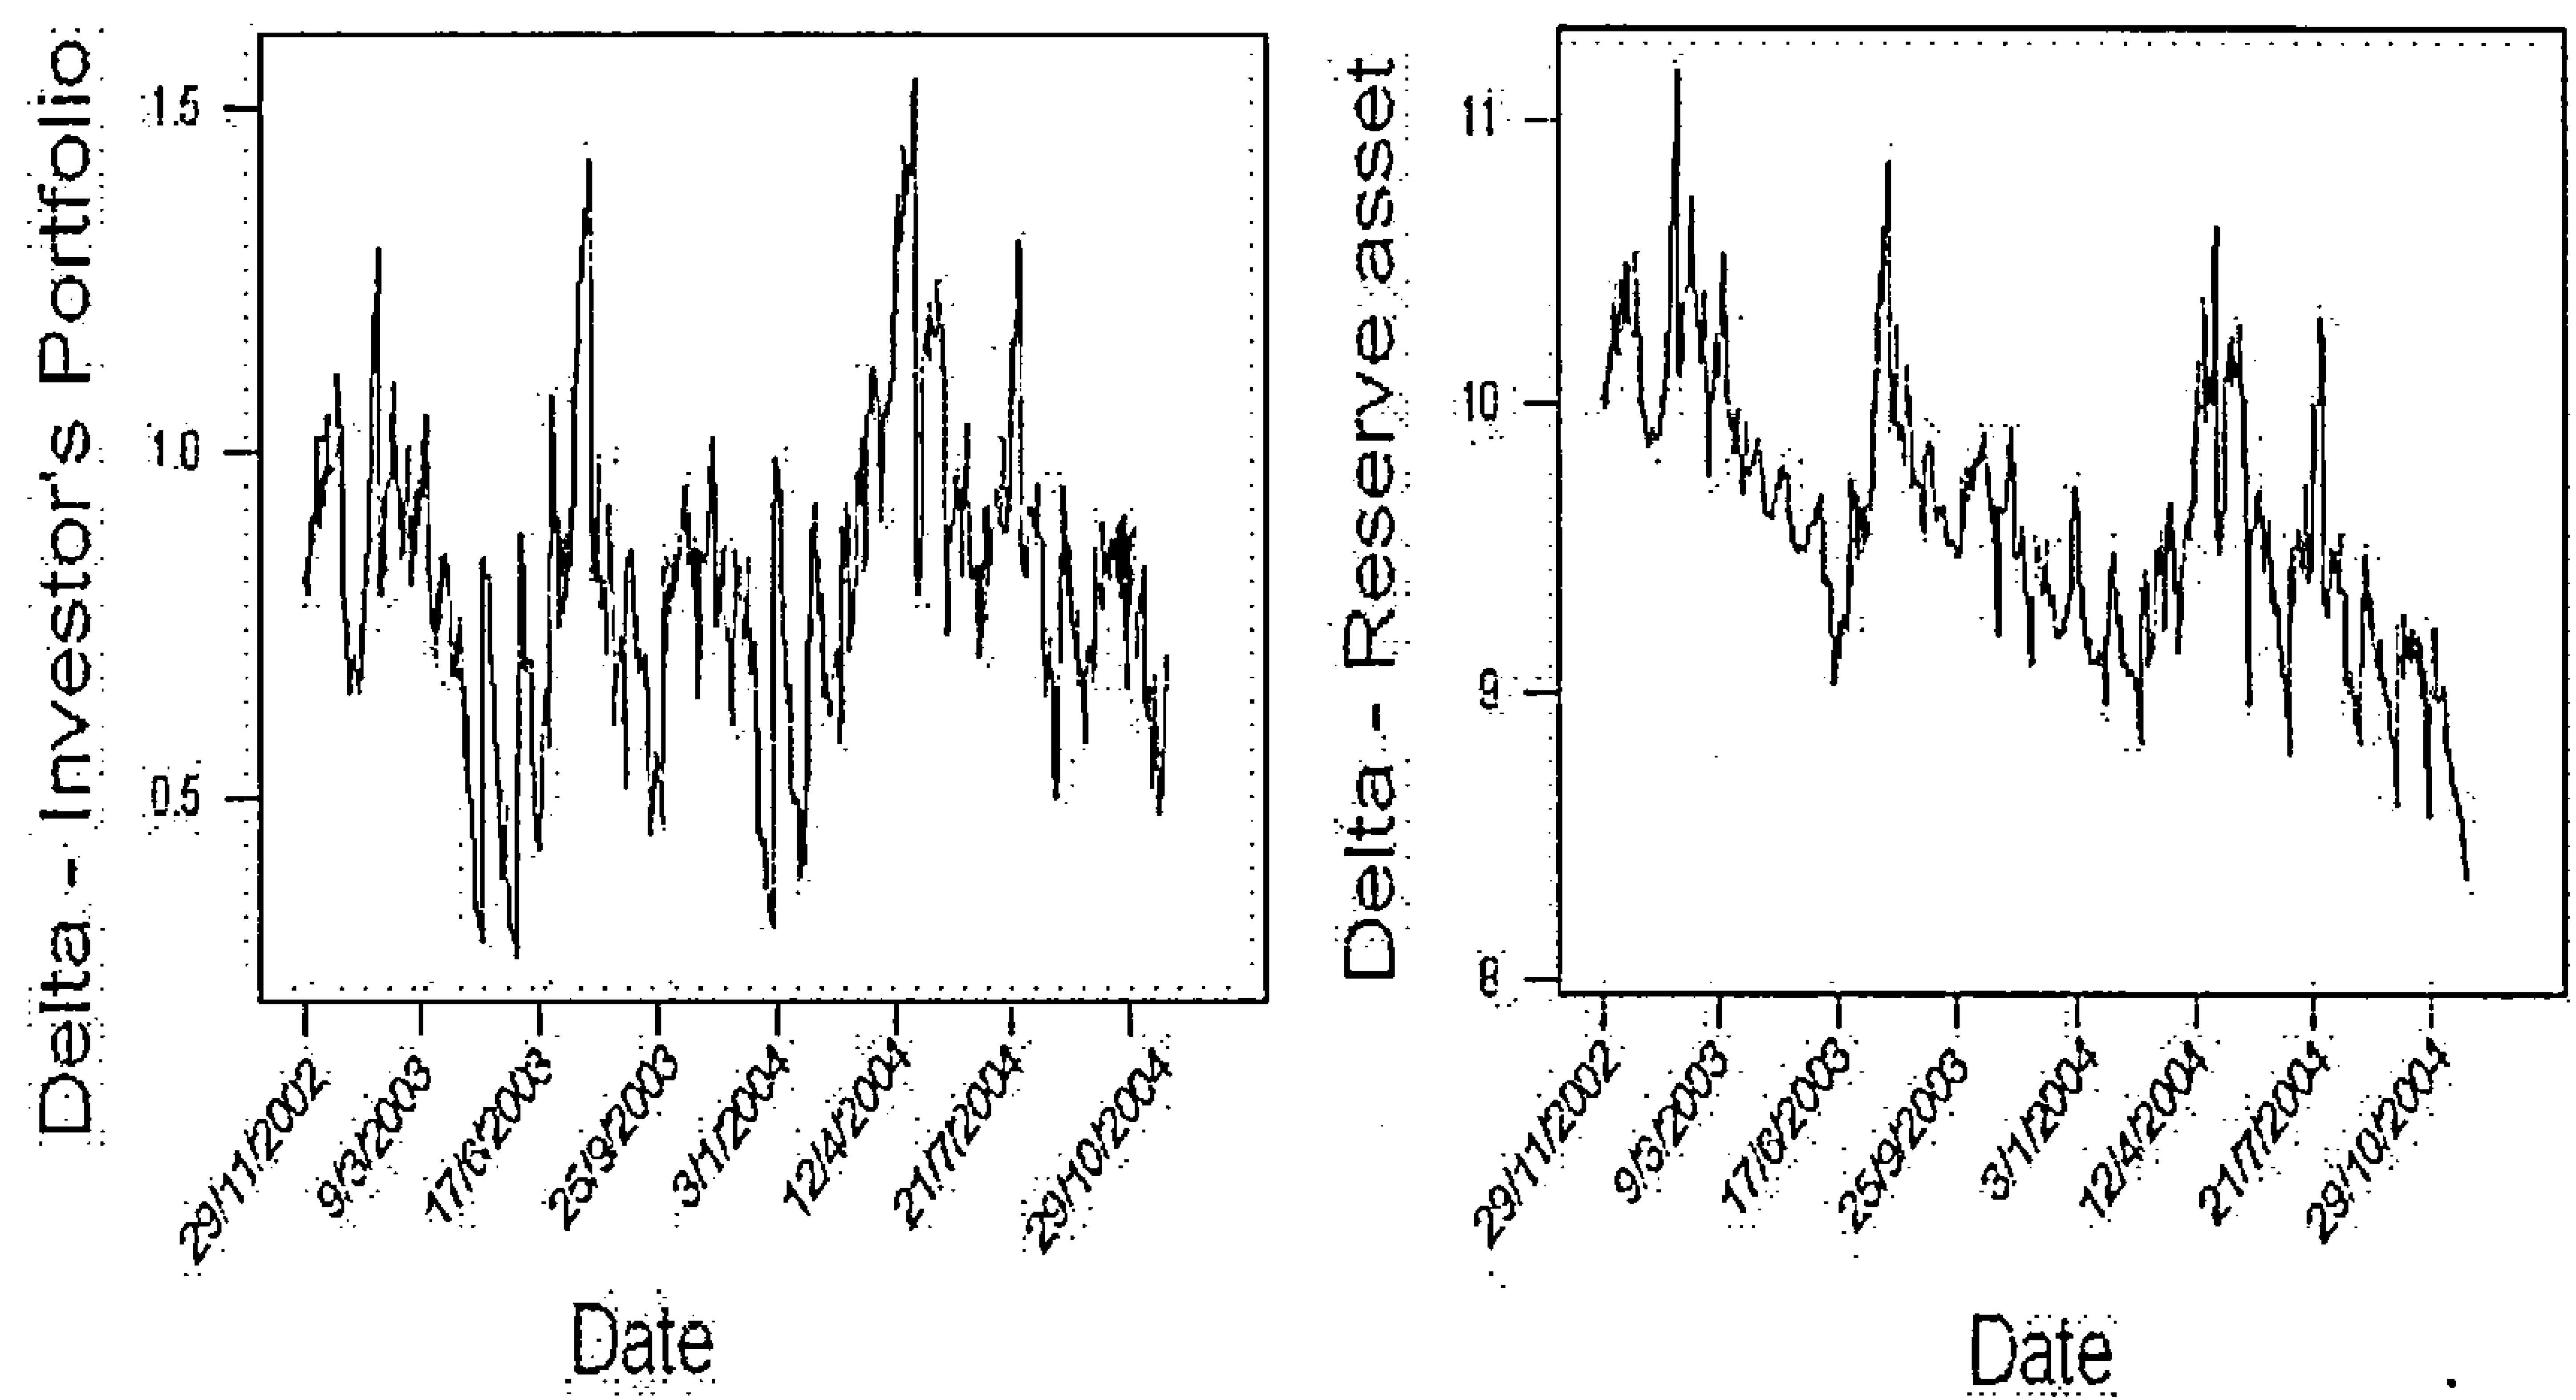

Figure 2.11. Evolution of controls Leveraged Capital Holdings return replication strategy, Dec. 2002 – Oct. 2004.

Figure 2.11 shows the evolution of the replication strategy's controls over the period Dec. 2002 – Oct. 2004[29]. The graph confirms that the replication strategy holds long positions in both the investor's portfolio and the reserve asset. It also shows that the number of units of the reserve asset held is much higher than for the investor's portfolio.

[29] The period Dec. 2002 – Oct. 2004 is representative for the period 1987 – 2004. A graph covering the full 1987-2004 test period would be too condensed to provide any worthwhile insights.

This is because the volatility of the Eurodollar future is quite low compared to that of LCH and the investor's portfolio. It therefore requires substantial 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. For example, during April 2004 the value of the investor's portfolio dropped by almost 4%. As a result, the number of units of the investor's portfolio to hold rose from 0.90 at the start to 1.54 at the end of the month. At the same time, the number of units of the reserve asset to hold rose from 9.34 to 10.63. At the beginning of May, however, the strategy was reset to its starting values, meaning that the allocation to the investor's portfolio dropped to 0.85 units and the allocation to the reserve asset to 9.58 units.

2.5.2 Calamos Multi-Strategy Fund L.P.

The second example is a convertible arbitrage fund. The Calamos Multi-Strategy Fund (CMSF) was established in 1989 by convertible bond experts John and Nick Calamos. For most of its life CMSF has pursued a convertible arbitrage strategy. Since 2004, however, CMSF has adopted a long/short equity strategy as well. Managed primarily for the personal accounts of the Calamos family and a small group of friends, the fund is relatively small with $14.1 million under management (TASS, October 2005)[30].

[30] Although the fund is only small, we decided to include it because the Calamos family is very well known for their work on convertibles and convertible arbitrage. See for example Calamos (1998) and Calamos (2003).

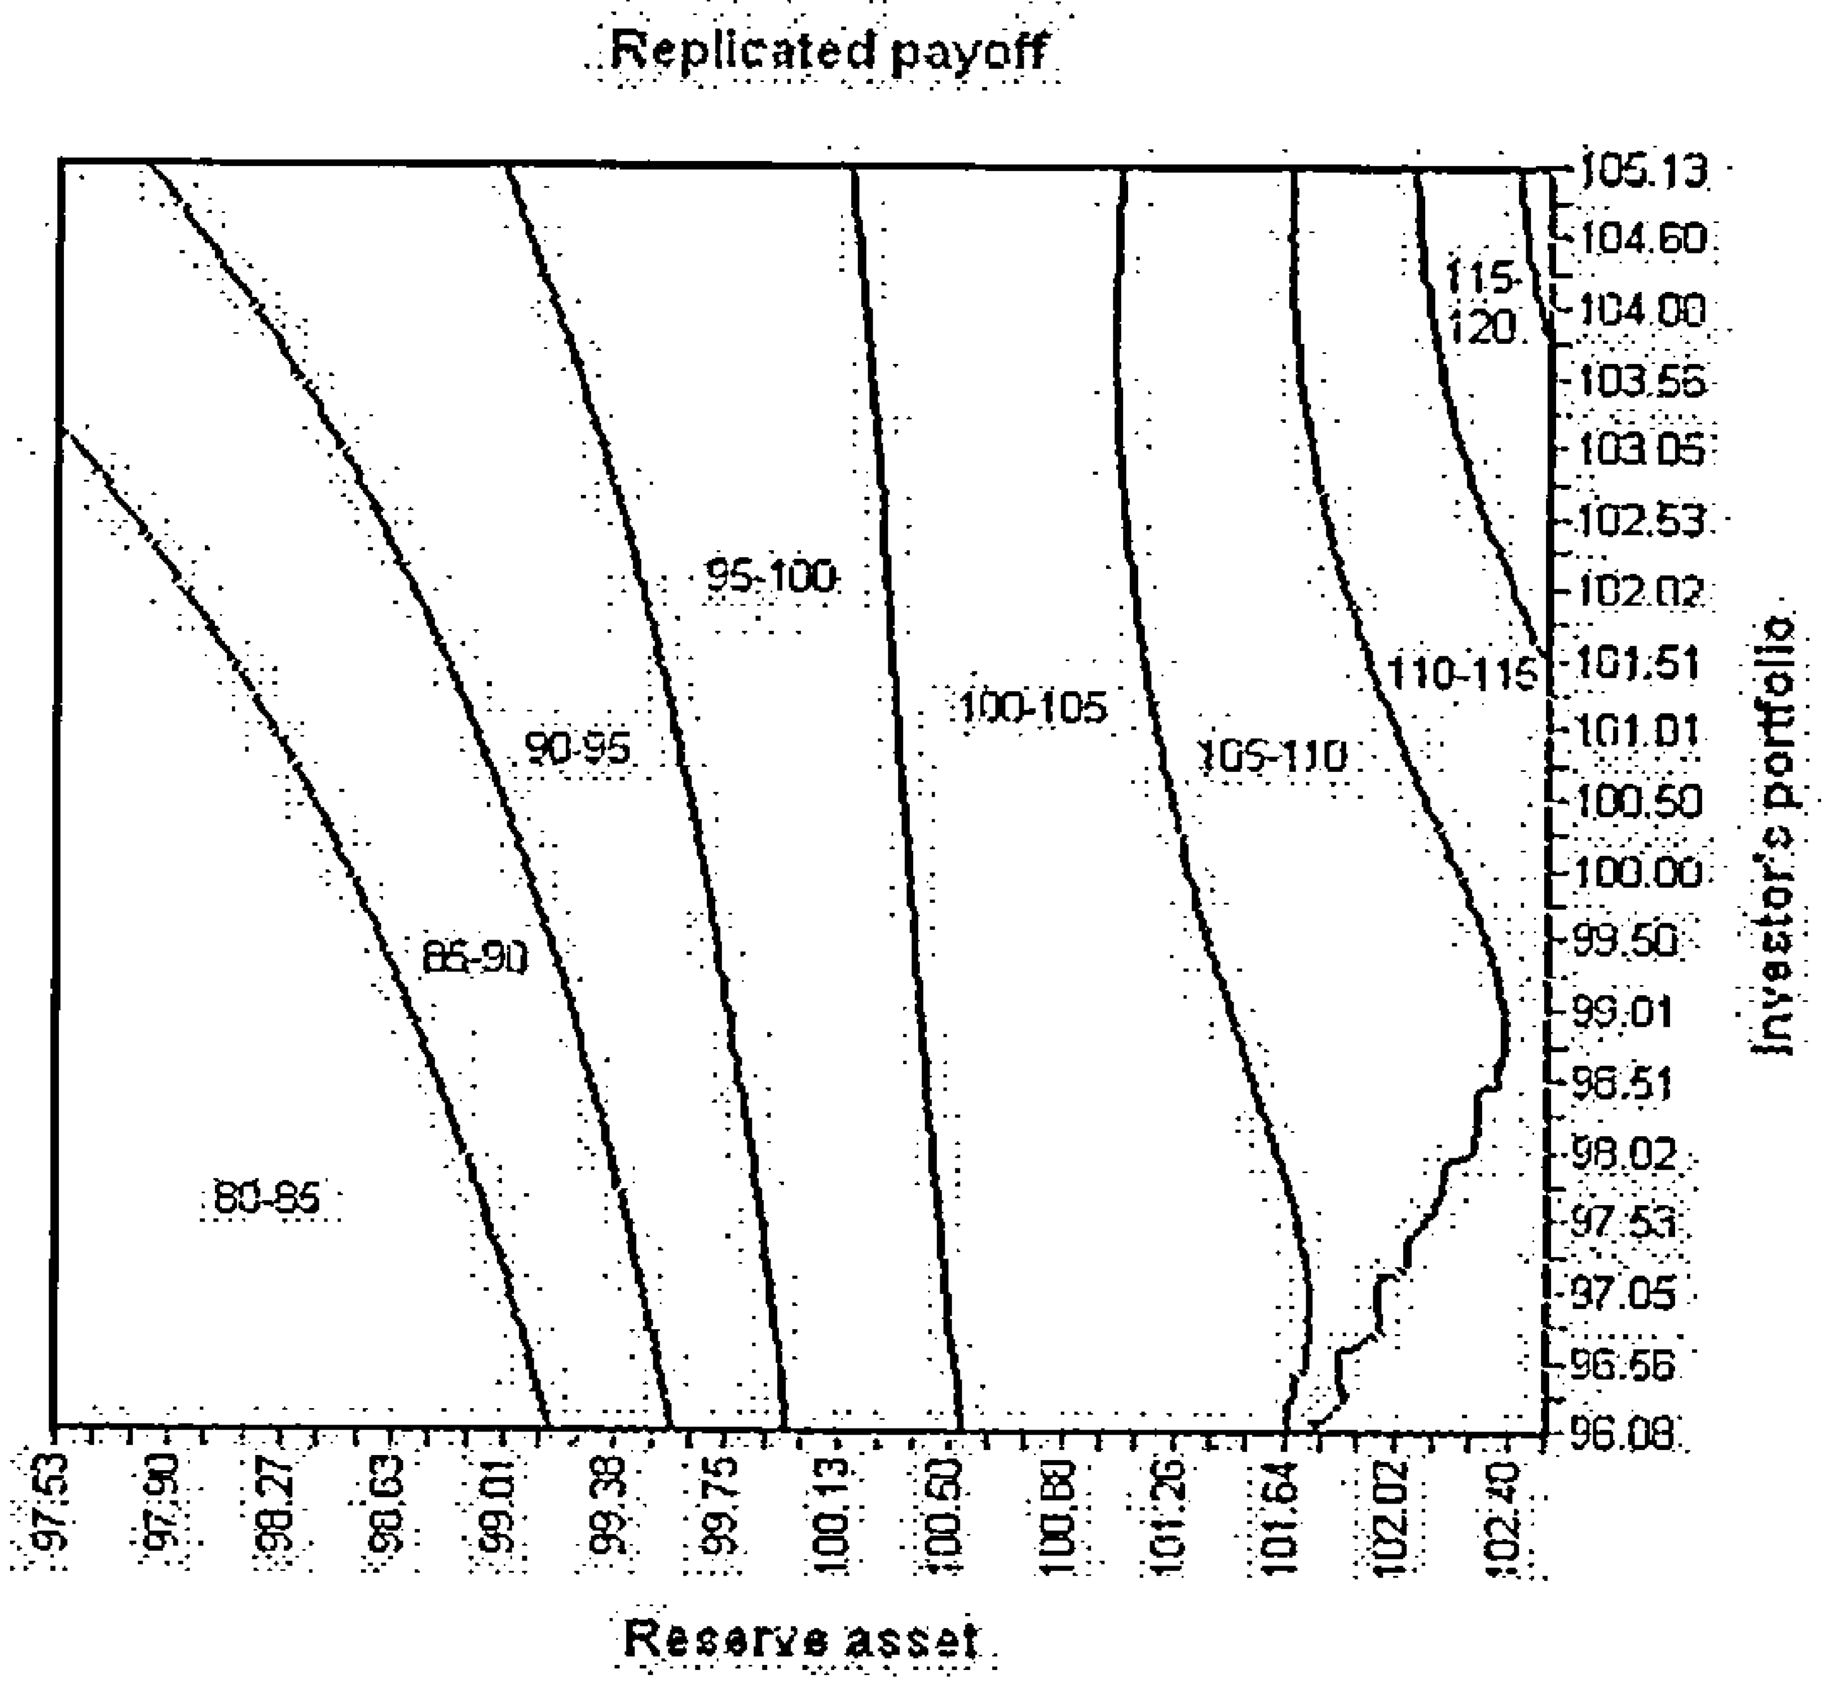

Figure 2.12. Contour plot payoff function Calamos Multi-Strategy Fund.

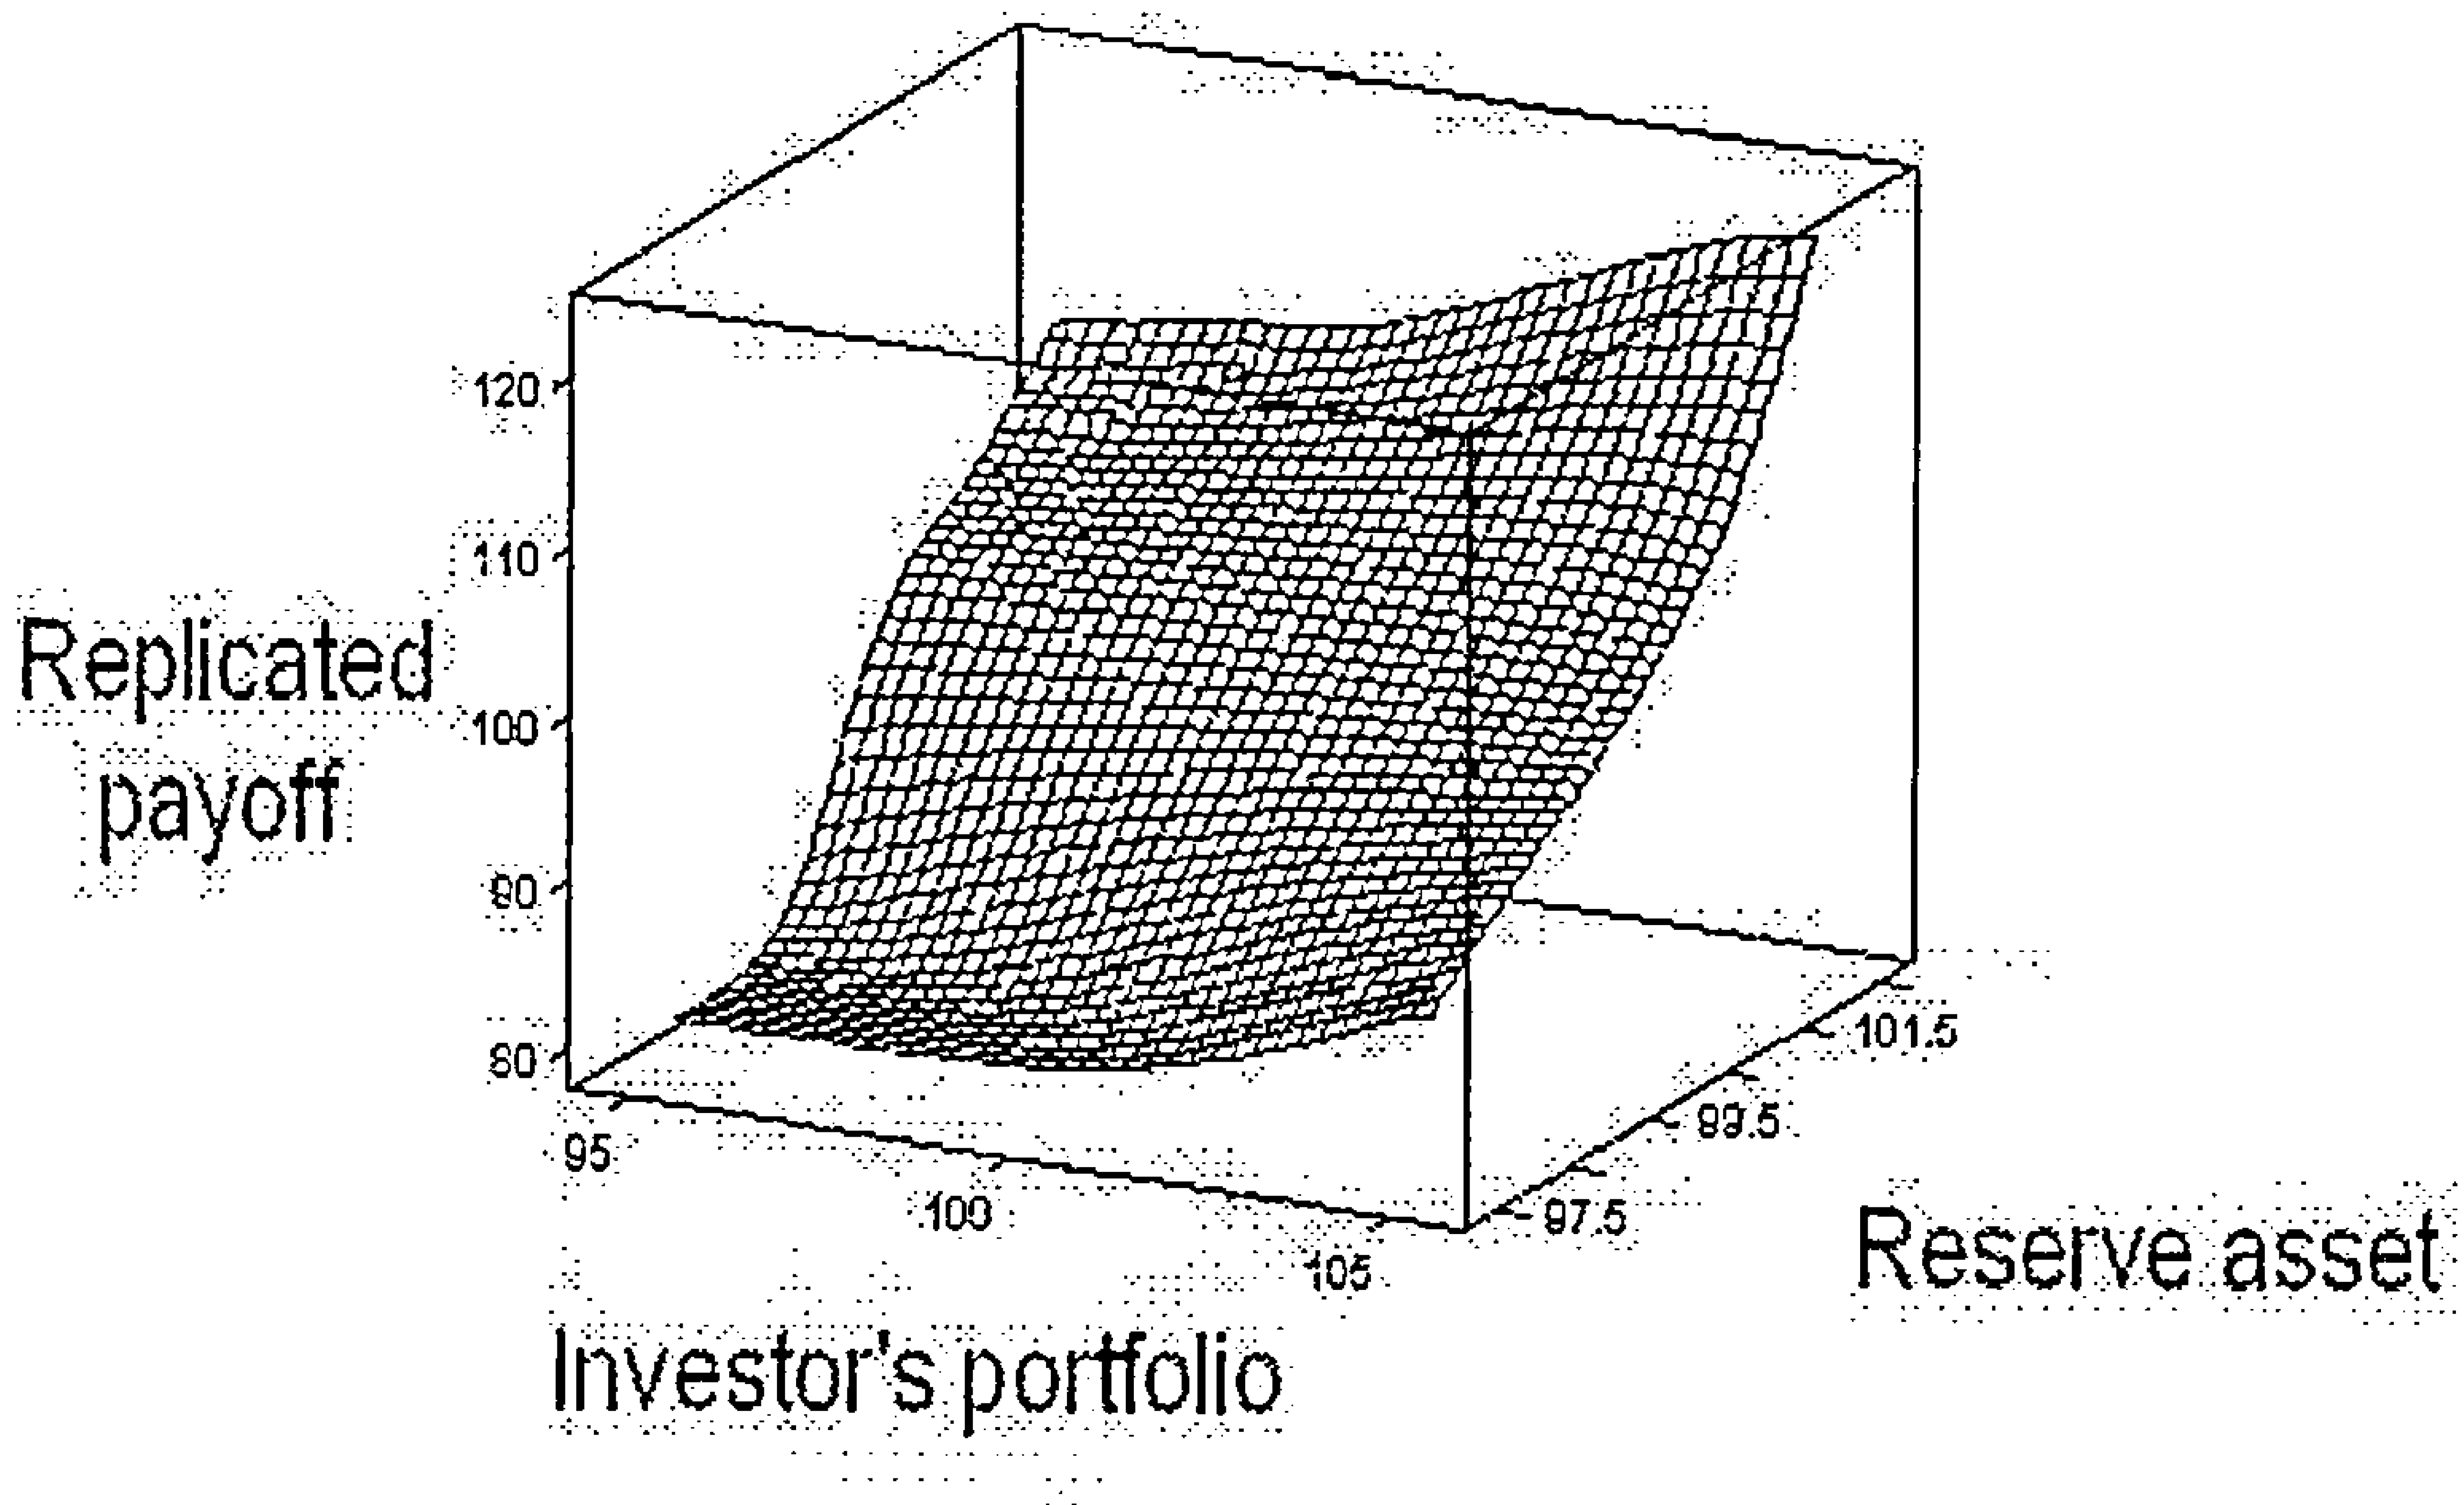

Figure 2.13. 3D plot payoff function Calamos Multi-Strategy Fund.

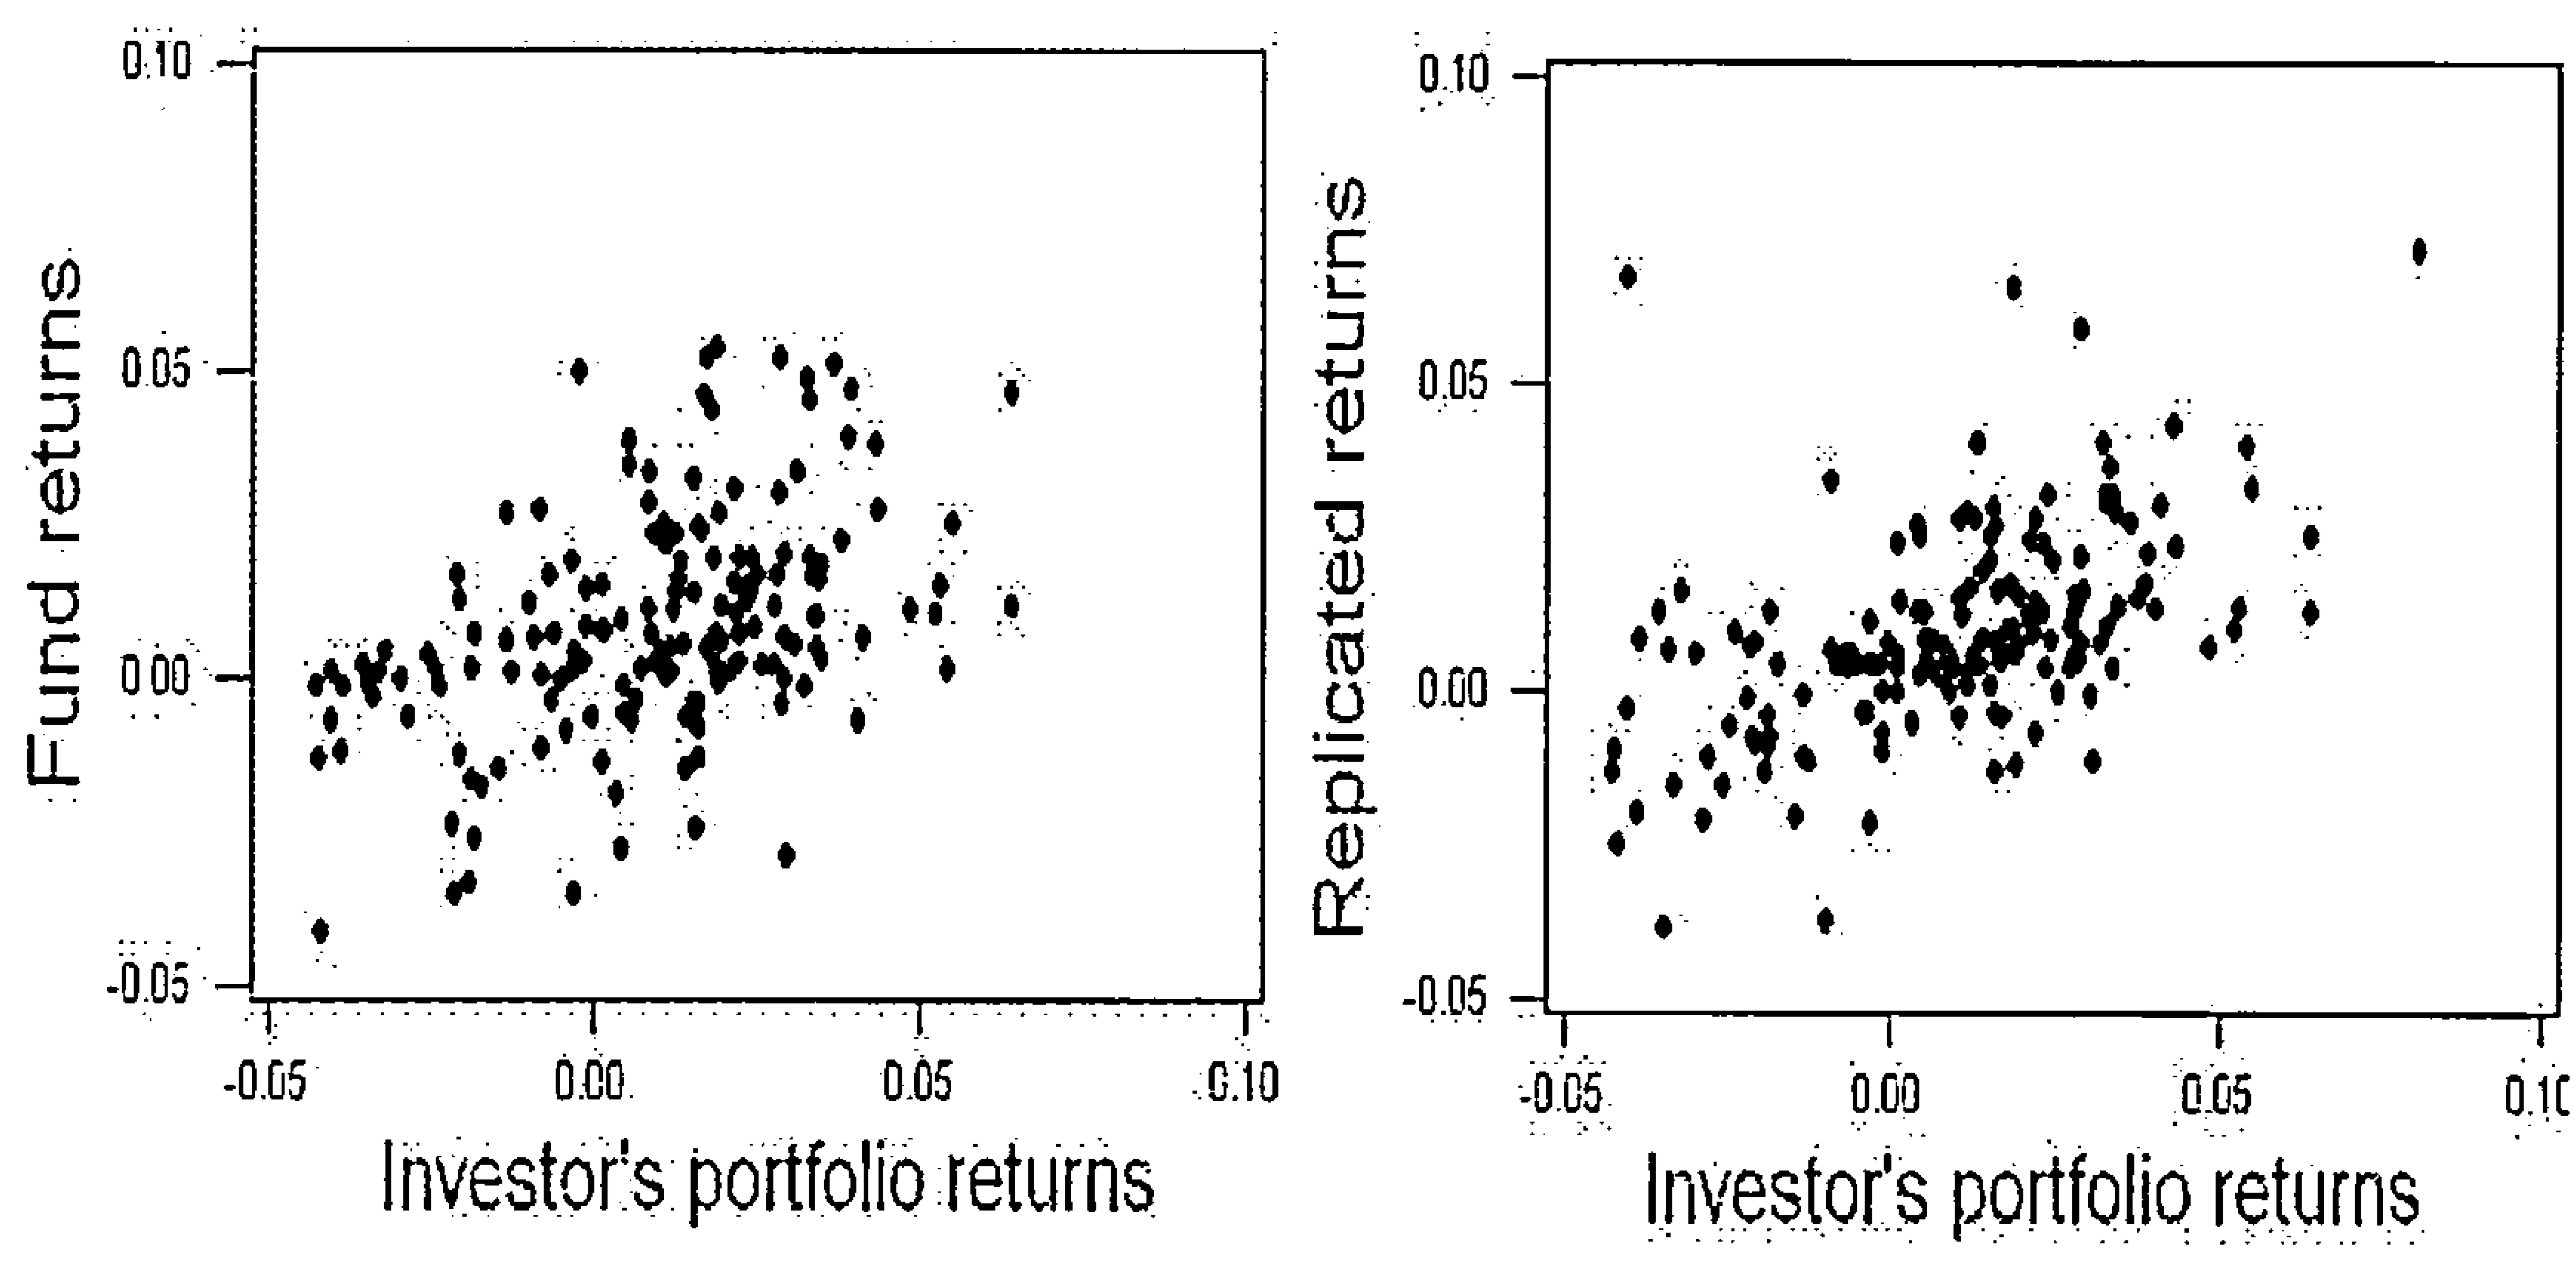

Figure 2.14. Scatter plot investor's portfolio returns versus Calamos Multi-Strategy Fund returns (left) and replicated returns (right), 1991 - 2004.

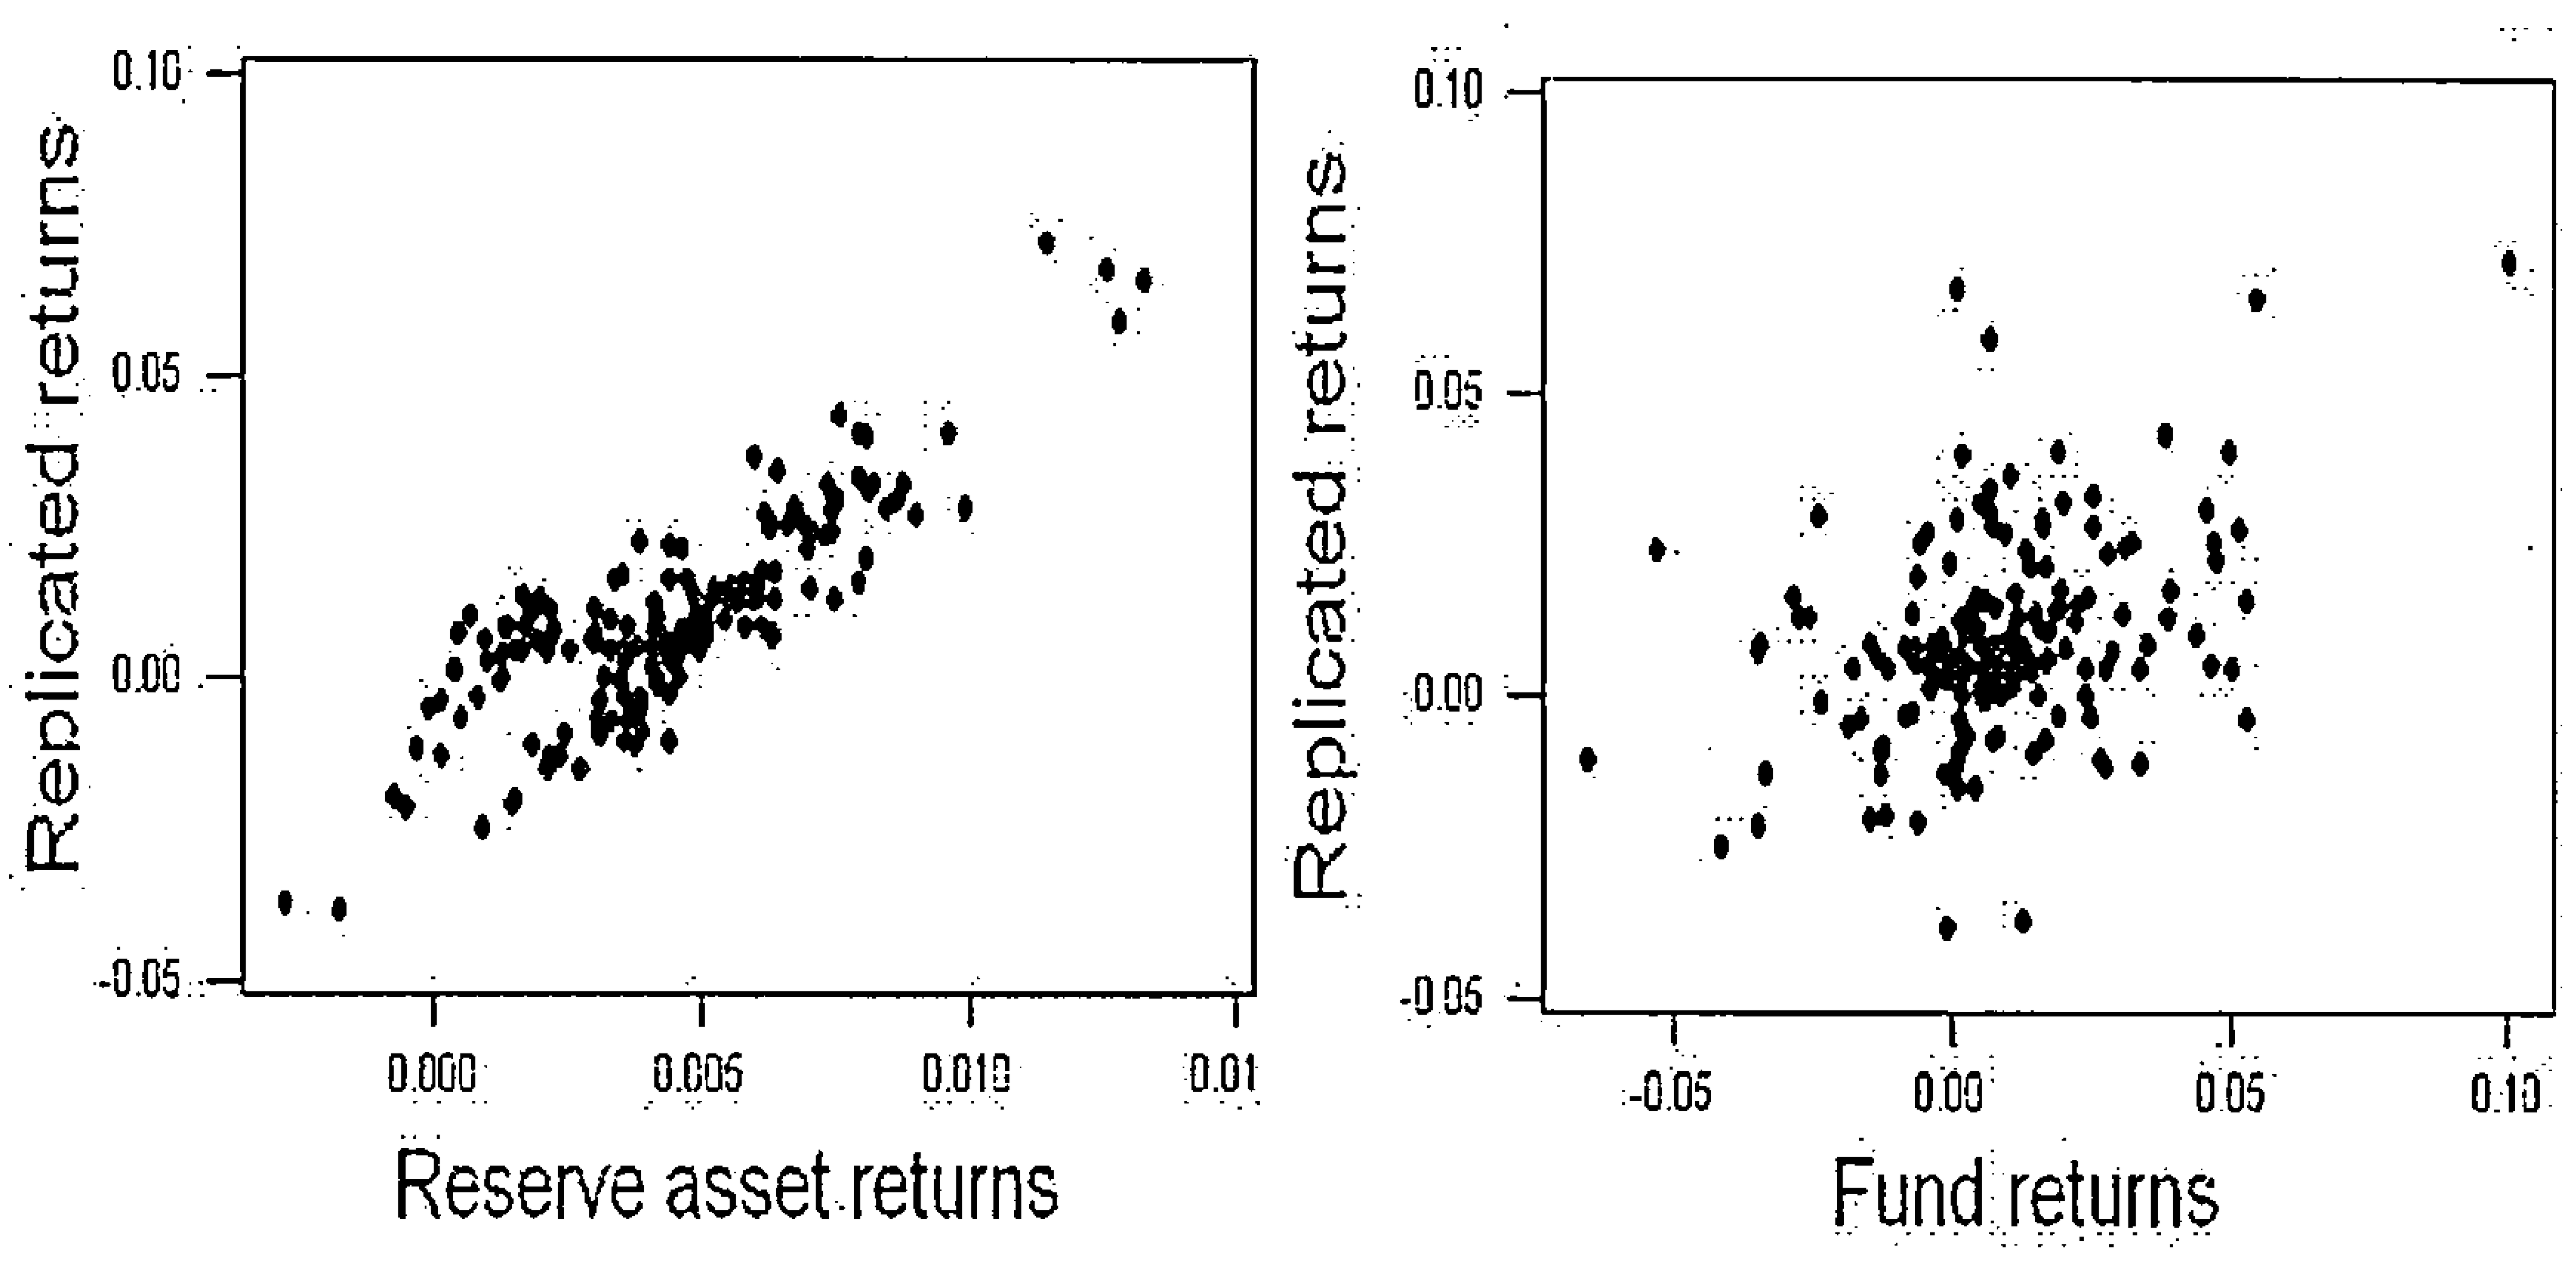

Figure 2.15. Scatter plot reserve asset returns versus replicated returns (left) and Calamos Multi-Strategy Fund returns versus replicated returns (right), 1991 - 2004.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt (robust)	Corr. with Portfolio	Kendall's Tau
CMSF	0.080	0.0213	0.2357	0.0154	2.6296	1.6937	0.509	0.337
Replica	0.094	0.0170	0.6656	0.0582	2.2128	0.9525	0.506	0.388
Univariate K-S Statistic = 0.103, (approximated) p-value = 0.322								
Bivariate K-S Statistic = 0.087, (approximated) p-value = 0.719								

Table 2.4. Monthly return statistics Calamos Multi-Strategy Fund and replication strategy, 1991 - 2004.

The desired payoff function for CMSF as per October 1^{st}, 2004 is shown in Fig 2.12 and 2.13. At first sight, it looks similar to that for LCH, but, as is easiest seen from the contour plot, there are some significant differences as well. Figure 2.14 shows the same scatter plots as in Figure 2.9. Comparing both plots, we again see that they are very similar, indicating the replication strategy performs quite well. This is confirmed by the entries in Table 2.4. As before, all parameters are very similar, including the means and the correlation with the investor's portfolio. Both the univariate and bivariate K-S test confirm that there is no significant difference between the original and replicated distributions. Figure 2.15 shows the same scatter plots as in Figure 2.10. We again see a positive relationship between the reserve asset returns and the replicated returns, confirming the efficiency of the replication strategy. The plot of the fund returns versus the replicated returns shows a random scatter, making it clear that although the replicated returns have similar statistical properties as CMSF, they come in a completely different order.

2.5.3 Rocker Partners L.P.

Most hedge funds' returns are positively correlated with the equity market. Our final example therefore concerns a dedicated short seller, the returns of which are likely to be negatively correlated with the stock market. Rocker Partners (RP) was started in 1985 by David Rocker. While RP maintains both long and short positions, the general focus is on short selling. The fund is therefore popular with investors as a hedge against their long biased investments. RP currently has $611.1 million under management (TASS, October 2005).

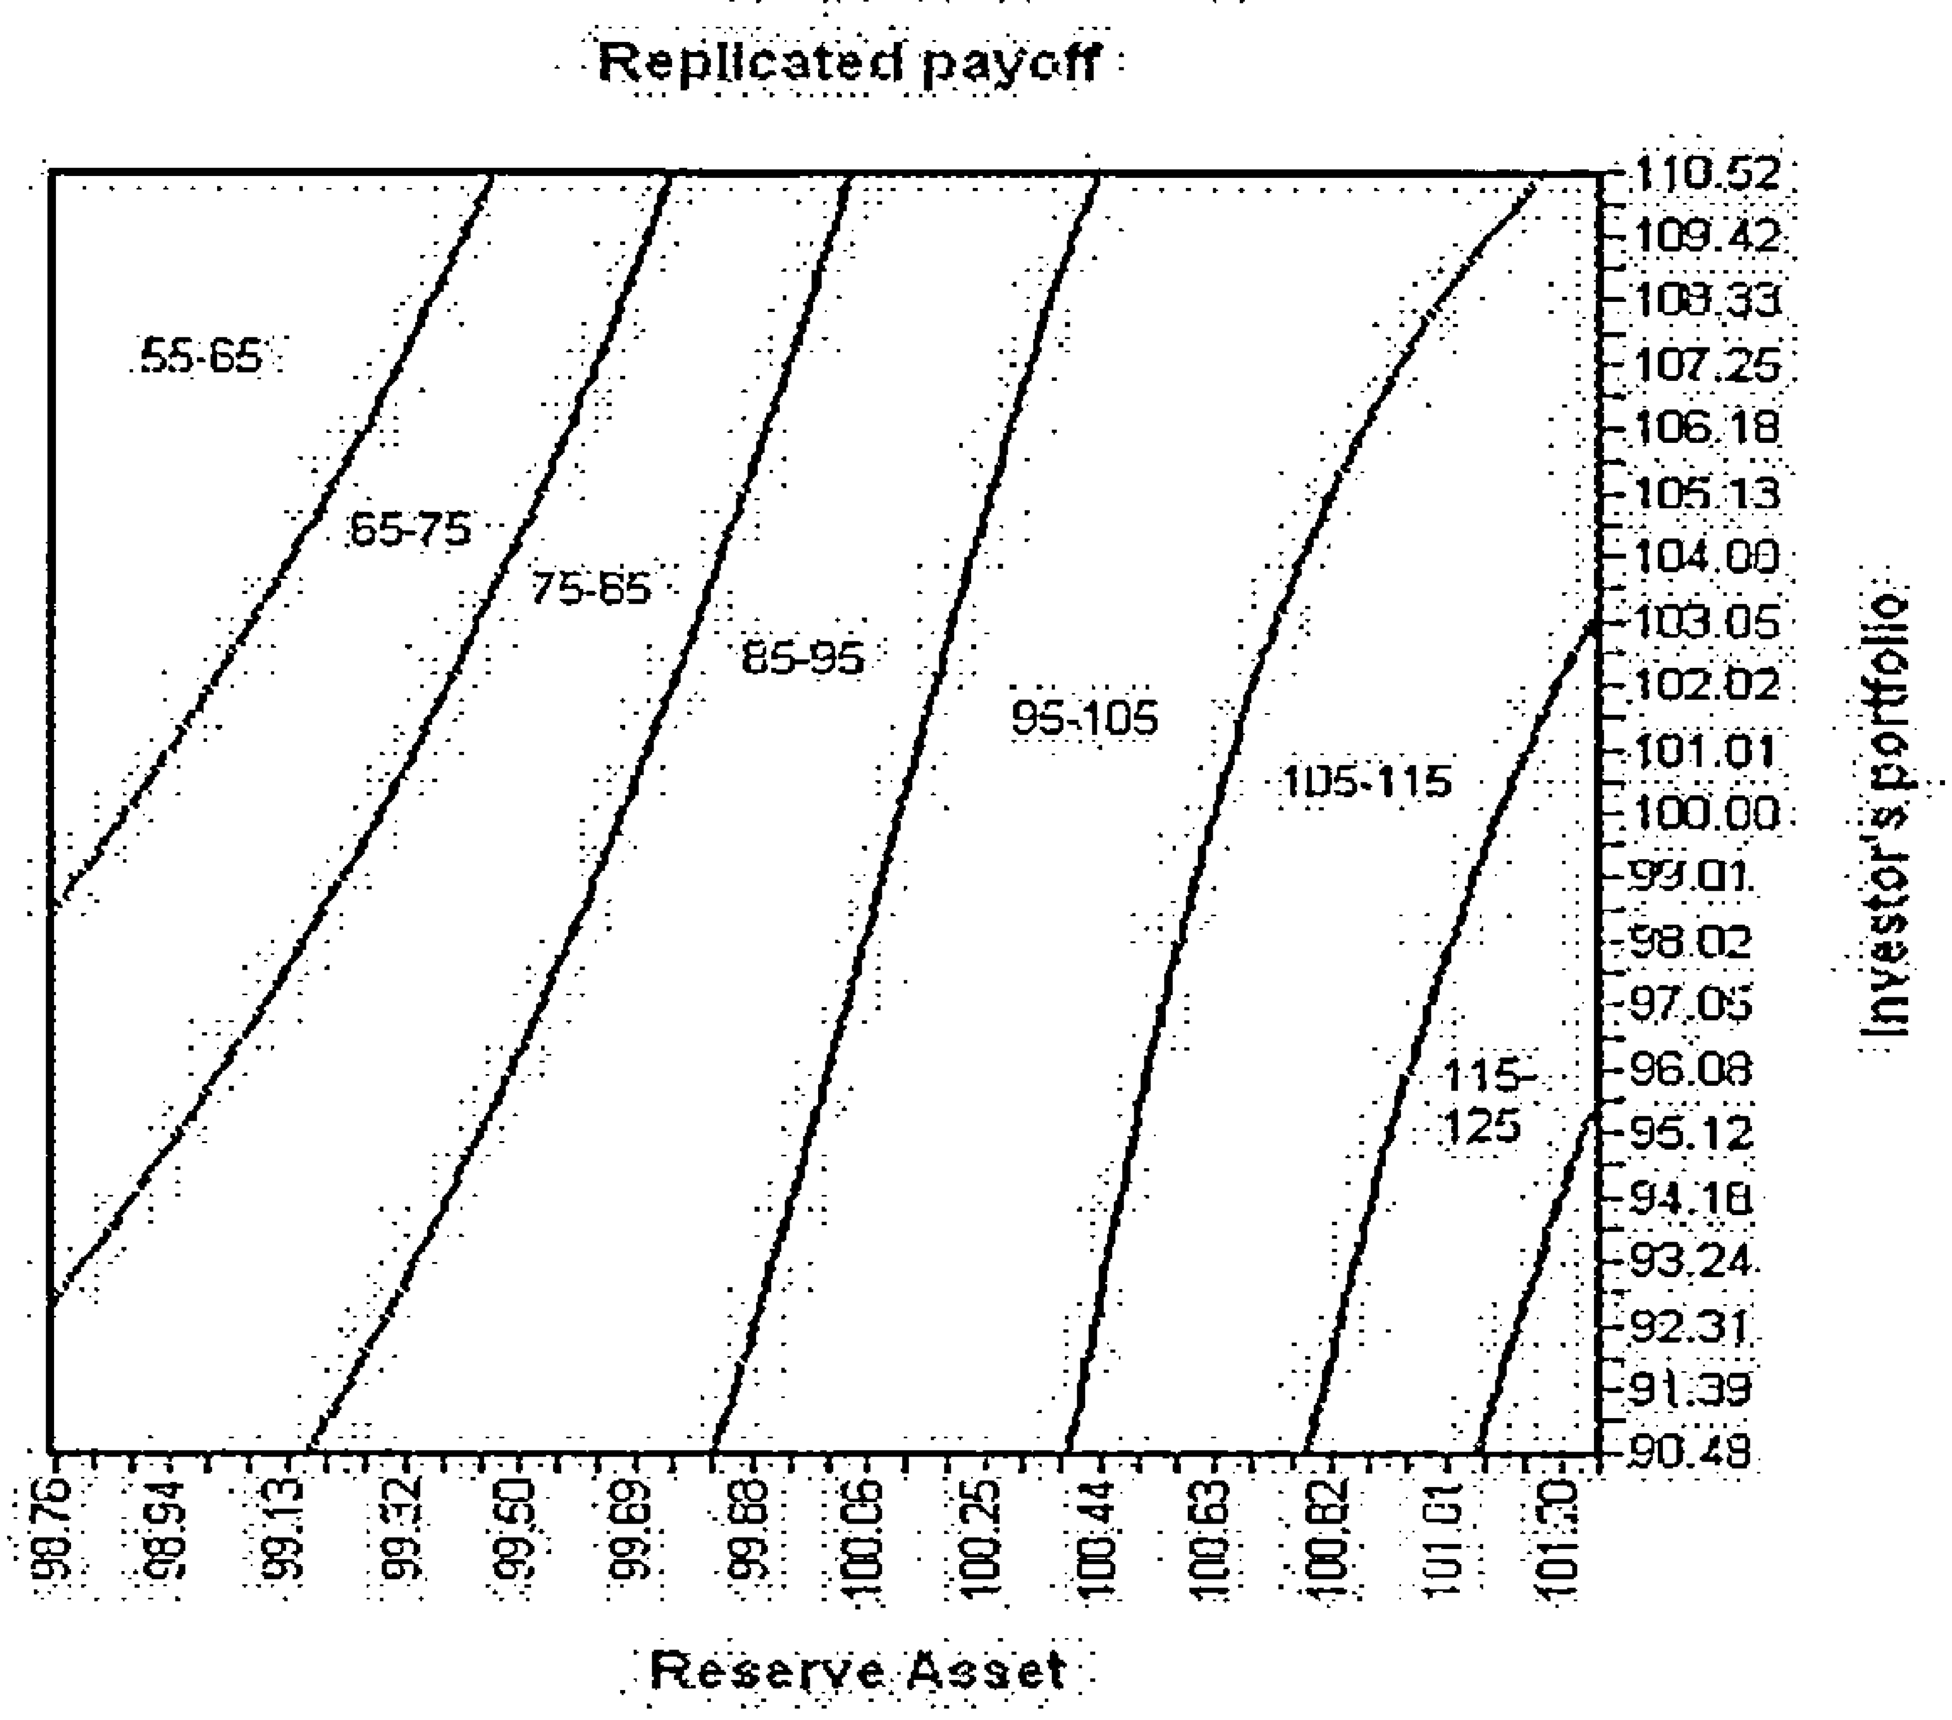

Figure 2.16. Contour plot payoff function Rocker Partners.

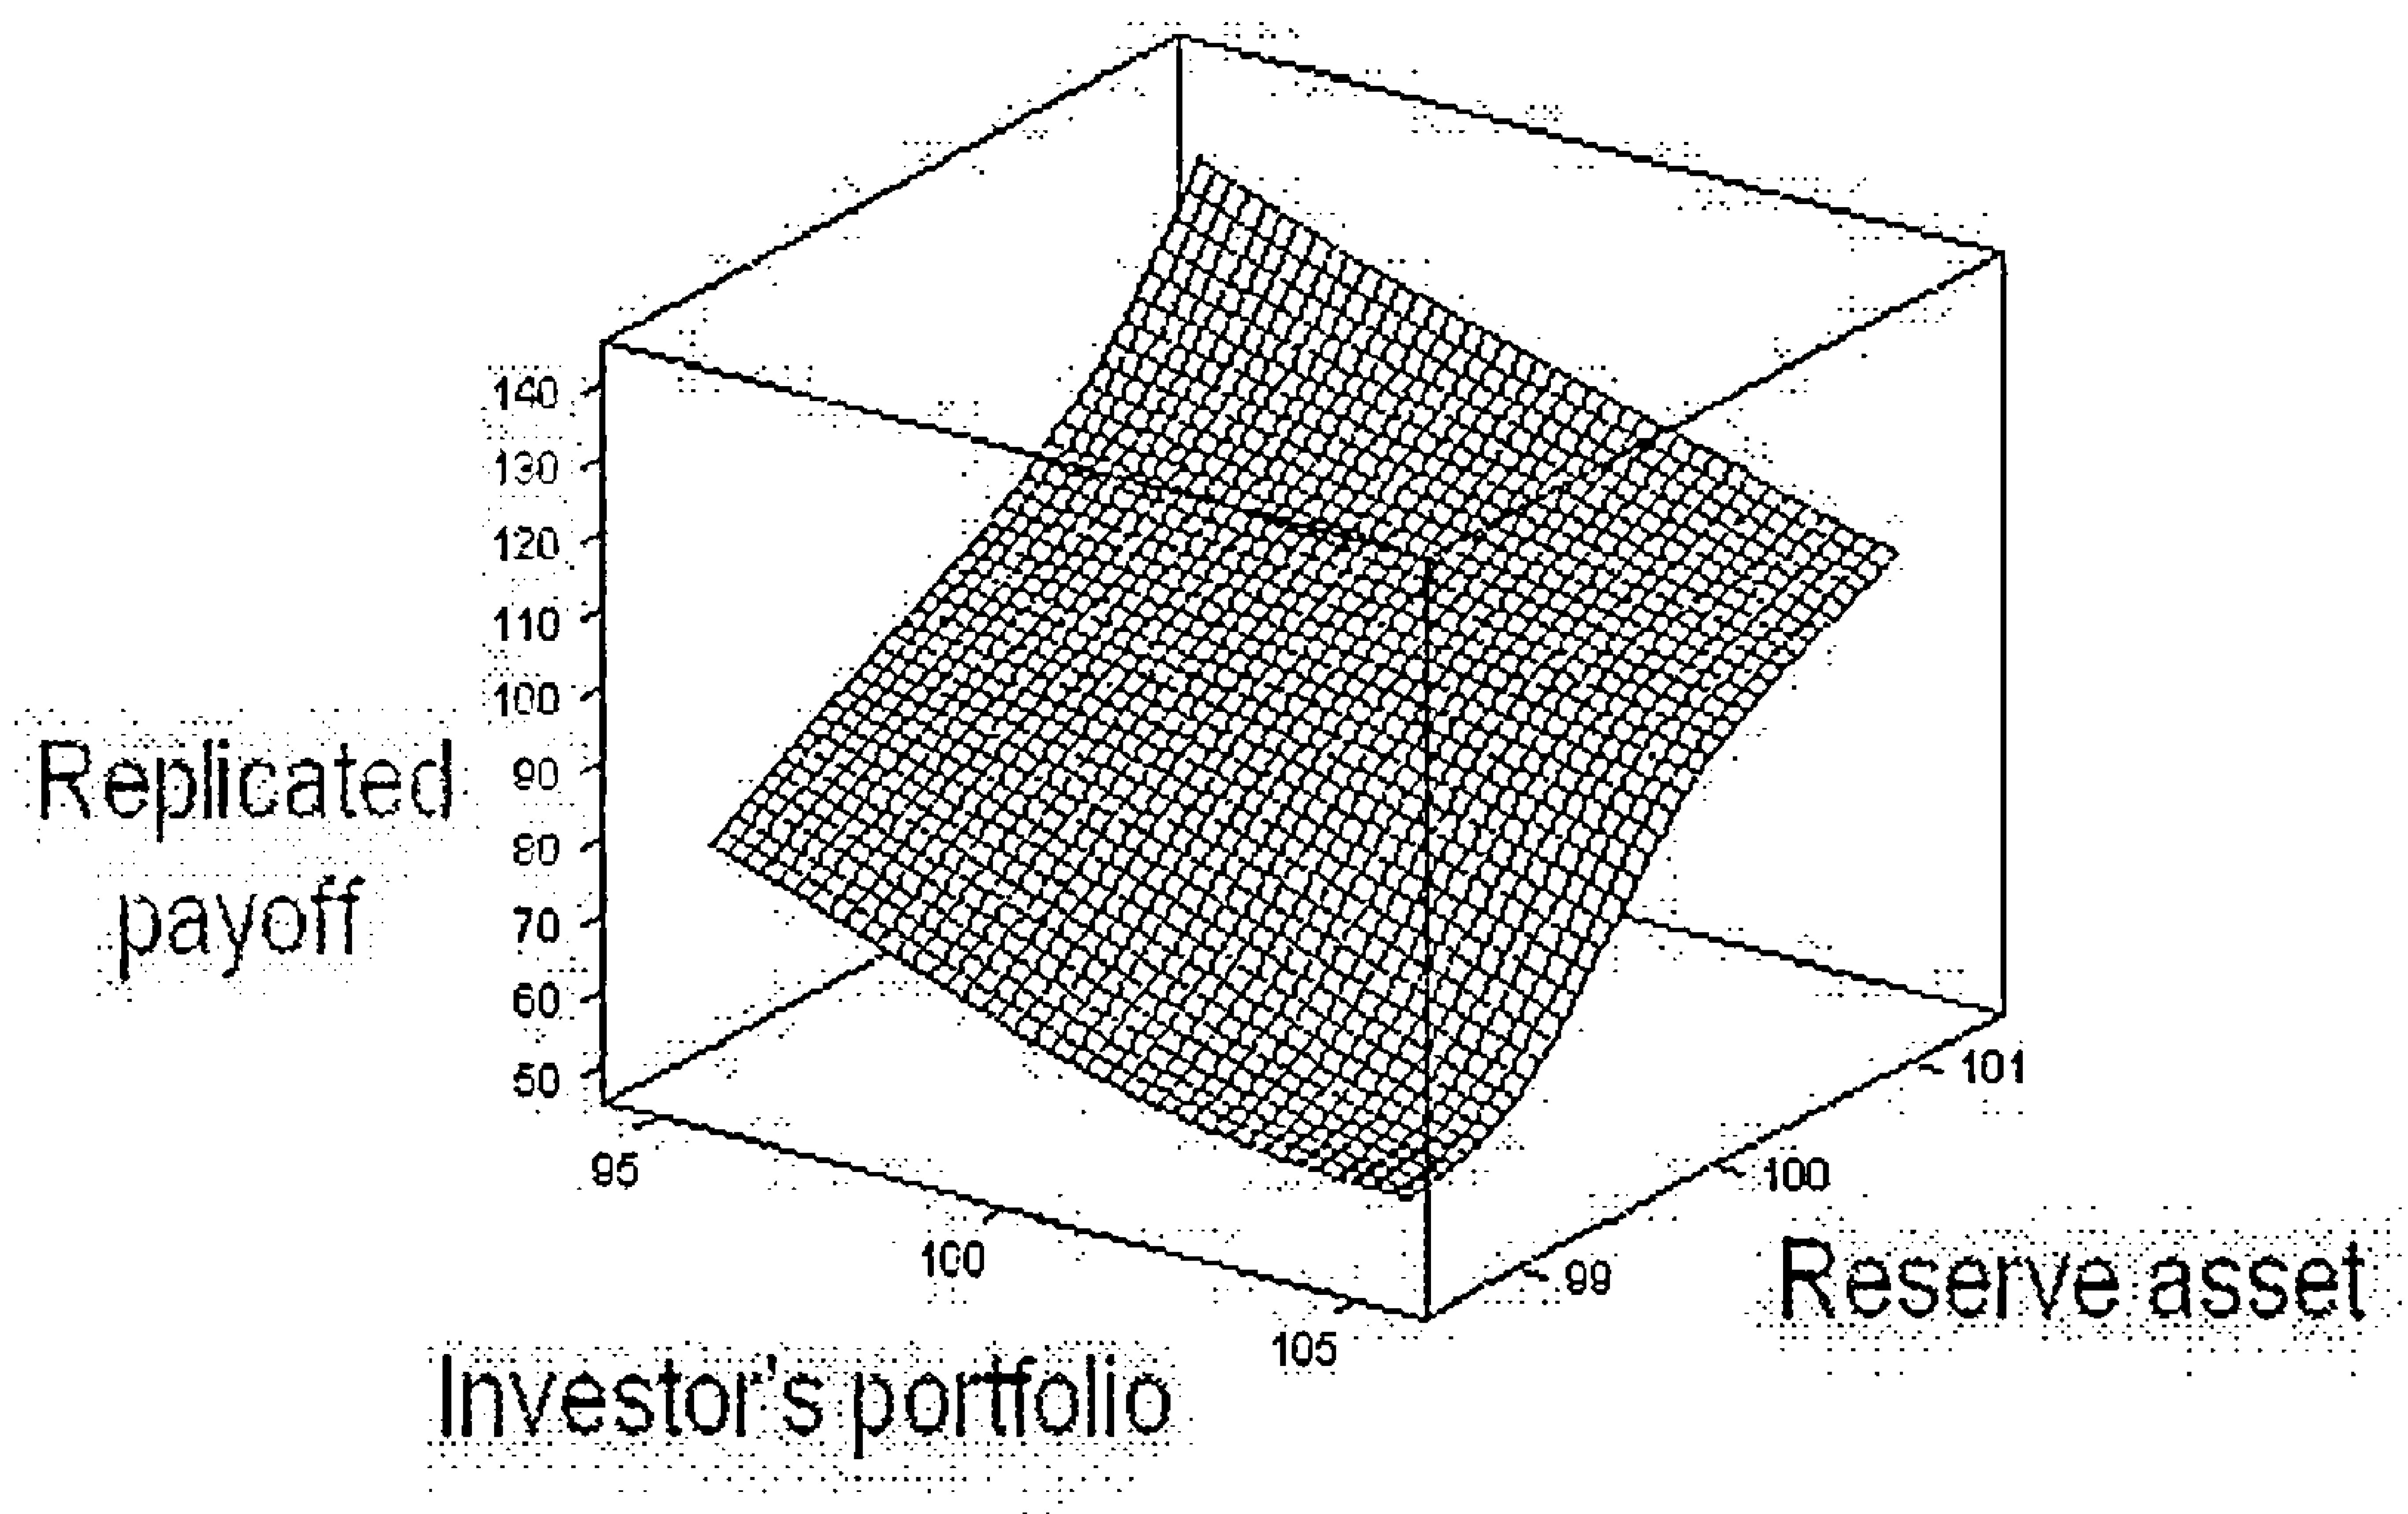

Figure 2.17. 3D plot payoff function Rocker Partners.

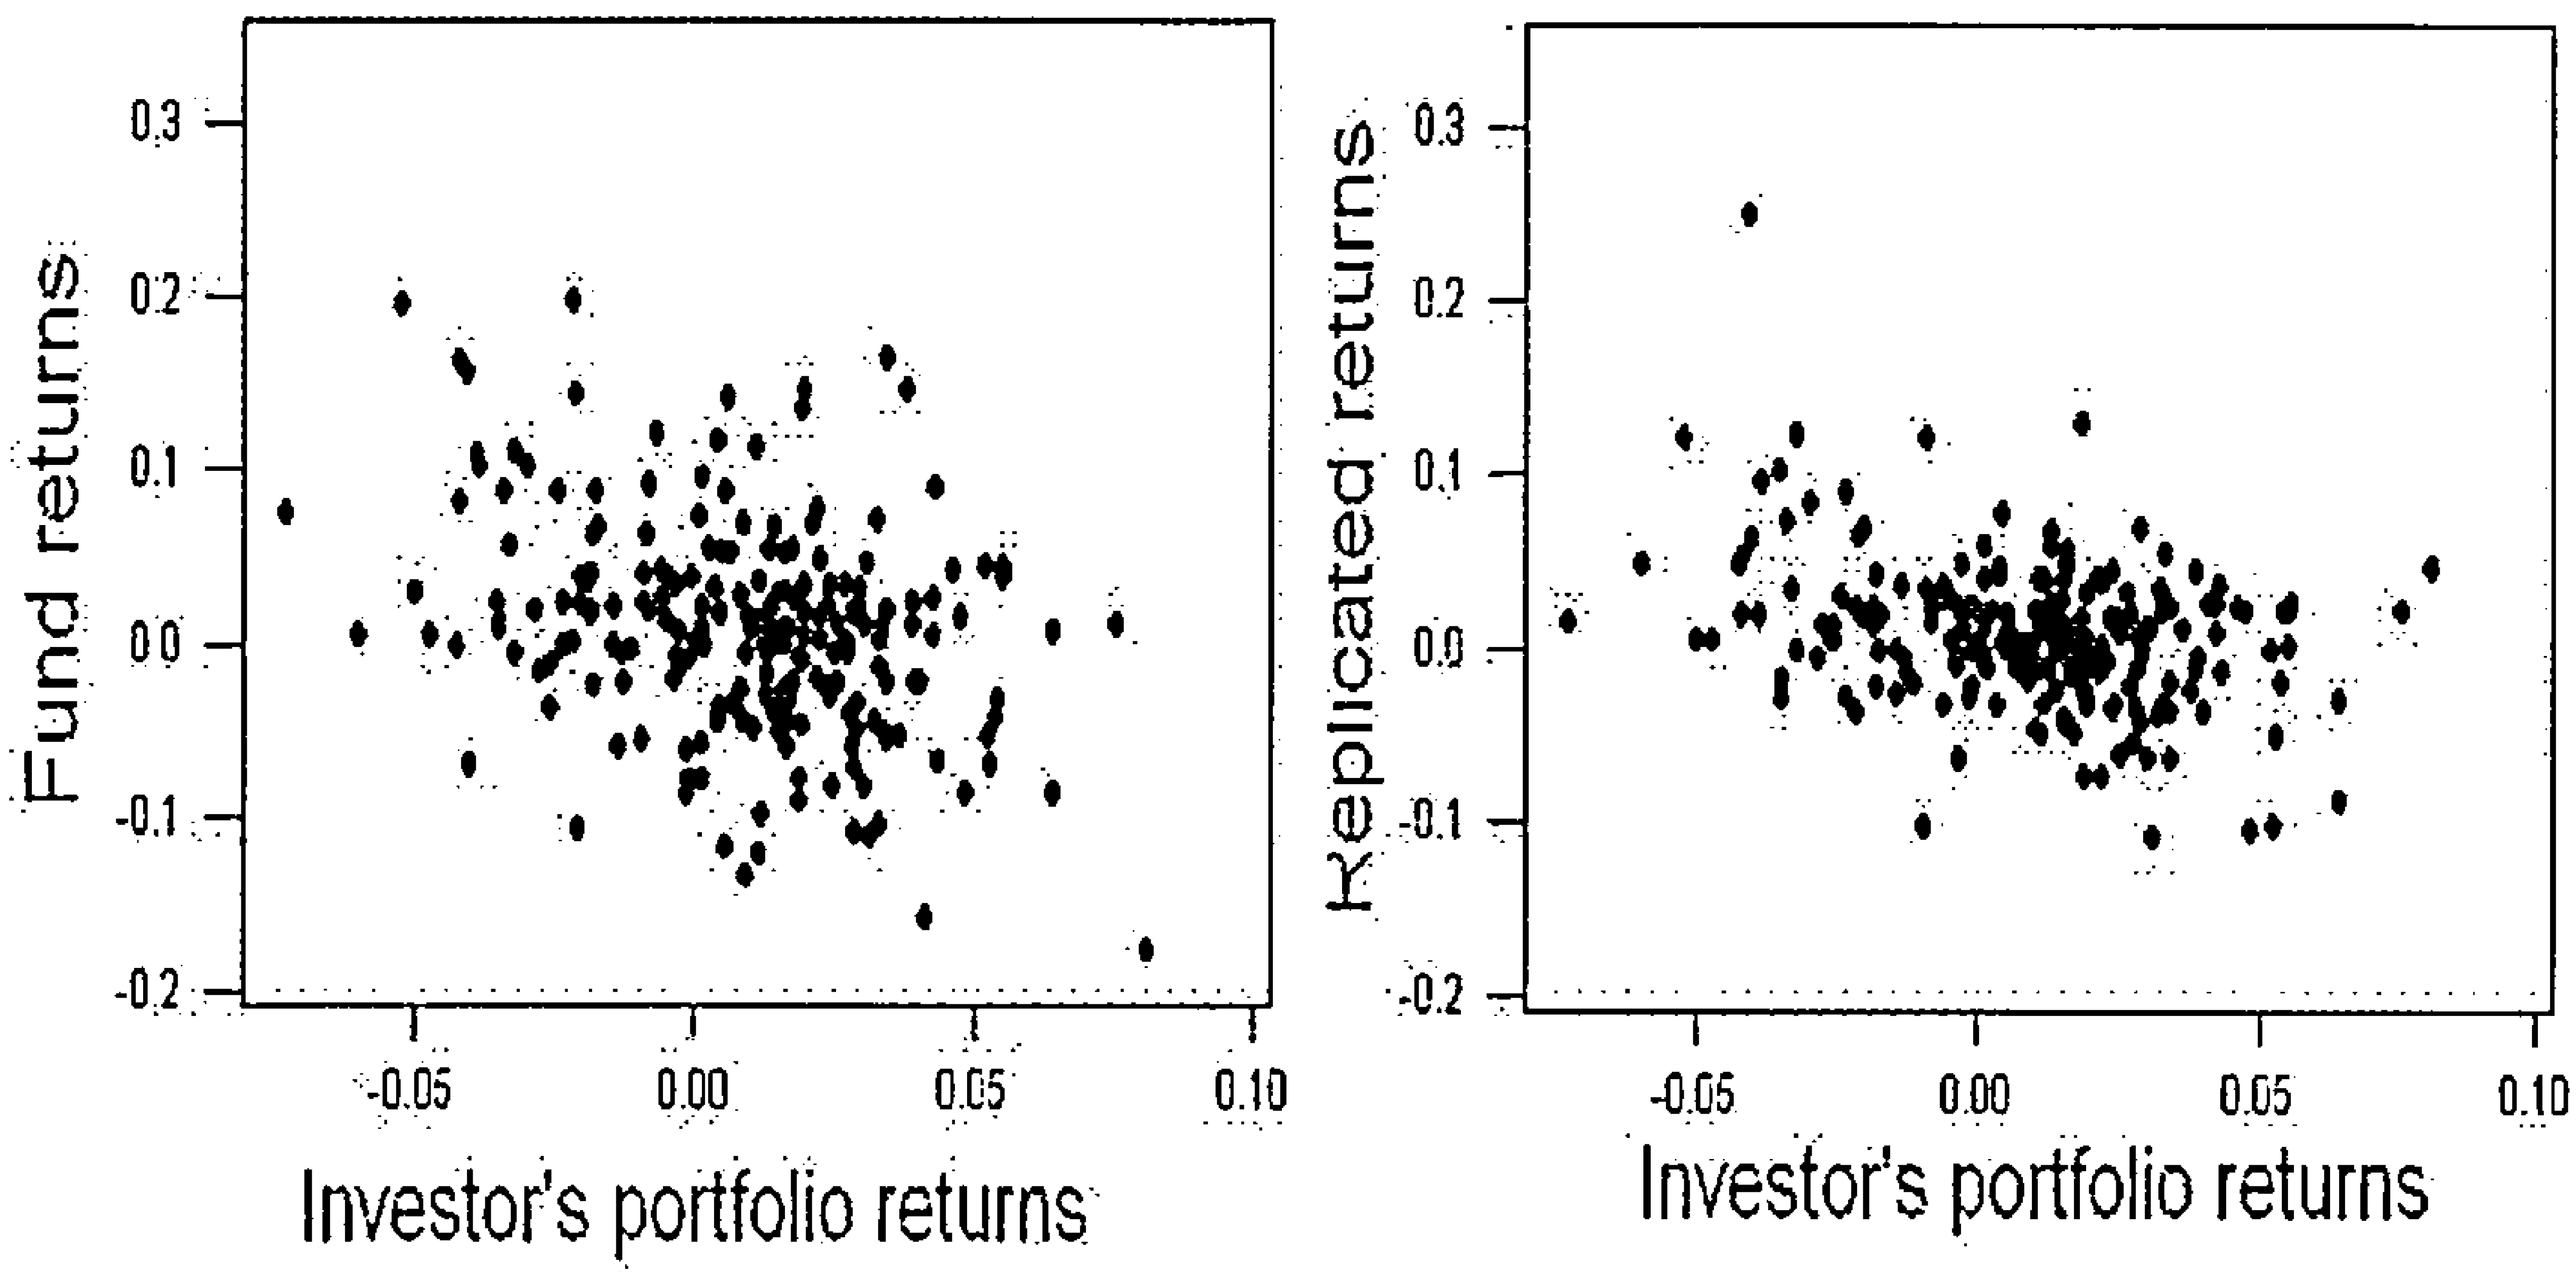

Figure 2.18. Scatter plot investor's portfolio returns versus Rocker Partners returns (left) and replicated returns (right), 1987 - 2004.

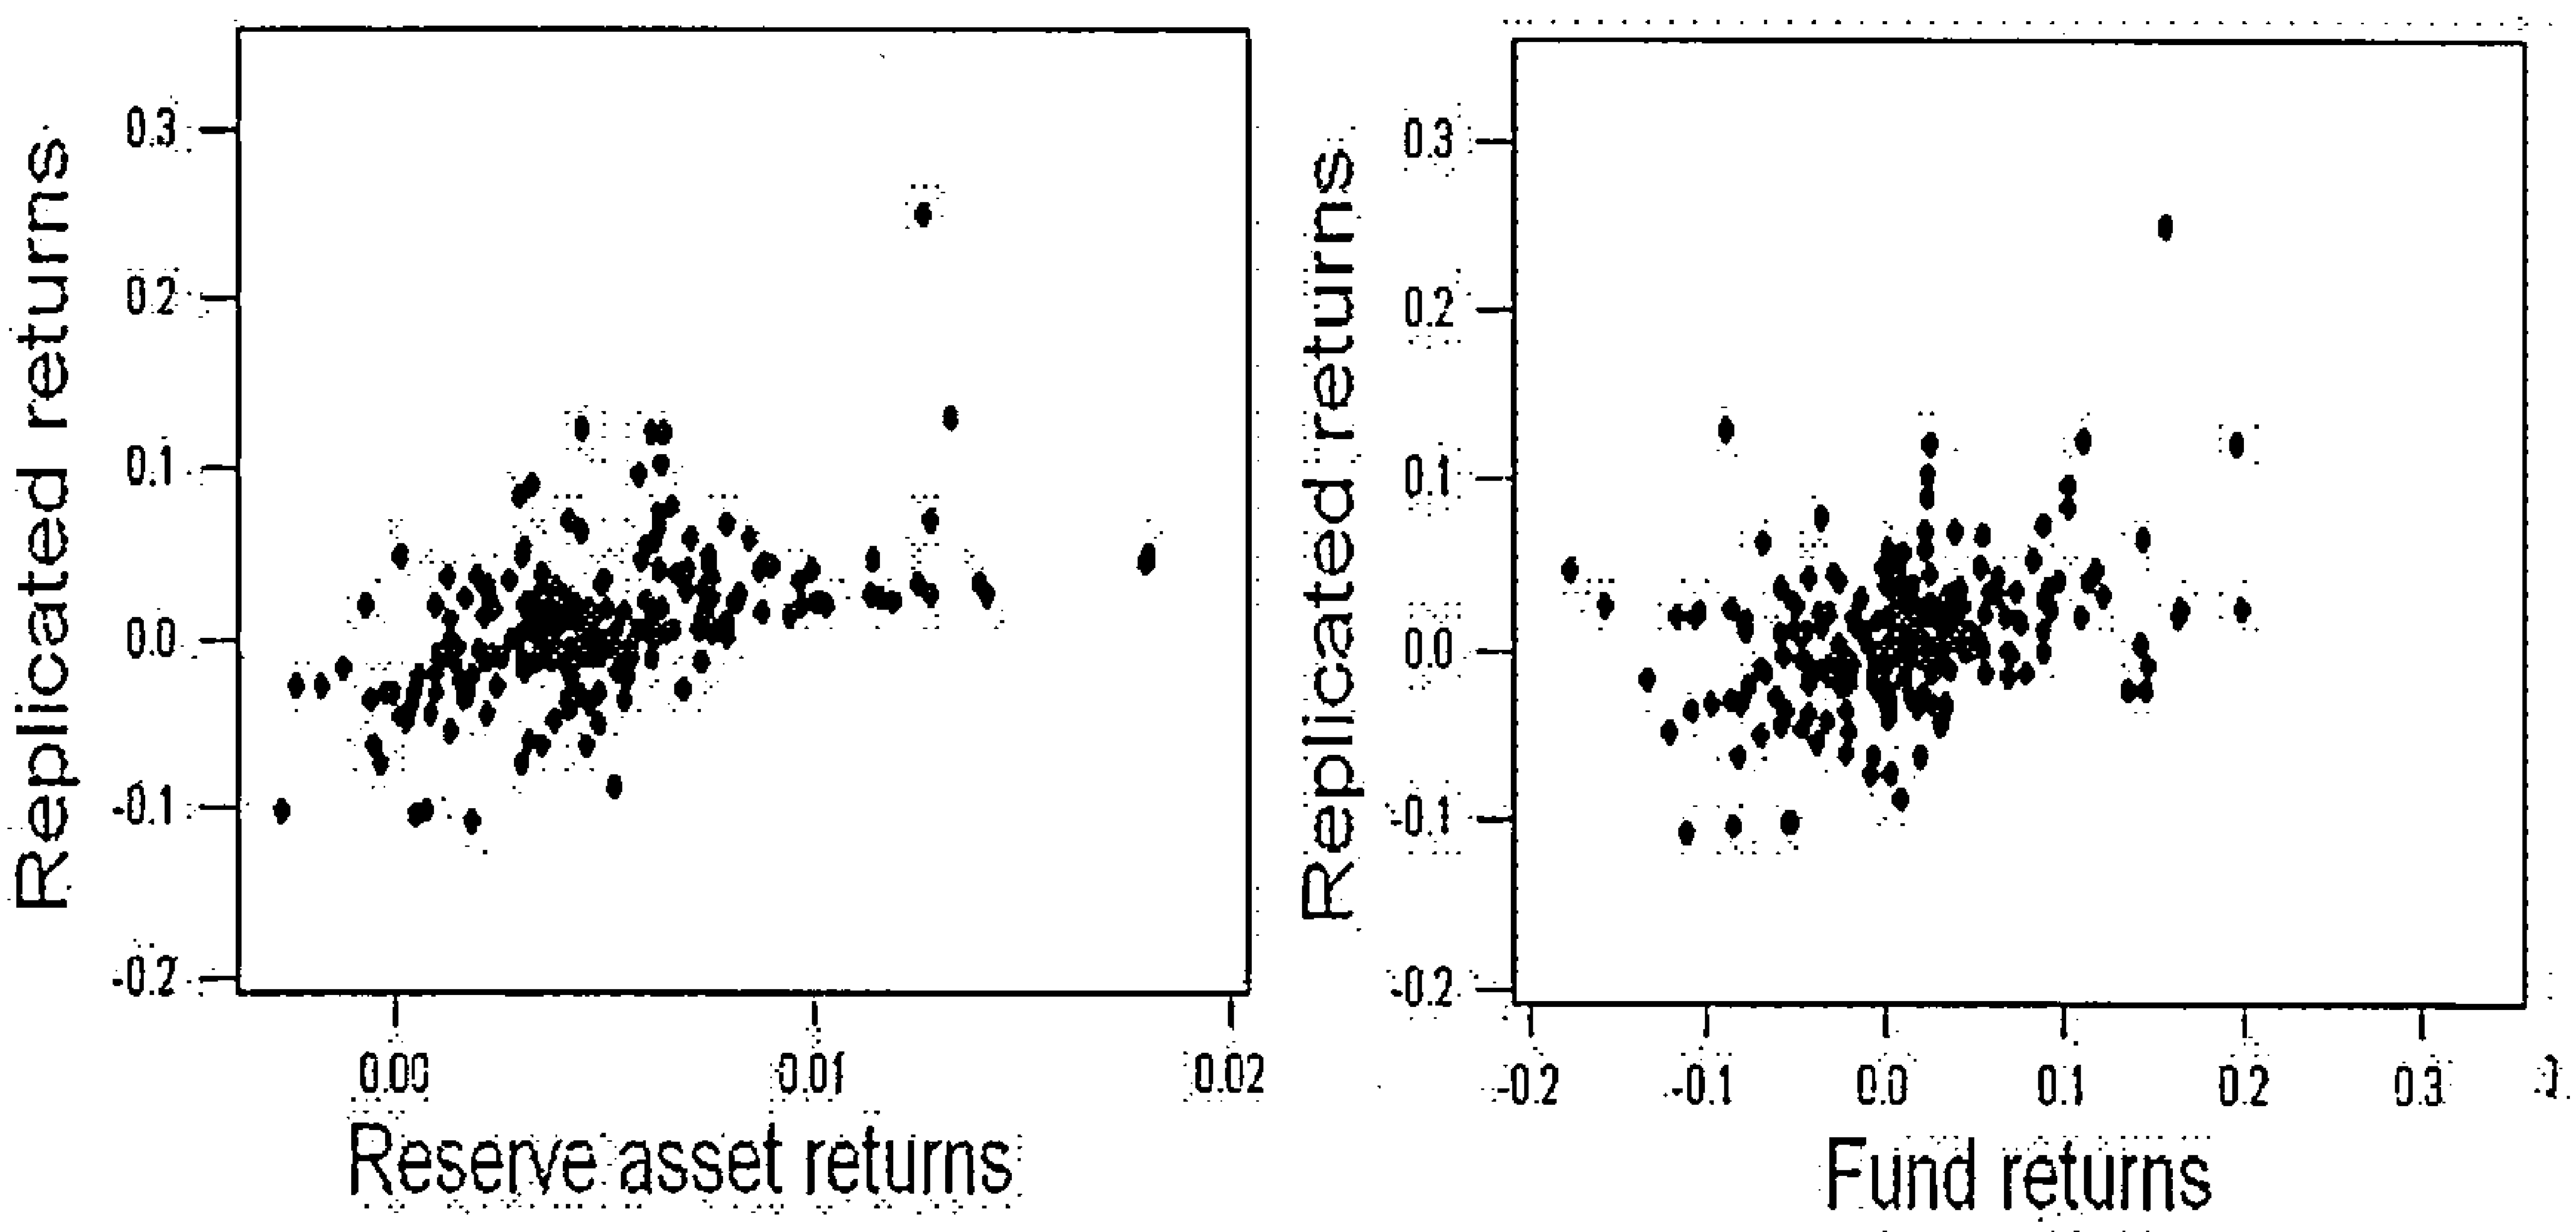

Figure 2.19. Scatter plot reserve asset returns versus replicated returns (left) and Rocker Partners returns versus replicated returns (right), 1987 – 2004.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt. (robust)	Corr. with Portfolio	Kendall's Tau
RP	0.0058	0.0684	-0.2456	-0.0992	1.5588	1.3862	-0.302	-0.179
Replica	0.0083	0.0430	0.8377	-0.0385	5.0043	1.5521	-0.346	-0.196
Univariate K-S Statistic = 0.117, (approximated) p-value = 0.101								
Bivariate K-S Statistic = 0.111, (approximated) p-value = 0.295								

Table 2.5: Monthly return statistics Rocker Partners and replication strategy, 1987 - 2004.

The desired payoff function for RP as per October 1st, 2004 can be found in Figure 2.16 and 2.17. From these graphs we see that the payoff function for RP is quite different from what we found for LCH and CMSF. Of course, the payoff is a positive function of the reserve asset. The RP payoff, however, is a negative function of the investor's portfolio. The replication strategy will therefore go long in the reserve asset, but short the investor's portfolio. This is of course what one would expect for a short seller, whose returns are likely to be negatively correlated with the market. From Figure 2.18 and Table 2.5 we see that the replication strategy performs in the same way as before. The replicated return statistics are again similar to those of the fund returns. Even the negative correlation with the investor's portfolio is closely replicated. Figure 2.19 paints a similar picture as Figures 2.10 and 2.15.

3. Evaluation of Funds of Funds and Hedge Funds[31]

3.1 Preliminary comments

All performance evaluation studies in finance follow the same general procedure. First, using a fund's track record and possibly some additional data over the same period as well, the fund return is characterized in some way. With the Sharpe ratio this is done by calculating the volatility of the fund return. With alphas this is done by estimating a fund's exposure to the relevant risk factors. Second, based on this characterization, a benchmark return is determined and compared with the actual average fund return over its track record. With the Sharpe ratio the benchmark return is derived from the average index return and the volatility of the index, while with alphas it derives from the average returns of the risk factors.

The replication-based evaluation procedure presented here is not different. It is just a different characterization. Where other approaches use volatility or factor loadings, this one uses the desired payoff function. Where other approaches use the average return on the index or the chosen risk factors, this one use the average interest rate, building block volatilities and correlation over a fund's track record to set a benchmark. What is different, however, is that there is no unrealistically strong assumptions concerning the exact nature of a fund's risk exposure or the behaviour of markets in general. As shown in Section 2.4, a fairly limited set of returns will often be enough to obtain a sufficiently good estimate of the desired distribution and the efficiency measure. As such, this procedure is quite robust.

Another point worth noting about the above evaluation procedure is the fact that it explicitly takes transaction costs into account by, instead of a Black-Scholes type option pricing model, using the Boyle and Lin (1997) model. In factor model based evaluations, transaction costs are typically ignored, despite the fact that maintaining the replicating portfolio's factor loadings at their desired levels is likely to require

[31] The material in this chapter appeared as Alternative Investment Research Center WP 40, Cass Business School.

significant periodic rebalancing. In addition, when dealing with hedge funds the risk factors used may be quite unusual and may therefore be accompanied by significant levels of transaction costs.

3.2 An Example

To clarify the above, let's look at a worked-out example. XYZ is a well-known fund of hedge funds, which started in 1985. Given XYZ's monthly, net-of-fee returns since 1985, the first step is to model the joint distribution of XYZ and the investor's portfolio, as well as the joint distribution of the investor's portfolio and the reserve asset. Before we can do so we need to decide what exactly the investor's portfolio and the reserve asset are, as well as unsmooth the raw fund return data.

Let's assume that the representative investor's portfolio consists of 50% S&P 500 and 50% long-dated US Treasury bonds. Let's also assume that all exposure management is done in the futures markets. Futures have several advantages over cash, in particular high liquidity and low transaction costs, which is extremely important given the dynamic nature of the KP replication strategies. We trade S&P 500 futures on the CME and T-bond futures on the CBOT. To keep things simple, we use nearby Eurodollar futures (CME) as the reserve asset.

	Standard Deviation	Skewness	Excess Kurtosis	1M Auto Correlation
XYZ smooth	0.0370	-1.726	11.505	0.138
XYZ unsmooth	0.0424	-1.746	11.581	0.008

Table 3.1. Risk statistics XYZ.

Table 3.1 shows the marginal risk characteristics of the raw and unsmoothed XYZ returns. From the table, we see that XYZ's raw returns exhibit negative skewness and positive autocorrelation. Application of the unsmoothing procedure eliminates the autocorrelation and produces returns with the same degree of skewness, but with a substantially higher volatility (annualised 14.7% vs. 12.8% for the raw returns).

We are now ready to infer the desired and the building block distribution. Using the same methodology as in Chapter 2, we find that the best fit (according to the AIC) is provided by the following set of marginals and copulas[32]:

XYZ: Student-t ($\mu = 0.0101$, $\sigma = 0.0406$, df $= 4.0544$)
Portfolio: Normal ($\mu = 0.0101$, $\sigma = 0.0282$)
Reserve: Johnson ($\xi = 0.0031$, $\lambda = 0.0046$, $\gamma = -0.60$, $\delta = 1.599$)
Copula (XYZ, portfolio): Normal ($\rho = 0.754$)
Copula (portfolio, reserve): Gumbel ($a = 1.3349$)

Given the above distributions, we can derive the desired payoff function following the methodology developed in Chapter 2. The result is depicted in Figure 3.1 and shows that the desired payoff is an increasing function of both the investor's portfolio and the reserve asset, implying that the replication strategy will take long positions in both assets. Subsequently, we price this payoff function using the Boyle and Lin (1997) model, assuming transaction costs in the futures markets are 1bp one-way. This produces a value for the KP efficiency measure of 99.53, meaning that, seen over the whole life of the fund, XYZ's returns are not as miraculous as many investors may have thought. Trading S&P 500, T-bond and Eurodollar futures, investors could have generated the same risk profile as XYZ and obtained a higher average return at the same time.

[32] Copula functions and marginal distributions defined in Appendix A.3.

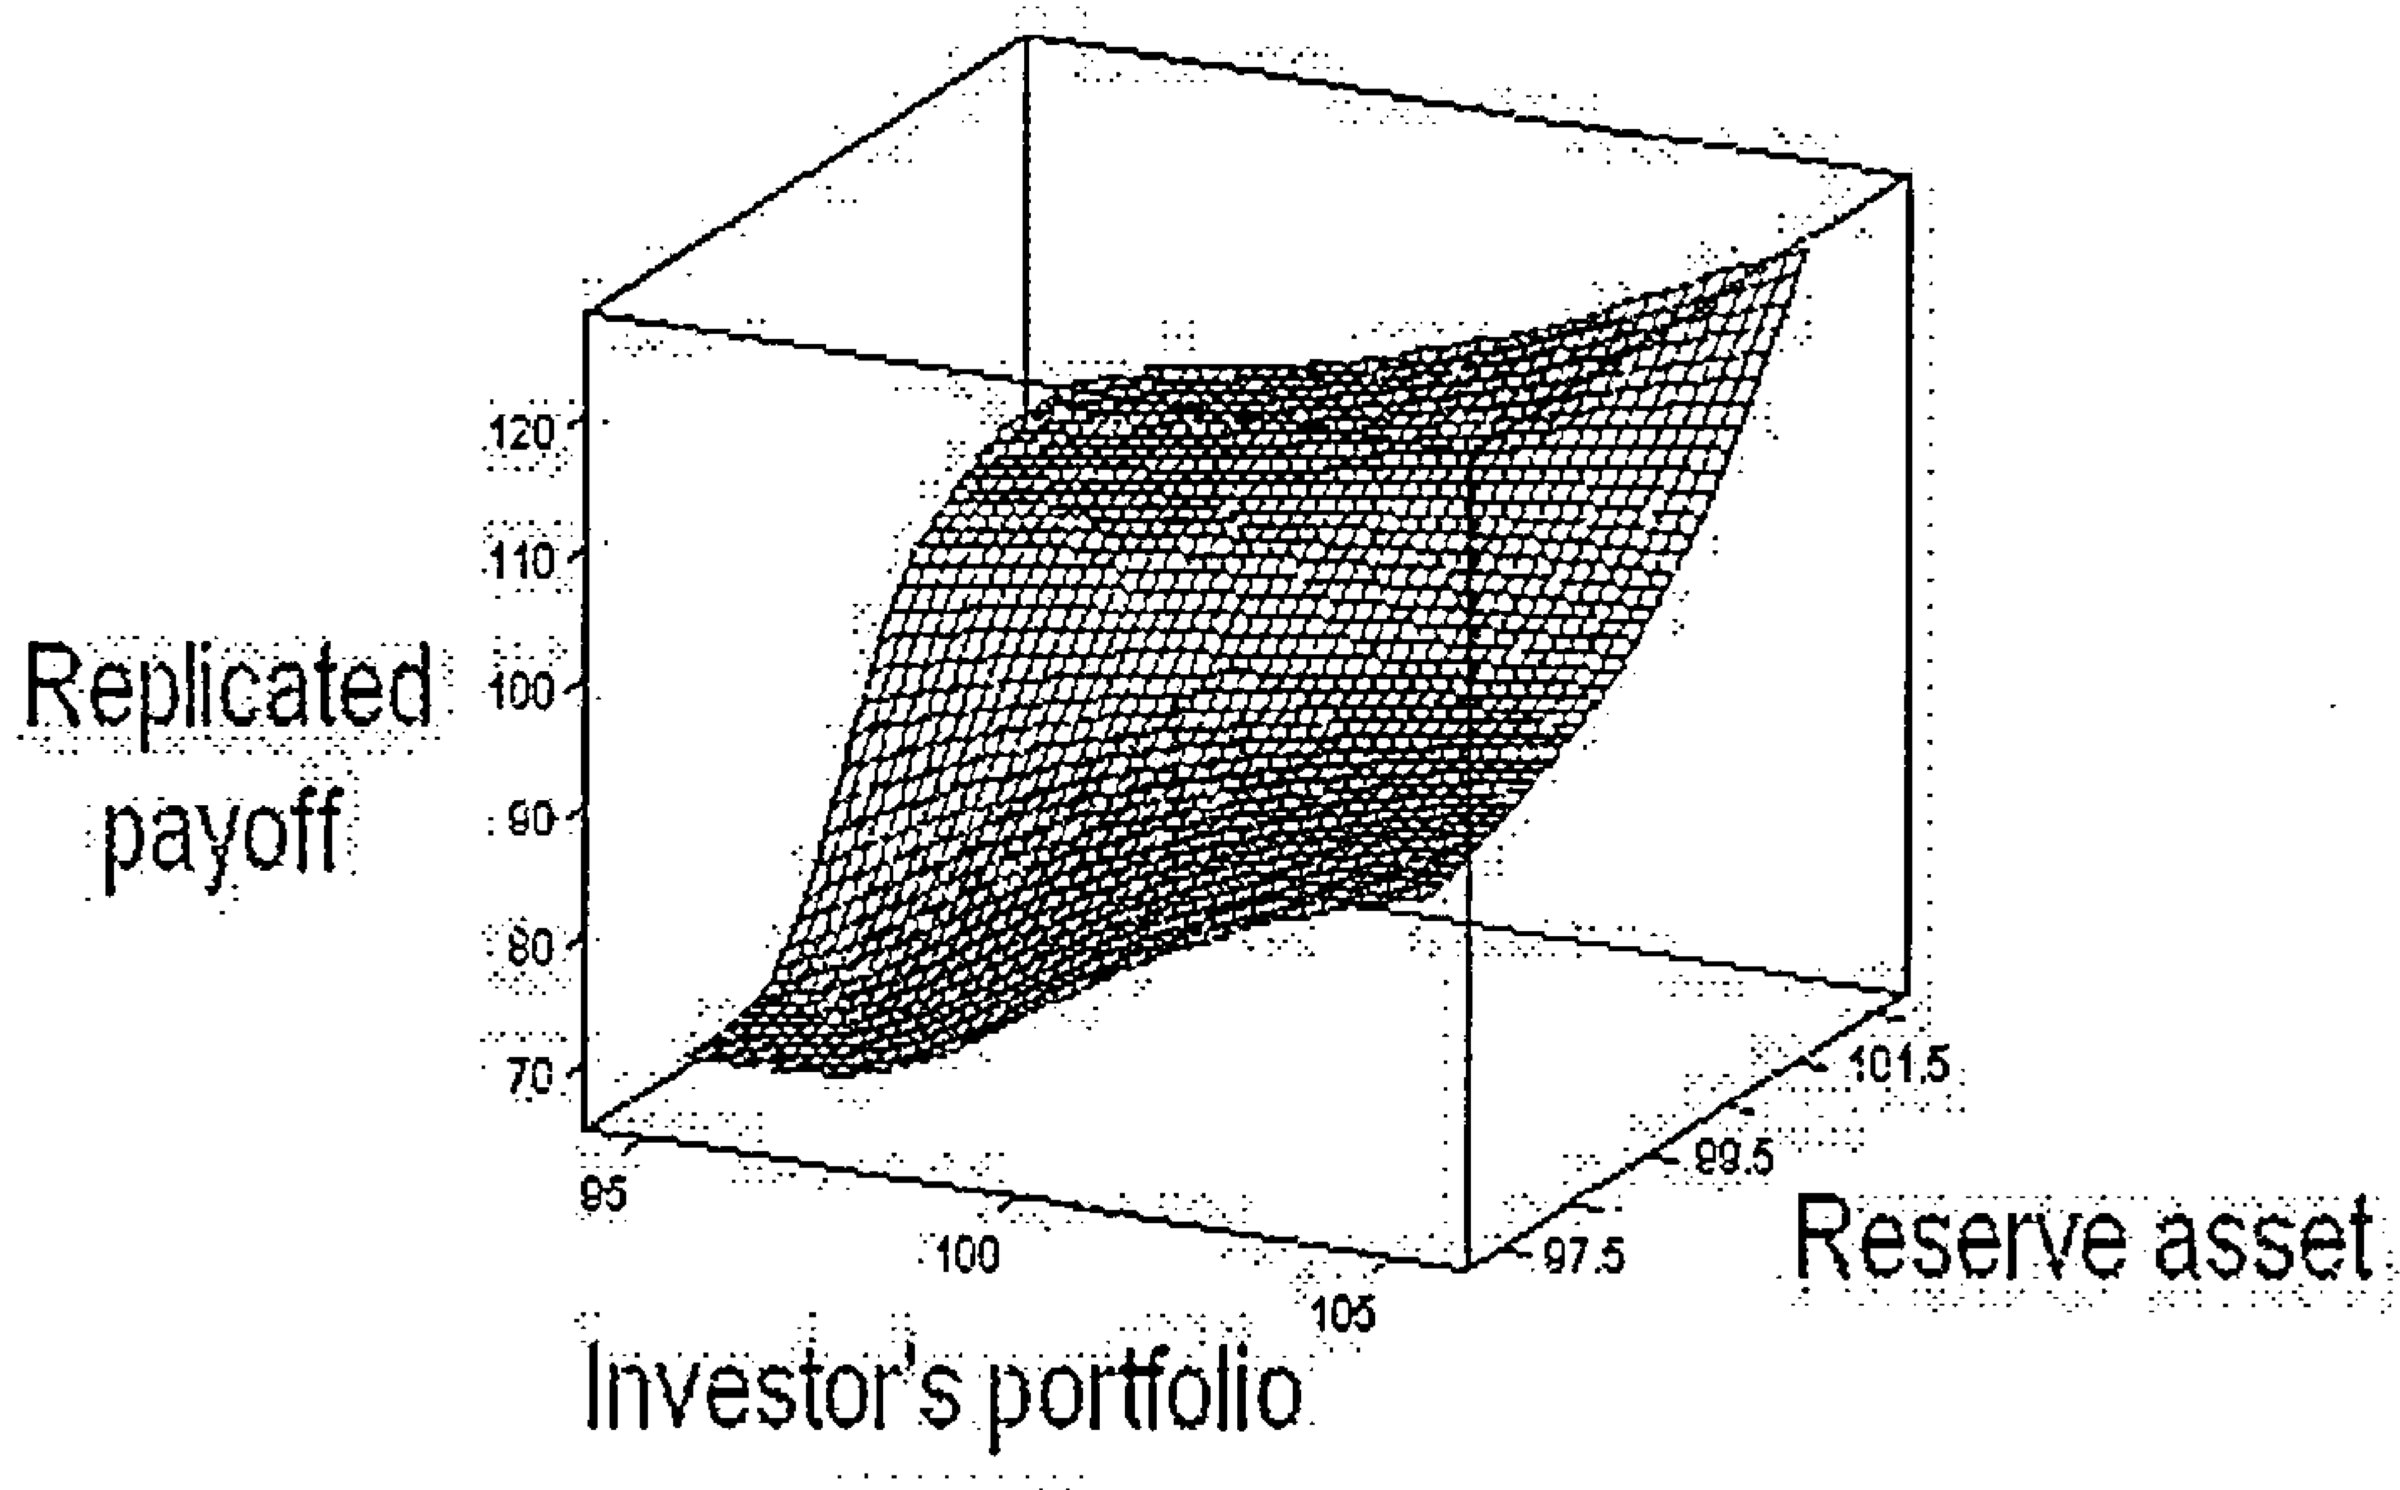

Figure 3.1. Desired payoff function for replication XYZ returns.

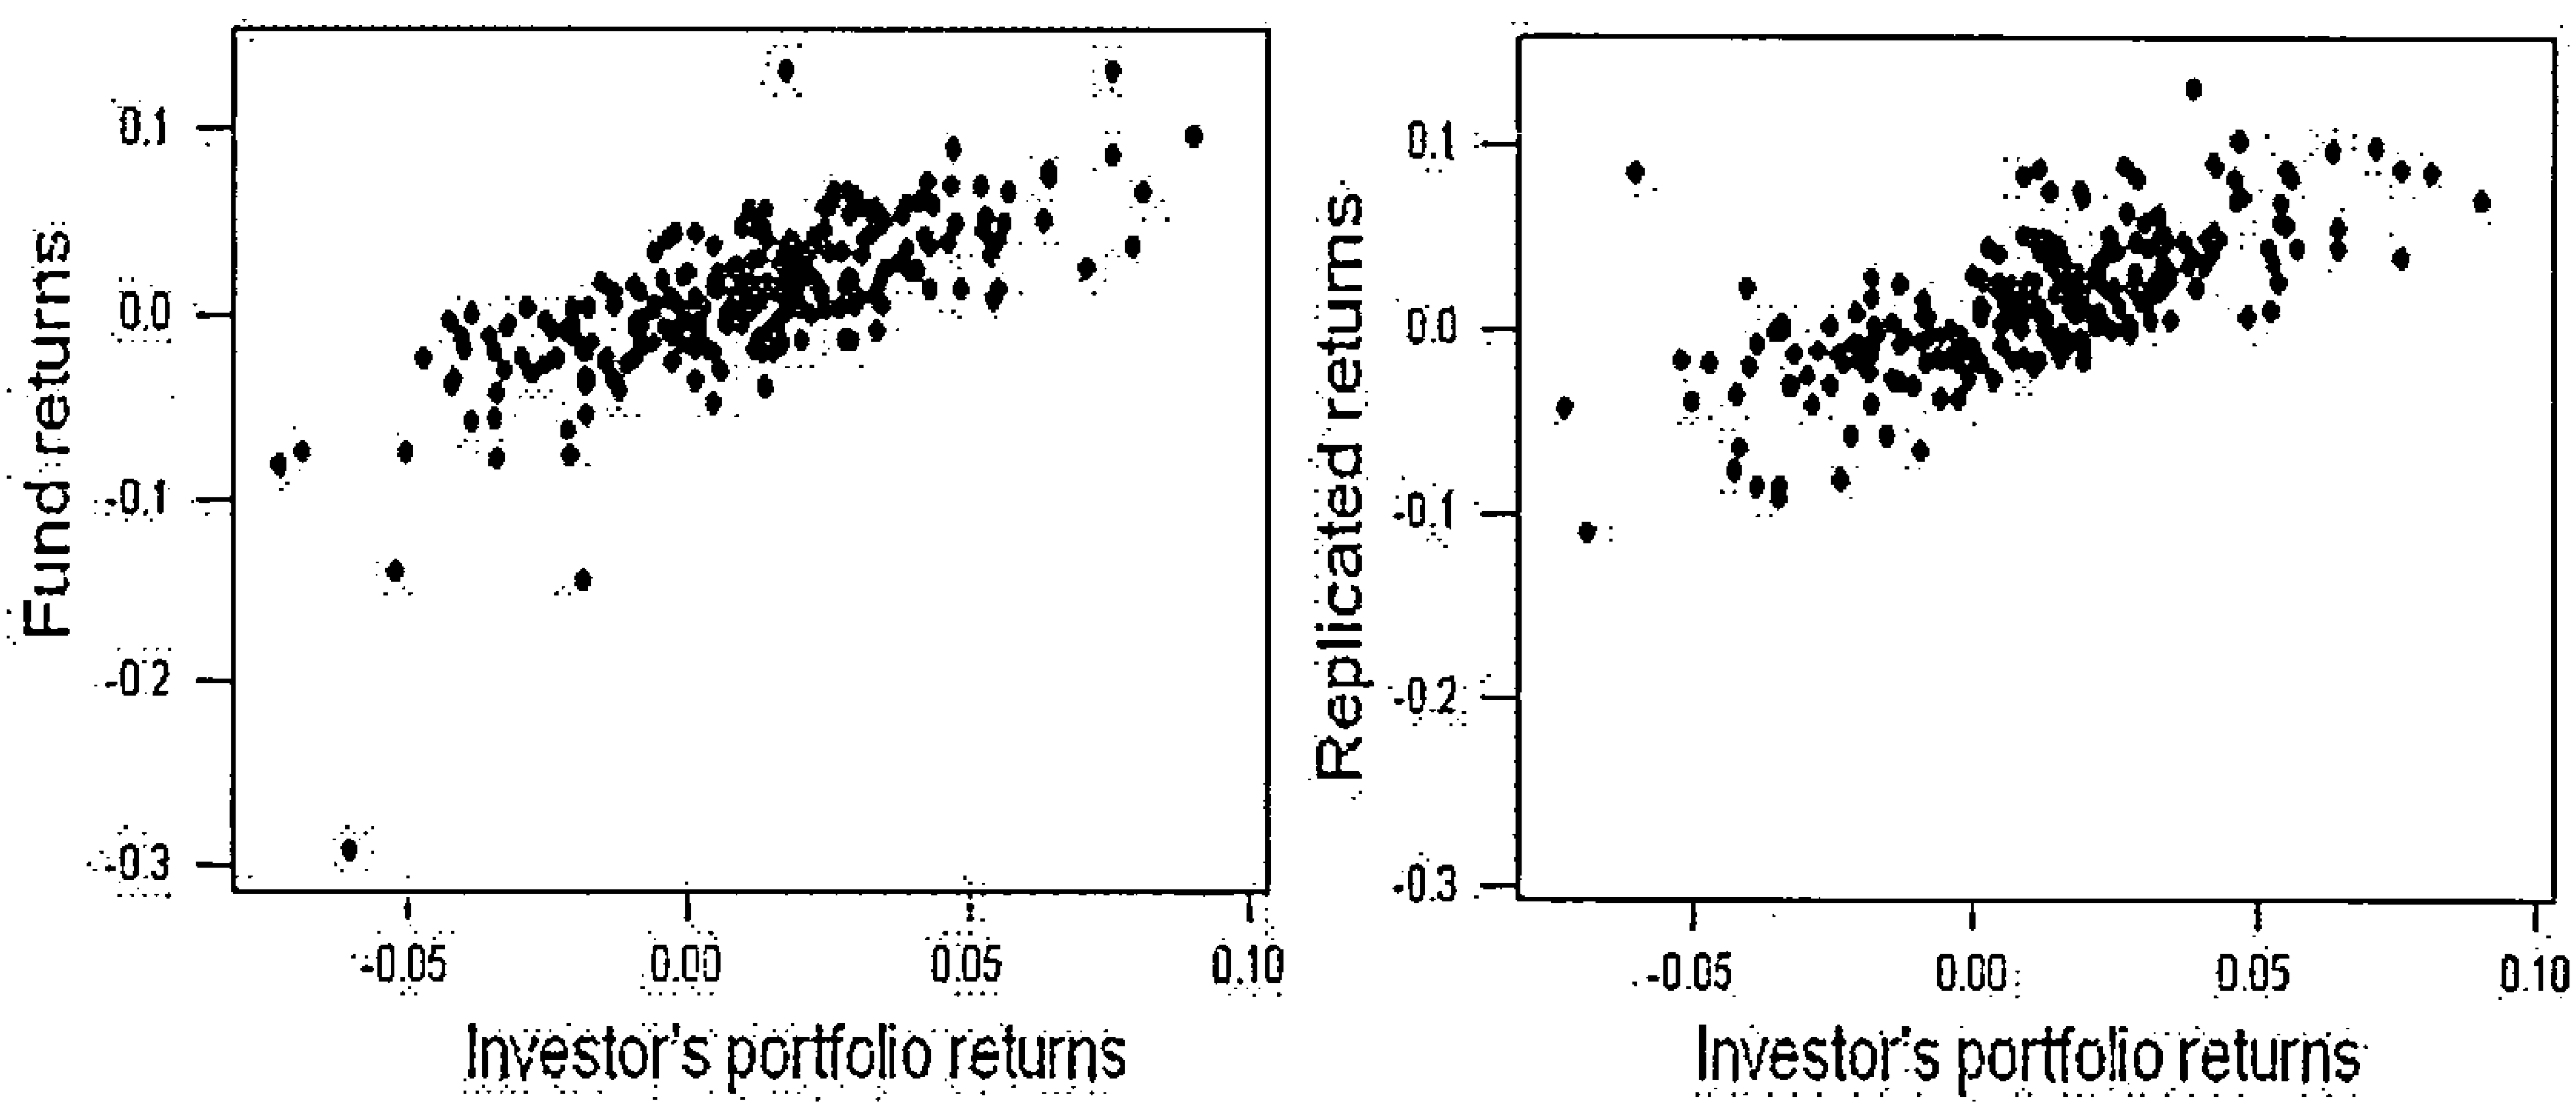

Figure 3.2. Scatter plot investor's portfolio returns versus XYZ returns (left) and replicated returns (right).

To see how well the derived payoff function succeeds in replicating the desired distribution, Figure 3.2 shows a scatter plot of the investor's portfolio return versus the XYZ return (left) as well as a plot of the portfolio return versus the replicated return (right). The two plots are very similar, suggesting that the replication has indeed been successful. We see that the replication strategy is unable to replicate the three large losses that XYZ reported during the sample period. This is not surprising as these are clearly outliers, which simply cannot be captured by a parametric model like ours.

	Mean	St. Dev	Skewness	Skewness (robust)	Excess Kurtosis	Ex. Kurt. (robust)	Corr. with Portfolio	Kendall's Tau
XYZ	0.0102	0.0424	-1.7463	-0.1600	11.5812	0.4366	0.714	0.540
Replica	0.0150	0.0388	0.1184	-0.1269	1.2691	0.6889	0.721	0.548
Univariate K-S Statistic = 0.054, (approximated) p-value = 0.862								
Bivariate K-S Statistic = 0.056, (approximated) p-value 0.924								

Table 3.2. Statistics XYZ and replicated returns.

A further indication of the accuracy of the replication strategy comes from comparing the mean, standard deviation, skewness and kurtosis of XYZ's returns with those of the replicated returns. The latter statistics can be found in Table 3.2, together with the correlation and Kendall's Tau with the investor's portfolio. Since the XYZ returns exhibit a few negative outliers, apart from the standard skewness and kurtosis measures, we also report a more robust skewness and kurtosis measure[33]. 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[34]. Comparing the entries in Table 3.2, it is clear that the statistical properties of XYZ's returns have been successfully replicated. The replication strategy has not only replicated the marginal distribution of XYZ's returns but also its relationship with the investor's portfolio. The same conclusion follows from both the K-S tests.

Having clarified the procedure to be used, we proceed with the evaluation of the performance of 875 funds of hedge funds and 2073 individual hedge funds. Since funds of funds are distinctly different from individual hedge funds, we do so in two separate parts. Section 3.3 deals with funds of funds. Section 3.4 deals with individual hedge funds.

[33] See Hinkley (1975) and Crow and Siddiqui (1967) for details.

[34] See Fasano and Franceschini (1987) for details.

3.3 Funds of Hedge Funds

3.3.1 Funds of Hedge Funds: Data Description

Our sample consists of 875 funds of hedge funds with a minimum of 4 years of history available. All data were obtained from The Barclay Group as per November 2006. Funds denominated in another currency than USD are converted to USD, i.e. the perspective taken is that of a USD-based investor. Table 3.3 provides some information on the start and end dates of the track records of the funds in our sample.

	Jan 1985	Jan 1988	Jan 1991	Jan 1994	Jan 1997	Jan 2000	Jan 2003	Jan 2006	Nov 2006
Start after	870	863	834	775	613	381	0	0	0
End before	0	0	0	0	0	0	46	122	218

Table 3.3. Fund of funds start and end date details.

Table 3.3 shows that, reflecting the increasing popularity of hedge funds in the second half of the 1990s, the majority of funds started after 1994. Most hedge fund databases, first started collecting data around the mid to late 1990s. As a result, they contain no funds that stopped reporting before that date. Out of the 875 funds in our database, 218 funds stopped reporting before November 2006. This confirms that, although lower than for individual hedge funds, the attrition rate in funds of funds is still quite high.

	4-5Y	5-6Y	6-7Y	7-8Y	8-9Y	9-10Y	10Y-11Y	11Y-12Y	12Y-14Y	14Y+
No. funds	211	137	106	83	66	52	61	53	48	58

Table 3.4. Length fund of funds track records.

Table 3.4 provides details on the length of the available fund of funds track records. Out of the 875 funds in the sample, only 220 have 10 or more years of history. This again reflects the fact that most funds of funds are still relatively young and attrition can be significant.

3.3.2 Funds of Hedge Funds: Distributional Analysis

A crucial stage in the evaluation procedure is the proper modelling of the distributional characteristics of the fund, the investor's portfolio and the reserve asset. This means that, although not explicitly designed to do so, the evaluations provide a wealth of information on the distributional properties of fund of funds returns. Table 3.5 summarizes how often (out of a total of 875 funds) a given marginal or copula was used in the evaluations for modelling the fund return marginal and the joint distribution of the fund and the investor's portfolio return.

Marginals	No.	Copulas	No.
Normal	40.7%	Normal	29.4%
Student-t	47.1%	Student	6.0%
Johnson	12.2%	Gumbel	8.0%
		SJC	9.6%
		Cook-Johnson	19.7%
		Frank	27.3%

Table 3.5. Distributional characteristics fund of funds returns.

Table 3.5 confirms that, despite an often substantial degree of diversification in the larger funds, the majority of fund of funds returns are far from normally distributed. Out of 875 funds, 519 (59.3%) funds' marginal return is better modelled by a Student-t or Johnson distribution than a normal distribution. In addition, for only 257 (29.3%) of the 875 funds is the relationship with the investor's portfolio (consisting of 50% S&P 500 and 50% T-bonds) best modelled by the normal copula. This emphasizes how important it is to evaluate fund of funds performance using a method, which does not rely on the assumption of normally distributed returns.

3.3.3 Funds of Hedge Funds: Replication and Evaluation Results

As in the example in Section 3.2, in the evaluations we assume that the representative investor's portfolio consists of 50% S&P 500 and 50% long-dated US Treasury bonds, with all exposure management done through collateralised nearby futures contracts. The reserve asset is taken to consist of a basket of nearby Eurodollar (CME), 5-year note (CBOT), 10-year note (CBOT), S&P 500 (CME), Russell 2000 (CME) and GSCI

(CME) futures. We chose this particular basket for no other reason than that is well-diversified over three asset classes and therefore contains relatively little uncompensated risk. To compensate for their low volatility, the Eurodollar and 5-year note futures are leveraged by a factor 5 and the 10-year note futures by a factor 4.

Transaction costs on all futures contracts are assumed to be 1bp one-way. Commission costs in futures tend to be extremely low, while quoted bid ask spreads in the most liquid contracts are typically not much higher than a few basis points. The quoted spread, however, is not necessarily a good approximation for the actual transaction costs experienced. The time-sensitivity of the required trades is typically quite low, as these trades are primarily exposure adjustments and are not meant to capture a specific short-term profit opportunity. It is therefore not necessary to trade on the first available bid or ask. Over time, competing with the market makers and trying to buy on the bid and sell at the offer could yield a very substantial saving in transaction costs. In fact, in managed futures trading programs, like Man's well-known AHL program for example, good execution tends to account for a very significant part of the bottom line P&L.

For the pricing of the payoff functions, we use average 1-month USD Libor as the relevant interest rate, while estimating the required volatilities and correlations over the period covered by the track record of the fund that is being evaluated. The interest rate data was obtained from Datastream, while the futures data was obtained from Commodity Systems Inc. (CSI).

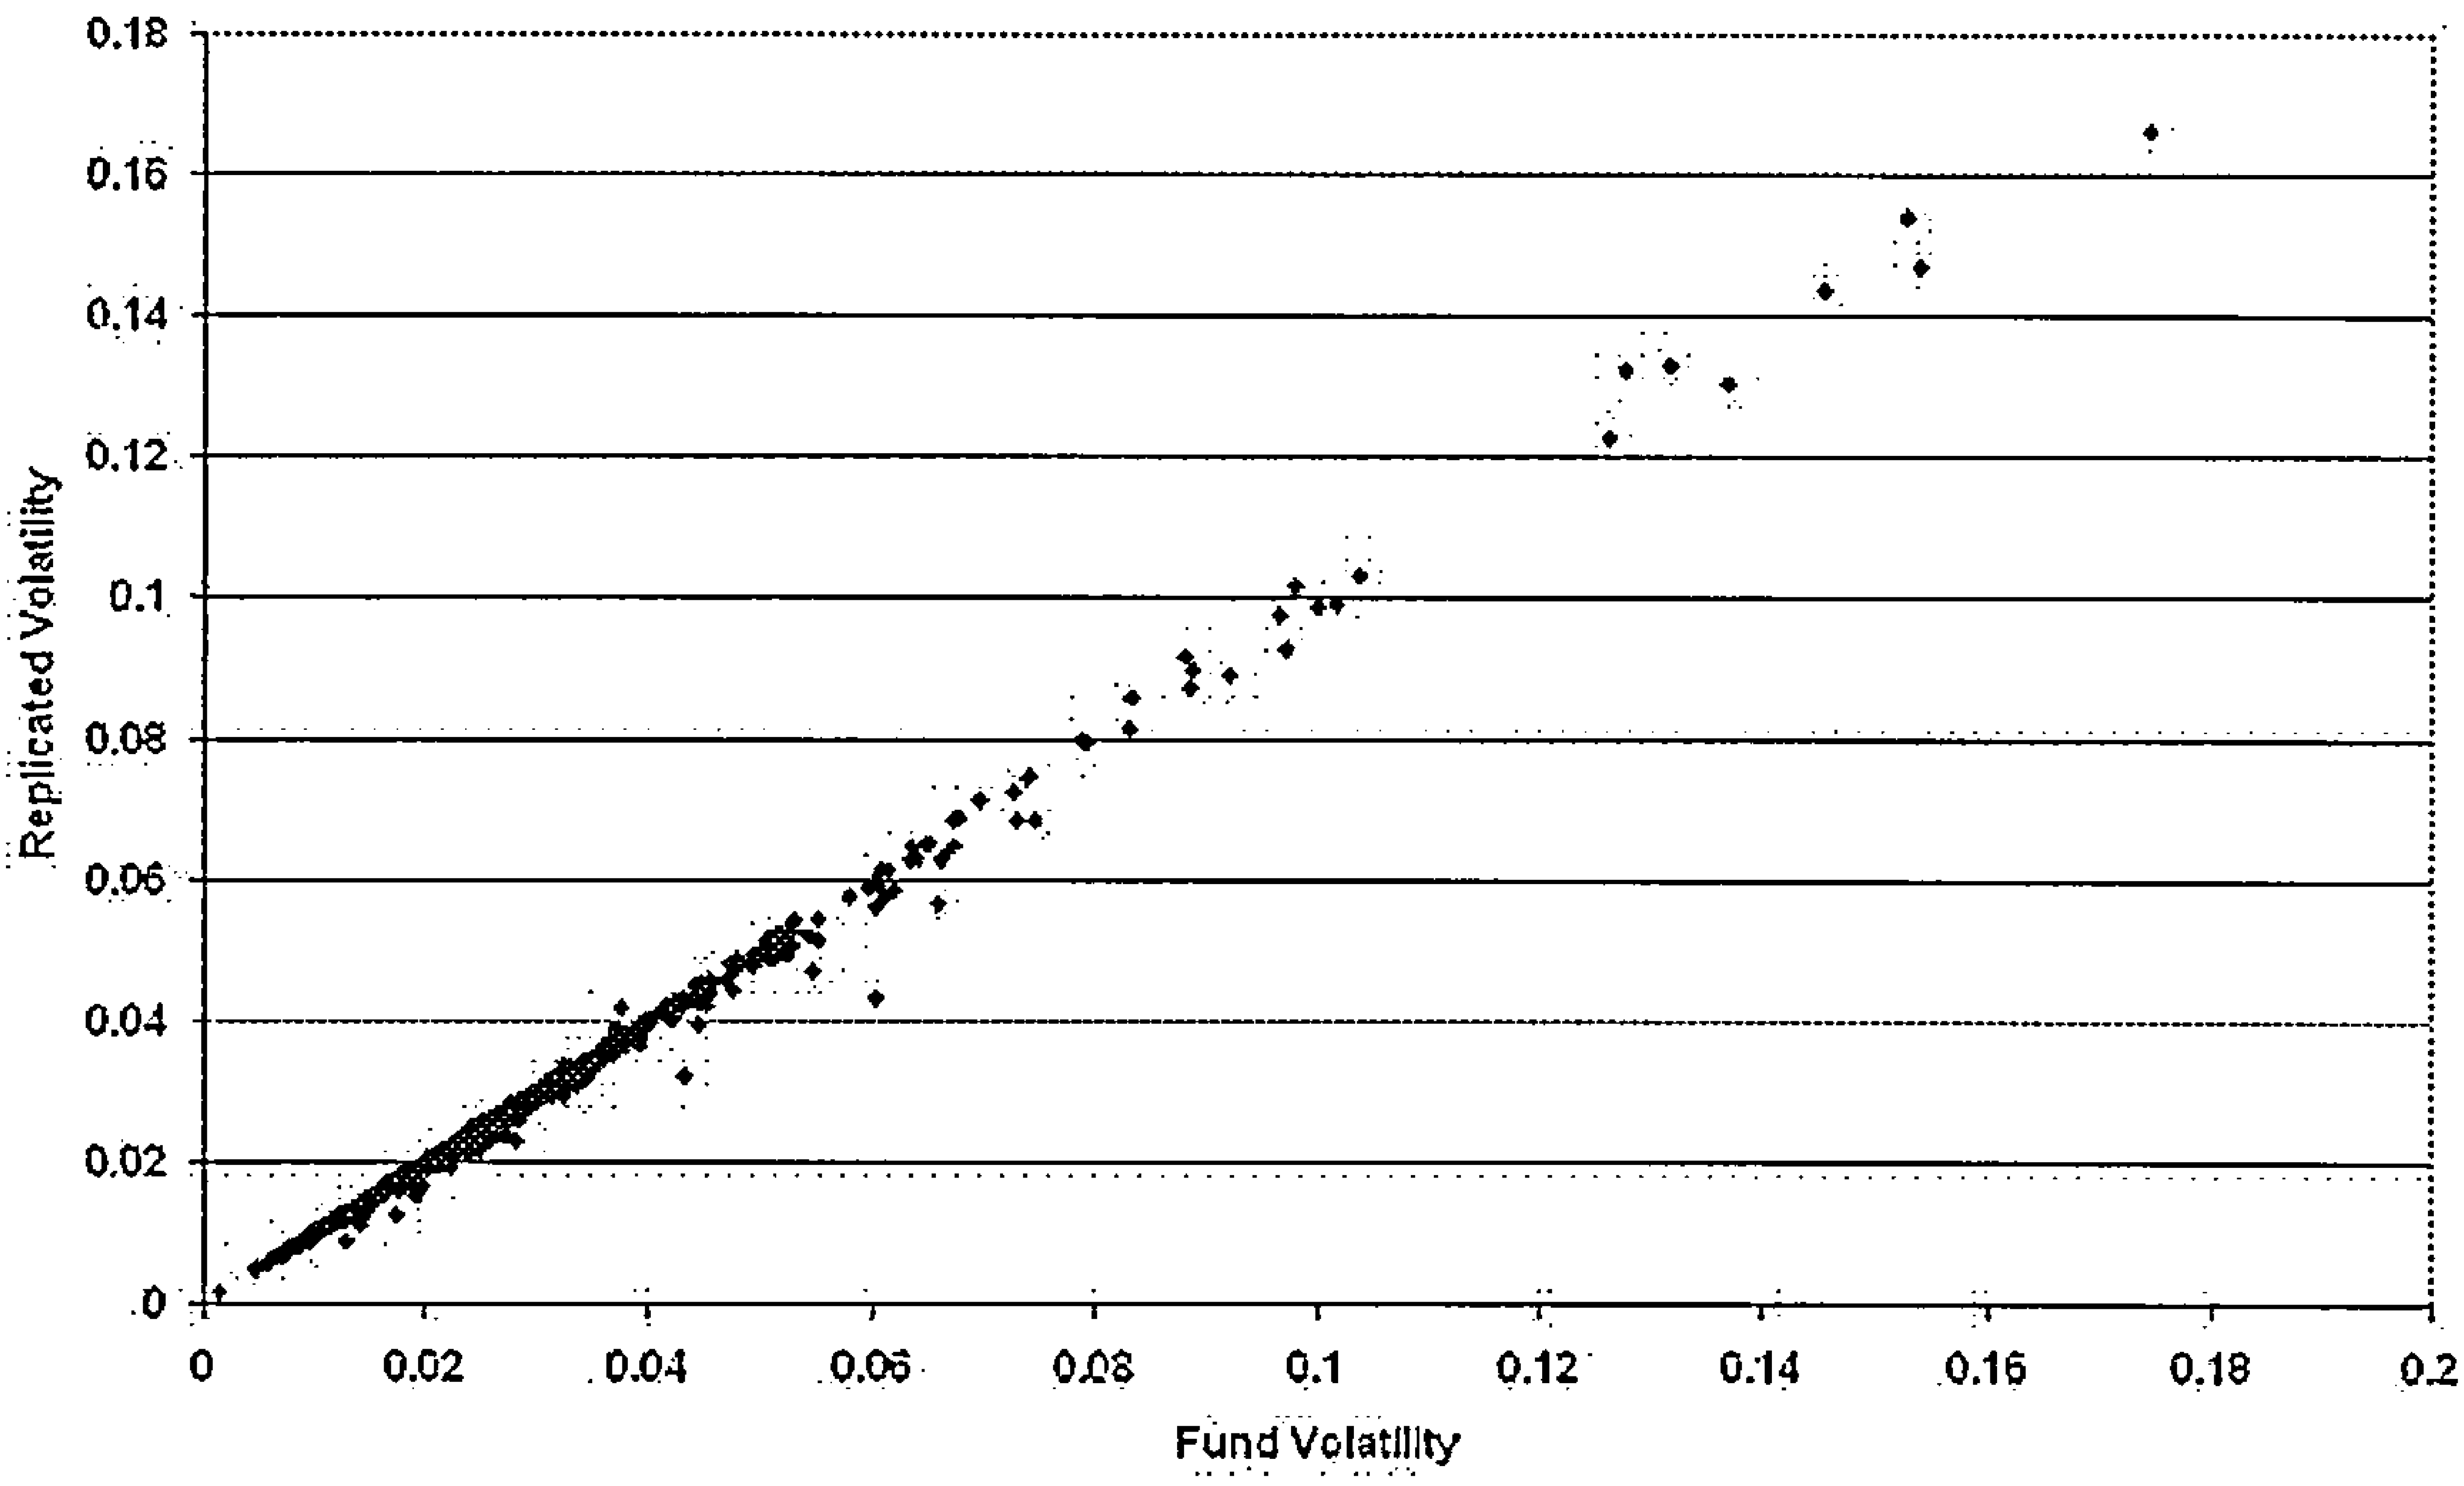

Figure 3.3. Scatter plot fund vs. replicated standard deviation.

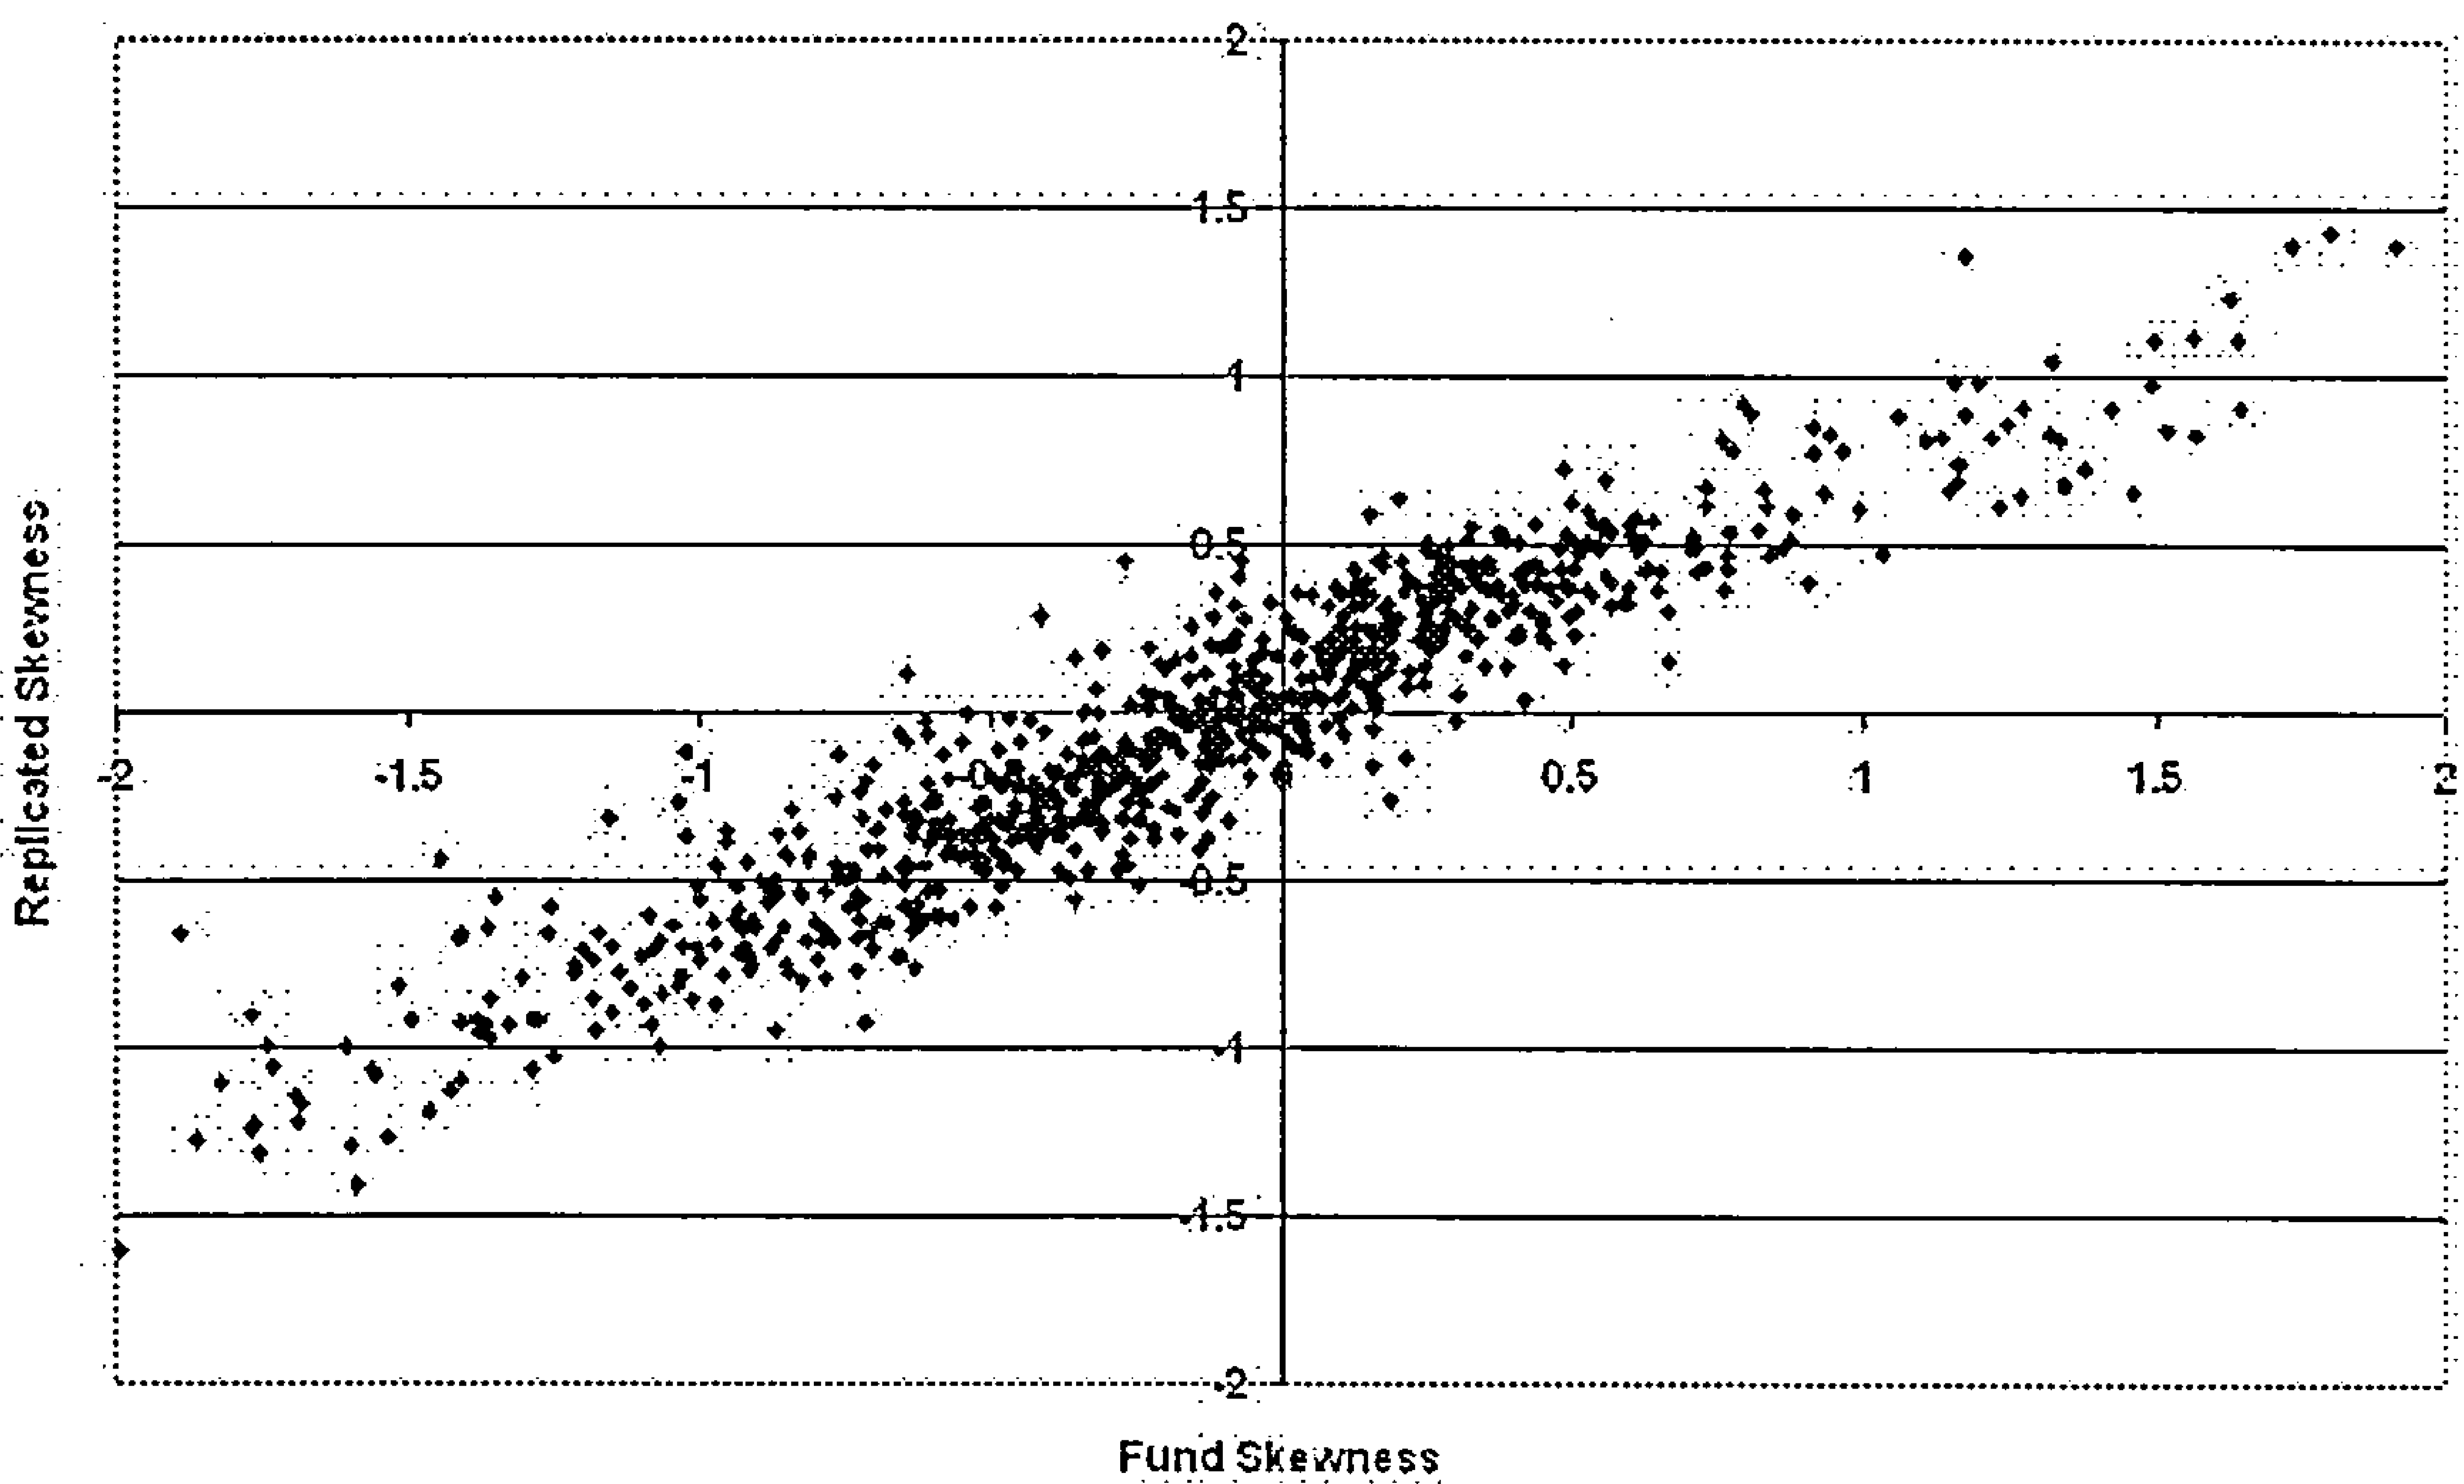

Figure 3.4. Scatter plot fund vs. replicated skewness.

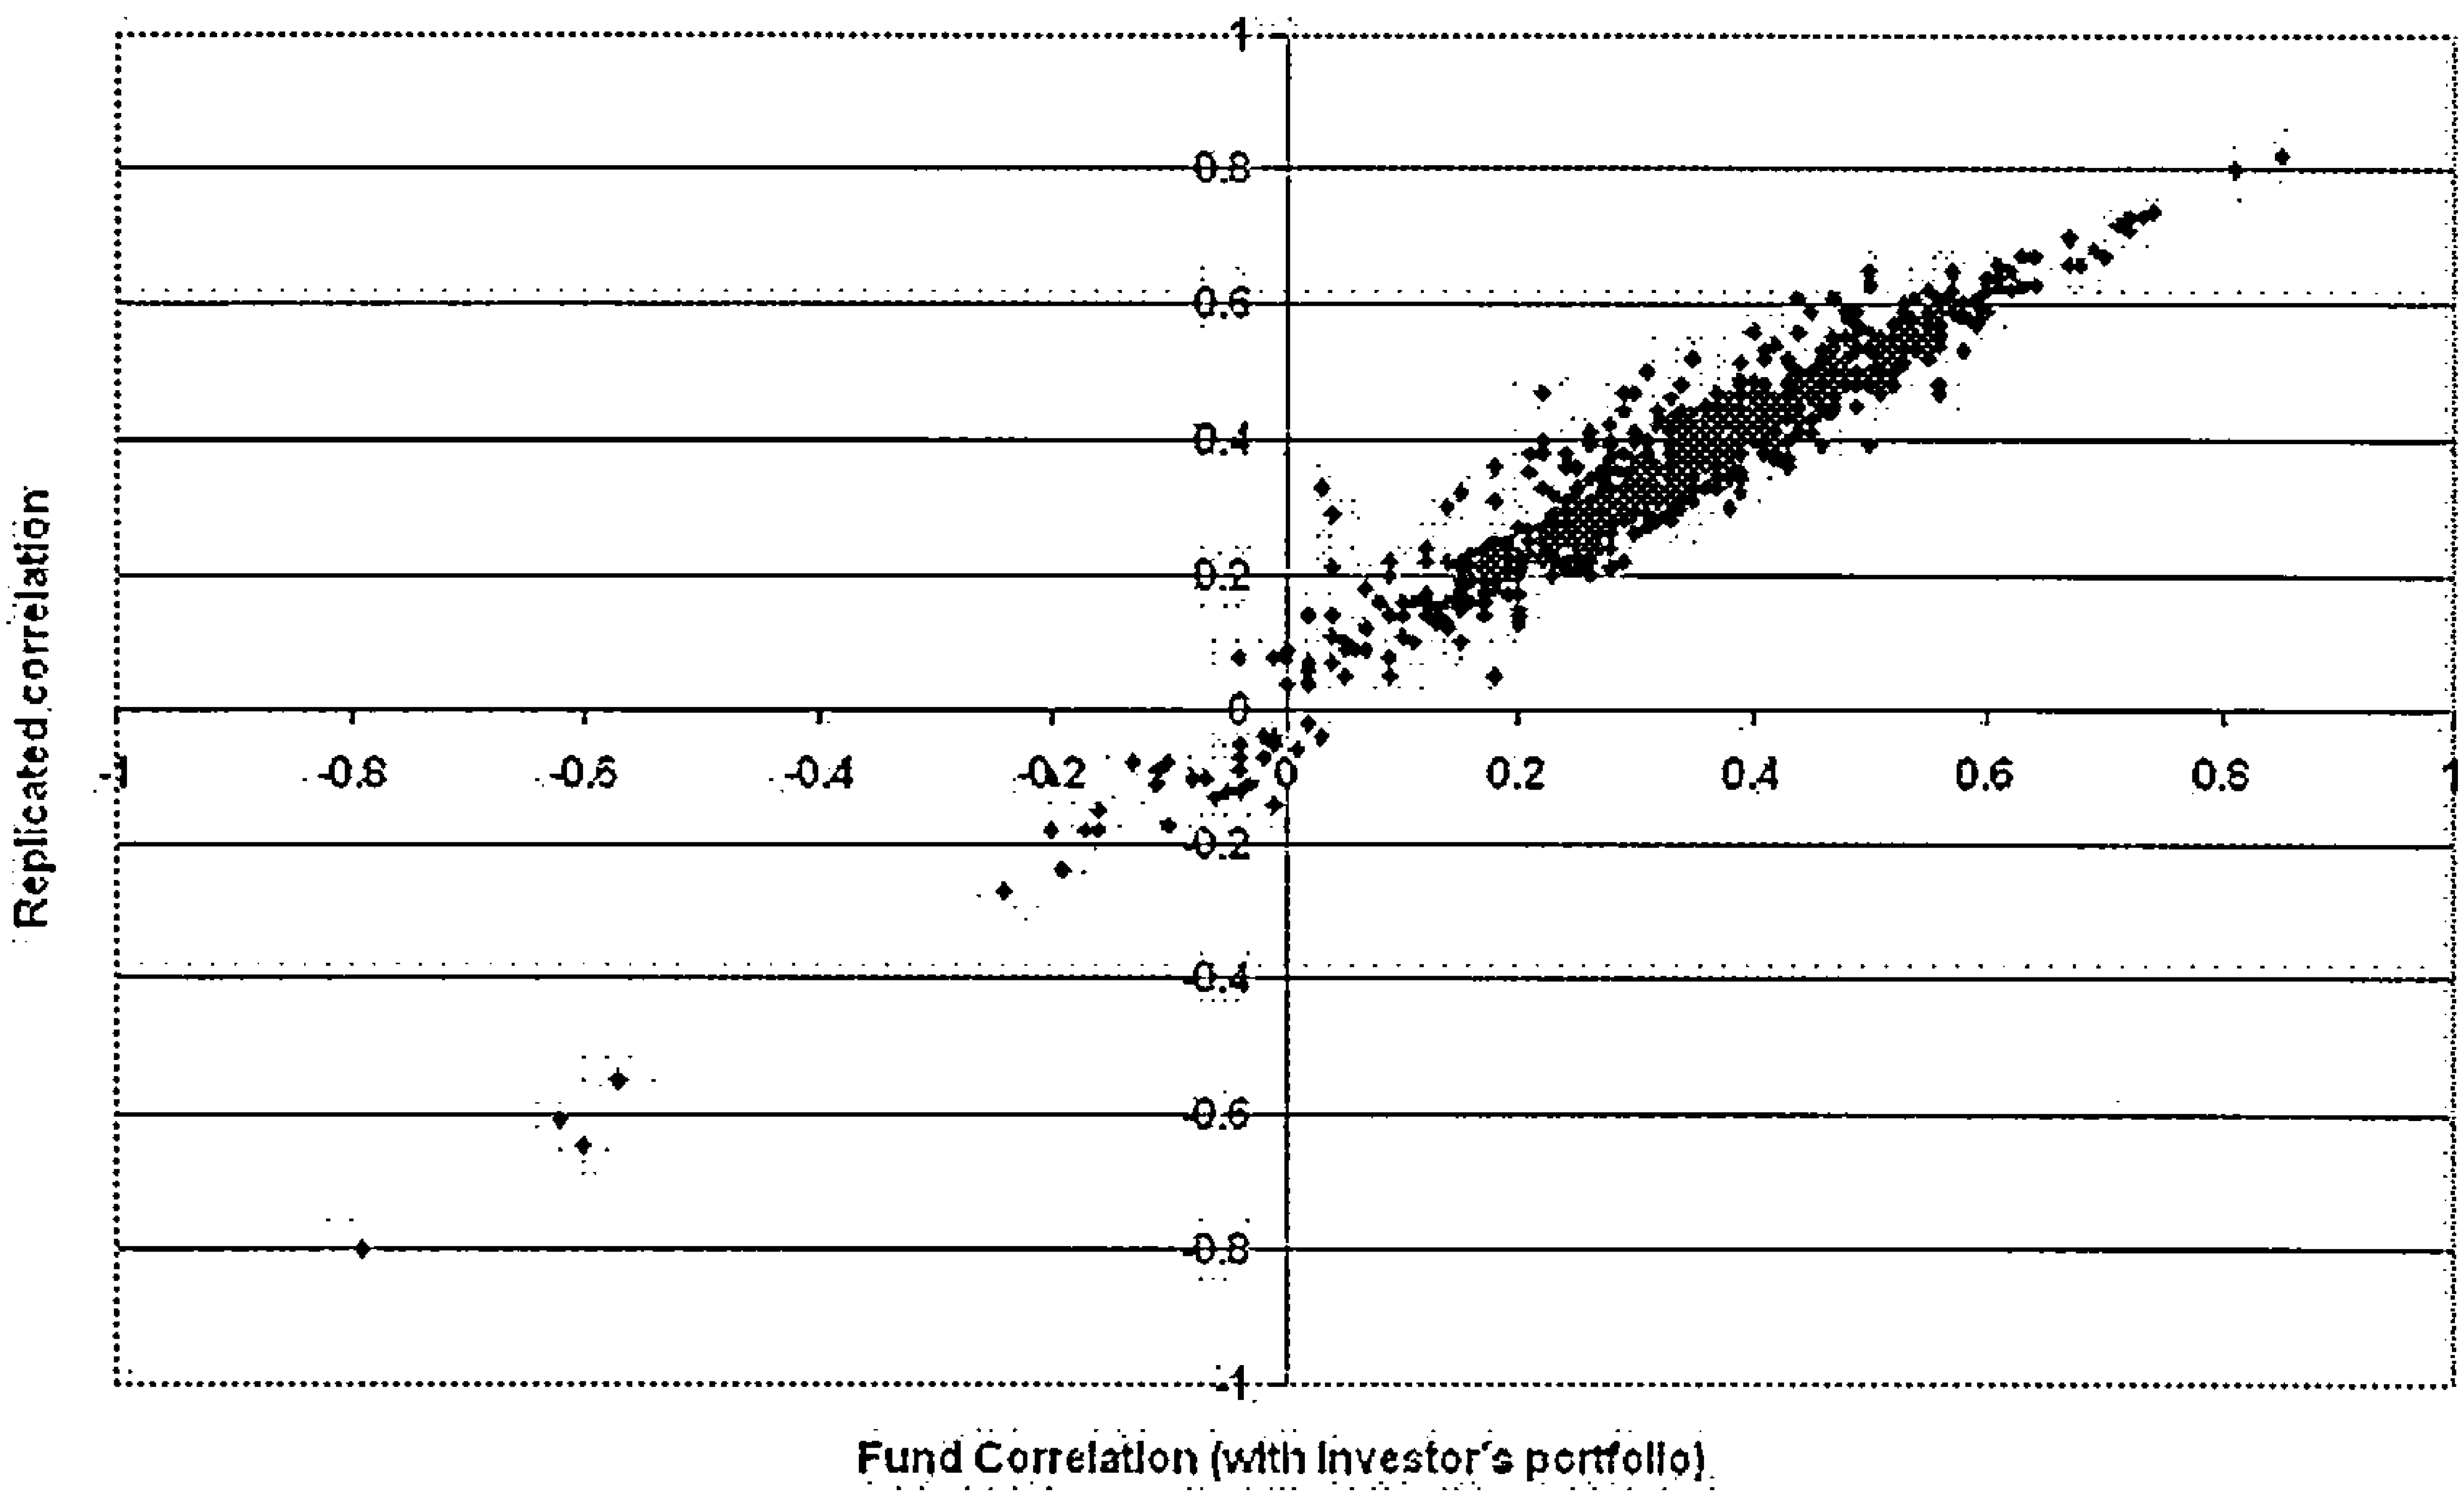

Figure 3.5. Scatter plot fund vs. replicated correlation with investor's portfolio.

To get an idea of the accuracy of the replication procedure, Figure 3.3 – 3.5 show scatter plots of the fund standard deviation (Fig. 3.3), standard skewness (Fig. 3.4) and correlation with the investor's portfolio (Fig. 3.5) versus the replicated values for all 875 funds. As is clear from these graphs, on average the replication of these parameters is unbiased and quite accurate. Not surprisingly, the replication of skewness can be difficult at times as fund returns may contain one or more outliers, which will have a major impact on the standard skewness statistic, but which cannot be replicated. We encountered this problem before in the example in Section 3.2.

Figures 3.3-3.5 also provide additional information on the risk-return profile of the funds in our sample. From Figure 3.4 for example, we see that for most funds of funds estimated skewness lies somewhere between –1 and +1. Likewise, from Figure 3.5 we see that the majority of funds are positively correlated with a portfolio of 50% stocks and 50% bonds. Most correlation coefficients lie between 0.2 and 0.6, indicating that many funds of funds' returns are a lot less 'market neutral' than is often suggested.[35]

[35] In this context it is important to note that at least for some of the more complex distributions encountered (see Table 3.5), the correlation coefficient will not be a particularly good measure of dependence and may underestimate the true level of dependence.

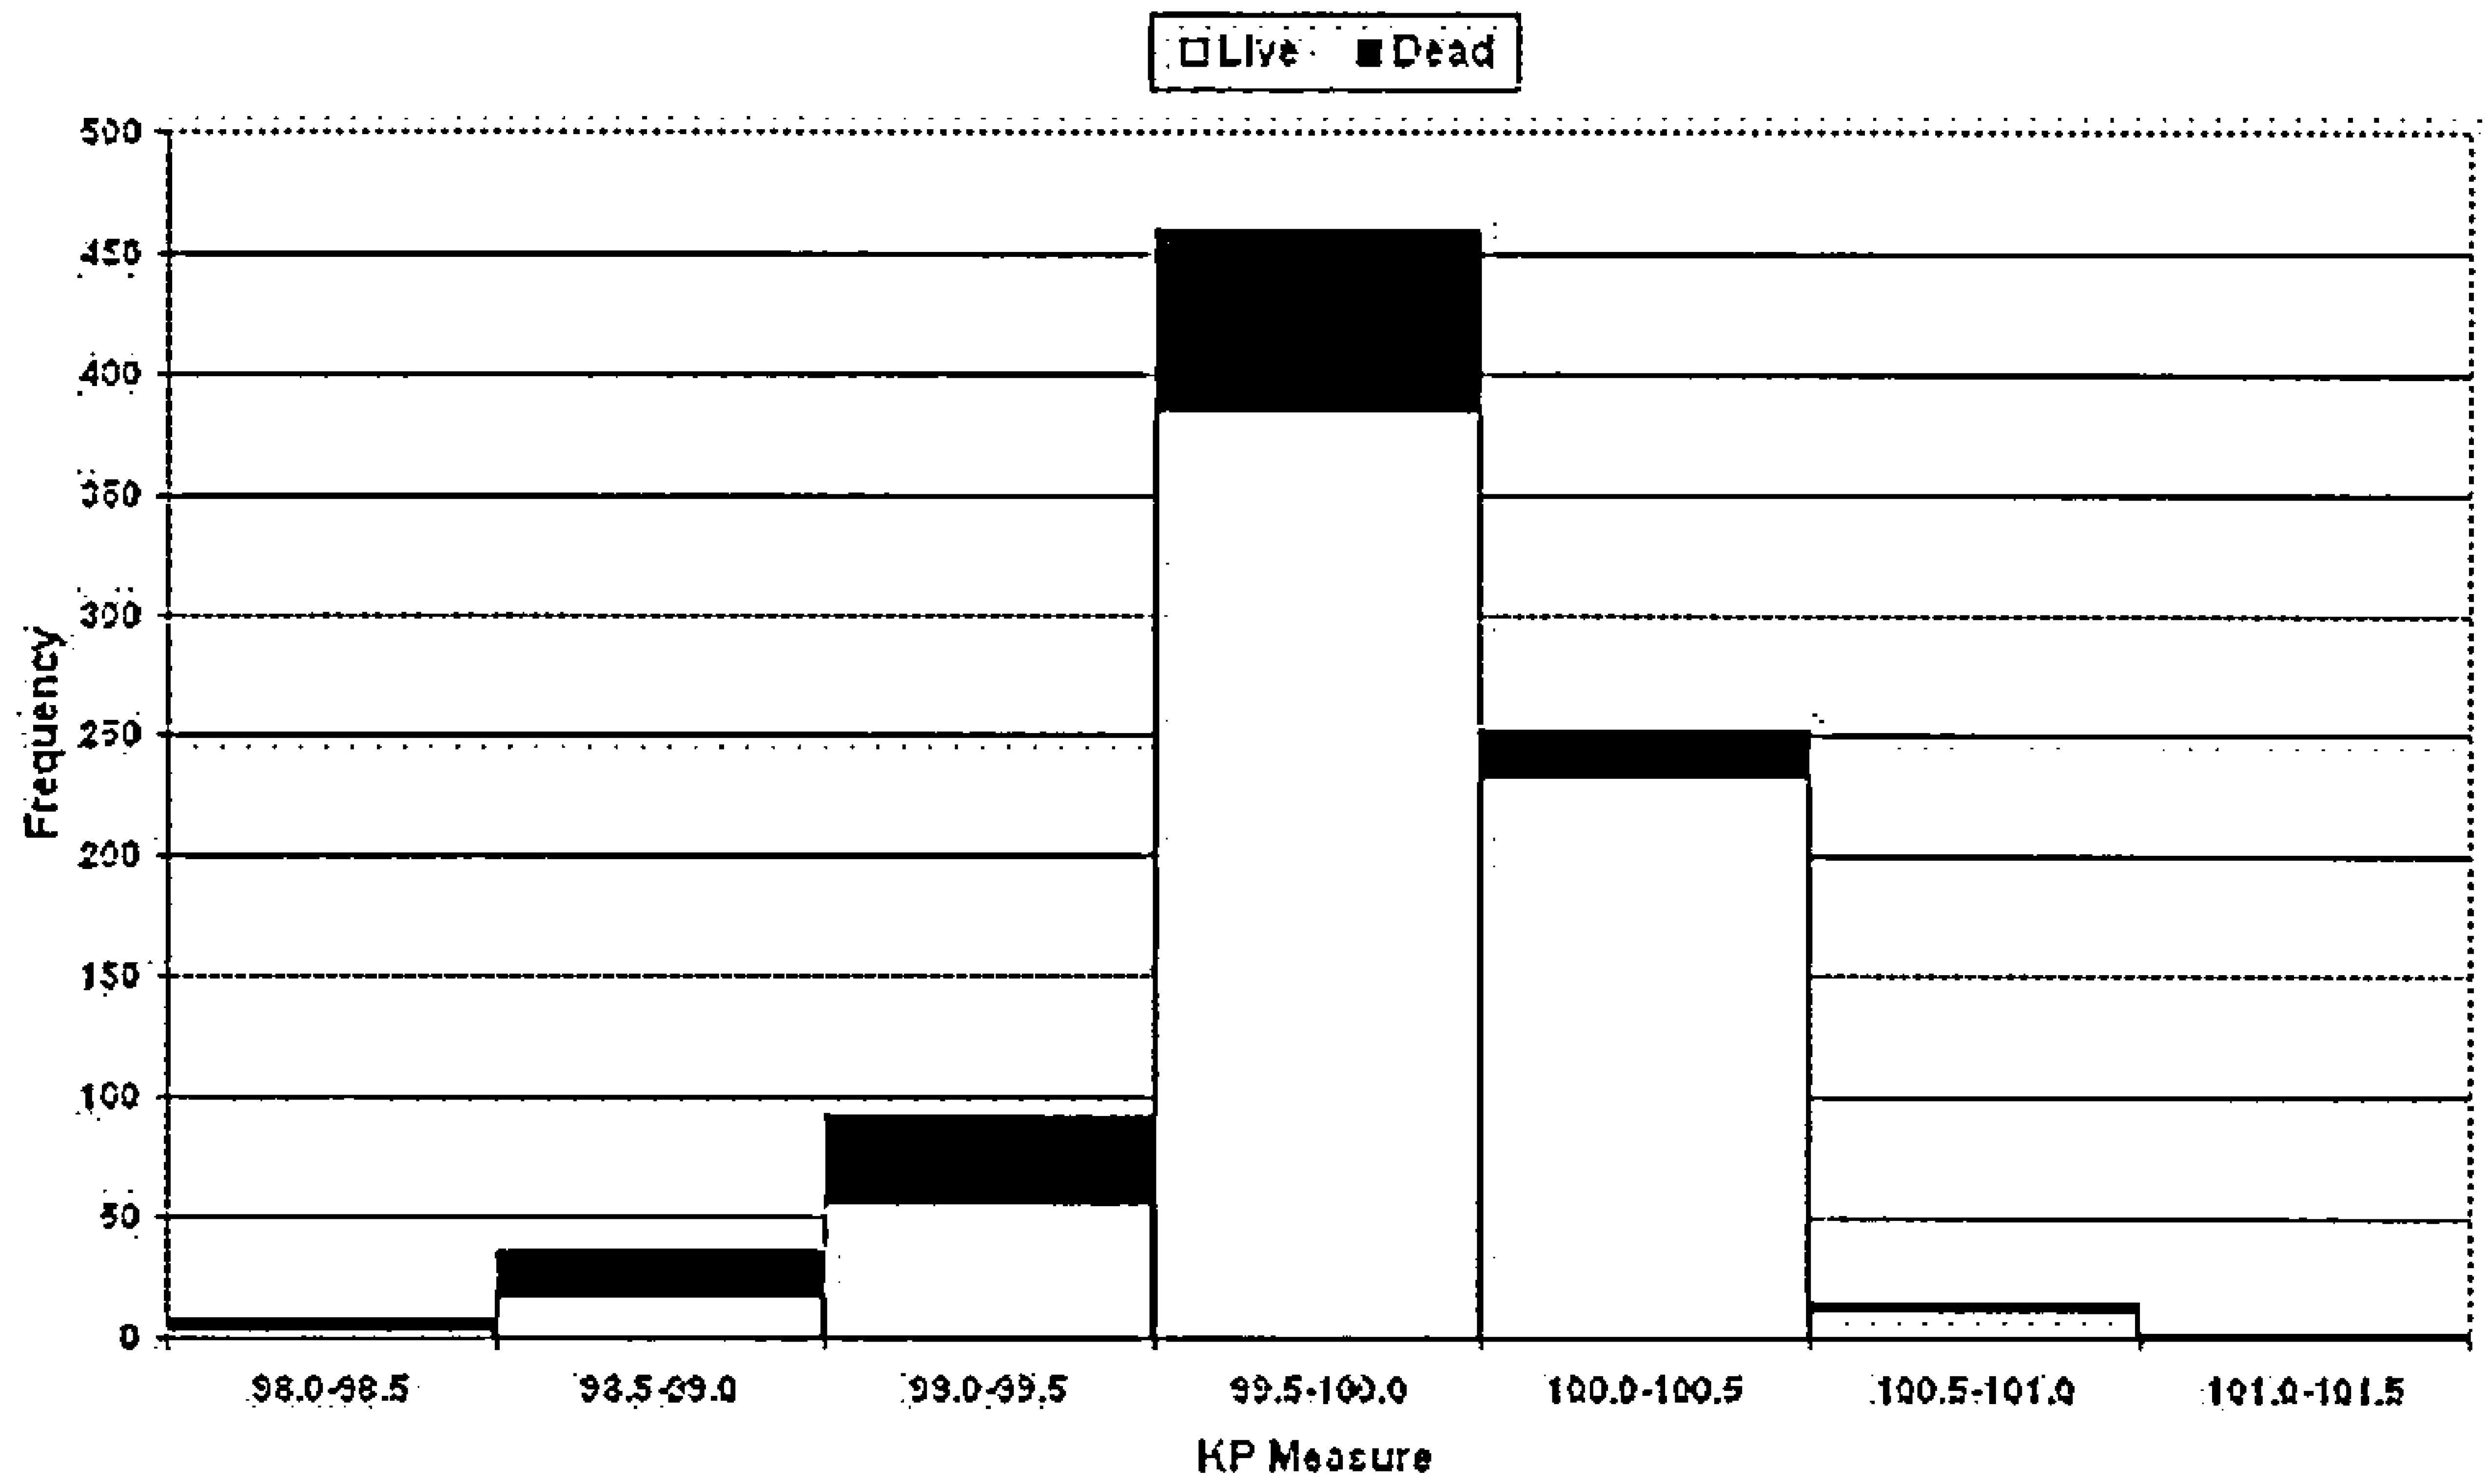

Figure 3.6. Histogram KP efficiency measure 875 funds of funds.

Figure 3.6 shows a histogram of the values of the KP efficiency measure obtained for the 875 funds of funds in our sample. From the graph we see that the majority of funds produce a value for the efficiency measure that is below 100. The average value for the KP efficiency measure over all 875 funds is 99.80. We tested the statistical significance of the above efficiency measure results by calculating bootstrapped 95% confidence intervals. Distinguishing between three cases, we obtained the following results:

- The confidence interval is entirely lower than 100 – **531 funds (60.7%).**
- The confidence interval contains 100 – **181 funds (20.7%).**
- The confidence interval is entirely higher than 100 – **163 funds (18.6%).**

This confirms that the majority of funds of hedge funds have not provided their investors with returns, which they could not have generated themselves in the futures market.

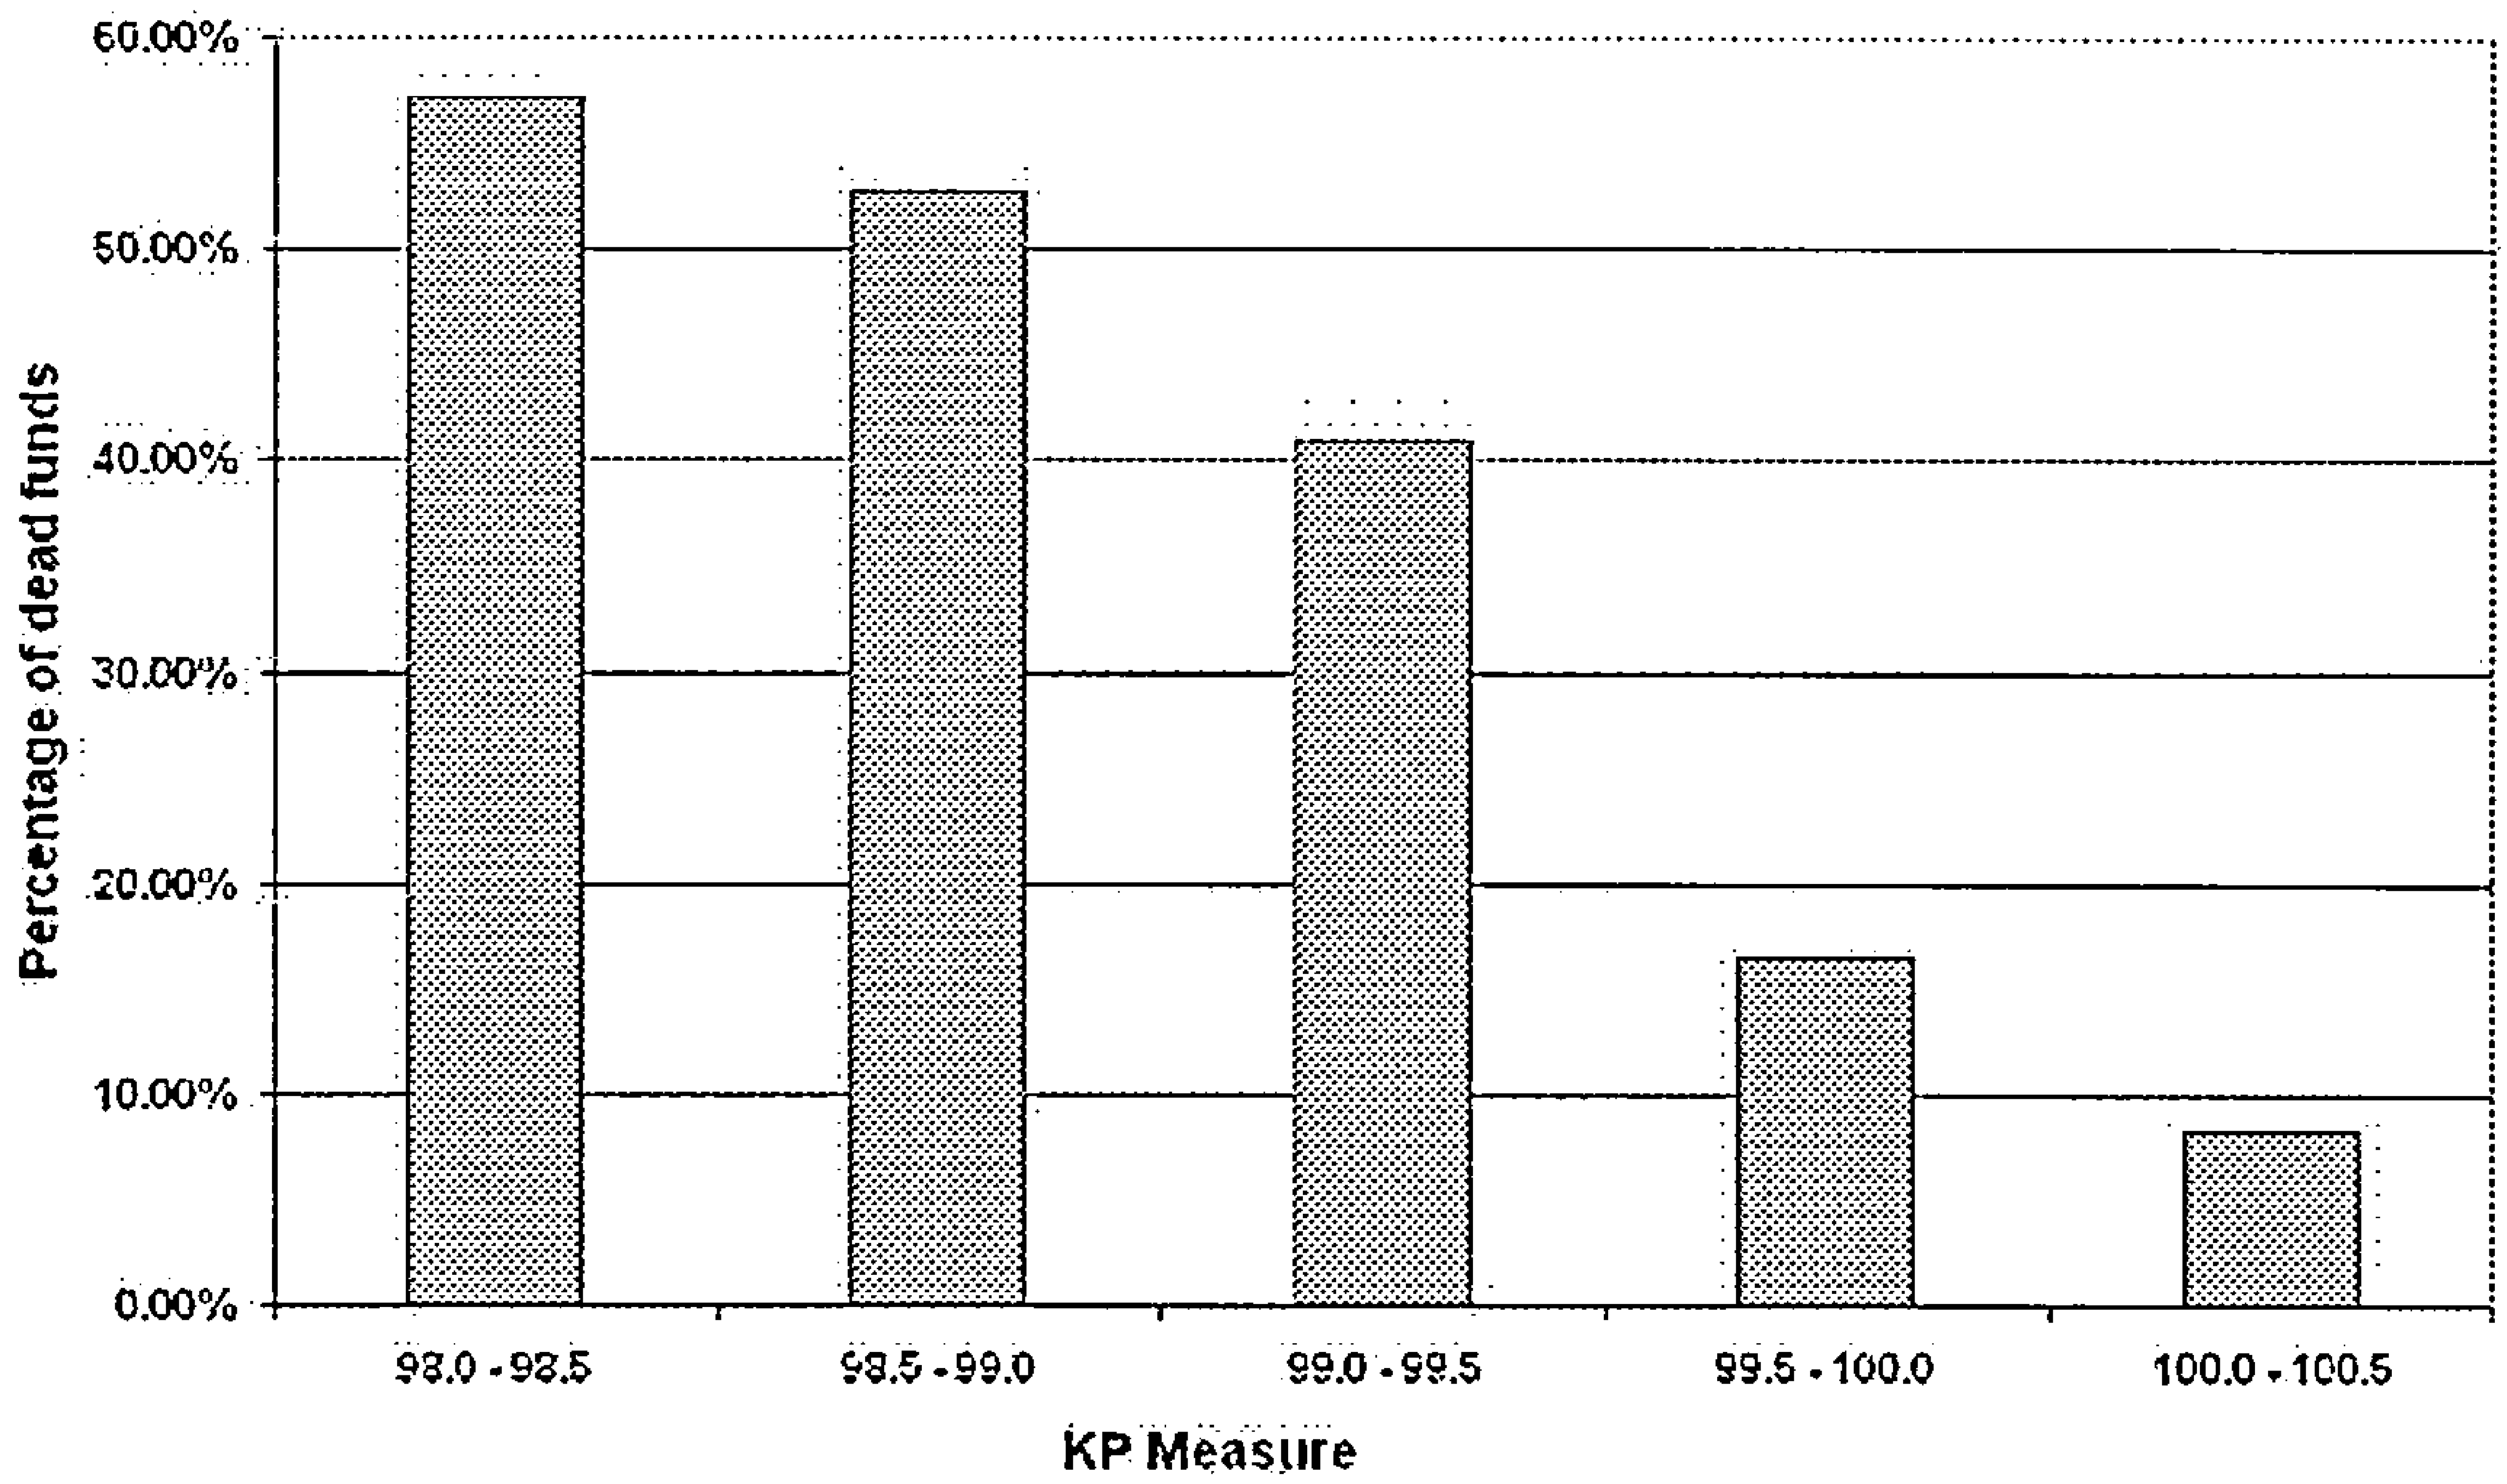

Figure 3.7. Percentage of funds of funds that stopped reporting before November 2006 as function KP efficiency measure.

Since lack of performance is one of the main reasons for funds to close down, Figure 3.7 shows the percentage of funds that stopped reporting to the database as a function of their KP measure. From the graph we see that there is a strong relationship. Out of the 48 funds with a KP measure below 99, no less than 24 (50%) stopped reporting. Out of the 273 funds with a KP measure higher than 100, only 27 (9.9%) did so. A similar relationship is observed in the average KP measures of live and dead funds. The average KP measure over the 218 dead funds is 99.64, while over the 657 funds still alive the average is 99.86.

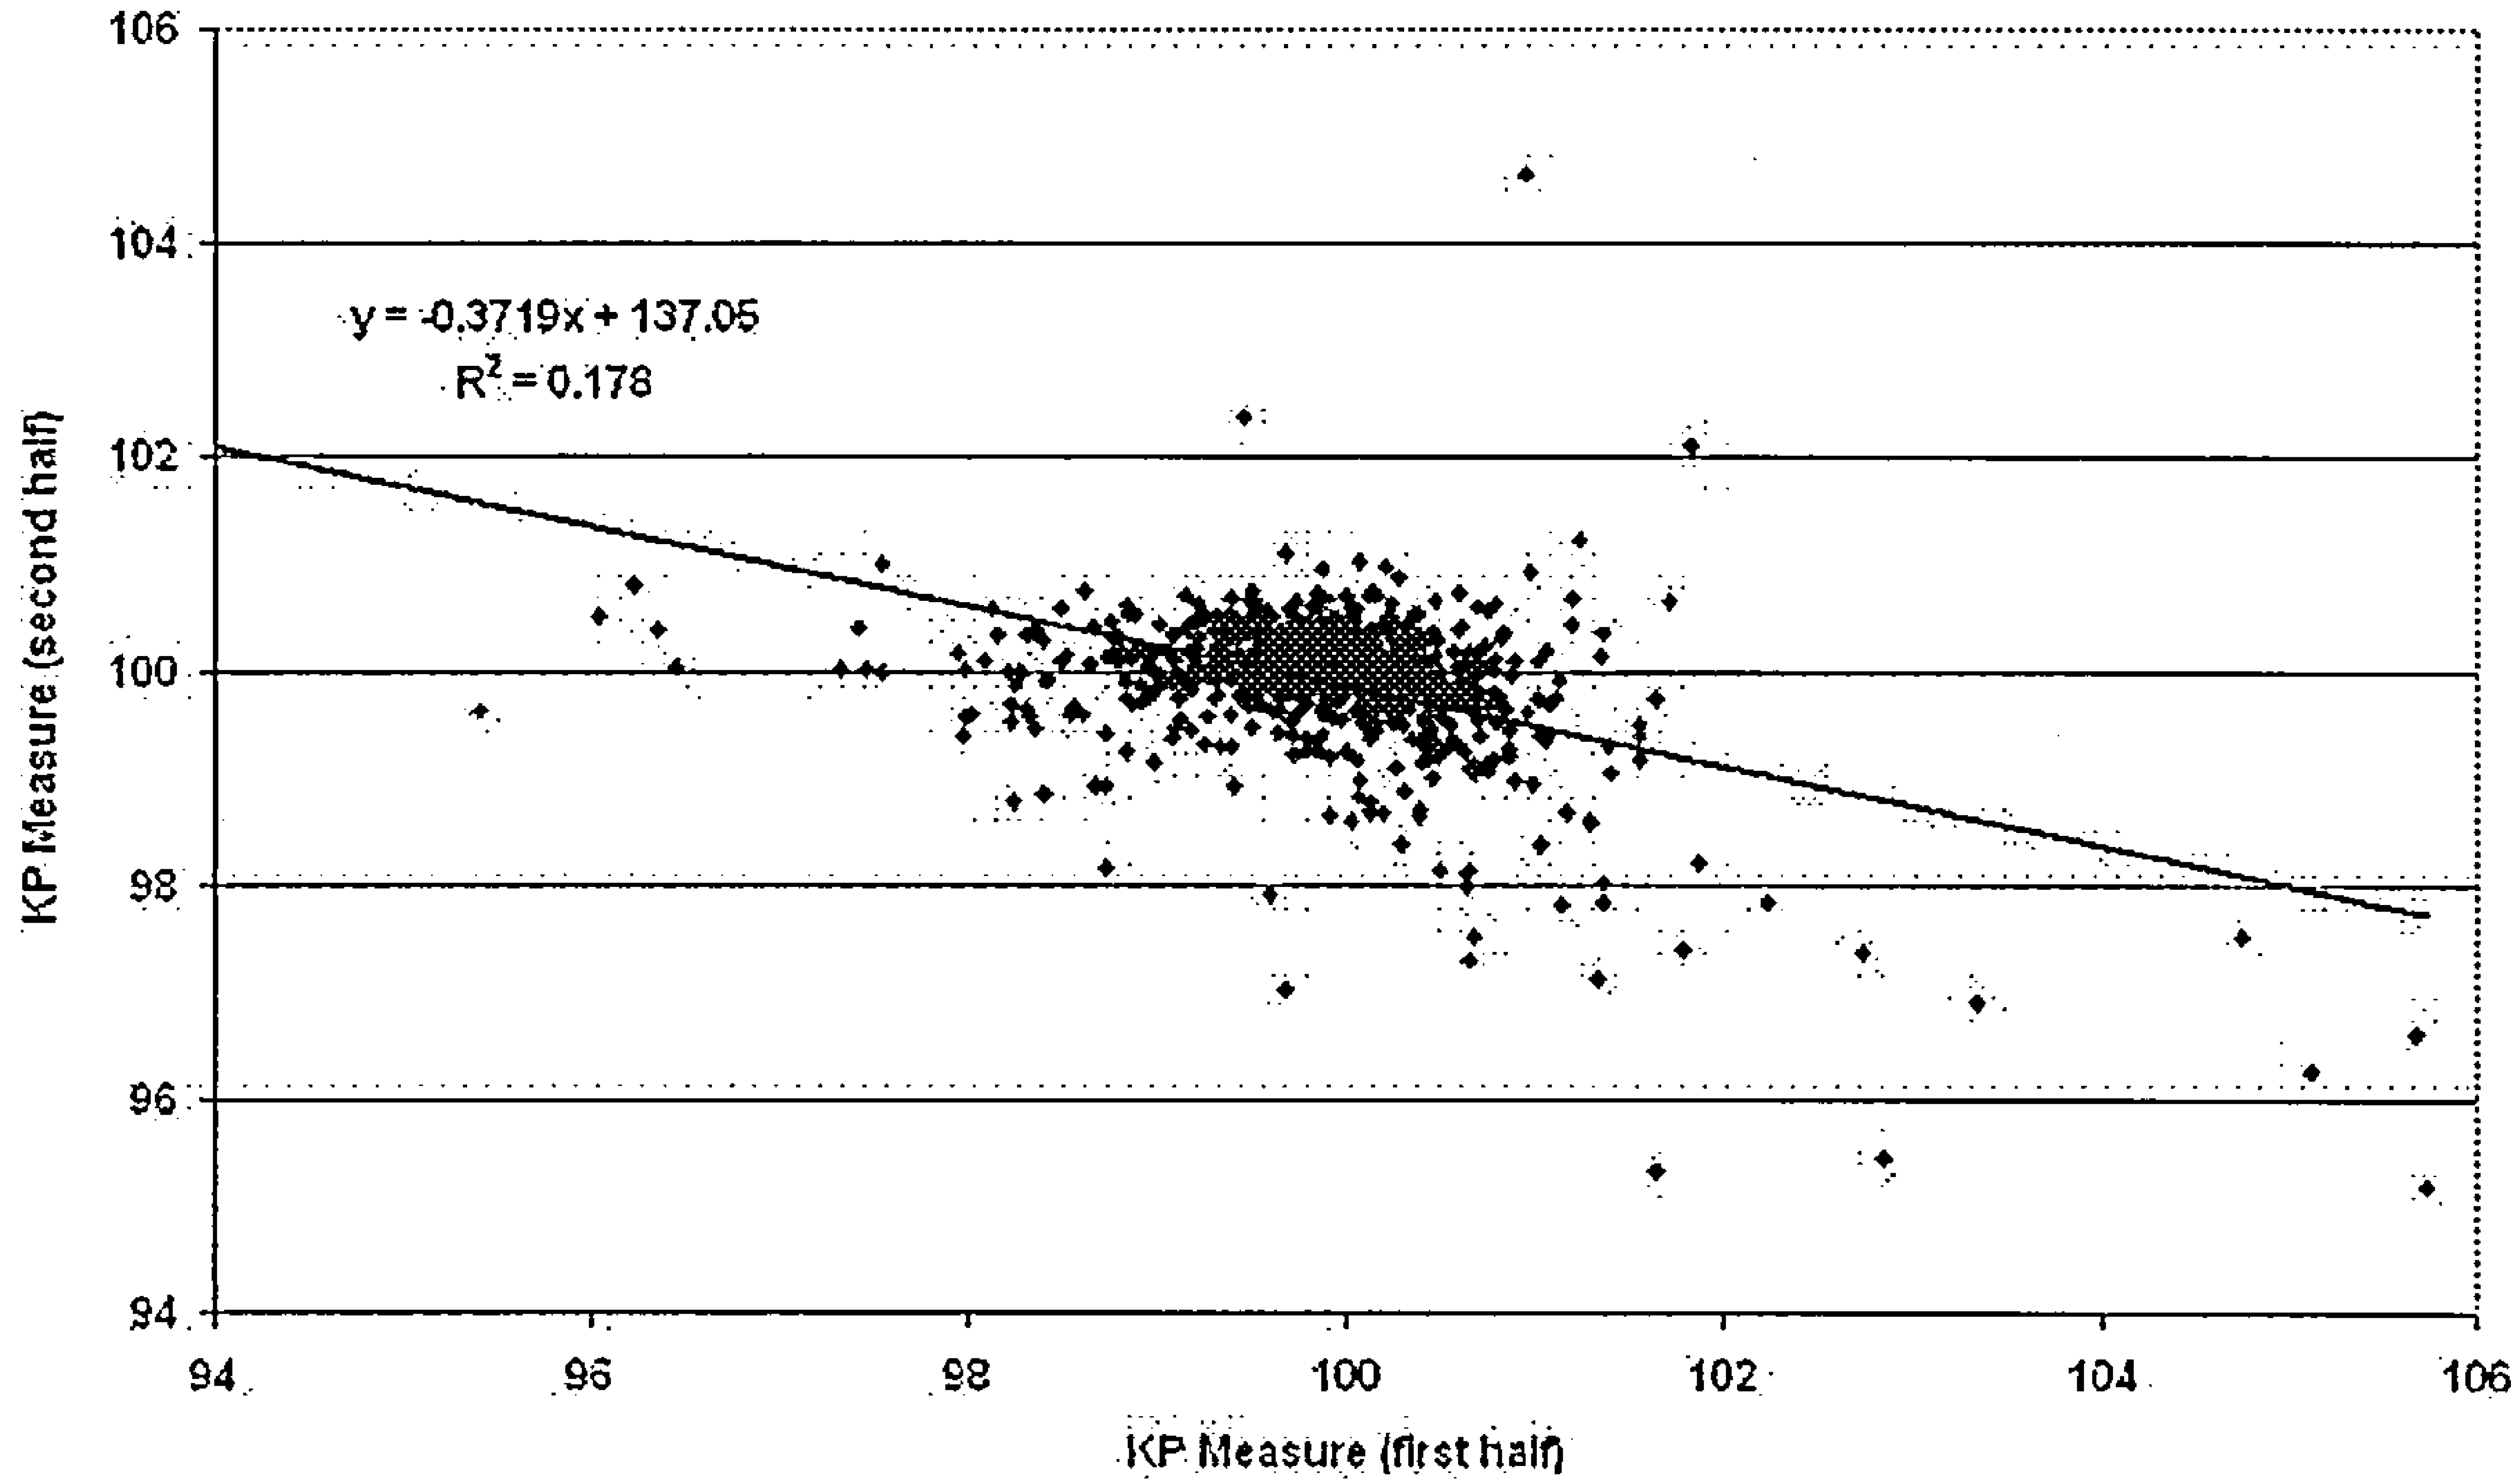

Figure 3.8. KP efficiency measures over split track record.

Another question concerns the performance of funds of funds through time. Especially after two years of somewhat disappointing results, it is often claimed that overall hedge fund performance is deteriorating, with the massive inflow of capital over recent years being the most obvious cause. We therefore split the track record of all funds of funds with 8 or more years of history in two equal parts and calculated the KP measure over each part. The result is plotted in Figure 3.8.

The average KP measure over the first period is 99.91, while over the second, more recent, period the average is 99.89. This suggests that on average fund of funds performance has changed little over time. This is not entirely true though as Figure 3.8 reveals a clear tendency for funds with a relatively high (low) KP measure in the first period to produce a relatively low (high) KP measure in the second. As the assets under management of funds that do well can be expected to grow substantially (organically as well as through additional inflows) and vice versa, this finding supports the idea that increased fund size has a negative impact on future performance.

3.4 Individual Hedge Funds

3.4.1 Individual Hedge Funds: Data Description

Our individual hedge funds sample consists of 2073 funds with a minimum of 4 years of history available. As before, all data were obtained from The Barclay Group as per November 2006 and funds denominated in another currency than USD are converted to USD. We study all funds together, as well as the various strategy groups separately, so we can detect possible differences between them. The strategy classification used and the number of funds within each strategy group can be found in Table 3.6. Table 3.7 and 3.8 provide some information on the start and end dates of the track records of the funds in our sample.

Category	No. funds
Convertible Arbitrage	80
Distressed Securities	55
Emerging Markets	182
Equity Long	309
Equity Long/Short	411
Equity Market Neutral	65
Equity Short Bias	20
Event Driven	96
Fixed Income Arbitrage	149
Global Macro	80
Merger Arbitrage	55
Multi-Strategy Arbitrage	104
Mutual Fund Timing	34
Other	242
Sector Funds	165
Statistical Arbitrage	26
Total	2073

Table 3.6. Strategy classification and number of hedge funds.

Start after:	Jan 1985	Jan 1988	Jan 1991	Jan 1994	Jan 1997	Jan 2000	Jan 2003
Convertible Arbitrage	79	79	76	69	50	24	0
Distressed Securities	55	55	51	48	36	24	0
Emerging Markets	181	178	173	159	114	53	0
Equity Long	306	298	279	251	186	96	0
Equity Long/Short	409	403	388	351	267	153	0
Equity Market Neutral	64	63	57	53	48	22	0
Equity Short Bias	20	20	18	14	9	2	0
Event Driven	94	92	85	73	50	31	0
Fixed Income Arb.	149	148	145	135	108	55	0
Global Macro	78	73	72	65	53	27	0
Merger Arbitrage	53	51	49	46	31	12	0
Multi-Strategy Arb.	104	104	101	91	66	35	0
Mutual Fund Timing	34	34	33	33	27	6	0
Other	239	235	224	183	82	30	0
Sector Funds	164	161	154	136	108	64	0
Statistical Arbitrage	25	25	25	22	14	6	0
Total	2054	2019	1930	1729	1249	640	0

Table 3.7. Hedge funds start date details.

End before:	Jan 2000	Jan 2003	Jan 2006	Nov 2006
Convertible Arbitrage	0	5	25	35
Distressed Securities	0	6	10	16
Emerging Markets	1	18	37	65
Equity Long	0	30	65	110
Equity Long/Short	0	57	138	183
Equity Market Neutral	0	3	20	27
Equity Short Bias	0	2	9	10
Event Driven	0	6	26	40
Fixed Income Arb.	0	21	48	62
Global Macro	0	12	23	27
Merger Arbitrage	0	8	25	29
Multi-Strategy Arb	0	8	28	43
Mutual Fund Timing	0	3	22	25
Other	0	192	200	207
Sector Funds	0	24	57	75
Statistical Arbitrage	0	6	12	14
Total	1	401	745	968

Table 3.8. Hedge funds end date details.

Table 3.7 shows that the majority of funds started after 1994. Only 143 funds started before January 1991, and only 19 before January 1985. Table 3.8 shows that out of 2073 funds, no less than 968 funds stopped reporting before November 2006. This confirms that the attrition rate in hedge funds is very substantial[36].

Length Track record:	4-5Y	5-6Y	6-7Y	7-8Y	8-9Y	9-10Y	10-12Y	12-14Y	14Y+
Convertible Arbitrage	22	19	15	9	3	4	5	2	1
Distressed Securities	9	11	7	7	5	4	5	4	3
Emerging Markets	50	28	27	23	8	12	16	6	12
Equity Long	63	62	52	31	29	23	24	18	7
Equity Long/Short	82	69	62	34	37	24	49	29	25
Equity Market Neutral	8	4	12	11	7	8	12	2	1
Equity Short Bias	4	1	2	3	2	3	1	3	1
Event Driven	17	12	14	8	6	9	7	7	16
Fixed Income Arbitrage	33	22	22	14	13	9	12	7	17
Global Macro	13	12	8	11	4	6	7	7	12
Merger Arbitrage	17	10	6	7	2	6	4	3	0
Multi-Strategy Arbitrage	24	7	16	11	7	6	19	6	8
Mutual Fund Timing	4	12	7	3	4	1	2	1	0
Other	40	37	32	32	21	10	35	17	18
Sector Funds	35	18	20	15	15	10	21	13	18
Statistical Arbitrage	5	4	4	1	1	2	3	2	4
Total	426	328	306	220	164	137	222	127	143

Table 3.9. Length of hedge funds track records.

Table 3.9 provides details on the length of the available hedge fund track records. Out of the 2073 funds in the sample, only 492 have more than 10 years of history. This again reflects the fact that most funds are still relatively young and attrition levels can be very significant.

3.4.2 Individual Hedge Funds: Distributional Analysis

A crucial stage in the evaluation procedure is the proper modelling of the distributional characteristics of the fund, the investor's portfolio and the reserve asset. This means that, although not explicitly designed to do so, the evaluations provide a wealth of

[36] Not all funds that stop reporting into a database do so because they close down. The majority does so, however. For more details on hedge fund and fund of funds attrition see Kat and Amin (2003a).

information on the distributional properties of fund of funds returns. Table 3.10 summarizes how often a given marginal distribution or copula was used in the evaluations for modelling the fund return marginal and the joint distribution of the fund and the investor's portfolio return.

	Marginal Distribution			Copula					
	Norm	Stud	John	Norm	Stud	Gumb	SJC	Cook-J	Frank
Convertible Arbitrage	46.3%	45.8%	7.9%	39.4%	15.0%	9.4%	10.0%	11.3%	15.0%
Distressed Securities	47.3%	46.1%	6.7%	33.6%	20.0%	10.0%	11.8%	13.6%	10.9%
Emerging Markets	45.6%	47.3%	7.1%	43.4%	15.4%	7.4%	9.6%	11.8%	12.4%
Equity Long	50.3%	44.2%	5.5%	40.3%	20.2%	8.9%	8.3%	7.9%	14.4%
Equity Long/Short	51.8%	43.4%	4.8%	36.0%	17.6%	13.4%	10.3%	7.4%	15.2%
Equity Market Neutral	50.3%	46.7%	3.1%	33.1%	25.4%	10.8%	10.8%	10.8%	9.2%
Equity Short Bias	45.0%	50.0%	5.0%	50.0%	17.5%	7.5%	5.0%	2.5%	17.5%
Event Driven	41.7%	51.4%	6.9%	28.6%	28.6%	7.8%	9.9%	13.0%	12.0%
Fixed Income Arb	48.3%	45.6%	6.0%	36.6%	17.4%	14.1%	11.7%	9.4%	10.7%
Global Macro	50.0%	42.1%	7.9%	31.3%	20.6%	13.1%	11.9%	10.6%	12.5%
Merger Arbitrage	46.7%	43.6%	9.7%	43.6%	17.3%	3.6%	10.9%	20.0%	4.5%
Multi-Strategy Arb	43.9%	48.1%	8.0%	41.3%	15.9%	7.7%	10.6%	13.9%	10.6%
Mutual Fund Timing	50.0%	41.2%	8.8%	45.6%	8.8%	4.4%	5.9%	7.4%	27.9%
Other	63.1%	30.9%	6.1%	49.0%	8.5%	4.1%	9.5%	15.1%	13.8%
Sector Funds	48.1%	47.5%	4.4%	35.2%	18.2%	14.2%	10.6%	7.3%	14.5%
Statistical Arbitrage	53.8%	44.9%	1.3%	51.9%	13.5%	11.5%	7.7%	5.8%	9.6%
Total	50.3%	43.8%	5.9%	39.2%	17.3%	9.9%	9.9%	10.3%	13.4%

Table 3.10: Overall distributional characteristics hedge fund returns.

Table 3.10 confirms that the majority of individual hedge fund returns are far from normally distributed. Out of 2073 funds, 1030 (49.7%) funds' marginal return is better modelled by a Student-t or Johnson distribution than a normal distribution. In addition, for only 813 (39.2%) of the 2073 funds is the relationship with the investor's portfolio of 50% S&P 500 and 50% T-bonds best modelled by the normal copula. This emphasizes once more how important it is to evaluate hedge fund performance using a methodology, which does not rely on the assumption of normally distributed returns.

3.4.3 Individual Hedge Funds: Replication and Evaluation Results

For the evaluation we make the same assumptions as before. The representative investor's portfolio consists of 50% S&P 500 and 50% long-dated US Treasury bonds and the reserve asset of nearby Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures. Transaction costs on all futures contracts are assumed to be 1bp one-way.

Figure 3.9-3.11 plot the standard deviation, standard skewness measure and correlation with the investor's portfolio for all 2073 funds. As is clear from these graphs, on average the replication of these parameters is unbiased and quite accurate. Taking into account the high sensitivity of the conventional skewness measure, the results in Figure 3.10 are quite satisfactory.

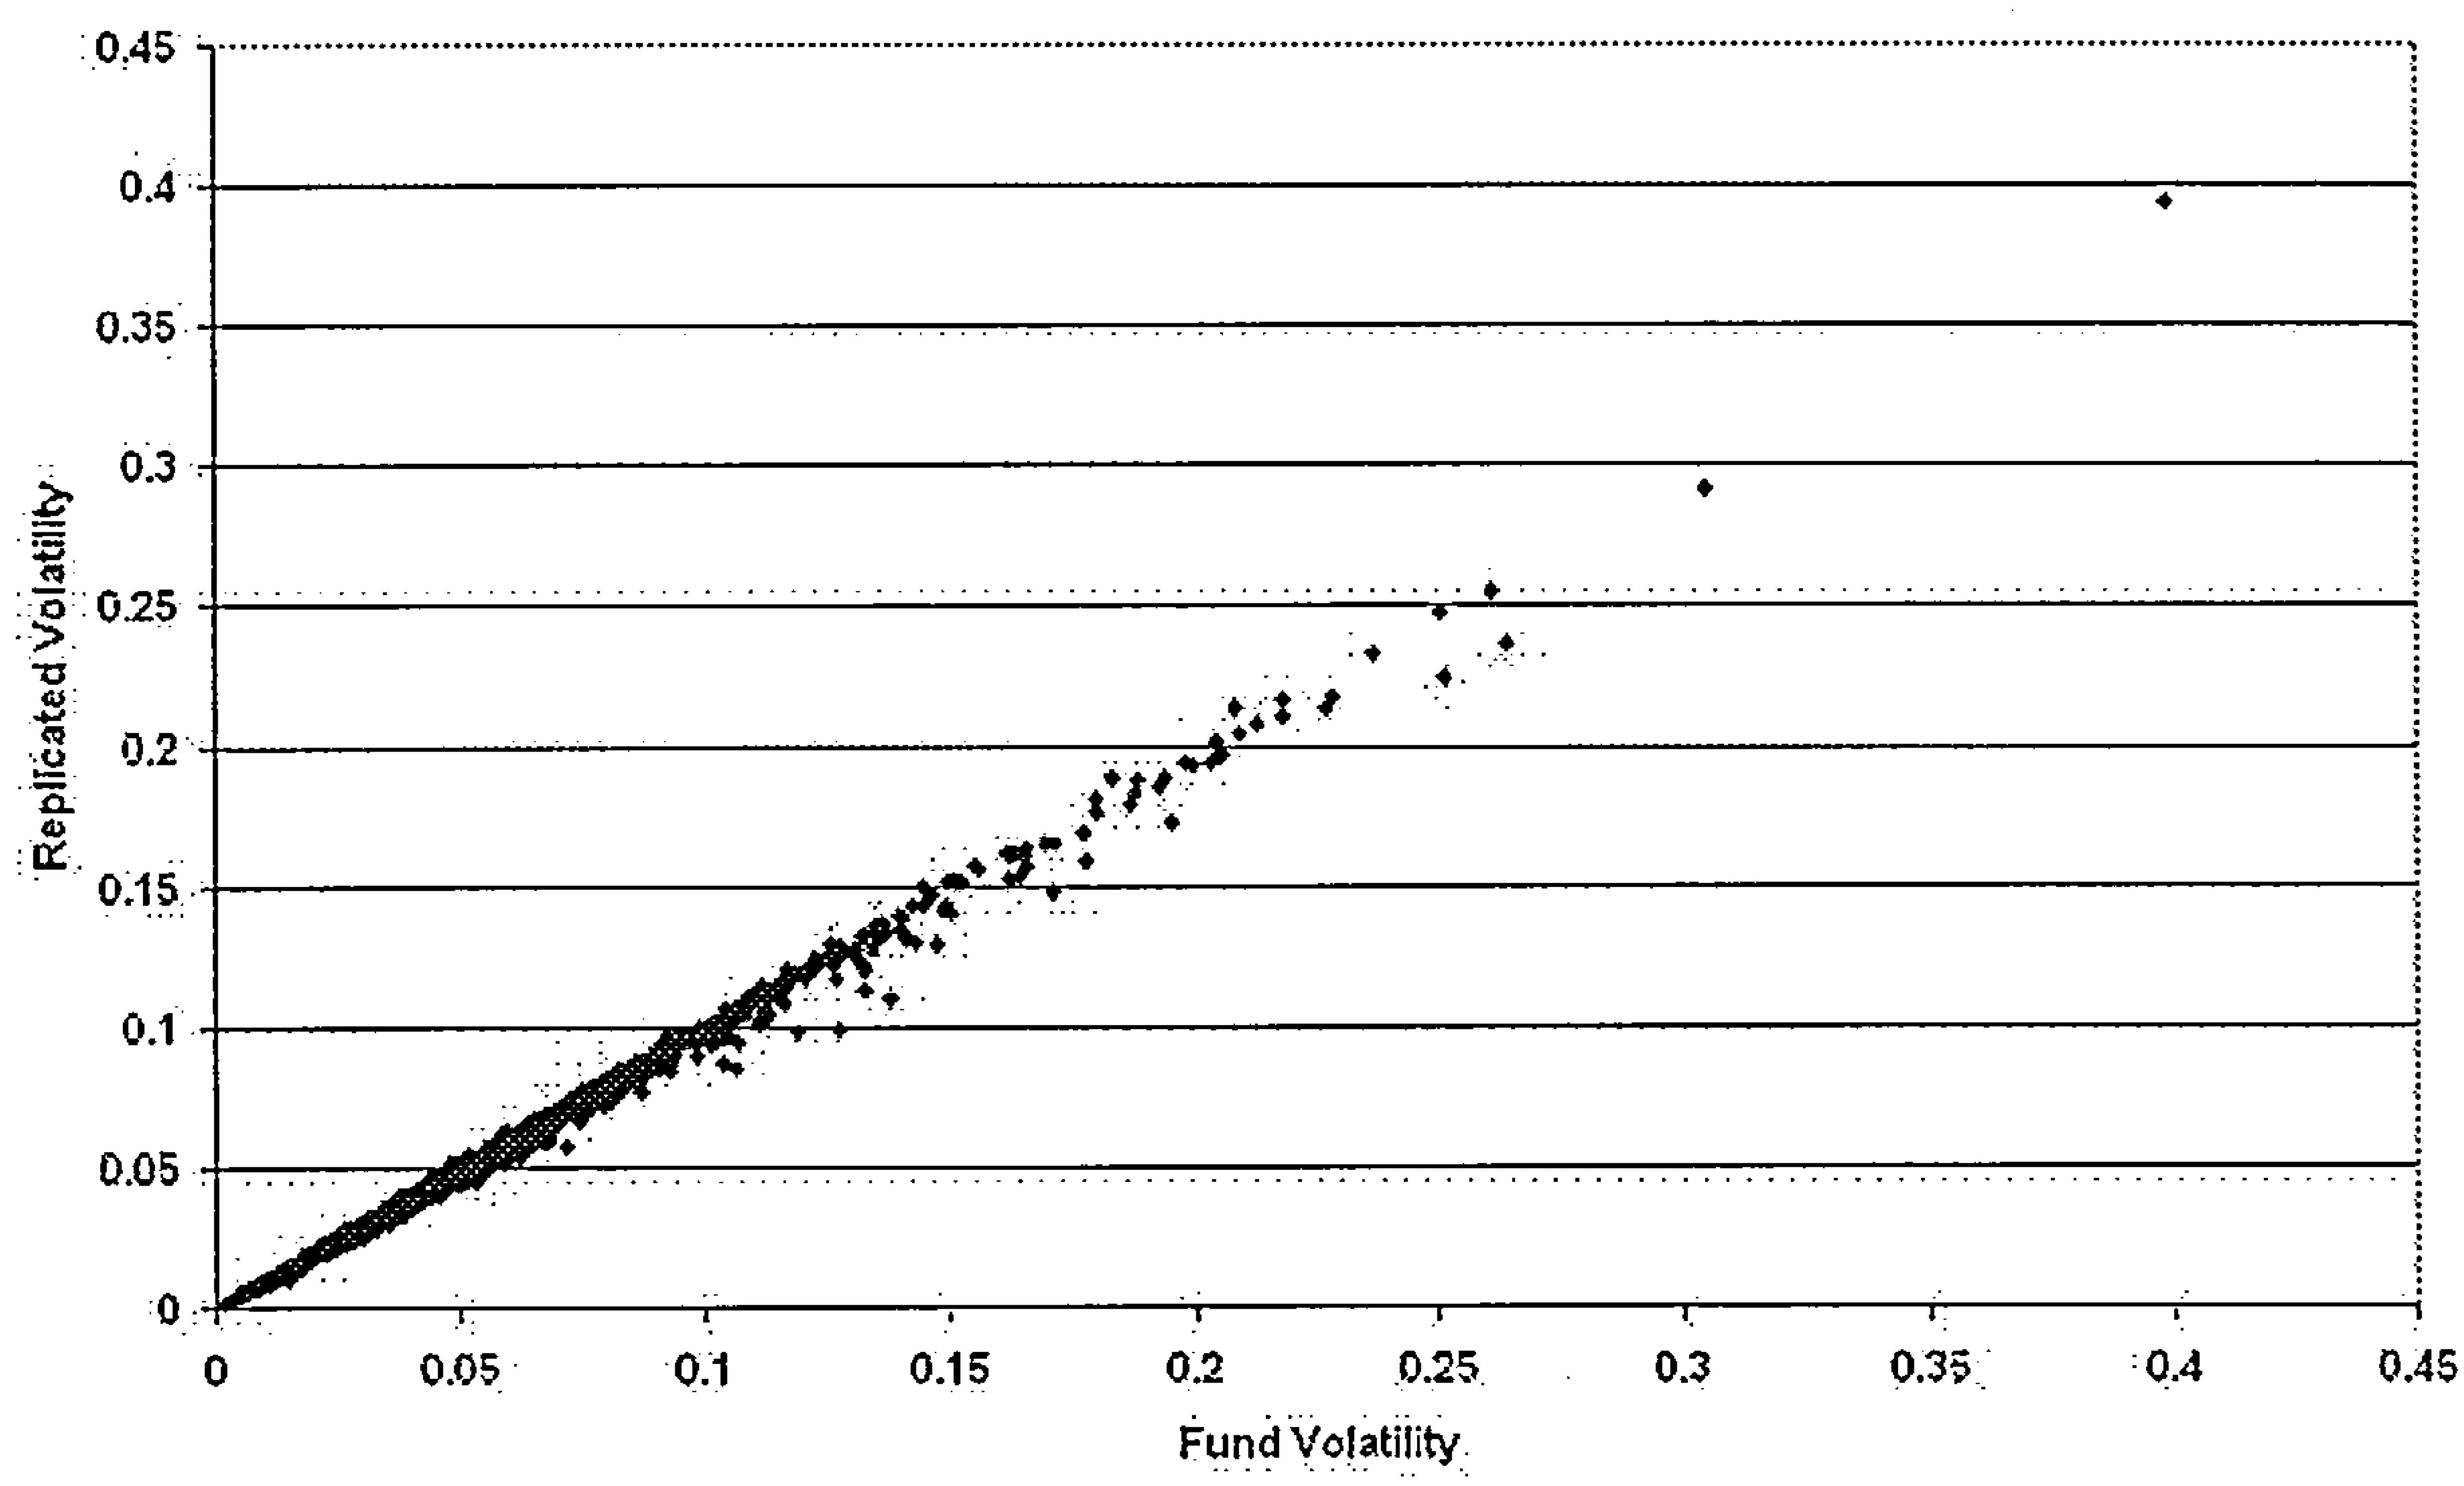

Figure 3.9. Scatter plot hedge fund vs. replicated standard deviation.

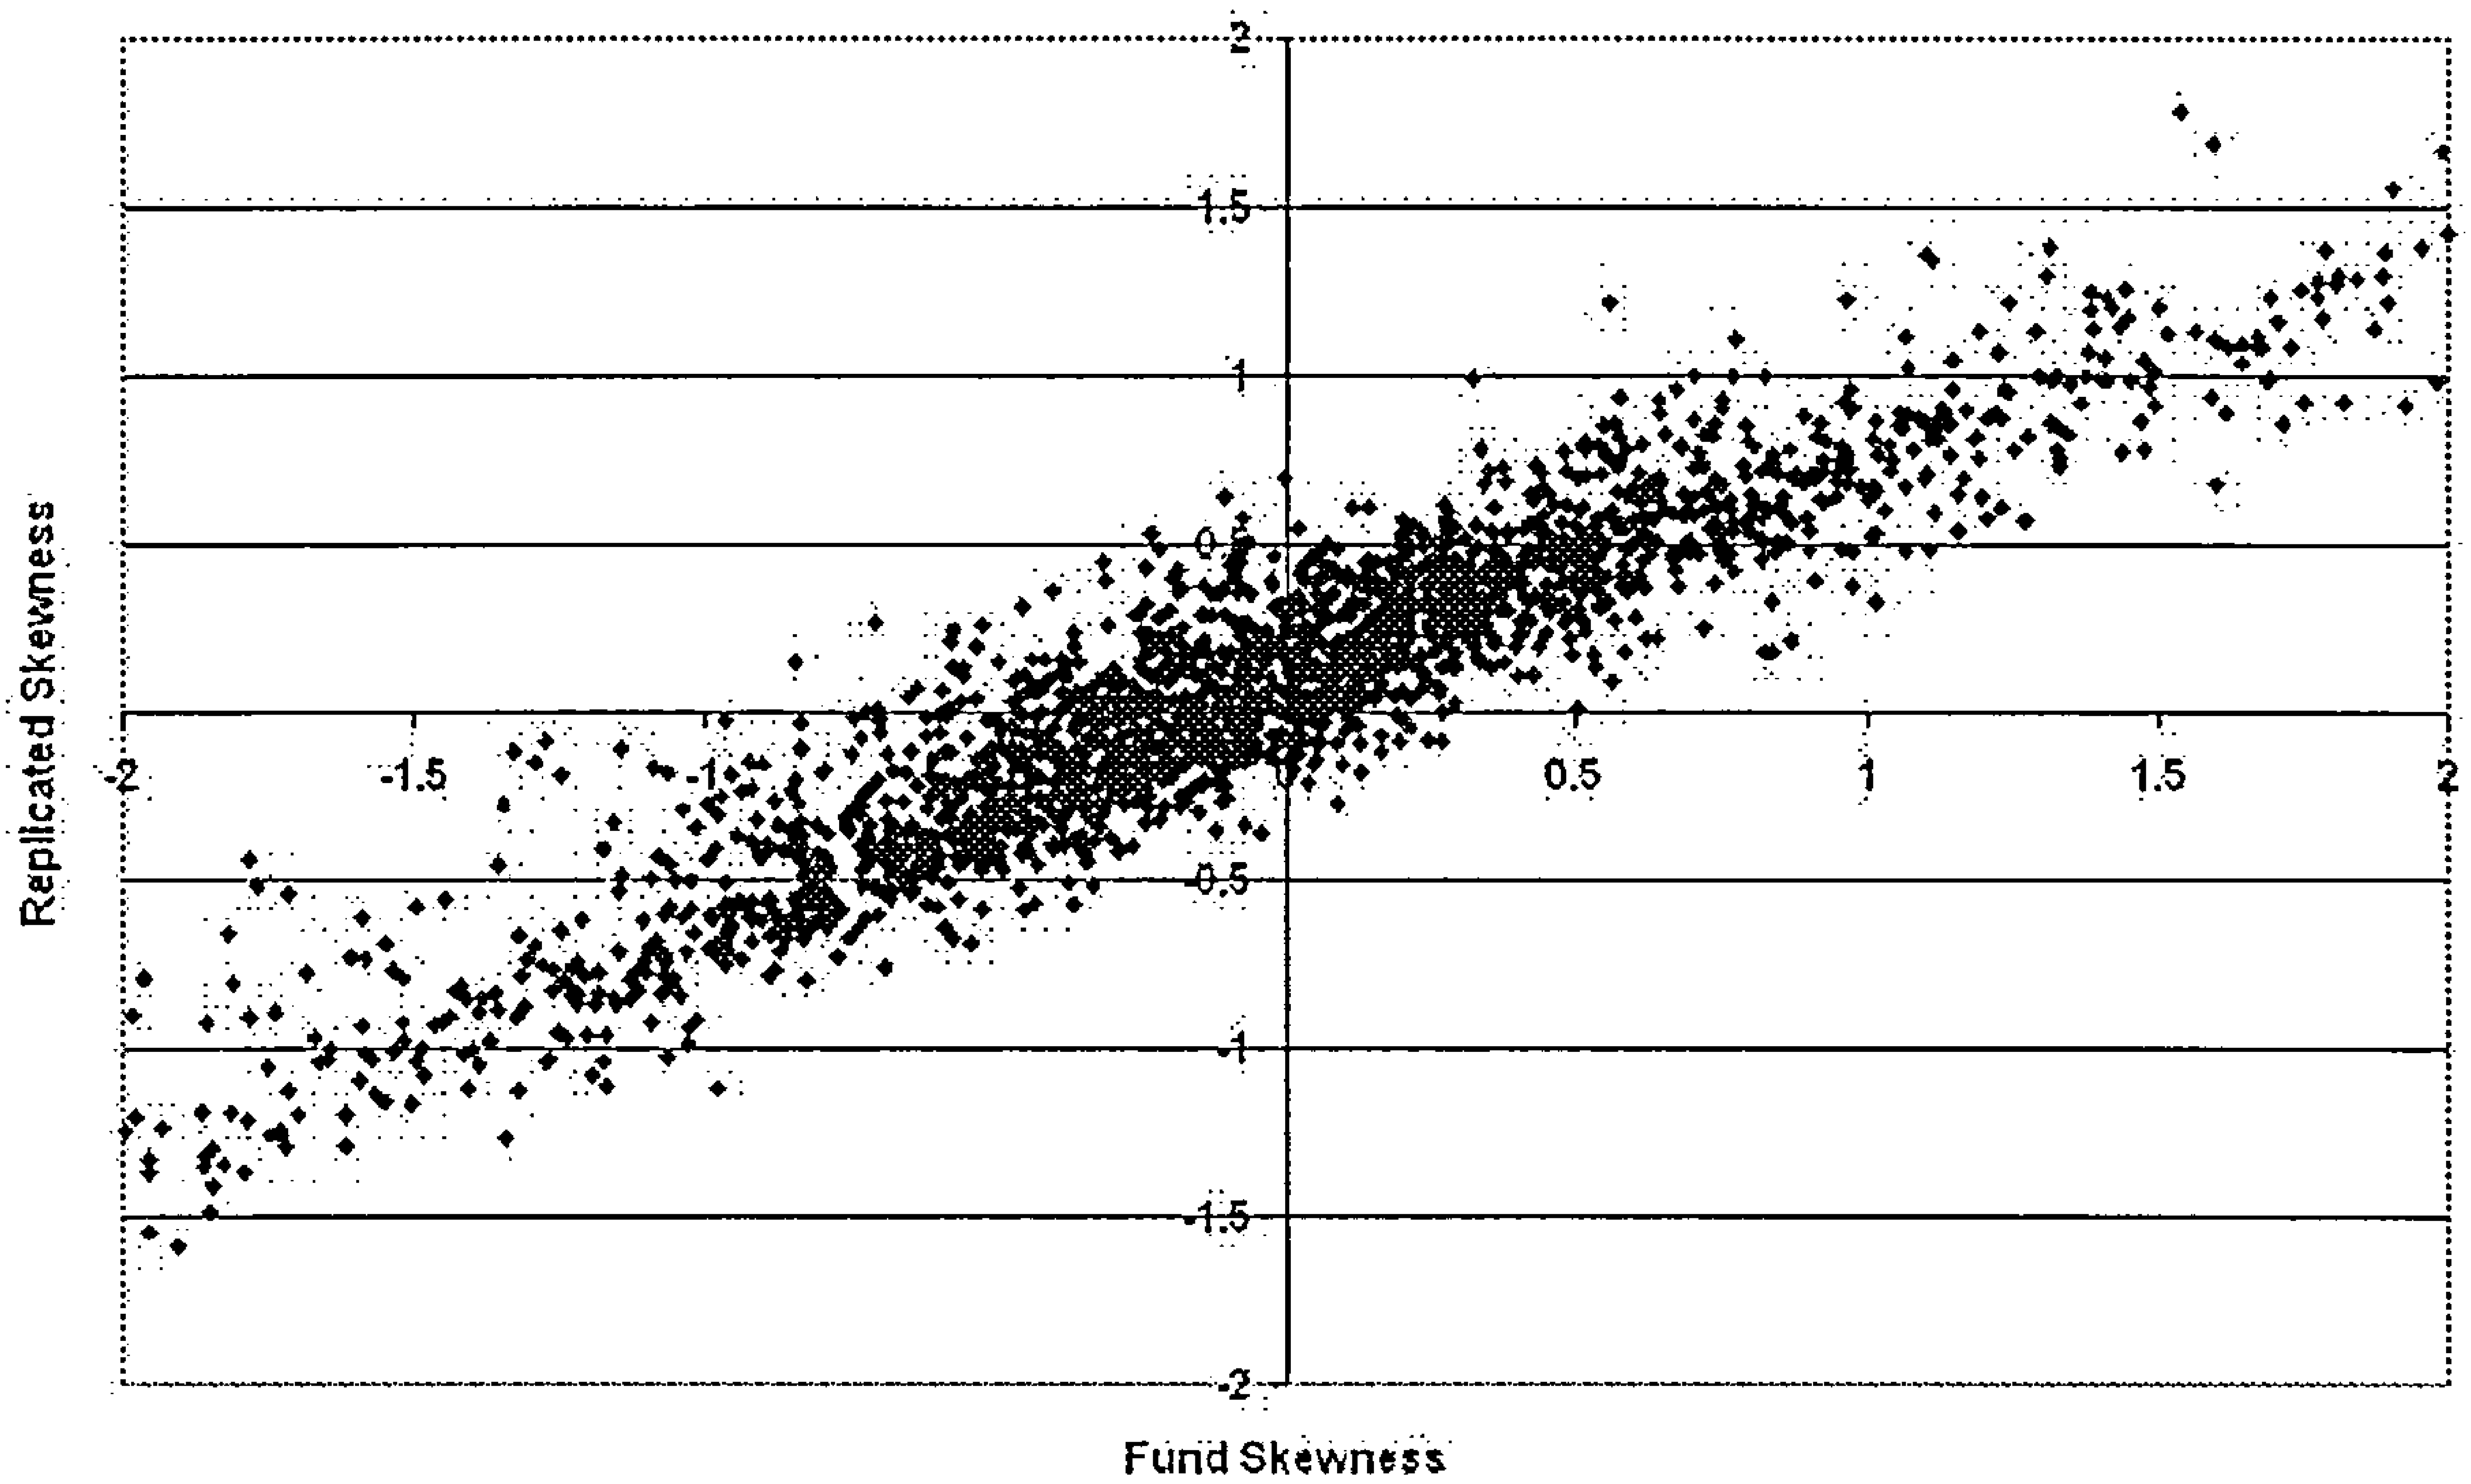

Figure 3.10. Scatter plot hedge fund vs. replicated (conventional) skewness.

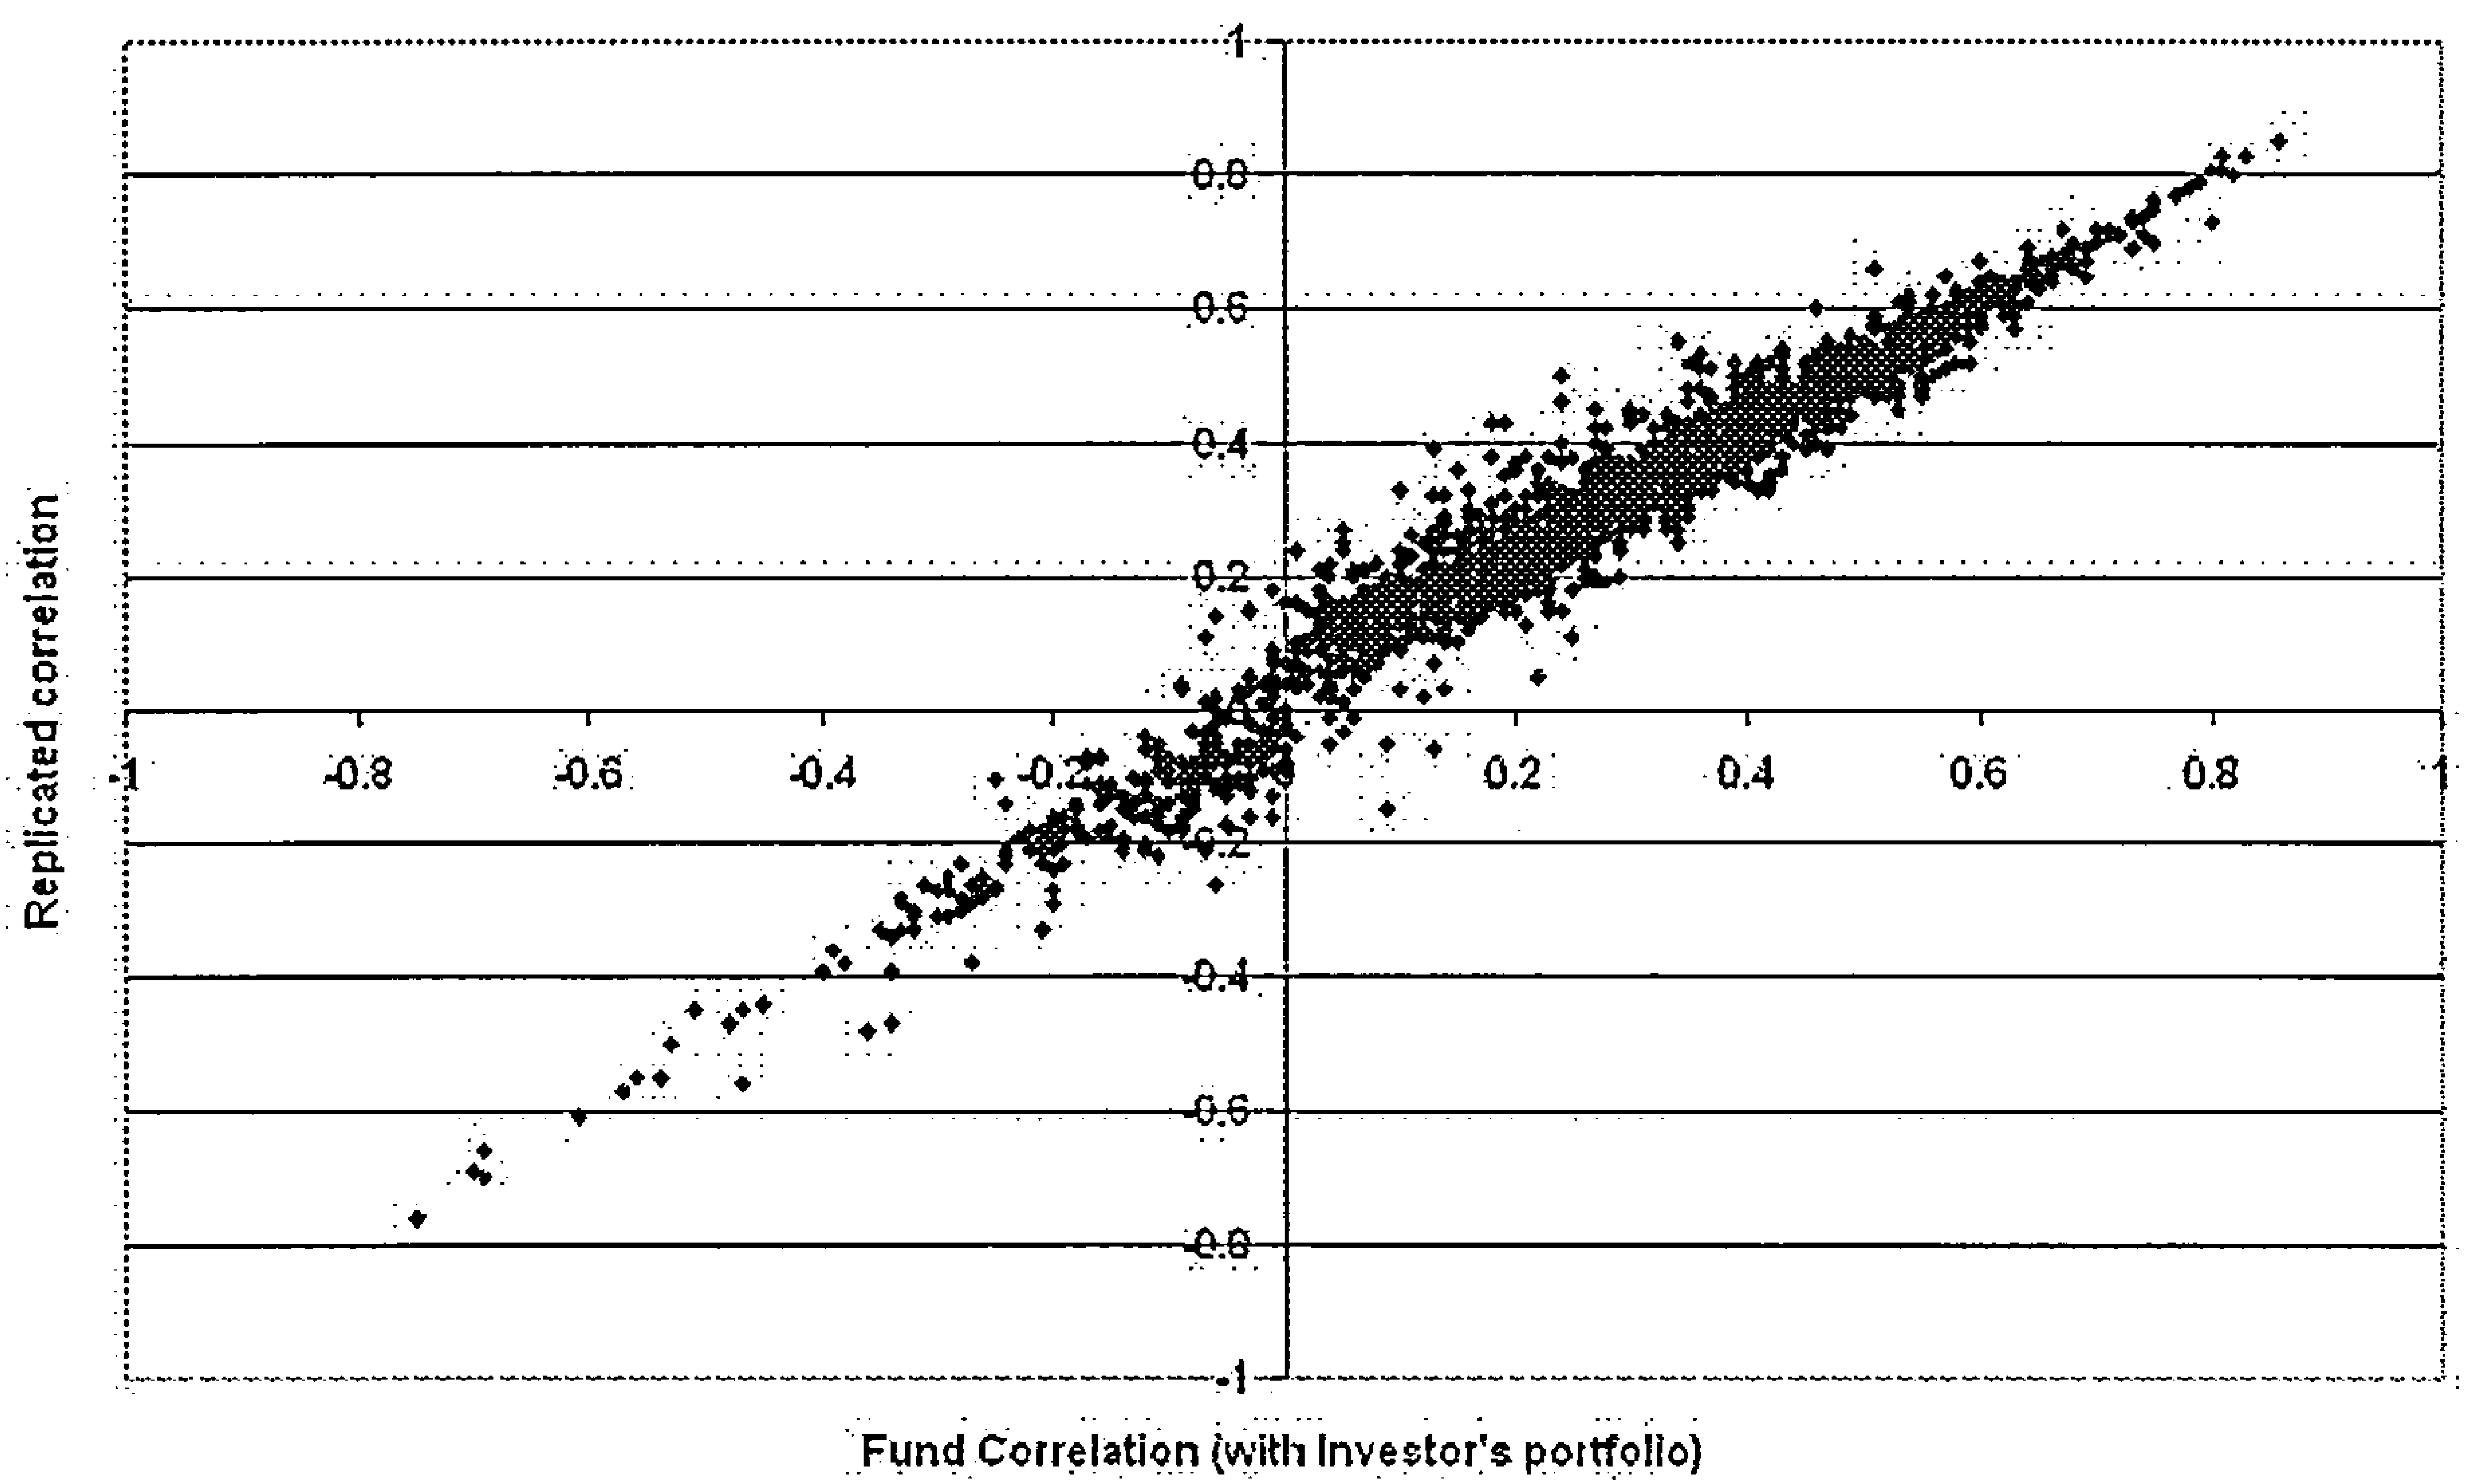

Figure 3.11. Scatter plot hedge fund vs. replicated correlation with investor's portfolio.

Figures 3.9-3.11 also provide additional information on the risk-return profile of the funds in our sample. From Figure 3.10 for example, we see that for most hedge funds estimated skewness lies between –1 and +1. Likewise, from Figure 3.11 we see that the majority of hedge funds are positively correlated with a portfolio of 50% stocks and 50% bonds. Most correlation coefficients lie between 0 and 0.6, again indicating that many hedge funds' returns are a lot less 'market neutral' than often suggested.

	No.	Inefficient	Equivalent	Efficient	Mean	StDev	Skew
Convertible Arbitrage	80	43 (53.8%)	16 (20.0%)	21 (26.2%)	99.816	0.447	-1.517
Distressed Securities	55	16 (29.1%)	12 (21.8%)	27 (49.1%)	100.064	0.585	-1.997
Emerging Markets	182	124 (68.1%)	30 (16.5%)	28 (15.4%)	99.422	1.010	-1.350
Equity Long	309	214 (69.2%)	58 (18.8%)	37 (12.0%)	99.472	0.878	-0.675
Equity Long/Short	411	271 (65.9%)	85 (20.7%)	55 (13.4%)	99.608	0.778	-1.365
Equity Market Neutral	65	49 (75.4%)	13 (20.0%)	3 (4.6%)	99.659	0.503	-0.804
Equity Short Bias	20	15 (75.0%)	5 (25.0%)	0 (0.0%)	99.592	0.656	0.506
Event Driven	96	37 (38.5%)	24 (25.0%)	35 (36.5%)	100.009	0.915	6.488
Fixed Income Arbitrage	149	82 (55.0%)	33 (22.2%)	34 (22.8%)	99.803	0.701	-0.565
Global Macro	80	58 (72.5%)	15 (18.8%)	7 (8.7%)	99.595	0.716	-0.547
Merger Arbitrage	55	22 (40.0%)	14 (25.4%)	19 (34.6%)	99.998	0.364	1.421
Multi-Strategy Arbitrage	104	36 (34.6%)	22 (21.2%)	46 (44.2%)	100.078	0.475	-0.988
Mutual Fund Timing	34	17 (50.0%)	8 (23.5%)	9 (26.5%)	99.670	1.025	-3.043
Other	242	72 (29.8%)	46 (19.0%)	124 (51.2%)	100.171	1.097	-0.088
Sector Funds	165	108 (65.4%)	39 (23.6%)	18 (10.9%)	99.437	1.110	-0.957
Statistical Arbitrage	26	17 (65.4%)	6 (23.1%)	3 (11.5%)	99.826	0.288	-0.066
Total	2073	1181 (57.0%)	426 (20.5%)	466 (22.5%)	99.816	0.447	-1.517

Table 3.11. Hedge fund evaluation results.

We tested the statistical significance of the KP efficiency measure results by calculating bootstrapped confidence intervals, distinguishing between three cases: (1) *Inefficient*, i.e. confidence interval entirely lower than 100, (2) *Efficient*, i.e. confidence interval entirely higher than 100, and (3) *Equivalent*, i.e. confidence interval contains 100. Table 3.11 summarizes the evaluation outcomes. From the table we see that for only 22.5% of all funds the KP measure is convincingly higher than 100. In other words, and similar to funds of funds, the majority of hedge funds have not provided their investors with returns, which they could not have generated themselves in the futures market.

The percentage of efficient funds varies considerably between the different strategy groups, with equity short bias producing the least (0%) and the category 'other' the most (51.2%) efficient funds. Distressed securities (49.1% efficient) and multi-strategy arbitrage (44.2% efficient) stand out as well. When interpreting these results, one has to keep in mind that the available dataset on hedge funds is limited and that most funds have relatively short track records. The idea behind the KP measure is that in the longer run investors receive a return that is fair compensation for the bottom-line risk that they have taken, irrespective of how that risk profile is obtained. For many hedge funds, however, we may not have enough data to be able to properly observe 'the longer run'. The shorter the track record, the more the efficiency measure may be influenced by sampling error[37], in both the fund and the assets traded in the replication strategy. The relatively high proportion of efficient funds in distressed securities and multi-strategy arbitrage for example may have been partly due to the combination of falling interest rates and shrinking credit spreads observed over recent years.

The last three columns of Table 3.11 show the mean, standard deviation and skewness of the frequency distribution of the efficiency measure values observed within each strategy group. For most strategies the distribution is negatively skewed, implying that within each group some funds have shown extremely bad performance relative to what could have been achieved trading a basket of liquid futures contracts. From the table we also see that especially the efficiency measures of strategies whose returns are known to be relatively volatile exhibit a relatively high standard deviation. Global macro and emerging markets for example exhibit relatively high standard deviations, while the opposite is true for convertible arbitrage, equity market neutral and merger arbitrage. Likewise, highly negative skewness is observed in exactly those strategies that are known to be most susceptible to shocks, such as distressed securities, convertible arbitrage, and emerging markets.

[37] Note that this applies to all performance evaluation procedures, not just the KP measure.

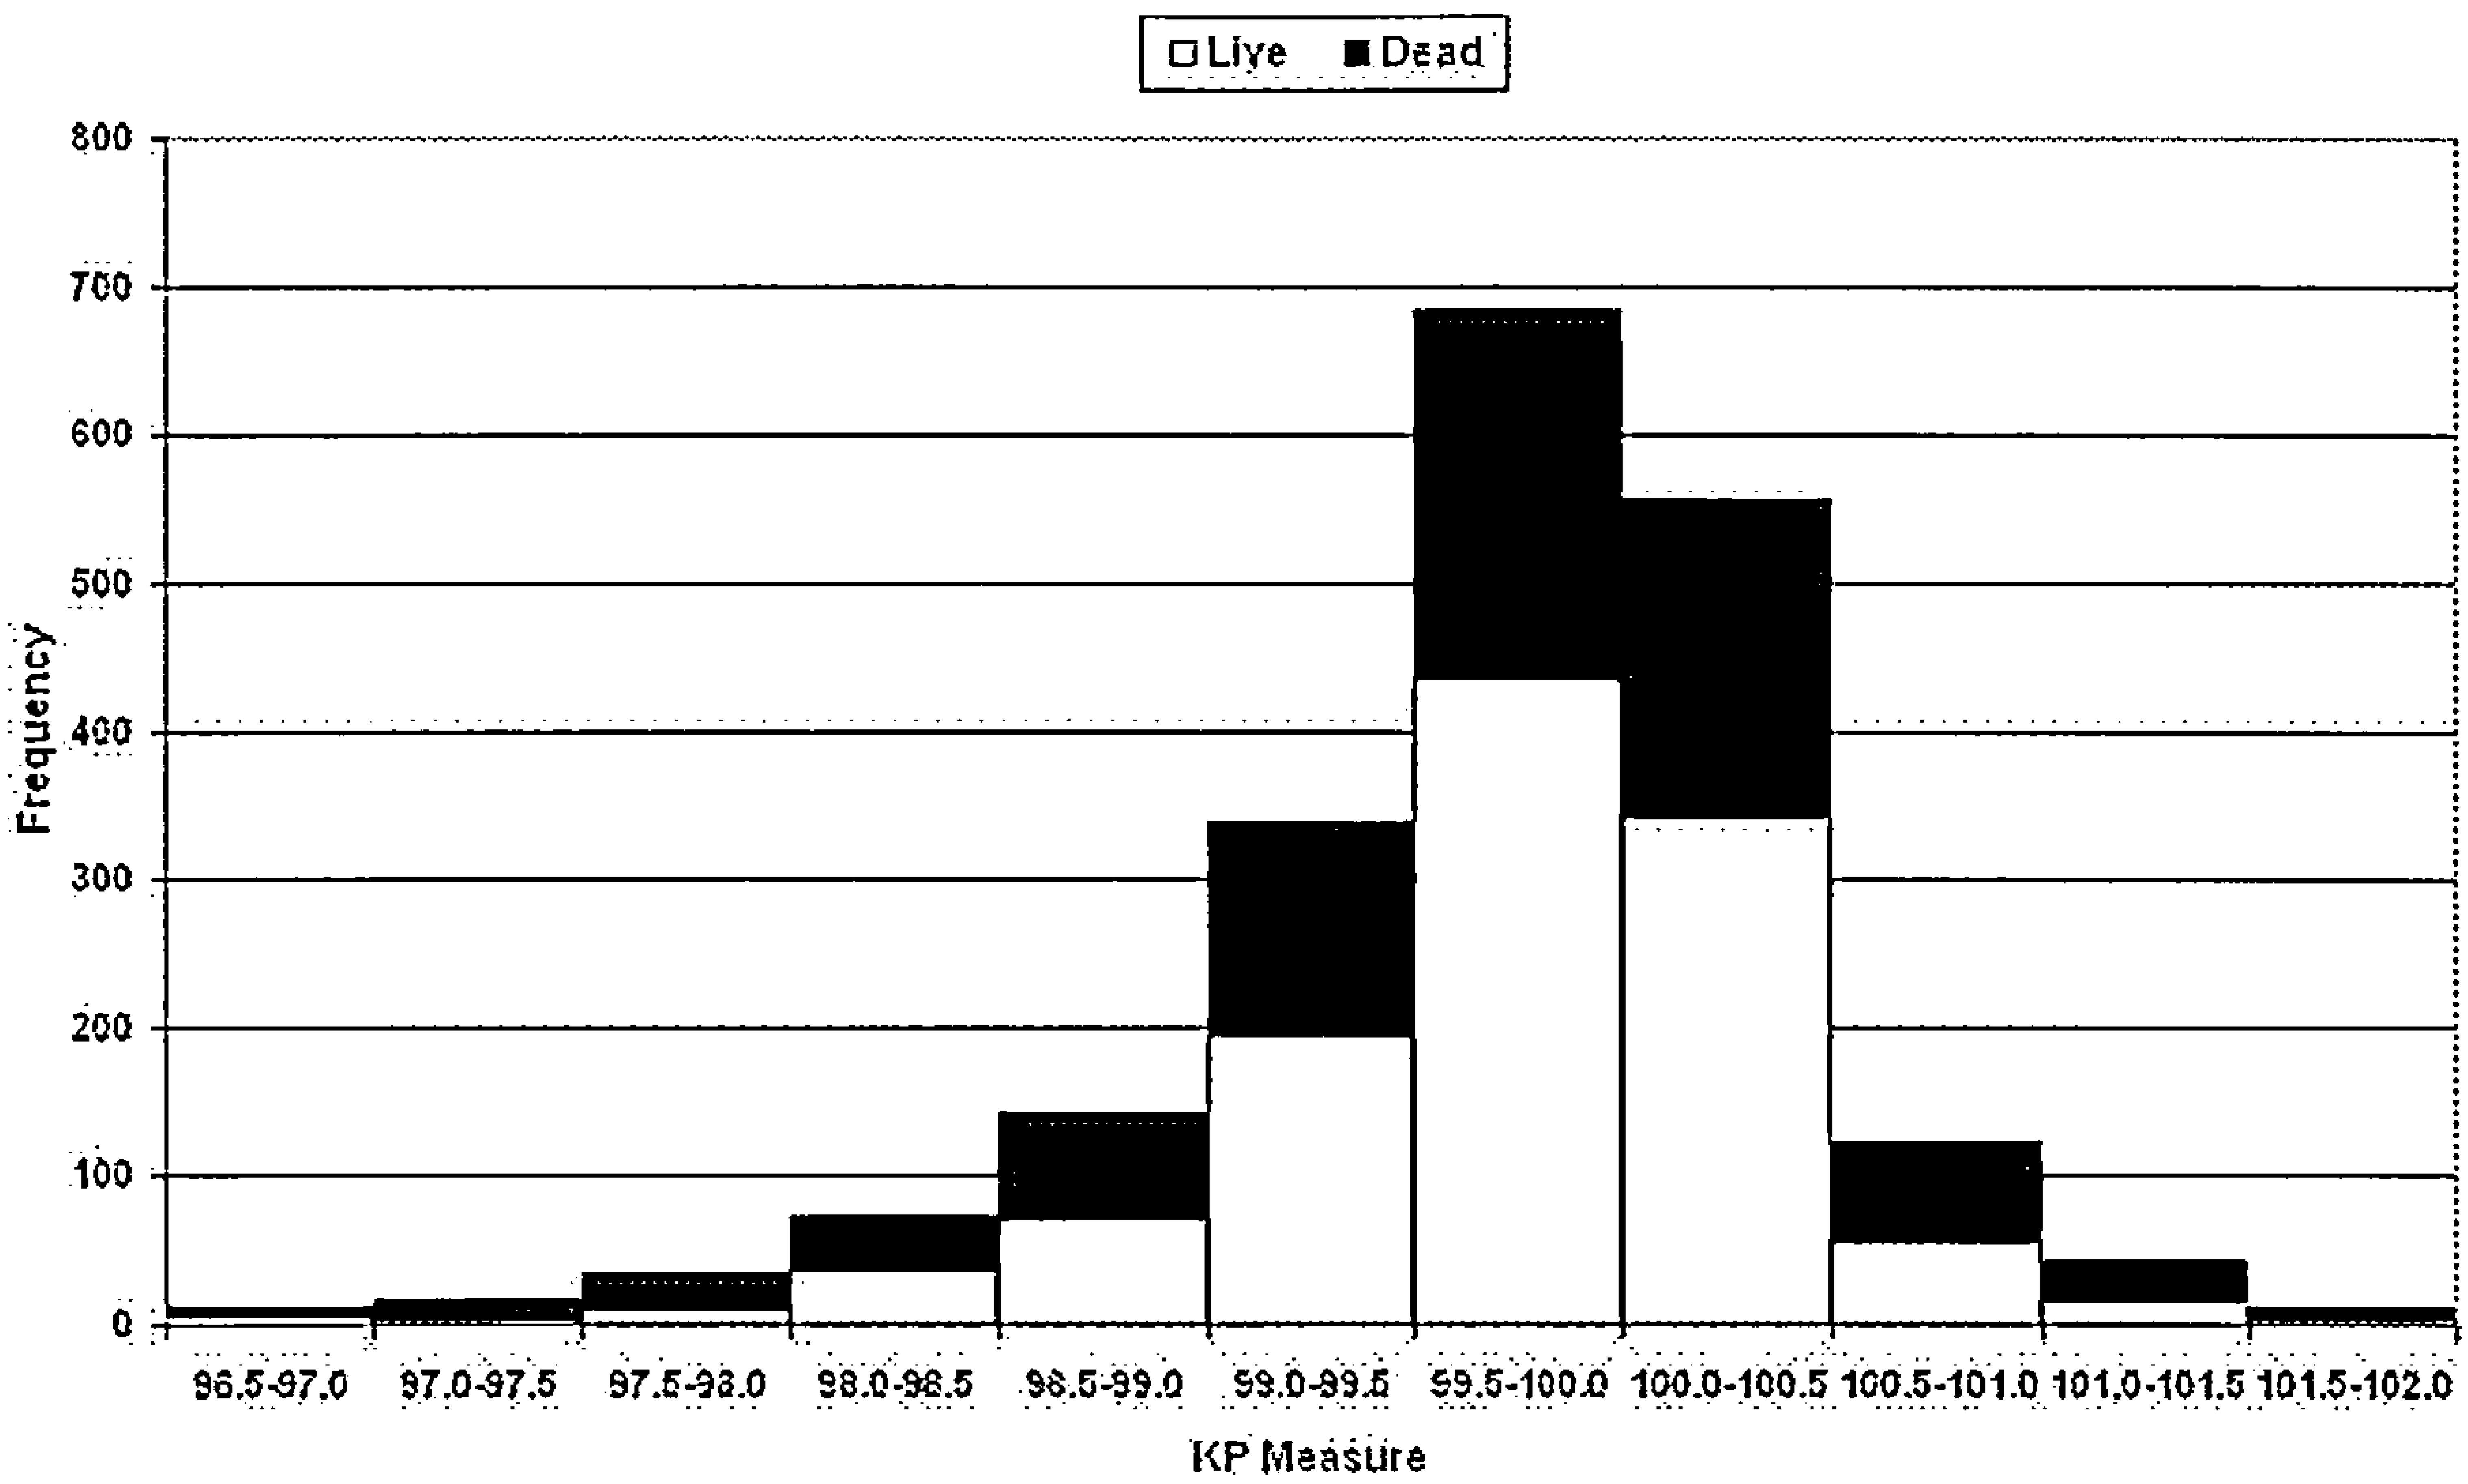

Figure 3.12. Histogram KP efficiency measure 2073 individual hedge funds.

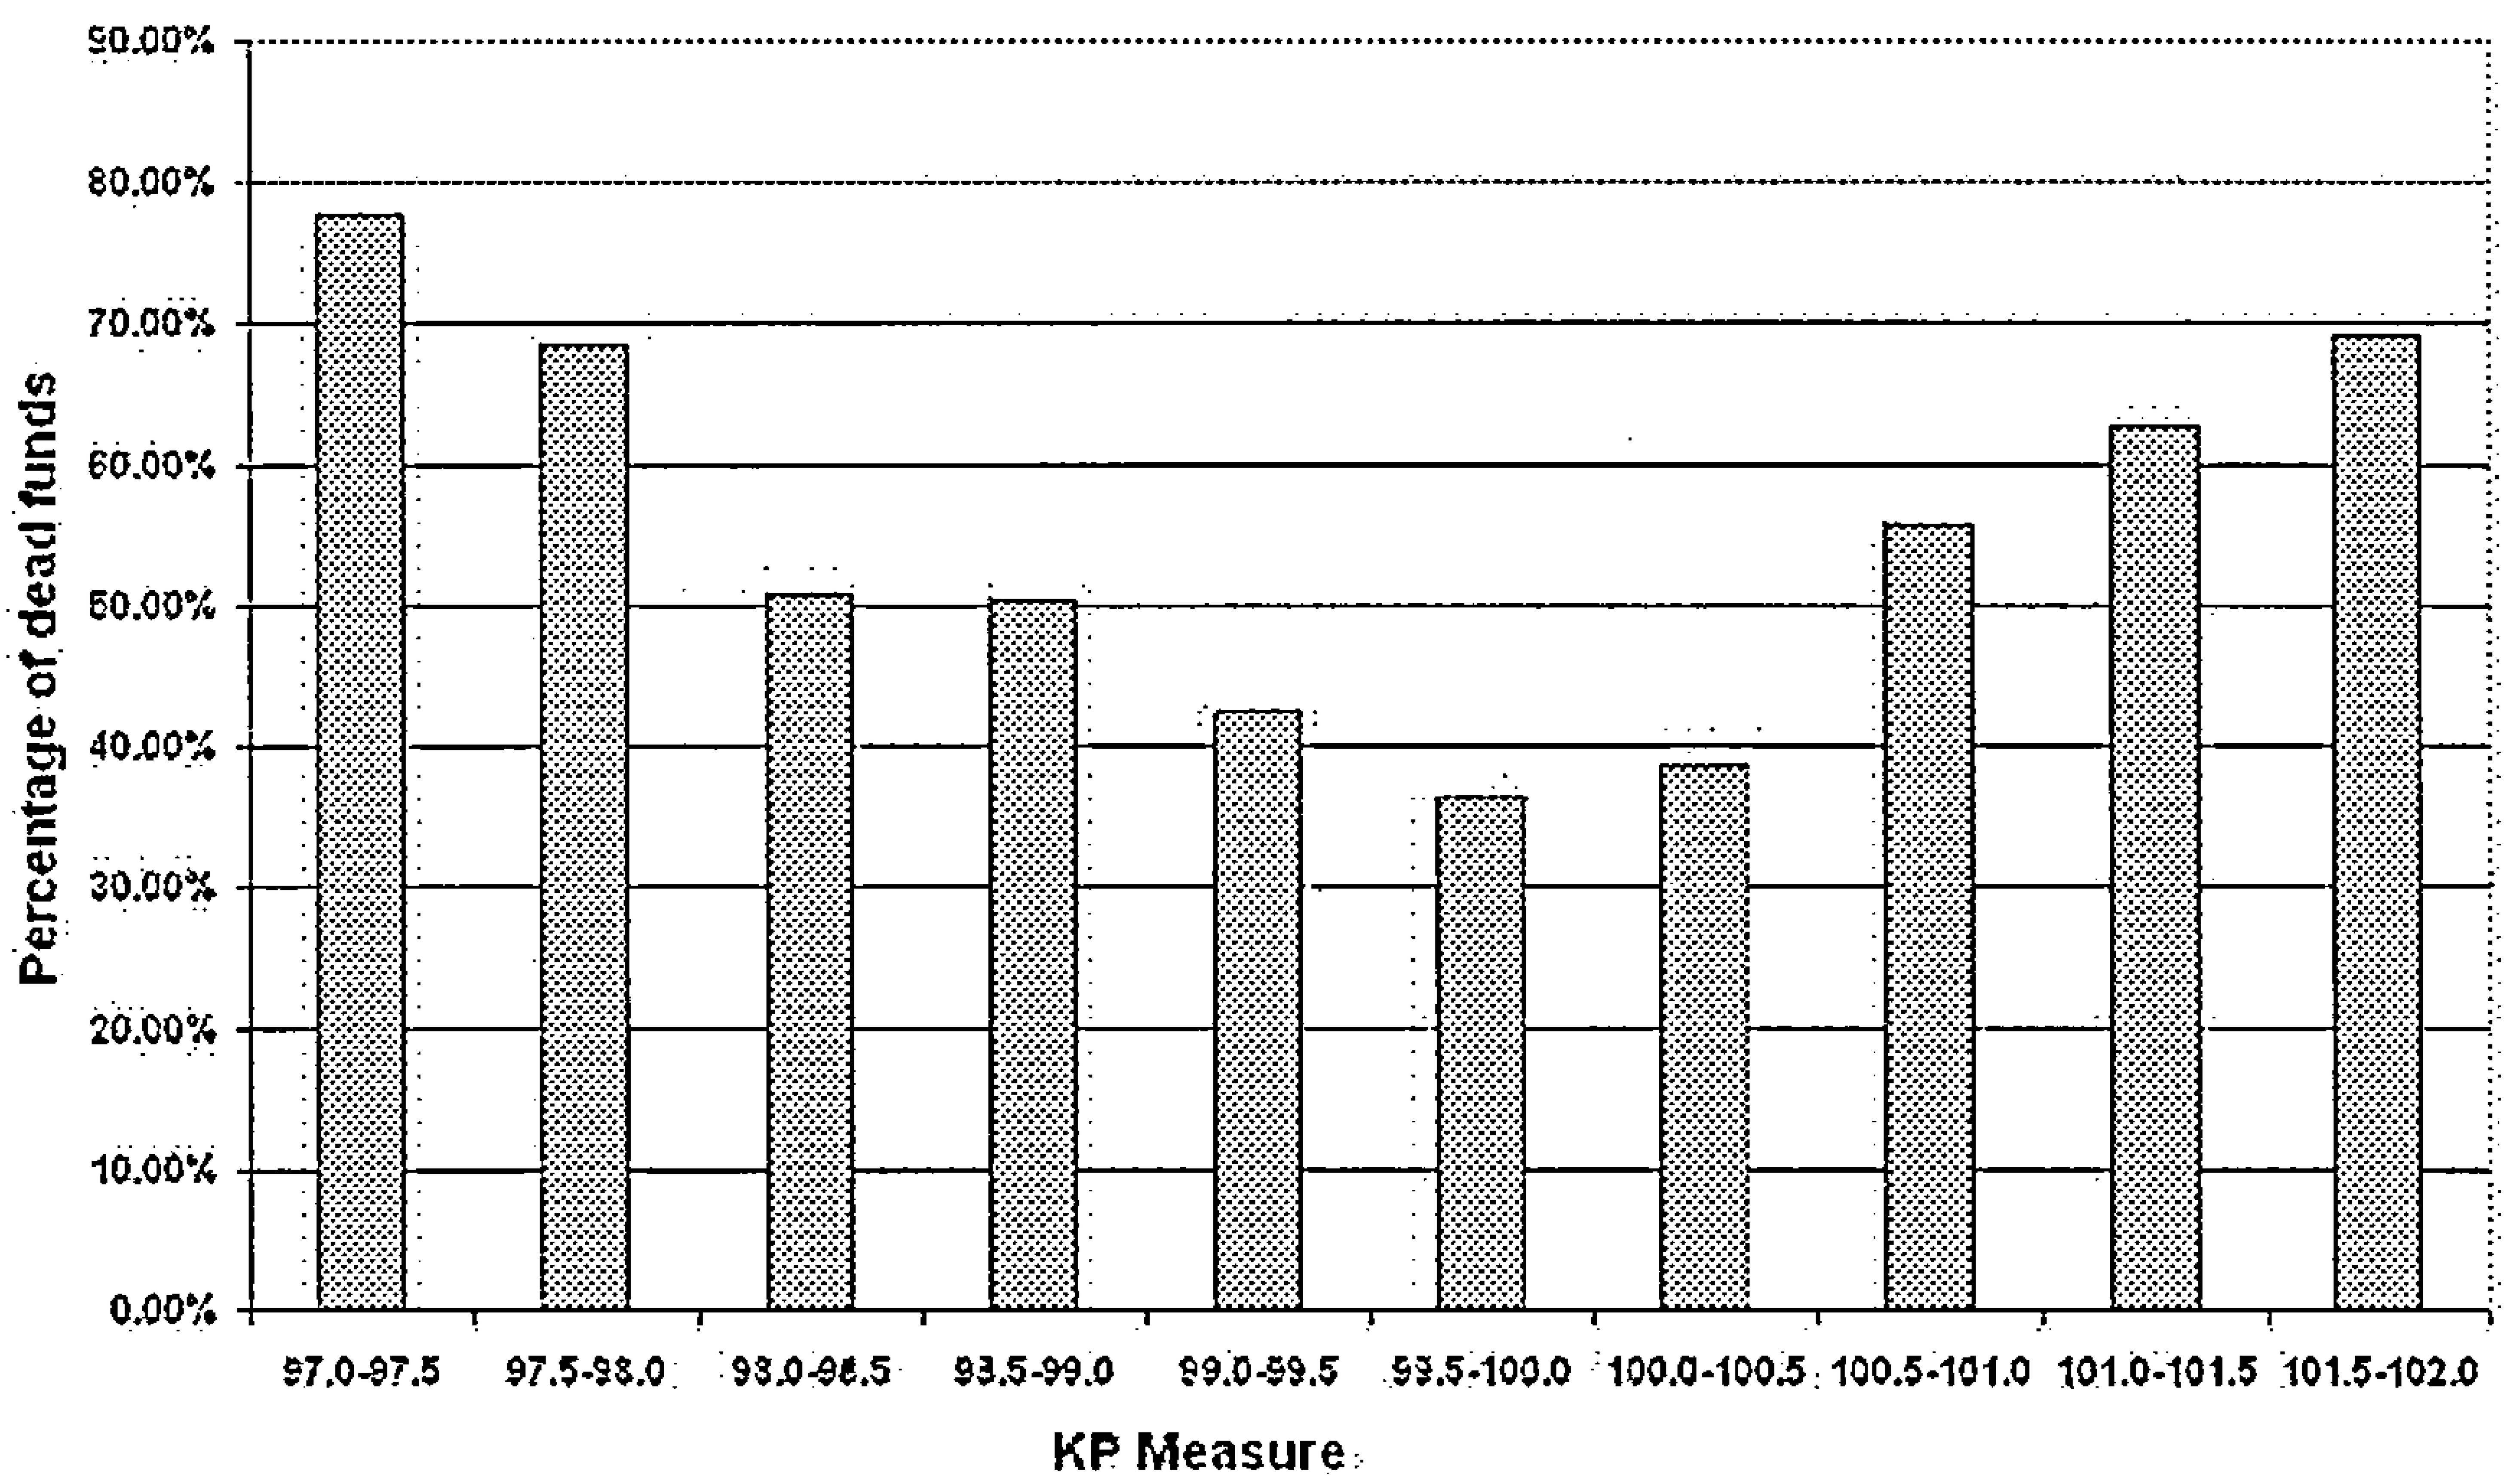

Figure 3.13. Percentage of individual hedge funds that stopped reporting before November 2006 as function KP efficiency measure.

Figure 3.12 shows the frequency distribution of the efficiency measure values found in all 2073 hedge funds[38], while Figure 3.13 separates out the percentage of funds that stopped reporting to the database before November 2006. From the graph we see that there is a strong relationship. Out of the 294 funds with a KP measure below 99, no less than 166 (56.5%) stopped reporting. Out of the 681 funds with a KP measure between 100 and 101, only 285 (41.8%) did so, with probably a significant number of these not really closing down but simply stopping reporting. A similar relationship is observed in the average efficiency measure values of live and dead funds. The average KP measure over the 892 dead funds is 99.66. If we assume that all funds with a KP measure above 100 did not really die, but simply stopped reporting, the average KP measure for dead funds drops to 99.13. Over the 1181 funds still alive on the other hand, the average is 99.75.

Track Record (months)	Mean	StDev	Skew	Kurt
48 - 71	99.80	1.05	0.12	9.62
72 - 95	99.56	0.95	-1.05	5.77
96 - 119	99.65	0.74	-1.26	3.24
120 - 143	99.80	0.56	-1.11	3.02
144+	99.79	0.55	-1.17	1.90

Table 3.12. Distribution KP measure as function length of track record.

To investigate whether there is any indication of older funds doing better than younger funds or vice versa, we sorted the funds in our sample on the length of their track record. Table 3.12 shows the statistics of the resulting frequency distributions of the KP measure. From the means we see that on average age has little or no impact on performance. The standard deviation, skewness, and kurtosis measures all drop when we move to funds with a longer track record, suggesting significant differences. On the other hand, with more data available, sampling error is less of an issue, which could well explain (at least part of) the declining dispersion of observed KP measure values.

[38] Histograms for the various strategies show a similar picture and are therefore not reported.

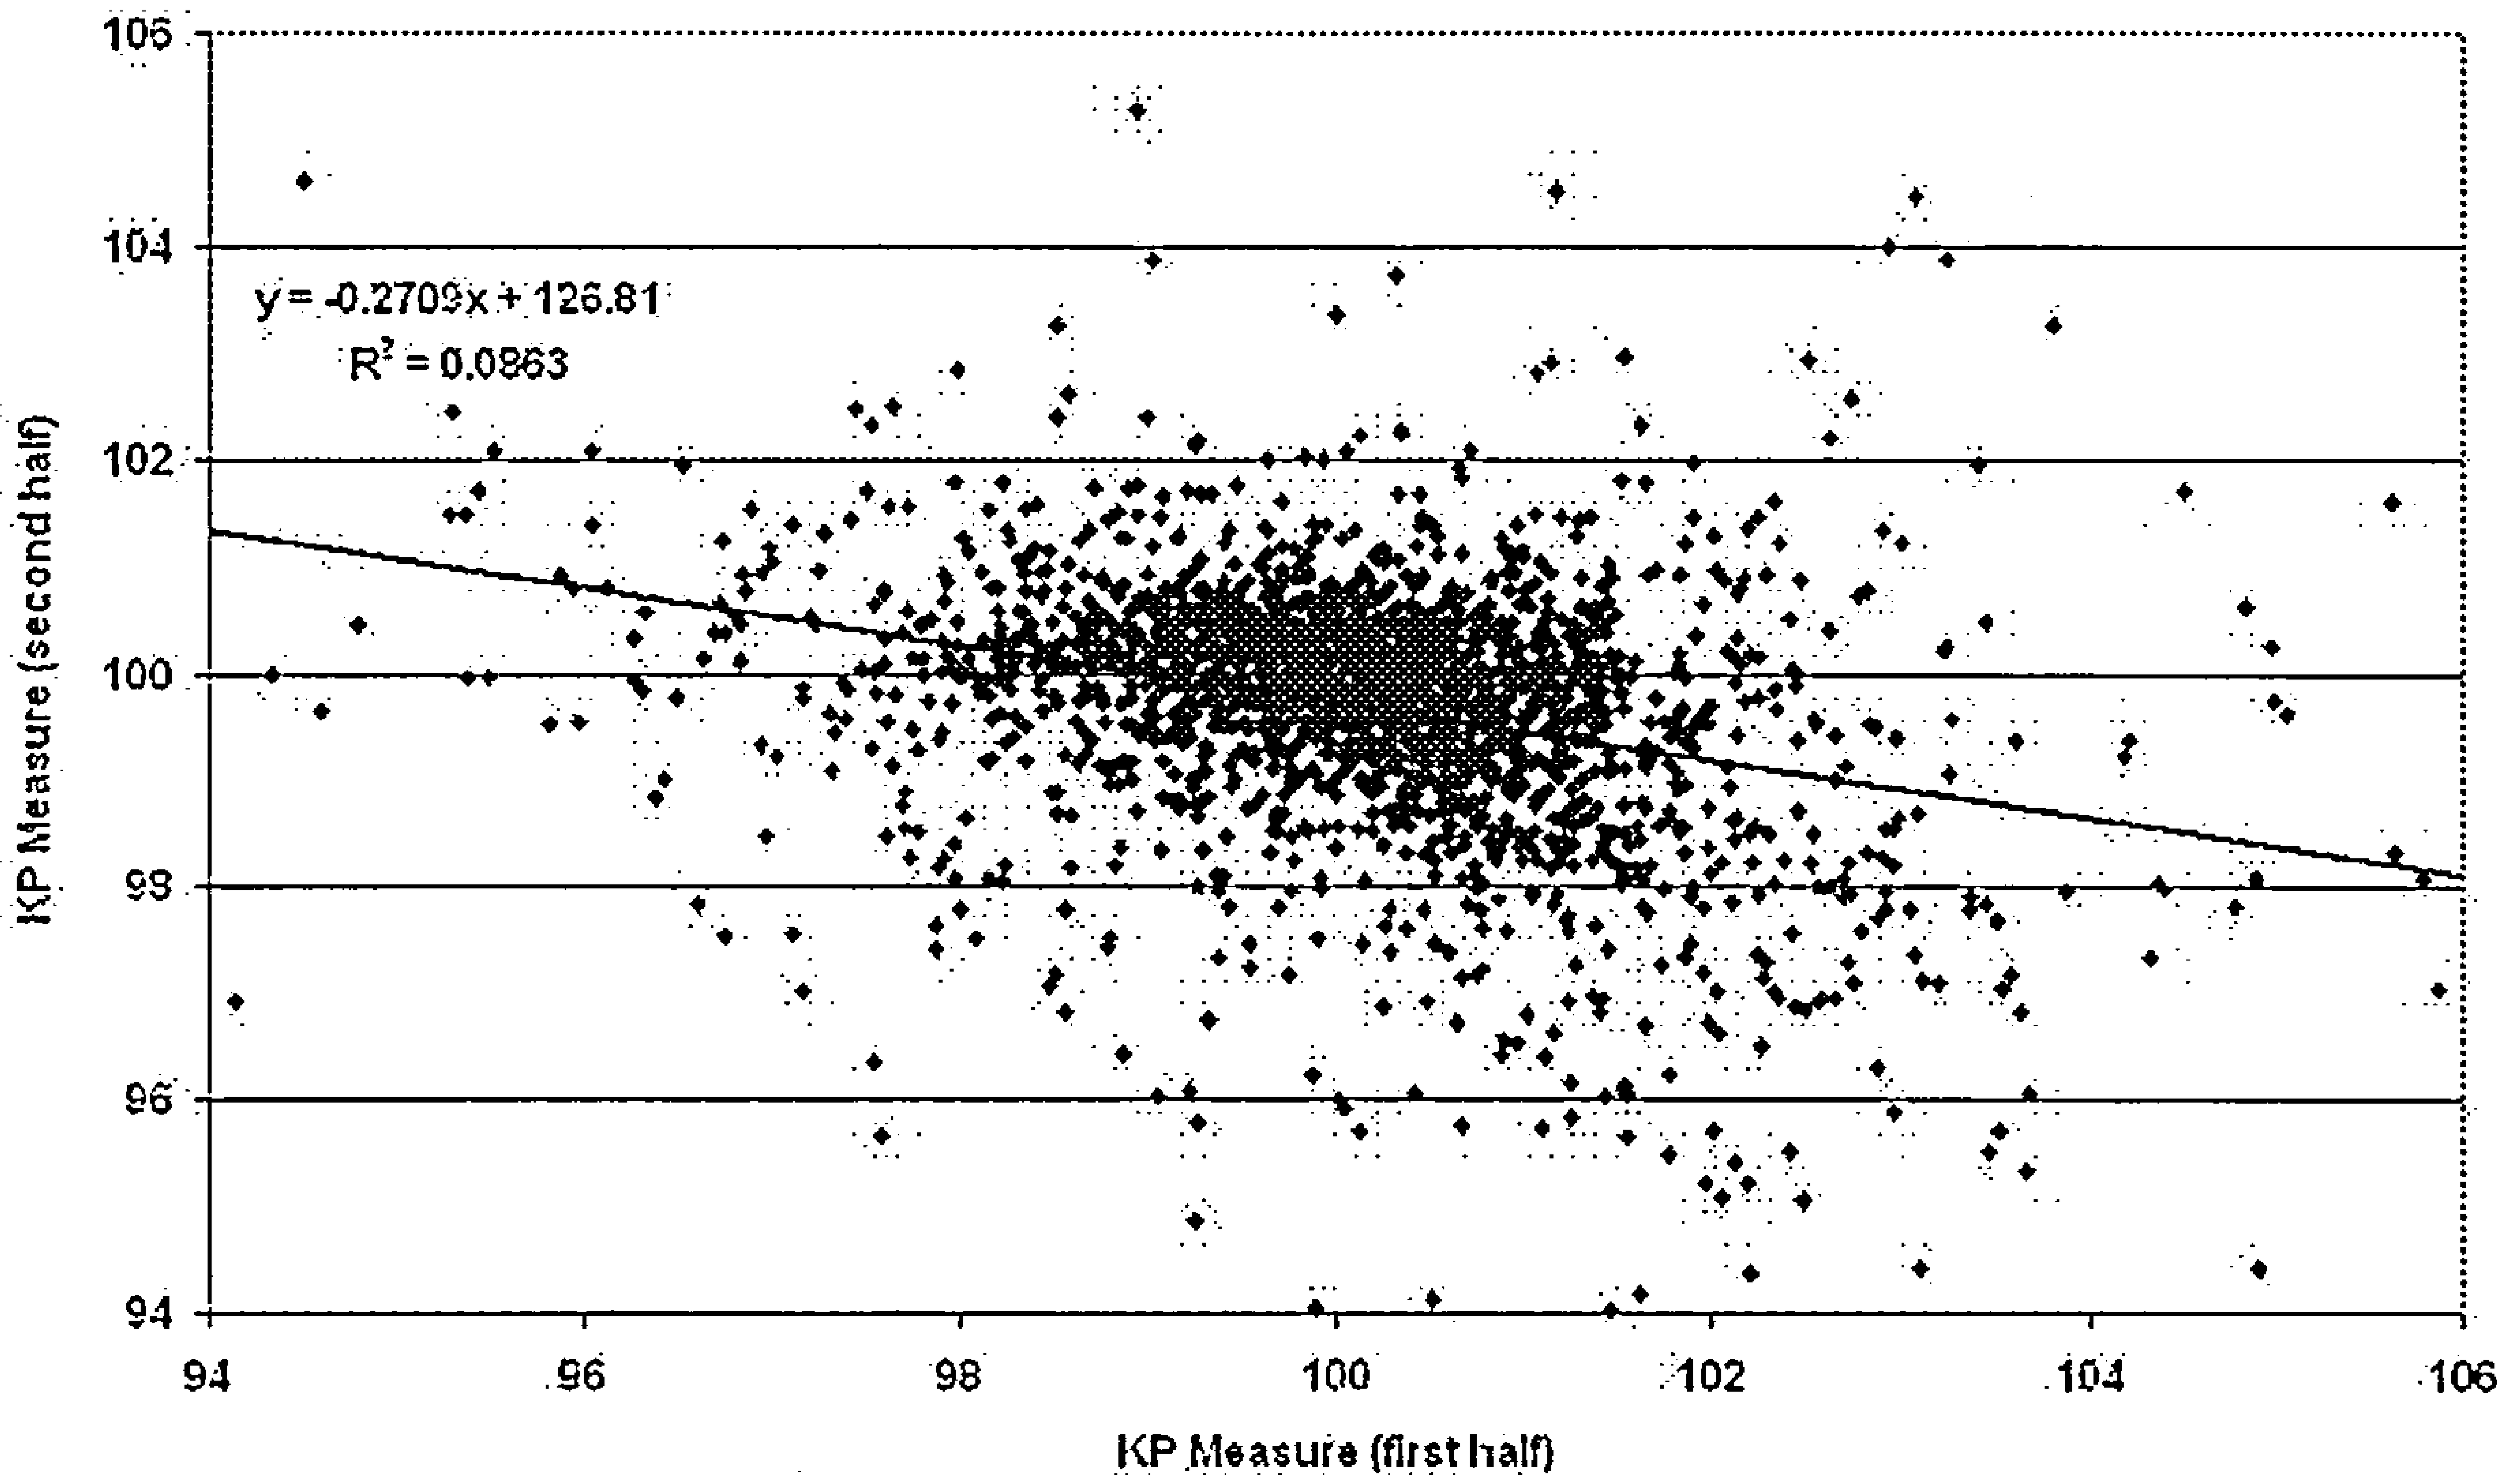

Figure 3.14: KP efficiency measures over split track record.

Finally, we split the track record of all funds with 8 or more years of history in two equal parts and calculated the KP measure over each part. The result is shown in Figure 3.14. Over all funds, the average over the first period was 100.18 while over the second, more recent, period the average was only 99.67. This indicates a very substantial deterioration in average hedge fund performance over time. As for funds of funds, Figure 3.14 reveals a tendency for hedge funds with a relatively high (low) KP measure in the first period to produce a relatively low (high) KP measure in the second. This again supports the hypothesis that increased fund size tends to have a negative impact on future performance.

4. Creation of synthetic funds[39]

As said in the Introduction, there is no reason why the replication technique developed in Chapter 2 could not be used to create completely new funds, providing investors with previously unavailable return characteristics. Some examples of the so-called 'synthetic funds' will be discussed in this chapter.

4.1 Out-of-Sample Tests

In this section we study the out-of-sample performance of four different synthetic funds, the details of which are shown in Table 4.1. Throughout we assume that the synthetic funds in question are created to further diversify a larger traditional portfolio consisting of 50% S&P 500 and 50% T-bonds.

	Volatility	Skewness	Excess Kurtosis	Correlation with Investor's Portfolio
Fund 1	12%	0.00	0.00	0.00
Fund 2	12%	2.00	10.00	0.00
Fund 3	12%	0.00	0.00	-0.50
Fund 4	Fund 1 with –5% floor on monthly return			

Table 4.1. Overview of the four synthetic funds studied.

The first case is quite straightforward. It concerns a synthetic fund that generates returns with a volatility of 12%, no significant skewness or kurtosis and zero correlation with the investor's existing portfolio of 50% stocks and 50% bonds. This risk profile is similar to that of a well-diversified portfolio of commodity futures[40]. Fund 2 is the same as fund 1 except that apart from zero correlation we also aim for a significant degree of positive skewness[41]. Fund 3 is also similar to fund 1, except that in this case we aim for even lower correlation with stocks and bonds. With a correlation of –0.5, this is similar to the risk profile of an investment in pure stock market volatility[42]. Finally, in fund 4

[39] The material in this chapter appeared as Alternative Investment Research Center WP 36, Cass Business School.

[40] See Kat and Oomen (2006a, 2006b) for details on the statistical behaviour of commodity futures investments.

[41] We have to raise excess kurtosis to 10 since it is very difficult to generate significant skewness without extra kurtosis.

[42] See Carr and Wu (2006) or Kat and Tassabehji (2006) for details on volatility investment.

we floor the monthly fund 1 return at –5%. This is similar to the risk profile of some of the hedge fund and commodity-linked notes that are offered by the main alternative product providers.

In the above we have not set a target for the expected fund return. The reason for this is that synthetic funds are not designed in isolation. Given interest rates, volatilities, correlations, etc., some parameter choices are feasible, while others are not, as the fund parameters have to be in line with the prevailing pricing environment in the global capital markets. Practically speaking, this means we can choose all parameters ourselves, except for one, which is subsequently determined by the capital markets. In all four cases studied we fully specify the funds' risk profiles, while leaving the expected return for the capital markets to determine. Once a fund's risk profile is specified and the accompanying dynamic trading strategy has been derived, we can of course calculate its expected return, but strictly speaking the latter is not part of the target.

With the reference portfolio given, the most important decision left is the composition of the reserve asset. Unfortunately, there is no universally optimal reserve asset[43]. What makes a good reserve asset depends very much on the composition of the reference portfolio and the expected return on the various asset classes. More specifically, and apart from liquidity, what should we be looking for in a good reserve asset? First, since it is the main building block of every trading strategy, its statistical properties need to be stable. This means a well-diversified portfolio will generally be preferred over a single asset. Second, since we'll always be long the reserve asset, it needs to have an attractive expected return relative to its risk level. In simple terms, the reserve asset needs to have a high Sharpe ratio. Note that this may be achieved in many different ways, ranging from a de-leveraged portfolio of high volatility assets to a highly leveraged portfolio of low volatility assets. Third, although not absolutely necessary, it helps when the reserve asset shares some of the skewness and kurtosis characteristics of the target, as a reserve asset without any skewness or kurtosis may have difficulty generating fund returns, which do display a significant degree of skewness or kurtosis.

[43] In Appendix A.4 the influence of the moments of the reserve asset distribution and its correlation with the reference portfolio on the KP measure is examined.

Since the outlook for the various asset classes as well as the composition of the investor's portfolio will change over time, in practice the choice of the reserve asset is a dynamic process, producing time-varying allocations. Unfortunately, the latter process is very difficult to simulate in a backtest without the suggestion of data mining. In what follows, we therefore assume that the composition of the reserve asset is fixed though time[44]. More specifically, we assume the reserve asset consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, T-bond, S&P 500, Russell 2000 and GSCI futures[45]. This captures three main asset classes. The resulting portfolio is quite well diversified, with, over the period Jan/1998-Feb/2007, an annualised volatility of 9.41%. Throughout, we trade the nearby futures contract, rolling into the next nearby contract on the first day of the expiry month, assuming transaction costs of 1bp one-way.

Before we look at the results of the backtests, it is important to note that there can be temporary discrepancies between the target parameters chosen and the sample parameters generated. We might be after a standard deviation of 12%, but when calculating the standard deviation from the returns actually generated we might find 11% or 13% instead. This is nothing unusual though. When tossing a coin, the chances of heads and tails are 50/50. This does not mean that when tossing a coin a limited number of times one will always find an equal number of heads and tails. In a small sample, heads may dominate tails or vice versa. When the number of observations increases, however, this is likely to be corrected as the sample becomes more representative for the distribution it is taken from.

In the above context, it is also important to note that over the last decade financial markets have exhibited some quite bizarre behaviour. Over 1995 – 1999 the S&P 500 rose by 212% and the Nasdaq by 447%. Subsequently, over 2000 - 2002, the S&P 500 fell by 40%, and the Nasdaq by no less than 68%. Short-term USD interest rates exhibited similar behaviour, dropping from 6.8% in 2000 to 1.1% in 2004, and rising

[44] Note that not allowing for tactical considerations in the selection of the reserve asset means that we may underestimate the returns that could have been achieved in practice. In reality, for example, it may not have been rational to be long interest rate futures when 1-month USD Libor stood at no more than 1.1%, as was the case in early 2004.

[45] Since the volatility of the various asset classes is quite different, before forming the portfolio, we leveraged the Eurodollar and 5Y note by a factor 5, and the 10Y note by a factor 4 to give these components a level of volatility more in line with stocks and commodities.

back to 5.3% in 2006. Commodity price rises caused the GSCI to rise by 138% over 2002 - 2006. Finally, the last decade also saw its fair share of crises: Russia, LTCM, 9/11, Iraq, etc. Obviously, all of this has a serious impact on our tests and should be taken into account when interpreting the results.

4.1.1 Fund 1

Let's assume a USD-based investor lived in January 1998 and started synthetic fund 1. Before we look at what kind of returns he would have generated over time, Figure 4.1 shows the payoff function, which the investor will be aiming to produce as per January 1998. From the graph we see that the desired fund payoff is an increasing function of the reserve asset, but a declining function of the investor's portfolio. Since the slope of the payoff function determines what positions to hold in the investor's portfolio and the reserve asset, this means that we'll be long the reserve asset and short the investor's portfolio. The reason for this is that the correlation between the reserve asset and the investor's portfolio exceeds the zero correlation that is targeted for the synthetic fund return. To reduce the correlation to the desired level, we therefore have to short the investor's portfolio[46].

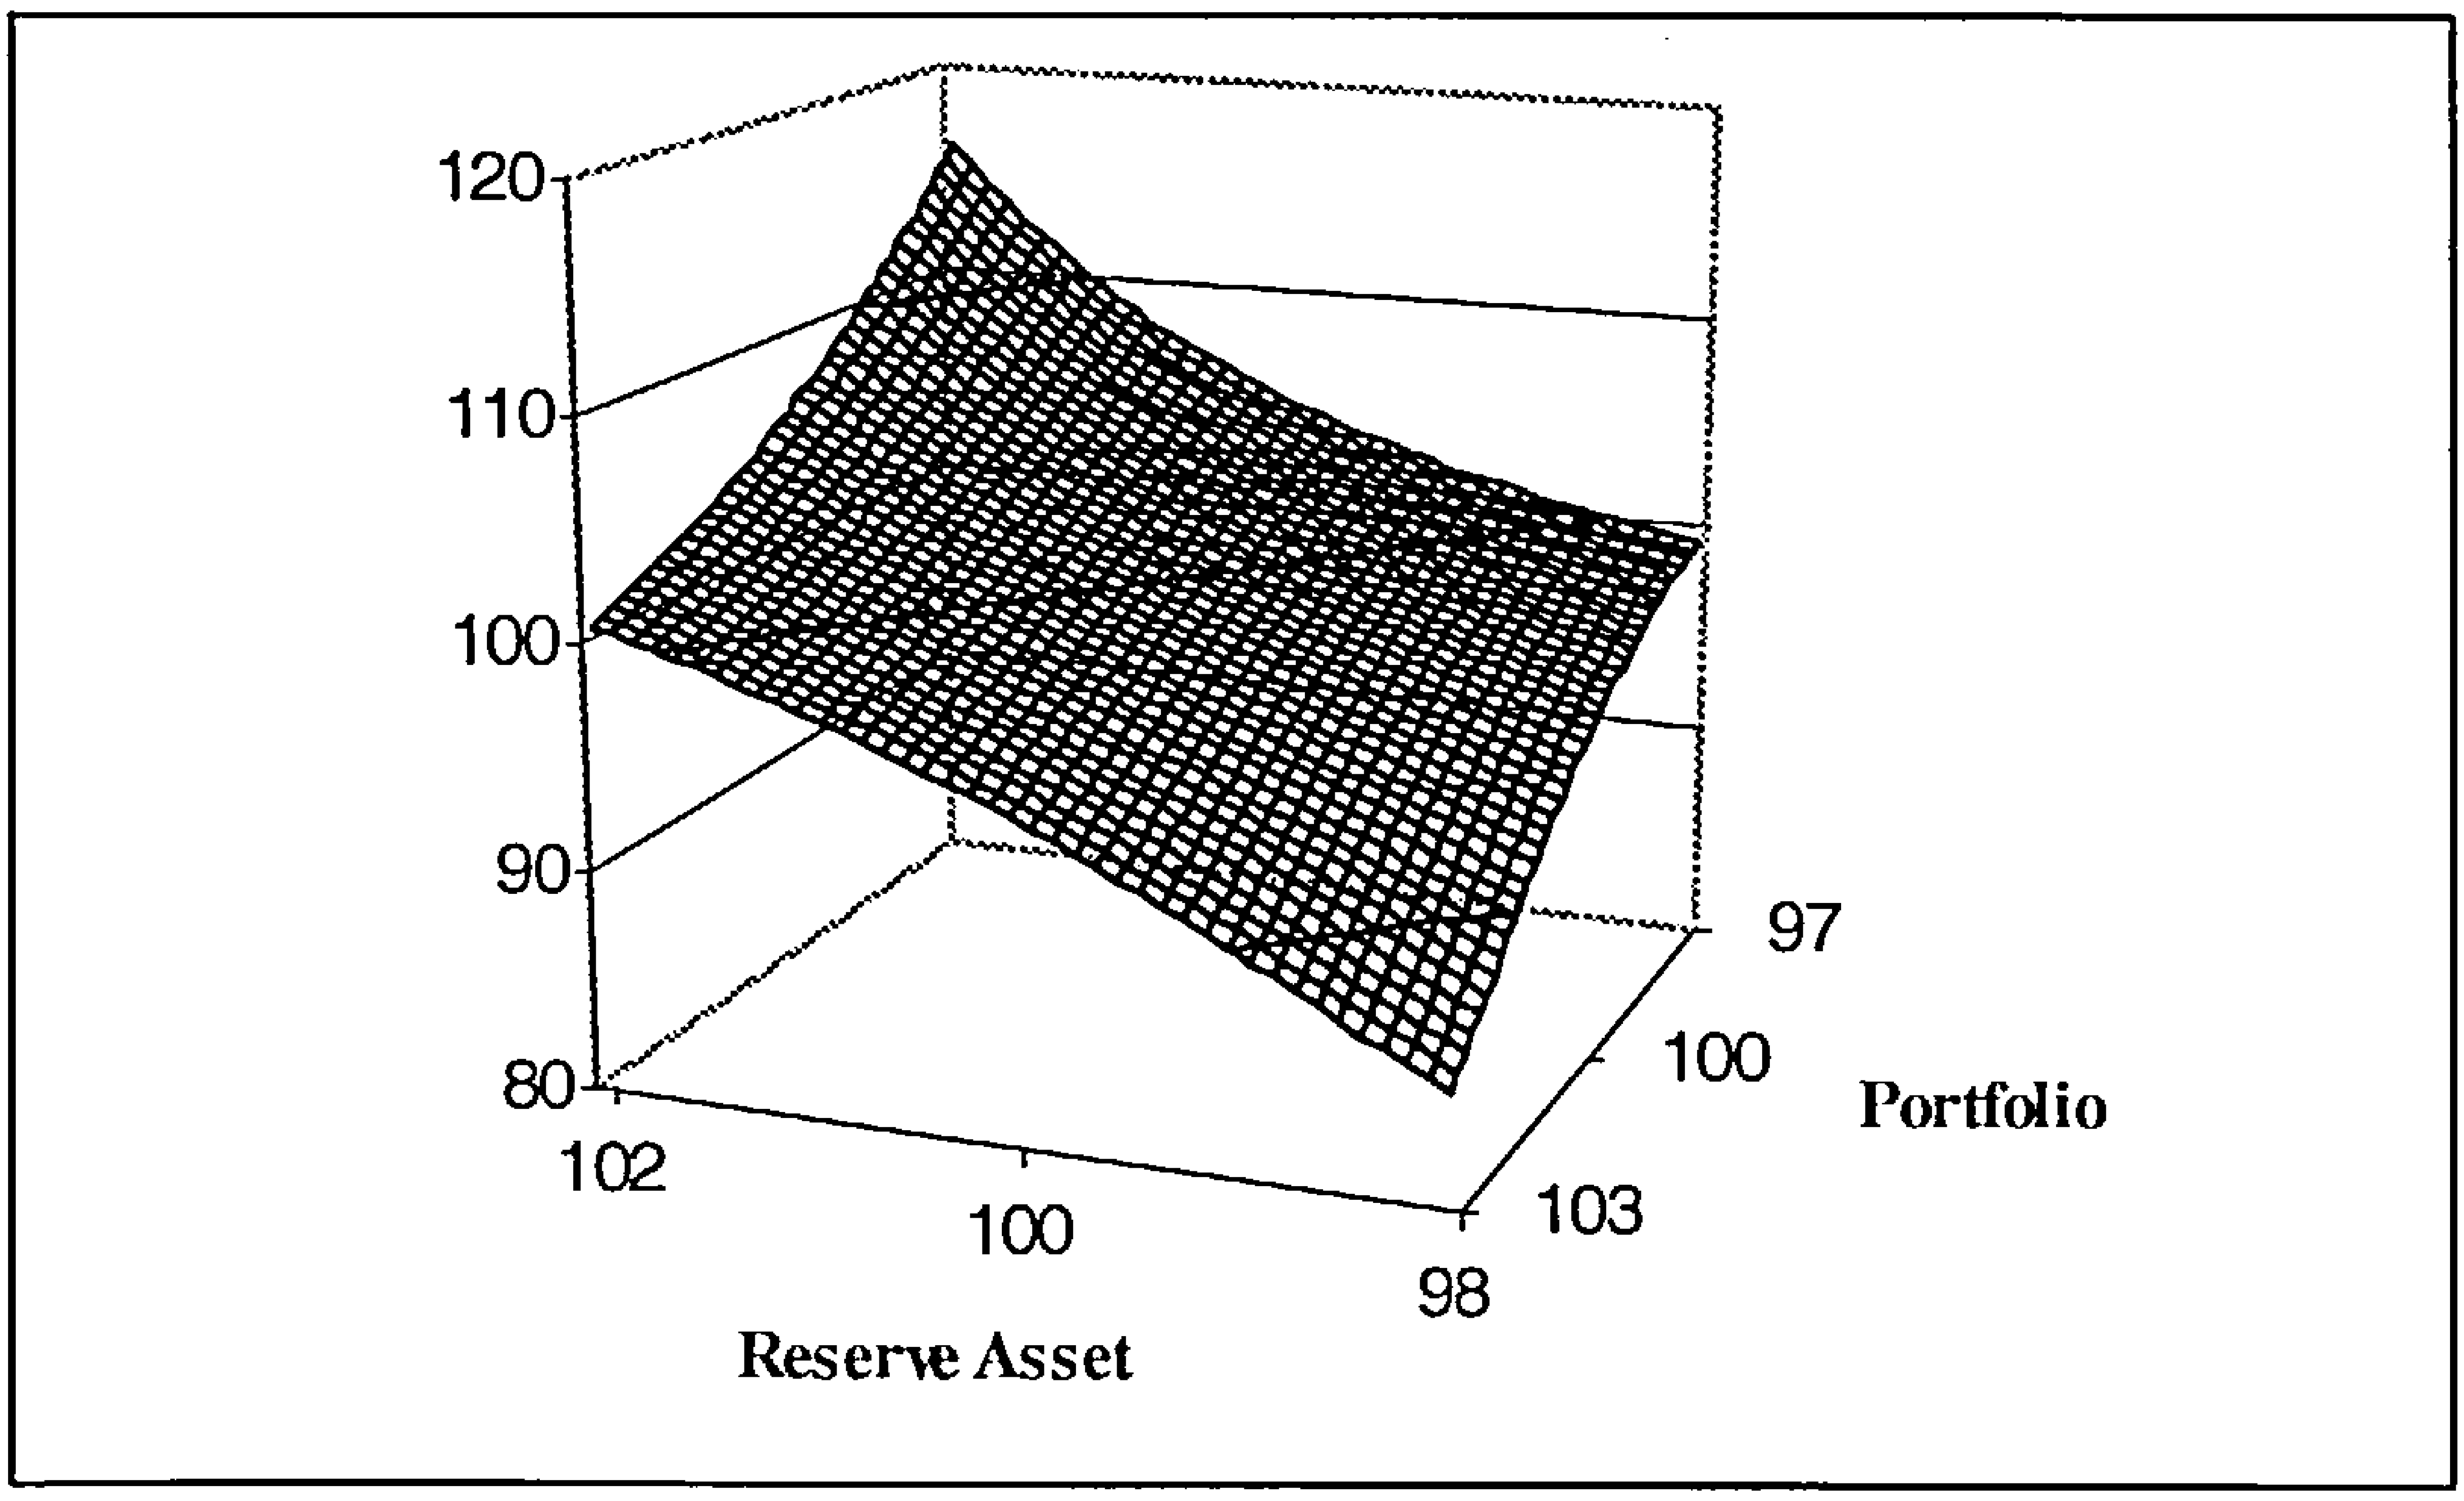

Figure 4.1. Target payoff synthetic fund 1, January 1998.

[46] Note that this makes the price of the reduction in correlation dependent on the risk premium on the investor's portfolio.

Figure 4.2 shows the evolution of the standard deviation of the synthetic fund return over the period January 2000 – February 2007[47], with the straight line representing the target value of 12%. The graph clearly shows that over the entire 7-year period the standard deviation of the synthetic fund return stayed close to the target value. There are a couple of small jumps, for example corresponding with the bursting of the NASDAQ technology bubble in March 2000, but these are quickly corrected over time.

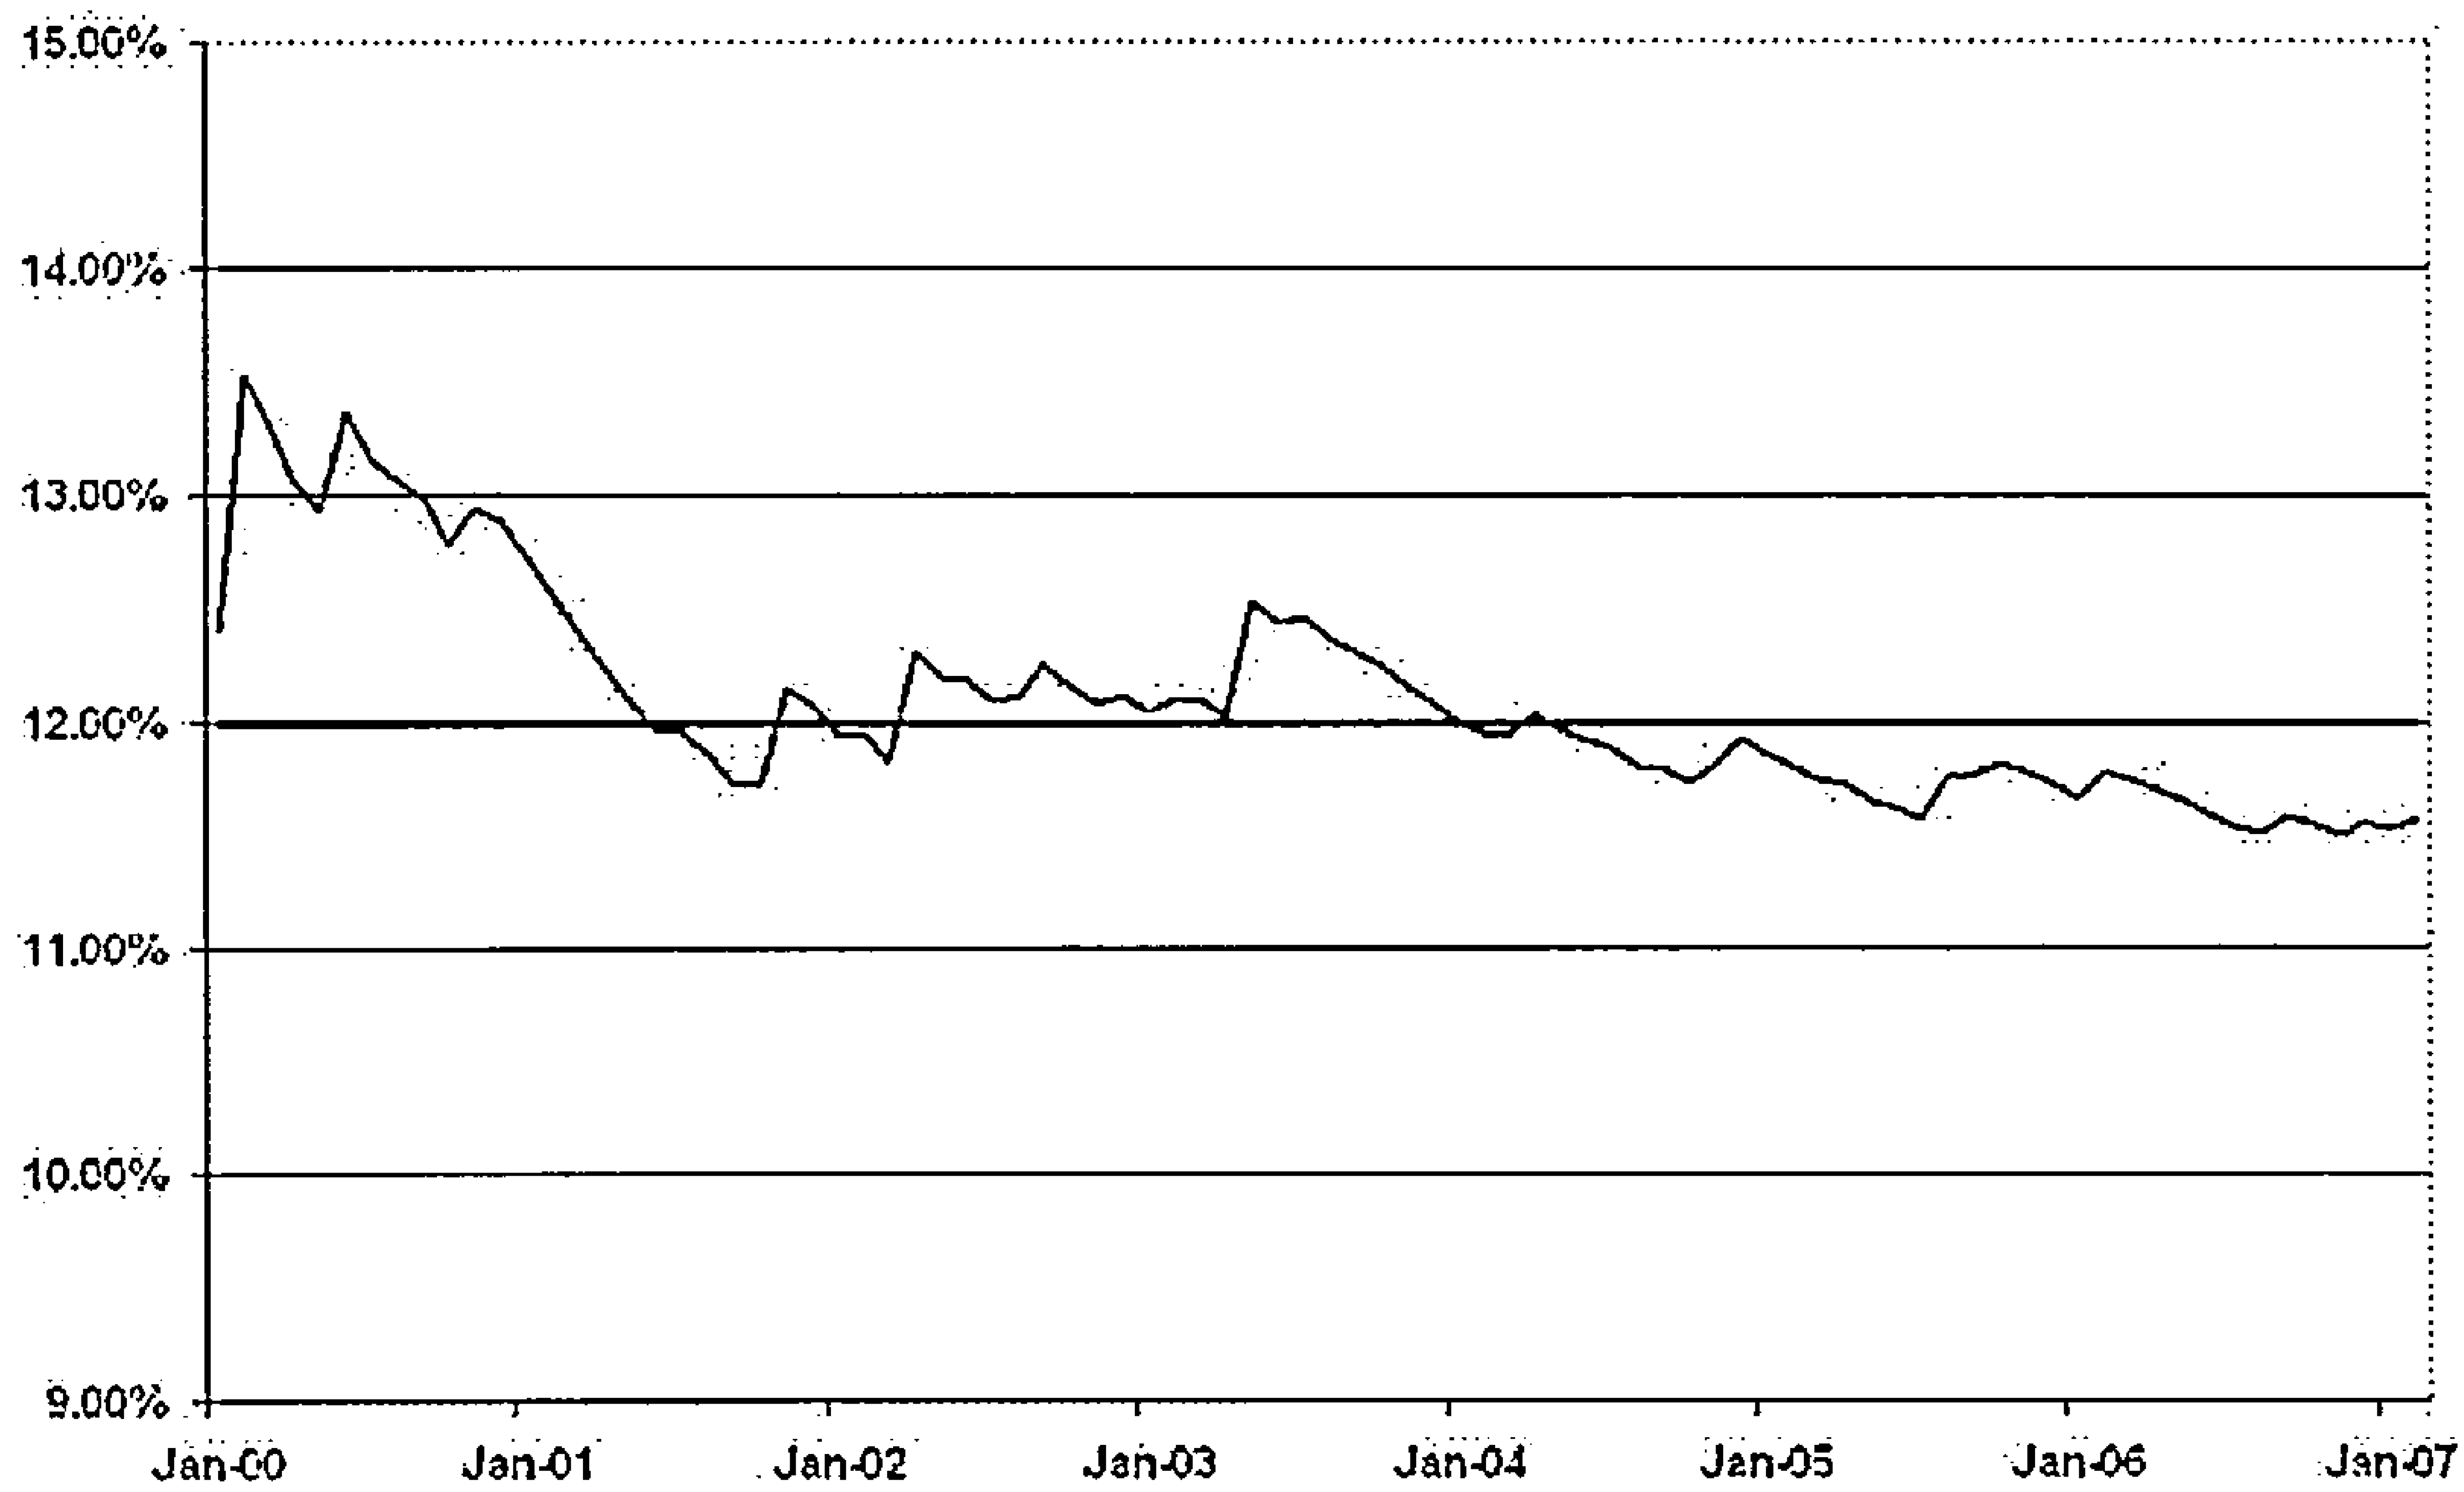

Figure 4.2. Standard deviation synthetic fund 1, January 2000 – February 2007.

Figure 4.3 shows the evolution of the skewness of the synthetic fund return over the same period, while Figure 4.4 shows the evolution of the correlation between the synthetic fund and the investor's portfolio. From these graphs it is clear that, as with the standard deviation, over the entire period studied the skewness and correlation of the synthetic fund return never deviated far from their target values. Given the at times tempestuous and erratic behaviour of markets, this is quite a remarkable achievement.

[47] Although the fund starts trading in January 1998, the graph in figure 1 (as well as the figures that follow) starts in January 2000 because to meaningfully estimate standard deviation we need at least 24 observations.

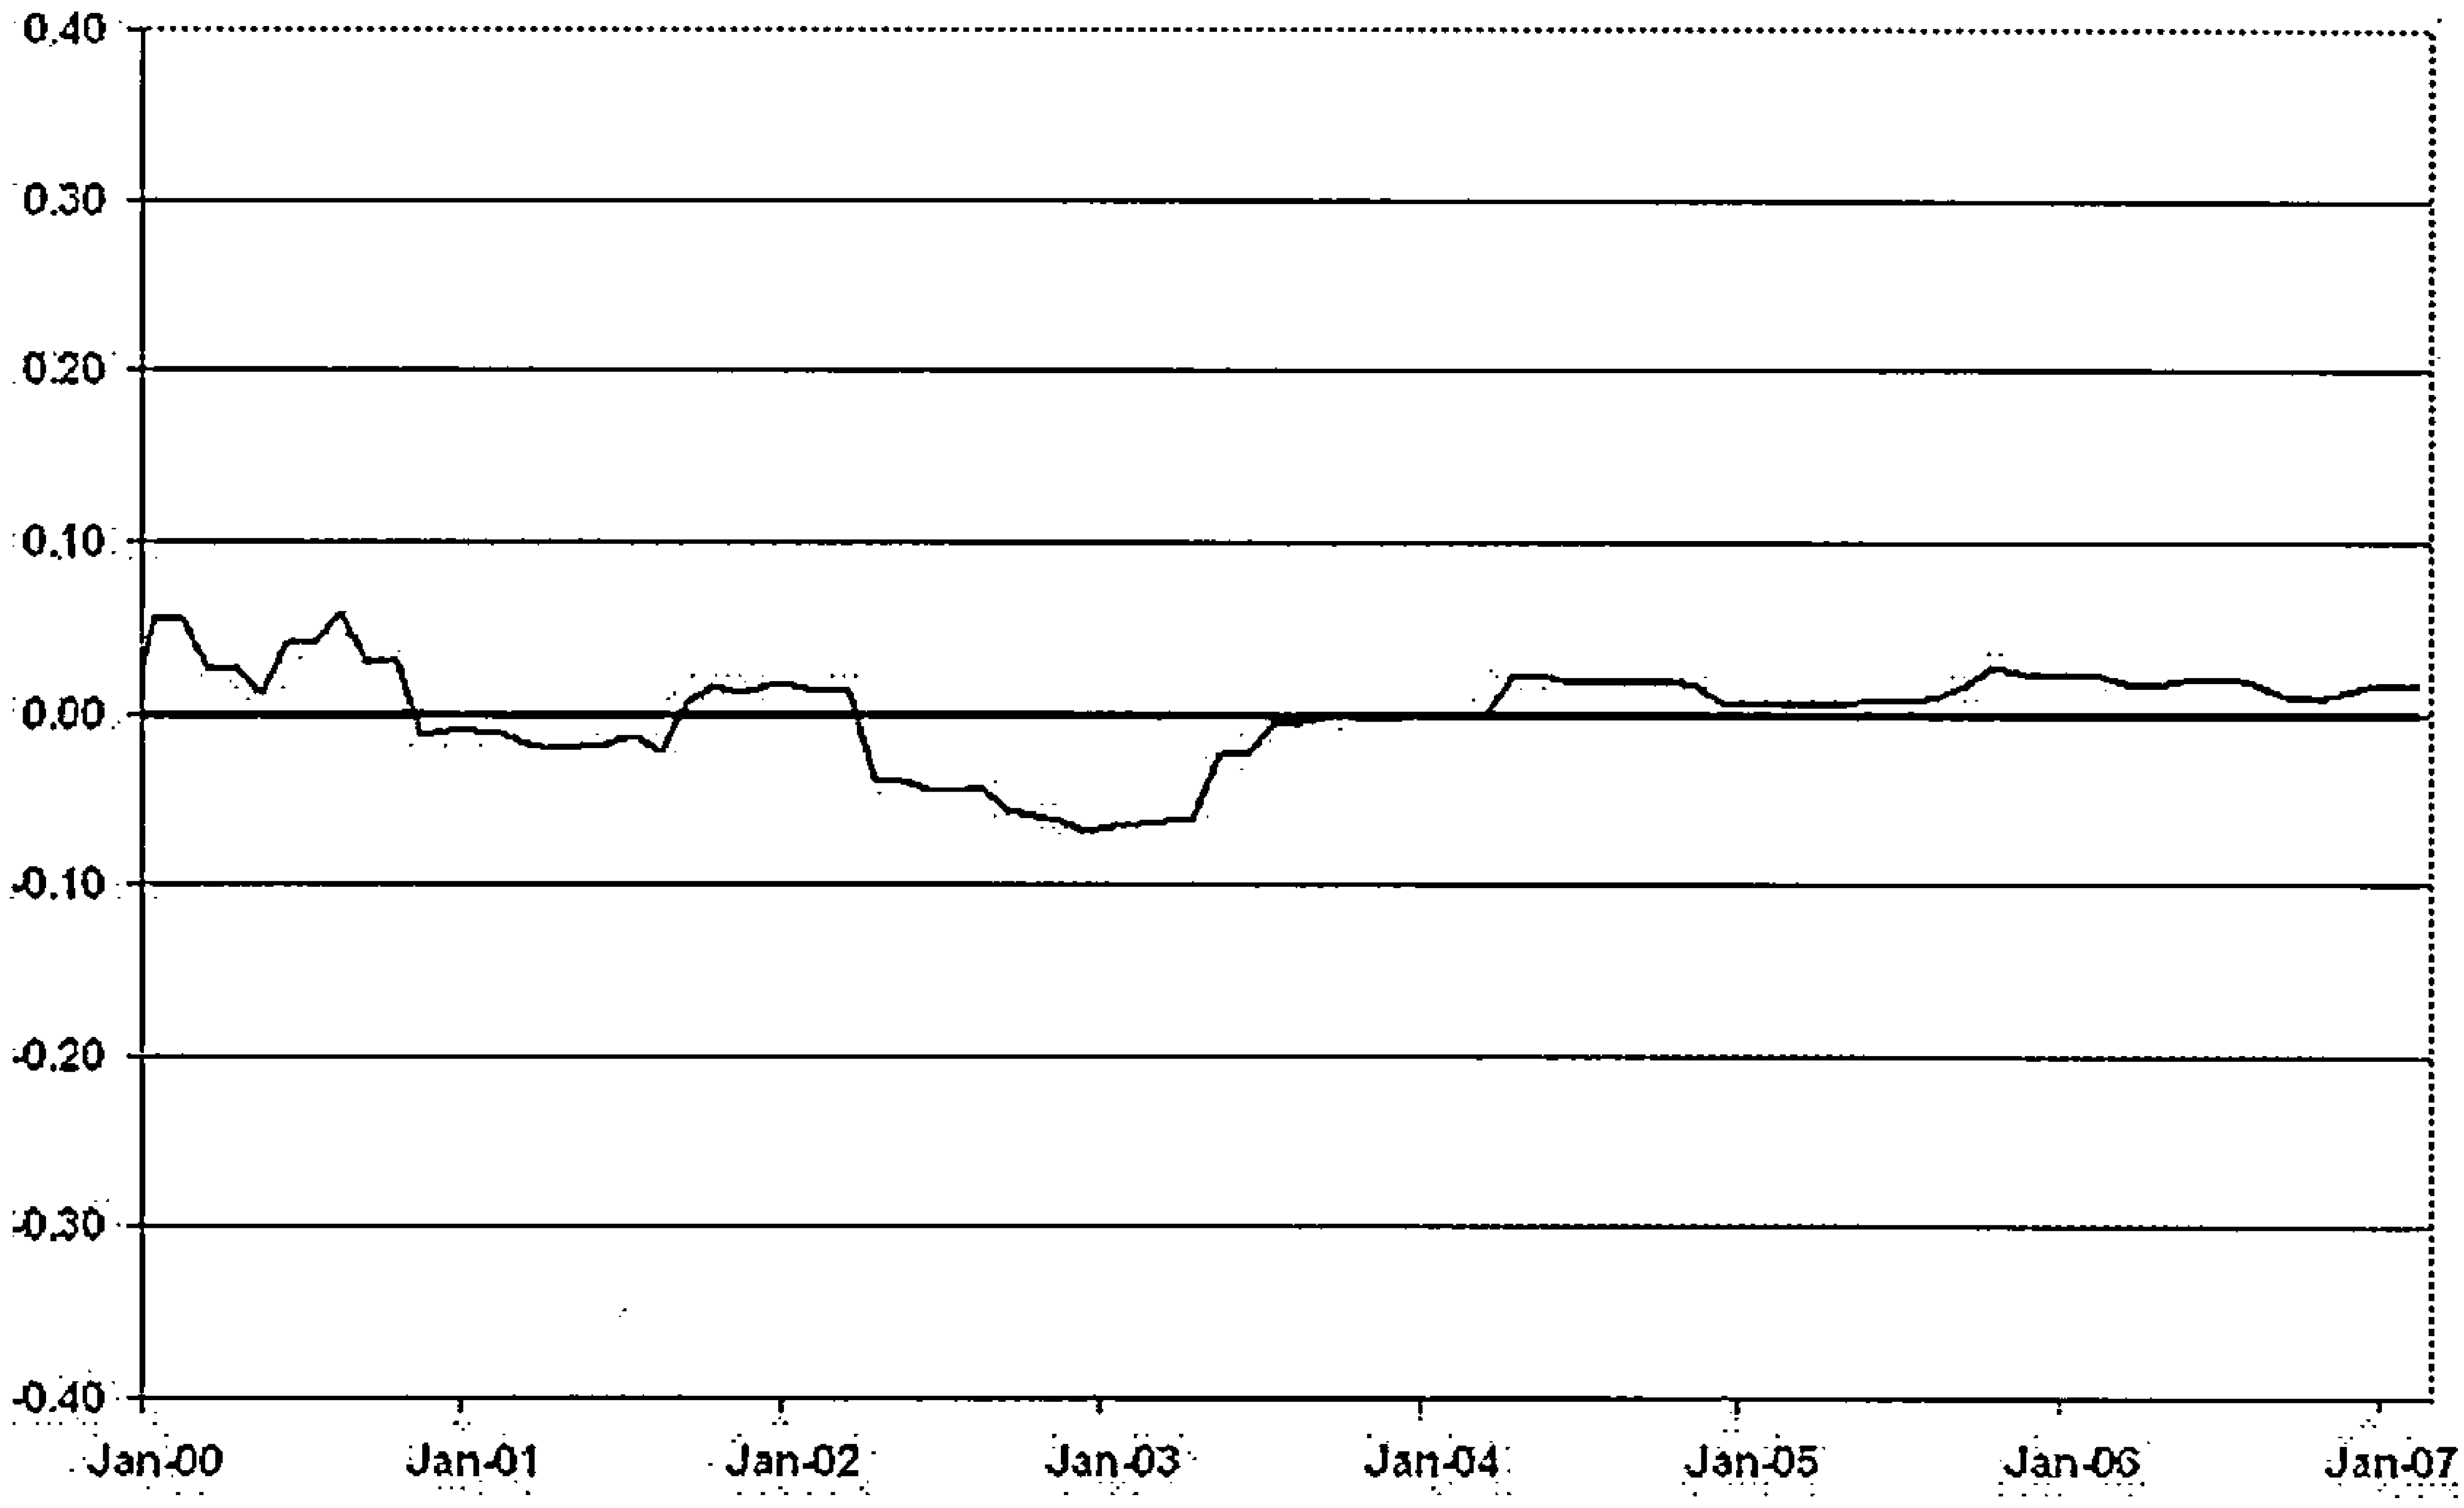

Figure 4.3. Skewness synthetic fund 1, January 2000 – February 2007.

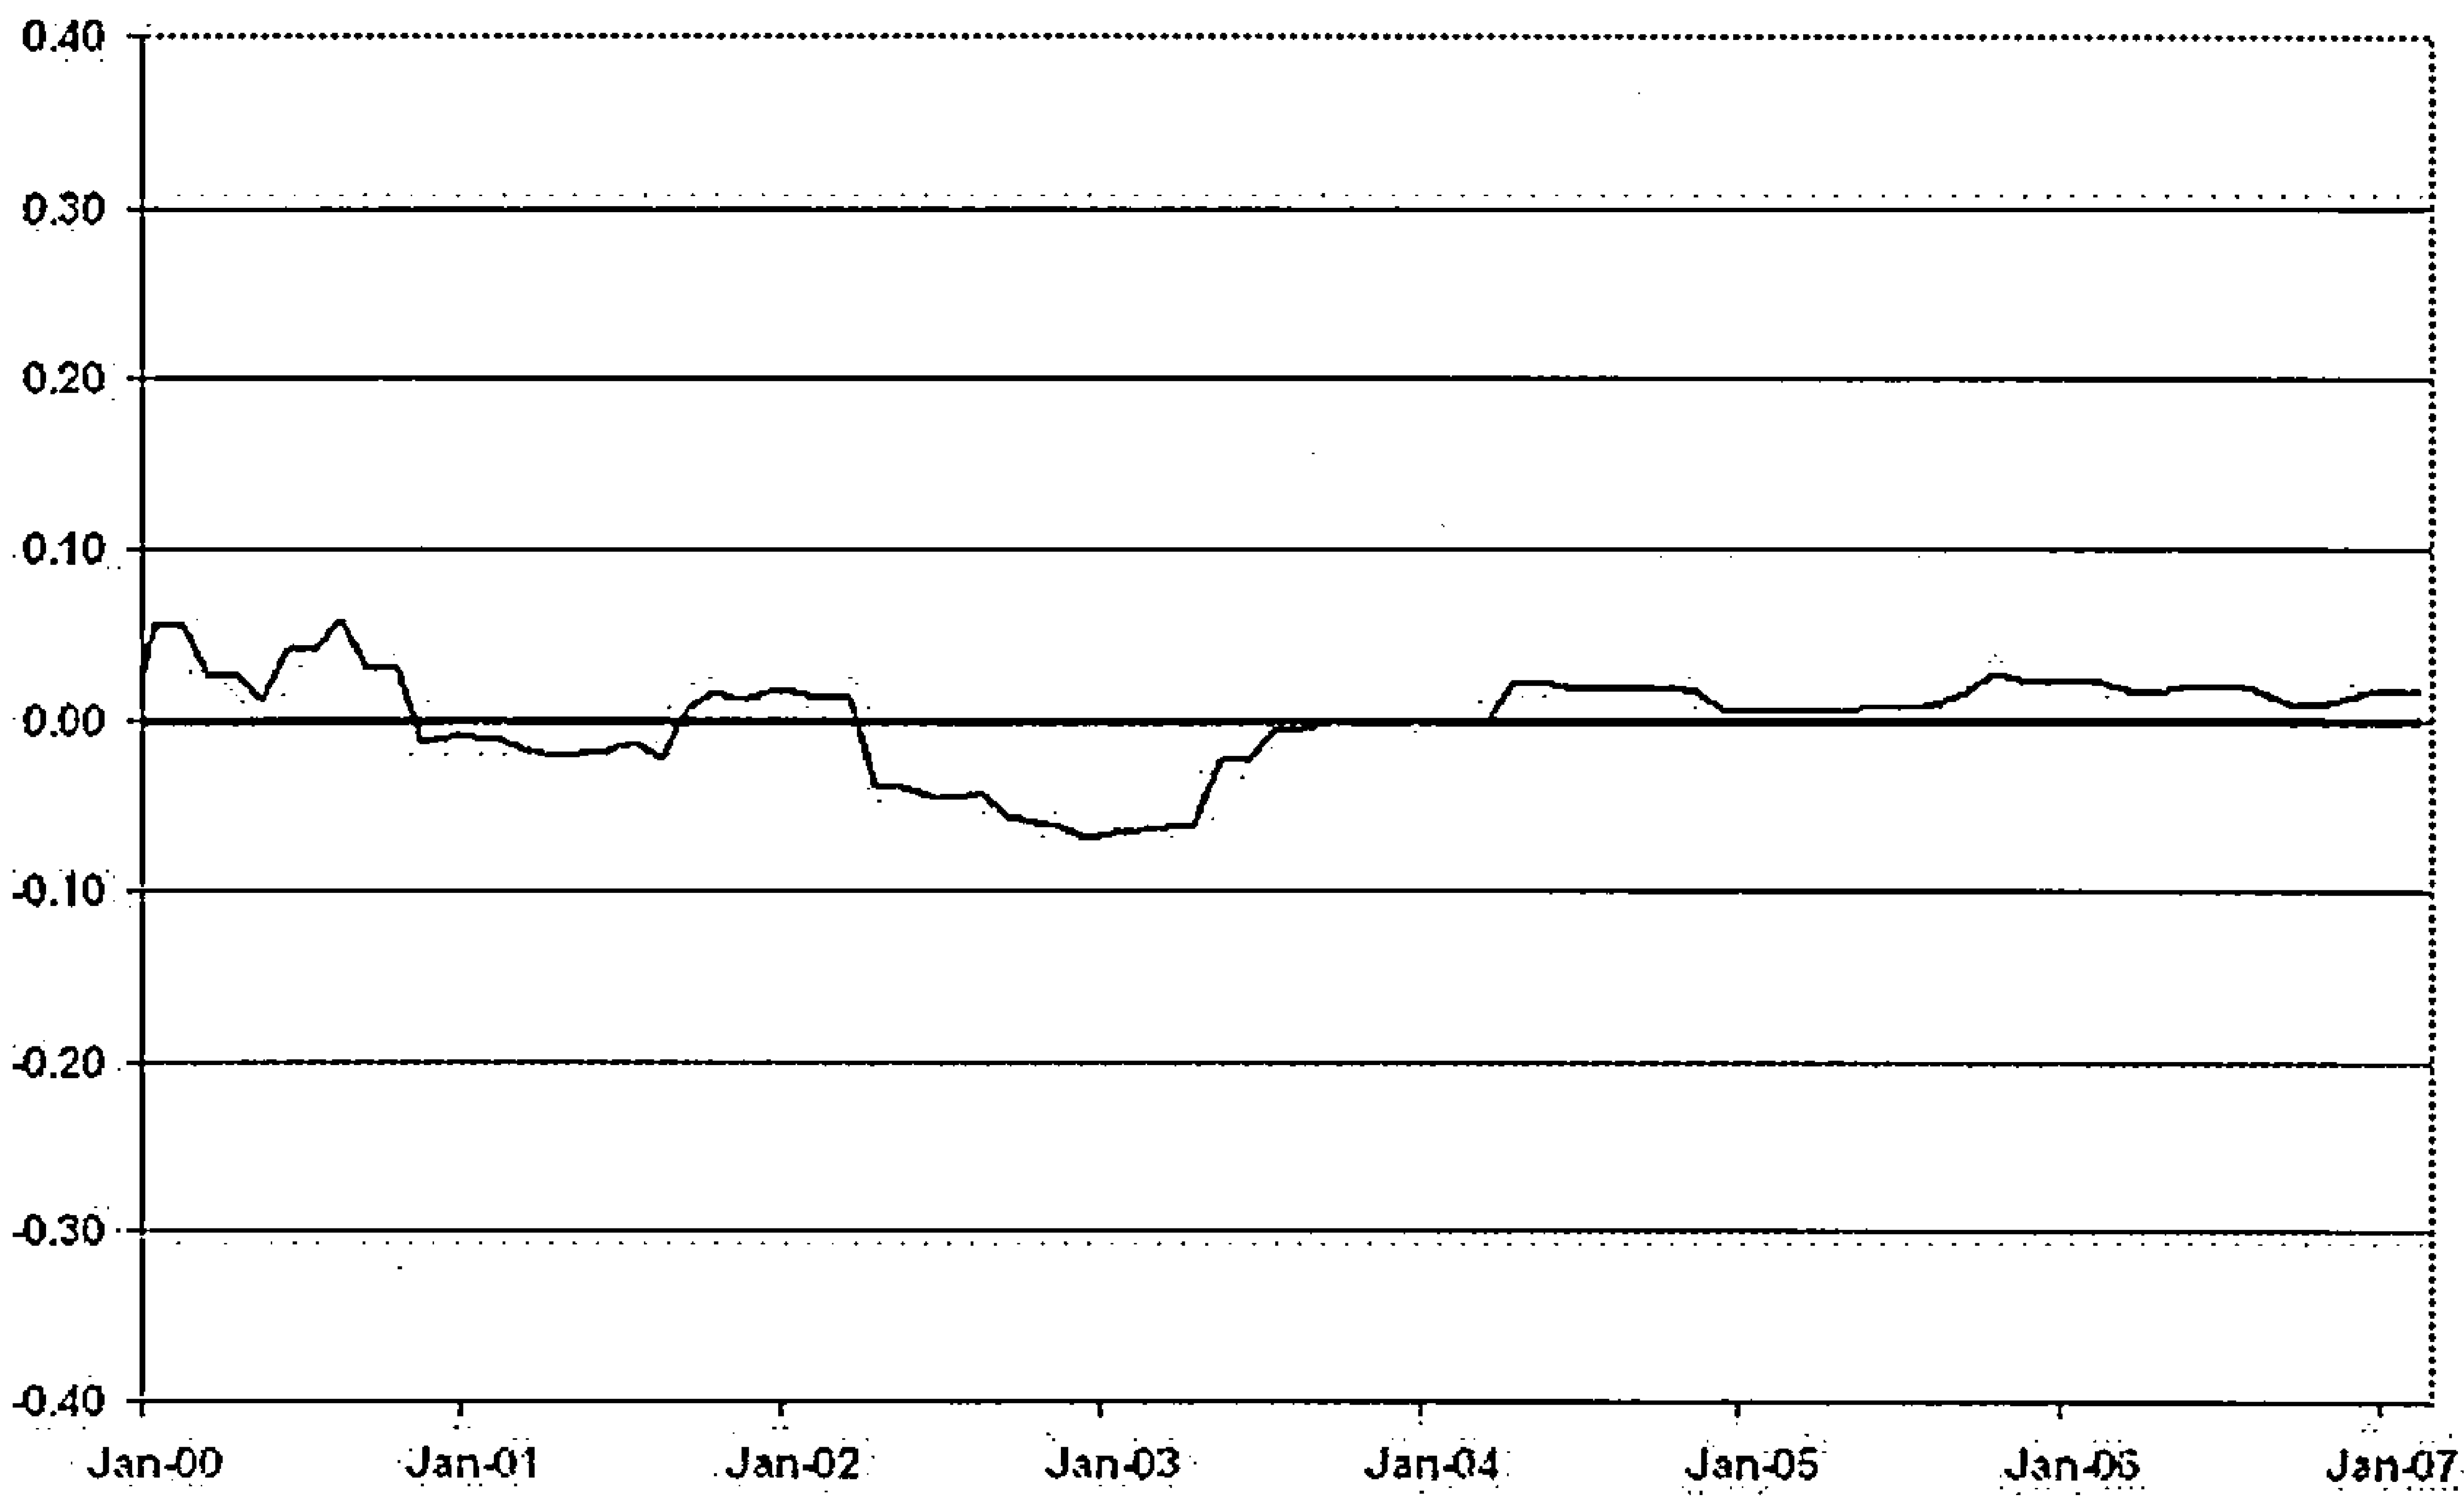

Figure 4.4. Correlation synthetic fund 1, January 2000 – February 2007.

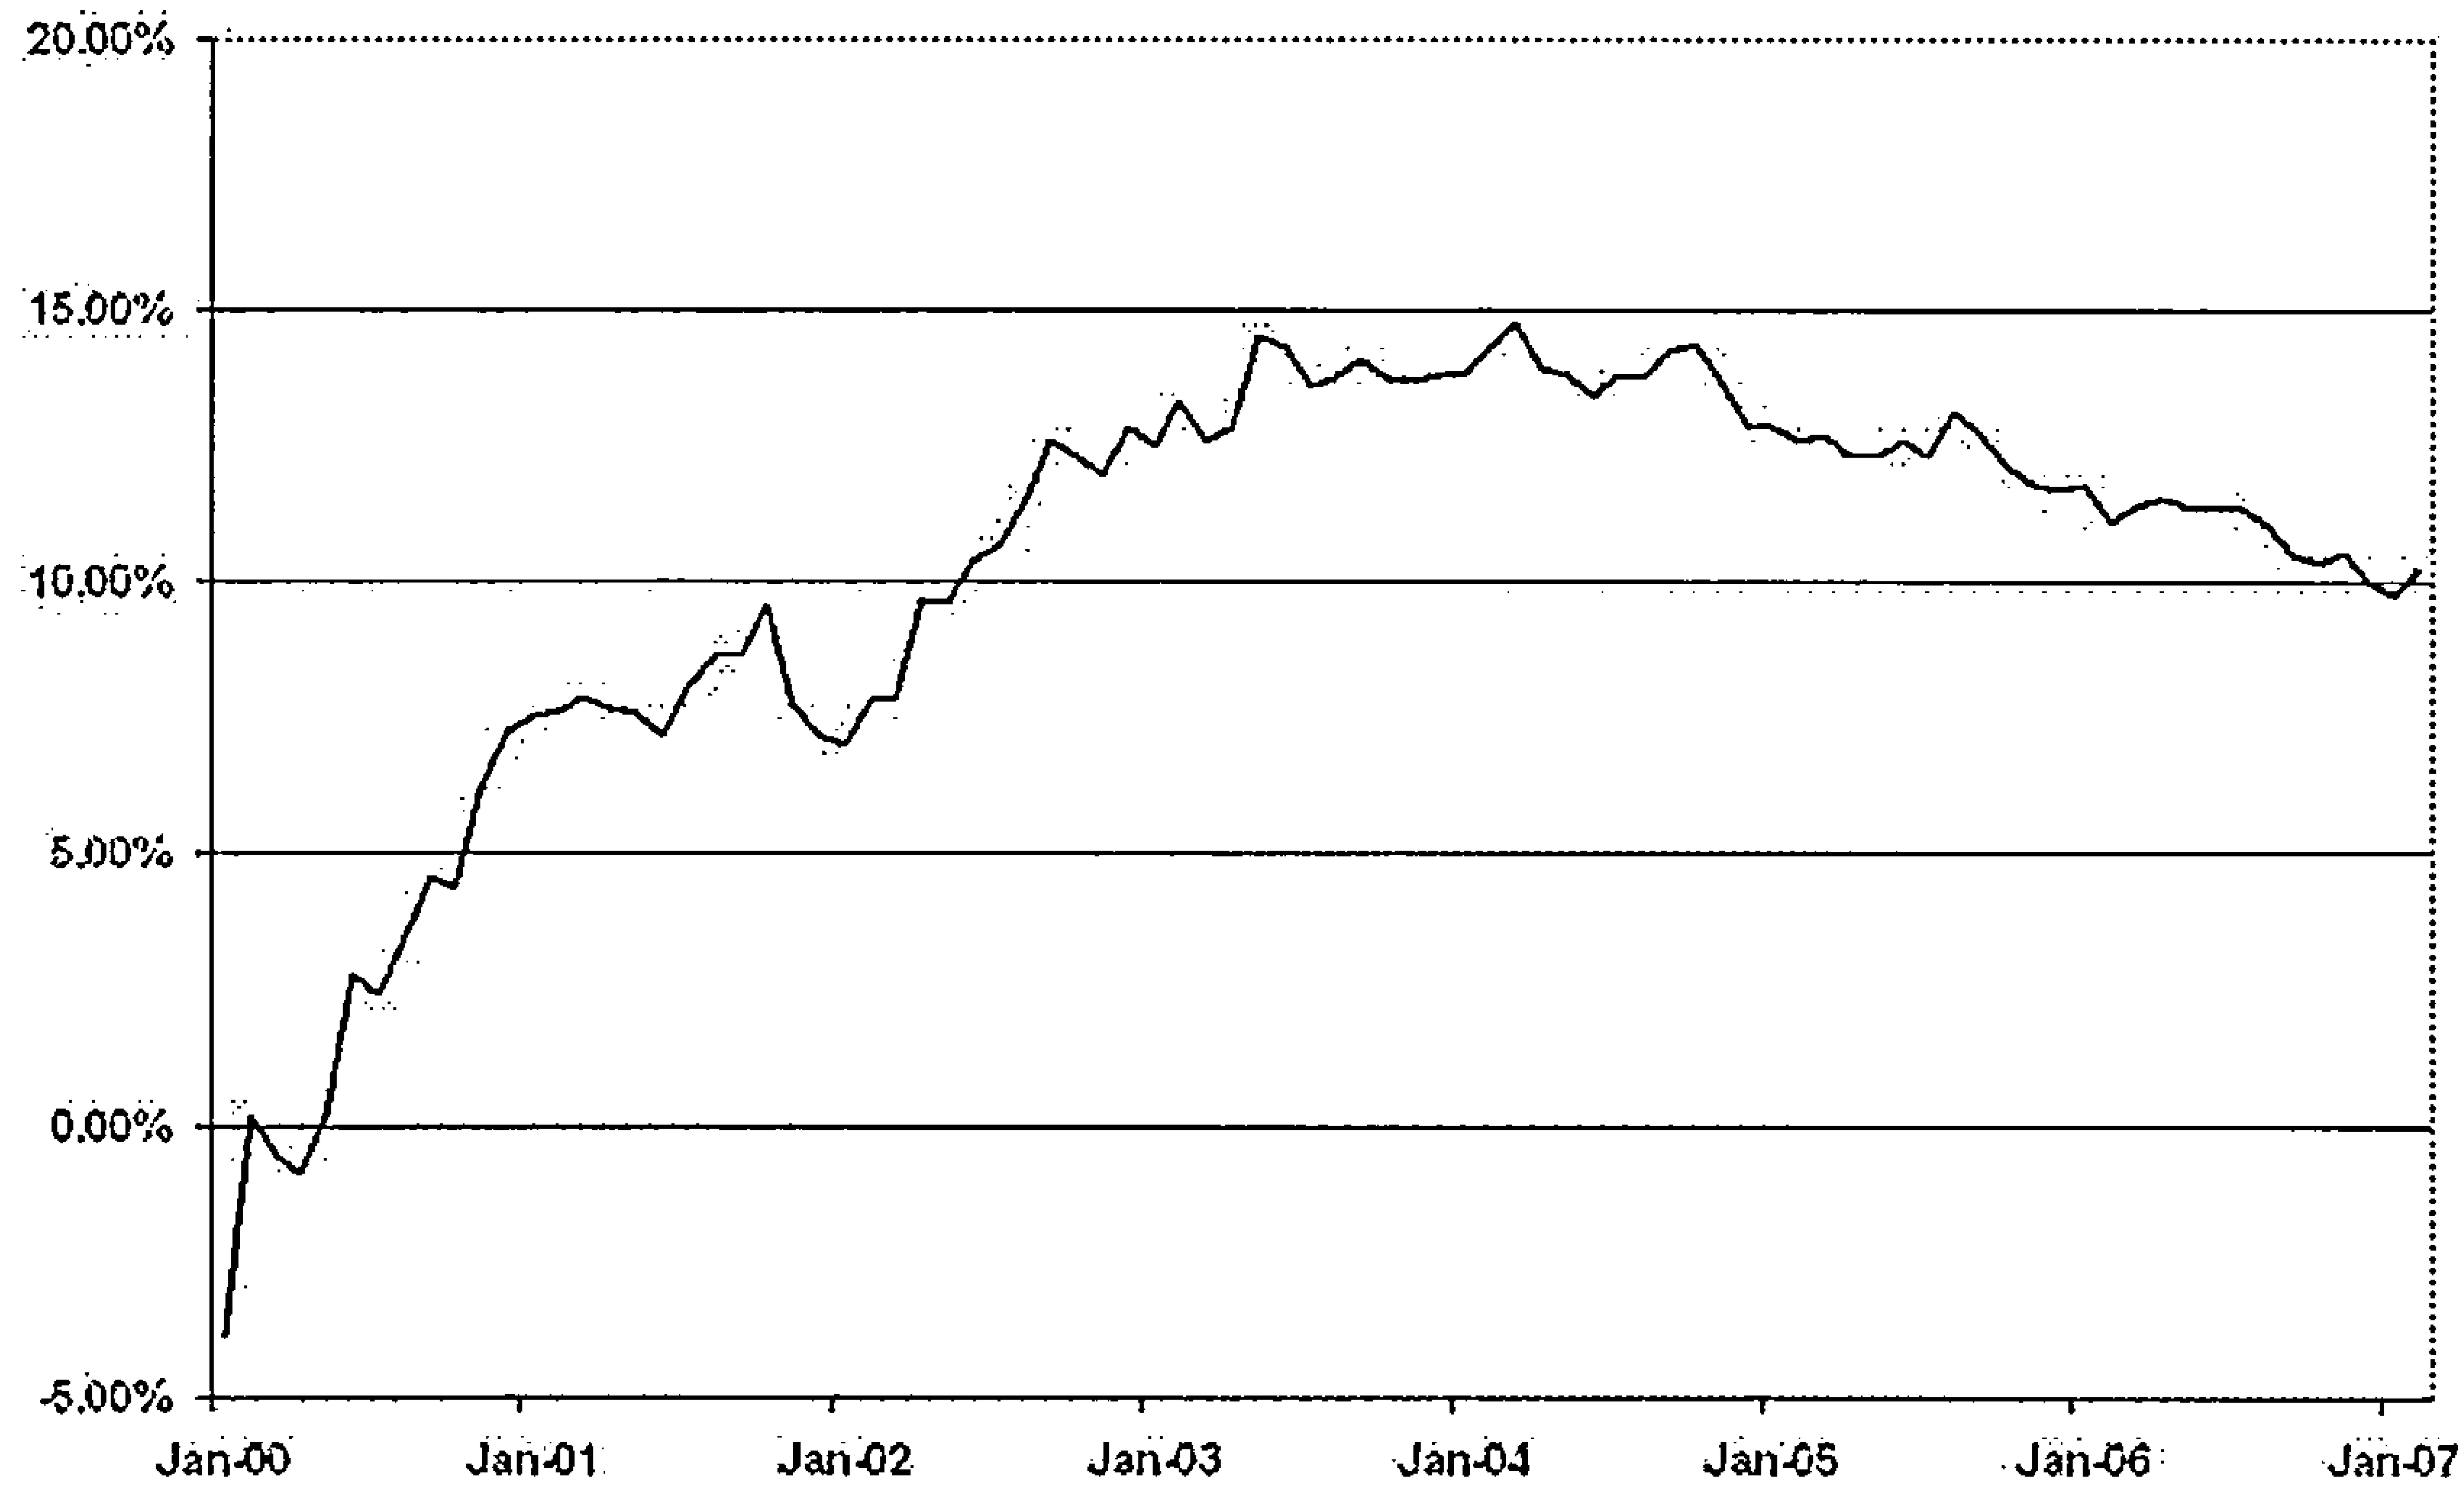

Figure 4.5. Average return synthetic fund 1, January 2000 – February 2007.

So far, we have not said anything about the synthetic fund's mean return. Given the relatively low correlation with the investor's portfolio, one might expect the fund to have provided a relatively low mean return. With the exclusion of the last couple of years, this is also the case with commodities for example. Figure 4.5 shows the evolution of the average return on the synthetic fund over the period January 2000 – February 2007. From the graph we see that over the period studied the fund's average return converged to around 10%, which is significantly more than what the average fund of hedge funds produced over the same period. We start with a negative average return in 1998-1999, which was quite a troublesome period with crisis in Russia followed by the near-collapse of LTCM. When interpreting the beginning of the graph we have to keep in mind that the fund's track record only starts in January 1998. By January 2000 we therefore only have 24 observations available, meaning that the negative returns experienced in October 1998 and the following months have a relatively strong impact on the mean.

	50/50	Fund 1	10% Fund	20% Fund	30% Fund
Mean Return	**7.07%**	**10.16%**	**7.39%**	**7.71%**	**8.02%**
Volatility	**7.65%**	**11.57%**	**7.02%**	**6.61%**	**6.47%**
Skewness	**-0.21**	**0.10**	**-0.15**	**-0.07**	**0.03**
Sharpe ratio	**0.44**	**0.56**	**0.52**	**0.60**	**0.66**

Table 4.2. Properties overall portfolio with varying allocations to synthetic fund 1, January 1998 – February 2007.

Since the synthetic fund is meant to be a diversifier for a larger, traditional portfolio, its performance should be evaluated in a portfolio context as well. One simple way to do so is by looking at the performance of the investor's portfolio with and without an allocation to the synthetic fund. Table 4.2 shows the properties of the investor's original 50% stocks – 50% bonds portfolio and the synthetic fund, as well as various mixes of the original portfolio and the fund. Comparing the fund with the investor's original portfolio, we see that over the period studied the fund produced a higher mean return, but with higher volatility. When mixing the original portfolio with the fund, due to their zero correlation, the volatility of the resulting portfolios drops substantially, producing Sharpe ratios that far exceed that of the investor's original portfolio. This confirms the attraction of the synthetic fund as a portfolio diversifier. *Although it may not make for the most attractive stand-alone investment, in a portfolio context the synthetic fund certainly delivers.*

The statistical properties of the synthetic fund returns over the period January 1998 – February 2007 have been very much in line with the target values set out at the start, but how much trading was required to accomplish this? Since futures have relatively short maturities and we are looking at monthly returns, there are three reasons for trading in our synthetic fund: (1) normal day-to-day exposure adjustment during the month, (2) resetting of all positions at the start of every new month, and (3) periodic rolling over of the nearby futures contract. Taking all three together, the second column in Table 4.3 shows the average daily trade size for the above synthetic fund over the period January 1998 – February 2007, assuming an initial fund value of $100 million. The third column shows the average daily trade size excluding the periodic rollovers. This gives an indication for the required trading volume if, instead of the nearby contract, we were to trade longer-dated futures contracts. From Table 4.3, we see that on average managing a $100m synthetic fund does not require very much trading at all. The numbers of contracts in Table 4.3 are only a very small fraction of the typical daily market volume. This confirms that liquidity problems are highly unlikely, even when the fund size was a lot larger than $100m.

Futures Contract	Average Daily Trade Size (Number of contracts)	Average Daily Trade Size (Excl. periodic rollover)
S&P 500	13	11
Russell 2000	8	6
Eurodollar	36	33
5-year Note	77	73
10-year Note	62	58
T-Bond	44	41
GSCI	28	25

Table 4.3. Average daily trade size synthetic fund 1, January 1998 – February 2007.

4.1.2 Fund 2

The second synthetic fund is similar to fund 1, except that apart from zero correlation with the reference portfolio we now also aim for a substantial degree of positive skewness in the fund's returns. Figure 4.6 shows the payoff function as per January 1998. Not surprisingly, it is quite similar to that of fund 1. The desired fund payoff is again an increasing function of the reserve asset and a declining function of the investor's portfolio. The payoff for a combination of a low value of the investor's portfolio and a high value of the reserve asset is much higher than before, however. The reason why especially this corner has been lifted is that the investor's portfolio exhibits some negative skewness itself, which makes it easiest to deliver the desired positive skewness in this way.

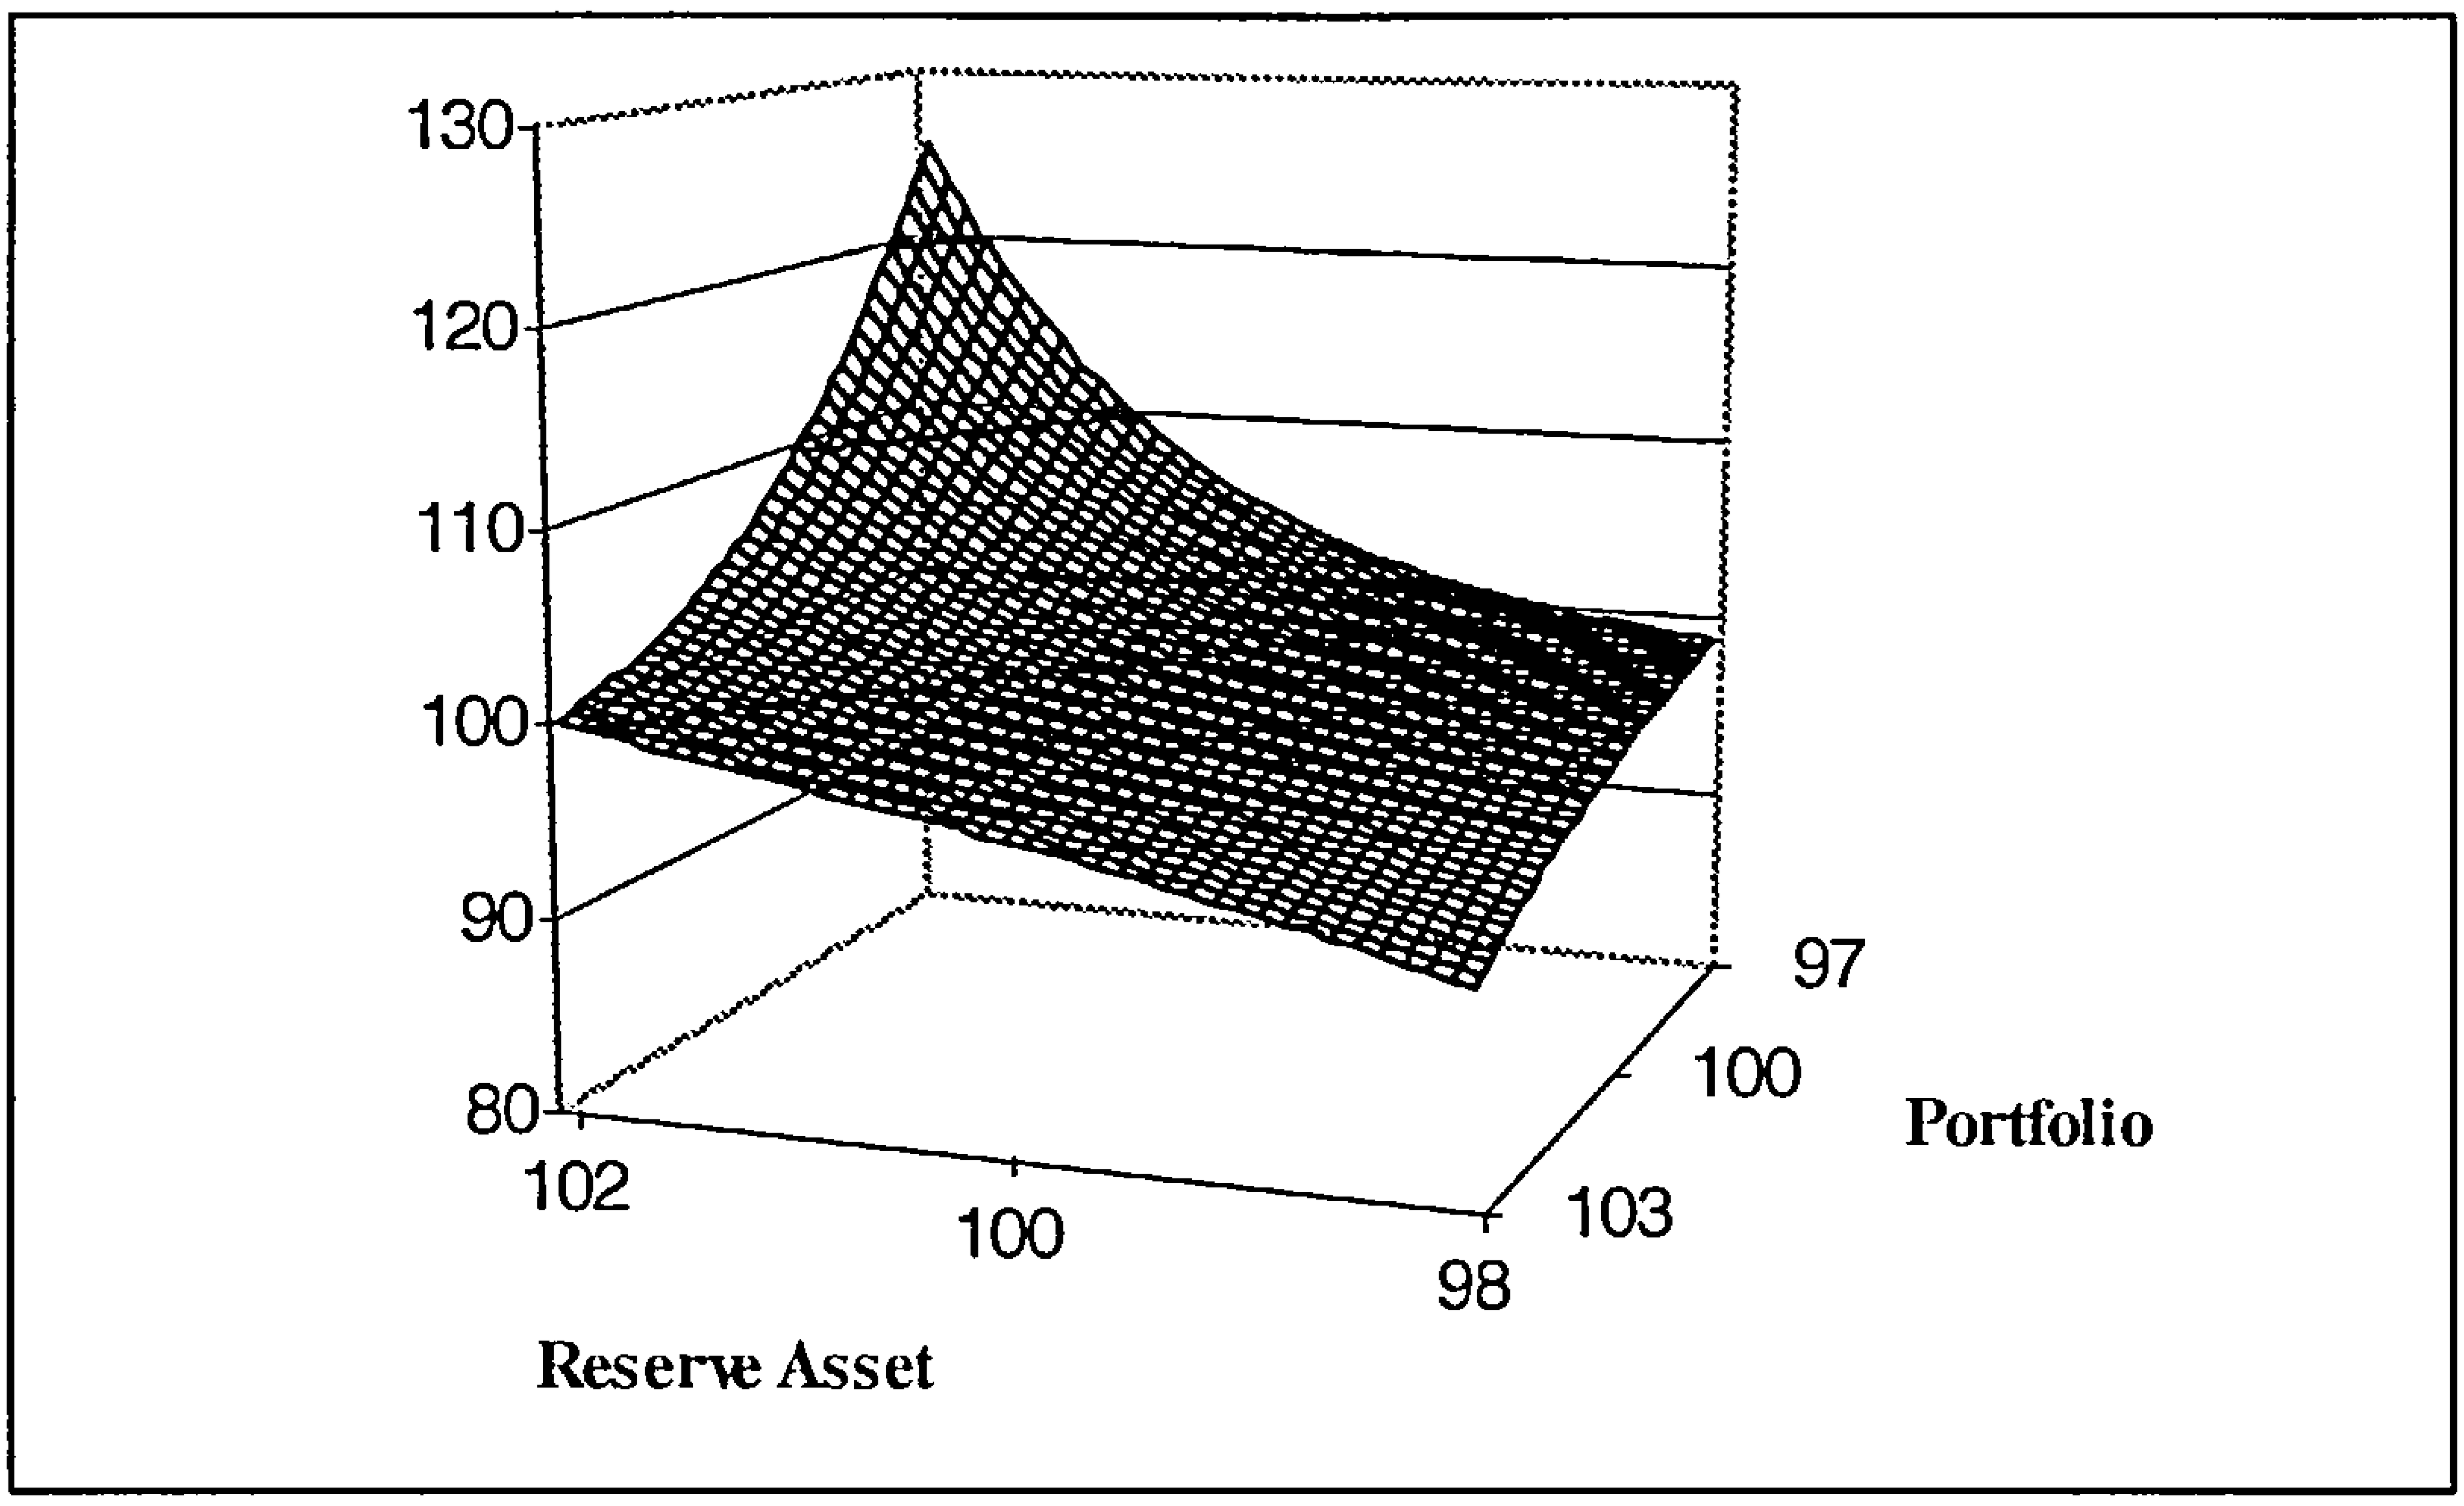

Figure 4.6. Target payoff synthetic fund 2, January 1998.

Since the volatility and correlation targets for fund 2 are the same as for fund 1, the volatility and correlation results are very similar as well. For brevity we therefore do not report these here. Figure 4.7, however, shows the evolution of the skewness of the synthetic fund return. It shows that in 1998-2000 we lose some skewness due to the equity bear market, but in 2002 we regain that thanks to the equity bull market. Over the entire period studied, the skewness of the synthetic fund return never deviates far from its target value. Figure 4.8 shows the evolution of the average return on the synthetic fund. The graph looks very similar to that in Figure 4.5, except that it begins just a little more negative and converges to a slightly lower level in the long run, which can be interpreted as the price paid for the improvement in skewness.

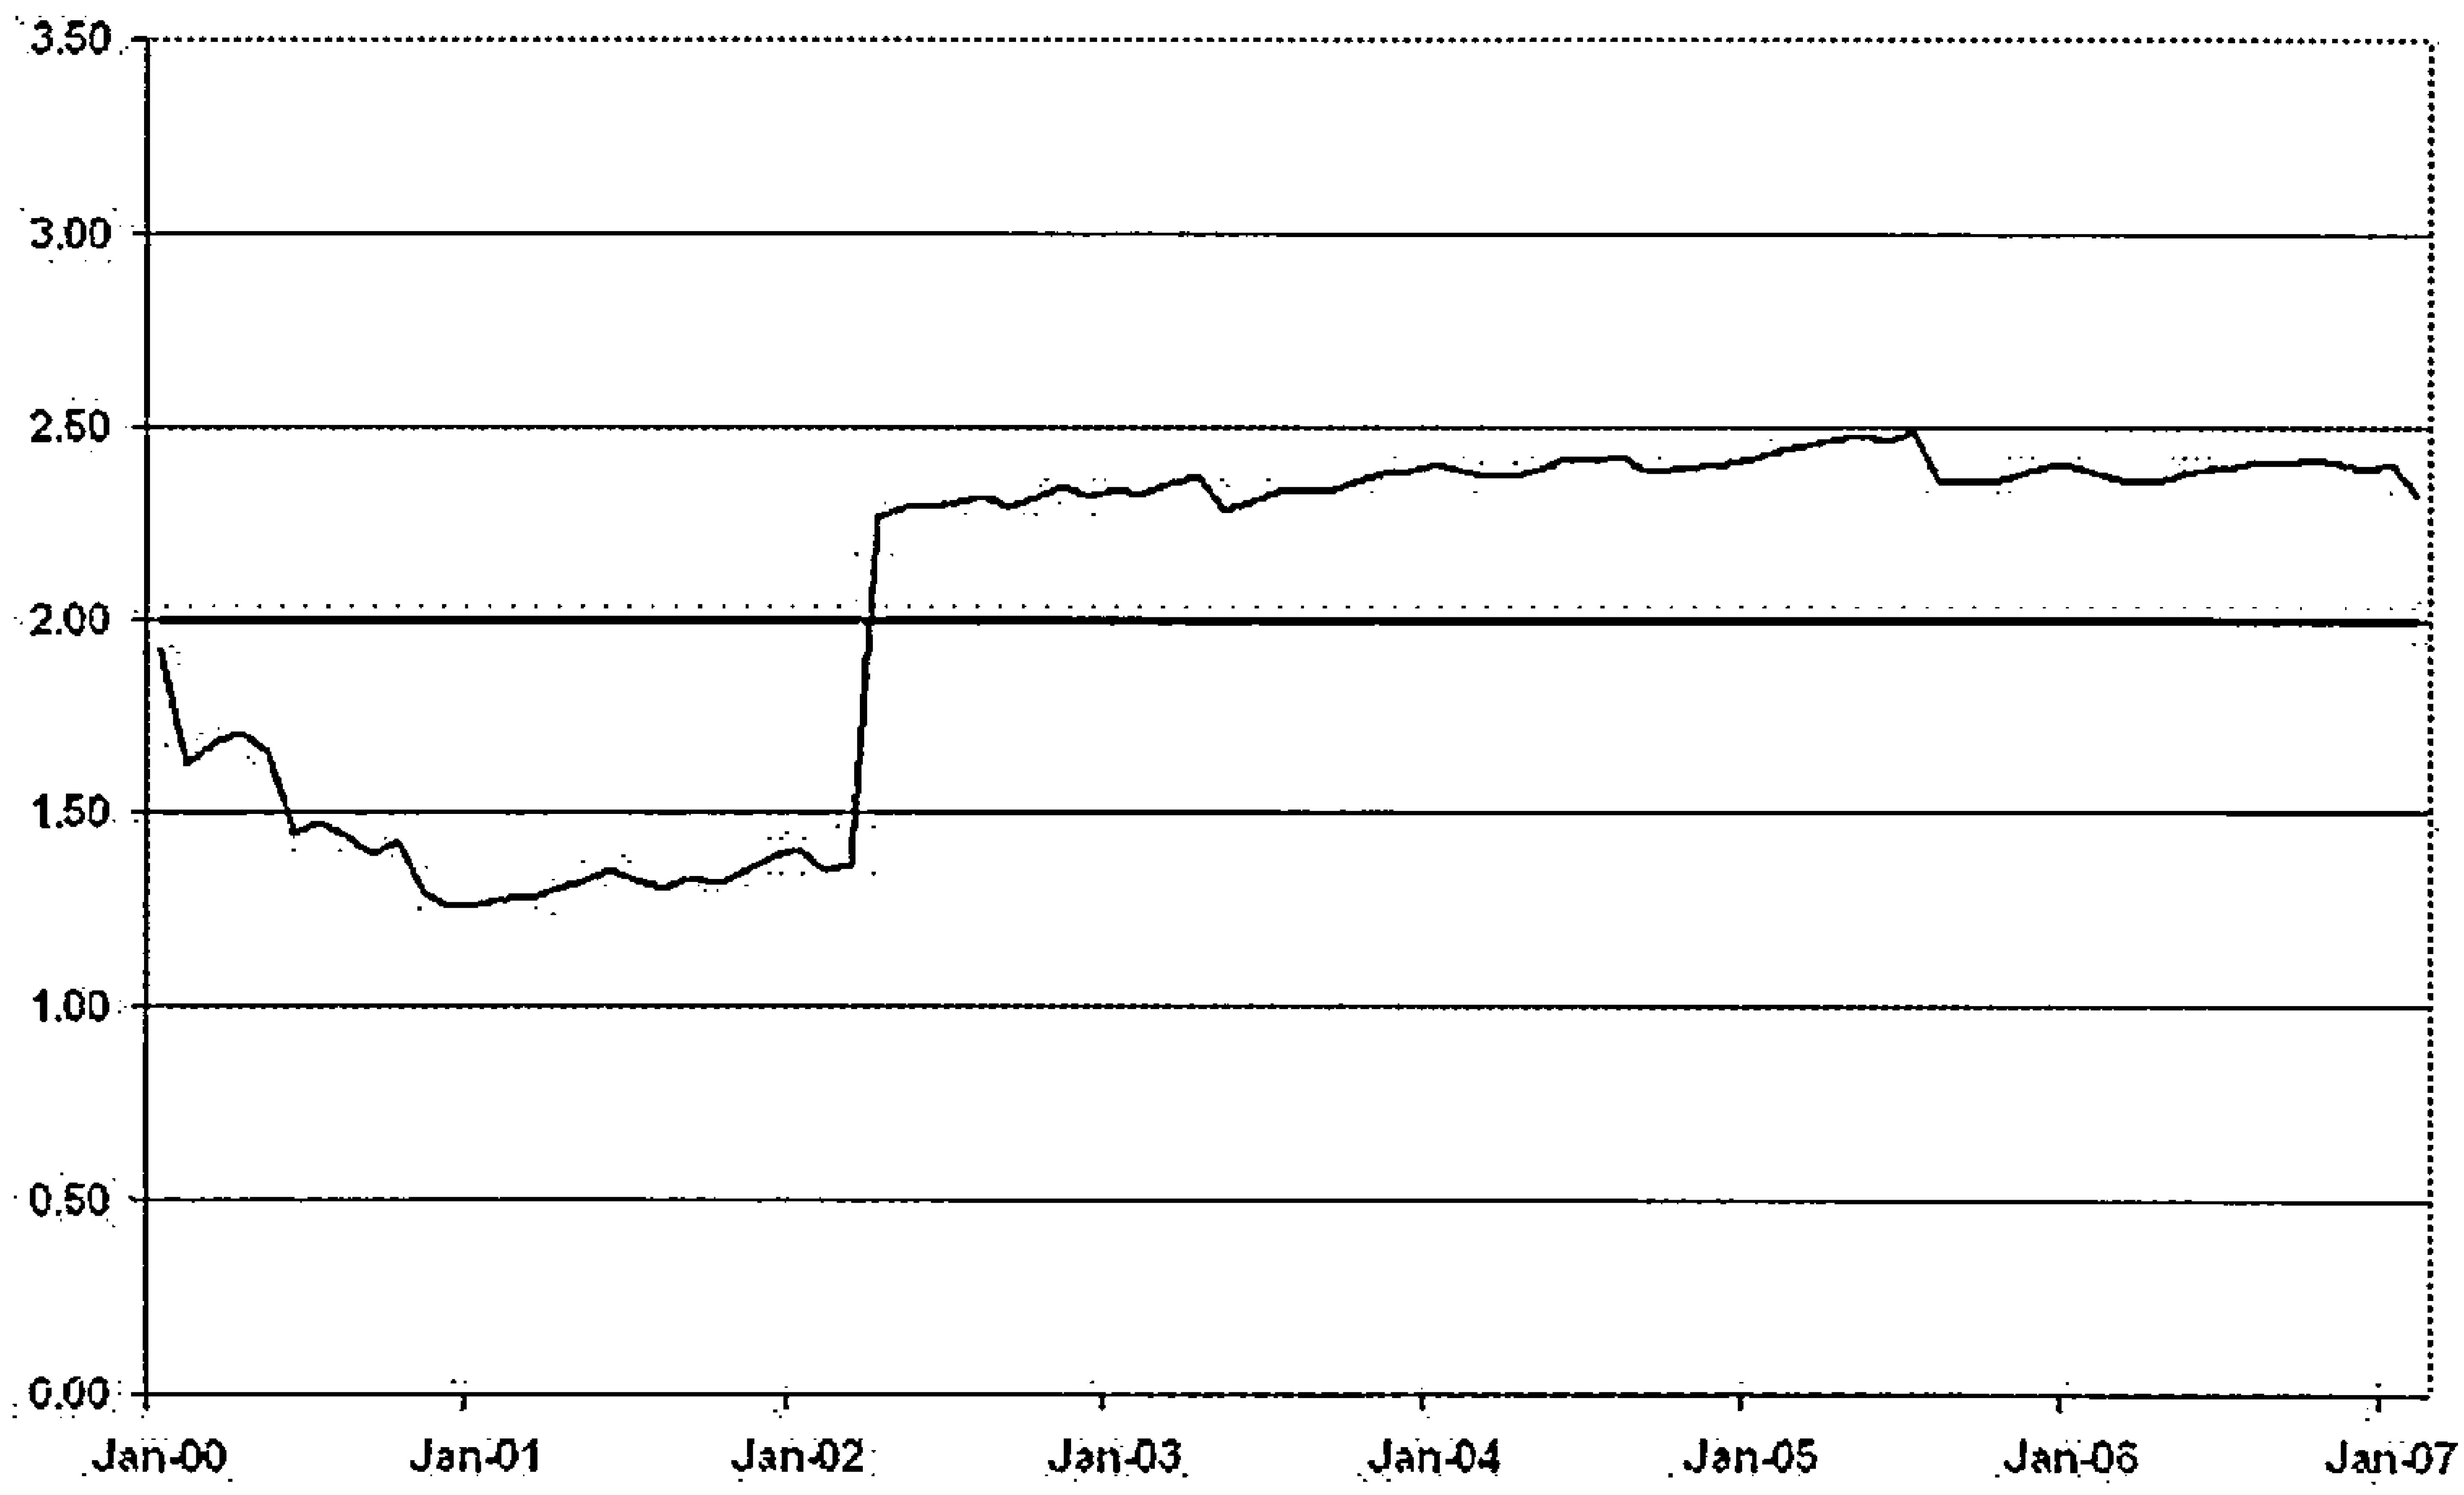

Figure 4.7. Skewness synthetic fund 2, January 2000 – February 2007.

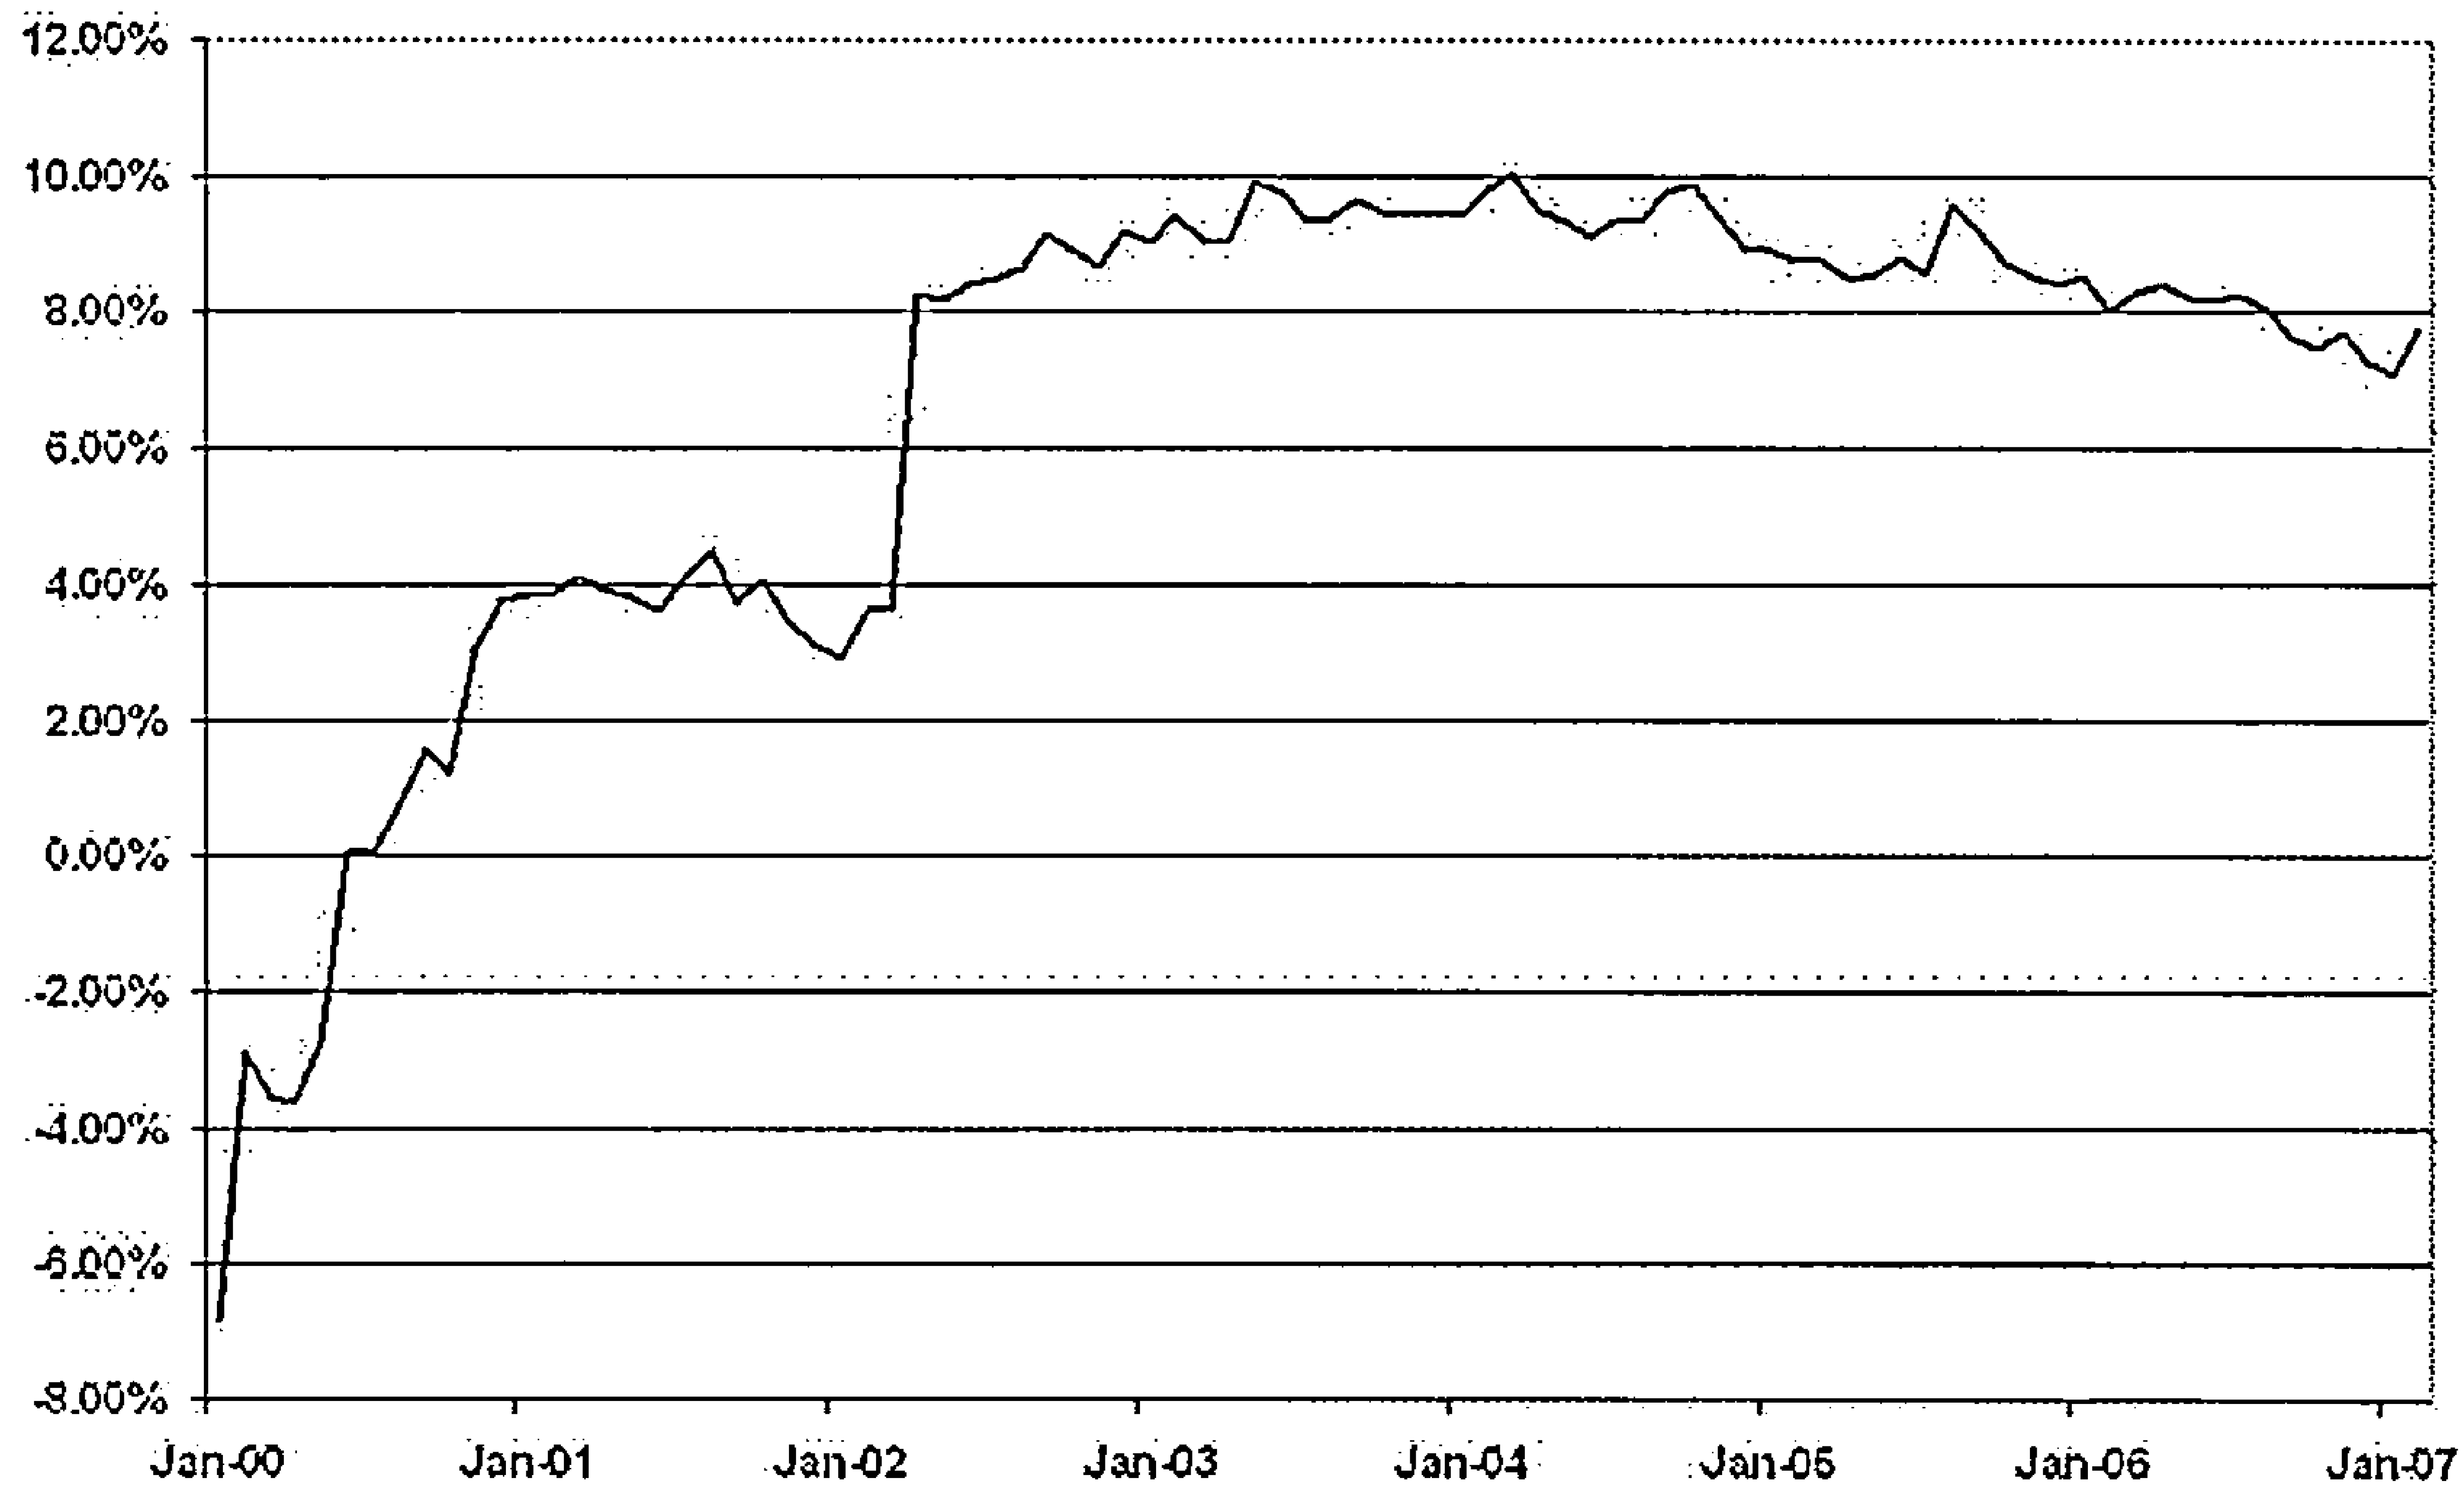

Figure 4.8. Average return synthetic fund 2, January 2000 – February 2007.

	50/50	Fund 2	10% Fund	20% Fund	30% Fund
Mean Return	7.07%	7.73%	7.14%	7.20%	7.27%
Volatility	7.65%	12.04%	7.01%	6.61%	6.51%
Skewness	-0.21	2.32	-0.11	0.14	0.60
Sharpe ratio	0.44	0.33	0.49	0.53	0.54

Table 4.4. Properties overall portfolio with varying allocations to synthetic fund 2, January 1998 – February 2007.

Table 4.4 shows how synthetic fund 2 performed in a portfolio context. We see that as a stand-alone investment the fund does not score very well. When mixed with the investor's original portfolio, however, it does much better. The overall portfolio's Sharpe ratio rises substantially and with a larger allocation it also eliminates the slight negative skewness found in the investor's original portfolio. The change in skewness is less than one might have expected, given the 2.23 skewness of the fund. For this to happen the allocation to the synthetic fund needs to be larger. A 50% allocation (not reported in Table 4.4) would produce a skewness of 1.70 for the overall portfolio.

4.1.3 Fund 3

Fund 3 is again similar to fund 1, except that this time we aim for seriously negative correlation. The payoff function for fund 3 as per January 1998 is shown in Figure 4.9. Comparing this graph with the payoff function for fund 1 as shown in Figure 4.1, we see that both are quite similar. This shows that it does not necessarily take a very large change to the payoff function to obtain significantly different results.

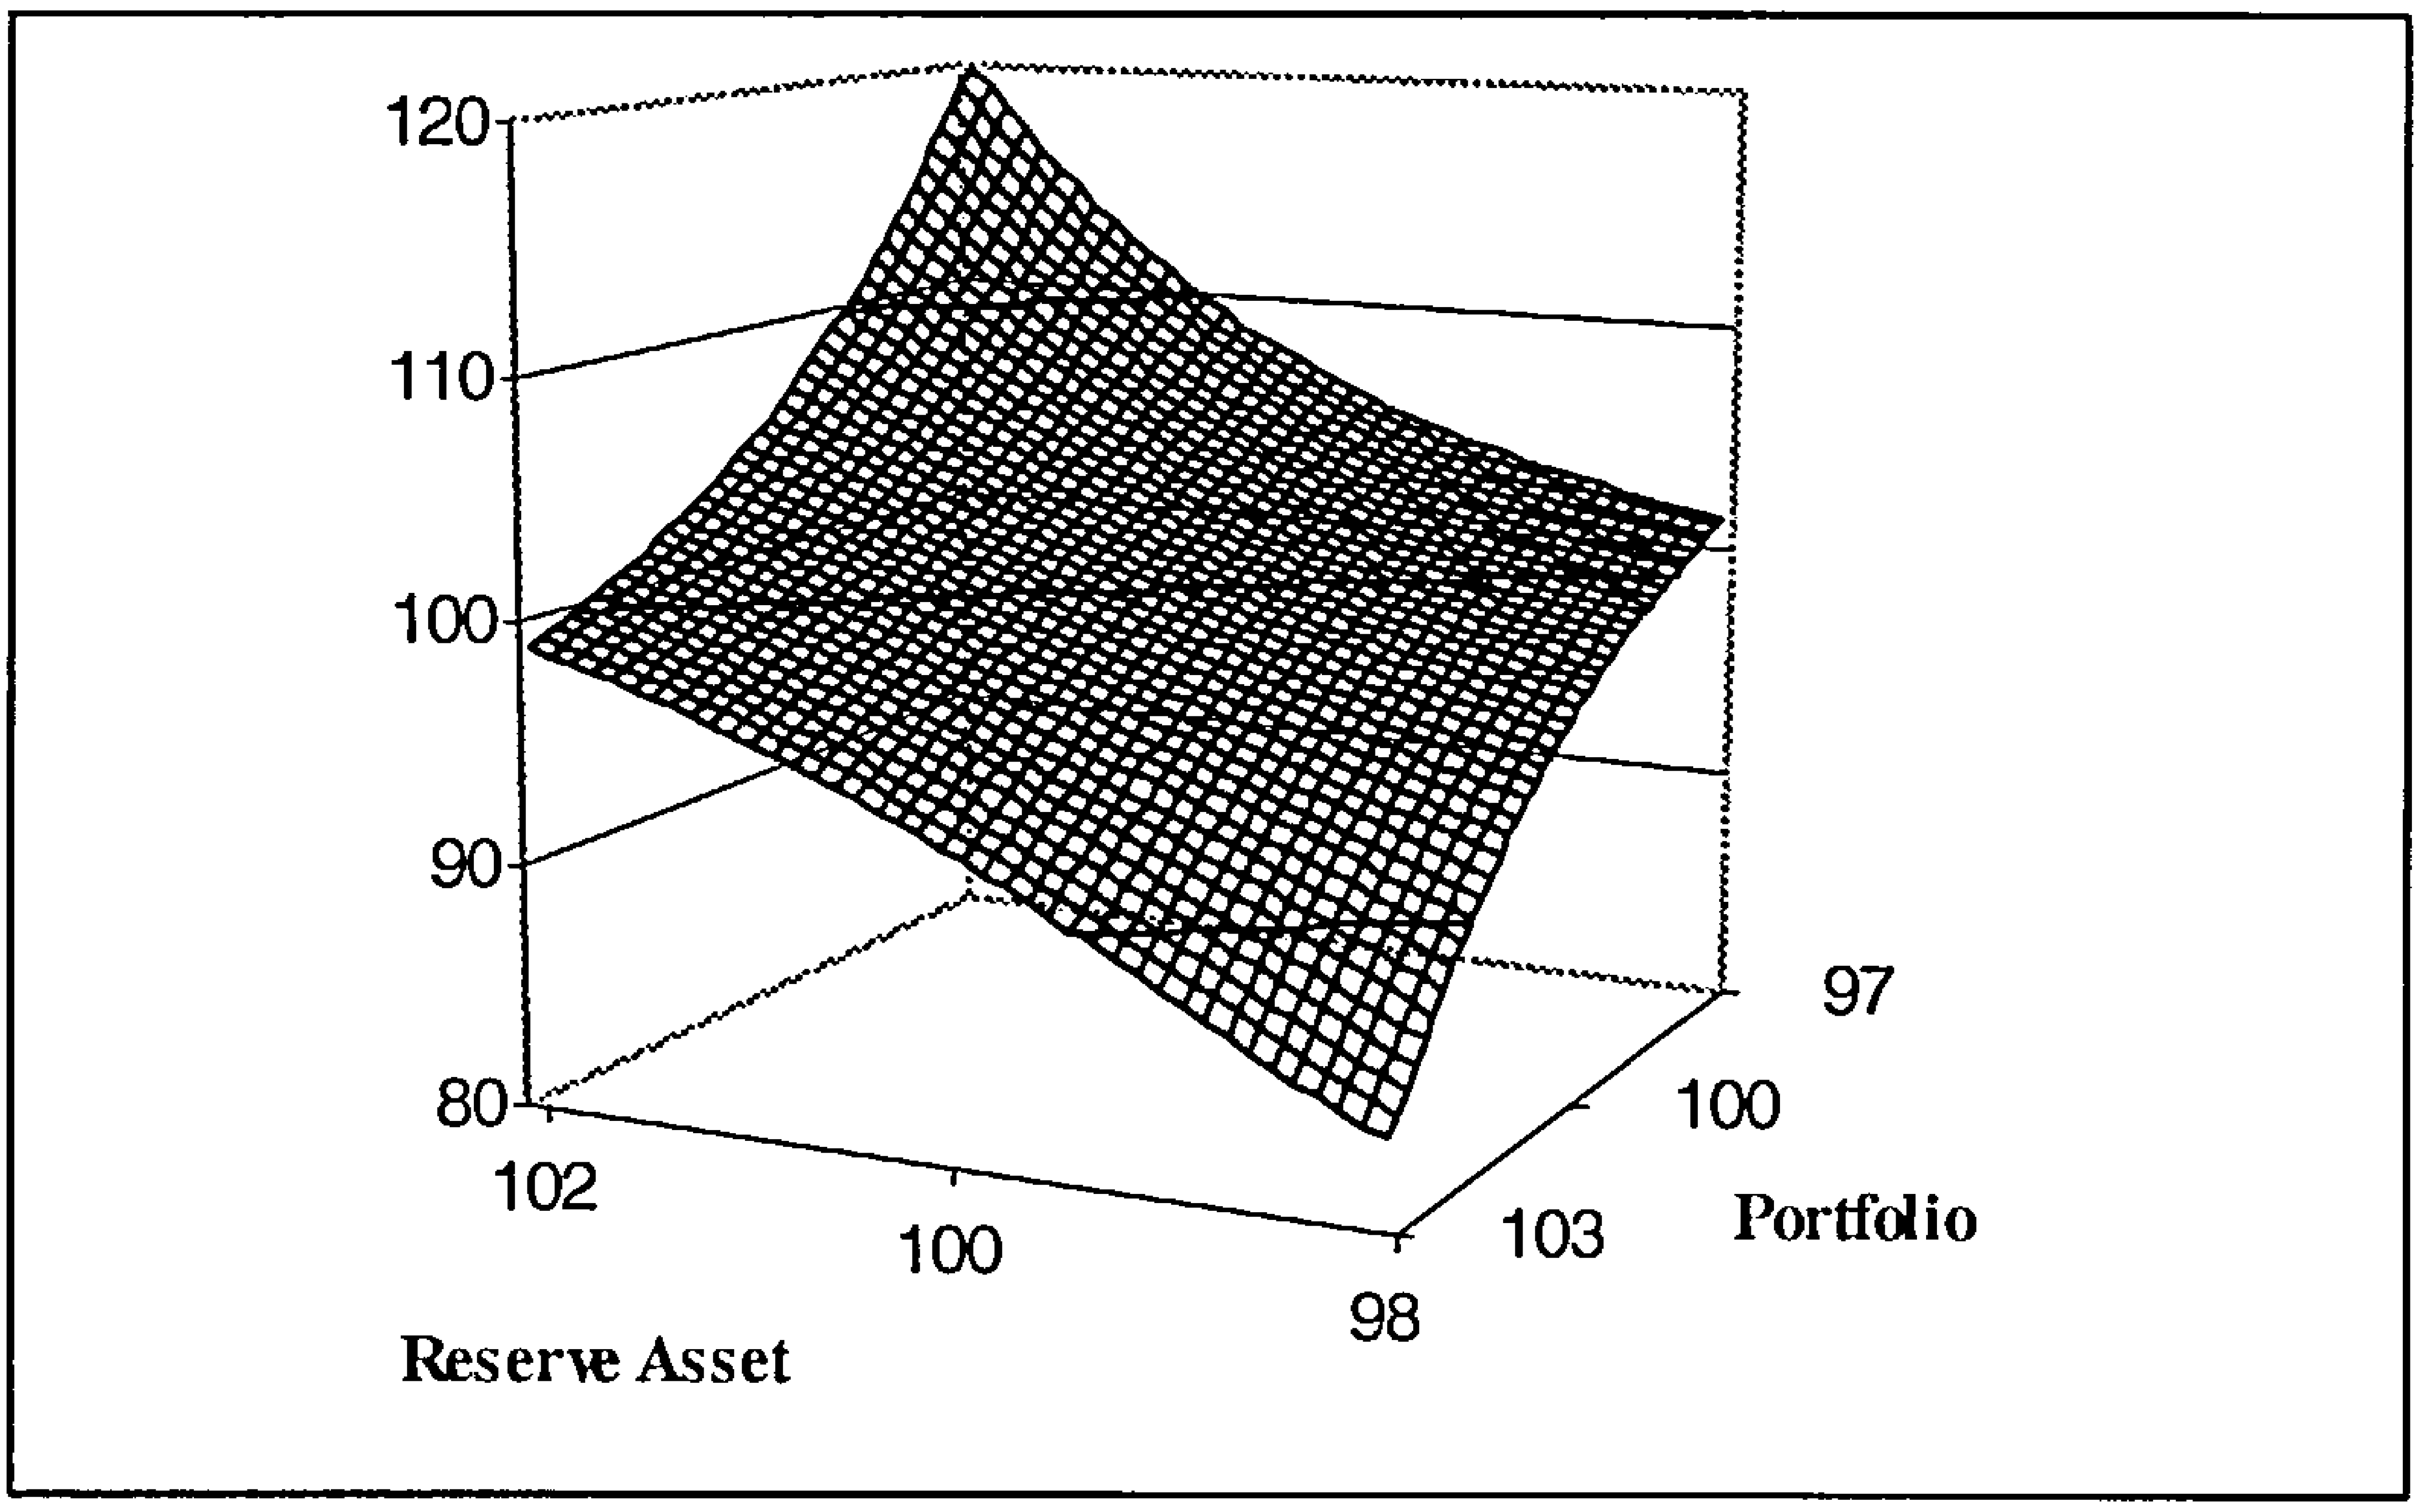

Figure 4.9. Target payoff synthetic fund 3, January 1998.

Figure 4.10 shows the evolution of the correlation between the synthetic fund and the investor's portfolio, while Figure 4.11 shows the average return. The graph in Figure 4.10 shows that the correlation of the synthetic fund return stayed close to its target value over the full 7-year period. Figure 4.11, however, shows that this does not come for free as the average return of the fund converges to no more than 6%. Intuitively, this is plausible. An asset, which has negative correlation with stocks and bonds, makes for a highly effective diversifier in a stock/bond portfolio. As a consequence, investor demand will be high, the asset's price will be high and its expected return correspondingly low. Of course, the expected return on our synthetic fund is not set directly by the market, but the expected return on the assets that are traded in the fund are, which is how the positive link between correlation and expected return filters in.

It is interesting to compare our synthetic fund with a direct investment in stock market volatility through the purchase of variance swaps. Carr and Wu (2006) show that over the period 1990-2005 such a strategy would have generated a highly negative mean excess return. A similar conclusion is found in Kat and Tassabehji (2006). Despite the fact that volatility returns tend to exhibit strong positive skewness, this makes our synthetic fund much more attractive than a long-only volatility investment strategy would have been.

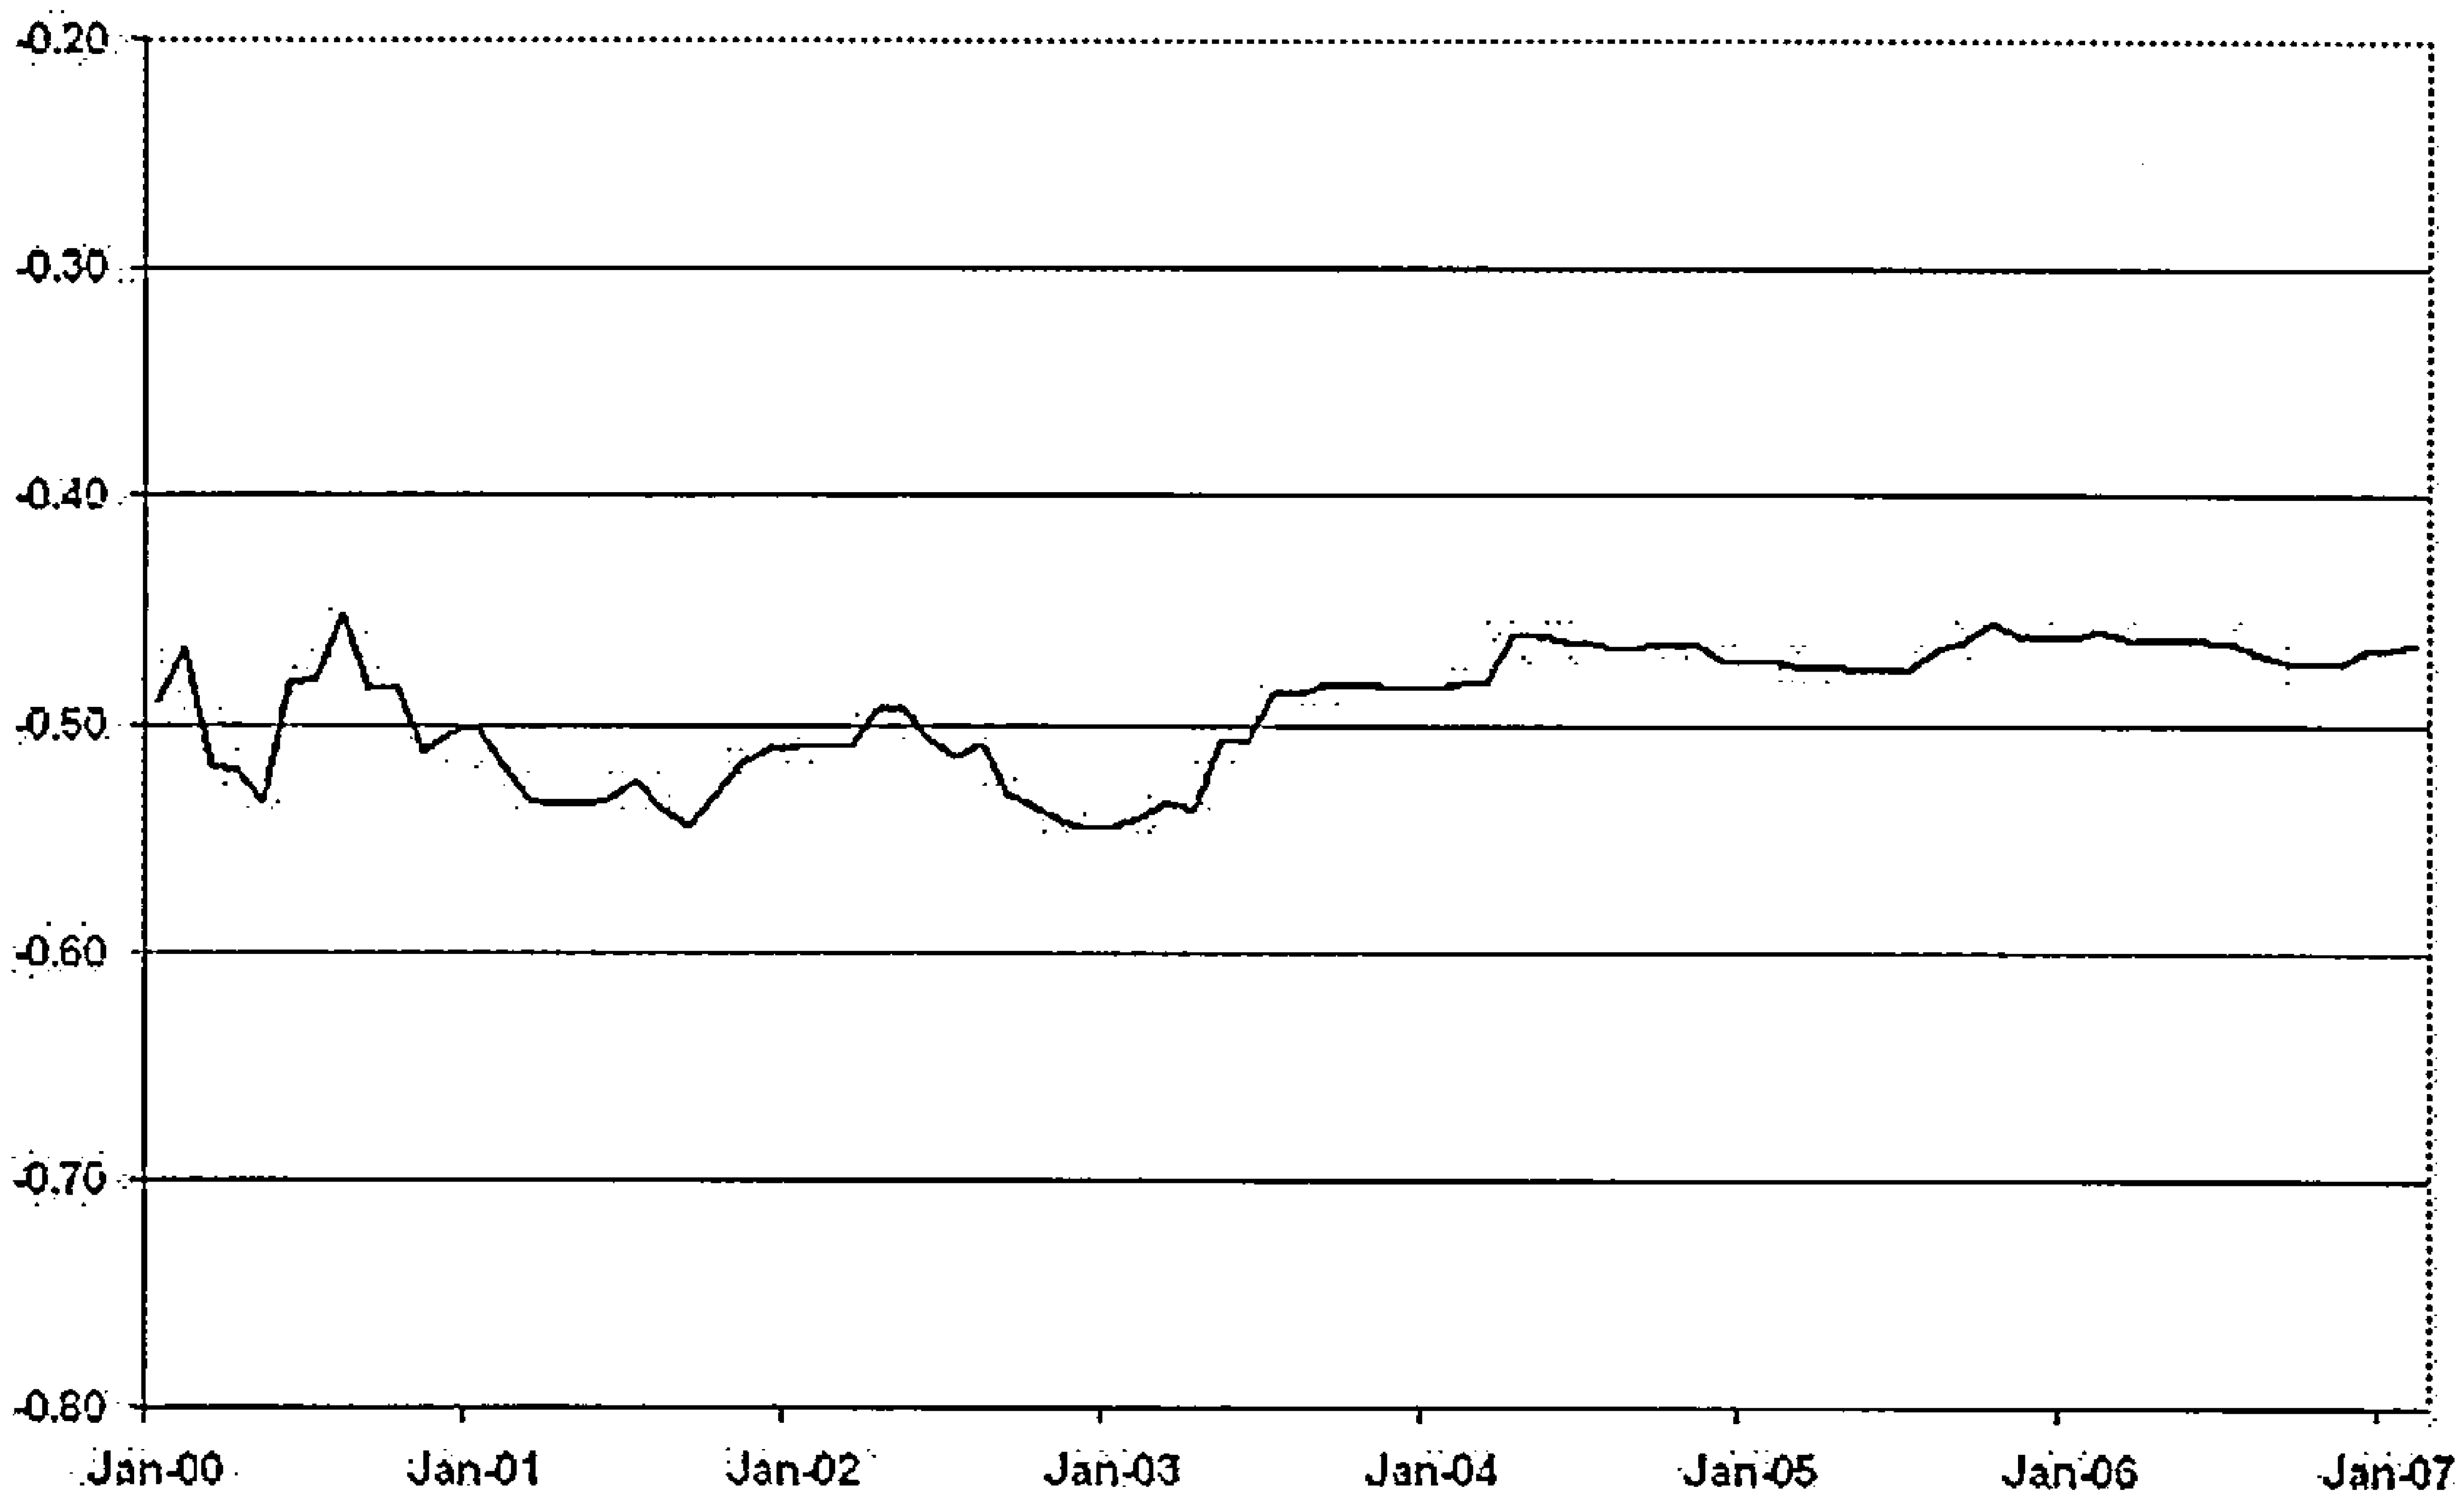

Figure 4.10. Correlation synthetic fund 3, January 2000 – February 2007.

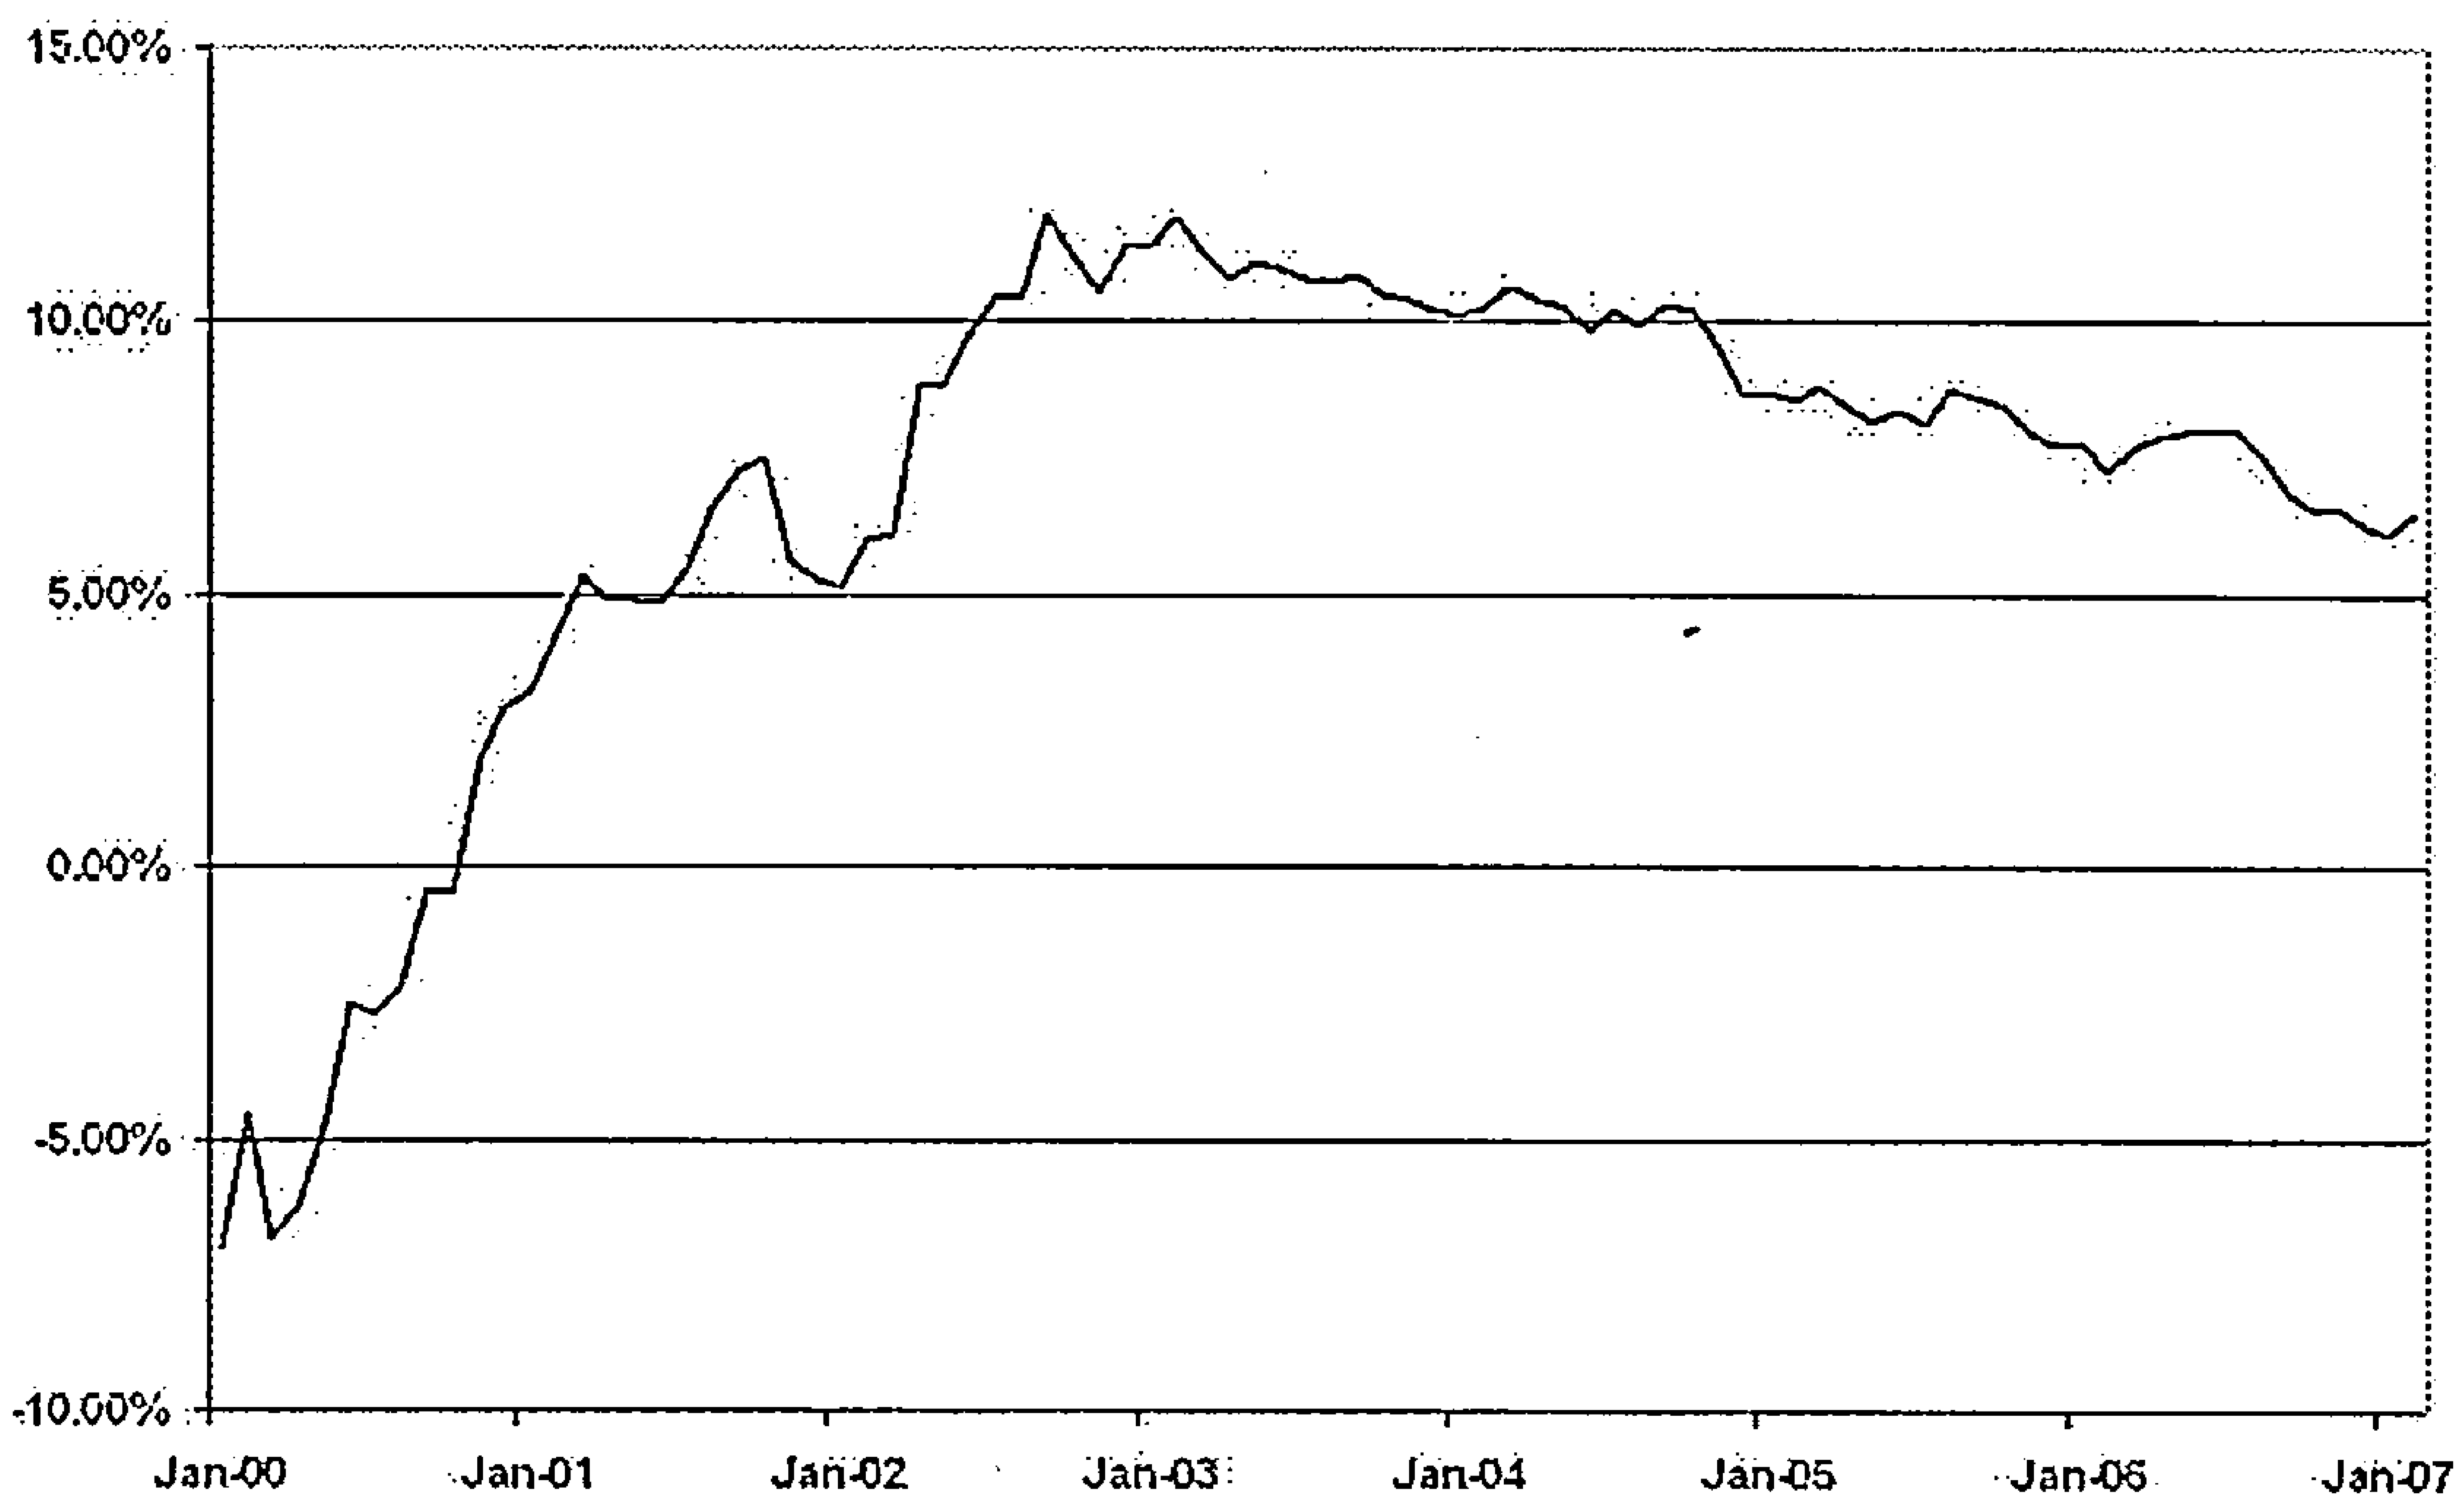

Figure 4.11. Average return synthetic fund 3, January 2000 – February 2007.

	50/50	Fund 3	10% Fund	20% Fund	30% Fund
Mean Return	7.07%	6.51%	7.01%	6.96%	6.90%
Volatility	7.65%	11.27%	6.44%	5.45%	4.82%
Skewness	-0.21	0.42	-0.16	-0.11	-0.05
Sharpe ratio	0.44	0.25	0.51	0.59	0.66

Table 4.5. Properties overall portfolio with varying allocations to synthetic fund 3, January 1998 – February 2007.

Due to the high price of negative correlation, the fund's mean return is low relative to its volatility, resulting in a Sharpe ratio of no more than 0.22. This makes fund 3 quite an unattractive investment on a stand-alone basis. Mixing the fund with the investor's original portfolio, as reported in Table 4.5, we see a familiar picture, however. Adding the synthetic fund to the investor's original portfolio, the overall portfolio's volatility drops sharply, but without a corresponding loss in mean return. As a result, the portfolio's Sharpe ratio rises very substantially. It is interesting to note that, as to judge from the resulting Sharpe ratios, fund 1 and 3 are equally effective in diversifying the investor's original portfolio. This confirms that the drop in mean excess return from lowering the synthetic fund's correlation with the investor's original portfolio is market-conform.

4.1.4 Fund 4

The last fund we study is again similar to fund 1, but this time we put a –5% floor under the monthly fund return. This is similar to buying an out-of-the-money put option. There are a number of important differences between buying real puts, and synthesizing puts through dynamic trading, however. An option is a legally binding contract between two counterparties that entitles the holder of the option to a specific payoff. As a result, apart from credit risk, buying puts provides a 'hard' floor, i.e. it fully protects against returns falling below the chosen floor level. Since we do not really buy puts, but simply integrate the hedging strategy for a put into the fund strategy instead, our floor is 'soft' in the sense that it could be breached if the market came down substantially over a short period of time. This may not sound good, but having a soft floor comes with a number of important benefits, which for a long-term investor will typically outweigh the downside of a soft floor. First, partly because of their 'hard' nature, options are expensive. The buyer of an option pays implied volatility, while when executing the accompanying hedging strategy, one pays spot volatility. The latter is typically a few percent lower than implied volatility. Synthesizing a put ourselves instead of buying one outright therefore helps to keep the fund's risk premium at an acceptable level. Second, since we work with bivariate payoffs, we will need a bivariate put as well, i.e. a put with a payoff depending on the investor's portfolio as well as the reserve asset. To buy such an option we will have to turn to the over-the-counter (OTC) options market, which implies paying additional margin to the investment bank that takes the other side, a high degree of illiquidity, and additional operational hassle.

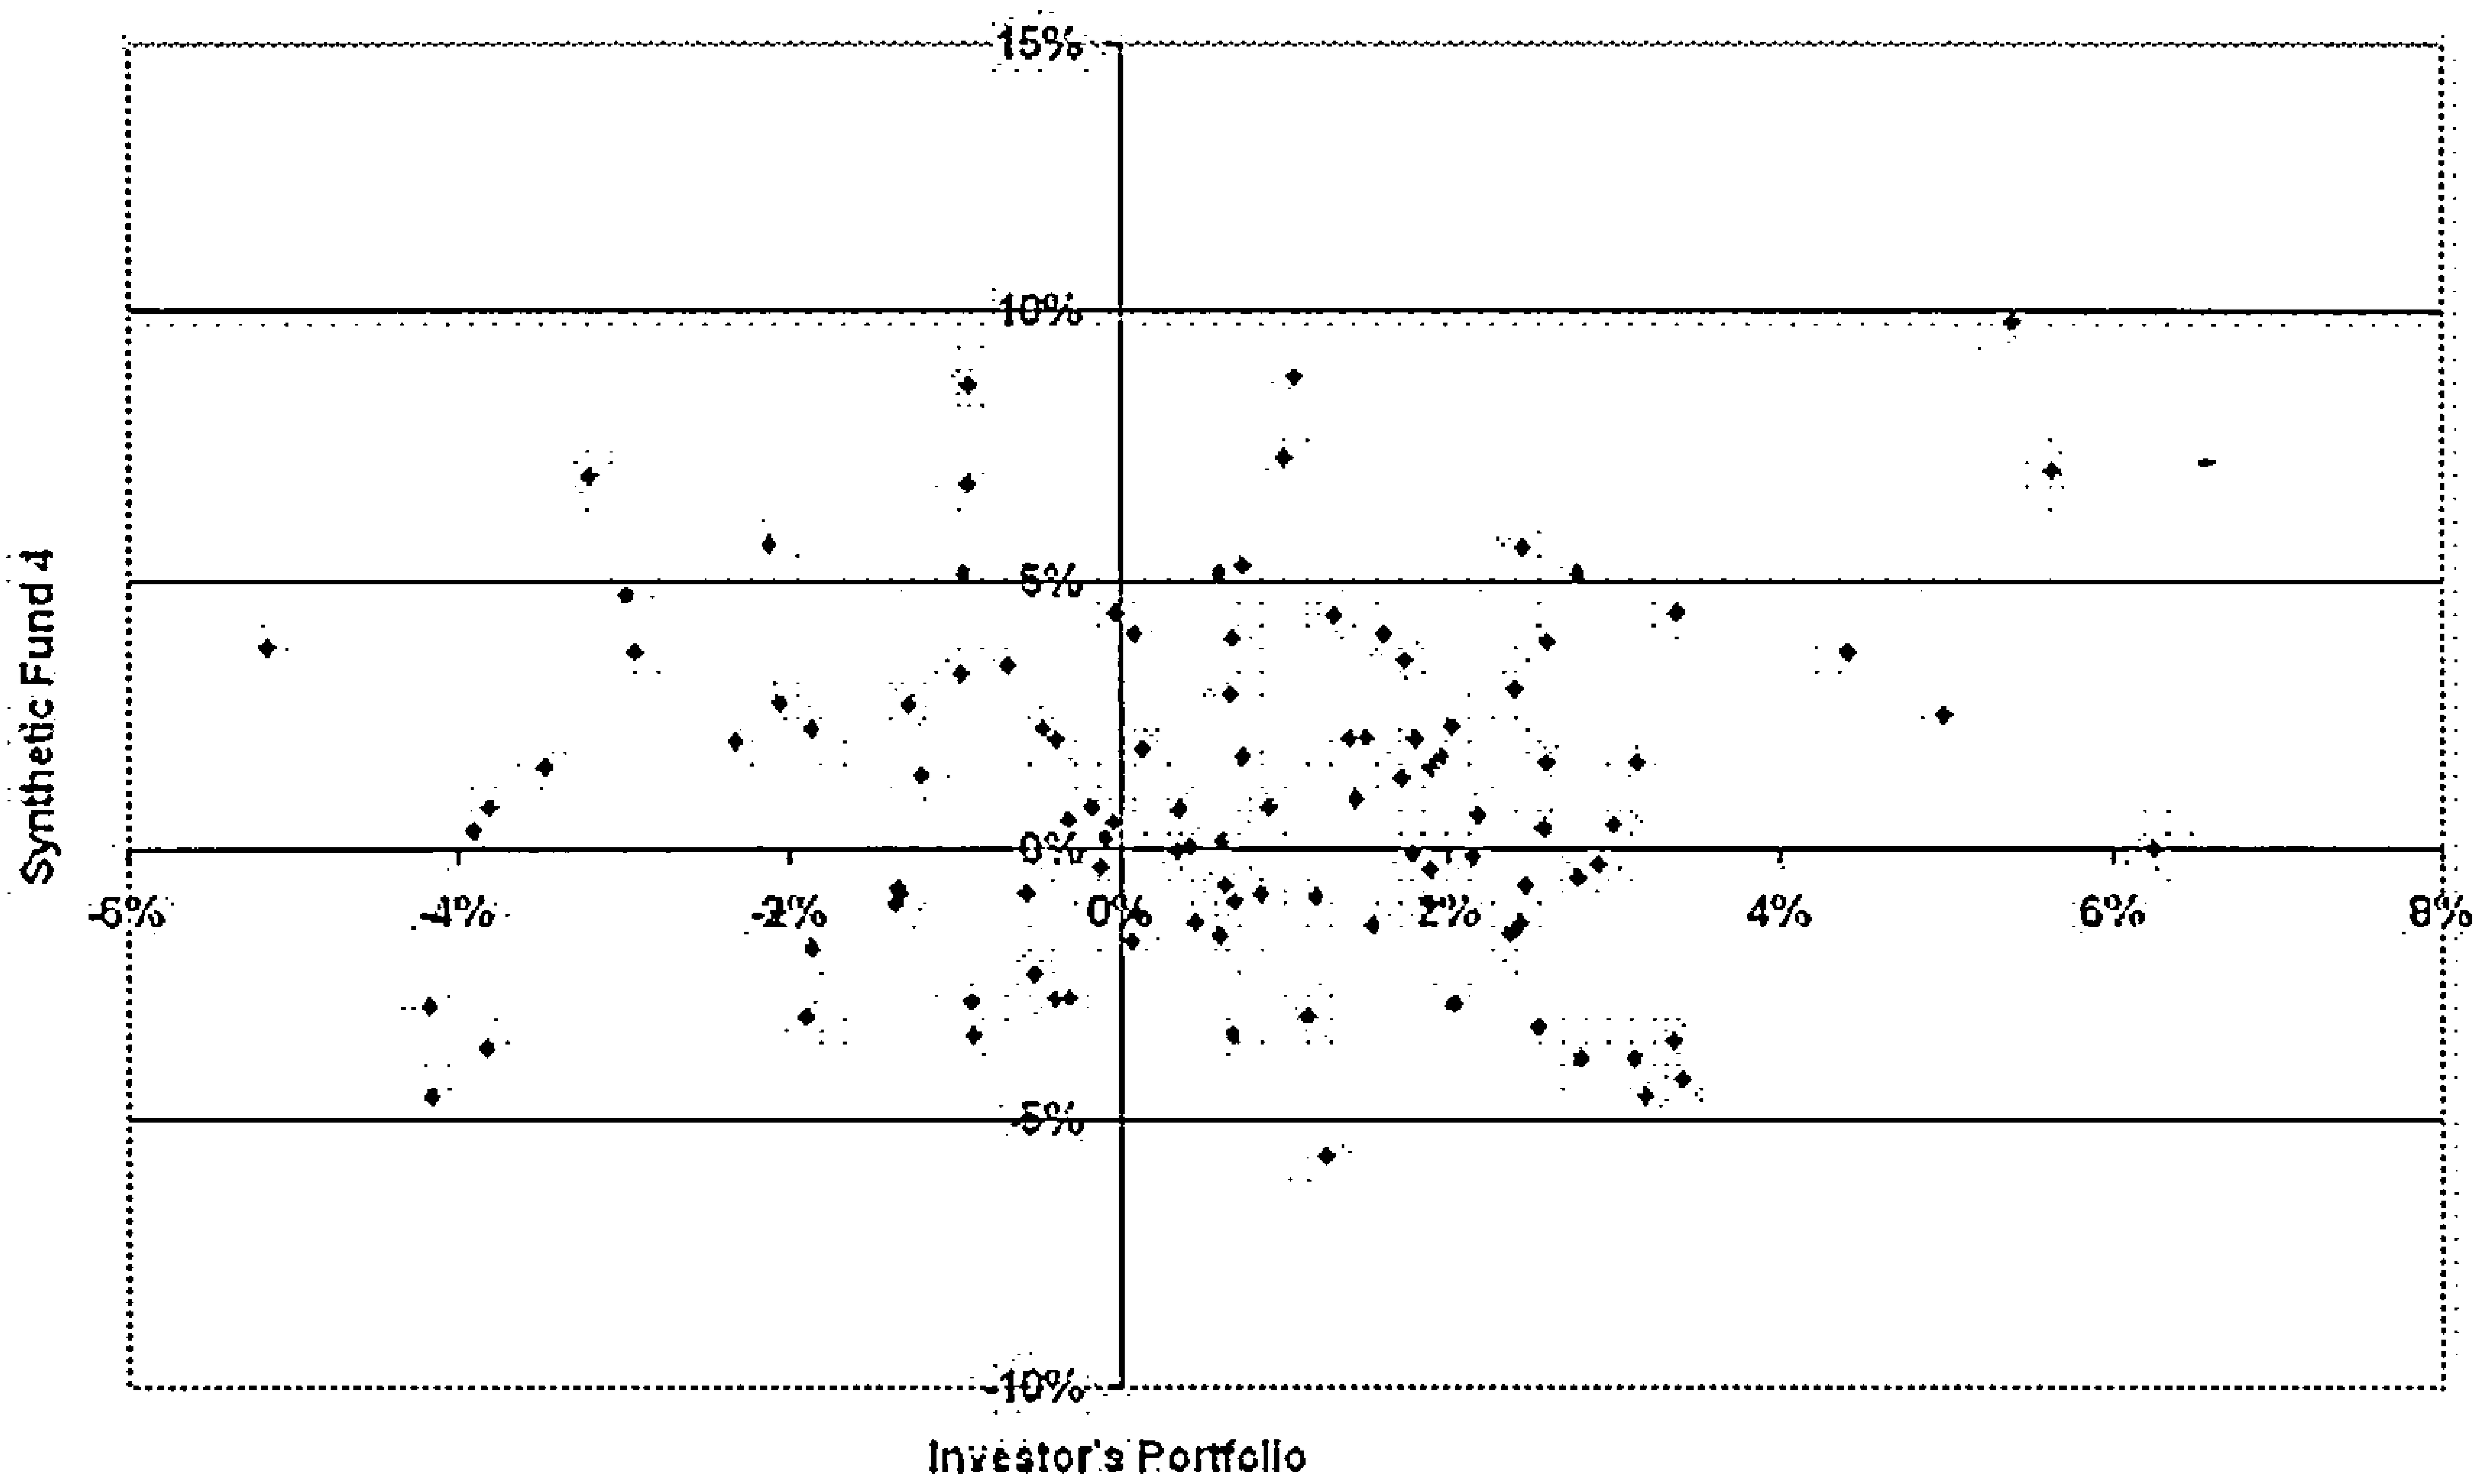

Figure 4.12. Scatterplot synthetic fund 4 return versus investor's portfolio return, January 2000 – February 2007.

Figure 4.12 shows a plot of the fund average return versus the average return on the investor's portfolio over the period January 1998 – February 2007. Apart from the random scatter that comes with the targeted zero correlation, the graph clearly shows the impact of the floor. It also shows that, despite the fact that the protection provided is 'soft', it is highly effective.

The few returns that do end up below –5% only do so to a limited extent. Another way to evaluate the workings of the floor is to compare the excess return on fund 4 with that on fund 1. This is done in Figure 4.13. The graph in Figure 4.13 confirms that without actually buying put options we have created a payoff profile, which closely resembles that of a portfolio protected with ordinary puts. On the upside the returns of fund 1 and 4 are very similar, but on the downside fund 4's losses are stopped out around the floor level.

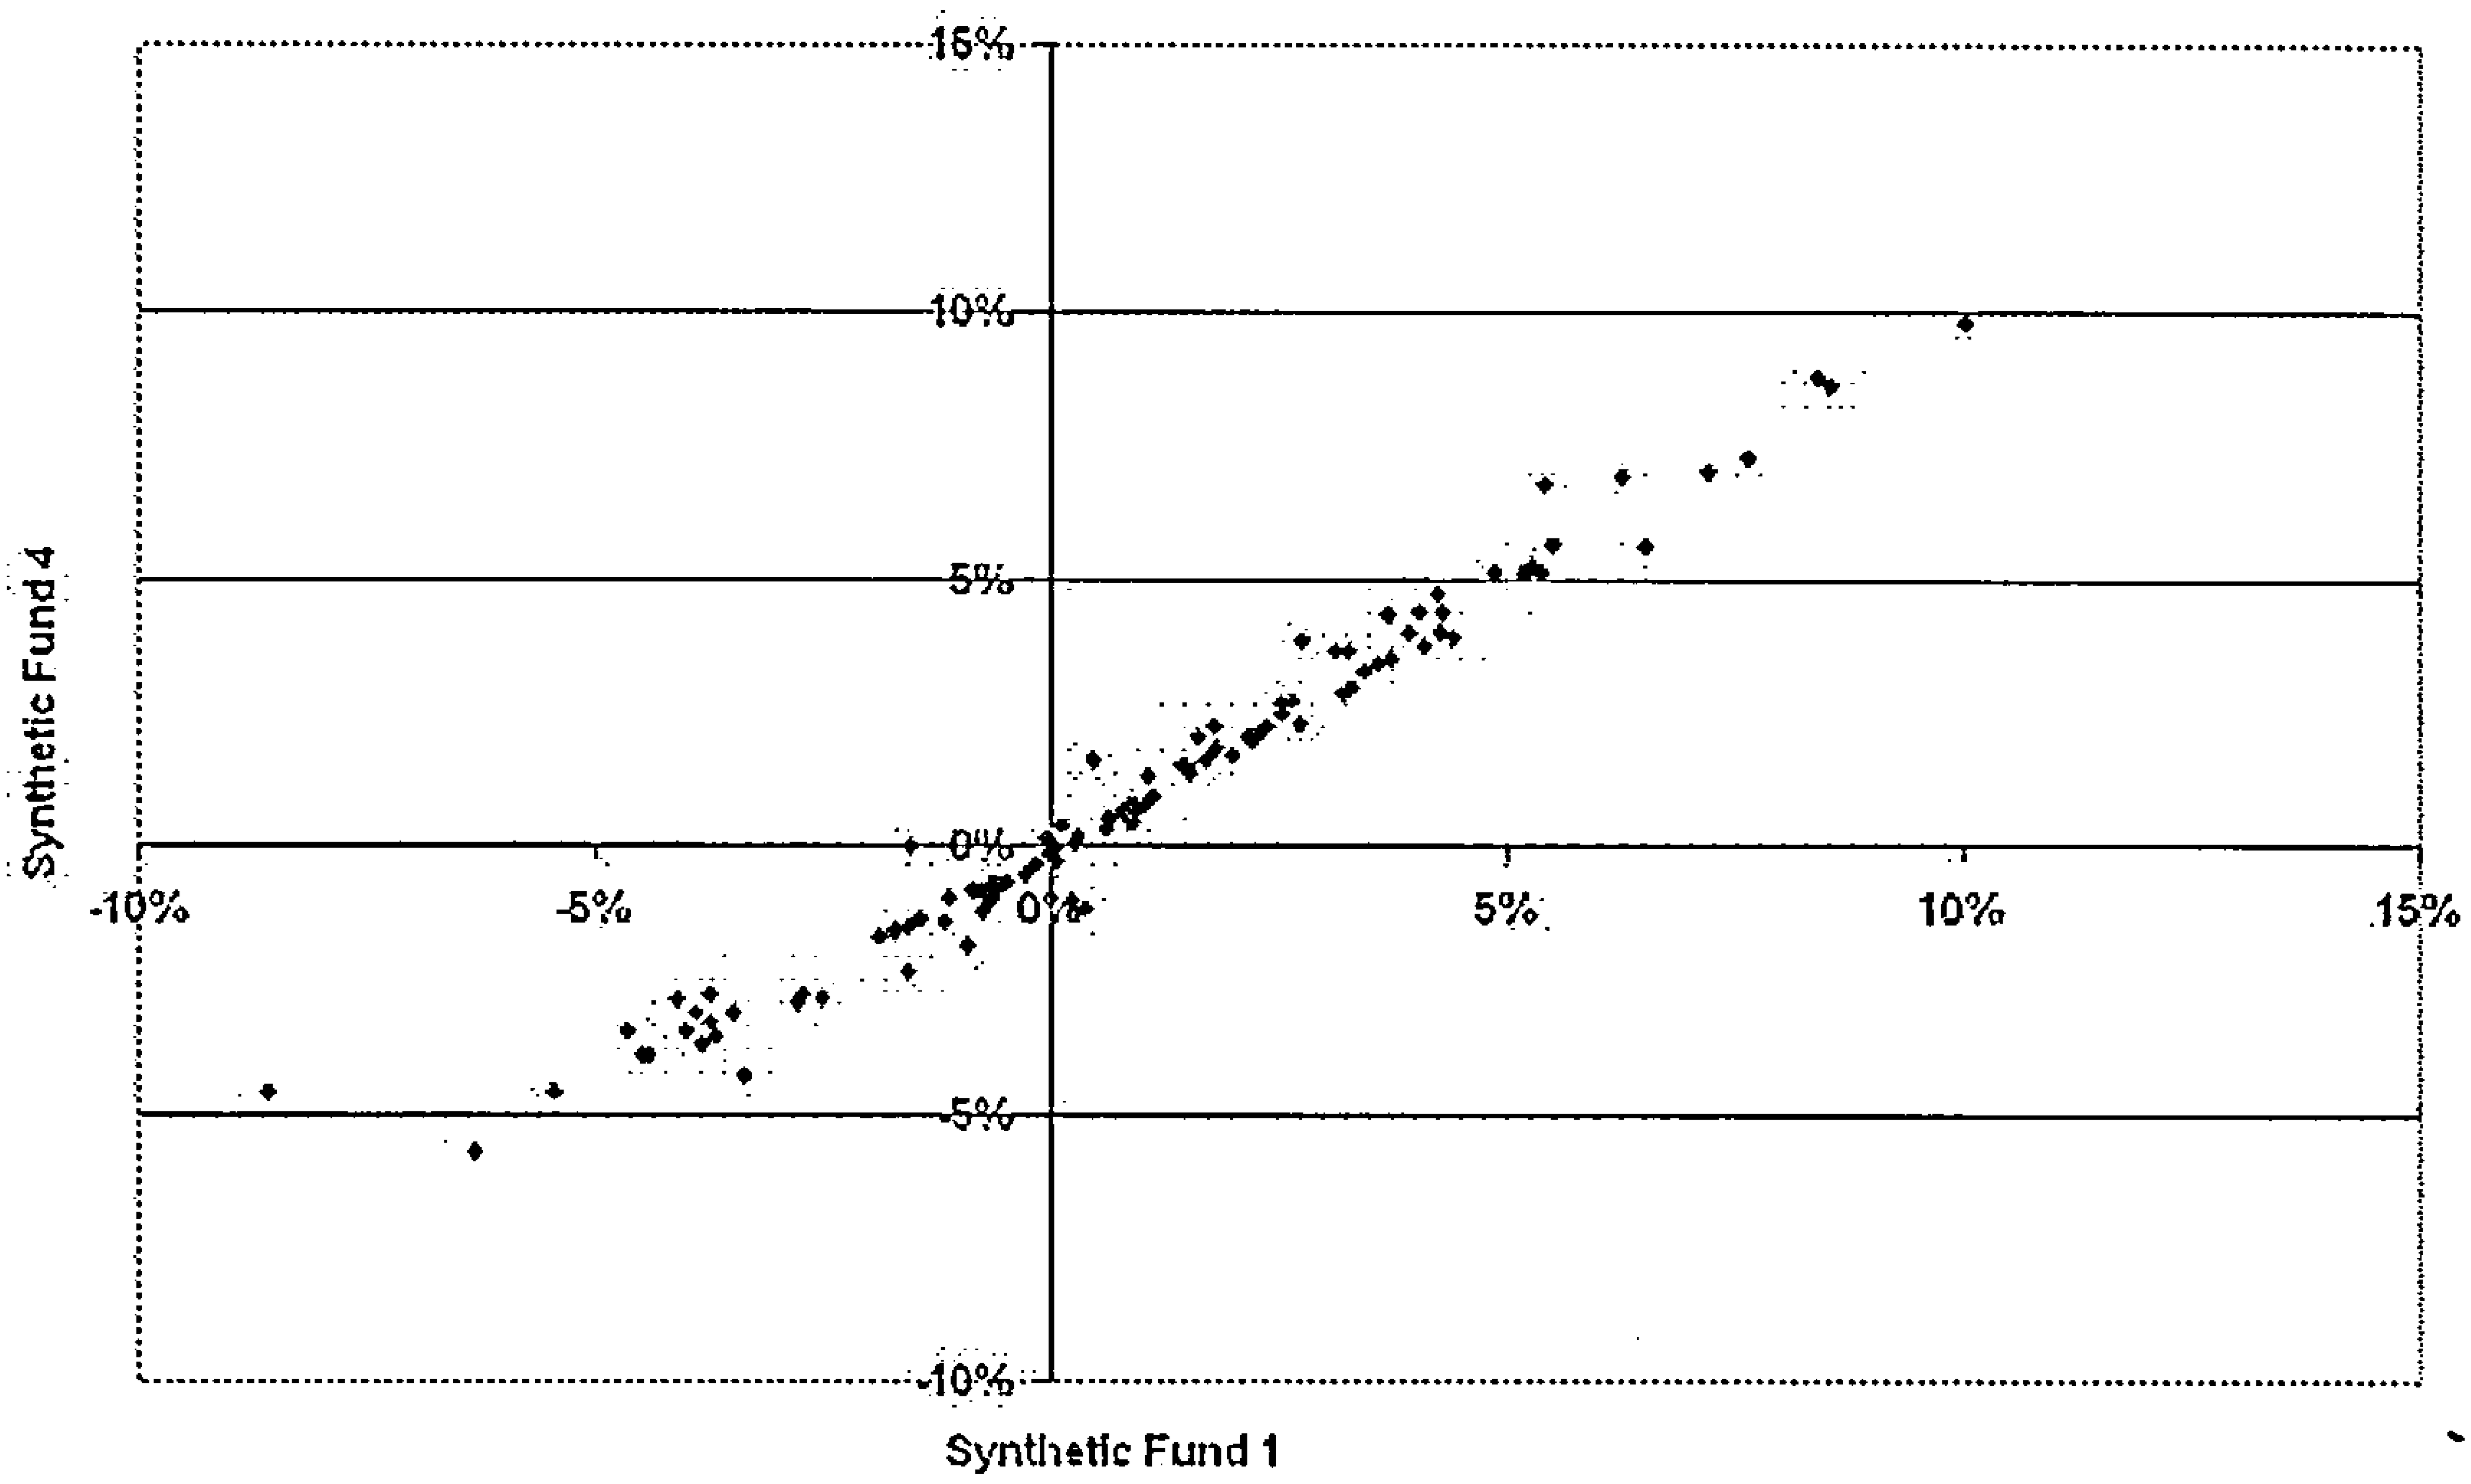

Figure 4.13. Scatterplot synthetic fund 4 return versus synthetic fund 1 return, January 2000 – February 2007.

Comparing Figure 4.14, which shows the evolution of the average return of fund 4, with Figure 4.5, we can see that the long-term average return for fund 4 lies around 1% higher than that for fund 1. Specially the period 1998-2000, in the beginning of the graph, was better for fund 4, with the protection of the floor.

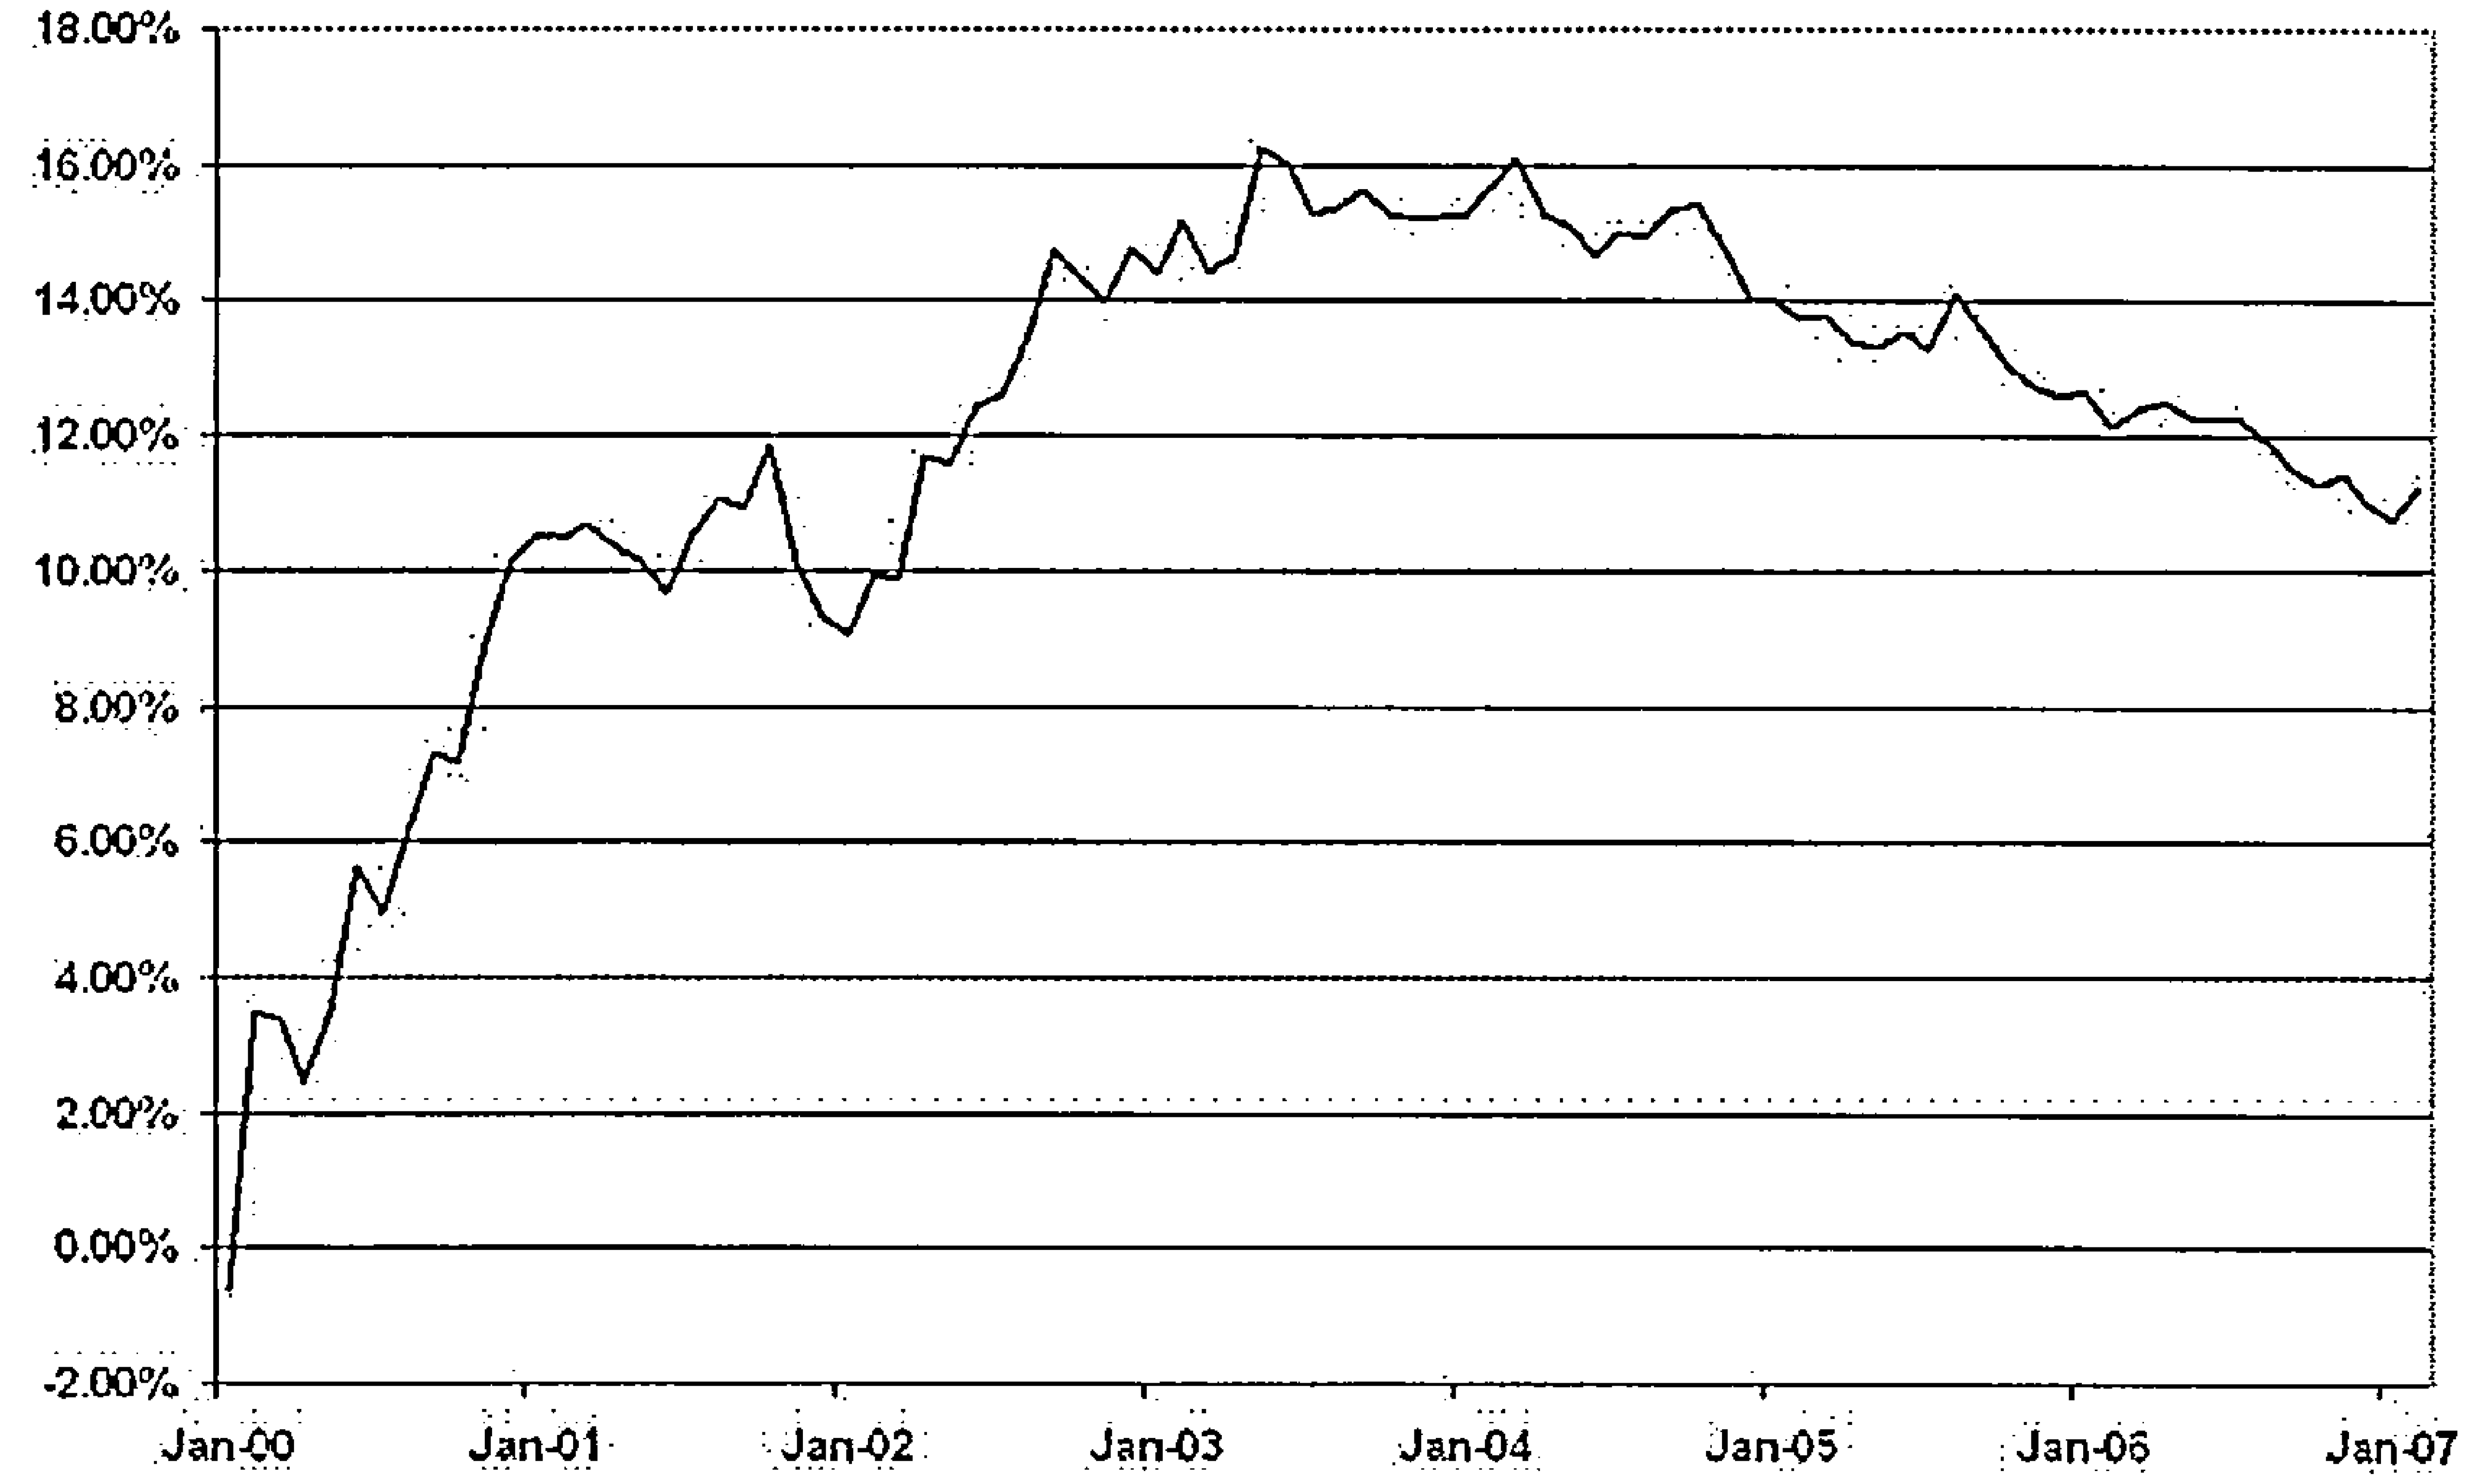

Figure 4.14. Average return synthetic fund 4, January 2000 – February 2007.

	50/50	Fund 4	10% Fund	20% Fund	30% Fund
Mean Return	7.07%	11.22%	7.49%	7.90%	8.32%
Volatility	7.65%	11.12%	7.01%	6.58%	6.40%
Skewness	-0.21	0.40	-0.12	0.00	0.16
Sharpe ratio	0.44	0.67	0.54	0.63	0.72

Table 4.6. Properties overall portfolio with varying allocations to synthetic fund 4, January 1998 – February 2007.

Table 4.6 places fund 4 in a portfolio context. Apart from the usual diversification benefits, it shows that, in terms of the parameters shown, the diversification properties of fund 4 are slight superior to those of fund 1 or fund 3. This has happened because of the protection which the floor provided in the troublesome period of 1998.

4.2 Sensitivity analyses

In this section is examined the influence of the transaction costs, the reference portfolio and the reserve asset on the replication results. Fund 1 is used to check the influence of these factors.

4.2.1 Transaction Costs

Depending on the targeted properties and the way markets behave, the replication technique introduced in Chapter 2 tend to be quite dynamic, so an obvious question to ask is whether the high turnover in the trading strategies is not excessively costly.

	Mean	StDev	Skew	Corr
Synthetic 0bps, daily	10.24%	11.57%	0.10	0.03
Synthetic 4bps, daily	10.08%	11.56%	0.10	0.03
Synthetic 14bps, daily	9.67%	11.56%	0.10	0.03
Synthetic 50bps, daily	8.12%	11.54%	0.10	0.03
Synthetic 0bps, 2 daily	10.01%	11.54%	0.08	0.03
Synthetic 4bps, 2 daily	9.93%	11.54%	0.08	0.03
Synthetic 14bps, 2 daily	9.74%	11.54%	0.08	0.03
Synthetic 50bps, 2 daily	8.96%	11.53%	0.08	0.03
Synthetic 0bps, 3 daily	9.92%	11.53%	0.07	0.03
Synthetic 4bps, 3 daily	9.85%	11.53%	0.07	0.03
Synthetic 14bps, 3 daily	9.70%	11.53%	0.07	0.03
Synthetic 50bps, 3 daily	9.20%	11.52%	0.07	0.03

Table 4.7. Sample properties zero correlation fund returns over the period January 1998 - February 2007.

Table 4.7 shows the sample properties of the returns on the Zero Correlation Fund for various bid-offer spreads (0bps, 4bps, etc.) and rebalancing frequencies (daily, once every 2 days, and once every 3 days) over the period January 1998 – February 2007. Note that the case of a 50bps spread is included for illustrative purposes only. In practice one would not expect to trade on such a wide spread very often.

There is no significant impact of transaction costs and/or the rebalancing frequency on the risk profile. For realistic cost levels, the results are still quite satisfactory, especially for lower rebalancing frequencies.

4.2.2 Reference Portfolio

The replication technique designs futures trading strategies that generate returns with predefined statistical properties, including the dependence with the so-called reference portfolio, i.e. the portfolio with respect to which dependence is measured. In most applications the reference portfolio will be the investor's existing portfolio. This brings up the question how much the bottom line result depends on the reference portfolio. In other words, to what extent do investors who start off with different reference portfolios end up with significantly different results?

Since the replication strategies will typically be used to diversify a larger, more traditional portfolio, the reference portfolio is taken to consist of three different mixes of S&P 500 and T-bond futures. More specifically, the portfolio weights 20/80, 50/50 and 80/20 are considered. As before, it is assumed that throughout the reserve asset consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures.[48]

	50/50 Reference Portfolio				**80B/20E Reference Portfolio**			
	Mean	**StDev**	**Skew**	**Corr**	**Mean**	**StDev**	**Skew**	**Corr**
Synthetic 0bps, daily	10.24%	11.57%	0.10	0.03	14.03%	11.83%	-0.05	0.09
Synthetic 4bps, daily	10.08%	11.56%	0.10	0.03	13.84%	11.83%	-0.05	0.09
Synthetic 14bps, daily	9.67%	11.56%	0.10	0.03	13.28%	11.82%	-0.04	0.09
Synthetic 0bps, 2 daily	10.01%	11.54%	0.08	0.03	14.07%	11.81%	-0.07	0.09
Synthetic 4bps, 2 daily	9.93%	11.54%	0.08	0.03	13.96%	11.80%	-0.07	0.09
Synthetic 14bps, 2 daily	9.74%	11.54%	0.08	0.03	13.68%	11.80%	-0.07	0.09
Synthetic 0bps, 3 daily	9.92%	11.53%	0.07	0.03	14.37%	11.81%	-0.07	0.09
Synthetic 4bps, 3 daily	9.85%	11.53%	0.07	0.03	14.30%	11.81%	-0.07	0.09
Synthetic 14bps, 3 daily	9.70%	11.53%	0.07	0.03	14.09%	11.81%	-0.07	0.09

Table 4.8: Sample properties Zero Correlation Fund returns over the period January 1998 - Febuary 2007.

[48] Since the volatility of the various asset classes is quite different, before forming the portfolio, we leveraged the Eurodollar and 5Y note by a factor 5, and the 10Y note by a factor 4 to give these components a level of volatility more in line with stocks and commodities. Note that this will increase the trading volume in these contracts.

Table 4.8 shows the results for the Zero Correlation Fund. The target in this case is to generate normally distributed returns with 12% volatility and zero correlation with the reference portfolio. Although the means for the two reference portfolios are somewhat different, the standard deviation, skewness and correlation results match the target very well in both cases. The 80/20 reference portfolio seems to require more trading than in the 50/50 case, as the mean drops slightly faster when the bid-ask spread increases. The results for a 20/80 mix are very similar and therefore not reported.

4.2.3 Reserve Asset

The reserve asset is the core portfolio of the trading strategy and therefore the main source of uncertainty. Although over time the strategy will move in and out of the reserve asset, it will never short it. Given its important role, the natural question to ask is how sensitive the replication results are to the choice of reserve asset.

Since hedge funds are typically used to diversify larger, more traditional portfolios, we will take the reference portfolio to consist of 50% S&P 500 and 50% T-bond futures. Although investors are completely flexible in their choice of reserve asset, it is important to keep in mind that the reserve asset is the core portfolio of the strategy. It should therefore offer a good ratio between expected return and risk. Unless the investor has strong views, it should be well diversified, i.e. not contain too much uncompensated risk, and offer a satisfactory risk premium. Put simply, a reserve asset with a good Sharpe ratio will also make for a good expected return on the resulting trading strategy.

In this subsection two different reserve assets are considered. Reserve Asset 1 consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures, where, to compensate substantial differences in volatility, we leveraged the Eurodollar and 5-year note by a factor 5, and the 10-year note by a factor 4. Reserve Asset 2 is much simpler and consists of an equally weighted portfolio of 1-month Libor, Russell 2000 and crude oil futures.

	Reserve Asset 1				Reserve Asset 2			
	Mean	**StDev**	**Skew**	**Corr**	**Mean**	**StDev**	**Skew**	**Corr**
Synthetic 0bps, daily	10.24%	11.57%	0.10	0.03	10.86%	11.45%	-0.01	-0.03
Synthetic 4bps, daily	10.08%	11.56%	0.10	0.03	10.76%	11.45%	-0.01	-0.03
Synthetic 14bps, daily	9.67%	11.56%	0.10	0.03	10.51%	11.45%	-0.01	-0.03
Synthetic 0bps, 2 daily	10.01%	11.54%	0.08	0.03	10.92%	11.45%	-0.01	-0.03
Synthetic 4bps, 2 daily	9.93%	11.54%	0.08	0.03	10.88%	11.45%	-0.01	-0.03
Synthetic 14bps, 2 daily	9.74%	11.54%	0.08	0.03	10.76%	11.44%	-0.01	-0.03
Synthetic 0bps, 3 daily	9.92%	11.53%	0.07	0.03	10.80%	11.44%	-0.02	-0.03
Synthetic 4bps, 3 daily	9.85%	11.53%	0.07	0.03	10.77%	11.44%	-0.02	-0.03
Synthetic 14bps, 3 daily	9.70%	11.53%	0.07	0.03	10.67%	11.44%	-0.02	-0.03

Table 4.9. Sample properties Zero Correlation Fund returns over the period January 1998 - February 2007.

Using Reserve Asset 1 and 2, Table 4.9 shows the sample properties of the replicated returns on the Zero Correlation fund for various bid-offer spreads (0bps, 4bps, etc.) and rebalancing frequencies (daily, once every 2 days, and once every 3 days) over the period January 1998 – February 2007. The risk parameters for both reserve assets are very close, with small differences in the mean return. Reserve Asset 2 does slightly better than Reserve Asset 1.

4.3 Synthetic Fund Alpha

Since the trading strategies are purely mechanical and do not involve any proprietary trading secrets, synthetic funds are not set up to generate alpha in the traditional sense, i.e. beat the market. Because of synthetic funds' mechanical nature, however, investors can do without expensive managers. Given the typical level of fees in alternative investments and the improbability of most managers being sufficiently skilled to make up for them, this means that although our synthetic funds' pre-fee returns may not be superior, their after-fee returns could very well be. *In the end, efficient risk management and cost control are much more certain routes to superior performance than trying to beat the market while paying excessive management and incentive fees.*

Although not explicitly designed to beat the market, synthetic funds do allow for tactical input through the choice of the reserve asset, which could therefore form a second source of alpha. Strictly speaking, the latter cannot be attributed to the fund, however, as it derives from inputs that are completely exogenous to the fund itself.

5. Hedge Fund Indexation[49]

In this chapter the replication technique introduce in Chapter 2 is used to test whether is possible the replication of hedge-fund indices of three major indices providers.

5.1 Synthetic Hedge Fund Index Returns

Having developed a potentially workable approach to hedge fund indexation, the big question of course is how it performs in practice. In this section we therefore take a detailed look at the returns from replicating a number of well-known hedge fund indices up to October 2006.

Having cleaned up the autocorrelation, we have to decide what futures contracts to trade to produce our synthetic hedge fund index returns with. In the replication terminology, this means we need to select our 'reference portfolio' and our 'reserve asset'. The 'reference portfolio' is the portfolio with respect to which we will measure correlation. Since hedge funds are typically used to diversify larger, more traditional portfolios, we will take the reference portfolio to consist of 50% S&P 500 and 50% T-bond futures. In the remainder of this chapter, when we talk about correlation, we mean the correlation between the index (or its replica) and this particular portfolio. The 'reserve asset' is the core portfolio of the replication strategy and therefore the main source of uncertainty. Since the outlook for the various asset classes will change over time, in practice the choice of the reserve asset is a dynamic process, producing time-varying allocations. Unfortunately, the latter process is very difficult to simulate in a backtest without the suggestion of data mining. In what follows we therefore assume that the composition of the reserve asset is fixed though time. More specifically, we assume the reserve asset consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures.[50]

[49] The material in this chapter appeared as Alternative Investment Research Center WP 38, Cass Business School.

[50] Since the volatility of the various asset classes is quite different, before forming the portfolio, we leveraged the Eurodollar, 5Y note by a factor 5, and the 10Y note by a factor 4 to give these components a level of volatility more in line with stocks and commodities.

Given the above choices, we replicated the returns on the Edhec, CISDM and HFRI indices. The CISDM and HFRI indices are calculated in the usual way, i.e. as portfolios of a (large) number of individual hedge funds.[51] The Edhec indices, however, are indices of indices, calculated as the first component of a principal component analysis of a large number of competing hedge fund indices.[52]

Table 5.1-5.3 show the sample properties of the monthly excess returns (over 1-month LIBOR) of the Edhec, CISDM and HFRI hedge fund indices as well as the synthetic funds designed to replicate them, over the period March 1999 – September (Edhec), May (CISDM), October (HFRI) 2006. When interpreting these results, it has to be kept in mind that the available hedge fund indices typically suffer from a variety of upward biases, including:

- **Self-reporting bias**: Only the more successful funds will report to a database.
- **Survivorship bias**: Some index providers remove funds that close down from the index's history.
- **Selection bias**: Database and index providers may have strict criteria to decide which funds to include in the database and/or index.
- **Backfill bias**: When a fund enters a database, typically its complete track record is included.

Although estimates very quite widely, it is generally thought that the combined upward bias in hedge fund indices due to the above adds up to 2-3% per annum. In this context two other observations are of interest as well. First, most hedge funds only report to one or two databases. This means that there is little overlap between databases and the indices derived from them. Second, it is not uncommon for recorded performance for the same fund to vary between databases. Together, this means that indices from different index providers may exhibit significantly different behaviour, with differences in monthly return sometimes exceeding 5% or even more!

[51] Although the CISDM indices reflect the performance of the median, instead of the average fund.

[52] Details on the composition and construction of the Edhec indices can be found in Edhec (2004).

	Index				Synthetic Fund			
Edhec Index	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
Convertible Arbitrage	9.22%	6.59%	0.03	0.23	8.68%	6.54%	0.12	0.23
Distressed Securities	13.40%	6.78%	0.40	0.44	10.79%	7.13%	0.59	0.46
Emerging Markets	16.23%	13.42%	0.04	0.65	18.37%	14.15%	1.13	0.66
Long/Short Equity	10.02%	8.64%	0.39	0.58	11.32%	8.74%	0.68	0.59
Equity Market Neutral	7.93%	2.52%	0.74	0.28	5.52%	2.29%	0.38	0.30
Fixed Income Arb.	6.83%	2.86%	-0.22	0.24	7.26%	3.07%	-0.21	0.24
Global Macro	8.82%	5.44%	0.81	0.40	9.41%	5.55%	1.07	0.42
Merger Arbitrage	8.03%	4.21%	-0.97	0.44	6.50%	4.05%	-0.25	0.43
Short Selling	2.10%	21.39%	0.00	-0.63	21.08%	23.85%	0.51	-0.60
CTA Global	6.11%	9.37%	0.07	-0.11	12.98%	9.60%	0.07	-0.10
Funds of Funds	8.96%	6.88%	1.06	0.48	10.09%	7.13%	1.37	0.48

Table 5.1. Sample properties Edhec index and synthetic fund returns over the period March 1999 - September 2006.

	Index				Synthetic Fund			
CISDM Index	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
Convertible Arbitrage	8.91%	5.45%	-0.20	0.24	6.83%	5.14%	-0.23	0.27
Distressed Securities	12.62%	6.21%	0.37	0.50	9.62%	6.34%	0.40	0.52
Emerging Markets	15.42%	11.47%	0.28	0.62	16.64%	12.11%	1.50	0.62
Long/Short Equity	10.19%	9.10%	0.83	0.56	11.85%	9.55%	1.17	0.60
Equity Market Neutral	7.76%	2.94%	0.65	0.34	6.27%	2.79%	0.18	0.34
Fixed Income Arb.	7.34%	1.56%	0.29	0.23	3.85%	1.60%	0.36	0.35
Global Macro	6.98%	4.15%	0.95	0.44	6.55%	3.93%	0.51	0.46
Merger Arbitrage	7.94%	3.96%	-0.14	0.44	6.14%	3.84%	-0.16	0.45
CTA (equal weighted)	6.97%	8.42%	0.36	-0.02	11.79%	8.48%	0.13	-0.07
Funds of Funds	7.53%	4.41%	0.23	0.48	6.99%	4.24%	0.99	0.49

Table 5.2. Sample properties CISDM index and synthetic fund returns over the period March 1999 - May 2006.

	Index				Synthetic Fund			
HFRI Index	**Mean**	**StDev**	**Skew**	**Corr**	**Mean**	**StDev**	**Skew**	**Corr**
Convertible Arbitrage	8.59%	5.53%	0.00	0.20	7.49%	5.42%	-0.01	0.22
Distressed Securities	13.27%	7.72%	0.32	0.48	12.36%	8.10%	0.43	0.49
Emerging Markets	**18.63%**	**16.94%**	**0.09**	**0.67**	**22.65%**	**18.14%**	**1.12**	**0.67**
Long/Short	**11.48%**	**10.68%**	**0.88**	**0.57**	**13.52%**	**11.26%**	**1.13**	**0.61**
Equity Market Neutral	**6.26%**	**2.92%**	**0.60**	**-0.01**	**6.77%**	**2.98%**	**0.22**	**-0.06**
Fixed Income Arb.	7.81%	3.61%	0.09	0.49	7.70%	4.00%	0.21	0.50
Global Macro	**9.05%**	**6.50%**	**0.53**	**0.30**	**10.38%**	**6.75%**	**0.60**	**0.30**
Merger Arbitrage	8.03%	4.10%	-0.67	0.38	7.05%	3.71%	-0.27	0.35
Short Selling	**4.15%**	**23.37%**	**0.05**	**-0.56**	**24.03%**	**26.64%**	**0.73**	**-0.53**
Funds of Funds	**8.40%**	**7.25%**	**0.66**	**0.50**	**11.63%**	**7.46%**	**1.50**	**0.49**

Table 5.3. Sample properties HFRI index and synthetic fund returns over the period March 1999 - October 2006.

To facilitate comparison, in the above tables the entries in bold are the cases where the average return on the synthetic fund exceeds that of the index. The entries in normal font are the cases where the average return on the index exceeds that of the synthetic fund.

From table 5.1-5.3, we can draw a number of very interesting conclusions:

- Our synthetic funds beat the Edhec indices 8 out of 11 times, the CISDM indices 3 out of 8 times, and the HFRI indices 6 out of 9 times. Given the upward bias present in these indices, this is an extremely good result. If we assumex the average index return is upwardly biased by 2.5%, then almost all synthetic funds would comfortably outperform the indices they are designed to replicate.

- Without exception, the volatility of the index return and the correlation with the 50/50 stock/bond portfolio are very accurately replicated.

- Keeping in mind the sensitivity of the traditional skewness measure for extreme observations and the fact that we only have 92 observations available, skewness is very satisfactory replicated as well.

- Across the three index families, the emerging markets, long/short, global macro, short selling and funds of funds indices are consistently dominated by their replicating synthetic funds.

The entries in table 5.1-5.3 are calculated as per the end of September (Edhec), May (CISDM), October (HFRI) 2006. However, the synthetic funds started in March 1999. It is therefore interesting to have a look at the evolution of the various sample parameters over time. The results for the Edhec funds of funds index are reported figure 5.1-5.4. All 4 graphs start in March 2001 as with less than 24 observations reliability would simply be too low.

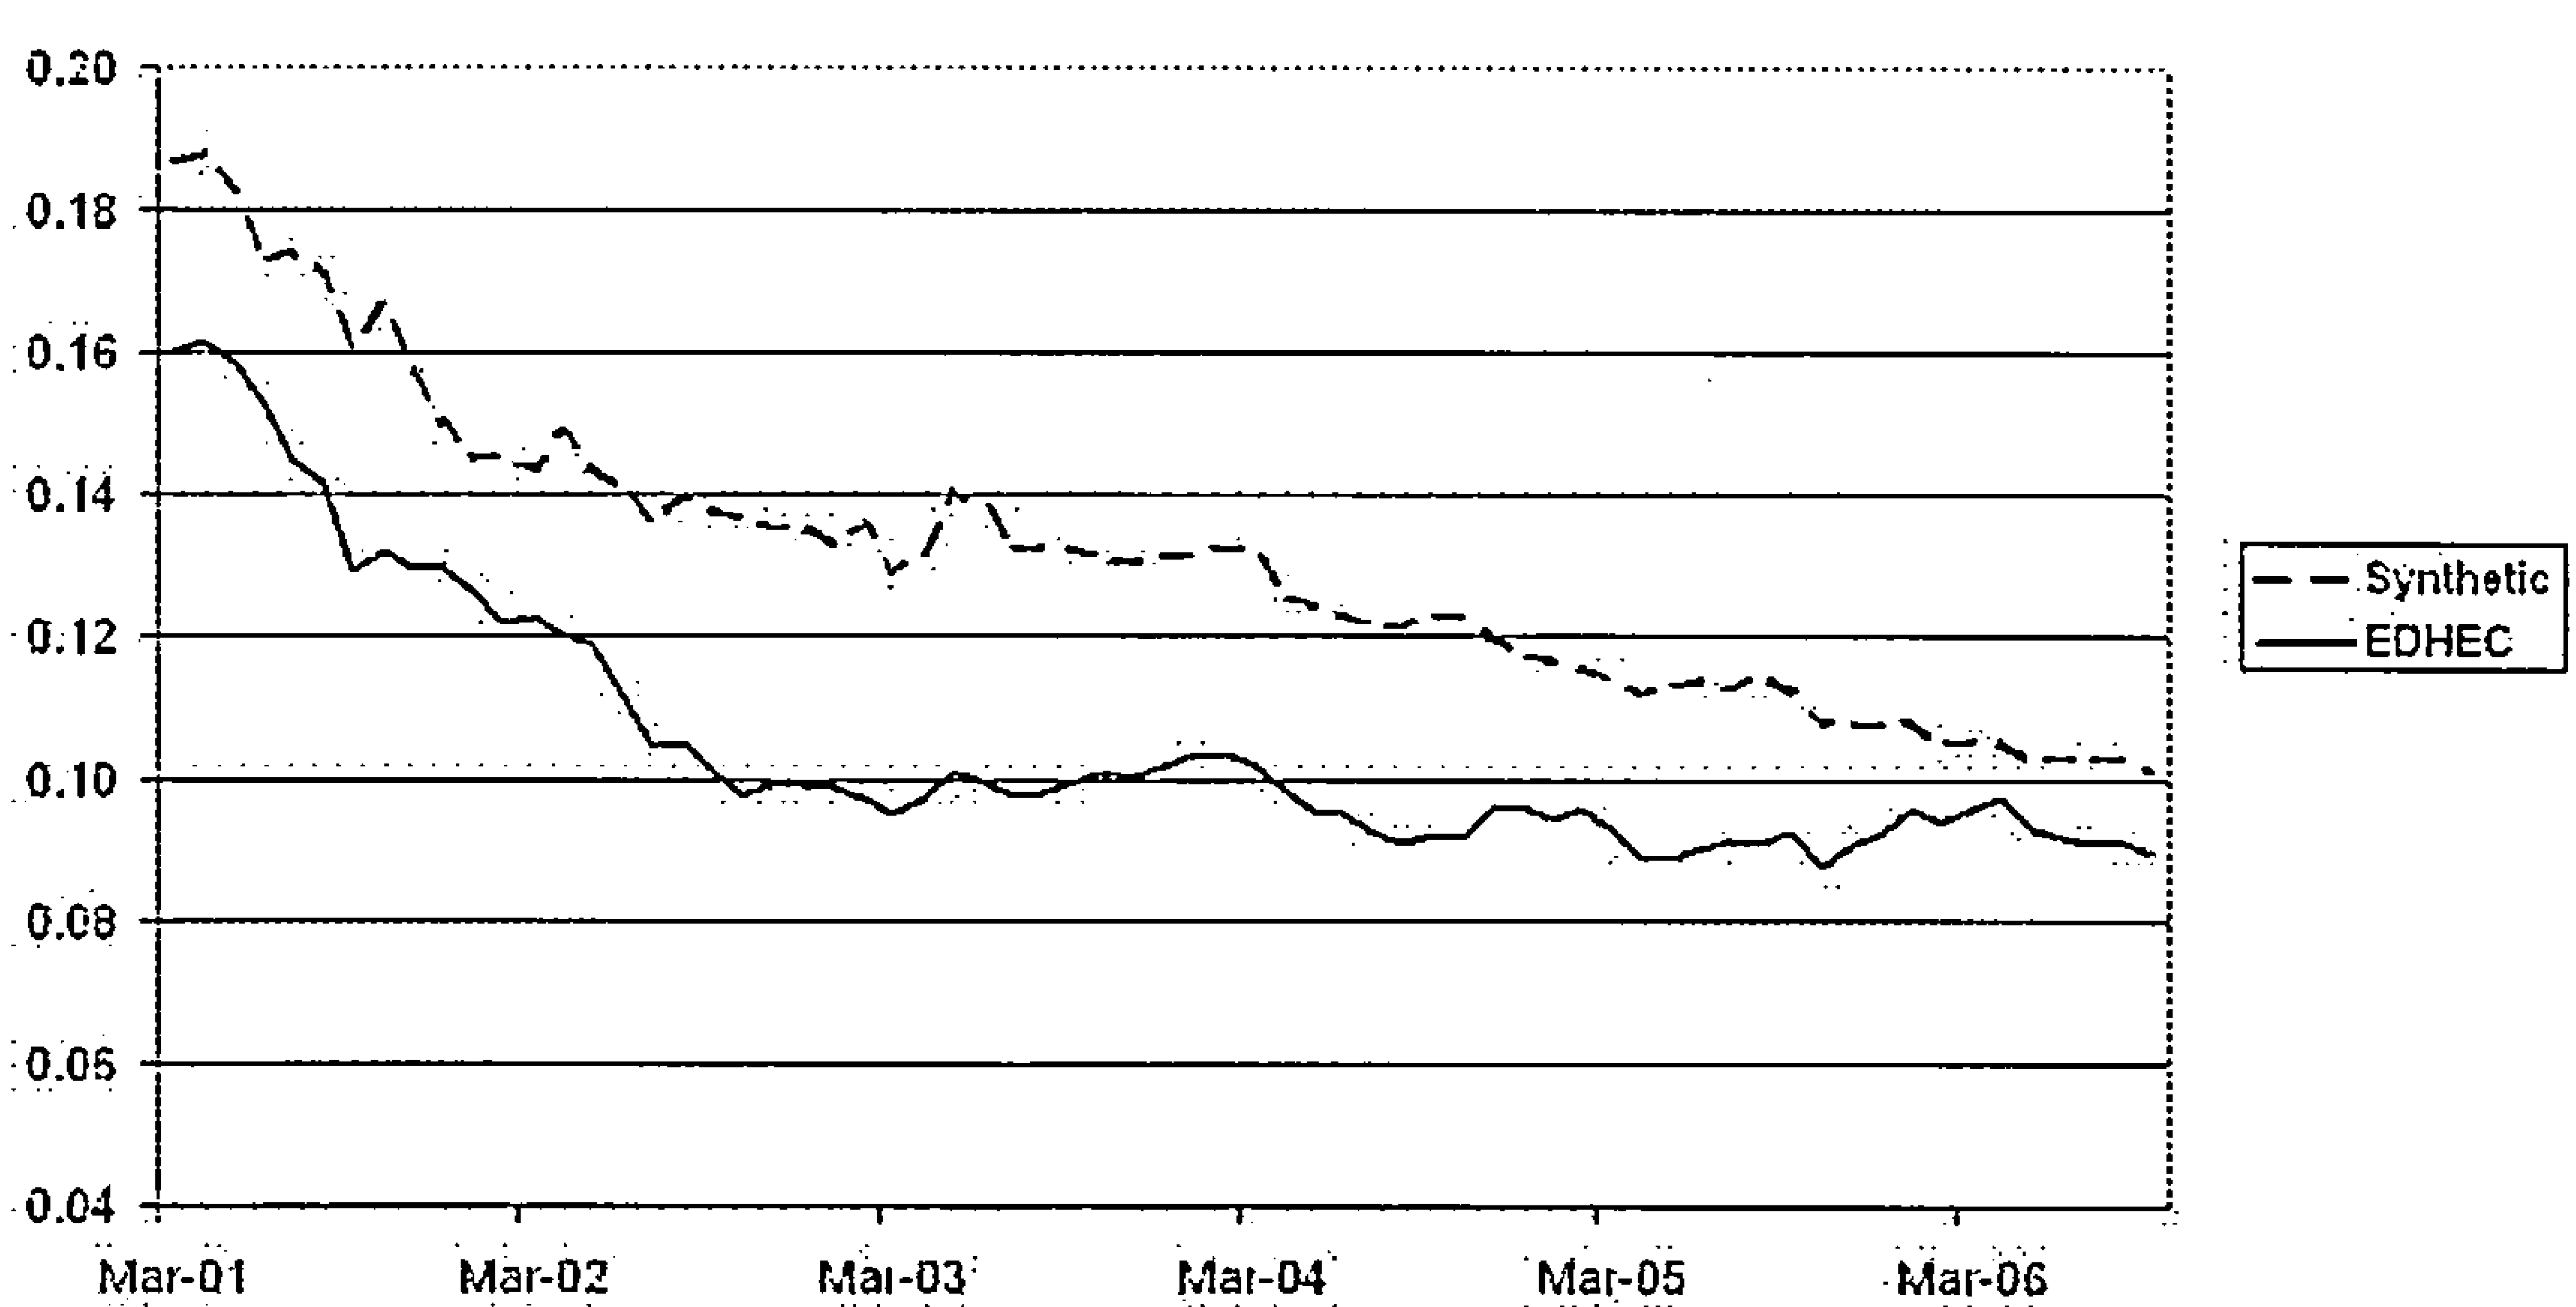

Figure 5.1. Average return Edhec funds of funds index and replicating synthetic fund, March 2001 – September 2006.

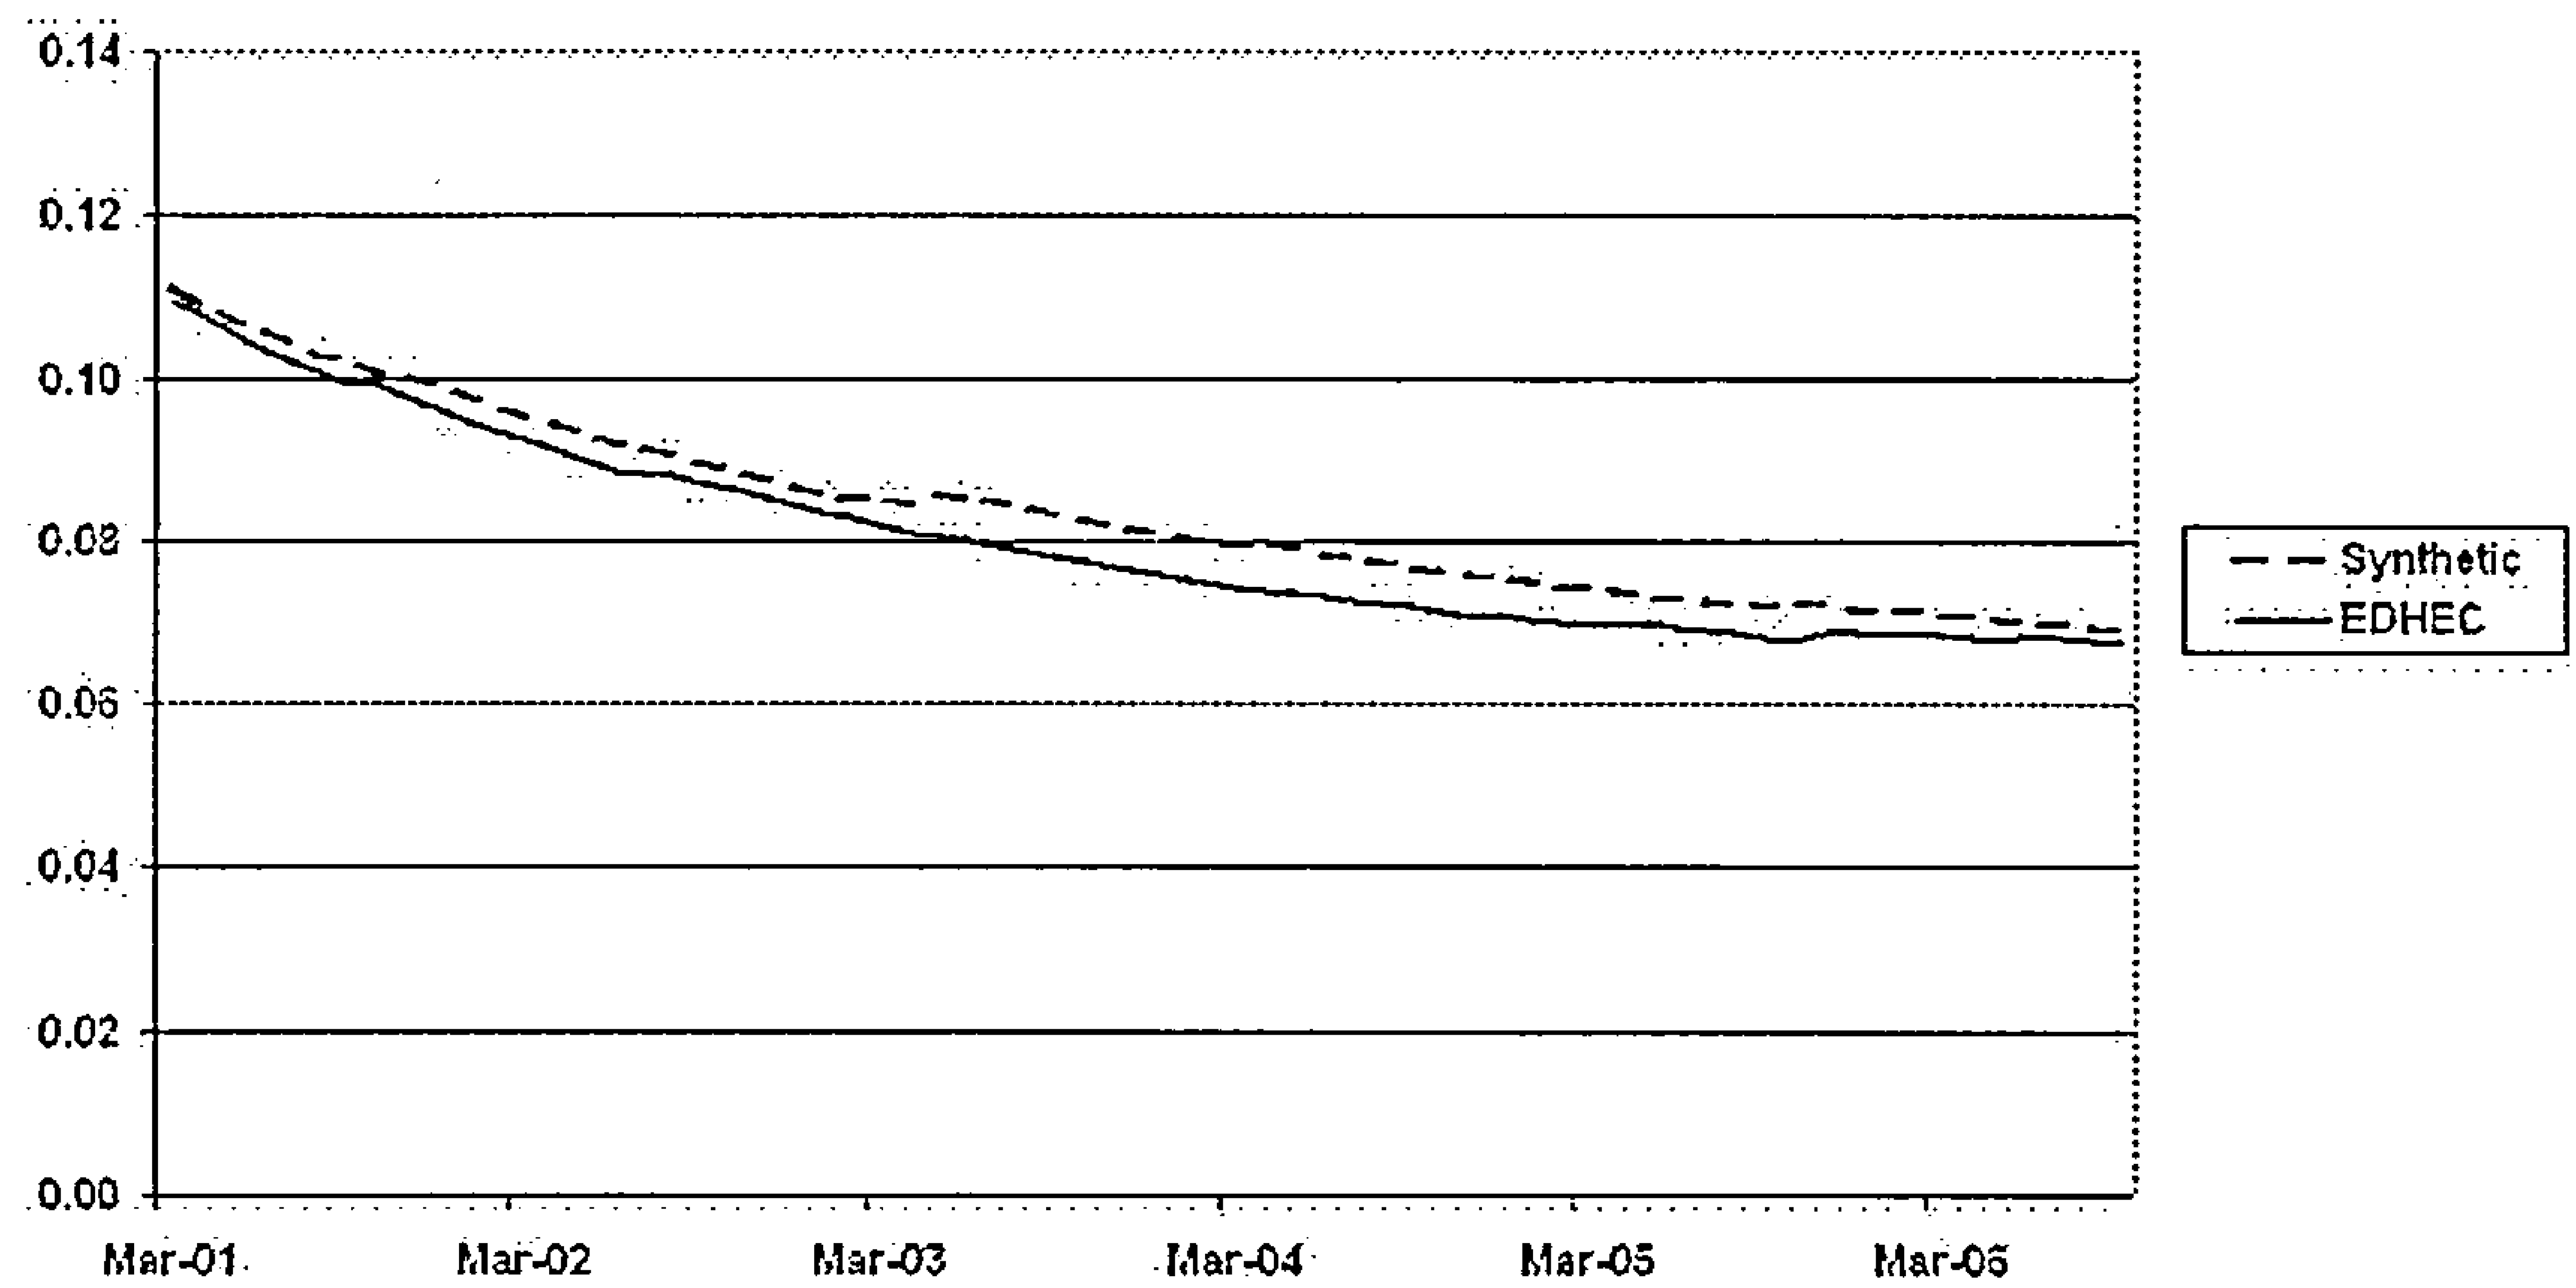

Figure 5.2. Volatility Edhec funds of funds index and replicating synthetic fund, March 2001 – September 2006.

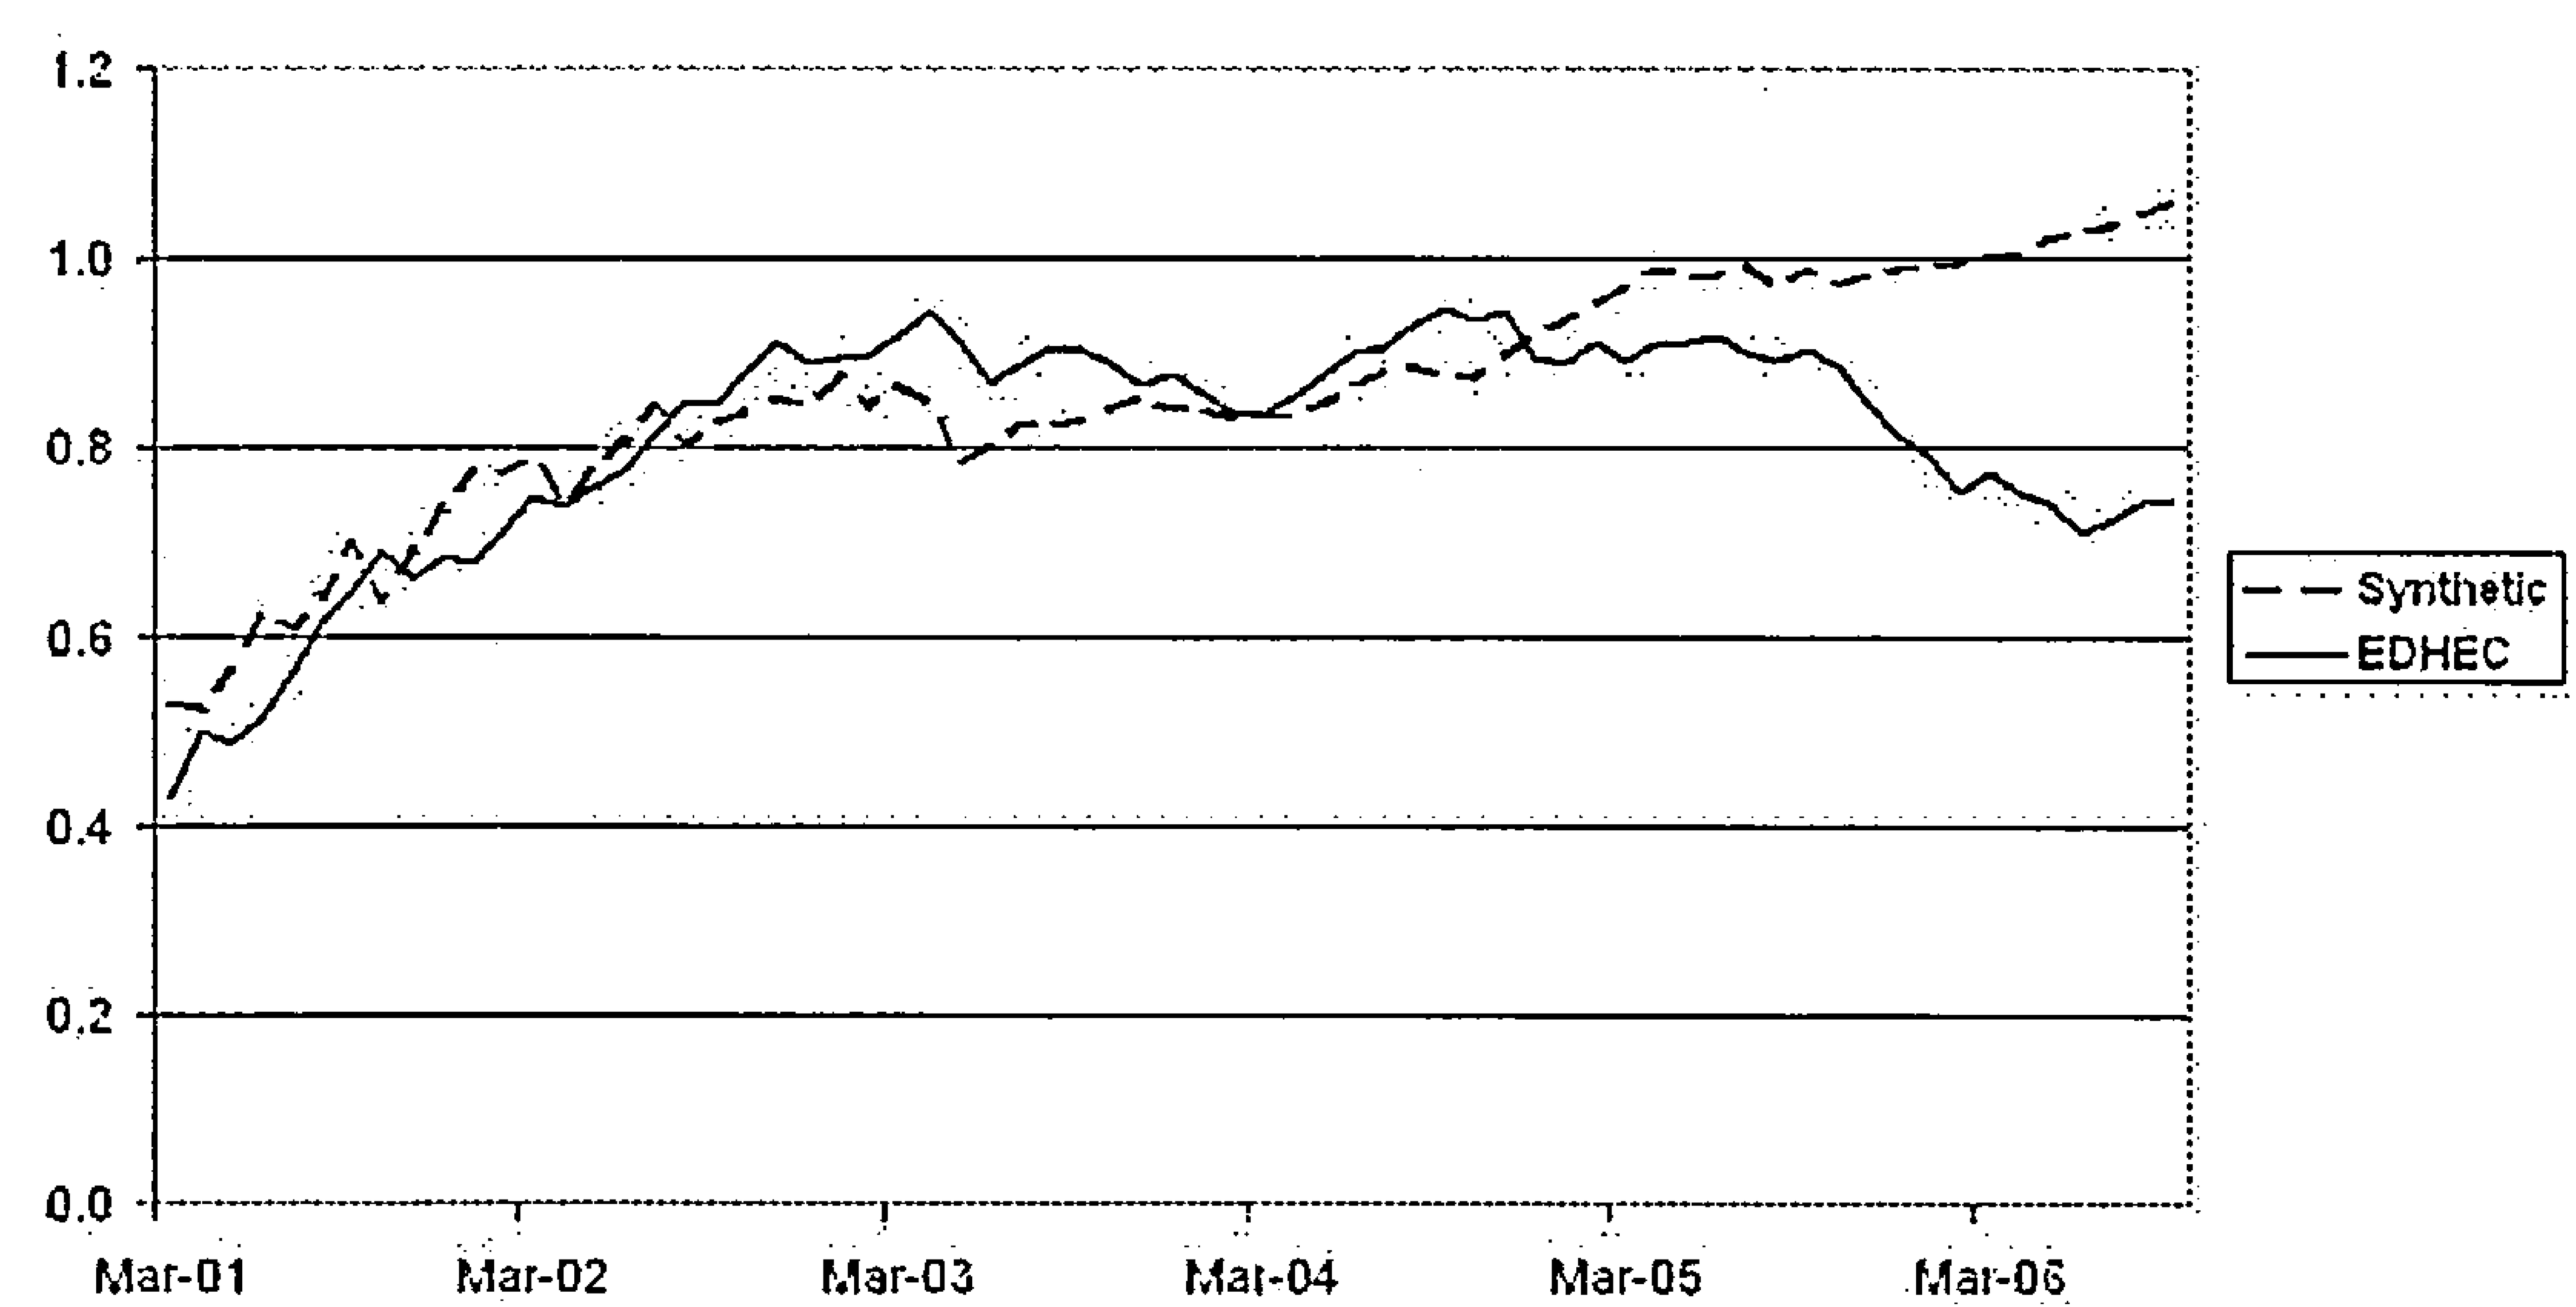

Figure 5.3. Skewness Edhec funds of funds index and replicating synthetic fund, March 2001 – September 2006.

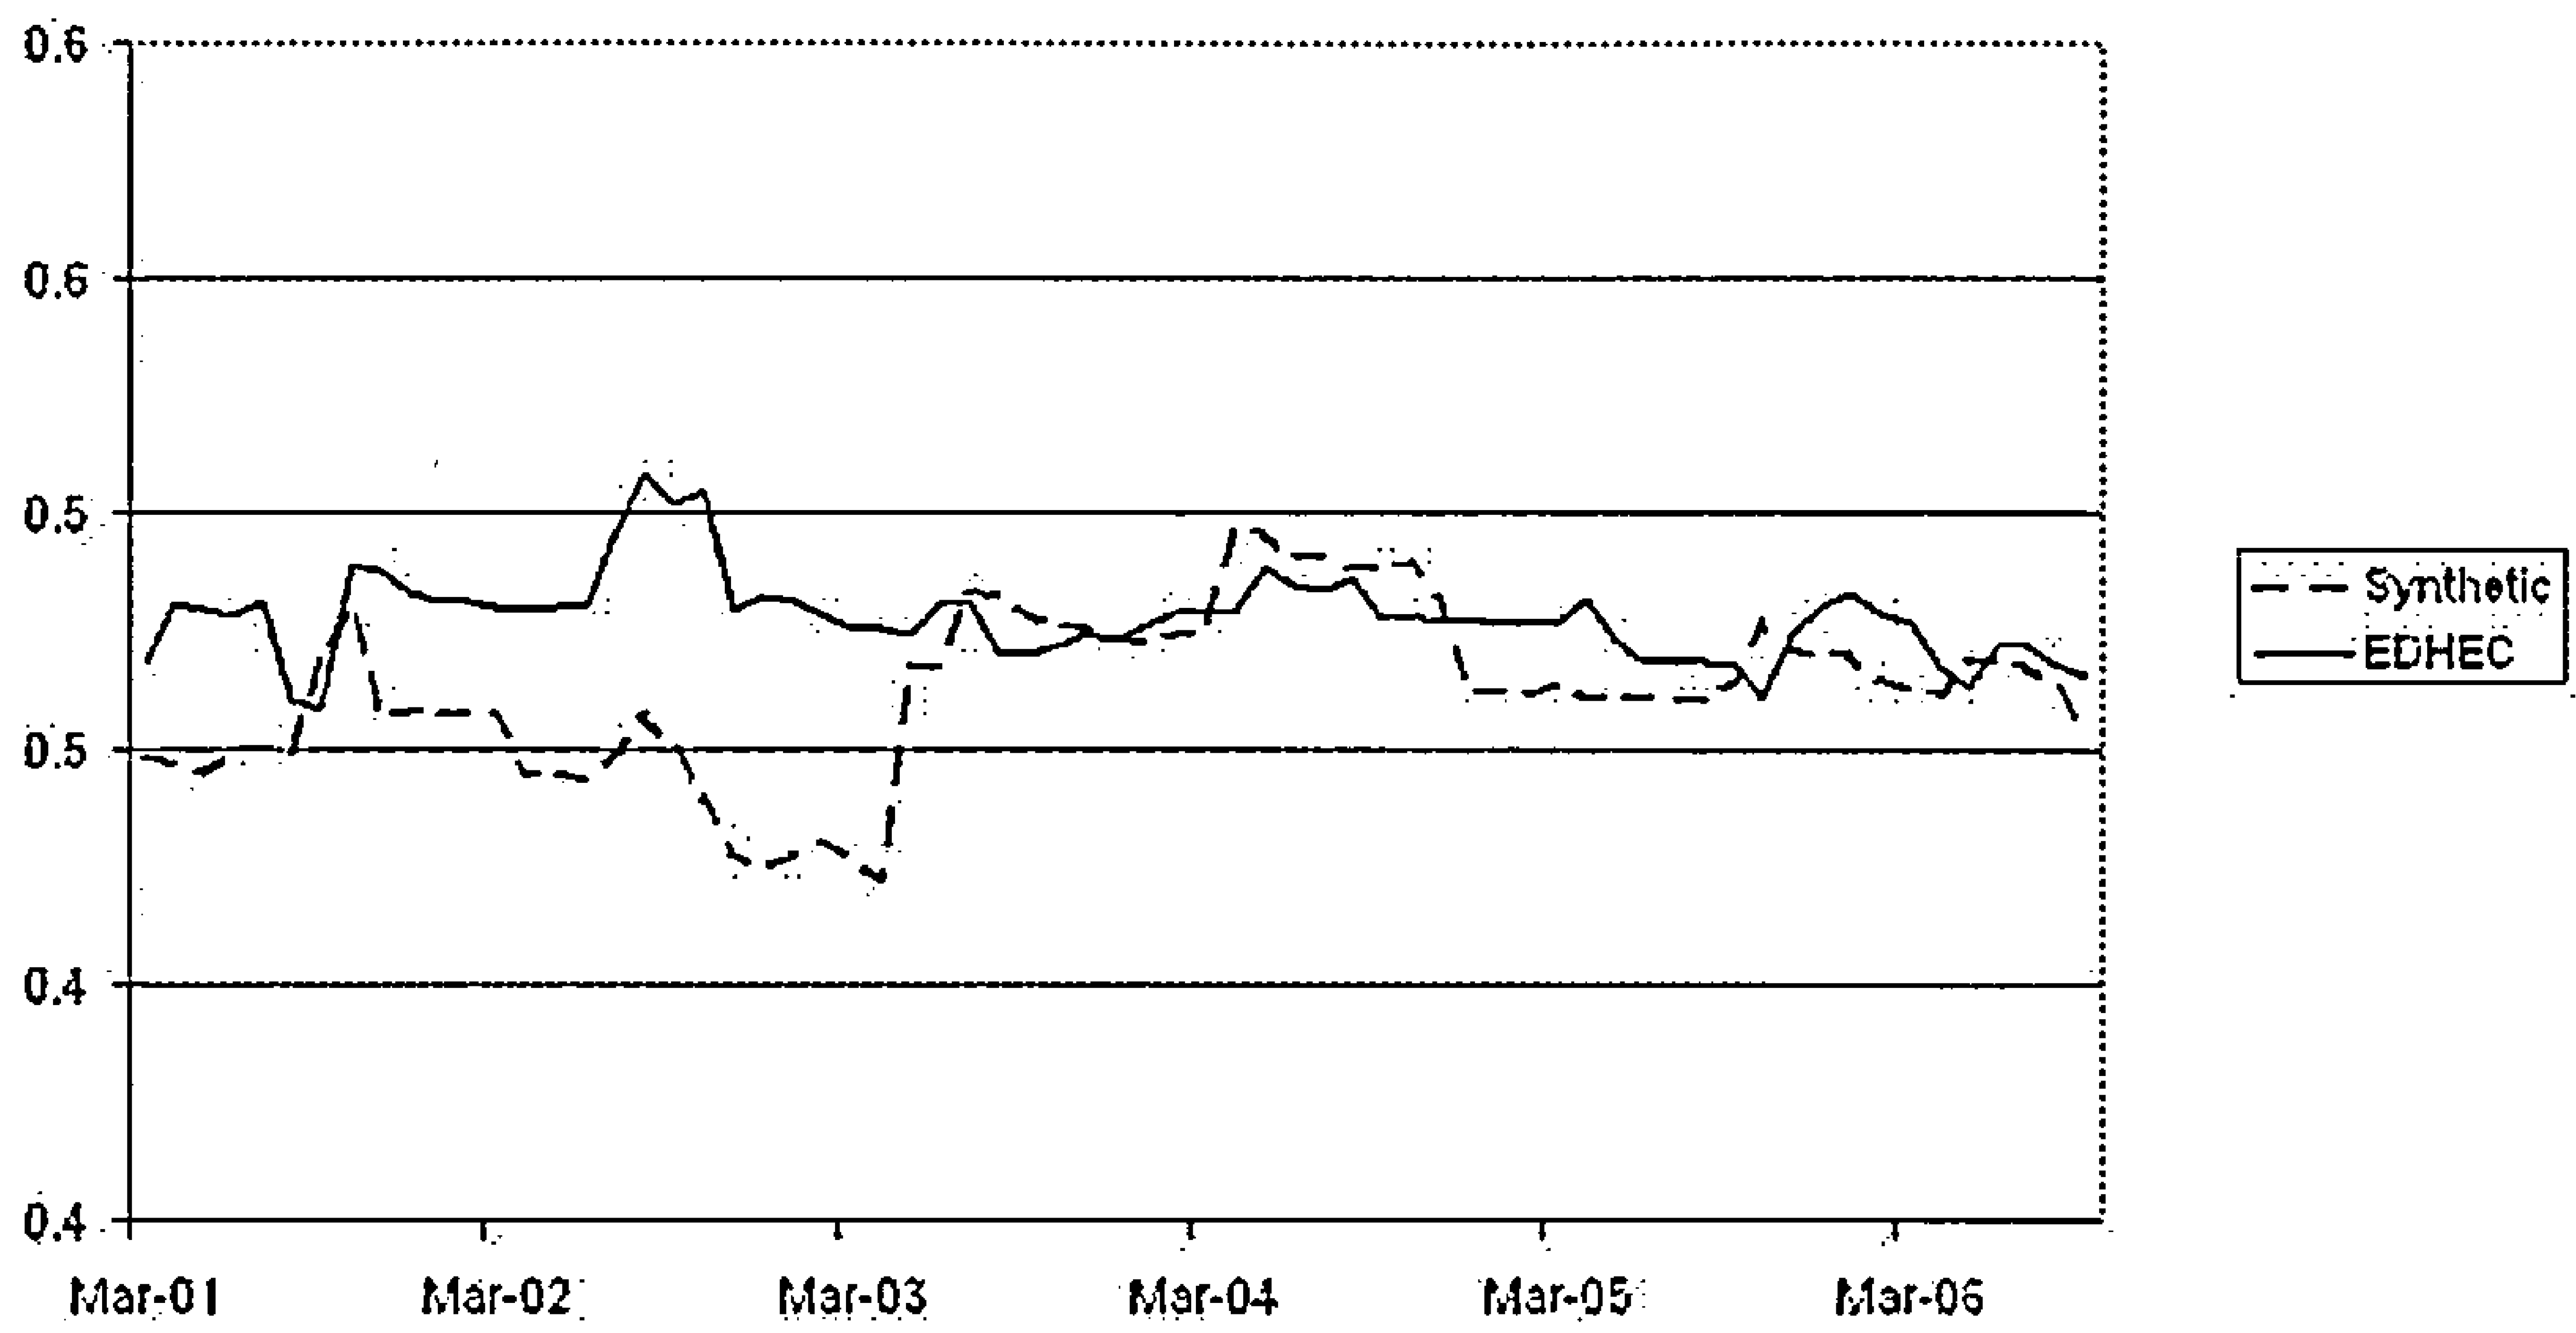

Figure 5.4. Correlation Edhec funds of funds index and replicating synthetic fund with portfolio of 50% S&P 500 and 50% T-bonds, March 2001 – September 2006.

The graphs in Figure 5.2-5.4 show that the synthetic fund matches the risk parameters of the Edhec funds of funds index directly from the start. The graphs also show that the synthetic fund is quite capable of following the changes in the index parameters over time. From Figure 5.3 we see that the difference in skewness reported in Table 5.1 only arose over the last year. Before that, the match was almost perfect. Finally, Figure 5.1 shows that over time the synthetic fund's average return has been consistently higher than that of the index. Similar results were obtained for all other indices.

From the above it is clear that for all 3 families of hedge fund indices the replication approach more than achieves its objective of providing investors with returns with the same characteristics as hedge fund indices, but without actually investing in hedge funds. *In all cases, risk profiles are very similar and, taking into account the upward bias present in these indices, the average return is markedly better.* In addition, for a moment assuming these indices are investable, the synthetic fund returns have no liquidity or transparency problems, which, as argued in Kat (2006), makes them even more valuable.

5.2 Sensitivity analyses

As in Section 4.2, in this section is also examined the influence of the transaction costs, the reference portfolio and the reserve asset on the replication results. To obtain a good impression of how these factors influence the bottom line, we study two significantly different cases, the HFRI Equity Market Neutral index and the HFRI Long-Short index. The former has low correlation with the market, and particularly with the reference portfolios considered, while the latter has not.

5.2.1 Transaction Costs

Depending on the targeted properties and the way markets behave, the replication technique introduced in Chapter 2 tend to be quite dynamic, so an obvious question to ask is whether the high turnover in replication trading strategies is not excessively costly.

Table 5.4 shows the sample properties of the replicated returns on the HFRI Equity Market Neutral index for various bid-offer spreads (0bps, 4bps, etc.) and rebalancing frequencies (daily, once every 2 days, and once every 3 days) over the period March 1999 – October 2006. Note that the case of a 50bps spread is included for illustrative purposes only. As discussed earlier, in practice one would not expect to trade on such a wide spread very often.

	Mean	StDev	Skew	Corr
HFRI EMN Index	6.26%	2.92%	0.60	-0.01
Synthetic 0bps, daily	6.79%	2.98%	0.22	-0.06
Synthetic 4bps, daily	6.75%	2.98%	0.20	-0.06
Synthetic 14bps, daily	6.65%	2.97%	0.19	-0.06
Synthetic 50bps, daily	6.30%	2.96%	0.11	-0.06
Synthetic 0bps, 2 daily	6.83%	3.00%	0.20	-0.07
Synthetic 4bps, 2 daily	6.81%	2.99%	0.19	-0.07
Synthetic 14bps, 2 daily	6.77%	2.99%	0.18	-0.07
Synthetic 50bps, 2 daily	6.60%	2.99%	0.16	-0.07
Synthetic 0bps, 3 daily	6.65%	2.98%	-0.01	-0.07
Synthetic 4bps, 3 daily	6.64%	2.98%	-0.01	-0.07
Synthetic 14bps, 3 daily	6.60%	2.98%	-0.02	-0.07
Synthetic 50bps, 3 daily	6.49%	2.97%	-0.04	-0.07

Table 5.4. Sample properties HFRI Equity Market Neutral index and synthetic fund returns over the period March 1999 - October 2006.

From the above table we can draw two interesting conclusions. First, the risk profile generated is largely independent of the level of transaction costs as well as the rebalancing frequency. Second, the impact of transaction costs on the mean is small and drops when rebalancing less often. Note that even with a 50bps spread the synthetic fund would still have beaten the index.

	Mean	StDev	Skew	Corr
HFRI L-S Index	11.48%	10.68%	0.88	0.57
Synthetic 0bps, daily	13.55%	11.26%	1.13	0.61
Synthetic 4bps, daily	13.48%	11.26%	1.13	0.61
Synthetic 14bps, daily	13.18%	11.25%	1.13	0.65
Synthetic 50bps, daily	12.14%	11.23%	1.13	0.61
Synthetic 0bps, 2 daily	13.10%	11.25%	1.15	0.61
Synthetic 4bps, 2 daily	13.04%	11.25%	1.15	0.61
Synthetic 14bps, 2 daily	12.86%	11.25%	1.15	0.61
Synthetic 50bps, 2 daily	12.34%	11.24%	1.15	0.61
Synthetic 0bps, 3 daily	12.78%	11.24%	1.14	0.61
Synthetic 4bps, 3 daily	12.73%	11.24%	1.14	0.61
Synthetic 14bps, 3 daily	12.65%	11.24%	1.14	0.61
Synthetic 50bps, 3 daily	12.29%	11.23%	1.15	0.61

Table 5.5. Sample properties HFRI Long-Short index and synthetic fund returns over the period March 1999 - October 2006.

Table 5.5 shows the sample properties of the replicated returns on the HFRI Long-Short index. Again, we see that the risk profile generated is independent of the level of transaction costs as well as the rebalancing frequency. The impact of transaction costs on the mean is slightly stronger than before, as the index being replicated is a lot more volatile. For realistic spreads, however, the results are still very satisfactory, with average costs again dropping with the rebalancing frequency. As in the previous case, the synthetic fund would have beaten the actual index even with a 50bps spread.

5.2.2 Reference Portfolio

As said in Section 4.2, in this section we study how much the bottom line result depends on the reference portfolio.

As before, the reference portfolio is taken to consist of three different mixes of S&P 500 and T-bond futures. More specifically, the portfolio weights 20/80, 50/50 and 80/20 are considered. It is assumed that throughout the reserve asset consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures.[53]

Table 5.6 shows the sample properties of the replicated returns on the HFRI Equity Market Neutral index for reference portfolios consisting of a 50/50 and 80/20 mix of T-bonds and S&P 500 over the period March 1999 – October 2006. The results for a 20/80 mix are very similar and therefore not reported.

[53] Since the volatility of the various asset classes is quite different, before forming the portfolio, we leveraged the Eurodollar and 5Y note by a factor 5, and the 10Y note by a factor 4 to give these components a level of volatility more in line with stocks and commodities. Note that this will increase the trading volume in these contracts.

	50/50 Reference Portfolio				80B/20E Reference Portfolio			
	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
HFRI EMN Index	6.26%	2.92%	0.60	-0.01	6.26%	2.92%	0.60	0.07
Synthetic 0bps, daily	6.79%	2.98%	0.22	-0.06	6.38%	2.88%	0.04	0.04
Synthetic 4bps, daily	6.75%	2.98%	0.20	-0.06	6.33%	2.88%	0.04	0.04
Synthetic 14bps, daily	6.65%	2.97%	0.19	-0.06	6.18%	2.88%	0.04	0.04
Synthetic 0bps, 2 daily	6.83%	3.00%	0.20	-0.07	6.43%	2.88%	0.07	0.04
Synthetic 4bps, 2 daily	6.81%	2.99%	0.19	-0.07	6.40%	2.88%	0.08	0.04
Synthetic 14bps, 2 daily	6.77%	2.99%	0.18	-0.07	6.36%	2.88%	0.07	0.04
Synthetic 0bps, 3 daily	6.65%	2.98%	-0.01	-0.07	6.43%	2.88%	0.07	0.05
Synthetic 4bps, 3 daily	6.64%	2.98%	-0.01	-0.07	6.41%	2.88%	0.07	0.05
Synthetic 14bps, 3 daily	6.60%	2.98%	-0.02	-0.07	6.36%	2.88%	0.06	0.05

Table 5.6. Sample properties HFRI Equity Market Neutral index and synthetic fund returns over the period March 1999 - October 2006.

From Table 5.6 we see that for either reference portfolio the volatility and correlation of the synthetic fund match those of the index almost exactly. The means are quite similar as well. Comparing the results for different bid-ask spreads it appears that the 80/20 case requires a little more trading than the 50/50 case though, as the mean drops slightly faster for the 80/20 than the 50/50 portfolio when the bid-ask spread increases.

	50/50 Reference Portfolio				80B/20E Reference Portfolio			
	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
HFRI L-S Index	11.48%	10.68%	0.88	0.57	11.48%	10.68%	0.88	0.15
Synthetic 0bps, daily	13.55%	11.26%	1.13	0.61	15.30%	11.31%	0.92	0.14
Synthetic 4bps, daily	13.48%	11.26%	1.13	0.61	15.08%	11.31%	0.92	0.14
Synthetic 14bps, daily	13.18%	11.25%	1.13	0.65	14.50%	11.29%	0.91	0.14
Synthetic 0bps, 2 daily	13.10%	11.25%	1.15	0.61	12.93%	11.03%	0.66	0.13
Synthetic 4bps, 2 daily	13.04%	11.25%	1.15	0.61	12.81%	11.02%	0.66	0.13
Synthetic 14bps, 2 daily	12.86%	11.25%	1.15	0.61	12.56%	11.02%	0.66	0.13
Synthetic 0bps, 3 daily	12.78%	11.24%	1.14	0.61	13.61%	11.10%	0.71	0.14
Synthetic 4bps, 3 daily	12.73%	11.24%	1.14	0.61	13.54%	11.10%	0.71	0.14
Synthetic 14bps, 3 daily	12.65%	11.24%	1.14	0.61	13.32%	11.09%	0.71	0.14

Table 5.7. Sample properties HFRI Long-Short index and synthetic fund returns over the period March 1999 - October 2006.

Table 5.7 shows the sample properties of the replicated returns on the HFRI Long-Short index. Despite the much higher volatility of this index, the results for both reference portfolios are again very similar. In both cases the standard deviation, skewness and correlation of the synthetic fund almost exactly match those of the index. Again, we see that the 80/20 case requires more trading than the 50/50 case, as the drop in mean due to a higher bid-offer spread is slightly more pronounced in the 80/20 case.

5.2.3 Reserve Asset

As in Subsection 4.2.3, we study how sensitive the replication results are to the choice of reserve asset.

As before, in this subsection two different reserve assets are considered. Reserve Asset 1 consists of an equally-weighted portfolio of 3-month Eurodollar, 5-year note, 10-year note, S&P 500, Russell 2000 and GSCI futures, where, to compensate substantial differences in volatility, we leveraged the Eurodollar and 5-year note by a factor 5, and the 10-year note by a factor 4. Reserve Asset 2 is much simpler and consists of an equally weighted portfolio of 1-month Libor, Russell 2000 and crude oil futures.

Using Reserve Asset 1 and 2, Table 5.8 shows the sample properties of the replicated returns on the HFRI Equity Market Neutral index for various bid-offer spreads (0bps, 4bps, etc.) and rebalancing frequencies (daily, once every 2 days, and once every 3 days) over the period March 1999 – October 2006. Comparing the entries at both sides of the table, it is clear that the differences are very minor. Reserve Asset 2 appears to be slightly better in replicating the positive skewness found in the HFRI EMN index. This is due to the fact that the returns on Reserve Asset 2 are slightly more skewed than those on Reserve Asset 1, which makes it easier to produce skewed synthetic fund returns.

	Reserve Asset 1				Reserve Asset 2			
	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
HFRI EMN Index	6.26%	2.92%	0.60	-0.01	6.26%	2.92%	0.60	-0.01
Synthetic 0bps, daily	6.79%	2.98%	0.22	-0.06	6.54%	2.84%	0.37	-0.09
Synthetic 4bps, daily	6.75%	2.98%	0.20	-0.06	6.50%	2.84%	0.36	-0.09
Synthetic 14bps, daily	6.65%	2.97%	0.19	-0.06	6.44%	2.84%	0.37	-0.09
Synthetic 0bps, 2 daily	6.83%	3.00%	0.20	-0.07	6.47%	2.84%	0.32	-0.09
Synthetic 4bps, 2 daily	6.81%	2.99%	0.19	-0.07	6.46%	2.84%	0.32	-0.09
Synthetic 14bps, 2 daily	6.77%	2.99%	0.18	-0.07	6.42%	2.84%	0.31	-0.09
Synthetic 0bps, 3 daily	6.65%	2.98%	-0.01	-0.07	6.68%	2.84%	0.44	-0.09
Synthetic 4bps, 3 daily	6.64%	2.98%	-0.01	-0.07	6.67%	2.84%	0.44	-0.09
Synthetic 14bps, 3 daily	6.60%	2.98%	-0.02	-0.07	6.64%	2.84%	0.44	-0.09

Table 5.8. Sample properties HFRI Equity Market Neutral index and synthetic fund returns over the period March 1999 - October 2006.

	Reserve Asset 1				Reserve Asset 2			
	Mean	StDev	Skew	Corr	Mean	StDev	Skew	Corr
HFRI L/S Index	11.48%	10.68%	0.88	0.57	11.48%	10.68%	0.88	0.57
Synthetic 0bps, daily	13.55%	11.26%	1.13	0.61	11.68%	10.93%	0.93	0.57
Synthetic 4bps, daily	13.48%	11.26%	1.13	0.61	11.58%	10.93%	0.92	0.57
Synthetic 14bps, daily	13.18%	11.25%	1.13	0.65	11.33%	10.92%	0.92	0.57
Synthetic 0bps, 2 daily	13.10%	11.25%	1.15	0.61	11.50%	11.11%	0.81	0.56
Synthetic 4bps, 2 daily	13.04%	11.25%	1.15	0.61	11.46%	11.10%	0.90	0.56
Synthetic 14bps, 2 daily	12.86%	11.25%	1.15	0.61	11.22%	11.09%	0.86	0.56
Synthetic 0bps, 3 daily	12.78%	11.24%	1.14	0.61	12.48%	10.99%	1.14	0.57
Synthetic 4bps, 3 daily	12.73%	11.24%	1.14	0.61	12.46%	10.99%	1.14	0.57
Synthetic 14bps, 3 daily	12.65%	11.24%	1.14	0.61	12.38%	10.99%	1.14	0.57

Table 5.9. Sample properties HFRI Long/Short index and synthetic fund returns over the period March 1999 - October 2006.

Table 5.9 shows the sample properties of the replicated returns on the HFRI Long-Short index. Despite the much higher target volatility, the results for both reserve assets match up very well. The mean returns for Reserve Asset 2 are slightly lower than for Reserve Asset 1, but the risk parameters are very close.

6. Conclusion

Much of investors' current interest in hedge funds derives from the fact that traditional asset classes seem to lack opportunity these days. 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 impressive and many investors seem more than happy to use this as a guide for future returns. Given today's low risk premiums, as well as the current size of the hedge fund industry itself, a repeat of the last 10-15 years is extremely unlikely, however.

Investing in alternatives comes with many drawbacks, including due diligence, liquidity, capacity, transparency and style drift 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, when reality kicks in and investors realize that hedge funds are no longer the money machines they once were (thought to be), their attitude will undoubtedly 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 has shown stellar performance over the last 3 years.

Although something cannot be created out of nothing, in Chapter 2 it was shown that it is possible to design dynamic trading strategies, which generate returns similar to those of individual hedge funds and funds of hedge funds. Since this is accomplished by trading (futures on) traditional assets only, these strategies avoid the typical drawbacks surrounding hedge fund and other alternative investments. As such, the synthetic hedge fund returns are clearly to be preferred over real hedge fund returns.

In Chapter 3 the hedge fund return replication technique was used to evaluate the net-of-fee performance of 875 funds of hedge funds and 2073 individual hedge funds. The results indicate that the majority of hedge funds and funds of hedge funds have not provided their investors with returns, which they could not have generated themselves by mechanically trading a basket of liquid futures contracts. Over time, it can be observed a substantial deterioration in overall hedge fund performance. In addition, it was found a tendency for the performance of successful funds also to deteriorate over time, which supports the hypothesis that increased fund size tends to hurt future performance.

Overall, only 22.5% of the 2073 individual funds and 18.6% of the funds of funds in the sample were able to beat the benchmark. This means that in terms of the KP measure individual hedge funds and funds of hedge funds are not too different. At first sight this may seem odd. With funds of funds putting on an additional layer of fees, one would expect the results for funds of funds to be substantially worse than for individual hedge funds. However, funds of funds diversify and given the low correlation between individual hedge funds, this means that the risk characteristics of fund of funds returns are typically a lot more conservative than those of individual hedge funds, which is reflected in the efficiency measure outcomes.

Compared with the various hedge fund performance evaluation studies that have been carried out over the last couple of years, these results are quite unusual. Often, the conclusion from hedge fund performance studies is that hedge funds generate superior returns, not inferior. This once again indicates how tricky factor model based performance evaluation can be. As long as one cannot be sure that all relevant risk factors are accounted for, it is impossible to know whether unexplained returns are indeed true alpha or just unexplained because one or more risk factors were left out or specified incorrectly. The methodology in Chapter 3 is more robust, as it relies on a simple principle: "if it can be replicated, it cannot be superior". Of course, some assumptions were necessary as well, but these are less crucial for the final outcome of the evaluation than the kind of assumptions required to make factor model based alphas work.

Should investors rush out to buy into those funds with the highest KP measures? Although tempting, the answer is no. The core problem of performance evaluation is separating luck and skill. With a limited set of data, however, it is impossible to make a clean cut, whatever the method used. The KP measure is founded on the idea that in the long run, risk and return are related, irrespective of how a given risk profile is obtained. When there are not enough data available to properly observe 'the long run', however, the efficiency measure becomes prone to sampling error. If the available dataset is limited, it is very hard to identify the presence of any extreme (but compensated) risks for example, since by definition extreme events only occur infrequently. A fund manager may have been taking the most horrific risks, but if so far he has been lucky, the premium collected for taking on those risks will show from his track record, but the risk will not. Likewise, one or more risk factors may have done extremely well over a prolonged period of time. This will bias the available sample, which in turn may have a significant impact on the outcome of the evaluation.

Since performance evaluations over relatively short time periods will always leave us with a considerable degree of uncertainty, a high KP measure should first and foremost be interpreted as a signal that further due diligence is warranted. One can only speak of truly superior performance if such follow-up research shows that the good evaluation outcome was not simply due to luck. In other words, that the manager in question has generated the observed excess return without taking any extreme risks and that all the relevant risk factors behaved in a more or less representative manner during the period under consideration. Questions like these can typically not be answered satisfactorily within a purely quantitative framework and require a thorough understanding of hedge fund strategies. No matter how sophisticated the econometrics, proper performance evaluation will therefore always remain a combination of science and art.

Finally, it has to be noted that although in terms of the returns delivered to investors, most funds of funds do not seem to add value, this does not mean there is no economic reason for funds of funds to exist. Most private and smaller institutional investors do not have the skills and/or resources required to perform the necessary due diligence that comes with hedge fund investment. In addition, given typical minimum investment requirements, small private investors will often lack sufficient funds to build up a well-diversified hedge fund portfolio. They therefore have no choice. If they want hedge funds, they will have to go through a fund of funds.

Large institutions do have a choice. Most of them, however, will still prefer to go the fund of funds route. This is quite surprising given the amount of fees that could be saved by skipping the middlemen. Apart from believing that fund of funds managers add enough value to justify their fees (which research has shown to be unlikely), part of the reason that many large institutions still go for funds of funds lies in the fact that the interests of institutional asset managers are typically not correctly lined up with the interests of those whose money they manage. As a result, job protection becomes an important consideration. By investing in a fund of funds, instead of picking hedge funds themselves, institutions avoid having to take responsibility for the bottom-line fund selection. In the end, all they can be held responsible for is the decision to invest in hedge funds and the selection of the fund of funds that they invested in; risks, which can easily be hedged by not making a move until others do, hiring a big name consultant and a big name fund manager, as most institutions do.

In Chapter 4, four out-of-sample tests were carried out of the Chapter 2 synthetic fund creation technique. The test results show that the resulting strategies are indeed capable of accurately generating returns with a variety of properties, including zero and even negative correlation with stocks and bonds. Under difficult conditions, the tests also yield impressive average excess returns for the synthetic funds studied. Combined with their liquid and transparent nature, this confirms that synthetic funds are an attractive alternative to direct investment in alternative asset classes such as (funds of) hedge funds, commodities, etc. Undoubtedly, investors will need time to come to grips with the concept, but given their benefits, there is no doubt synthetic funds have a bright future ahead of them.

Disappointing performance is leading hedge fund investors to look for cheaper alternatives. Hedge fund indexation has been suggested as a possible solution. Unfortunately, investable hedge fund indices are nothing more than funds of funds in disguise, with performance similar or even worse than real funds of funds. The core problem of hedge fund indexation is that as long as one still invests in hedge funds, the cost factor that indexation is meant to eliminate will still be there. In Chapter 5 it has been shown that it is possible to generate returns with statistical properties that are very similar to hedge fund indices, but without actually investing in hedge funds. The proposed strategies only trade liquid futures contracts and therefore not only offer investors an accurate replica, but at the same time solve many other problems typically surrounding hedge fund investments, such as illiquidity, lack of transparency, limited capacity, etc.

The effect of the transaction costs, the choice of the reference portfolio and the reserve asset was examined in Sections 4.2 and 5.2. There was almost no influence of these factors on the risk profile of the synthetic funds.

Although we use copulas, a non-parametric implementation of the replication technique approach is also possible. Instead of copulas, the two joint distributions could be estimated using a kernel approach for example, subsequently deriving the exotic option's payoff function using a numerical procedure instead of the parametric distributions. This is a possible way for future developments.

References

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

Akaike, H., 1973. Information Theory and an Extension of the Maximum Likelihood Principle. In B. Petrov and F. Csaki, ed. *Second International Symposium on Information Theory.* Academiae Kiado, Budapest, pp. 267-281.

Agarwal, V. and N. Naik, 2000a. Generalised Style Analysis of Hedge Funds, *Journal of Asset Management,* 1, pp. 93-109.

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

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

Amenc, N., L. Martellini, and M. Vaissi, 2003. Benefits and Risks of Alternative Investment Strategies, *Journal of Asset Management,* 4, pp. 96-118.

Amin, G. and H. Kat, 2003a. Welcome to the Dark Side: Hedge Fund Attrition and Survivorship Bias 1994-2001, *Journal of Alternative Investments,* pp. 57-73.

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

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

Arrow, K., 1964. The Role of Securities in the Optimal Allocation of Risk-Bearing, *Review of Economic Studies*, 31 (2), pp. 91-96 (originally published in French in 1953).

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

Baquero, G., J. Horst and M. Verbeek, 2004. Survival, Look-Ahead Bias and the Performance of Hedge Funds, forthcoming in the *Journal of Financial and Quantitative Analysis.*

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

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

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

Brown, S., W. Goetzmann, R. Ibbotson and S. Ross, 1992. Survivorship Bias in Performance Studies, *Review of Financial Studies,* 5, pp. 553-580.

Brown, S., W. Goetzmann, and J. Park, 1997. Conditions for Survival: Changing Risk and the Performance of Hedge Fund Managers and CTAs, Working Paper F-59, Yale School of Management.

Brown, S., W. Goetzmann, and R. Ibbotson, 1999. Offshore Hedge Funds: Survival and Performance 1989-1995, *Journal of Business,* 72, pp. 91-118.

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

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

Capocci, D. and G. Hubner, 2004. Analysis of Hedge Fund Performance, *Journal of Empirical Finance*, 11, pp. 55-89.

Carhart, M., 1997. On persistence in mutual fund performance, *Journal of Finance*, 52, pp. 57-82.

Carpenter, J. and A. Lynch, 1999. Survivorship Bias and Attrition Effects in Measures of Performance Persistence, *Journal of Financial Economics*, 54, pp. 337-374.

Carr, P. and L. Wu, 2006. A Tale of Two Indices, *Journal of Derivatives*, pp. 13-29.

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

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

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

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

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

Dybvig, P. and S. Ross, 2003. Arbitrage, State Prices and Portfolio Theory. *Handbook of the Economics of Finance*, 1 (2), pp. 605-637.

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

Edhec, 2004. Edhec Alternative Indexes, Working Paper.

Edwards, F. and M. Caglayan, 2001. Hedge Fund Performance and Manager Skill, *Journal of Futures Markets*, 21, pp. 1003-1028.

Eling, M., 2006. Autocorrelation, Bias, and Fat Tails – Are Hedge Funds Really Attractive Investments, *Derivatives Use, Trading & Regulation*, 12 (1), pp. 28-47.

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

Fung, W., and D. Hsieh, 1997. Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds, *Review of Financial Studies*, 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*, 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*, 14, pp. 313-341.

Fung, W., D. Hsieh, N. Naik and T. Ramadorai, 2007. Hedge Funds: Performance, Risk and Capital Formation, BNP Paribas Hedge Fund Centre Working Paper Series.

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

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

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

Gregoriou, G., and J. Gueyie, 2003. Risk-Adjusted Performance of Funds of Hedge Funds Using a Modified Sharpe Ratio, *Journal of Wealth Management,* 6 (3), pp. 77-83.

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

Hasanhodzic, J. and A. Lo, 2006. *Can Hedge Fund Returns Be Replicated? The Linear Case*, Working Paper MIT Laboratory for Financial Engineering.

He, H., 1990. Convergence from Discrete to Continuous-Time Contingent Claims Price, *Review of Financial Studies*, 3, pp. 523-546.

Hendricks, D., J. Patel and R. Zeckhauser, 1997. The J-Shape of Performance Persistence Given Survivorship Bias, *Review of Economics and Statistics*, 79, pp. 161-170.

Hill, J., B. Muller and V. Balasubramanian, 2004. The 'Secret Sauce' of Hedge Fund Investing – Trading Risk Dynamically, *Goldman Sachs Equity Derivative Strategy*, November 2.

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

Horst, J., T. Nijman and M. Verbeek, 2001. Eliminating Look-Ahead Bias in Evaluating Persistence in Mutual Fund Performance, *Journal of Empirical Finance*, 8, pp. 345-373.

Ibbotson, R., and P. Chen, 2006. The A, B, Cs of Hedge Funds: Alphas, Betas and Costs, Working paper 06-10, Yale ICF.

Jaeckel, P., 2002. *Monte Carlo Methods in Finance*, John Wiley & Sons.

Jaeger, L. and C. Wagner, 2005. *Factor Modelling and Benchmarking of Hedge Funds: Can Passive Investments in Hedge Fund Strategies Deliver?*, Journal of Alternative Investments, pp. 9-36.

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

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

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

Kat, H., 2003. The Dangers of Using Correlation to Measure Dependence, *Journal of Alternative Investments,* pp. 54-58.

Kat, H., 2006. How to Evaluate a New Diversifier with 10 Simple Questions, Alternative Investment Research Centre Working Paper 39, Cass Business School, 2006.

Kat, H. and J. Miffre, 2006. The Impact of Non-Normality Risks and Tactical Trading on Hedge Fund Alphas, Working Paper 26, Alternative Investment Research Centre, Cass Business School, City University London.

Kat, H and R. Oomen, 2006a. What Every Investor Should Know About Commodities Part I: Univariate Return Analysis, Working Paper 29, Alternative Investment Research Centre, Cass Business School, London.

Kat, H and R. Oomen, 2006b. What Every Investor Should Know About Commodities Part II: Multivariate Return Analysis, Working Paper 33, Alternative Investment Research Centre, Cass Business School, London.

Kat, H. and N. Tassabehji, 2006. So Now You Want to Invest In Volatility?, Alternative Investment Research Centre Working Paper, forthcoming, Cass Business School, City University London.

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

Kouwenberg, R., 2003. Do Hedge Funds Add Value to a Passive Portfolio, *Journal of Asset Management*, 3, pp 361-382.

Lhabitant, F., 2001. Assessing Market Risk for Hedge Funds and Hedge Fund Portfolios, *Journal of Risk Finance*, 2 (4).

Liang, B., 1999. On the Performance of Hedge Funds, *Financial Analysts Journal*, pp. 72-85.

Liang, B., 2000. Hedge Funds: The Living and the Dead, *Journal of Financial and Quantitative Analysis*, 35, pp. 309-326.

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

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

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

Patton, A., 2006a. Modelling Asymmetric Exchange Rate Dependence, *International Economic Review*, 47 (2), pp. 527-556.

Patton, A., 2006b. Estimation of Multivariate Models for Time Series of Possibly Different Lengths, *Journal of Applied Econometrics,* 21 (2), pp. 147-173.

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

Schneeweis, T., and R. Spurgin, 1996. Survivor Bias in Commodity Trading Advisor Performance, *Journal of Future Markets,* 16, 757-772.

Schneeweis, T., and R. Spurgin, 1998. Multi-Factor Analysis of Hedge Fund, Managed Futures, and Mutual Fund Return and Risk Characteristics, *Journal of Alternative Investments,* 1, 1-24.

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

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

Appendices

A.1 PDPM - Univariate Case

Suppose we lived in the standard model with one stock, one bond, perfect markets, a positive equity risk premium, and where the stock price follows a binomial tree with n steps, with time increments Δt. In one step, for each node, the stock price can move up by a factor $u = 1 + \mu\Delta t + \sigma\sqrt{\Delta t}$ or down by a factor $d = 1 + \mu\Delta t - \sigma\sqrt{\Delta t}$ with the same probability. After n steps, the initial price S_0 will have evolved into one of $n+1$ possible values $S_0\, u^n$, $S_0\, u^{n-1}d^1$, $S_0\, u^{n-2}d^2$, ..., $S_0\, d^n$, which we label as states 1, 2, ..., $n+1$ respectively. The bond returns $r\Delta t$ over each period, with r denoting the riskless rate.

What is the state price for each of the above $n+1$ states? Suppose we are at time $t=0$ and want to replicate the payoff of some particular investment. The stock price is S_0 and the bond value is B. Suppose the value of the investment one step ahead is either v_1 if the stock price goes up, or v_2 if it goes down. In order to replicate the investment, we need an investment of v_S shares and v_B bonds:

$$\begin{cases} S_0(1+\mu\Delta t+\sigma\sqrt{\Delta t})v_S + B(1+r\Delta t)v_B = v_1 \\ S_0(1+\mu\Delta t-\sigma\sqrt{\Delta t})v_S + B(1+r\Delta t)v_B = v_2 \end{cases}$$

Solving and reordering this result, we have the following one-period pricing relationship:

$$S_0 v_S + B v_B = \frac{1}{2(1+r\Delta t)}\left[1-\frac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right]v_1 + \frac{1}{2(1+r\Delta t)}\left[1+\frac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right]v_2 = p_1 v_1 + p_2 v_2,$$

where p_1 and p_2 are the state prices. Dividing the state prices by the probability of each state, which is ½, we obtain the following state-price density:

$$\rho_1 = \frac{1}{(1+r\Delta t)}\left[1-\frac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right] \text{ and } \rho_2 = \frac{1}{(1+r\Delta t)}\left[1+\frac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right]$$

Note that, assuming a positive risk premium, $S_0(1+\mu\Delta t+\sigma\sqrt{t}) > S_0(1+\mu\Delta t-\sigma\sqrt{t})$, we have $\rho_1 < \rho_2$. If we repeat this for all nodes in the tree, we find that for the last step the state-price density is inversely related to the terminal values of the stock. In state *1*, the stock has the highest value S_0u^n , but the state-price density assumes the lowest value, while in state *n+1* the stock has the lowest value $S_0\, d^n$ but the state-price density function assumes the highest value.

Now we are set to use Theorem 1 of Dybvig (1988a). By this theorem, the cheapest payoff function allocates terminal wealth as a non-increasing function of the state-price density. Combining this with the above, the cheapest payoff function should therefore *allocate terminal wealth as a non-decreasing function of the value of the stock.*

A.2 PDPM - Bivariate Case

Now assume that instead of one, we have two risky assets we can trade, which we will refer to as "the investor's portfolio" and "the reserve asset". The prices of both are denoted as S_P and S_R respectively. We denote the mean and standard deviation of the terminal wealth provided by the investor's portfolio as μ_P and σ_P. Likewise, we denote the mean and standard deviation of the terminal wealth provided by the reserve asset as μ_R and σ_R . ρ denotes the correlation coefficient between the two assets.

From assumption 1 (Section 2.1) we know that all investors are concerned about is the conditional distribution of the terminal wealth given S_P. As a direct consequence of this assumption, we need therefore to study the question how to allocate terminal wealth between states for each value $S_P = x$.

Supposing normality for the investor's portfolio and the reserve asset, we know that the distribution of R given S_P = x is a normal distribution with mean $\mu_R + \rho\frac{\sigma_R}{\sigma_P}(x-\mu_P)$ and standard deviation $\sigma_R\sqrt{1-\rho^2}$. We can use a binomial model to approximate this distribution. Doing so, we can perform the same analysis as in the univariate case to obtain the state-price density for the first step of the tree. Suppose that the initial price of the reserve asset is $S_{R,0}$. Then the prices for the two nodes are

$$S_{R,1} = S_{R,0}\left\{1+\left(\mu_R + \rho\frac{\sigma_R}{\sigma_P}(x-\mu_P)\right)\Delta t + \sigma_R\sqrt{1-\rho^2}\sqrt{t}\right\}$$

and

$$S_{R,2} = S_{R,0}\left\{1+\left(\mu_R + \rho\frac{\sigma_R}{\sigma_P}(x-\mu_P)\right)\Delta t - \sigma_R\sqrt{1-\rho^2}\sqrt{t}\right\}.$$

The state-price density is therefore given by:

$$\rho_1 = \frac{1}{(1+r\Delta t)}\left[1 - \frac{(\mu_R + \rho\frac{\sigma_R}{\sigma_P}(x-\mu_P) - r)\Delta t}{\sigma_R\sqrt{1-\rho^2}\sqrt{\Delta t}}\right]$$

and

$$\rho_2 = \frac{1}{(1+r\Delta t)}\left[1 + \frac{(\mu_R + \rho\frac{\sigma_R}{\sigma_P}(x-\mu_P) - r)\Delta t}{\sigma_R\sqrt{1-\rho^2}\sqrt{\Delta t}}\right]$$

Clearly, much depends on the value of ρ. Suppose $\rho > 0$. In that case $\rho_1 < \rho_2$ if $x > \mu_P + (r-\mu_R)\frac{\sigma_P}{\rho\sigma_R} = x_{\min}$, and $\rho_1 > \rho_2$ otherwise. In other words, the allocation rule for the cheapest payoff function will depend on the value of the investor's portfolio. If $x > x_{min}$, the rule is to allocate terminal wealth as a non-decreasing function of the value of the reserve asset, just as in the univariate case. If $x < x_{min}$, however, the rule is to allocate terminal wealth as a non-increasing function of the value of the reserve asset.

When $\rho < 0$, we see a similar phenomenon. In that case $\rho_1 < \rho_2$ if $x < \mu_P + (r-\mu_R)\frac{\sigma_P}{\rho\sigma_R} = x_{\max}$, and $\rho_1 > \rho_2$ otherwise. This means that if $x < x_{max}$, the rule for the cheapest payoff function is to allocate terminal wealth as a non-decreasing function of the value of the reserve asset. When $x > x_{max}$, however, the cheapest payoff function allocates terminal wealth as a non-increasing function of the value of the reserve asset.

A.3 Copulas and marginal distributions

In this appendix we list the six copula families used and also the moments of the marginal distributions.

For all copula definitions we assume $u \in [0,1]$, $v \in [0,1]$.

The normal copula is given by:

$$C_{Ga}(u,v\,;\rho) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{1}{2(1-\rho^2)}(s^2 - 2\rho st + t^2)\right) dtds\,,$$

where $\Phi^{-1}(.)$ is the standard normal inverse distribution and $\rho \in (-1,1)$ is the linear correlation coefficient

The Student-t copula is given by:

$$C_t(u,v\,;\rho,\vartheta) = \int_{-\infty}^{t_\vartheta^{-1}(u)} \int_{-\infty}^{t_\vartheta^{-1}(v)} \frac{1}{2\pi\sqrt{1-\rho^2}} \left(1+\frac{s^2 - 2\rho st + t^2}{\vartheta(1-\rho^2)}\right)^{-\left(\frac{\vartheta+2}{2}\right)} dtds$$

where t_ϑ^{-1} is the Student-t inverse distribution with ϑ degrees of freedom.

The Gumbel copula is given by:

$$C_G(u,v\,;a) = \exp\left(-\left[(-\ln(u))^a + (-\ln(v))^a\right]^{1/a}\right),\quad a > 1.$$

The Cook-Johnson copula is given by:

$$C_{CJ}(u,v\,;\alpha) = \left(u^{-\alpha} + v^{-\alpha} - 1\right)^{-1/\alpha},\quad \alpha > 0$$

The Frank copula is defined by:

$$C_F(u,v;\theta) = -\frac{1}{\theta}\log\left(1+\frac{(e^{-\theta u}-1)(e^{-\theta v}-1)}{e^{-\theta}-1}\right),\quad \theta \neq 0.$$

The SJC copula is based on the Joe-Clayton copula, which is given by:

$C_{JC}(u,v;\tau_{L,}\tau_U)=1-\left(1-\left\{[1-(1-u)^{\kappa}]^{-\gamma}+[1-(1-v)^{\kappa}]^{-\gamma}-1\right\}^{-1/\gamma}\right)^{1/\kappa}$, where

$\kappa=1/\log_2(2-\tau_U)$, $\gamma=-1/\log_2(\tau_L)$, and $\tau_U\in(0,1)$ and $\tau_L\in(0,1)$.

The SJC copula is then given by:

$$C_{SJC}(u,v;\tau_{L,}\tau_U)=\tfrac{1}{2}\left(C_{JC}(u,v;\tau_{L,}\tau_U)+C_{JC}(1-u,1-v;\tau_{U,}\tau_L)+u+v-1\right)$$

The probability density function of the normal distribution is given by:

$f_x(x)=\frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right)$, $x\in\Re$. The normal model is symmetrical around the mean and it has only two parameters, the mean μ and the standard deviation σ.

The probability density function of the Student-t distribution is given by

$f_x(x)=\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\sqrt{c\pi}}\left(1+\frac{(x-\mu)^2}{c^2}\right)^{-\left(\frac{\nu+1}{2}\right)}$, $x\in\Re$. The mean of the distribution is μ and its standard deviation is given by $c/\sqrt{\nu-2}$. The parameter ν is known as *degrees of freedom*. It controls the excess kurtosis of the distribution, which is given for $\nu>4$ by $k=\frac{E[(X-\mu)^4]}{E[(X-\mu)^2]^2}-3=\frac{6}{\nu-4}$. This distribution is also symmetrical.

The Johnson-SU family of distributions is defined by the following transformation: $X=\xi+\lambda\sinh\left(\frac{Z-\gamma}{\delta}\right)$, $X>\xi$, where Z is a random variable with standard normal distribution.

The moments of the distribution are given by:

$$E(X)=\xi+\frac{\lambda e^{\frac{1}{2\delta^2}}}{2}\left(e^{-\frac{\gamma}{\delta}}-e^{\frac{\gamma}{\delta}}\right)$$

$$V(X)=E(X-E(X))^2=\frac{\lambda^2}{4}\left\{e^{\frac{2}{\delta^2}}(e^{-\frac{2\gamma}{\delta}}+e^{\frac{2\gamma}{\delta}})-e^{\frac{1}{\delta^2}}(e^{-\frac{\gamma}{\delta}}-e^{\frac{\gamma}{\delta}})^2-2\right\}$$

$$E(X-E(X))^3=\frac{\lambda^3}{8}\{\,e^{\frac{9}{2\delta^2}}(e^{-\frac{3\gamma}{\delta}}-e^{\frac{3\gamma}{\delta}})-3e^{\frac{5}{2\delta^2}}(e^{-\frac{3\gamma}{\delta}}-e^{-\frac{\gamma}{\delta}}+e^{\frac{\gamma}{\delta}}-e^{\frac{3\gamma}{\delta}})+$$
$$2e^{\frac{3}{2\delta^2}}(e^{-\frac{3\gamma}{\delta}}-3e^{-\frac{\gamma}{\delta}}+3e^{\frac{\gamma}{\delta}}-e^{\frac{3\gamma}{\delta}})+3e^{\frac{1}{2\delta^2}}(e^{-\frac{\gamma}{\delta}}-e^{\frac{\gamma}{\delta}})\,\}$$

$$Skewness(X)=\frac{E(X-E(X))^3}{V(X)^{3/2}}$$

$$E(X-E(X))^4=\frac{\lambda^4}{16}\{\,e^{\frac{8}{\delta^2}}(e^{-\frac{4\gamma}{\delta}}+e^{\frac{4\gamma}{\delta}})-4e^{\frac{5}{\delta^2}}(e^{-\frac{4\gamma}{\delta}}-e^{-\frac{2\gamma}{\delta}}-e^{\frac{2\gamma}{\delta}}+e^{\frac{4\gamma}{\delta}})+$$
$$6e^{\frac{3}{\delta^2}}(e^{-\frac{4\gamma}{\delta}}-2e^{-\frac{2\gamma}{\delta}}+2-2e^{\frac{2\gamma}{\delta}}+e^{\frac{4\gamma}{\delta}})+e^{\frac{2}{\delta^2}}(-3e^{-\frac{4\gamma}{\delta}}+8e^{-\frac{2\gamma}{\delta}}-24+8e^{\frac{2\gamma}{\delta}}-3e^{\frac{4\gamma}{\delta}})$$
$$-3e^{\frac{1}{\delta^2}}(e^{-\frac{\gamma}{\delta}}-e^{\frac{\gamma}{\delta}})^2-12(e^{-\frac{2\gamma}{\delta}}-e^{\frac{2\gamma}{\delta}})+30\,\}$$

$$Excess\ Kurtosis(X)=\frac{E(X-E(X))^4}{V(X)^2}-3$$

It seems more complicated that it really is. If we specify a proper set of skewness and excess kurtosis, we can calculate the parameters γ and δ necessary to obtain these moments, since they do not depend on λ or ξ. After that, we can calculate λ specifying the variance of the distribution and finally we can obtain ξ specifying its mean.

A.4 Influence of the Reserve Asset distribution parameters

In this appendix we will study the influence of choice of the reserve asset in the price of the payoff function.

First we are going to set some notation:

S_P: investor's portfolio end-of-month payoff

S_I: fund end-of-month payoff

S_R: reserve asset end-of-month payoff

Log-returns: $X_P = \ln\left(\frac{S_P}{100}\right)$, $X_I = \ln\left(\frac{S_I}{100}\right)$ and $X_R = \ln\left(\frac{S_R}{100}\right)$

Moments: $E(X_P) = \mu_P$, $Var(X_P) = \sigma_P^2$

$E(X_I) = \mu_I$, $Var(X_I) = \sigma_I^2$

$E(X_R) = \mu_R$, $Var(X_R) = \sigma_R^2$

$Corr(X_P, X_R) = \rho_{PR}$, $Corr(X_P, X_I) = \rho_{PI}$

r: risk-free interest rate.

g^* is the cheapest payoff function such that:

$P(S_P \leq x, g^*(S_P, S_R) \leq y) = P(S_P \leq x, S_I \leq y), \forall x, y \in \Re^2$.

The KP measure is the price of the function g^*:

$\psi = \frac{1}{r} E_Q[g^*(S_P, S_R)]$,

where Q denotes de risk-neutral expectation, hence supposing that

$$\begin{pmatrix} X_r \\ X_R \end{pmatrix} \sim N\left(\begin{pmatrix} r - \frac{\sigma_P^2}{2} \\ r - \frac{\sigma_R^2}{2} \end{pmatrix}, \begin{pmatrix} \sigma_P^2 & \sigma_{PR} \\ \sigma_{PR} & \sigma_R^2 \end{pmatrix}\right).$$

The function g^* can be rewritten in terms of log-returns,

$g(x,y) = \ln\left(\frac{g^*(\exp(100x, 100y))}{100}\right)$ which

implies $g^*(x,y) = 100\exp\left(g\left(\ln\left(\frac{x}{100}\right), \ln\left(\frac{y}{100}\right)\right)\right)$.

By (9) in subsection 2.3.3, we know that the cheapest payoff function g is defined by

$$g(x,y) = F^{-1}_{X_I|X_P}(F_{X_R|X_P}(y \mid x) \mid x).$$

The price of this payoff function is given by

$\psi = \frac{100}{r} E_Q[\exp(g(X_P, X_R))] = \frac{100}{r} E_Q[h(X_P, X_R)]$, where $h(x,y) = \exp(g(x,y))$.

We know that a continuous function $m(x,y)$ can be written using Taylor approximation around a point x_0, y_0:

$$m(x,y) = m(x_0, y_0) + \sum_{i=1}^{\infty}\sum_{j=1}^{\infty} m^{(i)(j)}(x_0, y_0)\frac{(x-x_0)^i (y-y_0)^j}{i!\,j!}.$$

Using a first-order approximation for the function $h(X_P, X_R)$ around the risk-neutral means $(E_Q[X_P], E_Q[X_R])$, we have:

$$h(X_P, X_R) \cong h(E_Q[X_P], E_Q[X_R]) + \frac{\partial h}{\partial x}(X_P - E_Q[X_P]) + \frac{\partial h}{\partial y}(X_R - E_Q[X_R])$$

Taking expectations, we have $E_Q[h(X_P)] \cong h(E_Q[X_P], E_Q[X_R])$.

Substituting $E_Q[X_P] = r - \frac{\sigma_P^2}{2}$, $E_Q[X_R] = r - \frac{\sigma_R^2}{2}$ and

$h(x,y) = \exp\left[F^{-1}_{X_I|X_P}(F_{X_R|X_P}(y \mid x) \mid x)\right]$, we have:

$$E_Q[h(X_P, X_R)] \cong \exp(F^{-1}_{X_I|X_P}(F_{X_R|X_P}(r - \frac{\sigma_R^2}{2} \mid r - \frac{\sigma_P^2}{2}) \mid r - \frac{\sigma_P^2}{2}))$$

The price of the payoff function (KP Measure) is then approximated by:

$$\psi \cong \frac{100}{r}\exp(F^{-1}_{X_I|X_P}(F_{X_R|X_P}(r - \frac{\sigma_R^2}{2} \mid r - \frac{\sigma_P^2}{2}) \mid r - \frac{\sigma_P^2}{2})) \qquad (16)$$

ψ is a function of $\mu_P, \mu_I, \mu_R, \sigma_P, \sigma_I, \sigma_R, \rho_{PI}, \rho_{PR}, r$

We want to study the selection of the reserve asset, so we will concentrate on the parameters $\mu_R, \sigma_R,\ \rho_{PR}$ and assume the others as given. With this assumption, the function $F^{-1}_{X_I|X_P}$ will not change, so we only need to study the effect of $\mu_R, \sigma_R,\ \rho_{PR}$ on $F_{X_R|X_P}(r-\frac{\sigma_R^2}{2} \mid r-\frac{\sigma_P^2}{2}) \mid r-\frac{\sigma_P^2}{2})$.

Supposing normality of X_P and X_R , from probability theory we have:

$$X_R \mid X_P = r-\frac{\sigma_P^2}{2} \sim N\left(\mu_R + \rho_{PR}\frac{\sigma_R}{\sigma_P}\left(r-\mu_P-\frac{\sigma_P^2}{2}\right), \sigma_R^2(1-\rho_{PR}^2)\right)$$

So:

$$F_{X_R|X_P}(r-\frac{\sigma_R^2}{2} \mid r-\frac{\sigma_P^2}{2}) \mid r-\frac{\sigma_P^2}{2}) =$$

$$\Phi\left(\frac{r-\frac{\sigma_R^2}{2}-\mu_R-\rho_{PR}\frac{\sigma_R}{\sigma_P}(r-\mu_P-\frac{\sigma_P^2}{2})}{\sigma_R\sqrt{1-\rho_{PR}^2}}\right) \qquad (17)$$

where Φ is the standard cumulative normal distribution. Since Φ is strictly increasing, we only need to study its argument.

a) parameter μ_R (mean of the reserve asset)

If everything else is constant, the quantity in (17) is decreasing on μ_R, so the price of the payoff function (16) is also decreasing on μ_R. This is not a surprise. If we have a reserve asset with a higher mean, then it will be cheaper to build the required payoff function.

b) parameter σ_R (volatility of the reserve asset)

Considering everything else constant, the argument of (17) can be written as a function of σ_R :

$$\kappa(\sigma_R) = \left(\frac{r-\mu_R}{\sqrt{1-\rho_{PR}^2}}\right)\frac{1}{\sigma_R} - \left(\frac{1}{2\sqrt{1-\rho_{PR}^2}}\right)\sigma_R + \text{constant}$$

The derivative of this function is given by:

$$\kappa'(\sigma_R) = \frac{1}{\sqrt{1-\rho_{PR}^2}}\left(\frac{\mu_R - r}{\sigma_R^2} - \frac{1}{2}\right),$$ which is greater than zero if $\sigma_R^2 < 2(\mu_R - r)$

The conclusion is:

If $\sigma_R^2 < 2(\mu_R - r)$, then $\kappa(\sigma_R)$, and consequently the price of the payoff function ψ is increasing on σ_R.

If $\sigma_R^2 > 2(\mu_R - r)$, then $\kappa(\sigma_R)$, and consequently the price of the payoff function ψ is decreasing on σ_R.

Figure A.1 is a contour plot which shows the relationship between μ_R, σ_R and ψ. On the upper half of the graph, $\sigma_R^2 < 2(\mu_R - r)$ holds, so the price ψ is increasing on σ_R. On the bottom half of the graph, $\sigma_R^2 > 2(\mu_R - r)$ holds, so the price ψ is decreasing on σ_R.

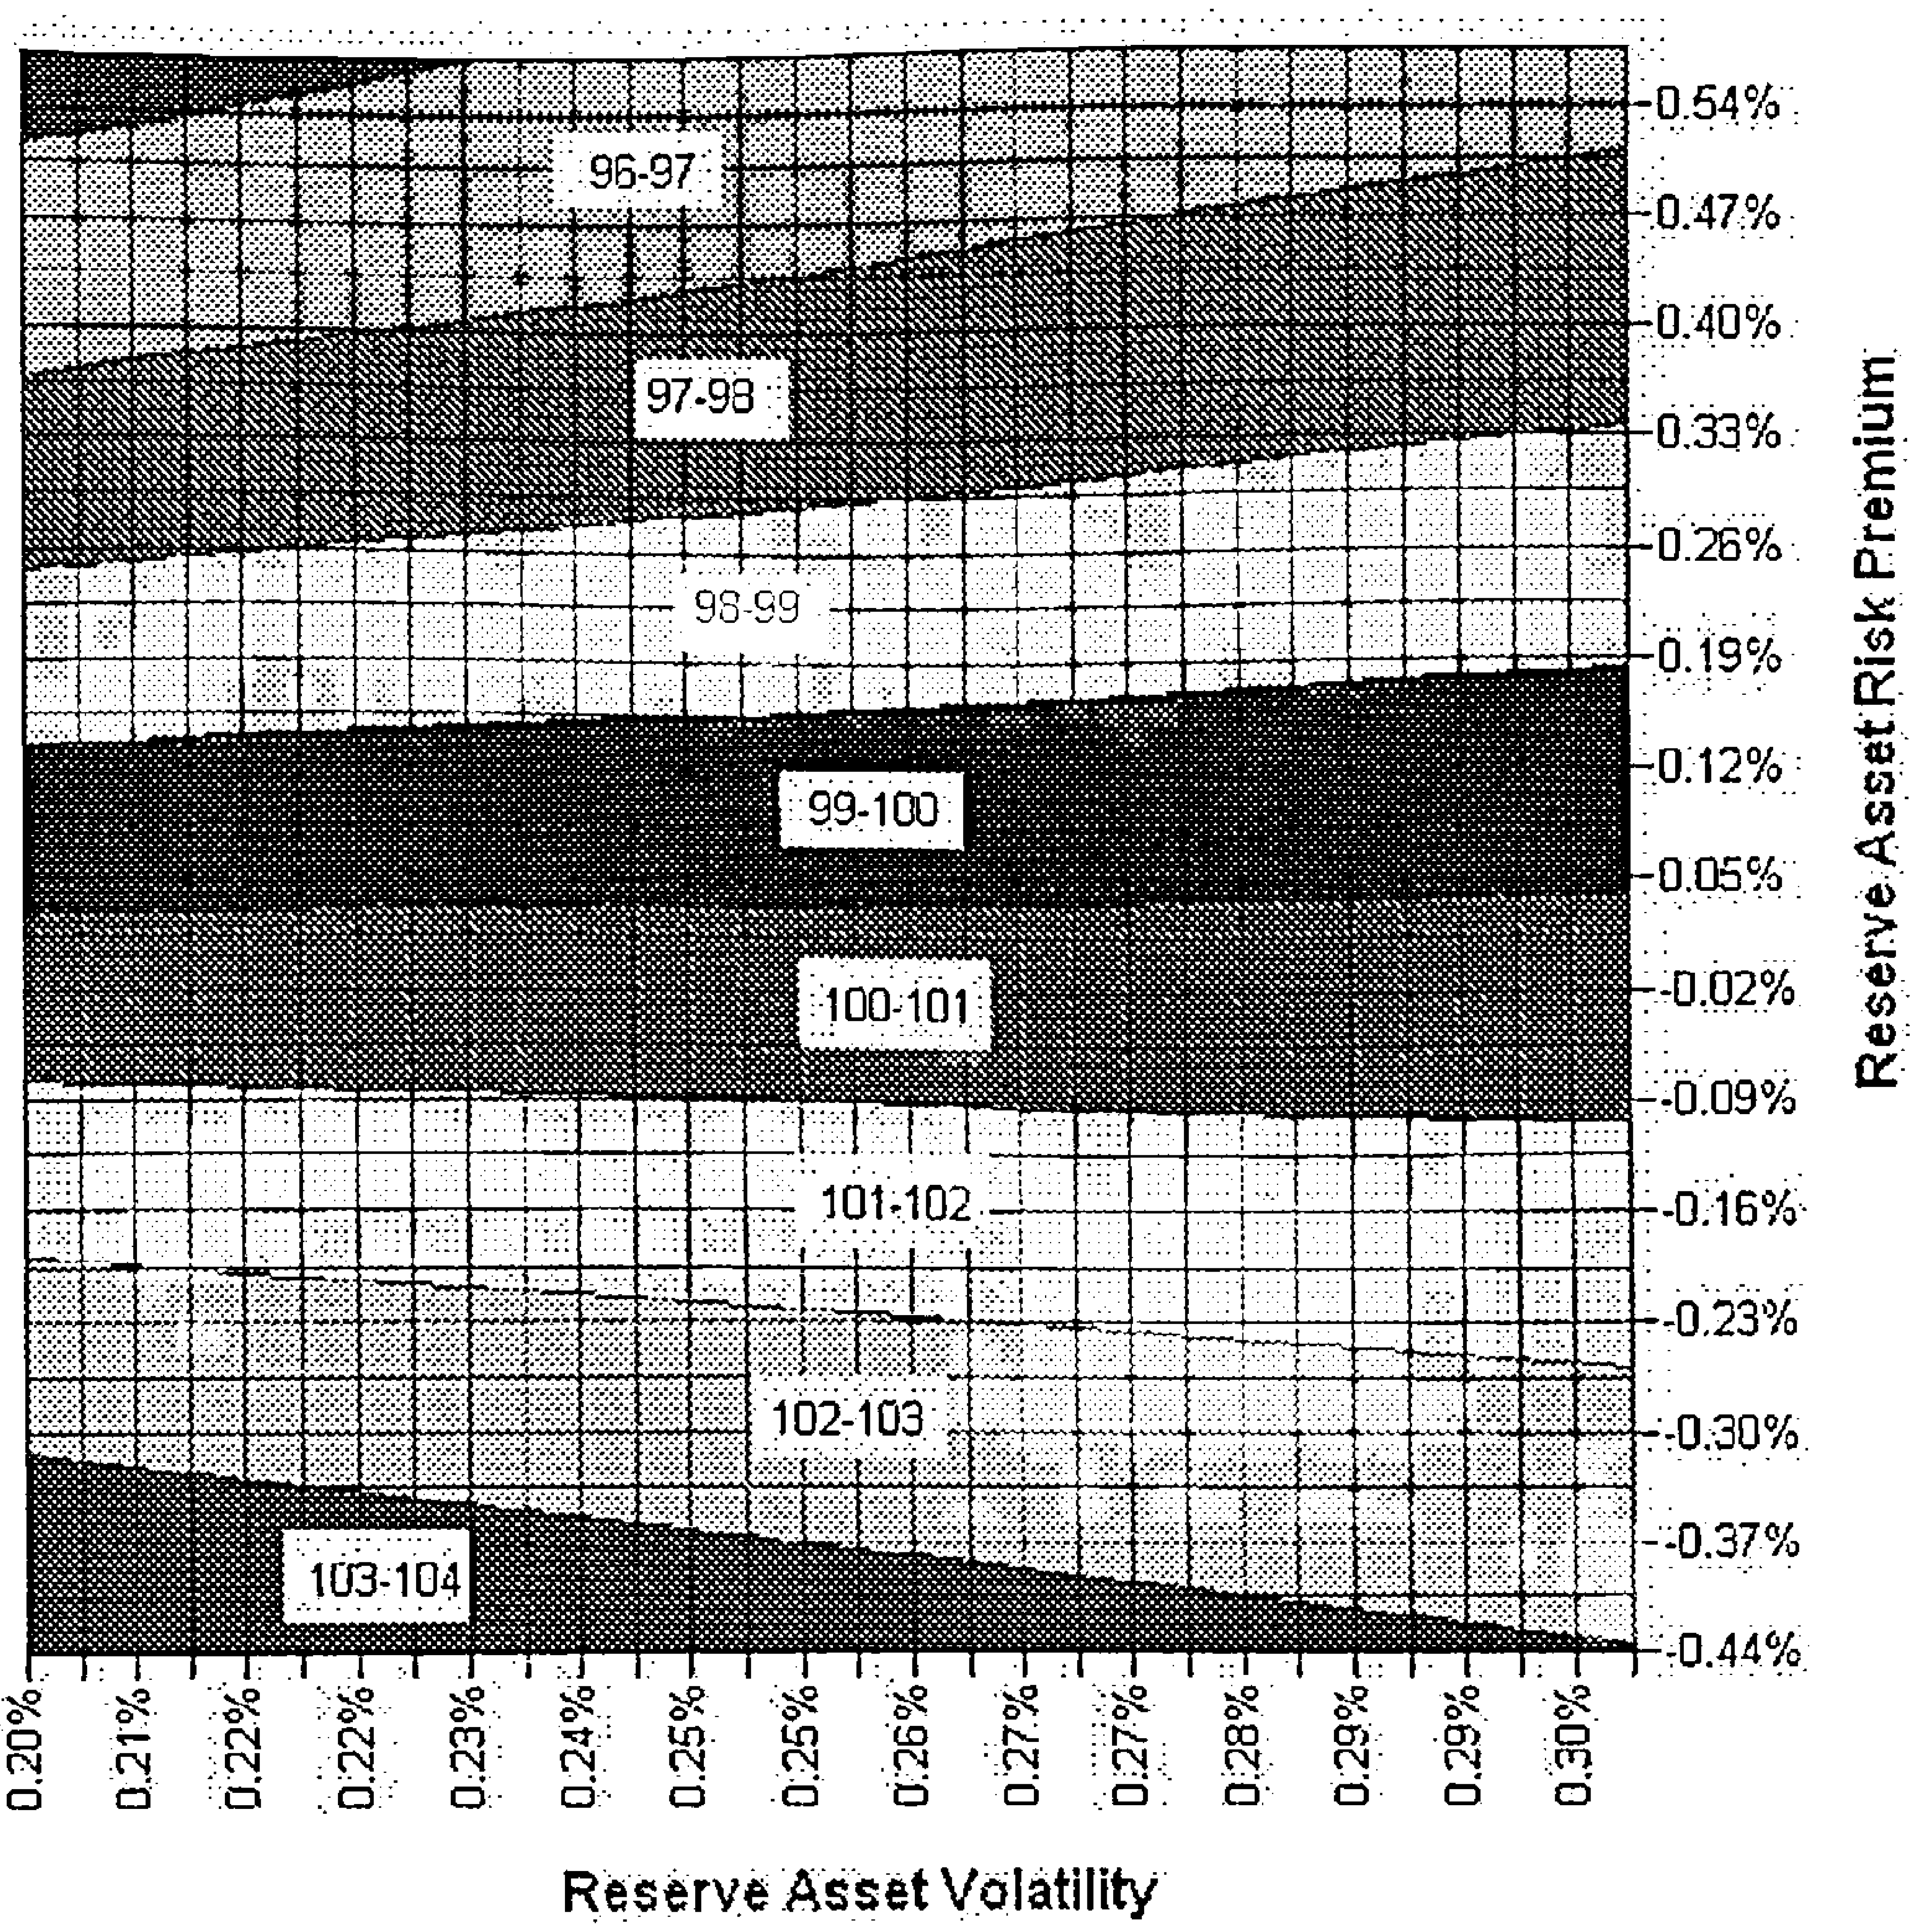

Figure A.1. Contour plot of the KP measure as function of the reserve asset risk premium and the reserve asset volatility.

If the risk premium is large enough (greater than $\sigma_R^2/2$), then we want the highest possible risk premium and the smallest possible volatility for the reserve asset, as it was also expected.

c) parameter ρ_{PR} (correlation between the reserve asset and the investor's portfolio)

Considering everything else constant, the argument of (17) can also be written as a function of ρ_{PR} :

$$\lambda(\rho_{PR})=\left(\frac{r-\mu_R}{\sigma_R}-\frac{\sigma_R}{2}\right)\frac{1}{\sqrt{1-\rho_{PR}^2}}-\frac{\rho_{PR}^2}{\sqrt{1-\rho_{PR}^2}}\left(\frac{r-\mu_P}{\sigma_P}-\frac{\sigma_P}{2}\right)$$

The derivative of this function is given by:

$$\lambda'(\rho_{PR})=\left(\frac{r-\mu_R}{\sigma_R}-\frac{\sigma_R}{2}\right)\frac{\rho_{PR}}{\sqrt[3]{1-\rho_{PR}^2}}-\left(\frac{r-\mu_P}{\sigma_P}-\frac{\sigma_P}{2}\right)\left(\frac{1}{\sqrt[3]{1-\rho_{PR}^2}}\right)$$

Which is positive if

$$\rho_{PR}<\frac{\left(\frac{r-\mu_P}{\sigma_P}-\frac{\sigma_P}{2}\right)}{\left(\frac{r-\mu_R}{\sigma_R}-\frac{\sigma_R}{2}\right)}=\frac{\left(\frac{\mu_P-r}{\sigma_P}+\frac{\sigma_P}{2}\right)}{\left(\frac{\mu_R-r}{\sigma_R}+\frac{\sigma_R}{2}\right)}$$, supposing positive risk premiums.

The conclusion is:

If $\rho_{PR}<\frac{\left(\frac{\mu_P-r}{\sigma_P}+\frac{\sigma_P}{2}\right)}{\left(\frac{\mu_R-r}{\sigma_R}+\frac{\sigma_R}{2}\right)}$, then the function $\lambda(\rho_{PR})$ is increasing on ρ_{PR}, and

consequently the price of the payoff function is ψ is increasing on ρ_{PR}.

If $\rho_{PR}>\frac{\left(\frac{\mu_P-r}{\sigma_P}+\frac{\sigma_P}{2}\right)}{\left(\frac{\mu_R-r}{\sigma_R}+\frac{\sigma_R}{2}\right)}$, then the function $\lambda(\rho_{PR})$ is decreasing on ρ_{PR}, and

consequently the price of the payoff function is ψ is decreasing on ρ_{PR}.

If either the denominator or numerator is negative, than the relation is first decreasing and then increasing.

Figure A.2 is a contour plot which shows the relationship between μ_R, ρ_{PR} and ψ. The risk premium of the investor's portfolio is positive. In the upper half of the graph, the denominator is also positive, so we can see that the price ψ is increasing on ρ_{PR} up to certain point (around $\rho_{PR} = 0.25$), and then decreasing on ρ_{PR}. In the bottom half of the graph the denominator is negative, so the inverse relation is observed (first decreasing on ρ_{PR} and then increasing on it).

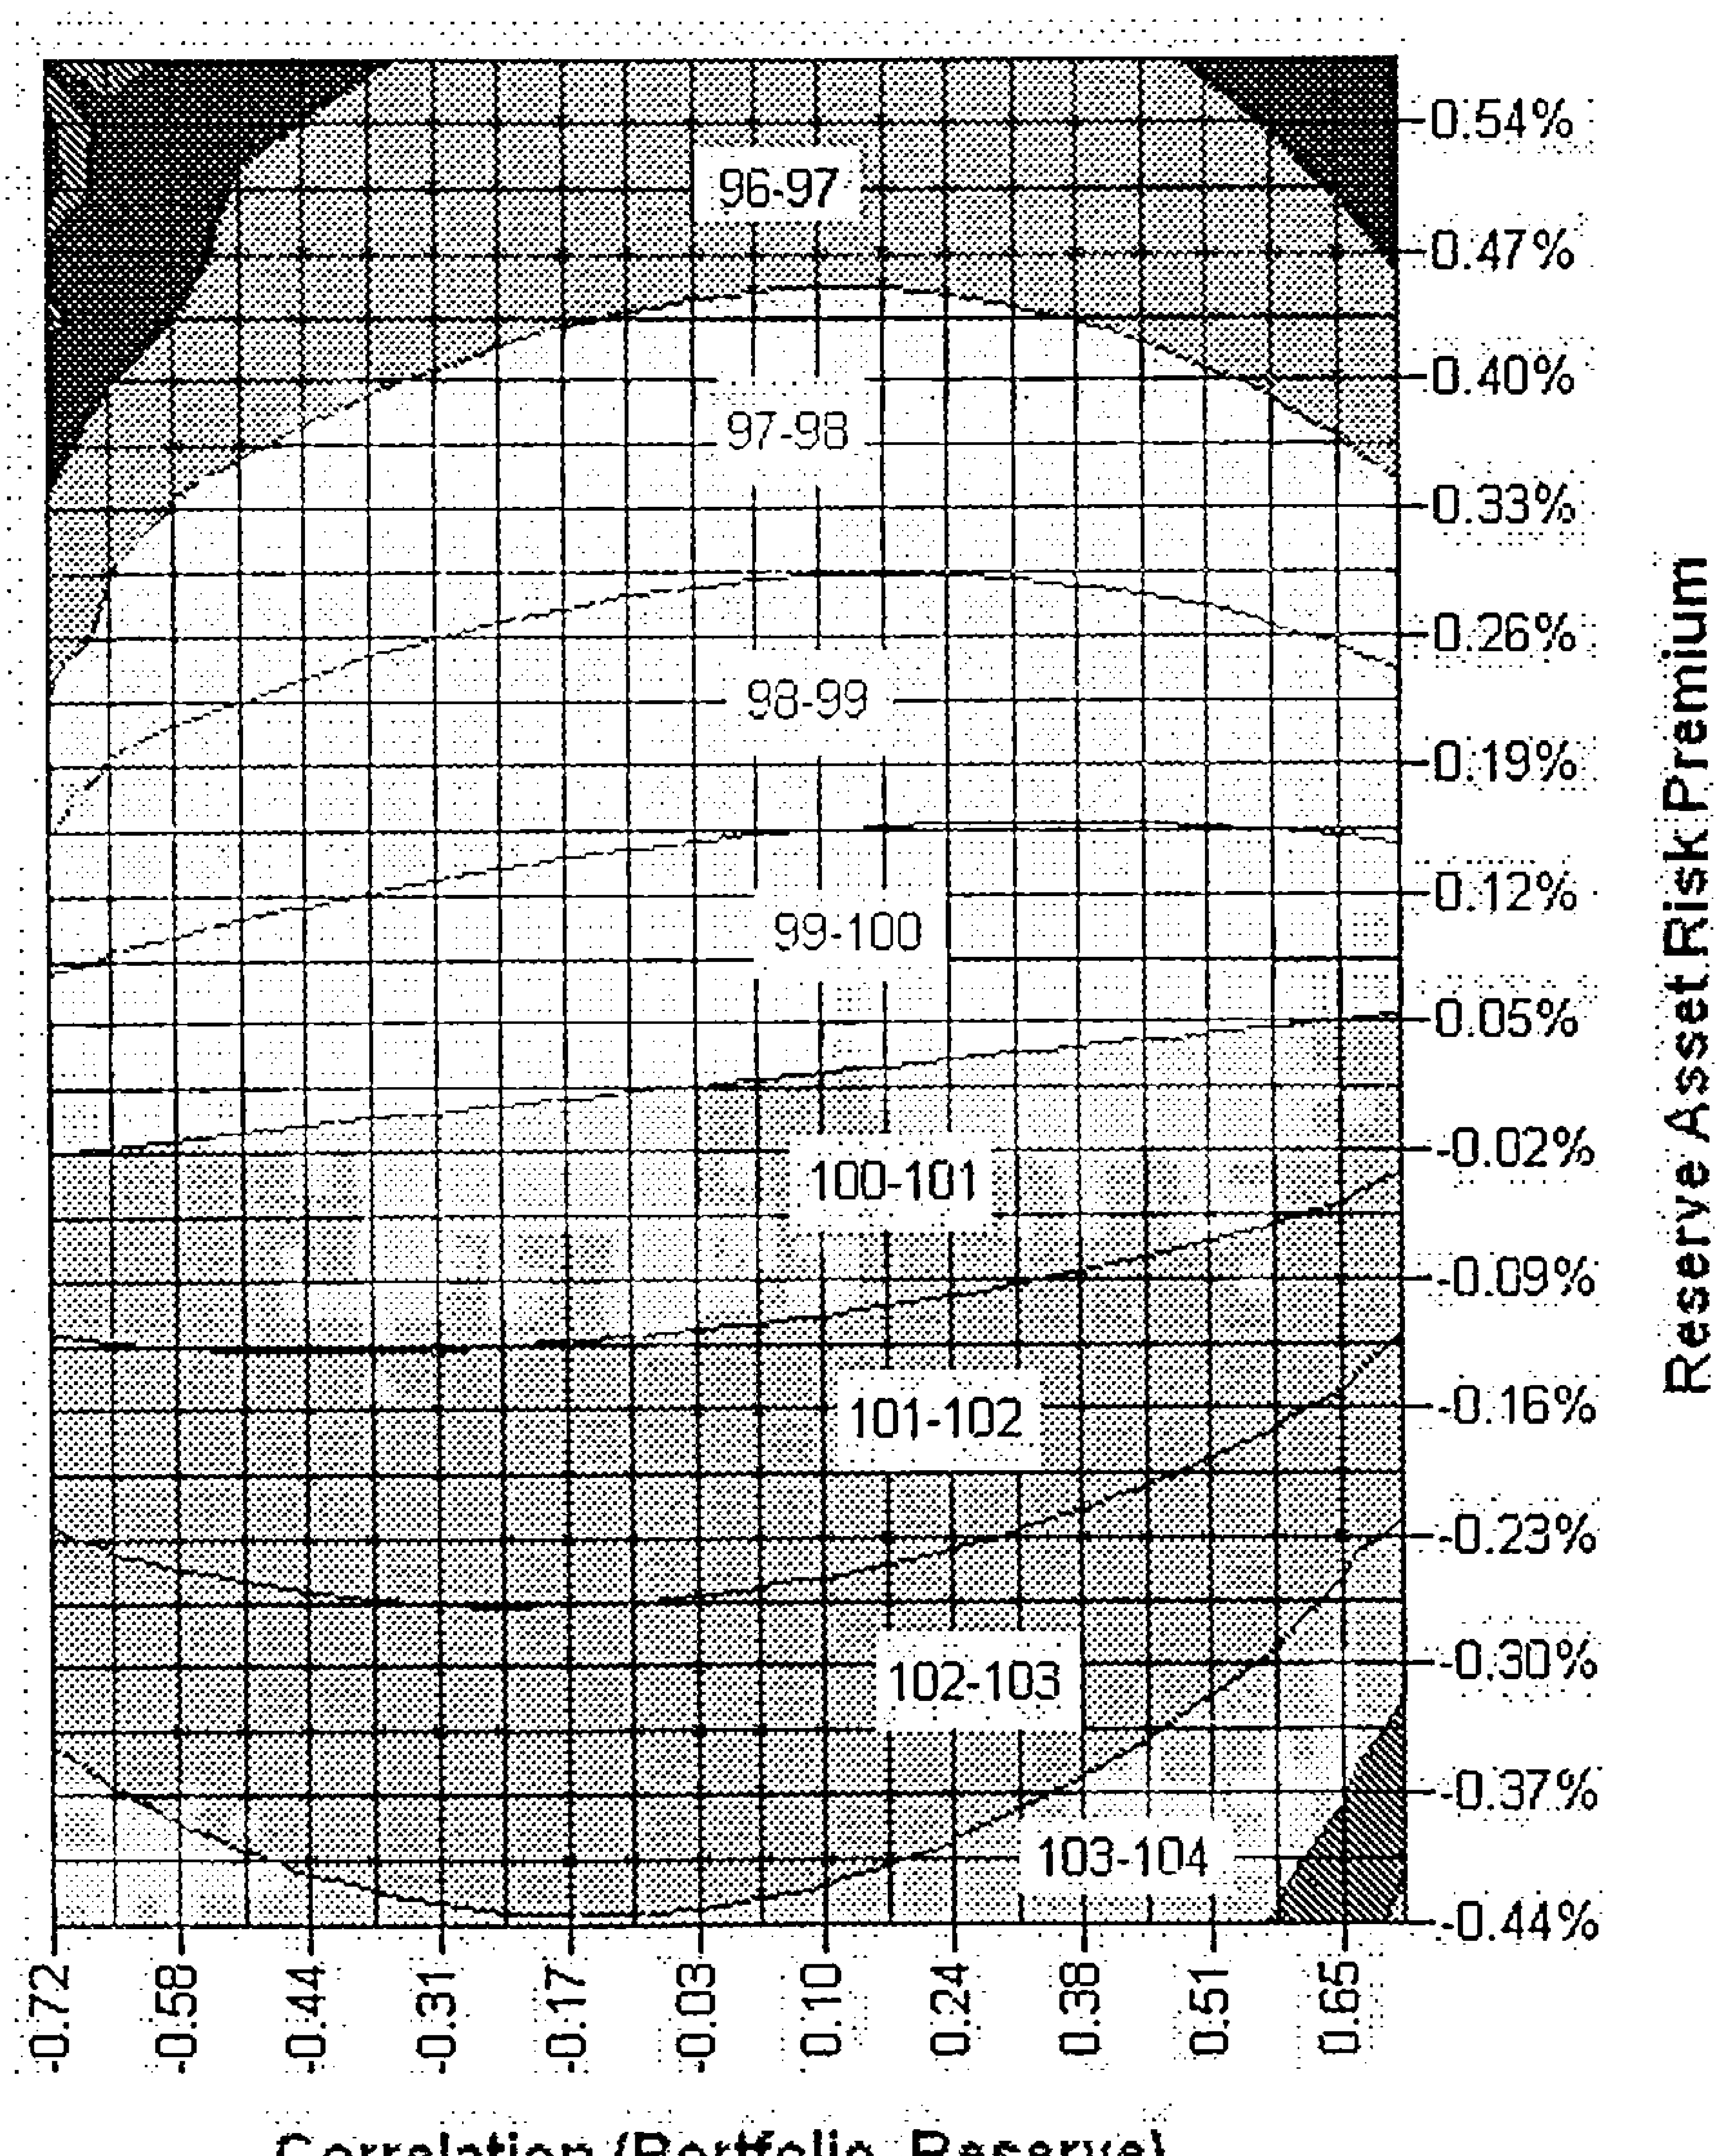

Figure A.2. Contour plot of the KP measure as function of the reserve asset risk premium and the correlation between the reserve asset and the reference portfolio.