

**ALTERNATIVE INVESTMENT RESEARCH
CENTRE WORKING PAPER SERIES**

Working Paper # 0036

**TELL ME WHAT YOU WANT, WHAT
YOU REALLY, REALLY WANT!**

An Exercise in Tailor-Made Synthetic Fund Creation

Harry M. Kat

Helder P. Palaro

**Alternative Investment Research Centre
Cass Business School, City University
106 Bunhill Row, London, EC2Y 8TZ
United Kingdom
Tel. +44.(0)20.70408677
E-mail: harry@airc.info
Website: www.cass.city.ac.uk/airc**

TELL ME WHAT YOU WANT, WHAT YOU REALLY, REALLY WANT!

An Exercise in Tailor-Made Synthetic Fund Creation

Harry M. Kat*

Helder P. Palaro**

This version: October 9, 2006

Please address all correspondence to:

Harry M. Kat
Professor of Risk Management and
Director Alternative Investment Research Centre
Cass Business School, City University
106 Bunhill Row, London, EC2Y 8TZ
United Kingdom
Tel. +44.(0)20.70408677
E-mail: harry@harrykat.com

*Professor of Risk Management and Director Alternative Investment Research Centre, Cass Business School, City University, London. **PhD Student, Cass Business School, City University, London. The authors like to thank CSI and Rudi Cabral for allowing us to access the CSI futures database.

TELL ME WHAT YOU WANT, WHAT YOU REALLY, REALLY WANT!

An Exercise in Tailor-Made Synthetic Fund Creation

Abstract

Recently, Kat and Palaro (2005) showed how dynamic trading technology can be used to create dynamic futures trading strategies (or ‘synthetic funds’ as we call them), which generate returns with predefined statistical properties. In this paper we put their approach to the test. In a set of four out-of-sample tests over the period March 1995 – April 2006 we show that the Kat and Palaro (2005) strategies are indeed capable of accurately generating returns with a variety of properties, including negative correlation with stocks and bonds and high positive skewness. Under difficult conditions, the synthetic funds also produce impressive average excess returns. Combined with their liquid and transparent nature, this confirms that synthetic funds are an attractive alternative to direct investment in popular alternative asset classes such as (funds of) hedge funds, commodities, etc.

Keywords: synthetic fund, dynamic trading, correlation, skewness, asset allocation.

JEL Classification: G11, G13, G23.

Introduction

Over the last 10 years, investors have had a lot to deal with. Stock markets went up and came down again in an unprecedented fashion. At the same time interest rates came down to historically low levels. As a result, many of today's investors have difficulty seeing profit potential in traditional assets. Stock markets are hesitant and bond prices will come down when interest rates rise again. With memories of double-digit returns still fresh, this is driving them towards 'alternative' investments.

Hedge funds have become extremely popular. Investing in hedge funds comes with many drawbacks, however, including the need for extensive due diligence, liquidity, capacity, transparency and style drift problems, excessive management and incentive fees, and possibly regulatory problems as well. As long as investors believe they will be rewarded with (close to) double-digit returns, they will take these problems for granted. However, given the low level of interest rates, shrinking risk premiums all across the board and a hedge fund industry that has grown 20-fold over the last 15 years, disappointment seems almost inevitable.

Commodities, emerging markets and credit-based structures have also become increasingly popular over the last couple of years. There has been little research into these asset classes' investment merits, however. As a result, their risk profile is not well understood. Over recent time, these asset classes have shown very good performance, which is attracting even more money. It is hard to say, however, what the future might bring and whether at current prices investment in these asset classes is still as good an idea as it might have been 5 years ago.

In their quest for new diversifiers, investors have systematically overlooked one alternative. Modern risk management techniques make it possible to obtain virtually any desired risk profile by dynamically trading traditional assets such as cash, stocks, bonds, etc. This means that investors do not necessarily have to venture into the great unknown of alternative investments to find new diversification opportunities. They can be found right on their doorstep. An additional benefit is that this form of diversification avoids the drawbacks of alternatives, as trading is mechanical and done exclusively in liquid markets.

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 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 we can design a dynamic trading strategy that, under all possible scenarios, provides the same payoff as the option, then, to prevent arbitrage, the option price must be equal to the amount required to start off the replication strategy.

Most applications of dynamic trading have been in option pricing and have therefore concentrated on generating specific payoff profiles. The work of Dybvig (1988a, 1988b) is the most notable exception. Instead of specific payoff functions, Dybvig concentrated on payoff distributions and showed that, again under certain assumptions, for a payoff to be efficient it has to allocate wealth as a non-decreasing function of the value of the underlying index. Dybvig's work did not attract much further attention, until it was rediscovered by Amin and Kat (2003). The latter developed dynamic trading strategies, trading the S&P 500 and cash, which generate returns with the same marginal distribution as the return of a given hedge fund or fund of hedge funds. With two ways, the fund and the strategy, to obtain the same distribution, this allowed them to evaluate hedge fund performance in a completely new way.

The Amin and Kat (2003) replication procedure was recently extended to a bivariate setting by Kat and Palaro (2005). Apart from the marginal distribution of the fund return, the latter also replicated the dependence structure between a fund and an investor's existing portfolio. This is a very significant step forward as, after several years of mediocre performance, most investors nowadays are not so much attracted to hedge funds because of the promise of superior returns, but primarily for their diversification potential.

Kat and Palaro (2005) used their replication technique to replicate and evaluate the returns of existing hedge funds¹. *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 to be available and guess what a fund's true risk-return profile is. Given an investor's existing portfolio, we can now structure a special tailor-made strategy (or 'synthetic fund' as we will call these strategies) 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.

We could even take the above idea one step further and, instead of creating a synthetic fund as an addition to an investor's existing portfolio, replace the investor's entire portfolio 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, we could simply design a synthetic fund that produced returns with exactly the characteristics they were after.

¹ The evaluation results can be found in Kat and Palaro (2006a, 2006b).

In this paper we put the above to the test and ask whether, in practice, it is really possible to create synthetic funds, which generate returns with predefined statistical properties. In the next section, we briefly discuss the Kat and Palaro (2005) technique and how it can be used to create synthetic funds with predefined return characteristics. In section 3, we put the technique to the test. We design four different synthetic funds with a variety of return characteristics and study how they perform out-of-sample. Section 4 contains a brief comment on synthetic funds' alpha. Section 5 concludes.

2. Fund Creation Methodology

We design our synthetic funds using the same technique as in Kat and Palaro (2005). Mathematical details can be found in the latter paper. The basic idea is quite straightforward, however. Obviously, the first step is to decide on the return characteristics of the fund to be created, including its relationship with the 'reference portfolio'. When the synthetic fund is meant to further diversify some portfolio, as will typically be the case, then the reference portfolio equals that portfolio, or a good proxy. The next step is the selection of the 'reserve asset'. The latter is the main source of uncertainty in the fund. Although allocations to the reserve asset will change over time, the strategy will never sell the reserve asset short. As such, it can be interpreted as the core portfolio of the fund. Next, we design an exotic option, which, given the bivariate distribution of the return on the reference portfolio and the reserve asset, has the exact same return characteristics as the fund we want to create. Finally, we derive a hedging strategy for the above option. Mechanical execution of this strategy will produce the desired returns.

In the above procedure there are many details that require attention. One of them is the design of the option to be replicated. The same statistical properties can be found in many different options. Fortunately, as shown in Kat and Palaro (2005, Appendix I), the cheapest of those options is easily identifiable. This means that our synthetic funds can be based on the most efficient trading strategies available and therefore offer investors the highest possible risk premium. Another important point concerns the derivation of the hedging strategy. Ideally, unlike most popular option pricing models, we would want to take transaction costs into account when doing so. Instead

of a Black-Scholes–Merton type model, we therefore use the multivariate option pricing model of Boyle and Lin (1997), which explicitly allows for transaction costs.

3. 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 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 1: Overview of 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 futures². Fund 2 is the same as fund 1 except that apart from zero correlation we also aim for a significant degree of positive skewness³. 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 volatility⁴. Finally, in fund 4 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.

² See Kat and Oomen (2006a, 2006b) for details on the statistical behaviour of commodity futures investments.

³ We have to raise excess kurtosis to 10 since it is very difficult to generate significant skewness without extra kurtosis.

⁴ See Carr and Wu (2006) or Kat and Tassabehji (2006) for details on volatility investment.

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. 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. The same argument does not apply to correlation though. Since low correlation tends to go together with a low expected return, it is typically preferable to select a reserve asset, which has a relatively high correlation with the reference portfolio and subsequently adjust the latter downwards via the trading strategy.

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 time⁵. More specifically, we assume the reserve asset consists of an equally-weighted portfolio of 3-month Eurodollar, 2-year note, 10-year note, T-bond, S&P 500, Russell 2000 and GSCI futures. This captures three main asset classes. The resulting portfolio is quite well diversified, with, over the period 1993-2006, an annualised volatility of 6.35%. 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 back to 5.3% in 2006. Commodity price rises caused the GSCI to rise by 138% over

⁵ 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.

2002 - 2006. Finally, the last decade also saw its fair share of crises: Thailand, 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.

Fund 1

Let's assume a USD-based investor lived in March 1995 and started synthetic fund 1. Before we look at what kind of returns he would have generated over time, figure 1 shows the payoff function, which the investor will be aiming to produce as per March 1995. 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 portfolio⁶.

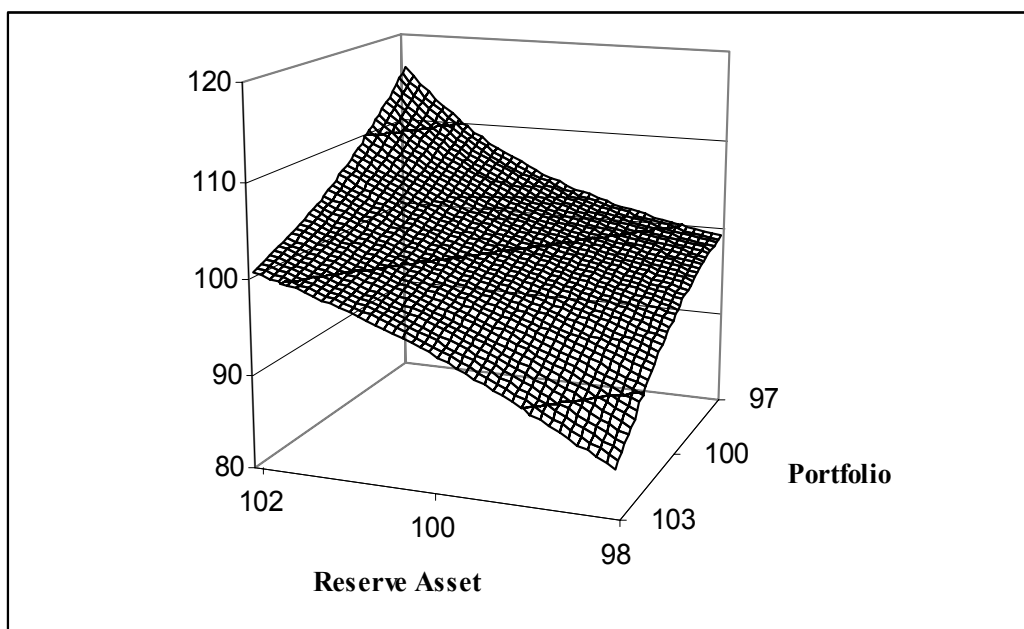


Figure 1: Target payoff synthetic fund 1, March 1995.

⁶ Note that this makes the price of the reduction in correlation dependent on the risk premium on the investor's portfolio.

Figure 2 shows the evolution of the standard deviation of the synthetic fund return over the period March 1997 – April 2006⁷, with the straight line representing the target value of 12%. The graph clearly shows that over the entire 9-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 2: Standard deviation synthetic fund 1, March 1997 – April 2006.

Figure 3 shows the evolution of the skewness of the synthetic fund return over the same period, while figure 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.

⁷ Although the fund starts trading in March 1995, the graph in figure 1 (as well as the figures that follow) starts in March 1997 because to meaningfully estimate standard deviation we need at least 24 observations.




Figure 3: Skewness synthetic fund 1, March 1997 – April 2006.




Figure 4: Correlation synthetic fund 1, March 1997 – April 2006.




Figure 5: Mean excess return synthetic fund 1, March 1997 – April 2006.

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 5 shows the evolution of the mean excess return (over 1-month USD Libor) on the synthetic fund over the period March 1997 – April 2006. From the graph we see that over the period studied the fund’s mean excess return converged to around 6%. With an average 1-month USD Libor rate of little over 4%, this implies a total return of about 10%, which is significantly more than what the average fund of hedge funds produced over the same period. We also see a substantial dip in 1997/98, which was quite a troublesome period with crises in Thailand and Russia followed by the near-collapse of LTCM. When interpreting 1997/98 we have to keep in mind that the fund’s track record only starts in March 1995. By October 1997 we therefore only have 30 observations available, meaning that the negative returns experienced in October 1997 and the following months have a relatively strong impact on the mean.

	50/50	Fund 1	10% Fund	20% Fund	30% Fund
Mean Return	9.71%	11.42%	9.88%	10.05%	10.22%
Volatility	8.34%	12.35%	7.55%	7.01%	6.77%
Skewness	-0.19	0.21	-0.25	-0.26	-0.18
Sharpe ratio	0.68	0.59	0.77	0.85	0.91

Table 2: Properties overall portfolio with varying allocations to synthetic fund 1, March 1995 – April 2006.

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 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. From their Sharpe ratios it appears that on a stand-alone basis the investor's original portfolio did better than the fund. Mixing in the fund, however, things change considerably. Due to the zero correlation between the fund and the original portfolio, 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 March 1997 – April 2006 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 3 shows the average daily trade size for the above synthetic fund over the period March

1995 – April 2006, 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 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 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	18	15
Eurodollar	15	13
2-year Note	18	15
10-year Note	35	31
T-Bond	60	53
GSCI	72	63

Table 3: Average daily trade size synthetic fund 1, March 1995 – April 2006.

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 6 shows the payoff function as per March 1995. 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.

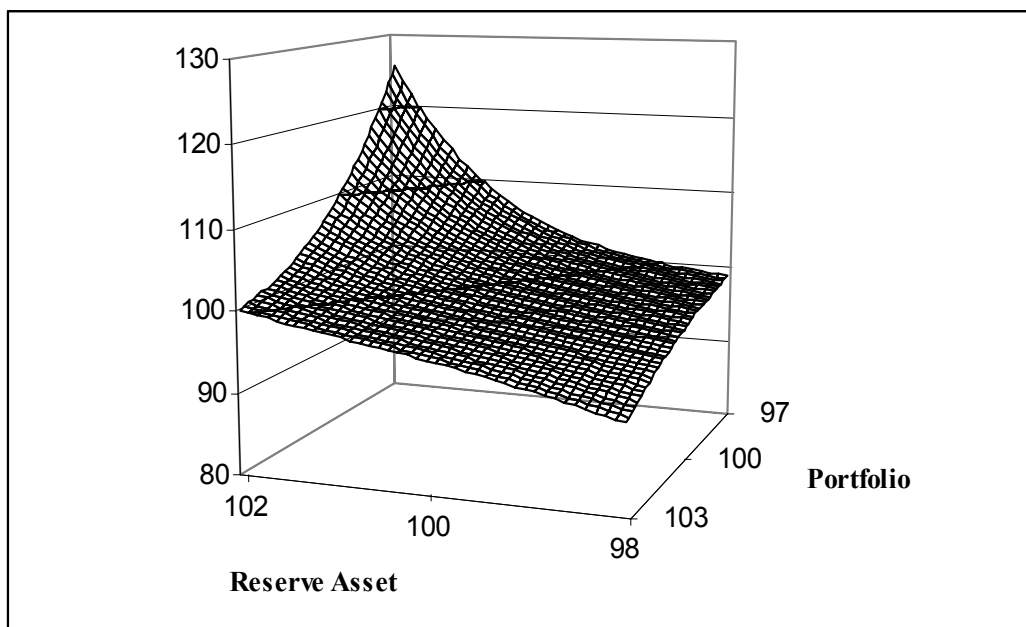


Figure 6: Target payoff synthetic fund 2, March 1995.

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 7, however, shows the evolution of the skewness of the synthetic fund return. It shows that in 1998-99 we lose some skewness due to the equity bull market, but in 2000-01 we regain that thanks to the equity bear market. Over the entire period studied, the skewness of the synthetic fund return never deviates far from its target value. Figure 8 shows the evolution of the mean excess return (again over 1-month USD Libor) on the synthetic fund. The graph looks very similar to that in figure 5, except that it dips just a little deeper in 1997/98 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 7: Skewness synthetic fund 2, March 1997 – April 2006.




Figure 8: Mean excess return synthetic fund 2, March 1997 – April 2006.

	50/50	Fund 2	10% Fund	20% Fund	30% Fund
Mean Return	9.71%	9.52%	9.69%	9.67%	9.65%
Volatility	8.34%	12.80%	7.59%	7.10%	6.93%
Skewness	-0.19	2.23	-0.23	-0.17	0.11
Sharpe ratio	0.68	0.43	0.74	0.79	0.81

Table 4: Properties overall portfolio with varying allocations to synthetic fund 2, March 1995 – April 2006.

Table 4 shows how synthetic fund 2 performed in a portfolio context. Again, and even more so than before, 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. Similar to variance, however, in a portfolio context it is not so much the skewness of the various portfolio components that matters, but much more the co-skewness between them. For more efficient skewness reduction, we could structure a fund with more appropriate co-skewness properties. Unfortunately, since this would seriously complicate the mathematics behind the fund strategy, this is outside the scope of the current paper.

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 March 1995 is shown in figure 9. Comparing this graph with the payoff function for fund 1 as shown in figure 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.

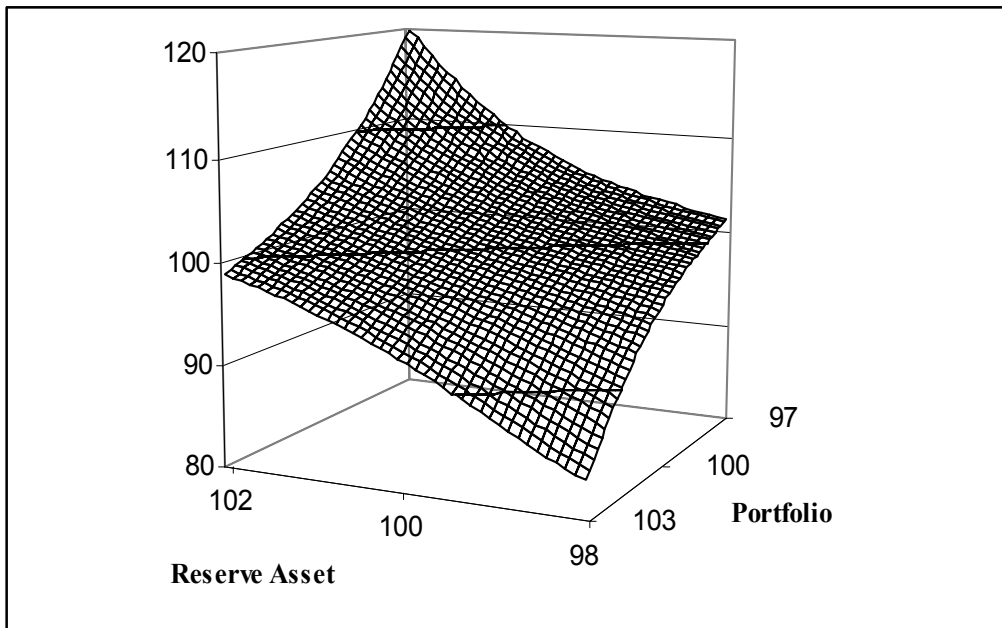


Figure 9: Target payoff synthetic fund 3, March 1995.

Figure 10 shows the evolution of the correlation between the synthetic fund and the investor's portfolio, while figure 11 shows the mean excess return. The graph in figure 10 shows that the correlation of the synthetic fund return stayed close to its target value over the full 9-year period. Figure 11, however, shows that this does not come for free as the mean excess return of the fund converges to no more than 2%. 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 10: Correlation synthetic fund 3, March 1997 – April 2006.




Figure 11: Mean excess return synthetic fund 3, March 1997 – April 2006.

	50/50	Fund 3	10% Fund	20% Fund	30% Fund
Mean Return	9.71%	6.81%	9.42%	9.13%	8.84%
Volatility	8.34%	12.21%	7.00%	5.89%	5.18%
Skewness	-0.19	0.21	-0.22	-0.21	-0.03
Sharpe ratio	0.68	0.22	0.76	0.86	0.92

Table 5: Properties overall portfolio with varying allocations to synthetic fund 3, March 1995 – April 2006.

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 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.

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 12: Scatterplot synthetic fund 4 excess return versus investor’s portfolio excess return, March 1995 – April 2006.

Figure 12 shows a plot of the fund excess return versus the excess return on the investor’s portfolio over the period March 1995 – April 2006. 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 13. The graph in figure 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.

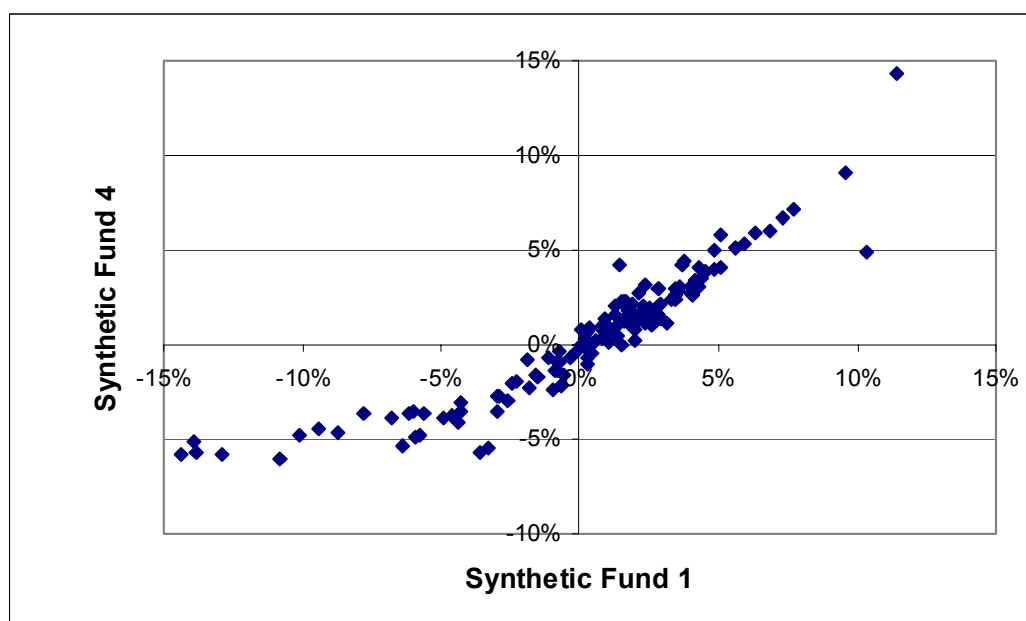


Figure 13: Scatterplot synthetic fund 4 versus synthetic fund 1 excess return, March 1995 – April 2006.

Although it does leave some residual risk, synthetic puts tend to work out substantially cheaper than real puts purchased in the OTC market. That doesn’t mean synthetic puts come for free though. This can be seen by comparing figure 14, which shows the evolution of the mean excess return of fund 4, with figure 5. This shows that the long-term mean excess return for fund 4 lies around 1% lower than that for fund 1. It also shows that after the 1997/98 dip, fund 4 is slower to recover than fund 1. The reason for this is that when the fund value approaches the floor, most of its market exposure is sold off. This prevents the fund value from falling further, but at the same time makes it more difficult to make up the loss. Finally, it should be noted that despite the floor we are unable to avoid the October 1997 – November 1998 dip

in mean excess return. This is because in none of the months that make up this period the fund value came down by more than 5%.




Figure 14: Mean excess return synthetic fund 4, March 1997 – April 2006.

	50/50	Fund 4	10% Fund	20% Fund	30% Fund
Mean Return	9.71%	10.41%	9.78%	9.85%	9.92%
Volatility	8.34%	11.79%	7.57%	7.03%	6.76%
Skewness	-0.19	0.51	-0.24	-0.26	-0.19
Sharpe ratio	0.68	0.54	0.75	0.82	0.86

Table 6: Properties overall portfolio with varying allocations to synthetic fund 4, March 1995 – April 2006.

Table 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 inferior to those of fund 1 or fund 3. The Sharpe ratio rises, but not as much as in case 1 or 3. This strongly suggests that, although intuitively attractive, *from an overall portfolio perspective explicitly flooring a diversifier's return may not be optimal as it may disrupt that diversifier's workings in a portfolio context.*

4. 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. Conclusion

In this paper we have carried out four out-of-sample tests of the Kat and Palaro (2005) synthetic fund creation technique. Our 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, our 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.

References

Amin, G. and H. Kat, Hedge Fund Performance 1990-2000: Do the Money Machines Really Add Value?, *Journal of Financial and Quantitative Analysis*, Vol. 38, No. 2, June 2003, pp. 1-24.

Arrow, K., The Role of Securities in the Optimal Allocation of Risk-Bearing, *Review of Economic Studies*, Vol. 31, No. 2, 1964, pp. 91-96 (originally published in French in 1953).

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 (eds.), *Advances in Futures and Options Research*, Vol. 9, Jai Press, pp. 111-127.

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

Dybvig, P., Distributional Analysis of Portfolio Choice, *Journal of Business*, Vol. 61, No. 3, 1988a, pp. 369-393.

Dybvig, P., Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market, *Review of Financial Studies*, Vol. 1, No. 1, 1988b, pp. 67-88.

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

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

Kat, H. and H. Palaro, Who Needs Hedge Funds? A Copula-Based Approach to Hedge Fund Return Replication, Alternative Investment Research Centre Working Paper No. 27, Cass Business School, City University London, 2005.

Kat, H. and H. Palaro, Replication and Evaluation of Fund of Hedge Fund Returns, Alternative Investment Research Centre Working Paper No. 28, Cass Business School, City University London, 2006a.

Kat, H. and H. Palaro, Superstars of Average Joes? A Replication-Based Performance Evaluation of 1917 Individual Hedge Funds, Alternative Investment Research Centre Working Paper No. 30, Cass Business School, City University London, 2006b.

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