Investor's Choice: An Investor-Driven, Forward-Looking Optimization Approach to Fund of Hedge Fund Construction

(appeared in: "Fund of Hedge Funds. Performance, Assessment, Diversification and Statistical Properties", Greg Gregoriou (Editor), Elsevier Finance, 2006)

Clemens H. Glaffig

Panathea Capital Partners GmbH & Co. KG, Freiburg, Germany

Abstract

This article proposes a method to optimize portfolios of hedge funds, taking investor preferences as the starting point to define an objective function that will be flexible enough to include general investor preferences. In particular, we can include path- and market-dependent objectives. The article also develops a method to produce a forward-looking data set on which the optimization can be based. Our method is also particularly apt for a case when a portfolio of hedge funds is used as an overlay. We compare this approach to classical optimizations on empirical data to highlight the effects of the additional degrees of freedom we have included.

Investor's Choice:
An Investor-Driven, Forward-Looking Optimization Approach
to Fund of Hedge Fund Construction

Introduction

Portfolio construction and optimization has attracted a vast amount of research through the last decades, and is now embedded into a well developed mathematical framework. The classical portfolio optimization approach is concerned with a universe of standard assets and how to combine them into a portfolio that maximizes investor utility under general constraints relative to some perception of risk. Most frequently, the result is a type of mean-variance optimization. The typical form of utility functions and the definitions of risk, however, need to conform to demands of general acceptance and ease of solvability. Little room remains for individual investor preferences and risk perceptions.

With the evolution of financial markets, however, standard assets are being supplemented by various alternative assets. Hedge funds in particular differ in content and behaviour from classical assets and funds. There is currently considerable interest among practitioners in the problem of portfolio construction and optimization based on these new assets and trading styles. With hedge funds as the portfolio constituents, any deficiencies of the classical approaches become even more pronounced.

In this article, we deviate from the traditional approach by making investors' own preferences the starting point of our optimization. Our set-up of objective function and constraints is fully flexible to reflect path- and market-dependent preferences. Specifically, we let the investors decide what constitutes risk, because not all forms of risk are equally undesirable. Often, perceived risk is a function of current wealth (or accumulated profit) and of the state of other markets. The price we pay for this flexibility is that a simple analytical solution to our optimization problem is typically not derivable. We will solve the optimization by using a heuristic search method, in this case, a genetic search.

The issue in simply extending methods of classical portfolio optimization to hedge fund portfolios is twofold. On the one hand, for portfolios of classical assets, assumptions of market behaviour translate easily into the corresponding behaviour of the assets to be included in the portfolio; this is not the case for alternative assets and hedge funds. Under a future return scenario for major markets, hedge fund returns are not at all clear.

In order to draw conclusions about the behaviour of hedge funds from specifications of major market behaviour, it is necessary to better understand the fund's strategies. In other words, we need to understand the fund's "mapping" from a specific market situation into trading action. As a consequence, obtaining reliable forward-looking data turns out very difficult. So producing a forward-looking data set on which the optimization can be performed is much more problematic than it is for classic portfolio optimization.

On the other hand, however, alternative assets and trading strategies depend in a non-linear and asymmetric way on underlying markets, so the classical mean-variance approach for portfolio optimization is inadequate. Risk is not properly reflected within this framework. Progress has been made by considering utility functions that depend on other measures of risk, such as VAR, CVAR, and downside risk (see, e.g., Morton, Popova, and Popova [2004] and De Souza and Gokcan [2004]). Or, because many hedge funds exhibit considerable skew and kurtosis, by using a four-factor, mean-variance skew kurtosis utility function (see Davies, Kat, and Lu [2004]).

However, as we have noted, it is much more important with alternative assets to allow for objective functions and definitions of risk that reflect investors' own risk preferences. For most investors, risk perception varies with market developments and with the performance of their other assets. None of the standard approaches reflects the necessary path and market dependence.

With regard to the data set underlying our optimization, we use a mixture of forward-looking and historic data, and formulate a market view on patterns of general market behaviour for the future (or a weighted partition of views). We then return to the historic data set and select those periods of time "closest" to that pattern. The difference is that, in the pure forward-looking approach, which is based on a regression against style factors or strategy learning algorithms, one has complete freedom with regard to the specification of future market developments. The fund behaviour in this situation comes as an approximation/optimization from the past, and hence can only be approximately correct.

In our approach, the data is a close approximation of our modelled/assumed scenario (as we search for past periods with market developments close to our specification), but the behaviour is precise. In spirit, our view is forward-looking by ensuring that the "mapping" problem (mapping market behaviour into hedge fund behaviour) is circumvented by having history solve it for us.

We split the optimization into two parts: 1) the data set, and 2) the methodology.

The Data Set: Defining Market Patterns

The data set underlying the optimization is of crucial importance for the results. As is well known from classical mean-variance optimization, slightly different data sets underlying the optimization can produce vastly different portfolio decompositions. In the case of alternative assets and hedge funds, one of the following four approaches for data sets underlying the optimization has typically been used.

1. Historic Data

Much of the fund of hedge fund optimization is performed on the basis of historical data. The advantage is that data for the behaviour of hedge funds are objectively available, so all the risk parameters and statistics can easily be computed. However, if history does not more or less precisely repeat itself, the optimization becomes useless for construction purposes. For example, stop loss or limit behaviour may not become obvious for the underlying data set, but it may dominate the fund's behaviour subsequently. Or, trends may have fed the performance in the past, but subsequent sideways market movements may steer the investment managers toward more risk.

2. Simulations

If we know the return distribution of a fund, we can perform a simulation. However, historically fitted return distributions are a statistical summary of past behaviour, an interpolated historical frequency count of returns. Given that hedge fund returns are highly non-stationary, this approach suffers from similar defects as historical data because future return distributions may look completely different. In addition, while it certainly makes sense to allow for heavy-tailed distributions, they should be truncated to reflect stop losses, which are "hard-wired" into the risk procedures of many hedge funds. This fact is often ignored in simulations.

3. Future Projections

We could use future projections of hedge fund returns conditioned on market behaviour in conjunction with a Black-Litterman type of approach. Here, the future behaviour of the primary markets influencing a given hedge fund is weighted according to its plausibility to the investor. The success of this approach depends on the way market behaviour can be related to hedge fund behaviour (the mapping problem), i.e., the degree to which one is able to extract the dynamics of the fund and the strategies employed by the fund.

4. Partitioning of History

The partitioning of history approach tries to combine the advantages of (1) and (3). It forms a collection of likely market developments, each attached to the investor's probability of occurrence. It then goes through history selecting periods of time that fit the individual market patterns and performs the optimization on the data set.

We use the "partitioning of history" method, which classifies historic data into predefined market patterns, for our optimization. We choose the pattern or weighted average of patterns that we (or the investor) deem most likely for the future period. The chosen patterns can include scenarios that may be problematic for the other assets the investor is holding. In that context, this method will be valuable if the portfolio being constructed is to be used as an overlay to an existing, classical portfolio.

To illustrate this method, we need to define market patterns. The more formal approach we take to describe a pattern is justified because a market pattern like "up-trend" can be defined in a variety of ways into a set of conditions on market prices over a specific time period. This again can lead to very different data sets for the same targeted pattern, and hence to very different optimal portfolios. We therefore define a pattern class, a set of time intervals over which the market behaviour is “close” to the desired pattern.

A pattern could be defined by specifying a market M and a pattern period T, and defining a set of rules that market M must fulfil starting from some initial time t and ending at t + T. The set of rules could be expressed via a product of indicator functions, expressing the fulfilment of the rules. For example, a continuous upward pattern could be defined by the product of indicator functions

$$1_{\{M(t)\le M(t+1)\}}\cdot 1_{\{M(t+1)\le M(t+2)\}}\cdots 1_{\{M(t+T-1\le M(t+T)\}} \tag{1}$$

More generally, we define a pattern by first specifying the number d of different markets, which will be part of the definition (e.g., if we use the S&P 500 index, the ten-year U.S. Treasury rate, and the USD-EUR exchange rates, it would be $d = 3$). The length of the pattern, T, is chosen to correspond to the period in the future that we construct the portfolio. We denote t as a time scale that is coarse, such as a weekly or monthly time scale, where s denotes smaller time scales such as daily or more frequent.

Patterns are defined by describing the behaviour of certain statistics on a coarse scale and aggregating information from small scales, e.g., a weekly maximum of a market is typically aggregated by taking the maximum over all tics of that market during a week. Here, t denotes one week, and s indicates a specific tic within a week.

We define a set of m statistics/functionals that we apply to the individual markets to enter the pattern definition:

$$X_k(i,t) = F_k(M_i(s); s \in (t-1,t]) \in \Re, \text{ with } i \in \{1,2,\ldots,d\}, k \in \{1,2,\ldots,m\} \tag{2}$$

Each $X_k(i, t)$ aggregates and maps the information of market M_i for time interval (t – 1, t] into a real number. Functions F_k determine the statistics of the markets to be used. Common examples include $X_k(i, t): = \max (M_i(s), s \in (t – 1, t])$ or $X_k(i, t): = \min (M_i(s), s \in (t – 1, t])$, the maximum or minimum of the i-th market during the time interval (t – 1, t].

A T-period market M_i pattern starting at time t_0 is defined by a map from $\Re^{mT} \to \{0,1\}$:

$$PM_i(t_o,T): (X_1(i,t), X_2(i,t), \ldots X_m(i,t), t \in \{t_0+1, t_0+2, \ldots, t_0+T\}) \to \{0,1\} \tag{3}$$

Definition: *A time interval* $[t_0, t_0 + T]$ *will have T-period pattern* PM_i, *if for* PM_i: $\Re^{mT} \to \{0,1\}$.

$$PM_i(t_0,T) = 1 \tag{4}$$

A general T-period pattern P starting at time t_0 for markets M_1 to M_d will be defined as follows.

Definition: *A time interval* $[t_0, t_0 + T]$ *will have d-market, T-period pattern P, if for P:* $\Re^{dmT} \to \{0,1\}$

$$P(t_0,T) := \prod_i PM_i(t_0,T) = 1 \tag{5}$$

The definition of P represents the rules that combine the functionals and statistics as defined by X_k into a temporal pattern.

Example

Let d = 1, m = 3 For $t \in \{t_0 + 1, t_0 + 2, \ldots, t_0 + T\}$,

$$
\begin{aligned}
X_1(1,t) &= F_1(M_1(s), s \in (t-1,t]) := M_1(t) \\
X_2(1,t) &= F_2(M_1(s), s \in (t-1,t]) := \max(M_1(s), s \in (t-1,t]) \\
X_3(1,t) &= F_3(M_1(s), s \in (t-1,t]) := \min(M_1(s), s \in (t-1,t])
\end{aligned} \tag{6}
$$

If we then let P be given by the indicator function

$$P(t_0, T) = PM_1(t_0, T) = \prod_k 1\{X_k(t_0+1) \le X_k(t_0+2) \le \ldots X_k(t_0+T)\} \tag{7}$$

we have defined a pattern by which the successive closing levels, the successive highs and the successive lows of the market M_1, are each increasing over the time interval $[t_0, t_0+T]$, which would be a definition of a solid upward trend of market M_1.

For practical purposes, it is also useful to define pattern classes as time periods, which may not have a predefined pattern, but are close. The use of pattern classes ensures that the set of time periods classified is not too small. It also serves the stability of solutions for our optimization method, as slightly different pattern definitions may lead to very different optimal portfolios.

We note that with each d-market pattern P, we can associate the d maps PM_i defining P, and to each of these maps, we associate a subset $\wp_i \subset \Re^{mT}$ by

$$\wp_i(t_0, T) = PM_i(t_0, T)^{-1}(\{1\}) \tag{8}$$

We then define

$$
\begin{aligned}
Z_i(t) &= (F_1(M_i(s), F_2(M_i(s), \ldots, F_m(M_i(s); s \in (t-1,t]); s \in (t-1,t]) \in \Re^m \\
&= (X_1(i,t), \ldots, X_m(i,t))
\end{aligned} \tag{9}
$$

and

$$Z_i(t_0, T) = Z_i(t_0+1), Z_i(t_0+2), \ldots, Z_i(t_0+T)) \in \Re^{mT} \tag{10}$$

We define pattern class P as those periods for which market behaviour, as encoded by $(Z_1(t_0T), \ldots, Z_d(t_0T))$, is close to the (dmT)-dimensional set ($\wp_1(t_0, T)$, $\wp_2(t_0, T)$, …, $\wp_d(t_0, T)) \subset \Re^{dmT}$.

Definition: *The time period $[t_0, t_0+T]$ is said to belong to ε-pattern class P, if*

$$\sum_i \left\| \wp_i(t_0, T) - Z_i(t_0, T) \right\| \le \varepsilon \tag{11}$$

where $\|\bullet\|$ corresponds to the Euclidian distance in $\Re^{mT}$.

The set of all time periods belonging to ε-pattern class P will be denoted by $[P]\varepsilon$. Note that:

- A time period can belong to more than one pattern class.
- If history must be partitioned into a family of pattern classes, we can easily use a kernel-based classification instead of the above definition.

The Methodology: Investor-Driven Objectives and the Optimization Algorithm

The methodology used for most traditional optimizations is largely influenced by the desire to solve problems analytically. The mathematical theory of constrained optimization is well developed and can accommodate the more obvious constraints like integer constraints for investment sizes as well as non-linear objective functions. But adding more customized preferences will require numerical procedures or search algorithms.

Our approach is driven by translating all the preferences of a specific investor into a quantitative framework, without any restriction on the functional form of these preferences. As the perception of risk is not unique, we will not limit ourselves to a predefined and dogmatic definition of risk via volatility, semivolatility, VAR, CVAR, or downside risk. We allow the investor to choose what he or she perceives as “risk” by essentially allowing him or her to decide which scenarios to avoid.

We will reflect strict boundary conditions via constraints on the optimization. We also build preferences into the objective function, penalizing scenarios to be avoided and rewarding desired scenarios. Specifically, we allow scenarios that depend on the return history of the portfolio, on the market environment, or on any other form of path dependence. Restrictions like the minimum number of portfolio members, leverage conditions, and maximum fund allocation can easily be accommodated and will be reflected in the constraints.

We consider portfolios

$$\text{Port}\ (\omega, t) = \sum w_i F_i(t) \tag{12}$$

where the $F_i(t)$ are the funds to be considered for inclusion valued at time t, and ω = (w_1, w_2, …, w_N) is the vector of weights, constant throughout the specific future investment period for the individual funds within the portfolio.

The Objective Function: The Sum of All Fears

The investor preferences will be reflected in the objective functions via rewards and penalties. The objective function to be maximized is

$$\text{Obj}(\omega, t) = \sum_i \alpha_i 1_{\{Si\}}(\omega, t) \tag{13}$$

where $1_{\{Si\}}$ is the indicator function for S_i, which denotes the i-th situation rewarded ($\alpha_i \geq 0$) or penalized ($\alpha_i \leq 0$). S_i may be time-, path-, and market-dependent. Given that any continuous function can be approximated by step functions, this form is more general than any of the classic utility functions.

Examples

For all examples, we assume the investment period is from t_0 to t_0 + T. Typical S_i include:

- S_i = {the portfolio return over the target period [t_0, t_0 + T] is between R_{i-1} and R_i}, with α_i, e.g., equal to a concave function of R_{i-1} or constant. Then:

$$\alpha_i 1_{\{Si\}}(\omega, t) = \alpha_i$$

if $(Port(\omega, t_0 + T) - Port(\omega, t_0))/Port(\omega, t_0) \in [R_{i-1}, R_i]$

and $t = t_0 + T$

0 otherwise (14)

- S_i = {negative tail co movement with market M over period $[t-1, t] \subset [t_0, t_0 + T]$}, characterized by:

$$\alpha_i 1_{\{Si\}}(\omega, t) = \alpha_i 1_{\{R(t) \le -x\}} 1_{\{|RM(t)| \ge y\}} \tag{15}$$

where $1_{\{R(t) \le -x\}}$ denotes the indicator function for event R(t) ≤ −x, and R(t) is the portfolio return over period [t – 1, t]. $1_{\{|RM(t)| \ge y\}}$ denotes the indicator function for the event |RM(t)| ≥ y, where |RM(t)| is the absolute value of the return of market M over the period [t – 1, t], x and y to be specified, so that

$\alpha_i 1_{\{Si\}}(\omega, t) = \alpha_i$ if $R(t) \le -x$ and $|RM(t)| \ge y$

0 otherwise (16)

- S_i = {the portfolio return over period [t – 1, t] ⊂ [t_0, t_0 + T] is between R_{i-1} and R_i, and the accumulated portfolio return over period [t_0, t] is larger than x}
-

$\alpha_i 1_{\{Si\}}(\omega, t) = \alpha_i(x)$ if $(Port(\omega, t) - Port(\omega, t-1))/Port(\omega, t-1) \in [R_{i-1}, R_i]$

and $(Port(\omega, t) - Port(\omega, t_0))/Port(\omega, t_0) \ge x$

0 otherwise (17)

The Constraints – the Aggregation of "Hard" Exclusions

The constraints will summarize the scenarios we exclude. We list only a few examples, which typically include:

- Non-negativity of weights, i.e., $w_i \ge 0$.
- Fully invested portfolio or no leverage, i.e., $\sum w_i = 1$ or $\sum w_i \le 1$.
- The maximum share of individual funds, i.e., max (w_i) ≤ A.
- The minimum number of funds in which the portfolio is invested.

- Any hard stop loss that is potentially contingent on other markets.
- Minimum investment sizes.

The Path to Optimization

The optimization proceeds as follows. Let $[t_1, t_n]$ denote the time period for which we have reliable historic data for all the funds we want to consider. Let T denote the period into the future for which the portfolio is to be constructed and over which the portfolio weights will not be changed. We first define a T-period pattern P for the market development in the future, according to the definitions above, by specifying a set of defining markets M_i and statistics F_k (we restrict ourselves to a single pattern, optimizations over a weighted average of patterns works analogous).

Within $[t_1, t_n]$, we consider all connected T-period time intervals, from which we select those in pattern class $[P]\varepsilon$ for pre specified ε. We allow different periods to overlap by a pre specified maximum number of time intervals. The members of this set of connected T-period time intervals will be close to the originally defined pattern P, and will serve as the data set on which our optimization is based.

The optimal portfolio is found by maximizing the sum of the objective functions over all T-period intervals belonging to pattern class [P] with respect to weight vector ω:

$$\mathrm{Max}(O(\omega)) := \mathrm{Max}_{\omega}\left(\sum_{[t,t+T]\in[P]} \sum_{u\in\{t+1,t+2,\dots,t+T\}} \mathrm{Obj}(\omega, u)\right) \tag{18}$$

where the inner sum runs over all time steps u within a given T-period, the outer sum runs over all T-periods in pattern class [P], and the constraints C are satisfied.

We use a genetic search algorithm to solve the problem. We recommend that readers not interested in the specifics of this algorithm skip the next section and move on to the empirical analysis section.

Search Algorithm

In our genetic search algorithm, each portfolio $\mathrm{Port}(\omega)$ is represented by its weight vector ω. The algorithm produces an initial population G of portfolios, each represented by its weight vector. All the portfolios within the population will compete against each other, and must fulfil the constraints.

Let F_i denote the i-th fund of the pool of admissible funds, and let $\omega(n, j)$ denote the weight vector, representing the n-th portfolio of the j-th population, $\omega(n, j) \in \Re^N$ with

$$\omega(n, j) = (w(1, n, j), w(2, n, j), \dots, w(N, n, j)) \tag{19}$$

where w(i, n, j) = the weight of the i-th fund F_i.

The first population of portfolios is generated with random $w(k, n, l)\,\forall k, n$. After the objective function of all portfolios corresponding to the $\omega(n, 1)$ of the first population is evaluated, we choose the two with the highest objective functions in order to generate the next population (ties are broken by random choice). In the literature, a number of alternative selections, e.g., tournament selection, are discussed. The next generation is generated from the top two portfolios (“the parents”) by using the following genetic operators:

Crossover: We randomly select a coordinate k. Then we generate a portfolio of the new generation by taking the first k portfolio weights from the first top portfolio and the remaining $N - k$ from the second top portfolio. A second new portfolio is generated the same way by interchanging the two top portfolios and applying the same procedure.

Mutation: We randomly change a portfolio weight with a predefined probability. The two top portfolios of a given generation will always be members of the next generations. Genetic operators are applied to the parents of a given population, until the number of the new population is $G - 2$. The new population then consists of the $G - 2$ new portfolios and the old parents. This procedure is repeated until a termination condition is fulfilled or after a fixed number of generations. The portfolio in the last population with the highest objective function is termed the “optimal portfolio.”

Note that a number of refinements to the genetic operators are useful: The algorithm should find a good balance between local search, i.e., local optimization, and exploration, i.e., exploring other regions of weight space to avoid being trapped with local maxima. Both previously mentioned operators can also be split into fine and broad search sections by restricting crossover to interchanging. For more details, see Banzhaf, Nordin, Keller and Francone [1998] or Weicker [2002].

Empirical Analysis: Exhibiting the New Degrees of Freedom

This section compares results for different classical approaches to the above approach. The emphasis here is not to determine superiority, but to illustrate the flexibility of the different approaches and how changes in optimization parameters influence the resulting optimal portfolio.

For all approaches, we let the obtained optimal portfolio run over a forward-looking period, i.e., we analyze out-of-sample performance. For our universe of hedge funds, we use a set of CSFB/Tremont hedge fund indices, which represent the following trading styles:

- Global Macro
- Convertible Arbitrage
- Equity Market-Neutral
- Distressed
- Event-Driven
- Fixed-Income Arbitrage
- Dedicated Short
- Long/Short Equity
- Managed Futures
- Emerging Markets
- Risk Arbitrage

Clearly, this universe of assets is not an optimal base from which to set up a realistic portfolio. However, for our analysis this is essentially unimportant. We want to emphasize the workings of our method and demonstrate the effects of its individual degrees of freedom. We use monthly data, even though most hedge funds nowadays provide daily performance data.

We construct our portfolio based on information before December 31, 2004. Our investment horizon is six months, i.e., a portfolio left unchanged from January 1, 2005-June 30, 2005. Over that six-month horizon, the various individual indices we consider as our funds have performed as follows:

	Global Macro	Convertible Arbitrage	Equity Market-Neutral	Distressed	Event-Driven	Fixed-Income Arbitrage
Return (%)	2.97	-6.34	1.45	4.09	3.60	-1.12
Max Mthly DD (%)	-0.25	-3.13	-0.34	-0.06	-0.64	-1.24
Max DD (%)	-0.25	-7.42	-0.56	-0.06	-0.64	-2.39

	Dedicated Short	Long/Short Equity	Managed Futures	Emerging Markets	Risk Arbitrage
Return (%)	13.42	0.86	-0.91	5.62	1.02
Max Mthly DD (%)	-5.91	-1.55	-5.39	-1.88	-0.54
Max DD (%)	-6.12	-2.69	-5.39	-1.88	-0.54

where Max Mthly DD denotes maximum monthly drawdown, and Max DD denotes maximum drawdown over the period.

We run three different classical optimizations, characterized by their respective objective functions:

$$\text{CL1: } \mathrm{R}/(\alpha \mathrm{Vol}) \tag{20}$$

where R = portfolio return over the period, Vol = $\sum|R_i|$, R_i = the return of the portfolio in the i-th month of the period, and α is a parameter.

$$\text{CL2: } \mathrm{R}/(\beta \mathrm{Semivol}) \tag{21}$$

where R = return of the portfolio over the period, Semivol = $\sum \min(R_i, 0)$, R_i is the i-th monthly return of the portfolio in the period, and β is a parameter.

$$\text{CL3: } \mathrm{R}/(\gamma \mathrm{MaxDD}) \tag{22}$$

where R = the portfolio return over the period, MaxDD = the maximum drawdown of the portfolio during the period, and γ is a parameter.

We run all these optimizations under the same set of constraints:

- Non-negativity of weights, that is., $w_i \geq 0$.
- No leverage, that is, $\sum w_i \leq 1$.
- The maximum share of individual funds must equal 25%, that is, max $(w_i) \leq 25\%$.
- The minimum investment size is 1% of the total portfolio.

We run the optimizations of CL1, CL2, and CL3 over the data set, which consists of the last sixty months before our starting point, that is, a purely historic data set.

The performance results over the six-month period immediately following the optimization period and the obtained weightings for the indices are as follows (all numbers in percent):

	Return	Max Mthly DD	Max DD	GM	CA	EMN	D	ED	FIA	DS	LSE	MF	EM	RA
CL1 α = 1	1.07	-0.45	-0.64	25	13	25	10	0	0	0	0	0	0	0
CL2 β = 1	1.85	-0.78	-0.99	11	11	25	13	14	16	7	0	1	1	1
CL3 γ = 1	2.11	-1.08	-1.09	15	13	22	0	21	16	11	0	0	2	0

Note: The remaining cash was invested at LIBOR flat.

In general, the specific penalty form of the classical approach -- here a form of volatility, semivolatility, and drawdown - weighs the historic behaviour over the entire history under that very specific penalty. For example, the index for distressed funds had a large drawdown in the fall of 1998 during the Russian crisis, which resulted in a high penalty and low weight specifically within CL3. Conversely, throughout time, the low volatility of that index led to a large weight under CL1.

Moreover, funds (or indices here), were excluded due to a bias in the data set. This may not at all reflect an investor’s expectations of future market behaviour. Additional preferences such as including or excluding certain sectors must be considered on an ad hoc basis by restricting the universe of admissible funds for the portfolio. Note that a number of pure (i.e., 100% weighted) “fund” investments perform better than any of CL1, CL2, and CL3 on a risk-adjusted basis over the out-of-sample period. These disadvantages of the classical models would also prevail in simulation approaches.

For comparison, we run various models under the new methodology. We group them first according to which six-month market patterns we want to use to optimize the portfolio. We define three simple patterns:

> **Pattern A:** *For a given six-month period, the S&P 500 has at least four months of successively lower monthly highs and lows (a downtrend).*
>
> **Pattern B:** *For a given six-month period, the S&P 500 has at least four months of successively higher monthly highs and lows (an uptrend).*
>
> **Pattern C:** *For a given six-month period, the credit spread of Euro BBB Corporate Bonds has three up and three down months.*

We run our optimizations on six-month time periods that each have the specified pattern (note that in this simple analysis we do not use pattern classes). We allow overlaps of up to three months, i.e., the starting month for two time periods of the same pattern must be at least three months apart. We use data from January 1, 1994-December 31, 2004 to select the six-month time periods.

Pattern A corresponds to a downtrend in equities. The market of our horizon period, from December 31, 2004-June 30, 2005, exhibits pattern A. Consequently, our horizon period does not have pattern B, which corresponds to an equity uptrend. Pattern C corresponds to a sideways movement in credit spreads. The horizon period also exhibits pattern C.

For each pattern, we run optimizations under different objective functions.

Pattern A:
Obj A1: The portfolio return over each six-month period in the data set, plus a penalty of 0.50% for months the portfolio return is below −4% while the S&P 500 return is larger than 3%.

Obj A2: The portfolio return over each six-month period in the data set, plus a penalty of 0.50% for months the portfolio return is below −4% while the BBB corporate bond spread increases more than 25%.

Pattern B:
Obj B1: The objective function is the same as in A1.

Pattern C:
Obj C1: The objective function is the same as in A2.

The objective functions are obviously somewhat randomly chosen. But our purpose is to demonstrate the effects of using different data sets to underlie the optimization and the effects of reflecting different preferences within the objective functions.

The A1 objective function balances pattern A. The pattern of an equity downtrend strongly favours a negative correlation to equity markets, but the penalty term of the objective function limits the correlation between negative portfolio return tails and positive equity return tails.

The A2 objective function limits the correlation between negative portfolio return tails and the tails of credit spread widening, and thus retains its corporate bond-friendly bias even during an equity downturn. While downturns in equity markets generally correspond to widening spreads, this is not necessarily true for “soft” downturns. In this sense, the A2 objective function favours soft equity downturns (downturns via the defined pattern, filtering out the soft downturns via the objective function). All of these optimizations will be run under the same set of constraints as before.

The following table gives the performance results over the six-month period January 1, 2005-June 30, 2005, as well as the obtained weightings for the indices (all numbers in percent):

	Return	Max Mthly DD	Max DD	GM	CA	EMN	D	ED	FIA	DS	LSE	MF	EM	RA
A1	4.27	-0.85	-0.85	25	0	0	2	0	0	25	0	25	0	0
A2	3.48	-1.49	-1.49	25	14	2	0	1	0	25	0	11	0	0
B1	2.72	-0.25	-0.25	25	0	0	25	2	0	0	15	0	7	1
C1	2.15	-0.26	-0.26	25	5	1	25	3	0	0	6	0	1	0

We make the following observations:

- Patterns matter. The right pattern versus the wrong pattern, A versus B, results in a clearly superior performance for A1 and A2 (although drawdown was not part of our objective in this example). It also confirms that the chosen portfolio reflects the pattern. A1, assuming an equity downtrend, weighs the dedicated short at the maximum allowed (25%), while B1 weighs it at its minimum allowed (0%). The classical approaches have weights somewhere in between because of drawdown reduction effects and the fact that the data set (the sixty months preceding December 31, 2004) displays extended periods of equity upturns and downturns.
- The long-short equity index has a 0% weight for downtrend patterns A1 and A2 due to its long bias. This is a part of the uptrend pattern in B1, but plays no part in classical approaches due to bias in the data set (more severe equity downturns overall).
- The distressed index carries full weight under B1 and C1, as their data sets exclude the problematic periods of that index and make use of their consistently good performance outside market downturns.
- Managed futures are included for A1 and A2, as they do relatively well in equity downtrends. That relative advantage diminishes, however, if more equity uptrends enter the data set.
- The added penalty in the objective function A1 results in the omission of some of the indices with high negative tail correlations. Changing the objective function from A1 to A2 to reflect a corporate bond-friendly bias results in consideration of the convertible arbitrage index, which contains a corporate credit feature.

These are only a few observations about some variations in the full flexibility of our new approach. However, it is clear that our approach greatly increases the number of degrees of freedom compared to classical approaches. Different assumptions about future market developments translate directly into different portfolio decompositions, while the use of pattern classes ensures that similar future pattern definitions produce similar or the same optimal portfolios. Changes in the objective function that allow us to consider individual investor preferences are reflected in the optimal portfolio.

Conclusions

Our methodology for constructing a portfolio of hedge funds starts from investor preferences and aversions and defines a fully flexible objective function and set of constraints. We have thus addressed some of the deficiencies of previous approaches, in which individual preferences and risk perceptions are subordinated to general definitions of objectives and risk. This flexibility allows investors to base risk/return objectives on path- and market-dependent measures, and to express preferences about markets or trading styles to be included in the optimization, rather than having them arbitrarily imposed ex ante. Thus the translation of investor preferences into portfolio weights is inherent within the optimization. This feature is what gives our approach a clear edge over classical approaches.

Moreover, we have combined a forward-looking approach with historical information on conditional hedge fund behaviour to produce a data set underlying the optimization. Investors can thus base their optimizations on a pattern of future market development for which reliable data on hedge fund behaviour exists, without the problematic use of historic data (pure or in simulations via historically derived distributions), or of the unreliable projections of conditional future hedge fund behaviour.

However, the price of this flexibility is the need for more elaborate search algorithms to solve the optimization problem. We use a genetic search here for its ease of implementation and its effectiveness in expressing constraints, but any other successful search algorithm will also do (see Schlottmann and Seese [2004] and references therein). The disadvantage of our data set is that we can only use past or close to past patterns. Reliable methods for predicting conditional hedge fund behaviour would not suffer from this restriction. But practitioners must decide which method is preferable to produce the data set underlying the optimization: using patterns that are close to past patterns and thus have more reliable data, or using the full flexibility to define a future pattern with potentially unreliable predictions about future hedge fund behaviour.

References

Abernathy, J., and A. Weisman. “The Dangers of Historical Hedge Fund Data.” In L. Rahl, ed. *Risk Budgeting*. London: Risk Books, 2000.

Alexander, C., and A. Dimitriu. “The Art of Investing in Hedge Funds: Fund Selection and Optimal Allocation.” Working paper, University of Reading, 2004.

Amenc, N., P. Malaise, L. Martellini, and D. Sfeir. “Using Index Options to Improve the Performance of Dynamic Asset Allocation Strategies.” Working paper, EDHEC Business School, 2004.

Amenc, N., and L. Martellini. “Portfolio Optimization and Hedge Fund Style Allocation Decisions.” *Journal of Alternative Investments*, 5, 2, 2002, pp. 7-20.

Banzhaf, W., P. Nordin, R. Keller, and F. Francone. *Genetic Programming*. *An Introduction*. San Francisco: Morgan Kaufmann Publishers, Inc., 1998

Davies, R., H. Kat, and S. Lu. “Fund of Hedge Fund Portfolio Selection: A Multiple-Objective Approach.” Working paper, Case Business School, 2004.

De Souza, C., and S. Gokcan. “Allocation Methodologies and Customizing Hedge Fund Multi-Manager Multi-Strategy Products.” *Journal of Alternative Investments*, 6, 4, 2004, pp. 7-21.

Devroye, L., L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition.* New York: Springer Verlag, 1996.

Fung, W., and D. Hsieh. “The Risk of Hedge Fund Strategies: Theory and Evidence from Trend Followers.” *Review of Financial Studies*, 14, 2, 2001, pp. 313-341.

Fung, W., and D. Hsieh. “Hedge Fund Benchmarks: A Risk-Based Approach.” *Financial Analysts Journal*, September/October 2004, pp. 65-80.

Glaffig, C. “Risk Management of Hedge Funds Using Neuro-Fuzzy and Genetic Algorithms.” Working paper, Panathea Capital Partners, Freiburg, Germany, 2004.

He, G., and R. Litterman. “The Intuition behind Black-Litterman Model Portfolios.” Working paper, Goldman Sachs & Co., 1999.

Lamm, R. McFall Jr. “Asymmetric Returns and Optimal Hedge Fund Portfolios.” *Journal of Alternative Investments*, 6, 2, 2003, pp. 9-21.

Lee, W. *Theory and Methodology of Tactical Asset Allocation*. New Hope, PA: Frank J. Fabozzi Associates, 2000.

Morton, D., E. Popova, and I. Popova. “Efficient Hedge Fund Construction under Downside Risk Measures.” Working paper, University of Texas, Austin, 2004.

Popova, I., D. Morton, and E. Popova. “Optimal Hedge Fund Allocation with Asymmetric Preferences and Distributions.” Technical report, University of Texas, Austin, 2003.

Schlottmann, F., and D. Seese. “Modern Heuristics for Finance Problems.” In S. Rachev, ed., *Handbook of Computational and Numerical Methods in Finance*. Boston: Birkhäuser, 2004.

Schlottmann, F., and D. Seese. “Finding Constrained Downside Risk-Return Efficient Credit Portfolio Structures Using Hybrid Multi-Objective Evolutionary Computation.” In G. Bol, G. Nakhaeizadeh, S.T. Rachev, T. Ridder, and K.-H. Vollmer, eds., *Credit Risk.* Heidelberg: Springer, 2003.

Terhaar, K., R. Staub, and B. Singer. “An Appropriate Policy Allocation for Alternative Investments.” *Journal of Portfolio Management*, 29, 3, 2003, pp. 101-111.

Weicker, K. *Evolutionäre Algorithmen*. Stuttgart: Teubner, 2002.