OPTIMAL HEDGE FUND ALLOCATION WITH ASYMMETRIC PREFERENCES AND DISTRIBUTIONS

Ivilina Popova[1], David Morton[2], Elmira Popova[3], Jot Yau[4]

[1]Corresponding author, Assistant Professor, Albers School of Business and Economics, Seattle University, 901-12th Ave., Seattle, WA 98122, Tel. (206)-296-5736, e-mail: popovai@seattleu.edu

[2]Associate Professor, Graduate Program of Operations Research, The University of Texas at Austin, 1 University Station C2200, Austin, TX 78712, Tel. (512)-471-4104, e-mail: morton@mail.utexas.edu

[3]Associate Professor, Graduate Program of Operations Research, The University of Texas at Austin, 1 University Station C2200, Austin, TX 78712, Tel. (512)-471-3078, e-mail: elmira@mail.utexas.edu

[4]Professor, Albers School of Business and Economics, Seattle University, 901-12th Ave., Seattle, WA 98122, Tel. (206)-296-5639, e-mail: jyau@seattleu.edu

1 Abstract

Hedge funds typically have non-normal return distributions marked by significant positive or negative skewness and high kurtosis. Mean-variance optimization models ignore these higher moments of the return distribution, and thus fail to convince investors who care about the unwanted skewness and kurtosis that hedge funds may work well in a portfolio. We use a new method which incorporates Monte Carlo simulation and optimization to solve for a variety of investment objectives and address the special issues of hedge fund allocation. We applied the new optimization model to examine the effects of semi-variance, conditional third and fourth moments on portfolio allocation with hedge funds. We show that conditional on the investor's objective, a substantial allocation to hedge funds is justified even with consideration for the highly unusual skewness and kurtosis.

2 Introduction

In the last decade, the investment styles collectively known as hedge funds or absolute return strategies have grown in number, assets and specialization. High-net-worth individuals as well as endowments and foundations discovered early the value of including hedge funds in their portfolios. There has been a shift in pension plan investment strategies towards alternative investments or hedge funds. For example, the Massachusetts Pension Reserve Investment Management in 2004 allocated 5% of its $35 billion plan into alternative investments. One reason for allocating a greater percentage of funds to alternative investments is the sharp correction in the equity markets between 2000 and 2003. There are several studies that address the issue of including hedge funds in standard institutional portfolios. For example, Kat[11], Till[22], Amenc, Giraud, Martellini and Vaissie [2], Coleman and Mansour[8], Amenc and Martellini [1], Brunel[7], Lamm[13], all have considered the problem of including hedge funds in a standard portfolio in mean-variance portfolio optimization.

The Markowitz mean-variance approach was introduced in the early 1950s as a rational tool to help guide the decision on optimal portfolio allocation. One of its basic assumptions is that the investor's objective is defined as a trade-off between risk and return. This methodology has been used to determine what is often considered the baseline neutral asset allocation: 60% stocks and 40% bonds. The portfolio represents the maximum Sharpe ratio portfolio in a mean-variance optimization where the input statistics come from the long run averages computed from the monthly returns for US equity (STOCKS), US long-term bonds (BONDS) and US 30 days T-bills (CASH) for the period say, from 1926 to 2000:

	STOCKS	BONDS	CASH
Mean	13%	6%	3.85%
Volatility	20.20%	8.70%	0.92%

With this example, we can see that if the investor has the traditional objective of seeking to maximize returns per unit of risk, the investor can use the mean-variance approach for portfolio allocation. The framework is less useful, however, if the return series is not normally distributed as in the case of hedge funds, or if the investor's preference is different from the mean-variance criterion.

Hedge funds typically have non-normal return distributions marked by significant positive or negative skewness and high kurtosis, see Amin and Kat [12] or Lo [14]. Mean-variance models ignore these higher moments of the return distribution, and thus fail to convince investors who care about the unwanted skewness and kurtosis how well hedge funds work in a portfolio. In simple terms, mean-variance models penalize funds that occasionally surprise on the upside while they also underestimate the risk of the funds that have asymmetric downside risk. While it is largely recognized that hedge fund returns are not normally distributed, many funds still use the mean-variance optimization framework. Amenc, Giraud, Martellini and Vaissie [2] report that only 2% of the European multi-managers pay attention to skewness and kurtosis; only 13% of European multi-managers have integrated an extreme risk measure and scenarios on extreme market conditions. However, 84% of multi-manager funds consider that volatility is of major concern to their clients and 82% consider Sharpe ratio as an important indicator.

The return distributions of traditional asset classes, e.g., stocks and bonds, have been studied extensively . Figures 1 and 2 indicate that the traditional asset return series are close to normal, and thus they can be defined well simply by their mean and standard deviation.

Insert Figure 1 Here

Insert Figure 2 Here

In contrast, hedge fund performance is not so easily described. The monthly return series for hedge funds demonstrate skewness and kurtosis. Two examples are shown in Figures 3 and 4 for large hedge funds (assets over $1 billion with history of greater than five years) .

Insert Figure 3 Here

Insert Figure 4 Here

For hedge fund indices, the measurement of their performance present complications beyond the non-normal distribution as for individual hedge funds. The source of other problems stems from the lack of consistency and rigor in the hedge fund indices. First, returns are likely to be biased upward as a result of "instant histories", survivor, and self-selection biases inherent in the hedge fund indices. For many indices, the composition of the index is solely dependent on which funds choose to have their data listed in a database. When a fund chooses to list, it typically brings 18 to 24 months of instant history into the index, and most likely a good history. On the other hand, if the performance of the fund has not been good, the fund may choose to stop reporting to the database. Should these funds be dropped from the database, the index return would be upwardly biased since the surviving funds in the index are the better performed funds. On the other hand, some of the bigger and successful funds may choose not to list for they no longer want to accept new investors and their information is not included in the index. Thus, the performance of the index may not accurately be reflected in the index due to these survivor and self-selection biases. Survivor bias is estimated to add 1.5% to 3.0% to annual returns, see Amin and Kat [12]. Moreover, if the better performed funds are closed to new investors but still included in the index, the index itself may not be truly investable. As such, the potential benefit of hedge funds as represented by hedge fund indices in asset allocation can be sheer inaccurate.

While various factors increase observed returns, volatilities are likely to be reported lower than the actual as a result of autocorrelation. The combination of higher observed returns and lower observed standard deviations is the primary driver of the optimal solution in mean-variance optimization which results in a higher Sharpe ratio. Adding in the skewness and kurtosis characteristics, mean-variance models can be expected to over-allocate to hedge funds. Regarding investors criteria for optimal portfolio, several alternative objectives have been proposed, including the minimax ratio, the mean absolute deviation ratio, the Sortino ratio, (see Sharpe [20], Young [24], Sortino [21] and Uryasev [23]). Some authors have considered higher moments of the return distribution in the portfolio optimization setup. For instance, Lamm[13] uses the Cornish-Fisher approximation in a value-at-risk context where the objective function is modified value-at-risk that incorporates asymmetry explicitly. It penalizes portfolios whose return distributions have negative skewness and excess kurtosis.

Although the previous studies have attempted to adjust for the impact of higher moments, different alternative objectives, stale pricing, and interval of rebalancing on the optimal allocation to hedge funds, many restrictive assumptions used in the mean-variance optimization remain, albeit unnecessary. With the advent in optimization methodologies used in the fields of transportation, communications and manufacturing in the last decade, combined with Monte Carlo simulation techniques and the rapid advances in computing power in recent years have created new powerful means to help make optimal decisions under uncertainty (an excellent source in this area is Pflug [18]). In this paper, we use a new methodology that combines simulation and optimization to solve for a variety of investment objectives that are of interest to portfolio managers. The simulation techniques discussed in section 4 that provide input -simulated paths from the hedge fund returns distribution- to the model, allow the optimization model to capture the information contained in the skewness and kurtosis of the non-normal return distributions. Thus, in contrast to the approaches used in previous studies we use not only the sample moments but the whole distribution. This flexibility allows us to better address the special issues of hedge fund allocation within a standard portfolio. We demonstrate that allocation to hedge funds within a traditional institutional portfolio may be optimal and rational even when the individual hedge funds have undesirable skewness and kurtosis. In addition, we show that the mean-variance objective under the assumption of normal return distribution is a special case of the new methodology.

The rest of the paper is organized as follows. We present the optimization model for utility maximization in section 3. We discuss the methodology in section 4. We demonstrate the convergence of the approximate solution in section 5 and the impact of different preferences on the solution in section 6. Empirical results and discussion of the inclusion of hedge funds in institutional portfolios are presented in section 7 and 8. We conclude the paper in section 9.

3 Optimization Model for Utility Maximization

The most general preference function is often expressed as maximizing expected utility, that is, maximizing the value to the investor. While the traditional objective of maximizing risk-adjusted returns is useful in many contexts, investors often have other preferences. For instance, an institution may want to increase the probability that it can achieve a certain benchmark return in order to meet the liabilities of its pension plan or, it may want to minimize the need to make additional contributions to a plan.

Problems that find optimal allocation with respect to a set of constraints by taking into account the uncertainty of the underlying asset returns are known as stochastic programs. We formally introduce such a problem as follows.

Maximizing expected utility is typically expressed as: $U(X, \bar{\omega})$. An investor in a world of uncertain returns can choose from a universe of assets $X = \{X_1, X_2, \ldots, X_n\}$. The assets are assumed to have random returns. The asset allocation problem is to determine the proportion of funds to be invested in each asset. Portfolio weights are denoted by $\bar{\omega} = \{\bar{\omega_1}, \bar{\omega_2}, \ldots, \bar{\omega_n}\}$, $\bar{\omega} \in W$ which might be subjected to linear, quadratic, or other types of constraints, such as cardinality. These might arise when the investor would like to choose a subset of the original universe but still maintain the initial constraints.

Putting the components together, the generic problem (optimizing over one period) can be expressed as follows:

$$\begin{aligned} z^* \quad &= \quad \max_{\bar{\omega} \in W} E_X \left[U \left(X, \bar{\omega} \right) \right] \\ \text{s.t.} \quad & \quad \begin{array}{l} \textit{Risk Constraint} \\ \textit{Additional Constraints} \end{array} \end{aligned}$$

This generic formula can be further refined to express more specific objectives or preferences. For example, if we are interested in maximizing the probability of outperforming a benchmark, r, the problem can be expressed in terms of maximizing expected utility as follows:

$$z^* = \max_{\bar{\omega} \in W} E_X \left[I \left(\bar{\omega}^T X > r \right) \right]$$

where I denotes the indicator function which takes a value of 1 if the event in the brackets occurs or 0 otherwise.

The formula can also be expressed easily in terms of maximizing probability:

$$z^* = \max_{\bar{\omega} \in W} P \left(\bar{\omega}^T X > r \right) \tag{1}$$

A different objective might be to minimize expected shortfall, that is to minimize the average size of any shortfall in those instances when the portfolio fails to outperform the benchmark r. In utility terms, this is expressed as follows:

$$z^* = \min_{\bar{\omega} \in W} E_X \left[r - \bar{\omega}^T X \right]^+ \tag{2}$$

Compound objectives can also be established for the model. Combining (1) and (2) above, the investor's objective might be to maximize the probability of outperforming a benchmark , and at the same time wanting to limit the average amount by which she would underperform benchmark . This combination objective can be formulated as:

$$z^* = \max_{\bar{\omega} \in W} P \left(\bar{\omega}^T X > r_1 \right) - \lambda E_X \left[r_2 - \bar{\omega}^T X \right]^+ \tag{3}$$

Formula (3) represents the objective of an investor, who has two benchmarks,

1. r_1 is the benchmark she wants to outperform, for example, the return of the S&P500

2. r_2 is the benchmark she does not want to fall below, for example zero or the risk-free rate.

The value of λ represents the priority or the weight placed on the second objective. By varying that weight, we can obtain a frontier in much the same way as one constructs the mean-variance efficient frontier by varying the desired expected return or the risk aversion parameter of the investor. However, in this frontier the axes are probability of outperformance and size of the average conditional shortfall.

4 Methodology

Problems (1), (2) and (3) are one-period stochastic programming problems. They can be solved exactly if X has a small number of scenarios, or if X is from a normal distribution in which case, the problems can be simplified to deterministic nonlinear programming problems.

If X is from a non-normal distribution and the number of scenarios is either large or continuous, then it is usually not possible to simplify the stochastic problem and solve it exactly. The only available approach is to solve the problem approximately. Asset allocation with hedge funds present, which generally involves both non-normal distributions and a continuous range of potential returns, is such a stochastic programming problem. Thus, we need to solve the problem approximately.

One standard approach for approximately solving stochastic programming problems is to use Monte Carlo sampling procedures. This method uses computer programs to generate randomly a number of observations, N, of some variable, X, from some distribution of potential values, and then solve the approximate stochastic program using that information. This solution with the sampled data represents one possible realization of the problem. By doing this many times, a series of approximate solutions can be constructed. These solutions can be used to construct an upper bound of the objective function. It will be an "upper bound" since we have perfect information about future returns and create the optimal portfolio weightings with that perfect information. That is, we see the future" and cheat" when making a decision regarding our portfolio allocation. This upper bound is known as a "wait-and-see" bound.

We can also compute a "lower bound" of the objective function. For any portfolio weights, $\bar{\omega}$, generate a large number of observations of the monthly returns for each asset class, and estimate the value for the objective function. This bound is known as "here-and-now" bound, since we start with a feasible weighting and then simulate what may take place with that weighting.

By employing these two bounding ideas together, a confidence interval on the optimality gap (the difference between the optimal value of the objective function, z^*, and the value for a given solution vector, $\bar{\omega}$) can be constructed. This confidence interval represents a statement of the quality of the solution $\bar{\omega}$. This methodology was originally developed in Mak, Morton and Wood [15], where they applied it to a variety of operations research problems. For a detailed description of the approximating model and new approaches for solving it see Popova, Morton and Popova [19].

5 Convergence of the Approximate Solution

Stochastic programs address decision-making with uncertainty around the variables. As we discussed, some of them can be solved only approximately. To test such methodologies, we run a large number of scenarios for a case with specified distributions, for example, a normal distribution. The results of the sequence of tests are then compared with an exact solution found using another proven methodologies. Mak, Morton and Wood (MMW) [15] demonstrated that, as the number of scenarios increase the lower and upper bounds on the objective function should converge toward the true optimal value.

To demonstrate the methodology and the convergence of the approximate solution using the MMW methodology, we look at a universe of 6 asset classes. We assume that the underlying true distribution is normal $N(EX, \Sigma)$ and we know the expected return, EX (the vector with the expected values of the asset classes) and the covariance matrix, Σ . For this particular example we use the historical returns and covariance matrix for 6 asset classes:

$$
\begin{aligned}
EX &= (15.8\%, 8.3\%, 6.5\%, 44.3\%, 9.1\%, 15.2\%) \\
\Sigma &= \begin{bmatrix}
1.91\% & 0.21\% & 0.15\% & 1.68\% & 1.33\% & 0.26\% \\
0.21\% & 0.16\% & 0.12\% & 0.03\% & 0.09\% & 0.04\% \\
0.15\% & 0.12\% & 0.40\% & -0.06\% & 0.01\% & -0.01\% \\
1.68\% & 0.03\% & -0.06\% & 6.40\% & 2.04\% & 0.55\% \\
1.33\% & 0.09\% & 0.01\% & 2.04\% & 2.29\% & 0.28\% \\
0.26\% & 0.04\% & -0.01\% & 0.55\% & 0.28\% & 0.21
\end{bmatrix}
\end{aligned}
$$

When the returns come from a multivariate normal distribution with known parameters and the objective function is to maximize probability of outperforming a fixed benchmark, the optimal portfolio is the one that maximizes the information ratio. For our example, we make typical assumptions: the fixed benchmark is 10%, short selling is not allowed and the sum of the weights must equal one. Solving the above quadratic programming problem, the resulting optimal portfolio is:

ASSET CLASS	WEIGHT
CLASS 1	0%
CLASS 2	0%
CLASS 3	0%
CLASS 4	22.62%
CLASS 5	0%
CLASS 6	77.38%
PROBABILITY OF OUTPERFORMING 10%	92.88%

Since we are assuming normality, the probability of outperformance also can be computed "exactly", see Popova, Morton and Popova [19] as follows:

$$
\begin{aligned}
P\left(\bar{\omega}^T X > 10\%\right) &= 1 - \Phi\left[\frac{\left(10\% - \bar{\omega}^T EX\right)}{\sqrt{\bar{\omega}^T \Sigma \bar{\omega}}}\right] \\
&= 92.88\%
\end{aligned}
$$

where, Φ denotes the standard normal cumulative function.

We then use our methodology to solve two sequences of approximate stochastic programs using the same data to establish the upper and lower bounds. We test with 100 scenarios, 250 scenarios, 500 scenarios and then 1,000 scenarios. Results are presented in the table below.

	True Probability	Lower Bound	Upper Bound	99% CI on the Optimality Gap
MAX INFORMATION RATIO	92.88%			
100 SCENARIOS		92.20%	94.87%	[0.0%; 2.9%]
250 SCENARIOS		92.52%	94.00%	[0.0%; 1.7%]
500 SCENARIOS		92.55%	93.45%	[0.0%; 1.1%]
1000 SCENARIOS		92.62%	93.16%	[0.0%; 0.8%]

The lower bound and the upper bound both converge on the "true" probability as the number of scenario increase. In addition, the last column shows that the size of the 99% confidence interval on the optimality gap decreases from almost 3% to less than 1% as the number of scenario increase.

Since all the weights are bounded between 0 and 1, one can expect to see convergence not only in the value function, that is, the probability of outperformance, but also in the solution space. The following table illustrates this convergence behavior. As the number of scenario increase, the optimal allocation converges toward the "true" solution.

ASSET CLASS	Analytical solution	100	250	500	1000
CLASS 1	0.00%	2.41%	0.98%	0.81%	0.55%
CLASS 2	0.00%	2.07%	2.17%	1.20%	0.61%
CLASS 3	0.00%	2.23%	0.80%	0.62%	0.37%
CLASS 4	22.70%	25.31%	23.39%	25.08%	22.12%
CLASS 5	0.00%	0.51%	0.04%	0.06%	0.00%
CLASS 6	77.30%	67.47%	72.63%	72.24%	76.36%

The above demonstration shows that this approach produces very close to the same results as the "true" optimal results under known cases.

6 Effect of Preferences

Next, we show the impact of having different preferences and compare the optimal results obtained by using our methodology to those obtained from the standard mean-variance approach.

Illustration

Suppose we have two assets whose distributions have the same mean and variance but different higher moments. Let the monthly mean be 1% and the monthly standard deviation 4% (numbers are consistent with the performance of S&P 500 over 1990-2000). The next table shows the descriptive statistics for the two assets. Figure (5) plots their densities. Also assume that the two assets are not highly correlated. (Asset 1 is S&P 500 and Asset 2 is generated to have the first 2 moments as Asset 1.)

	ASSET 1	ASSET 2
MEAN	1%	1%
ST. DEVIATION	4%	4%
SKEWNESS	-0.44	-7.71
KURTOSIS	1.15	74.74
MIN	-14%	-40%
MAX	11%	4%
CORRELATION	0.05	

Insert Figure 5 Here

Since the correlation coefficient is so small, from a mean-variance optimization objective point of view, the efficient frontier will be approximately 50% Asset 1 and 50% Asset 2. Suppose now that the investor's objective is to maximize the probability of having positive returns, and in the same time she is very averse towards having returns below -5%. If she chooses to use the standard mean-variance approach, independent of her benchmark objectives, her optimal allocation will be always 50-50. Just by looking at the assets densities, it is clear that depending on the objective, one should hold either more of Asset 1 or Asset 2. If, for example, the objective is to have positive returns, then Asset 2 is the natural choice. On the other hand, if the objective is to be protected from big negative returns, the natural choice is Asset 1 (Asset 2 can produce large negative returns with a very small probability). The question is: how to make the optimal decision. Our approach can identify, depending on the objective, the optimal combination of assets. Using the methodology described in the previous section, where $r_1 = 0\%$ and $r_2 = -5\%$, by varying the values for λ, we can produce a frontier. Figure 6 presents the frontier.

Insert Figure 6 Here

Note that the lowest point on the frontier corresponds to a portfolio that is very averse towards falling below -5%, hence the optimal allocation puts more weight on Asset 1, i.e. avoiding the possibility of a large negative return that could occur when holding Asset 2. The highest point on the frontier corresponds to a portfolio that is very aggressive in achieving positive returns, hence higher allocation in Asset 2. The mean-variance efficient portfolio 50-50 is also on the frontier. It is clear that the mean-variance allocation is a special case of our approach. We are enhancing it in two directions by considering asymmetric distribution and asymmetric utility function.

7 Adding Hedge Funds to a Standard Portfolio

In this section, we consider the optimal asset allocation of hedge funds to a traditional portfolio. We use our methodology to capture the full value of the information contained in both tails of the non-normal distribution of hedge fund returns.

For US Equity and US Fixed Income we again use the long run statistics that generated the 60%-40% portfolio discussed earlier. For hedge funds, we use monthly aggregate index returns for the period 1995-2000. The index selected, the EACM 100, represents a broadly diversified portfolio of hedge funds. This particular index was selected because its returns are significantly less than those of the "database" indices. This index is based on a relatively stable set of established funds. In contrast, indices based on databases have higher returns because new funds are frequently added with months or years of instant history and funds are frequently dropped when they stop reporting to the database or close to new investors. The result is database indices tend to carry instant history and survivor bias that is estimated at 1.5% to 3.0%. The historical characteristics for the 3 asset classes are:

	STOCKS	BONDS	HEDGE FUND
Mean	13.0%	6.0%	12.7%
Volatility	20.2%	8.7%	5.0%
Sharpe(at 5% risk-free rate)	0.4	0.1	1.5

CORRELATIONS	STOCKS	BONDS	HEDGE FUND
STOCKS	1		
BONDS	0.3	1	
HEDGE FUND	0.5	0.08	1

Before applying our methodology, consider a simple mean-variance optimization with the three asset classes. As expected, the mean-variance is not helpful with the inclusion of hedge funds. The high Sharpe ratio for the hedge funds causes them to dominate the allocation.

STOCKS	BONDS	HEDGE FUND
0%	1.4%	98.6%

While mathematically correct, a 99% allocation to hedge funds would not be considered rational. In fact, by only considering the mean and variance for each asset class, we are ignoring much valuable information contained in the series of monthly returns.

As we saw earlier, the normality assumption is reasonable for the US equity and fixed income returns but it is far from appropriate for the hedge fund returns. In addition, hedge fund index data presents substantial problems with regard to survivor bias (Fung and Hsieh [10]) that inflates returns and autocorrelations that reduce observed volatility (Brooks and Kat [5]). Therefore, we systematically build a less optimistic return sample for hedge funds that incorporates the higher moments for the hedge fund asset class and uses more reasonable parameters for mean return and volatility. This new data set incorporates a healthy pessimism to challenge the highly optimistic allocations seen in the mean-variance analysis.

For our distribution, we assume a mean return for hedge funds of 9.1%. This adjusts for the full 3% survivor bias found in hedge fund index returns by Fung and Hsieh plus a significant downward revision to compensate for the strong bull market and associated capital market activity during the late 1990s. While hedge funds tend to have a low beta with the equity markets, they do tend to perform better when there is significant capital market activity in the form of merger activity, IPO issuance, restructuring, etc.

We also target a standard deviation of 9.0%. Brooks and Kat [5] suggest eliminating autocorrelation effects would increase volatility by up to 60%. Our own Bayesian estimate of risk suggests an even higher standard deviation which is closer to 9.0%. Looking at historical data of various indices, we targeted skewness of -1 and kurtosis of 6. We use a mixture of two normal distributions to create a distribution that has the above four target moments. Solving a nonlinear system of equations gives the values for the parameters of the mixture. The corresponding two normal distributions for the hedge fund and the resulting density plots for the three asset classes are:

Insert Figure 7 Here

Insert Figure 8 Here

Figure 7 shows the two normal distributions. Distribution 1 represents "steady state" returns when trading strategies are basically working as intended. It is characterized by a normal distribution with mean 14.9% and volatility of 6.2%, metrics similar to the HFR index data. Distribution 2 represents event risks or shocks to the steady state. It is a normal distribution with mean -15.5% and volatility of 14.2%. The probability of choosing Distribution 1 is 81%.

Figure 8 shows the combined density for the simulated hedge fund data. It is created by assigning an 81% probability to returns coming from Distribution 1 and 19% from Distribution 2. This could be thought of as if one month in five, we see a strategy failure, e.g., deal breaks for merger arbitrage, equity market corrections for long-short strategies, or extraordinary interest rate volatility.

Figure 9 shows the densities for the three asset classes. We use normal distributions for the US Equity and Long-Term Bonds and the simulated distribution for the Hedge Fund.

Insert Figure 9 Here

With the return distributions of these three asset classes, we now consider incorporating hedge funds into a traditional portfolio. For our case here, we look at an investor who does not only seek to exceed a 12% benchmark return but also to avoid any losses for the portfolio overall. In our model, the investor's true objective is to maximize the probability of outperforming a benchmark, $r_1 = 12\%$ and her risk measure is the expected shortfall with respect to a benchmark,

$r_2 = 0\%$. In our formulation, the investor wants to solve the following problem:

$$\max_{\bar{\omega}\in W} P\left(\bar{\omega}^T X > r\right) - \lambda E_X \left[r - \bar{\omega}^T X\right]^+$$
$$r_1 = 12\%$$
$$r_2 = 0\%$$

By varying the value of λ, we can increase or decrease the investor's aversion towards risk, which in this case is expressed in terms of having negative returns. If λ is made a very large number, the investor is only interested in avoiding losses (by analogy, the minimum variance portfolio in mean-variance optimizations). In this case, our methodology solves the optimal risk-averse portfolio to be:

C: 10% Stocks, 70% Bonds and 20% Hedge Fund.

By increasing λ, we plot an efficient frontier with a vertical axis showing the probability of outperforming r_1 and a horizontal axis of the risk term, in this case, the expected shortfall or the first partial moment in Figure 10. The risk term is scaled such that the largest value is equal to one.

Insert Figure 10 Here

The portfolio that is represented in red is 60% Stocks and 40% Bonds. It is clear that one can improve portfolio efficiency in two directions. First, the efficient portfolio that has the same probability of outperformance as the 60%-40% portfolio, but with less risk is portfolio:

A: 30% Stocks, 28% Bonds, 42% Hedge Fund.

Portfolio A preserves the probability of outperformance, but reduces the risk measure by 50%.

Second, the portfolio with the same risk but higher probability of outperformance is portfolio:

B: 27% Stocks, 1% Bonds and 72% Hedge Fund.

8 Who will NOT invest in Hedge Funds?

Next, let's look at risk more broadly than simply aversion to losses, which is the first partial moment. Let's consider a more general risk term that is the sum of the first K partial moments. In other words, assume that the investor is worried not only about the expected loss, but also about semivariance (placing a high penalty on any losses), conditional third moment (related to skewness), and conditional fourth moment (related to kurtosis - fear of extreme negative event).

$$\max_{\bar{\omega}\in W} P\left(\bar{\omega}^T X > r\right) - \lambda E_X\left(\left[r - \bar{\omega}^T X\right]^+\right)^K$$

The first case, that is for $K = 1$, was presented in the previous section. Using simulations, we can approximately identify the efficient frontiers for $K \geq 2$. The next three figures show the frontiers for K equal to 2, 3 and 4. From the figures you can see that as K increase the investor becomes more concerned about events far out in the left tail. As a consequence the allocation to hedge funds decreases. This is visually depicted by the position of the 60%-40% portfolio. As K increase, 60%-40% portfolio moves closer to the efficient frontier. One can speculate that there will exist a value of K, probably very large, for which the 60%-40% portfolio will become efficient.

Insert Figure 11 Here

Insert Figure 12 Here

Insert Figure 13 Here

If we continue to expand the possible definitions of risk, it becomes clear that only an investor who seeks to avoid any extreme event with hedge funds will prefer the 60%-40% portfolio with no hedge funds to any portfolio with hedge funds. However, such an investor is willing to accept the variance inherent in the 60% allocation to stocks. Further, by making this choice she will reduce her chances of beating her benchmark of 12% per year.

The following tables show the allocation in the three asset classes for different risk measures. The portfolios in the first table are chosen from the set of efficient portfolios that have the same probability of outperforming 12% as the standard 60%-40% portfolio. We can see that moving from the stocks to hedge funds reduces volatility substantially. As the risk measure changes ($K > 1$), the allocation to hedge funds decreases. However, in the case of $K = 4$, the allocation is 21%, still a substantial allocation. Thus, we conclude that as long as the investor is comfortable with being protected against the first four partial moments, an allocation to hedge funds is efficient. Of course, if the investor wants sure protection against rare events, the hedge fund asset class will not be an appropriate investment.

Risk = $E_X\left(\left[r - \bar{\omega}^T X\right]^+\right)^K$	Bonds	Stocks	Hedge Fund	% Reduction in Risk
$K = 1$	28%	30%	42%	58.1%
$K = 2$	31%	34%	35%	52.1%
$K = 3$	33%	39%	28%	46.1%
$K = 4$	38%	41%	21%	34.7%

In summary, our analysis shows that by transferring half of the allocation from Stocks to Hedge Fund, risk decreases relative to the risk of the 60%-40% portfolio.

The portfolios in the second table are chosen from the set of efficient portfolios that have the same value as the risk measure for the standard 60%-40% portfolio. We see similar pattern in the allocation of hedge funds as K increase. The starting allocation in hedge fund $K = 1$ is much greater that the starting allocation in the previous table. This is due to the fact that now we are considering portfolios that will bring us closer to the goal of outperforming 12% per year. Hence, our risk aversion is low. Notice that this time the transformation of weight is from the fixed income to the hedge fund. The last row of the table, $K = 4$, shows that the allocation in hedge fund is 38%. This is a point on the efficient frontier that corresponds to a smaller value of the risk aversion parameter compared to the value of the risk aversion parameter for the last portfolio in the previous table. At that point, the allocation to hedge funds is 21%. The increase in the probability of outperformance is mainly due to the increase in the allocation to hedge funds.

Risk = $E_X\left(\left[r-\bar{\omega}^T X\right]^+\right)^K$	Bonds	Stocks	Hedge Fund	% Increase in Probability
$K = 1$	1%	27%	72%	9.0%
$K = 2$	8%	35%	57%	6.0%
$K = 3$	10%	43%	47%	5.0%
$K = 4$	14%	48%	38%	4.0%

In summary, by transferring almost half of the allocation from fixed income to hedge funds, the probability of outperforming a 12% benchmark compared to the probability of the 60%-40% portfolio increase.

9 Conclusion

The return distributions of hedge funds and their indices are marked by their non-normal properties. This presents difficulties to traditional mean-variance models for asset allocation. Such models will tend to be overly optimistic in their allocations to hedge funds. In this paper, we introduce a more effective means of dealing with hedge fund portfolios. Using procedures from stochastic programming, this model is able to recognize and use the information embedded in the unusual return distributions of hedge funds.

We applied the model to hedge funds to examine the effects of semi-variance, conditional third and fourth moments on portfolio allocations to hedge funds. The results indicate a substantial allocation to hedge funds is justified even with consideration for the highly unusual kurtosis and skewness.

We will extend our analysis by using individual hedge funds in the near future. The analysis will be complete by the end of this summer.

References

[1] Amenc N. and L. Martellini, "Portfolio Optimization and Hedge Fund Style Allocation Decisions, The Journal of Alternative Investments, 2002, Vol. 5., No. 2, pp. 7-20.

[2] Amenc N., Giraud J., Martellini L. and M. Vaissie, "Taking a close look at the European fund of Hedge Funds Industry: Comparing and Contrasting Industry Practices and Academic Recommendation," Journal of Alternative Investments, Winter 2004, 59-69 .

[3] Amin, G.S., and Kat, H., "Hedge Fund Performance 1990-2000: Do the Money Machines Really Add Value? Journal of Financial and Quantitative Analysis, 2003, Vol. 38, No.2, pp. 251-274.

[4] Biglova A., Ortobelli S., Rachev S. and S. Stoyanov, "Different Approaches to Risk Estimation in Portfolio Theory," The Journal of Portfolio Management, Fall 2004, 103-112.

[5] Brooks, C. and Kat, H., "The Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors," The Journal of Alternative Investments, Fall 2002, Vol. 5 Issue 2, p26, 19p.

[6] Browne, S., The Risk and Rewards of Minimizing Shortfall Probability, The Journal of Portfolio Management, Summer 1999, 76-85.

[7] Brunel J., "Revisiting the Role of Hedge Funds in Diversified Portfolios," The Journal of Wealth Management, Winter 2004, 35-48.

[8] Coleman M. and A. Mansour, "Real Estate in the Real World: Dealing with Non-Normality and Risk in an Asset Allocation Model," Journal of Real Estate Portfolio Management, 2005, Vol. 11, No. 1, 37-53.

[9] Cremers, J., Kritzman M. and S. Page, "Optimal Hedge Fund Allocation," Journal of Alternative Investments, Spring 2005, 70-81.

[10] Fung, W. and Hsieh, D., "Is Mean-Variance Analysis Applicable to Hedge Funds?" Working Paper, November 1997.

[11] Kat, H., "Integrating Hedge Funds into the Traditional Portfolio," The Journal of Wealth Management, Spring 2005, 51-57.

[12] Kat, Harry M. and Amin, Gaurav S., "Welcome to the Dark Side: Hedge Fund Attrition and Survivorship Bias over the Period 1994-2001," December 11, 2001, http://ssrn.com/abstract=293828.

[13] R. McFall Lamm, Jr, "The answer to your dreams? Investment implications of positive asymmetry in CTA returns," Journal of Alternative Investments, Spring 2005, 22-32.

[14] Lo, A.W. Risk Management for Hedge Funds: Introduction and Overview. Financial Analysts Journal, 2001, Vol. 57, No. 6, pp.16-33.

[15] Mak, W. K., Morton, D. and K. Wood, "Monte Carlo Bounding Techniques for Determining Solution Quality in Stochastic Programs," Operations Research Letters, 24, 1999, 47-56.

[16] Markowitz, H., Portfolio Selection, John Wiley & Sons, New York, NY, 1959.

[17] Martin, George, "Making Sense of Hedge Fund Returns: What Matters and What Doesn't," Derivatives Strategy, 2001.

[18] Pflug, G., Optimization of Stochastic Models, The Interface Between Simulation and Optimization, Kluwer Academic Publishers, 1996.

[19] Popova, I., Morton, D. and Popova, E., "Efficient Fund of Hedge Fund Construction Under Downside Risk Measures," Journal of Banking and Finance, forthcoming.

[20] Sharpe, W. F. "The Sharpe Ratio", The Journal of Portfolio Management, Fall 1994, pp. 45-58.

[21] Sortino, F. A. "Upside-Potential Ratios Vary by Investment Style", Pensions and Investments, 28 (2000), pp. 30-35.

[22] Till, H., "On the Role of Hedge Funds in Institutional Portfolios," The Journal of Alternative Investments, Spring 2005, 77-89.

[23] Uryasev, S. P. Probabilistic Constrained Optimization Methodology and Applications, Kluwer Academic Publishers, 2000.

[24] Young, M. R. "A MiniMax Portfolio Selection Rule with Linear Programming Solution", Management Science, 44(1998), pp. 673-683.

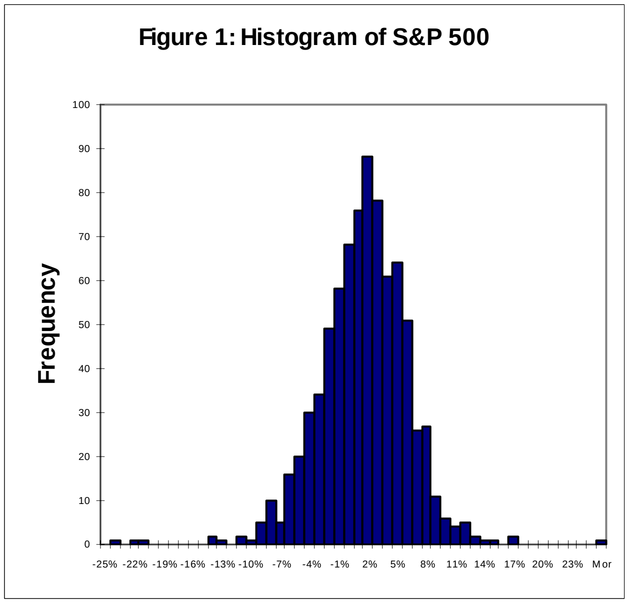

Figure 1 shows the histogram for S&P 500 monthly returns for the period 1996-2001.

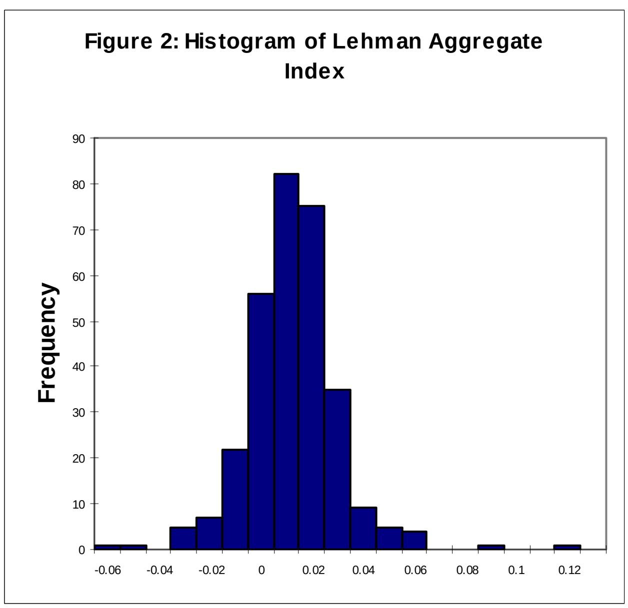

Figure 2 shows the histogram for the Lehman Aggregate Bond Index monthly returns for the period 1996-2001.

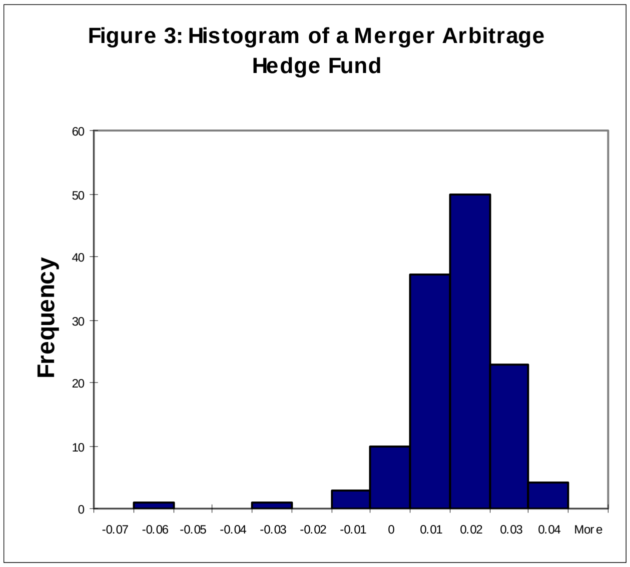

Figure 3 shows the histogram of Merger arbitrage hedge fund monthly returns with assets over $1 billion and with history of greater than five years.

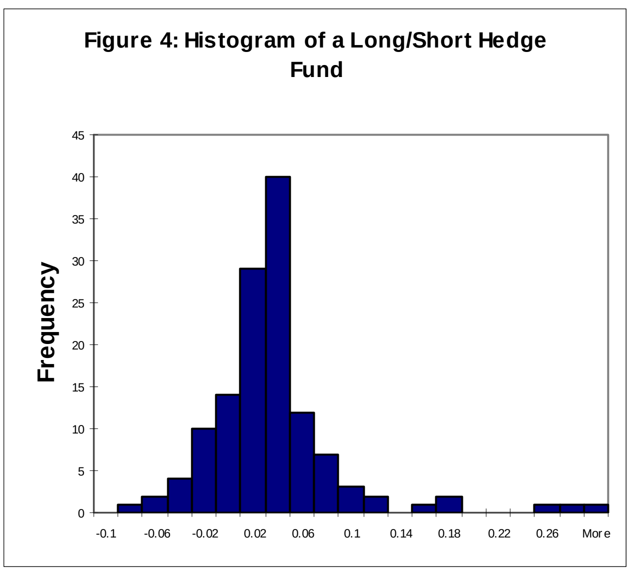

Figure 4 shows the histogram of Long/Short hedge fund monthly returns with assets over $1 billion and with history of greater than five years.

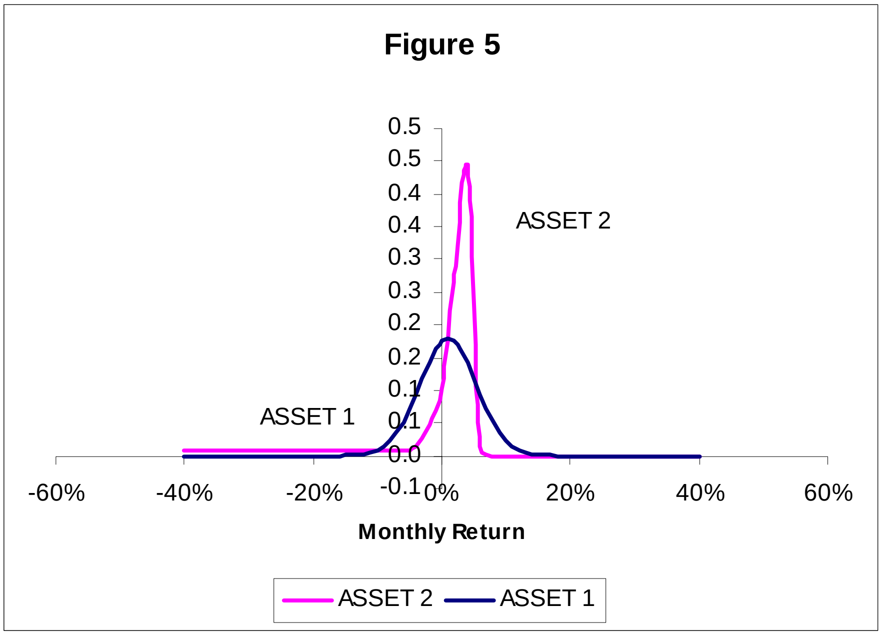

Figure 5 plots the densities for the two assets for the Example. Asset 1 has a monthly mean of 1% and a monthly standard deviation of 4%; Asset 2 has the same mean and standard deviation, additionally Asset 2 has a negative skewness and high kurtosis.

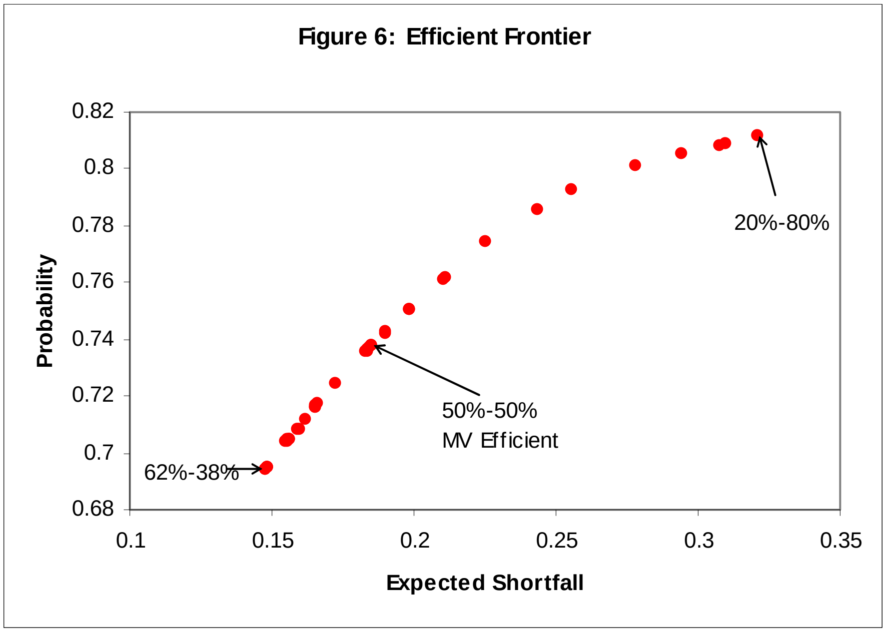

Figure 6 shows the efficient frontier produced for the Example. The x-axis is the Expected Shortfall with respect to (-5%), the y-axis is the Probability of outperforming 0%.

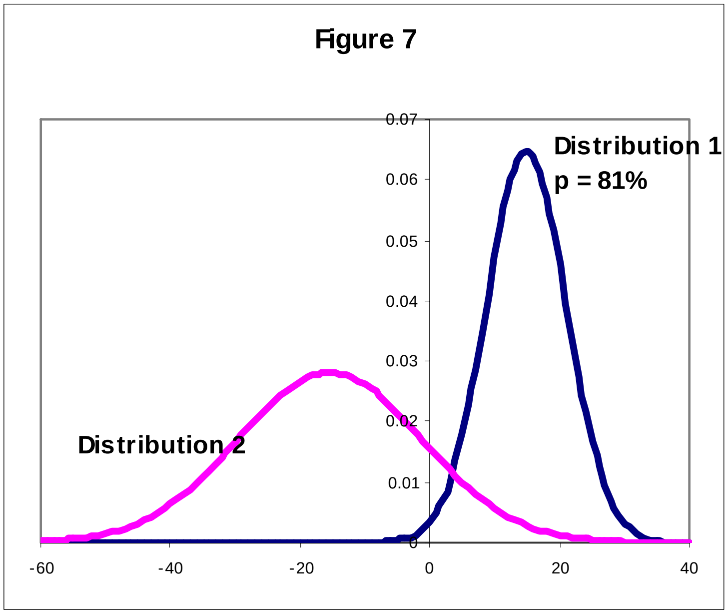

Distribution 1 represents "steady state" returns when trading strategies are basically working as intended. It is characterized by a normal distribution with mean 14.9% and volatility of 6.2%, metrics similar to the HFR index data. Distribution 2 represents event risks or shocks to the steady state. It is a normal distribution with mean -15.5% and volatility of 14.2%. The probability of choosing Distribution 1 is 81%.

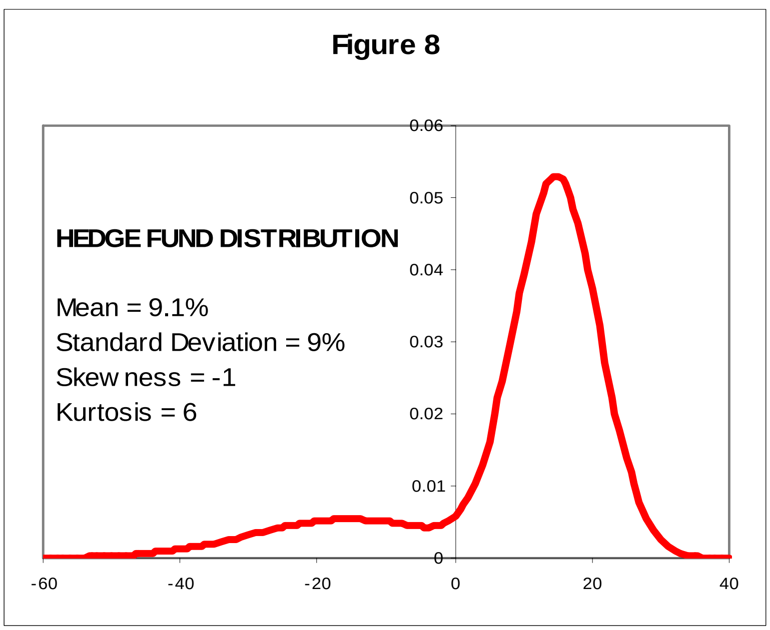

Figure 8 shows the combined density for the simulated hedge fund data. It is created by assigning an 81% probability to returns coming from Distribution 1 and 19% from Distribution 2.

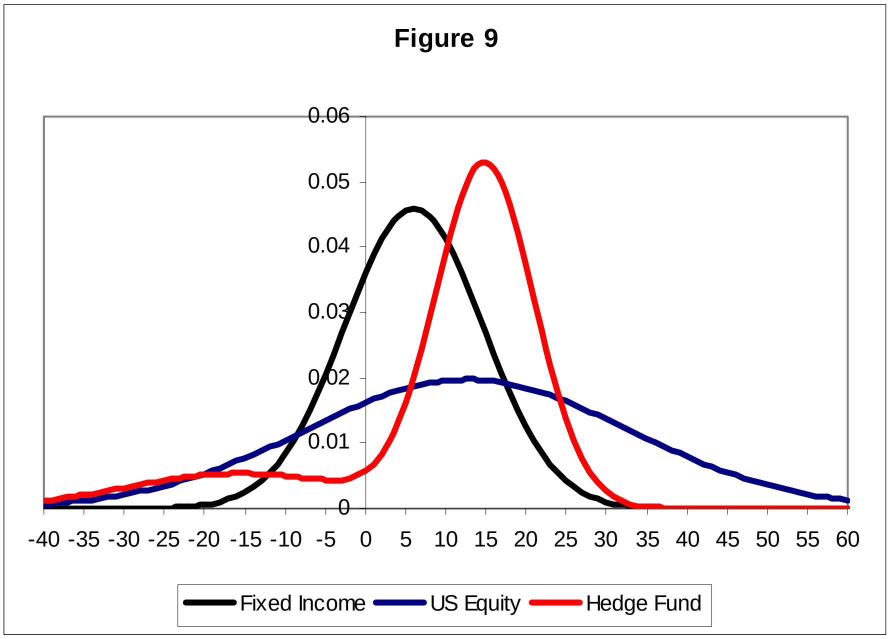

Figure 9 shows the densities for the three asset classes. We use normal distributions for the US Equity and Long-Term Bonds and the mixture of two Normal distributions for the Hedge Fund.

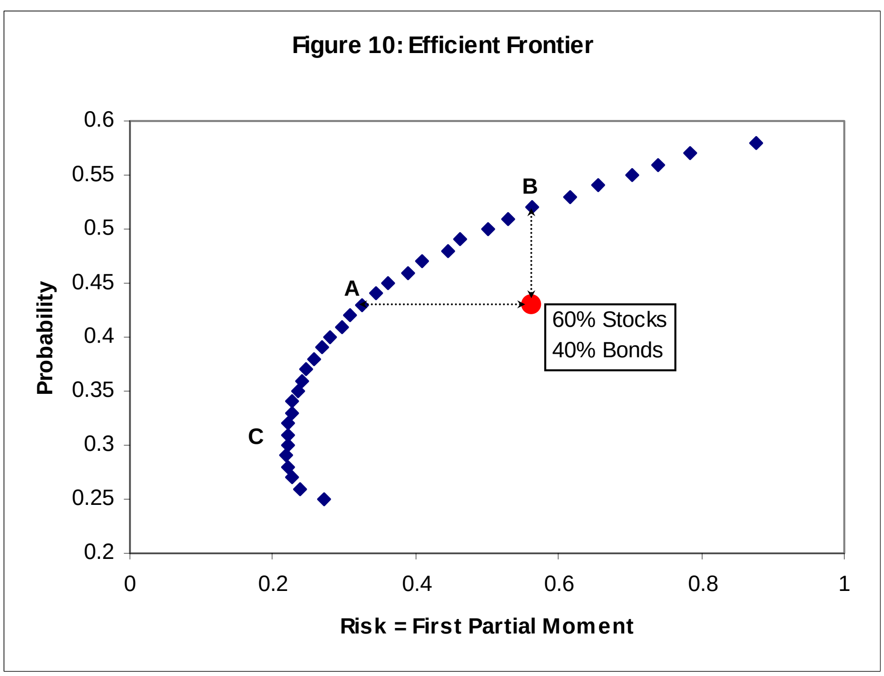

Figure 10 plots an efficient frontier with a vertical axis showing the probability of outperforming r_1 and a horizontal axis of the risk term, in this case, the expected shortfall or the first partial moment. The risk term is scaled such that the largest value is equal to one.

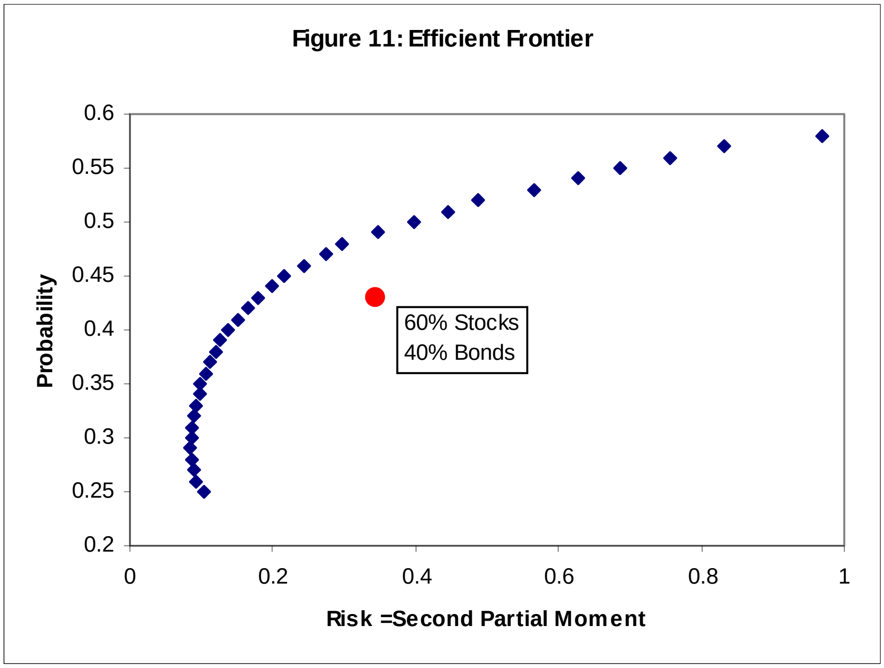

Figure 11 plots an efficient frontier with a vertical axis showing the probability of outperforming r_1 and a horizontal axis of the risk term, in this case, the second partial moment. The risk term is scaled such that the largest value is equal to one.

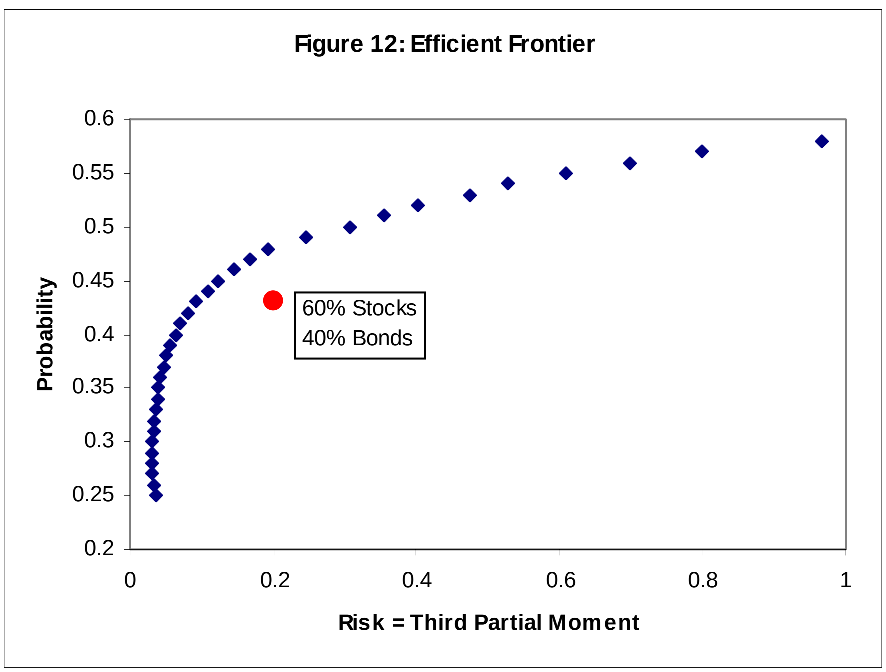

Figure 12 plots an efficient frontier with a vertical axis showing the probability of outperforming r_1 and a horizontal axis of the risk term, in this case, the third partial moment. The risk term is scaled such that the largest value is equal to one.

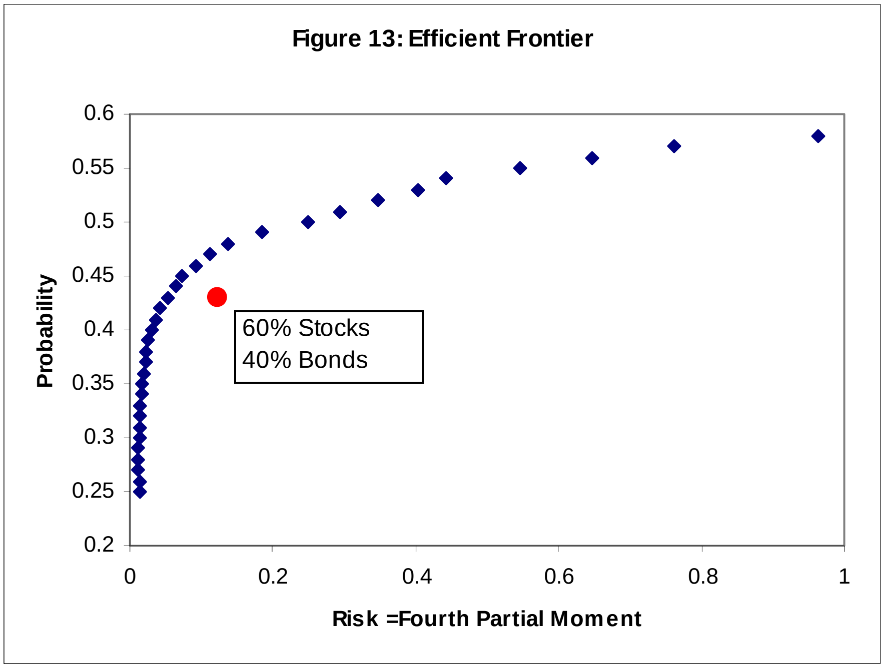

Figure 13 plots an efficient frontier with a vertical axis showing the probability of outperforming r_1 and a horizontal axis of the risk term, in this case, the fourth partial moment. The risk term is scaled such that the largest value is equal to one.