Estimating the true embedded risk management cost of total return strategies

Francesco Sandrini[1], David Edelman[2], and Dermot Ryan[1]

[1] Pioneer Investments Management,
1 George's Quay Dublin 2, Dublin
{Francesco.Sandrini, Dermot.Ryan}@pioneerinvestments.com
http://www.pioneerinvestments.com
[2] Michael Smurfit Business School,
UCD, Dublin
davide@ucd.it
http://www.ucd.ie

Abstract. We propose to measure the value added by periodic portfolio rebalancing in actively managed strategies. Using Monte-Carlo simulation and dynamic stochastic programming we simulate the pay-off of an actively managed strategy. We seek to replicate this pay-off using a static investment based on the same Monte-Carlo scenarios and the same investment timeframe, but including in the static portfolio some derivative strategies not available to the active manager. We contend that the allocation to the derivative strategies quantifies the value added by active management. We then test the sensitivity of the solution to various parameters of the problem.

Keywords. Stochastic optimization, option hedging, simulated annealing.

M.S.C.classification. 47N10, 91G20, 91G60.

J.E.L. classification. C61, C63, G11.

1 Introduction

The goal of this paper is to simulate an actively managed total return portfolio as a pseudo ALM problem of managing a portfolio of assets against a simulated LIBOR rate and to try to quantify the value added by the active manager under different circumstances. We do this by attempting to replicate the managers final payoff using an optimal static portfolio including options.

Our approach to ALM problems is based on multi-stage stochastic optimization; see for example Mulvey *et al.* [11] for an introduction to ALM problems and Dempster et al. 2002 [5] for a general and complete description of the stochastic optimization approach. See Grinold and Kahn [6] for a discussion of active portfolio management with risk controls, or Scherer [17] for a treatment of option replicating strategies such as portfolio insurance. For details on modern portfolio management and risk budgeting techniques see [14] or [15].

1.1 Description of the approach

We model two simplified investment strategies. In the first strategy an active manager manages a portfolio consisting of a "safe" money market index and one or more "risky" assets (equities, fixed income or credit), against a benchmark of the money market index, re-balancing the strategy periodically. In the second a passive manager makes a static investment in the same instruments plus a number of derivative instruments whose underlying assets are the risky assets, benchmarked against the money market index. The exposure to option instruments in some sense quantifies the value added by the active management. Using such a framework we can measure the difficulty of the active manager's job under different hypothesis.

The modelling is performed within a Monte Carlo framework. We generate a large number of forward trajectories for each of the assets under consideration in order to generate an expected distribution of final wealth for the dynamic strategy, and then attempt to match that pay-off as closely as possible with a static portfolio based on the same scenarios.

This problem presents several numerical challenges. The dynamic problem requires solving a large linear optimization problem with a large number of variables and constraints (see [21]). While the static optimization problem involves a much smaller number of variables and constraints, the utility function we will consider is extremely badly behaved, with many local minima and discontinuities, making it completely unsuited to variance-based optimization methods and thus we must consider alternative optimization methods.

1.2 An actively managed portfolio

We mimic the active manager's decision process using multi stage optimization and stochastic programming. Consider for example a situation in which a manager chooses an initial allocation between a money market index and an equity index to be held for one period, at which stage the portfolio may be rebalanced to a new allocation to be held for a further period. The target of the manager is to beat the benchmark (the money market index) at the end of the second period.

The factors which influence the manager's decision include the allowed budget of risk (defined, for example, in terms of a maximum probability of shortfalling the target) and the expected returns on the assets but also the knowledge that he will be free to change his allocation in the future in response to the actual performance of the assets during the first period. For example, in a scenario in which the equity underperforms expectations in the first period, thereby increasing the probability of a shortfall at the end of period two, the manager may decide to increase his holding of the money market index for the second period in order to minimize the expected downside. Thus the manager can make more efficient use of the risk budget with his initial allocation, in the knowledge that he will be able to rebalance after one period.

This decision is essentially a two-stage linear stochastic program, where a decision maker takes some action at the first stage, after which a random even occurs, affecting the outcome of the first stage decision. A recourse decision can then be made in the second stage that compensates for any bad effects that might have been experienced as a result of the first stage decision. The optimal policy from such a model is a first stage policy and a collection of recourse decisions (a decision rule) defining which second stage action should be taken in response to each random outcome.

We model this procedure using a Monte Carlo tree which branches at the start of the first period and then again in each scenario at the start of the second period. Then we seek a strategy consisting of an allocation at the start of the first period and an allocation for each possible scenario at the start of period two, which is optimal in some sense. The exact nature of the solution will be determined by the choice of utility function used to evaluate different strategies, which must reflect both the investment style under consideration and the budget of risk available to the manager.

1.3 A static portfolio

Next we try to replicate the payoff achieved by an active manager using a static "buy and hold" strategy. The ability of the active manager to rebalance his portfolio in response to portfolio performance, moving 100% into the money market index if necessary in order to hedge against any downside, introduces a measure of optionality into the active portfolio's expected payoff. Therefore, including a put option on the equity index with strike price equal to the expected final value of the benchmark in the static portfolio increases the ability of the passive manager to replicate the active payoff. We can also allow the active manager to buy a call on the equity with a much higher payoff, allowing us to mimic the ability of the active manager to increase risk in well performing scenarios.

We model the static portfolio by using the same scenarios for the underlying assets as above, and seeking the allocation between the assets and a number of optional strategies on the assets which produces a payoff that "best" (in a sense to be described below) matches the payoff of the dynamic strategy.

1.4 Overview of the paper

In the next section we will present a more thorough mathematical description of the two stages of the problem. We will then enumerate the various options in the way we tackle the problem. In Section 3 we will present the results of a control case and then we will examine the effect of varying each of the options listed in Section 2.

2 Problem description

2.1 Mathematical formulation of the optimization problem

The forward asset price scenarios are arranged in a tree structure as shown in Figure 1. Letting T denote the time horizon of the simulation in years and f denote the number of sampling times per year, then the set T_S of sampling times is given by $T_S = \{0, \frac{1}{f}, \frac{2}{f}, ..., 1, ..., 2, ..., T - \frac{1}{f}, T\}$. Let A denote the set of assets we are considering; for each scenario we will simulate the price of each asset $a \in A$ at each time $t \in T_S$. In Figure 1 the vertical dimension represents time and there is a node on each scenario at each sampling time.

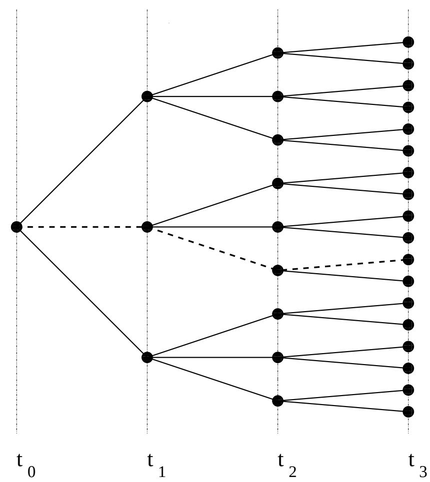

Fig. 1. Scenario tree: the broken line represents a single scenario

Let $T_D \subset T_S$ denote the portfolio decision times at which rebalancing may occur (T_D must include $t = 0$). As explained above, each scenario branches at each time $t \in T_D$, and this branching behaviour is regular in the sense that the same number of new scenarios is the same in each scenario at a given $t \in T_D$. It should be clear that the structure of the tree is completely specified by T, f, T_D and by the number of new branches at each $t \in T_D$.

We also use the following notation; N denotes the set of all nodes of the tree, N_D the set of all decision nodes and n_0 the unique node at time $t = 0$. To each node n we assign a probability, $P(n)$, which is the reciprocal of the number of scenarios at that time and for each node $n \in N - \{n_0\}$ we denote by $p(n) \in N$ the unique node preceding n on the same scenario.

Let $a' \in A$ be the target asset (i.e. the money market rate). For each node we simulate an asset price for each asset and thus for each node $n \in N - n_0$ we can define $r_{n,a}$ to be the return on asset a between nodes $p(n)$ and n.

Let $x_{n,a}$ be the amount of asset a held at node n and let $x^+_{n,a}$ and $x^-_{n,a}$ denote respectively the amount of a bought and sold at node $n \in N_D$. Let $W(n)$ denote the total portfolio wealth at node n and let $B(n)$ denote the target at node n.

The following constraints completely define the optimization problem;

$$x_{n,a} \geq 0 \text{ for all } a \in A, n \in N, \tag{1}$$

$$x^+_{n,a} \geq 0, x^-_{n,a} \geq 0 \text{ for all } a \in A, n \in N_D, \tag{2}$$

$$\sum_{a \in A} (x^+_{n,a} - x^-_{n,a}) = 0 \text{ for all } n \in N_D, \tag{3}$$

$$x_{n,a} = x_{p(n),a}(1 + r_{n,a}) \text{ for all } n \in N - N_D, \tag{4}$$

$$x_{n,a} = x_{p(n),a}(1 + r_{n,a}) + x^+_{n,a} - x^-_{n,a} \text{ for all } n \in N_D - \{n_0\}, \tag{5}$$

$$W(n) = \sum_{a \in A} x_{n,a} \text{ for all } n \in N, \tag{6}$$

$$B(n) = B(p(n))(1 + r_{n,a'}) \text{ for all } n \in N - n_0, \tag{7}$$

and

$$W(n_0) = B(n_0) = 100. \tag{8}$$

For the static problem we also have the extra constraint

$$x^+_{n,a} = 0, x^-_{n,a} = 0 \text{ for all } a \in A, n \in N_D, \tag{9}$$

i.e. no rebalancing is permitted.

2.2 Dynamic Setup

The choice of objective function will determine the shape of the distribution of final wealth resulting from the optimal allocation strategy - thus we need to find an objective function which reflects the investment style of total return strategies as well as the available budget of risk.

Typically the simulated pay-off distribution from a total return strategy assumes a quite skewed shape. The manager trades off a little decrease in expected absolute returns for active downside protection (perhaps specified as a pre-determined probability of a shortfall relative to the benchmark). In the examples we will use the *Semivariance Utility function*, defined as follows;

$$U(\{x_{n,a}\}) = \sum_{n \in N_F} P_n(\beta W^-(n) - (1 - \beta)W(n)) \tag{10}$$

where

$$W^-(n) = \begin{cases} B(n) - W(n) & \text{if W(n) } < \text{ B(n)} \\ 0 & \text{otherwise.} \end{cases}$$

The relative risk aversion of utility as a function of wealth is defined as

$$R(W) = \frac{-WU''(W)}{U'(W)}$$

(see for example [7]). Thus for $W > B$ this utility displays constant relative risk aversion ($R(W) = 0$) meaning that the relative allocation to risky assets does not change as wealth increases on the upside. For $W < B$ we find that

$$R(W) = \frac{-1}{1 - \frac{1-\beta}{2\beta W}}$$

which is negative for $\beta > \frac{1}{2W+1}$, meaning that the utility displays decreasing relative risk aversion on the downside. Thus the allocation to risky assets will decrease as the wealth decreases leading to a truncated left tail in the expected final wealth distribution.

The degree of downside risk aversion is determined by the value of β. In practice we choose $\beta \in [0, 1)$ to reflect the risk budget of the portfolio. For values of β close to 1, the objective function penalizes a shortfall more heavily than it rewards excess return, while decreasing β allows the manager more latitude for tolerating a shortfall in some scenarios if it is rewarded by a higher excess return in other scenarios.

By only applying the objective function at the final nodes we also allow the manager more freedom to choose a risky initial allocation than if the objective were applied at all times, or even just the rebalancing times. Fig. 2 illustrates how the desired right skewed payoff emerges over time, in this case for a three stage optimization problem.

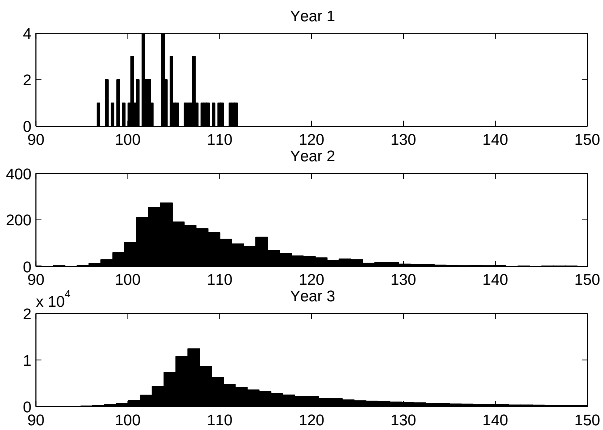

Fig. 2. Distribution of dynamic wealth: Results after three re-balancing stages.

2.3 Static Setup

For the static portfolio we broaden the investment universe to include a number of option strategies based on the risky asset(s) and seek an allocation that matches as closely as possible the payoff produced by the active management.

Thus in this optimization problem the benchmark is the distribution of final wealth of the active strategy and the objective is to find the static allocation strategy for which the distribution of final wealth is as close as possible to the dynamic distribution.

We consider strategies consisting of long and short positions in put options and call options. In some cases we also allow the optimizer to determine the optimal strike price of each option so as to be able to achieve the closest possible fit.

Let F_A denote the cumulative density function of the active strategy. Let S represent any choice of static asset allocation strategy (including the choice of strike prices) and let F_S denote the cumulative density function of the associated final NAV. Then the utility function is defined as

$$U(S) = max_{x \in (-\inf, \inf)}(|F_A(x) - F_S(x)|), \tag{11}$$

that is the test statistic of the two-sample Kolmogorov-Smirnov test.

This problem sets a numerically challenging problem. The objective function has a large number of discontinuities and local minima and is completely unsuited to traditional variance based optimization methods. In addition the fact that the strike prices may be variables of the problem makes it quite different to traditional asset allocation optimizations. Thus in this case we use the Adaptive Simulated Annealing technique [3], a guided random search engine suitable for problems with many local minima.

2.4 Scenario generation

We present two different processes for generating future price scenarios for the assets; Geometric Brownian Motion and multivariate GARCH.

Geometric Brownian Motion The process is defined as:

$$\begin{aligned} \tfrac{dS(t)}{S(t)} &= \mu dt + \sigma dW(t) \\ dW(t) &\overset{D}{\approx} \sqrt{dt} Z \end{aligned} \tag{12}$$

where Z is a standard Gaussian random variable. The drift μ and volatility σ is set equal to the historical risk premium and volatility for the sake of simplicity. The correlations are based upon historical correlations and are modelled via Cholesky decomposition.

GARCH DCC Time varying correlations are often estimated with Multivariate Garch models that are linear in squares and cross products of the data. A new class of multivariate models called dynamic conditional correlation (DCC) models was proposed by Engle [8] 1999,2002).This family of models is a very useful way to describe the evolution over time of the correlation matrix of large systems. The DCC-GARCH model overcomes the computational constraints of the multivariate GARCH models by adopting a two-step procedure. In the first step, a set of univariate GARCH models is estimated for each asset return. In the second stage, a simple specification is used to model the time-varying correlation matrix, which is obtained using the standardised residuals from the first stage. A particularly appealing feature of the model is that it preserves the simple interpretation of univariate GARCH models, while providing a consistent estimate of the correlation matrix. The inefficiency inherent in the two stage estimation process is coped with by modifying the asymptotic covariance of the correlation estimation parameters. Correlations are critical inputs for many of the common tasks of financial management, including risk management. Hedges require estimates of the correlation between the returns of the assets in the hedge. If the correlations and volatilities are changing, then the hedge ratio should be adjusted to account for the most recent information. By using this specification we aim to account for this feature and the variability induced by the dynamic re-balancing activity. Similarly, structured products such as rainbow options that are designed with more than one underlying asset, have prices that are sensitive to the correlation between the underlying returns. A forecast of future correlations and volatilities is the basis of any pricing formula.

We'll use a standard GARCH DCC framework to generate Monte Carlo scenarios in replacement of the GBM described in the general set up. We follow the general approach as described by Engle (2002) [8].

$$
\begin{aligned}
r_t \,|\Im_{t-1}| &\approx N(0, D_t R_t D_t) \\
D_t^2 &= diag(\omega_i) + diag(\kappa_i) \circ r_{t-1} r_{t-1}^{'} + diag(\lambda_i) D_{t-1}^2 \\
\epsilon_t &= D_t^{-1} r_t \\
Q_t &= S \circ (u^{'} - A - B) + A \circ \epsilon_{t-1} \epsilon_{t-1}^{'} + B \circ Q_{t-1} \\
R_t &= diag(Q_t)^{-1} Q_t diag(Q_t^{-1})
\end{aligned}
\qquad (13)
$$

Where R is a correlation matrix containing the conditional correlations. It's not the goal of the paper to provide a detailed description of the model, please refer to Engle (2002). The assumption of normality in the first equation gives rise to a likelihood function. Without this assumption, the estimator will still have the QML interpretation. The second equation simply expresses the assumption that each of the assets follows a univariate GARCH process.

2.5 Objective and risk neutral measures

Consider put or call options on a given underlying asset with different strikes but the same expiration. If we obtain market prices for those options, we can apply the Black-Scholes (1973) model to back-out implied volatilities. Intuitively, we might expect the implied volatilities to be identical. In practice, it is likely that they will not be.

Most derivatives markets exhibit persistent patterns of volatilities varying by strike. In some markets, those patterns form a smile. In others, such as equity index options markets, it is more of a skewed curve. This has motivated the name volatility skew. In practice, either the term "volatility smile" or "volatility skew" (or simply skew) may be used to refer to the general phenomena of volatilities varying by strike.

There are various explanations for why volatilities exhibit skew. Different explanations may apply in different markets. In most cases, multiple explanations may play a role. Some explanations relate to the idealized assumptions of the Black-Scholes approach to valuing options. Almost every one of those assumptions - log normally distributed returns, return homoskedasticity, etc. - could play a role. For example, in most markets, returns appear more leptokurtic than is assumed by a log normal distribution. Market leptokurtosis would make way out-of-the-money or way in-the-money options more expensive than would be assumed by the Black-Scholes formulation. By increasing prices for such options, the volatility smile could be the markets' indirect way of achieving such higher prices within the imperfect framework of the Black-Scholes model. Other explanations relate to relative supply and demand for options. In equity markets, the volatility skew could reflect investors' fear of market crashes which would cause them to bid up the prices of options at strikes below current market levels.

It is generally considered best practice to use implied volatility when pricing equity options (see for example [18]). We will compare the effects of using historical volatility versus implied volatility. In fact, as we will see, the outcome of our experiment is relatively insensitive to the method used to price the options, and in most cases we have considered historical volatility for the sake of objectivity and expediency.

3 Results

Below we will present the results of the experiments. We performed a basic version of the job as a control case and then observed the sensitivity of the results to changing different options. The variables of the problem which we vary are;

- the risk budget of the active manager,
- the asset return generating process (i.e. GBM or GARCH DCC),
- the investment universe,
- the dimensions of the tree, i.e. the number of rebalancings and the number of scenarios,
- the use of risk neutral or objective measures in pricing the options.

3.1 Case 1: The control case

In the control version of the problem we use the following parameters. The simulation horizon is set at 3 year, ($T = 3$), with rebalancing at 0, 1 and 2 years

($T_D = \{0, 1, 2\}$) with branching factors of respectively 50,50 and 40, giving 100,000 terminal scenarios. The sampling frequency is weekly ($f = 52$) and the scenarios are generated using Geometric Brownian Motion.

We consider an investment universe consisting of two asset classes, money market and equities. To simulate these assets we have used JP Morgan 3M Cash Index and the Dow Jones EuroStoxx index respectively (both denominated in Euro), based on a 5 year sample of weekly data(2001-2006). The money market index also acts as the benchmark.

Of course, it is not possible to directly invest in an index such as the EuroStoxx (although one can replicate such an investment using index futures or exchange traded funds). However the intention is that the return and volatility of the index could be considered as a proxy for the return and vitality one could expect of an investment in the relevant asset class.

We have used the SemiVariance loss objective function described in 2.2, applied at terminal nodes and we have set $\beta = 0.3$. For the static portfolio we have included two derivatives.

In this case we do not allow the optimizer to vary the strike prices, but instead we have chosen two option strategies which we believe should allow the static manager to better replicate the active strategy.

The first is a put option on the equity index with strike price 108.2, which is the expected value at $t = 3$ of a portfolio invested 100% in the money market index. This is intended to mimic the active manager's ability to create downside protection by increasing the allocation to the safe asset if the risk of a shortfall become too high.

The second is a call option on the equity with strike price equal to 140.0, which gives us access to more upside in the scenarios in which the equity performs particularly strongly. This is analogous to the ability of the active manager increase the equity allocation when the risk of a shortfall is low.

For simplicity we have used flat historical volatility in pricing the options. We will see the effect of risk-neutral pricing in a later section.

The results of the dynamic and static portfolio allocation decisions are shown in Table 1. Fig 3 displays the distribution of final wealth across 100,000 scenarios of the two strategies.

3.2 Effect of varying the risk parameter

We consider the effect on the outcome of our problem of changing the value of the risk parameter β. We consider two different versions of the control case with tighter risk control - a medium risk case with $\beta = 0.7$ and a low risk case with $\beta = 0.9$.

Case 2(a): $\boldsymbol{\beta = 0.7}$. Table 2 summarizes the results of the medium risk problem. In comparison to the control case, two differences are notable. Firstly the lower allocation to derivatives reflects the lower allocation to equities in the dynamic portfolio and the lower level of dynamic hedging performed by the

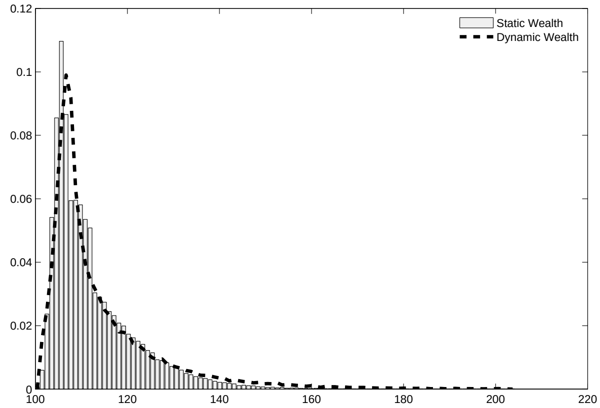

Fig. 3. Case 1: Control Case

Table 1. Case 1: Control Case

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	80.7%	74.8%	73.9%	18.21%
Equity Index	19.3%	25.2%	26.1%	80.25%
Put Option (Equity, strike 108.3)				1.12%
Call Option (Equity, strike 140.0)				0.42%
Final result	Dynamic			Static
Benchmark wealth	108.3			108.3
NAV	114			112.5
Standard deviation NAV	13.3			9.3
Probability NAV > benchmark	0.56			0.58
ks-statistic				4.26%

active manager. Secondly the lower value of the ks-statistic indicated that a more accurate replication job is possible in comparison to the higher risk case.

Table 2. Case 2(a): Medium Risk

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	96.3%	95.1%	94.6%	95.4%
Equity Index	3.7%	4.9%	5.4%	4.28%
Put Option (Equity, strike 108.3)				0.29%
Call Option (Equity, strike 140.0)				0.03%
Final result	Dynamic			Static
Benchmark wealth	108.28			108.28
NAV	109.12			109.38
Standard deviation NAV	2.62			1.78
Probability NAV > benchmark	0.56			0.58
ks-statistic				4.03%

Case 2(b): $\beta = 0.9$. Table 3 summarizes the results of the low risk case. Note again that, in comparison to the $\beta = 0.3$

Table 3. Case 2(b): Low Risk

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	98.6%	98.2%	97.9%	98.13%
Equity Index	1.4%	1.8%	2.1%	1.75%
Put Option (Equity, strike 108.3)				0.098%
Call Option (Equity, strike 140.0)				0.012%
Final result	Dynamic			Static
Benchmark wealth	108.3			108.3
NAV	108.7			108.6
Standard deviation NAV	0.74			0.89
Probability NAV > benchmark	0.60			0.60
ks-statistic				2.93%

3.3 Changing the data generating process

Next we examine the sensitivity of the solution to the data generating solution. We compare the results of the control case with the solution to an identical problem with only the scenario generating process changed. In Table 4 we note the larger allocation to options compared to the control case. This reflects the extra value added by active management in the context of time varying correlations and volatilities. Also the ks-statistic is slightly higher than in the control case.

Table 4. Case 3: Garch scenarios

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	66.1%	57.4%	52.3%	58.34%
Equity Index	33.9%	42.6%	40.7%	39.62%
Put Option (Equity, strike 108.2)				1.92%
Call Option (Equity, strike 140.0)				0.12%
Final result	Dynamic			Static
Benchmark wealth	108.3			108.3
NAV	117.5			115.7
Standard deviation NAV	17.5			13.4
Probability NAV > benchmark	0.63			0.65
ks-statistic				4.61%

3.4 Changing the investment universe

Next, we increase the size of the investment universe to 4 assets by adding a government bond index and a corporate bond index, calibrated to the JP Morgan EMU Government Bond Index and the MSCI EMU Credit (Investment Grade) Index (2001-2006) (see Fig 4).

A total return manager would typically utilize a much broader universe of asset classes in order to exploit the benefit of diversification. Unfortunately in this case we are constrained by the numerical complexity of the optimization problems. We consider the four asset case for the purposes of comparison with the two asset case.

Table 5 summarizes the results. In comparison with the control case we see a much lower allocation to options. This indicates the greater ease with which a manager can beat his benchmark in this context (the dynamic portfolio displays a far higher Sharpe Ratio). The ks-statistic is slightly lower than the control case as the extra asset classes make it easier to replicate a given strategy. One could extrapolate from these results that the effect of adding further assets would allow us to better replicate the dynamic strategy.

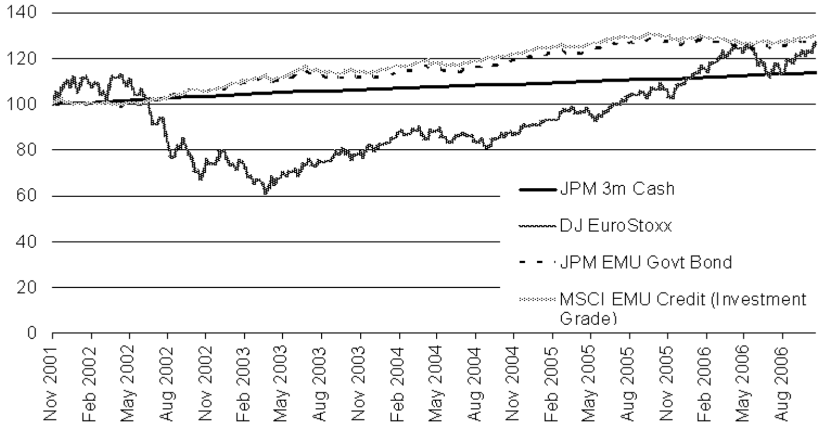

Fig. 4. The historical performance of the indices used to calibrate the 4 assets classes

Table 5. Case 4: 4 assets

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	0%	0%	0.1%	6.29%
Equity Index	41.8%	28%	39.7%	39.13%
Bond Index	0%	1.2%	3.4%	0.64%
Corp Bond Index	58.2%	70.8%	56.8%	53.36%
Put Option (Equity, strike 108.3)				0.41%
Call Option (Equity, strike 140.0)				0.17%
Final result	Dynamic			Static
Benchmark wealth	108.3			108.3
NAV	123.6			122
Standard deviation NAV	18.8			15.3
Probability NAV > benchmark	0.84			0.82
ks-statistic				3.62%

3.5 Changing the rebalancing frequency

In this section we change the frequency of rebalancing, comparing the results of the control case to the results obtained using a tree with branching at times 0, $\frac{1}{2}$, 1, 1 $\frac{1}{2}$ and 2 years and using 300,000 ($20 \times 15 \times 10 \times 10 \times 10$) scenarios.

The extra rebalancing times produces a notable effect of the dynamic payoff, with an expected NAV of 121.5 after three years. This presents a much harder challenge for the static portfolio to achieve. To facilitate the task we include two extra options on the equity index, so that the static portfolio now includes two puts and two calls. Furthermore, we allow the optimizer to chose the strike prices of all but one of the options (the first put option, as before, has a strike of 108.3, being the expected value of a 100% money market portfolio).

However, even with these extra degrees of freedom, the optimal static portfolio is a much poorer fit to the dynamic portfolio compared to any of the cases considered so far, as evidenced by the ks-statistic of 7.58%. This, in spite of the much higher overall exposure to derivatives (Table 6).

Unfortunately the numerical difficulty of the optimization problem increases exponentially with the addition of extra rebalancing times, but certainly the evidence of the extra value added by more active management is clear.

Table 6. Case 5: "bushy tree"

	Dynamic					Static
Asset Allocation	t = 0	t = 0.5	t = 1	t = 1.5	t = 2	t = 0
Money Market Index	81.8%	60.4%	63.2%	59.4%	58.5%	74.74%
Equity Index	18.2%	39.6%	36.8%	40.6%	41.5%	26.04%
Equity Put 1 (strike 108.3)						- 0.98%
Equity Put 2 (strike 84.01)						- 1.0%
Equity Call 1 (strike 122.2)						-0.29%
Equity Call 2 (strike 141.7)						1.49%
Final result	Dynamic					Static
Benchmark wealth	108.3					108.3
NAV	121.5					117.5
Standard deviation NAV	23.6					19.7
Prob NAV > benchmark	0.69					0.64
ks-statistic						7.58%

3.6 Changing from objective to risk neutral pricing

In this section we test the effects of changing from flat historical volatilities to market implied volatilities in pricing the options in the static portfolio. Fig. 5 displays the data sample we used of to calibrate our model.

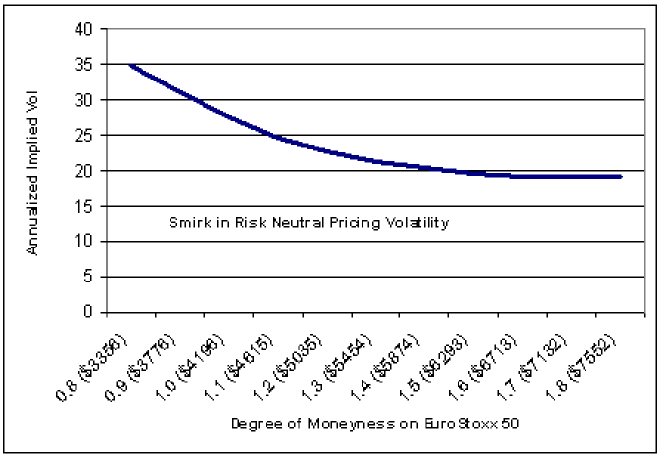

Fig. 5. The chart displays smirk effect on Dow Jones Eurostoxx Equity Index Options: DJ Index Options 3 years horizon as of November the 22nd 2007.

Table 7 displays the results. The change from the control case is slight. The effect of using the volatility smile is to increase the price of the options, consequently we can afford less exposure. The effect is that while the option exposure is only slightly reduces (1.41% compared to 1.54% in the control case), the fit is significantly poorer (5.55% compared to 4.26%).

Table 7. Case 6: Volatility smile

	Dynamic			Static
Asset Allocation	t = 0	t = 1	t = 2	t = 0
Money Market Index	80.7%	74.8%	73.9%	76.36%
Equity Index	19.3%	25.2%	26.1%	22.23%
Put Option (Equity, strike 108.2)				1.33%
Call Option (Equity, strike 140.0)				0.08%
Final result	Dynamic			Static
Benchmark wealth	108.3			108.3
NAV	114			112.3
Standard deviation NAV	13.3			9
Probability NAV > benchmark	0.56			0.6
ks-statistic				5.55%

4 Conclusion

We have attempted to quantify the value of active portfolio management in a total return context and to examine the effect on the quantification of various choices considered in our setup. Table 8 summarizes our findings.

Table 8. Summary

Case	Total option exposure p.a.	Goodness of fit
Control	0.51%	4.26%
Medium Risk	0.11%	4.03%
Low Risk	0.04%	2.93%
Garch Scenarios	0.68%	4.61%
4 Assets	0.20%	3.62%
More Rebalancing	1.25%	7.58%
Volatilty Smile	0.47%	5.55%

We interpret the first column as the value added by the active management compared to a static strategy. Thus the results indicate;

- The value added by active management increases with increasing risk budget. For the medium risk strategy which delivers approximately 27 bps p.a. over the benchmark the value added is of the order of 11 bps. In comparison, the control case delivers approximately 180 bps over the benchmark with a value added of 51 bps.
- The value added by the manager increases when a more realistic asset return generating process (GARCH) is used.
- The value added by the manager decreases with a broadening of the investment universe. Increasing the investment universe allows the manager to exploit the benefits of diversification.
- The value added by the manager (and the difficulty of replicating the results statically) increase dramatically with increased rebalancing.
- The results are relatively insensitive to a change from objective to risk neutral option pricing.

References

1. Alexander, C.: Measuring and modelling financial risk. Wiley series in Financial Engineering, Volume 1, New York (1998)
2. Baxter, M., Rennie, A.: Financial calculus: an introduction to derivatives pricing. Cambridge University Press, Cambridge (1996)
3. Chen, S., Luk, B.L.: Adaptive simulated annealing for optimization in signal processing applications. Signal Processing **79** (1999) 117–128

4. Consigli, G., Dempster, M.A.H.: The CASM stochastic programming model for dynamic Asset-Liability Management. In Ziemba, W.T., and Mulvey, M.J.: Worlwide assets and liabilities models. Cambridge University Press, Cambridge (1998)
5. Dempster, M.A.H, Germano, M., Medova, E., Sandrini, F., Scrowston M.: Managing guarantees. Journal of Portfolio Management (1997) 51–61
6. Grinold, R, Kahn, R.: Active portfolio management: a quantitative approach for producing superior returns and controlling risk. McGraw-Hill, New York (1999)
7. Elton, E., Gruber, M., Brown, S., Goetzmann, W.: Modern portfolio theory and investment analysis. Wiley, New York (2003)
8. Engle, R., and Kroner, F.: Multivariate simultaneous generalized ARCH. Econometric Theory **11** (1995) 122–150
9. Golub, B.W, and Tilman, L.M.: Risk management: approaches for fixed income markets. Wiley Frontiers in Finance, New York (2000)
10. Knight, J., Satchell, S.: Forecasting volatility in the financial markets. Butterworth-Heinemann Finance, London (2002)
11. Mulvey, J.M., Koray, D., Simsek, Fabozzi, Frank J. et al: Assisting Defined benefit pension plans, Revision number 4, (2007)
12. Nefci, S.N.: Principles of Financial Engineering, Elsevier Academic Press, San Diego (2004)
13. Mulvey, John M.: An Asset-Liability Investment System, Interfaces 24 (1994) 22-33
14. Rahl. L.: Risk Budgeting A New Approach to Investing, Risk Books, London (2000).
15. Satchell, S., Scrowcroft, A.: Advances in Portfolio Construction and Implementation, Butterworth-Heinemann, Oxford (2005).
16. Scherer, B.: Asset and Liability Management Tools, a Handbook of Best Practice, Risk Books, London (2003)
17. Scherer, B.: Portfolio Construction and Risk Budgeting (3rd Edition) Risk Books, London (2007).
18. Schoutens, W.: Levy Processes in Finance: Pricing Financial Derivatives, Wiley, New York (2003).
19. Wilkie, A.D: More On a Stochastic Asset Model for Actuarial Models , Journal of Portfolio Management, (1995)
20. Zenios, S.: Asset-Liability Management Under Uncertainty for Fixed Income Securities, Annals of Operational Research 59 (1998) 77-98
21. Zenios, S., Nielson, S., Consiglio, A.: Practical Financial Optimization: Decision Making for Financial Engineers, Wiley Blackwell, New York (2009)
22. Ziemba, W.T., Mulvey, M.J.: Worlwide Assets and Liabilities Models, Cambridge University Press, Cambridge (1998)