

The performance of stochastic dynamic and fixed mix portfolio models^{*}

Stein-Erik Fleten^{*}, Kjetil Høyland¹, Stein W. Wallace²

Department of Industrial Economics and Technology Management, Alfred Getz v. 1, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway

Abstract

The purpose of this paper is to demonstrate how to evaluate stochastic programming models, and more specifically to compare two different approaches to asset liability management. The first uses multistage stochastic programming, while the other is a static approach based on so-called constant rebalancing or fixed mix. Particular attention is paid to the methodology used for the comparison. The two alternatives are tested over a large number of realistic scenarios created by means of simulation. We find that due to the ability of the stochastic programming model to adapt to the information in the scenario tree, it dominates the fixed mix approach.

Key words: Simulation, stochastic programming, portfolio selection, asset liability management, performance measurement, nonlinear programming.

1 Introduction

The purpose of this work is to demonstrate how to quantitatively compare two models of the same underlying decision problem, and more specifically to compare two different approaches to portfolio management. The first is a

^{*} Accepted for publication in European Journal of Operational Research.

^{*} Corresponding author. Tel.: +47-73 59 12 96; fax: +47-73 59 36 03.

Email addresses: Stein-Erik.Fleten@iot.ntnu.no (Stein-Erik Fleten), Kjetil.Hoyland@gjensidige.no (Kjetil Høyland), Stein.Wallace@hiMolde.no (Stein W. Wallace).

¹ Present address: Gjensidige Asset Management ASA, P.o. box 276, NO-1324 Lysaker, Norway

² Present address: Molde University College, Servicebox 8, NO-6405 Molde, Norway

multistage stochastic program, while the other is based on the simple decision rule of constant rebalancing, also called fixed mix. The comparison is done in a fair and realistic way. Realistic here means a situation that is close to the practical use of the models, where the models are rerun at regular intervals. An out-of-sample test procedure like the one applied in this paper is suitable for this kind of comparison. The test methodology is standard, but this paper is one of the few attempts to use it on a real problem.

In the literature, there are infrequent reports on empirical testing of the performance of stochastic programming models. Building stochastic simulation models to see how a stochastic programming model performs in practice is a complex task, because it involves solving a large number of stochastic programs, which in itself is difficult.

The test methodology is explained in the context of a specific portfolio problem, but the conclusions are general enough to be used in other areas. We find that a dynamic stochastic approach dominates a fixed mix approach, but that the degree of domination is much smaller when the models are compared out-of-sample than when they are compared in-sample. The reason for this is that in an out-of-sample context, where the random input data is (at least) structurally different from the in-sample scenarios, the stochastic programming model loses its advantage in optimally adapting to the information available in the scenario tree. Furthermore, the performance of the fixed mix approach will improve, because the asset mix is updated at every stage.

Kusy & Ziemba (1986) test a two-stage simple recourse model and compare this to a stochastic decision tree model. However, a new scenario tree is not generated each time the horizon is rolled forward. Cariño & Turner (1998) compare the fixed mix-rule with a true dynamic stochastic optimization-based model. The model horizon is not rolled forward, and no out-of-sample data are used in the test procedure. Cariño, Myers & Ziemba (1998) show the historical performance of an asset liability model used by Yasuda Kasai, a Japanese insurance company. Vassiadou-Zeniou & Zenios (1996) and Zenios, Holmer, McKendall & Vassiadou-Zeniou (1998) also do validation back-testing based on historical data. Independently, Kouwenberg (2001) has developed and tested a pension fund asset liability management model using rolling horizon simulations.

In Section 2 we describe the stochastic dynamic approach and the fixed mix approach. Section 3 shows the simulation methodology that is applied to compare the two approaches. In Section 4 we present the numerical results before the conclusions are given in Section 5.

2 The stochastic dynamic and the fixed mix approaches

We use a multi period, stochastic asset liability management model developed for the Norwegian mutual life insurance company Gjensidige Spareforsikring. A mathematical description is provided in Høyland & Wallace (2001a). The portfolio selection problem is modeled at the strategic level, where capital is allocated among a few aggregated asset classes such as stocks and bonds. The objective is to maximize the expected portfolio value at the end of the horizon net of costs, subject to rebalancing and legal constraints on balance figures. The costs are composed of transaction costs³ and imputed costs associated with the violation of the legal constraints. These imputed costs are used as the measure of risk in the model.

Section 2.1 describes the two approaches to the model, while Section 2.2 briefly describes the scenario generation process.

2.1 How the two approaches differ

Both the fixed mix and the dynamic versions of the model require decisions to be made at discrete points in time and discrete probability distributions for the uncertain variables. The dynamic model can be explained by considering a scenario tree, of which a two period (three stage) example is shown in Figure 1. The top node represents the decisions today and the nodes further down




Fig. 1. The scenario tree. The nodes represent decisions, while the arcs represent realizations of the uncertain variables.

represent conditional decisions at later stages. The arcs linking the nodes represent realizations of the uncertain variables. This approach will capture the dynamics of decision making, since decisions are adjusted to the realizations of the uncertainties.

In the fixed mix model we assume the following decision rule: The portfolio is rebalanced to fixed proportions (for instance 70% bonds and 30% stocks)

³ The presence of transaction costs, and the tradeoff between short term risk and return and long term risk and return, makes a dynamic stochastic approach necessary.

at all future decision nodes. This means that at all decision nodes, assets are bought and sold so that the fixed mix is maintained. The optimization problem is to find the proportions that maximize the objective function. Note the difference between a fixed mix model and a fixed mix investment strategy. In the fixed mix investment strategy, the proportions are kept constant over a long investment horizon, while applying a fixed mix model means that the proportions are changed every time the model is rerun. In this paper we compare the results of a fixed mix model with the results of a stochastic, dynamic model. See Perold & Sharpe (1988) for a description of the fixed mix and other decision rules.

The fixed mix approach does not require the tree structure, but allows the uncertainties to be described in terms of streams of outcomes (or scenarios), as illustrated in Figure 2. However, we have applied the same scenario tree




Fig. 2. Possible description of the uncertainties with the fixed mix approach. In this approach the model is unable to exploit the perfect foresight implied after the first stage in this tree.

for both approaches. Doing this ensures that we get comparable results and means that the only difference between the two model formulations are the constraints in the fixed mix model assuring that the portfolio is rebalanced to the fixed mix at every decision node. If the approaches were to be compared with respect to both solution time and solution quality, this procedure might disfavour the fixed mix approach⁴. However, this problem does not arise since this paper is concerned with comparing the quality of the solutions given the quality on the input data, and not solution times.

2.2 Scenario generation

Generation of scenarios can be based on simulation or construction. The applied scenario generation method is a combination of the two. We let the decision maker specify the market expectations by any statistical properties that are considered relevant for the problem to be solved, and construct the

⁴ For the fixed mix approach we have chosen to include the transaction costs incurred in every state at every stage due to rebalancing to the fixed mix. Thus we have as many recourse variables as for the dynamic approach and the solution time for the fixed mix approach is at least as large as for the dynamic approach in our case.

tree so that these statistical properties are preserved. This is done by letting (random) outcomes and probabilities in the scenario tree be decision variables in a nonlinear optimization problem where the objective is to minimize the square distance between the statistical properties specified by the decision maker and the statistical properties of the constructed tree (Høyland & Wallace 2001*b*).

Generally, expectations for financial markets will depend on the state of the economy. Some statistical properties are clearly state dependent, while others might be specified independently of the state. As an example of state dependency, consider the volatility of stock returns. Empirical studies have shown the effect called volatility clustering, meaning that the volatility increases after a period of extreme returns. We model this effect by letting the volatility (standard deviation) in period $t + 1$ be a function of the outcomes' deviation from the mean in period t . For more details of the scenario generation system, see Section 4.1.

3 Comparing the performance of the stochastic dynamic and the fixed mix models

This section illustrates how we test the quality of the solutions obtained from each approach. In practice, only the first stage solution will be used for actual decision making, whatever approach is used. The conditional decisions at stages two and onward are only made in order to find the right incentives for the first stage decisions. If the model system is run for example each quarter, a new instance of the model will be generated and solved at the end of a quarter.

We want to test how good each method is on average, for a large number of realistic economic scenarios, denoted *simulation scenarios*. A single simulation scenario consists of realizations of the uncertain variables in each simulation period. To test the solutions of the two approaches we proceed as follows: At the beginning of the first period, our decision model tells us what to do, and at the end of the first period we see the consequences of our decisions. Given that information and the new state of the economy, we make a new model instance and obtain the decision for the beginning of the second period. Based on the outcomes in the second period, we calculate the consequences of this decision. The process continues for the third period, and in principle, indefinitely. Keep in mind that we do not use the information on what will happen when we make the decisions. Each time we make a decision, the future is stochastic, and the information we do use is in the form of scenario trees. These trees are created solely based on past and current information, and not on future outcomes in the simulation scenarios, i.e. presently unavailable information. Denote the scenarios contained in these trees *optimization scenarios*.

3.1 The test procedure

Our testing methodology follows a sequence of four major steps, assuming that the difficult task of specifying market expectations is already done:

- (1) Based on the market expectations, generate optimization scenarios and use the two approaches to obtain the present decision.
- (2) Based on the same market expectations, generate a high number of simulation scenarios.
- (3) For each outcome in each simulation scenario, generate an optimization scenario tree and solve the problem using both the dynamic and the fixed mix approaches.
- (4) Based on the consecutive decisions and outcomes, calculate the economic performance of both approaches.

Because our scenario generation system is such that optimization scenario trees can be generated conditional on sampled outcomes, we can generate sufficient input to the two optimization approaches for each outcome in each simulation scenario (since we assume that our decisions do not influence the stochasticity, we could in principle have generated all scenario trees in advance). Of course, for the very first stage we have no prior stage-information, so the scenario tree will be generated based only on the current market expectations.

The procedure is illustrated in the Figures 3–5 below.

3.2 Potential error sources

If the simulation scenarios differ from the optimization scenarios then a potential bias toward one of the approaches is introduced in the test procedure. What could happen is that the difference between the information in the optimization scenarios and the simulation scenarios makes the dynamic approach a priori worse off. The dynamic approach uses the information in the optimization scenarios better than the fixed mix approach, but it could be using misleading information so that it consistently makes relatively bad decisions.

Modeling the problem with a finite horizon is another potential error source. With few periods, the potential gain for a dynamic approach over a static one is small.

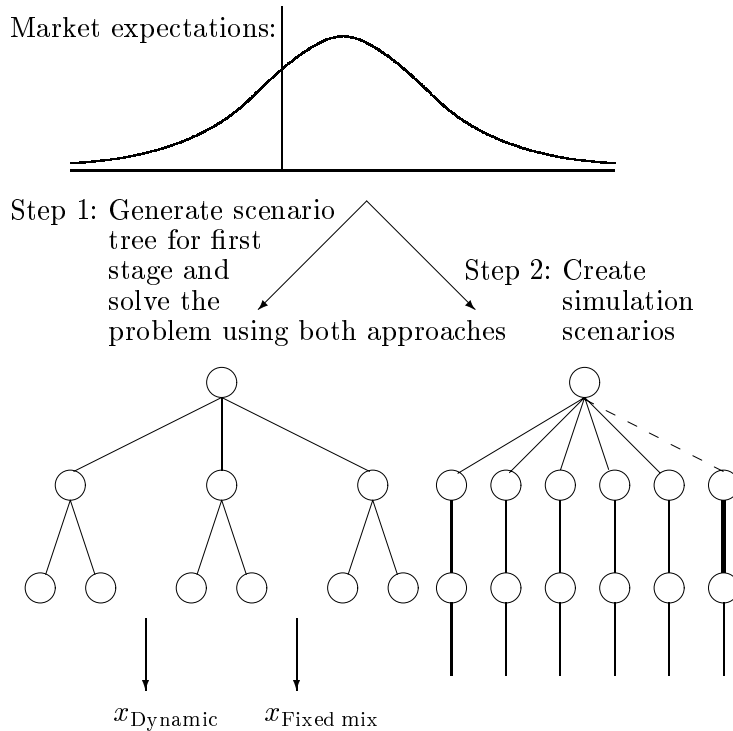


Fig. 3. Steps 1 and 2 of the test procedure. In step 1 we construct optimization scenarios that are consistent with the market expectations, and optimize using both approaches to obtain first stage decisions, denoted x_{Dynamic} and $x_{\text{Fixed mix}}$. In step 2 we generate a high number of simulation scenarios based on the same market expectations. Consult Section 4 and Appendices A and B to see how the simulation scenarios are generated.

4 Numerical example

This section presents the numerical results. We study a mutual life insurance company with a portfolio similar to that of Gjensidige Spareforsikring anno 1997, using a somewhat distorted data set. Assume that the company invests in the money markets (i.e. bonds with less than a year to maturity) and stocks. We let the money markets be represented by the three month NIBOR (Norwegian InterBank Offering Rate) and stocks be represented by the Oslo Stock Exchange Total index. The starting values of assets and liabilities are given in Table 1. The liabilities are the basis for the calculation of the legal requirements, and the risk is modeled by penalizing the violations of the legal requirements. For details, see the model formulation in Høyland & Wallace (2001a).

We employ a three period (four stage) model with a total of 1250 scenarios to obtain the present solutions from the dynamic approach and the fixed mix approach, corresponding to step 1 in Section 3.1. Each period is one year, so that the total planning horizon is three years. To obtain the future solutions,

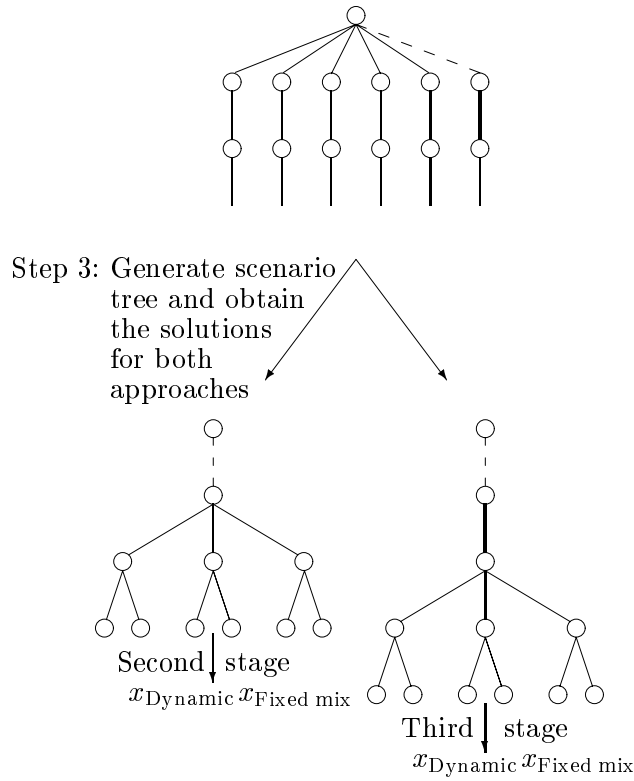


Fig. 4. Step 3: Generating conditional trees and new solutions. After the outcome represented by the dotted line, new information is received, and the model horizon is rolled forward. A new scenario tree is generated conditional on the information in that outcome. The model is solved to obtain a new present solution, in the figure called the second stage solution. The next outcome in the simulation scenario is represented by the thick line, which in a similar way gives rise to the third stage solution. Thus for each scenario, we obtain the solutions for consecutive stages after generating scenario trees conditional on previous outcomes.




Fig. 5. Outcomes and decisions for one scenario. Since decisions and outcomes are known for this scenario, in step 4 we can calculate the objective function value at the end of the third period. So each simulated scenario gives rise to one performance number for the dynamic stochastic approach, and another for the fixed mix approach.

Table 1
Starting balance.

Assets		Liabilities	
Money markets (3 month NIBOR)	70	Insurance fund	87
Stocks	30	Supplemental reserves	4
		Equity	6
		Primary capital	2
		Other liabilities	1
Sum assets	100	Sum liabilities	100

corresponding to step 3 in Section 3.1, two-period (three stage) models of 72 scenarios are used. Although this way the third stage decisions are made based on a model whose horizon is beyond the horizon of the simulation scenarios, the fixed mix and dynamic approaches are treated alike, thus this is not a significant source of bias.

The model is programmed in AMPL (Fourer, Gay & Kernighan 1993) and solved using MINOS 5.4 (Murtagh & Saunders 1983).

Section 4.1 illustrates the process of generating scenarios, both for simulation and optimization. The two versions of the model are solved for different levels of risk aversion, and the results from comparing the quality of the solutions are given in Section 4.2.

4.1 Scenario generation

This section explains how the market expectations are specified, including state dependencies, and how optimization scenarios and simulation scenarios are generated.

4.1.1 Specification of market expectations

The basis for both the simulation scenarios and optimization scenario trees are user supplied percentiles for the first period marginal distributions of stock returns and interest rates, the correlation between the asset classes and the definition of the state dependent statistical properties.

The percentiles supplied by the user are given in Table 2. Marginal distribution functions are fitted to the percentiles, as explained in Appendix A.

The correlation between stocks and interest rates is assumed to be 0.2 in all

Table 2

Percentiles of the marginal distributions. For the money market, the market views are expressed in terms of expectations for the interest rates, while for stocks the expectations are given in terms of total returns.

Percentile	0	0.05	0.25	0.5	0.75	0.95	1
Short term interest rates	1.5	2.4	3.6	3.9	4	5.8	7
Stocks	-29	-23	3	15	17	19	59

periods.

As explained in Section 2, some statistical properties are modeled as state dependent, while others are assumed independent of the state of the economy. The state dependencies generally depend upon the characteristics of the asset class. In this example we have modeled state dependent expected returns and volatilities for both asset classes. The other statistical properties are assumed state independent, meaning that they are the same in all states of the economy at a certain point in time.

In order to capture the volatility clustering effect, the state dependent standard deviation (or volatility) in period $t > 1$ is given by

$$\sigma_{it} = (1 - \nu_i)\bar{\sigma}_{it} + \nu_i |x_{i,t-1} - \mu_{i,t-1}|, \quad (1)$$

where $i \in \{s, b\}$ is the asset class index, either stocks (s) or interest rates (b), $\nu_i \in [0, 1]$ is the volatility clustering parameter for asset class i (a large ν_i leads to a large degree of volatility clustering), $\bar{\sigma}_{it}$ is the average standard deviation of the outcome of asset class i in period t , x_{it} is the outcome of asset class i in period t , and μ_{it} is the expected outcome of asset class i in period t . This way the volatility will increase after extreme returns, and decrease after more normal returns, in line with empirical observations.

For interest rates we model a mean reversion effect⁵, and let the expected value of the interest rate at the end of period $t > 1$ be given by

$$\mu_{bt} = \kappa\alpha + (1 - \kappa)x_{i,t-1}, \quad (2)$$

where κ is the mean reversion factor (a high κ leads to a large degree of mean

⁵ Mean reversion means that interest rates tend to revert to an average level. When interest rates are high, the economy slows down and interest rates tend to fall, and when interest rates are low, the economy booms and interest rates tend to rise. The mean reversion effect usually needs more than two or three years to manifest itself to a significant degree, and a low value for the mean reversion factor is chosen in the numerical example.

reversion), α is the mean reversion level and x_{bt} is the interest rate at the end of period t .

For stocks we assume that the expected total return in each period is given by

$$\mu_{st} = x_{b,t-1} + \theta_t \sigma_{st}, \quad (3)$$

where x_{bt} is the state dependent short term interest rate in period t , σ_{st} is the state dependent standard deviation of return on stocks in period t and θ_t is a risk premium constant in period t .

4.1.2 Generating optimization scenarios

For generating optimization scenario trees we have assumed that the relevant statistical properties are the first four moments of the marginal distributions, the correlation and the description of the state dependent properties. From the marginal distribution functions, which are derived as explained in Appendix A, we can calculate all marginal moments. Table 3 contains specifications of the marginal distribution properties for period 1 and the state independent marginal distribution properties for periods 2 and 3. The volatility clustering parameter, ν_i , in Equation (1) is set to 0.3 for both assets, the mean reversion level and factor (α and κ) in Equation (2) are set to 4% and 0.2 respectively, while the risk premium constant, θ_t , in Equation (3) is set to 0.3 in all periods.

Table 3

Market expectations used for generation of optimization scenarios. The expectations for the first period are derived from the marginal distributions, see Appendix A for details. While the skewness and kurtosis⁶ are assumed state independent, the expected value and the standard deviation are assumed state dependent.

Asset class	Distribution	(End of)	(End of)	(End of)
	Property	Period 1	Period 2	Period 3
Money market—3 months	Exp. spot rate	0.04	State dep	State dep
	Standard dev.	0.01	State dep	State dep
	Skewness	0.5	0.5	0.5
NIBOR	Kurtosis	3.0	3.0	3.0
Domestic stocks	Exp. return	0.085	State dep	State dep
	Standard dev.	0.15	State dep	State dep
	Skewness	-0.5	-0.5	-0.5
	Kurtosis	4.0	4.0	4.0

For the present optimization problem (corresponding to step 1 in Section 3.1), a three period (four stage) scenario tree is generated. This has 50 outcomes in the first period, and 5 in the second and the third, yielding a total of 1250 scenarios. The scenario tree is consistent with the market expectations given in Table 3 and the correlation and state dependent statistical properties defined in Section 4.1.1. For the future optimization problems, we generate two-period (three stage) scenario trees with 12 outcomes in period 1 and 6 outcomes in period 2, leading to a total of 72 scenarios. The first period specifications in these future optimization problems will be dependent on the sampled outcome. The state dependencies are specified in the same way as before. The generated trees are consistent with these specifications, in addition to the state independent specifications in Table 3 and the correlation defined in Section 4.1.1.

4.1.3 Generating simulation scenarios

For generating the simulation scenarios we do not use the calculated moments in 3, but sample from the fitted cumulative distribution functions directly. To capture the correlation, interest rates are sampled conditional on each stock return. For each stock return, the conditional interest rate distribution from which to sample is found, and an interest rate sample is drawn from this distribution. For details, see Appendix B.

For the subsequent periods, the means and the standard deviations are updated for mean reversion and volatility clustering according to Equations (1), (2) and (3). In the end we have a number of simulation scenarios containing subsequent stock return and interest rate pairs for three periods, satisfying the same statistical properties as the first-stage optimization scenarios, including the state dependencies.

4.2 Numerical results

Note that since the fixed mix formulation leads to a non-convex optimization model, we would usually need global optimization routines to be sure that the optimal solution is found. However, testing shows that for our data sets, the problem is convex⁷.

⁶ The normal distribution has a kurtosis of 3.0. A kurtosis larger than 3.0 means that the distribution is more peaked around the mean and have fatter tails than the normal distribution.

⁷ Many model instances were generated, varying the starting balance, the market expectations and the degree of risk aversion. Using many different starting values for each instance, the optimization always converged to the same solution.




Fig. 6. Optimization results of step 1 in Section 3.1. The efficient frontiers show the tradeoff between the expected return and the risk, and is obtained by solving the models for different required rates of expected returns. The numbers in parentheses show the initial investment in stocks for the dynamic and the fixed mix approaches, respectively.

The goal is to minimize risk subject to a minimum target expected portfolio return. Risk is measured in terms of shortfalls relative to legal requirements, and is given by the expected accumulated quadratic shortfalls:

$$F = \min \sum_{s \in \mathcal{S}} P_s \left[\sum_{j \in \mathcal{J}} c_j \sum_{t \in \mathcal{T}} \rho_{st} x_{s jt}^2 \right], \quad (4)$$

where \mathcal{S} is the set of all scenarios, \mathcal{J} is the set of shortfall types, P_s is the probability that scenario s occurs, c_j is a weight parameter allowing the decision maker to weigh the relative importance of different shortfall types, ρ_{st} is a path dependent discount factor⁸, depending on the scenario $s \in \mathcal{S}$ and the time period $t \in \mathcal{T}$, $x_{s jt}$ is the shortfall of type $j \in \mathcal{J}$, in time period $t \in \mathcal{T}$ under scenario $s \in \mathcal{S}$. The objective is minimized for different required levels of minimum expected portfolio returns.

Figure 6 shows the results of solving the first stage models (referring to step 1 in Section 3.1) for both approaches for different levels of minimum expected portfolio return. We see that in this *in-sample* comparison the dynamic approach clearly dominates the fixed mix approach since the target expected returns are achieved for lower levels of risk. For instance, for an expected return of 15.0%, the (square root of the) risk, \sqrt{F} , is 0.42 for the dynamic approach and 1.84 for the fixed mix approach. Observe that the fixed mix approach must choose a more aggressive portfolio (i.e. more stocks) than the

⁸ It is path dependent because it depends on short term interest rates, which are stochastic.




Fig. 7. Results from the simulation procedure. Each point on the efficient frontiers are found by solving 600 optimization programs, one for each period in each simulation.

dynamic approach to achieve the target expected in-sample return.

Figure 6 does not provide a fair comparison since it does not take into account that the models will be rerun in the future. The fixed mix approach suffers under such assumptions due to the lack of dynamic decision making. To make a fair comparison the testing procedure in Section 3 is applied and 200 simulation scenarios of three periods are generated as described in Section 4.1.3. The reason for not increasing the sample size is that a single run (of which there are 8 in Figure 7) involves solving 600 stochastic programs and requires more than 10 hours solution time on a Sun Ultra 2 (with 200 simulation scenarios of three periods).

Figure 7 shows the results of this out-of-sample comparison and we see that the dynamic approach still dominates the fixed mix approach, but to a smaller degree.

Due to the apparent small difference, statistical tests are performed. Note that the expected returns from the fixed mix approach are generally higher than for the dynamic approach, but the dynamic approach yields lower risk. Thus it is not possible to test whether the dynamic approach has a higher expected return at a given level of risk, or that it gives lower expected shortfall costs compared to the fixed mix approach for a given level of expected return. The tests must compare risk adjusted expected values, such as a convex combination of expected return and risk.

Given that the dynamic approach gives lower expected return and risk, a *conservative* test, i.e. one that is not biased in favor of the dynamic approach, can be constructed using a convex combination where there is a relatively

large weight on expected return. The following utility function is employed:

$$\max U = \omega \text{Expected portfolio value} - (1 - \omega) \text{Risk} \quad (5)$$

where $\omega \in [0, 1]$ is the weight parameter governing the degree of risk aversion, and where $\text{Risk} = F$ is our measure of risk in this problem.

To test the hypothesis that the dynamic stochastic approach dominates the fixed mix approach, the following statistic is used:

$$z = \omega(V_{\text{Dynamic}} - V_{\text{Fixed mix}}) - (1 - \omega)(F_{\text{Dynamic}} - F_{\text{Fixed mix}}) \quad (6)$$

where V is sample return, and F is sample risk. The probability that the mean of z is positive corresponds to the probability of dominance and is found via bootstrapping (Efron & Tibshirani 1993), since the distribution of z is not normal. A bootstrap of z is a vector of the same length as z with elements picked at random (with replacement) from the elements of z . The probability of the mean of z being positive is estimated as counting the number of bootstrap z having positive mean and dividing by the number of bootstrap runs.

The test compares the performance of the two approaches for the pairs (DS2, FM2) and (DS3, FM3) displayed in Figure 7. There will be a weight ω for each pair. Weights are chosen such that utility maximization yields the points FM2 and FM3 respectively, given that one must choose a point on the fixed mix efficient frontier⁹.

The result of the test is shown in Table 4. The probability that the dynamic approach dominates the fixed mix approach is higher than 50%, however the difference between the performances is not statistically significant. If we choose a weight corresponding to the average of the slopes of the efficient frontiers at FM2 and DS2, we get a probability of dominance of 0.8650. For (FM3, DS3), we get $p = 0.8531$.

⁹ Utility maximization means that the marginal rate of substitution of risk for return equals the negative of the slope of the fixed mix frontier:

$$-\frac{dEV}{dF} = \frac{\partial U / \partial F}{\partial U / \partial EV} = -\frac{1 - \omega}{\omega}$$

where EV is the expected portfolio value and F is the risk along the frontier in Figure 7. Thus $\omega = 1 / (1 + \widehat{\frac{dEV}{dF}})$, where $\widehat{\frac{dEV}{dF}}$ is an estimate of the derivative of the fixed mix efficient frontier, e.g. at FM2. This means that we can find a weight that corresponds to a utility function whose maximum is attained at FM2. Since we choose the fixed mix frontier, the test results will not be biased in favor of the dynamic approach.

Table 4

Test of whether the dynamic approach dominates the fixed mix approach. A hundred thousand bootstrap runs were used. The p-value reported is the (bootstrap) probability of dominance. The number of observations is 200.

(DS2,FM2)	$\omega = 0.17645$	$\bar{z} = 0.04143$	$\hat{\sigma}_z = 0.8418$	p=0.7456
(DS3,FM3)	$\omega = 0.47974$	$\bar{z} = 0.01450$	$\hat{\sigma}_z = 1.976$	p=0.5251

The reason that the dynamic approach dominates to a smaller degree, is that it exploits information in the optimization scenarios that is not present in the simulation scenarios. When comparing the approaches in-sample, the dynamic approach has the advantage of the ability to perfectly adapt to the information given. In the out-of-sample case, the information given is no longer perfect, so an advantage of stochastic programming over the alternative has been removed. The performance of the fixed mix approach will be closer to the performance of the dynamic approach because in the simulations the asset mix is allowed to be adjusted after each stage. In other words, since we are using rolling horizon simulations and allow the “fixed” mix to change at each stage and in every state, the fixed mix model actually becomes relatively dynamic.

The simulation scenarios and the optimization scenarios are necessarily different, for example regarding the structure of the evolution of information. In the optimization scenarios the number of outcomes per stage increases exponentially, but in the simulation scenarios the number is constant at 200 for stages 2 and 3. In particular, the simulation scenarios do not contain so-called worst case scenarios where asset class returns are negative in all subsequent periods. Both the simulation scenarios and the optimization scenarios satisfy the specified statistical properties, but for the optimization scenarios, after the worst case outcome in the first and second period, one of the subsequent outcomes also has a low return. In the simulation scenarios there is no reason why a low return outcome should follow a low return outcome in the previous period. These worst case scenarios have a significant effect on the perceived risk of the portfolio approaches as seen in Figure 6, due to the quadratic nature of risk. Although the probability of these worst case scenarios is very small, the shortfall costs accumulated are so large that they affect the overall expected shortfall cost. This can be seen from the difference in the scale of risk in Figures 6 and 7, where the risk in the out-of-sample case is much smaller¹⁰.

Further comparing Figures 6 and 7 we see that the out-of-sample expected return of the dynamic approach is roughly the same as in-sample, while for the fixed mix approach the expected return is higher. This is because the

¹⁰ It is hard to judge whether it is the optimization scenarios or the simulation scenarios that are more realistic; both sets match the specified statistical properties including state dependencies. If the number of simulation scenarios were increased, a worst case scenario might occur, causing the risk in Figure 7 to increase.

fixed mix approach generally has a higher share of stocks than the dynamic approach. At the second and third stages in the simulation, the fixed mix approach will choose more aggressive portfolios in order to fulfill the return requirement. This explains why the efficient frontier for the fixed mix approach is shifted not only to the left, but also upwards from Figure 6 to Figure 7.

5 Conclusion and future work

In this paper we have compared the performance of two alternative versions of a portfolio model. The comparison is severely complicated by the fact that the portfolio selection process involves dynamic decision making under uncertainty, so particular attention has been paid to the design of the out-of-sample simulation test. The results show that the stochastic dynamic approach weakly dominates the fixed mix approach. We expect that the degree of dominance would increase if the number of stages in the decision model is increased, since three stages probably is too low to fully take advantage of dynamic decision making.

The applied simulation procedure takes into account that a new instance of the decision model will be rerun when new information is available. With the new information, a new description of the uncertainties is generated and the decision models are resolved. The procedure involves solving thousands of stochastic programs.

This work has not focused on numerical efficiency. However, the structure of the simulation program is ideal for parallel implementation. Increased numerical efficiency also enables testing more realistic models with more asset classes and decision stages, and allows the sample size to be increased. We leave this for future work.

Another future research area is to improve the fixed mix strategy by creating more dynamic decision rules, so that the model behaves more like the true dynamic model.

Comparing the robustness of the approaches with respect to errors in the specification of market expectations as well as errors in the discrete approximation of the distributions, would be another interesting extension of this work.

References

Cariño, D. R. & Turner, A. L. (1998), Multiperiod asset allocation with derivative assets, *in* W. T. Ziemba & J. M. Mulvey, eds, 'Worldwide Asset and Liability

Modeling', Cambridge University Press, Cambridge, U.K., pp. 182–204.

- Cariño, D. R., Myers, D. H. & Ziemba, W. T. (1998), 'Concepts, technical issues and uses of the Russell-Yasuda Kasai financial planning model', *Operations Research* **46**(4), 450–462.
- Efron, B. & Tibshirani, R. J. (1993), *An Introduction to the Bootstrap*, Chapman & Hall, New York.
- Fourer, R., Gay, D. M. & Kernighan, B. W. (1993), *AMPL: A Modeling Language for Mathematical Programming*, The Scientific Press, San Francisco.
- Høyland, K. & Wallace, S. W. (2001a), 'Analyzing legal restrictions in the Norwegian life insurance business using a multistage asset liability management model', *European Journal of Operations Research* **134**(2), 65–80.
- Høyland, K. & Wallace, S. W. (2001b), 'Generating scenario trees for multistage decision problems', *Management Science* **47**(2), 295–307.
- Kouwenberg, R. R. P. (2001), 'Scenario generation and stochastic programming models for asset liability management', *European Journal of Operational Research* **134**(2), 51–64.
- Kusy, M. I. & Ziemba, W. T. (1986), 'A bank asset and liability management model', *Operations Research* **34**(3), 356–376.
- Lurie, P. M. & Goldberg, M. S. (1998), 'An approximate method for sampling correlated random variables from partially-specified distributions', *Management Science* **44**(2), 203–218.
- Murtagh, B. A. & Saunders, M. A. (1983), MINOS 5.4 user's guide, Technical Report SOL 83–20R. Revised Feb 1995, Department of Operations Research, Stanford University.
- Perold, A. F. & Sharpe, W. F. (1988), 'Dynamic strategies for asset allocation', *Financial Analysts Journal* **44**(1), 16–27.
- Vassiadou-Zeniou, C. & Zenios, S. A. (1996), 'Robust optimization models for managing callable bond portfolios', *European Journal of Operational Research* **91**(2), 264–273.
- Zenios, S. A., Holmer, M. R., McKendall, R. & Vassiadou-Zeniou, C. (1998), 'Dynamic models for fixed-income portfolio management under uncertainty', *Journal of Economic Dynamics & Control* **22**(10), 1517–1541.

A Specifying market expectations

Expressing market expectations can be done in many ways. We have chosen to let the decision maker supply percentiles for the marginal cumulative probability distribution functions for all uncertain variables (this is how the asset

allocation managers in the life insurance company prefer to express the market expectations), see Figure A.1. An approximating cumulative distribution

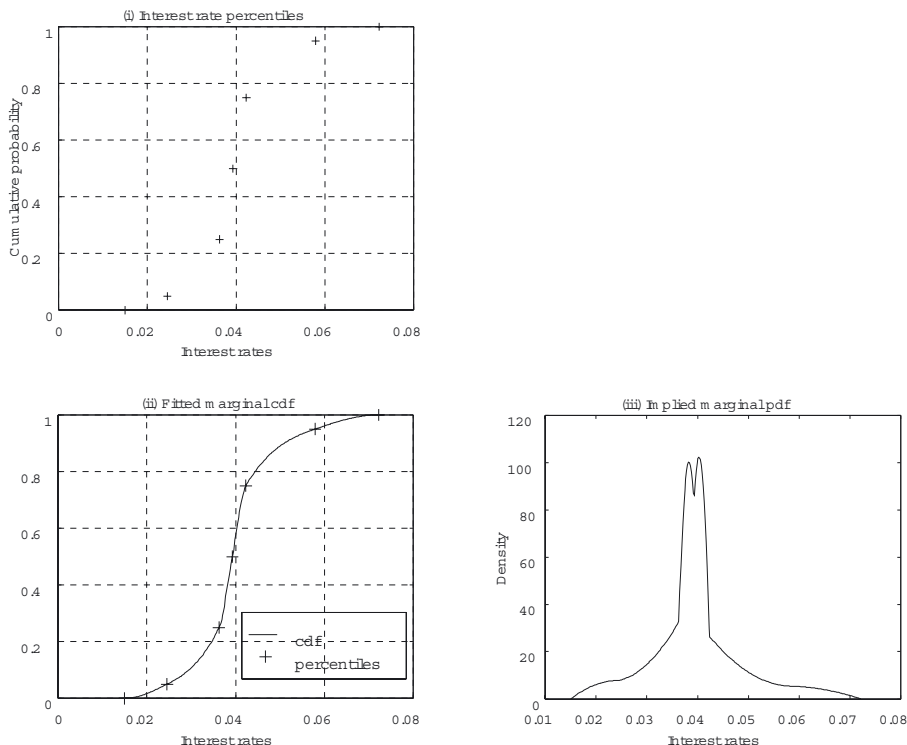


Fig. A.1. (i) User-supplied percentiles for interest rates at the end of the first period. (ii) A cumulative distribution function is fitted to the percentiles. (iii) The probability distribution function for interest rates.

function is fitted to these percentiles. The properties that are listed in Table 3 are calculated from the function that is fitted to the percentiles.

In addition, we let the user specify the correlation between all stochastic variables, and define the state dependent properties.

The approximating cumulative distribution functions are found using a NAG (Numerical Algorithms Group) C Library routine for interpolating data. This method does not guarantee that the second derivative changes sign only once, in the case of Figure A.1 causing a somewhat peculiar form near the top of the distribution. However, the resulting function is monotonic, so we are guaranteed that the curve will have the properties of a cumulative distribution function, and that the user specified percentiles are fit exactly (including the 0% and 100% points).

B Generating the scenarios that are input to the simulation

This section explains how simulation scenarios for interest rates and return on stocks are generated. The method works fast, so the time spent on this procedure is negligible compared to the actual simulation.

The simulation scenarios are the ones we use to test the performance of the two approaches, and they are the basis for generating optimization scenarios for later stages (but not for the first stage, as explained in the paper). They are sampled from the probability distributions that are fitted to the user supplied percentiles.

The stock returns are sampled first, using so-called inversion sampling. Uniform random numbers are sampled and input to the inverse of the cumulative distribution function to yield the random returns. Since we have specified a certain correlation between interest rates and stock returns, the interest rates are sampled conditional on the stock returns. The distribution function for interest rates conditional on a stock return represents a distribution with different mean and standard deviation, and it is found by means of a linear transformation of the percentiles. This transformation is found using the formulas for the conditional expectation and variance of a bivariate normal distribution. The skewness or kurtosis of the distribution is not changed significantly when the percentiles are adjusted, so this represents a feasible sampling process (An alternative way of sampling from marginal distributions with correlated random variables is given by Lurie & Goldberg (1998).)

The following notation is used:

ρ_{sb}	correlation between stock returns and interest rates
$x_s!$	a given return on stocks
μ_s	expected return on stocks
μ_b	expected interest rate
$\mu_{s b}$	expected interest rate given a return on stocks
σ_s	standard deviation of stocks returns
σ_b	standard deviation of interest rates
$\sigma_{b s}$	standard deviation given a return on stocks
$x_{qb,i}$	percentile of interest rates
$x_{qb s,i}$	percentile of interest rates given a return on stocks

The new mean, given the return on stocks, is

$$\mu_{s|b} = \mu_b + \rho_{sb} \frac{\sigma_b}{\sigma_s} (x_s! - \mu_s). \quad (\text{B.1})$$

The new standard deviation is

$$\sigma_{b|s} = \sigma_b \left(1 - \rho_{sb}^2\right). \quad (\text{B.2})$$

So the linear transformation of the percentiles is given by

$$x_{qb|s,i} = \mu_{s|b} + \frac{\sigma_b}{\sigma_{b|s}} (x_{qb,i} - \mu_b) \quad (\text{B.3})$$

for percentile i for interest rates, $x_{qb,i}$.

Applying this linear transformation to each percentile, and then creating a new distribution function yields a distribution that has the new mean and standard deviation, and unchanged skewness and kurtosis. The interest rate is then sampled from this distribution.

For the second period, the means and the standard deviations are updated for mean reversion and volatility clustering according to Equations (1), (2) and (3). The correlation between stock returns and interest rates is assumed constant for all periods. The above scheme is reapplied for period three.