

Dynamic asset allocation under uncertainty for pension fund management¹

by

Georg Ch. Pflug and A. Świętanowski

Department of Statistics, OR and Computer Science,
University of Vienna, Universitätsstr. 5,
A-1010 Wien, Austria

e-mail: georg.pflug@univie.ac.at, swietanowski@bigfoot.com

Abstract: Decision making in managing the asset and liability structure of a pension fund can be supported by stochastic dynamic optimization. We discuss our model, which is based on data analysis and forecast for the asset-side as well as a simulation model for the liability side.

The core of our decision support system consists of the following building blocks: a set of securities, a pricing module based on a multifactor Markov model to derive expected returns of securities, a simulation-based model for liabilities, a carefully chosen objective function suitable for the pension fund and a stochastic optimization problem solver. We consider the use of different objectives in the model and decomposition techniques to solve the stochastic portfolio optimization problem. Our final goal is to design an efficient parallel implementation.

Keywords: asset-liability management, pension fund management, financial modeling, stochastic dynamic optimization

1. Introduction

The growing importance of pension funds has boosted the need for methodologically sound principles for asset allocation. Whereas the economical side of pension fund management has been addressed by some authors (e.g., Haberman, 1994; Zimbidis and Haberman, 1993; Haberman, 1993; Dufresne, 1986), the pertaining decision problem, as a problem of optimization under uncertainty has not yet been discussed thoroughly in literature.

The characteristics of decision making for pension fund asset allocation are:

¹This research is part of the Special Research Program SFB F011 AURORA supported by

- the planning horizon is long (10 - 30 years);
- the liabilities are determined by the out-payments, which in turn depend on the mortality of the population;
- the rules of operation of pension funds are often complicated, since they must determine how investment gains or losses are distributed among the participants;
- in particular, through these rules, asset performance influences the stream of liabilities: bad asset performance allows the fund to reduce out-payments, whereas good performance leads to an increase;
- legal (risk limiting) and operational constraints restrict the possible decisions.

In this paper, we describe the AURORA optimization model for managing pension funds. The model consists of sub-models for the asset side, the liability side, allows to specify the objective and constraints and contains a solver for the large scale linear or nonlinear program. The software is being written in Fortran 90 and High Performance Fortran (HPF) for parallel execution.

2. Modeling the assets

The pension fund may invest in several asset categories, like national bonds, international bonds, national equities, international equities, etc.

We will take here the example of a large Austrian pension fund, which decides how much to invest into two asset categories:

1. national bonds,
2. foreign bonds and stocks.

After this decision is made, the further execution is passed to two portfolio managers (one for each category), who make the particular investments. Since the pension fund does not directly manage the variety of assets, we may assume that there are only two assets visible to the fund: asset 1 is the national bond portfolio and asset 2 is the other asset (stocks and foreign bonds).

Asset 1 has lower return and lower variability in comparison with asset 2. Fig. 1 shows monthly returns of the two asset categories in the last two years.

For the optimal allocation decision, the future possible developments of the assets must be modeled as a discrete stochastic process, in particular a discrete time discrete state Markov process. Since the computational complexity of the optimization problem is determined by the arc-degree of the transition graph, the number of successors of each state should be as small as possible. This is the reason why Markov processes with only two or three successors (birth-and-death processes) are popular. The simplest model is a random walk on the line with only two successors (the neighbors) of each state. We call this a binary lattice.

For each asset category, we construct a binary lattice which describes the future returns of this category. Let us briefly describe, how the lattice is estimated from the historical data




Figure 1. Monthly return of asset 1 (solid line) and asset 2 (dashed line).

The monthly returns $R_i^{(m)}$, which we have observed, are assumed to be i.i.d. and stem from a lognormal distribution, i.e., $\log R^{(m)} \sim N(\mu, \sigma^2)$. The parameters μ and σ^2 can be estimated from data. We are interested in yearly returns $R^{(y)}$, i.e.,

$$R^{(y)} = \exp\left(\sum_{i=1}^{12} \log R_i^{(m)}\right) = \prod_{i=1}^{12} R_i^{(m)}.$$

Clearly

$$\log R^{(y)} \sim N(12\mu, 12\sigma^2)$$

and by the well known exponential moments of the normal distribution

$$\begin{aligned} E(R^{(y)}) &= \exp(12\mu + 6\sigma^2) \\ E([R^{(y)}]^2) &= \exp(24\mu + 24\sigma^2). \end{aligned}$$

We want to find a two-point distribution $D^{(y)}$, which approximates the distribution of $R^{(y)}$. To this end, we have to find constants a and b , such that the distribution $D^{(y)}$ has the same first and second moments as $R^{(y)}$.

moments as $R^{(y)}$. The determining equations are

$$a = E(R^{(y)}) - \sqrt{E(R^{(y)2}) - [E(R^{(y)})]^2}$$

$$b = E(R^{(y)}) + \sqrt{E(R^{(y)2}) - [E(R^{(y)})]^2}.$$

Example. For the bond portfolio shown in Fig. 1, we have $\hat{\mu} = 0.0026$, $\hat{\sigma}^2 = 0.0000291$, which gives $a = 1.0130$, $b = 1.0516$, whereas for the stock portfolio we have $\hat{\mu} = 0.0077$, $\hat{\sigma}^2 = 0.0004253$ and $a = 1.0214$, $b = 1.1789$.

The pension fund makes investment decisions every 6 months. Therefore, the square roots of the factors a and b are the final modeled 6-month returns. The constructed lattice for asset category 1 is shown in Fig. 2. Note that the numbers at the nodes of the lattice represent the return *accumulated* from now to the time at which a given node appears.




Figure 2. Example: lattice of accumulated asset return over time.

3. Modeling the liabilities

The liabilities of a pension fund are determined by the total amount of pension to be paid to the beneficiaries in one period. To be more precise, we are interested in the net cash flow resulting from paying out the pensions and at the same time receiving contributions from the customers who continue to work. Both quantities depend on the demographic development of the corresponding risk group, where a risk group is defined as a group of individuals (numbering a thousand or more people) whose pension contributions are managed together and who share a common reserve capital.

Risk factors for the liability side are:

- longevity of the beneficiaries,
- insufficient number of new contributors entering,
- stopping of contribution due to economic difficulties of the contributors.

The risk contained in these uncertainties is also modeled by a binary lattice. However, construction of this lattice cannot be based on historical data, since these data are not available or, if available, are not very relevant. The best way of dealing with the liability risk factors is by simulation.

Let us take again the example of a typical Austrian pension fund. Each customer of the fund can be seen as being in one possible state of a discrete Markov chain. The customers randomly change state according to their demographic status.

Fig. 3 shows an example of a transition graph for customers. In each state, payments flow from the customer to the fund or vice versa. The transition probabilities as well as the (sex, age and position dependent) money flows between customer and fund can be estimated from historic data of the fund, from its operational rules and from general demography.

By simulating independently each of the present customers together with the possible generation of new customers, we may get a picture of the future distribution of total money flows, i.e. the liabilities.

The scenario lattice is then defined by taking the independent product of the three lattices: the lattice for asset category 1, the lattice for asset category 2 and the lattice for the liabilities. The full scenario lattice is an octal lattice: each node has $8 = 2^3$ successors. The size of the resulting history tree grows rapidly with the number of decision periods (see Table 1).

| periods | # nodes | # variables | # equations | sparsity |
|---------|---------|-------------|-------------|-----------|
| 1 | 9 | 81 | 45 | 0.337E-01 |
| 2 | 73 | 657 | 365 | 0.425E-02 |
| 3 | 585 | 5,265 | 2,925 | 0.532E-03 |
| 4 | 5,545 | 49,905 | 27,725 | 0.561E-04 |
| 5 | 67,273 | 605,457 | 336,365 | 0.462E-05 |
| 6 | 951,305 | 8,561,745 | 4,756,525 | 0.327E-06 |

Table 1. Sizes of the optimization problem for the underlying octal lattice.

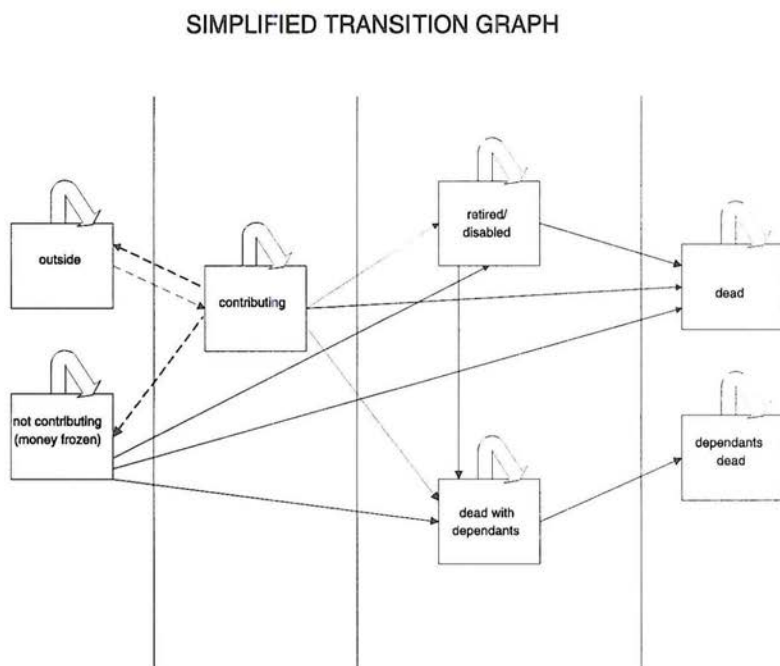


Figure 3. Pension fund customer state transition graph.

4. Problem specific objectives and constraints

The objective of a pension fund is to guarantee a high return on the participants' contributions. But the customers also depend on the pension fund to actually provide for their needs in the future. Therefore, safety of the portfolio is of paramount concern.

There are institutional and legal rules regarding pension fund operation in Austria. Within the constraints set out by the legislators, the managers of the pension fund are responsible for balancing between the increased return and increased security of their asset decisions.

4.1. Decision constraints

One decision aspect typical in pension funds (as well as, e.g., trust funds) is the presence of legal and organizational constraints on decisions. They are designed to restrict the risk of losses, which would adversely affect the pensioners. In the case of an Austrian pension fund the *legal* constraints require that:

- no less than 40% of all assets be held in Austrian bond funds and cash,
- no more than 40% of assets be held in equities and options,
- no more than 20% in real estate, etc.

Due to further details in the legal definitions, in the real life example we use throughout, up to 50% of assets may be held in equities.

Within those legal constraints, the fund managers of a particular risk group may decide to limit their decisions further by, e.g., requiring at least 60% of assets in Austrian bonds. Thus, the general portfolio management problem with random external cash flows (liabilities and customers' contributions) has to be extended to include such constraints. Typically, they can be written down in form of linear inequalities.

4.2. Objective functions

Since we are designing a decision support system, our main task is to present a *decision proposal* together with some representation of its possible consequences. We cannot simply perform one optimization and present the results to the pension fund management as the only *right* decision. Furthermore, the decision maker has to have the ability to influence the resulting decision, preferably by adjusting one or a few easily understandable parameters of the model.

In decision making under uncertainty the result of a decision can no longer be well represented by, say, a single number. Instead, a distribution (or density) of a random variable should be used.

We visualize the terminal wealth distribution by drawing it together with a box-plot (as shown in Figure 4). Such risk profile can effectively and intuitively

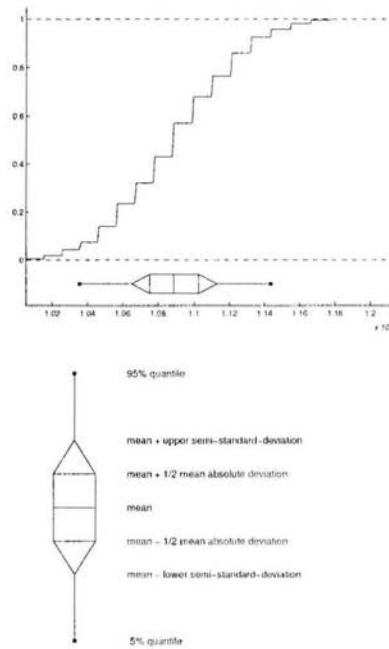


Figure 4. Graphical representation of terminal wealth distribution.

Since the decision maker wants and needs to see the risk profile as the result of the decision, his (or her) attitude towards risk has to be made a part of the objective.

In the following we shall shortly discuss several possible approaches to accurate representation of decision maker's goals by the model.

4.2.1. (Risk adjusted) expected terminal wealth

The operating rules do not allow the fund to make profit, since all profits from asset allocation have to be redistributed among the customers. However, a steady growth of fund's capital brings reputation and new customers to the fund. Therefore, at first it seems reasonable to optimize the total expected wealth of the fund at the end of the planning horizon, even though this is the wealth of the customers and not of the company.

Terminal wealth is a random variable, since it depends on the scenario. At each terminal node n of the scenario tree, the wealth variable $W(n)$ is known together with the probability $p(n)$ that this node is reached. An example below and Fig. 5 show how two decisions (resulting from two fixed mix strategies) are effectively compared.

Example. Following the fixed mix strategy, the portfolio is rearranged at each decision moment to meet the prespecified proportions of the asset categories. As a consequence we have to sell parts of a well performing asset category and buy a badly performing one.

The terminal wealth distribution with 50% and resp. 25% of the riskier asset under these strategies turned out to be as shown in Fig. 5. Expectedly, the larger proportion of the riskier but more profitable asset allows a higher expected terminal wealth, but also causes larger variation of results across scenarios.

Clearly, maximization of expected terminal wealth is not amenable to decision maker's manipulation, unless, e.g., some measure of his/her risk aversion is incorporated into the model (which thus becomes a so called *mean risk model*). Portfolio safety is not easily quantifiable. However, there exist quite a few formulations of the stochastic optimization problem objective that incorporate some risk penalty. Thus, the maximized objective could be composed of a sum of the expected wealth and the (negative) penalty for risk. Some of the possible objective function forms are:

- risk measured by mean absolute deviation (MAD):

$$E(W) - \rho E(|W - E(W)|) = \sum_n W(n)p(n) - \rho \sum_n |W(n) - \sum_m W(m)p(m)|p(n),$$

where $0 < \rho \leq \frac{1}{2}$ measures the degree of risk aversion.

The mean absolute deviation $E(|W - E(W)|)$ is a good measure of the risk associated with the decision. It is consistent with the expected utility theory.




Figure 5. Terminal wealth distribution with fixed mix strategies: proportion of the riskier asset is 25% and 50%, resp.

tic dominance (SSD) in a sense that a unique optimal solution of a problem with such objective is SSD efficient (see Ogryczak and Ruszczyński, 1999, for this and related results). Also, the upper bound on ρ can be increased and the above result can be substantially strengthened in the case of symmetric distributions.

The risk aversion factor ρ allows to adapt the objective to the specific needs of the decision maker. Moreover, MAD preserves linearity: if the terminal wealth W is linear in the decision variables x , then the objective is also linear in x .

- lower-semivariance as risk criterion:

$$E(W) - \rho \sqrt{E\{([W - E(W)]_{-\infty}^0)^2\}}.$$

where $0 < \rho \leq 1$ (hereafter $[\cdot]_a^b$ denotes projection on a closed interval (a, b)).

Again, such objective is consistent with SSD in the same sense as above.

Also, it is convex in the decisions x , if W is linear in x .

Note that since MAD is equivalent to two times mean absolute *semideviation*, both those risk measures can be seen as penalties for *downdside* risk only.

4.2.2. Objective with a target level of terminal wealth

An approach related to the terminal wealth maximization is to set the target terminal wealth \hat{W} and penalize underachievement

$$\max E\left(\varepsilon[W(T) - \hat{W}]_0^{+\infty} - [W(T) - \hat{W}]_{-\infty}^0\right),$$

where $\varepsilon > 0$ will encourage attempts to gain more than \hat{W} whenever possible. Manipulation of \hat{W} allows the decision maker to steer between solutions with lower and higher expected returns.

The resulting objective may be seen as nothing more than a concave (piecewise linear) utility function. However, using the simple formulation above (and its equally simple interpretation) we steer clear of the difficulties of identifying the decision maker's utility.

4.2.3. Disadvantages of a terminal wealth objective function

For individual investors, expected terminal wealth or risk-adjusted terminal wealth seems to be the right objective. For pension funds, however, it appears at least debatable. The main goal is, doubtlessly, to guarantee a secure and steadily growing outpayment of individual pensions.

Let us have a closer look at the operational details. The wealth of the fund is at all times divided into two parts: capital and (bounded) risk reserve. The pensions are paid from and are proportional to the capital. The risk reserve is

- when the fund's earnings are poor (less than some return rate r_*), the reserve is used to make up for the losses, and when it runs out, the outpayments are reduced,
- high earnings (above some return rate r^* , $r^* > r_*$) allow to fill the reserve up to its upper limit; any earnings above that are immediately added to the capital and cause increased outpayments.

It is easy to see that in boundary situations, i.e., when the reserve is empty or full, increased outpayments reduce wealth while decreased outpayments increase it. Thus, under such circumstances maximization of wealth is exactly the opposite of what the pensioners and the fund management may wish for.

4.2.4. Objective with a target level of return

One possibility of avoiding the problems with mean terminal wealth objectives is to set the target return rate \hat{r} and penalize underachievement while encouraging earnings higher than \hat{r}

$$\max \sum_{t=1}^T E \left(\varepsilon [R_t(x_t) - \hat{r}]_0^{+\infty} - [R_t(x_t) - \hat{r}]_{-\infty}^0 \right),$$

where $\varepsilon > 0$. Note that in the above formulation we take into account the returns in all periods. This reflects the pensioners' desire for steady and predictable income throughout the time of taking pensions.

As it was in the case of target wealth, manipulation of \hat{r} allows the decision maker to influence solutions and their risk profiles.

4.2.5. Objectives based on aspiration and reservation levels

A related but slightly more advanced model takes the random returns $R_t(x_t)$ and compares them to the two levels r_* and r^* . Below, the decision maker's expected satisfaction (measured from 0 to 1) in all periods is maximized

$$\max \sum_{t=1}^T E \left(\frac{1}{r^* - r_*} [R_t(x_t) - r_*]_0^{r^* - r_*} \right).$$

More generally, any two return rates $r_1 < r_2$ can be used. The lower one represents the lowest acceptable performance, the upper one is the return rate we aspire to achieve under favorable circumstances. In multicriteria optimization such levels are known as *reservation* and *aspiration* levels, respectively. Varying them, the decision maker may easily influence the optimal solution. It is important to note that their interpretation is clear to the decision maker.

This leads to a non-concave problem. The single period problem

$$\max E \left([\tau R^{(1)} + (1 - \tau) R^{(2)}] r^* \right)$$

is non-concave in general. Its solution satisfies

$$E \left(\mathbf{1}_{\{r_* \leq xR^{(1)} + (1-x)R^{(2)} \leq r^*\}} (R^{(2)} - R^{(1)}) \right) = 0.$$

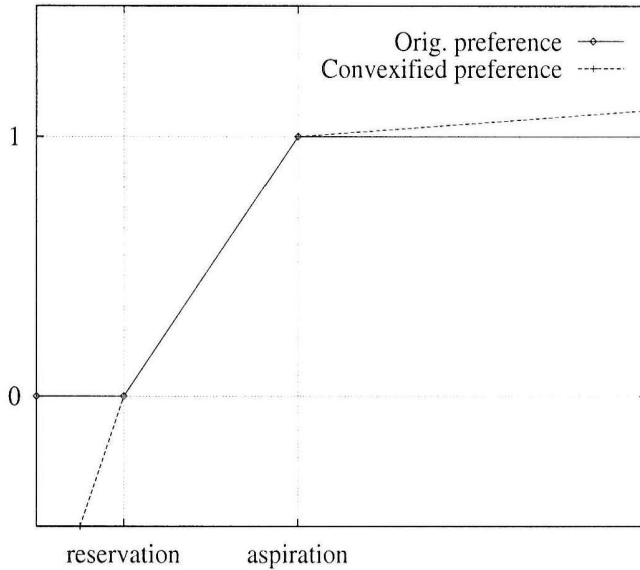


Figure 6. Objective function based on aspiration and reservation levels.

Following the common practice, we may replace the function $x \rightarrow [x]_a^b$ with a concave function

$$\Psi_{a,b}(x) = \begin{cases} a - (1 + \varepsilon)(a - x) & x < a \\ x & a \leq x < b \\ b + \varepsilon(x - b) & b \leq x. \end{cases}$$

for some $\varepsilon > 0$ (see Fig. 6 for illustration). This formulation makes the objective piecewise linear concave and is usable even in a linear program.

4.2.6. Empirical comparison of solutions with different objectives

In Fig. 7 and Table 2 the risk profiles for decisions resulting from three different strategies are presented. Again, returns accumulated over the whole planning horizon are compared.

A fixed mix strategy (with 50% of each asset) is clearly neither the safest nor the most profitable with respect to the expected terminal wealth maximization. The risk adjusted MAD strategy (even with a very small risk aversion factor $\alpha = 0.01$) offers a better performance than the fixed mix strategy. The MAD strategy

| Decision type | 5% quantile | Exp. val. | 95% quantile |
|------------------|-------------|-----------|--------------|
| Fixed mix | 1012.5310 | 1074.9568 | 1139.8300 |
| MAD risk measure | 1024.6790 | 1060.0931 | 1096.0450 |
| Target return | 1022.5470 | 1085.2518 | 1150.3970 |

Table 2. Comparison of some solutions (values in thousands of Austrian schillings).




Figure 7. Comparison of risk profiles for decisions taken with different objective functions (values in thousands of Austrian schillings).

The target following optimal strategy with a high target of ATS 1.12 billion in the last period fights for gains quite aggressively (and falls short of the target in most scenarios).

The mentioned examples are just a small sample of all possible comparisons that a decision maker may have to make. Choice of the objective function and its parameters (risk aversion factors, target levels, etc.) is an important part of the decision process. Visual representation of the crucial characteristics of the proposed solutions is necessary for the decision support system to be an acceptable tool for pension fund asset management.

5. The stochastic dynamic optimization problem

On the level of abstraction appropriate for considering the optimization methods, all uncertainties are treated in a uniform manner. We have a discrete time, discrete state stochastic vector process $\mathbf{X}(t)$, $t = 1, 2, \dots, T$ as the process of uncertainties. This vector process models all economic events which are the source of uncertainty and risks for the pension fund management, which in our case are asset returns as well as random contribution and liability streams.

With the process $\mathbf{X}(t)$ we associate the history process

$$\mathcal{X}(t) = (\mathbf{X}(1), \mathbf{X}(2), \dots, \mathbf{X}(t)).$$

Since the process \mathbf{X} has finitely many states, one may arrange all states of the history process in a finite tree, called the scenario, or history, tree (see Fig. 8). Its root is the starting state $\mathcal{X}_1 = \mathbf{X}(1)$. The nodes of the tree may be numbered $n \in \mathcal{N} = \{1, \dots, N\}$ so that each node number n corresponds in a one-to-one manner to a history of the process $\mathbf{X}(\cdot)$ up to the time at which this node occurs. The (unconditional) probability of reaching node n is denoted by p_n .

At each node of the tree a decision is taken. The depth of each node corresponds to the time for this decision. The decision is how much assets should be bought and sold in order to be always able to meet the liabilities (pay the pensions).

The linearly constrained stochastic dynamic optimization problem *with decomposable objective* can be expressed most directly using the tree structure:

$$\begin{aligned} \min \quad & \sum_{n \in \mathcal{N}} p_n f_n(x_n) \\ \forall (n \in \mathcal{N}) \quad & \begin{cases} T_n x_{p(n)} + A_n x_n = b_n \\ x_n \in X_n, \end{cases} \end{aligned} \tag{1}$$

where $p(n)$ denotes the predecessor of node n , $\forall (n \in \mathcal{N})$ $A_n \in \mathbb{R}^{m_n \times n_n}$, $T_n \in \mathbb{R}^{m_n \times n_{p(n)}}$. To avoid treating the root node as a special case, we define $X_{p(r)} = \{x_{p(r)}\}$ and $x_{p(r)} = \text{const}$. For all practical purposes $x_{p(r)}$ may be treated as an initial state and the constant term $T_r x_{p(r)}$ may be subtracted from the right hand side of the constraint. Copyright © 2008 John Wiley & Sons, Inc.




Figure 8. An example of a scenario tree with three periods.

If functions $f_n(x_n)$ are linear, i.e., $f_n(x_n) = c_n^T x_n$, and the decomposable constraints $x_n \in X_n$ have a simple form like, e.g., $x_n \geq 0$, then the problem (1) becomes a large scale linear problem.

In the case of non-decomposable objectives, e.g., in the mean risk models, the formulation is slightly different:

$$\begin{aligned} \min \sum_{n \in \mathcal{N}} p_n f_n(x_n) + F(x_i : i \in \mathcal{M} \subseteq \mathcal{N}) \\ \forall (n \in \mathcal{N}) \quad \begin{cases} T_n x_{p(n)} + A_n x_n = b_n \\ x_n \in X_n. \end{cases} \end{aligned} \tag{2}$$

Although the change seems very small it has wide ranging consequences for the optimization procedure. The tree-oriented expression is no longer as easily amenable to decomposition as the form (1) is.

Adding the non-decomposable term $F(x_i : i \in \mathcal{M} \subseteq \mathcal{N})$ to the objective creates an implicit all-to-all logical link between all nodes $i \in \mathcal{M}$. E.g., in the mean risk models \mathcal{M} is typically the set of terminal nodes of the tree. In a case of decomposition methods this forces a lot of additional synchronization of results and in fact might require major rethinking of those methods. In direct solution algorithms, like the interior point method mentioned in the following section, it causes unwelcome structural changes of the problem that also add a lot to the practical difficulty of solution.

6. Selected optimization methods

It is neither possible nor desirable to present or even name all optimization methods that can be applied to the problem at hand. As this part of work is still in progress we shall outline only those methods that are considered candidates for implementation and application within the AURORA Financial Management System.

The huge size (see Table 1 for a calculation of dimensions of a real life problem) as well as computational effort needed to solve the realistic optimization problems of this class call for solution methods which will exploit to the limits the most modern and powerful computers, especially the parallel ones. Thus, methods amenable to efficient parallel implementation are in the focus of our attention. Again, detailing the methods for parallel implementation is beyond this paper’s intended scope.

At present all optimization methods listed below are either tested for suitability to our problem, or in some stage of development.

As noted in the previous section, introduction of risk terms causes significant and so far unresolved difficulties for all of the optimization methods currently considered. In the discussion below we will therefore consider only the decomposable models:

- expected terminal wealth objective (with no risk term),
- objective with a target level of terminal wealth

- objective with a target level of return,
- objective based on aspiration and reservation levels.

6.1. Direct solution by an interior point method

The literature of interior point methods (IPMs) already includes a large amount of excellent reviews, monographs and even textbooks. We shall therefore refer the reader to Wright (1997) for a thorough introduction of IPMs, while we shall only remark briefly on one aspect of efficiency in the context of pension fund asset liability management problem. See also, e.g., Birge and Qi (1990), Birge and Holmes (1992), Gassmann (1991), Czyzyk et al. (1994), Ruszczyński (1993a), Berger et al. (1991) and references therein for more information on IPMs in the context of stochastic programming.

Thanks to a special structure of our problem (1), a primal-dual IPM may be successfully used for solving rather large problem instances. In IPMs the efficiency of a symmetric (possibly semi- or quasi-) definite factorization of a so-called normal matrix determines the efficiency of the whole method. It turns out that our particular constraint matrix structure is very favorable when employing the primal normal equations approach. Fig. 9 presents an example of a constraint matrix and the corresponding normal equations matrix. It is easy to see that the normal matrix is relatively sparse. In fact, both the formation of the normal matrix and its factorization turn out to cause very little fill in, which is a precondition for a successful and efficient solution of the problem. Even without further specialization of the algorithm relatively large problems were already solved successfully. Consult Dackner et al. (1998) for the precise statement of the optimization problem.

6.2. Augmented Lagrangian decompositions

All the augmented Lagrangian decomposition methods mentioned below are based on the same principles:

- the idea of relaxing inconvenient constraints (in our case — those linking decomposable subproblems) and introducing, instead, a form of penalty for their violation,
- Lagrangian augmentation which enables the use of a simple iterative method, the so called multiplier algorithm, for coordination by means of adjustments of penalties (Bertsekas, 1982).

Detailed statements are of necessity left out of this work. The reader may wish to consult Ruszczyński (1995) for a full account of the methods, including some discussion of parallel implementation issues.

6.2.1. Scenario decomposition

Scenarios (distinct paths from the root to one of the terminal nodes) are the units of decomposition. They are linked by the so-called non-anticipativity constraints

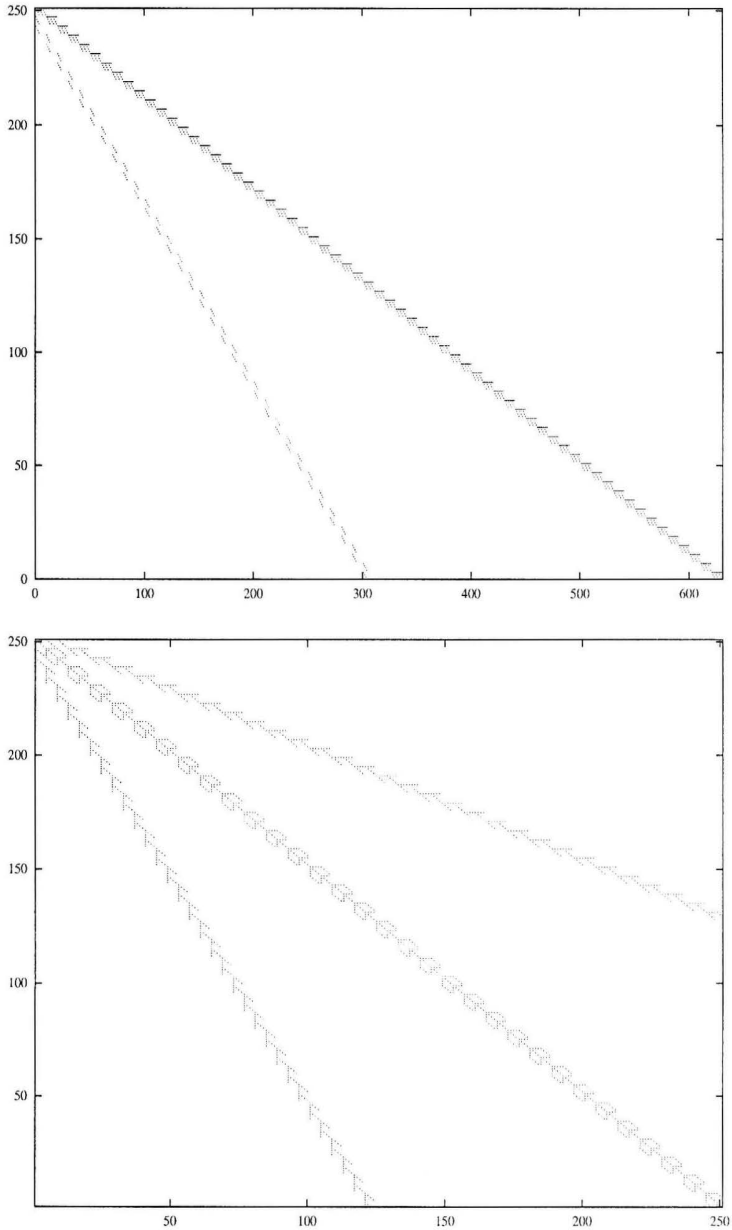


Figure 9. The structures of the global constraint matrix and the normal system matrix for a model based on a six stage binary tree.

which require that decisions in two scenarios should remain the same throughout all those periods during which the scenarios are indistinguishable. See Fig. 10 for an illustration. Thus, the stochastic problem can be defined as optimization of decisions for separate scenarios with an additional non-anticipativity constraint. This constraint is subsequently relaxed and penalized.




Figure 10. Scenario tree example from Figure 8 expanded into separate scenarios linked by nonanticipativity constraints.

6.2.2. Node decomposition

In this approach the tree structured problem (1) is seen as a collection of separate node problems with local constraints $x_n \in X_n$ linked by the transfer of decisions x_n along the branches. The constraints $T_n x_{p(n)} + A_n x_n = b_n$ which define the dynamics of the decision process, also link the subproblems. They are relaxed by placing them in the augmented Lagrangian.

6.3. Nested Benders or regularized decomposition

Nested Benders (possibly regularized) decomposition (Benders, 1962; Birge, 1985; Ruszczyński, 1992b) is another node oriented decomposition method for

solution of our tree structured problem (1). At each nonterminal node n a synchronization function $Q(x_n)$ is appended to the original objective. It represents the lower bound on the “cost to go”, i.e., the expected objective value increase after applying the current decision x_n as the initial state in all the subtrees starting at successors of node n . In subsequent iterations the approximation $Q(x_n)$ is improved until it is locally accurate and the global minimum is attained.

The proposed value of x_n is passed from the current stage to the next one. Given x_n , the next stage subproblem find the best decisions x_k for all k such that $p(k) = n$ and pass back the price (dual) information that allows to refine $Q(x_n)$. When this process is viewed globally on the scenario tree, the solution process passes in waves from one stage to another and synchronizes between the stages. There are many possibilities of directing the flow of currently solved node subproblems, however certain amount of synchronization is always required.

Nested regularized decomposition (Ruszczynski, 1993b) develops ideas of nested Benders decomposition (Birge, 1985) and two the stage regularized decomposition (Ruszczynski, 1986). Unlike its predecessors, it allows asynchronous parallel execution of both master and slave problems at all nodes of the tree, thus greatly diminishing the scalability concerns caused by the need to synchronize for information passing. A complete description of the solution method may be found in Ruszczynski (1993b).

7. Conclusions

We have shown how the management of a pension fund can be based on stochastic dynamic optimization. The quality of the decisions found by the optimization algorithms depends heavily on the accuracy of the stochastic scenario model of the future. The octal scenario tree presented in this paper seems to be relatively simple (due to rather coarse discretization) but already leads to very large scale optimization problems. The future development lies in finer modeling (with possible inclusion of expert opinion in the scenarios) together with high performance parallel computing.

Objective functions equipped with understandable and user-tunable parameters (like risk aversion, target level of wealth, aspiration level for return, etc.) allow the decision maker to learn about the decision problem at hand and adjust the optimization problem according to his/her risk preferences, which is especially important in the setting of portfolio management of a pension fund. By varying one or two parameters, a variety of solutions can be produced and compared with each other as well as with the outcomes of any fixed mix strategy or other *ad hoc* strategy. The comparison of a whole distribution of solutions is facilitated with an easy to understand graphical representation (the box plot).

It must be understood that we present a decision *support* tool and not a decision *making* tool: the final decision must always be taken by a responsible manager. Our tool's major task is to aid the decision maker's understanding of

References

- BENDERS, J.F. (1962) Partitioning procedures for solving mixed-variable programming problems. *Numerische Matematik*, **4**, 238–252.
- BERGER, A., MULVEY, J. and VANDERBEI, R.J. (1991) Solving multistage stochastic programs with tree dissection. Technical Report, Statistics and Operations Research, Princeton University.
- BERTSEKAS, D.P. (1982) *Constrained optimization and Lagrange multiplier methods*. Academic Press, New York.
- BIRGE, J.R. (1985) Decomposition and partitioning methods for multistage stochastic linear programs. *Operations Research*, **33**, 989–1007.
- BIRGE, J.R. and HOLMES, D.F. (1992) Efficient solution of two-stage stochastic linear programs using interior point methods. *Computational Optimization and Applications*, **1**, 245–276.
- BIRGE, J.R. and QI, L. (1990) Computing block-angular Karmarkar projections with applications to stochastic programming. *Management Science*, **34**, 12, 1472–1479.
- CZYZYK, J., FOURER, R. and MEHROTRA, S. (1994) A study of the augmented system and column-splitting approaches for solving two-stage stochastic linear programs by interior point methods. Technical Report 93–05, Industrial Engineering and Management Sciences, Northwestern University, Evanston.
- DOCKNER, E., MORITSCH, H., PFLUG, G.CH., ŚWIĘTANOWSKI, A. (1998) The AURORA financial management system. Technical Report AURORA TR1998-08, Vienna University.
- DUFRESNE, D. (1986) *Pension Funding and Random Rates of Return*. Volume 171 of NATO Adv. Sci. Inst. Ser. C: Math Phys. Sci. Reidel, Dordrecht.
- GASSMANN, H.I. (1991) Interior point methods in stochastic linear programming: Reducing fill-in. Technical Report WP–91–4, School of Business Administration, Dalhousie University, Halifax.
- HABERMAN, S. (1993) Pension funding: the effect of changing the frequency of valuations. *Insurance. Math. Econom.*, **13**, 3, 263–270.
- HABERMAN, S. (1994) Autoregressive rates of return and the variability of pension contributions and fund levels for a defined benefit pension scheme. *Insurance. Math. Econom.*, **14**, 3, 219–240.
- OGRYCZAK, W. and RUSZCZYŃSKI, A. (1999) From stochastic dominance to mean-risk models: Semideviations as risk measures. *European Journal of Operational Research*, **116**, 33–50.
- RUSZCZYŃSKI, A. (1986) A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming*, **35**, 309–333.
- RUSZCZYŃSKI, A. (1993a) Interior point methods in stochastic programming. Working Paper WP–93–8, International Institute for Applied Systems Analysis, Laxenburg.
- RUSZCZYŃSKI, A. (1993b) Parallel decomposition of multistage stochastic pro-

- gramming problems. *Mathematical Programming*, **58**, 201–228.
- RUSZCZYŃSKI, A. (1995) On convergence of an augmented Lagrangian decomposition method for sparse convex optimization. *Mathematics of Operations Research*, **20**, 634–656.
- WRIGHT, S.J. (1997) *Primal Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, Philadelphia.
- ZIMBIDIS, A. and HABERMAN, S. (1993) Delay, feedback and variability of pension contributions and fund levels. *Insurance. Math. Econom.*, **13**, 3, 271–285.

