# An asset liability management system for a Swedish life insurance company

Fredrik Altenstedt
Department of Mathemathics
Chalmers University of Technology
412 96 Göteborg, Sweden

October 8, 2003

**Abstract**

When managing assets used for life insurance, an investor must make a trade-off between the long-term goal of high return over the customers' lifetime, and the short-term goal of fulfilling government regulations. In order to facilitate making this trade-off we have developed an asset liability management model for the Swedish set of laws and regulations. The most significant difference between our model and other models from the literature is that in the Swedish case a company's actions will not only influence the asset allocation, but also the reserve requirements. As the problem is dynamic in nature, we employ multi-stage stochastic programming to optimize the asset allocation. In order to determine the usefulness of the system developed, we perform tests using rolling horizon simulations, where the stochastic programming approach is compared to a fix-mix benchmark. We further test different shapes of the scenario tree. The results from these tests suggest that the practice of aggregating stages in an ALM model in order to tackle problems with many stages will give the solution a significant bias.

**Keywords:** Stochastic programming; asset liability management; Optimization; Finance; Life insurance

## 1 Introduction

The basic retirement benefit system in Sweden is tax-funded. Each employer pays a fraction of the salary of their employees to the government retirement funds, which provide benefits for the retired work force. Private insurance offered by life insurance companies exists as a complement to the public system. On the Swedish market there exist different kinds of private pension insurances, but the only one treated in this paper is the traditional type, where the company carries the main part of the investment risk. Under such a traditional plan, the customer pays an amount each month, and upon retirement he/she receives a monthly sum depending on the total amount paid and the success of the company's investments. However, there is a lowest level of return that the company must pay to the customer; hence a traditional pension insurance policy may be viewed as a defined contribution plan with an element of a targeted money purchase plan.

In most applications of financial asset management the investor tries to choose investments giving good returns while minimizing the risk, where risk is measured as the variance or downside moments of investment returns. For an insurance company, such risk measures ignore the purpose of the investments, which is to be able to honor the company's liabilities. Hence, a more appropriate risk measure for an insurance company is the expected deficit, measuring the risk that the company does not own enough assets to cover the customers' claims when they occur. *Asset liability management* (ALM) deals with the unified administration of assets and liabilities, and may provide a strategy more suitable for a life insurance company than the separate management of assets and liabilities.

In this paper we develop an asset liability management system for a Swedish life insurance company. The system consists of a mathematical model of the company's assets and liabilities, a mathematical model of the surrounding economy and a stochastic programming solver capable of solving large instances of the model. In Section 2 we give an overview of stochastic models for ALM, with a focus on models used for retirement benefits. In Section 3 we provide an overview of the model of the company and the surrounding economy. Finally, we present tests of the model in Section 4, and conclusions are drawn from these tests in Section 5.

# 2 Overview of stochastic programming models for asset liability management

Ever since Dantzig [11] first considered extensions of ordinary mathematical programming problems to accommodate stochasticity in the data, a number of researchers have used stochastic programming to address problems for which decisions have to be made from incomplete information. In most of these models, the underlying stochasticity has a continuous distribution. As continuous distributions are hard to handle computationally, the distributions are discretized in order to obtain a finite number of possible outcomes. If this discretization is to accurately mimic the underlying continuous distribution, a large number of discretization points are needed. Hence the resulting problems tend to be large and require significant computational resources to find optimal solutions. This is especially true for multi-period problems, for which decisions and observations are interspersed; the size of such problems grows exponentially with the number of stages. This computational burden made applications of stochastic programming few and far in between until computing power became more ubiquitous and affordable.

As risk and the stochasticity underlying risk is most central to the financial industry a large number of stochastic programming models have been developed for financial applications. Bradley and Crane [2] developed a model for managing bond portfolios as well as a decomposition solver for the (at the time) large resulting stochastic linear programs. Kallberg, White, and Ziemba [21] created a two-stage model to manage short-term cash flows; the model was later extended to several time periods by Kusy and Ziemba [23]. Other uses of stochastic programming for financial problems include [6, 5, 7], where the authors use stochastic programming to improve the asset liability management of a Japanese insurance company, as well as a number of models dealing with ALM for pension funds and life insurance companies, described further below.

## 2.1 Swedish life insurance rules

The problem faced by a life insurance company is to optimally manage its assets while at all times complying with the current laws and regulations. Hence, a short description of these laws is needed in order to formulate the company's problem mathematically.

### 2.1.1 Rates of return

Under Swedish law there is a lowest rate of return that the holder of a life insurance policy must receive on his/her savings. This guaranteed rate of return is set by the regulating authorities; its level is however influenced by the long-term market interest rates. In addition to this lowest rate of return, excess funds are distributed among the customers using a bonus rate of return, which is set by the company. Hence, each contract may be viewed as an ordinary bank account accruing interest given by the bonus rate. This means that the company itself influences its future reserve requirements by setting their bonus rate. Since there is an element of guaranteed return, a Swedish life insurance policy may be viewed as a target money purchase plan. The policy is however not entirely consistent with such a plan, as the guaranteed rate of return may change during the plan's lifetime.

### 2.1.2 Reserves

In order to guarantee the solvency of a life insurance company, laws require that a life insurance company owns assets in excess of two reserve levels. The *retrospective reserve* is computed as the sum of all payments made to the company, increased with the bonus rate of return accrued so far. If a customer of age $x$ makes a one-time deposit of $D_0$ at time 0, the retrospective reserve $V$ for this deposit at time $t$ will be given by

$$V(t) = D_0(1-\rho)\exp\left(\int_0^t \delta + (\alpha\mu(x+\tau) - \beta)\, d\tau\right), \tag{2.1}$$

where $\delta$ is the bonus rate intensity, $\mu(x)$ is the death intensity of a customer (the probability that a customer will die in a short interval $t$ to $t+\Delta t$ is given by $\mu(x)\Delta t$), $\rho$ is a deposit fee and $\alpha$ and $\beta$ are security factors. When the reserves are calculated, an insurance policy is viewed as a collection of one-time payments, and the reserve requirement for the company is found by summing the reserve contributions from all payments to the company by their customers.

The *prospective reserve* is given by the present value of the lowest allowed future payments given by the guaranteed rate of return. In order to illustrate this reserve, we assume that we have a customer of age $x$ who has a policy giving a guaranteed continuous stream of 1 SEK per year. The stream will be paid out $m$ years from now and and continue for $r$ years, given that the customer is still alive. The present value of this stream will be

$$K(x, m, r) = \int_m^{m-r} e^{-\bar{\delta}}\tau P(T_x > \tau)\, d\tau,$$

where $\bar{\delta}$ is the interest rate intensity of the guaranteed rate of return offered to the customer, and $P(T_x > \tau)$ is the probability that the remaining lifespan of a person of age $x$ is larger that $\tau$. If our customer now makes a one-time deposit of $D_0$, this will buy her a guaranteed continuous stream of $g$ SEK per year, with $g$ being defined by

$$g := \frac{D_0}{(1+\rho)K(x, m, r)},$$

although once she reaches her retirement age, the actual payment stream received may be higher, as given by the bonus rate above. Here, again, $\rho$ is a deposit fee. The prospective reserve for this one time deposit will be

$$S = gK(x, M, r),$$

but now the present value $K$ has to be computed using the lowest allowed guaranteed rate of return. The lowest allowed guaranteed rate of return is set by the regulating authority *Finansinspektionen* (Swedish for the Finance Inspection), and it is affected by the long-term market rates. The promise made to the customer, $g$, is fixed and independent of the rate used to compute it, which means that if the guaranteed rate of return is changed, the present value of a future payment stream, $K(x, m, r)$ will change, and so will the prospective reserve. Hence there is a coupling between the asset prices and the liabilities via the long-term market rates.

### 2.1.3 Reserve requirements

The consolidation of the company is defined as the sum of the value of all assets divided by the retrospective reserve (defined above). The regulating authorities require that a life insurance company keeps its consolidation between 100% and 120%. These limits are not strictly enforced: consolidation levels below 100% are accepted for shorter periods of time. However, if this requirement is violated during long periods of time, or by a significant amount, authorities may require the company to retroactively lower the bonus rate of return. Retroactively reducing the bonus is perfectly legal, but very undesirable from a public relations point of view.

In addition to covering the retrospective reserve, the company is required to own assets in excess of the prospective reserve. Furthermore, just covering this reserve is not enough; there are limits on which type of assets may be used to cover the reserve. For example, a company may cover at most 25% of the reserve using equity. (This does not limit the amount that a company may invest in equity, as long as the company owns at least 75% of the value of the prospective reserve in other assets.) As the prospective reserve deals with guaranteed benefits, the requirements related to this reserve are enforced in a more stringent way than the requirements related to the retrospective reserve: breaking these rules may result in the liquidation of the company.

## 2.2 Comparison to models under foreign rules

In order to highlight the differences and similarities between the model presented in this paper and other models used for ALM, a short overview of rules and regulations of other countries is given next.

### 2.2.1 Dutch conditions

The Netherlands has a pension system based largely on funded assets. As a consequence, large assets are held by Dutch pension funds (an estimated 25% of all the pension assets in the EU, according to Boender [1]) and the massive funds held by Dutch funds have inspired much research into ALM solutions under the Dutch set of rules.

Dert [13, 14] has developed an asset liability model for a Dutch pension fund, with chance constraints regulating the probability of under-funding. The problem addressed is that of managing assets for a fund set up by a group of companies (the sponsors of the fund) in order to provide retirement benefits for the sponsoring companies' employees. The benefits are dependent on the employees' final and average salary, and the job of the manager of the fund is to keep the fund sufficiently solvent while keeping the contributions low and predictable. In order to do so the objective is to minimize the expected present value of the contributions required to keep the fund sufficiently solvent, with an added penalty for remedial contributions. The resulting model is a mixed integer linear problem, which is solved via a heuristic method. Dert models the economy via a VAR (Vector Auto Regressive) model, and the status of the beneficiaries of the fund is modeled by Markov chains.

In [17], Kouwenberg and Gondzio use the same framework, but replace the probability constraints by penalties for under-funding, and impose restrictions on the yearly variations of payments made to the fund. Again, the objective is to minimize the present value of the contributions, while penalizing remedial contributions.

The same framework is used once again by Kouwenberg [22], where different scenario generation models are examined. The model is identical to the one described by Gondzio and Kouwenberg [17] with the exception of applying a quadratic penalty on the remedial contributions.

The most notable difference compared to the Swedish case is the fact that the Dutch funds administer a defined benefits scheme (that is, participants receive a pension based on their final salary, and the contributions vary in order to achieve this), while the problem we address is mainly a defined contribution scheme (the contributions are defined by the customer, and the customer gets a pension dependent on the fund's development).

### 2.2.2 The Norwegian problem

Høyland and Wallace [20, 19] have treated the asset liability problem of a life insurance company in Norway. It is similar to the Swedish case: individuals save up for their own retirements, as opposed to the Dutch case, where the funding is provided by the employer. The case considered is hence a defined contributions scheme. The major difference from the Swedish setting is that all reserve requirements are considered to be exogenous variables (that is, they are not affected by the company's actions). Norway has a requirement of a lowest guaranteed rate of return on the assets, similar to the Swedish case. These rules do however differ from the Swedish ones as the guaranteed rate of return has to be given annually in Norway, as opposed to over the lifetime of the contract in Sweden. Requiring annual yields is a very limiting constraint, severely reducing the possible investments. Furthermore, Norwegian life insurance customers has the right to move their contract to a competing firm by paying a nominal fee (i.e. the contract contains a surrender option), a right the Swedish customers lack. Taking these rules into account, Høyland models the problem by constructing a multi-stage stochastic linear programming model, with a piecewise linear convex objective function.

#### 2.2.3 Other nationalities

Asset liability management problems for pension management have been addressed by a number of other researchers. Pflug et al. [25, 27, 26] consider the problem in an Austrian setting. The model is a multi-stage linearly constrained problem, with different objective functions in different versions. It is solved using interior point methods, taking advantage of the special structure of the constraint matrix for efficient factorization.

A similar problem for a British fund is addressed by Dempster and Consigli [8], who formulated the general purpose CALM model. This model differs from the other models described mainly by making it possible to add capital gains taxes, which requires the model to keep track of not only how much of each asset is owned, but also when it was purchased. Another difference is the absence of reserve requirements; liabilities need only be covered when they occur. The model is a multi-stage stochastic linear programming model, although versions with quadratic penalties exist.

An alternative to using multi-stage stochastic programming is explored by Brennan and Schwartz [3], who treat investment for index-linked policies with guarantees, a situation slightly similar to ours. In their framework, the payments made to the customers may be modeled as a combination of a guaranteed amount and a call-option on the index, with the guaranteed amount as the strike-price. under this framework, investment strategies are found by replicating the option. However, this framework is not very appropriate for our problem, as the company in our case influences the amount payed to a large degree, by setting the bonus rate of return.

## 3 The ALM model

In this section we describe the mathematical model of the company and the surrounding economy. In order not to explicitly denote all dependencies of random information, parameters and variables which vary between scenarios will be indicated with a boldface font.

### 3.1 Modeling of exogenous variables

We have modeled seven asset classes for the company to invest in, chosen based on our partner company's current asset classification. These assets are Swedish and foreign bonds, Swedish and foreign stock, Swedish treasury bills, Swedish real interest rate bonds, and Swedish real estate.

**Interest rates and bond prices** In our model of a Swedish life insurance company, the interest rates influence both the yield on bonds and the reserve requirements (see Section 2.1.1). This means that the interest rates have a high impact on which decision is optimal. Hence, bond prices and interest rates must be modeled in a consistent way. This is done by employing a version of the Brennan–Schwarts interest rate model [4]. This model has two state variables: the instantaneous interest rate and the rate of return on a console bond (a bond which never matures). These rates of returns are simulated by time-stepping a stochastic differential equation, and the price of any zero-coupon bond may be obtained by numerically solving the bond pricing equation (see Appendix B for details). This procedure is used to obtain prices of treasury-bills and bonds consistent with the interest rates.

A two-factor model is chosen since preliminary studies indicated that the one-factor models tested (the Vasicek model [9] and the Cox-Ingersoll-Ross model [10]) yielded too high correlations between long and short interest rates. As console bonds are not traded on the Swedish market, we have used the interest rate of the longest government bond on the Swedish market when fitting the model to market data. The parameter values obtained are given in Table 6 in Appendix A.

**Other asset prices** Inspired by the capital asset pricing method, we assume that the expected yield on assets other than Swedish bonds and treasury-bills are given by the risk-free rate of return to which a risk premium is added. The only exception is long-term foreign bonds where the expected yield is taken as the same as the expected yield on Swedish bonds, in effect assuming that the real long-term interest rate in and outside of Sweden are comparable. We generate the returns for these assets conditionally on the returns given on Swedish bonds and Swedish treasury bills, in order to get the desired correlation between the different asset classes. The parameters used for scenario generation in this study are given in Tables 7 and 8 in Appendix B.

The parameter values used in the simulation are a mixture of historical data and reasonable guesses, as the parameters used should not reflect the past, but the company managers' expectations about the future. Hence, historical data is used only to estimate correlations. In this article we only compare differences in performance between different configurations of the model itself. If the model is to be used commercially, or if a comparison is to be made versus our partner company's historical behavior, a better estimation of these parameters is needed. Especially, the values concerning Swedish real interest bonds and Swedish real estate are given by estimates from people within our partner company, and not based on real-world data. This is currently not a serious problem, since company policies prohibit trade in these assets, and hence trade in these assets is prohibited also in our model. The reason why these assets are not entirely excluded from the model is that these assets may still be used to cover the reserves, hence influencing which other asset allocations we may be interested in making.

**Customers and liabilities** As mentioned before in Section 2.1.3, the company needs to consider two reserves, one influenced by the bonus rate of return, and one entirely exogenous. A complication is that the retrospective reserve is not a linear function of the bonus rate, as may be seen from equation (2.1).

In order to be able to retain a linear ALM model, the calculation of the retrospective reserve needs to be slightly simplified. We do this by linearizing the retrospective reserve around an assumed value of the bonus rate. We believe that linearizing the reserve computations in this way will not greatly impact the solution, as the nonlinearity of the reserve is rather weak. As an example, we use one value of the bonus rate for the first 6 months, and another value for the next 18 months. In this case, the maximum difference between the linear approximation and the true reserve value will be less than 0.3% when the bonus rates are kept between 3% and 15%. (In previous trials, we have determined the reserves in an iterative fashion, by solving the ALM-problem in order to obtain the bonus rate of return for each node in the scenario tree, which have been used to update the expansion, after which the problem was resolved. The solutions obtained in this fashion did not significantly deviate from the solutions obtained without iterating.)

The prospective reserve is only influenced by the regulating authorities, and may be computed in a straightforward way. It is still dependent on the state of the economy, as the size of this reserve is determined by the guaranteed rate of return, which is influenced by the long-term interest rates.

For this study, we have used real-world data of the current stock of customers, augmented with our partner company's assumptions on the future inflow of customers. The customers have been aggregated according to year of birth, and the number of customers is assumed to be so large that the fraction of customers dying each year is exactly given by the mortality model.

A further assumption needed to make the model linear is that the customers do not react to the actions of the company, although the savings level of the customers may be related to other exogenous economic variables. It would be possible to treat the effects of customer behavior by linearizing how the consumer collective reacts to a change in the bonus rate of return, but as we have no data on the customers' reaction to a change in bonus rate, no such effort has been made.

**Implementation of the economy model** One of the goals when implementing the model of the economy was to make it easy to change or replace it if it was found to be insufficient, for instance by replacing the Brownian motions driving the model by a database of historical samples. In order to make this possible, the stochastic model is implemented in a black-box fashion as a pair of subroutines. The first of these routines take the two interest rates and a time length as input and return a pair of antithetic sample-paths of the desired time length, where the starting state is given by the two interest rates. The second routine returns the first four moments and the covariances of how the random variables develop over the given time length. If these moments are not explicitly available, they are found by simulation, by using a large number of samples from the first routine.

Antithetic samples are used as we wish to lower the errors in the sample means. As described earlier, the scenarios are generated using a stochastic differential equation, driven by an uncorrelated multivariate Brownian motion $X$. If we have a sample-path $x$ of this Brownian motion, used to generate a sample-path $y(x)$ of the economic variables, an equally probable outcome of $X$ is $-x$. We may use this sample to generate another sample-path $y(-x)$ from the distribution of economic variables. As $y(x)$ and $y(-x)$ are negatively correlated the sample mean gets a lower variance. The effects of using sample means for scenario tree generation have been investigated by Higle [18], who report that the method provided significantly better solution stability than ordinary random sampling.

## 3.2 Constraints

The model of the company is implemented as a multi-stage stochastic linear programming problem, having the following constraints:

**Time linking constraints** The development of the assets held are given by the time linking constraints, the assets owned in the first stage are given by what is held initially, corrected according to the first buy and sell decision:

$$\boldsymbol{x}_i^0 = \boldsymbol{y}_{+i}^0 - \boldsymbol{y}_{-i}^0 + \bar{x}_i, \quad i \in I. \tag{3.1}$$

Here, the different asset classes into which we may invest are given by $I$, and the assets held at time 0, $\boldsymbol{x}_i^0$, are given as the initial assets $\bar{x}_i$ to which we add what is bought, $\boldsymbol{y}_{+i}^0$, and deduct what is sold, $\boldsymbol{y}_{-i}^0$. In the same fashion, the asset inventory at later stages are given by

$$\boldsymbol{x}_i^t = \boldsymbol{y}_{+i}^t - \boldsymbol{y}_{-i}^t + (1 + \eta_i^t)\boldsymbol{x}_i^t, \quad i \in I, t \in T \setminus \{0\}, \tag{3.2}$$

where the price development of asset class $i$ since the last trading time is given by $\eta_i^t$.

In order to simplify the constraints regarding the reserves, we also introduce a variable giving the total value of what the company owns

$$\boldsymbol{x}_{tot}^t = \sum_{i \in I} \boldsymbol{x}_i^t, \quad t \in T. \tag{3.3}$$

**Cash balance** Naturally, the sum of all transactions in a period must add up to 0, which is guaranteed by the cash balance constraint:

$$\begin{aligned} &\boldsymbol{P}_{\text{in}} - \boldsymbol{P}_{\text{out}} - \boldsymbol{\theta}^t \boldsymbol{x}_{tot}^t \\ &+ \sum_{i \in I} \left( \boldsymbol{y}_{-i}^t (1 - \gamma_i) - \boldsymbol{y}_{+i}^t (1 + \gamma_i) \right) + \sum_{i \in I} \boldsymbol{\rho}_i^t \boldsymbol{x}_i^{t-1} = 0, \quad t \in T. \end{aligned} \tag{3.4}$$

Here, the payments to the customers, $\boldsymbol{P}_{\text{out}}$, the payments from customers, $\boldsymbol{P}_{\text{in}}$, the tax payments and the net of the transactions and direct income (given by $\boldsymbol{\rho}_i^t$) must add up to 0. The taxes paid by a life insurance company in Sweden is proportional to the assets owned and a benchmark interest rate, the state borrowing rate (a weighted average of long-term government bonds). According to Swedish rules, each year the company pays 15% of what they would have earned as capital gains, had their assets been invested at the state borrowing rate. This is modeled by the factor $\boldsymbol{\theta}^t$, which is defined as the length of the period leading up to time $t$ times 0.15 times the average state borrowing rate during this period. Furthermore, transaction costs for asset $i$ are given by $\gamma_i$, and transaction costs are assumed to be proportional to the amount bought or sold. Furthermore, the direct yield of asset $i$ during the period leading up to time $t$ is given by $\rho_i^t$.

**Bonus rate** The bonus rate given to the customers is the primary way in which the different life insurance companies compete. Hence the company whish to hold this rate high, and preferably even over time. Right now we do not enforce the requirement that the rate should be even, but only try to avoid low rates of bonus return. As having a life insurance policy should be an attractive alternative compared to using a high interest account, we define offsets from the long-term interest rate as $\Delta r_{\text{ref},a}, a \in A$, and add penalties for having bonus rates below these levels:

$$\boldsymbol{r}^t + \boldsymbol{z}_a^t \geq \boldsymbol{r}_{ref}^t + \Delta r_{\text{ref},a}, \quad t \in T, a \in A. \tag{3.5}$$

**Prospective reserve requirements** The rules for a life insurance company specify that it must at all times be able to cover the prospective reserve using the correct types of assets, as described in Section 2.1.3. There is a limit to how much of the reserve may be covered using assets of a specified type. For instance, a maximum of 25% of the reserve may be covered using stock. In order to capture these requirements, we introduce variables $\hat{\boldsymbol{x}}_i^t$ stating how much of each asset class is used to cover the prospective reserve, (naturally we have $\hat{\boldsymbol{x}}_i^t \leq \boldsymbol{x}_i^t$). Right now only constraints on two sets of asset classes are relevant for the model, but we formulate these requirements more generally should the company decide to expand the model later on (for instance by including corporate bonds). Here $K$ represents the set of rules. For each rule $k$ there is a set $I_k$ of assets affected by this rule, and a maximum fraction of the reserve, $c_k$, which may be covered by these assets. We hence get the upper bound on assets used to cover the reserve as

$$\sum_{i \in I_k} \hat{\boldsymbol{x}}_i^t \leq c_k \boldsymbol{S}^t, \quad t \in T, k \in K^t, \tag{3.6}$$

where $\boldsymbol{S}^t$ is the value of the prospective reserve at time $t$.

Failure to cover the reserve using the correct assets is a grave violation of the regulations, and will lead to the liquidation of the company. As liquidation is an alternative, although not a pleasant one, we do not strictly enforce this rule. We do instead add a steep penalty for failing to cover the reserve using the correct assets, a penalty given proportionally to the violation $\boldsymbol{z}_p^t$:

$$\sum_{i \in I} \hat{\boldsymbol{x}}_i^t + \boldsymbol{z}_p^t \geq \boldsymbol{S}^t, \quad t \in T. \tag{3.7}$$

As failing to cover the prospective reserve is a grave violation of the rules, almost failing to cover the prospective reserve is undesirable. We hence add a set of increasing penalties for almost failing to cover the reserve. In order to enforce avoiding violating this reserve requirements, we introduce a set of security levels $f_q, q \in Q$, where $f_q$ ranges from 1.15 down to 0.9. We add levels below 1 even though the company goes bankrupt if the condition is violated, as the company may ask for extra contributions from the owners in order to avoid bankruptcy, and these contributions should be held small. The degree to which the security levels are violated are given by $\boldsymbol{z}_q^t$, defined by

$$\boldsymbol{x}_{tot}^t + \boldsymbol{z}_q^t \geq f_q \boldsymbol{S}^t, \quad t \in T, q \in Q. \tag{3.8}$$

Note that this expression uses the total assets $\boldsymbol{x}_{tot}^t$, not the sum of $\hat{\boldsymbol{x}}_i^t$. If the assets $\hat{\boldsymbol{x}}_i^t$ were used in the expression (3.8) this would force us to cover 115% of the reserve *using the correct asset classes*, something which is not required by law.

**Retrospective reserve requirements** The retrospective reserve is defined in order to guarantee that the company is on track to live up to the bonus rate promises given to the customers, as described in Section 2.1.3. We model failing to cover the retrospective reserve in the same fashion as the prospective reserve, using progressively steeper penalties. Hence we define a set of lower limits, $\kappa_{\min,m}$, $m \in M$ for a set $M$ of penalties, and define the violation as previously by:

$$\boldsymbol{x}_{tot}^t + \boldsymbol{z}_m \geq \kappa_{\min,m} \boldsymbol{V}^t, \quad t \in T, m \in M. \tag{3.9}$$

In addition to preventing a too low consolidation, the regulating authorities wish a company to keep the consolidation under 120%, although this is not strictly enforced. (The idea is that gains made should be distributed to the policy-holders, and not retained in the company.) Hence we add a small penalty for having a too high consolidation as

$$\boldsymbol{x}_{tot}^t - \boldsymbol{z}_k \leq \kappa_{\max} \boldsymbol{V}^t, \quad t \in T. \tag{3.10}$$

**Trading restrictions** Our partner company has a policy of not actively trading real interest bonds and real estate. As for other assets, the company is small enough not to affect the market, and hence no trading restrictions are needed for other assets. However, for model generality, we still add upper bounds $u_i$ for the trade of all assets, and set maximum trade high for all assets but real estate and real interest bonds.

$$\boldsymbol{y}_{+i}^t \leq u_i, \quad i \in I, t \in T, \tag{3.11}$$

$$\boldsymbol{y}_{-i}^t \leq u_i, \quad i \in I, t \in T. \tag{3.12}$$

As the rules for life insurance companies does not allow short selling, we add lower limits for all assets.

$$\boldsymbol{x}_i^t \geq 0, \quad i \in I, t \in T. \tag{3.13}$$

**Expansions** The value of the retrospective reserve, and the sum paid to the customers, are affected by past values of the bonus rate of return. As mentioned previously, we linearize the reserve and the payments, giving us the following expressions:

$$\boldsymbol{V}^t = \bar{\boldsymbol{V}}^t + \sum_{\tau=0}^{t-1} \frac{\partial \bar{\boldsymbol{V}}^t}{\partial \bar{\boldsymbol{r}}^\tau} (\boldsymbol{r}^\tau - \bar{\boldsymbol{r}}^\tau), \quad t \in T, \tag{3.14}$$

$$\boldsymbol{P}_{\text{out}}^t = \bar{\boldsymbol{P}}_{\text{out}}^t + \sum_{\tau=0}^{t-1} \frac{\partial \bar{\boldsymbol{P}}_{\text{out}}^t}{\partial \bar{\boldsymbol{r}}^\tau} (\boldsymbol{r}^\tau - \bar{\boldsymbol{r}}^\tau), \quad t \in T \tag{3.15}$$

In this expression, $\bar{\boldsymbol{V}}^t$ is the value the retrospective reserve would have if the bonus rates at previous periods had been $\bar{\boldsymbol{r}}^\tau, \tau < t$. In the same fashion, $\bar{\boldsymbol{P}}_{\text{out}}^t$ is the payments to the customers computed using the assumed values of the bonus rate of return.

## 3.3 Objective of the optimization

According to a recently abolished law, all Swedish life insurance companies had to be mutual companies, dividing all their profits among the customers. As our partner company does not plan to change their status from being a mutual company, we may assume that this restriction still applies. Since the company is mutual, a reasonable view is that the customers own all the money inside the company. Hence it is appropriate to maximize the expected value of the assets held by the company, as it is the customers' money. To the value of the company, benefit payments made to the customers must be added. In order not to skew the results and favor high inflation scenarios over low inflation scenarios, all results are discounted using the rate of inflation. In addition to maximizing the present value of the participants' money, the management of the company must take into account the rules and regulations imposed by the authorities. This is captured in the model by penalizing the deviations from certain goals, as described above.

As expressed above, there are rules that the assets owned must cover two reserve levels, with increasing penalties for increasing violations. Hence the objective will be a concave function of the total assets owned, making the company act in a risk-adverse fashion.

Thus we get the objective function, shown in equation (3.16), as the expected present value of all payments made to the customers, plus the expected terminal value of the company's assets reduced with penalties for violating different rules and regulations. The objective of the optimization is to maximize

$$\begin{aligned} w(x,z,r) \quad &:= \mathsf{E}[\boldsymbol{d}^{\mathrm{T}}\boldsymbol{x}_{tot}^{T} + \sum_{t\in T}\boldsymbol{d}^{t}\boldsymbol{P}_{\mathrm{out}}^{t} + \sum_{t\in T} l^{t}\boldsymbol{d}^{t}(-s_p\boldsymbol{z}_p^{t} - \sum_{a\in A} s_a\boldsymbol{z}_a^{t}\bar{\boldsymbol{V}}^{t} - \\ & s_k\boldsymbol{z}_k - \sum_{m\in M} s_m\boldsymbol{z}_m - \sum_{q\in Q} s_q\boldsymbol{z}_q)]. \end{aligned} \tag{3.16}$$

In this expression, the parameter $l^t$ gives the length of the period following time $t$, which is used to scale the penalties. Penalties for the constraint violations defined in equations (3.5)–(3.10) above are given by $s_p, s_a, s_k, s_m, s_q$. Now the optimization problem may be stated as the maximization of (3.16) under the constraints (3.1)–(3.15).

# 4 Numerical tests

## 4.1 Rolling horizon simulations

Naturally, we wish to determine if our proposed model gives useful advice on how to run a life insurance company. As observing objective values does not give any information on how good a solution process is when used iteratively, we use rolling horizon simulations similar to Kouwenberg [22], Fleten, Høyland and Wallace [15], and Golub et al. [16], to test the performance of the method. The major difference between our tests and the ones performed in the cited articles is that they all treat tests performed where the number of stages in the test scenarios and the scenario trees are the same. (Kouwenberg uses a five period (six stage) tree whereas Fleten et al. use a three period tree). We use test scenarios with ten periods, forcing us to aggregate time-stages in the ALM model. This makes it possible for us to test if there is any performance difference between trees with few stages and many branches per stage, and trees with many stages but few branches at each stage.

Rolling horizon simulations work by applying the decision given by the optimization routine to our model of the company, letting some simulated time pass, an apply optimization again to obtain a new decision. The process is repeated and thus simulates how well the company actually fares when it is governed by the optimization model. These simulations are carried out as follows: First we generate a set of sample-paths which will be used for the evaluation, with all these sample-paths originating in a common state of the world. In order to run a scenario of length $T$, where the solution is evaluated each $n$ time-steps, we apply the following procedure to each path.

0: Set $t = 0$ and go to step 2.

1: Use the sample-path and the state of the company just after the decision at time $t-n$ to generate the state of the company just before a decision is made at time $t$.

2: Use the state of the world of the current sample-path at time $t$ to generate a scenario tree. Note that this tree is independent of the future realization of the sample-path, and hence that information does not leak into the scenario generation procedure.

3: Optimize over the tree, generating a decision for the company at time $t$. Use this decision to determine the state of the company after the decision is made at time $t$. Store information about the state of the company just after the decision.

4: if $t < T$, set $t = t + n$ and go to 1.

5: Use the stored states of the company to determine total penalties and the terminal value of the company, which are used to evaluate the success of a method on a particular scenario.

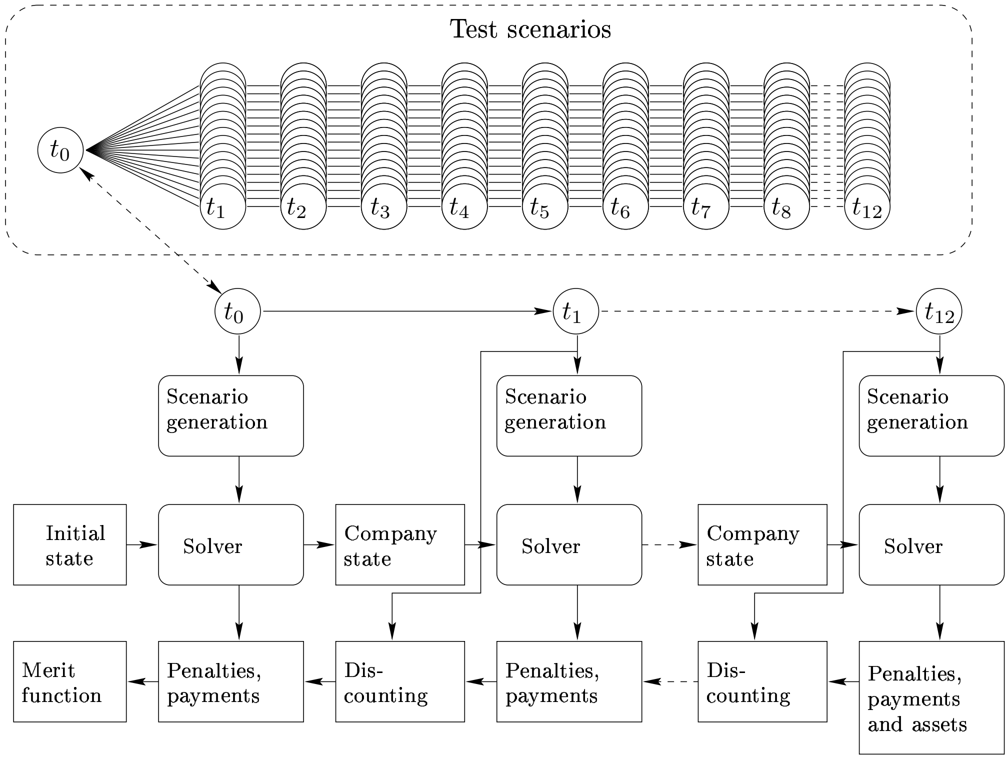


Figure 1: Testing system.

This procedure is repeated for each strategy tested, and the results are compared on a scenario to scenario basis. An illustration of the testing procedure is given in Figure 1.

The success of the method for a specific scenario is measured by taking the terminal value of the assets held by the company, to which we add all payments made to the customers, and subtract the penalties incurred. All values are discounted using the rate of inflation. Comparing the performance of two methods for one scenario will give us one sample of a random variable defined as the difference between using the two methods on a random scenario. As each scenario will give us one sample, we may use these samples to test whether the average value of the difference significantly differs from 0.

## 4.2 Scenario tree generation

In this work we employ two techniques for generating scenario trees: fitting and random sampling. When random sampling is used, the tree is generated recursively by sampling possible outcomes for the children of a node, using the model of the economy. In order to improve the sampling procedure, we use antithetic sampling and translate the samples obtained so that they will have the same expected mean as the sampled distribution.

When we do fitting the tree is again generated recursively, but by simultaneously optimizing the values of interest rates, asset prices and probabilities of the children of a specified node, in order to match the statistical properties of the underlying distribution. The variables fitted are the long and short interest rate, as well as asset prices for all assets except Swedish bonds and treasury-bills. Prices for Swedish long and short bonds are later given by the values of the interest rates. When using an optimized scenario tree generation, we fit the mean, variance, skewness and kurtosis of each variable, as well as the covariances between the different variables.

## 4.3 Optimization of fix-mix strategies

The current method used by our partner company to determine their operating strategy is hard to describe in mathematical terms. The board meets with even intervals to determine which asset mix the company should hold, and how high the bonus rate should be. The asset mix is given as a lower and upper bound on the fraction of the total wealth invested in different asset classes. When deciding on which asset mix to choose one naturally considers the consequences of keeping this asset mix for an extended period of time. Hence an approximation of this strategy is to choose the fixed mix of assets which will give the best yield over the period in question.

### 4.3.1 Fix-mix evaluation

In order to evaluate the effects of keeping one fixed asset mix, we add constraints to the model described in Section 3 and solve it. These constraints are

$$\alpha_i x^t_{\text{traded}} = x^t_i, \qquad i \in \bar{I}, \qquad i = 1, \ldots, T \tag{4.1}$$

where we use the following notation.

$x^t_{\text{traded}}$ total assets owned at time $t$, except assets invested in real estate and real interest bonds,

$\bar{I}$ Set of all assets except real interest bonds, real estate and Swedish bonds,

$\alpha_i$ fix-mix fraction of asset $i$, $\alpha_i \geq 0, \sum_{i \in \bar{I}} \alpha_i \leq 1$.

No prescribed fraction is set for Swedish bonds in order to avoid linearly dependent constraints. With this setup, Swedish bonds will absorb the remainder of the funds invested, whenever $\sum_{i \in \bar{I}} \alpha_i < 1$.

This extended model is solved using the same procedure as the ordinary ALM model. Naturally, there are more efficient methods to evaluate a fixed mix. Our method however has the advantage of using the same method to determine the bonus rate of return in both models.

The fix-mix asset fractions are optimized using a gradient descent algorithm. Gradients with respect to asset fractions are obtained from the dual variables corresponding to the constrains (4.1), in the manner described in Dantzig and Thapa [12] on page 196. These gradients may only be obtained if $x^t_{\text{traded}}$ is a basic variable. The total assets owned will most probably always be positive, and we indeed encountered no case for which $x^t_{\text{traded}} = 0$ occurred (which otherwise would have resulted in the evaluation being aborted).

Note that finding a fixed asset mix is not necessarily a convex problem (see Maranas et al. [24]), but in our case the method converges to the same solution for all starting points, when a large number of starting points are used.

## 4.4 Questions

In order to estimate the usefulness of the model, we ask a number of questions which we will try to answer using numerical experiments.

| Case | Split | Length, months | Size (rows/cols/scenarios) |
|---|---|---|---|
| Case 1, 2 | 30*10*10 | 6,12,24 | 164179/197823/3000 |
| Case 1, fix | 30*10*10 | 6,12,24 | 180834/201154/3000 |
| Case 3:1 | 10*8*8*8 | 6,12,12,12 | 294639/353883/5120 |
| Case 3:2 | 10*6*4*4*4 | 6,6,6,12,12 | 264079/316463/3840 |
| Case 3:3 | 8*6*4*4*4 | 6,6,6,12,12 | 211273/253183/3072 |
| Case 4, fix | 8*6*4*4*4 | 6,6,6,12,12 | 231718/257272/3072 |

Table 1: Scenario tree sizes.

As the current strategy of our partner company is similar to a fixed mix strategy, we would like to investigate if using dynamic asset allocation in the model will improve its performance. Although dynamic asset allocation have been found to outperform a fix-mix approach (See Kouwenberg [22] and Fleten, Høyland and Wallace [15]), we still perform these tests in order to make sure that this is true also for our chosen methods, as well as for establishing a reference value against which performance differences depending on scenario tree shape may be compared.

Other researchers (Høyland and Wallace [15], Kouwenberg [22], and Higle [18]) have made clear that the way scenario trees are generated is most important for the performance of the system. As the generation of scenario trees by fitting of the stochastic properties is rather expensive (creating a scenario tree takes significantly longer than solving the resulting problem), we would like to see if this difference is significant for our problem.

Finally we perform a study to determine if changing the shape of the scenario tree (making it longer but more narrow) will affect the performance.

#### 4.4.1 Fix-mix versus SP

In order to determine if the stochastic programming solution fares better than the fix-mix approach we apply both methods to 150 scenarios 5 years in length, with the portfolio being re-balanced every 6 months. All scenario trees for this test is generated by fitting the statistical properties of the underlying distribution. We report results from these simulations in Table 3 as Case 1. We report the difference and the standard deviation of the difference, as a large portion of the variability in the solutions is explained by differences between the scenarios. The scenarios are generated using antithetic sampling, in order to reduce the variance of the sampled mean, and we hence report two values for the standard deviation: the value obtained from all samples, and the value obtained when we average over each antithetic pair before calculating the standard deviation. If we assume that the two methods of generating scenario trees are equivalent, each scenario (or pair of scenarios) is a sample of the random variable defined as the merit function difference of the two methods used on a random scenario. As the number of runs is fairly large, we may assume that the mean of these variables has a normal distribution with standard deviation

$$\sigma_{\text{mean}} = \frac{\sigma_{\text{sample}}}{\sqrt{n}}$$

with $n$ being the number of samples. Observing that $0.8198/\sqrt{75} = 0.0947$ we see that the mean difference of the two methods equals approximately 1.6 times the standard deviation. More specifically we see that the probability of getting a larger deviation from 0 given that none exists is 10.8%, which is the strength of the test of our claim. The resulting asset fractions are reported in Table 2 together with their standard deviations (only the actively traded assets are reported.

| | Swedish bonds | Foreign bonds | Swedish stock | Foreign stock | Swedish t-bill |
|---|---|---|---|---|---|
| Fix-mix | 0.62/0.0 | 0.04/0.0 | 0.13/0.0 | 0.09/0.0 | 0.005/0.0 |
| SLP | 0.67/0.04 | 0.0/0.0 | 0.12/0.03 | 0.09/0.03 | 0.0/0.0 |

Table 2: Asset fractions owned in Case 1 (fraction/std.).

| Method | Mean (BSEK) | Std (BSEK) |
|---|---|---|
| Case 1 | | |
| SP | 22.2606 | 4.0778/1.9555 |
| Fix-mix | 22.1085 | 3.3445/1.4798 |
| Difference | 0.1521 | 1.1525/0.8198 |
| Case 2 | | |
| Fitted tree | 22.2606 | 4.0778/1.9555 |
| Random tree | 22.2641 | 4.179/1.88719 |
| Difference | -0.0035 | 1.2585/0.8981 |
| Case 3 | | |
| 30*10*10 | 22.2641 | 4.179/1.88719 |
| 10*8*8*8 | 22.3193 | 4.8865/2.1128 |
| Difference | -0.0552 | 2.4121/1.8296 |
| 30*10*10 | 22.2641 | 4.179/1.88719 |
| 10*6*4*4*4 | 22.6856 | 4.5646/1.9961 |
| Difference | -0.4215 | 2.1907/1.5004 |
| 30*10*10 | 22.2641 | 4.179/1.88719 |
| 8*6*4*4*4 | 22.5202 | 4.9785/2.1256 |
| Difference | -0.2561 | 2.4331/1.5880 |
| Case 4 | | |
| SP | 22.5202 | 4.9785/2.1256 |
| Fix-mix | 22.1168 | 3.4093/1.5139 |
| Difference | 0.3489 | 2.8009/1.7033 |

Table 3: Merit function values.

#### 4.4.2 Random versus optimized scenarios

Partly because the price of constructing scenario trees by optimization is rather high (constructing the scenario tree takes significantly longer than solving the resulting SLP with our implementation), we wish to determine if constructing scenario trees by fitting the properties of the underlying random distribution will give a significantly better performance compared to constructing trees by anthitetic random sampling with adjusted means. We run 150 scenarios using both randomly generated scenario trees and fitted trees. The results from this test are reported in Table 3, as Case 2. Using the same methods as earlier, we see that the probability of getting a larger deviation from 0, given that no difference exists, is 97%, and we may hence not draw the conclusion that a difference exists between the methods. Hence the rest of the trials will be carried out using scenario trees generated by random antithetic sampling with correction.

If we look at the asset fractions obtained by the two methods, as reported in Table 4, we see (not surprisingly) that the first-stage solutions obtained from randomly generated trees have a larger standard deviation. More important is the fact that there exists a statistically significant bias: when using a random tree, a larger fraction of the wealth is invested in stock.

| | Swedish bonds | Foreign bonds | Swedish stock | Foreign stock | Swedish t-bill |
|---|---|---|---|---|---|
| rand. tree | 0.64/0.08 | 0.005/0.02 | 0.13/0.5 | 0.11/0.05 | 0.0/0.0 |
| opt. tree | 0.67/0.04 | 0.0/0.0 | 0.12/0.03 | 0.09/0.03 | 0.0/0.0 |

Table 4: Asset fractions owned in Case 2 (fraction/std.).

#### 4.4.3 Tree size

In this study we have used rather wide and short trees compared to other studies, mainly to compensate for the larger number of assets. In order to see if this choice affects the performance of the method, we try other shapes of trees, which are more narrow and deep, as this kind of shape is more common in other researchers' tests. For instance, Dempster and Consigli [8], having 5 assets, uses a 10 stage tree divided as $7, 3, 2^7$ (the first stage splits into seven nodes, each splitting in three in the next stage, each splitting in 2 in the next seven stages). Kouwenberg [22], having 4 assets, uses a configuration of $10, 6^2, 4^2$ , and Høyland and Wallace [20], having 4 assets, uses a configuration of $60, 6^2$.

In order to test if changing the shape of the tree will affect the efficiency of our method, we test three sizes of trees. The trees of each size are all generated by antithetic random sampling with corrected means. The sizes of these trees are given in Table 1, as Case 3. The length of the stages in the trees are adjusted to give the tree the same overall length in time, as we otherwise may not rule out that a difference in performance stems from different time horizons. We use 3 different trees, one with five stages and two with six stages in addition to the base case with four stages. As may be seen from Table 3, the two longer trees used, having 6 stages, perform significantly better than the base case with 4 stages. Also well worth noting is that the trees with 6 stages both outperform the tree with 5 stages, despite having a lower total number of scenarios. If we perform the same statistical test as before, we see that the probability of getting a larger deviation from zero when comparing Case 3:3 to the base case is 16%, suggesting that using longer and deeper trees does actually improve performance. In this comparison the trees have essentially the same number of scenarios, 3000 vs. 3072, but the size of the deterministic equivalent problems differ, as a longer tree will give more variables and constraints. When comparing Case 3:1 and 3:2, the probability of obtaining a larger difference given that none exists is 8%, again suggesting that deeper and more narrow trees perform better. Worth noting in this case is that the longer tree has significantly fewer scenarios (3840 versus 5120), although the difference in the size of the deterministic equivalent is smaller.

The fractions of available assets invested into different asset classes at time-stage 0 is given in Table 5. As may be seen from this table, the longer but more narrow scenario trees yields solutions with a higher fraction of the wealth invested in stock. This means that aggregating time-stages in the scenario tree will give the solution a bias, and that this bias is large enough to make longer but more narrow trees outperform shorter and wider trees, despite a significantly larger instability in the solutions.

| | Swedish bonds | Foreign bonds | Swedish stock | Foreign stock | Swedish t-bill |
|---|---|---|---|---|---|
| base case | 0.64/0.08 | 0.005/0.02 | 0.13/0.5 | 0.11/0.05 | 0.0/0.0 |
| case 3:1 | 0.60/0.12 | 0.02/0.05 | 0.15/0.09 | 0.12/0.09 | 0.003/0.02 |
| case 3:2 | 0.58/0.12 | 0.03/0.08 | 0.18/0.09 | 0.10/0.08 | 0.002/0.02 |
| case 3:3 | 0.51/0.18 | 0.07/0.13 | 0.19/0.09 | 0.11/0.09 | 0.01/0.07 |

Table 5: Asset fractions owned (fraction/std.).

### 4.4.4 Fix-mix versus SP revisited

As the performance seems to improve when using deeper, more narrow trees, we again try the fix-mix version of our model, using the trees given in Case 3:3. This test is reported as Case 4 in Table 3. Here, the difference between the two methods is bigger compared to test Case 1, and the probability of getting such a big difference given that none exists is approximately 8%. This is not surprising, as being required to keep the same asset mix is more of a handicap in a deeper tree, which has a higher number of opportunities to re-balance the portfolio.

## 5 Conclusions

In this paper we have developed an ALM model for a Swedish life insurance company, including a model of the surrounding economy. Two procedures for constructing scenario trees from the economy model are implemented and tested. We perform rolling horizon simulations to compare the different scenario generation techniques, and compare if having different shapes of the scenario trees impacts the performance of the model, when applied to out of sample scenarios.

By comparing different shapes of the scenario tree, we see that trees with many stages but few branches at each stage clearly outperform shorter, wider trees, despite the fact that the solutions obtained using these trees are very unstable (the fractions invested in different asset classes vary significantly between different scenario trees). In fact, the gain in out of sample performance gained by going from a four stage scenario tree to a six stage scenario tree was significantly larger than the gain from going from a fixed-mix approach to a stochastic programming approach using four stages.

In this work we have used anthitetic random sampling with correction as the method to generate scenarios, as our tests indicated no difference in performance between the two scenario tree generation methods used. However, the solution obtained by using random-corrected trees not only had a larger variability in the solutions, these solutions also had a larger investment in stock than the solutions obtained using fitted trees. Hence random sampling will not only give the solutions a larger variability, it will introduce a bias as well. When comparing different shapes of the scenario tree, we found that long,narrow trees performed better, while investing even more in stock. Hence it seems as if the relatively good performance of the random sampling technique (compared to Kouwenbergs findings [22]) may be explained by the fact that the random sampling introduces a bias which counteracts the bias from using scenario trees with aggregated stages.

## 6 Acknowledgments

This work was done in cooperation with Nordea Life & Pension, who also partially funded the project.

## References

[1] G. C. Boender, *A hybrid simulation/optimisation scenario model for asset/liability management*, European Journal of Operational Research, 99 (1997), pp. 126–135.

[2] S. P. Bradley and D. B. Crane, *A dynamic model for bond portfolio management*, Management Science, 19 (1972), pp. 139–151.

[3] M. J. Brennan and E. S. Schwartz, *Alternative investment strategies for the issuers of equity linked life insurance policies with an asset value guarantee*, The Journal of Business, 1 (1979), pp. 63–93.

[4] M. J. Brennan and E. S. Schwartz, *An equilibrium model of bond pricing and a test of market efficiency*, Journal of Financial and Quantitative Analysis, 17 (1982), pp. 301–329.

[5] D. R. Cariño, D. H. Myers, and W. T. Ziemba, *Concepts, technical issues, and uses of the Russell–Yasuda Kasai financial planning model*, Operations Research, 46 (1998), pp. 450–462.

[6] D. R. Cariño and W. T. Ziemba, *Formulation of the Russell–Yasuda Kasai financial planning model*, Operations Research, 46 (1998), pp. 433–449.

[7] D. R. Cariño, W. T. Ziemba, T. Kent, D. H. Myers, C. Stacey, M. Sylvanus, A. L. Turner, and K. Watanabe, *The Russell–Yasuda Kasai model: An asset/liability model for a Japanese insurance company using multistage stochastic programming*, Interfaces, 24 (1994), pp. 29–49.

[8] G. Consigli and M. Dempster, *Dynamic stochastic programming for asset–liability management*, Annals of Operations Research, 81 (1998), pp. 131–161.

[9] J. Cox, J. Ingersoll, and S. Ross, *An intertemporal general equilibrium model of asset prices*, Econometrica, 53 (1985), pp. 363–384.

[10] J. Cox, J. Ingersoll, and S. Ross, *A theory of the term structure of interest rates*, Econometrica, 53 (1985), pp. 385–407.

[11] G. B. Dantzig, *Linear programming under uncertainty*, Management Science, 1 (1955), pp. 197–206.

[12] G. B. Dantzig and M. N. Thapa, *Linear Programming*, vol. 1, Springer-Verlag, New York, 1997.

[13] C. Dert, *Asset liability management for pension funds; A multistage chance constrained programming approach*, PhD thesis, Erasmus University Rotterdam, 1995.

[14] C. L. Dert, *A dynamic model for asset liability management for defined benefit pension funds*, in Worldwide Asset and Liability Modeling, W. T. Ziemba and J. M. Mulvey, eds., Cambridge University Press, 1998, ch. 20, pp. 501–536.

[15] S.-E. Fleten, K. Høyland, and S. W. Wallace, *The performance of stochastic dynamic and fixed mix portfolio models*, European Journal of Operational Research, 140 (2002), pp. 37–49.

[16] B. Golub, M. Holmer, R. McKendall, L. Pohlman, and S. Zenios, *A stochastic programming model for money management*, European Journal of Operational Research, 85 (1995), pp. 282–296.

[17] J. Gondzio and R. Kouwenberg, *High performance computing for asset liability management*, Operations Research, 49 (2001), pp. 879–891.

[18] J. L. Higle, *Variance reduction and objective function evaluation in stochastic linear programs*, INFORMS Journal on Computing, 10 (1998), pp. 236–247.

[19] K. Høyland, *Asset liability management for a life insurance company: A stochastic programming approach*, PhD thesis, Department of Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway, 1998.

[20] K. Høyland and S. W. Wallace, *Analyzing legal regulations in the Norwegian life insurance business using a multistage asset-liability management model*, European Journal of Operational Research, 134 (2001), pp. 293–308.

[21] J. Kallberg, W. R.B, and W. Ziemba, *Short term financial planning under uncertainty*, Management Science, 28 (1982), pp. 670–682.

[22] R. Kouwenberg, *Scenario generation and stochastic programming models for asset liability management*, European Journal of Operational Research, 134 (2001), pp. 279–292.

[23] M. Kusy and W. Ziemba, *A bank asset and liability management model*, Operations Research, 34 (1986), pp. 356–376.

[24] C. Maranas, C. Androulakis, C. Floudas, A. Berger, and J. Mulvey, *Solving long term financial planning problems via global optimization*, Journal of Economic Dynamics and Control, 21 (1997), pp. 1405–1425.

[25] G. Pflug, *Scenario tree generation for multiperiod financial optimization by optimal discretisation*, Mathematical Programming, Series B, 89 (2001), pp. 251–271.

[26] G. Pflug, A. Świętanowski, D. Dockner, and H. Moritch, *The AURORA financial management system: model and parallel implementation design*, Annals of Operations Research, 99 (2000), pp. 189–206.

[27] G. C. Pflug and A. Świętanowski, *Dynamic asset allocation under uncertainty for pension fund management*, Tech. Rep. TR1998–15, Vienna University, 1998.

# A Bond pricing

In our version of the Brennan–Schwartz two-factor model, two state variables are used: the return on a console bond (a bond that never matures), and the instantaneous rate of return. The movements of these returns are given by the stochastic differential equation

$$\begin{aligned} dr =& \alpha_r(l - s - r)dt + \sigma_r r dz_r, \\ dl =& \alpha_l(\hat{l} - r)dt + \sigma_l l dz_l, \end{aligned}$$

where

| | |
|---|---|
| $r$ | instantaneous interest rate, |
| $l$ | console rate, |
| $\alpha_{r,l}$ | mean reversion strength for the two processes, |
| $\sigma_{r,l}$ | standard deviation parameter, |
| $\bar{l}$ | mean reversion level of console rate, |
| $s$ | difference between long and short rate, |
| $\rho$ | instantaneous correlation of $dz_l$ and $dz_r$, |
| $\lambda$ | market price of short-term interest rate risk, |
| $c$ | continuous coupon payed by bond, |
| $\tau$ | time to maturity, |
| $B(l, r, c, \tau)$ | price of bond as a function of the state variables and bond properties. |

Bonds are priced by numerically solving the partial differential equation

$$\begin{aligned} &B_{rr}\sigma_r^2 r^2/2 + B_{rl}\rho\sigma_r r\sigma_l l + B_{ll}\sigma_l^2 l^2/2 + B_r(\alpha_r(l - s - r) - \lambda\sigma_r r) + \\ &B_l(\sigma_l^2 l\alpha_l^2(\hat{l} - l)^2 + l^2 - rl) - B_\tau + c - Br = 0, \end{aligned}$$

with the boundary condition $B(l, r, c, 0) = k$, where $k$ is the face value of the bond (stating that the bond must yield its face value at maturity). For a derivation of this formula, see [4].

# B Data

| | |
|---|---|
| $\alpha_r$ | 1.2492 |
| $\alpha_l$ | 0.1884 |
| $\sigma_r$ | 0.1555 |
| $\sigma_l$ | 0.1874 |
| $\bar{l}$ | 0.0523 |
| $s$ | 0.0131 |
| $\rho$ | 0.5808 |
| $\lambda$ | -0.4 |

Table 6: Parameter values for interest rate model.

| | Risk premium | Standard deviation |
|---|---|---|
| Swedish bonds (SB) | - | 0.0644 |
| Swedish T-bills (ST) | - | 0.0212 |
| Swedish stock (SS) | 0.07 | 0.2487 |
| Foreign bonds (FB) | - | 0.0951 |
| Foreign stocks (FS) | 0.06 | 0.1805 |
| Swedish real estate (ES) | 0.07 | 0.1805 |
| Swedish real bonds (RB) | 0.03 | 0.0355 |

Table 7: Means and standard deviations for asset model (yearly).

| | SB | ST | SS | FB | FS | ES |
|---|---|---|---|---|---|---|
| ST | 0.54 | | | | | |
| SS | 0.4643 | 0.1130 | | | | |
| FB | 0.2600 | 0.1848 | 0.2313 | | | |
| FS | 0.3332 | 0.0641 | 0.6914 | 0.5801 | | |
| ES | 0.6 | 0.4 | 0.2 | 0.1 | 0.05 | |
| RB | 0.2 | 0.1 | 0.2 | 0.2 | 0.1 | 0.2 |

Table 8: Correlations for asset model.