Stein-Erik Fleten

Portfolio management emphasizing electricity market applications

A stochastic programming approach

Stein-Erik Fleten
Section of Managerial Economics and Operations Research
Department of Industrial Economics and Technology Management
Norwegian University of Science and Technology
N–7491 Trondheim
Stein-Erik.Fleten@iot.ntnu.no
http://www.iot.ntnu.no/~fleten

Dr. ing. thesis
January 2000

Thesis advisor:
Stein W. Wallace
Co-advisor:
Olav Fagerlid

Evaluation committee:
Prof. Werner Römisch, Humboldt-Universität zu Berlin
Prof. Kurt Jörnsten, Norwegian School of Economics and Business Administration
Prof. Bjørn Nygreen, NTNU

Written in LaTeX
NTNU Doktor ingeniøravhandling 2000:16
ISBN 82-7984-037-0
ISSN 0802-3271

Preface

This thesis has been prepared at the Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology (NTNU) in partial fulfillment of the requirements for the doktor ingeniør degree. The work has been carried out during the period from January 1996 to January 2000 with Professor Stein W. Wallace as the thesis advisor.

Three years of work were dedicated to the completion of the doctoral degree and one year's work to duties set up by the industry sponsors.

The latter-mentioned year was financed through the Norwegian Electricity Association (Enfo), while the Norwegian Research Council (NFR) financed the work on the doctoral degree. Four Enfo members were represented in the project's steering committee that was responsible for setting up "duty" work and projects: Norsk Hydro, Aust-Agder Kraftverk, Vestfold Kraft Energi and Hafslund. These were also the most active sponsors.

This thesis spans a range of subjects, but the common theme is stochastic programming and portfolio management, with an emphasis on problems in the electricity industry. The thesis has two parts. The first is an introduction. It also outlines the contents of the papers in part two and provides discussions framing the papers. Part two consists of five papers of which one is an exercise in writing theoretical work in a popular form.

Acknowledgements

I am profoundly grateful for the excellent supervision from Professor Stein W. Wallace. He has provided very exciting research tasks and useful contacts through his extensive network of top researchers, and has always been a discerning discussion partner and source of good advice. I warmly recommend him as a supervisor. Further thanks are due to my co-supervisor Olav Fagerlid, my co-authors Alexei A. Gaivoronski, Kjetil Høyland, Trond Jørgensen and Petter E. de Lange, and other colleagues at the Department of Industrial Economics and Technology Management.

During this project I had the pleasure of staying at the University of British Columbia for the academic year 1996. I am most grateful to Professor William T. Ziemba for inviting me and for making the stay very inspiring both professionally and socially. Thanks to professors Ziemba and Wallace I was invited to present my work in the Workshop on Applications of Planning under Uncertainty in Palo Alto, California, July 1999. Three Nobel laureates were present during my talk, of which two were active in the discussion. This was a highlight in my academic career.

Thanks to all my sponsors, and to P. Longva and A. Grundt at Norsk Hydro for project initiation, data, interesting discussions and active involvement.

Last but not least, thanks to my family for their endurance and support.

Trondheim, January 28, 2000

Stein-Erik Fleten

Abstract

Using a stochastic programming approach, we consider portfolio management problems in the electricity and insurance businesses.

Traditional portfolio management models assume that the markets in which the manager operates are perfectly competitive. There is reason to question this in the case of deregulated electricity markets, which are often dominated by large vertically integrated firms. Employing a two-stage stochastic Cournot-type game model of the Scandinavian electricity market, we investigate the potential for use of market power by large producers. The model takes into account the commitment effect of hydroelectric generation and forward contract decisions. We find that Statkraft, the largest pure hydro producer, has no market power on the seasonal level due to the aggregate hydro capacity of the large number of smaller producers. However, Vattenfall, the largest producer, has incentives to withhold thermal capacity in order to raise prices.

Leaving the potential problems related to market power aside, we consider a price-taking hydropower producer facing uncertainty both in prices and reservoir inflow. Taking the view of a risk averse producer, we propose a portfolio management model for the purpose of hedging the financial risks using electricity contracts and the flexibility of the production assets. This is a multistage stochastic programming model, and studying a case with real data from Norway, we find that such a model has the potential to improve the portfolio management processes currently used.

Model verification is a challenging task in stochastic programming. For a portfolio management problem in a Norwegian insurance company, we compare two alternative model approaches. The first is a stochastic programming model, and the second is one where the asset mix is assumed constant, called the fixed mix approach. We explain how such a comparison can be done, considering the fact that in actual use, the models will be rerun before each decision is made. We find that the stochastic programming approach performs only slightly better than the fixed mix approach.

With a stochastic programming approach it is possible to model portfolio management problems where the investment universe can contain derivative assets. We have applied this in an analysis of a casualty insurer's problem

where we investigate the use of financial reinsurance, through selected options, in cases where the insurance company considers bearing catastrophe risks. Particular attention has been paid to the prices of the derivatives used in this model, using both arbitrage and equilibrium concepts.

Contents

1 Introduction

This thesis focuses on portfolio management issues using a dynamic stochastic approach. It consists of this introductory part and five papers on the subject. It is not on stochastic programming applications to portfolio management in general; an emphasis is put on problems arising in deregulated electricity markets, but we are also concerned with applications to insurance portfolios. Most analyses have paid little attention to complexity and numeric implementation issues; instead, we have focused on modeling aspects.

The thesis represents a continuation of my sivilingeniør thesis (Fleten 1995). In that work, the Norwegian electricity market is described, market failure issues are discussed, the portfolio management problem is thoroughly introduced and described, and a number of common models and systems used for market analysis and risk management are reviewed. Some topics, such as a description of the market and current risk management models, will not be repeated in this introduction, and only to the extent necessary in the papers.

The thesis is organized as follows. Section 2 introduces the papers. In Section 3, we discuss portfolio theory both in general, but also with an emphasis on electricity problems. Beginning with the original work of Markowitz, we move on to asset pricing theory and electricity contract pricing. We explain how the electricity derivative prices can guide hydro scheduling in the light of this theory, which implies that managers should not be concerned about the variance of profit. However, the assumptions of the classical theories do not hold in practice, and we argue that decision-makers should try to hedge the variability of profits. For special cases one can call on a separation theorem implying that production decisions can be made independently from hedging in the contract market. In the electricity market this is not the case. Since there are transaction costs on contracts and the production scheduling problem is inherently dynamic, a stochastic programming approach is called for.

Section 4 discusses an implementation of the model presented in Paper 2, Section 5 explains the scientific contribution of the thesis, and Section 6 states the conclusions and indicates areas for future research. Finally, the five papers are included in their entirety.

The most important part of this introductory part is Section 2, 5 and 6.

2 The papers

This section explains the development of each paper and identifies my contribution to the extent necessary[1]. The thesis supervisor, professor Stein W. Wallace, has been involved in most papers playing the role of mentor and discussion partner, and has been a vital source of ideas and direction.

Paper 1: A two-stage game model of the Nordic electricity market, is joint work with Tjing T. Lie, School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore. He visited NTNU on his sabbatical from January to June 1999. During that time the basic ideas of the paper were developed. The model was implemented (at NTNU) and the paper written during the winter of 1999/2000. We plan to submit it to an electricity-oriented journal.

Paper 2: Hedging electricity portfolios via stochastic programming, is my main work, and the co-authors, Stein W. Wallace and William T. Ziemba (University of British Columbia), have both played the role of supervisors. It has been under development since 1996, with conference presentations underway (e.g. Fleten, Wallace & Ziemba (1997), Fleten & Wallace (1998), Fleten, Wallace & Ziemba (1999)). We have benefited from collaborative work with A. Grundt and colleagues at Norsk Hydro and B. Mo at SINTEF Energy Research. After publication of the report by Grundt, Dahl, Fleten, Jenssen, Mo & Sætness (1998), the development of the model was bifurcated. It has been implemented as a prototype and is currently in use for testing and decision support at Norsk Hydro; Mo, Gjelsvik & Grundt (2000) reports on this other line of development. Paper 2 is currently in the stochastic programming e-print series (SPEPS) available at `http://dochost.rz.hu-berlin.de/speps`, and it is submitted to the IMA Volumes in Mathematics and its Applications, Springer Verlag. A referee report has not been received.

Paper 3: The performance of stochastic dynamic and fixed mix portfolio models, is joint work with Kjetil Høyland, Gjensidige Asset Management,

[1] I will use the form "I" and "my" when it is necessary to distinguish my work from the other authors or from the general research community.

and Stein W. Wallace. The paper was presented at the 21st meeting of the EURO Working Group on Financial Modelling, Venice, Italy, October 1997, after the ideas were developed but before the first draft was finished. The paper then appeared in Høyland (1998), and was submitted to the European Journal of Operational Research. After receiving comments from two anonymous referees, I have made a major revision according to the recommendations of the referees. The most important changes regards the stated focus of the paper, the explanation of the numerical results and inclusion of a statistical test on the hypothesis that the dynamic stochastic portfolio management approach dominates the fixed mix approach.

Paper 4: Modeling financial reinsurance in the casualty insurance business, is based on joint work with Petter E. de Lange, Gjensidige NOR Spareforsikring, and Alexei A. Gaivoronski (NTNU), who was de Lange's dr. ing. supervisor. Large parts is de Lange's work, I helped with model debugging. The development of the derivative pricing model within the stochastic programming framework was my contribution. The paper appeared in de Lange (1999), and was submitted to a journal, but a referee report has not been received. I have done minor revisions to the abstract and the presentation of the results.

Paper 5: Real options and managerial flexibility, is joint work with Trond Jørgensen, NTNU[2], and Stein W. Wallace. It was published in Telektronikk, Fleten, Jørgensen & Wallace (1998), and appeared in Jørgensen (1999). It is an exercise in writing popular science.

2.1 Why are these papers in one thesis?

Paper 1 is on "portfolio management" in a market where some producers have market power. We analyze the electricity market and learn how large producers make decisions on both production and forward contracts when they take into account the effect the decisions have on price. Such analyses should be performed as a part of the portfolio management process, for example when generating scenarios.

A model for joint hedging and hydroelectric generation scheduling for a price-taking firm is set up in Paper 2, which is the core paper of this thesis.

[2]Jørgensen is now affiliated with UUNET, USA.

Paper 3 is on validating stochastic programming applications to portfolio management, and the particular case studied is from the insurance business. Although it may have been advantageous to use a case from the electricity industry, some validation was performed in Paper 2, and the analysis of Paper 2 should be sufficiently general to warrant its inclusion in the thesis.

The fixed mix approach described in Paper 3 has a parallel in the electricity industry. Many companies organize their hedging around a "fixed position". The position at any future time is defined as the expected total generation at that time minus what is sold on contracts for that week. The companies are regularly revising their estimates of expected generation as new information on future spot prices and inflow becomes available, and some companies buy or sell in the forward market to get a zero position immediately after such a revision. In a separate note I investigate the profitability and risk properties of such fixed mix approaches to hedging (Fleten 1998). I find that such decision rules will both increase risk and lower expected profit compared to the model in Paper 2. For hydro producers there is usually a negative correlation between the inflow to reservoirs and spot price, which means that these firms have a natural hedge in their production portfolio. Thus it is not surprising that I find that a decision rule where the firm rebalances to a position that corresponds to selling 50% of expected generation forward, performs much better than the rule where 100% is sold (the zero-position rule).

Paper 4 is seemingly an odd one, very specific to the casualty insurance business. However, it is an application of stochastic programming to portfolio management, and my contribution, the derivative pricing model, is more generally applicable than just to the insurance business. It demonstrates a knowledge of financial theory necessary to make good dynamic stochastic portfolio models[3].

NTNU requires that doctoral candidates write a popular science paper to demonstrate such communication skills. Paper 5 accommodates this within the general subject area of the thesis.

In the next section we discuss how some issues in financial theory relate to portfolio management, with an emphasis on electricity applications.

[3] See, e.g. Christiansen & Wallace (1998) on stochastic programming and option theory and Klaassen (1998) on portfolio management and asset pricing.

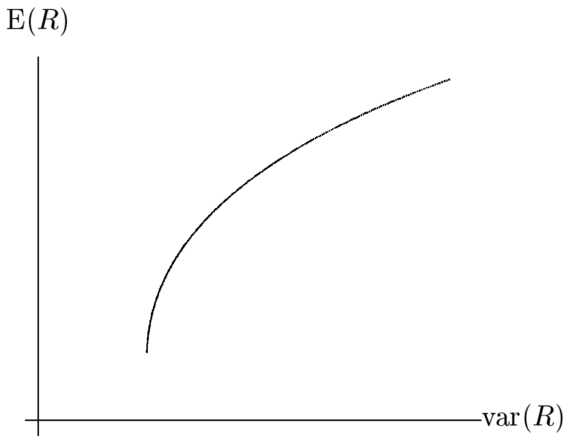

Figure 1: Efficient frontier.

3 Portfolio theory

3.1 The origin: Markowitz' mean-variance model

Modern portfolio theory began with Harry Markowitz and his 1952 publication on the mean-variance portfolio problem:

$$\begin{aligned}&\min\ \mathrm{var}(R)\\&\text{s.t.}\quad \mathrm{E}(R)=\rho\\&R \text{ is a portfolio return}\end{aligned}\tag{1}$$

where ρ is a specified scalar. When this quadratic programming problem is solved for a range of ρ's, we are able to trace out the efficient frontier as seen in Figure 1.

This model has the following disadvantages, due to its single period nature:

- There is no tradeoff between short and long term
- transaction costs are ignored

Even when there are no transaction costs, single period portfolio optimization is optimal only for special cases, specifically when the investor has a logarithmic utility function, see Mossin (1968) and Hakansson (1971).

From the mean-variance model one can derive the *capital asset pricing model* (CAPM), see e.g. Sharpe (1964). The portfolio selection problem can be formulated more generally, for example using other utility functions and more general asset distributions. Asset pricing relationships are then found by aggregating the first-order condition of the portfolio optimization problem on the part of individual investors and relating that to aggregate consumption or wealth.

3.2 Electricity contracts and asset pricing theory

Here we discuss some issues relating to the pricing of electricity term contracts. For now, we assume that markets are frictionless and that there are no market failures. In this subsection we disregard dynamic aspects of the electricity market.

Asset pricing models such as the CAPM can be modified to be applied to electricity forwards and futures. The CAPM relates the expected future spot prices to the current forward price of a contract for future delivery of electricity through a *risk premium*, to be defined shortly.

A *forward contract* is an agreement between a buyer and a seller on future delivery of electricity at an agreed price. In the Nordic market, this usually does *not* mean that electricity is physically delivered and received; it is a financial contract where the payment from the buyer to the seller is the difference between the agreed contract price and the spot price in the delivery period of the contract. The spot price is usually defined as the Nord Pool system price, and settlement occurs at regular intervals during delivery, e. g. monthly. In other power systems such contracts are termed "contracts for differences".

A *futures contract* is also an agreement on future financial delivery or purchase of electricity. It is a more standardized product, traded on the Nordic Power Exchange Nord Pool. Nord Pool is always the contractual partner, and takes on all credit risk. Futures contracts are settled daily through the use of margin accounts. This is the main difference between futures and forwards, and since many players seem to prefer the forward type of settlement, Nord Pool also provides forward products having delivery over seasons and years. These particular forwards are most frequently traded over the counter via independent brokers. For these OTC contracts the parties have the option to buy a clearing service, such that the credit risk is borne by a third party.

Both futures and forward contracts have a specified delivery period over which the energy is delivered at a constant power level. We will assume that interest rates are constant, and in such cases the forward and futures contract prices are equal (Cox, Ingersoll & Ross 1981).

A decreasingly popular contract is the load factor contract, having flexibility in when to draw energy from the contracted energy volume. There is an upper limit on the power level at which the buyer can withdraw energy, and a specified delivery period. In addition, there are often clauses stating that a fraction of the energy is to be delivered during certain parts of the delivery period. The most common terms in such a contract is a delivery period of one year and an energy level corresponding to 5000 hours of maximum power.

Increasingly popular are option contracts, both European and Asian types. An European call option gives the buyer the right, but not the obligation, to receive the difference between the price of the underlying contract and the option strike price. These options are usually written on futures or forward contracts, thus they are futures options. An Asian call option gives the buyer a payoff equal to the difference, if positive, between the average spot price during a specified "delivery" period and the strike price. Puts are also traded, paying off when the underlying contract or commodity has a price at maturity lower than the strike price.

If p_t^f is the current forward price for delivery in period t and $\mathrm{E}(\pi_t)$ is the expected time-average price in the delivery period, then the risk premium r_t^p is defined as:

$$r_t^p = \mathrm{E}(\pi_t) - p_t^f \tag{2}$$

At delivery, the forward/futures price equals the spot price, so r_t^p is also the expected change in the futures price up to delivery. Black (1976) showed that the CAPM for term contracts is

$$r_t^p = \beta \left[\mathrm{E}(r_m) - r \right], \tag{3}$$

where r_m is the return on the market portfolio, r is the risk free interest rate, and β is the "kroner" beta of the electricity price, defined as

$$\beta = \frac{\text{cov}\,(\text{E}(\pi_t), r_m)}{\text{var}\,(r_m)}.$$

We can now write

$$p_t^f = \text{E}(\pi_t) - \beta\,[\text{E}(r_m) - r]\,, \tag{4}$$

for the CAPM price of a futures contract.

This pricing relation should not be taken as a guide for daily trading, it is at best a long-term theoretical equilibrium, since it ignores risk aversion as a driving force for forward trading. On the contrary, we believe that most of the current participants in the electricity term markets are motivated by hedging production or consumption, or are speculating on the various derivative markets and the spot market, without much regard to the correlation with the rest of the capital market.

To illustrate what the CAPM *does* capture, consider the following two examples, adapted from Salahor (1998). Future spot prices are uncertain, and they are linked to world energy prices.

1. Suppose that there is a 50% probability that the price of electricity is above 140 NOK/MWh one year from now. Since world energy prices are positively correlated with the economy as a whole, so is the electricity price [4]. Now consider two option contracts. One will pay 1000 NOK if and only if the electricity price is above 140 NOK/MWh; this is a call option. The other will pay 1000 NOK if and only of the price is below 140 NOK; this is a put option. The expected payoff on each option is 500 NOK, but what will be the market price of these options?

   The call option will pay off in a situation where it is more likely that the world economy is going well, and the put option will pay off when it is likely that the economy is poor. Since most people are risk averse, the price of the call would be lower than the price of the put. The options have a nondiversifiable risk, so the call option price will be discounted for risk, and the put option price will have a markup for risk.

[4] Empirical tests on Norwegian electricity prices counter this fact. The linkage to the world economy is weak, so far, but we expect it to increase as Norway is increasingly connected to the thermal systems in the European Union.

2. Inflow to hydro reservoirs is uncertain. Suppose there is a 50% chance that the accumulated energy inflow to the hydro plants in a river system over the next year is above 1 TWh. Suppose further that the owner of the plants is a diversified investor, so that no matter what the inflow outcome is, it will have virtually zero effect on the wealth of the investor. Again consider two options. The first, call it a weather call, will pay 1000 NOK if and only if the accumulated energy inflow is above 1 TWh. The other, a weather put, will pay off 1000 NOK if and only if it is below 1 TWh. The market prices of these two options are likely to be very near 500 NOK, the expected payoff. The risk in these options is diversifiable, because inflow is not correlated with the world economy, and its effect on diversified investors will be negligible.

In summary, although CAPM is not well suited to pricing options[5], it is able to capture the pricing of risk. In the second example there was a zero risk premium, but in the first example, the uncertainty was assumed to be correlated with the state of the general economy. So there was a nonzero risk premium, because the risk could not be removed completely by portfolio diversification.

Assume that we want to find the value of an asset that delivers electrical energy into the spot market at a constant rate. There is no flexibility here, so the value would be

$$\text{NPV} = \text{Energy Quantity*Energy Value - Discounted Costs} \qquad (5)$$

We assume that the costs are either deterministic, or they have only diversifiable risk. The costs would then be discounted at the risk free rate. The energy value is equal to the electricity forward price discounted at the risk free rate.

Some may wonder why the energy value not is connected to the expected spot price in the delivery period. Suppose the forward price is 100, while the expected future spot price is 110. If someone uses 110 for the value of future delivery, they would be willing to buy at a price that is 110 or lower. In such a case, an arbitrageur would buy a forward at 100 and

[5]For options it is better to use arbitrage principles to find its value, however, the equilibrium (CAPM) analysis is still valid because the pricing of risk is embedded in the option price through the underlying commodity.

agree to deliver to that someone for between 100 and 110. At the time of delivery, the arbitrageur pays 100 for the forward, delivers the power and receives more than 100 for it. The message is that the forward price, not the expected future spot price, reflects the economic value of future delivery.

We next turn to the issue of modeling electricity prices in continuous time. Our motivation is that the better the portfolio manager captures the stochastic dynamic behavior of the electricity prices, the higher quality he is able to put into the decision process. The proposed model can be used for pricing electricity derivatives, and is an alternative to the approach presented in Paper 2, more in line with modern asset pricing theory. The presentation is somewhat compact and requires some knowledge of continuous time asset pricing theory. The reader may skip this subsection without loss of continuity.

3.3 A two-factor model of electricity prices

The electricity commodity market in the Nordic region is reasonably well functioning, see e.g. Paper 1, Amundsen & Bergman (1998), Amundsen, Bergman & Andersson (1998) and Johnsen, Verma & Wolfram (1999). This section deals with the stochastic behavior of electricity prices, which play a central role in models for evaluating investments in new and existing plants, portfolio management models, and for valuing financial contracts on electricity. We assume that there are no producers exercising market power or exploiting possible information asymmetries.

Electricity prices exhibit seasonality due to the fact that river inflow, which is an important input factor to electricity supply in this region, is large in the summer, and low in winter. Also, load is high in winter and low in summer.

This note will concentrate on a model that captures such seasonality, so it can be used, for example, in connection with models for trading, risk management, and to some extent plant valuation. The model is not intended to capture within-week seasonality of prices, which could be useful in some contexts, for example in valuation of upgrading capacity in existing plants[6].

[6]This represents a possible extension of this work. However, using a finer resolution

We expect that when prices are high, high-cost producers will enter the market, and some consumers will substitute other energy sources for electricity, putting a downward pressure on prices. Conversely, some supply will exit the market, and flexible consumers will increase demand when prices are relatively low. See Figure 1 in Paper 1.

Ordinary Black-Scholes models imply that the volatility of futures prices equal the volatility of spot prices, and the variance of future spot prices is not finite when the horizon increases, which does not seem realistic.

The model proposed is a continuous time two-factor stochastic process. It adds seasonality to the two-factor model in Schwartz (1997), as suggested by Hvarnes (1998). The first factor is the spot price, which has a drift dependent on the second factor, the convenience yield. The second factor is mean reverting and has annual seasonality.

Convenience yield can be interpreted as "the flow of services accruing to the holder of the spot commodity but not to the owner of a futures contract"[7]. Since it is impossible to "hold" electricity, we must interpret it as a measure of the balance between the available supply opportunities and demand. There is a positive correlation between spot price and convenience yield due to the impact of reservoir levels on spot and future prices. When reservoir levels decrease, the spot price increases since energy is scarce and the convenience yield should also increase since futures prices will not increase as much as the spot price, and vice versa when inventories increase. Information about electricity consumption preferences and generation technology in the market is embedded in the equilibrium convenience yield.

3.3.1 Modeling considerations

Let S_t be the spot price. Think of this as the average price during a week. To derive a fair market value of an electricity contract it is necessary to deduce the market risk premium of the process of S_t. This is analogous to finding the time-varying beta of the spot price. We assume here that the

than a week the modeling should take into account the possibility of large spikes in the spot price.

[7]We assume that there is no storage costs.

spot price risk can be hedged away using futures markets[8].

A plot of the nearest term futures contract is shown in Figure 2. The weekly

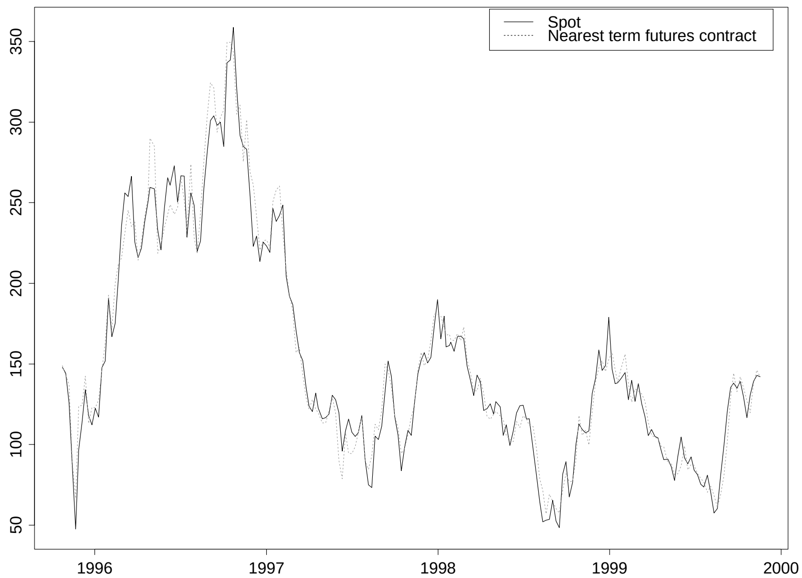

Figure 2: Nearest term contract on electricity vs. realized spot prices. Source: Nord Pool.

auto-correlations of this time series (see Figure 3) reveal oscillations which signal seasonality. Moreover, the first order auto-correlation coefficient for weekly data is around 0.9, signaling mean reversion. We could possibly include more than one seasonal component with different periods. However, the risk of over-fitting is high, and common sense dictates that the seasonality should be annual.

The joint stochastic process we propose for the spot price and the convenience yield δ is

$$dS = (\mu - \delta)Sdt + \sigma_1 Sdz_1 \quad (6)$$

[8]The alternative would be to split the price in two components: Market (hedgeable) component and basis component, i.e. which cannot be hedged away using the term market.

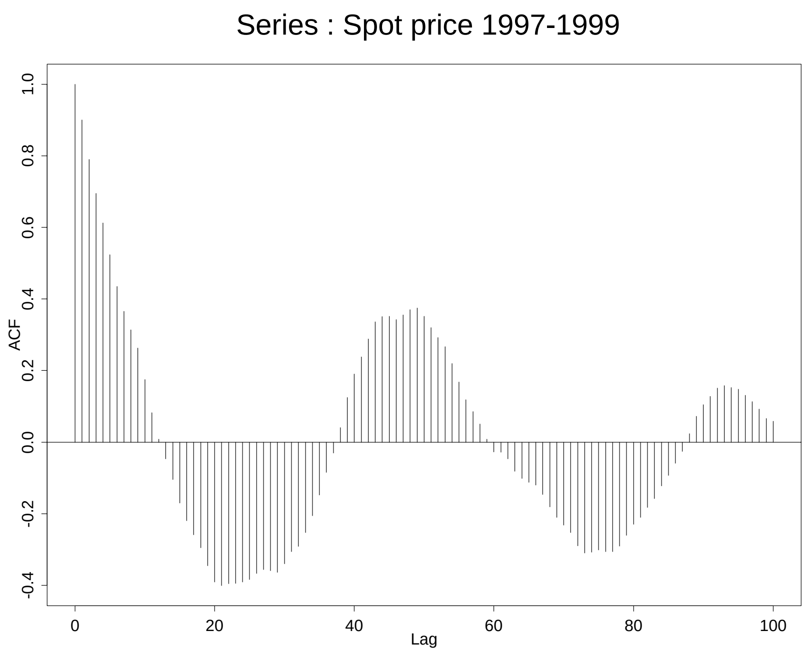

Figure 3: Weekly auto-correlations for the nearest term contract.

$$d\delta \;=\; \kappa(\alpha(t) - \delta)dt + \sigma_2 dz_2, \tag{7}$$

where the dz are random shocks (Wiener processes), and where these shocks are correlated with:

$$dz_1 dz_2 = \rho dt. \tag{8}$$

In a stochastic process equation, the term multiplying the dt is the infinitesimal trend in the processes. For spot price, the trend is composed of μ, the expected rate of price changes, and δ, the current rate of convenience yield. For the convenience yield, the trend is composed of a mean reverting component, κ, the long-run mean yield $\alpha(t)$, and the yield itself. The parameter κ, or rather its reciprocal $1/\kappa$, can be interpreted roughly as time it takes before the convenience yield is back to its mean. The variance of the change in the convenience yield is σ_2^2, and σ_1^2 is the variance of proportional price changes. We let

$$\alpha(t) = a + b\cos(2\pi t - \phi), \tag{9}$$

where a is a constant term and b is the seasonal peak, which peaks annually according to a phase given by ϕ. Note that prices can not become negative in this model, but the convenience yield can, as is reasonable. Mean reversion and seasonality enters the spot price through the drift term and through correlation between price and convenience yield.

As in Schwartz (1997), this model is an arbitrage model in which the stochastic behavior of prices and convenience yields is exogenously given. The value of any contingent claim on electricity can then be derived as a function of these primitives, imposing the condition that no arbitrage profits exist in perfect markets. There is debate over the question of whether it is possible to construct a no-arbitrage portfolio that hedges the electricity price. What is certain is that electricity can not be "held", so it can not enter the hedge portfolio directly. However, the requirement that the commodity must be a part of a replicating portfolio turns out not to be necessary. A recent thesis by Deng (1999) on financial methods in electricity markets solves this by considering the electricity price as a state variable on which derivatives are written, and assumes the existence of a risk-neutral probability measure on the state variable. Such a measure exists only if there are no arbitrage in all available financial instruments, and it ensures that pricing is easy (derivatives are priced using expectation under the new measure) and consistent (nobody will disagree on prices). Another possibility is to assume that the *consumption capital asset pricing model* (CCAPM) holds, and employ arguments used by Sick (1995). The result is the same.

3.3.2 Futures prices

We assume that interest rates are constant, so that futures prices equal forward prices. Let r be the riskless interest rate, and $F(S, \delta, t, T)$ be the futures price when the current spot price is S, convenience yield is δ, current time is t and the maturity date T. The contingent claims approach of e.g. Gibson & Schwartz (1990), invoking standard perfect market assumptions and absence of arbitrage, leads to the following differential equation for a futures price

$$\frac{1}{2}\sigma_1^2 S^2 F_{SS} + \sigma_1\sigma_2\rho S F_{S\delta} + \frac{1}{2}\sigma_2^2 F_{\delta\delta} \qquad (10)$$

$$+(r-\delta)SF_S + (\kappa(\alpha(t)-\delta)-\lambda)F_\delta + F_t = 0 \qquad (11)$$

with boundary condition

$$F(S_t, \delta_t, t, t) = S_t. \tag{12}$$

Since convenience yield is non-traded, the differential equation (11) depends on investor risk preferences embedded in the market price of risk for convenience yield, λ. For tractability, this is assumed constant. The risk-neutralized processes can be expressed as:

$$dS = (r - \delta)Sdt + \sigma_1 S dz_1^* \tag{13}$$
$$d\delta = [\kappa(\alpha(t) - \delta) - \lambda]\, dt + \sigma_2 dz_2^*. \tag{14}$$

Adjusting the price distribution for risk and then discounting cash flows of a contract at the riskless rate of interest is what is instructed by modern financial theory. As is common, the adjustment is only made in the drift term (however not in the mean reversion parameter κ).

Bjerksund (1991) and others have shown that the solution to the differential equation is ($\tau = T - t$):

$$F(S, \delta, t, T) = S \exp\left[-H(\tau)\delta + A(\tau) + r\tau\right]. \tag{15}$$

Or, in log form:

$$\ln F(S, \delta, t, T) = \ln S - H(\tau)\delta + A(\tau) + r\tau, \tag{16}$$

where, when α is constant

$$A(\tau) = \frac{[H(\tau) - \tau]\left(\kappa\rho\sigma_1\sigma_2 - \kappa\lambda - \sigma_2^2/2 + \kappa^2\alpha\right)}{\kappa^2} - \frac{\sigma_2^2 H(\tau)^2}{4\kappa}$$
$$H(\tau) = \frac{1 - e^{-\kappa\tau}}{\kappa} \tag{17}$$

When $\alpha(t)$ is time dependent as in Equation (9), the $A(\tau)$ is replaced by

$$A(t,T) = \frac{[H(\tau) - \tau]\left(\kappa\rho\sigma_1\sigma_2 - \kappa\lambda - \sigma_2^2/2\right)}{\kappa^2} - \frac{\sigma_2^2 H(\tau)^2}{4\kappa} + B(T) - B(t)e^{-\kappa\tau} - \frac{b}{\omega}\left(\sin(\omega T - \phi) - \sin(\omega t - \phi)\right) \tag{18}$$

where

$$B(s) = \frac{b}{\kappa^2 + \omega^2}\left(\kappa\cos(\omega s - \phi) + \omega\sin(\omega s - \phi)\right), \tag{19}$$

and $\omega = 2\pi$. We have set $a = 0$ here, and instead let that parameter be absorbed in the estimation of λ.

3.3.3 Nord Pool contracts

The futures contract defined here is not the same as the forward contracts traded in the Nordic market, which could be specified such that the buyer of one MWh receives the difference between the spot price S_t and the predetermined contract price y for each instant during delivery of the contract. Assuming that payments are deferred until the end of the delivery period and no interest is earned on funds accumulated, the ex-post payment to the buyer is

$$Y^{t_2} = \int_{t_1}^{t_2} (S_t - y)dt. \tag{20}$$

The futures/forward price is equal to the contract price giving a zero value of the above formula:

$$G(S, \delta, t, t_1, t_2) = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} F(S, \delta, t, T)dT \tag{21}$$

This integral must be computed numerically[9]. However, the difference between this more realistic future price and the one derived in Equation (15) is small when $t_2 - t_1$ is small, and in the remainder of the note we will work with $F(S, \delta, t, T)$.

3.3.4 Volatility structure

The volatility of futures prices is important because it is used as input to other analyses, for example futures option pricing. This volatility depends on the volatility of spot price and convenience yield, correlation between the two factors, the speed of adjustment of convenience yield, and time to maturity. For derivation, see Hilliard & Reis (1998):

$$\sigma(\sigma_1, \sigma_2, \rho, \kappa, \tau) = \sqrt{\sigma_1^2 + \sigma_2^2 H(\tau)^2 - 2\sigma_1\sigma_2\rho H(\tau)} \tag{22}$$

This represents the term structure of the volatility of futures price. It does not depend on the level of convenience yield or the level of spot prices, a fact that simplifies futures option pricing.

[9] Note that this is the average value of F, thus it should be easy to develop approximate formulas.

3.3.5 Parameter estimation

Since the purpose of this subsection is to demonstrate a way to capture the economic information in the derivatives markets, only preliminary estimation has been performed. The estimation of parameters in the model is complicated by the fact that convenience yield is unobservable. The extended Kalman filter (Harvey 1989) can be used for estimating the parameters and convenience yield simultaneously, as explained by Schwartz (1997). In our preliminary estimation, we choose a simpler approach. Gibson & Schwartz (1990) show a way to find a proxy for convenience yield, and we use a variant of that. For each trading date t, a cosine function was fitted to the term structure via nonlinear least squares, giving an estimate of the futures price $F(T)$ for a contract maturing at time T. The convenience yield was then found using the following formula:

$$\delta_t = r_t - \frac{\partial}{\partial T}[\ln F(T)]\Big|_{T=t} \tag{23}$$

where r_t is short term interest rate[10]. The intuition behind this is as follows. The term structure of futures prices represents the expected value of spot prices in a risk neutral world. The risk-neutral growth rate of electricity prices is thus $ln\dot{F}$. On the other hand, when δ is net convenience yield, electricity prices behave like a traded security paying a dividend return of δ, thus its risk neutral growth is $r - \delta$.

The data used is daily prices on Nord Pool futures and forwards, from October 1995 to October 1999. Every Friday was selected for estimation. Each trading date has 17–27 different contracts, i.e. delivery periods. To estimate the value of a single week in a block, season or year, we picked the middle value or the week numbers expected to have a value in the average of the weeks stacked.

The parameters were found by minimizing the squared error of observed vs. model futures prices. Figure 4 displays the estimated convenience yield on the secondary axis, and the aggregate Norwegian reservoir level in percent on the primary axis. The strong negative correlation confirms classical theories of storage, as developed by Kaldor (1939) and Working (1949).

Table 1 displays the estimated parameters.

[10] Thanks to Gjensidige Asset Management for supplying interest rate data, three month NIBOR (Norwegian InterBank Offering Rate).

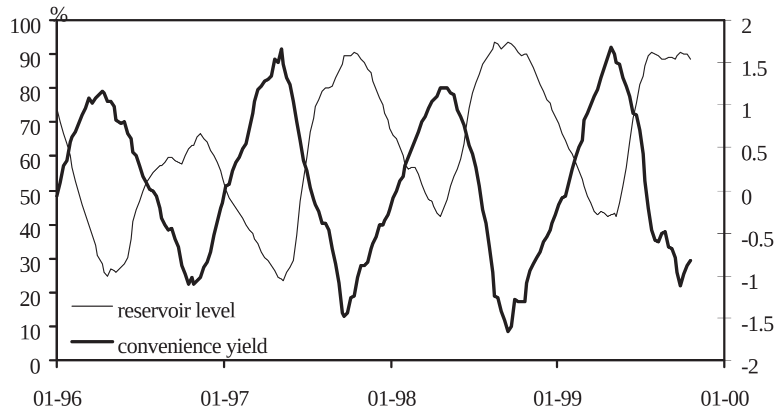

Figure 4: Estimated convenience yield of electricity prices, thick line, on right hand side axis. Also shown is the aggregate Norwegian reservoir level, thin line, on the lefth hand axis.

Table 1: Parameters estimated by minimizing the sum of square errors over selected trading dates and all contracts.

Parameter	κ	λ	σ_1	σ_2	ρ	b	ϕ
Estimated value	57.2	5.12	0.442	0.100	0.167	1.38	2.11

Figure 5 displays observed vs. modeled futures prices.

3.3.6 Options

European options can be valued by solving

$$\frac{1}{2}\sigma_1^2 S^2 V_{SS} + \sigma_1\sigma_2\rho S V_{S\delta} + \frac{1}{2}\sigma_2^2 V_{\delta\delta}$$
$$+(r-\delta)SV_S + (\kappa(\alpha(t)-\delta)-\lambda)V_\delta + V_t = rV \tag{24}$$

subject to the boundary condition $V(S,\delta,t,t) = \max(S_t - K, 0)$, where V is the (call) option price and K is the exercise price. This will probably have to be done numerically, however, this type of contract is not traded in the market today. Still, some contracts or real assets may be viewed as a series of such options.

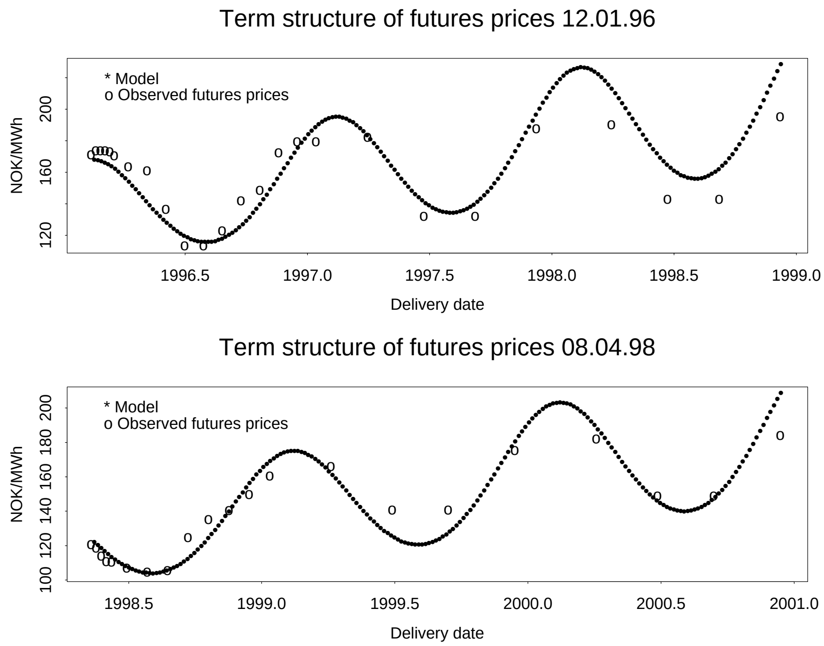

Figure 5: Observed term structure of futures prices vs. model prediction given spot price and convenience yield at two representative dates. The first date is out of sample, while the second is in sample. Source: Nord Pool.

For valuation of European futures options, see Hilliard & Reis (1998), where an explicit formula is given.

3.3.7 General contracts and real options

For contracts or real assets where the owner must make decisions continuously, we recommend discretizing (in time and space) the stochastic processes and model and solve the joint valuation and decision problem using stochastic programming techniques. The traditional approach in finance is to use either a binomial or trinomial tree, and subsequently take a dynamic programming approach. However, other choices are also possible.

For Asian options, we recommend simulation from the risk-neutralized version of the joint stochastic process.

3.3.8 Two-factor model summary

A model for the stochastic behavior of electricity prices has been developed. This has been done through adding seasonality to the long run mean of the convenience yield of a model presented by Schwartz (1997).

The strength of such a model is that it captures the economic information in the market prices. The challenge lies in combining such a model with stochastic programming models such as the ones in Paper 2, 3 and 4, see e.g. Christiansen & Wallace (1998), Klaassen (1998) and Paper 4.

3.4 Hydro scheduling in perfect asset markets

Now we turn to the issue of using the futures prices as signals for investment and operating decisions. Deregulated electricity industries often develop markets for contracts for future delivery of electricity. This facilitates price transparency and risk sharing. Regardless of the reason for the presence of such markets, if competitive, the equilibrium forward prices will reflect the economic value of future delivery.

Traditional medium-term hydro scheduling assumes a regulated deterministic price of energy, and often a stochastic inflow of water to the reservoirs. In a deregulated market, a single hydro producer is likely to face uncertain spot market prices also. When there is a well-functioning financial futures market for electricity available, the modern asset pricing theory will guide how to discount cash flows in such a model.

Assume that the firm we are studying operate in a deregulated electricity market with access to competitive forward-futures prices, and whose owners have diversified portfolios or want their managers to maximize the expected long-term value of the firm without regard to diversifiable risk. In such a case the decision-maker should use the futures and forward market prices as signals guiding the hydro scheduling. Assume there is available a distribution of spot prices and inflows available in the form of a scenario tree. One way of incorporating the economic information in the futures prices is to discount the spot prices in the tree in such a way as to equate the discounted term structure of expected spot price with the term structure of futures prices. That is, choose the discount rate γ_t for price-related cash flows so that

$$\frac{\mathrm{E}(\pi_t)}{1+\gamma_t} = \frac{p^f}{1+r_t} \quad \forall t. \tag{25}$$

The disadvantage of this approach is that it does not recognize that the riskiness of future operating income is different in each node, calling for a stochastic discount rate.

An alternative approach is to combine a continuous-time model such as the one presented in Section 3.3 with existing hydro scheduling models. This issue is left for future work.

3.5 Frictionless markets?

3.5.1 Market imperfections lead to risk aversion

The Modigliani & Miller (1958) analysis of capital structure implies that it is not necessary to hedge at the corporate level, since investors can do that on their own account. Thus if risk management is costly, firms should not do it, but leave it to the owners. This is in contrast with conventional views such as "hedging is necessary to reduce risk", or "it is necessary to take risks to earn money", which do not take into account the pricing of risk in the market.

The Modigliani-Miller theorem applies to situations where there are no transaction costs, no taxes and no information asymmetries. In practice, such "imperfections" are present, and corporations should hedge if the benefit to the owners is greater than the cost, and if hedging is cheaper for the company than for the owners.

Since there are fixed fees, information costs etc. associated with participating in the electricity term markets, it is usually cheaper for the company to hedge. I.e. there are economies of scale to hedging (Mian 1996).

A majority of the owners in the Scandinavian electricity industry are governmental, either in the form of the state, counties or municipalities. These can be seen as independent economic decision-makers, who are generally not diversified in the general capital market. It is also natural to consider them risk averse, either due to the tightness of public budgets, or risk aversion on the part of managers.

Regardless of ownership structure, the usual reasons for corporations to be risk averse apply. These are "market imperfections" affecting firm's profits in ways that can not be offset by individual investors, such as convex taxation (Smith & Stulz 1985), bankruptcy costs and financial distress (Brealey & Myers 2000), agency problems (Stulz 1990) and effectiveness of incentive schemes. Financial hedging also improves the informativeness of corporate earnings as a signal of management ability and project quality by eliminating irrelevant noise (DeMarzo & Duffie 1995). Further, if external financing sources are more costly to corporations than internally generated funds, there will typically be a benefit to hedging: hedging adds value to the extent that it helps ensure that a corporation has sufficient internal funds available to take advantage of attractive investment opportunities (Froot, Scharfstein & Stein 1993). Last, but perhaps most importantly, hedging permits greater leverage and thus tax advantages of debt (Ross 1996).

The combination of risk aversion on the part of both owners and managers, and the fact that it is cheaper for a firm to operate in the electricity derivative markets than for the owners, make the case for hedging at the corporate level in the electricity industry.

3.5.2 Electricity market failure

In the case of the electricity markets, there may also be other market failures making it difficult to apply portfolio theory. These are market power, asymmetric information, externalities and public goods. In Paper 1 we examine the first of these, the others are left for future work[11]. Particularly interesting is the issue of asymmetric information in the Nordic electricity market because possible information advantages are obviously very profitable when dealing in the term market.

3.6 Hedging and joint production

Sandmo (1971) pioneered the analysis of firms' behavior under output price risk. He showed that the presence of such risk reduces output. Another well known result, potentially important for our purposes, is that under certain assumptions (no production uncertainty or basis risk) production

[11] Some of these issues were discussed in Fleten (1995).

planning can be done independently from hedging. Production decisions depend only on the futures price, whereas hedging/speculative decisions also depend on the producer's subjective beliefs. This *separation theorem* was stated by Holthausen (1979) and others.

In the electricity industry, the assumptions needed to make this separation work are not met, due to the presence of basis risk and production uncertainty.

When there is basis risk, production must be determined jointly with hedging (Anderson & Danthine 1980). The basis is the difference between the commodity underlying the derivative contract used for hedging, and the spot commodity that the company has to sell or buy. For example, the commodity underlying Nord Pool futures is the Nord Pool system price, while a particular distributor may have customers only in Northern Norway, which often is treated as a separate price area. There is basis risk when there is uncertainty about the basis in the delivery period. Currently, Nord Pool does not offer products that hedge against spatial price risk.

In our work we have abstracted from the basis risk problem. Still, there are other factors causing breakdown of separation, such as input cost uncertainty being correlated with the output price uncertainty. See e.g. Viaene & Zilcha (1998). The following example illustrates the breakdown of separation.

3.6.1 A two stage example

In this subsection we regard the hydro scheduling problem in a model with two stages. There is a single reservoir. The marginal value κ of the water in the reservoir x_t at the end of the model horizon is lower than the lowest possible value of the spot price p_t at that stage. One decides on discharge u_t after learning what the spot price and inflow ξ_t are. Thus at stage two, the problem is

$$Q(x_1; p_2, \xi_2) = \max_{u_2, x_2} \{p_2 u_2 + \kappa x_2\} \tag{26}$$

$$\begin{aligned} \text{s.t.} \quad x_2 + u_2 + s_2 &= x_1 + \xi_2 \\ 0 \leq x_2 &\leq \overline{x}_2 \\ 0 \leq u_2 &\leq \overline{u}_2 \end{aligned} \tag{27}$$

Table 2: Stage two payoffs.

Cash flow name	Value now	Low price	High price
Future operating income	$Q(x_1) = ?$	$Q(x_1; p_2^a, \xi_2^a)$	$Q(x_1; p_2^b, \xi_2^b)$
Risk free bond	$(1+r)^{-1}$	1	1
Forward contract	0	$p_2^a - K$	$p_2^b - K$

Spill is denoted s_t.

The solution to this problem is

$$Q(x_1; p_2, \xi_2) = p_2 \min(x_1 + \xi_2, \overline{u}_2) + \kappa \min(\overline{x}_2, (x_1 + \xi_2 - \overline{u}_2)^+) \tag{28}$$

where the notation "a^+" means the positive part of a.

The overall problem then becomes

$$\begin{aligned} \max_{u_1, x_1} \Big\{ p_1 u_1 + (1+\rho)^{-1} \mathrm{E}\Big[& \tilde{p}_2 \min(x_1 + \tilde{\xi}_2, \overline{u}_2) \\ & + \kappa \min(\overline{x}_2, (x_1 + \tilde{\xi}_2 - \overline{u}_2)^+) \Big] \Big\} \end{aligned}$$

$$\begin{aligned} \text{s.t.} \quad x_1 + u_1 + s_1 &= x_0 + \xi_1 \\ 0 \le x_1 &\le \overline{x}_1 \\ 0 \le u_1 &\le \overline{u}_1, \end{aligned} \tag{29}$$

which transforms into

$$\begin{aligned} \max_{x_1} \Big\{ p_1 x_1 + p_1(x_0 + \xi_1) + (1+\rho)^{-1} \mathrm{E}\Big[& \tilde{p}_2 \min(x_1 + \tilde{\xi}_2, \overline{u}_2) \\ & + \kappa \min(\overline{x}_2, (x_1 + \tilde{\xi}_2 - \overline{u}_2)^+) \Big] \Big\} \end{aligned} \tag{30}$$

$$\text{s.t.} \quad (x_0 - \overline{u}_1 + \xi_1)^+ \le x_1 \le \min(\overline{x}_1, x_0 + \xi_1) \tag{31}$$

Note that the future profit function was found in the traditional way using discounted expectation.

Assume now that there are only two outcomes, and that there exists a forward contract that pays off $\tilde{p}_2 - K$ at the second stage. See Table 2.

According to the principles of modern asset pricing, we find the risk neutral probabilities so that under the risk neutral measure, the expected value of the forward contract is zero. Thus the risk neutral probability p of a low price is

$$0 = p(p_2^a - K) + (1-p)(p_2^b - K) \tag{32}$$

$$\text{risk neutral probability} = p = \frac{p_2^b - K}{p_2^b - p_2^a} \tag{33}$$

Thus according to asset pricing theory the value of the future operating income is

$$Q(x_1) = pQ(x_1; p_2^a, \xi^a) + (1-p)Q(x_1; p_2^b, \xi^b) \tag{34}$$

and the overall problem becomes

$$\max_{x_1} \Big\{ p_1 x_1 + p_1(x_0 + \xi_1) + (1+r)^{-1}\hat{\mathrm{E}}\Big[\tilde{p}_2 \min(x_1 + \tilde{\xi}_2, \overline{u}_2) + \kappa \min(\overline{x}_2, (x_1 + \tilde{\xi}_2 - \overline{u}_2)^+)\Big]\Big\} \tag{35}$$

$$\text{s.t.} \quad (x_0 - \overline{u}_1 + \xi_1)^+ \leq x_1 \leq \min(\overline{x}_1, x_0 + \xi_1) \tag{36}$$

where the expectation $\hat{\mathrm{E}}$ is taken with respect to the risk neutral probability measure, and r is the risk free rate of return. Notice that in this case, the empirical probabilities are not needed.

So far, all we have done is to show how modern asset pricing theory guides valuation for this two-stage example. In the words of Dixit & Pindyck (1994), the future contract acts as a "spanning asset". For the remainder of this subsection, we assume that the forward exists and is fairly priced, so that we can safely take the value of the future operating income as given by the modern asset pricing approach.

Is it possible to go one step further, and use the deterministic forward price instead of the uncertain future spot price in the model? This is what is implied by the separation theorem. We show that the answer is no when the natural assumption is made that inflow is correlated with spot prices. The difference between the future operating income function with stochastic inflow and spot price,

$$Q(x_1) = \frac{p_2^b - K}{p_2^b - p_2^a} Q(x_1; p_2^a, \xi^a) + \frac{K - p_2^a}{p_2^a - p_2^b} Q(x_1; p_2^b, \xi^b), \tag{37}$$

and the future operating income function with the forward price K substituted for uncertain spot price,

$$Q(x_1) = \frac{p_2^b - K}{p_2^b - p_2^a} Q(x_1; K, \xi^a) + \frac{K - p_2^a}{p_2^a - p_2^b} Q(x_1; K, \xi^b), \tag{38}$$

is

$$\frac{(p_2^b - K)(K - p_2^a)}{p_2^b - p_2^a} \left(\min(x_1 + \xi_2^a, \overline{u}_2) - \min(x_1 + \xi_2^b, \overline{u}_2) \right). \tag{39}$$

This is not zero because in any reasonable setting, $K \in \langle p_2^a, p_2^b \rangle$ and $\xi_2^a \neq \xi_2^b$ when inflow is imperfectly correlated with price. It is thus not correct to substitute in forward prices for uncertain spot prices, because the future operating income function changes if you do so, thus the decisions may also be wrong.

Thus we have shown by an example that production uncertainty causes a breakdown of the separation theorem in the case of hydropower scheduling.

3.7 The case for stochastic programming

Although forward markets reveal a deterministic number for the value of future delivery of electricity, the uncertain prices should be taken into account when planning production. This is because in general, in a decision-making process under uncertainty, inserting deterministic values for the random parameters does not lead to optimal decisions (Kall & Wallace 1994). In that case, the futures market should be used for hedging the uncertain hydro revenues, and this hedging should be determined jointly with the production schedule in a dynamic stochastic model. Such an approach is described in Paper 2.

As mentioned in Section 3.1, for general portfolio selection problems, single period optimization will not lead to optimal dynamic portfolio policies when there are transaction costs or when the objective can not be expressed as a logarithmic utility function. In such cases a dynamic approach is called for, for example via stochastic programming. See Ziemba & Mulvey (1998) for a survey and collection of financial portfolio management models.

4 Implementation

As mentioned in Section 2, a prototype of the model presented in Paper 2 has been specified. We chose to use the so-called stochastic dual dynamic programming (SDDP) algorithm, see Grundt et al. (1998) for a discussion. This was implemented in a contractual research project by SINTEF Energy Research for Norsk Hydro (Mo et al. 2000) and is currently in use on a trial basis for decision support at Norsk Hydro concurrently with further development.

4.1 Stochastic programming algorithms

Multistage stochastic programming problems are solved using a number of different algorithms. A simple alternative is to formulate the deterministic equivalent mathematical program, and solve this using standard deterministic techniques. The traditional alternative is to formulate the model so as to fit the framework of dynamic programming (SDP), and solve the problem using stochastic dynamic programming. Also well known is the nested Benders decomposition algorithm. This section will at some level of detail describe a more recent algorithm called stochastic dual dynamic programming, which can be seen as a combination of SDP and nested Benders.

Dynamic programming was pioneered by Bellman (1957). It is a powerful algorithm provided the state space is small. The size of the state space can roughly be seen as the minimum number of variables needed to fully describe the model information that is necessary to carry over from one stage to the next. These state variables can be decision variables in the mathematical program, representing such information as the level of inventory. The state variables can also represent trend level of stochastic model input parameters. The fact that dynamic programming becomes inefficient as the number of state variables is larger than, say three, is called the curse of dimensionality.

The curse of dimensionality is not present in nested Benders decomposition. However, the computational effort is proportional to the total number of stochastic events in the problem. This is equal to the number of nodes in the event tree that describes the evolution of information in our model. This number grows exponentially when the number of stages increases, so nested Benders decomposition can only be used when the number of stages is small. In summary, SDP can handle a large number of stages, but not a large state space, and nested Benders can handle a large state space but not a large number of stages.

Nested Benders decomposition was first used on multistage stochastic programs by Birge (1985). Pereira (1989) developed stochastic dual dynamic programming (SDDP). The basic idea of the algorithm is to store the future cost function of dynamic programming in the form of nested Benders cuts instead of in a table, which is usual in SDP. This overcomes the curse of dimensionality. To overcome the stage-dimensionality problem of Benders decomposition, it is required of the structure of the problem that all state variables relating to the trend level of stochastic model parameters are included in the state space, like in SDP. Since a function has to be convex in order to be approximated via Benders cuts, this means that the problem must be convex in all state variables.

The reason for explaining some aspects of this algorithm is that it is little known in the research community and is not described in textbooks. Further, the algorithm was chosen in a commercial prototype implementation.

4.2 The stochastic dual dynamic programming algorithm

We study multistage stochastic programming problems that arise in hydroelectric scheduling. In that problem, the state variables are the hydro reservoir levels and the trend in stochastic inflow and spot market price. At any stage, all state variables except price are related to each other through linear functions. Thus the future cost function of the previous stage is convex in these state variables. However, the price state variable is related to reservoir levels and inflow through a product term making the overall future cost function for this stage nonconvex.

When trend levels of stochastic parameters are to be included in the state space, then we must include equations governing the stochastic processes of these variables in the description of the problem. This will transform the problem from one in which more information (on the trend levels) is needed in order to find the value of the future cost function, into one where the value can be found using only information about the current state. In order to give these ideas a firmer footing, an example from hydro scheduling is discussed.

Let x_t be the decision variables that also are state variables. Let u_t be other decision variables. Let ξ_t be stochastic parameters that also are state variables, and let ϵ_t be other stochastic parameters. At stage t, the problem is:

$$\begin{aligned} \phi_t(x_{t-1}, \xi_{t-1}) = \min_{u_t, x_t} \mathrm{E}_{\xi_t, \epsilon_t}(r_t(u_t, x_t, \xi_t) + f_t(x_t, \xi_t)) \\ \text{s.t.} \quad x_t = g(x_{t-1}, u_t, \xi_t, \epsilon_t) \\ \underline{x}_t \leq x_t \leq \overline{x}_t \\ \underline{u}_t \leq u_t \leq \overline{u}_t \\ \text{and} \quad \xi_t = \beta_t \xi_{t-1} + \epsilon_t \end{aligned} \tag{40}$$

Here f_t is represented in terms of cuts. These cuts are generated in the course of the algorithm, thus when one of the $\phi_t(x_{t-1}, \xi_{t-1})$ problems (for a given ξ_t and ϵ_t) is solved f_t is only an approximation.

Note that those subproblems are independent.

Let x_t be the reservoir content, u_t the reservoir discharge, ν_t the inflow, π_t the price and r_t spill. In a given state at stage t, a subproblem for a simplified hydro scheduling model may look like:

$$\begin{aligned} \phi_t(x_{t-1}, \pi_{t-1}, \nu_{t-1}) = \\ \max_{u_t, x_t} \{\pi_t u_t + f_t(x_t, \pi_t, \nu_t)\} \\ \text{s.t.} \quad u_t + x_t + r_t = x_{t-1} + \nu_t \\ \underline{x}_t \leq x_t \leq \overline{x}_t \\ \underline{u}_t \leq u_t \leq \overline{u}_t \end{aligned} \tag{41}$$

The solution of this subproblem for all outcomes of the uncertain variables will give information for the previous stage in the form of a cut. It will give information to the next stage in the form of values for the state variables. A cut is a linear function in the state variables that gives the following information: At the "reference point" (the current value of the state variables), what is the future cost? And with changes in the state variables from this reference point, how will the future cost change? We assume that such a cut is a lower bound approximation of the true future cost function. With θ_{t-1} as a variable representing the future cost at stage $t-1$, we can thus write the cut:

$$\theta_{t-1} \leq \hat{f}_{t-1}(x_{t-1}, \pi_{t-1}, \nu_{t-1}),$$

where $\hat{f}_{t-1}$ is the linear approximation function (cut). All parameters for this function can be found from primal and dual information of (41). The level parameter is the objective function value for the current value of states, $\mathrm{E}\phi_t(\hat{x}_{t-1}, \hat{\pi}_{t-1}, \hat{\nu}_{t-1})$. Information about the gradient parameters for reservoir level and inflow is represented by the dual variables. The gradient parameter for price is the optimal release.

As is known from linear programming theory, the future cost function will be convex in the right hand side, but concave in the objective coefficients. To further see that including price as a state variable may lead to trouble, note that if we substitute out optimal release u_t, we get an objective function that looks like the following:

$$\max_{x_t} \pi_t \left(-x_t - r_t + x_{t-1} + \nu_t\right) + f_t()$$

Formally, we can check whether ϕ_t is concave ($-\phi_t$ is convex) by examining the Hessian. Remember that we must require convexity at all stages and at any iteration of the algorithm. This is obviously not the case, since there are product terms of price and inflow, and price and reservoir level.

Stochastic dual dynamic programming is organized around two major iterations, called forward simulation and backward recursion. In the forward simulation phase, we start by solving the first stage problem and then solve one stage at a time. Primal information about the state variables is always stored and passed to the next stage. In the other major iteration, called backward recursion, we start at the last stage and move back solving one stage at a time. Dual information in terms of cuts is stored and passed to the previous stage. The values for the state variables as found in the forward simulation are used as estimates for the current state values.

The algorithm described thus far is essentially the same as nested Benders decomposition imposed with a particular ordering rule, namely depth first.

In contrast to nested Benders decomposition not all combinations of the random variables are considered in the forward simulation. Instead, a number of scenarios are sampled. In the backward recursion, within each stage, all possible combinations of the random variables are considered. However, as starting point for the state variables the scenario values are used.

Without the sampling scheme, the algorithm will converge as does nested Benders decomposition and SDP. That the algorithm converges when sampling is used, is an unproved conjecture. However, a similar algorithm by Chen & Powell (1999) has been proved to converge.

Why does this work? A natural prerequisite is that the scenarios are a good representation of possible future outcomes. Among other things, this means that it should cover the support of the random variables well. Thus the scenarios should lead to a realistic range of state variables. One may think of this range of different state variables as an analogue to the discrete grid of state variables used in ordinary SDP. This will in turn lead to a good description of the future cost functions, and eventually also to optimal decisions.

To avoid the problems connected to having price as a state variable, price is defined as a "super" state, i.e. it is discretized and separate SDDP cuts are build up for each discrete price state at each stage. Price is assumed to be stochastically independent from inflow. Transition between price states is governed by a sequence of Markov models (see Gjelsvik & Wallace (1996)).

4.3 Test case

The sizes of the stochastic programs that have been solved using the Norsk Hydro prototype are considerable by today's standards. This subsection describes a case where the entire Norsk Hydro portfolio is studied. Hydro has 38 reservoirs and 27 power stations along five Norwegian rivers. There are 21 different futures contracts that can be traded initially, and the portfolio also includes 20 load factor contracts (without the opportunity to resell or buy more). There are 2 periods in which profit is measured and penalized, the periods being years. This gives a total of roughly 90 state variables. Five discrete price states are used for each stage.

The planning period is two years with a weekly time resolution, giving a total of 104 stages. 60 price/inflow scenarios are used as the sample scenarios in the forward simulation phase.

Using a DEC Alpha 4100 computer, 20 backward-forward iterations take about 8 hours.

The tests show promising results in terms of reduced downside variance. See Figure 6. However, some problems have been identified. The long term price dependencies inherent in the 60 price/inflow scenarios are not captured in the Markov price models, causing a spurious expected gain from forward trading[12]. The forward price is biased downward in high spot price states, and is biased upward in low spot price states. The model will generally recommend to sell forward when spot prices are low and buy when prices are high, causing the spurious hedging gain.

Moreover, the forward prices for maturities greater than six months from now, derived from the Markov price models, are insensitive to the current price state. This means that there is no uncertainty in those forward contract prices, and consequently the model sees no gain in trading in those contracts at the current stage; it might as well wait.

The fact that price is assumed to be independent of inflow is also likely to cause biases, since the model does not fully capture the "natural hedge" aspect of the risk of a hydropower portfolio[13].

In summary, by moving to a specialized algorithm, we have both gained and lost. The sizes of the problems that are solvable within reasonable time are formidable. On the other hand, the assumption that price is stochastically independent of inflow will probably lead to systematic errors in the recommended decisions, especially regarding contract positions.

[12]Since contracts are priced at the expected future spot price, the transaction costs should make expected profit from hedging negative.

[13]Price and inflow are of course correlated in the 60 scenarios used in the forward simulation, but this dependency will not be reflected in the construction of the cuts.

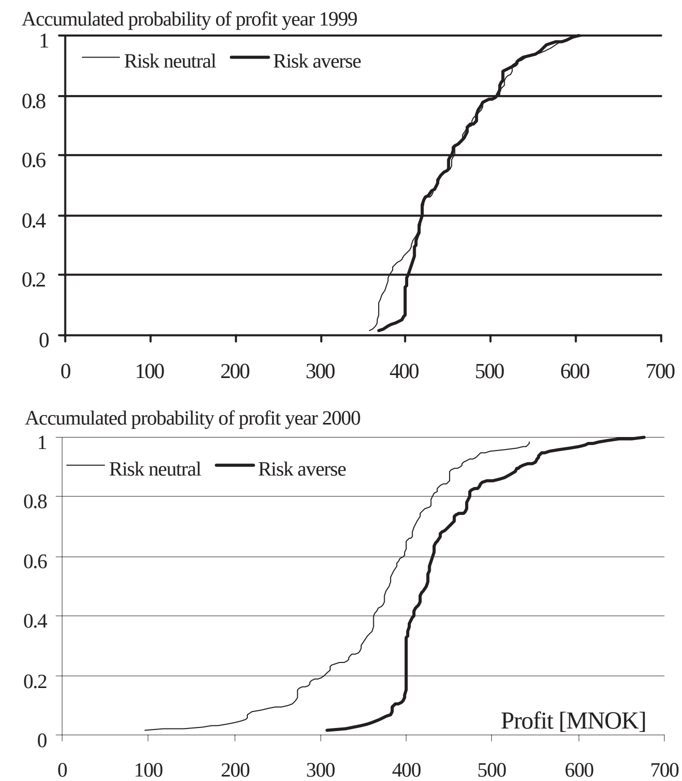

Figure 6: Cumulative probability functions for profit resulting from a test case at Norsk Hydro.

5 Scientific contribution

In general, this thesis puts known ideas into new contexts and applications. The new aspect in Paper 1 on electricity market gaming, as seen in relation to, e.g., Bushnell (1998), Halseth (1998), is the inclusion of the commitment aspect of hydro and forward contracts through a closed loop analysis.

In Paper 2, the contribution lies in the combination of hydro scheduling and contract management using stochastic programming, and in the discussions. Many of the ideas have been presented before, for example in the finance literature. Anderson & Danthine (1980) presented a mean-variance framework for integrated risk management and production planning. Assuming that the production cost function is deterministic and separable (which does not apply to the electricity industry), they outline a theory for this general problem. In his master's thesis, Ranatunga (1995) introduces a stochastic dynamic model for the joint problem of operating a single thermal generation plant and dynamic hedging via forward contracts. That work, and Paper 2 can be seen as stochastic programming extensions to the mean-variance models for hedging and joint production. However, on the basis of the research reported in Paper 2 we dare take credit for changing how Norwegian energy researchers in particular, and other stochastic programmers in general, think about risk aversion in the context of hydro scheduling and contract management. Inducing such changes is part of what research is all about.

For Paper 3 on comparing stochastic models it can be argued that there is little new since the test methodology is generally known, however we communicate qualitative aspects of the difference between the dynamic stochastic and the fixed mix approaches, and report a quantitative result on the performance difference for a given case.

Regarding Paper 4, I contributed to the derivative pricing model, which is new, but which nevertheless uses known theoretical results.

Paper 5 does not contain research contributions.

6 Conclusions and future work

This thesis has focused on portfolio management using stochastic programming, with an emphasis on joint hedging and electricity scheduling in deregulated markets.

Paper 1 investigates the potential for use of market power in the Scandinavian electricity market. We find that the market is functioning reasonably well. Paper 2 explains a new model for the portfolio management problem of hydropower producers using both futures/forwards, options and the flexibility of the physical generation resources. Paper 3 explains how to compare stochastic decision models via simulation, and applies this to two approaches to an insurance portfolio problem; a dynamic stochastic approach and a fixed mix approach. Since the models will be rerun at regular intervals, the performance of the fixed mix approach is found to be closer to the dynamic approach when compared via simulation than when comparing the outcome of the optimization at the first stage. Paper 4 introduces a portfolio model for a casualty insurer that considers bearing catastrophe risks. It is advantageous for such insurers to use financial reinsurance, i.e. derivatives that pay off following catastrophes, and we develop a new model to price such derivatives.

The portfolio management model developed in Paper 2 can be generalized to other businesses. Consider for example a risk averse metal smelter, facing uncertainty both in the price of its main input factor, electricity, and in the price of its output products. If there is some flexibility in the production process, there can be gains from coordinating the production planning and the risk management function, which usually mitigate risk through the financial metal markets.

The scenario generation process and the pricing of derivatives could be made even more consistent with economic theory. An alternative is to use the price model such as the one suggested in Section 3.3 together with a model for the hydro inflow processes and the other stochastic parameters, find the risk neutral version of these processes (e.g. using futures price data), and then proceed with your favorite stochastic programming/sceduling method using these new processes, appropriately discretized.

References

Amundsen, E. S. & Bergman, L. (1998), The performance of the deregulated electricity markets in Norway and Sweden: A tentative assessment, Working paper No. 0798, Dept. of Economics, Univ. of Bergen, Norway.

Amundsen, E. S., Bergman, L. & Andersson, B. (1998), Competition and prices on the emerging Nordic electricity market, Working Paper Series in Economics and Finance No. 217, Stockholm School of Economics.

Anderson, R. W. & Danthine, J.-P. (1980), 'Hedging and joint production: Theory and illustrations', *Journal of Finance* **35**(5), 487–498.

Bellman, R. (1957), *Dynamic Programming*, Princeton University Press, Princeton, N.J.

Birge, J. R. (1985), 'Decomposition and partitioning methods for multistage stochastic linear programs', *Operations Research* **33**, 989–1007.

Bjerksund, P. (1991), Contingent claims evaluation when convenience yield is stochastic: Analytic results, Working Paper, Norwegian School of Economics and Business Administration.

Black, F. (1976), 'The pricing of commodity contracts', *Journal of Financial Economics* **3**, 167–179.

Brealey, R. A. & Myers, S. C. (2000), *Principles of Corporate Finance*, 6th edn, McGraw-Hill, Boston.

Bushnell, J. (1998), Water and power: Hydroelectric resources in the era of competition in the Western US, University of California Energy Institute, POWER Working Paper-056.

Chen, Z.-L. & Powell, W. (1999), 'Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse', *Journal of Optimization Theory and Applications* **102**, 497–524.

Christiansen, D. S. & Wallace, S. W. (1998), 'Option theory and modeling under uncertainty', *Annals of Operations Research* **82**, 59–82.

Cox, J. C., Ingersoll, J. E. & Ross, S. A. (1981), 'The relation between forward prices and futures prices', *Journal of Financial Economics* **59**, 321–346.

de Lange, P. E. (1999), Asset Liability Management in the Insurance Business. Using Stochastic Programming, PhD thesis, Norwegian University of Science and Technology, Trondheim.

DeMarzo, P. M. & Duffie, D. (1995), 'Corporate incentives for hedging and hedge accounting', *Review of Financial Studies* **8**(3), 743–771.

Deng, S. (1999), Financial Methods in Competitive Electricity Markets, PhD thesis, University of California, Berkeley.

Dixit, A. K. & Pindyck, R. S. (1994), *Investment under uncertainty*, Princeton University Press, Princeton, NJ.

Fleten, S.-E. (1995), Portfolio management in the electricity industry, Master's thesis, The Norwegian Institute of Technology, University of Trondheim.

Fleten, S.-E. (1998), Evaluering av beslutningsregler for vannkraftprodusenters sikringsporteføljer, Unpublished note, NTNU.

Fleten, S.-E. & Wallace, S. W. (1998), Power scheduling with forward contracts, *in* A. Løkketangen, ed., 'Proceedings, fifth meeting of the Nordic section of the Mathematical Programming Society', Molde, Norway.

Fleten, S.-E., Jørgensen, T. & Wallace, S. W. (1998), 'Real options and managerial flexibility', *Telektronikk* **94**(3/4), 62–66.

Fleten, S.-E., Wallace, S. W. & Ziemba, W. T. (1997), Portfolio management in a deregulated hydropower-based electricity market, *in* E. Broch, D. Lysne, N. Flatabø & E. Helland-Hansen, eds, 'Hydropower '97', Balkema, Rotterdam, Trondheim, Norway, pp. 197–204.

Fleten, S.-E., Wallace, S. W. & Ziemba, W. T. (1999), A dynamic portfolio approach to hedging in electricity generation, *in* M. Dessouky, S. Eid & E. Shayan, eds, 'the 26'th International Conference on Computers and Industrial Engineering', Melbourne, Australia, pp. 322–327.

Froot, K. A., Scharfstein, D. S. & Stein, J. C. (1993), 'Risk management: Coordinating corporate investment and financing policies', *Journal of Finance* **48**(5), 1629–1658.

Gibson, R. & Schwartz, E. S. (1990), 'Stochastic convenience yield and the pricing of oil contingent claims', *Journal of Finance* **45**(3), 959–976.

Gjelsvik, A. & Wallace, S. W. (1996), Methods for stochastic medium-term scheduling in hydro-dominated power systems, Report EFI TR A4438, Norwegian Electric Power Research Institute, Trondheim.

Grundt, A., Dahl, B. R., Fleten, S.-E., Jenssen, T., Mo, B. & Sætness, H. (1998), 'Integrert risikostyring (integrated risk management)'. Energiforsyningens Fellesorganisasjon-Pub. nr. 255, Oslo.

Hakansson, N. H. (1971), 'On optimal myopic policies with and without serial correlation in yields', *Journal of Business* **44**, 324–334.

Halseth, A. (1998), 'Market power in the Nordic electricity market', *Utilities Policy* **7**(4), 259–268.

Harvey, A. C. (1989), *Forecasting, structural time series models and the Kalman filter*, Cambridge University Press, Cambridge.

Hilliard, J. E. & Reis, J. (1998), 'Valuation of commodity futures and options under stochastic convenience yields, interest rates, ans jump diffusions in the spot', *Journal of Financial and Quantitative Analysis* **33**(1), 61–86.

Holthausen, D. M. (1979), 'Hedging and the competitive firm under price uncertainty', *American Economic Review* **69**, 989–995.

Høyland, K. (1998), Asset liability management for a life insurance company. A stochastic programming approach, PhD thesis, Norwegian University of Science and Technology, Trondheim.

Hvarnes, H. (1998), Portfolio management for TrønderEnergi, Master's thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology.

Johnsen, T. A., Verma, S. K. & Wolfram, C. D. (1999), Zonal pricing and demand-side bidding in the Norwegian electricity market, University of California Energy Institute, POWER Working Paper-063.

Jørgensen, T. (1999), Project Scheduling as a Stochastic Dynamic Decision Problem, PhD thesis, Norwegian University of Science and Technology, Trondheim.

Kaldor, N. (1939), 'Speculation and economic stability', *Review of Economic Studies* **7**, 1–27.

Kall, P. & Wallace, S. (1994), *Stochastic Programming*, Wiley, Chichester.

Klaassen, P. (1998), 'Financial asset pricing theory and stochastic programming models for asset/liability management: A synthesis', *Management Science* **44**(1), 31–48.

Markowitz, H. (1952), 'Portfolio selection', *Journal of Finance* **7**, 77–91.

Mian, S. (1996), 'Evidence on corporate hedging policy', *Journal of Financial and Quantitative Analysis* **31**, 419–439.

Mo, B., Gjelsvik, A. & Grundt, A. (2000), Integrated risk management of hydro power scheduling and contract management, Paper submitted to IEEE Transactions on Power Systems.

Modigliani, F. & Miller, M. H. (1958), 'The cost of capital, corporation finance, and the theory of investment', *American Economic Review* **48**, 261–297.

Mossin, J. (1968), 'Optimal multiperiod portfolio policies', *Journal of Business* **41**, 215–229.

Pereira, M. V. F. (1989), 'Optimal stochastic operations scheduling of large hydroelectric systems', *International Journal of Electrical Power & Energy Systems* **11**(3), 161–169.

Ranatunga, R. A. S. K. (1995), Risk averse operation of an electricity plant in an electricity market, Master's thesis, School of Electrical Engineering, University of New South Wales. ME dissertation.

Ross, M. P. (1996), Corporate hedging: What, why and how?, Working Paper, University of California, Berkeley.

Salahor, G. (1998), 'Implications of output price risk and operating leverage for the evaluation of petroleum development projects', *Energy Journal* **19**(1), 13–46.

Sandmo, A. (1971), 'On the theory of the competitive firm under price uncertainty', *American Economic Review* **61**, 65–81.

Schwartz, E. S. (1997), 'The stochastic behavior of commodity prices: Implications for valuation and hedging', *Journal of Finance* **52**(3), 922–973.

Sharpe, W. F. (1964), 'Capital asset prices: A theory of market equilibrium under conditions of risk', **19**, 425–442.

Sick, G. (1995), Real options, *in* R. Jarrow, V. Maksimovic & W. T. Ziemba, eds, 'Finance', Vol. 9 of *Handbooks in Operations Research and Management Science*, Elsevier, Amsterdam, pp. 631–691.

Smith, C. W. & Stulz, R. M. (1985), 'The determinants of firm's hedging policies', *Journal of Financial and Quantitative Analysis* **20**(2), 391–405.

Stulz, R. M. (1990), 'Managerial discretion and optimal financing policies', *Journal of Financial Economics* **26**, 3–27.

Viaene, J.-M. & Zilcha, I. (1998), 'The behavior of competitive exporting firms under multiple uncertainty', *International Economic Review* **39**(3), 592–609.

Working, H. (1949), 'The theory of the price of storage', *American Economic Review* **39**, 1254–1262.

Ziemba, W. T. & Mulvey, J. M., eds (1998), *Worldwide Asset and Liability Modeling*, Cambridge University Press, Cambridge, U. K.

Paper 1:

Fleten, S-E., Lie, TT. (2000). Market power in the Nordic electricity market - the effect of hydro and contracts. Paper presented at the *Annual European Energy Conference*, Bergen, 31.8-1.9.2000

Abstract:
http://wgate.bibsys.no/gate1/SHOW?objd=r01007372&base=FORSKPU2&lang=N&bibk=a

Paper 2:

Fleten, S-E, Wallace, S.W & Ziemba W.T. (2002). Hedging electricity portfolios via stochastic programming. In C. Greengard and A. Ruszczynski (Eds), *Decision Making Under Uncertainty: Energy and Power*, vol. 128 of IMA Volumes on Mathematics and Its Applications, pp. 71-93. New York: Springer-Verlag.

Abstract:
http://wgate.bibsys.no/gate1/SHOW?objd=r02014936&base=FORSKPU2&lang=N&bibk=a

Paper 3

The performance of stochastic dynamic and fixed mix portfolio models

The performance of stochastic dynamic and fixed mix portfolio models

S.-E. Fleten[a] K. Høyland[b] S. W. Wallace[a]

[a]*Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, N–7491 Trondheim, Norway*
Stein-Erik.Fleten@iot.ntnu.no *Stein.Wallace@iot.ntnu.no*

[b]*Gjensidige Asset Management, N-1324 Lysaker, Norway*
Kjetil.Hoyland@gjensidige.no

Abstract

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

Keywords: Mathematical programming, asset liability management, stochastic programming, simulation.

1 Introduction

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

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

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

Kusy & Ziemba (1986) test a two-stage simple recourse model and compare this to a stochastic decision tree model. However, a new scenario tree is not generated each time the horizon is rolled forward. Cariño & Turner (1998) compare the fixed mix-rule with a true dynamic stochastic optimization-based model. The model horizon is not rolled forward, and no out-of-sample data are used in the test procedure. Cariño, Myers & Ziemba (1998) show the historical performance of an asset liability model used by Yasuda Kasai, a Japanese insurance company. Vassiadou-Zeniou & Zenios (1996) and Zenios, Holmer, McKendall & Vassiadou-Zeniou (1998) also do validation backtesting based on historical data. Birge (1982) compares a class of stochastic linear programs with the corresponding expected value problem. Independently, Kouwenberg (1998) has developed and tested a pension fund asset liability management model utilizing rolling horizon simulations.

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

2 The stochastic dynamic and the fixed mix approaches

We use a multi period, stochastic asset liability management model developed for the Norwegian mutual life insurance company Gjensidige Spareforsikring. A mathematical description is provided in Høyland & Wallace (1999*a*). The portfolio selection problem is modeled at the strategic level, where capital is allocated among a few aggregated asset classes such as stocks and bonds. The objective is to maximize the expected portfolio value at the end of the horizon net of costs, subject to rebalancing and legal constraints on balance figures. The costs are composed of transaction costs[1] and imputed costs associated with the violation of the legal constraints. These imputed costs are used as the measure of risk in the model.

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

2.1 How the two approaches differ

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

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

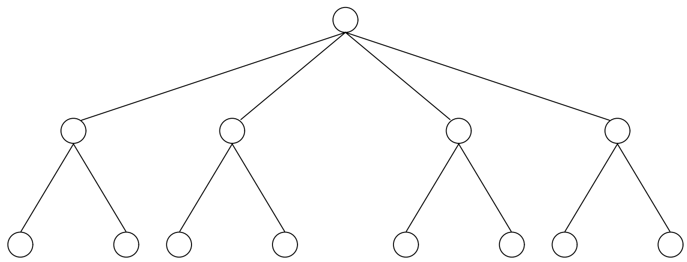

Figure 1: The scenario tree. The nodes represent decisions, while the arcs represent realizations of the uncertain variables.

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

In the fixed mix model we assume the following decision rule: The portfolio is rebalanced to fixed proportions (for instance 70% bonds and 30% stocks) at all future decision nodes. This means that at all decision nodes, assets are bought and sold so that the fixed mix is maintained. The optimization problem is to find the proportions that maximize the objective function. Note the difference between a fixed mix model and a fixed mix investment strategy. In the fixed mix investment strategy, the proportions are kept constant over a long investment horizon, while applying a fixed mix model means that the proportions are changed every time the model is rerun. In this paper we compare the results of a fixed mix model with the results of a stochastic, dynamic model. See Perold & Sharpe (1988) for a description of the fixed mix and other decision rules.

The fixed mix approach does not require the tree structure, but allows the uncertainties to be described in terms of streams of outcomes (or scenarios), as illustrated in Figure 2. However, we have applied the same scenario tree for both approaches. Doing this ensures that we get comparable results and means that the only difference between the two model formulations are the constraints in the fixed mix model assuring that the portfolio is rebalanced to the fixed mix at every decision node. If the approaches were to be compared with respect to both solution time and solution quality, this procedure might disfavour the fixed mix approach[2]. However, this problem does not arise since this paper is concerned with comparing the quality of the solutions given the quality on the input data, and not solution times.

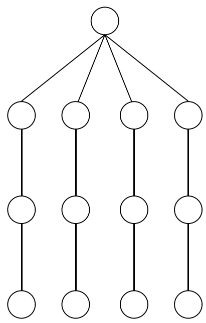

Figure 2: Possible description of the uncertainties with the fixed mix approach.

2.2 Scenario generation

Generation of scenarios can be based on simulation or construction. The applied scenario generation method is a combination of the two. We let the decision maker specify the market expectations by any statistical properties that are considered relevant for the problem to be solved, and construct the tree so that these statistical properties are preserved. This is done by letting (random) outcomes and probabilities in the scenario tree be decision variables in a nonlinear optimization problem where the objective is to minimize the square distance between the statistical properties specified by the decision maker and the statistical properties of the constructed tree (Høyland & Wallace 1999*b*).

Generally, expectations for financial markets will depend on the state of the economy. Some statistical properties are clearly state dependent, while

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

others might be specified independently of the state. As an example of state dependency, consider the volatility of stocks. Empirical studies have shown the effect called volatility clumping, meaning that the volatility increases after a period of extreme returns. We model this effect by letting the volatility (standard deviation) in period $t+1$ be a function of the outcomes' deviation from the mean in period t. For more details of the scenario generation system, see Section 4.1.

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

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

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

3.1 The test procedure

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

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

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

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

3.2 Potential error sources

If the simulation scenarios differ from the optimization scenarios then a potential bias toward one of the approaches is introduced in the test procedure. What could happen is that the difference between the information in the optimization scenarios and the simulation scenarios makes the dynamic approach a priori worse off. The dynamic approach uses the information in

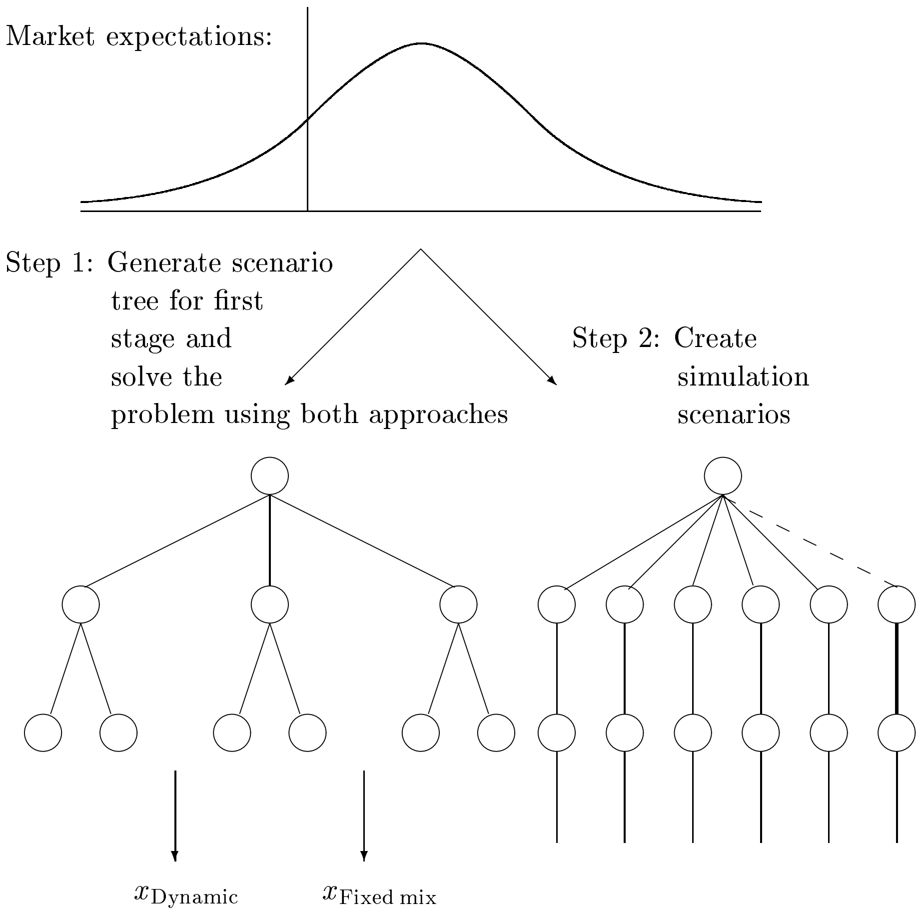

Figure 3: Steps 1 and 2 of the test procedure. In step 1 we construct optimization scenarios that are consistent with the market expectations, and optimize using both approaches to obtain first stage decisions, denoted x_{Dynamic} and $x_{\text{Fixed mix}}$. In step 2 we generate a high number of simulation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

the optimization scenarios better than the fixed mix approach, but it could be using misleading information so that it consistently makes relatively bad decisions.

Modeling the problem with a finite horizon is another potential error source.

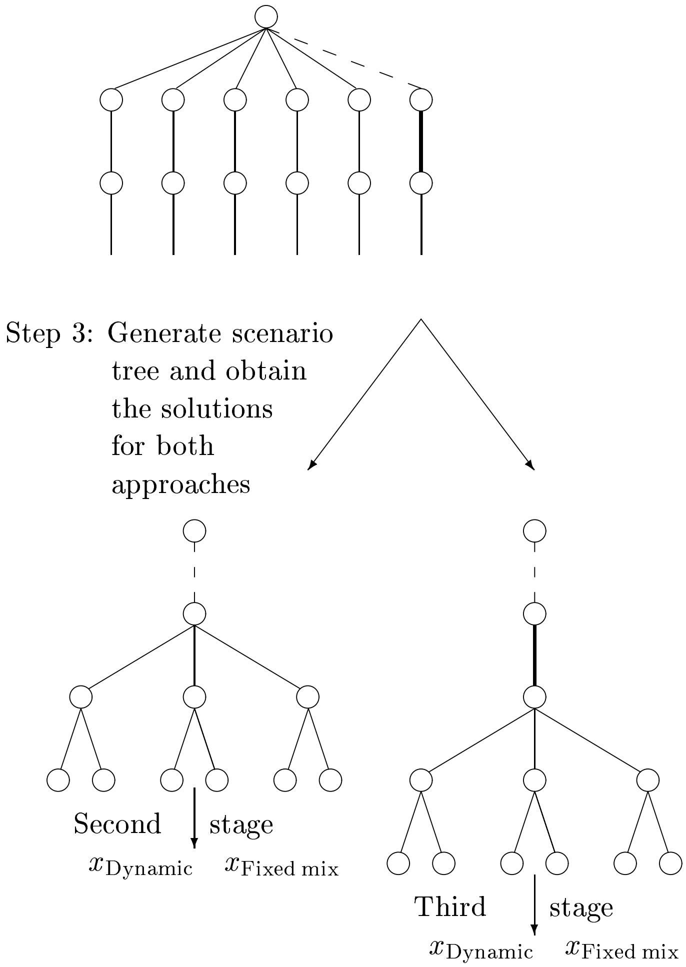

Figure 4: Step 3: Generating conditional trees and new solutions. After the outcome represented by the dotted line, new information is received, and the model horizon is rolled forward. A new scenario tree is generated conditional on the information in that outcome. The model is solved to obtain a new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

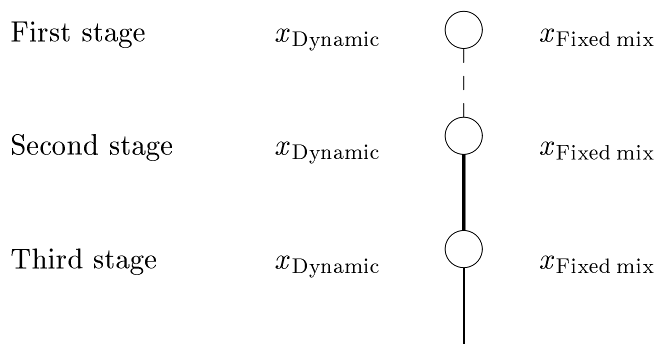

Figure 5: Outcomes and decisions for one scenario. Since decisions and outcomes are known for this scenario, in step 4 we can calculate the objective function value at the end of the third period. So each simulated scenario gives rise to one performance number for the dynamic stochastic approach, and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

With few periods, the potential gain for a dynamic approach over a static one is small.

4 Numerical example

This section presents the numerical results. We study a mutual life insurance company with a portfolio similar to that of Gjensidige Spareforsikring anno 1997, using a somewhat distorted data set. Assume that the company invests in the money markets (i.e. bonds with less than a year to maturity) and stocks. We let the money markets be represented by the three month NIBOR (Norwegian InterBank Offering Rate) and stocks be represented by the Oslo Stock Exchange Total index. The starting values of assets and liabilities are given in Table 1. The liabilities are the basis for the calculation of the legal requirements, and the risk is modeled by penalizing the violations of the legal requirements. For details, see the model formulation in Høyland & Wallace (1999*a*).

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

Table 1: Starting balance.

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

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

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

4.1 Scenario generation

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

4.1.1 Specification of market expectations

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

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

Table 2: Percentiles of the marginal distributions. For the money market, the market views are expressed in terms of expectations for the interest rates, while for stocks the expectations are given in terms of total returns.

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

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

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

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

$$\sigma_{it} = (1 - \nu_i)\overline{\sigma}_{it} + \nu_i \left| x_{i,t-1} - \mu_{i,t-1} \right| , \tag{1}$$

where $i \in \{s, b\}$ is the asset class index, either stocks (s) or interest rates (b), $\nu_i \in [0, 1]$ is the volatility clumping parameter for asset class i (a large ν_i leads to a large degree of volatility clumping), $\overline{\sigma}_{it}$ is the average standard deviation of the outcome of asset class i in period t, x_{it} is the outcome of asset class i in period t, and μ_{it} is the expected outcome of asset class i in period t. This way the volatility will increase after extreme returns, and decrease after more normal returns, in line with empirical observations[3].

For interest rates we model a mean reversion effect[4], and let the expected

[3]There is some evidence (Billio & Pelizzon 1997) that volatility increases more when prices go down, however, for simplicity we have chosen to model this symmetrically.

[4]Mean reversion means that interest rates tend to revert to an average level. When interest rates are high, the economy slows down and interest rates tend to fall, and when interest rates are low, the economy booms and interest rates tend to rise. The

value of the interest rate at the end of period $t > 1$ be given by

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

where κ is the mean reversion factor (a high κ leads to a large degree of mean reversion), α is the mean reversion level and x_{bt} is the interest rate at the end of period t.

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

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

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

4.1.2 Generating optimization scenarios

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

For the present optimization problem (corresponding to step 1 in Section 3.1), a three period (four stage) scenario tree is generated. This has 50 outcomes in the first period, and 5 in the second and the third, yielding a total

---

mean reversion effect usually needs more than two or three years to manifest itself to a significant degree, and a low value for the mean reversion factor is chosen in the numerical example.

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

Table 3: Market expectations used for generation of optimization scenarios. The expectations for the first period are derived from the marginal distributions, see Appendix A for details. While the skewness and kurtosis[6]are assumed state independent, the expected value and the standard deviation are assumed state dependent.

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

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

4.1.3 Generating simulation scenarios

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

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

4.2 Numerical results

Note that since the fixed mix formulation leads to a non-convex optimization model, we would usually need global optimization routines to be sure that the optimal solution is found. However, testing shows that for our data sets, the problem is convex[7].

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

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

where $\mathcal{S}$ is the set of all scenarios, $\mathcal{J}$ is the set of shortfall types, P_s is the probability that scenario s occurs, c_{jt} is a weight parameter allowing the decision maker to weigh the relative importance of different shortfall types, ρ_{st} is a path dependent discount factor[8], depending on the scenario $s \in \mathcal{S}$ and the time period $t \in \mathcal{T}$, x_{sjt} is the shortfall of type $j \in \mathcal{J}$, in time period $t \in \mathcal{T}$ under scenario $s \in \mathcal{S}$. The objective is minimized for different required levels of minimum expected portfolio returns.

Figure 6 shows the results of solving the first stage models (referring to step 1 in Section 3.1) for both approaches for different levels of minimum expected portfolio return. We see that in this *in-sample* comparison the dynamic approach clearly dominates the fixed mix approach since the target expected returns are achieved for lower levels of risk. For instance, for an expected return of 13.0%, the (square root of the) risk, $\sqrt{F}$, is 0.18 for

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

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

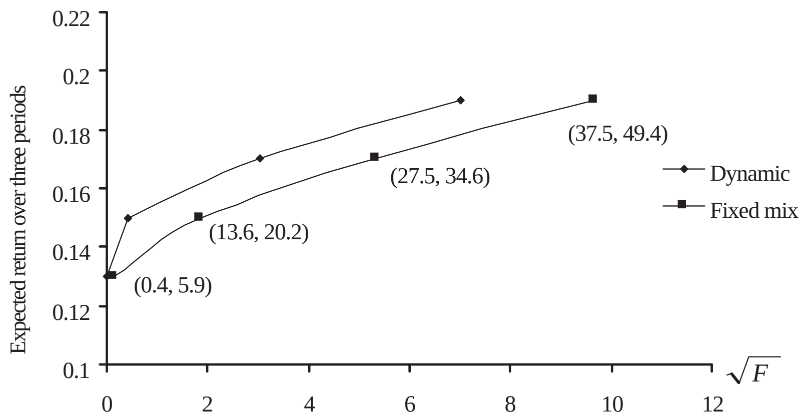

Figure 6: Optimization results of step 1 in Section 3.1. The efficient frontiers show the tradeoff between the expected return and the risk, and is obtained by solving the models for different required rates of expected returns. The numbers in parantheses show the initial investment in stocks for the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

the dynamic approach and 1.84 for the fixed mix approach. Observe that the fixed mix approach must choose a more aggressive portfolio (i.e. more stocks) than the dynamic approach to achieve the target expected in-sample return.

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

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

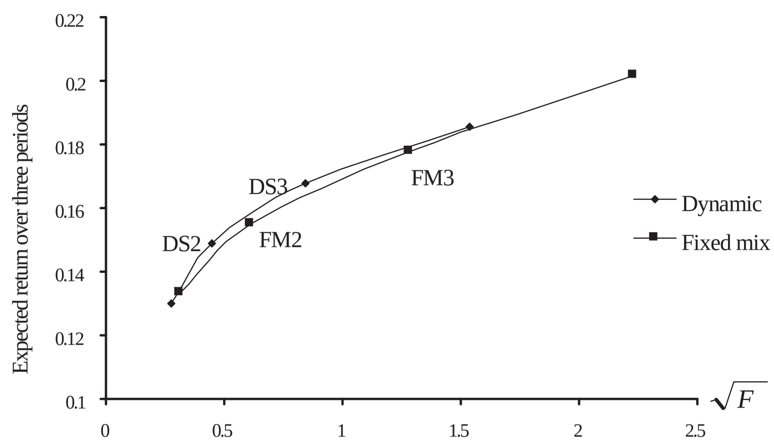

Figure 7: Results from the simulation procedure. Each point on the efficient frontiers are found by solving 600 optimization programs, one for each period in each simulation.

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

Given that the dynamic approach gives lower expected return and risk, a *conservative* test, i. e. one that is not biased in favor of the dynamic approach, can be constructed using a convex combination where there is a relatively large weight on expected return. The following utility function is employed:

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

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

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

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

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

The test compares the performance of the two approaches for the pairs (DS2, FM2) and (DS3,FM3) displayed in Figure 7. There will be a weight ω for each pair. Weights are chosen such that utility maximization yields the points FM2 and FM3 respectively, given that one must choose a point on the fixed mix efficient frontier[9].

The result of the test is shown in Table 4. The probability that the dynamic

Table 4: Test of whether the dynamic approach dominates the fixed mix approach. A hundred thousand bootstrap runs were used. The p-value reported is the (bootstrap) probability of dominance. The number of observations is 200.

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

approach dominates the fixed mix approach is higher than 50%, however the difference between the performances is not statistically significant. If

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

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

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

we choose a weight corresponding to the average of the slopes of the efficient frontiers at FM2 and DS2, we get a probability of dominance of 0.8650. For (FM3, DS3), we get $p = 0.8531$.

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

The simulation scenarios and the optimization scenarios are necessarily different, for example regarding the structure of the evolution of information. In the optimization scenarios the number of outcomes per stage increases exponentially, but in the simulation scenarios the number is constant at 200 for stages 2 and 3. In particular, the simulation scenarios do not contain so-called worst case scenarios where asset class returns are negative in all subsequent periods. Both the simulation scenarios and the optimization scenarios satisfy the specified statistical properties, but for the optimization scenarios, after the worst case outcome in the first and second period, one of the subsequent outcomes also has a low return. In the simulation scenarios there is no reason why a low return outcome should follow a low return outcome in the previous period. These worst case scenarios have a significant effect on the perceived risk of the portfolio approaches as seen in Figure 6, due to the quadratic nature of risk. Although the probability of these worst case scenarios is very small, the shortfall costs accumulated are so large that they affect the overall expected shortfall cost. This can be seen from the difference in the scale of risk in Figures 6 and 7, where the risk in the out-of-sample case is much smaller[10].

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

Further comparing Figures 6 and 7 we see that the out-of-sample expected return of the dynamic approach is roughly the same as in-sample, while for the fixed mix approach the expected return is higher. This is because the fixed mix approach generally has a higher share of stocks than the dynamic approach. At the second and third stages in the simulation, the fixed mix approach will choose more aggressive portfolios in order to fulfill the return requirement. This explains why the efficient frontier for the fixed mix approach is shifted not only to the left, but also upwards from Figure 6 to Figure 7.

5 Conclusion and future work

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

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

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

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

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

References

Billio, M. & Pelizzon, L. (1997), Pricing options with switching volatility, WP 97.07, Department of Economics, University of Venice.

Birge, J. R. (1982), 'The value of the stochastic solution in stochastic linear programming with fixed recourse', *Mathematical Programming* **24**, 314–325.

Cariño, D. R. & Turner, A. L. (1998), Multiperiod asset allocation with derivative assets, *in* W. T. Ziemba & J. M. Mulvey, eds, 'Worldwide Asset and Liability Modeling', Cambridge University Press, Cambridge, U.K., pp. 182–204.

Cariño, D. R., Myers, D. H. & Ziemba, W. T. (1998), 'Concepts, technical issues and uses of the Russell-Yasuda Kasai financial planning model', *Operations Research* **46**(4), 450–462.

Efron, B. & Tibshirani, R. J. (1993), *An Introduction to the bootstrap*, Chapman & Hall, New York.

Høyland, K. & Wallace, S. W. (1999*a*), Analyzing legal restrictions in the Norwegian life insurance business using a multi-period asset liability model, Accepted for publication in European Journal of Operations Research.

Høyland, K. & Wallace, S. W. (1999*b*), 'Generating scenario trees for multi stage decision problems', *Management Science.* To appear.

Kouwenberg, R. R. P. (1998), Scenario generation and stochastic programming models for asset liability management, Accepted for publication in European Journal of Operational Research.

Kusy, M. I. & Ziemba, W. T. (1986), 'A bank asset and liability management model', *Operations Research* **34**(3), 356–376.

Lurie, P. M. & Goldberg, M. S. (1998), 'An approximate method for sampling correlated random variables from partially-specified distributions', *Management Science* **44**(2), 203–218.

Perold, A. K. & Sharpe, W. F. (1988), 'Dynamic strategies for asset allocation', *Financial Analysts Journal* **January**, 16–27.

Vassiadou-Zeniou, C. & Zenios, S. A. (1996), 'Robust optimization models for managing callable bond portfolios', *European Journal of Operational Research* **91**, 264–273.

Zenios, S. A., Holmer, M. R., McKendall, R. & Vassiadou-Zeniou, C. (1998), 'Dynamic models for fixed-income portfolio management under uncertainty', *Journal of Economic Dynamics & Control* **22**(10), 1517–1541.

A Specifying market expectations

Expressing market expectations can be done in many ways. We have chosen to let the decision maker supply percentiles for the marginal cumulative probability distribution functions for all uncertain variables (this is how the asset allocation managers in the life insurance company prefer to express the market expectations), see Figure 8. An approximating cumulative distribution function is fitted to these percentiles. The properties that are listed in Table 3 are calculated from the function that is fitted to the percentiles.

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

The approximating cumulative distribution functions are found using a NAG (Numerical Algorithms Group) C Library routine for interpolating data. This method does not guarantee that the second derivative changes sign only once, in the case of Figure 8 causing a somewhat peculiar form near the top of the distribution. However, the resulting function is monotonic, so we are guaranteed that the curve will have the properties of a cumulative distribution function, and that the user specified percentiles are fit exactly (including the 0% and 100% points).

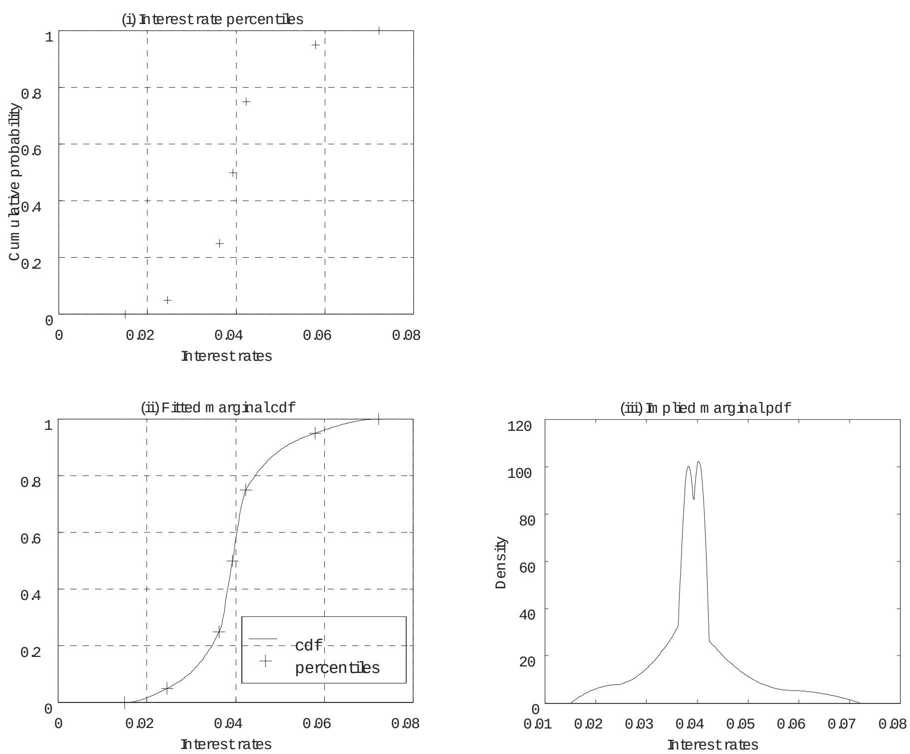

Figure 8: (i) User-supplied percentiles for interest rates at the end of the first period. (ii) A cumulative distribution function is fitted to the percentiles. (iii) The probability distribution function for interest rates.

B Generating the scenarios that are input to the simulation

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

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

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

The following notation is used:

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

$x_{qb|s,i}$ percentile of interest rates given a return on stocks

The new mean, given the return on stocks, is

$$\mu_{s|b} = \mu_b + \rho_{sb}\frac{\sigma_b}{\sigma_s}\left(x_{s'} - \mu_s\right). \tag{7}$$

The new standard deviation is

$$\sigma_{b|s} = \sigma_b\left(1 - \rho_{sb}^2\right). \tag{8}$$

So the linear transformation of the percentiles is given by

$$x_{qb|s,i} = \mu_{s|b} + \frac{\sigma_b}{\sigma_{b|s}}\left(x_{qb,i} - \mu_b\right) \tag{9}$$

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

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

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

Paper 4

Modeling financial reinsurance in the casualty insurance business

Modeling financial reinsurance in the casualty insurance business

S.-E. Fleten[a] A. A. Gaivoronski[a] P. E. de Lange[b]

[a] *Department of Industrial Economics and Technology Management*
NTNU, N–7491 Trondheim, Norway
Stein-Erik.Fleten@iot.ntnu.no *Alexei.Gaivoronski@iot.ntnu.no*

[b] *Gjensidige NOR Spareforsikring, N–7469 Trondheim, Norway*
Petter-E.de-Lange@gjensidige.no

Abstract

In this paper we examine the rationale for financial reinsurance in the casualty insurance business. This concept refers to an investment strategy that uses the financial markets to hedge insurance risk. The casualty insurer takes positions in derivatives on underlying assets, whose prices are highly positively or negatively correlated with specific insurance risks. We formulate the asset liability management problem for a casualty insurer, in the context of a dynamic, stochastic portfolio selection model. The optimal solution is a portfolio of financial assets earning a return that, including premium income from written policies, will guarantee compliance with the legal statutes in all but a few extreme states of nature. In this context we compare properties of optimal portfolios with and without the possibility of financial reinsurance. We let an alleged representative policyholder, endowed with a linear plus negative exponential utility function, evaluate the various optimal portfolios. We find that, in a regulated environment and when policyholders' utility functions exhibit reasonable levels of risk aversion, portfolios reflecting financial reinsurance dominate portfolios that are not financially reinsured.

Keywords Financial reinsurance, risk neutral valuation, arbitrage pricing, stochastic optimization.

1 Introduction

Using derivatives to hedge portfolios of financial securities against the downside volatility of security price processes, is by now common in practical portfolio management. Also textbooks on derivative assets often devote a whole chapter to trading strategies and market risk management (Hull 1997, Cox & Rubinstein 1985). When constructing portfolios by means of option hedging, investors seek to reduce downside portfolio risk, while retaining a share of the upside potential. The random sequence of return from the hedged portfolio is assumed to fit the structure of risk preferences embedded in the investors' utility functions. One might term this approach to portfolio selection *asset management*, because no explicit attention is paid to the nature of the investors liabilities.

An asset liability management (ALM) approach to portfolio selection, explicitly models the link between risk preferences and investors perception of actual business risk. When the investor is a casualty insurance company, risk perception is connected to the structure of his liabilities, i.e. the loss payments agreed to pay the policyholders with the occurrence of a qualifying accident. Both life insurers and casualty insurers are subject to regulations when they take positions in the financial markets. Allegedly authorities enforce regulations to protect policyholders interests from too risky investment strategies, which they assume will result from leaving the insurance business free to invest at its own discretion.

In this paper we examine the asset liability management problem as it applies to casualty insurers when these are viewed as portfolio managers. We assume that what the casualty insurance industry perceive as business risk, is failing to comply with the legally enforced regulations. Put differently we might say that the authorities have forced a specific utility function upon the casualty insurers. This function is concave, reflecting risk averse investment strategies. In the next section we define the ALM problem that must be solved by a casualty insurer. This portfolio selection problem is conceptually different and more complex than the analogous portfolio selection problem of a life insurer. A distinguishing characteristic of the former optimization problem is the sequence of *random* claims, which a casualty insurer is obliged to pay its customers. The claims distributions have large variances. Therefore it would be unsatisfactory to represent claims in terms of expected values, i.e. deterministically; our model must be able to evaluate the economic performance of the insurance company in carefully chosen possible realizations in the tail of the claims distributions. It is the random nature of claims, which provides a rationale for introducing the concept of financial reinsurance into the ALM problem of a casualty insurer. In contrast, a life insurer does not make a serious mistake in regarding his liabilities as being deterministic, when formulating his ALM problem.

The concept of financial reinsurance designates an investment strategy that uses derivatives to hedge insurance risks. This presupposes the existence of underlying assets whose price processes are highly correlated with the distribution of specific insurance risks. The underlying assets need not be traded, as long as one can trade derivatives written on them. The key feature of financial reinsurance is the property that the derivatives will pay off in states where the insurance policies demand huge loss payments to policyholders. We note that this hedging scheme is not confined to non-life insurers. It applies equally well to any portfolio manager with stochastic liabilities, whose payoff structure is highly correlated with financial securities (or other phenomena) on which derivative assets can be bought.

We use a multistage stochastic programming model to solve the portfolio selection problem of a casualty insurer. The model is set up in discrete time on a discrete event space. The main contribution of our paper is to demonstrate that financial reinsurance offers an opportunity, at a certain cost, for casualty insurers to select portfolios whose random payoff structure more accurately fits the sequence of future random liabilities. Thus financial reinsurance reduces the volatility of net cash flows. We then let the individual policyholders evaluate financial reinsurance. We assume the existence of a representative policyholder endowed with a linear plus negative exponential utility function, and let him evaluate the utility of wealth generated by the optimal portfolios selected by the insurance company. Underlying the analysis is the assumption that the policyholders own either the insurance company, or the funds being managed by the casualty insurer. Our findings support intuition in that when policyholders' utility functions exhibit reasonable levels of risk aversion, financially reinsured portfolios dominate non-reinsured regulated portfolios. This paper also offers an illustration of an approach to model derivative assets in stochastic optimization models.

Other asset liability models are described by Dert (1995), Consigli & Dempster (1998), Gaivoronski & de Lange (1998), Gaivoronski, Høyland & de Lange (1998), Gaivoronski & Stella (1998), Kusy & Ziemba (1986), Dupačová, Bertocchi & Moriggia (1998), Cariño & Ziemba (1998), King (1993), Mulvey & Vladimirou (1992), Høyland & Wallace (1999*a*), Cariño & Turner (1998) and Zenios, Holmer, McKendall & Vassiadou-Zeniou (1998). These models are designed to facilitate an investment strategy, yielding an uncertain inflow of money, capable of supporting a sequence of future (possibly uncertain) liability payments. Cariño and Turner include derivative assets in their ALM-model. In order to avoid computational instabilities due to the large range of returns to the funds invested in the derivative assets, they use options in combination with some other asset class like cash. In Section 4.1 we describe our approach to derivative asset pricing, which is consistent with financial theory. We formulate an optimization model that not only produces arbitrage free asset prices, it also selects the appropriate *equilibrium* returns. It turns out that these distributions do not cause problems with computational stability.

The rest of this paper is organized as follows: Section 2 further discusses the issue of financial reinsurance. In Section 3 we present the ALM-model. The derivative pricing method is explained in Section 4. In Section 5 we construct a numerical example, by which we examine the impact of financial reinsurance on the portfolio decisions of a casualty insurer. Section 6 concludes.

2 Problem outline

We define the ALM problem confronting a casualty insurer as that of constructing a dynamic portfolio of financial assets earning a return that, including premium income from written policies, will guarantee compliance with the legal statutes in all but a few rarely occurring states of nature. The insurance company is averse to the risk of such occurences, but the goal of minimizing such risk must be traded off against the desire to achieve maximum expected return. By a dynamic portfolio we mean an investment strategy over time that, in our context, has been designed to meet a sequence of future *uncertain* liability payments. In practice, it is the current time $t = 0$ portfolio decision we are after, but future portfolio revision possibilities affect the current decision.

Because of the high volatility of potential claims associated with specific risks that a casualty insurer might underwrite, the possibility to hedge his earnings is crucial to his business strategy. A standard hedging scheme consists in ceding a portion of the premium income to other insurers, who in return commit themselves to paying a fraction of future realized claims. This kind of *regular* reinsurance is used by most casualty insurers. However, when dealing with more extreme insurance risks, additional means of risk sharing may be desirable. Through the legal regulations the casualty insurer is obliged to match his earnings and liabilities so closely, that claims will be redeemed almost surely, and without the need to borrow. This could mean that some potential lucrative risks are not insured, or rather not insured enough, because there is a small (empirically estimated) probability that huge claims could get the company into financial distress. However, the financial markets with their growing share of derivatives provide an additional hedging opportunity. As mentioned above financial reinsurance designates an investment strategy which uses financial derivatives to hedge insurance risks. The derivative assets are derived from underlying securities, whose price and dividend processes are highly correlated with specific insurance risks. The derivatives pay off in states where the insurance contracts demand huge loss payments. With this hedging opportunity casualty insurers may take on relatively more insurance risk and still comply with the legal regulations.

In this paper we consider a mutual casualty insurer, which by construction is owned by the policyholders. How do the policyholders feel about financial reinsurance? If they had exactly the same objective, i.e. utility function as the casualty insurer, obviously they would approve of whatever is optimal to the company. But presumably they do not think in terms of legal regulations. Their concern is (except paying as little premium as possible!) that the insurer pays out the agreed amount if an accident occurs. We imagine the existence of a representative policyholder who has a linear plus negative exponential utility function. This objective is often used in financial portfolio selection because of its plausible risk aversion properties (Bell 1995, Huang & Litzenberger 1988). Assuming that the linear plus negative exponential utility function is a fair representation of policyholders preferences we ask: Does this function contradict the objective function employed by the insurance company? Put differently, to what extent do the (enforced) incentives of the portfolio managers comply with the interests of the policyholders? These questions are addressed in Section 5.2.

3 Model description

Decisions in the model are made on a state space which is a finite or countable set, whose time index set is denoted $T = (0, 1, 2, \ldots, t, \ldots, \tau - 1)$. The decision process can be visualized by the scenario tree presented in Figure 1. Nodes in the tree are associated with decisions, whereas arcs, connecting the nodes, account for realizations of random variables. The top node reflects the current investment decision, which is the one we are looking for. Subsequent nodes represent conditional decisions, which will be made according to newly revealed information. The scenario tree displays the dynamic feature of the decision making process. The fact that we are able to rebalance the portfolio, incurring transaction costs, at future dates when more information is available is reflected in the current investment decision, which should be less conservative than a static portfolio decision. Transaction costs counteract, but do not offset, this effect. In the particular scenario tree depicted in Figure 1, there are two time periods (three stages) and 4 times 4 = 16 scenarios.

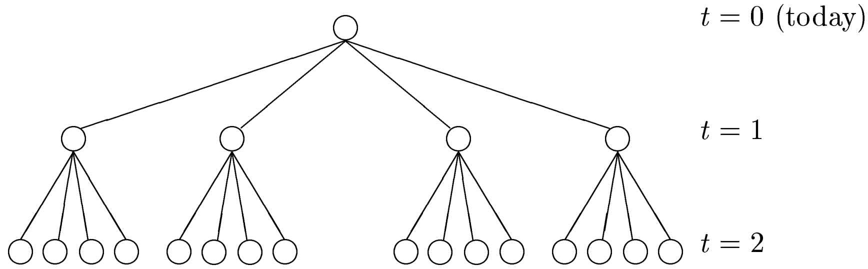

Figure 1: Scenario tree.

The company is initially endowed with a portfolio of assets, enabling it to generate cash flows for an infinite span of time. But the planning period within which decisions are made is limited, and the infinite aspect of the cash flows is represented by the continuing value of the company. In our formulation this is the value, at the horizon, of an infinite sequence of constant payments. Each conditional decision will be made on the basis of the following information:

1. Premium contributions from written insurance policies over the interval $(t-1, t)$.

2. Realized dividends and coupon payments from the portfolio of asset classes over the interval $(t-1, t)$.

3. Price gain/loss on the asset classes over the interval $(t-1, t)$.

4. Income from financial reinsurance, i.e. return on the derivatives over the interval $(t-1, t)$.

5. Interest rates over the interval $(t, t+1)$.

6. Insurance claims directed against the company by the policyholders over the interval $(t-1, t)$.

Then the company will decide how to revise its portfolio of securities, and whether or not to help finance the procurement of new assets through borrowing. The optimal investment policy maximizes the market value of the company adjusted for a penalty term reflecting the costs of violating the legal restrictions. This concave objective function (1) is formulated and explained in the subsequent section.

3.1 Mathematical formulation

In this subsection we formulate the basic equations of the portfolio selection model, which describe the chain of events constituting the ALM process. The collection of equations (2)–(26) applies to each scenario s. The scenarios are associated with probabilities, which appear in the objective (1). The model entities are defined on the following sets:

$I = [1, \ldots, \kappa]$:	the set of all asset classes.
$S = [1, \ldots, \omega]$:	the set of all scenarios.
$\Theta = [1, \ldots, n]$:	the set of all nodes in the scenario tree.
$T = [1, \ldots, \tau - 1]$:	the set of discrete decision points.
, $= [1, \ldots, \tau]$:	the set of discrete time points, including the horizon.
H:	the set of shortfall types.

Beyond time $t = 0$, most variables and parameters are indexed by s, indicating that they are scenario dependent. Recall that a scenario s constitutes a path , i.e. a sequence of nodes, through the scenario tree depicted in Figure 1. To formulate the model, we need the following additional notation:

Scenarios and nodes

$s_v \equiv (\theta_r^v)_{r=0}^{\tau}$: scenario s defined as a sequence v of nodes.
θ_r^v: a node defined as an element of the particular sequence v at time t.

The subscript v in will be omitted unless it is necessary to refer to a specific scenario as a sequence v of nodes. Also, for simplicity, we shall sometimes denote a node just θ.

Endogenous variables

V_{its}^{-}: the value of asset class i before portfolio revision at time t.
V_{its}^{+}: the value of asset class i after portfolio revision at time t.
V_{ts}^{-}: the value of all asset classes before portfolio revision at time t.
V_{ts}^{+}: the value of all asset classes after portfolio revision at time t.
V_{ts}^{R}: the risk adjusted value of the asset classes at time t.
V_{ts}^{L}: the liquidity weighted value of the asset classes at time t.
X_{ts}: transaction costs from trading the asset classes at time t.
E_{ts}: transaction costs on the derivative asset at time t.
K_{ts}: amount of money invested in derivative assets at time t, maturing at time $t+1$.
D_{ts}: income from financial reinsurance (derivatives) at time t.
W_{ts}: interest costs on outstanding debt at time t.
J_{ts}: taxes paid at time t.
A_{ts}: accumulated loans at time t.
C_{ts}: net cash flows earned between time $t-1$ and t.
F_{ts}: earnings in period t, i.e. between time $t-1$ and t.
Y_{ts}: profit to equity earned in period t, i.e. between time $t-1$ and t.
B_{ts}: the value of the company's equity at time t.
H_{ts}: technical reserves at time t.
Z_{ts}: shortfalls of type h at time t.
G_{ts}: income from regular reinsurance at time t.

R_{ts}: total monetary price gain on the portfolio between time $t-1$ and t.
C_s^P: perpetual post horizon expected annual cash flows.
R_s^P: perpetual post horizon expected annual price gain.
Q: the casualty insurer's objective function.

Decision variables

x_{its}: the amount of money traded in asset class i at time t.
n_{ts}: the number of derivatives bought at time t.
l_{ts}: loans obtained at time t.
b_{ts}: the fraction of incurred claims at time t that the company decided at time $t-1$ to cover by reinsurance.
e_{ts}: allocations to technical reserves at time t.

Scenario independent parameters

α_{it}: the fraction of the amount of money traded in each asset class i that the company has to pay as transactions costs at time t.
ι_t: the fraction of the amount of money invested in the derivative asset that the company has to pay as transactions costs at time t.
Ξ: premium income from written policies at time t.
λ: a proportional factor used to calculate the cost of reinsurance as a fraction of premium income.
γ_{it}: upper bound on the asset mix at time t.
δ_{it}: lower bound on the asset mix at time t.
ν_{it}^U: upper bound on the permissible trading volumes at time t.
ν_{it}^L: lower bound on the permissible trading volumes at time t.
ϑ_i^L: a liquidity weight assigned to asset class i.
ϑ_i^R: a risk weight assigned to asset class i.
T_t: the appropriate tax rate at time t.
ψ: a constant associated with the solvency margin.
ϕ: a constant associated with the solvency adequacy requirement.
Ψ^R: a constant associated with the capital adequacy requirement.
Ψ^L: a constant associated with the liquidity constraint.
Λ: the functional symbol for the shortfall function.
Λ_h: a cost coefficient associated with shortfalls of type h.

Scenario dependent parameters

ρ_{ts}: a risk adjusted, path dependent discount factor.
β_{ts}: a risk free, path dependent discount factor.
$P^i_{tk}(s)$: the price at time t in scenario s of a European put option on security i with strike price S_i maturing at time $t+1$ (further modeled in Section 4.1).
μ_{its}: scenario generated ex-dividend price gain on asset class i between time $t-1$ and t.
ϵ_{its}: dividend payment at time t on each asset class i, including coupon payment on bonds and rental income on real estate.
$\varphi^i_{tj}(s)$: payoff from derivatives on security i at time t in outcome j associated with scenario s (further modeled in Section 4.1).
ϖ_s: the companys weighted average cost of capital at the horizon.
π_s: the subjective (empirical) probability that scenario s occurs.
χ_{1ts}: claims of type 1 due to policyholders at time t.
χ_{2ts}: claims of type 2 (originating from underwriting hazardous risks) due to policyholders at time t.
$r_{tj}(s)$: the risk free interest rate between time t and $t+1$, associated with scenario s.
ζ_{ts}: a ratio associated with the solvency margin.

The ALM model can now be formulated as follows:

Maximize

$$\begin{aligned} Q \;=\; & \sum_{s\in S} \pi_s \Big\{ \sum_{t\in\Gamma} \rho_{ts}\left(C_{ts}+R_{ts}\right) + \rho_{\tau s}\left(C^P_s + R^P_s\right)/\varpi_s - \rho_{\tau s} A_{\tau s} \\ & - \sum_{t\in\Gamma} \Lambda\left(\beta_{ts} Z_{ts}\right) \Big\} \end{aligned} \tag{1}$$

Subject to

$$V^+_{its} = V^-_{its} + x_{its}, \quad \forall i \in I, t\in T, s\in S \tag{2}$$

$$V^-_{its} = V^+_{it-1s}\left(1+\mu_{its}\right), \quad \forall i\in I, t\in\, ,\, , s\in S \tag{3}$$

$$X_{ts} = \sum_{i\in I} \left|\alpha_{it} x_{its}\right|, \quad \forall t\in T, s\in S \tag{4}$$

$$E_{ts} = \iota_t K_{ts}, \quad \forall t\in T, s\in S \tag{5}$$

$$
\begin{aligned}
K_{ts} &= n_{ts}P^i_{tk}(s), \quad \forall t \in T, s \in S && (6)\\
D_{ts} &= n_{t-1s}\varphi^i_{tj}(s), \quad \forall t \in T, s \in S && (7)\\
A_{ts} &= \sum_{k=0}^{t} l_{ks}, \quad \forall t \in , , s \in S && (8)\\
C_t &= l_t - K_t, \quad t = 0 &&\\
C_{ts} &= (1-\lambda b_{t-1s})\,\Xi_t + l_{ts} + \sum_{i\in I} V^+_{it-1s}\epsilon_{its} + D_{ts} + G_{ts} && (9)\\
&\quad -K_{ts} - \chi_{1ts} - \chi_{2ts} - W_{ts} - X_{ts} - J_{ts}, &&\\
&\quad \forall t \in (1,\dots,\tau), s \in S &&\\
C_{ts} &= \sum_{i\in I} x_{its}, \quad \forall t \in T, s \in S && (10)\\
G_{ts} &= b_{t-1s}\chi_{1ts}, \quad \forall t \in (1,\dots,\tau), s \in S && (11)\\
W_{ts} &= r_{tj}(s)l_{ts}, \quad \forall t \in (1,\dots,\tau), s \in S && (12)\\
F_{ts} &= (1-\lambda b_{t-1s})\,\Xi_t + \sum_{i\in I} V^+_{it-1s}\,(\mu_{its} + \epsilon_{its}) + D_{ts} + G_{ts} &&\\
&\quad -K_{ts} - \chi_{1ts} - \chi_{2ts} - W_{ts} - X_{ts} - e_{ts}, && (13)\\
&\quad \forall t \in (1,\dots,\tau), s \in S &&\\
J_{ts} &= \max\left[0, T_tF_{ts}\right], \quad \forall t \in (1,\dots,\tau), s \in S && (14)\\
Y_{ts} &= F_{ts} - J_{ts}, \quad \forall t \in (1,\dots,\tau), s \in S && (15)\\
B_{ts} &= B_{t-1s} + Y_{ts}, \quad \forall t \in (1,\dots,\tau), s \in S && (16)\\
H_{ts} &= H_{t-1s} + e_{ts}, \quad \forall t \in (1,\dots,\tau), s \in S && (17)\\
\delta_{it}V^+_{ts} &\leq V^+_{its} \leq \gamma_{it}V^+_{ts}, \quad \forall i \in I, t \in T, s \in S && (18)\\
\nu^L_{it} &\leq x_{its} \leq \nu^U_{it}, \quad \forall i \in I, t \in (1,\dots,\tau), s \in S && (19)\\
x_{its_k} &= x_{its_v} \quad \text{if} \quad \theta^k_r = \theta^v_r \quad \forall r : 0 \leq r \leq t, i \in I, t \in T, s_k, s_v \in S && (20)\\
Z_{1ts} &= \max\left[0, \psi\Xi_t\zeta_{ts} - (B_{ts} + D_{ts} + \lambda b_{t-1s}\Xi_t + \phi H_{ts})\right], &&\\
&\quad \forall t \in (1,\dots,\tau), s \in S && (21)\\
V^R_{ts} &= \sum_{i\in I} \vartheta^R_i V^-_{its}, \quad \forall t \in (1,\dots,\tau), s \in S && (22)\\
Z_{2ts} &= \max\left[0, \Psi^R V^R_{ts} - (B_{ts} + D_{ts} + \lambda b_{t-1s}\Xi_t)\right], &&\\
&\quad \forall t \in (1,\dots,\tau), s \in S && (23)\\
Z_{3ts} &= \max\left[0, -(C_{ts} + R_{ts} - l_{ts})\right], \quad \forall t \in (1,\dots,\tau), s \in S && (24)\\
Z_{4ts} &= \max\left[0, \chi_{1ts} + \chi_{2ts} - \Psi^L V^L_{ts} - D_{ts}\right], &&\\
&\quad \forall t \in (1,\dots,\tau), s \in S && (25)
\end{aligned}
$$

$$V_{ts}^{L} = \sum_{i \in I} \vartheta_i^L V_{its}^-, \quad \forall t \in (1, \ldots, \tau), s \in S \tag{26}$$

The three first bracketed terms of the objective function (1) gives the net present value of the company in each scenario s, including its scenario dependent continuing value. Λ is a convex cost function through which shortfalls are penalized. Observe that the (regulated) companys risk preferences are reflected in this cost function, so that shortfalls are discounted to the present using risk free, time dependent interest rates. The risk adjusted discount factor ρ_{ts} embodies the business risk that the market assigns to the operations of the company. In a capital asset pricing model framework[1], this factor would account for the fact that unsystematic risk has been diversified away from the investors portfolios. We define ρ_{ts} as the companys cost of capital, and interpret this to be the average return on equity that a mutually owned company would have to offer its pool of risk averse insurance takers.

Constraints (2) define portfolio revision. The x_{its} variable is positive or negative according to whether or not asset class i is bought or sold.

Constraints (3) define how asset values evolve between time $t-1$ and t.

The nonlinear constraints (4) define transaction costs that accrue each time an asset class is traded. When implemented, these constraints have been linearized according to a standard procedure.

Constraints (5) define transaction costs that accrue when the derivative asset is bought. It is assumed that the derivatives are held to maturity, i.e. never sold.

Constraints (6) and (7) are related to financial reinsurance. Constraints (6) calculate the value of derivative holdings at time t. Equivalently, this is the cost at time t of financial reinsurance against insurance risks at time $t+1$. Constraints (7) calculate the payoff at time t from derivatives obtained at time $t-1$. The derivatives are alive for one period only, and, as we emphasized above, are always held to maturity. Observe that option prices $P_{tk}^i(s)$ and option payoffs $\varphi_{tj}^i(s)$ are treated as data in the ALM model. These parameters are obtained from the derivative pricing model presented

[1] A textbook presentation of the Capital Asset Pricing Model is offered by Huang & Litzenberger (1988).

in Section 4.

Loans accumulate according to constraints (8). It is assumed for computational simplicity that loans are to be accumulated and repaid at the horizon.

At each time t, the company calculates the scenario dependent net cash flow, which is to be allocated among the asset classes, according to constraints (9) (At time $t = 0$, there is of course only one set of variables that is shared by all scenarios.)

The assumption that all net cash flows are reinvested is reflected in the budget constraints (10).

The amount of money received from (regular) reinsurance at time t in scenario s is calculated as a fraction of the associated claims payment, as defined in constraints (11). We assume that there is a proportionate relationship between (state dependent) received reinsurance at time t and premium income ceded at time $t - 1$.

Interest costs on loans are made consecutively and these are calculated according to constraints (12).

Constraints (13) show how earnings are calculated at each time t. We need this item in order to calculate taxes.

Having calculated earnings, we may calculate taxes according to constraints (14). In conformity with the tax equations, these constraints were linearized by a standard procedure. Observe that the two sets of equations (13) and (14) imply that, in this model, no distinction is made between realized and non-realized return on the asset classes.

Now we may calculate profits available for equity at each time t according to constraints (15).

The value of the company's equity between time $t - 1$ and t is updated according to constraints (16).

Constraints (17) defines allocations to technical reserves at time t.

Constraints (18) impose lower and upper bounds on the asset mix according to legislative regulations. V_{ts}^{+} denotes the portfolio value (excluding derivative holdings) after revision at time t, i.e. $V_{ts}^{+} = \sum_{i \in I} V_{its}^{+}$.

Typically insurance companies are large traders in financial markets, and must consider the impact of their trading policies on asset prices. Consequently constraints (19) impose trading limits on the portfolio revision.

Constraints (20) impose non-anticipativity constraints on the scenario dependent decisions at each node in the scenario tree. We note that s_k and s_v are arbitrary scenarios. A similar set of constraints apply to each type of decision variables in the model. These constraints are necessary in order for the decisions, associated with a particular scenario s_k, not to anticipate future realizations of uncertain parameters. Since each node is shared by many scenarios, each of which is associated with a set of decision variables, typically there are many decision variables representing the same decision at each node. Constraints (20) imply that all but one of these decision variables are redundant. See Rockafellar & Wets (1991) for a discussion of the modeling and solution of general stochastic programs with non-anticipativity constraints. When we implemented the model we avoided the redundancy problem associated with scenario dependent decisions, by means of defining the model entities in terms of nodes rather than scenarios.

Constraints (21)-(26) are related to the definition of shortfall variables, which are indexed over the set H. We have linearized all nonlinear constraints in the actual implementation of the model. If correctly defined the shortfall variables reflect the legal restrictions with witch a casualty insurer attempts to comply. The legislation affects the company's economic activity both through the balance sheet and the profit and loss account. A casualty insurer must adapt to capital adequacy requirements, solvency requirements and restrictions on borrowing, the latter implying, among other things, a concern for the liquidity of the asset classes. Also technical reserves of different categories must not fall below specific minimum requirements. Finally, the portfolio composition at any decision point cannot be determined at the company's discretion, there are upper limits on the amount of money which can be held in the different asset classes as a fraction of technical reserves.

If the company at any time t should fail to comply with any of the above regulations, it is penalized in the objective function through the shortfall variables Z_{hts}, which then attain positive values.

Constraints (21) defines the shortfall variables associated with the solvency requirements. The basis for calculating the solvency margin is either current period premium income or average loss payments incurred over the last three periods, depending on which calculation yields the highest amount of money. We have based our definition of the solvency margin on current period premium income. In (21) the solvency margin is calculated as a percentage ψ of current period premium income Ξ_t multiplied by a ratio denoted ζ_{ts}. The parameter ζ_{ts} is the ratio between current period loss payments, net of received reinsurance, and last period gross loss payments in scenario s. Thus a more risky portfolio of policy owners "on average" increases the solvency requirement through higher loss payments, whereas investments in reinsurance decreases the margin. Next, the solvency capital is compared to the margin. The solvency capital is defined by the last term on the right hand side (21). It consists of equity, income from financial reinsurance (derivatives) at time t, investment in regular reinsurance at time $t-1$ and a fraction ϕ of a model aggregate representing technical reserves. Whenever the solvency capital falls short of the required margin, the shortfall variable Z_{1ts} attains positive value. If the solvency margin requirement is satisfied Z_{1ts} is set to zero because of the costs associated with a positive shortfalls.

Constraints (22) define a risk adjusted value of the asset classes at time t. This value is needed in order to define the capital adequacy requirement. Each asset class is assigned a weight according to its alleged financial risk.

In constraints (23) the capital adequacy requirement is calculated as a percentage Ψ^R of the risk adjusted value of the portfolio. The second term on the right hand side of (23) is the capital adequacy requirement. The items to be measured against this requirement is found both on and off the balance sheet. In this model, in order to avoid shortfall costs at any time t, the sum of equity, income from financial reinsurance and investment in regular reinsurance, the latter measured at time $t-1$, must be greater than or equal to the capital adequacy requirement.

Constraints (24) are funding requirements. A casualty insurer is concerned that its earnings at any time t is sufficient to cover the incurred loss payments to policyholders, in all but a few extreme states of nature. This preference implicitly reflects the fact that the insurance company is limited to restrictive borrowing. Only exceptionally will it be permitted to obtain loans. Thus whenever it does choose to obtain loans to cover part of its liabilities, it incurs shortfall costs. If the numerical value of the bracketed expression in (24) is negative, the sum of premium income less administrative expenses and return on the portfolio of asset classes was insufficient to cover claims directed against the company, at that particular point in time. Consequently the associated shortfall variable attains positive value.

The fact that a casualty insurer is not generally permitted to obtain loans to facilitate its operations, has an impact on the extent to which it cares about the liquidity of its assets. This connection is reflected in constraints (25). Generally a casualty insurer is required to cover its liabilities immediately as they occur and might have to sell off assets, in which case the liquidity of the assets is of vital importance. Consequently there is a link between its earnings and the portfolio composition. The connection is as follows: In states of nature where earnings are sufficient to cover loss payments, the company will rebalance in accordance with its (possibly revised) market expectations. In such states it may keep a relatively large portion of its funds in non liquid assets, such as stocks, because it will only trade at its own discretion. However, this is not necessarily true in states of nature where its earnings are insufficient to cover its liabilities due to huge loss payments. Then it might have to sell off assets fast to cover the loss payments. In such states the company would be better off with a more liquid portfolio, i.e. a relatively large portion of its funds should be allocated to cash. In order to force the company to maintain a relatively liquid portfolio, it incurs positive shortfalls Z_{4ts} whenever the loss payments exceed a percentage of the liquidity weighted value of the portfolio.

Observe that the above implementation of the legal restrictions reflects the fact that derivative hedging actually reduces the company's risk exposure, and therefore contributes to compliance with the legal statutes.

Constraints (26) define the liquidity weighted value of the portfolio at time t. Income from the derivatives has a liquidity weight equal to one because it is cash.

We mentioned above that, when developing the portfolio selection model, we implicitly regarded option prices as input data to the model, and the value of the option portfolio, K_{ts}, is equal to the number of option contracts times their price. We need to have a model for option prices because it is necessary to know the prices on options in future states, not only for the first stage—for which it is possible to use market prices on options. In the next section, we set up a model which selects a discrete probability measure with which the value of risky cash flows, including option prices, easily can be computed. This probability measure satisfies the important economic no arbitrage condition.

4 Pricing derivatives in stochastic programming models

In their famous 1973 paper, Fisher Black and Myron Scholes reported the discovery of a closed form formula for the price of a European call option on a non-dividend paying stock. Black and Scholes set their analysis in continuous space and time, under a relatively strict set of assumptions. Their most important discovery was the fact that their differential equation, which represents the change in the price of the derivative asset with respect to change in the price of the underlying asset and time, does not contain risk preferences. This led to the principle of risk neutral valuation, which is one of the most important concepts of financial theory. In essence this principle states that since the differential equation governing the evolution of the option price is not affected by risk preferences, neither is its solution, i.e. the option price, and consequently when pricing derivatives we may assume that investors are risk neutral. The prices we obtain, however, are valid in all worlds.

Over the last three decades a substantial body of literature, dedicated to derivative pricing, has shown how the Black-Scholes analysis can be extended and how their assumptions can be relaxed. Numerous different pricing methods and contexts have been suggested, but the principle of risk neutral valuation continues to apply. In the next section we shall make explicit use of this principle when pricing a put index option. We shall maintain the most important of the assumptions underlying the Black-Scholes analysis: that there are no arbitrage opportunities. An arbitrage opportunity is a sequence of portfolio decisions with an initial cost of zero, maintaining non-negative value in all future states, and which has a strictly positive possibility of ending up in a state with a positive value. Loosely speaking the principle of no arbitrage means that in order to make a profit in excess of the risk free rate of return, you need to incur risk. This condition reflects the basic economic assumption about financial markets that well informed arbitrageurs would exploit arbitrage opportunities immediately if such existed. Therefore in practice arbitrage opportunities can only persist for a very short period of time, and in theoretical models arbitrage should not be present at all.

It is well known that the no-arbitrage condition is equivalent to saying that ex-dividend price processes and dividend processes satisfy the martingale property. A security price process and its accumulated dividends are a martingale if, at any time t, the conditional expectation of their discounted sum at any future date s equals the sum of their value at time $t < s$. The expectation should be taken under a specific probability measure and if there exists a risk free investment opportunity, the discount rate should be equal to the risk free rate of return. The main task of the next section is to construct this probability measure in terms of a discrete probability distribution.

4.1 Constructing arbitrage free derivative prices

It is a necessary condition for the martingale probability measure to be unique in a discrete time and space economy that, at each time t, the number of linearly independent long lived security (prices) be equal to the maximum number of conditional realizations of those prices at time $t + 1$, i.e. the maximum number of branches leaving a node in the scenario three at time t. This condition is not satisfied in our model economy, where the number of possible outcomes in each period, Ω_t, exceeds the number of securities. However, when there are no arbitrage opportunities, there exists a *set* of probability measures under which all securities are martingales. In this section we formulate a set of constraints that must be satisfied by security prices in order to avoid arbitrage. The martingale probability measure we thus obtain is not unique.

A crucial implication of the above result is the fact that derivative prices satisfying the no-arbitrage condition are not unique. As we use option contracts to hedge a portfolio of financial assets against insurance risks, this raises two questions: (i) Is the optimal portfolio much affected by the choice of a specific martingale probability measure, implying a specific vector of option prices? (ii) If the answer to (i) is yes, how should we select an appropriate martingale measure? Below we present a method which produces discrete martingale probability measures. The method is similar to that of Jackwerth & Rubinstein (1996), although they are able to constrain the problem more than we are, because they use observed bid and ask prices for derivatives in the market.

As noted above, according to the principle of risk neutral valuation, when valuing derivatives we may assume that investors are risk neutral. In such a world the expected return on all securities must be the risk free rate of return. The set of security prices (implicitly) generated as data to our portfolio selection model, obviously reflects risk preferences, and so does the associated (subjective or empirical) probability measure. If we can find a way to move probability mass around, so that the generated set of security prices yields a *net* expected rate of return equal to the risk free interest rate, then we have a candidate for our martingale measure in terms of a risk neutral probability distribution.

Provided the following set of equations are satisfied by security prices when investors are risk neutral, there are no arbitrage opportunities:

$$\begin{aligned} \sum_{j=1}^{\Omega_t} \Pi_{t-1j}(s) M^i_{tj}(s) &= (1 + r_{t-1j}(s))\, M^i_{t-1k}(s), \quad \forall i \in I, t \in \,,\, , k \in \Omega_{t-1} \\ \sum_{j=1}^{\Omega_t} \Pi_{t-1j}(s) &= 1, \quad \forall t \in T \end{aligned} \tag{27}$$

We note that each outcome, with the obvious exception of last period outcomes, is associated with several scenarios s. $M^i_{tj}(s)$ is the price of security i at time t in outcome j. (Outcome j occurred between $t-1$ and t.) There are Ω_t possible outcomes at time t, one of which will occur. The Π_{tj}'s are the risk neutral probabilities at time t of outcome j occurring at time $t+1$, which constitutes the martingale distribution we are looking for. The risk free rate of return is $r_{tj}(s)$ between time t and $t+1$. As mentioned above the security prices are data to our portfolio selection model, and these are regarded as parameters in (27). We note that all states must occur with a strictly positive probability. This technical condition is needed in order for the martingale measure to guarantee absence of arbitrage opportunities. Once a martingale measure has been selected, the put prices can be computed according to the following formula:

$$P_{tk}^{i}(s) = \sum_{j=1}^{\Omega_t} \frac{\Pi_{tj}(s) \max\left[0, S_i - M_{t+1j}^{i}(s)\right]}{1 + r_{tj}(s)}, \quad \forall i \in I, t \in T, k \in \Omega_{t-1} \quad (28)$$

where $P_{tk}^{i}(s)$ is the price at time t of a European *put* option on security i with strike price S_i maturing at time $t+1$. Observe that the price is contingent upon outcome k occurring at time t. When the number of states j, at any time t, exceeds the number of securities, there exists a continuum of solutions to (27) in terms of the probabilities $\Pi_{tj}(s)$.

In order to find a solution we shall optimize an objective with respect to the $\Pi_{tj}(s)$'s, subject to the set of equations (27). Since in our case the arbitrage principle fails to provide unique derivative prices, we utilize *equilibrium* principles. We maintain our requirement that prices must be arbitrage-free, but in addition we require that prices are as close as possible to reasonable equilibrium prices. Our objective is thus to minimize a distance measure, i.e. the distance between equilibrium and absence of arbitrage.

In equilibrium valuation models it is often assumed that the market participants can be aggregated into a representative agent[2]. The prices that can be inferred from the first-order condition of the portfolio-consumption choice problem of this agent are the desired equilibrium security prices. An example of such a condition is that the marginal rate of substitution (MRS) between present and future state contingent consumption, is equal to the price of an *elementary claim* that pays off \$1 if a particular state occurs and nothing otherwise, i.e. the price of a *state-contingent claim*, O_{tj}:

$$\pi_{tj}(s)\mathrm{MRS}_{tj}(s) = O_{tj}(s), \quad \forall t \in \ , \ , j \in \Omega_t \quad (29)$$

The $\pi_{tj}(s)$'s are the empirical time t probabilities that outcome j will occur at time $t+1$. The state-contingent claims have the property:

$$M_{t-1k}^{i}(s) = \sum_{j=1}^{\Omega_t} O_{t-1j}(s) M_{tj}^{i}(s), \quad \forall i \in I, t \in \ , \ , k \in \Omega_{t-1} \quad (30)$$

[2]An outline of this theory is found in Huang & Litzenberger (1988).

Since by (27)

$$M_{t-1k}^{i}(s) = \frac{1}{1 + r_{t-1j}(s)} \sum_{j=1}^{\Omega_t} \Pi_{t-1j}(s) M_{tj}^{i}(s), \quad \forall i \in I, t \in , , k \in \Omega_{t-1} \tag{31}$$

we see that prices of state-contingent claims correspond to discounted martingale probabilities of (27):

$$\pi_{tj}^{\text{eqm}}(s) = O_{tj}(s)\left(1 + r_{tj}(s)\right) = \pi_{tj}(s)\text{MRS}_{tj}(s)\left(1 + r_{tj}(s)\right), \quad \forall t \in , , k \in \Omega_{t-1} \tag{32}$$

This means that we can find an equilibrium martingale probability measure $\pi_{tj}^{\text{eqm}}(s)$ using the marginal rate of substitution of a representative agent, the risk free interest rate and the empirical probability. These probabilities should reflect a reasonable equilibrium, i.e. one in which the prices of the state-contingent claims, or state prices, are such that the value of receiving \$1 in an abundant state of nature is less than that of receiving \$1 in a poor state of nature.

As a proxy for future stage contingent consumption of the representative agent, i.e. the extent to which the state of the world is "rich" or "poor", we use the index level $M_{tj}^{1}(s)$. The utility function of that agent, which determines the MRS $\left(M_{tj}^{1}(s)\right)$, is a matter of choice. A good choice is one that represents the aggregate market well, Jackwerth (2000) suggests "a power utility function of moderate risk aversion". Thus we obtain our risk neutral probability measure as the solution to the following set of conditional optimization problems: At each node θ_t^v in the scenario tree solve;

$$\min_{\Pi_{tj}(s)} \sum_{j=1}^{\Omega_t} \left[\Pi_{tj}(s) - \pi_{tj}^{\text{eqm}}(s)\right]^2, \quad \forall t \in T, \theta_t^v \in \Theta \tag{33}$$
$$\text{s.t. } (27) \text{ and } (32)$$

i.e. the summation is taken over all arcs leaving node θ_t^v. When the put options associated with the martingale obtained from (33) are priced according to (28), they do possess the desired return properties; since a put option pays off in poor states of nature, their expected return should be low or even negative. We demonstrate the puts significance to portfolio hedging through a numerical example in Section 5.

Jackwerth (2000) suggests that out of the money puts empirically tend to be slightly overpriced. Computing the maximum put price vector satisfying the no-arbitrage condition, yields conservative results in terms of the portfolio share of puts, i.e. the level of reinsurance, in the optimal portfolios selected by the ALM-model of Section 3. The objective (34) applied to (27) at each node implements this approach.

$$\max_{\Pi_{tj}(s)} \sum_{j=1}^{\Omega_t} \Pi_{tj}(s)\varphi^i_{tj}(s), \quad \forall t \in T, \theta^k_t \in \Theta \tag{34}$$
$$\text{where} \quad \varphi^i_{tj}(s) = \max\left[0, S_i - M^i_{tj}(s)\right]$$

We have further tried the objective (35) on the set of constraints (27).

$$\min_{\Pi_{tj}(s)} \sum_{j=1}^{\Omega_t} \left[\Pi_{tj}(s) - \pi_{tj}(s)\right]^2, \quad \forall t \in T, \theta^k_t \in \Theta \tag{35}$$

Applying (35) to (27) yields a put price vector in the interior of the set of allowable put price vectors. However, neither (34) nor (35) produced reasonable put returns, as did the objective in (33).

In the next section we construct a numerical example, illustrating the importance of financial reinsurance to the ALM problem of a casual insurer. The derivatives we use are put index options, which have been priced according to (28) after solving the system (33).

5 Numerical example

We have implemented and solved a deterministic equivalent of the stochastic program associated with the model presented in Section 3. The ALM model is developed in a highly aggregated setting, and can be used to decide about distribution of funds between classes of securities. It has three asset classes, 4 stages, which is equivalent to three time periods, and ten outcomes in each time period. This amounts to a total of 1000 scenarios, i.e. paths through the scenario tree depicted in Figure 1. By means of a numerical example we shall illustrate the concept of financial reinsurance. As pointed out in Section 2 and 3 the risk preferences embedded in the objective function (1) reflects the companys effort to comply with the legal regulations. These in turn are enforced to protect policyholders interests against too risky investment strategies, which the authorities think might result from unregulated market behavior. Therefore we calculate the utility of wealth that an alleged representative risk averse policy holder, endowed with a linear plus negative exponential utility function, gains from the insurance companys optimal portfolio decisions.

We study a casualty insurer which may allocate his funds to cash, stocks and put options on a stock index whose price process is highly negatively correlated with the claims payments associated with a specific insurance risk in his portfolio of insurance contracts. The initial balance sheet for this company is shown below in Table 1.

Table 1: Initial balance sheet.

Assets		Liabilities	
Cash	160	Loans	14
Stocks	40	Technical reserves	126
Put options	0	Equity	60
Total assets	200	Total liabilities	200

When deciding which asset classes should be available for our insurance company, we only distinguish between highly liquid, non-risky assets and non-liquid, risky assets. The initial distribution of funds across assets is supposed to reflect that casualty insurers empirically are highly liquid. The next section briefly explains how we generate data for our portfolio selection model.

5.1 Scenario generation

Input data to the portfolio selection model outlined in Section 3.1 are random returns on the asset classes and stochastic loss payments due to policyholders. These are the stochastic parameters of the model. In order to provide these data, we generate a limited number of discrete outcomes with statistical properties that match those of a subjective multivariate continuous probability distribution, specified in terms of the first M central moments and correlations. For this purpose, we use a scenario generation model, employing a method based on nonlinear optimization. This model was developed by Høyland & Wallace (1999*b*) to provide input for an ALM-model of a life insurance company.

The implemented scenario generation routine works as follows: The user specifies his or her subjective market expectations, in terms of percentiles of marginal distributions for asset returns and claims, and also a correlation matrix to capture the interdependencies of the random parameters. Ideally, we would like to construct a scenario tree retaining as many as possible of the statistical properties of the continuous marginal distributions, in addition to the correlation matrix. From a computational point of view this would however result in a prohibitively large number of discrete outcomes. The crucial question is: Which are the economically significant statistical properties of the original marginal distributions, in view of the problem at hand? Put differently this question could be phrased: Which statistical properties does (or should) the decision-maker care about? When this question has been decided, we generate a scenario tree, i.e. a set of discrete probability distributions, which is consistent with the relevant statistical properties. In our case, we believe that the relevant properties are expectation, variance, skewness, kurtosis and correlations.

The scenario generation model can take account of inter-temporal dependencies between the variables, and we have modeled mean reversion for interest rates, and volatility clumping for all asset classes. The method is explained by Høyland & Wallace (1999*b*). The market specifications, underlying our numerical example, are reported in Table 2. "Claims 1" accounts for the main bulk of insurance risks that the company has underwritten. "Claims 2" is associated with hazardous risks, which constitute a small fraction of total liabilities.

We have also specified a state independent correlation matrix, which is reported in Table 3.

5.2 Experiment and results

We now describe the method we use to examine the rationale for financial reinsurance. A basic assumption is that, under the prevailing statutory regulations, the casualty insurer is risk averse. We assume that he has to comply with the following statutory regulations: (i) capital adequacy requirements, (ii) solvency requirements, (iii) credit rationing, i.e. borrow-

Table 2: Subjective market specifications. A * means that the specification is state dependent.

Asset classes incl claims	Distribution Property	(End of) Period 1	(End of) Period 2	(End of) Period 3
Cash	Exp. value spot rate	0.04	*	*
	Standard deviation	0.01	*	*
	Skewness	0.0	0.0	0.0
	Kurtosis	3.0	3.0	3.0
Stocks	Expected value return	0.08	*	*
	Standard deviation	0.15	*	*
	Skewness	-1.0	-1.0	-1.0
	Kurtosis	4.0	4.0	4.0
Index	E. value gross return	0.08	*	*
	Standard deviation	0.18	*	*
	Skewness	-1.0	-1.0	-1.0
	Kurtosis	4.0	4.0	4.0
Claims 1	Expected value	-23.0	*	*
	Standard deviation	5.0	*	*
	Skewness	-2.0	-2.0	-2.0
	Kurtosis	8.0	8.0	8.0
Claims 2	Expected value	-10.0	*	*
(hazardous	Standard deviation	10.0	*	*
risks)	Skewness	-2.0	-2.0	-2.0
	Kurtosis	8.0	8.0	8.0

Table 3: Specified correlations in all periods.

	Cash	Stocks	Index	Claims 1	Claims 2
Cash	1.0	-0.6	-0.4	0.0	0.0
Stocks	-0.6	1.0	0.3	0.0	0.0
Index	-0.4	0.3	1.0	0.0	0.7
Claims 1	0.0	0.0	0.0	1.0	0.0
Claims 2	0.0	0.0	0.7	0.0	1.0

ing is in practice prohibited. As explained when presenting the model in Section 3, the legal requirements are not implemented as hard constraints. Instead each legal constraint is associated with a costly recourse activity, which may be chosen such as to compensate its violation, if any. This is equivalent to assuming that risk, as perceived by the insurance company, is failing to comply with the legal requirements. The cost function assigning penalty costs to positive shortfalls, is quadratic:

$$\Lambda\left(\beta_{ts} Z_{ts}\right) = \sum_{h \in H} \Lambda_h \beta_{ts} Z_{hts}^2. \tag{36}$$

Here β_{ts} is a risk free discount factor, for period t in scenario s, Z_{hts} denotes actual shortfalls of type h and Λ_h is a constant. This function penalizes large shortfalls relatively more than small shortfalls, reflecting the assumption of convex recourse costs. In these experiments the constant Λ_h is assigned the numerical value 0.5 for all h.

Associated with insurance risks are claims payment and premium income. For most risk categories in the casualty insurance business, expectations of claims payments plus operating costs are equal to or even slightly higher than expectations of premium income. This implies that the company must earn its profits from investing policyholders funds in the financial markets. We shall however assume that expected premium income associated with hazardous risks are higher than the expectation of the corresponding highly volatile claims ("Claims 2") plus associated operating costs. This seems like a plausible condition to be satisfied by hazardous risks, if these are to be attractive to risk averse casualty insurers.

In this numerical example, there exits a stock index having random payoffs that are positively correlated with the claims payments denoted "Claims 2" in Table 3. The correlation between the two variables is assumed to be 0.7. Recall from Table 3 that the distributions of claims have *negative* expectation. Consider a realization of claims with a negative deviation from its expectation. If the expected value is -10 the observation could be -15, i.e. a higher loss payment. The assumption that the index price process is positively correlated with the claims, implies that at the same time we expect to observe a realization of the index *below* its expected value. In practice this means that when claims payments peak, we expect to realize payoffs on the puts to help finance the loss payments.

The situation described above could fit many real world phenomena. For instance imagine a casualty insurer who contracts to insure individuals and firms against earthquakes in Japan. If a serious earthquake occurs, he will have to pay huge loss payments. In the same event, many Japanese companies will be seriously hurt. Consequently a long position in put options on a suitable Japanese stock index, will (imperfectly) insure the company against claims generated by this specific insurance risk.

When we construct our numerical example, we have normalized all initial security prices to 1. At each time t, put index options of varying exercise prices expiring at time $t+1$ are available. Intuitively we would expect *out of the money puts*, contrary to in the money- or at the money puts, to be best suited to provide insurance against worst case scenarios for the following two reasons: (i) They are the cheaper ones, being the most distant from the payoff region. (ii) Their strike price has been selected such as to yield positive payoffs when the index reaches some threshold, below which the company needs insurance. Therefore our first trial is an out of the money put option with strike price 0.9. Of course what is considered a worst case scenario must be decided by each casualty insurer. We have also analyzed an at the money put option with strike price 1.0. Note that the puts are alive for one period only. This simplifies the implementation of the no arbitrage condition. We note that whether or not puts that are bought beyond time $t = 0$ are in the money, at the money, or out of the money is state dependent in this model.

In Table 4 below we display the optimal time zero portfolios selected by the casualty insurer. The portfolios are characterized as follows:

Portfolio A:	Unregulated portfolio including hazardous risks.
Portfolio B:	Unregulated portfolio excluding hazardous risks.
Portfolio C:	Regulated financially noninsured portfolio including hazardous risks.
Portfolio D(D*):	Regulated financially reinsured portfolio including hazardous risks.
Portfolio E:	Regulated portfolio excluding hazardous risks.

Portfolio A and B are selected in an unregulated market, in which borrowing is possible. In the absence of regulations, some policymakers seem to believe that a mutual casualty insurer would not consider its owners

Table 4: Optimal time zero portfolios and associated objective value for the casualty insurer.

Port-folio	Cash	Stocks	Strice price	Put value	Loans	St.dev of $\Delta V^-_{\tau s}$	Objective value
A	0	210			10	63	330
B	0	210			10	60	281
C	171.8	28.2			0	19.9	86
E	152.3	47.7			0	11	103
D	164.6	33.4	0.9	2.0	0	27	107
D*	172.2	26.8	1.0	1.0	0	17	98

risk preferences, i.e. the policyholders risk preferences, in the sense that in a CAPM framework he does not care about unsystematic risk. When two-fund separation applies, he invests to maximize expected payoffs and allocates all his funds to the security yielding the highest expected (rate of) return. In our example this is stocks. Accordingly, portfolio A and B consists of stocks only. This also holds for all future conditional decisions that are not shown in Table 4. The remaining portfolios are optimal when the insurer is required to comply with the legal statutes, for a given degree of risk aversion ($\Lambda_h = 0.5$). Portfolio A yields higher objective function value to the casualty insurer than does portfolio B, implying that the hazardous risk does provide positive expected cash flows.

The regulated portfolios C, D and D* both includes hazardous risks, and in terms of shortfall risk portfolio D* is the least risky. The difference between porfolio D and D* is the strike price of the put, as shown in Table 4. In this particular example, the put options have lower *expected* payoff than the risk free rate of return. This is because of the way we have selected our martingale probability measure. Under the risk neutral probability measure the expected return on *all* securities must be the risk-free rate of return. To accomplish this, in our scenario tree probability mass is shifted from abundant to poor outcomes. Poor outcomes include outcomes in which the puts pay off. The put prices we obtain from risk neutral valuation are valid in all worlds. However, in the scenario tree the puts pay off less frequently than under the risk neutral distribution. It follows that for given put prices under the empirical distribution, i.e. on our scenario tree, the puts must yield expected return less than the risk-free rate of return.

Portfolios D and D* yield higher objective function value to the risk averse casualty insurer than portfolio C. This is because the former portfolios are less constrained, i.e. derivative hedging is possible and beneficial. We observe that it is optimal to take positions in put options at time $t = 0$. Conditional time $t = 1$ derivative positions in portfolio D are displayed in Table 5. There are 10 conditional decision points at $t = 1$.

Table 5: Conditional time $t = 1$ put holdings. Strike price 0.9.

Outcome	Put holdings
1	0
2	3.59537
3	0
4	1.68444
5	0
6	0
7	31.5668
8	6.74603
9	3.21446
10	0

The more interesting comparison is between portfolios D (D*) and E. Observe from Table 4 that a risk averse casualty insurer can actually increase his utility (objective function value) from insuring hazardous risks, *provided* his earnings can be successfully financially reinsured. In our example puts with exercise price 0.9 provides satisfactory insurance. Conditional on this exercise price, portfolio D clearly outperforms portfolio E, in terms of objective value. The financial markets provide a possibility to hedge the casualty insurers income against claims that are due to specific hazardous risks in his *insurance* portfolio. He may thus be able to insure lucrative but hazardous risks, while incurring acceptable shortfall costs in order to maintain compliance with the legal statutes.

As mentioned in the introduction to this paper, we assume that the policyholders own either the insurance company or the funds being managed by the insurance company. We further assume that there exists a well informed representative policyholder, who is capable of identifying the companys financial position. This policyholder realizes that a highly volatile portfolio of stocks might be a threat to its ability to pay the incurred claims, if a worst case scenario occurs.

In order to examine the policyholders' evaluation of financially reinsured hazardous risks, we would like to focus on a distinguishing property of derivatives, namely their ability to hedge *downside* risk. It is of course the downside potential of insurance fund return that worries risk averse policyholders. This means that we should endow the representative policyholder with a utility function that explicitly penalizes downside variance. Therefore we have represented policyholders non-satiated desire for wealth by the linear plus negative exponential utility function, which has this property:

$$E\left[U\left(\Delta V_{\tau s}^{-}\right)\right] = \sum_{s \in S} \pi_s \left(\Delta V_{\tau s}^{-} - b_1 e^{-b_2 \Delta V_{\tau s}^{-}}\right) \tag{37}$$

$\Delta V_{\tau s}^{-}$ is the change in the value of the portfolio between time $t = 0$ and $t = \tau$ in each scenario s. The parameter b_2 controls aversion against downside variance, and b_1 controls the aversiveness of this risk. We let $b_1 = 1$. Increasing b_2 implies that the policyholders more strongly resent extreme downside outcomes. The linear plus negative exponential utility function is further examined by Bell (1995). Table 6 shows the utility function (37) evaluated at portfolios A through E.

Slightly risk averse policyholders (b_2 =0.01) prefer portfolio A to all the other portfolios (as they should). However, when risk aversion increases beyond 0.01 the unregulated portfolios A and B are the least desired. Consider now the regulated portfolios C, D and D*, which all include hazardous risks. When the risk aversion parameter b_2 in the representative policyholders utility function is greater than 0.01, the financially non-reinsured portfolio C is never preferred to the reinsured portfolios D and D* which includes put options. Observe that puts with strike price 1 provides better insurance against down side variance than otherwise identical puts with strike price 0.9. But the former are more expensive and are only preferred when $b_2 \geq 0.3$.

Figure 2 shows the effect on the cumulative distribution function (CDF) for the change in wealth, $\Delta V_{\tau s}^{-}$, of introducing put options into the portfolio.

Since the put options have negative expected returns, the expected (positive) change in wealth is smaller after including puts. However, a careful examination of the lower left tails of the distributions reveals why financial reinsurance is preferred. The downside risk has become smaller, i.e. in this

Table 6: Evaluating the policyholder's expected linear plus exponential utility of wealth, for varying levels of risk aversion.

Risk aversion parameter b_2	Portfolio A	Portfolio B	Portfolio C
0.01	43.69	34.03	19.53
0.1	$-1.505 \cdot 10^{13}$	$-4.889 \cdot 10^{11}$	-47715
0.2	$-1.614 \cdot 10^{19}$	$-1.574 \cdot 10^{17}$	$-1.236 \cdot 10^{8}$
0.3	$-2.453 \cdot 10^{31}$	$-1.784 \cdot 10^{28}$	$-1.600 \cdot 10^{15}$
0.5	$-1.046 \cdot 10^{56}$	$-2.584 \cdot 10^{50}$	$-5.382 \cdot 10^{29}$
0.8	$-1.897 \cdot 10^{93}$	$-5.223 \cdot 10^{83}$	$-4.429 \cdot 10^{51}$
0.9	$-5.25 \cdot 10^{105}$	$-6.728 \cdot 10^{94}$	$-9.027 \cdot 10^{58}$
Risk aversion parameter b_2	Portfolio E	Portfolio D Strike 0.9	Portfolio D* Strike 1.0
0.01	11.82	19.04	16.28
0.1	11.37	-39635	-41225
0.2	5.91	$-9.301 \cdot 10^{7}$	$-9.552 \cdot 10^{7}$
0.3	-1030	$-8.746 \cdot 10^{14}$	$-8.724 \cdot 10^{14}$
0.5	$-1.414 \cdot 10^{8}$	$-1.507 \cdot 10^{29}$	$-1.385 \cdot 10^{29}$
0.8	$-2.340 \cdot 10^{16}$	$-5.300 \cdot 10^{50}$	$-4.470 \cdot 10^{50}$
0.9	$-1.688 \cdot 10^{19}$	$-8.221 \cdot 10^{57}$	$-6.722 \cdot 10^{57}$

area the CDF associated with the distribution which include options is the lower curve. The lower left tails of the distributions are displayed in Figure 3.

Consider now the regulated portfolio E. Recall that this portfolio is selected by a *risk averse* casualty insurer, when he is not permitted to underwrite hazardous risks. From Table 6 we see that as far as the policyholders are concerned, this portfolio clearly dominates all other portfolios when risk aversion is greater than 0.01. Recall from Table 4 that the regulated insurance company prefers portfolio D including puts with strike price 0.9. How can we interpret this result?

It seems as if we are looking at a conflict of interests between the casualty insurer and his policyholders. But at this point we must distinguish between two different situations. First, in regard to the earth quake example given above, consider a Norwegian casualty insurer whose policyholders are not exposed to earthquake risk. Then ask if he should insure Japanese

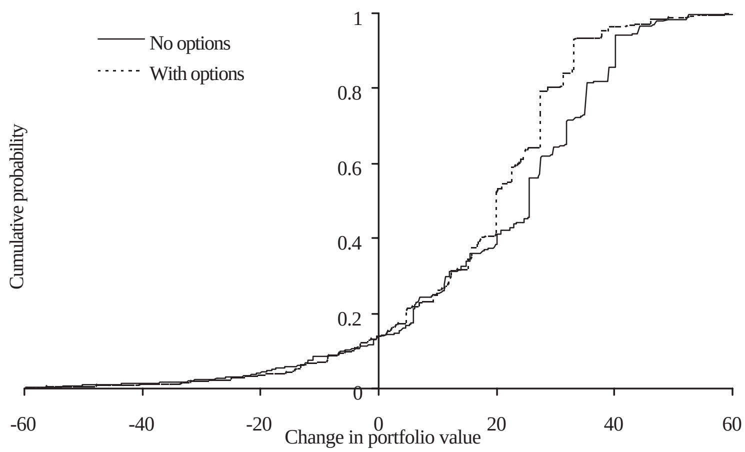

Figure 2: Cumulative distribution functions for the change in wealth $\Delta V^{-}_{\tau s}$, between time $t = 0$ and $t = \tau$.

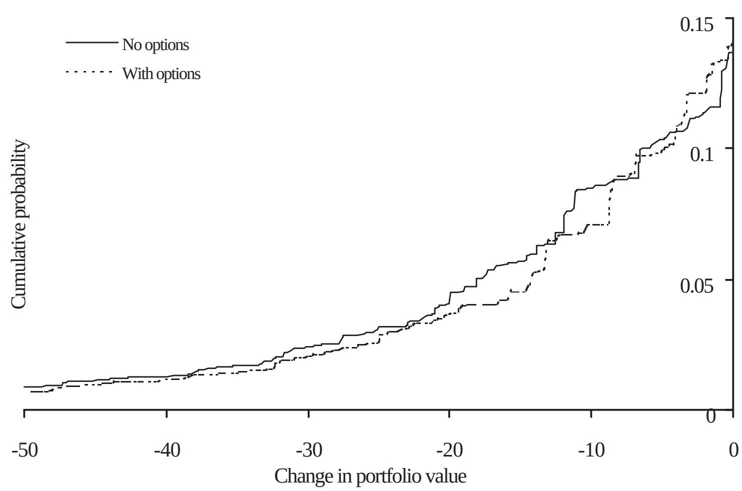

Figure 3: Lower left tail of cumulative distribution functions for the change in wealth $\Delta V^{-}_{\tau s}$, between time $t = 0$ and $t = \tau$.

citizens against earthquakes. From the perspective of Norwegian policyholders according to our analysis the answer is no. Norwegian citizens have no desire to buy insurance against earthquakes. Therefore they resent the idea that the insurance companies, from which they buy casualty insurance, underwrite hazardous earthquake risk.

Now consider a casualty insurer located in an area frequently exposed to earthquakes. People in this area *need* earthquake insurance. The insurance companies in this area already have earthquake risks in their insurance portfolios. In fact, in as much as the policyholders are the true owners of the insurance company, these firms are owned by policyholders who have bought earthquake insurance. Looking again at Table 6, in this case the correct portfolios to compare are portfolio C and D. The policyholders and the insurance company certainly agree on these portfolios. The financially reinsured portfolio D is clearly preferred. We conclude that the financial markets do provide a beneficial opportunity to hedge hazardous insurance risks. With this opportunity the insurance company may assume greater insurance risk. The derivatives will help finance potential huge claims payments. Not even in worst case scenarios will the company seriously violate the legal requirements. Therefore by incurring relatively modest recourse costs, it will be able to maintain compliance with the statutes.

Our results depend on the type of catastrophe risk that is insured and the existence of financial options that are sensitive to such risk. However, we consider our choice of correlation between claims due to catastrophes and the chosen stock index, 0.7, to be low enough to cover many cases. Further the results depend on the relationship between premiums and claims, but in practice that ratio does not vary too much. Finally, the price of financial reinsurance is also important, i.e. the price of the put options. We believe our asset pricing model does not select too low prices. Still, there are fixed costs and information costs connected to engaging in financial reinsurance, implying that there are economies of scale in such endeavours. Thus our results applies best to companies having a certain size in related activities. In summary we believe our results apply to a wide range of casualty insurance cases.

6 Conclusion

We have demonstrated that financial derivatives provide an effective means to hedge hazardous insurance risks. A distinction must be made between casualty insurers that are basically owned by policyholders who need insurance against hazardous risks, and those who are not. The former already have hazardous risks in their insurance portfolio and are well advised to consider buying financial reinsurance. For these companies the question is not whether or not to take on hazardous risks, but how to manage these risks.

In regard to casualty insurers whose policyholders are not exposed to specific hazardous risks, our analysis suggests not to underwrite such risks. Even if these risks are attractive to the company in terms of expected net cash flows, when policyholders do not need insurance against hazardous risks neither do they want such risks to be part of underwriters insurance portfolios.

References

Bell, D. (1995), 'Risk, return and utility', *Management Science* **41**, 23–30.

Black, F. & Scholes, M. (1973), 'The pricing of options and corporate liabilities', *Journal of Political Economy* **81**, 637–659.

Cariño, D. R. & Turner, A. L. (1998), Multiperiod asset allocation with derivative assets, *in* W. T. Ziemba & J. M. Mulvey, eds, 'Worldwide Asset and Liability Modeling', Cambridge University Press, Cambridge, U.K., pp. 182–204.

Cariño, D. R. & Ziemba, W. T. (1998), 'Formulation of the Russell-Yasuda Kasai financial planning model', *Operations Research* **46**(4), 443–449.

Consigli, G. & Dempster, M. A. H. (1998), 'Dynamic stochastic programming for asset-liability management', *Annals of Operations Research* **81**, 131–162.

Cox, J. C. & Rubinstein, M. (1985), *Option Markets*, Prentice-Hall, Englewood Cliffs, N. J.

Dert, C. (1995), Asset Liability Management for Pension Funds, A Multistage Chance Constrained Programming Approach, PhD thesis, Erasmus University, Rotterdam, The Netherlands.

Dupačová, J., Bertocchi, M. & Moriggia, V. (1998), Postoptimality for scenario based financial planning models with an application to bond portfolio management, *in* W. T. Ziemba & J. M. Mulvey, eds, 'Worldwide Asset and Liability Modeling', Cambridge University Press, Cambridge, U.K., pp. 263–285.

Gaivoronski, A. A. & de Lange, P. E. (1998), An asset liability model for casualty insurers: Complexity reduction vs. parameterized decision rules, Accepted for publication in Annals of Operations Research.

Gaivoronski, A. A. & Stella, F. (1998), Nonstationary optimization approach for finding universial portfolios, Working Paper, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway, To appear in Annals of Operations Research.

Gaivoronski, A. A., Høyland, K. & de Lange, P. E. (1998), Statutory regulations of casualty insurance companies: An example from the norwegian casualty insurance industry, Working Paper, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway.

Høyland, K. & Wallace, S. W. (1999*a*), Analyzing legal restrictions in the Norwegian life insurance business using a multi-period asset liability model, Accepted for publication in European Journal of Operations Research.

Høyland, K. & Wallace, S. W. (1999*b*), 'Generating scenario trees for multi stage decision problems', *Management Science.* To appear.

Huang, C. & Litzenberger, R. H. (1988), *Foundations of Financial Economics*, North-Holland, New York.

Hull, J. C. (1997), *Options, Futures and Other Derivatives*, 3rd edn, Prentice Hall International, London.

Jackwerth, J. C. (2000), 'Recovering risk aversion from option prices and realized returns', *Review of Financial Studies* **13**(2), Forthcoming.

Jackwerth, J. C. & Rubinstein, M. (1996), 'Recovering probability distributions from option prices', *Journal of Finance* **51**(5), 1611–1631.

King, A. J. (1993), ‘Asymmetric risk measures and tracking models for portfolio optimization under uncertainty’, *Annals of Operations Research* **45**, 165–177.

Kusy, M. I. & Ziemba, W. T. (1986), ‘A bank asset and liability management model’, *Operations Research* **34**(3), 356–376.

Mulvey, J. M. & Vladimirou, H. (1992), ‘Stochastic network programming for financial planning problems’, *Management Science* **38**(11), 1642–1664.

Rockafellar, R. T. & Wets, R. J.-B. (1991), ‘Scenarios and policy aggregation in optimization under uncertainty’, *Mathematics of Operations Research* **16**(1), 119–147.

Zenios, S. A., Holmer, M. R., McKendall, R. & Vassiadou-Zeniou, C. (1998), ‘Dynamic models for fixed-income portfolio management under uncertainty’, *Journal of Economic Dynamics & Control* **22**(10), 1517–1541.

Paper 5:

Fleten, S-E, Jørgensen, T. & Wallace, S. W. (1998). Real options and managerial flexibility. *Telektronikk* 94(3/4), 62-66.

Abstract:
http://wgate.bibsys.no/gate1/SHOW?objd=r99003819&base=FORSKPU2&lang=N&bibk=a