COMISEF Computational Optimization Methods in Statistics, Econometrics and Finance

- Marie Curie Research and Training Network funded by the EU Commission through MRTN-CT-2006-034270 -

COMISEF WORKING PAPERS SERIES

WPS-040 03/06/2010

Multistage Stochastic Portfolio Optimisation in Deregulated Electricity Markets Using Linear Decision Rules

P. Rocha
D. Kuhn

Multistage Stochastic Portfolio Optimisation in Deregulated Electricity Markets Using Linear Decision Rules

Paula Rocha* Daniel Kuhn†

03 June 2010

Abstract

The deregulation of electricity markets increases the financial risk faced by retailers who procure electric energy on the spot market to meet their customers' electricity demand. To hedge against this exposure, retailers often hold a portfolio of electricity derivative contracts. In this paper, we propose a multistage stochastic mean-variance optimisation model for the management of such a portfolio. To reduce computational complexity, we perform two approximations: stage-aggregation and linear decision rules (LDR). The LDR approach consists of restricting the set of decision rules to those affine in the history of the random parameters. When applied to mean-variance optimisation models, it leads to convex quadratic programs. Since their size grows typically only polynomially with the number of periods, they can be efficiently solved. Our numerical experiments illustrate the value of adaptivity inherent in the LDR method and its potential for enabling scalability to problems with many periods.

Keywords: OR in energy, electricity portfolio management, stochastic programming, risk management, linear decision rules

1 Introduction

Over the recent decades the energy industry has been undergoing liberalisation and deregulation. As a result, state-owned utilities are being privatised, and vertically integrated companies are being replaced by firms specialised in generation, transmission, distribution, or retail sale of energy. In many countries, this deregulation process culminates in the emergence of competitive spot markets, along with forward and derivatives markets.

Under this new setting, firms shift their focus from reliable and cost-efficient energy supply to more profit-oriented goals, competing to provide energy at the price set by the market. Therefore, traditional optimisation methods aimed at minimising expected costs without accounting

*Corresponding author, Department of Computing, Imperial College London, London SW7 2AZ, United Kingdom, e-mail: `paula.martins-da-silva-rocha08@imperial.ac.uk`

†Department of Computing, Imperial College London, London SW7 2AZ, United Kingdom, e-mail: `d.kuhn@imperial.ac.uk`

for risk and market behaviour are now redundant. This has lead to a surge in publications attempting to address the need for models adapted to the deregulated environment. The focus of the academic literature has been primarily on the perspective of the producer, specifically on power generation scheduling and bidding problems of generating companies. However, much less attention has been paid to the procurement of electric energy by retailers.

In a deregulated market, utility companies are more exposed to financial risk. Due to the non-storability of electricity and inelastic electricity demand, the electricity spot price is one of the most volatile commodity prices [29]. Under the regulated regime, electricity providers were able to pass external fuel price shocks onto consumers through regulated electricity prices. However, in the deregulated environment, such cost recovery is unlikely. Since the electricity price charged to the final consumer is usually fixed long before consumption occurs, electricity providers who purchase electric energy in the spot market absorb the entire risk of volatile spot prices. Therefore, electricity retailers usually seek protection against this uncertainty by managing a portfolio of financial and/or physical electricity derivative contracts (see [10] for a survey of popular financial instruments), in part to lock in the future price of electric energy.

Portfolio optimisation dates back to the seminal work of Markowitz [26], who proposes a methodology to construct efficient portfolios based on a trade-off between expected return of a portfolio and its associated risk measured in terms of the portfolio variance. Since this approach is static, that is, rebalancing of the portfolio is not envisaged, it fails to capture two important aspects of portfolio management: the trade-off between short-term and long-term consequences of an investment strategy based on the evolution of the random parameters, and the presence of transactions costs that affect portfolio holdings over time. Hence, this methodology may lead to short-sighted strategies, if applied repeatedly over subsequent periods, as the model does not account for the value of waiting for new information [37]. In contrast, a multistage stochastic programming approach enables the modelling of portfolio rebalancing at multiple future time points, in each case based on the information available up to that particular time point. For a comprehensive overview on multistage stochastic programming, see, e.g., [4, 19, 32]. A review of mean-variance portfolio models is provided in [34].

The application of portfolio theory to construct multistage stochastic optimisation models for electricity firms is relatively recent. One of the earliest contributions is due to Fleten et al. [14], who suggest that production planning and financial risk management should be integrated in order to maximise expected profit at some acceptable level of risk. Multistage stochastic models for the electricity procurement of utility companies have been proposed in [13, 16, 18, 21]. These papers consider mean-risk optimisation models (see, e.g., [32, Chapter 6]), which encompass several ways of procuring electric energy (for instance, via bilateral volume contracts, power derivative contracts, spot contracts and self-production) to satisfy the customers' electricity demand. The trading of futures at intermediate periods is envisaged in [13, 21], whereas the acquisition of energy derivatives and the signing of bilateral contracts occur at the beginning of the planning horizon only in [16, 18]. To improve model tractability, electricity demand is assumed to be deterministic in [13, 18]. All of these models use some variant of the Conditional

Value-at-Risk to quantify risk.

Stochastic programming provides a powerful mechanism for modelling dynamic portfolio selection problems. However, the arising optimisation models are notoriously difficult to solve. Only recently, this common perception has received a theoretical underpinning. Dyer and Stougie prove that two-stage stochastic programming problems are #P-hard [12]. A rather pessimistic verdict is also given by Shapiro and Nemirovski who demonstrate that multistage stochastic programs "generically are computationally intractable already when medium-accuracy solutions are sought" [33]. Complexity results of this type indicate that, for fundamental reasons, stochastic programming problems need to undergo some simplification in order to gain computational tractability. Note that analytical solutions are only available for unrealistically simple stochastic programming models.

The classical approach to make stochastic programming models amenable to numerical optimisation algorithms is to replace the underlying process of the random parameters by a discrete stochastic process, which is representable as a finite scenario tree. This tree ramifies at all time points when new random data becomes observable. Scenario tree approaches to stochastic programming have been studied extensively over the past decades, see, e.g., the survey paper [11]. Scenario trees are popular because they support intuition and lead to accurate results when having many branches [28]. Their disadvantage is that the solution time of the underlying optimisation model scales with the size of the tree, while the tree grows exponentially with the number of decision stages. Sometimes it is possible to reduce the number of branches of an overly bushy tree via scenario reduction techniques [17]. Frequently, however, the number of branches emanating from each node must be larger than the number of random parameters observed at that node. Otherwise, arbitrage opportunities would be built into the tree that render the underlying optimisation model unbounded [22]. In a scenario tree framework, the exponential growth of these models cannot be avoided, and it may even be hard to find a solution that is feasible (let alone optimal).

The recourse decisions associated with a stochastic program represent *decision rules*, that is, measurable functions of the observable random parameters. Instead of approximating the data process (as is done in tree-based methods), one can alternatively simplify the functional form of the decision rules. Focussing on *linear* decision rules (LDR), for instance, converts the original stochastic program to a semi-infinite program. Only with the advent of modern robust optimisation techniques in the last few years, it has been recognised that this semi-infinite program is equivalent to a conic optimisation problem that can be solved efficiently, i.e., in polynomial time [3]. The striking advantage of the LDR approximation is that it permits scalability to multistage models: in a linear decision framework, the problem size grows only polynomially with the number of decision stages. The LDR approximation has successfully been used to solve supply chain problems with more than 70 decision stages [2, 24], network design problems involving hundreds of random variables [1], or robust control problems involving 12 state variables and 20 time stages [15].

An application of LDRs to financial portfolio optimisation is due to Calafiore [7], who later extended it to account for transactions costs [8]. The author proposes a multiperiod version of the mean-variance Markowitz model, subject to constraints on the expected portfolio composition at each intermediate period. By restricting the form of the portfolio rebalancing decisions to affine functions of the past periods' returns, the problem is then converted into a finite-dimensional convex quadratic program. To the best of our knowledge, the LDR approximation has not yet been applied in the context of electricity portfolio optimisation.

In this paper, we present a multistage mean-variance model for the management of a hedging portfolio of electricity derivatives from the viewpoint of a price-taking retailer that procures electric energy to satisfy its customers' electricity demand. To reduce computational complexity, we carry out a stage-aggregation (see, e.g., [4, Chapter 11.2] and [23]) and, in addition, we apply a LDR approximation. Both of these simplifications lead to a conservative approximation of the original problem and thus underestimate the retailer's flexibility. We show that the resulting problem can be reformulated as a tractable convex quadratic program. Since this approximate problem grows only polynomially with the number of periods, it can be solved efficiently. Moreover, it only requires information about the support and the first four moments of the uncertain parameters — a desirable feature considering that the full joint distribution of the random parameters is rarely available. In a series of numerical experiments, we provide insight into the sensitivity of the optimal value to a selection of input parameters and illustrate the value of adaptivity inherent in the LDR approximation. We also evaluate the accuracy of the stage-aggregation approximation and highlight its potential for reducing computational time. Finally, to assess the scalability and the accuracy of the LDR approach, we compare it to a sample average approximation (SAA) that consists of constructing a scenario tree via conditional sampling [30]. Our tests indicate that the LDR method offers superior scalability as well as accuracy.

The remainder of the paper is organised as follows. Section 2 specifies the retailer's electricity procurement problem. Section 3.1 presents the electricity portfolio optimisation model, which is formulated as a multistage stochastic program in Section 3.2. Section 4 approximates the exact problem by a numerically tractable problem via stage-aggregation and LDRs. Section 5 reports on numerical results, and conclusions are drawn in Section 6.

2 Problem Specification

A price-taking electricity retailer must meet the electricity demand of its customers over a given planning horizon which is subdivided into time intervals indexed by $t \in \mathcal{T} := \{1, \dots, T\}$. Without loss of generality, we assume that interval t starts at time $(t-1)\Delta$, where Δ represents the interval length. The amount of electric energy demanded at period t is denoted by D_t. Assuming that the retailer has no generation capability, it can satisfy this demand by purchasing electric energy for immediate consumption on the spot market, at price S_t per unit of energy. Here, S_t denotes the average spot price in interval t. Working with average prices is justified if the demand volume D_t is consumed at a constant rate within interval t.

Relying solely on the spot market to satisfy demand is known to be very risky due to occasional spikes in spot prices [29]. In order to hedge against spot price risk the retailer can purchase different types of electricity forward contracts for physical delivery, indexed by $i \in \mathcal{I} := \{1, \ldots, I\}$. A forward contract constitutes an obligation to buy (or sell) a prescribed volume of electric energy during a certain delivery period in the future, at a pre-established price per unit of energy. Note that energy derivative prices are typically quoted per unit of energy rather than per contract. The forward contract types differ with respect to their delivery period (e.g., monthly, quarterly, or annual) and their load profile, which specifies the delivery rate during the delivery period. Commonly traded load profiles are base load and peak load. Base load provides a constant delivery rate during every hour of the delivery period, whereas peak load provides a constant delivery rate from 8am to 8pm on any weekday within the delivery period. For forward contracts of type i, let $B(i)$ denote the first interval and $E(i)$ the last interval in the delivery period, and $\mathcal{T}_i$ the set of time intervals in which electricity is delivered. The volume of electric energy supplied by a contract of type i during interval $t \in \mathcal{T}_i$ is v_t^i, so the total volume of such a contract is $v^i = \sum_{t \in \mathcal{T}_i} v_t^i$. For this type of contract, the forward price specified at the start of interval t, which is to be paid during the delivery period, is F_t^i per unit of energy. For ease of exposition, we assume that trading of any forward contract ceases at the start of its delivery period.

Apart from entering into forward contracts, the retailer may also acquire different types of European call options, indexed by $j \in \mathcal{J} := \{1, \ldots, J\}$. A European call option of type j gives the retailer the right to buy a forward contract of type $i(j) \in \mathcal{I}$ at maturity $M(j)$, at the strike price K^j per unit of energy. In exchange for this right, the retailer pays the premium C_t^j per unit of energy at period t when the call option is negotiated. We assume that options are financially settled, that is, the price difference between the agreed strike price and the market price of the underlying forward contract is settled in cash at the maturity time of the call option.

For the sake of a transparent exposition, it is assumed that no transaction costs are incurred in trading and that no discounting takes place over time. Note that these assumptions may easily be relaxed at the cost of additional notation.

Our assumptions are inspired by the structure and regulations of real electricity markets such as the European Energy Exchange or Nord Pool. At present, base and peak load forwards, futures, and European-style options are traded in Nord Pool's financial market. Forward contracts are listed for each calendar month, quarter and year, with a delivery rate of 1 MW. Forward contracts are traded until the day before delivery starts and are settled against the spot price throughout the delivery period. The options' underlying instruments are either quarterly or annual forward contracts. Options can only be exercised on the expiry day, which is set to a few days before the delivery period of the underlying forward contract starts. Although no physical delivery of power contracts takes place in the Nord Pool financial market, there are other markets, such as the European Energy Exchange, where this possibility is envisaged. Note that our model formulation in Section 3.1 is, nonetheless, consistent with cash settlement of forward contracts.

3 Model Formulation

3.1 Portfolio Optimisation Model

The retailer aims to determine a cost-efficient mix of electricity derivative contracts over a medium-term planning horizon, given that the customers' electricity demand must be met uninterruptedly. Let $x^i_{\mathrm{f},t}$ represent the number of forward contracts of type i bought (if $x^i_{\mathrm{f},t} \geq 0$) or sold (if $x^i_{\mathrm{f},t} < 0$) by the retailer at the beginning of period t, and let $x^i_{\mathrm{F},t}$ denote the retailer's position in forward contracts of the same type in interval t after portfolio rebalancing. In addition, let $x^j_{\mathrm{c},t}$ denote the number of European call options of type j traded by the retailer at the start of period t, and let $x^j_{\mathrm{C},t}$ be the retailer's position in type-j options in interval t after portfolio rebalancing. Note that in order to obtain a tractable optimisation model, we assume that fractional numbers of contracts may be held.

The retailer faces four types of costs in any period $t \in \mathcal{T}$, which are related to different financial activities:

Spot Market Transactions: The volume of electric energy supplied through the forward contracts of type i in period t is $v^i_t x^i_{\mathrm{F},t}$ if $t \in \mathcal{T}_i$ and zero otherwise. Any gap between the energy supplied through the entire portfolio of forward contracts and the customers' electricity demand D_t is covered through transactions in the spot market. Hence, any surplus on electric energy provided through forward contract agreements is sold on the spot market at price S_t. Conversely, if a shortage of energy arises, it will lead to purchases of electric energy on the spot market. The resulting cash outflow amounts to

$$z_{\mathrm{s},t} = S_t \Bigg(D_t - \sum_{\substack{i \in \mathcal{I}: \\ t \in \mathcal{T}_i}} v^i_t x^i_{\mathrm{F},t} \Bigg).$$

Forward Trading: Buying $x^i_{\mathrm{f},t}$ contracts of type i in period t incurs a total cost of $F^i_t x^i_{\mathrm{f},t} v^i$. Recall that v^i stands for the total volume of a forward contract of type i. Note that trading of these forward contracts ceases at $B(i)$, so that we can disregard forward contracts whose delivery period has started by period t.

$$z_{\mathrm{f},t} = \sum_{\substack{i \in \mathcal{I}: \\ B(i) > t}} F^i_t x^i_{\mathrm{f},t} v^i$$

Call Option Trading: In exchange for a payment $C^j_t x^j_{\mathrm{c},t} v^{i(j)}$, the retailer obtains the right to purchase $x^j_{\mathrm{c},t}$ forward contracts of type $i(j)$, at the pre-established price K^j per unit of energy, at maturity. Since no options of type j are exchanged after $M(j)$, we can disregard options that have matured by period t.

$$z_{\mathrm{c},t} = \sum_{\substack{j \in \mathcal{J}: \\ M(j) > t}} C^j_t x^j_{\mathrm{c},t} v^{i(j)}$$

Exercise of Call Options: A European call option is exercised only if its strike price is exceeded by the market price of the underlying forward at maturity. Since we assume that options are financially settled, the resulting payoff per unit of energy amounts to $\max(F^{i(j)}_{M(j)} - K^j; 0)$. Notice that solely options that mature in period t can be exercised. The overall cost from exercising options in interval t is thus given by

$$z_{\mathrm{e},t} = -\sum_{\substack{j\in\mathcal{J}:\\ M(j)=t}} \max(F^{i(j)}_t - K^j; 0)\, x^j_{\mathrm{C},t} v^{i(j)}.$$

The retailer's aim is then to find a policy for the management of a portfolio of forward contracts and European call options that minimises the total cost

$$\sum_{t=1}^{T} \left\{ z_{\mathrm{s},t} + z_{\mathrm{f},t} + z_{\mathrm{c},t} + z_{\mathrm{e},t} \right\}.$$

The retailer's decisions are subject to the following constraints:

Budget Constraints: We impose the following budget restrictions at any $t \in \mathcal{T}$.

$$\begin{aligned} x^i_{\mathrm{F},t} &= x^i_{\mathrm{F},t-1} + x^i_{\mathrm{f},t}, \ i \in \mathcal{I} \\ x^j_{\mathrm{C},t} &= x^j_{\mathrm{C},t-1} + x^j_{\mathrm{c},t}, \ j \in \mathcal{J} \end{aligned}$$

These constraints guarantee that the position in derivatives of a certain type in interval t equates the respective position in interval $t-1$ adjusted by the transactions at the start of period t.

No-short-selling Constraints: We assume that short-selling of forwards and call options is not allowed at any $t \in \mathcal{T}$ since electricity retailers usually use energy derivatives for hedging and not for speculation.

$$\begin{aligned} x^i_{\mathrm{F},t} &\geq 0, \ i \in \mathcal{I} \\ x^j_{\mathrm{C},t} &\geq 0, \ j \in \mathcal{J} \end{aligned}$$

No-trading Constraints: We impose the following constraints at any $t \in \mathcal{T}$ to ensure that the respective trading volume of contracts no longer exchanged in period t is equal to zero.

$$\begin{aligned} x^i_{\mathrm{f},t} &= 0, \ i \in \mathcal{I} : B(i) \leq t \\ x^j_{\mathrm{c},t} &= 0, \ j \in \mathcal{J} : M(j) \leq t \end{aligned}$$

Note that our model is flexible enough to accommodate additional linear constraints on portfolio adjustments and composition.

3.2 Multistage Stochastic Program

For notational convenience, we work henceforth with an abstract formulation of the portfolio optimisation problem described in Section 3.1. We denote by $u_t \in \mathbb{R}^n$ the control variable comprising the trading decisions $x^i_{\mathrm{f},t}, i \in \mathcal{I}$, and $x^j_{\mathrm{c},t}, j \in \mathcal{J}$, while $s_t \in \mathbb{R}^n$ is a state variable that comprises the position variables $x^i_{\mathrm{F},t}, i \in \mathcal{I}$, and $x^j_{\mathrm{C},t}, j \in \mathcal{J}$. The cost vectors $c_t \in \mathbb{R}, c_{\mathrm{u},t} \in \mathbb{R}^n$ and $c_{\mathrm{s},t} \in \mathbb{R}^n$ are defined in such a way that $c_{\mathrm{u},t}^\top u_t = z_{\mathrm{f},t} + z_{\mathrm{c},t}$ and $c_t + c_{\mathrm{s},t}^\top s_t = z_{\mathrm{s},t} + z_{\mathrm{e},t}$ hold.

The budget, the no-short-selling and the no-trading constraints are equivalent to $s_t = s_{t-1} + u_t$, $s_t \geq 0$ and $G_{\mathrm{u},t} u_t = 0$, respectively. Here, $G_{\mathrm{u},t}$ denotes a truncation operator that eliminates from u_t the decisions on contracts which are still traded in interval t.

Stochasticity appears in the portfolio optimisation model in the form of uncertain electricity demands D_t, spot prices S_t, and derivative prices $F_t^i, i \in \mathcal{I} : B(i) \geq t$, and $C_t^j, j \in \mathcal{J} : M(j) \geq t$, which are revealed sequentially at periods $t \in \mathcal{T}$. Some of these random parameters, in particular spot and derivative contract prices, are typically highly correlated. Therefore, we assume that it is possible to represent the uncertain parameters revealed in interval t as functions of a smaller set of risk factors $\zeta_t \in \mathbb{R}^{k_t}$. In other words, we assume that the variability in all random parameters of period t is completely explained by the variability in the risk factors ζ_t. Note that the dependence between the uncertain parameters and the risk factors may be *non-linear.* For technical reasons related to Section 4.2, we introduce the vector $\xi_t \in \mathbb{R}^{p_t}$ which is formed by appending to ζ_t enough random parameters perfectly dependent on ζ_t to guarantee that $c_t, c_{\mathrm{u},t}$ and $c_{\mathrm{s},t}$ are representable as linear functions of $\xi^t := (\xi_1^\top, \ldots, \xi_t^\top)^\top \in \mathbb{R}^{p^t}$, where $p^t := \sum_{s=1}^t p_s$. Note that a ξ_t with these properties always exists; for instance, we are free to define $\xi_t := (\zeta_t^\top, c_t, c_{\mathrm{u},t}^\top, c_{\mathrm{s},t}^\top)^\top$. For an example, we refer to Section 5.1. We denote the history of risk factors up to period t by $\zeta^t := (\zeta_1^\top, \ldots, \zeta_t^\top)^\top \in \mathbb{R}^{k^t}$, where $k^t := \sum_{s=1}^t k_s$. Moreover, we set $\zeta := \zeta^T$, $\xi := \xi^T$, $k := k^T$ and $p := p^T$. For technical reasons related to Section 4.2, the support Z_t of ζ_t is assumed to be representable as a non-empty compact polyhedron and to span $\mathbb{R}^{k_t}$. In contrast, the support of ξ_t, which contains ζ_t as a subvector, is typically non-convex. Without loss of generality, we require that $k_1 = 1$ and $Z_1 = \{1\}$. Thus, ζ_1 is a degenerate random variable governed by a Dirac distribution centered at 1. This specification allows us to represent affine functions of the non-degenerate risk factors $(\zeta_2^\top, \ldots, \zeta_t^\top)^\top$ in a condensed manner as linear functions of ζ^t.

In practice, the decisions $u_1, s_1, \ldots, u_T, s_T$ are not pre-committed at the start of the planning horizon. Instead, they are selected sequentially in time and are, therefore, allowed to adapt to the available information. Consequently, u_t and s_t are interpreted as decision rules, i.e., functions that map the observation history ζ^t of the risk factors to decisions $u_t(\zeta^t)$ and $s_t(\zeta^t)$, respectively. The space of decision rules $\mathcal{X}_{k^t,n}$ is the space of all measurable bounded functions from $\mathbb{R}^{k^t}$ to $\mathbb{R}^n$. Stipulating that decisions depend solely on the history of risk factors is a reasonable assumption since the random parameters can be uniquely explained by the risk factors. Indeed, observing perfectly dependent random variables does not provide any additional information.

Using the notation introduced so far, the portfolio optimisation problem may be formulated abstractly as the following multistage stochastic program

$$\begin{array}{ll} \min & \mathbb{F}\left(\sum_{t=1}^{T} c_t(\xi^t) + c_{\mathrm{u},t}(\xi^t)^\top u_t(\zeta^t) + c_{\mathrm{s},t}(\xi^t)^\top s_t(\zeta^t) \right) \\ \text{s.t.} & u_t, s_t \in \mathcal{X}_{k^t,n} \quad \forall t \in \mathcal{T} \\ & \left. \begin{array}{l} s_t(\zeta^t) = s_{t-1}(\zeta^{t-1}) + u_t(\zeta^t) \\ s_t(\zeta^t) \geq 0 \\ G_{\mathrm{u},t} u_t(\zeta^t) = 0 \end{array} \right\} \mathbb{P}\text{-a.s. } \forall t \in \mathcal{T} \end{array} \tag{$\mathcal{SP}$}$$

where $\mathbb{F}$ is a probability functional (with respect to the distribution $\mathbb{P}$ of the random vector ξ) that maps the random overall costs to a real number.

4 Approximations

The stochastic program $\mathcal{SP}$ is a functional optimisation problem over an infinite-dimensional space of policies. Thus, it is computationally intractable. LDRs may be used to overcome this obstacle. Once this approximation is applied, the resulting multistage optimisation problem is, in principle, amenable to polynomial-time solution procedures. However, this problem may still contain a large number of decision stages and, consequently, decision variables, possibly leading to unacceptable computation times. In order to set up an approximate portfolio optimisation problem that can be efficiently solved, we thus successively perform two approximations based on stage-aggregation and LDRs.

4.1 Stage-Aggregation

To speed up computation, we establish a new optimisation problem with fewer decision stages. The planning horizon $\mathcal{T} = \{1, \ldots, T\}$ is subdivided into a number of macroperiods indexed by $m \in \mathcal{M} := \{1, \ldots, M\}$. For each $m \in \mathcal{M}$, let t_m be the first interval belonging to macroperiod m. We always require $t_1 = 1$. Moreover, for notational convenience, we define $t_{M+1} := T+1$. We require that each macroperiod covers one or more normal periods, which implies $|\mathcal{M}| \leq |\mathcal{T}|$. We assume that electricity prices and demand are no longer observed at all intervals $t \in \mathcal{T}$ but only at periods $t \in \tilde{\mathcal{T}} := \{t_m : m \in \mathcal{M}\}$. Thus, decisions taken during macroperiod m only rely on the history of risk factors at the beginning of macroperiods, $\tilde{\zeta}^m := (\zeta_{t_1}^\top, \ldots, \zeta_{t_m}^\top)^\top \in \mathbb{R}^{\tilde{k}^m}$, where $\tilde{k}^m := \sum_{m'=1}^{m} k_{t_{m'}}$. There is no incentive to rebalance the portfolio of electricity derivatives if no new information is observed. Therefore, all trading decisions within macroperiod m are taken at the beginning of period t_m only. Hence, we can set $u_t(\tilde{\zeta}^m) = 0$ at $t \in \{t_m+1, \ldots, t_{m+1}-1\}$. Due to the budget constraints, the positions in the different derivative contracts remain constant at $s_{t_m}(\tilde{\zeta}^m)$ throughout the entire macroperiod m. Consequently, the no-short-selling restrictions are redundant at $t \in \mathcal{T} \setminus \tilde{\mathcal{T}}$. It is implicit that any excess (or shortage) of electric energy to meet the customers' demand is sold (or acquired) in the spot market at all periods $t \in \mathcal{T}$. Also, call options may be exercised at any $t \in \mathcal{T}$, since their maturities do not necessarily coincide with the start dates of the macroperiods.

By suppressing trading at periods $t \in \mathcal{T} \setminus \tilde{\mathcal{T}}$, the feasible set of problem $\mathcal{SP}$ is reduced. In addition, the information that underlies the trading decisions has been limited, since only observations of risk factors at periods $t \in \tilde{\mathcal{T}}$ affect the decisions. For these two reasons, the stage-aggregated optimisation problem constitutes a conservative approximation to $\mathcal{SP}$ in the sense that any policy feasible in the approximate problem can be extended to a policy feasible in $\mathcal{SP}$ with the same objective value, but the converse is not true.

Expressing the approximate problem in terms of decisions at $t \in \tilde{\mathcal{T}}$ only, we arrive at the following aggregated multistage stochastic program

$$
\begin{aligned}
\min \quad & \mathbb{F}\Bigg(\sum_{m=1}^{M} \tilde{c}_m(\xi^{t_{m+1}-1}) + c_{\mathrm{u},t_m}(\xi^{t_m})^\top u_{t_m}(\tilde{\zeta}^m) + \tilde{c}_{\mathrm{s},m}(\xi^{t_{m+1}-1})^\top s_{t_m}(\tilde{\zeta}^m)\Bigg) \\
\text{s.t.} \quad & u_{t_m}, s_{t_m} \in \mathcal{X}_{\tilde{k}^m,n} \quad \forall m \in \mathcal{M} \\
& \left.\begin{aligned}
& s_{t_m}(\tilde{\zeta}^m) = s_{t_{m-1}}(\tilde{\zeta}^{m-1}) + u_{t_m}(\tilde{\zeta}^m) \\
& s_{t_m}(\tilde{\zeta}^m) \geq 0 \\
& G_{\mathrm{u},t_m} u_{t_m}(\tilde{\zeta}^m) = 0
\end{aligned}\right\} \mathbb{P}\text{-a.s. } \forall m \in \mathcal{M}
\end{aligned}
\tag{$\mathcal{ASP}$}
$$

where

$$
\tilde{c}_m(\xi^{t_{m+1}-1}) := \sum_{t=t_m}^{t_{m+1}-1} c_t(\xi^t) \quad \text{and} \quad \tilde{c}_{\mathrm{s},m}(\xi^{t_{m+1}-1}) := \sum_{t=t_m}^{t_{m+1}-1} c_{\mathrm{s},t}(\xi^t).
$$

Problem $\mathcal{ASP}$ inherits some useful properties from problem $\mathcal{SP}$. By construction, the cost coefficients may be written as non-anticipative linear functions of the random parameters, that is, $\tilde{c}_m(\xi^{t_{m+1}-1}) = \tilde{c}_{\mathrm{c},m}^\top \xi^{t_{m+1}-1}$ for some vector $\tilde{c}_{\mathrm{c},m} \in \mathbb{R}^{p^{t_{m+1}-1}}$, $c_{\mathrm{u},t_m}(\xi^{t_m}) = C_{\mathrm{u},t_m}\xi^{t_m}$ for some matrix $C_{\mathrm{u},t_m} \in \mathbb{R}^{n\times p^{t_m}}$, and $\tilde{c}_{\mathrm{s},m}(\xi^{t_{m+1}-1}) = \tilde{C}_{\mathrm{s},m}\xi^{t_{m+1}-1}$ for some matrix $\tilde{C}_{\mathrm{s},m} \in \mathbb{R}^{n\times p^{t_{m+1}-1}}$. By the assumptions in Section 3.2, the support $\tilde{Z} := \times_{m=1}^{M} Z_{t_m}$ of the risk factors $\tilde{\zeta} := \tilde{\zeta}^M$ is representable by a non-empty compact polyhedron of the form

$$
\tilde{Z} = \{\tilde{\zeta} \in \mathbb{R}^{\tilde{k}} : W\tilde{\zeta} \geq h\}
$$

for some matrix $W \in \mathbb{R}^{l\times\tilde{k}}$ and a vector $h \in \mathbb{R}^l$, where $\tilde{k} := \tilde{k}^M$. Recall that we assumed that $\zeta_1 = 1$ $\mathbb{P}$-a.s. in Section 3.2. Thus, we require that the inequalities $W\tilde{\zeta} \geq h$ imply $\zeta_1 = e_1^\top\tilde{\zeta} = 1$, where e_1 denotes the first standard basis vector in $\mathbb{R}^{\tilde{k}}$.

4.2 Linear Decision Rule Approximation

The stage-aggregated problem $\mathcal{ASP}$ remains computationally intractable since it constitutes an optimisation problem over an infinite-dimensional function space. To gain numerical tractability, we apply a LDR approximation, that is, we restrict the functional form of the decision rules to those that are representable as

$$
u_{t_m}(\tilde{\zeta}^m) = \tilde{U}_m\tilde{\zeta}^m \quad \text{and} \quad s_{t_m}(\tilde{\zeta}^m) = \tilde{S}_m\tilde{\zeta}^m \tag{4.1}
$$

for some matrices $\tilde{U}_m, \tilde{S}_m \in \mathbb{R}^{n\times\tilde{k}^m}, m \in \mathcal{M}$. By considering only decision rules of the type (4.1) and taking the linearity of the cost coefficients in the history of the random data into account, one arrives at the following approximate problem.

$$
\begin{aligned}
\min \quad & \mathbb{F}(\xi^\top V\xi) \\
\text{s.t.} \quad & V \in \mathbb{R}^{p\times p}, \tilde{U}_m, \tilde{S}_m \in \mathbb{R}^{n\times\tilde{k}^m} \ \forall m \in \mathcal{M} \\
& V = \sum_{m=1}^{M} P_{t_{m+1}-1}^\top \tilde{c}_{\mathrm{c},m} e_1^\top Q_M + P_{t_m}^\top C_{\mathrm{u},t_m}^\top \tilde{U}_m Q_m + P_{t_{m+1}-1}^\top \tilde{C}_{\mathrm{s},m}^\top \tilde{S}_m Q_m \\
& \left.\begin{aligned}
& \tilde{S}_m R_m\tilde{\zeta} = \tilde{S}_{m-1}R_{m-1}\tilde{\zeta} + \tilde{U}_m R_m\tilde{\zeta} \\
& \tilde{S}_m R_m\tilde{\zeta} \geq 0 \\
& G_{\mathrm{u},t_m}\tilde{U}_m R_m\tilde{\zeta} = 0
\end{aligned}\right\} \mathbb{P}\text{-a.s. } \forall m \in \mathcal{M}
\end{aligned}
\tag{$\mathcal{ASP}^u$}
$$

Here, we used the truncation operators $P_t, t \in \mathcal{T}$, $Q_m, m \in \mathcal{M}$, and $R_m, m \in \mathcal{M}$, defined through

$$\begin{aligned} &P_t : \mathbb{R}^p \mapsto \mathbb{R}^{p^t}, \xi \mapsto \xi^t, \\ &Q_m : \mathbb{R}^p \mapsto \mathbb{R}^{\tilde{k}^m}, \xi \mapsto \tilde{\zeta}^m, \\ &R_m : \mathbb{R}^{\tilde{k}} \mapsto \mathbb{R}^{\tilde{k}^m}, \tilde{\zeta} \mapsto \tilde{\zeta}^m, \end{aligned}$$

and the fact that $e_1^\top Q_M \xi = \zeta_1 = 1$ $\mathbb{P}$-a.s. Since $\mathcal{ASP}^u$ was obtained by restricting the underlying feasible set, it provides an upper bound to problem $\mathcal{ASP}$. Notice that $\mathcal{ASP}^u$ involves only finitely many decision variables (the entries of the matrices $\tilde{U}_m, \tilde{S}_m, m \in \mathcal{M}$, and V). As the almost sure constraints in $\mathcal{ASP}^u$ are continuous in $\tilde{\zeta}$, they hold for all $\tilde{\zeta}$ in the support $\tilde{Z}$. Therefore, $\mathcal{ASP}^u$ exhibits semi-infinite constraints parameterised by $\tilde{\zeta} \in \tilde{Z}$ and appears to be intractable. However, it is possible to re-express this semi-infinite constraint system in terms of a finite number of linear constraints. The equality constraints in $\mathcal{ASP}^u$ imply that the linear hull of $\tilde{Z}$ belongs to the null space of the linear operators $\tilde{S}_m R_m - \tilde{S}_{m-1} R_{m-1} - \tilde{U}_m R_m$ and $G_{\mathrm{u},t_m} \tilde{U}_m R_m$. Given that $\tilde{Z}$ spans the whole of $\mathbb{R}^{\tilde{k}}$, we may equivalently require that $\tilde{S}_m R_m = \tilde{S}_{m-1} R_{m-1} + \tilde{U}_m R_m$ and $G_{\mathrm{u},t_m} \tilde{U}_m R_m = 0$. To simplify the semi-infinite inequality constraints, we use the following proposition, which can be seen as a special case of a major result in robust optimisation (cf., Theorem 3.2 in [3]):

Proposition 1. *For any $u \in \mathbb{R}^{\tilde{k}}$ the following statements are equivalent:*

(i) $u^\top \tilde{\zeta} \geq 0 \quad \forall \tilde{\zeta} \in \tilde{Z} = \{\tilde{\zeta} \in \mathbb{R}^{\tilde{k}} : W\tilde{\zeta} \geq h\}$;
(ii) $\exists \lambda \in \mathbb{R}^l$ *with* $\lambda \geq 0$, $W^\top \lambda = u$, *and* $h^\top \lambda \geq 0$.

Letting $u_i^\top$ denote the i-th row of the matrix $\tilde{S}_m R_m$, Proposition 1 allows us to replace the semi-infinite constraints $u_i^\top \tilde{\zeta} \geq 0$ for all $\tilde{\zeta} \in \tilde{Z}$ by a finite number of linear constraints involving a new decision vector $\lambda_i \in \mathbb{R}^l, i = 1, \ldots, n$. By interpreting $\lambda_i^\top$ as the i-th row of a matrix $\Lambda_m \in \mathbb{R}^{n \times l}$, we can replace the semi-infinite inequality constraints in $\mathcal{ASP}^u$ by the linear constraints $\Lambda_m W = \tilde{S}_m R_m$, $\Lambda_m h \geq 0$, and $\Lambda_m \geq 0$. Thus, $\mathcal{ASP}^u$ is equivalent to

$$\begin{array}{ll} \min & \mathbb{F}\big(\xi^\top V \xi\big) \\ \text{s.t.} & V \in \mathbb{R}^{p\times p}, \tilde{U}_m, \tilde{S}_m \in \mathbb{R}^{n\times \tilde{k}^m}, \Lambda_m \in \mathbb{R}^{n\times l} \ \forall m \in \mathcal{M} \\ & V = \displaystyle\sum_{m=1}^{M} P_{t_{m+1}-1}^\top \tilde{c}_{\mathrm{c},m} e_1^\top Q_M + P_{t_m}^\top C_{\mathrm{u},t_m}^\top \tilde{U}_m Q_m + P_{t_{m+1}-1}^\top \tilde{C}_{\mathrm{s},m}^\top \tilde{S}_m Q_m \\ & \left.\begin{array}{l} \tilde{S}_m R_m = \tilde{S}_{m-1} R_{m-1} + \tilde{U}_m R_m \\ G_{\mathrm{u},t_m} \tilde{U}_m R_m = 0 \\ \Lambda_m W = \tilde{S}_m R_m \\ \Lambda_m h \geq 0 \\ \Lambda_m \geq 0 \end{array}\right\} \forall m \in \mathcal{M} \end{array} \tag{4.2}$$

In mainstream stochastic programming the probability functional $\mathbb{F}$ is often chosen to be the expected value. A common approach to reflect risk averse preferences in optimisation problems is to let $\mathbb{F}$ be a mean-risk functional (see, e.g., [32, Chapter 6]), which constitutes a weighted average of the expected value and some measure of dispersion that quantifies the uncertainty of the costs. The advantage of this approach is that it allows for a trade-off between minimising the expected costs and their risk. Here, we use the variance as the dispersion measure — a popular choice which was first advocated by Markovitz in the context of financial portfolio optimisation [26]. For a given weight $\gamma \in [0, 1]$ assigned to the variance, we can express the objective function of problem (4.2) in terms of the second order moment matrix $\Phi := \mathbb{E}(\xi\xi^\top)$ and the fourth order moment tensor $\Psi := \mathbb{E}(\xi\xi^\top \otimes \xi\xi^\top)$ of the random vector ξ under the probability measure $\mathbb{P}$.

$$
\begin{aligned}
\mathbb{F}\big(\xi^\top V\xi\big) =& \gamma\, \mathbb{V}\mathrm{ar}\big(\xi^\top V\xi\big) + (1-\gamma)\,\mathbb{E}\big(\xi^\top V\xi\big)\\
=& \gamma \left[\mathbb{E}\Big\{\big(\xi^\top V\xi\big)^\top \big(\xi^\top V\xi\big)\Big\} - \Big(\mathbb{E}\{\xi^\top V\xi\}\Big)^2\right] + (1-\gamma)\,\mathbb{E}\big(\xi^\top V\xi\big)\\
=& \gamma \left[\mathbb{E}\Big(\mathrm{tr}\Big\{(\xi^\top \otimes \xi^\top)(V^\top \otimes V)(\xi \otimes \xi)\Big\}\Big) - \Big(\mathbb{E}\big(\mathrm{tr}\{V\xi\xi^\top\}\big)\Big)^2\right]\\
&+ (1-\gamma)\,\mathbb{E}\big(\mathrm{tr}\{V\xi\xi^\top\}\big)\\
=& \gamma \left[\mathbb{E}\Big(\mathrm{tr}\Big\{(V^\top \otimes V)(\xi \otimes \xi)(\xi^\top \otimes \xi^\top)\Big\}\Big) - \big(\mathrm{tr}\{V\Phi\}\big)^2\right] + (1-\gamma)\mathrm{tr}\{V\Phi\}\\
=& \gamma \left[\mathrm{tr}\big\{\big(V^\top \otimes V\big)\Psi\big\} - \big(\mathrm{tr}\{V\Phi\}\big)^2\right] + (1-\gamma)\mathrm{tr}\{V\Phi\} \qquad (4.3)
\end{aligned}
$$

Here, "tr" denotes the trace operator and "$\otimes$" denotes the Kronecker product. The equalities in the third and fifth rows follow from the the mixed-product property of the Kronecker product. Substituting (4.3) into (4.2) yields the following tractable convex quadratic program with linear constraints.

$$
\begin{array}{ll}
\min & \gamma \left[\mathrm{tr}\big\{\big(V^\top \otimes V\big)\Psi\big\} - \big(\mathrm{tr}\{V\Phi\}\big)^2\right] + (1-\gamma)\mathrm{tr}\{V\Phi\}\\
\text{s.t.} & V \in \mathbb{R}^{p\times p}, \tilde{U}_m, \tilde{S}_m \in \mathbb{R}^{n\times \tilde{k}^m}, \Lambda_m \in \mathbb{R}^{n\times l}\ \forall m \in \mathcal{M}\\
& V = \displaystyle\sum_{m=1}^{M} P_{t_{m+1}-1}^\top \tilde{c}_{\mathrm{c},m} e_1^\top Q_M + P_{t_m}^\top C_{\mathrm{u},t_m}^\top \tilde{U}_m Q_m + P_{t_{m+1}-1}^\top \tilde{C}_{\mathrm{s},m}^\top \tilde{S}_m Q_m\\
& \left.\begin{array}{l}
\tilde{S}_m R_m = \tilde{S}_{m-1} R_{m-1} + \tilde{U}_m R_m\\
G_{\mathrm{u},t_m} \tilde{U}_m R_m = 0\\
\Lambda_m W = \tilde{S}_m R_m\\
\Lambda_m h \geq 0\\
\Lambda_m \geq 0
\end{array}\right\} \forall m \in \mathcal{M}
\end{array} \qquad (4.4)
$$

The size of (4.4) is polynomial in $\tilde{k}$, l, M, n and p. Under the reasonable assumption that $\tilde{k}$, l, n and p are of the order $O(M)$ in realistic problem instances, the size of problem (4.4) grows only polynomially with M. Thus, it can be efficiently solved with standard quadratic programming solvers.

Furthermore, (4.4) only requires information about the support $\tilde{Z}$ of the risk factors $\tilde{\zeta}$ and the first four moments of the uncertain parameters ξ. Since the full joint distribution of ξ is rarely available, this is an attractive feature of the model. Moreover, the user is free to compute the moments and the support applying his or her favourite estimation technique.

5 Numerical Example

To validate the outlined mean-variance model and the underlying approximations, we present the results of a large number of experiments based on the following scenario. A price-taking Scandinavian retailer must meet the electricity demand of its customers over a planning horizon of 28 days, split into daily intervals, indexed by $t \in \mathcal{T} := \{1, \ldots, 28\}$. In the electricity markets, three baseload forward contracts, indexed by $i \in \mathcal{I} := \{1, 2, 3\}$, with delivery rate of 1 MW are tradable. Their delivery periods start at the beginning of days 2, 11, and 20 and terminate at the end of days 10, 19, and 28, respectively. Each of these forward contracts covers a delivery period of 9 days and has, therefore, a volume of 216 MWh. These baseload contracts serve as underlying instruments for one European call option each, indexed by $j \in J = \{1, 2, 3\}$, which has a strike price of 115 NOK/MWh and matures at the beginning of the delivery period of the underlying forward contract. The retailer is assumed to have no initial holdings in forward and call option contracts. All optimisation problems were solved using ILOG CPLEX 11.2, on a Linux workstation with dual 2.66 GHz Intel core processors with 4 GB RAM.

5.1 Uncertainty Modelling

As the true moments of the uncertain parameters are unknown, they have to be estimated from historical data. Since most electricity markets are relatively immature, long histories of liquid spot and derivatives prices do not exist. Hence, there is a lack of sufficient data for estimating stable multiperiod moments based exclusively on historical data (i.e., estimation errors might be large), especially if the planning horizon covers several periods. Therefore, we estimate a parametric model for the electricity prices and the demand, from which we estimate the support $\tilde{Z}$ and obtain the moments via sampling.

Uncertain Parameters: We assume that the electricity spot price and the electricity demand are the explanatory risk factors in each period $t \in \mathcal{T}$, i.e., $\zeta_t = (S_t, D_t)^\top$. Electricity derivative prices and payoffs are representable as functions of the spot prices. Since the trading of a derivative contract ceases at the start of its delivery period, the dimension of ξ_t is non-increasing in t. For example, at day $t = 2$ we set

$$\xi_2 = (S_2, D_2, F_2^1, F_2^2, F_2^3, \max(F_2^1 - X^1; 0), C_2^2, C_2^3, S_2 D_2)^\top$$

whereas at day $t = 3$, because forwards and options of type 1 are no longer traded,

$$\xi_3 = (S_3, D_3, F_3^2, F_3^3, C_3^2, C_3^3, S_3 D_3)^\top.$$

Spot Price Modelling: The unique characteristics of electricity, such as its limited storability, grid-bound nature and inelastic demand, distinguish it from other commodities and financial assets [29]. Thus, electricity prices do not follow martingale processes but exhibit seasonality, mean-reversion, stochastic volatility as well as spikes. Following Lucia and Schwartz [25], we assume that the logarithm of the spot price can be described by an Ornstein-Uhlenbeck process with seasonality.

$$
\begin{aligned}
S(\tau) &= \exp(f(\tau) + X(\tau)) \\
\mathrm{d}X(\tau) &= -\alpha^s X(\tau)\mathrm{d}\tau + \sigma^s \mathrm{d}W(\tau),
\end{aligned} \tag{5.1}
$$

where $\alpha^s > 0$, and $W(\tau)$ is a standard Brownian motion process. The seasonal component $f(\tau)$ is considered to be completely predictable. The deseasonalised component of the logarithm of the spot price $X(\tau)$ follows a mean-reverting process with constant mean-reversion rate α^s, zero long-run mean and constant volatility σ^s.

Derivative Pricing: Hedging derivative contracts with the underlying asset or commodity requires the ability to store the underlying. However, electricity cannot be efficiently stored. Thus, traditional storage-based no-arbitrage methods for valuing derivatives cannot be directly applied. Nonetheless, based on standard arbitrage arguments with two derivative assets it is possible to find a risk neutral probability measure $\mathbb{Q}$, under which the current value of any derivative asset is equal to the discounted expected value of its future payoffs [5]. It has been shown in [25] that the process $X(\tau)$ obeys the stochastic differential equation

$$
\mathrm{d}X(\tau) = \alpha^s(\tilde{\mu}^s - X(\tau))\mathrm{d}\tau + \sigma^s \mathrm{d}\tilde{W}(\tau), \tag{5.2}
$$

where λ denotes the market price of risk, $\tilde{\mu}^s := -\lambda\sigma^s/\alpha^s$, and $\tilde{W}(\tau) := W(\tau) + \lambda\tau$ is a standard Brownian motion under $\mathbb{Q}$. For the sake of analytical tractability, λ is assumed to be constant.

Forward Price: The forward price at time τ for the delivery of 1 MWh at time $\tau' \geq \tau$ is chosen in such a way that the contract is worthless at time τ. By solving the stochastic differential equation (5.2), it can be shown that this instantaneous-delivery forward price is given by

$$
\begin{aligned}
\mathbb{E}_\tau^{\mathbb{Q}}[S(\tau')] = \exp\Big(& f(\tau') + \big[\ln(S(\tau)) - f(\tau)\big] e^{-\alpha^s(\tau'-\tau)} + \tilde{\mu}^s(1 - e^{-\alpha^s(\tau'-\tau)}) \\
& + \frac{\sigma^{s2}}{4\alpha^s}(1 - e^{-2\alpha^s(\tau'-\tau)})\Big).
\end{aligned} \tag{5.3}
$$

If delivery spans a finite interval, the price of a zero-cost forward depends on the settlement specification. As we assume that settlement takes place at the end of the delivery period, the price of a forward contract with a finite delivery period is equal to the arithmetic average of the instantaneous-delivery forward prices in the delivery period.

European Call Option Premium: To determine the premium of a European call option at time τ, the risk-neutral distribution of the underlying forward price at the maturity time of the option is required. By applying Itô's Lemma to (5.3) and using (5.2), it can be shown that instantaneous-delivery forward prices are lognormally distributed under $\mathbb{Q}$. Thus, the risk-neutral distribution of the price of a forward contract with a finite delivery period at the maturity time of the corresponding option is an arithmetic average of lognormal distributions. Although this distribution does not possess an analytical representation, it can be reasonably approximated by a lognormal distribution. We therefore calculate its first two moments and subsequently fit a lognormal distribution to these moments [35]. This approximation allows us to price European options on electricity forwards via the Black Scholes formula [6].

Electricity Demand Modelling: The electricity demand is modelled in a similar fashion as the stochastic spot price since it typically exhibits mean reversion and seasonality [29]. We assume that the retailer's demand of electricity evolves according to

$$\begin{array}{rcl} D(\tau) & = & \exp(g(\tau) + Y(\tau)) \\ \mathrm{d}Y(\tau) & = & -\alpha^d Y(\tau)\mathrm{d}\tau + \sigma^d \mathrm{d}W^d(\tau), \end{array} \tag{5.4}$$

where $\alpha^d > 0$, $g(\tau)$ is the seasonal component, and $W^d(\tau)$ is a standard Brownian motion process, which is independent of $W(\tau)$. Thus, $Y(\tau)$ follows a stationary mean-reverting process with a zero long-run mean and a speed of adjustment α^d. Notice that the electricity demand and the spot price are independent as a consequence of the independence of $W^d(\tau)$ and $W(\tau)$. This is justified by the inelasticity of the demand to the spot price and the retailer being a price-taker. Moreover, empirical studies show that the correlation between the spot price and the demand is weak in electricity markets [27].

Moment Estimation: For each $t \in \mathcal{T}$, we set the daily average spot price $S_t := S((t-1)\Delta)$ and the electricity demand $D_t := D((t-1)\Delta)$, with seasonal components

$$\begin{array}{rcl} f(t) & = & c^s + \beta^s \text{ workday}_\text{t} + \delta^s \cos\left((t + \omega^s)\frac{2\pi}{365}\right), \\ g(t) & = & c^d + \beta^d \text{ workday}_\text{t} + \delta^d \cos\left((t + \omega^d)\frac{2\pi}{365}\right), \end{array}$$

respectively. To estimate the moments of the random parameters, we generated sample trajectories of the electricity spot price and demand by explicitly solving (5.1) and (5.4), respectively. In addition, we calculated for each sample the corresponding trajectories of the remaining random parameters as afore-described. The estimates of the moments of ξ were then obtained via Monte Carlo sampling. The parameters used in our numerical example are displayed in Table 1. For electricity spot and derivative prices, we adopt the parameters estimated in [25] based on daily data from the Nord Pool market. In addition, we assume that $S_1 = 110$ NOK/MWh and $D_1 = 4000$ MWh.

	c	β	δ	ω	α	σ	λ
Spot price	4.867	-0.09	0.306	0.836	0.016	0.086	0.033
Demand	8.48364	-0.09978	0.27601	0.836	0.07	0.06	

Table 1: Parameters of numerical example (where time is measured in days)

Support Estimation: From the explicit solution of (5.1) and (5.4) we have that S_t and D_t follow lognormal distributions under the real world probability measure $\mathbb{P}$ and are thus supported on $[0, \infty)$. However, the LDR approximation may be weak if the support of the uncertain parameters is unbounded. In extreme cases, some LDRs can be forced to become constant in order to obey the constraints on the whole support. One way to overcome this problem would be to employ, e.g., piecewise linear decision rules [9]. However, the tractability of the optimisation model deteriorates with the use of more complex decision rules. Thus, we choose to adhere to LDRs but to work with a truncated support that covers most of the mass of the original probability distribution. We assume the support $\tilde{Z}$ to be the box uncertainty set defined from 99.9% marginal confidence intervals of S_t and D_t at $t \in \tilde{\mathcal{T}}$. We remark that the truncation of the support has a negligible impact on the moments.

Sample Size: Based on the estimated moments and support, an approximation of (4.4) is obtained by replacing the real inputs with their estimates. Solving the problem for 100 different independent sample sets, we find that a sample size of 100,000 is sufficient to guarantee a 1.8% accuracy with a confidence level of approximately 99%.

5.2 Sensitivity Analysis

Unless otherwise indicated, a pure risk minimisation framework ($\gamma = 1$) is adopted in this section. Moreover, the duration of each macroperiod is assumed to be two days.

To assess the value of adaptivity, we compare the optimal value of (4.4) with the optimal value of the approximate program obtained using constant decision rules (CDR), i.e., decision rules that do not depend on the random data. CDRs are appropriate to model a retailer that precommits to a portfolio strategy at the start of the planning horizon and implements the corresponding decisions irrespective of the future market behaviour. Clearly, these inflexible portfolio strategies are outperformed by LDRs, which can adapt to changing market conditions. Since the class of CDRs is covered by the class of LDRs, the CDR approximation constitutes an upper bound to (4.4).

5.2.1 Efficient Frontier

Solving the quadratic program (4.4) for different values of the risk aversion coefficient γ yields a parametric family of optimal portfolio strategies. Plotting the expected value against the standard deviation of the corresponding overall costs for each $\gamma \in [0, 1]$ generates an efficient frontier.

The left chart of Figure 1 depicts two approximate efficient frontiers obtained from the LDR and the CDR approximations, each one based on 20 different values of γ in the range $[5 \times 10^{-8}, 1]$. For the same expected overall cost, the risk of the LDR solution is lower than the risk of the CDR solution. This confirms our intuition that incorporating adaptivity into the decision model is beneficial, in particular when the decision maker is risk-averse ($\gamma > 0$).

For $\gamma = 0$, the expected cost minimisation problem can be solved analytically. A particular forward contract is bought if and only if its cost is smaller than the expected cost (with respect to $\mathbb{P}$) of purchasing electric energy with the same load profile in the spot market during the delivery period of the forward contract. Similarly, to determine the optimal positions in the call options, the retailer compares the option premium with the expected payoff of the option at maturity, under the probability measure $\mathbb{P}$. If there is no risk premium ($\lambda = 0$), then both alternatives are equally expensive. The retailer is then indifferent between purchasing forward contracts or buying electric energy in the spot market at the time of delivery, as well as being indifferent

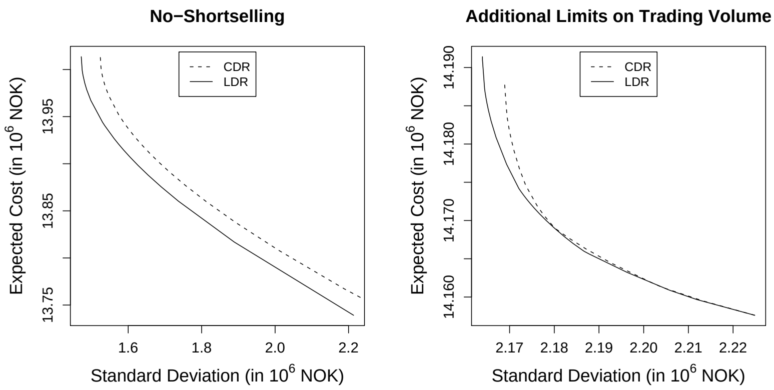

Figure 1: Efficient frontier

between purchasing call options on forward contracts or not. When the electricity market is in contango ($\lambda < 0$) the retailer must pay a risk premium to the suppliers for purchasing forward contracts. In this case, the forward contracts are more expensive, in expectation, than buying electric energy with the same load profile in the spot market during their delivery period, so a risk-neutral retailer will not buy any forwards. Similarly, the expected payoff at maturity falls short of the option premium, and, consequently, the retailer will refrain from purchasing any call options. During backwardation ($\lambda > 0$) a risk-neutral retailer prefers to buy forwards since, in expectation, they are cheaper than purchasing electric energy in the spot market at the time of delivery. Likewise, the retailer opts to acquire call options since the expected payoff at maturity exceeds the corresponding premia. Ideally, the retailer would buy as many forwards and call options as possible at each macroperiod and later sell the provided energy in the spot market. If no limits on the trading volume of forwards and options are imposed, the retailer can achieve an infinite expected profit through this strategy. In this case, problem (4.4) becomes unbounded — an effect that has been confirmed in our numerical experiments.

In conclusion, for $\gamma = 0$ the optimal decisions depend solely on the sign of the market price of risk and can be precommitted at the beginning of the planning horizon. The revelation of new information at later stages will provide no incentive to revise the original decisions. Therefore, no value is added to the decision process through the use of adaptive decision rules. To illustrate this point, we solve problem (4.4) repeatedly for $\gamma \in [0, 1]$, subject to additional portfolio constraints that limit the trading volume of derivative contracts[1] to avoid unboundedness of the optimisation problem. The resulting efficient frontier together with the corresponding CDR frontier are shown in the right chart of Figure 1. For $\gamma = 0$, the optimal solutions of the two approximations coincide. However, as risk aversion increases, the value of adaptivity, that is, the benefit from using LDRs increases, and it is highest when the sole objective is to minimise the risk.

[1]We limit the trading volume of derivatives of any given type to 50 in the first macroperiod, and we require the positions to change by less than 20% over each subsequent macroperiod.

5.2.2 Volatility

Figures 2 and 3 show the impact of the spot price and the demand volatility on the optimal objective value, respectively. If the price volatility is zero, the retailer can anticipate the prices of spot and electricity derivatives over the entire planning horizon. Hence, the electricity demands are the only uncertain parameters in the portfolio optimisation problem. Under these circumstances, it does not matter how these demands are satisfied if the aim is to minimise the overall risk. If prices are volatile, then rebalancing the hedging portfolio at later periods in light of new information on the risk factors should lead to an increased performance. The higher the volatility σ^s, the higher the uncertainty and the more substantial the benefit from using LDRs instead of rigid CDRs that cannot adapt to new information, see Figure 2. Moreover, we observe that for higher levels of σ^s there is a considerable gain from employing LDRs.

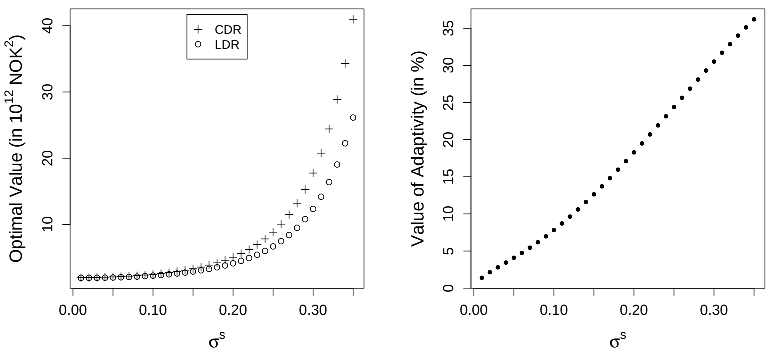

Figure 2: Impact of spot price volatility

If the electricity demand over the planning horizon is deterministic, there is almost no advantage of using LDRs instead of CDRs. In this case, the retailer can purchase forwards in the first period (in which prices are deterministic and thus exhibit no variance) to cover the known future demands, thereby substantially reducing uncertainty. This effect is most prominent if instantaneous-delivery forwards are available in the market, or if the demand is constant over the delivery period. As volatility σ^d increases, the variance of the overall costs rises, and the benefit of using decision rules that allow for adjustments in the portfolio in response to new information increases. However, the relative outperformance of the LDRs with respect to the CDRs slowly decreases as σ^d increases since new information becomes less important for predicting future variances.

5.2.3 Mean Reversion Rate

Figures 4 and 5 depict the impact of the mean reversion rates of the spot price and the demand on the optimal value of problem (4.4), respectively. As the speed of adjustment increases, spot prices revert faster to their mean level, and, therefore, their variance decreases, explaining the

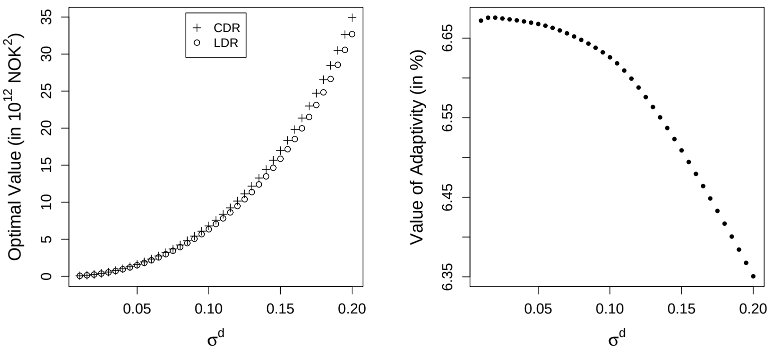

Figure 3: Impact of demand volatility

decline in the variance of the overall costs. As α^s increases, current prices play a less significant role in explaining future expected prices and their variance. Consequently, the benefit of using adaptive decision rules decreases with α^s. If the mean reversion rate tends to infinity, spot prices revert instantaneously to their long-term mean level, and thus become deterministic. Then, the only random parameters remaining in the optimisation problem are the demands. Finding an optimal trading policy that minimises the overall variance becomes redundant, since spot and derivative prices over the whole planning horizon are known with certainty.

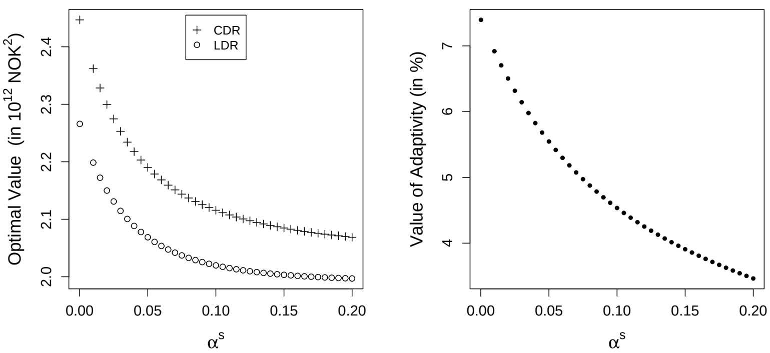

Figure 4: Impact of mean reversion in spot price

Similar effects are observed when we vary the mean reversion rate of the demand. As α^d increases, the electricity demand reverts faster to its mean level, so its variance and, consequently, the risk decline. An increased speed of mean reversion renders future expected demands and their variance less dependent on current and past loads. Hence, the benefits of rebalancing the portfolio in response to new information on electricity demand is smaller. As α^d tends to infinity, the retailer's demand reverts instantaneously to its equilibrium level, and so is, de facto, deterministic. Under these circumstances, the gain from using LDRs instead of CDRs becomes (practically) non-existent.

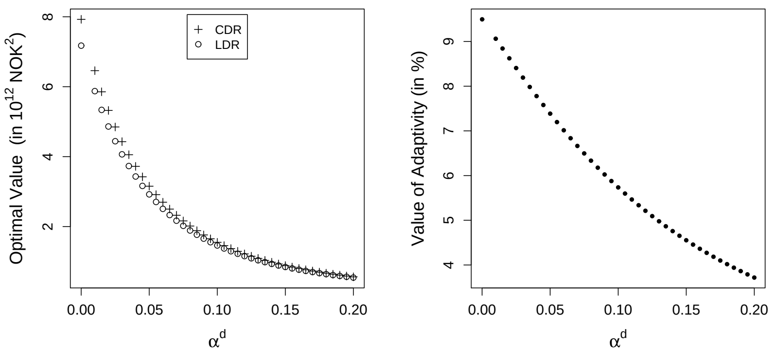

Figure 5: Impact of mean reversion in demand

5.2.4 Market Price of Risk

Figure 6 shows the impact of the market price of risk on the optimal objective value. A change in λ does not impact the optimal objective value when CDRs are employed, but it has a major impact when LDRs are used. The optimal objective value is lower for larger positive or negative market prices of risk. In those cases, the gain from applying LDRs instead of CDRs can be substantial, see Figure 6. The stochastic program takes into account the discrepancy between the cost of each forward contract and the corresponding expected cost of the same volume of electric energy (with the same load profile) in the spot market during the delivery period of the forward contract. Similarly, it considers the disparity between the premium of each call option and the corresponding expected payoff at maturity. Consequently, the differences between the (co)variances of both alternatives are taken into consideration. For higher (positive or negative) market prices of risk these discrepancies will be larger, making the possibility of revising decisions at later stages to reflect new information even more relevant. Note that differences between the (co)variances exist even if no risk premium is required. For example, spot market transactions occur after their respective forward transactions, so their variance is larger when $\lambda = 0$.

5.2.5 Number of Macroperiods

Figure 7 visualises the optimal value of problem (4.4) as a function of the number of macroperiods. We observe a near-monotonic convergence from above as the number of decision stages increases. This behaviour is consistent with the fact that the stage-aggregation discussed in Section 4.1 provides an upper bound on the optimal objective value. The saturation of the optimal objective value supports our hypothesis that the approximation is accurate. Furthermore,

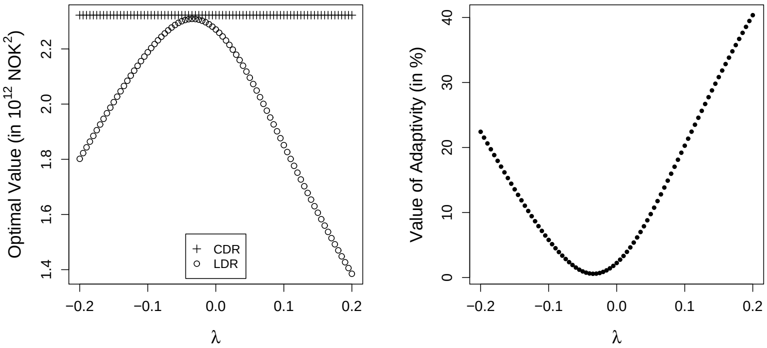

Figure 6: Impact of market price of risk

the near-monotonic behaviour reflects our intuition that the approximate portfolio optimisation problem provides an increasingly accurate upper bound on the optimal value of the original problem as the number of macroperiods approaches the number of normal periods of the original problem.

Figure 7 indicates that an approximation based on 14 macroperiods is reasonably accurate since the relative improvement in the optimal objective value from adding further decision stages is close to zero. In fact, the optimal value of this approximation overestimates the optimal value with 28 periods by merely 0.6%. However, it reduces the solution time dramatically: the runtime of the original problem is approximately 11.5 seconds, whereas the solution time of the approximated problem with 14 decision stages lies below 1 second. As can be seen from Figure 7, precision may be improved by increasing the number of decision stages at the expense of additional runtime.

Notice that the optimal objective value of the CDR approximation barely changes as the number of effective decision stages increases. Since decisions are fixed at the beginning of the planning horizon, an increased number of effective decision stages will not lead to a noticeable improvement. Consequently, as the number of decision stages increases, so does the benefit from using the LDR as opposed to the CDR approximation.

5.3 Comparison with Sample Average Approximation

The standard approach to solve problems of type $\mathcal{ASP}$ numerically is to discretise the underlying probability space. The process of selecting a discrete probability distribution that approximates the true distribution of the risk factors well is known as scenario generation. For a survey of scenario generation techniques, we refer to [20]. In order to assess the accuracy and the scalability of the decision rule approach advocated in this paper, we compare it to a SAA approach that replaces the true distribution of the risk factors with a discrete scenario tree constructed via conditional sampling [30]. Note that, since scenario trees branch when new

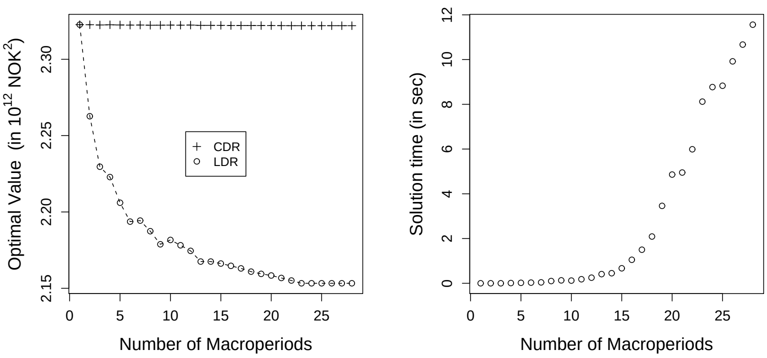

Figure 7: Impact of number of macroperiods

information is observed, the resulting tree ramifies only at the start of each macroperiod $m \in \mathcal{M}$. The number of branches emanating from each tree node (that is, the branching factor) is kept constant throughout the tree, and we assign equal conditional probabilities to each branch. While one may argue that a more refined scenario generation technique should be employed, for the purpose of comparing the scalability of the LDR approximation with that of the scenario generation approach the SAA technique suffices.

In Table 2 we compare the LDR with the SAA approximation for different choices of the branching factor and the number M of macroperiods. Due to run time restrictions associated to the SAA problems, M is limited to a maximum of 10, while the branching factor is fixed to 2 (SAA2), 3 (SAA3), 4 (SAA4), 6 (SAA6), 11 (SAA11), 35 (SAA35) and 1200 (SAA1200) branches per node. Each SAA problem is solved for 20 statistically independent scenario trees. Table 2 displays the average and the coefficient of variation (CV) of the optimal objective values, where the coefficient of variation is defined as the standard deviation expressed as a percentage of the mean. In addition, the average run times (CPU) are reported in seconds. Missing entries (n/a) indicate that the corresponding approximate problems could not be solved in less than one day.

Since SAA problems grow exponentially with M, the branching factor of problems with more than a few macroperiods must be small enough to guarantee that the corresponding stochastic program can be solved in a reasonable time. However, if the number of scenarios is not sufficiently large, the associated tree may not approximate the true probability distribution reliably. Table 2 shows that, for SAA problems with a small branching factor, the dispersion of the optimal objective values around their mean is very high, indicating that these problems provide poor approximations for $\mathcal{ASP}$. Focusing on the case $M = 1$, the SAA estimator for the optimal objective value achieves a reasonable degree of precision, provided the sample size is very large, say 1200. In the SAA1200 case, it is clear that the SAA method is consistent with the LDR approximation. Conversely, for a low branching factor, the SAA estimator with $M = 1$ exhibits a very low degree of precision and accuracy, both of which improve as the branching factor rises. These findings are in line with the statistical properties of the SAA estimator for the optimal objective value (see, e.g., [32, Chapter 5]). This estimator is known to be downwards biased, providing a valid statistical lower bound to the true problem. Furthermore, its bias converges to zero as the sample size tends to infinity [36]. In general, we observe that the variability of the optimal value estimates diminishes and the average objective value increases as the branching factor increases. Therefore, we conjecture that, except for SAA problems with a very high branching factor, the SAA estimators with $M > 1$ are severely downwards biased – a statement substantiated by the uncertainty surrounding the estimates of the optimal objective value. Shapiro [31] reports that typically the bias and dispersion of the SAA optimal value estimator grow fast with M, rendering the corresponding statistical lower bounds inaccurate already for a small number of decision stages.

Nonetheless, Table 2 reveals the heavy computational burden of solving larger SAA problems. As the number of decision stages or the branching factor increases, the run time of the SAA problems can rise substantially. Comparing the average run times of both methods, it is evident that the LDR method exhibits superior scalability. While the LDR problem with $M = 28$ can be solved in under 12 seconds, a solution to the corresponding SAA problem with $M > 10$ could not be located in less than a day, even for a branching factor as low as 2. Moreover, to achieve an adequate degree of accuracy, a prohibitive number of scenarios is required for $M > 1$, leading to SAA problems that could not be solved in less than a day.

6 Conclusions

In this article, we examine a multistage mean-variance portfolio optimisation model for an electricity retailer. To convert the exact model into a tractable quadratic program, we perform two approximations: we aggregate periods into macroperiods, and we restrict the decision rules to those affine in the history of the risk factors. The resulting approximate problem provides an upper bound on the optimal value of the exact problem. Since the size of the approximate problem grows only polynomially with the number of macroperiods, it is amenable to efficient solution. Moreover, the probability distribution of the random parameters affects this problem only through its first four moments and through the support of the risk factors.

Our numerical experiments support our expectation that the approximation based on stage-aggregation is accurate. Moreover, they illustrate the potential for significantly reducing the solution time without sacrificing much precision. Our tests indicate that incorporating adaptivity in the form of LDRs into the portfolio optimisation model is beneficial, especially in a risk minimisation framework. Adaptivity appears to be particularly valuable in the presence of high spot price volatility or large (positive or negative) market prices of risk.

With the aim of evaluating the accuracy and scalability of the LDR method, we compared it with a SAA approximation. Due to the severe bias and dispersion of the SAA optimal value estimators, no meaningful comparison between these and the corresponding LDR estimators could be established. Nonetheless, our numerical tests highlight the heavy computational burden of solving SAA problems with many periods and the superiority of the LDR approach in enabling scalability to multistage models.

Although LDRs are very effective at conferring tractability to multistage models, they may lead to a non-negligible loss of accuracy. A method for systematically estimating this loss has been recently proposed in [24]. The method consists of applying LDRs both to the primal and a dual version of the exact stochastic program, and determining the gap between the optimal values of the corresponding tractable programs. Future research will be concerned with determining the optimality gap of the mean-variance portfolio optimisation model.

Acknowledgements: We thank Fundação para a Ciência e a Tecnologia and EPSRC (under grant EP/H0204554/1) for financial support.

References

[1] A. Atamturk and M. Zhang. Two-stage robust network flow and design under demand uncertainty. *Operations Research*, 55(4):662–673, 2007.

[2] A. Ben-Tal, B. Golany, A. Nemirovski, and J.P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.

[3] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[4] J.R. Birge and F. Louveaux. *Introduction to stochastic programming.* Springer Verlag, 1997.

[5] T. Björk. *Arbitrage theory in continuous time.* Oxford University Press, USA, 2004.

[6] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

[7] G.C. Calafiore. Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473, 2008.

[8] G.C. Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM Journal on Control and Optimization*, 48(4):2254–2274, 2009.

[9] X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357, 2008.

[10] S.J. Deng and S.S. Oren. Electricity derivatives and risk management. *Energy*, 31(6-7):940–953, 2006.

[11] J. Dupačová, G. Consigli, and S.W. Wallace. Scenarios for multistage stochastic programs. *Annals of Operations Research*, 100(1):25–53, 2000.

[12] M. Dyer and L. Stougie. Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3):423–432, 2006.

[13] A. Eichhorn and W. Römisch. Mean-risk optimization models for electricity portfolio management. In *Proceedings of PMAPS*, 2006.

[14] S.E. Fleten, S.W. Wallace, and W.T. Ziemba. Hedging electricity portfolios via stochastic programming. *IMA Volumes in Mathematics and its Applications*, 128:71–94, 2002.

[15] P.J. Goulart, E.C. Kerrigan, and D. Ralph. Efficient robust optimization for robust control with constraints. *Mathematical Programming*, 114(1):115–147, 2008.

[16] A.R. Hatami, H. Seifi, and M.K. Sheikh-El-Eslami. Optimal selling price and energy procurement strategies for a retailer in an electricity market. *Electric Power Systems Research*, 79(1):246–254, 2009.

[17] H. Heitsch and W. Römisch. Scenario reduction algorithms in stochastic programming. *Computational Optimization and Applications*, 24(2):187–206, 2003.

[18] R. Hochreiter, G. Pflug, and D. Wozabal. *Multi-stage stochastic electricity portfolio optimization in liberalized energy markets*, volume 199 of *IFIP International Federation for Information Processing*, pages 219–226. Springer, 2006.

[19] P. Kall and S.W. Wallace. *Stochastic programming*. Wiley, 1994.

[20] M. Kaut and S.W. Wallace. Evaluation of scenario generation methods for stochastic programming. *Pacific Journal of Optimization*, 3(2):257–271, 2007.

[21] J. Kettunen, A. Salo, and D.W. Bunn. Optimization of Electricity Retailer's Contract Portfolio Subject to Risk Preferences. *IEEE Transactions on Power Systems*, 25(1):117–128, 2010.

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

[23] D. Kuhn. Aggregation and discretization in multistage stochastic programming. *Mathematical Programming*, 113(1):61–94, 2008.

[24] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 2009. In press.

[25] J.J. Lucia and E.S. Schwartz. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5(1):5–50, 2002.

[26] H. Markovitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.

[27] T.D. Mount, Y. Ning, and X. Cai. Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, 28(1):62–80, 2006.

[28] T. Pennanen. Epi-Convergent Discretizations of Multistage Stochastic Programs. *Mathematics of Operations Research*, 30(1):245–256, 2005.

[29] D. Pilipovic. *Energy risk: Valuing and managing energy derivatives.* McGraw-Hill, 2007.

[30] A. Shapiro. Inference of statistical bounds for multistage stochastic programming problems. *Mathematical Methods of Operations Research*, 58(1):57–68, 2003.

[31] A. Shapiro. On complexity of multistage stochastic programs. *Operations Research Letters*, 34(1):1–8, 2006.

[32] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory.* SIAM, 2009.

[33] A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. *Applied Optimization*, 99:111–146, 2005.

[34] M.C. Steinbach. Markowitz revisited: Mean-variance models in financial portfolio analysis. *SIAM Review*, 43(1):31–85, 2001.

[35] S.M. Turnbull and L.M. Wakeman. A quick algorithm for pricing European average options. *Journal of Financial and Quantitative Analysis*, 26(3):377–389, 1991.

[36] R. Wets. Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *The Annals of Statistics*, 16(4):1517–1549, 1988.

[37] W.T. Ziemba and J.M. Mulvey. *Worldwide asset and liability modeling.* Cambridge University Press, 1998.

	M	1	2	3	4	5	6	7	8	9	10
LDR	Mean (10^{12} NOK^2)	230.61	224.67	221.55	220.95	219.23	218.10	218.14	217.41	216.87	216.99
	CV (%)	0.79	0.74	0.77	0.73	0.76	0.72	0.74	0.75	0.76	0.74
	CPU (sec)	0.00	0.00	0.00	0.01	0.02	0.03	0.04	0.10	0.13	0.12
SAA2	Mean (10^{12} NOK^2)	1.31	1.55	1.98	11.61	8.51	4.67	12.42	11.65	7.55	14.90
	CV (%)	447.21	329.18	165.15	193.23	98.35	53.36	44.98	24.28	26.96	31.59
	CPU (sec)	0.00	0.00	0.00	0.01	0.03	0.20	2.37	34.16	515.61	8065.25
SAA3	Mean (10^{12} NOK^2)	5.54	24.76	11.43	28.99	31.48	23.90	n/a	n/a	n/a	n/a
	CV (%)	409.83	144.51	79.42	67.98	43.79	38.17	n/a	n/a	n/a	n/a
	CPU (sec)	0.00	0.00	0.03	0.85	57.15	4651.30	n/a	n/a	n/a	n/a
SAA4	Mean (10^{12} NOK^2)	15.04	27.20	22.77	40.56	59.22	n/a	n/a	n/a	n/a	n/a
	CV (%)	210.80	80.39	90.33	30.50	27.46	n/a	n/a	n/a	n/a	n/a
	CPU (sec)	0.00	0.01	0.35	68.57	18233.15	n/a	n/a	n/a	n/a	n/a
SAA6	Mean (10^{12} NOK^2)	38.34	62.78	56.70	79.87	n/a	n/a	n/a	n/a	n/a	n/a
	CV (%)	134.35	74.05	51.52	30.47	n/a	n/a	n/a	n/a	n/a	n/a
	CPU (sec)	0.00	0.03	15.27	18432.47	n/a	n/a	n/a	n/a	n/a	n/a
SAA11	Mean (10^{12} NOK^2)	118.92	106.58	117.48	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CV (%)	61.89	39.43	38.33	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CPU (sec)	0.00	1.62	21355.83	n/a	n/a	n/a	n/a	n/a	n/a	n/a
SAA35	Mean (10^{12} NOK^2)	167.53	178.02	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CV (%)	30.08	19.41	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CPU (sec)	0.04	15392.05	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
SAA1200	Mean (10^{12} NOK^2)	230.45	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CV (%)	3.36	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	CPU (sec)	14067.26	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a

Table 2: LDR vs SAA: Impact of number of macroperiods on optimal objective value