universität wien

DISSERTATION

Titel der Dissertation

„Behavioral pricing of energy swing options:

A tailored solution algorithm for a multi-stage stochastic bilevel program with applications in energy markets“

Verfasser

DI Peter Gross

angestrebter akademischer Grad

Doctor of Philosophy (PhD)

Wien, 2014

Studienkennzahl lt. Studienblatt:	A 094136
Dissertationsgebiet lt. Studienblatt:	Statistik und Operations Research
Betreuerin / Betreuer:	o. Univ.-Prof. Mag. Dr. Georg Pflug

Contents

1. Energy derivatives **7**

1.1. Preliminaries . . . 7
1.1.1. Energy portfolio modeling from a financial point of view . . . 7
1.1.2. Risk factors . . . 8
1.1.3. Notation . . . 8
1.2. Valuation and pricing . . . 10
1.2.1. Linear case: Risk neutral valuation by expectation . . . 10
1.2.2. Nonlinear case: Risk averse valuation and acceptability functionals . . . 13
1.2.3. Pricing of rigid contracts . . . 14
1.2.4. Long flexible contracts: buyer's problem . . . 15
1.2.5. Short flexible contracts: seller's problem . . . 16
1.3. The behavioral pricing problem . . . 18
1.3.1. Seller's problem: Indifference pricing variant . . . 18
1.3.2. Seller's problem: Risk neutral variant . . . 18

2. Bilevel problems **21**

2.1. Introduction . . . 21
2.1.1. Origin . . . 21
2.1.2. Further applications . . . 22
2.2. Basic properties of bilevel programs . . . 23
2.2.1. Optimistic and pessimistic formulations . . . 23
2.2.2. Geometric illustration . . . 24
2.3. Optimistic bilevel programs: Single level reformulations . . . 26
2.3.1. Single level reformulation using optimality conditions: MPECs 27
2.3.2. Optimality conditions of the KKT-reformulation under Guignard CQ . . . 29
2.3.3. Value function reformulation . . . 30
2.3.4. Optimality conditions for the value-function-reformulation under partial calmness . . . 30
2.3.5. Optimality conditions of the KKT-reformulation under partial calmness . . . 32
2.4. Solution methods . . . 33
2.4.1. Computational complexity . . . 33
2.4.2. MPEC reformulations . . . 34
2.4.3. Descent methods . . . 35

2.4.4. Parametrization of lower level solution 36
2.4.5. Techniques from semi infinite optimization 36

3. Problem analysis and tailored algorithm 39
3.1. The behavioral pricing problem: Properties 39
3.1.1. Assumptions on [UL] and [LL] 40
3.1.2. The problem for real data 45
3.2. Properties of the [UL] problem 47
3.2.1. Properties of the [UL] feasible region 48
3.2.2. Implicit treatment of [UL] portfolio optimization 49
3.2.3. Properties of the [UL] acceptability $\Phi(k)$ 49
3.2.4. Approximation of the [UL] acceptability $\Phi(k)$ 52
3.3. A Newton type algorithm exploiting the [UL] problem structure . . . 56
3.4. Replacing $\hat{\mathbf{\Phi}}_{\mathbf{i}}(\mathbf{k})$ by a linear model 59

4. Application in the natural gas market 63
4.1. Modeling of a natural gas portfolio 63
4.1.1. Modeling details 64
4.2. Stochastic modeling of risk factors 67
4.2.1. Data 67
4.2.2. A model linking oil and temperature to natural gas prices . . 69
4.2.3. Quality of the hybrid model 72
4.3. Deterministic approximation of the data process 73
4.3.1. Introduction 73
4.3.2. A numerical procedure for tree approximation of stochastic processes 75
4.3.3. Generation of the tree used in Section 4.4 76
4.4. Relative pricing of swing options 79
4.4.1. General model properties 79
4.4.2. Optimal setting of swing-exercise price: different model specifications 80
4.4.3. Conclusion 86

A. Fundamental concepts and technologies used in previous chapters 89
A.1. Quantifying risk and acceptability 89
A.1.1. Risk and acceptability 89
A.1.2. Polyhedral risk/acceptability functionals 90
A.1.3. The average value-at-risk 91
A.2. Basic concepts from convex analysis 93
A.2.1. Convex geometry 93
A.2.2. Geometric characterization of optimality 94
A.2.3. Constraint qualifications 94
A.3. Distances between stochastic processes 96
A.3.1. Distance concepts 96

A.3.2. Lipschitz property of the value function for nested and transportation distance 97

B. Parameters used in Chapter 4 **101**

C. Kurzzusammenfassung **103**

D. Abstract **105**

Bibliography **109**

List of symbols

The following table lists the most important symbols used in the text, in order of their introduction.
Note that indexes and arguments of functions may be omitted whereever clear from the context.

Chapter 1

P_t	spot price at time t
k_t^i	paramter for contract i at time t, in particular exercise price k
y_t^i	exercise strategy for contract i at time t
$S_t^i(k^i, P; y^i)$	cash flow from energy derivative i at time t
$\Pi(y)$	cumulative payoff of a portfolio of derivatives with joint exercise strategy y
$\mathbb{E}$	mathematical expectation
$\mathcal{A}$	acceptability functional
$z_t \lhd \mathfrak{F}_t$	z_t measurable with respect to sigma algebra $\mathfrak{F}_t$
$y^\star(x)$	optimal response of [LL]
[LL]	lower level decision maker in a bilevel problem (buyer)
[UL]	upper level decision maker in a bilevel problem (seller)

Chapter 2

$Y(x)$	[LL] feasible set
$\nabla_x F(x, y)$	gradient of $F(x, y)$ with respect to x
$N_{Y(x)}(y)$	normal cone to the set $Y(x)$ at $y \in Y(x)$
$\varphi(x)$	optimal value of [LL] problem for given [UL] decision x

Chapter 3

$AV@R_\alpha$	average value-at-risk
$X_t(y^\star)$	[UL] feasible set at stage t
y_l, y_u	lower and upper bound on [LL] consumption per stage
y_L, y_U	lower and upper bound on [LL] cumulative consumption
$\mathcal{S}$	set of scenarios
$\mathcal{N}$	set of nodes
$\mathcal{N}_s$	set of nodes belonging to scenario s
p_s	probability of scenario s
P_n,$\tilde{P}_n$	spot prices at node n

Chapter 3 (continued)

$\mathcal{P}_1$	matrix of elements $p_{i,n} = \begin{cases} 1 & \text{if } n \in \mathcal{N}_i \\ 0 & \text{else} \end{cases}$
$\mathcal{P}_2$	matrix of elements $p_{i,n} = \begin{cases} P_n & \text{if } n \in \mathcal{N}_i \\ 0 & \text{else} \end{cases}$
$\bar{x} \in \mathbb{R}^{\lvert\mathcal{N}\rvert}$	vector of net spot market purchase by [UL] per node
m_{depth}	market depth
ρ	threshold for acceptability
$\Phi(k)$	[UL] acceptability level for given strike price k
$\hat{\Phi}_i(k)$	overestimator for [UL] acceptability level, given [UL] and [LL] solution at k_i
$\lambda_\zeta(k_0)$	dual variable for [UL] portfolio optimization for k_0 (payoff constraint)
$\lambda_y(k_0)$	dual variable for [UL] portfolio optimization for k_0 (market depth constraint)
$k^\star$	optimal strike price
$\hat{k}_i$	underestimator of optimal strike price at iteration i
Δ_i	slope of linear model for $\hat{\Phi}_i(k)$

Chapter 4

$\tilde{T}_t$	weighted average of temperatures
$l_t(\tilde{T}_t)$	synthetic natural gas demand
O_t	oil spot price at time t
$\tilde{O}_t$	moving average of oil prices

Preface

Swing options are flexible energy supply contracts between two parties[1] which give one of them the right to receive energy from the other for a fixed price per unit (strike price). The receiving party (called buyer) can choose the amount of delivery at each instance according to some constraints, specified in the contract.

This kind of contract plays a central role in the energy industry since long before the liberalization of the energy markets, because they almost replicate energy generation units (such as power plants or natural gas production sites) and are flexible enough to absorb uncertain demand. Therefore, they can be used quite efficiently as risk management tools.
In spite of the rise of power and gas exchanges, at least in some energy markets these bilateral, nonstandard contracts have kept their attractivity compared to exchange traded standard products. Possible reasons for their popularity are their greater flexibility (they can be tailor-made for the buyer) and the absence of market-liquidity risks (there is no danger that no sellers are in the market when energy is needed).

In the literature, the standard approach to pricing swing options is based on the buyer's point of view. For a fixed strike price, the maximal expected[2] payoff from selling the contract on the spot market is computed. The resulting price consists of a fixed payment equaling the expected payoff of the contract plus the strike price per consumed unit (i.e. depending on the amount of consumption). Thus, this value resembles an objective market price, shared by all market participants which also excludes arbitrage opportunities for the buyer.
Usually pricing of energy swing options stops at this point, i.e. the value of a contract is also used as a price for the contract.

The drawback of this standard approach is that it relies on assumptions on the market (completeness and the no arbitrage principle) which may not be satisfied for energy markets. Relaxing these assumptions, valuations are now subjective i.e. depending on risk preferences and in particular on the agents' portfolios.

An alternative approach applicable in a more general setting was introduced by [PB09] as the so called *behavioral* model. This approach takes a game theoretic point of view: In the corresponding game, the seller takes the role of the "leading" player, the buyer the role of the "following" player. The task of the leader is to set

[1] considered to have access to the market, i.e. they can sell (and buy) energy on short notice at market prices

[2] usually under an equivalent martingale measure

the exercise price for the swing contract and after that, the follower can optimally exercise the contract during its period of validity. The objective of the seller is to set the lowest feasible strike price whereas the buyer maximizes the profits from exercising the contract; the actions of both players are constrained due to contract specifications and the acceptable risk (with respect to their payoffs). Apparently, the optimal decision of the leader is taken under anticipation of its influence on the follower's exercise strategy. Therefore, the optimal decision of the leader is the solution to an optimization problem, which is constrained by the solution of another, dependent optimization problem (decisions of the follower). The class of this kind of hierarchical mathematical optimization problems is called bilevel problem.

To date the behavioral model is one of only few attempts aiming at selecting an optimal *strike price*[3] for energy derivatives. Especially in markets where not one unique price level exists (i.e. generators or importers of energy have access to alternative, often cheaper sources than the market) this pricing approach seems promising.

Previous works on the behavioral pricing model ([PB09, KP13a]) considered a setting where the two players faced different prices to model their different roles in the markets. In this thesis the existing bilevel pricing model is extended by including an explicit model of the contract seller's portfolio. This means, that in addition to the anticipated buyer's exercise strategy, the seller's decision accounts for the own generation and hedging strategy. This allows to model the different decision problems of buyer and seller in a realistic way.
An additional feature of this extended model is that the optimal price is relative to the content of the current portfolio of the seller, since it influences the possibilities and the risks involved in supplying energy. Relative prices can for example be used to analyze portfolio effects (i.e. the benefits from size and contents in terms of a seller's portfolio), a topic which - to our knowledge - has received no attention in the literature yet.

The main objective of this thesis was the development of a solution algorithm tailored to the behavioral pricing model, which is capable of solving a problem of realistic size. The algorithm was implemented in MATLAB for testing and visualization and in AIMMS for a case study.

The scientific contribution is to show that, when a linear portfolio model is used, it is possible to perform behavioral pricing including the seller's portfolio, with little more effort than optimizing the seller's portfolio. This is remarkable since the problem class containing the behavioral pricing problem (stochastic multistage bilevel problems) is widely considered to be too difficult to solve for actual applications. Furthermore, the application of the algorithm for a case study illustrates the natural emergence of relative prices in the indifference pricing framework.
I consider the thesis as applied work, i.e. the development of a problem specific solution algorithm and its application to a realistic problem instance. Beside the design of the solution algorithm, a lot of effort was put in modeling and programming.

[3]usually, for pricing the strike price is assumed to be a known constant

These two aspects are the pillars of the thesis; consequently, related theoretical concepts from the literature (concerning bilevel problems, risk measures and tree generation) are treated in a rather short manner and have been largely shifted to the Appendix A.

The structure of the thesis is as follows:

The first chapter contains an introduction to the mathematical modeling of energy portfolios, highlighting the differences between linear and nonlinear valuation and rigid and flexible contracts. At the end of the first chapter, the bilevel pricing problem is introduced in a natural way as the indifference price of a short contract.
In the second chapter, a small overview of bilevel programming is provided, including illustrations of problem properties, one level reformulations and algorithmic approaches.
The third chapter deals with the properties of the bilevel pricing problem for real data and the tailored solution algorithm.
The topic of the fourth chapter is the application of the algorithm for the pricing of energy swing options. It contains short discussions of the actual modeling of a natural gas portfolio and its risk factors, deterministic approximation of the stochastic optimization problem and some numerical results from the AIMMS model.

At this point I want to acknowledge the support from various sides which made this thesis possible:

First of all, I want to thank Georg Pflug for giving me the opportunity to spend some time in academia, which was a period of great freedom and intellectual excitement for me.

In particular, I am very grateful for being accepted as a student of the Initiativkolleg Computational Optimization. The courses offered by this PhD school covered all areas of modern optimization and were held by top researchers in their fields. Especially for me, as a student of stochastic programming, it was a privilege to be taught by both Roger Wets and Terry Rockafellar.

Finally, I want to thank my colleagues from ISOR and the IK. In particular, my thanks goes to Raimund Kovacevic for his constant support during my first years, Bita Analui for inspiring discussions, and Eric Laas-Nesbitt for being a perfect roommate.

Peter Gross

Vienna, May 2014

1. Energy derivatives

In the following chapter, the notation and our point of view on energy derivatives for analysis purposes will be introduced. Ideas fundamental for the subsequent work will be provided and put in context.

1.1. Preliminaries

1.1.1. Energy portfolio modeling from a financial point of view

In the following section, energy derivatives and energy portfolios will be studied. For this we assume a financial point of view: this simply means that each objective will be formulated, and all risks will be measured in terms of money.

Using terminology from finance, energy derivatives are simply derivatives where the underlying is energy. To be more precise, they are contracts specifying the physical delivery of energy and/or cash transactions which depend on energy prices.
Energy portfolios are collections of energy derivatives or assets for generation, transformation, storage and transport of energy.

For notational convenience, we consider spot market products as the smallest building blocks for energy. This means that we assume the portfolio (generation, delivery, etc.) leads to constant energy flow on intervals (days or hours) for which spot products are traded. Real time or balancing markets are ignored for the sake of simplicity.

Under this assumption, all physical energy flows from the portfolio (generation, consumption, transportation etc.) can be offset by trading on spot markets. Consequently, holding an energy portfolio results in cash flows which depend solely on energy spot prices.
Therefore we will not distinguish between elements of the portfolio which are purely financial, and elements which result in physical delivery of energy. In particular, since all assets can be analyzed by the (net) cash flow generated by quantifying energy flows at spot market prices, we will refer to all elements of an energy portfolio, including assets, as energy derivatives (or as a synonym, contracts).

Although the financial point of view allows us to neglect the difference between "physical" and "financial" energy, there is still one distinction to be made:

- *rigid* contracts deliver a pre-specified amount of energy

- *flexible* contracts allow the holder to decide about the amount of energy to be delivered

The main example for rigid contracts are futures contracts, which specify the delivery of a constant load of energy over a fixed time period for a fixed price. Flexible contracts are (swing) options where the holder can decide about exercise (and the amount to be delivered) on short notice. Notice that in the case of our wording also physical assets such as power plants, or gas storage are flexible energy derivatives.

1.1.2. Risk factors

Future energy spot prices are not known in advance. For mathematical modeling, they will be described by random variables. The cash flows from energy derivatives are influenced by spot prices; in particular, adverse spot price movements lead to losses. Consequently, spot prices are a source of risk (or a risk factor) for energy derivatives.
For cash flows from rigid contracts, spot prices are the only risk factors. In contrast, flexible contracts depend on the exercise behavior of the contract holder in addition to the spot price risk.
Notice the fundamental difference between flexible contracts which have been bought (i.e. long contracts) and which have been sold (short contracts): a long flexible contract generates a cash flow which can be controlled to a certain extent (by the own exercise strategy) whereas a short flexible contract generates a cash flow with an additional risk factor: the behavior of the buyer.

In our setting[1], all risk factors considered are either (spot) prices or the exercise behavior of the holder of a flexible contract.

1.1.3. Notation

At this point we can introduce some notation for the analysis of energy derivatives:

$(S_t^i)_{t\in 1,\dots,T}$ is the sequence of *net payments* S_t^i generated by realizing an energy derivative $i \in \{1,\dots,I\}$ at the spot market during its delivery period $t = 1,\dots,T$, which will be referred to as the *cash flow* generated by the energy derivative. The *payoff* of a derivative or a portfolio of derivatives is its cumulative cash flow, i.e. $S^i = \sum_{t=1}^{T} S_t^i$.
In general, $S_t^i = S_t^i\left(k^i, P, y^i\right)$, i.e. it may depend on *spot prices* $P = (P_t)_{t=1,\dots,T}$, additional *parameters* $k^i = (k_t^i)_{t=1,\dots,T}$, an *exercise strategy* $y^i = (y_t^i)_{t=1,\dots,T}$ (which is fixed in the case of rigid contracts) and a discount factor (which we in general omit for brevity). However, for simplicity of notation we will only explicitly state the

[1]For example, we omit physical breakdowns of infrastructure or changes in feasible decisions due to political developments, etc.

dependencies when necessary. Also, the index i referring to one of several contracts in a portfolio will be omitted if possible.

Under the *linear market model*, the net payments of a flexible contract depend linearly on the exercise strategy, i.e. $S_t(y) = (P_t - k_t)\, y_t$ for some parameters k_t and spot prices P_t. In this model, the exercise behavior does not influence the market prices which means that the holder of the contract is a price taker.

An energy derivative resulting in delivery of energy (or equivalently from the financial point of view, variable payments) will be referred to as *long contract*, otherwise it will be called *short contract*. In the case of flexible derivatives, the holder of the long contract can decide about the exercise (i.e. y_t are decision variables) whereas the holder of a short flexible contract (seller) has to satisfy the demand of the buyer (i.e. y_t are random variables).

Finally, we introduce the variable $\Pi(y)$ for the cumulative payoff of a whole portfolio:

$$\Pi(y) = \sum_{i=1}^{I} \sum_{t=1}^{T} S_t^i(y_t^i)$$

Notice that this random variable depends on the joint exercise policy $y = (y_t^i)$, $i = \{1, \dots I\}$, $t = \{1, \dots, T\}$. For simplicity, this is not reflected in the notation wherever it is clear from the context. Also, the indirect dependencies of $\Pi(y)$ on the spot price process P and parameters k will not be stated explicitly.

Example. Futures contract:
A long monthly futures contract binds its holder to receive a constant load of energy (say natural gas) during a whole month $(1, \dots, T)$. Both price and load-level of delivery are fixed in advance, say at a level 10MWh/h $(= y_t \, \forall t)$ and a price of $F(= k_t \, \forall t)$. For our analysis we map the physical delivery onto the spot market, i.e. we look at the cash flow induced by selling the whole energy at delivery. Using the notation from above, the resulting cash flow is given by $S_t = 10(P_t - F)$.
Notice that this is equivalent to the cash flow from the floating leg of a swap contract which exchanges fixed payment for the varying spot price.

The most prominent example for flexible contracts in energy markets are swing options, which can be constructed from futures by allowing y_t to be chosen (under some constraints) by the buyer on short notice.

Example. Swing option:
A swing option gives the holder the right to receive energy deliveries $y = (y_t)_{t=1,\dots,T}$ at a fixed strike price k which has been specified at contract closure. The amount y_t can be specified by the holder on short notice but is usually subject to constraints

$$y_l \le y_t \le y_u, \;\; y_L \le \sum_{t=1}^{T} y_t \le y_U$$

In the linear market model, the cash flow from a swing option in terms of spot prices P_t is given by the sequence

$$S_t(k; y) = (k - P_t)\, y_t\,.$$

A specific type of swing option widely used in natural gas markets are index-coupled flexible delivery contracts. For this contract type, a time dependent delivery price k_t is generated by coupling the fixed price k to a price-index P_t^f of a competing fuel (often oil distillates). Using a weight κ and an initial index price P_0^f, this leads to formulas such as $k_t = k + \kappa \left(P_t^f - P_0^f\right)$. In this case, the formula for the payoff stays the same, however it now depends on a two dimensional stochastic process (delivery prices k_t and spot prices P_t) and additionally the delivery decisions y_t of the holder.

1.2. Valuation and pricing

The general principle of valuation of a derivative is to assign to it one single real number, which simply makes it possible to compare the "favourableness" of contracts (many characteristics) by comparing the assigned numbers. Notice that these numbers need not to represent monetary units as for example in the case of valuation by utility functionals.
This number can for example be the net present value of the payoff; however more general concepts such as utility can be used. As observed in the previous section, from the financial point of view, the analysis of energy derivatives can be reduced to the analysis of the generated cash flow, which is modeled as a sequence of random variables. Consequently, valuation is the mapping of a sequence of random variables onto the real line by the valuation functional.
The choice of the valuation functional is a fundamental modeling decision. The model of interest in this work will rely on nonlinear, risk averse valuation. To put the choice in context we will also discuss linear (risk neutral) valuation.

Valuation and pricing are connected to each other by the concept of the fair market price, which identifies the price of an asset with its market value. This approach seems obvious as long as objective valuations are achievable.
For risk averse valuation however, prices depend on the existing portfolios (and risk preferences) of agents i.e. no objective market values exist for derivatives. As a result, the concept of the fair price of a contract has to be generalized to the price which makes the seller/buyer indifferent between having and not having the contract. This so called "indifference pricing" approach leads to *relative prices*, which depend on the current portfolio (and the specific valuation functional) of the seller/buyer.

1.2.1. Linear case: Risk neutral valuation by expectation

A very common approach both in theory and application is to valuate random cash flows by the expected value of their sum (often also discounted by interest rates),

which leads to the expected net present value of a cash flow. In the wording introduced above, the expected value of the cumulative cash flow is chosen as valuation functional.

This valuation approach has the advantage of being relatively easy to compute which is mainly due to the linearity of the expected value. This property often results in the possibility to compute separately valuation and optimization of the different portfolio elements:

Fact. *Let a portfolio of I derivatives with net payments $S_t^i(y^i)$ be valuated by its cumulative expected cash flow. If there are no coupling constraints between decision variables, i.e. the constraints on y^i can be written as*

$$f(y^1, \dots, y^I) \leq 0 \Leftrightarrow \begin{cases} f_1(y^1) \leq 0 \\ \vdots \\ f_I(y^I) \leq 0 \end{cases},$$

then in the linear market model, summation and the maximization of the expectation can be interchanged, i.e.

$$\max_{y^1,\dots,y^I} \mathbb{E}\left\{ \sum_{i=1}^{I} \sum_{t=1}^{T} S_t^i\left(y^i\right) | f_i(y_i) \leq 0 \right\} = \sum_{i=1}^{I} \sum_{t=1}^{T} \max_{y^1,\dots,y^I} \mathbb{E}\left\{ S_t^i\left(y^i\right) | f_i(y_i) \leq 0 \right\} \quad (1.1)$$

and

$$\operatorname*{argmax}_{y^1,\dots,y^I} \mathbb{E}\left\{ \sum_{i=1}^{I} \sum_{t=1}^{T} S_t^i\left(y^i\right) | f_i(y_i) \leq 0 \right\} = \left(\operatorname*{argmax}_{y^i} \mathbb{E}\left\{ S_t^i\left(y^i\right) \middle| f_i(y^i) \leq 0 \right\} \right)_{i=1,\dots,I} \quad (1.2)$$

As a consequence, exercise strategies of portfolio elements and valuation of the resulting cash flows can be computed independently from each other which is a significant simplification. Notice that coupling constraints such as for example coupling of production capacities (as in the case of hydro reservoirs) or limited market depth etc. can destroy this structural advantage.
Ultimately, this separation property leads to "objective" valuations: Since valuation of a derivative is independent of the present portfolio of the holder, every market participant will arrive at the same result of valuation (provided they rely on the same assessment of distributions of underlying risk factors). As an immediate consequence, in the linear setting, the fair market price for any derivative is given by its value in cash.

Another characteristic property of using the expected value as valuation functional is its ignorance of risk: losses and gains are treated in exactly the same way. The obvious benefit of this property is that risk preferences of agents do not have to be specified or considered, since they play no role at all.

The absence of risk averseness can be justified by a law of large number if the payoff is generated repeatedly under the same probability distribution. This argument is sometimes used in applications in the setting of short term power production. In particular for the valuation of power plants with high flexibility one can argue that the revenues generated at single days are independent from each other. Thus, random revenues resulting from the power plant dispatch converge on average to the expected value of their revenue - if a sufficiently large time period is observed. In this case the expected value of "observed" (also called real world probability measure) revenues is chosen as valuation functional (cf. [TB02]).

A different approach resulting in the use of the expected value as valuation functional which is widely used in practice, relies on replication arguments which have been developed in mathematical finance, dating back to the Black-Scholes framework for pricing European options.
In this approach, the price of a derivative is set equal to the required initial investment in a replicating portfolio, which is a portfolio generating an exactly identical cash flow as the (shorted) derivative.
Underlying assumptions are the existence of a replication portfolio (i.e. completeness of the market) and that two derivatives producing identical cash flows have the same value (i.e. no arbitrage). Under these assumptions[2], it can be shown that this price is unique.
The crucial property of this approach is stated in the fundamental theorem of asset pricing (see for example [DS94]): For the assumptions above, there exists a unique probability measure under which the value of any derivative equals the expected value of its payoff. As a consequence, valuation can be performed using the expected payoff of the derivative under the so called equivalent martingale (or risk neutral) measure which has to be derived from the original real world probability measure by a change of measure.
Since the risk is assumed to be hedged away, this approach suggests one "objective" value which will be assigned to a derivative by any market participant - independent of the own risk preferences. Notice that by linearity of the expected value, the price is also independent of the portfolio of any market participant.
In practice, replication is the most popular approach for valuation of energy derivatives. For this, starting at a stochastic process identified to fit observed prices (real world probabilities), as a first step the spot price process under the equivalent martingale measure (risk neutral process) has to be identified. This is usually done by recalibrating the real world price process to fit (as closely as possible) the quotes of the derivatives in the market which are assumed to be priced under the equivalent martingale measure. The value of a derivative is then given by its expected payoff under the risk neutral price process.

If the price process of the underlying is known, the expected payoff of a rigid contract

[2] additional assumptions such as frictionless trading, unconstrained market depth no taxes are often used for obtaining closed forms for derivative pricing for specific price processes, cf. [Gem05] p 83

can be computed immediately analytically or by Monte Carlo integration. A huge amount of literature regarding replication pricing exists, including adaptations to special properties of energy markets. Treatments can be found for example in the books [EW03] and [BGS07]. A rigorous modeling framework is discussed in [BBK08].

1.2.2. Nonlinear case: Risk averse valuation and acceptability functionals

Although its apparent benefits, risk neutral linear valuation may not be justified: in particular for energy markets, completeness and liquidity are highly questionable. In a certain sense, even the concept of no arbitrage does not seem as reasonable as in financial markets due to the presence of generation units which may enable their owner to arbitrage between different energies.

Since in this setting risks cannot be hedged away completely, they have to be accounted for explicitly. Risk averse objective functions can be constructed by modifying the objective to model risk preferences of the decision maker. For this purpose, risk measures have been introduced which map random variables onto the real line. Two distributions can now be compared by comparing their images under the risk measure. A large number of different of measures for risk can be found in the literature.

For the moment[3], the only interesting property of objective functions incorporating risk measures is that they are in general nonlinear in the payoff variable.

In the following, nonlinear valuation functionals will be referred to as acceptability functionals $\mathcal{A}(\cdot)$, which map profit distributions onto the real line. A random profit U_1 is preferred over a random profit U_2 if $\mathcal{A}(U_1) > \mathcal{A}(U_2)$. Our models will be formulated as maximization of profit distributions under the concave acceptability functionals. Notice that this setting is equivalent to the minimization of convex risk measures often encountered in literature, where the loss distribution[4] is mapped onto the real line.

Remark. The connection between penalization and replication: As noted above, one of the central reasons for having to explicitly account for risk is that exact replication of derivatives is infeasible due to market incompleteness.
One possibility to repair this infeasibility is to relax the replication assumption: Instead of requiring the replication portfolio to offset the cash flow of the derivative almost surely, deviation is now allowed as long as it does not exceed a certain residual risk, which can be quantified by a risk measure. The price of the derivative is then the minimal initial investment in a portfolio subject to the residual risk being acceptable.

[3] For the tailored algorithm, the objective function will be required to be polyhedral. See Appendix A.1 for a more detailed treatment of measuring risk and acceptability.

[4] which equals the negative profit distribution

The constraint on the residual risk can be lifted to the objective function (by introducing a penalty), which results in the nonlinear valuation approach discussed above. From duality theory it follows that for convex problems (i.e. convex feasible sets and concave objectives for maximization) there exists a penalty which gives the same result as the constraint formulation.

From a computational viewpoint, nonlinearity of the valuation functional has drastic consequences: The portfolio optimization problem does not automatically decompose into the separate optimization of its single components. The exercise of operating decisions for the joint portfolio in general leads to superior results than exercise optimal for the components.

In a similar way, the value of any derivative cannot be viewed independently of the rest of the portfolio. Thus, when nonlinear valuation functionals are used, values are subjective in two ways: they are dependent on the risk preferences which differ between agents and in particular they depend on the portfolio of each agent.

1.2.3. Pricing of rigid contracts

The linear case In the case of a rigid contract, its price can be readily derived from its value: an upfront payment of the expected payoff will offset the value of a derivative, thus it is the fair market price. Note that in theory, by linearity of the valuation measure, any cash flow having the same expected value as the derivative can be used as a price.

The nonlinear case: Acceptability and indifference pricing The fact that in the risk averse case the values of derivatives are subjective makes the concept of a fair market price inapplicable. To incorporate the portfolio and risk preferences of an agent into pricing, the concepts of acceptability pricing and indifference pricing have been introduced (cf. [KP13b], [Car09] or [Pen12]).
From the acceptability pricing point of view, a market participant is willing to accept any deal which puts the acceptability of the payoff from the new portfolio (i.e. after the deal) above some threshold.
Determining this threshold by the value of the initial portfolio (i.e. before the deal) leads to the concept of indifference pricing. The smallest price which makes the value of the new portfolio equal the value of the old portfolio, is called the indifference price. At this price the market participant is indifferent between the two portfolios and would thus agree to buy (sell) the derivative in exchange for the payment.

If the valuation functional $\mathcal{A}(\cdot)$ is translation equivariant (i.e. $\mathcal{A}(U + c) = \mathcal{A}(U) + c$, see Definition A.2), then also any deterministic payment c amounting to the difference of the value of the portfolio after accepting the deal and the acceptability threshold ρ is acceptable (or an indifference price). Thus, for an initial portfolio of

rigid contracts with cumulative cash flow $\bar{S}^i = \sum_{t=1}^{T} S_t^i$, the price c for selling a rigid contract $S(k,P)$ for (random) prices P and parameter k can be obtained by

$$c = \mathcal{A}\left[\left(\sum_{i=1}^{I} \bar{S}^i - S(k,P)\right)\right] - \rho\,.$$

In the particular case of $\rho = \mathcal{A}\left(\sum_{i=1}^{I} \bar{S}^i\right)$ being the acceptability of the initial portfolio, c is the indifference price.
However, nondeterministic payments lead to significantly more complex pricing problems as will be discussed in the following subsection.

Remark. In the special case of a linear valuation functional, the indifference price equals the fair market price. Therefore the indifference price can be seen as a generalization of the fair market price for illiquid markets, where risks cannot be priced in an objective way.

1.2.4. Long flexible contracts: buyer's problem

Once the valuation functional has been chosen, the valuation of rigid contracts can readily be performed. In the case of flexible contracts however there still persists the issue of the exercise strategy. A usual choice for valuation is the "optimal" exercise strategy, i.e. the strategy optimizing the expected value of the derivative (portfolio). For a long flexible contract $S(k,P,y)$, the optimal exercise policy $y^\star$ is given by

$$y^\star \in \underset{y=y_1,\dots,y_T}{argmax} \left\{ \mathbb{E}\left[S(k,P;y)\right] \middle| s.t. \begin{array}{l} y \in Y \\ y_t \lhd \mathfrak{F}_t \end{array} \right\} \qquad \text{(risk neutral)} \tag{1.3}$$

We refer to problem (1.3) as the *buyer's problem.* The optimal exercise policy depends on the price process $P = (P_t)_{t=1,\dots,T}$ of the underlying energy and on the parameters of the option k. Also, physical or contractual constraints on the decisions are denoted by $y \in Y$, and information constraints by $y_t \lhd \mathfrak{F}_t$. The latter encode the sole dependence of decisions on informations available at time $t-1$, since in our setting, consumption decisions have to be taken already one day before delivery.

Using this identity for the exercise behavior, valuation and pricing can be performed as described in the previous sections.

Depending on the concrete formulation, the resulting optimization problem can be solved for example by multistage stochastic programming (cf. [HK09]), stochastic dynamic programming (cf. [TB02],[BCC12]) or PDEs/finite element methods (cf. [Dah05]).

Notice that, in the spirit of the discussion of the previous subsection, the assumption of risk neutral buyer can be relaxed by formulating the buyer's problem as a risk

averse optimization problem. For the case of an initial portfolio of I *rigid*[5] derivatives with cumulative cash flows $\bar{S}^i = \sum_{t=1}^{T} S_t^i$

$$y^\star \in \underset{y=y_1,\ldots,y_t}{argmax} \left\{ \mathcal{A}\left[\left(\sum_{i=1}^{I} \bar{S}^i + S\left(k, P; y\right)\right)\right] \middle| s.t. \begin{array}{l} y \in Y \\ y_t \lhd \mathfrak{F}_t \end{array} \right\} \text{ (risk averse)} \qquad (1.4)$$

Notice that similar to the previous subsection, this problem takes into account the cash flow $\sum_{i=1}^{I} \bar{S}^i$ and the specific risk preferences of the buyer encoded in $\mathcal{A}(\cdot)$. The resulting problem (1.4) will also be covered by the tailored algorithm but only appear shortly in Subsection 3.1.1 and Section 4.1.

1.2.5. Short flexible contracts: seller's problem

For distinguishing seller's and buyer's problem we introduce $x = x_t^i,\ i \in \{1, \ldots I\}, t \in \{1, \ldots, T\}$ as exercise decisions of the seller for long flexible contracts in the portfolio.
Also, in this subsection, we will highlight the dependency of the exercise policy of the buyer $y^\star$ on the (contract parameter) strike price k by $y^\star(k)$ wherever one exists.

Deterministic payment For the seller[6] of a flexible contract, the exercise behavior of the buyer[7] is an additional risk factor, which is given by the solution of the stochastic optimization problem (1.3) or (1.4) (buyer's problem).

Having this in mind, an acceptable or indifference price can be obtained exactly as in the case of a rigid contract. Assuming that the valuation functional $\mathcal{A}(\cdot)$ is translation equivariant, the price required for selling a specific derivative can be obtained by an upfront cash (or deterministic) payment of the difference between the seller's initial portfolio including the short derivative and the threshold required ρ (or the value of the initial portfolio). For initial portfolio $\Pi(x)$, the price $c^\star$ for a short derivative[8] $S(P; y^\star)$ is given by

$$c^\star = \underset{x}{max} \mathcal{A}\left[\Pi(x) - S\left(P; y^\star\right)\right] - \rho\ .$$

Note the use of $y^\star$, which is the optimal exercise strategy of the buyer of $S(P; y)$. Trivially, this can also be written as the solution of the optimization problem

$$c^\star = \underset{c}{min} \left\{ c \middle| \underset{x}{max} \mathcal{A}\left[\Pi(x) - S\left(P; y^\star\right)\right] + c \geq \rho \right\} \qquad (1.5)$$

[5] the concept naturally extends to the notationally inconvenient case of additional flexible contracts in the initial portfolio
[6] i.e. the holder of a short flexible contract
[7] i.e. the holder of a long flexible contract
[8] with fixed contract parameters k, which are ignored in the notation

Flexible contracts with coupled payments and exercise: Swing options

A conceptually entirely different problem arises, when instead of deterministic payments of a (market or subjective) value, contract parameters for flexible contracts are involved. This problem emerges at the design of swing options, which are contracts where the buyer has the right to determine the amount received and pays a fixed price per consumed unit (strike price, see Example 1.1.3). Usually the reward for the seller is not provided by an upfront cash (or deterministic) payment from the buyer, but by the cash flow determined by the consumption of the buyer at the price specified when the agreement is concluded. In this case, the acceptability level for the seller is given by

$$\mathcal{A}\left(\Pi(x) - S(k; y^\star(k))\right)$$

where the cumulative cash flow $S(k; y^\star(k))$ depends on the buyer's optimal exercise strategy $y^\star(k)$ which is the solution of the buyer's problem (1.3) (or (1.4) respectively).
The difficulty of this problem arises from the fact, that the exercise strategy of the buyer now depends on the fixed price k. This leads to a complicated situation: the seller has to determine an exercise price, which influences the exercise behavior of the buyer, which in turn determines the value of the portfolio.

This problem can be formulated as an optimization problem, which we call the *seller's problem*. It is to find the "optimal strike price" k, such that there exists an exercise strategy $x = x_t^i$, $i = \{1, \dots, I\}$, $t = \{1, \dots, T\}$, for the initial portfolio $\Pi(x)$, under which the portfolio including the short flexible contract stays acceptable:

$$\begin{aligned} \max_{k;x} \; & F(x, k, y^\star(k)) \\ s.t. \quad & \mathcal{A}\left(\Pi(x) - S(k; y^\star(k))\right) \geq \rho \\ & x \in X,\, x_t^i \lhd \mathfrak{G}_t,\, k \lhd \mathfrak{G}_0 \end{aligned}$$

The strike price k is assumed to be set before the delivery period at time 0, therefore it may only depend on the information $\mathfrak{G}_0$. The portfolio-balancing decisions of derivative i, x_t^i, are assumed to be taken on short notice - thus they may depend on $\mathfrak{G}_t$, the information available at time t.

Notice that the exercise behavior of the buyer has to be the solution of the buyer's problem (1.3)

$$y^\star(k) \in \underset{y=y_1,\dots,y_T}{argmax} \left\{ \mathbb{E}\left[\tilde{S}\left(k; y\right)\right] \middle| s.t. \; \begin{array}{l} y \in Y \\ y_t \lhd \mathfrak{F}_t \end{array} \right\}$$

or (1.4) respectively. Here $\tilde{S}\left(k; y\right)$ denotes the cumulative cash flow from the contract for the holder, which emphasizes the fact that the payoff distribution of the buyer is not necessary identical to the cost distribution of the seller. The fact that one of the constraints of the buyer's problem is the solution of another optimization problem makes the buyer's problem a stochastic bilevel problem.

1.3. The behavioral pricing problem

The objective function of the seller, $F(x, k, y^\star)$, has not yet been specified. In the light of the discussion in the previous section, we will present an indifference pricing variant and a risk neutral pricing variant of the problem.

1.3.1. Seller's problem: Indifference pricing variant

The indifference pricing concept introduced in Subsection 1.2.3 can immediately be used to determine the strike price required by a risk averse seller: The objective is to choose the minimal (strike) price, such that the acceptance of the portfolio including the contract is greater or equal to the acceptance of the initial contract, anticipating the reaction of the holder:

$$\left|\begin{array}{lll} & \min\limits_{k;x} k & \\ [UL] & & \mathcal{A}\left(\Pi(x) - S(k; y^\star)\right) \geq \rho \\ & & x_t \in X_t,\ x_t \lhd \mathfrak{G}_t,\ k \lhd \mathfrak{G}_0 \\ [LL] & y^\star \in & argmax \left\{ \mathbb{E}\left(\tilde{S}(k; y)\right) \,\middle|\, \begin{array}{l} y_t \in Y_t \\ y_t \lhd \mathfrak{F}_t \end{array} \right\} \end{array}\right. \tag{1.6}$$

The seller's problem (1.6) is a stochastic multistage optimization problem where one source of uncertainty is given by the solution of another stochastic multistage optimization problem. Due to the acceptance constraint which can incorporate risk preferences of the seller, it will be referred to as *(risk averse) behavioral pricing problem*.
Notice that in this indifference pricing model, the strike price assigned to a contract depends on the portfolio of the seller; i.e. on the possibility to adapt to the anticipated consumption of the buyer.
In principle, problem (1.6) is just an extension of the setting of a deterministic indifference price (problem (1.5)). The complexity in the extension arises from the incorporation of the interdependency of the exercise price k and the resulting random net payments $S(k; y)$.

An algorithm tailored to the solution of problems of the type of (1.6) will be presented in Chapter 3, followed by numerical illustrations in Chapter 4.
Problem (1.6) can also be formulated in a similar way for the risk averse buyer's problem. It will be shown in Chapter 3 that this leads to an identical problem structure and the algorithm developed can be applied.

1.3.2. Seller's problem: Risk neutral variant

As observed above, in the case of a linear valuation functional, the portfolio of the seller does not influence its value. In the absence of [UL] constraints, minimizing

the strike price k makes little sense.

One interesting problem is to choose $F(k,y) = \mathbb{E}\left(S(k;y)\right)$, i.e. to look for the strike price k maximizing expected [UL] payoff, which potentially differs from the [LL] payoff $\tilde{S}$. To keep the offer reasonable for [LL], a constraint on the minimal payoff requested by [LL] is included by $\mathbb{E}\left(\tilde{S}(k;y)\right) \geq \nu_0$.
Thus, the seller's problem for the risk neutral case will be

$$\left| \begin{array}{l} \max\limits_{k} \mathbb{E}\left(S(k,y^\star)\right) \\ \qquad s.t. \qquad\quad y^\star \in argmax \left\{ \mathbb{E}\left(\tilde{S}(k;y)\right) \;\middle|\; \begin{array}{c} \mathbb{E}\left(\tilde{S}(k;y)\right) \geq \nu_0 \\ y_t \in Y_t,\ y_t \lhd \mathfrak{F}_t \end{array} \right\} \end{array} \right. , \tag{1.7}$$

the *risk neutral behavioral pricing problem*. Notice that in the case of equal probability distributions for [UL] payoff and [LL] cost, the solution has to satisfy $-\mathbb{E}\left(S(k;y)\right) = \mathbb{E}\left(\tilde{S}(k;y)\right) \geq \nu_0$, thus the [UL] optimal expected payoff will be $-\nu_0$ at most.

However, the case of different distributions for [UL] payoff and [LL] cost might be more interesting to analyze. Although this setting may seem unconventional at first glance, in practice companies often value their resources with internal transaction prices[9] which may differ from spot prices. Thus, internal prices used by the seller may well differ from spot prices determining the buyer's behavior, which leads to different payoff/cost distributions. Then, under risk neutral valuation, the problem of designing an exercise price amounts to solving problem (1.7).

It is easier to solve problem (1.7) than problem (1.6): Since the valuation in the risk neutral case is independent of the portfolio of the seller, the seller's only decision variable is the strike price k which is unconstrained. However, due to the fundamentally different objective functions, the tailored algorithm for problem (1.6) cannot be used for solving problem (1.7).
Since we also are convinced that the risk neutral variant does not appropriately capture the real problem, we will not devote any further attention to problem (1.7) (except one single short remark on the *non*-applicability of some optimality conditions in Chapter 2). For the solution of a similar problem by iterative linear programming we refer to [KP13a].

[9] for example reflecting own production or sourcing costs

2. Bilevel problems

The behavioral pricing problem belongs to the class of stochastic bilevel problems. A short overview over bilevel problems and solution methods is presented in this chapter.

2.1. Introduction

2.1.1. Origin

Originally, the first bilevel problem was stated in the course of a game theoretic analysis of oligopolistic markets in the 1930's by Heinrich Freiherr von Stackelberg. To describe real market situations, von Stackelberg introduced a framework where actions are taken *sequentially* by agents with *different objectives* depending on the decisions of all agents.

The simplest example is the case of two firms, which have to determine the optimal level of their production. Since the total level of supply determines the market price, the decision of each firm influences the payoff of the other. By committing itself to a certain level, the firm acting first (leader) sets the stage for the second firm (follower) which can maximize its profit by solving a classical problem of optimization. Thus the leader has to solve an optimization problem which is constrained by another optimization problem.

Using a little more general formulation, this problem can be stated as

$$\left|\begin{array}{ll} [UL] & \max\limits_{x} F(x,y) \\ & \quad s.t. \quad G(x,y) \leq 0 \\ [LL] & y \in \underset{z}{argmin}\ (f(x,z) | g(x,z) \leq 0) \end{array}\right. . \tag{2.1}$$

We denote by [UL] the so called upper-level problem, which is to be solved by the leader or upper-level player and by [LL] the lower level problem, to be solved by the lower level player or follower.

2.1.2. Further applications

An important field for bilevel optimization are traffic control problems and toll setting. An example is [PR02], where upper level variables can be the design of the traffic network, signals and tolls whereas the lower level consists of a traffic equilibrium problem.
Another application of bilevel programming is the determination of optimal chemical equilibria ([Dem02]). The basic question is how to compose a mixture of chemical substances such that the required substance is actually generated by the reaction. Apparently, this can be modeled as a convex optimization problem and is considered as lower level problem. The upper level problem arises from maximizing the amount of the substance generated and/or imposing the additional requirement of minimizing unwanted byproducts. The same principle can also be used in bio engineering, see for example [BPM03].

The origin of bilevel problems from strategic decision making also leads to several applications in energy markets. Bilevel formulations have been used to study the economics of energy markets and market designs from a game theoretic point of view as for example in [BJ05].
A different application of bilevel problems in energy markets is the analysis of the interdependent decision problems of several actors along value chains. The particular case of natural gas transportation is the topic of the series of papers ([KRM06],[KPVTK10] and [KDPVK11]). In the setting analyzed, a shipping company extracts natural gas at some fields and wants to sell it at some markets maximizing the profit. Transportation is done by pipelines which are operated by a different[1] company, having to satisfy physical constraints on operation. The bilevel structure arises from the coupling of the decision problems of the two agents with different objectives.
As another application of bilevel programs in energy markets, in [CAC09] a bilevel formulation was used to describe the problem of setting of retail prices in an oligopolistic power retail market. In their model, the producer ([UL] player) has to determine a delivery price and hedging/production decisions, while the the end-consumer ([LL]-player) decides about accepting the delivery contract offered at the conditions set by [UL]. The [UL] problem is modeled as stochastic multistage portfolio optimization problem, whereas the [LL] problem is deterministic (since the consumption of retail customers behavior does not depend on wholesale spot prices). Notice the difference to the behavioral pricing problem which is analyzed in this thesis, where both [UL] and [LL] are multistage stochastic programs, which leads to a completely different optimization problem.

Further applications of bilevel problems can be found for example in the book [Dem02] and the survey papers [Dem03] and [CMS07]. Many applications can also be found in the literature on mathematical programs with equilibrium constraints

[1] vertical integration of production and transportation companies is often prohibited in energy markets

(MPEC), which are closely connected to bilevel problems (for details, see Subsection 2.3.1).

2.2. Basic properties of bilevel programs

2.2.1. Optimistic and pessimistic formulations

Definition. The *optimal response* of the lower level player is the set valued mapping assigning to every feasible [UL] decision the set of [LL] solutions, i.e.

$$y^\star(x) := \underset{z}{argmin}\ (f(x,z)|g(x,z) \leq 0)$$

Remark. A problem which superficially seems related to the optimal response in this context is the the recourse action[2] in two stage stochastic programming. Their distinguishing feature is that in bilevel programs the lower level problem influences the upper level problem via *optimal decisions* whereas in two stage stochastic programs the first stage decision is only influenced by the *optimal value* of the second stage. The specific difficulty of bilevel problems is rooted in the usually very complex interdependency of upper level and lower level optimal decisions.

Since $y^\star(x)$ is in general set-valued, it may not always be clear for the upper level player how the lower level player will respond. To ensure that problem (2.1) is well defined, two different solution concepts have been introduced for bilevel problems:

Definition. The *optimistic solution* of a bilevel problem is the case where [LL] selects from the set of optimal solutions the most suitable for [UL] i.e.

$$\left|\begin{array}{ll} [UL] & \underset{x}{max}\ \left(\underset{y}{sup} F(x,y)\right) \\ & \qquad s.t. \qquad G(x,y) \leq 0 \\ [LL] & y \in \underset{z}{argmin}\ (f(x,z)|g(x,z) \leq 0) \end{array}\right. , \tag{2.2}$$

the *pessimistic solution* is the case where [LL] selects from the set of optimal solutions the worst case for [UL] i.e.

$$\left|\begin{array}{ll} [UL] & \underset{x}{max}\ \left(\underset{y}{inf} F(x,y)\right) \\ & \qquad s.t. \qquad G(x,y) \leq 0 \\ [LL] & y \in \underset{z}{argmin}\ (f(x,z)|g(x,z) \leq 0) \end{array}\right. . \tag{2.3}$$

Clearly, in case [LL] has a unique solution, no such distinction is necessary.

In the particular case where [LL] is convex and finite dimensional, we have that the set of optimal solutions of [LL] is compact and thus the *sup* in problem (2.2) and the *inf* in problem (2.3) are attained.

[2] i.e. the action taken for revealed information after the first stage action has already been taken under uncertainty

2.2.2. Geometric illustration

To illustrate some properties of bilevel problems we look at the following example:

$$\left|\begin{array}{ll} [UL] & \max\limits_{x} x(y_1+y_2) \\ & \quad s.t. \qquad 0 \le x \le 30 \\ [LL] & y \in \underset{y_1,y_2}{argmin} \left((30-x)y_1 + (20-x)y_2 \,\middle|\, \begin{array}{c} a \le y_i \le b \\ A \le y_1+y_2 \le B \end{array} \right) \end{array}\right. \tag{2,4}$$

A sketch containing the optimization problem of the lower level player is depicted in Figure 2.1. The gray area describes the feasible set, the directions of the gradient of the objective function are marked by arrows for the [UL] values of $x = 0$ and $x = 20$. In this context, the choice of the [UL] player amounts to determining the direction of the gradient of the objective function for the [LL] player, i.e. setting the angle between gradient and the coordinate axis.
It is easy to see that for a given value of the [UL] decision variable, the [LL] maximization leads to optimal solutions depicted in Figure 2.1 as the thick gray line at the border of the feasible set. Explicitly, the optimal response $y^\star(x)$ can be represented as

$$y^\star(x) = \begin{cases} \begin{pmatrix} b-\theta \\ B-b+\theta \end{pmatrix}, \theta \in [0, 2b-B] & \text{for } x = 0 \\ \begin{pmatrix} B-b \\ b \end{pmatrix} & \text{for } x \in (0,20) \\ \begin{pmatrix} B-b-\theta \\ b \end{pmatrix}, \theta \in [0, B-b] & \text{for } x = 20 \\ \begin{pmatrix} 0 \\ b \end{pmatrix} & \text{for } x \in (20,30) \\ \begin{pmatrix} 0 \\ b-\theta \end{pmatrix}, \theta \in [0, b-A] & \text{for } x = 30 \end{cases}$$

Note that for $x = 0$, $x = 20$ and $x = 30$ the solutions to the [LL] problem are the whole facets of the feasible set which are orthogonal to the gradient of the objective. At these points, $y^\star(x)$ is nonunique; for every other value of x in the feasible region it is a single point.

To solve the upper level problem, the maximal value of $x(y_1+y_2)$ over $y^\star(x)$ has to be found. Using the explicit form of $y^\star(x)$, the [UL] objective can be computed as

$$F\left(x, y^\star(x)\right) = \begin{cases} Bx & \text{for } x \in [0,20) \\ (B-\theta)x\,, \theta \in [0, B-b] & \text{for } x = 20 \\ bx & \text{for } x \in (20,30) \\ \theta x\,, \theta \in [b, A] & \text{for } x = 30 \end{cases}$$

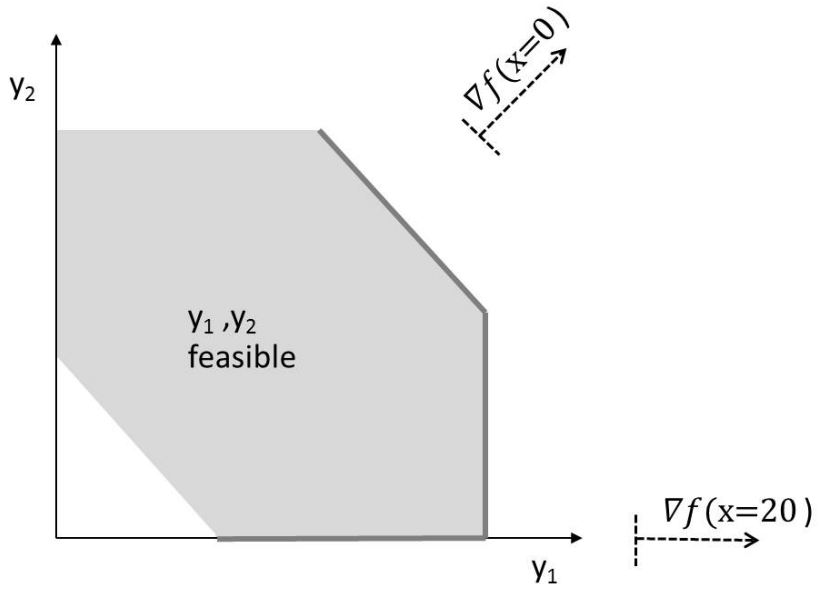


Figure 2.1.: Geometric representation of a bilevel problem (schematic).
The thick gray line represents [LL] responses to the [UL] decisions. The gradient of the [LL] objective function, $\nabla f(x)$, is represented for two particular choices of x at which the [LL] response is not unique (see Figure 2.2 for the effect on the [UL] objective).

As addressed in Subsection 2.2.1, this initial optimization problem is not well posed since its objective is a multifunction. To avoid this problem, we now look at the optimistic (F_o) and the pessimistic (F_p) problem:

$$F_o\left(x, y^\star(x)\right) = \begin{cases} Bx & \text{for } x \in [0, 20] \\ bx & \text{for } x \in (20, 30] \end{cases} \qquad F_p\left(x, y^\star(x)\right) = \begin{cases} Bx & \text{for } x \in [0, 20) \\ bx & \text{for } x \in [20, 30) \\ Ax & \text{for } x = 30 \end{cases}$$

where $A < b$ and $b < 2/3B$ was assumed. Figure 2.2 illustrates the [UL] objective function and the different versions: The black line shows the function $F\left(x, y^\star(x)\right)$, where the dotted parts show the image intervals at $x = 20$ and $x = 30$. The optimistic problem formulation uses the objective F_o, which arises from replacing the intervals by their uppermost point (marked gray). The pessimistic formulation uses F_p, which is generated by replacing them by their lowest point (marked black). Clearly, the optimal value of the optimistic problem occurs at $x = 20$. Notice however, that an optimal value of the pessimistic problem does not exist. More generally, observe that [UL] is a nonconvex optimization problem with discontinuous objective function. In the special case where [UL] and [LL] are linear programs when the variables of the other problem are fixed, the [UL] objective is a piecewise linear function.

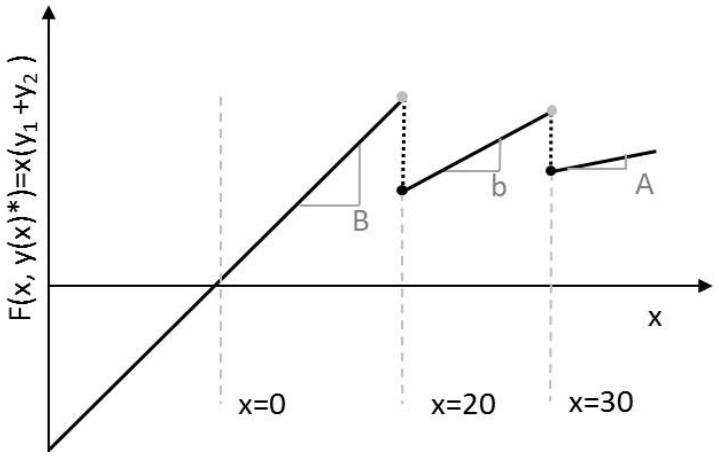


Figure 2.2.: [UL] objective for the optimal response. Dotted lines (black) mark nonunique function correspondences, gray points the optimistic, and black points the pessimistic version.

To investigate an additional aspect, consider a modification of problem (2.4), where a constraint containing [LL] variables is added to the [UL] problem:

$$\left|\begin{array}{lll} & \max\limits_{x} x(y_1+y_2) & \\ [UL] & s.t. & 0 \le x \le 30 \\ & & c_1y_1 + c_2y_2 \le 1 \\ [LL] & y \in \underset{y_1,y_2}{argmin} \left((30-x)y_1 + (20-x)y_2 \,\middle|\, \begin{array}{c} a \le y_i \le b \\ A \le y_1+y_2 \le B \end{array} \right) & \end{array}\right. \tag{2.5}$$

For suitable choice of c_1, c_2 (in particular, $c_1 > c_2$), the new constraint intersects with the feasible region of the [LL] problem as illustrated in Figure 2.3. Since the constraint does not influence the [LL] optimal decisions, it results in the [UL] infeasibility of some of the optimal responses. Thus, indirectly constraints are induced on the [UL] decision variable, which may lead to a disconnected feasible set.
This effect can be observed in Figure 2.3, where the set of [LL] solutions $y^\star(x)$ feasible for [UL] is restricted by the dashed line, thus separating the feasible choices of x (here corresponding to the angles of the gradient of the [LL] objective, which are controlled by x) into two parts.

2.3. Optimistic bilevel programs: Single level reformulations

One widespread approach to analyze or solve optimistic bilevel problems is to reformulate them as single level problems, which makes it possible to apply analytic

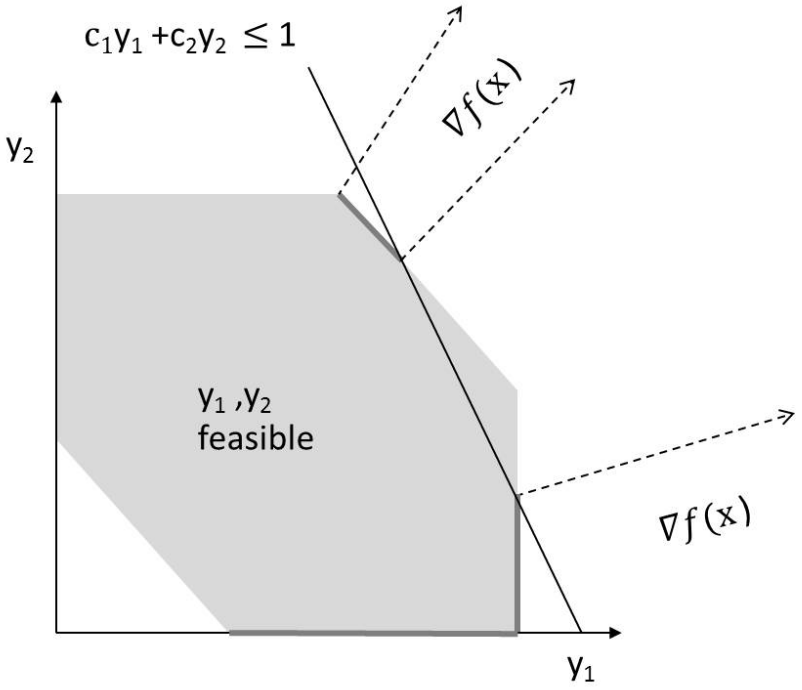


Figure 2.3.: Schematic illustration of the effect of [UL] constraints containing [LL] variables (coupling constraint).

tools developed for classical optimization theory.

The subsequent presentation of the optimality conditions is largely based on the recent papers by Dempe and Zemkoho, [DZ11] and [DZ12].

Following to their presentation, we assume the [UL] constraint only to depend on the [UL] variable x i.e.

$$\left|\begin{array}{ll} [UL] & \max\limits_{x} F(x,y) \\ & \quad s.t. \qquad G(x) \leq 0 \\ [LL] & y \in \underset{z}{argmin}\ (f(x,z) | g(x,z) \leq 0) \end{array}\right. . \tag{2.6}$$

We confine ourselves to the case where [UL] and [LL] are optimization problems of convex type, and the functions $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ and $G : \mathbb{R}^n \to \mathbb{R}^k$, $g : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^p$ are continuously differentiable.

2.3.1. Single level reformulation using optimality conditions: MPECs

Our starting point is to replace the [LL] optimization problem with a constraint forcing the [LL] decision variable to satisfy first order optimality conditions.

In particular, assuming that the [LL] program is convex, and the [LL] objective function is continuous in the [LL] decision variables (i.e. treating [UL] variables as parameters), the application of the first order optimality conditions (Appendix, Theorem A.2) leads to the following single level program:

$$\begin{aligned} \max_{x,y} & F(x,y) \\ s.t. \quad & G(x) \leq 0 \\ & 0 \in \nabla_y f(x,y) + N_{Y(x)}(y) \end{aligned} \tag{2.7}$$

where $Y(x) := \{y | g(x,y) \leq 0\}$ denotes the [LL] feasible set for a particular choice of the [UL] variable x and $N_{Y(x)}(y)$ the normal cone[3] to $Y(x)$ at some $y \in Y(x)$. This optimization problem is equivalent to (the optimistic variant of) problem (2.6) as observed by various authors, for example in [DZ11].

Using the convexity of the [LL] problem, the normal cone can be expressed by the variational inequality

$$N_{Y(x)}(y) = \left\{ v \in \mathbb{R}^n \,\middle|\, v^T(u-y) \leq 0 \,, \forall u \in Y(x) \right\}.$$

Plugging this formula into problem (2.7) leads to an optimization problem with variational inequality constraints.

Assume in addition that a constraint qualification[4] holds for the [LL] problem. Then, for $y \in Y(x)$ (i.e. $g(x,y) \leq 0$) the normal cone can be explicitly computed by a linear combination of the gradients of active constraints:

$$N_{Y(x)}(y) = \left\{ \nabla_y g(x,y)^T \lambda \,\middle|\, \lambda \geq 0, \lambda^T g(x,y) = 0 \right\}$$

Using this formula, necessary optimality conditions for feasible y can be described by the well known KKT conditions

$$\exists \lambda, \lambda \geq 0,\ \lambda^T g(x,y) = 0 \,\middle|\, \nabla_y f(x,y) + \lambda^T \nabla_y g(x,y) = 0 \ .$$

Together with requiring [LL] feasibility by $g(x,y) \leq 0$, this formula for the normal cone can be used to to obtain a single level reformulation of optimistic bilevel problem with the following structure:

$$\begin{aligned} \max_{x,y} \ & F(x,y) \\ s.t. \quad & G(x) \leq 0 \\ & \nabla_y f(x,y) + \lambda^T \nabla_y g(x,y) = 0 \\ & g(x,y) \leq 0, \lambda^T g(x,y) = 0 \\ & \lambda \geq 0 \end{aligned} \tag{2.8}$$

[3] see Definition A.8

[4] for example LICQ, see Appendix, Definition (A.10)

which is referred to as the KKT-reformulation of the bilevel problem. From the KKT-conditions, problem (2.8) inherits complementarity conditions and consequently is a mathematical program with complementarity constraints (short MPCC or MPEC[5]).

The connections between problem (2.8) and the original problem (2.6) were thoroughly studied in [DD10] and the following connection was established (for [LL] satisfying Slater conditions): the set of local optimal solutions of (2.6) are local optima of (2.8) but local optima of the KKT-reformulation are not necessarily local optima of the original problem. However, global solutions of problem (2.8) coincide with global solutions of the original problem (2.6).

Not very surprisingly, also the one-level problem (2.8) stays difficult to solve. Due to the complementarity constraints, it is difficult to find constraint qualifications weak enough to use characterizations of the type of (A.2).
Classical constraint qualifications such as the Mangasarian-Fromowitz constraint qualification[6] (MFCQ) (and as a consequence a series of more restrictive constraint qualifications) are violated at any feasible point of an MPEC and thus also for the KKT-reformulation as a special case (see for example [Fle05]).

2.3.2. Optimality conditions of the KKT-reformulation under Guignard CQ

As a concept, the characterization of optimality by first order conditions used above can be directly applied to one level reformulations of bilevel programs. However, difficulties arise when "the feasible set is complicated" which in this context means that the constraint qualifications fail.
As we just noted, standard constraint qualifications tend to fail for one level reformulations of bilevel problems. However, the weaker Guignard constraint qualification[7] (GCQ) may still be satisfied for the one level reformulation, which leads to the following necessary optimality conditions:

Theorem 2.1. *Optimality conditions for KKT-reformulation: [[DZ11], Th. 3.5.] Assume that the lower level problem is a convex one and the MFCQ is satisfied at all feasible y ($y \in Y(x)$) and feasible x ($x \in X$). Furthermore let $(\bar{x}, \bar{y}, \bar{u})$ be a local optimal solution to problem* (2.8) *and the GCQ be satisfied at $(\bar{x}, \bar{y}, \bar{u})$. Then, there*

[5] from Mathematical Program with Equilibrium Constraint, which gives a hint at the origin of the complementarity constraints encountered in many MPCCs
[6] see Appendix, Definition (A.11)
[7] see Appendix, Definition (A.12)

exists $(\lambda, \alpha, \gamma, \nu)$ *in* $\mathbb{R}_+ \times \mathbb{R}^p \times \mathbb{R}^k \times \mathbb{R}^m$ *such that:*

$$\begin{aligned}
\nabla_x F(\bar{x},\bar{y}) + \nabla_x g(\bar{x},\bar{y})^T(\alpha - \lambda\bar{u}) + \nabla G(\bar{x})\gamma + \nabla_x \mathcal{L}(\bar{x},\bar{y},\bar{u})^T\nu &= 0 \\
\nabla_y F(\bar{x},\bar{y}) + \nabla_y g(\bar{x},\bar{y})^T(\alpha - \lambda\bar{u}) + \nabla_y \mathcal{L}(\bar{x},\bar{y},\bar{u})^T\nu &= 0 \\
\nabla_y f(\bar{x},\bar{y}) + \nabla_y g(\bar{x},\bar{y})^T\bar{u} &= 0 \\
\alpha \geq 0, \alpha^T g(\bar{x},\bar{y}) &= 0 \\
\bar{u} \geq 0, \bar{u}^T g(\bar{x},\bar{y}) &= 0 \\
\nabla_y g(\bar{x},\bar{y})\nu - \lambda g(\bar{x},\bar{y}) \geq 0, \bar{u}^T(\nabla_y g(\bar{x},\bar{y})\nu) &= 0
\end{aligned}$$

where $\mathcal{L}(x,y,u) := \nabla_y f(x,y) + \nabla_y g(x,y)^T u$ *denotes the Lagrangian of the lower level problem.*

2.3.3. Value function reformulation

A very general approach to reformulation (attributed to [Out88]), which does not rely on convexity of the lower level problem, is the so called value-function-reformulation. Interpreting the [LL] problem as parametric optimization problem, the optimal value of the [LL] problem can be written as a function of the [UL] decision variables

$$\varphi(x) = \min_y \{ f(x,y) \,|\, g(x,y) \leq 0 \} .$$

Since the set of [LL] optimal solutions can be described by

$$\{ y \,|\, f(x,y) - \varphi(x) \leq 0,\ g(x,y) \leq 0 \} ,$$

a single level problem equivalent to the optimistic version of problem (2.6) is given by

$$\begin{aligned}
&\max_{x,y} F(x,y) \\
&\quad s.t. \quad G(x) \leq 0 \\
&\qquad\quad f(x,y) - \varphi(x) \leq 0 \\
&\qquad\quad g(x,y) \leq 0 .
\end{aligned} \tag{2.9}$$

Notice that although this approach is very flexible with respect to the [LL] problem structure, the constraint containing the [LL] value function leads to difficulties at deriving optimality conditions.

2.3.4. Optimality conditions for the value-function-reformulation under partial calmness

The main difficulty in the optimization problem (2.9) lies in the constraint $f(x,y) - \varphi(x) \leq 0$. Since this strongly impedes the description of the feasible set, one step

towards deriving optimality conditions is to lift the constraint to the objective function. To guarantee that the optimal solutions to the modified problem coincide with the original optima, the concept of partial calmness has been introduced in the seminal paper [YZ95].

Definition 2.2. Problem (2.9) is partial calm at $(\bar{x}, \bar{y})$ if there is a number $\mu > 0$ and a neighborhood U of $(\bar{x}, \bar{y}, 0)$ such that

$$F(x, y) - F(\bar{x}, \bar{y}) + \mu\,|u| \geq 0$$

for all $(x, y, u) \in U$ feasible to the partially perturbed problem

$$\begin{aligned} \min\ & F(x, y) \\ s.t.\quad & f(x, y) - \varphi(x) + u = 0\ . \\ & G(x) \leq 0 \\ & g(x, y) \leq 0 \end{aligned}$$

Theorem 2.3. *[YZ95]: Let $(\bar{x}, \bar{y})$ be a local optimal solution of problem (2.9). This problem is partially calm at $(\bar{x}, \bar{y})$ if and only if there exists $\lambda > 0$ such that $(\bar{x}, \bar{y})$ is a local optimal solution of the partially penalized problem to*

$$\begin{aligned} \min\ & F(x, y) + \lambda(f(x, y) - \varphi(x)) \\ & s.t.\quad G(x) \leq 0 \\ & \qquad g(x, y) \leq 0 \end{aligned}$$

The following theorem results from [DZ11], and characterizes optimality conditions under partial calmness.
In this case, convexity of both, the [UL] and [LL] problem are assumed. Notice that the value function $\varphi(x)$ does not appear in the optimality conditions.

Theorem 2.4. *[[DZ11], Th. 3.4. and Th. 4.1.] Let $(\bar{x}, \bar{y})$ be a local optimal solution to problem (2.9) and assume that problem (2.9) is partial calm (with respect to $f(x, y) - \varphi(x)$) at $(\bar{x}, \bar{y})$. We also let $\varphi(x)$ be convex and finite around $\bar{x}$. Furthermore, assume that the functions G and g are convex and there exists $(\tilde{x}, \tilde{y})$ such that $G(\tilde{x}) < 0$ and $g(\tilde{x}, \tilde{y}) < 0$. Then, there exists $(\lambda, \alpha, \beta, \gamma) \in \mathbb{R}_+ \times \mathbb{R}^p \times \mathbb{R}^p \times \mathbb{R}^k$ such that:*

$$\begin{aligned} \nabla_x F(\bar{x}, \bar{y}) + \nabla_x g(\bar{x}, \bar{y})^T(\alpha - \lambda\beta) + \nabla G(\bar{x})^T\gamma &= 0 \\ \nabla_y F(\bar{x}, \bar{y}) + \nabla_y g(\bar{x}, \bar{y})^T(\alpha - \lambda\beta) &= 0 \\ \nabla_y f(\bar{x}, \bar{y}) + \nabla_y g(\bar{x}, \bar{y})^T\beta &= 0 \\ \alpha \geq 0,\ \alpha^T g(\bar{x}, \bar{y}) &= 0 \\ \beta \geq 0,\ \beta^T g(\bar{x}, \bar{y}) &= 0 \\ \gamma \geq 0,\ \gamma^T G(\bar{x}) &= 0 \\ \lambda &> 0 \end{aligned}$$

In general, partial calmness is hard to check. However, a useful special case has been reported in [DZ11]:
Let the [LL] optimization problem be linear in x and linear in y, i.e. be given by

$$\begin{aligned} \min_{y} \; & a(x)^T y + b(x) \\ s.t. \quad & C(x)y \leq d(x) \end{aligned} \tag{2.10}$$

where $a(x)$, $b(x)$, $C(x)$ and $d(x)$ are linear functions of x. Consider the optimal value reformulation

$$\begin{aligned} \min_{x,y} \; & F(x,y) \\ s.t. \quad & a(x)^T y + b(x) \leq \varphi(x) \\ & C(x)y \leq d(x) \end{aligned} \tag{2.11}$$

of a bilevel problem, where the objective $F(x,y)$ is Lipschitz continuous.

Theorem 2.5. *[[DZ11], Th. 4.2.] Let $(\bar{x},\bar{y})$ be an optimal solution of problem (2.11), the function F be Lipschitz continuous and $Y(x) \neq \emptyset \; \forall x$. Then, problem (2.11) is partially calm at $(\bar{x},\bar{y})$.*

Remark. At the first glance, Theorem 2.5 seems appropriate to justify the usage of the optimality conditions from Theorem 2.4 for the *risk neutral* bilevel pricing problem (1.7) without any further optimality conditions.
Notice however that the condition $Y(x) \neq \emptyset$ will be eventually violated at some strike price due to the presence of [LL] constraint enforcing a minimal expected payoff ($\mathbb{E}_{LL}\left(\tilde{S}(k;y)\right) \geq \nu_0$), which makes Theorem 2.5 inapplicable for this case.

2.3.5. Optimality conditions of the KKT-reformulation under partial calmness

The reason for the failure of constraint qualifications for the KKT-reformulation lies in the complementarity slackness condition. As in the case of the value-function reformulation, the concept of partial calmness can also be used to lift the complementarity constraint to the objective function, thus generating a feasible set which has much nicer properties.
In [DZ12], partial calmness of the KKT-reformulation with respect to the complementarity constraints is used to justify the lifting of the complementarity constraints to the objective. Optimality conditions similar to the ones in Theorem 2.1 can be derived for the resulting reformulation.

2.4. Solution methods

The subsequent Subsection provides a short discussion of several algorithmic approaches to the solution of bilevel problems. Over the last decades a lot of effort was put in this topic; the section below is nonexhaustive by far.
Apart from the approaches discussed below, other well-known techniques from various fields of optimization have been applied at bilevel programs including vertex enumeration, penalty methods, bundle algorithms, trust region methods and various heuristics. For further information, we again refer to the overviews in [CMS07], [Dem02] and [Dem03] and the references therein.

Each of the methods introduced below performs well for a particular class of problems. Nevertheless, each of them appears to be not suitable for solving the risk averse behavioral pricing problem (1.6) arising from practical application, simply due to the mere size of the problem.

2.4.1. Computational complexity

Computational complexity distinguishes between easy and difficult computational problems by the number of elementary operations that are needed in the worst-case, to solve a mathematical program with respect to the size of its input data.
"Easy" problems can be solved by a number of operations which can be bounded by a polynomial of the input data size. The class of these problems is referred to as the polynomially solvable problems and denoted by $\mathcal{P}$. One prominent member of this family are linear programs.
In this context, "difficult problems" belong to the class of $\mathcal{NP}$ which are solvable by polynomial non deterministic algorithms. These algorithms require that the sole testing if a given point is the solution of the problem is possible in polynomial time; finding of candidates is assumed to be a random operation.
The subclass of $\mathcal{NP}$-hard problems consists of "the most difficult" problems: all elements of $\mathcal{NP}$ can be transformed into a $\mathcal{NP}$-hard problem, such as the satisfiability problem, in polynomial time. Showing that a $\mathcal{NP}$-hard problem can be solved in polynomial time would imply that $\mathcal{P} = \mathcal{NP}$, which is widely believed not to be true, but still an unsolved question. For a more detailed introduction to computational complexity we refer to Appendix B in [WS10].

Bilevel programs are very difficult to solve. Even bilevel problems of the simplest type, where all functions involved are linear (jointly in [UL] and [LL] variables), are contained in $\mathcal{NP}$. This is illustrated by the following example, in which a mixed binary program (belonging to the class $\mathcal{NP}$) is formulated as a linear bilevel program:

Example. ([CMS07], [Dem02]) Consider the mixed binary program:

$$\begin{aligned} \max_{x,u}\ & cx + eu \\ s.t.\quad & Ax + Eu \leq b \\ & x \geq 0,\ u \in \{0,1\}^{n_u} \end{aligned}$$

where $c \in \mathbb{R}^{n_x}, e \in \mathbb{R}^{n_u}, A \in \mathbb{R}^{m \times n_x}, E \in \mathbb{R}^{m \times n_u}, b \in \mathbb{R}^m$. Denoting a vector where each component equals one as $\mathbf{1}$, the binary requirement on u can be expressed as

$$0 = min\{u, \mathbf{1} - u\},$$

which leads to the following bilevel program

$$\left|\begin{array}{ll} & \max\limits_{c,u,y}\ cx + eu \\ [UL] & \quad s.t. \quad\ Ax + Eu \leq b \\ & \qquad\quad x \geq 0 \\ & \qquad\quad y^\star = 0 \\ [LL] & y^\star \in \underset{y}{argmax} \left\{ \sum_{i=1}^{n_u} y_i \,\middle|\, \begin{array}{c} y \leq u \\ y \leq \mathbf{1} - u \end{array} \right\} \end{array}\right.$$

where integrality constraints have been expressed by an optimization problem.

Solving bilevel problems in polynomial time is only possible in certain degenerate cases (cf. [WTKR13]).

In general, nonlinear bilevel problems have been shown to be $\mathcal{NP}$-hard ([Dem02]).

2.4.2. MPEC reformulations

For algorithmic purposes KKT-reformulations are a rather popular method to solve bilevel problems. As observed above, under some regularity conditions, an optimistic bilevel problem can be transformed to a MPEC which has the same global optima.

MPECs have been extensively researched, see for example [Dem03],[Fle05],[FC08] and the references therein.

The MPEC resulting from the KKT-reformulation can be solved by transformation of the complementarity constraints into binary constraints: For two variables $x, y \in \mathbb{R}^n$, for which componentwise $0 \leq x, y \leq M$ holds, the constraint $x^T y = 0$ can be equivalently written by

$$\begin{aligned} 0 \leq\ & x\ \leq zM \\ 0 \leq\ & y\ \leq (1-z)M \end{aligned}$$

where $z \in \{0,1\}^n$ is an auxiliary variable with binary components. Using this so called "big M" method the problem can be solved using standard mixed integer

solvers.
This method has for example been applied in [CAC09] for solving the rather large deterministic equivalent of stochastic bilevel problems of a specific structure.

However, in general this procedure may be numerically unstable, in particular when no reasonable bounds can be derived on dual variables[8]. This problem can be avoided by designing branch and bound algorithms to directly branch at the complementarity constraints, omitting the detour of a MIP reformulation (c.f. [CMS07]). As one example, the YALMIP[9] optimization toolbox for MATLAB includes a generic bilevel solver for small problems based on a branch and bound procedure.

A different approach was pursued by [FL10] for the solution of a transportation network design problem, formulated as a two stage stochastic programming problem with equilibrium constraints. The Progressive Hedging algorithm was used as a heuristic for scenariowise decomposition, allowing the subproblems (still containing equilibrium constraints) of moderate size to be solved by tailored nonconvex optimization algorithms.

Apart from tackling stochastic MPECs by solving their deterministic approximations, it is also possible to solve the original problem by sampling. For example in [SX08], the SAA algorithm was used to obtain estimates of stationary points of a two stage stochastic problem with recourse and equilibrium constraints, which has a structure similar to problem (2.7).

2.4.3. Descent methods

One possibility to obtain stationary points[10] of bilevel problems is by descent algorithms, which iteratively improve the objective function value. For example, the following principle can be used (cf. [CMS07]):
Assuming uniqueness of the [LL] solution, it can be expressed as $y(x)$, an implicit function of the [UL] variable. Where $\nabla_x y(x)$, the gradient of the [LL] solution with respect to the [UL] decision variable is well defined, using the chain rule a direction of descent is given by

$$\nabla_x F(x, y(x)) = \nabla_x F(x, y) + \nabla_y F(x, y) \nabla_x y(x)$$

which can be used to improve the objective value by moving inside the feasible set. In practice however, obtaining the derivatives at feasible points and not leaving the feasible region at update steps poses major challenges.

[8] see [Löf04], documentation of the bilevel solver in YALMIP under http://users.isy.liu.se/johanl/yalmip/pmwiki.php?n=Commands.Solvebilevel
[9] [Löf04]
[10] note that in general constraint qualifications have to be satisfied to link the stationarity to local optimality

One particularly interesting work was done by [Wer04], who designed a gradient descent algorithm to bilevel problems with [UL] coupling constraints and a two-stage stochastic setting.
As illustrated in Subsection 2.2.2, coupling constraints can lead to disconnected feasible regions, which is particularly problematic for gradient descent methods. In [Wer04], this problem was dealt with by separation of the [UL] feasible set into connected convex sets, which were then searched by the algorithm.

2.4.4. Parametrization of lower level solution

The lower level problem can be interpreted as a parametric optimization problem, where the upper level variables are the parameters.
In the particular case of a low-dimension of these parameters, it may be possible to compute the lower level solution explicitly and to use the result in the upper level problem.

This idea was used in [PB09] for solving a behavioral pricing problem (the ancestor of problem (1.6)): The authors observed that the solution of the lower level problem is a piecewise constant function of the upper level parameter. Consequently, to determine the optimal solution of the lower level problem for the whole range of upper level prices, the lower level problem only has to be solved once for a certain price interval (upper level decision variable). In a second step, the upper level optimal solution is computed for each of the intervals for the respective lower level solutions. From the obtained values, the best is selected in a final step.

2.4.5. Techniques from semi infinite optimization

In the previous section, single level reformulations of *optimistic* bilevel programs with *convex lower level problems* have been introduced which allow the application of tools from nonlinear optimization to derive stationary points and also can be used for solution procedures.

An entirely different approach for solving *pessimistic* bilevel programs for *nonconvex lower level problems* has been developed in [WTKR13] and [TWR09]. For these problems, the authors use the connection of bilevel programs and semi-infinite optimization problems (cf. [Ste12] for a treatment of the latter) to derive a one level optimization problem.

The authors consider problems of the type

$$\begin{aligned} \min_x \; & F(x) \\ s.t. \quad & G(x, y^\star) \leq 0 \\ & \forall y^\star \in argmin\, \{f(x,y) | y \in Y\} \\ & x \in X \end{aligned} \tag{2.12}$$

where $F : X \mapsto \mathbb{R}$, and $G, f : X \times Y \mapsto \mathbb{R}$ are continuous.

The authors show that pessimistic bilevel problems of the type of (2.3) can be formulated as an instance of (2.12); however it is necessary that the [LL] feasible region is independent of the [UL] decision variable x.
Since the feasible region of a pessimistic bilevel problem may not be closed (cf. Subsection 2.2.2) the authors introduce ϵ-approximations of the original problem with closed feasible region. The approximate problems can be formulated as semi-infinite optimization problem:

$$\begin{aligned} \min_{x} \; & F(x) \\ s.t. \quad & G(x, y^\star) \leq 0 \\ & x \in X \\ & \forall y^\star \in \mathcal{Y}_\epsilon(x) = \{ y \in Y \,|\, f(x,y) < f(x,y') + \epsilon \;\; \forall y' \in Y \} \end{aligned} \tag{2.13}$$

The authors show that for $\epsilon \to 0$ the solutions of problem (2.13) converge to the optimal value of problem (2.12) if an optimal value is attained[11]. Numerically the ϵ approximations are solved by a tailored discretization scheme which is applied to another reformulation of problem (2.13).

[11] If the optimal value is not attained, more subtle concepts have to be employed for describing of optimality of a solution.

3. Problem analysis and tailored algorithm

The previous chapter illustrated the general properties and severe difficulties of bilevel programs. Despite the efforts put in the design of solution procedures, all standard approaches seem to be unsuitable for the pricing problem at hand: already the mere size, but in particular the nonlinearity of the (risk averse) behavioral pricing problem are serious obstacles for numerical solutions.

Consequently, a tailored solution algorithm is designed which exploits the specific structure of the behavioral pricing problem (1.6) and only relies on the solution of a sequence of linear programs.

In the subsequent chapter, this tailored algorithm will be introduced, preceded by an analysis of some of the structural properties of the optimization problem to be solved.

The text is illustrated by figures generated in MATLAB, partly relying on code for tree manipulation and bilevel problem generation written by Georg Pflug.

3.1. The behavioral pricing problem: Properties

Recall the risk averse stochastic behavioral pricing problem (1.6) introduced in Chapter 1 in symbolic notation

$$\left|\begin{array}{lll} & \min\limits_{k;x} k & \\ [UL] & & \mathcal{A}\left(\Pi(x) - \tilde{S}(k;y^\star)\right) \geq \rho \\ & & x_t \in X_t(y^\star),\, x_t \lhd \mathfrak{G}_t,\, k \lhd \mathfrak{G}_0 \\ [LL] & y^\star \in & argmax \left\{ \mathbb{E}\left(S(k;y)\right) \,\middle|\, \begin{array}{l} y_t \in Y_t \\ y_t \lhd \mathfrak{F}_t \end{array} \right\} \end{array}\right. \tag{3.1}$$

To tackle the problem computationally, we make several modeling assumptions which will be described in detail in the following subsection.

3.1.1. Assumptions on [UL] and [LL]

The probability distributions involved (and the generated filtrations) are finite. In general it is computationally infeasible to solve arbitrary stochastic programs with continuous distributions. A very common approach for computations is to solve deterministic approximations[1] of the original problems. These approximations rely on finite probability distributions and scenario trees to represent the original distributions and stochastic processes. The case study in the last chapter will also follow this path, containing a description of the deterministic approximation used for our application in Subsection 4.3.2. For additional information about the concept of scenario trees we refer to [PP14].
Concerning notation we will use the following conventions: $\mathcal{N}$ will represent the set of nodes of the tree, $\mathcal{S}$ the set of scenarios and $\mathcal{N}_s \subseteq \mathcal{N}$ the nodes belonging to a specific scenario $s \in \mathcal{S}$. $|\mathcal{N}|$, $|\mathcal{S}|$ and $|\mathcal{N}_s|$ denote cardinalities of the respective sets, i.e. the number of nodes, the number of scenarios and the number of nodes in scenario s. By n_- we will denote the predecessor of a given node n. The root node, which belongs to every scenario, will be excluded from the set $\mathcal{N}_s$ by definition.

$\mathcal{A}(\cdot)$ is polyhedral Risk measures and acceptability functionals map (sequences of) random variables onto the real numbers. They are identical in every but one aspect: Risk measures are usually applied on loss distributions whereas acceptability functionals are applied on profit distributions. From this context, it is clear that risk measures are used in connection with minimization whereas acceptability functionals are used mainly in maximization frameworks.

In this work, we chose the context of maximization, applying acceptability functionals on the cumulative payoff distributions. Notice that identical results arise from a minimization setting where risk measures are applied on the cumulative loss distributions.

As will become clear later, the key property of the structure of [UL] exploited by the tailored solution algorithm is that it can be represented by a convex optimization problem, where the the [LL] decision appears only at the right hand side of the constraints.
This is not a strong restriction. The following representation of any convex combination of $AV@R_\alpha$ and $\mathbb{E}$ as linear optimization problem with the required structure (which will be used later) already provides a sufficiently rich class of objective functions for our purposes.

Proposition 3.1. *For discrete probability distribution with masses p_n and values Y_n, the acceptability functional $\gamma\mathbb{E}(Y) + (1-\gamma)AV@R_\alpha(Y)$ can be written as the*

[1] often referred to as deterministic equivalents

solution of the optimization problem

$$\begin{array}{ll} \underset{a,z_i^+,z_i^-}{max} & a + \sum_{n=1}^{N} p_n \left(\gamma z_n^+ - \frac{1+(\alpha-1)\gamma}{\alpha} z_n^- \right) \\ s.t. & Y_n - a = z_n^+ - z_n^- \\ & z_n^+ \geq 0, z_n^- \geq 0,\ a \in \mathbb{R} \end{array}$$

Proof. Using the identities

$$\begin{aligned} Y - a &= [Y-a]^+ - [Y-a]^- \\ AV@R_\alpha(Y) &= \max_a \left\{ a - \frac{1}{\alpha} \mathbb{E}\left[Y-a\right]^- \right\} \end{aligned}$$

we get the reformulations

$$\begin{aligned} &\gamma \mathbb{E}Y + (1-\gamma) AV@R_\alpha(Y) = \\ &= \max_a \left\{ \gamma \mathbb{E}Y + (1-\gamma) \left(a - \frac{1}{\alpha} \mathbb{E}\left[Y-a\right]^- \right) \right\} \\ &= \max_a \left\{ \gamma \mathbb{E}\left([Y-a]^+ - [Y-a]^- \right) + \gamma a + (1-\gamma) \left(a - \frac{1}{\alpha} \mathbb{E}\left[Y-a\right]^- \right) \right\} \\ &= \max_a \left\{ a + \gamma \mathbb{E}\left([Y-a]^+ \right) - \frac{1+(\alpha-1)\gamma}{\alpha} \mathbb{E}\left[Y-a\right]^- \right\} \end{aligned}$$

Setting $z^+ = [Y-a]^+$ and $z^- = [Y-a]^-$ and using discrete distributions leads to the required result. □

A more general criterion for a suitable structure of the [UL] problem is the polyhedrality of $\mathcal{A}(\cdot)$. In the finite setting considered, polyhedral acceptability functionals can be represented by linear optimization problems, where the random variable appears only on the right hand side of a linear equality constraint (for a precise definition, see Definition A.3 in the Appendix), thus guaranteeing the applicability of the tailored solution algorithm.
Proposition 3.1 is only a special case of the more general result that any convex combination of expectation and a polyhedral acceptability functional (such as the $AV@R_\alpha$) is again a polyhedral acceptability functional (cf. [PR07]). For additional information on risk measures see Appendix A.1 and the references therein.

The [UL] constraint $\mathbf{x_t} \in \mathbf{X_t}(\mathbf{y}^\star)$ can be written as linear inequality in $(\mathbf{x_t}, \mathbf{y}^\star)$. Beside the risk constraint discussed above, the only constraint coupling upper- and lower level variables arises from limited market depth (cf. Subsection 4.1.1). For our purposes it suffices to consider the simple type of constraints $y_t^\star + Ax_t \geq b$ for some matrix A and vector b of suitable dimensionality. Again, the crucial feature is the sole appearance of the lower level solution as an additive term in an upper level constraint.

Linear market model. As discussed in Chapter 1, under the linear market model prices are not influenced by the size of market transactions. As one consequence, the cumulative spot market value of a contract $\tilde{S}(k;y)$, is a linear function of the sold quantities. In particular, this assumption also allows us a to formulate linear [UL] portfolio model. Denote the discrete random variable $\tilde{S}(k;y) \in \mathbb{R}^{|\mathcal{S}|}$ by the vector of its realizations for the different scenarios, and with y_n and $\tilde{P}_n$ decision variables and spot prices at node n, i.e. specific components of the vectors y and $\tilde{P}$. Using this, we can write $\tilde{S}(k;y)$ by the following matrix-vector multiplication:

$$\tilde{S}(k;y) = \begin{bmatrix} \sum_{n\in\mathcal{N}_1} \left(\tilde{P}_n - k\right) y_n \\ \vdots \\ \sum_{n\in\mathcal{N}_{|\mathcal{S}|}} \left(\tilde{P}_n - k\right) y_n \end{bmatrix} =: \left(\mathcal{P}_2 - k\mathcal{P}_1\right) y \tag{3.2}$$

where $\mathcal{P}_1, \mathcal{P}_2 \in \mathbb{R}^{|\mathcal{S}|\times|\mathcal{N}|}$ for a total number of scenarios $|\mathcal{S}|$ and a total number of nodes $|\mathcal{N}|$. The n-th entry of the s-th row of $\mathcal{P}_2$ contains the respective spot price if $n \in \mathcal{N}_s$ (i.e. node n belongs to scenario s), and is zero else. Notice that in the case of electricity prices, the entries of $\mathcal{P}_2$ are not necessarily positive. $\mathcal{P}_1$ has the same structure as $\mathcal{P}_2$ except of spot prices being replaced by ones.
A possible extension of the model in the framework introduced consists of prices depending linearly on quantities, which would lead to (convex) quadratic optimization problems for upper and lower level.

$\Pi(\mathbf{x})$ is linear in the real valued decision variable $\mathbf{x}$. Mainly for computational simplicity, the distribution of the [UL] portfolio payoff is assumed to be a linear function of the [UL] decision variables, which are assumed to be real valued. We think that the algorithm introduced in this chapter can be easily extended to the case of a convex portfolio model. Note however that the algorithm does not extend to a nonconvex problem; the classical unit commitment problem encountered in optimal dispatch of power plants for example cannot be treated.

The assumptions of finite probability spaces, polyhedral acceptability measures and linear portfolio model make it possible to formulate the [UL] portfolio optimization problem as linear program.
Analogous to Subsection 1.1.3, we use $\Pi x \in \mathbb{R}^{|\mathcal{S}|}$ as the discrete random variable describing the payoff of the initial portfolio for exercise strategy x. For keeping the notation as simple as possible, we assume nonanticipativity constraints $x_t \lhd \mathfrak{G}_t$ to be already encoded in this variable.
Furthermore, we denote with $a \cdot \mathbf{1}$ a vector in $\mathbb{R}^{|\mathcal{S}|}$where each component equals a and with p_s the probability of scenario s.
Then, as observed in Proposition 3.1, the particular choice of $\mathcal{A} = \gamma\mathbb{E} + (1 -$

$\gamma)AV@R_\alpha$ and using the notation from equation (3.2) leads to the linear program

$$\begin{aligned} \mathcal{A}\left(\Pi(x) - S(k;y^\star)\right) = \underset{a,z^+,z^-}{max} \quad & a + \sum_{s\in\mathcal{S}} p_s \left(\gamma z_s^+ - \frac{1+(\alpha-1)\gamma}{\alpha} z_s^-\right) \\ s.t. \quad & \Pi x - a\cdot\mathbf{1} - \left(\mathcal{P}_2 - k\mathcal{P}_1\right) y^\star = z^+ - z^- \\ & a\in\mathbb{R}, z^+\in\mathbb{R}_+^{|\mathcal{S}|}, z^-\in\mathbb{R}_+^{|\mathcal{S}|} \end{aligned} \tag{3.3}$$

The choice of $\mathcal{A} = AV@R_\alpha$ can be written as[2]

$$\begin{aligned} AV@R_\alpha\left(\Pi(x) - S(k;y^\star)\right) = \underset{a,z}{max} \quad & a - \frac{1}{\alpha}\left(\sum_{s\in\mathcal{S}} p_s z_s\right) \\ s.t. \quad & a\cdot\mathbf{1} - \Pi x - \left(k\mathcal{P}_1 - \mathcal{P}_2\right) y^\star \le z \\ & a\in\mathbb{R}, z\in\mathbb{R}_+^{|\mathcal{S}|} \end{aligned} \tag{3.4}$$

[LL] can be stated as linear program. In the basic model for [LL] (the risk neutral buyer's problem (1.3)), we assumed the [LL] player to maximize the expected payoff from the swing option held, that is, to be risk neutral, which results in a linear optimization problem.
Using the finiteness of the underlying probability distribution, for spot prices P_n the [LL] problem can be explicitly stated as

$$\begin{aligned} \underset{y_0,\dots,y_{|\mathcal{N}|-1}}{max} \quad & \textstyle\sum_{s\in\mathcal{S}} p_s \left(\sum_{n\in\mathcal{N}_s} y_{n_-}(P_n - k)\right) \\ s.t. \quad & y_l \le y_n \le y_u \;\forall n\in\mathcal{N} \\ & y_L \le \textstyle\sum_{n\in\mathcal{N}_s} y_n \le y_U \;\forall s\in\mathcal{S} \end{aligned} \quad . \tag{3.5}$$

Notice that in this formulation the consumption at any node n is determined by the decision variable y_{n_-} at the predecessor n_- of n. This explicitly reflects the information constraint $y_t \lhd \mathfrak{F}_t$, which enforces all consumption decisions to be taken one stage in advance. An equivalent formulation of the [LL] problem (3.5) can be obtained by directly using variables y_n to describe the consumption at a node n, and explicitly enforcing nonanticipativity by the constraint

$$y_n = y_m \;\forall n,m \;:\; n_- = m_-$$

i.e. [LL] consumption at nodes with the same predecessor must be equal. By additionally defining

$$\left(c_1 + kc_2\right)^T y := \sum_{s\in\mathcal{S}} p_s \left(\sum_{n\in\mathcal{N}_s} y_n(P_n - k)\right) \tag{3.6}$$

[2] see for example [RU00]

for a simpler notation, we arrive at a [LL] problem of the form

$$\begin{array}{ll} \underset{y_1,\dots,y_{|\mathcal{N}|}}{max} & (c_1 + kc_2)^T y \\ s.t. & y_l \leq y_n \leq y_u \, \forall n \in \mathcal{N} \\ & y_L \leq \sum_{n \in \mathcal{N}_s} y_n \leq y_U \, \forall s \in \mathcal{S} \\ & y_n = y_m \, \forall n, m \, : \, n_- = m_- \end{array} . \tag{3.7}$$

Remark. Omitting the nonanticipativity constraint in (3.7) means that consumption y is determined for revealed prices, thus in real time. This setup resembles the model for the natural gas storage facility used in Subsection (4.1.1).

Formulating the [LL] problem by (3.7) has the benefit that its decision variables y can be used as the [LL] consumption, which is used as input in the [UL] problem. Thus, we will stick with this notation for the rest of this chapter.
Notice that by arguments analogous to the [UL] case, for polyhedral acceptability functionals also the risk averse buyer's problem (1.6) can be stated as a linear [LL] optimization problem.

The [LL] optimal consumption $\mathbf{y}^\star(\mathbf{k})$ is unique for all choices of $\mathbf{k}$. This assumption is again made for simplicity. Of course, also in the linear programming setting chosen for the [LL] problem, the solution can be non-unique. In this case it might be necessary to derive *pessimistic* solutions for the [LL] problem. However, this situation is *improbable*, as shown in [Dem02]:

Theorem 3.2. *([Dem02], Th. 7.1) Consider a linear optimization problem*

$$min\{c^T x | Ax = b, x \geq 0\}$$

having a parametric objective function. Then, for Lebesgue almost all $c \in \mathbb{R}^n$ this problem has a unique optimal solution.

We think this argument suffices for assuming uniqueness of the [LL] solution for the analysis of real world problems.
Note that from a conceptional point of view, the algorithm developed in this chapter can be easily extended to cover the pessimistic case: In the cases where parametric linear programming is used, all of the lower level solutions have to be determined. The face algorithm from [ZYL12] seems a promising method for this specific problem. An alternative would be to consider perturbations of the [LL] solutions at every step.
In practice, checking for multiple solutions may lead to a severe increase of computational effort.

3.1.2. The problem for real data

Under the assumptions above, for real data both [UL] and [LL] can be stated as large scale linear programs - provided the other player's decision is constant. To see this, we make the following observations:

The [UL] problem contains constraints of three different structures:

- Constraints on portfolio management (such as production or market actions) containing only [UL] decision variables, can be encoded by standard inequality constraints

$$\tilde{A}x \leq \tilde{b}\,.$$

- For modeling limited market depth (i.e. constraints on the maximal cumulative quantity which can be traded at spot market prices), at every node n the sum of the net spot purchase $\bar{x}_n$ of [UL] and the [LL] consumption y_n is constrained[3] above by some quantity $\bar{m}$ and below by some $\underline{m}$. Denote with $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_{|\mathcal{N}|})$ the cumulative [UL] consumption per node, i.e. a negative value of $\bar{x}_n$ amounts to a long and a positive value of $\bar{x}_n$ to a short position in respective spot market at node n. Then, the sum of this [UL] purchase $\bar{x}$ and the [LL] consumption, $y^\star = (y^\star_1, \ldots, y^\star_{|\mathcal{N}|})$, gives the net spot market trading[4] for [UL] which is constrained by $\underline{m} \leq \bar{x} + y^\star \leq \bar{m}$.
  However, for notational convenience, we will only use the one sided expression

$$\bar{x} + y^\star \leq m_{depth}\ , \tag{3.8}$$

  which will be sufficent to illustrate the mechanism of the tailored algorithm. Clearly, this naturally extends to the case of an additional constraint on spot market trading from below, which is used in the implementations of the case study in Subsection (4.4.2).

- As observed above, for a polyhedral risk measure, the acceptability of a portfolio can be written as linear program.

For further structural investigations, we assume the objective to be the average value-at-risk, $AV@R_\alpha$, which leads to representations as linear programs of the type

[3] This approach to modeling market depth has the benefit of staying in the linear-programming framework. However, volume dependent prices (i.e. a nonlinear market model) or penalty approaches may be more accurate descriptions of the market situation, leading to a significantly more complicated model.

[4] For additional illustration consider the following example: For a given node n, let [UL] produce $x_{n,1} = 30$ units of energy, store $x_{n,2} = 10$ units and spend additional $x_{n,3} = 15$ units on existing supply obligations. Using the conventions above, the cumulative consumption is then given by $\bar{x}_n = -x_{n,1} + x_{n,2} + x_{n,3} = -5$. Let the [LL] consumption at n be $y^\star_n = 10$. Then, the net spot market trading of [UL] amounts to $\bar{x}_n + y^\star_n = 5$ units of energy, to be bought at the spot market.

of problem (3.4). The problem structure for the case of an objective given by a combination of $AV@R_\alpha$ and $\mathbb{E}$ can be similarly derived based on the description (3.3).

Using the notation from problem (3.4) and including additional constraints on portfolio management and market depth, the linear program giving the maximal acceptability is

$$\max_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x} + y^\star \leq m_{depth} \\ a \cdot \mathbf{1} - \Pi x - (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star \leq z \end{array} \right\} \tag{3.9}$$

for $\tilde{c} = -1/\alpha \left[p_1, \ldots, p_{|\mathcal{S}|}\right]^T$. Requiring a minimal level of acceptability, ρ, amounts to constraining this optimal value from below. By virtue of the relation

$$\max_x c^T x \geq \rho \;\Leftrightarrow\; \exists x :\; c^T x \geq \rho$$

the acceptability constraint can be transformed into linear inequality constraints. Thus [UL] can be written as

$$\begin{aligned} &\min_k \left\{ k \;\middle|\; \max_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x} + y^\star \leq m_{depth} \\ a \cdot \mathbf{1} - \Pi x - (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star \leq z \end{array} \right\} \geq \rho \right\} \\ &\quad = \min_{k;x;a,z} \left\{ k \;\middle|\; \begin{array}{l} a + \tilde{c}^T z \geq \rho \\ \tilde{A}x \leq \tilde{b} \\ \bar{x} + y^\star \leq m_{depth} \\ a \cdot \mathbf{1} - \Pi x - (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star \leq z \end{array} \right\} \end{aligned} \tag{3.10}$$

.

For suitable choice of matrices A_1, A_2, A_3 and vector b , all of these constraints can be expressed by the vector valued constraint

$$A_1 \begin{pmatrix} a \\ x \\ z \end{pmatrix} + A_2 y^\star + k A_3 y^\star \leq b \,.$$

The lower level constraints encode the requirements on short term (i.e. daily, hourly) and cumulative consumption, thus containing entirely linear constraints which are independent of [UL] decisions.

Based on these preliminaries, the risk averse behavioral pricing problem for real data can be written as

$$\left| \begin{array}{ll} [UL] & \min\limits_{k;a,x,z} k \\ & \qquad A_1 \begin{pmatrix} a \\ x \\ z \end{pmatrix} + A_2 y^\star + k A_3 y^\star \leq b \\ [LL] & \quad y^\star = \operatorname*{Argmax}\limits_{y=y_1,\ldots,y_{|\mathcal{N}|}} \left\{ (c_1 + k c_2)^T y \;\middle|\; \begin{array}{c} y_l \leq y_n \leq y_u \,\forall n \in \mathcal{N} \\ y_L \leq \sum_{n \in \mathcal{N}_s} y_n \leq y_U \,\forall s \in \mathcal{S} \\ y_n = y_m \,\forall n, m \,:\, n_- = m_- \end{array} \right\} \end{array} \right. \tag{3.11}$$

Observations: As mentioned above, the real-world problem (3.11) can be seen as multistage stochastic program [UL], where one of the risk factors is given by the solution of a multistage stochastic optimization problem, [LL]. In particular, both [UL] and [LL] are large linear programs when the decision variables from the other level are fixed. The problem as a whole is a nonlinear bilevel program which is $\mathcal{NP} - hard$ (cf. Subsection 2.4.1).
A particular difficulty arises from [UL] coupling constraints from the acceptability requirement (3.9) and possibly also from enforcing the constraints on market transactions by (3.8). These constraints can lead to disconnected set of feasible strike prices which can indeed be observed in practice (see Figure 3.1 as an example). Amongst others, this property is a central drawback for descent algorithms as presented in Subsection 2.4.3.

As a contrast to these difficulties, problem (3.11) provides one significant benefit: the [LL] problem is only connected to the [UL] problem by a single one dimensional variable, the strike price k.

This property is emphasized by the following equivalent representation of (3.11):

$$\left|\begin{array}{lll} & \min\limits_{k} k & \\ [UL] & \underset{x;a,z}{max}\left\{ a+\tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x}+y^\star \leq m_{depth} \\ a\cdot\mathbf{1}-\Pi x-(k\mathcal{P}_1-\mathcal{P}_2)\,y^\star \leq z \end{array} \right\} \geq \rho & \\ [LL] & y^\star = \underset{y=y_1,\ldots,y_{|\mathcal{N}|}}{Argmax}\left\{ (c_1+kc_2)^T y \;\middle|\; \begin{array}{c} y_l \leq y_n \leq y_u \,\forall n\in\mathcal{N} \\ y_L \leq \sum_{n\in\mathcal{N}_s} y_n \leq y_U \,\forall s\in\mathcal{S} \\ y_n = y_m \,\forall n,m \,:\, n_- = m_- \end{array} \right\} & \end{array}\right. \tag{3.12}$$

The key observation in this representation is that the problem can be lead back to a one dimensional optimization problem in k. To check feasibility, the optimal *value* of one linear program ([UL]) has to be computed which requires the optimal *solution* of a second linear program ([LL]). The algorithm designed for solving the problem for real data consists of iterative linear approximations of the optimal value of [UL] as a function of k.

3.2. Properties of the [UL] problem

The following section highlights some properties of the [UL] problem which will be used to design the solution algorithm.
For purposes of illustrating and testing the algorithm, small[5] test problems were solved in MATLAB. The implementation uses YALMIP[6] and CPLEX 12.2.0.0.

[5] The tree considered consists of around 50 stages and 200 scenarios. The [UL] portfolio is not explicitly modeled for these test examples. Instead, [UL] and [LL] are subjected to different prices.

[6] [Löf04]

In the setup used for computations, the [UL] portfolio consists of the purchase of energy futures contracts, delivering a constant amount of energy over several stages at a fixed price, which can be traded in the root node only. The spot price processes for [UL] and [LL] differ; the swing option specifies constraints on as well the daily as the cumulative [LL] consumption above and below by nonzero constants.

3.2.1. Properties of the [UL] feasible region

Figure 3.1 illustrates the central aspects of the behavioral pricing problem.

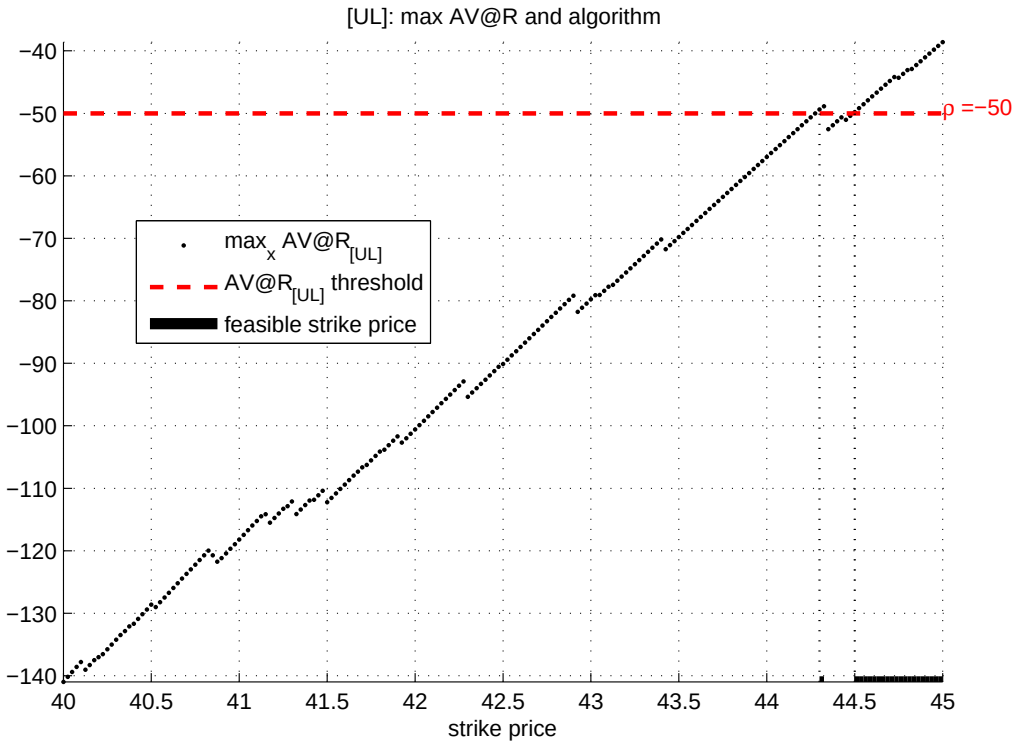


Figure 3.1.: One dimensional representation of the bilevel problem. Notice the disconnected set of feasible strike prices marked by a thick black line on the x-axis.

To generate the figure, for fixed strike price the corresponding lower level solution is computed, based on which the best possible [UL] decision is computed, giving the optimal [UL] acceptability. The acceptability (here $AV@R_\alpha$) of the profit distribution resulting from this decision is depicted as black dot. Iterating this computation for various choices of k shows the [UL] acceptability as a function of the strike price k.
Strike prices are considered [UL] feasible where the acceptability level is above the

required threshold ρ (dashed red line). The distinctive nonmonotonicity of the acceptability level in k in this case leads to a disconnected set of feasible strike prices. In this representation, the solution can be found simply by choosing the leftmost strike price which is feasible.
Notice the computational difficulties for obtaining this one dimensional expression: each black dot requires the solution of two linear programs, leading to a total effort of solving about 800 linear programs to generate Figure 3.1.

3.2.2. Implicit treatment of [UL] portfolio optimization

For getting a handle on the [UL] acceptability directly as a function of the strike price k, we introduce some additional notation:

The acceptability of the [UL] portfolio is given by the optimal value of an optimization problem,

$$\Phi(k) := \max_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x} + y^\star(k) \leq m_{depth} \\ a \cdot \mathbf{1} - \Pi x - (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) \leq z \end{array} \right\} \tag{3.13}$$

which is stated as a function of the strike price k. Note that this value depends on the corresponding [LL] solution $y^\star(k)$ (a discrete random variable).
With this notation, problem (3.11) is equivalent to

$$\left| \begin{array}{ll} [UL] & \min\limits_{k} k \\ & \Phi(k) \geq \rho \\ [LL] & y^\star = \underset{y=y_1,\ldots,y_{|\mathcal{N}|}}{Argmax} \left\{ (c_1 + kc_2)^T y \;\middle|\; \begin{array}{c} y_l \leq y_n \leq y_u \, \forall n \in \mathcal{N} \\ y_L \leq \sum_{n\in\mathcal{N}_s} y_n \leq y_U \, \forall s \in \mathcal{S} \\ y_n = y_m \, \forall n, m \,:\, n_- = m_- \end{array} \right\} \end{array} \right. . \tag{3.14}$$

Observe that (3.14) implies that the set of feasible k is connected for all choices of ρ if and only if the functional $\Phi(k)$ is quasiconcave.

3.2.3. Properties of the [UL] acceptability $\Phi(k)$

$\Phi(k)$ can be read as the optimal value of an optimization problem as a function of a perturbation of its constraints, which is controlled by its argument k. Our ultimate goal will be to construct a sequence of approximations of $\Phi(k)$, which we will use to design our tailored algorithm.
Before this, we present a short discussion of the properties of $\Phi(k)$.

As the solution of a maximization problem depending on its right hand side, it is immediately clear that $\Phi(k)$ is nondecreasing in k if the feasible region increases in

k, and nonincreasing if it decreases in k. Recalling the definition of $\Phi(k)$ by the optimization problem (3.13), we observe that k changes the feasible region in two ways:

- directly, by the appearance of k in the term $(k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star$, which is the vector of net [UL] cash flows from the contract for each single scenario. For constant $y^\star$, this expression influences the linear parts of $\Phi(k)$.
- indirectly, by influencing the [LL] exercise strategy $y^\star(k)$, which in turn influences the jumps of $\Phi(k)$.

Since $y^\star(k)$ is the solution of a linear program, whose objective coefficients depend on k, it is a piecewise constant function of k . This implies that $\Phi(k)$ is piecewise linear, piecewise concave and nondecreasing: since the elements of $\mathcal{P}_1$ are either 0 or 1, and assuming that $y_l \geq 0$, for a constant [LL] solution an increase in k directly translates to an increase of the feasible region of the optimization problem (3.13). As the optimal value of (3.13), $\Phi(k)$ is thus piecewise linear, piecewise concave and nondecreasing on intervals of constant [LL] solution (cf. Figure (3.1)).

The intervals of piecewise linear $\Phi(k)$ are separated by jumps, which are triggered by changes in the [LL] solution $y^\star(k)$. Motivated by Figure (3.1), one would expect that the decreasing consumption for increasing strike price k leads to negative jumps only.
However, the analysis of the jumps of $\Phi(k)$ is more complicated than that:

A decrease in the [LL] consumption does in general not lead to a decrease in the size of the feasible region of problem (3.13) due to the following properties:

- *[LL] scenario coupling: constraints on nonanticipativity and cumulative consumption:* The third constraint in the optimization problem (3.13) contains the term $(k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k)$, i.e. the vector of net [UL] profits from the contract for different scenarios. The structure of (3.13) shows that any decrease in $(k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k)$ leads to a decrease of the feasible region and thus to a downward jump in $\Phi(k)$.
  [LL] however maximizes the expectation over all scenarios and is subject to nonanticipativity constraints. As a result, the [LL] decision for consumption at two given scenarios, which are coupled by nonanticipativity or cumulative consumption constraints, may be to disclaim a (relatively low or improbable) profit in one in exchange for avoiding a (relatively high or probable) loss in the other scenario. Whereas the corresponding *expected* loss is reduced for [LL], the [UL] profit is increased in one and decreased in another *scenario*, therefore again both increasing and decreasing the feasible set. Consequently, a decrease in [LL] consumption is not directly linked to the size of the feasible region of (3.13).
- *Different spot price processes: $\tilde{P}_n$ for [UL] and P_n for [LL]*: A similar argument as for the nonanticipativity constraint can be made for different spot prices for [UL] and [LL]: In this case it may well be optimal for [LL] to reduce con-

sumption at a given node which simultaneously relieves [UL] from unfavorable delivery obligations, again increasing the feasible region in problem (3.13).

Moreover, an increase in the strike price does in general not even imply a decrease in [LL] consumption:

- The [LL] consumption-plan at a given single node n, $y_n^\star(k)$, is in general not monotonous in k. For an illustration of this effect see Figure 3.2, which depicts the [LL] consumption for the same problem used for other illustrations in this chapter.
  The reason for this behavior probably results from "asymmetries" (different probabilities / and or different number of successors) in the scenario tree. If changes in k change the [LL] objective coefficients in a uniform manner (i.e. keeping the order from largest to smallest constant, but only increasing the number of coefficients with negative signs), then the relative cost coefficients do not change their order. Consequently, changes in the value of a decision variable appear if and only if the sign of the associated cost coefficient changes, resulting in nonincreasing $y_n^\star(k)$ in k.
  If however not just the signs but also the order of the relative cost coefficients is changed by k (which can happen when the single decisions carry different "weights"), then the solution $y_n^\star(k)$ shows more complex behavior.

Altogether, we collected several structural properties which contradict the assumption that the feasible region of problem (3.13) decreases with increasing k. This complicates an exact description of the behavior of $\Phi(k)$ at its jumps.
Notice that these properties imply that $\Phi(k)$ is in general neither quasiconvex nor quasiconcave in k and therefore the feasible set of strike prices is disconnected.

However, for the special case where these conditions are removed from (3.14), we can immediately derive the following properties:

Lemma 3.3. *Consider the modification of decision problem* (3.14)*, using the original expressions for the [LL] objective (equation* (3.6)*) and [UL] constraints (equation* (3.2)*),*

$$\left|\begin{array}{ll} & \min\limits_{k} k \\ [UL] & \Phi(k) = \max\limits_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ a \cdot \mathbf{1} - \Pi x - \begin{pmatrix} \sum_{n\in\mathcal{N}_1} y_n(k-P_n) \\ \vdots \\ \sum_{n\in\mathcal{N}_{|\mathcal{S}|}} y_n(k-P_n) \end{pmatrix} \leq z \end{array} \right\} \geq \rho \\ [LL] & y^\star = \underset{y=y_1,\ldots,y_{|\mathcal{N}|}}{Argmax} \left\{ \sum_{s\in\mathcal{S}} p_s \left(\sum_{n\in\mathcal{N}_s} y_n(P_n - k) \right) \;\middle|\; y_l \leq y_n \leq y_u \; \forall n \in \mathcal{N} \right\} \end{array}\right. \tag{3.15}$$

where [UL] and [LL] are assumed to use the same spot price processes, market depth constraints as well as nonanticipativity and cumulative consumption constraints have

been omitted.
Then, $\Phi(k)$ is piecewise linear, continuous and concave.

Proof. Without nonanticipativity and cumulative consumption constraints, [LL] consumption at each single node of the tree can be optimized independently of the decisions on other nodes. In particular, the [LL] of the modified problem (3.15) can be decomposed into $|\mathcal{S}|$ separate deterministic optimization problems:

$$\begin{aligned} & \max_{y} \left\{ \sum_{s\in\mathcal{S}} p_s \left(\sum_{n\in\mathcal{N}_s} y_n(P_n - k) \right) \,\middle|\, y_l \le y_n \le y_u \,\forall n \in \mathcal{N} \right\} \\ = \; & \sum_{s\in\mathcal{S}} p_s \max_{(y_n)_{n\in\mathcal{N}_s}} \left\{ \sum_{n\in\mathcal{N}_s} y_n(P_n - k) \,\middle|\, y_l \le y_n \le y_u \,\forall n \in \mathcal{N} \right\} \end{aligned}$$

Each of these problems amounts to maximizing $\sum_{n\in\mathcal{N}_s} y_n(P_n - k)$, i.e. the objective of the [LL] payoff for a single scenario. As a result, the optimal value $\sum_{n\in\mathcal{N}_s} y_n^\star(k)(P_n - k)$ is nonincreasing, piecewise linear and continuous in k.
If [UL] and [LL] use identical spot price processes, each of these objective values corresponds to the negative of net [UL] profits from the contract at a given scenario s (i.e. $(\sum_{n\in\mathcal{N}_s} y_n^\star(k)(k - P_n))$), which then must be a nondecreasing, piecewise linear continuous function of k. As a result, an increase in k amounts to a linear perturbation of the right hand side of constraints in problem (3.15). As the solution of this optimization problem, $\Phi(k)$ is piecewise linear increasing, continuous and concave in k.

□

3.2.4. Approximation of the [UL] acceptability $\Phi(k)$

As the solution of a parametric optimization problem, $\Phi(k)$ is numerically difficult to deal with. The following propositions describe an approximation from above, which is much simpler to handle:

Proposition 3.4. *Let $\Phi(k)$ be the solution of* (3.13)*. Then, for every choice of k and the corresponding [LL] solution $y^\star(k)$, an overestimator is given by*

$$\hat{\Phi}_0(k) := \Phi(k_0) + \begin{bmatrix} \lambda_\zeta(k_0) \\ \lambda_y(k_0) \end{bmatrix}^T \begin{bmatrix} (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) - (k_0\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k_0) \\ y^\star(k_0) - y^\star(k) \end{bmatrix} \ge \Phi(k) \tag{3.16}$$

and $\lambda_\zeta(k_0)$ and $\lambda_y(k_0)$ are given by the Lagrange multipliers associated to the constraints

$$\begin{aligned} a\cdot\mathbf{1} - \Pi x - (k_0\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k_0) \le z \quad & \textit{(multiplier} \quad \lambda_\zeta(k_0)\textit{)} \\ \bar{x} + y^\star(k_0) \le m_{depth} \quad & \textit{(multiplier} \quad \lambda_y(k_0)\textit{)} \end{aligned}$$

in the optimum of problem (3.13) *at k_0 and $y^\star(k_0)$.*

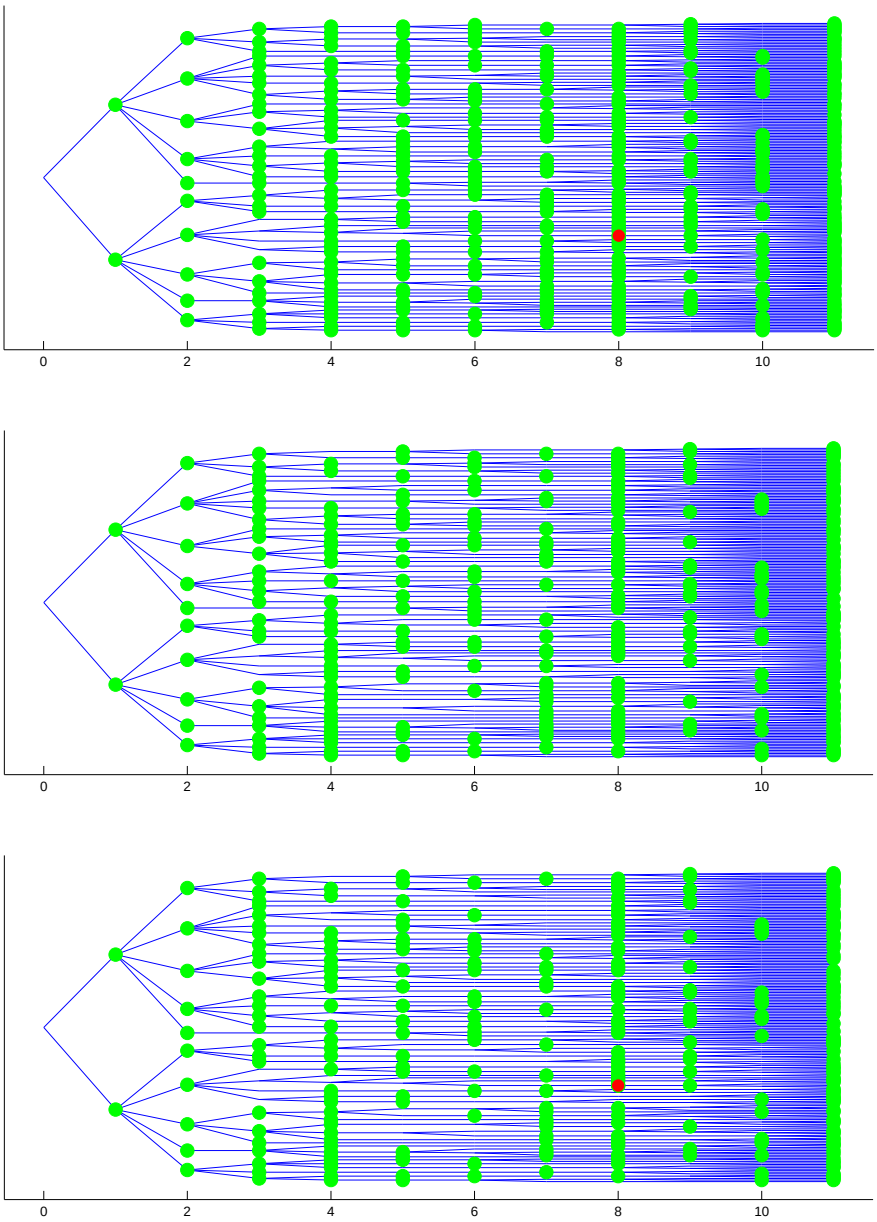


Figure 3.2.: Illustration of the nonmonotonicity of $y_n^\star(k)$.
Top: [LL] consumption at $k = 40.7$, Mid: [LL] consumption at $k = 44$, Bottom: consumption at $k = 45$. The dot size marks the consumption at a specific node. The consumption at a specific node at stage eight is highlighted with red color. Notice that while consumption at a low price exists, it vanishes at a medium strike price but appears again at a high strike price.

Proof. Recall the optimization problem defining $\Phi(k)$ by

$$\Phi(k) := \max_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x} + y^\star(k) \leq m_{depth} \\ a \cdot \mathbf{1} - \Pi x - (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) \leq z \end{array} \right\}$$

Introducing the additional function $\zeta(k) = (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k)$, we can rewrite the same problem as

$$v(\zeta(k), y^\star(k)) = \max_{x;a,z} \left\{ a + \tilde{c}^T z \;\middle|\; \begin{array}{l} \tilde{A}x \leq \tilde{b} \\ \bar{x} - m_{depth} \leq -y^\star(k) \\ a \cdot \mathbf{1} - \Pi x - z \leq \zeta(k) \end{array} \right\}$$

where $v(\zeta(k), y^\star(k))$ denotes the optimal value function with respect to the right hand side $\zeta(k)$ and $y^\star(k)$ of the respective constraints. Notice that by construction, the value function equals the [UL] acceptability level for all choices of k i.e. $v(\zeta(k), y^\star(k)) = \Phi(k)\, \forall k$.
Consequently, a change of a given k_0 to some k triggers deviations $(\zeta(k) - \zeta(k_0))$ and $(-y^\star(k) + y^\star(k_0))$, which can be interpreted as perturbations of the right hand side of inequality constraints of a convex program. Since $v\left(\zeta(k), y^\star(k)\right)$ is the solution of a convex program and equals $\Phi(k)$, the upper subgradient inequality (3.16) holds. In particular, the upper subgradient $[\lambda_\zeta(k_0)\ \lambda_y(k_0)]^T$ is given by the dual variables associated to the perturbed constraints in the optimum ([RW98], Theorem 10.13). □

This approximation is exact locally around k_0, which is established next.

Proposition 3.5. *Under the assumptions of Proposition 3.4, $\forall k_0$ there exists a closed interval $[\kappa_a, \kappa_b] \ni k_0\,, |\kappa_a - \kappa_b| > \delta$ on which $\Phi(k)$ is continuous, nondecreasing and piecewise linear in k and $y^\star(k)$ is constant.*

In particular, there exists a subinterval $[k_a, k_b] \subseteq [\kappa_a, \kappa_b]$, where $\hat{\Phi}_0(k) = \Phi(k)\, \forall k \in [k_a, k_b]$.

Proof. Since $y^\star(k)$ is the solution of a parametric linear program, it is piecewise constant in k on closed intervals $[\kappa_1, \kappa_2]$ with disjoint interiors. Since by assumption $y^\star(k)$ is unique, $\kappa_1 \neq \kappa_2$, thus $\exists \delta : |\kappa_1 - \kappa_2| > \delta$.
Let $y^\star \in [\kappa_1, \kappa_2]$ be constant. Then, $\Phi(k)$ is the optimal value of a perturbation in k of the right hand side of an inequality constraint of a linear program. Therefore, $\Phi(k)$ is continuous, monotonous and piecewise linear in k on $[\kappa_1, \kappa_2]$. From piecewise linearity of $\Phi(k)$, it follows that there exists a decomposition into a finite number of intervals of nonzero length on each of which $\Phi(k)$ is linear. In particular, for every choice of k_0 there exists an interval $[k_a, k_b]$ containing k_0, where $\Phi(k)$ is linear and therefore any linearization $\hat{\Phi}_0(k)$ is exact for $k \in [k_a, k_b]$. □

In the following, we will investigate the behavior of a sequence of upper approximations $\hat{\Phi}_i(k)$ of $\Phi(k)$, which are generated by replacing k_0 by a sequence of strike prices k_i, i.e.

$$\hat{\Phi}_i(k) := \Phi(k_i) + \begin{bmatrix} \lambda_\zeta(k_i) \\ \lambda_y(k_i) \end{bmatrix}^T \begin{bmatrix} (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) - (k_i\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k_i) \\ y^\star(k_i) - y^\star(k) \end{bmatrix} .$$

Definition 3.6. Define with $k^\star$ the solution of the bilevel pricing problem, i.e. the minimal strike price which is [UL] acceptable:

$$k^\star := min\left\{v \,|\Phi(v) \geq \rho\right\} ,$$

and with $\hat{k}_{i+1}$ the minimal strike price, where the approximated acceptability $\hat{\Phi}_i\,(k)$ is above the threshold, i.e.

$$\hat{k}_{i+1} := min\left\{v \,\middle|\hat{\Phi}_i(v) \geq \rho\right\} .$$

Notice that for every i, $\hat{k}_i$ is a lower bound on the optimal strike price, since due to the inequality $\Phi(k) \leq \hat{\Phi}_i(k)$ it holds that $\hat{k}_i \leq k^\star$.
Finally, we derive a criterion for coincidence of $\hat{k}_i$ and $k^\star$:

Proposition 3.7. *If $\hat{k}_{i+1} = \hat{k}_i$, then $\hat{k}_i$ is the minimal strike price $k^\star$, i.e.*

$$\hat{k}_{i+1} = \hat{k}_i \Rightarrow \hat{k}_i = k^\star$$

In other words, the optimal strike price $k^\star$ is given by any fix point of the mapping

$$\kappa \mapsto min\left\{v \,\middle|\Phi(\kappa) + \begin{bmatrix} \lambda_\zeta(\kappa) \\ \lambda_y(\kappa) \end{bmatrix}^T \begin{bmatrix} (v\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(v) - (\kappa\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(\kappa) \\ y^\star(\kappa) - y^\star(v) \end{bmatrix} \geq \rho\right\} .$$

Proof. Assume $\hat{k}_{i+1} = \hat{k}_i$. From Proposition 3.5 we know that

$$\exists\,[\kappa_1, \kappa_2] \ni \hat{k}_i :\ \Phi(k) = \hat{\Phi}_i(k)\,\forall k \in [\kappa_1, \kappa_2] .$$

$\hat{k}_{i+1}$ is on the one hand determined by the intersection of $\hat{\Phi}_i(k)$ with the acceptability threshold, on the other hand it is in the interval where the approximation is exact since $\hat{k}_{i+1} = \hat{k}_i$. Therefore, we have that $\rho = \hat{\Phi}_i(\hat{k}_{i+1}) = \Phi(\hat{k}_{i+1})$. □

3.3. A Newton type algorithm exploiting the [UL] problem structure

The tailored algorithm computes a sequence of solutions of auxiliary problems

$$\hat{k}_{i+1} := min\left\{v \,\middle|\, \hat{\Phi}_i(v) \geq \rho\right\}$$

which are based on a sequence of overestimators $\hat{\Phi}_i(k)$. The solutions $\hat{k}_i$ are lower bounds for the minimal acceptable strike price $k^\star$, which can be iteratively improved by Algorithm 3.1.

Interpreting the subgradient inequality (3.16) as first order approximation of the function $\Phi(k) - \rho$, the algorithm resembles the Newton algorithm for finding zeros.

Algorithm 3.1 Newton-type algorithm for solving problem (3.14)

```
set initial strike price k_1, set i = 1
compute y*(k_1) and Φ(k_1), λ_ζ(k_1), λ_y(k_1)
set Φ̃ = Φ(k_1)
while ρ − Φ̃ > err
    while Φ̂_i(k) ≤ ρ
        k̃ = k
        k = update(k)
        compute y*(k), Φ̂_i(k)
    end while
    obtain k_{i+1} by solving equation (3.17) for k
    compute y*(k_{i+1}) and Φ(k_{i+1}), λ_ζ(k_{i+1}), λ_y(k_{i+1})
    set Φ̃ = Φ(k_{i+1}), i = i + 1
end while
return k_i
```

Algorithm (3.1) is centered around two nested while loops:

- The outer loop updates the overestimator $\hat{\Phi}_i(k)$ with respect to i, using the exact [UL] solution, the [LL] solution and the subgradients at a specific strike price k_i.
- The inner loop finds the intersection of the overestimator with the threshold i.e. the solution of the problem

$$min\left\{k \,\middle|\, \Phi(k_i) + \begin{bmatrix}\lambda_\zeta(k_i)\\ \lambda_y(k_i)\end{bmatrix}^T \begin{bmatrix} (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) - (k_i\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k_i) \\ y^\star(k_i) - y^\star(k) \\ k \geq k_i \end{bmatrix} \geq \rho \right\}.$$

This is achieved by iteratively testing if the overestimator is above the acceptability threshold. If this happens, the intersection of $\hat{\Phi}_i(k)$ with ρ must have been in the interval where $y^\star(k) = y^\star(\tilde{k})$. Therefore, $\hat{k}_{i+1}$ can be obtained by solving the linear equation

$$\Phi(k_i) + \begin{bmatrix}\lambda_\zeta(k_i)\\ \lambda_y(k_i)\end{bmatrix}^T \begin{bmatrix} (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(\tilde{k}) - (k_i\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k_i) \\ y^\star(k_i) - y^\star(\tilde{k}) \end{bmatrix} = \rho \quad (3.17)$$

for k.
Otherwise a new strike price is computed by the procedure $update(k)$. To guarantee optimality of the solution, the function $update(k)$ has to increase k by the smallest amount necessary to change the [LL] solution, i.e.

$$update(k) = max\left\{\kappa | \hat{\Phi}_i(\kappa) = \hat{\Phi}_i(k)\right\} + \epsilon$$

where ϵ denotes the tiny disturbance needed to avoid the old basis still being optimal.

Numerical values for $update(k)$ can be obtained from sensitivity analysis which can usually be performed by standard linear programming software. However in practice, it may be sufficient to rely on constant increments such as $k = k + 0.01$, in particular if the objective is a price which usually has a minimal increment, or tick size.
Note that iterative evaluation of $update(k)$ in the inner loop amounts to a crude parametric programming algorithm. Essentially, the algorithm as presented in (3.1) requires the solution of [LL] as parametric linear program for parameter values between k_1 and the optimal $k^\star$, which puts the main computational burden to the inner loop. Figure 3.3 illustrates steps performed at an implementation of Algorithm 3.1, and the rather high number of [LL] solutions required.

Proposition 3.8. *Given any starting value $k_1 \leq k^\star$ and a sufficiently small choice of the tolerance err, the Newton-type algorithm 3.1 terminates in a finite number of steps at the optimal strike price $k^\star$.*

Proof. Proposition 3.4 implies that $\hat{\Phi}_1(k) \geq \Phi(k)$ and also $k_1 \leq k^\star$; Proposition 3.5 that there exists an interval of exact approximation, $[a_1, b_1] \ni k_1$, where $\hat{\Phi}_1(k) = \Phi(k)$ and $\Phi(k)$ is nondecreasing.
Recall that $\Phi(k)$ is piecewise linear in k, thus the search interval $[k_1, k^\star]$ consists of a finite number of intervals on which any linear approximation is exact. We will show that the algorithm systematically eliminates these intervals until $k_i = k^\star$.
Assume that $k^\star \in [a_1, b_1]$. Then, $\hat{k}_2 = min\left\{v \,\middle|\, \hat{\Phi}_1(v) \geq \rho\right\} = min\left\{v \,|\Phi_1(v) \geq \rho\right\} = k^\star$, i.e. the algorithm terminates at the exact solution $k^\star$.
Otherwise, monotonicity of $\hat{\Phi}(k)$ on $[a_1, b_1]$ implies that $\hat{k}_1 > k_1$ and $\hat{k}_1 \notin [a_1, b_1]$. Thus, at least the interval of exact approximation (i.e. on which $\Phi(k)$ is exactly approximated by $\hat{\Phi}_1(k)$) is removed from the search interval $[k_1, k^\star]$.
Repeating this step again either results in obtaining the optimal value or in cutting

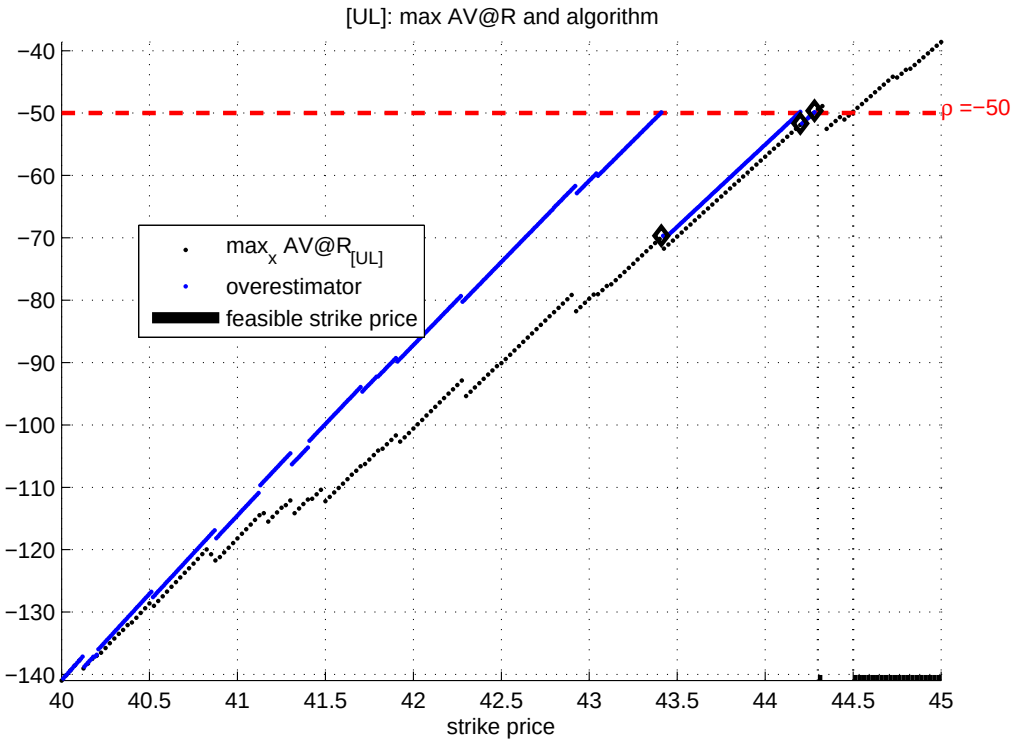


Figure 3.3.: Illustration of the Newton-type algorithm. For given strike price, black diamonds represent evaluations of the whole problem ([UL] and [LL]) at iterations $2, 3$ and 4, returning the subderivative of the [UL] acceptability. Blue dots mark evaluations of the [LL] problem needed for obtaining $\hat{k}_i$. In the depicted case, convergence was reached after 430 solutions of [LL] and 4 solutions of [UL].

away at least one additional inerval of exact approximation from the search interval. Since the initial interval $[k_1, k^\star]$ consists of only a finite number of intervals of exact approximation, the algorithm will arrive at the optimal value in a finite number of steps.

□

Note that to some extent, the complications of the problem have been shifted to the mapping

$$\begin{aligned} \mathbb{R}^+ &\to \mathbb{R}^{|\mathcal{S}|+|\mathcal{N}|} \\ k &\mapsto \begin{pmatrix} (k\mathcal{P}_1 - \mathcal{P}_2)\, y^\star(k) \\ y^\star(k) \end{pmatrix}, \end{aligned}$$

which has to be computed for each evaluation of $\hat{\Phi}_i(k)$. Since $y^\star(k)$ is the solution of a linear program, obtaining the exact optimal solution by Algorithm 3.1 (without

modifications or approximations) requires the solution of the parametric [LL] program on the interval $[k_1, k^\star]$.
The effort necessary for this strongly depends on the concrete structure of the program. In the worst case, the number of different values taken by the solution may be exponential in the dimension of the problem (cf. [BT97], p 220). The modification in Section 3.4 focuses on designing computationally cheap methods to find a starting value k_1 which is already very close to the optimal solution.
In practice, using warm starts[7] for the iterative solutions of [LL] in the inner loop of Algorithm 3.1 turns out to result in rather efficient computations, since consecutive problems are very similar.
Note that convergence speed is also influenced by the size of intervals of constant [UL] solution (with respect to k), since it determines the size of the area where the approximation $\hat{\Phi}_i(k)$ is exact. Nevertheless, the solution of the whole problem has to be determined only once (to get $\Phi(k)$,$\lambda_y(k)$ and $\lambda_\zeta(k)$) per cut. This makes the method particularly suitable for application at large [UL] problems.
The main conceptual benefit of this algorithm is that it operates on the infeasible set (computing the minimal strike price as the lower boundary of the feasible set). Placing the cuts in the infeasible set avoids troubles with the disconnectedness of the feasible set.

3.4. Replacing $\hat{\mathbf{\Phi}}_\mathbf{i}(\mathrm{k})$ by a linear model

Although the dependency between k and $y^\star(k)$ is rather difficult to model and the functional $\hat{\Phi}_i(k)$ shows discouraging properties, it can be approximated from above by a linear function quite well.
The linear approximation can be obtained in two steps: First, slopes Δ_j are fitted between the computed exact value $\Phi(k_i)$ and evaluations of the overestimator at several probes p_j, $\hat{\Phi}_i(p_j)$. The set of probes at a given iteration i, $probes_i = \{p_1, \ldots, p_J\}$, consists of several strike prices which are above k_i. For example, the implementation used for generating Figure 3.4 generated the probes at each iteration i by $probes_i = \{k_i + 0.1, k_i + 0.5, k_i + 1\}$. More subtle strategies would place the probes depending on the distance between $\Phi(k_i)$ and ρ.
To obtain a conservative linear overestimator, the worst case (i.e. steepest) slope is used for a linear model of $\hat{\Phi}_i(k)$. Intersecting the linear model with the acceptability threshold results in a new candidate strike price k_{i+1}. Evaluating [LL] at k_{i+1} and then in turn [UL] at $y^\star(k_{i+1})$ leads to another estimate $\hat{\Phi}_{i+1}(k)$ which can be approximated linearly again. The procedure can be iterated until the [UL] optimal value $\Phi(k_m)$ reaches a certain distance to the threshold (denoted as *tol* in Algorithm 3.2). The resulting candidate k_m for the strike price can then be used either directly as solution or further driven towards the optimum by applying Algorithm 3.1.

[7] i.e. storing the basis of the solution between solving two instances

Pseudocode for the procedure can be found in Algorithm 3.2 and the corresponding computational illustration is depicted in Figure 3.4.

Algorithm 3.2 Starting value using linear approximation of $\hat{\Phi}_i(k)$

```
set initial strike price k_1, set i = 1
compute y*(k_1) and Φ(k_1), λ_ζ(k_1), λ_y(k_1)
set Φ̃ = Φ(k_1)
while Φ̃ + tol < ρ

    compute steepest slope for probes

        Δ_i = max { (Φ̂_i(p_j) − Φ(k_i)) / p_j  |  p_j ∈ probes_i }

    new strike price by extrapolation k_{i+1} = (ρ − Φ(k_i)) / Δ_i + k_i
    compute y*(k_{i+1}) and Φ(k_{i+1}), λ_ζ(k_{i+1}), λ_y(k_{i+1})
    set Φ̃ = Φ(k_{i+1})

end while
return k_{i+1} as starting value for Algorithm 3.1
```

In practice, this simple method shows remarkably good results, even for rather small[8] values of *tol*, i.e. requiring convergence close to the true optimum. To get an idea of the magnitude of the improvement, note that in the example used for illustrations in this chapter, the number of [LL] instances to be solved for convergence was reduced by a factor of more than 30 by introducing linear approximations, while keeping the number of [UL] evaluations constant.
Numerical experiments suggest that the quality of the linear approximation even increases with the size of the linear programs to be solved. The reason for this is probably an increasing smoothness of the optimal value function $\Phi(k)$ (and its estimate $\hat{\Phi}_i(k)$) for large problems. Therefore, an approximation of the optimal value by Algorithm 3.2 followed by Algorithm 3.1 serves as the workhorse for all applications in Chapter 4.

Of course, the modification turns the algorithm into a heuristic, i.e. the guaranty to converge towards the optimal value is lost. As usual for approaches which work on local approximations of the original problem, the region where the approximating model is assumed to be valid has to be specified in advance by choice of the parameter *tol*.

Recall that for the special problem structure in (3.15), where [LL] constraints on cumulative consumption, nonanticipativity and market depth are omitted and [UL] and [LL] are subject to the same prices, we proved concavity of $\Phi(k)$ (c.f. Lemma

[8] here, about 1% of ρ

(3.3)). As a direct consequence, every tangent on $\Phi(k)$ is automatically a global overestimator, which makes the algorithm in this subsection converge to the exact solution.

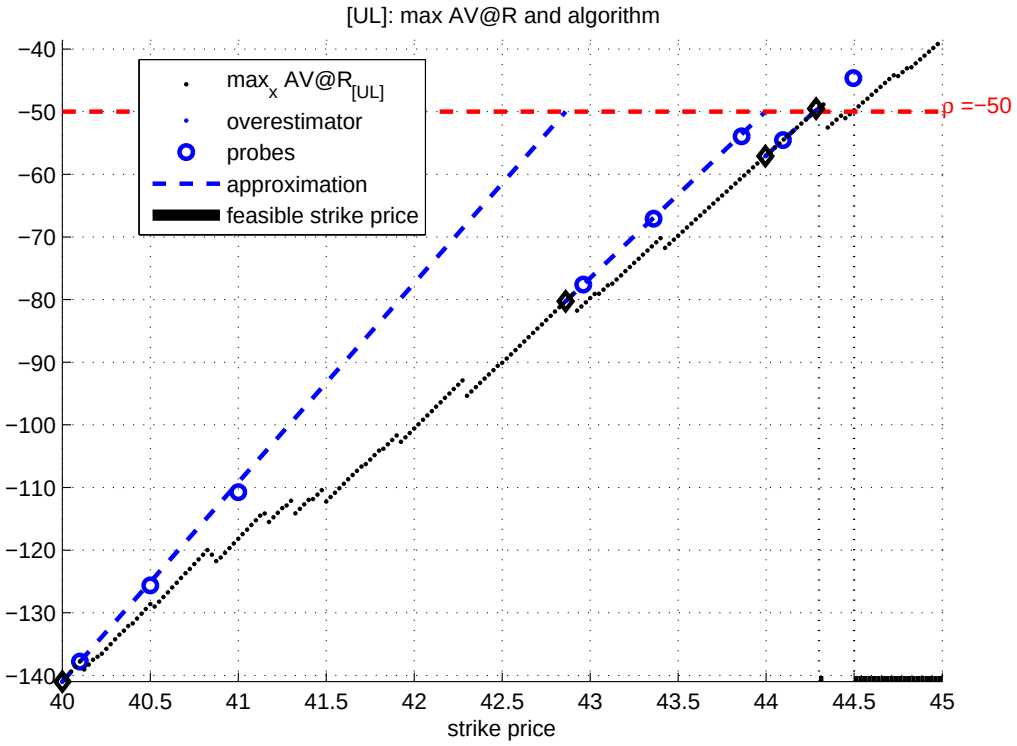


Figure 3.4.: Iterative construction of linear models for $\Phi_i(k)$ (dashed).
Blue circles correspond to evaluations of [LL] at "probes" to obtain the slope of the linear models. Black diamonds show evaluations of [UL]. The overestimator $\hat{\Phi}_i(k)$ is depicted by blue dots. In the case shown, convergence was reached after solving a total of 13 instances of [LL] and 4 instances of [UL].

Remark. Another approxach to improve initial values for Algorithm 3.1 is to design overestimators using simplifications of the [LL] problem. One possibility is to simplify the description of the [LL] feasible set. Omitting cumulative consumption constraints makes it possible to solve [LL] by a greedy algorithm; omitting the nonanticipativity constraints results in scenariowise decomposability. The resulting lower bounds for the optimal strike price $k^\star$ can be used as starting value k_1 if also the [LL] objective function is modified in an adequate way. After this kind of preliminary step the algorithm is still guaranteed to converge to the true solution.
This approach was also implementated in MATLAB and tested for several setups. However, it turned out that overestimators relying on simplified [LL] problems performed poorly compared to approximations obtained from the linear model.

4. Application in the natural gas market

The algorithm described in the previous section was implemented in AIMMS to calculate exercise-prices for energy swing options relative to given energy portfolios.

In the following chapter, price setting for natural gas swing options relative to a natural gas portfolio is chosen as an application for the following reasons:

- In continental Europe, exchange based natural gas markets are still not very mature. Besides justifying the application of indifference pricing, this gives us the opportunity to compare futures and flexible supply contracts as hedging instruments.
- Intra day decisions (trading and generation) are, partly due to system inertia, not as important as in the cases of an electric power portfolio. This makes it possible to use a daily time grid and simpler models for the risk factors of the portfolio.
- As a contrast to electricity generation portfolios, natural gas portfolios consisting of financial contracts (spot and futures) and physical contracts (flexible delivery and storage) can be quite realistically modeled as linear programs.

In the first section, a detailed description of the modeling of the portfolios of seller and buyer will be given. The second section contains a description of the stochastic modeling of the risk factors underlying both portfolios followed by a discussion of their discretization in the third section. The fourth section contains a discussion of some numerical results of the relative pricing problem.

4.1. Modeling of a natural gas portfolio

The natural gas portfolio of the seller's problem is intended to capture the situation of an importing or redistribution company. The portfolio consists of infrastructure (storage), a long flexible supply contract (oil indexed, exercised day ahead) and stochastic demand attributed to one or several shorted contracts. In the first stage, the seller has access to futures markets for natural gas and oil which makes it possible to hedge the long supply contract to some extent. Additionally, the seller has access to a spot market for short term portfolio balancing.

In the first step, the natural gas portfolio of the buyer ([LL]) consists solely of the long swing option (exercised day ahead) bought from the seller. This is due to the fact, that under valuation via expectation the exercise strategy of the swing option does not depend on the rest of the portfolio (equation (1.2)). In an extension step (page 85), the objective of the buyer will be altered to include risk aversion, leading to an incorporation of stochastic demand similar to the seller's portfolio.
The buyer is assumed only to use spot markets[1]; again, linear valuation makes the holder indifferent to using futures contracts if they are priced at the expected spot price.

The models for the seller [UL] allow the possibility to incorporate market-frictions (bid-ask spreads) for natural gas spot and natural gas futures markets as well as limits on the market depth (i.e. limited size of transactions).

For computational ease we incorporate only one of each instruments. In principle, extension to several instruments is straightforward but demands higher modeling and computational efforts and makes analysis of results more difficult.

4.1.1. Modeling details

Decision structure For modeling the decisions of the problem at hand, a distinction between the initial date (stage 0) and further stages has to be made: in the initial date, the strike price is set (negotiated) and hedging decisions (futures purchases) can be made. At all consecutive stages ($t = 1, \dots, T$), the swing option is exercised by the buyer and the seller performs portfolio balancing by trading, using natural gas storage and exercising the own supply contract.
Note that the time steps between stages, in particular between $t = 0$ and $t = 1$, are not necessarily equidistant. Especially in the context of price setting for an energy delivery contract, it is reasonable to negotiate a contract several months before delivery starts. The resulting difference of available information at price setting and portfolio balancing is captured by the tree representations of the filtrations $\mathfrak{G}_0, \mathfrak{G}_1, \dots, \mathfrak{G}_T$.
As already mentioned in the introduction, during the delivery period both the buyer and the seller have to exercise their long contract day ahead (i.e. for uncertain prices) whereas the buyer can operate the storage in real time (i.e. for revealed prices). This decision structure is shown schematically in Figure 4.1.

Objective functions of [UL] and [LL] The seller's objective will be modeled as $AV@R_\alpha(Y)$ of the random cumulative payoff Y, since the resulting payoff distributions can be interpreted rather easily.

The buyer's objective will be modeled as $\mathbb{E}(Y_{LL})$ of the cumulative payoff Y_{LL}. As an extension, to capture a wider range of [LL] behavior, the assumption of risk neu-

[1] for reasons of simplicity, here no transaction costs are included

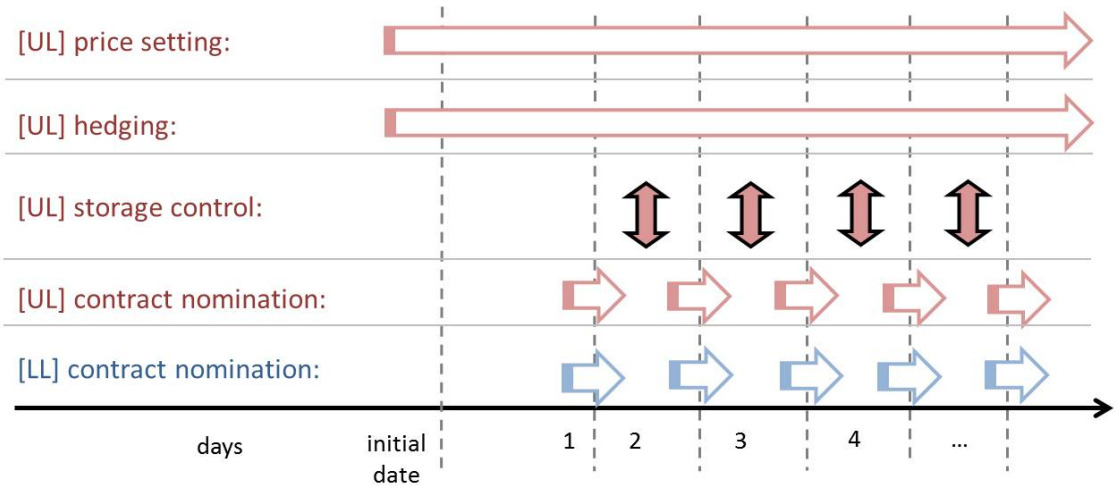


Figure 4.1.: Decision structure for buyer's problem [UL] and seller's problem [LL]

trality will be relaxed[2] by replacing $\mathbb{E}(Y_{LL})$ in the [LL] objective with $AV@R_\alpha(Y_{LL})$ for α close to 1.

[UL] contract: flexible contract with index coupling As already mentioned in Example 1.1.3, exercise prices for natural gas swing options are often coupled to a price index of a fuel competing to natural gas.

Recall from Example 1.1.3, that for this type of contracts, the exercise price of the delivery contract (for the seller) is coupled to an index price by the following formula:

$$P_t = AP_0 + \kappa(I_t - P_0)$$

i.e. P_t, the delivery price at time t, is given by a fixed component AP_0 plus a flexible component $(I_t - P_0)$ which is weighted by a linking factor κ. The flexible component consists of the deviation of a price index I_t from an initial value P_0. Usually, the parameters AP_0 and P_0 are calibrated to natural gas forward prices. For the applications below, these parameters were chosen to match the expected spot prices for natural gas during the delivery period.

A common choice for an index price I_t is to fix the exercise price for several months using the mean index price over a certain time period. In our example we look at a total time horizon of four months. The average oil price of the first three months (stage 0) is used to determine the index price for the fourth month (delivery period, stages $1, \ldots, 30$). Since for our application, the oil price has been scaled down to the level of natural gas prices due to numerical properties, we use a rather high linking factor $\kappa = 0.8$.

[2]For a discussion of the connection between average value at risk and expected value we refer to Appendix, A.1.3 and the related remark.

Exercise decisions are taken day ahead (i.e. consumption at time t is decided at $t-1$). Consumption is constrained by daily upper and lower limits as well as upper and lower limits on cumulative consumption over the whole delivery period.

Stochastic demand In general, the [UL] portfolio contains short contracts, which are represented by a stochastic demand (which has to be satisfied at spot market prices). This demand is assumed to be largely based on retail customers, thus it is generated by a nonlinear transformation of the temperature. The transformation is a sigmoid function, as described in Subsection 4.2.2, which has been slightly modified to increase loads at extreme temperatures (see Figure 4.3 for an illustration).

Storage facility The storage facility is defined by ramp rates for input and output and a storage-volume. Costs arise from filling (and possibly emptying) the storage, but not for keeping natural gas in stock. Notice that in the model, input and output decisions can be taken in real time i.e. for revealed prices. Our standard parameter choices are roughly based on real life storage operation contracts[3] and allow an inflow speed of half the outflow speed at variable costs of 0.5 €/MWh for inflow only.

Futures market In the initial stage ($t = 0$), the seller can open long or short positions in natural gas futures for hedging against spot price fluctuations in the delivery period. Also, positions in oil-index futures for hedging against contract delivery price fluctuations in the first time step can be opened.

The futures prices are given by the expected spot price plus transaction costs which have to be paid for both purchases and sales. This transaction cost should reflect a premium which appears for both buyers and sellers in the market, avoiding arbitrage in the model and modeling illiquidity or risk praemias.

Spot market During the delivery period ($t = 1, \ldots, T$), the seller can trade at the spot market. Similar to the case of futures markets, deviations from perfect liquidity are modeled by bounds on the market depth and bid ask spreads[4].
Limited market depth amounts to constraints on the size of net spot transactions which are given by the sum of [LL] consumption, the stochastic demand, the net flow from spot market to storage, futures positions and finally spot trading for each single period[5]. The computational results in the last part of this section contain restrictions on the spot transactions of several sizes.
The buyer of the swing option has access to the same spot prices; however without neither any restrictions on market depth nor transaction costs.

[3] see [EWE13] from http://www.ewe-gasspeicher.de/speicher-ruedersdorf.php
[4] for the spot market, absolute cost in €/MWh are used
[5] cf. [OG08] for a similar modeling approach regarding market depth

A overview of the input parameters of the standard model setup can be read from the screenshot in Appendix B.1.

4.2. Stochastic modeling of risk factors

As can be seen from the description of the portfolio in the previous section, the risk factors we have to model are natural gas spot prices, the index price of the seller's supply contract and the demand in the seller's portfolio.

As a first step, we develop a discrete time model for the joint time series of these three variables using continuous distributions. In a second step, this model is approximated by a discrete model (i.e. a tree), which is the topic of Section 4.3. Notice that although our analysis of risk factors is based on real data, its main purpose is to illustrate methods for setting up a portfolio model and the challenges encountered at this. Although we think that the model of the risk factors used generates a realistic input for our optimization problem in terms of structure (i.e. a high dimensional non-Markovian process with nontrivial dependencies between the variables) we are far from claiming that the stochastic model as presented below performs well enough to be used for real-life applications.

4.2.1. Data

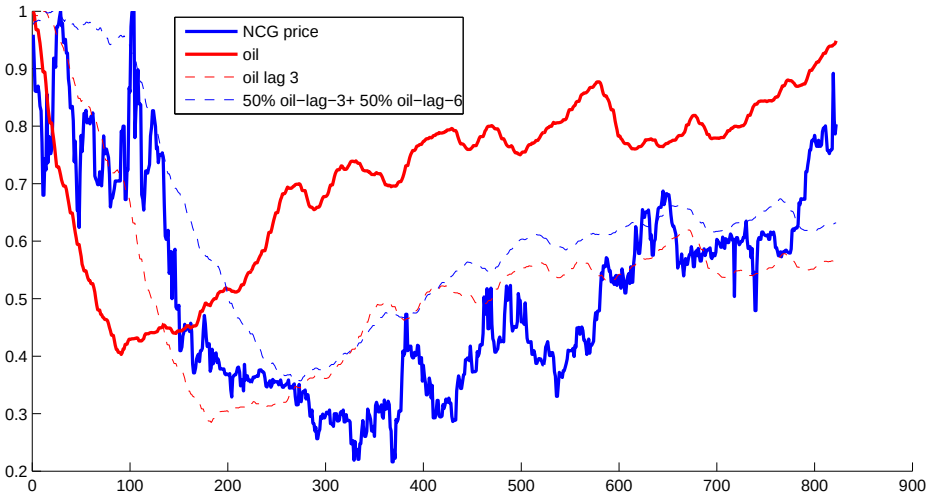


Figure 4.2.: Natural gas and oil prices
Dashed lines: oil prices lagged for three months (red) and linear combination of three month and six month lagged oil prices (blue)

For natural gas spot prices we use EEX-NCG spot prices. These prices refer to exchange traded products, specifying the delivery of a constant load of natural gas over all 24 hours of the subsequent[6] day in a German market area. The data used is publicly available at the webpage of NCG[7].
The data used for temperatures and crude oil spot prices are from the webpage of David Wozabal[8]. The temperatures consist of a population weighted mean of the temperature over all cities with more than 10000 inhabitants in Germany[9]. Both datasets can be used to obtain daily values for prices and temperatures, including Saturdays, Sundays and other holidays.
The time period examined ranges from October 1st 2008 to January 1st 2011.

Oil prices As a first step we investigate the dependency of the natural gas price and oil price in Figure 4.2. In central Europe, linking natural gas delivery prices to lagged, average oil prices was a widespread practice at the time period analyzed. Consequently, we expect some link between the price levels of (lagged) oil spot prices and natural gas spot prices which can indeed be observed.
Notice in particular the effect of the collapse of energy prices caused by the financial crisis in 2008, which can be observed at the beginning of the oil time series and - with a little time delay - at natural gas spot prices around observation 150.

Natural gas demand Since our dataset does not contain natural gas consumption it is impossible for us to directly explore the dependency between natural gas spot prices and the level of consumption. As a workaround, we construct synthetic natural gas consumption data l_t by transformation of temperature data T_t via the following function, the so-called sigmoid function:

$$l_t = \frac{A}{1 + \left(\frac{B}{T_t - 40}\right)^C} + D \tag{4.1}$$

In practice, this formula is used for day ahead estimates of the quantity to be supplied by natural gas providers for certain types of their customers. For more details and the various choices of the coefficients we refer to [Bun07]. For our model we choose the parameters $A = 3, B = -37, C = 5.6, D = 0.08$, which roughly corresponds to a standard domestic home. To incorporate the inertia of natural gas consumption for heating purposes, instead of the daily average a moving average over several days is used as temperature value in the formula, replacing T_t by the weighted average

$$\tilde{T}_t = \frac{\sum_{i=0}^{3} \left(1/2\right)^i T_{t-i}}{\sum_{i=0}^{3} \left(1/2\right)^i} .$$

[6] or on the closest previous trading day in case of the day after a holiday
[7] http://datenservice.net-connect-germany.de/
[8] http://homepage.univie.ac.at/david.wozabal/EEXdata.csv
[9] cf. [KW14]

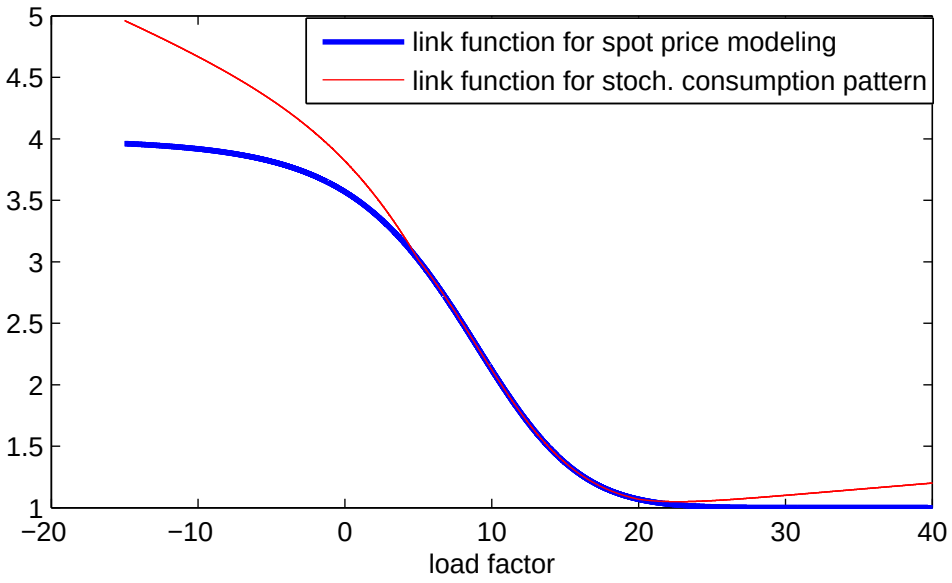


Figure 4.3.: Generation of synthetic load profiles via sigmoid functions.
Blue: Link function used for generating the synthetic load time series
Red: Corrected link function accounting for the effect of power production at extreme temperatures, used for generating stochastic demand in the [UL] and the risk averse [LL] portfolio model (see page 66)

The resulting time series $l(\tilde{T}_t)$ will be referred to as synthetic load.

At this point we want to stress the fact that we generated the synthetic load by not much more than educated guessing: we did not fit the coefficients of the sigmoid function either to any consumption data or to match the observed spot prices.

4.2.2. A model linking oil and temperature to natural gas prices

Capturing the interdependence structure between risk factors is one of the key challenges at the modeling of energy portfolios.
We assume the load to be temperature driven; therefore we will explicitly model oil prices and temperature as risk factors and use a transformation of them to obtain the natural gas prices and the portfolio load.
The resulting model is able to capture some of the features of interest; however it is far from delivering satisfactory results for real life application.
Nevertheless, this model illustrates the central benefit of the approximation procedure for stochastic processes, used for tree generation (Section 4.3): The method works for arbitrary distributions of the underlying risk factors, including high dimensional models with nontrivial dependencies between the different risk factors; the

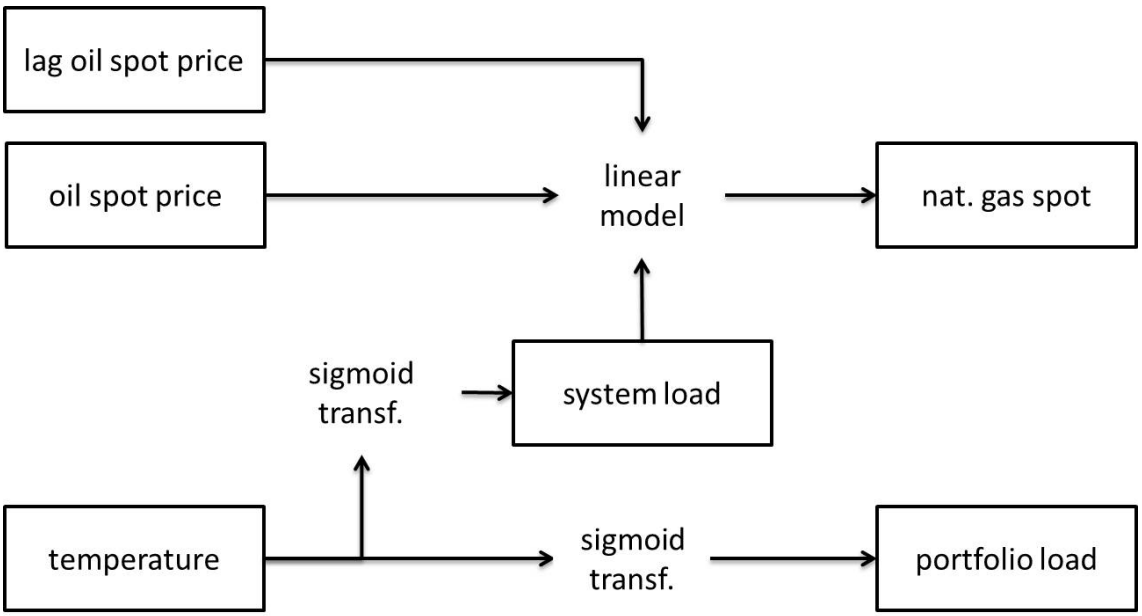


Figure 4.4.: Interdependency between risk factors

only requirement is the possibility to take samples from the conditional distributions of the process.

At this point we want to emphasize that - in spite of the questionable quality of our model - the principal approach of transformation of several "easy to describe" risk factors to obtain the joint distribution of risk factors of the whole portfolio seems very promising (see for example "hybrid models" in [EW03] or "structural models" in [CCS13], both as applications to electricity markets).

Natural gas as a function of oil and temperature We determine the dependency between the logarithm of the natural gas spot price P_t, oil spot prices O_t and synthetic system load $l(\tilde{T}_t)$ (depending on moving-average temperature) by fitting a linear model.

In addition to the oil spot price we include a lagged variable O_{t-6M}, consisting of the oil spot price of six months ago. This is motivated by the fact that oil indexed natural gas delivery contracts usually rely on lagged oil prices as well. To account for a possible impact of extreme demands, we include a quadratic load-term into the model. For numerical reasons at tree fitting, we scale oil prices and load factors to the level of natural gas prices. The interdependency of the variables is depicted in Figure 4.4.

Calibration by ordinary linear regression results in the following model for logarithms

of natural gas prices:

$$log(P_t) = \begin{bmatrix} 1 & O_{t-6M} & O_t & l(\tilde{T}_t) & l(\tilde{T}_t)^2 \end{bmatrix} \cdot \begin{bmatrix} 0.6950 \\ 0.0608 \\ 0.0577 \\ -0.0372 \\ 0.0014 \end{bmatrix} + \epsilon_t \tag{4.2}$$

where the error term is $\epsilon_t \sim N(0, 0.0211)$. This configuration leads to a coefficient of determination of $R^2 \sim 0.81$.

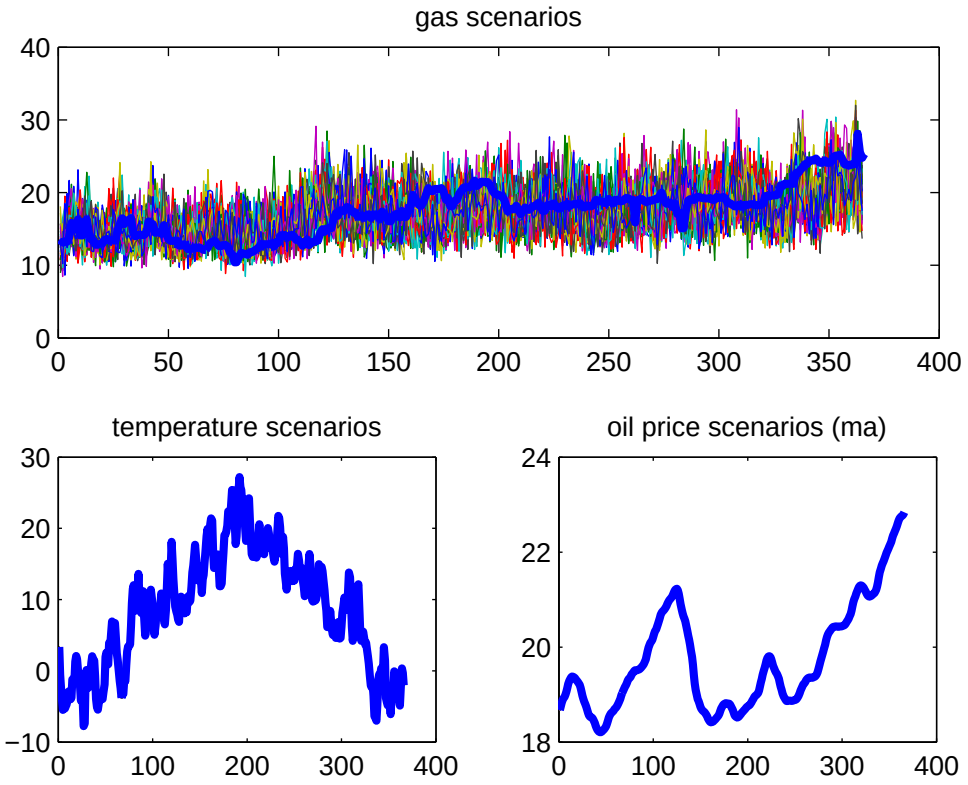


Figure 4.5.: Simulations of the natural gas spot price for historic temperature and oil price data

The benefit of this model is the incorporation of a significant connection between the risk factors. However, it turns out that in particular short term fluctuations of the natural gas price are not adequately captured as can be seen in Figure 4.5.

Oil price model A standard approach to modeling oil prices is to fit geometric Brownian motion or geometric mean reverting processes by investigating the

log-returns of the prices[10]. However, in our case the log-returns appear to be of doubtable quality, since our data contains a high amount of identical prices at consecutive days [11]. Since essentially we are only interested in the connection (presumably present rather for long term levels than for short term shocks) of natural gas prices and long term averages of oil prices (for the index of the delivery contract) we replace the oil price $\tilde{O}_t$ by its 5-day moving average $O_t = \frac{\sum_{i=0}^{4} \tilde{O}_{t-i}}{5}$. This modified data is also used as input in the regression model.
The log-returns of $\tilde{O}_t$ are modeled by an $AR(3)$ process with errors from a normal distribution, where model selection was based on the AIC.

Temperature model A simple model for temperature, consisting of a deterministic seasonality (constant plus sine and cosine with periodicity of 365 days) and an $AR(3)$ model (with normally distributed error terms) was fitted. Again, the order of the auto-regressive process was based on the AIC.

4.2.3. Quality of the hybrid model

The first assessment of our linear model is to simulate natural gas prices for the observed oil price and temperature time series, and to compare them to the observed natural gas price time series (Figure 4.5). It is immediately clear that the volatility of the natural gas prices is overestimated since the explanatory variables carry too little information to capture the short term dynamics of the natural gas prices.

Out of sample trajectories (Figure 4.6) show that the models for temperature and 5-day oil price moving averages behave reasonably well.
Notice in particular the red and turquoise scenario, in which high price scenarios for oil translate into high price scenarios for natural gas. Also the histogram of prices occurring during all simulations (Figure 4.7) roughly fits to the values actually observed. However the two distinctive peaks in the density of natural gas prices are clearly not captured in the fitted model. Also the distribution of simulated prices shows too fat tails.
Joint estimation of all model parameters would immediately improve the fit. Also, additional explanatory variables such as for example the relative storage level[12] should improve the quality of the model significantly, as well as using more sophisticated models for sources of risk (for example seasonal volatility for temperature).

[10] see for example [CS00]
[11] recall that our data also includes prices for holidays which appear to be extrapolations of prices at close dates
[12] c.f. [OG09] for an examination of the effect of scarcity on natural gas prices

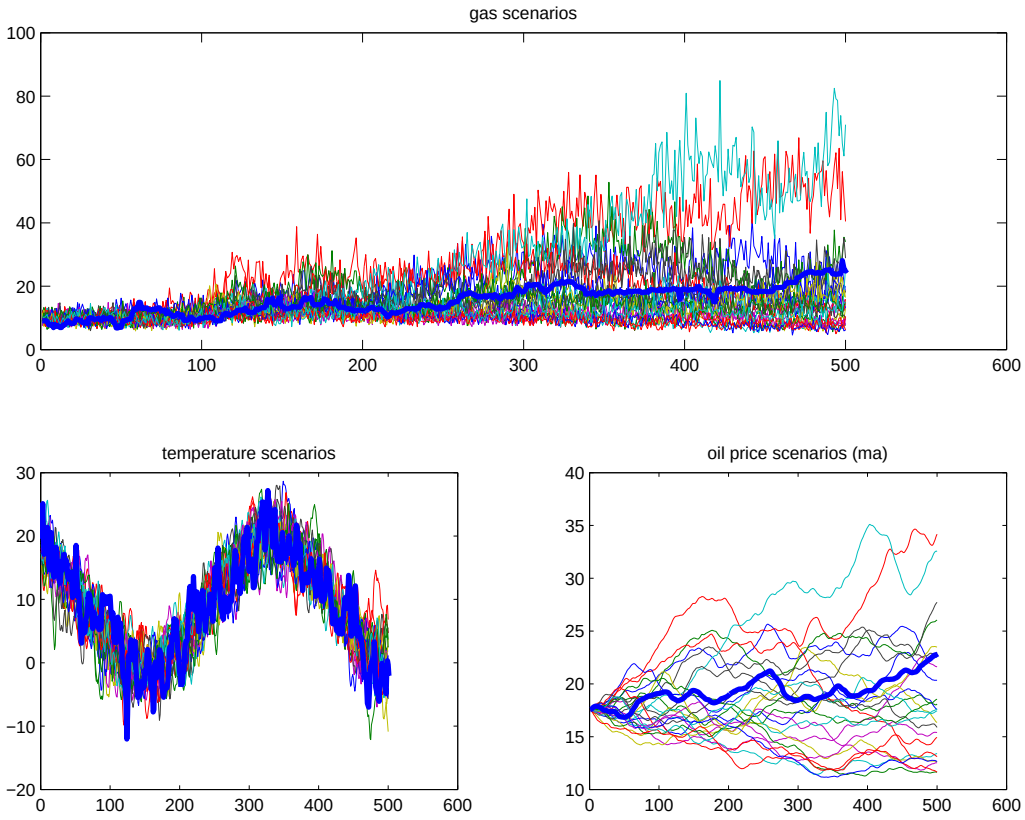


Figure 4.6.: 25 scenarios of temperature, oil spot prices and natural gas spot prices

4.3. Deterministic approximation of the data process

4.3.1. Introduction

Stochastic optimization problems are infinite dimensional optimization problems, since their decision variables are functions of the random information. To solve this type of problem in practice, usually a simplification of the original problem is necessary.
If however the random information can be described by discrete random variables, the corresponding stochastic optimization problem is automatically finite dimensional. Consequently, for computational purposes, usually the original stochastic process $\mathbb{P}$, originally estimated from data, is approximated by a discrete process $\tilde{\mathbb{P}}$. For an illustration of how to design a "good" approximation we investigate a value function of the type

$$v_{\mathbb{P}}(x) = \{\mathbb{E}_{\mathbb{P}} f(x, \xi_{\mathbb{P}}) \,|g(x, \xi_{\mathbb{P}}) \leq 0\}$$

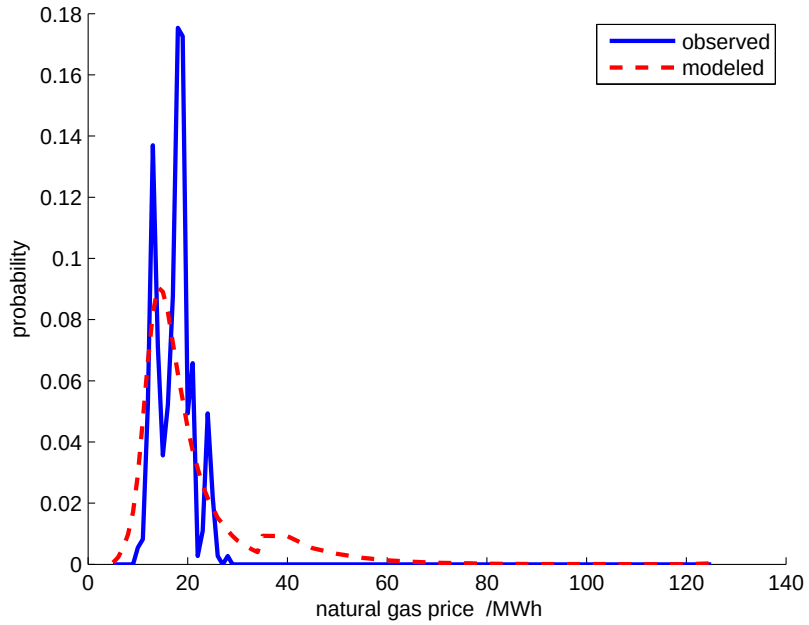


Figure 4.7.: Histograms of simulated prices and observed prices

where x is the decision variable and $\mathbb{P}$ is the stochastic process generating the value process $\xi_{\mathbb{P}}$. Let the solutions of the original and the approximate stochastic optimization problem be given by

$$x^\star \in argmin\ v_{\mathbb{P}}(x) \text{ and } \tilde{x} \in argmin\ v_{\tilde{\mathbb{P}}}(x)$$

respectively. Basically, we want to design the approximate problem in such a way that its *solution* $\tilde{x}$ also performs well for the original problem. Depending on the usage of our model, it may already be sufficient to just obtain a good estimate of the optimal *value* of the original problem.

Subsequently, we are going to measure the appropriateness of the approximation by the difference of the objective *values* of the original and the approximate problem by requiring the approximation error

$$\epsilon := |v_{\tilde{\mathbb{P}}}(\tilde{x}) - v_{\mathbb{P}}(x^\star)| \tag{4.3}$$

to be small. Note that this approach does neither guarantee that the optimal solutions of the optimization problems are close to each other nor that the approximate solution performs well in the original problem. To see the difference to the latter requirement, observe that

$$|v_{\mathbb{P}}(x^\star) - v_{\mathbb{P}}(\tilde{x})| \leq \underbrace{|v_{\mathbb{P}}(x^\star) - v_{\tilde{\mathbb{P}}}(\tilde{x})|}_{\epsilon} + \underbrace{|v_{\tilde{\mathbb{P}}}(\tilde{x}) - v_{\mathbb{P}}(\tilde{x})|}_{\delta} .$$

A bound on the nested distance between true process and its approximation extends to a bound on ϵ by the Lipschitz property of the value function (see Appendix, Theorem A.15). The algorithm used for generating the trees used for our application is one way to find approximations leading to small values of ϵ. Although we expect a solution providing a small value of ϵ to also perform well for the original problem, we have no way of assessing the size of δ in practice without resorting to simulation.

In the following sections we will sketch the practical procedure of approximating a stochastic process. The computations were performed in MATLAB, based on a dynamic tree generation procedure proposed and implemented by Georg Pflug.
For more information on the distance concepts for stochastic processes used below and the connection to stability of stochastic optimization problems we refer to Appendix A.3; for a rigorous treatment to [PP14].

4.3.2. A numerical procedure for tree approximation of stochastic processes

A procedure for dynamic tree generation For tree construction we use a dynamic tree generation algorithm proposed in [PP14], which simultaneously determines structure, values and probabilities of a tree which approximates a stochastic process.

The purpose of the algorithm is to determine a tree $\tilde{\mathbb{P}}$ which is as close as possible to the original stochastic process $\mathbb{P}$ in terms of the (multi-stage) nested distance $\mathbf{d}_N(\mathbb{P}, \tilde{\mathbb{P}})$ with respect to the one-norm on the observation space (Appendix, Definition A.14). By virtue of the Lipschitz property of the value function with respect to the nested distance (Theorem (A.15)), this results in a "small" value of ϵ as defined in equation (4.3).
Since it is computationally infeasible to optimize the nested distance between stochastic processes directly, instead of $\mathbf{d}_N(\mathbb{P}, \tilde{\mathbb{P}})$ an upper bound of the type of Theorem A.17 is used as a proxy. This bound is decomposable into single-stage transportation distances between probability measures and can therefore be computed much more efficiently.
The basic functionality of the dynamic tree generation algorithm is to iteratively generate successor nodes to a given node, until the single stage transportation distance between the successors of the node and the conditional distribution to be approximated is below some prespecified threshold θ. Due to stagewise decomposability of the bound, this procedure implies that the nested distance itself will also be small.
The MATLAB implementation by Georg Pflug used for computations, uses three inputs: the number of stages of the tree T, a tolerable error for stagewise approximation θ and a simulator which generates conditional samples of the original process. For the actual computations, θ was determined dynamically depending on the probability of a given node.

The output of the implementation is a tree approximation of the original process, containing values, tree structure and probabilities. The implementation relies on Monte Carlo Simulations for computation of the single stage transportation distances and also for determining the probabilities of single nodes. Values for nodes are obtained as the minimizers of the single stage transportation distances.
The usage of simulations instead of analytical integration provides a high flexibility both in the choice of the stochastic process to be modeled and also in the dimension of the process.
For more details we refer to [PP14], Section 4.3 .

4.3.3. Generation of the tree used in Section 4.4

Our specific problem contains the additional feature of different length (in time) of the stages in our model. This is due to the time delay between agreement of delivery price and hedging (initial period, $t = 0$), and the actual exercise of the contract (delivery period, $t = 1, \dots, T$), which is also illustrated in Figure 4.1. The long time period in the first stage (consisting of the 90 time steps between price setting/hedging and start of the delivery period) is solely interesting for determining the price of the [UL] supply contract. Since this price depends on the average oil price over this time window, the approximation of the tree does not have to have a high quality at the first 90 periods.

For modeling this situation, we start out with a 120 stage tree (90 stages between price setting and delivery and 30 stages of exercise). Lower approximation quality at the first 90 stages is achieved by enforcing constraints on the number of successor nodes at early stages. The resulting tree is then collapsed into a 31 stage tree.

However, enforcing a very coarse approximation in early stages results in an explosive branching behavior as soon as the constraints on the number of successor nodes are removed.

Consequently, in addition to enforcing low branching in the first stages, we also constrained the maximal number of successor nodes during the delivery period by two, resulting in a tree where each node has either one or two successors.

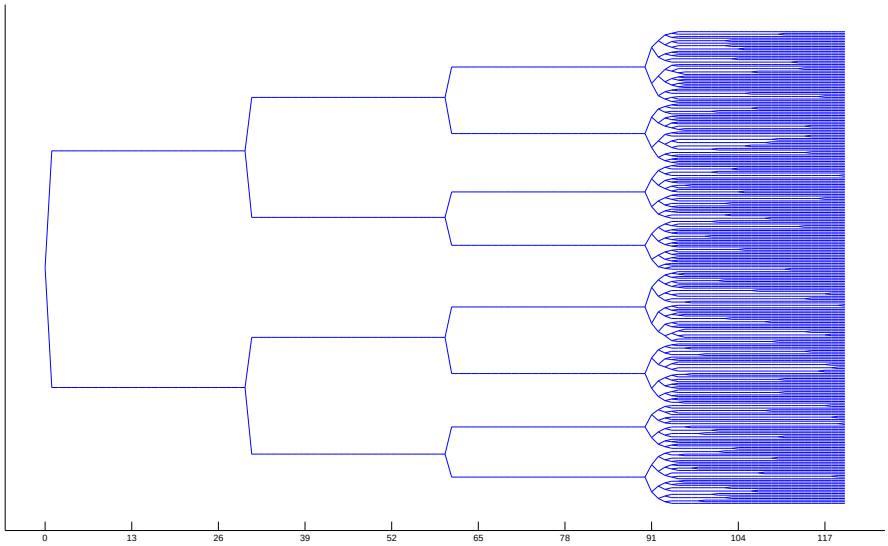


Figure 4.8.: Structure of the fitted tree used for the pricing problem.
A stage dependent constraint on the number of successors was introduced, allowing splitting only at stages $\{0, 29, 69, 90, \ldots 119\}$.
In the model discussed below, the first 90 stages of the depicted tree are treated as one stage.

The resulting tree structure is shown in Figure 4.8, the value structure in Figure 4.9.
Red rectangles in Figure 4.9 highlight the values of the stochastic processes which are directly used in the model: (averages of) oil prices over the first 90 days, temperatures and natural gas prices over the last 30 days. However, due to their influence on natural gas prices, indirectly oil prices of all stages are required. The quite erratic placement of natural gas prices wherever a node is split (i.e. at nodes 29, 69 and somewhat less visible after stage 90) arises from the very high short term volatility of the natural gas spot price model (4.2). Reflecting this variability with only a very small number of nodes leads to strong repulsion of nodes having the same predecessor.

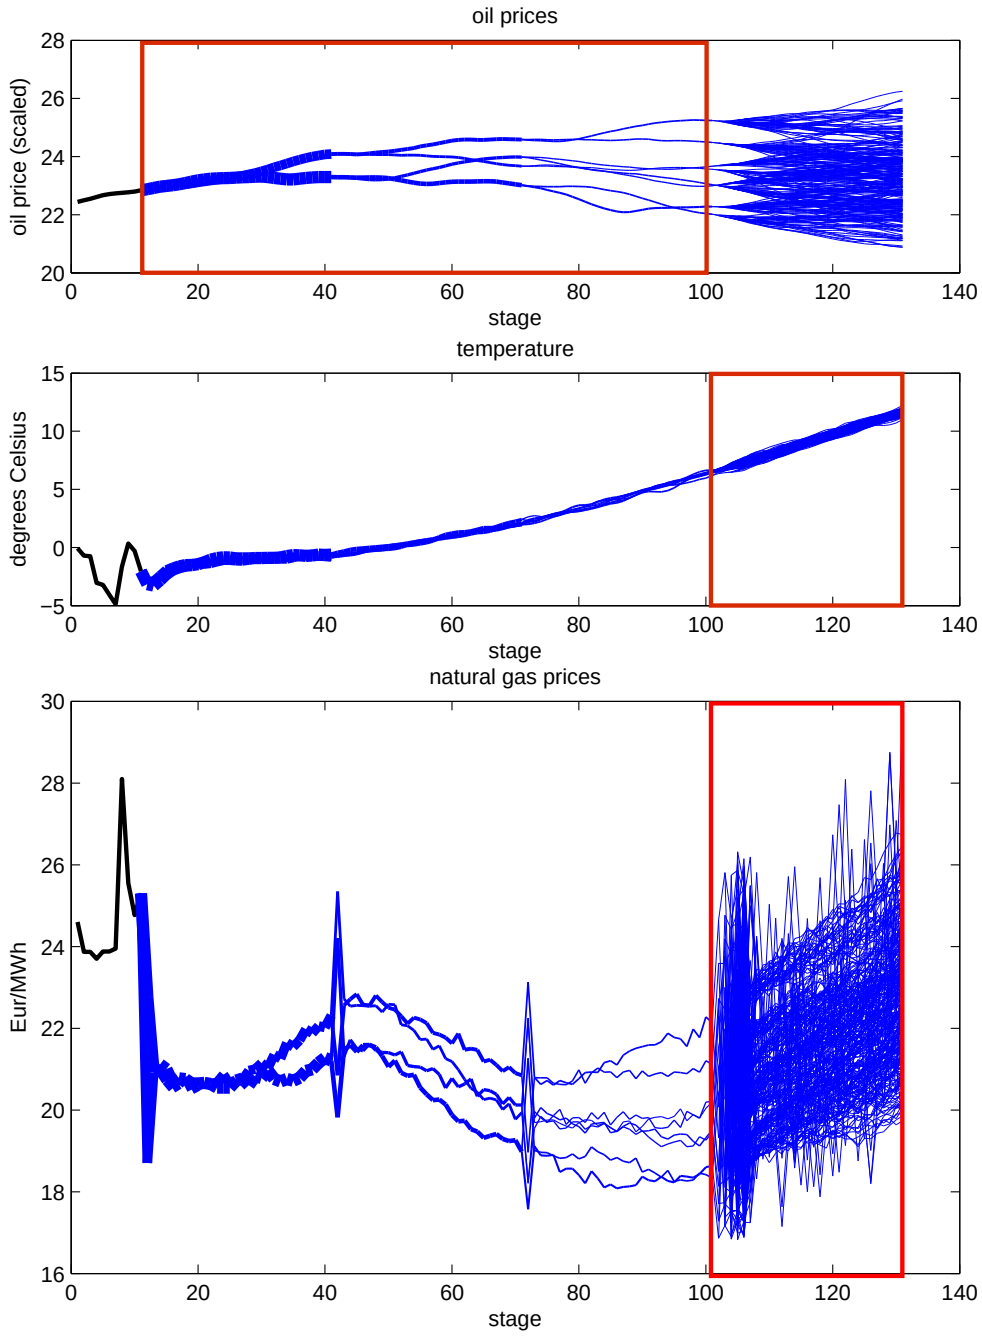


Figure 4.9.: Values of the estimated tree.
Black: historic data (used as initial values), blue: discretized stochastic processes. Red rectangles highlight the values which are directly used in the model. Width of line corresponds to probability of the scenario. (In the model, the first 90 stages are treated as one stage.)

4.4. Relative pricing of swing options

In this final section we present some numerical results from several configurations of the bilevel pricing problem which were implemented in AIMMS.

4.4.1. General model properties

For computational convenience, we chose a rather small tree consisting of 288 scenarios and 6776 nodes for 30 stages (where the first stage consists of 90 days; the subsequent of one day each). Using the GUROBI 5.1 solver, solving an instance of [UL] or [LL] in this setup took only a few seconds. Employing the linear approximation of $\hat{\Phi}_i(k)$ (Algorithm 3.2), a strike price leading to an acceptability within 1000€ of the required threshold ρ is found after solving roughly 20 instances of [UL] and [LL] combined. The resulting overall solution time for the bilevel problem is less than one minute.
This allowed an examination of the model's behavior for a wide range of different parameters, which is presented in Section 4.4.2.

The structure of the tree used is illustrated in Figure 4.8; its values are shown in Figure 4.9. This setting represents the situation where a contract for delivery over one month is closed (and also hedged) three months in advance. In our setup, the index price for the long [UL] contract consists of the average oil spot price over the three months between contract closure and start of delivery.
As a standard setting for the seller ([UL]), we choose an initial portfolio consisting of a long flexible index coupled supply contract with minimal consumption of 25% (relative to as well a daily as a cumulative maximal amount), and a storage unit.
Notice that the swing option to be priced amounts to only a small portion of the initial [UL] portfolio: the maximal cumulative amount of energy to be consumed by [LL] is about 17% of the maximal amount of energy to be consumed by [UL] from the supply contract in the initial [UL] portfolio.
The holder of the [UL] portfolio also has access to the natural gas spot market and futures markets for hedging the natural gas price in the delivery period and the index of the long supply contract. Trading in the futures market occurs 90 days before delivery; trading comes at a cost of 1% of the natural gas futures or index price[13].
The [UL] objective is to maximize $AV@R_{.15}(\cdot)$ of the profit from operation of the portfolio against the spot market.
The exact setting of the standard model parameters subsequently used can be found in Appendix B.1.

[13] which corresponds to a bid-ask spread of 2%

	transaction cost (nat. gas spot)	market depth	bid-ask spread (nat. gas futures)
high liquidity	0 €/MWh	10000 MWh/h	0
medium liquidity*	0.1 €/MWh	1000 MWh/h	2%
low liquidity**	0.5 €/MWh	500 MWh/h	10%

Table 4.1.: Parameter values for market liquidity settings
$\star$: See Figure 4.10
$\star\star$: See Figure 4.11

4.4.2. Optimal setting of swing-exercise price: different model specifications

The influence of "natural gas market liquidity" We use the term market liquidity to describe both costs for executing trades and market depth. Transaction costs can be interpreted either as bid-ask spreads or as risk premia, prohibiting to trade at the expected spot price. Market depth simply describes the maximal size of orders which can be placed at the market over a given day. To allow efficient comparison of results, we introduce three standard setups - low, medium and high - for natural gas market liquidity as described in Table 4.1. Notice that market liquidity does not affect index futures or the buyer's problem.

For the moment, we focus on the pricing of a swing option, whose holder is required to consume between 20% and 100% of a prespecified daily and also of a prespecified cumulative amount of energy.
In this paragraph we will compare the cases of medium and low market liquidity as illustrated in Figures 4.10 and 4.11.
Compared to an expected natural gas price of 20.4€/MWh for the delivery period, for medium market liquidity, the optimal strike price is 21.2€/MWh, which decreases for low liquidity to 20.3€/MWh. Note that in the latter case, an exercise price is offered at about the expected spot price i.e. without demanding any additional risk premium.

Figure 4.10 gives a more detailed picture of the behavior of [UL] and [LL] for the medium market liquidity setting:

- The [UL] payoff distributions are compared at the top. The leftmost vertical line illustrates the value of the acceptability functional $AV@R_\alpha(\cdot)$ and, further to the right, the two corresponding $V@R_\alpha(\cdot)$ levels for the respective CDFs. The horizontal line marks $\alpha = 0.15$ used for the acceptability functionals. Notice that here the CDF of the portfolio after selling the swing option (blue) is dominated by the CDF of the initial portfolio (red).
- The [UL] cumulative consumption CDFs associated to the payoff distributions are shown in the mid picture. The vertical lines show the cumulative amount of natural gas hedged on the natural gas futures market illustrating the roughly

10% lower hedging rate in case of selling the swing option; the range of the x-axis corresponds to max and min cumulative quantities (y_U and y_L) specified in the contract. A slight tendency to increased consumption can be observed for the case of the shorted swing option.

- The [LL] cumulative consumption can be found at the bottom, showing a consumption level close to the lower bound on cumulative consumption (of roughly 20000 MWh) in about 50% of the scenarios.

Figure 4.11 shows the results for the low market liquidity setting. Notice that in this case, the overall acceptability is much lower than in the reference case depicted in Figure 4.10:

- Again, the top figure shows the payoff distribution. Compared to the case of medium market liquidity, the situation has visibly changed. Now the payoff CDF in the case of the shorted swing option all but dominates the initial portfolio.
- Compared to the medium market liquidity setting, [UL] cumulative consumption has changed, which is especially pronounced in the case of extremely low and extremely high consumption. Notice in particular the significant drop in hedging with natural gas futures due to the higher transaction cost.
- Comparing the [LL] consumption at the bottom of the two figures shows the drastic influence of the strike price on the [LL] behavior: An exercise price at spot price level (Figure 4.11) leads to a roughly 25% chance of maximal consumption and 25% of minimal consumption at a given day. Increasing the price by about one Euro to 21.2 €/MWh (Figure 4.10) increases the chance of minimal consumption to about 50% whereas the probability of maximal consumption drops to roughly 10%.

Summarizing, the effect of decreasing market liquidity on the [UL] behavior is lower consumption from the own supply contract combined with less hedging with natural gas futures contracts. Whereas this effect was to be expected, the behavior of [LL] is more interesting: In the case of low market liquidity, the probability of high consumption of [LL] is increased significantly. [UL] provoked this behavior of [LL] by setting very low strike prices, apparently to avoid trading on natural gas spot and futures markets, which comes at relatively high cost.

This pattern can also be observed at the hedging strategies for different liquidity settings as shown in Table 4.2:
Due to the lower spot price risk, in the case of the shorted swing option the amount of natural gas spot price risk hedged with natural gas futures is lower than for the reference portfolio.
However, a higher quantity of index futures is purchased when a swing option is sold. This reflects that overall consumption of [UL] from the own supply contract, and therefore risk from index-volatility, is higher when an additional supply obligation

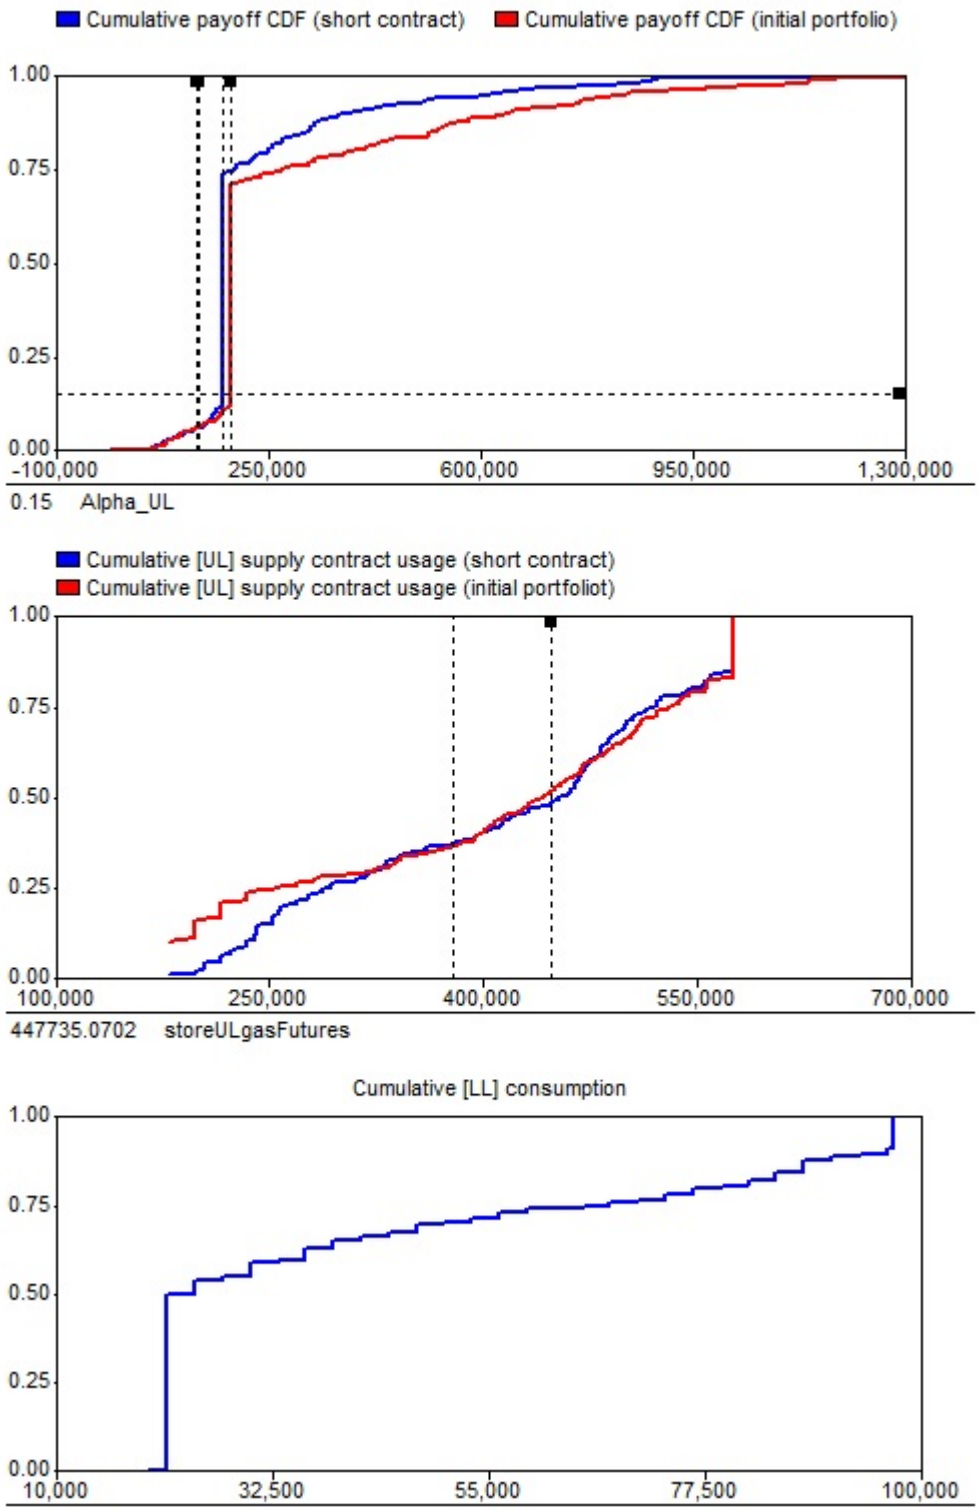


Figure 4.10.: Results of optimization for medium market liquidity (AIMMS screenshot).
The optimal strike price is 21.2€/MWh (expected value of natural gas spot price at 20.4€/MWh)

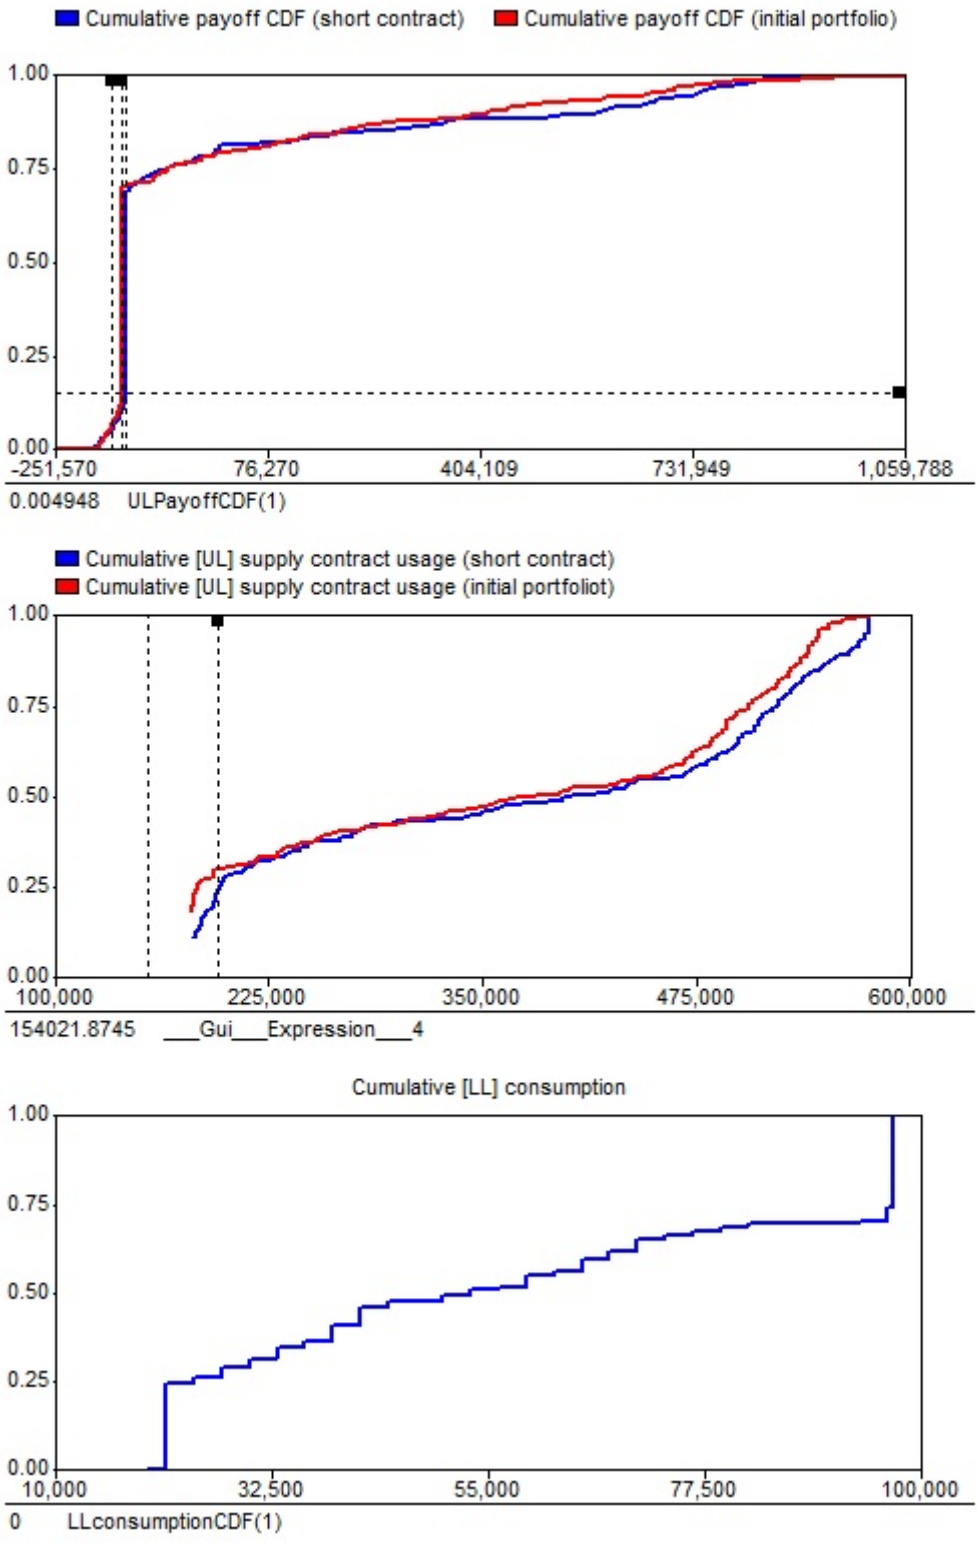


Figure 4.11.: Results of optimization for low market liquidity (AIMMS screenshot) The optimal strike price is 20.3€/MWh (expected value of natural gas spot price at 20.4€/MWh)

	nat. gas. Futures (short)		index Futures (long)	
	reference PF	short swing	reference PF	short swing
high liquidity	568 MWh/h	492 MWh/h	141 MWh/h	145 MWh/h
medium liquidity*	397 MWh/h	332 MWh/h	102 MWh/h	109 MWh/h
low liquidity**	271 MWh/h	214 MWh/h	70 MWh/h	80 MWh/h

Table 4.2.: Hedging activity for different market liquidity (rounded to integers).
$\star$: See Figure 4.10
$\star\star$: See Figure 4.11

is entered by [UL].[14] Clearly, selling a swing option shifts risk from natural gas spot markets (hedged by futures contracts) to the price-index market and the random [LL] behavior. In our model, market illiquidity does not concern hedging at the index markets; besides, requiring only a low exercise price guarantees a predictable [LL] exercise pattern, also reducing risk. Apparently at some level of illiquidity, hedging spot price risk with natural gas futures contracts is simply more expensive than selling a swing option at very low prices.

The effect of supply obligations in the reference portfolio Table 4.3 contains the optimal[15] exercise prices for several parameter configurations additional to market liquidity:

- The percentage ranges in the leftmost column provide a measure for the short term flexibility of the swing option offered by [UL]. The given percentage specifies the allowed daily consumption of [LL] relative to the maximal daily consumption.
- Initial demand quantifies (uncertain) supply obligations already in the [UL] reference portfolio (and thus to be covered in addition to the demand from the sold swing option). The given percentage specifies the relative size of maximal demand from the supply obligation and the maximal consumption from the [UL] supply contract.

Clearly, there is a strong influence between the [LL] flexibility and the strike price. This is not surprising, since for the extremal case of requiring 100% consumption, the swing option resembles a futures contract, which in our model is valued at the average natural gas spot price (at least in the most liquid case).

[14] The link between index and [UL] gas exercise price is set at 0.8, which means that about 80% of the volatility of the index are transmitted to the contract holder. Together with the overall lower volatility of oil prices, this explains the significantly lower hedging of the contract than of the natural gas spot.

[15] Notice the deviation of ca. 1% between the results in the previous section and comparable results in Table 4.3, top left. These result from the higher [LL] flexibility of 20% in the previous section and the rounding of the strike prices.

initial demand	0%, 25%			75%			100%		
liquidity	low	med	high	low	med	high	low	med	high
25% - 100%	20.1	21	21.2	21.2	21.2	21.2	22.4	21.5	21.2
50% - 100%	19.6	20.5	20.8	20.8	20.7	20.7	21.9	21.0	20.7
75% - 100%	19.4	20.2	20.4	20.7	20.6	20.5	21.6	20.7	20.5

Table 4.3.: Optimal strike price (€/MWh) for several configurations of market liquidity and [LL] minimal consumption requirements for [LL] objective $\mathbb{E}(\cdot)$. Expected natural gas spot price during delivery period is 20.4€/MWh.
The percentages in the leftmost column specify the allowed daily consumption of [LL] relative to the maximal daily consumption.

The level of the initial demand has a rather interesting effect on the optimal exercise price of the swing option: for low initial demand, the strike price decreases with decreasing market liquidity. This resembles the increasing cost associated to trading at the spot market and hedging the respective price risk, as was observed in the previous paragraph. For high initial demand however, the situation flips: In this situation, the additional demand from shorting a swing option leads to an increase of spot activity which is more expensive in the case of low market liquidity. Consequently now the exercise price *increases* with decreasing market liquidity.

The storage capacity has virtually no influence on the optimal strike price, thus we chose to omit the results from the table. Nevertheless, the storage parameters strongly influence the *value* of the portfolio. From this we conclude that storage operation can be largely decoupled from the rest of the portfolio optimization, probably due to the real time operation.

The effect of a risk averse buyer [LL] As already discussed in Subsection 3.1.1 our framework also allows us to perform bilevel pricing for the case of a risk averse buyer. Table 4.4 contains the strike prices under a risk averse contract holder [LL], maximizing $AV@R_{0.15}(\cdot)$ of the payoff. Since at nonlinear valuation the exercise of the contract by [LL] also depends on any additional content of the portfolio, we included stochastic demand into the [LL] portfolio in a similar manner (but scaled down) as for [UL] in the previous paragraph, however without adding market liquidity constraints to [LL].

As a general pattern, we observe the strike price to be decreasing with increasing [LL] risk aversion as illustrated in Table 4.4. Note that the [LL] behavior appears to be quite robust with respect to perturbations of the buyer's risk preferences: perturbing $\mathbb{E}$ to $AV@R_\alpha$ by setting $\alpha = 0.85$ ("close to 1", which was the motivation for this setup), the results remain unchanged to a setup where the [LL] decision maker is risk neutral.

initial demand	25%			75%			100%		
liquidity	low	med	high	low	med	high	low	med	high
25% - 100%	19.8	20.5	20.7	20.9	20.9	20.8	21.9	21.1	20.8
50% - 100%	19.7	20.4	20.6	20.8	20.7	20.6	21.8	20.9	20.6
75% - 100%	19.5	20.2	20.4	20.6	20.5	20.4	21.6	20.7	20.4

Table 4.4.: Optimal strike price (€/MWh) for several configurations of market liquidity and [LL] minimal consumption requirements for [LL] objective $AV@R_{0.15}(\cdot)$. Expected natural gas spot price during delivery period is 20.4€/MWh.
The percentages in the leftmost column specify the allowed daily consumption of [LL] relative to the maximal daily consumption.

4.4.3. Conclusion

The previous section illustrated results from a bilevel pricing approach for swing options in the natural gas market. The following two final paragraphs contain a concluding summary.

Observations about the model We observed that the optimal exercise prices obtained by the model depend on the concrete specification of the portfolio of the seller in a reasonable way: a portfolio long in natural gas leads to a smaller strike price than an already balanced portfolio. This effect is rather obvious. Notice however the more subtle effect in our specific setup, that storage operation does not influence the strike price, since this portfolio element can be used equally efficient with or without shorting a swing option.

We also see that [LL] flexibility and risk aversions directly influence the strike price: As expected the strike price increases with the "amount of [LL] flexibility" (i.e. minimal consumption requirements), whereas [LL] risk aversion generally decreases the optimal strike price - although in our specific example the strike price turned out to be rather robust with respect to [LL] risk preferences.

Finally, we saw a strong dependence between the market model (i.e. settings for market liquidity) and the optimal strike price. In our setup of a risk averse seller, we observe that futures trading is replaced by the sale of a swing option at a very competitive exercise price. Thus ensuring steady consumption by the buyer, the swing option replaces natural gas futures as hedging instrument to some extent.

Observations about the algorithm The tailored algorithm developed in Section 3.3 (i.e. the combination of the Newton-type Algorithm 3.1 and the initial value search in Algorithm 3.2) is able to solve the bilevel pricing problem quite efficiently by solving only few of the separated [UL] and [LL] problems.

We did not perform systematic examination of the performance of the heuristic for different sizes of the stochastic programs involved. However, our experience shows that the performance improves with increasing number of nodes of the deterministic approximation. This is probably due to the increase in smoothness of the [UL] value function.

Notice that although the specific bilevel problem turns out to be solvable rather efficiently, it remains very difficult to solve the bilevel pricing problem for a real world setup. The bottleneck turns out to be the solution of the single level stochastic programs, [UL] and [LL]: Real world pricing problems typically involve contracts with a delivery horizon of one or several years. Due to its short term flexibility, it is necessary to use at least a daily time grid. This would require us to perform long term portfolio optimization with high time resolution, which seems still to be too demanding for current computational power.
Again, we want to stress the fact that the difficulty arising from the bilevel structure has been circumvented by the algorithm proposed in a satisfying way. However, the mere size of the "classical" stochastic multistage portfolio optimization problems to be solved remains a challenge.

A. Fundamental concepts and technologies used in previous chapters

A.1. Quantifying risk and acceptability

If not explicitly stated otherwise, the results presented in this section can be found in [PR07]. Following their notation, we will avoid the terminus risk measure and instead refer to risk (and acceptability) functionals to emphasize the difference between risk measures and probability measures.

A.1.1. Risk and acceptability

An axiomatic approach to risk measures by [ADEH99] resulted in the coining of the term coherent risk measures which amounts to the following, now widespread requirements on risk functionals:

Definition A.1. *Coherent risk measure*: A mapping ρ defined on some linear space $\mathcal{Y}$ of real valued random variables on $(\Omega, \mathcal{F}, P)$ is called a coherent risk (capital) functional, if the following properties are satisfied for all $Y \in \mathcal{Y}$:

1. Translation antivariance

   $$\rho(Y + c) = \rho(Y) - c$$

   for constant c.

2. Convexity

   $$\rho(\lambda Y^{(1)} + (1 - \lambda)Y^{(2)}) \leq \lambda\rho(Y^{(1)}) + (1 - \lambda)\rho(Y^{(2)})$$

   for $0 \leq \lambda \leq 1$.

3. Positive homogeneity

   $$\rho(\lambda Y) = \lambda\rho(Y)$$

   for $\lambda \geq 0$.

4. Pointwise antimonotonicity

$$Y^{(1)} \leq Y^{(2)} \Rightarrow \rho(Y^{(1)}) \geq \rho(Y^{(2)})$$

ρ is a coherent risk capital functional if and only if $\mathcal{A}(Y) = -\rho(Y)$ is a positively homogeneous acceptability functional:

Definition A.2. *Acceptability functionals.* A real valued mapping $\mathcal{A}$ defined on a linear space $\mathcal{Y}$ of random variables on $(\Omega, \mathcal{F}, P)$ is called acceptability functional if the following properties are true for all $Y \in \mathcal{Y}$:

1. Translation equivariance.

$$\mathcal{A}(Y + c) = \mathcal{A}(Y) + c$$

   for all constants c.

2. Concavity

$$\mathcal{A}\left(\lambda Y^{(1)} + (1-\lambda)Y^{(2)}\right) \geq \lambda \mathcal{A}\left(Y^{(1)}\right) + (1-\lambda)\mathcal{A}\left(Y^{(2)}\right)$$

   for $0 \leq \lambda \leq 1$.

3. Monotonicity

$$Y^{(1)} \leq Y^{(2)}\, a.s. \text{ implies that } \mathcal{A}\left(Y^{(1)}\right) \leq \mathcal{A}\left(Y^{(2)}\right)$$

A.1.2. Polyhedral risk/acceptability functionals

Many risk (acceptability) functionals can be expressed by optimization problems (cf. [RU13],[Eic07]). Risk (acceptability) functionals, which can be written as certain linear two-stage stochastic optimization problems where the random variable only appears at the right hand side, are called *polyhedral risk (acceptability)* functionals:

Definition A.3. *Polyhedral risk (acceptability) functional*: Let $\mathcal{Y} = L_p(\mathfrak{F}, \mathbb{R})$ be the Banach space of real random variables on some probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ for some $p \in [1, \infty)$. A probability functional $\mathcal{R} : \mathcal{Y} \mapsto \bar{\mathbb{R}}$ is called polyhedral if there exists $k_0, k_1 \in \mathbb{N}$, $c_i, w_i \in \mathbb{R}^{k_i}, i \in \{0,1\}$ and nonempty polyhedral sets $V_i \subset \mathbb{R}^{k_i}$, $i \in \{0,1\}$, such that

$$\mathcal{R}(Y) = sup\left\{ \langle c_0, v_0 \rangle + \mathbb{E}\left(\langle c_1, v_1 \rangle\right) \,\middle|\, \begin{array}{c} v_1 \in L_p\left(\mathfrak{F}, \mathbb{R}^{k_1}\right) \\ v_i \in V_i, i \in \{0,1\} \\ \langle c_0, w_0 \rangle + \langle c_1, w_1 \rangle = Y \end{array} \right\}$$

Remark. Any convex combination of a polyhedral risk (acceptability) functional and the expectation is again polyhedral.

A.1.3. The average value-at-risk

Definition A.4. *The average value-at-risk*: The average value-at-risk at level α, $0 < \alpha \leq 1$ of Y is defined as

$$AV@R_\alpha(Y) = \frac{1}{\alpha}\int_0^\alpha G^{-1}(u)\,du$$

Proposition A.5. *[[PR07], Section 2.2] Properties of the average value-at-risk:*

1. The dual representation of $AV@R_\alpha(Y)$ is

   $$AV@R_\alpha(Y) = \min\left\{\mathbb{E}\left(YZ\right)|\mathbb{E}\left(Z\right) = 1, 0 \leq Z \leq {}^1\!/\!_\alpha\right\}$$

   where the minimum is attained. Thus $AV@R_\alpha$ is a concave u.s.c. functional. In particular, the expected value is a special case of the average value-at-risk.

2. Average value-at-risk as optimization problem:

   $$AV@R_\alpha(Y) = \max\left\{x - \frac{1}{\alpha}\mathbb{E}\left(\left[Y - x\right]^-\right) : x \in \mathbb{R}\right\}$$

   (cf. [RU00]) where the maximum is attained.
   In particular, the average value-at-risk is polyhedral.

3. The average value-at-risk as conditional tail expectation:
   If the CDF $G(u)$ is continuous in u,

   $$AV@R_\alpha(Y) = \mathbb{E}\left[Y|Y \leq G^{-1}(\alpha)\right] = \mathbb{E}\left[Y|Y \leq V@R_\alpha(Y)\right]$$

   This formula can be modified to deal with jumps in the CDF.

4. The average value-at-risk is a concave minorant of the value-at-risk

   $$V@R_\alpha(Y) = G^{-1}(\alpha) \geq \frac{1}{\alpha}\int_{-\infty}^\alpha G^{-1}(p)dp = AV@R_\alpha(Y)$$

   In particular, the average value-at-risk is concave which does not hold for the value-at-risk.

Remark. The first property listed in proposition (A.5) is the central motivation for the way the relaxation of risk neutrality of the [LL] decision maker is modeled in Paragraph 4.4.2 on page 85. In the dual representation of the average value-at-risk in Proposition A.5.2, it is evident that Z has the properties of a density. In particular, assuming Y has the probability distribution P, Z can be interpreted as the Radon-Nikodym derivative of some probability measure Q with respect to P. Noting that

$$\mathbb{E}_P\left(YZ\right) = \int Y\frac{dQ}{dP}dP = \mathbb{E}_Q\left(Y\right)$$

the average value-at-risk can also be represented as

$$AV@R_\alpha(Y) = \min_{Q \in \mathcal{Q}} \{\mathbb{E}_Q(Y)\} \quad \text{where } \mathcal{Q} = \{Q \,|^{dQ}/_{dP} \leq {}^1/_\alpha\}$$

i.e. the worst-case expectation among a set of probability measures, the size of which is controlled by α. Varying the size of this set gives us a natural way to control risk aversion of the decision maker.

A.2. Basic concepts from convex analysis

The following section provides some definitions used in Section 2.3. The content is taken from [Fle05] and [RW98], which we also refer to for more details.

A.2.1. Convex geometry

Let an optimization problem be given by

$$\begin{aligned} \min\ & f(x) \\ s.t.\quad & g_i(x) \leq 0,\ i = 1, \ldots, p \\ & h_j(x) = 0,\ j = p+1, \ldots q \end{aligned} \tag{A.1}$$

where the functions $f : \mathbb{R}^n \to \mathbb{R}$, $g_i : \mathbb{R}^n \to \mathbb{R}$ and $h_j : \mathbb{R}^n \to \mathbb{R}$ are continuously differentiable. We will use C for the set of feasible x and $\mathcal{I}_g$ will be used to describe the set consisting of the indices of active constraints at a point x, i.e.

$$\mathcal{I}_g := \{i \,|g_i(x) = 0\} \ .$$

Definition A.6. *(Bouligand) tangent cone*: For any $\bar{x} \in C$ the tangent cone of C at $\bar{x}$ is defined as

$$T_C(\bar{x}) := \left\{ x \in \mathbb{R}^n \,|\exists c_k \in C, \exists t_k \downarrow 0 : c_k \to \bar{x} \text{ and } \frac{c_k - \bar{x}}{t_k} \to x \right\}$$

and the linearized tangent cone of C at $\bar{x}$ as

$$T_C^{lin}(\bar{x}) := \left\{ x \in \mathbb{R}^n \,\middle|\, \nabla g_i(\bar{x})^T x \leq 0 \,\forall i \in \mathcal{I}_g,\ \nabla h_j(\bar{x})^T x = 0 \right\} \ .$$

Remark. The linearized tangent cone is polyhedral and therefore in particular closed and convex. Further note that $T_C(\bar{x}) \subseteq T_C^{lin}(\bar{x})$ (c.f. *[Fle05]).*

Definition A.7. *Dual cone*: For a given, arbitrary cone $\mathcal{C}$, its dual cone $\mathcal{C}^\star$ is defined as

$$\mathcal{C}^\star := \left\{ v \in \mathbb{R}^n \,\middle|\, v^T d \geq 0,\ \forall d \in \mathcal{C} \right\} \ .$$

Note that in the literature, the term polar cone is often used as the negative of the dual cone.

Definition A.8. *Normal cone*: Let $C \subset \mathbb{R}^n$ and $\bar{x} \in C$. A vector v is normal to C at $\bar{x}$ in the regular sense, or a regular normal, written $v \in \hat{N}_C(\bar{x})$, if

$$v^T(x - \bar{x}) \leq o\left(|x - \bar{x}|\right) \quad \text{for } x \in C \ .$$

It is normal to C at $\bar{x}$ in the general sense, or simply a normal vector, written $v \in N_C(\bar{x})$, if there are sequences $x^v \to \bar{x}$, $x^v \in C$ and $v^v \to v$ with $v^v \in \hat{N}_C\left(x^v\right)$. The cone consisting of all normal vectors will be referred to as the normal cone.

The normal cone $N_C(\bar{x})$ is closed; the regular normal cone $\hat{N}_C(\bar{x})$ is convex and closed. If C is convex, all normal vectors are also regular normal vectors.

A.2.2. Geometric characterization of optimality

This geometric characterization of the feasible set can already be used to derive basic first order conditions for optimality:

Theorem A.9. *[[RW98] Th 6.12 basic first-order conditions for optimality]: Consider a problem of minimizing a differentiable function f_0 over a set $C \subset \mathbb{R}^n$. A necessary condition for $\bar{x}$ to be locally optimal is*

$$\nabla f_0(\bar{x})^T w \geq 0 \ \forall w \in T_C(\bar{x})$$

which is the same as $-\nabla f_0 \in \hat{N}_C(\bar{x})$ and implies

$$-\nabla f_0(\bar{x}) \in N_C(\bar{x}), \quad or \ \nabla f_0(\bar{x}) + N_C(\bar{x}) \ni 0 \tag{A.2}$$

When C is convex, these tangent and normal cone conditions are equivalsent and can be written also in the form

$$\nabla f_0 \left(\bar{x}\right)^T (x - \bar{x}) \geq 0 \ \ for \ all \ x \in C$$

When f_0 is convex too, the equivalent conditions are sufficient for $\bar{x}$ to be globally optimal.

A.2.3. Constraint qualifications

To actually check if a point satisfies first order optimality conditions (such as (A.2)), analytic expressions for the geometric characterization described above have to be used. The applicability of those relies on constraint qualifications as presented below:

Definition A.10. *LICQ* (Linear Independence Constraint Qualification): Let $x \in C$ be a feasible point of the program (A.1). The linear independence constraint qualification, or LICQ, holds at x if the gradient vectors

$$\nabla g_i(x), \forall i \in \mathcal{I}_g,$$
$$\nabla h_j(x), \forall i = p+1, \ldots, q$$

are linearly independent.

The MFQ is weaker than the LICQ:

Definition A.11. *MFCQ* (Mangasarian Fromowitz Constraint Qualification): Let $x \in C$ be a feasible point of the program (A.1). The Mangasarian-Fromowitz constraint qualification, or MFCQ, holds at x if the gradient vectors

$$\nabla h_j(x) \ \forall j = p+1, \ldots, q$$

are linearly independent and there exists a vector $d \in \mathbb{R}^n$ such that

$$\nabla g_i(x)^T d < 0, \ \forall i \in \mathcal{I}_g,$$
$$\nabla h_j(x)^T d = 0, \forall i = p+1, \ldots, q \ .$$

An even weaker constraint qualification is given by the GCQ ([Fle05], Proposition 2.8):

Definition A.12. *GCQ* (Guignard Constraint Qualification): Let $x \in C$ be a feasible point of the program (A.1). We say that the Guignard constraint qualification, or GCQ, holds at x if

$$T_C(x)^\star = T_C^{lin}(x)^\star.$$

A.3. Distances between stochastic processes

This short overview of selected results on distances between stochastic processes and the stability of solutions of stochastic programs collects results from the papers [PK], [PP11] and [MP07].

A.3.1. Distance concepts

Transportation distance The transportation distance, also known as Wasserstein distance or Kantorovich distance, can be used to quantify the distance between two probability distributions.

Definition A.13. Given two probability spaces $\mathcal{P} := (\Xi, \Sigma, P)$ and $\mathcal{P}' := (\Xi', \Sigma', P')$ and a convex function $d : \Xi \times \Xi' \to \mathbb{R}$, the transportation distance of order $r \geq 1$ is the optimal value of the optimization problem

$$\begin{aligned} \mathbf{d}_{trans}(P, P') := \quad & min_\pi \left(\int d(\xi, \xi')^r \, \pi(d\xi, d\xi') \right)^{\frac{1}{r}} \\ s.t. \quad & \pi(M \times \Xi') = P(M) \;\; (M \in \Sigma) \\ & \pi(\Xi \times N) = P'(N) \;\; (N \in \Sigma') \end{aligned} \tag{A.3}$$

where the minimum in (A.3) is among all bivariate probability measures $\pi \in \mathcal{P}(\Xi \times \Xi')$ which are measures on the product sigma algebra $\Sigma \otimes \Sigma'$.

For simplicity, we restrict ourselves to the case of $r = 1$.

The solution can be seen as transportation plan, which requires the smallest effort to transport the mass of the one marginal probability distribution to the other in terms of the transportation cost d on the observation space. Typically, d is a distance function; for our application, it will be the one-norm on the observation space, i.e. $d(\xi, \xi') = |\xi - \xi'|$.

Nested distance The nested distance is a generalization of the transportation distance to stochastic processes. For this, the nested distance has to take into account the dynamics of information revelation (i.e. sequences of generated filtrations) of the stochastic processes to be compared.

Let $(\xi_t)_{t \in \{0,\dots,T\}}$ be a sequence of random variables $\xi_t : (\Omega, \mathcal{F}) \to (\Xi_t, \Sigma_t)$ and denote with $\Xi := \Xi_1 \times \dots \Xi_T$ the product space equipped with the product sigma algebra $\Sigma := \sigma(\Sigma_1 \otimes \dots \otimes \Sigma_T)$.

Definition A.14. For nested distributions $\mathbb{P} := (\Xi, \Sigma, P)$, $\mathbb{P}' := (\Xi', \Sigma', P')$ and a real-valued, convex function $d : \Xi \times \Xi' \to \mathbb{R}$ the nested distance of order $r \geq 1$ is

the optimal value of the optimization problem

$$\begin{aligned} \mathbf{d}_N(\mathbb{P},\mathbb{P}') := \quad & min_\pi \left(\int d\left(\xi,\xi'\right)^r \pi\left(d\xi, d\xi'\right) \right)^{\frac{1}{r}} \\ s.t. \quad & \pi\left(M \times \Xi' | \Sigma_t \otimes \Sigma_t'\right) = P(M|\Sigma_t) \ \ (M \in \Sigma_T,\, t \in \{0,\ldots,T\}) \\ & \pi\left(\Xi \times N | \Sigma_t \otimes \Sigma_t'\right) = P'(N|\Sigma_t') \ \ (N \in \Sigma_T',\, t \in \{0,\ldots,T\}) \end{aligned} \tag{A.4}$$

where the infimum in problem (A.4) is among all bivariate probability measures $\pi \in \mathcal{P}(\Xi \times \Xi')$, which are measures on the product sigma algebra $\Sigma_T \otimes \Sigma_T'$.

Again we restrict ourselves to the case of $r = 1$.

The particular usefulness of the nested distance for approximating stochastic optimization problems is rooted in the preservation of regularity properties (Lipschitz and Hölder continuity) of the objective function of multistage stochastic programs.

A.3.2. Lipschitz property of the value function for nested and transportation distance

Consider the multistage stochastic optimization problem

$$\begin{aligned} v = \quad & \min_{x_1,\ldots,x_T} \mathbb{E}(H(\xi_1, x_1, \ldots, \xi_T, x_T)) \\ s.t. \quad & x \in C, x_t \lhd \Sigma_t \end{aligned}$$

Theorem A.15. *[[PP12] Lipschitz property of the value function] Let $\mathbb{P}, \mathbb{P}'$ be two nested distributions. Assume that C is convex and the profit function H is convex in x for any ξ fixed,*

$$H(\xi, (1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)H(\xi, x_0) + \lambda H(\xi, x_1)\,.$$

Moreover let H be uniformly Hölder continuous ($\beta \leq 1$) with constant L_β that is,

$$\left|H(\xi, x) - H(\bar{\xi}, x)\right| \leq L_\beta \cdot \left(\sum_{t\in T} d_t(\xi_t, \tilde{\xi}_t)\right)^\beta$$

for all $x \in C$.
Then the value function v inherits the Hölder constant with respect to the multistage distance, that is,

$$\left|v(\mathbb{P}) - v(\tilde{\mathbb{P}})\right| \leq L_\beta \cdot \mathbf{d}_N(\mathbb{P}, \tilde{\mathbb{P}})^\beta$$

From this stability result we conclude that bounds on the distance between the "true" stochastic process and its deterministic approximation translate to bounds on the optimal *value* of the corresponding stochastic optimization problem and the optimal *value* of the deterministic approximation. Notice that this does not necessarily guarantee that *solutions* of the approximated problems perform well for the original problem.

Unfortunately, using the concept of nested distance for the deterministic approximation (i.e. determining tree structure, values and probabilities) of a given stochastic process leads to an optimization problem which is too difficult to be solved for trees of realistic sizes. As a consequence, actual computations have to rely on simpler concepts. This can be done by utilizing overestimators of the nested distance, which rely on stagewise decomposable distance concepts.
As a prerequisite, let $\mathbb{P}$ be a nested distribution, P its multivariate distribution, which is dissected into the chain of transition probabilities $P = P_1 \circ P_2 \circ \cdots \circ P_T$ of conditional distributions.

Definition A.16. $(K_2, \ldots, K_T)$ - Lipschitz property: Let P be a probability on $\mathbb{R}^{MT}$, dissected into transition probabilities $P_1, \ldots, P_T$. We say that P has the $(K_2, \ldots, K_T)$-Lipschitz property, if for $t = 2, \ldots, T$,

$$\mathbf{d}_{trans}(P_t(\cdot|u), P_t(\cdot|v)) \leq K_t \, d(u, v)$$

for the probability distribution $P_t(\cdot|u)$, conditional on u. $P_t(\cdot|u)$ is referred to as transition probability.

Under the assumptions from above, an upper bound on the nested distance which only relies on single stage transportation distances can be constructed as follows:

Theorem A.17. *[Lemma 3.11 [PP11]] If P has the $(K_2, \ldots, K_T)$- Lipschitz property, then*

$$\mathbf{d}_N(\mathbb{P}, \tilde{\mathbb{P}}) \leq \sum_{t=0}^{T-1} \underset{\xi}{sup} \left[\mathbf{d}_{trans} \left(P_{t+1}(\cdot|\xi), \tilde{P}_{t+1}(\cdot|\xi) \right) \right] \cdot \Pi_{s=t+1}^{T} (K_s + 1) \tag{A.5}$$

This theorem provides the rationale behind the algorithm introduced in Section 4.3.2 which works solely on single stage transportation distances.

Notice that while the theorems above provide a fundament to derive reasonable deterministic approximations for computing optimal values of the portfolio optimization problems [UL] and [LL], the statements are not directly applicable to the optimal value of the bilevel pricing problem itself. This is due to the fact, that the problem is nonconvex and thus violates the assumptions of the theorems; besides, no Lipschitz properties of the optimal value function (here: the strike price) with respect to the underlying probability measure are evident.
Although it seems plausible that, with improving accuracy of the [LL] solution, the

solution of the bilevel problem will eventually converge to the true solution, more advanced convergence theory is necessary to rigorously describe the behavior of the deterministic approximation of the bilevel problem.
For an investigation of the stability of variational inequalities (covering also the bilevel pricing problem) we refer to [JW09] and the references therein.

B. Parameters used in Chapter 4

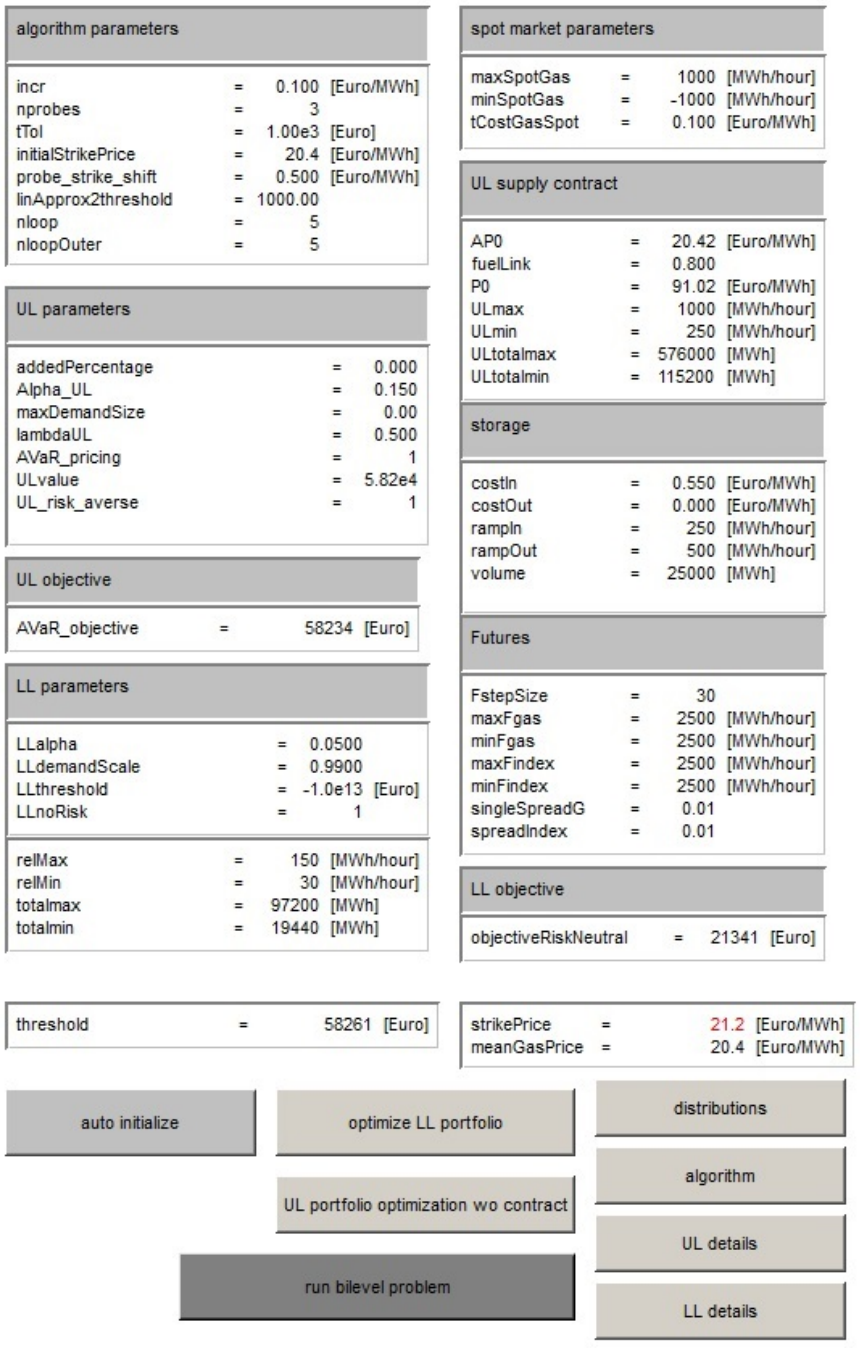


Figure B.1.: Parameters of the medium market liquidity model in Figure 4.10 (Screenshot from AIMMS)

C. Kurzzusammenfassung

Energie Swing Optionen sind Energieversorgungsverträge deren Halter (Energie Empfänger) die Liefermengen kurzfristig festlegen kann, wobei der Lieferpreis jedoch bereits im Vorhinein fixiert wurde. Aufgrund der speziellen Struktur von Energiemärkten ist die vollständige Absicherung derartiger Verträge durch am Markt gehandelte Produkte (d.h. Replikation) nicht möglich. Deshalb müssen von Energieproduzenten Preis- und Mengenrisiken aus dem Verkauf derartiger Versorgungsverträge besonders berücksichtigt werden. Das Ausmaß derartiger Risiken kann jedoch durch die Preissetzung in gewissem Umfang beeinflusst werden: Das Preisrisiko (und zum Teil auch das Mengenrisiko) kann durch bestehende Energieportfolios von Produzenten aufgefangen werden, während die Wahl des Lieferpreises das Bezugsverhalten von Energieempfängern und damit ebenfalls das Mengenrisiko beeinflusst.
Unter Berücksichtigung dieser Dynamik kann die Bestimmung des optimalen Lieferpreises als mehrstufiges stochastisches Bilevel Optimierungsproblem formuliert werden, in dem die optimale Entscheidung von Verkäufern (Setzen des Lieferpreises und Portfoliooptimierung) unter Antizipation der optimalen Entscheidung von Käufern (optimale Vertragsnutzung) getroffen wird. Dieser Zugang zur Bepreisung von Energie Swing Optionen ist einer der wenigen die die optimale Preissetzung und nicht die Bewertung zum Ziel haben. Insbesondere hängen die so ermittelten Preise von der konkreten Zusammensetzung der Erzeugungsportfolios der jeweiligen Verkäufer ab. Eine weitere Besonderheit dieses Modells ist, dass keine Annahmen über den Energiemarkt (wie Vollständigkeit, oder keine Arbitrage) getroffen werden müssen. Der Nachteil dieses Bepreisungsmodells ist die extreme Schwierigkeit der Lösung des Problemtyps der mehrstufigen stochastischen Bilevel Optimierung.
In der vorliegenden Arbeit wird ein maßgeschneiderter Algorithmus entwickelt, der eine effiziente Lösung des Bepreisungsproblems für lineare Portfoliomodelle ermöglicht. Eine Fallstudie belegt die praktische Anwendbarkeit der Methode für Probleme von realistischer Größenordnung. Dabei wird insbesondere der Einfluss des Erzeugungsportfolios von Verkäufern sowie die Auswirkung von geringer Marktliquidität auf den optimalen Lieferpreis untersucht.

D. Abstract

Energy swing options are energy delivery contracts whose holder can specify the amount of energy to be delivered on short notice at a prespecified delivery price per unit. Due to the special structure of energy markets replication of this kind of contracts by exchange traded products is not possible. Therefore, risks emanating from unknown prices and demand have to be considered specifically when selling this derivative. Whereas some of the risks may be absorbed by the energy portfolio of the producer to some extent, the size of these risks can also be influenced by the delivery price: The delivery price influences the exercise behavior of the buyer and by this also the risk exposure of the seller.
Incorporating these dynamics, determining the optimal delivery price leads to a multistage stochastic bilevel optimization problem: The optimal decision of the seller (setting the delivery price to be offered and optimization of the own generation portfolio) is determined anticipating its effect on the resulting optimal contract exercise policy of the buyer. The goal of this approach to price setting is to calculate a delivery price, instead of computing a cash value of the whole contract as usually done at derivative pricing. Also, widespread assumptions on the market such as completeness or no-arbitrage conditions are absent. Moreover, prices determined in the bilevel-model are depending on the initial generation portfolio of the seller, reflecting the seller's capabilities to absorb the buyer's consumption.
The drawback of the bilevel pricing model is the extreme difficulty of solving multistage stochastic bilevel problems.
In this thesis, a tailored algorithm is developed, which allows an efficient solution of the bilevel pricing problem for linear portfolio models. A case study illustrates the applicability of the method for problems of realistic size. In particular, the study investigates the influence of the composition of the seller's portfolio and market liquidity on the optimal choice of the delivery price.

DI. Peter Gross
Email: peter.gross@univie.ac.at
http://homepage.univie.ac.at/peter.gross/

Curriculum Vitae DI Peter Gross

Personal Information

Name	Peter Gross
Nationality	Austrian
Place of birth	Steyr

Education

Since Februar 2011	**PhD student at Universität Wien** (PhD Program in Statistics and Operations Research)
October 2003 – June 2009	**Technical Mathematics at TU Wien** Concentration on mathematical methods in economics Topic of Master thesis: Valuation of fossil fueled power plants with respect to operational constraints

Selected talks

IFIP , Klagenfurt, September 2013	The optimal choice of the strike price for energy swing options
ICSP ,Bergamo, July 2013	Price setting for energy swing options via stochastic bilevel problems
VOCAL, Veszprem, December 2012	Pricing energy swing options via stochastic bilevel programs
ISMP, Berlin, August 2012	Risk averse bilevel problems in energy markets

Publications

P. Gross, R. M. Kovacevic, G. Ch. Pflug (2013). Energy Markets,
In: Handbook of Risk Management for Energy Production and Trading (Kovacevic, Pflug and Vespucci Ed.), International Series in Operations Research and Management Science, Springer

Professional experience

July 2011 – June 2014	**Universität Wien**, Initiativkolleg Computational Optimization College assistant
September 2009 – June 2011	**Wagner, Elbling & Company** Advisor, natural gas portfolio management

Bibliography

[ADEH99] Phillippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9: 203–228, 1999.

[BBK08] Fred Espen Benth, Jurate Saltyte Benth, and Steen Koekebakker. *Stochastic Modelling of Electricity and Related Markets*. World Scientific, 2008.

[BCC12] J. Frédéric Bonnans, Zhihao Cen, and Thibault Christel. Energy contracts management by stochastic programming techniques. *Annals of Operations Research*, 200: 199–222, 2012.

[BGS07] Markus Burger, Bernhard Graeber, and Gero Schindlmayr. *Managing Energy Risk*. Wiley, 2007.

[BJ05] Mette Bjørndal and Kurt Jørensten. The Deregulated Electricity Market Viewed as a Bilevel Programming Problem. *Journal of Global Optimization*, 33: 465–475, 2005.

[BPM03] Anthony P. Burgard, Priti Pharkya, and Costas D. Maranas. Optknock: A bilevel programming framework for identifying gene knockout strategies for microbial strain optimization. *Biotechnology and Bioengineering*, 84: 647–657, 2003.

[BT97] Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, 1997.

[Bun07] Bundesverband der deutschen Gas- und Wasserwirtschaft (BGW). *Praxisinformation P 2007 / 13 Gastransport / Betriebswirtschaft*, 2007.

[CAC09] Miguel Carrión, Jose M. Arroyo, and Antonio J. Conejo. A bilevel stochastic programming approach for retailer futures market trading. *IEEE Transactions on Power Systems*, 24: 1446–1456, 2009, http://www.uclm.es/area/gsee/Archivos

[Car09] René Carmona. *Indifference Pricing: Theory and Application*, chapter 7 Applications to Weather Derivatives and Energy Contracts, pages 241–265. Princeton University Press, 2009.

[CCS13] Michael Coulon, Rene Carmona, and Daniel Schwarz. Electricity price modeling and asset valuation: A multi-fuel structural approach. *Mathematics and Financial Economics*, 7: 167–202, 2013. arXive.org @ 24.09.2012.

[CMS07] Benoit Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153: 235–256, 2007.

[CS00] Les Clewlow and Chris Strickland. *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, 2000.

[Dah05] M. Dahlgren. A continuous time model to price commodity-basedswing options. *Review of Derivatives Research*, 8: 27–47, 2005, http://www.springerlink.com/content/j600204722r04757/.

[DD10] Stephan Dempe and J. Dutta. Is bilevel programming a special case of a mathematical program with complementarity constraints? *Mathematical Programming*, 131: 37–48, 2010.

[Dem02] Stephan Dempe. *Foundations of Bilevel Programming*. Kluver Academic Publishers, 2002.

[Dem03] Stephan Dempe. Annotated bibliography on bilevel programming and mathematical programs with equilibrium constraints. *Optimization*, 52: 333–359, 2003.

[DS94] Freddy Delbaen and Walter Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300: 463–520, 1994.

[DZ11] S. Dempe and A. B. Zemkoho. The bilevel programming problem: reformulations, constraint qualifications and optimality conditions. *Math. Program.*, 2011.

[DZ12] S. Dempe and A. B. Zemkoho. On the Karush-Kuhn-Tucker reformulation of the bilevel optimization problem. *Nonlinear Analysis*, 75: 1202–1218, 2012.

[Eic07] Andreas Eichhorn. *Stochastic Programming Recourse Models: Approximation, Risk Aversion, Applications in Energy*. PhD thesis, Humboldt-Universität zu Berlin, 2007.

[EW03] Alexander Eydeland and Krzysztof Wolyniec. *Energy and power risk management: new developements in modeling, pricing and hedging*. John Wiley & Sons, 2003.

[EWE13] EWE GASSPEICHER GmbH, Moslestraße 7, 26122 Oldenburg. *Speichervertrag UGS Rüdersdorf*, 02 2013. accessed at 19.12.2013.

[FC08] J. David Fuller and William Chung. Benders decomposition for a class of variational inequalities. *European Journal of Operational Research*, 185: 76–91, 2008.

[FL10] Yueyue Fan and Changzheng Liu. Solving Stochastic Transportation Network Protection Problems Using the Progressive Hedging-based Method. *Networks and Spatial Economics*, 10: 193–208, 2010.

[Fle05] Michael Flegel. *Constraint Qualifications and Stationarity Concepts for Mathematical Programs with Equilibrium Constraints*. PhD thesis, Julius-Maximilian-Universität Würzburg, 2005.

[Gem05] Hélyette Geman. *Commodities and Commodity Derivatives, Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, 2005.

[HK09] Gido Haarbrücker and Daniel Kuhn. Valuation of electricity swing options by multistage stochastic programming. *Automatica*, 45: 889–899, 2009, http://www.sciencedirect.com/science/article/pii/S0005109808005621.

[JW09] Alejandro Jofré and Roger J.-B. Wets. Variational convergence of bivariate functions: lopsided convergence. *Math. Program. Ser. B*, 116: 275–295, 2009.

[KDPVK11] Vyacheslav V. Kalashnikov, Stephan Dempe, Gerardo A. Pérez-Valdés, and Nataliya I. Kalashnykova. Natural gas bilevel cash-out problem: Convergence of a penalty function method. *European Journal of Operations Research*, 215: 532–538, 2011.

[KP13a] Raimund M. Kovacevic and Georg Ch. Pflug. Electricity swing option pricing by stochastic bilevel optimization: a survey and new approaches. *European Journal of Operational Research*, (0): –, 2013, http://www.sciencedirect.com/science/article/pii/S0377221713010138.

[KP13b] RaimundM. Kovacevic and GeorgCh. Pflug. Pricing of energy contracts: From replication pricing to swing options. In Raimund M. Kovacevic, Georg Ch. Pflug, and Maria Teresa Vespucci, editors, *Handbook of Risk Management in Energy Production and Trading*, volume 199 of *International Series in Operations Research & Management Science*, pages 387–411. Springer US, 2013.

[KPVTK10] Vyacheslav V. Kalashnikov, Gerardo A. Pérez-Valdés, Asgeir Tomasgard, and Nataliya I. Kalashnykova. Natural gas cash-out problem: Bilevel stochastic optimization approach. *European Journal of Operational Research*, 206: 18–33, 2010.

[KRM06] Vyacheslav V. Kalashnikov and Roger Z. Rios-Mercado. A natural gas cash-out problem: A bilevel programming framework and a penalty function method. *Optim Eng*, 7: 403–420, 2006.

[KW14] Raimund M. Kovacevic and David Wozabal. A semiparametric model for electricity spot prices. *IIE Transactions*, 46: 344–356, 2014. to appear in IIE Transactions, accepted 2013.

[Löf04] J. Löfberg. YALMIP: A Toolbox for Modeling and Optimization in MATLAB. In *Proceedings of the CACSD Conference* , Taipei, Taiwan , 2004. http://users.isy.liu.se/johanl/yalmip .

[MP07] Radoslava Mirkov and Georg C. Pflug. Tree approximations of dynamic stochastic programs. *SIAM Journal of Optimization*, 18: 1082–1105, 2007.

[OG08] Steve Ohana and Helyette Geman. Time-consistency in managing a commodity portfolio: A dynamic risk measure approach. *Journal of Banking and Finance*, 32: 1991–2005, 2008.

[OG09] Steve Ohana and Hélyette Geman. Forward curves, scarcity and price volatility in oil and natural gas markets. *Energy Economics*, 31: 576–585, 2009.

[Out88] Jiři V. Outrata. A note on the usage of nondifferentiable exact penalties in some special optimization problems. *Kybernetica*, 24: 251–258, 1988.

[PB09] Georg C. Pflug and Nikola Broussev. Electricity swing options: Behavioral models and pricing. *European Journal of Operational Research*, 197: 1041–1050, 2009.

[Pen12] Teemu Pennanen. Indifference pricing in illiquid markets. http://www.mth.kcl.ac.uk/ teemu/pricing.pdf acessed at 28.04.2014, 2012.

[PK] Alois Pichler and Raimund Kovacevic. Scenario trees - a process distance approach. www.optimization-online.org, accessed at 30.04.2013.

[PP11] Georg Ch. Pflug and Alois Pichler. *Stochastic Optimization Methods in Finance and Energy*, chapter 15: Approximations for Probability Distributions and Stochastic Optimization Problems, pages 343–387. Springer, 2011.

[PP12] Georg Ch. Pflug and Alois Pichler. A distance for multistage stochstic optimization models. *SIAM Journal on Optimization*, 22: 1–23, 2012.

[PP14] Georg Ch. Pflug and Alois Pichler. *Multistage Stochastic Optimization*. Springer, to appear 2014.

[PR02] M. Patriksson and R. T. Rockafellar. A mathematical model and descent algorithm for bilevel traffic management. *Transportation Science*, 36: 271–291, 2002.

[PR07] Georg Ch. Pflug and Werner Römisch. *Modeling, measuring, and managing risk*. World Scientific, 2007.

[RU00] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2: 493–517, 2000.

[RU13] R. Tyrrell Rockafellar and Stan Uryasev. The fundamental risk quadrangle in risk management, optimization and statistical estimation. *Surveys in Operations Research and Management Science*, 18: 33–53, 2013.

[RW98] R. Tyrrell Rockafellar and Roger J-B. Wets. *Variational Analysis*. Springer, 1998.

[Ste12] Oliver Stein. How to solve a semi-infinite optimization problem. *European Journal of Operational Research*, 223: 312–320, 2012.

[SX08] Alexander Shapiro and Huifu Xu. Stochastic mathematical programs with equilibrium constraints, modelling and sample average approximation. *Optimization*, 57: 395–418, 2008.

[TB02] Chung-Li Tseng and Graydon Barz. Short-term generation asset valuation: A real options approach. *Operations Research*, 50: 297–310, 2002.

[TWR09] Angelos Tsoukalas, Wolfram Wiesemann, and Berc Rustern. *Lectures on Global Optimization*, chapter 10, pages 215–244. AMS, 2009.

[Wer04] Adrian S. Werner. *Bilevel stochastic programming problems: analysis and application to telecommunications*. PhD thesis, Norwegian University of Science and Technology, 2004.

[WS10] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2010.

[WTKR13] Wolfram Wiesemann, Angelos Tsoukalas, Polyxeni-Margarita Kleniati, and Berc Rustem. Pessimistic bilevel optimization. *SIAM Journal of Optimization*, 23: 353–380, 2013.

[YZ95] J. J. Ye and D. L. Zhu. Optimality conditions for bilevel programming problems. *Optimization: A Journal of Mathematical Programming and Operations Research*, 33: 9–27, 1995.

[ZYL12] Lei-Hong Zhang, Wei Hong Yang, and Li-Zhi Liao. On an efficient implementation of the face algorithm for linear programming, 2012. www.optimization-online.org, accessed at 14.04.2014.