

Behavioral pricing of energy swing options by stochastic bilevel optimization

Peter Gross¹ · Georg Ch. Pflug^{2,3}

Received: 20 January 2015 / Accepted: 12 January 2016 / Published online: 2 February 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Holders of energy swing options are free to specify the amounts of energy to be delivered on short notice, paying a fixed price per unit delivered. Due to the complexity of potential demand patterns, risk elimination by replication of these contracts at energy exchange markets is not possible. As a consequence, when selling delivery contracts, the energy producer has to explicitly consider the risk emanating from fluctuations in supply cost. The impact of these risk factors can be mitigated by the contract seller, who is an energy producer, to a certain extent: Supply cost fluctuations can be absorbed by the own generation portfolio whereas demand uncertainties can be influenced by the choice of the strike price, implicitly changing the buyer's behavior. Considering this, the determination of the optimal strike price can be formulated as a stochastic bilevel problem where the optimal decision of upper level player (price setting and production) depends on the optimal decision of a lower level player (demand depending on the price). We present a solution algorithm tailored to the resulting specific stochastic bilevel problem. We illustrate the effects of the behavioral pricing approach by studying behavioral price setting for natural gas swing options, highlighting in particular the influence of the seller's production and contract portfolio as well as of the market liquidity on optimal exercise prices.

✉ Georg Ch. Pflug
georg.pflug@univie.ac.at
Peter Gross
peter.gross@univie.ac.at

¹ Computational Optimization Project, University of Vienna, Vienna, Austria

² Department of Statistics and Operations Research, University of Vienna, Vienna, Austria

³ International Institute of Applied System Analysis, Laxenburg, Austria

Keywords Stochastic bilevel program · Swing option pricing · Energy portfolio optimization · Multistage stochastic optimization

Mathematics Subject Classification 97M40 Operations research, economics · 90C15 Stochastic programming · 91A65 Hierarchical games

1 Introduction

1.1 Energy portfolios from a financial point of view

In this paper, energy derivatives and energy portfolios are studied. We assume a financial point of view: objectives will be formulated, and all risks will be measured in terms of money. Energy derivatives are contracts specifying the physical delivery of energy and/or cash transactions which depend on energy prices. Energy portfolios are collections of energy derivatives and/or assets for generation, transformation, storage and transportation of energy.

For convenience, we consider spot market products as the smallest building blocks for energy and assume that generation and delivery 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. 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, with the most prominent example being futures contracts (or swaps). On the other hand, *flexible* contracts allow the holder to decide about the amount of energy to be delivered. Examples for flexible contracts are all types of options, in particular swing options on which we will put our focus in the rest of the paper.

1.2 Swing options

Energy swing options are flexible contracts which give their holder the right to order energy from the seller on short notice for a pre-specified price per unit delivered. Usually, these contracts constrain the consumption for the holder both from above and below, i.e. they restrict the feasible delivery decisions.

These tailor-made contracts can be used quite efficiently by their holders to absorb both fluctuations in demand (quantity risks) and spot prices (price risks); on the other hand they allow the seller to sell not only energy but also the flexibility of the own generation portfolio.

Contrary to simpler contracts where the price is determined by a deterministic upfront payment, we consider the exchange of two random cash flows: one of them is the random value of the consumed energy quantified at spot prices, the other is the value of the consumed energy at the pre-specified fixed exercise prices.

1.3 Pricing of swing options in “illiquid” markets

In the literature, the standard approach to the problem of pricing swing options is based on the buyer’s point of view. For a fixed strike price, the value of a contract is computed as the maximal expected payoff (under an equivalent martingale measure) from selling any energy sourced from the contract on the spot market. The corresponding price consists of a fixed payment equaling the value of the contract plus the payment of the strike price per consumed unit. This value resembles an objective market price for the contract, shared by all market participants which also excludes arbitrage opportunities for the buyer.

The drawback of objective market prices is that they rely on assumptions on the market (completeness and the no-arbitrage principle) which are in general not satisfied for energy markets. Relaxing these assumptions, valuations become subjective, i.e. dependent on risk preferences and especially on the agent’s portfolios. In particular, the incompleteness of energy markets implies that risks have to be explicitly considered when trading with this kind of contracts. In mathematical terms, the explicit incorporation of risk aversion leads to a replacement of the expected value by nonlinear functionals for computing the value of a contract.

As an immediate consequence of introducing this nonlinearity, the “value” of a whole portfolio is not equal to the sum of the values of its single contracts. Therefore, prices cannot be objective due to different portfolios of the agents and possibly due to their different risk preferences.

To deal with pricing in this framework of incomplete markets the indifference pricing framework can be used (cf. [6,8]). Here the optimal exercise price is the lowest price for which the the risk-adjusted value of the seller’s portfolio before and after selling the contract is equal; i.e. the seller is indifferent between selling and not selling the contract. In a first step, the level of risk of the seller’s actual portfolio has to be determined. In this paper we assume that this step was already done, resulting in a risk level of, say ρ . In a second step, a constraint is formulated, namely that the risk out of adding the contract to the portfolio must be limited by ρ .¹ Notice that in this pricing approach for swing options the focus is put on determining the optimal offered price (the strike price) and not on the value of the swing options as in the classical pricing approach.

1.4 The behavioral pricing problem

Clearly, the complexity of pricing problems is increased by including an illiquid market setting for both rigid and flexible contracts: The valuation of the whole portfolio of the

¹ In fact, we prefer here to use acceptability level, which can be seen as negative risk levels.

seller is necessary to determine indifference prices. However, in particular for flexible contracts, the problem is even more complicated due to the following interdependency of the strike price and the acceptability of the seller's portfolio:

Strike price \Rightarrow buyer's behavior \Rightarrow seller's risk \Rightarrow seller's acceptability.

The strike price offered by the seller directly influences the exercise behavior of the buyer. This in turn determines the risk emanating from the contract for the seller. Since in the illiquid market setting the risk from the contract has to be explicitly measured, this substantially influences the acceptability of the portfolio.

The problem of choosing the minimal acceptable strike price for the seller (i.e. the indifference price) under explicit consideration of the buyer's reaction will be referred to as the *behavioral pricing problem* (cf. Sect. 2.3.3 for a mathematical formulation).

The choice of the minimal strike price subject to optimal portfolio exercise is a stochastic optimization problem relying on the exercise policy of the holder, which is another (multistage) stochastic optimization problem. The holder's decision is determined by the seller's choice of exercise price; therefore the problem of determining the optimal exercise price is a stochastic multistage bilevel problem, also called a stochastic Stackelberg game. This problem class is NP-hard and extremely difficult to solve with standard methods, see for example [2–4] for overviews.

We call the seller's problem the *upper level [UL] problem* and the buyer's problem the *lower level [LL] problem*.

1.5 Scientific contribution and structure of the article

The behavioral pricing problem as described in Sect. 1.4 was introduced by [8] and additionally treated in [6]. In both papers, algorithmic approaches for solving the behavioral pricing problem are suggested. Numerical results are presented for rather stylized problems, where the different portfolios of the seller and the buyer of the swing option to be priced are represented by different price processes for both parties.

In this paper, which relies on the PhD thesis of Peter Gross [5], we extend these previous results by explicitly including the seller's portfolio in the pricing problem. To deal with this more complex model we introduce a tailored solution algorithm which reduces the original bilevel problem to a series of standard optimization problems. In a second step, this algorithm is extended by a heuristic which allows us to obtain high-quality solutions very fast. To analyze the dynamics modelled by the behavioral pricing approach and to illustrate the efficiency of our algorithm we perform a case study. The numerical results highlight the effects of market liquidity and the influence of the seller's portfolio on the indifference price for a natural gas swing option.

Additional to this introduction, the article consists of three sections: In the following Sect. 2, the mathematical model for the behavioral pricing problem and some notation will be introduced, essentially casting the problem already verbally described in Sect. 1.4 into mathematical formulation.

Section 3 contains some structural properties of the pricing problem and the tailored solution algorithm; the last Sect. 4 presents the numerical results.

2 The behavioral pricing problem for swing options

2.1 Preliminaries: risk factors, random variables and trees

As usual in stochastic optimization models, unknown quantities are modelled as random variables. In our case these are the spot prices $S(t)$, for stages $t = 1, \dots, T$ and the existing flexible energy supply obligations. Notice that, as functions of the available information, also decision variables in stochastic programs are random variables. Since we are in particular interested in the numerical study of behavioral pricing, our modeling is constrained by computational tractability. In particular, we assume all random variables are discrete and consequently, stochastic processes can be identified with scenario trees. Roughly speaking, this amounts to a representation of a stochastic process by several single scenarios (trajectories) and the flirtations by the branching of the scenarios (cf. [10] for details).

Concerning the notation for scenario trees we will use the following conventions: \mathcal{N} will represent the set of all nodes of the tree, \mathcal{S} the set of scenarios, i.e. the set of all path from the root to the leaves of the tree; each scenario is uniquely determined by its leaf node. For a given leaf node s let $\mathcal{N}_s \subseteq \mathcal{N}$ the sequence of its predecessor nodes, including s itself, but excluding the root. $|\mathcal{N}|$ and $|\mathcal{S}|$ denote cardinalities of the respective sets, i.e. the number of nodes and the number of scenarios. The cardinality of \mathcal{N}_s is $|\mathcal{N}_s| = T$. By $n-$ we denote the predecessor of a given node n .

For simplicity of notation we will use subindices both for referring to scenarios and nodes: For example, the probability of any scenario s is given by p_s , the probability of a node n by p_n . In the finite tree model, spot prices are associated to each node n of the tree and denoted by S_n .

2.2 The buyer’s problem

A swing option gives the holder the right to receive energy deliveries $y = (y(t))_{t=0, \dots, T-1}$ at a fixed strike price k which has been specified at contract closure. The required amount $y(t)$ can be communicated by the holder on short notice but is usually subject to (local and global) constraints

$$y^{(l)} \leq y(t) \leq y^{(u)}, \quad y^{(L)} \leq \sum_{t=0}^{T-1} y(t) \leq y^{(U)}.$$

Using a linear market model where agents are price takers, the profit of the [LL] resulting from the swing option is given by the sequence $(S(t) - k) \cdot y(t - 1)$. In the finite tree model, this expression is written as $(S_n - k) \cdot y_{n-}$ at node n . Notice that 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, which enforces all consumption decisions to be taken one stage in advance.

Remark 1 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 $k(0)$ to a price-index $I(t)$ of a competing fuel (often oil distillates). Using a weight κ and an initial index value $I(0)$, this leads to formulas such as $k(t) = k + \kappa (I(t) - I(0))$. In this case, the formula for the payoff stays the same, however it now depends on a two dimensional stochastic process (spot prices $S(t)$ and index values $I(t)$) and additionally the delivery decisions $y(t)$ of the holder.

For determining the payoff distribution of a contract, there still remains the issue of selecting the actual exercise strategy. Since we adopt here the standpoint of the contract seller, we typically do not know the risk profile and risk aversion of the contract holder. Therefore, a default choice is to assume an exercise strategy which maximizes the expected value of the derivative.

Using the finiteness of the underlying probability distribution, for spot prices S_n the [LL] problem can then be explicitly stated as

$$\max_{y_0, \dots, y_{|\mathcal{N}|-1}} \left\{ \begin{array}{l} \sum_{s \in \mathcal{S}} P_s (\sum_{n \in \mathcal{N}_s} y_n - (S_n - k)) \\ \text{s.t. } \left\{ \begin{array}{l} y^{(l)} \leq y_n \leq y^{(u)} \quad \forall n \in \mathcal{N} \\ y^{(L)} \leq \sum_{n \in \mathcal{N}_s} y_n \leq y^{(U)} \quad \forall s \in \mathcal{S} \end{array} \right. \end{array} \right\}. \tag{1}$$

Remark 2 Problem (1) resembles a knapsack problem, but notice that the decision has to be taken one period before the spotprices S_n become available.

Remark 3 Choosing the exercise policy which maximizes the expected value amounts to risk neutral valuation, which we will avoid for the seller, but accept for the buyer. As a matter of fact, risk neutral valuation is used only for simplicity of analysis. Relaxation of this assumption is straightforward by replacing the expected value by another acceptability functional; however in this case the buyer’s portfolio has also to be considered at determining the optimal exercise strategy. Numerical experiments suggest that the strike prices computed are rather robust with respect to risk preferences of the seller (cf. [5] for details).

2.3 The seller’s problem

2.3.1 Linear portfolio model

Mainly for computational simplicity, we assume that the payoffs of the portfolios under consideration are linear functions of the upper level decision variables x , such as trading or production decisions. This will be expressed by using a linear mapping Πx , where Π represents the random portfolio payoff of the seller; in addition linear constraints $Ax \leq b$ may be present for the decision variables.

To stay in the linear programming framework, we also use a linear market model, i.e. the actions of the agents considered do not influence market prices. Illiquidity will be modeled by constraints on market depth. Alternatively, one may represent market illiquidity by a nonlinear market model. This may also be treated in our framework as long as we stay in the framework of convex optimization.

Finally, also the risk sensitive functional for indifference pricing will be chosen in a way such that linear programming is applicable. That is, we use the average value-at-risk (conditional value-at-risk) as measure of acceptability of a risky portfolio.

Using the notation from the previous subsection, the payoff variable for the seller with initial portfolio $\Pi x \in \mathbb{R}^{|S|}$ for given decisions x and a sold swing option described by Problem (1) is given by

$$Y = \Pi x - \begin{pmatrix} \sum_{n \in \mathcal{N}_1} y_n^*(k) \cdot (S_n - k) \\ \vdots \\ \sum_{n \in \mathcal{N}_{|S|}} y_n^*(k) \cdot (S_n - k) \end{pmatrix},$$

where $y_n^*(k)$ denote the optimal exercise behavior of the buyer, i.e. the solutions of Problem (1), where k is the strike price.

2.3.2 The seller’s problem

Using the average value-at-risk at level α ($AV@R_\alpha$) as valuation functional, the acceptability of the portfolio with return Y can be determined by a linear optimization problem²

$$AV@R_\alpha(Y) = \max \left\{ a + \frac{1}{\alpha} \mathbb{E}[\min(Y - a, 0)] : a \in \mathbb{R} \right\}.$$

For a given optimal exercise behavior $y^*(k)$ corresponding to a strike price k , and using π_s to denote the s -th row of the matrix Π , the [UL] acceptability level can be formulated in the following way:

$$\max_{x; a, z} \left\{ \begin{array}{l} a - \frac{1}{\alpha} (\sum_{s \in \mathcal{S}} p_s z_s) \\ s.t. \begin{cases} a - \pi_s x + (\sum_{n \in \mathcal{N}_s} y_n^*(k) \cdot (S_n - k)) \leq z_s \quad \forall s \in \mathcal{S} \\ Ax \leq b \\ a \in \mathbb{R}, z = (z_s) \in \mathbb{R}_+^{|S|} \end{cases} \end{array} \right\} =: \Phi(k) \tag{2}$$

Here, the artificial variables $a \in \mathbb{R}, z_s \in \mathbb{R}_+$ have been introduced to lift the nonlinear valuation problem in x to a linear program in a, z and x . Our main interest will concern the dependency of the value of this acceptability on the exercise price k . For this reason we introduce notation $\Phi(k)$ for the [UL] acceptability level in dependency on k .

² See for example [13].

This formulation for acceptability of the seller’s portfolio can now be extended to describe the seller’s problem: Determining the smallest strike price k which results in an acceptability level above some threshold ρ can be written as

$$\min_k \left\{ k \mid \max_{x; a, z} \left\{ \begin{array}{l} a - \frac{1}{\alpha} (\sum_{s \in \mathcal{S}} p_s z_s) \\ a - \pi_s x + \left(\sum_{n \in \mathcal{N}_s} y_n^*(k) \cdot (S_n - k) \right) \leq z_s \quad \forall s \in \mathcal{S} \\ Ax \leq b \\ a \in \mathbb{R}, z_s \in \mathbb{R}_+ \end{array} \right. \right\} \geq \rho$$

or, by virtue of the relation $\max_{x \in X} cx \geq \rho$ iff $\exists x \in X : cx \geq \rho$ as

$$\min_{k; x; a, z} \left\{ \begin{array}{l} k \\ \text{s.t.} \left\{ \begin{array}{l} a - \frac{1}{\alpha} (\sum_{s \in \mathcal{S}} p_s z_s) \geq \rho \\ a - \pi_s x + \left(\sum_{n \in \mathcal{N}_s} y_n^*(k) \cdot (S_n - k) \right) \leq z_s \quad \forall s \in \mathcal{S} \\ Ax \leq b \\ a \in \mathbb{R}, z_s \in \mathbb{R}_+ \end{array} \right. \end{array} \right\}. \tag{3}$$

Notice that also the [UL] Problem (3) would be an LP, if $y^*(k)$ did not depend on k . But since this dependency exists, the problem is a hard one and approximative algorithms are needed for the practical calculation.

Remark 4 The [LL] Problem (1) may have a non unique solution for finitely many k ’s. In our formulation, it is sufficient that one element $y^*(k)$ of the solution set fulfills (3), i.e. we adopt the pessimistic and optimistic view of bilevel programming (At least one solution of [LL] is what the [UL] decision considers).

2.3.3 The behavioral pricing problem

The crucial part of the overall optimization problem is the dependency of the buyer’s exercise behavior y^* on the strike price k as stated in (1). Thus, putting the seller’s portfolio problem and the buyer’s optimal consumption problem together, we arrive at the bilevel model for the behavioral pricing problem:

$$\left[\begin{array}{l} \min_k \\ [UL] \quad \max_{x; a, z} \left\{ \begin{array}{l} a - \frac{1}{\alpha} (\sum_{s \in \mathcal{S}} p_s z_s) \\ a - \pi_s x + \left(\sum_{n \in \mathcal{N}_s} y_n^*(k) \cdot (S_n - k) \right) \leq z_s \quad \forall s \in \mathcal{S} \\ Ax \leq b \\ a \in \mathbb{R}, z_s \in \mathbb{R}_+ \end{array} \right\} \geq \rho \\ [LL] \quad y^*(k) = \underset{y=y_1, \dots, y_{|\mathcal{N}|}}{\text{Argmax}} \left\{ \begin{array}{l} \sum_{s \in \mathcal{S}} p_s \left(\sum_{n \in \mathcal{N}_s} y_n \cdot (S_n - k) \right) \\ y^{(l)} \leq y_n \leq y^{(u)} \quad \forall n \in \mathcal{N} \\ y^{(L)} \leq \sum_{n \in \mathcal{N}_s} y_n \leq y^{(U)} \quad \forall s \in \mathcal{S} \\ y_n = y_m \quad \forall n, m : n_- = m_- \end{array} \right\} \end{array} \right. \tag{4}$$

Problem (4) is the deterministic equivalent of a stochastic multistage bilevel program.

Recalling (3), it is immediately clear that as well [UL] as [LL] can be written as linear programs if the decision variables of the other level are fixed. In the full bilevel formulation, the variables appear in bilinear form. It is well known that even completely linear bilevel problems are NP-hard. In our case, however, the [LL] problem is only influenced by a single one-dimensional [UL] decision variable, the exercise price k . This property will be exploited to design an algorithm to specifically solve the behavioral pricing problem in the subsequent section.

Remark 5 More general setups are possible for the bilevel pricing problem, but they have been omitted from the presentation for simplicity. The results in the article hold for polyhedral acceptability functionals (see [11]), which - in the finite setting considered - can be written as standard linear programs where the random variable only appears at the right hand side of the constraints. An interesting example are linear combinations of average value-at-risk and the expected value.

Another detail omitted from the presentation for simplicity are additional constraints used for modeling market depth, which also couple [UL] and [LL] decision variables. Conceptually, these can be treated in a similar manner as the acceptability constraint. For more details we refer to [5].

3 An algorithm tailored for solving the behavioral pricing problem

3.1 Properties of the [UL] feasible region

Figure 1 illustrates the central aspects of the behavioral pricing problem.

To generate the figure, for a fixed strike price the corresponding [LL] consumption 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 .

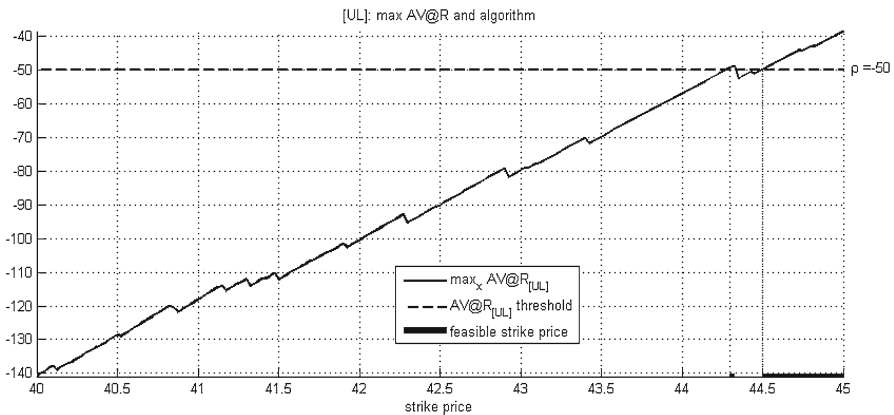
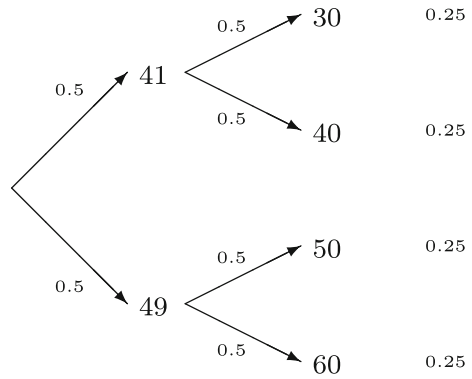


Fig. 1 One dimensional representation of the behavioral pricing problem. Notice the disconnected set of feasible strike prices marked by a thick black line on the x -axis.

Fig. 2 A simple example tree



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 Fig. 1.

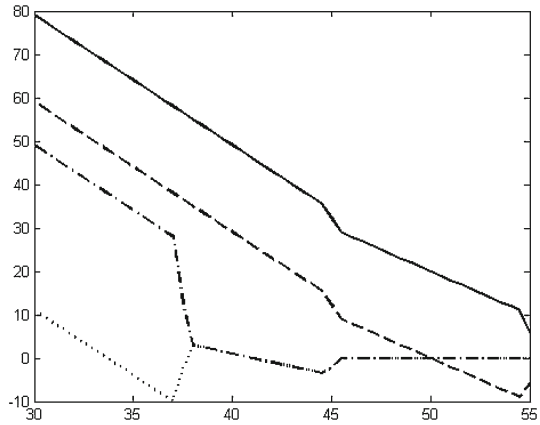
Remark 6 One might think that the [LL] profit distribution is monotonically decreasing (in the sense of first order stochastic ordering) w.r.t the strike price k . Notice that the same variable - with interchanges sign - describes the [UL] cost distribution before hedging. Nonmonotonicity of the [LL] profit distribution w.r.t. the strike price k is a caused by the global constraints $y^{(L)}$ and $y^{(U)}$. Without global constraints, the optimal [LL] decisions are trivial

$$y_n^* = \begin{cases} y^{(l)} & \text{if } \sum_{m \in n^+} p_m S_m \leq k \\ y^{(u)} & \text{if } \sum_{m \in n^+} p_m S_m > k \end{cases}$$

and the [LL] profit is clearly monotonically decreasing with increasing k . As an illustration of the case with global constraints, consider the simple scenario tree for the spot prices (S_n) shown in Fig. 2. For the choice $y^{(l)} = 0$ resp. $y^{(u)} = 2$ for all nodes and $y^{(L)} = 0$ resp. $y^{(U)} = 3$, the profit distributions as functions of k are depicted in Fig. 3. There are four scenarios, each occurring with probability 0.25, and the profits for the [LL] decision maker do not monotonically decrease, neither in the pointwise nor in the first order dominance sense.

An additional nonmonotonicity for the [UL] is caused by the fact that the cost variables for the [UL] can be well hedged by available hedging instruments for some specific k 's and for others not. The possible risk-reduction depends more on the structure of the cost variable than on the strike price k and this is another cause for nonmonotonicity. Thus the nonconnectedness of the [UL] feasible set is a fundamental property, even if the problem is formulated with continuous spot price distributions.

Fig. 3 The profit distributions of the lower level in dependence of the strike price k for the 4 scenarios (solid, dashed, dotted, dash-dotted). The distributions are not monotonic in the first order sense w.r.t. k



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

The maximal value $\Phi(\cdot)$ of the parametric optimization Problem (2) is numerically difficult to deal with. The following propositions describe an approximation from above, which is much simpler to handle:

Proposition 1 *Let $\Phi(k)$ be the maximal value function of (2). Then, for every choice of k_0 and the corresponding [LL] solution $y^*(k_0)$, a majorant of $\Phi(\cdot)$ is given by*

$$\hat{\Phi}_{k_0}(k) := \Phi(k_0) + \lambda_{\zeta}(k_0)^T \begin{pmatrix} \sum_{n \in \mathcal{N}_1} (y_n^*(k) \cdot (k - S_n) - y_n^*(k_0) (k_0 - S_n)) \\ \vdots \\ \sum_{n \in \mathcal{N}_{|S|}} (y_n^*(k) \cdot (k - S_n) - y_n^*(k_0) (k_0 - S_n)) \end{pmatrix} \geq \Phi(k) \tag{5}$$

where $\lambda_{\zeta}(k_0)$ is the vector of Lagrange multipliers associated to the constraints

$$a - \pi_s x - \left(\sum_{n \in \mathcal{N}_s} y_n^*(k) \cdot (k - S_n) \right) \leq z_s \quad \forall s \in \mathcal{S}$$

in the optimum of Problem (2) at fixed k_0 and $y^*(k_0)$.

Proof Recall the optimization problem defining $\Phi(k)$ by Eq. (2). Introducing the function

$$\zeta(k) = \begin{pmatrix} \sum_{n \in \mathcal{N}_1} y_n^*(k) \cdot (k - S_n) \\ \vdots \\ \sum_{n \in \mathcal{N}_{\mathcal{S}}} y_n^*(k) \cdot (k - S_n) \end{pmatrix},$$

we can rewrite Eq. (2) as

$$\Phi(k) = \max_{x: a, z} \left\{ a - \frac{1}{\alpha} \left(\sum_{s \in \mathcal{S}} p_s z_s \right) \mid Ax \leq b, a \cdot \mathbf{1} - \Pi x - z \leq \zeta(k) \right\}.$$

Consequently, a change of a given k_0 to some k triggers deviations $(\zeta(k) - \zeta(k_0))$, which can be interpreted as perturbations of the right hand side of inequality constraints of a convex program. Since $\Phi(k)$ is the maximal value of a convex program, the upper subgradient inequality (5) holds. In particular, the subgradient $\lambda_\zeta(k_0)$ is given by the dual variables associated to the perturbed constraints in the optimum (see for example [12], Theorem 10.13).

The majorant $\hat{\Phi}_{k_0}$ coincides with Φ locally around k_0 , which is established next. \square

Proposition 2 *Under the assumptions of Proposition 1, $\forall k_0$ there exists a closed interval $[\kappa_1, \kappa_2] \ni k_0, |\kappa_2 - \kappa_1| > 0$ on which $y^*(k)$ is constant and hence $\Phi(k)$ is continuous, nondecreasing and piecewise linear in k on this interval.*

In particular, there exists a subinterval $[k_1, k_2] \subseteq [\kappa_1, \kappa_2]$, where $\hat{\Phi}_0(k) = \Phi(k) \forall k \in [k_1, k_2]$.

Proof Since $y^*(k)$ is the solution of a parametric linear program, it is piecewise constant in k on closed intervals $[\kappa_1, \kappa_2]$ with disjoint interiors, see e.g. [7]. One may assume that $y^*(k)$ is unique in (κ_1, κ_2) and $|\kappa_2 - \kappa_1| > 0$.

Let $y^* \in [\kappa_1, \kappa_2]$ 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 $[k_1, k_2]$ containing k_0 , in which $\Phi(k)$ is linear and therefore any linearization $\hat{\Phi}_0(k)$ is exact for $k \in [k_1, k_2]$. \square

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

$$\hat{\Phi}_{k_i}(k) := \Phi(k_i) + \lambda_\zeta(k_i)^T \begin{pmatrix} \sum_{n \in \mathcal{N}_1} (y_n^*(k)(k - S_n) - y_n^*(k_i)(k_i - S_n)) \\ \vdots \\ \sum_{n \in \mathcal{N}_{|S|}} (y_n^*(k)(k - S_n) - y_n^*(k_i)(k_i - S_n)) \end{pmatrix}.$$

Denote by k^* the solution of the bilevel pricing problem, i.e. the minimal strike price which is [UL] acceptable

$$k^* := \min \{k \mid \Phi(k) \geq \rho\}. \tag{6}$$

Let k_1 be such that $\Phi(k_1) < \rho$ and set k_{i+1} as the minimal strike price, where the approximated acceptability $\hat{\Phi}_{k_i}(k)$ is above the threshold, i.e.

$$k_{i+1} := \min \left\{ k \mid \hat{\Phi}_{k_i}(k) \geq \rho \right\}.$$

Notice that for every i , k_i is a lower bound on the optimal strike price, since due to the inequality $\Phi(k) \leq \hat{\Phi}_{k_i}(k)$ it holds that $k_i \leq k^*$.

Using this notation, we derive a criterion for equality of k_i and k^* :

Proposition 3 *If $k_{i+1} = k_i$, then k_i is the optimal solution (the minimal acceptable strike price) k^* .*

Proof Assume $k_{i+1} = k_i$. From Proposition 2 we know that

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

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 $k_{i+1} = k_i$. Therefore, we have that $\rho = \hat{\Phi}_{k_i}(k_{i+1}) = \Phi(k_{i+1})$. □

3.3 The tailored algorithm

The tailored algorithm solves a sequence of auxiliary problems

$$k_{i+1} := \min \left\{ k \mid \hat{\Phi}_{k_i}(k) \geq \rho \right\}$$

which are based on a sequence of majorants $\hat{\Phi}_{k_i}(k)$. The solutions k_i are lower bounds for the minimal acceptable strike price k^* , which can be iteratively improved by the algorithm in Table 1.

Interpreting the subgradient inequality (5) as first order approximation of the function $\Phi(k) - \rho$, the algorithm resembles the Newton algorithm for finding zeros.

The algorithm in Table 1 is centered around two nested while loops:

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

$$\Phi(k_i) + \lambda_\zeta(k_i)^T \begin{pmatrix} \sum_{n \in \mathcal{N}_1} \left(y_n^*(\tilde{k})(k - S_n) - y_n^*(k_i)(k_i - S_n) \right) \\ \vdots \\ \sum_{n \in \mathcal{N}_{|S|}} \left(y_n^*(\tilde{k})(k - S_n) - y_n^*(k_i)(k_i - S_n) \right) \end{pmatrix} = \rho \quad (7)$$

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.

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

Table 1 Newton-type algorithm for solving Problem (6)

```

set initial strike price  $k_1$ , set  $i = 1$ 
compute  $y^*(k_1)$  and  $\Phi(k_1), \lambda_\zeta(k_1)$ 
set  $\tilde{\Phi} = \Phi(k_1)$ 
while  $\rho - \tilde{\Phi} > err$ 
  while  $\hat{\Phi}_i(k) \leq \rho$ 
     $\tilde{k} = k$ 
     $k = update(k)$ 
    compute  $y^*(k), \hat{\Phi}_{k_i}(k)$ 
  end while
  obtain  $k_{i+1}$  by solving equation (7) for  $k$ 
  compute  $y^*(k_{i+1})$  and  $\Phi(k_{i+1}), \lambda_\zeta(k_{i+1})$ 
  set  $\tilde{\Phi} = \Phi(k_{i+1}), i = i + 1$ 
end while
return  $k_i$ 

```

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. In practice however, 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 (1) requires the solution of [LL] as parametric linear program for parameter values between k_1 and the optimal k^* , which puts the main computational burden to the inner loop.

Figure 4 illustrates steps performed at an implementation of the algorithm in Table 1, and the rather high number of [LL] solutions required.

Proposition 4 *Given any starting value $k_1 \leq k^*$ and a sufficiently small choice of the tolerance err , the Newton-type algorithm 1 terminates in a finite number of steps at the optimal strike price k^* .*

Proof Proposition 1 implies that $\hat{\Phi}_{k_1}(k) \geq \Phi(k)$ and also $k_1 \leq k^*$; Proposition 2 that there exists an interval of exact approximation, $[a_1, b_1] \ni k_1$, where $\hat{\Phi}_{k_1}(k) = \Phi(k)$ and $\Phi(k)$ is nondecreasing.

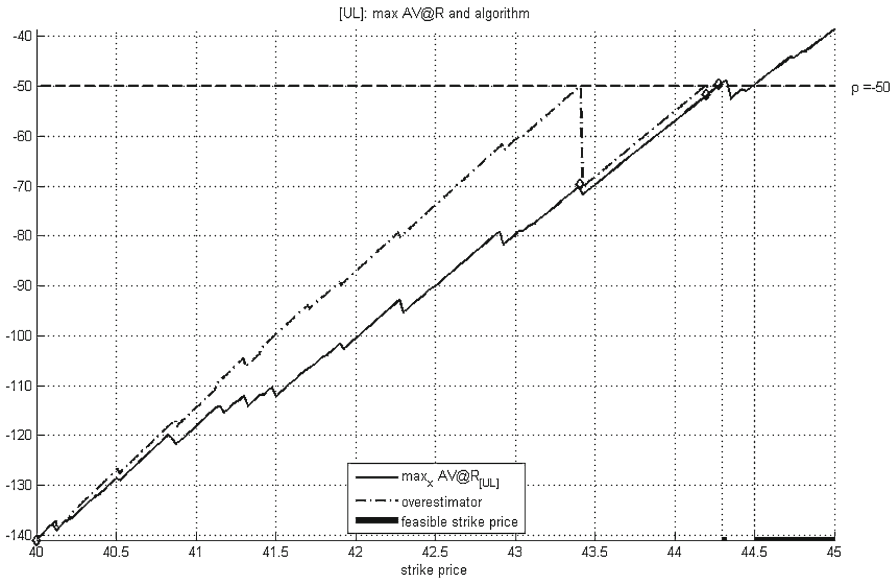


Fig. 4 Illustration of the Newton-type algorithm. For given strike price, diamonds represent evaluations of the whole problem ([UL] and [LL]) at iterations 2, 3 and 4, returning the subderivative of the [UL] acceptability. Dots mark evaluations of the [LL] problem needed for obtaining k_i . In the depicted case, convergence was reached after 430 solutions of [LL] and only 4 solutions of [UL]

Recall that $\Phi(k)$ is piecewise linear in k , thus the search interval $[k_1, k^*]$ 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 $k_i = k^*$.

Assume that $k^* \in [a_1, b_1]$. Then, $k_2 = \min \{k \mid \hat{\Phi}_{k_1}(k) \geq \rho\} = \min \{k \mid \Phi_{k_1}(k) \geq \rho\} = k^*$, i.e. the algorithm terminates at the exact solution k^* .

Otherwise, monotonicity of $\hat{\Phi}_{k_1}(k)$ on $[a_1, b_1]$ implies that $k_2 > k_1$ and $k_2 \notin [a_1, b_1]$. Thus, at least the interval of exact approximation (i.e. on which $\Phi(k)$ is exactly approximated by $\hat{\Phi}_{k_1}(k)$) is removed from the search interval $[k_1, k^*]$.

Repeating this step again either results in obtaining the optimal value or in cutting away at least one additional interval of exact approximation from the search interval. Since the initial interval $[k_1, k^*]$ 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 determination of $y^*(k)$ as the solution of the parametric [LL] program on the interval $[k_1, k^*]$.

The effort necessary for this strongly depends on the concrete structure of the linear program. In the worst case, the number of different values taken by the solution may be exponential in the dimension of the problem (cf. [1], p. 220). The modification in the following Sect. 3.4 focuses on designing a computationally cheap method to find a starting value k_1 which is already very close to the optimal solution. Nevertheless,

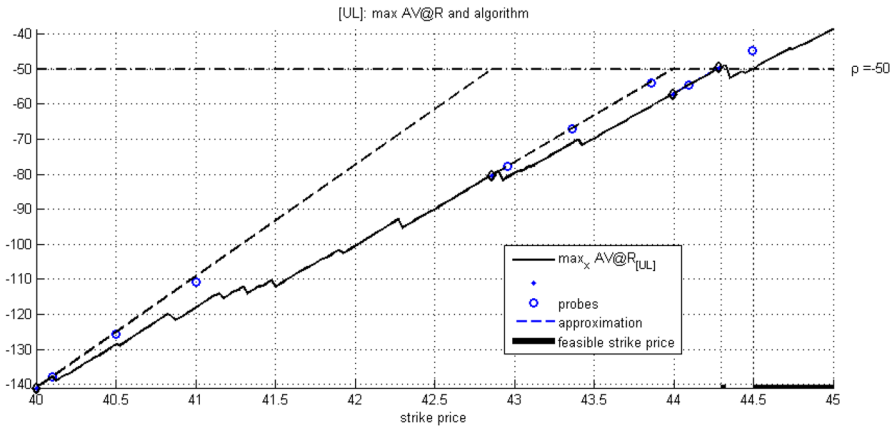


Fig. 5 Iterative construction of linear models for $\Phi_{k_i}(k)$ (dashed). Circles correspond to evaluations of [LL] at “probes” to obtain the slope of the linear models. Diamonds show evaluations of [UL]. In the case shown, convergence was reached after solving a total of 13 instances of [LL] and 4 instances of [UL]

the solution of the whole problem has to be determined only once per cut (to obtain $\Phi(k)$ and $\lambda_\zeta(k)$). 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 $\hat{\Phi}_{k_i}(\cdot)$ by a linear model

Although the dependency between k and $y^*(k)$ is rather difficult to model, the function $\hat{\Phi}_{k_i}(k)$ itself can be approximated from above quite well by a linear function.

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 probe values p_j for k , $\hat{\Phi}_i(p_j)$. The set of probes at a given iteration i , $probes_i = \{p_1, \dots, p_J\}$, consists of several strike prices which are above k_i . For example, the implementation used for the computations for Fig. 5 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 in a flexible way, 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 $y^*(k_{i+1})$ 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 the algorithm in Table 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 the algorithm in Table 1.

Table 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  $\Phi(k_1), \lambda_\zeta(k_1), \lambda_y(k_1)$ 
set  $\tilde{\Phi} = \Phi(k_1)$ 
while  $\tilde{\Phi} + tol < \rho$ 
    compute steepest slope for probes:

$$\Delta_i = \max \left\{ \frac{\hat{\Phi}_{k_i(p_j)} - \Phi(k_i)}{p_j} \mid p_j \in probes_i \right\}$$

    new strike price by extrapolation  $k_{i+1} = \frac{\rho - \Phi(k_i)}{\Delta_i} + k_i$ 
    compute  $y^*(k_{i+1})$  and  $\Phi(k_{i+1}), \lambda_\zeta(k_{i+1})$ 
    set  $\tilde{\Phi} = \Phi(k_{i+1})$ 
end while
return  $k_{i+1}$  as starting value for the algorithm in Table 1

```

Pseudocode for the procedure can be found in the algorithm in Table 2 and the corresponding computational illustration is depicted in Fig. 5.

In practice, this simple method shows remarkably good results, even for rather small³ 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 section, 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 the algorithm in Table 2 followed by the algorithm in Table 1 serves as the workhorse for all applications in the subsequent section.

Of course, the modification turns the algorithm into a heuristic, i.e. the guarantee 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 .

Remark 7 Another approach to improve initial values for the algorithm in Table 1 is to design overestimators using simplifications of the [LL] problem. One possi-

³ Here, about 1 % of ρ .

bility 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 non-anticipativity constraints results in scenariowise decomposability. The resulting lower bounds for the optimal strike price k^* 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 implemented 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 Numerical results

4.1 A sketch of applied modeling

4.1.1 Natural gas portfolio model

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.

The buyer is assumed only to use spot markets;⁴ 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.

As can be seen from the description of the portfolio, 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.

⁴ For reasons of simplicity, here no transaction costs are included.

4.1.2 Stochastic modeling of risk factors

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).

We use a hybrid model for our risk factors, linking oil prices and natural gas consumption to natural gas prices. The parameters of the nonlinear link function were identified by linear regression. Natural gas consumption is artificially generated from temperature data by a sigmoid function.

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, $t = 1, \dots, T$). The long time period in the first stage (consisting of 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. A MATLAB implementation of a dynamic tree generation algorithm proposed in [10], Sect. 4.3 was used for actual computations.

4.2 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 3. Notice that market liquidity does not affect index futures or the buyer’s problem.

Table 3 Parameter values for market liquidity settings

	Transaction cost (nat. gas spot), €/MWh	Market depth, MWh/h	Bid-ask spread (nat. gas futures)
High liquidity ^a	0	10000	0
Medium liquidity ^b	0.1	1000	2 %
Low liquidity ^c	0.5	500	10 %

^a See Fig. 6

^b See Fig. 7

^c See Fig. 8

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.

We analyze the case where the seller is long in natural gas (i.e. has more consumption obligations than supply obligations) and is risk averse ($AV @ R_{0.15}$ is used as measure for acceptability). Excess natural gas has to be sold either by entering futures contracts (hedging with exchange traded products), by selling a swing option or by trading at the spot market.

4.2.1 The case of high market liquidity (Fig. 6)

The initial payoff distribution (top picture in Figure 6) is practically replicated when a swing option is sold. The same applies for both the seller's ([UL], mid picture) and the buyer's ([LL], bottom picture) supply contract usage.

This setup leads to a [LL] exercise behavior of about 50 % minimal consumption and 10 % maximal consumption. This is the same situation as in the case of med market liquidity.

Cumulative consumption is nearly equal between the reference portfolio and the case of shorted swing option. The probability of maximal consumption from [UL] is high, at roughly 50 %.

The only significant difference of the setups compared in Figure 6 is the hedging behavior indicated in the vertical lines in the mid picture: the rightmost line (roughly 400,000 MWh) represents the cumulative amount hedged in the initial portfolio and the leftmost line (roughly 350,000 MWh) represents the cumulative amount hedged in case a swing option is sold. Recalling that the [UL] profit and exercise distributions in both cases are the identical, we conclude that selling a swing option partly substitutes hedging with natural gas futures.

4.2.2 The case of medium market liquidity (Fig. 7)

Due to the practically unchanged optimal strike price, the [LL] behavior remains identical to the case of high market liquidity.

However, significant differences in the [UL] results can be observed: Overall profits decrease as can be observed in the topmost picture in Fig. 7 (notice the different scales). The same is true for [UL] contract usage, depicted in the mid picture, where the probability of maximal consumption decreased from about 50 % to roughly 20 %. Both effects are probably due to transaction costs. Notice that in this case the payoff distribution and the consumption behavior of [UL] is very different if the portfolio is liquidated by trading at markets and in the case of selling the swing option.

4.2.3 The case of low market liquidity (Fig. 8)

The case of low market liquidity now leads to a significant change for the [UL] player: both, payoff and cumulative consumption drastically drop.

For the initial portfolio, the probability of minimal [UL] consumption has increased to roughly 20 %, whereas the probability of maximal consumption is practically zero.

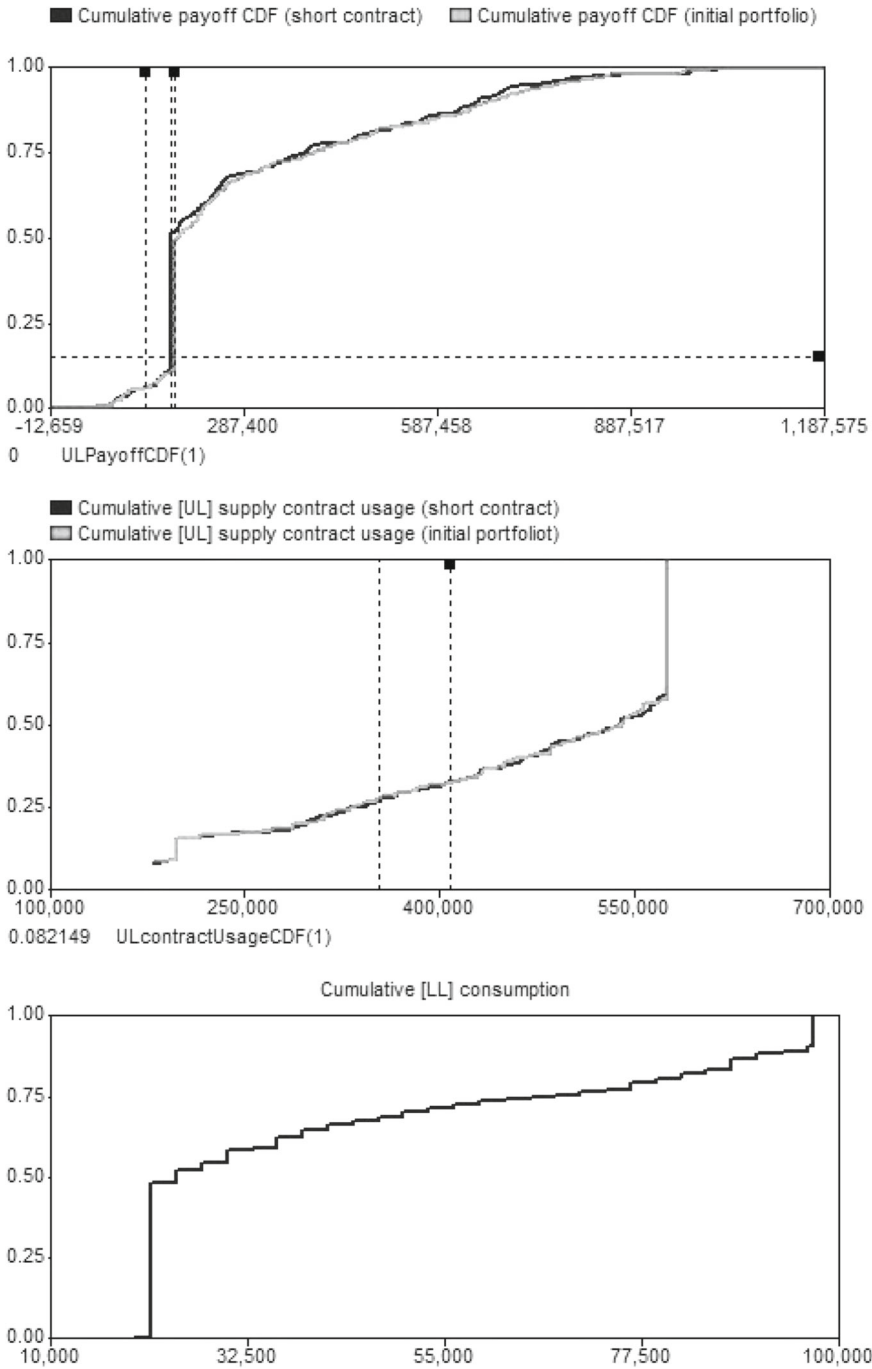


Fig. 6 Results of optimization in the reference case of high market liquidity (AIMMS screenshot). The optimal strike price is 21.4 €/MWh (expected natural gas price at 20.4 €/MWh)

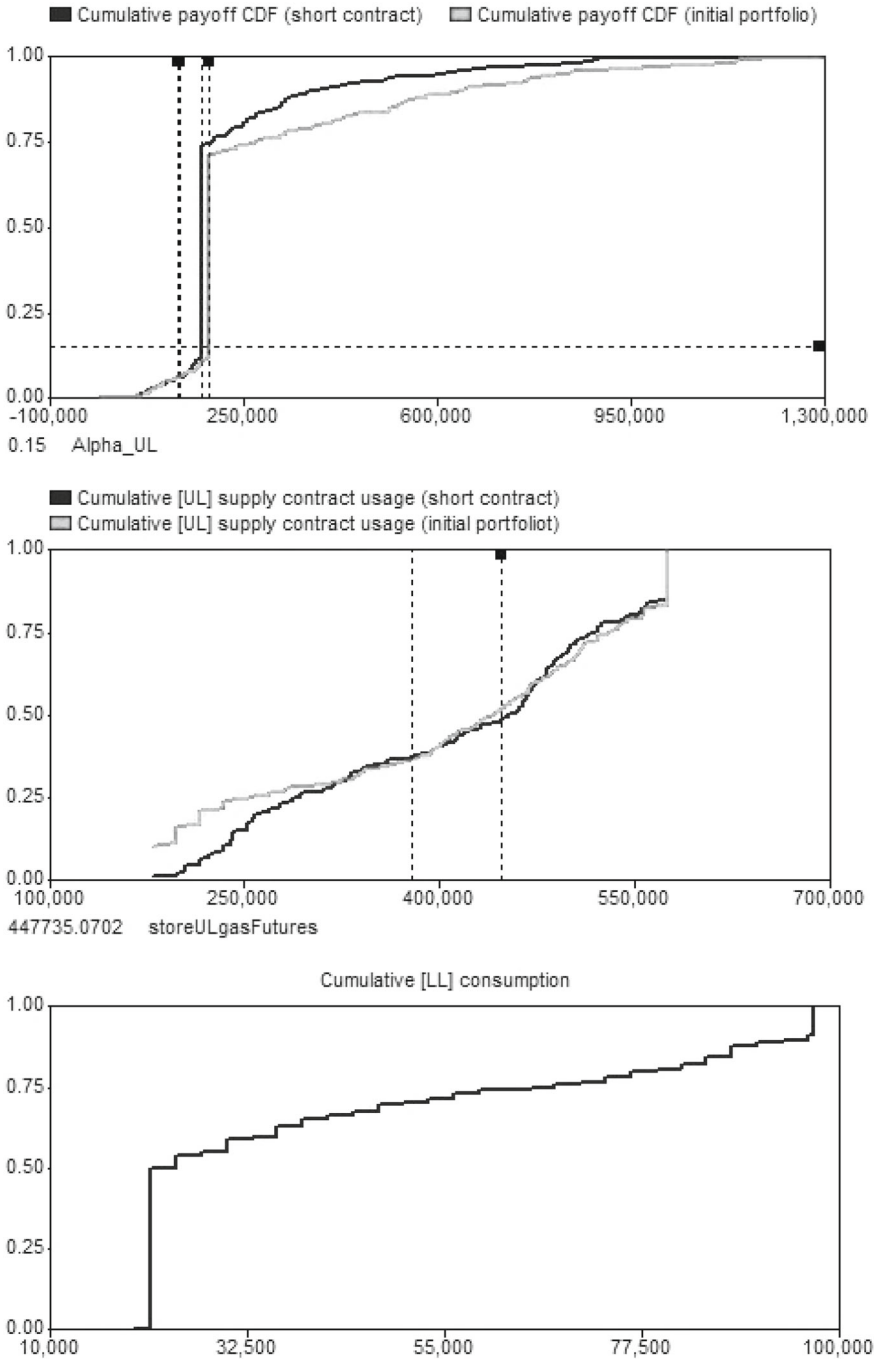


Fig. 7 Results of optimization for medium market liquidity (AIMMS screenshot). The optimal strike price is 21.2 €/MWh (expected value of natural gas spot price at 20.4 €/MWh)

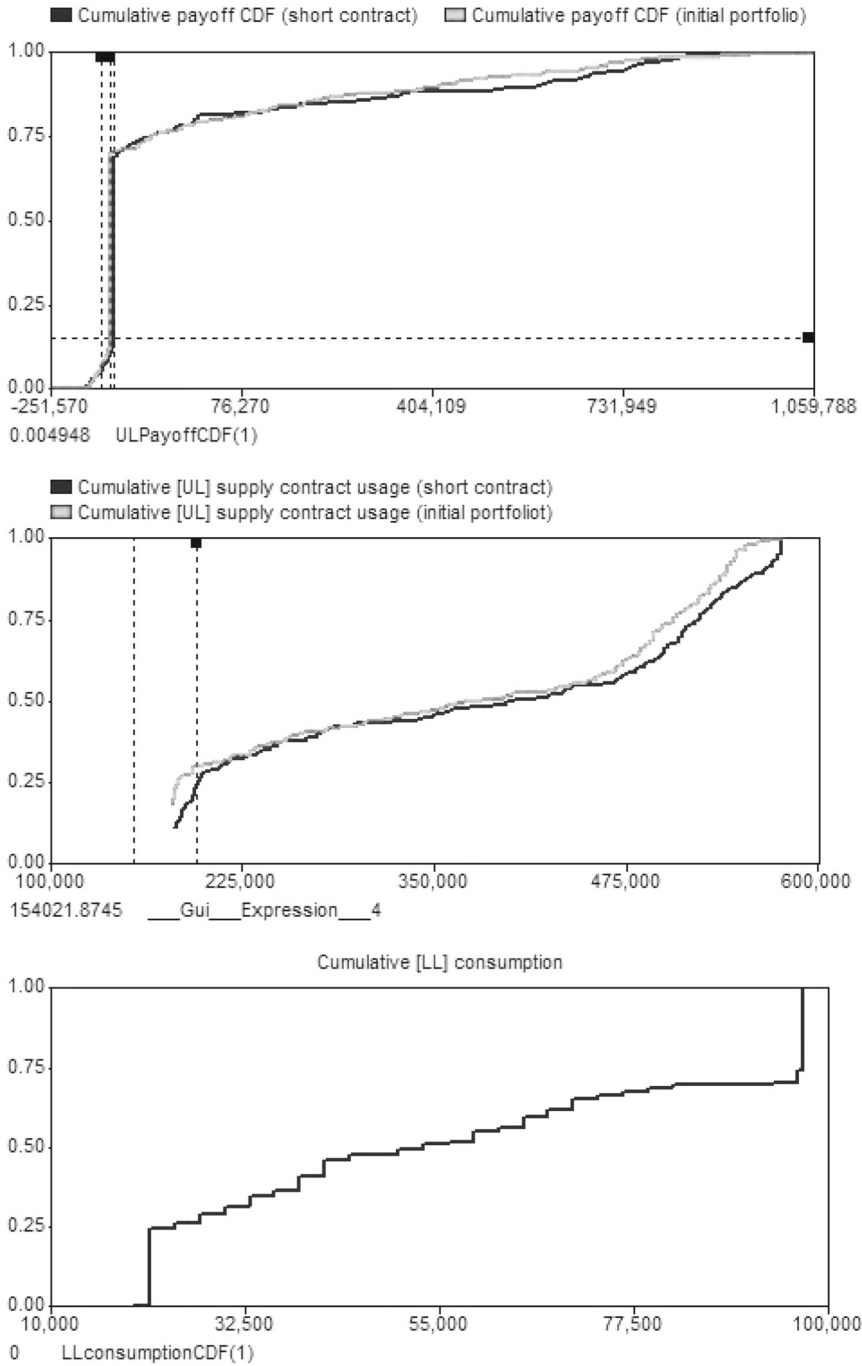


Fig. 8 Results of optimization for low market liquidity (AIMMS screenshot) The optimal strike price is 20.3 €/MWh (expected value of natural gas spot price at 20.4 €/MWh)

Table 4 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.

Initial demand liquidity	0 %, 25 %			75 %			100 %		
	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

The percentages in the leftmost column specify the allowed daily consumption of [LL] relative to the maximal daily consumption

Compared to this, when a swing option is sold, the probabilities of extremely low consumption is slightly decreased whereas the probability of consumption above average is increased.

When the swing option is sold, the indifference strike price essentially amounts to the average spot price, i.e. virtually no risk premium is required. The effect of this can be in the lowest picture of Fig. 8: compared the cases discussed above, now the probability of minimal [LL] consumption halved to 25 % whereas the probability of maximal consumption increased to about 25 %.

4.3 Dependency of prices and the portfolio of the seller

Table 4 contains the optimal exercise prices for several parameter configurations additionally 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. Note that for the extremal case of requiring 100 % consumption, the swing option resembles a futures contract. The lower the minimal consumption required, the more the swing option deviates from a futures contract.

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.

4.4 Conclusion

The tailored algorithm presented in this paper is able to compute the solution of the bilevel pricing problem quite efficiently by solving only few of the separated [UL] and [LL] problems.

Notice that although the specific bilevel problem can be solved 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.

In a case study, we studied the impact of market properties and the seller's portfolio on the optimal delivery price. We observed that in our model, a seller with excess natural gas requires a smaller strike price than a seller with an already balanced portfolio. We also see that [LL] flexibility directly influences the strike price: As expected the strike price increases with the "amount of [LL] flexibility" (i.e. minimal consumption requirements).

Finally, we observed a strong interdependence between the market model (i.e. settings for market liquidity), the optimal strike price and the seller's portfolio.

Acknowledgments Open access funding provided by [University of Vienna]. The authors gratefully acknowledge the support from the WWTF project ENERGY POLICIES AND RISK MANAGEMENT FOR THE 21st CENTURY.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Bertsimas, S., Tsitsiklis, J.N.: Introduction to linear optimization. Athema Scientific, USA (1997)
2. Colson, B., Marcotte, P., Savard, G.: An overview of bilevel optimization. *Ann OR* **153**, 235–256 (2007)
3. Dempe S.: Foundations of Bilevel Programming. Kluwer Academic Publishers (2002)
4. Dempe, S.: Annotated bibliography on bilevel programming and mathematical programs with equilibrium constraints. *Optimization* **52**, 333–359 (2003)
5. Gross P.: Behavioral pricing of energy swing options: A tailored solution algorithm for a multi-stage stochastic bilevel problem with applications in energy markets. PhD thesis, University of Vienna (2014)

6. Kovacevic, R.M., Ch, Pflug G.: Electricity swing option pricing by stochastic bilevel optimization: a survey and new approaches. *Euro J OR* **237**(2), 389–403 (2014)
7. Nozicka, F., Guddat, J., Hollatz, H., Bank, B.: *Theorie der linearen Optimierung*. Akademie Verlag, Berlin (1974)
8. Pennanen T.: Indifference pricing in illiquid markets. <http://www.math.kcl.ac.uk/teemu/pricing>
9. Brousseau, N., Pflug Ch, G.: Electricity swing options. *Euro J OR* **197**, 1041–1050 (2009)
10. Pflug G Ch., Pichler A. *Multistage Stochastic Optimization*. Springer, 2014
11. Eichhorn, A., Römisch, W.: Polyhedral risk measures in stochastic programming. *SIAM J. Opt.* **16**, 69–95 (2005)
12. Rockafellar R T., Wets R J B. *Variational Analysis*, Springer 1998
13. Rockafellar, R.T., Uryasev, S.: Optimization of conditional value-at-risk. *J Risk* **2**, 493–517 (2000)