DTU Library

Optimal Price-Energy Demand Bids for Aggregate Price-Responsive Loads

Saez Gallego, Javier; Kohansal, Mandi; Sadeghi-Mobarakeh, Ashkan; Morales, Juan M.

Published in:
I E E E Transactions on Smart Grid

Link to article, DOI:
10.1109/TSG.2017.2677974

Publication date:
2018

Document Version
Peer reviewed version

Optimal price-energy demand bids for aggregate price-responsive loads

Javier Saez-Gallego, Mahdi Kohansal, *Student Member, IEEE,* Ashkan Sadeghi-Mobarakeh, *Student Member, IEEE,* and Juan M. Morales, *Senior Member, IEEE*

***Abstract*—In this paper we seek to optimally operate a retailer that, on one side, aggregates a group of price-responsive loads and on the other, submits block-wise demand bids to the day-ahead and real-time markets. Such a retailer/aggregator needs to tackle uncertainty both in customer behavior and wholesale electricity markets. The goal in our design is to maximize the profit for the retailer/aggregator. We derive closed-form solutions for the risk-neutral case and also provide a stochastic optimization framework to efficiently analyze the risk-averse case. In the latter, the price-responsiveness of the load is modeled by means of a non-parametric analysis of experimental random scenarios, allowing for the response model to be non-linear. The price-responsive load models are derived based on the Olympic Peninsula experiment load elasticity data. We benchmark the proposed method using data from the California ISO wholesale electricity market.**

***Index Terms*—Price-energy bidding, demand response, electricity market, smart grid, data-driven.**

NOTATION

The main notation used throughout the paper is stated below for quick reference. Other symbols are defined as required.

A. Indexes and sets

- t Time period $t \in \{1, 2, \dots 24\}$.
- b Bidding block $b \in \{1, 2, \dots B\}$.
- w Realization of the stochastic variables, represented as scenarios $w = \{1, 2, \dots N\}$.

B. Input stochastic processes

- X Load.
- Λ^D Day-ahead price.
- Λ^R Real-time price.
- Π Retail price.

C. Decision variables

- X^D Stochastic process representing scheduled energy in the day-ahead market.
- $x^D_{t,w}$ Scheduled energy in the day-ahead market for time t and scenario w.
- $u_{t,b}$ Price bid for time t and block b.

Remark: a subscript t under the stochastic processes indicate the associated random variable for time t.

D. Parameters

- ϕ_w Probability of each scenario w.
- $x_{t,w}$ Load at time t and scenario w.
- $\lambda^D_{t,w}$ Day-ahead price at time t and scenario w.
- $\lambda^R_{t,w}$ Real-time price at time t and scenario w.
- $\pi_{t,w}$ Retail price at time t and scenario w.
- C_b Width of energy block b.
- L Fraction of the load that must be purchased in the day-ahead market.
- β Probability of occurrence of chance constraint.

I. INTRODUCTION

With the increasing deployment of smart grid technologies and demand response programs, more markets around the world are fostering demand bids that reflect the response of the consumers to changing electricity prices [1], [2]. In this paper, we consider the case of a retailer who procures energy to a pool of consumers in a typical two-settlement electricity market, as for example, the California wholesale electricity market CAISO [3]. The retailer submits price-energy demand bids to the day-ahead market, and only energy quantity bids to the real-time market in order to counterbalance the deviations from the scheduled day-ahead energy market to the actual load. The possibility of arbitrage is indirectly allowed depending on the submitted bid to the day-ahead market and the realization of the stochastic processes affecting the problem.

We assume that the load is price-responsive, in the sense that it may change depending on the price of electricity during the considered period. The retailer passes the retail price onto her consumers, who react accordingly. We do not make any assumption about the means that the consumers use to adjust their consumption based on the retail price, because the proposed methodology relies on historical data of aggregate load and retail price to estimate the relationship between them. Also because of this, we do not need to make any assumption on the nature of the load that the retailer aggregates. Furthermore, we consider that the retail price is directly linked to the market price and prefixed beforehand, for example, as a prearranged percentage of the day-ahead price. Therefore, the retail price is out of the retailer's control (in the short term at least). Finally, note that the communication flow between the retailer and her consumers is one-directional: the

J. Saez-Gallego, J. M. Morales (corresponding author) are with the Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark (email addresses: {jsga, jmmgo}@dtu.dk), and their work is partly funded by DSF (Det Strategiske Forskningsråd) through the CITIES research center (no. 1035-00027B) and the iPower platform project (no. 10-095378). M. Kohansal and A. Sadeghi are with the Department of Electrical and Computer Engineering, University of California, Riverside, USA. E-mail: {mkohansal, asadeh004}@ece.ucr.edu. Their work is supported in part by the United States National Science Foundation grants 1253516, 1307756, and 1319798.

price is communicated to the consumers by the retailer, who, in turn, observes the aggregate load.

The contributions of the paper are summarized as follows:

- An analytic solution to the problem of finding optimal block-wise price-energy demand bids in the day-ahead market when risk is not considered. Moreover, we propose a mixed-integer linear programming solution approach to the risk-averse case.
- The dynamic price-responsive behavior of consumers is modeled based on scenarios. The conditional probability of the load given a certain retail price trajectory is estimated using a non-parametric approach.
- We assess the practicality of the proposed methodology by using data from a real-world experiment.

The estimation of demand bids has been extensively studied in the past years [4, ch. 7]. Several papers share the common goal of estimating price-energy bids relative to specific types of load, for example, time-shiftable loads [5], electric vehicles [6] and thermostatically-controlled loads [7]. Our methodology differs with those in the fact that we do not make any assumption on the nature of the load. Methodologies based on forecasting tools [8], [9] generally do not make assumptions on the type of price-responsive load either, but, on the other hand, do not tackle the bidding problem.

Besides the works on forecasting, another group of papers focus on finding the optimal bid for generic loads. The work in [10] elaborates on a robust bidding strategy against procurement costs higher than the expected one, considering uncertainty in the prices only. Uncertain prices and demand are taken into account in [11] but minimizing imbalances and disregarding the economic side of the bidding. Our approach resembles that of [12] with the main differences being that we use data to estimate the price-response of the dynamic load, and that we consider energy-block bidding in a one-price balancing market as the US CAISO [3]. Authors of [13] consider, from the theoretical point of view, the problem of allocating a deterministic load by deciding which fraction should be purchased in the day-ahead market and which in the real-time market. Finally, authors of [14] study demand curves in an arbitrage- and risk-free situation by using a game theory.

Regarding the generation of scenarios of the stochastic processes, our methodology is inspired from [15]–[17]. From the application point of view, our approach differs in the final goal, as they deal with wind energy production. To our knowledge, there is no previous work that characterizes the dynamic price-responsive load with a set of scenarios. From the methodological point of view, our approach differs with the existing literature in the estimation of the conditional distribution of the price-responsive load, taking into account the full trajectory of the day-ahead price. This enables us to capture the full dynamics of the load across the hours of the next operational day. The real-time price is modeled in an analogous manner. In both cases, we model their distributions using a non-parametric approach that allows for non-linear responses to a given day-ahead price trajectory.

The paper is structured as follows. In Section II we introduce the retailer's bidding problem. Section III provides the analytic solution to the risk-neutral case. In Section IV we formulate the stochastic optimization model for solving the bidding problem with risk constraints. Section V elaborates on the scenario-generation technique. Next, in Section VI we analyze results from the bidding problem under the generated scenarios. Finally, in Section VII we draw conclusions and implications.

II. Problem Formulation

Consider a utility retailer/aggregator that seeks to maximize its profit based on the revenue that it collects from its loads, the payments it makes to the day-ahead market, and the payments it makes or receives in the real-time market. Mathematically speaking, we need to solve the following optimization problem:

$$\underset{X_t^D, u_{t,b}}{\text{Maximize}} \;\; \mathbb{E}\Bigg\{ \sum_{t=1}^{24} \Big(\Pi_t X_t - \Lambda_t^D X_t^D - \Lambda_t^R \big(X_t - X_t^D \big) \Big) \Bigg\} \tag{1a}$$

subject to

$$X_t^D = \sum_{b=1}^{B} C_b \; \mathbb{I}(u_{t,b} \geq \Lambda_t^D) \qquad \forall t, b \tag{1b}$$

$$u_{t,b+1} \leq u_{t,b} \qquad \forall t, b = 1 \dots B-1 \tag{1c}$$

$$\mathrm{P}\left(X_t^D \in [(1-L)\, X_t, (1+L) X_t] \right) \geq \beta \qquad \forall t \tag{1d}$$

$$\underline{\lambda} \leq u_{t,b} \leq \overline{\lambda} \qquad \forall t, b \tag{1e}$$

where $\mathbb{I}(\cdot)$ is the 0-1 indicator function.

The objective function (1a) is the expected total daily profit, composed of three terms. The first term represents the revenue that the retailer makes form selling energy to the consumers at the retail price. The second term represents the cost of purchasing energy from the day-ahead market. The third term accounts for the cost/revenue of purchasing/selling energy from/to the real-time market. The energy purchased or sold in the real-time market is equal to the difference between the purchased quantity at the day-ahead market and the realized load, i.e., $X_t - X_t^D$.

Constraint (1b) defines the scheduled energy in the day-ahead market to be equal to the sum of the width of the blocks of energy which have a price-bid higher than the market price. In other words, blocks of energy will be purchased if their price-bid is higher or equal to the day-ahead price. Note that $u_{t,b}$ is the decision variable which determines the shape of the submitted bidding curve to the day-ahead market.

Constraint (1d) models the risk-aversion of the retailer through two parameters. Parameter L represents the maximum fraction of the load that can be procured in the real-time market. This parameter could be defined by the retailer, but could also be constrained by the ISO as a way to avoid putting too much pressure on the real-time market, this way safeguarding and prioritizing the security and stability of the power system. Values of L close to 1 indicate that the full amount of the load can potentially be bought in the real-time market. On the other hand, as L decreases, we give priority to purchasing energy in the day-ahead market. Parameter β

indicates the minimum probability with which the constraint (1d) must be fulfilled. Values of β close to 1 indicate a hard constraint, while lower values of β indicate that the constraint is loose. The parameter β can be interpreted as the aversion of the retailer towards purchasing a certain fraction of the load in the day-ahead market. Low values of β can be interpreted as a sign that the retailer seeks to profit from arbitrage rather than from serving the load. As we show in the case study, higher values of β yield lower expected profit but also lower risk. Note that, for large L and small β, constraint (1d) becomes irrelevant, indicating the neutrality of the retailer towards risk.

Constraint (1c) ensures that the estimated bidding curve is monotonically decreasing which is a typical requirement in electricity markets. Finally, constraint (1e) set lower and upper bounds to the price bids, which are given by the market rules [18]. All in all, the expected profit depends on the decision variable "price-bid" and also on the realization of the input stochastic variables.

The maximum number of blocks that is allowed depend on the market rules [18] as well. The width of each block C_b must be set by the retailer depending on the magnitude of the load.

As in practice, here we assume that the retail price is given exogenously, in other words, it is not a decision variable of the retailer. The main driver for this consideration is the fact that the retail price must, to a certain extent, represent the true cost of electricity. This might not always be the case if the retail price is subject to the will of the retailer. As a consequence, the retailer's bidding strategy does not directly affect the behavior of the load, since the behavior of the load depends on the retail price and other factors such as the weather conditions. Another implication is that only the profit of the retailer is affected by her bidding strategy and the realized market prices.

III. Closed-Form Analytical Solution in Absence of Risk Constraints

In this subsection we elaborate on the closed-form analytic solution to problem (1), when the risk constraint (1d) is disregarded, or equivalently, when $L \to \infty$ and/or $\beta = 0$.

The retailer's bidding problem (1) can be decomposed by time period, so that 24 smaller optimizations problems can be solved instead, one for each time t.

In the risk-neutral case, each of these smaller optimization problems writes as follows:

$$\underset{X_t^D, u_{t,b}}{\text{Max.}} \left(\mathbb{E}\Big\{ X_t \left(\Pi_t - \Lambda_t^R \right) \Big\} - \mathbb{E}\Big\{ X_t^D \left(\Lambda_t^D - \Lambda_t^R \right) \Big\} \right) \quad (2)$$

subject to (1b), (1c) and (1e). The advantage of reformulation (2) is that we can perform simpler optimization problems in parallel. Note that the first term of (2) is constant with respect to the decision variables $u_{t,b}$ and X_t^D, whereas the last term is not. Hence, both the stochastic load X_t and the retail price Π_t can be dropped out of the optimization problem (1) in this case. Interestingly, this implies that, in the risk-neutral case, the retailer's optimal bidding strategy is not affected by the price-responsive nature of the load.

Next we analyze the case when Λ^D and Λ^R are statistically independent. Results are presented in Theorem 1. For ease of reading, and given that the maximization problem (2) can be decomposed per time period, we drop the time index t in the remaining of this section.

Theorem 1: The optimal price bid u_b^* in problem (2), when the day-ahead and real-time prices are independent, is equal to the expected value of the real-time price.

The proof of Theorem 1 is given in Appendix A. Theorem 1 also shows that, given the risk-neutral setup and independent prices, we do not obtain extra benefit from bidding a curve instead of a single price-quantity bid.

The assumption of statistically independent prices is not necessarily fulfilled in practice (see, for example, [19, Fig. 1]). For this reason, in Theorem 2 below, we provide the analytic solution to problem (2) when Λ^D and Λ^R are statistically dependent.

Theorem 2: A global optimum solution to problem (2) satisfies that the price bids for all blocks is equal to u^*. Moreover, u^* is equal to either $\underline{\lambda}$, $\overline{\lambda}$, or $\mathbb{E}\left\{\Lambda^R | \Lambda^D = u^*\right\}$ with $\frac{d}{du}\mathbb{E}\left\{\Lambda^R | \Lambda^D = u^*\right\} < 1$ in the latter case.

The proof of Theorem 2 is given in Appendix B. One could interpret the result of Theorem 2 in the following way: the optimal price bid will be the one for which price consistency holds, namely, for which the expected real-time price is equal to the day-ahead price. A second conclusion drawn from Theorem 2 is that the maximum profit is achieved with the same price-bid for each block. If there is more than one price bid that maximizes the expected profit (i.e., several global maxima), then the price bid for each block can be chosen indistinctly between them. Similarly as with Theorem 1, we do not obtain extra benefit from bidding a curve when prices are dependent.

From a practical point of view, Theorem 1 and 2 allow us to simplify the demand curve to a simple price-quantity bid. By taking into account this implication, we can obtain the optimal price bid in the case when the distributions of prices are discrete, which allow us to compute the optimal price bid when the uncertainty is modeled by scenarios. The optimal price bid can be chosen by evaluating the profit in the local maxima, which are characterized according to the following remark:

Remark 2: Given a discrete set of scenarios for the day-ahead and real-time prices, let us consider the re-ordered pair of terms $\{\lambda_w^D, \mathbb{E}\left\{\Lambda^R | \Lambda^D = \lambda_w^D\right\}\}$ such that $\lambda_w^D \leq \lambda_{w+1}^D$. Local maxima[1] are achieved at the stationary points $u^* = \lambda_w^D$ such that $\lambda_w^D \leq \mathbb{E}\left\{\Lambda^R | \Lambda^D = \lambda_w^D\right\}$ and $\lambda_{w+1}^D > \mathbb{E}\left\{\Lambda^R | \Lambda^D = \lambda_{w+1}^D\right\}$.

Note that, due to market rules, the price bid have a maximum and minimum allowed values. In practice, one needs to check also if the maximum profit is achieved when the price bid is equal to one of its bounds. Using Remark 2 one can find the optimal price-bid by just performing a finite set of simple calculations.

As a final remark, it is noteworthy to say that the results from Theorem 1 and 2 show that the solution to (2) does not depend on the retail price, neither on the load. From a practical point of view this means that the risk-neutral retailer acts as

[1] The proof is available upon request.

a financial trader, making profit by selling and buying energy in both markets.

IV. Scenario-based Solution in Presence of Risk Constraints

In this section we present a solution to problem (1) using a scenario-based approach. The input for every time t is a set of N scenarios, each one characterized by a realization of the retail price $\pi_{t,w}$, the day-ahead price $\lambda^D_{t,w}$, the real-time price $\lambda^R_{t,w}$, and the load $x_{t,w}$. Each scenario has a probability of occurrence of ϕ_w.

We reformulate constraint (1b) by adding a binary variable $y_{t,w,b}$. Then, constraint (1b) is replaced by:

$$\begin{aligned} x^D_{t,w} &= \sum_b y_{t,w,b} C_b && \forall t, w \\ u_{t,b} - \lambda^{DA}_{t,w} &\leq \mathrm{M} y_{t,w,b} && \forall t, w, b \\ -u_{t,b} + \lambda^{DA}_{t,w} &\leq \mathrm{M}(1 - y_{t,w,b}) && \forall t, w, b \\ y_{t,w,b} &\in \{0, 1\} && \forall t, w, b \end{aligned} \tag{3}$$

where M is a large enough constant. The equations above imply that $y_{t,w,b} = 1$ if $u_{t,b} \geq \lambda^D_{t,w}$ and 0 otherwise.

Next, we reformulate constraint (1d) by adding two extra binary variables. We first define $\underline{z}_{t,w} = 1$ if $x^D_{t,w} \leq (1 - L)x_{t,w}$, and $\underline{z}_{t,w} = 0$ otherwise. Secondly, we define $\overline{z}_{t,w} = 1$ if $x^D_{t,w} \geq (1 + L)x_{t,w}$, and $\overline{z}_{t,w} = 0$ otherwise. Consequently, the chance constraint (1d) can be replaced by the following set of equations:

$$\begin{aligned} & x^D_{t,w} - (1 - L)x_{t,w} \leq M(1 - \underline{z}_{t,w}) && \forall w \\ & -x^D_{t,w} + (1 - L)x_{t,w} \leq M \underline{z}_{t,w} && \forall w \\ & x^D_{t,w} - (1 + L)x_{t,w} \leq M \overline{z}_{t,w} && \forall w \\ & -x^D_{t,w} + (1 + L)x_{t,w} \leq M(1 - \overline{z}_{t,w}) && \forall w \\ & \frac{1}{N} \sum_w \left(\overline{z}_{t,w} + \underline{z}_{t,w} \right) \leq 1 - \beta. \end{aligned} \tag{4}$$

All in all, taking into consideration the reformulations presented above, the optimal price-bid is found by maximizing, for every time t,

$$\underset{x^D_{t,w}, u_b}{\text{Maximize}} \quad \sum_w \phi_w \bigg(\pi_{t,w} x_{t,w} - \lambda^D_{t,w} x^D_{t,w} - \lambda^R_{t,w} (x_{t,w} - x^D_{t,w}) \bigg) \tag{5}$$

subject to (1c), (1e), (3), and (4).

V. Scenario Generation

In this section we elaborate on the modeling of the stochastic variables by scenarios. The proposed approach to generate scenarios has several advantages. First, we do not need to make any assumption on the type of price-responsive load we model. The response of the load to the price is directly observed in the data and modeled by a non-parametric distribution. For this very same reason, the response of the load to the price is allowed to be non-linear. Second, it is a fast approach, hence, big datasets can be quickly processed. Finally, the proposed approach is adequate for bidding purposes, since forecasting the load is not the main goal of the paper but rather account for its uncertainty in order to make an informed decision.

Each scenario is characterized by a 24-long sequence of day-ahead prices, real-time prices, retail prices and observed load. The proposed method to approximate their joint distribution is summarized as follows. First of all, we model the marginal distribution of the day-ahead price. Note that the day-ahead price is not dependent on the real-time price, neither on the bid of a small price-taker consumer. Second, we model the distribution of the load conditioned on the retail price using a non-parametric approach. Lastly, we model the distribution of the real-time price conditioned on the day-ahead price. The real-time price depends on the day-ahead price, but not on the load of a price-taker retailer.

The rest of this section is organized as follows. First, in Section V-A, we briefly elaborate on the technique to generate scenarios of day-ahead price. Then, for each scenario of day-ahead price, we generate conditional scenarios of real-time price and load in Section V-B.

A. Day-ahead Price Scenarios

The first step in the scenario generation procedure is to model the day-ahead price using an Autoregressive Integrated Moving Average model (ARIMA). We choose the most adequate model according to the AICc criteria [22]. Using the estimated model, we draw scenarios using the methodology explained in [15]. Because the scenarios are used in day-ahead trading, they are generated in a rolling horizon manner everyday at 12:00 with a lead time of 13 to 36 hours.

B. Load and Real-time Price Scenarios

In this section, we elaborate on the proposed methodology to draw scenarios from the distribution of load conditioned on the retail price. The methodology to generate conditional real-time price scenarios is analogous, hence, we omit it for brevity.

For this subsection, we consider a scenario of day-ahead prices $\tilde{\lambda}^D = \{\lambda^D_1, \ldots, \lambda^D_{24}\}$ that is generated using the methodology explained in Section V-A. Under the considered setup, as explained in the sections above, the retail price is given exogenously. In the case study, we assume the retail price to be proportional to the day-ahead price, that is, $\Pi = k\lambda^D$. Therefore a scenario of retail price is directly specified from a scenario of day-ahead price.

The procedure outlined next allows us to weigh the historical trajectories, such that trajectories with a retail price "closer" to the given retail price $\tilde{\pi}$ weigh more. These weights are used later in this section to compute the conditional density function of the load, given $\tilde{\pi}$. To begin with, we define $\pi^{(j)}$ as the 24-long vectors of retail price, with each element referring to an hour of the day, and with j referring to the index of the historical day considered. Then, we compute the Euclidean distance $d^{(j)} = ||\pi^{(j)} - \tilde{\pi}||$. In this way, we "summarize" each historical price trajectory $\pi^{(j)}$ with a single value, so that trajectories "closer" to the given retail price $\tilde{\pi}$ have a lower distance. Next, we use a Gaussian kernel to weight trajectories,

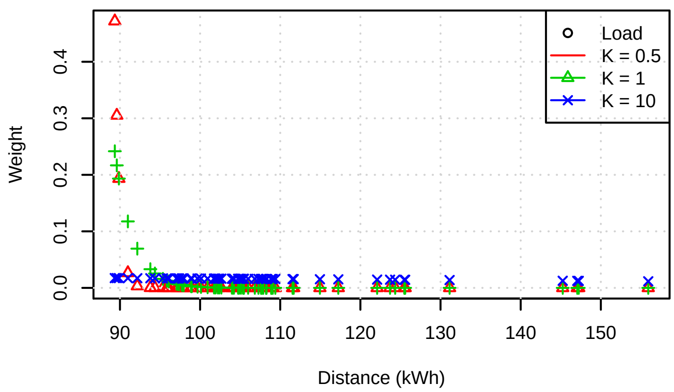

Fig. 1. The weights of the historical retail price trajectories are shown against their distance to the price reference.

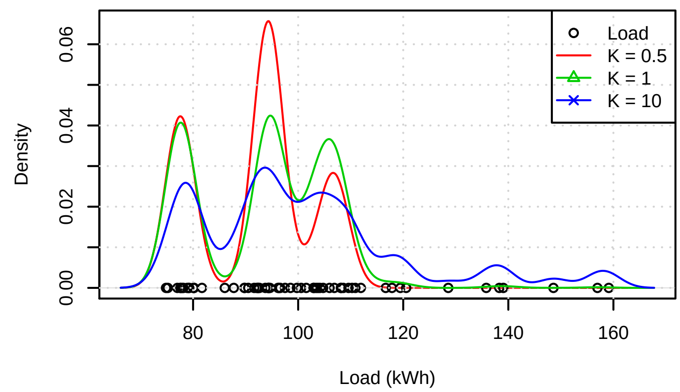

Fig. 2. The estimated conditional distribution of the load given the retail price is shown for different values of the bandwidth parameter K.

such that the weights are equal to $w^{(j)'} = f(d^{(j)})$, where f is the probability density function of a normal distribution with mean 0 and standard deviation σ_f. For the case study, we used $\sigma_f = K\sigma_d$, meaning that the standard deviation for f is equal to the standard deviation of the distances σ_d, multiplied by a *bandwidth* parameter K. Finally, we normalize the weight $w^{(j)} = \frac{w^{(j)'}}{\sum w^{(j)'}}$ so that their sum is equal to 1.

The effect of the bandwidth parameter K over the weights can be seen in Fig. 1. On its x-axis, we represent $d^{(j)}$ and on the y-axis the weights $w^{(j)}$. A smaller bandwidth penalizes price references further away. This is the reason why, when $K = 0.5$, there are few scenarios with a weight significantly greater than zero. On the other hand, when $K = 10$, all scenarios weigh similarly.

The procedure to generate each scenario is inspired from [15] and [16]. In short, we first transform the load data to a normal distribution using a non-parametric transformation. Then, we compute its covariance, and finally, generate random correlated Gaussian errors that are transformed back to the original distribution. The procedure consists of the following seven steps:

1) For each hour of the day, we compute a non-parametric estimation of the density of the price-responsive load [23] conditional on a retail price trajectory $\tilde{\pi}$. We do this by computing the kernel density estimator at hour t with the weights $w^{(j)}$ in the following way:
$$\hat{f}_t(x|\tilde{\pi}) = \frac{1}{J}\sum_{j=1}^{J} w^{(j)} G_h(x - x_t^{(j)}), \qquad (6)$$
where $G_h(x)$ is a kernel (non-negative function that integrates to one and has zero mean), h is its bandwidth, and $x_t^{(j)}$ is the observed load at time t and day j. An example of a estimated density using a Gaussian kernel is shown in Fig. 2, for different values of K and same h. For K close to zero ($K = 0.5$ in the case study), the weighting gives relatively high importance to few observations, therefore, the estimated density is more localized around them.
2) Using $\hat{f}_t(x|\tilde{\pi})$ from Step 1, we compute the cumulative density function, called $\hat{F}_t(x|\tilde{\pi})$.
3) The transformed load values $y_t^{(j)} = \hat{F}_t(x_t^{(j)}|\tilde{\pi})$ for every hour t follow a uniform distribution $U(0,1)$. Then, we normalize the load data through the transformation $z_t^{(j)} = \Phi^{-1}(y_t^{(j)})$, where $\Phi^{-1}(Y)$ is the probit function. Consequently, $(z_t^{(1)}, \ldots, z_t^{(J)}) \equiv Z_t \sim N(0,1)$.
4) We estimate the variance-covariance matrix Σ of the transformed load Z, relative to the 24 hours of the day. One could do it recursively as in [15].
5) Using a multivariate Gaussian random number generator, we generate a realization of the Gaussian distribution $\tilde{Z} \sim N(0, \Sigma)$.
6) We use the inverse probit function to transform $\tilde{Z}$ to a uniform distribution, that is, $\tilde{Y} = \Phi(\tilde{Z})$.
7) Finally, we obtain a scenario of load by transforming back $\tilde{Y}$ using the inverse cumulative density function from step 2, that is, $\tilde{x}_t = \hat{F}_t^{-1}(\tilde{Y}_t|\tilde{\pi}), \forall t$. Numerically, we use a smoothing spline to interpolate $\hat{F}_t^{-1}(\tilde{Y}_t|\tilde{\pi})$.

The procedure outline above generates a scenario of load conditioned on the retail price. Steps 5 to 7 are repeated as many times as needed if more scenarios of load per retail price are desired.

VI. Case Study

In this section we first introduce the datasets and the generated scenarios using the methodology from Section V. Then, in Section VI-B, we analyze in detail the solution of the bidding model with and without considering risk. Afterwards, in Section VI-C, we benchmark the performance of the proposed models and present the final conclusions.

A. The Data and Practical Considerations

The scenarios of day-ahead and real-time prices are generated using historical hourly values from CAISO [3]. We use three months of training data, from August to October 2014. The test period spans over November 2014. For the retail price and price-responsive load, we use data from the Olympic Peninsula experiment [21]. In this experiment, the electricity price was sent out every fifteen minutes to 27 households that participated in the experiment. The price-sensitive controllers and thermostats installed in each house decided when to turn

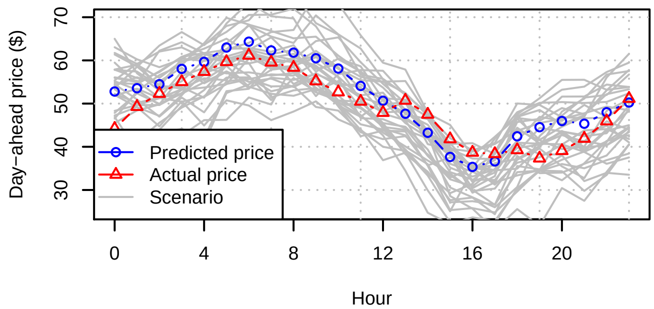

Fig. 3. Actual price, point forecast and generated scenarios for the day-ahead price.

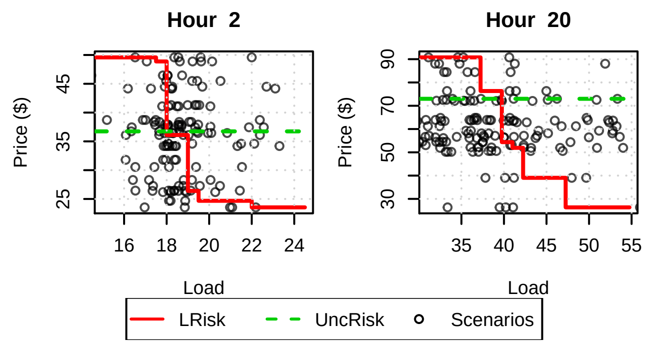

Fig. 4. The left figure is relative to the 2nd of November, while the right figure is relative to the 1st of November.

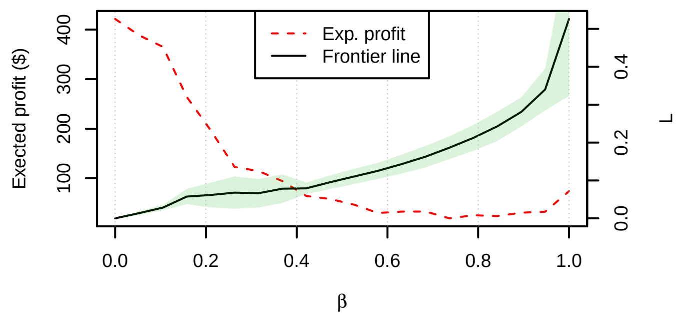

Fig. 5. On the left axis, in dashed red, the expected profit for every β, with L being in the feasibility frontier line. On the right axis, the feasibility frontier line is shown for combinations of β and L.

on and off the appliances, based on the price and on the house owner's preferences. The training and test months are the same as for the CAISO data, but relative to year 2006.

Some practical considerations need to be addressed. Firstly, that the day-ahead price and the retail price come from two different datasets. For this reason, prices are normalized. The second practical consideration is that we assume $\Pi = k\Lambda^D$ with $k = 1$ even though it is not fulfilled in practice. However, this does not affect the comparison of the proposed models, due to the fact that we use the same set of scenarios for the benchmark and for all the models. This issue could be solved in future work when data from new experiments becomes available.

Throughout the case study, we set a total of 20 blocks, where the width is equality distributed between a maximum and a minimum bidding quantities, set to be equal to the historical range of the scenarios at every hour. They are represented by the dotted lines in Fig. 6(b).

For the case study we use a total of 150 scenarios. For the estimation of the densities, we use a Gaussian kernel with a bandwidth h given by Silverman's rule of thumb [23]. Also, the bandwidth parameter is set to $K = 0.5$. For the model of the day-ahead price we use an ARIMA(3,1,2)(1,1,1) with a seasonal period of 24 hours. The Root Mean Square Error (RMSE) for the model of the day-ahead price (13 to 36 lead hours) is, on average, \$3.22, which is in line with the forecasting performance that other authors have achieved using similar methods [24].

A subset of the generated scenarios of day-ahead price is given in Fig. 3. By graphical inspection we conclude the scenarios of day-ahead price are a plausible representation of the actual day-ahead price and its uncertainty.

B. Model Analysis

To begin with, we discuss the results from the risk-neutral model (2). The solution to this model, for a given set of scenarios, is calculated either using Remark 2 or by solving (5) with $\beta = 0$. In Fig. 4, we show the scenarios of retail price and load in dots for hour 20 of November 1st, and hour 2 of November 2nd. The estimated bidding curve is displayed as a dashed green line. In accordance with Theorem 2, the resulting bidding curve is flat.

Next, we discuss the results from the risk-averse model (5). We start by analyzing the effect of the risk parameters L and β on the expected profit and the feasibility of the problem. In Fig. 5 we show, on the right axis, the feasibility frontier plot for L and β and its standard deviation in a shadowed area. We calculated it empirically, using data relative to the 1st of November. The combinations of L and β shown *below* the displayed dark line result, on average, in an infeasible solution. The frontier line is dependent on the scenarios of load: higher variability in the scenarios of load will require a greater value of L for the problem to be feasible. On the left axis of Fig. 5, we show the expected profit for the risk-averse problem, with the combination of β and L that lay on the frontier line. Naturally, the highest profits are achieved for low values of β, that is, when the retailer is less risk averse. From now on, we set the risk parameter β to 0.8. The value of L is chosen from the frontier plot, to be as small as possible.

Fig. 6(a) shows the scenarios of day-ahead price (continuous lines), together with the estimated optimal price bids (horizontal segments), for each hour. We observe that the magnitude of the price bid depends on the scenarios of day-ahead price. In Fig. 6(b), we show the amount of energy bought in the day-ahead market for each scenario and the span of the bidding blocks in dashed lines. In Fig. 6(c), we show the scenarios of load. On average, in day-ahead market we buy approximately the expected value of the load.

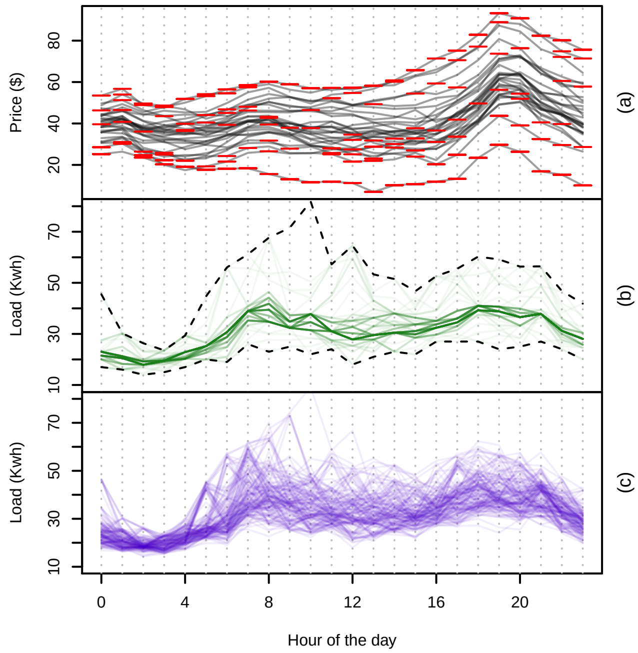

Fig. 6. In (a), the day-ahead scenarios (lines) are shown together with the estimated price-bids (horizontal segments). The day-ahead purchase for each scenario, and the load in each scenario, are shown in (b) and (c), respectively.

The estimated price bid by the risk-averse model is represented by the continuous red line in Fig. 4. Note that, at hour 2, the estimated price-responsiveness is much smaller than during hour 20. The reason is that, according to the scenarios of load, the load shows a lower variation during the early morning than during the early night.

C. Benchmark: Results in November and December

In this subsection, we benchmark the following models:

ExpBid: Single block model, where E_1 is equal to the expected value of the load, and the price-bid of the single block is equal to infinity. In other words, we always buy the expected load in the day-ahead market. No optimization is needed as the solution is trivial.

LRisk: Risk-averse model (5) with 20 bid blocks. The price-bid for each block is optimized.

UncRisk: Unconstrained risk model (2). The solution can be obtained by using Remark 2 or by solving (5) with $\beta = 0$.

In order to reproduce the real-time functioning of the markets, we validate the models using a rolling horizon procedure. Everyday at 12:00, we generate scenarios for the next operational day, and afterwards obtain the optimal bidding curve for all the benchmark models. The data from the last two months is used in the scenario-generation procedure, and the process is repeated daily all over the months of November and December.

In Table I we show the mean (1st column) and the standard deviation (2nd) of the profit for the three benchmark models, during November and December. We observe that the simple model *ExpBid* under-performs the rest of the models and, indeed, delivers a negative expected profit. The risk-optimized

TABLE I
MEAN AND STANDARD DEVIATION OF THE PROFIT FOR THE BENCHMARKED MODELS DURING NOVEMBER AND DECEMBER.

	Mean	Std. dev.
ExpBid	-1.78	34.52
LRisk	22.26	45.22
UncRisk	188.82	259.62

problem *LRisk* yields positive expected profit, with a variance greater than the *ExpBid* model but substantially lower than for the *UncRisk* problem. The risk-unconstrained model *UncRisk*, as anticipated, provides the highest mean returns.

VII. CONCLUSION

In this paper we consider the bidding problem of a retailer that buys energy in the day-ahead market for a pool of price-responsive consumers. Under the considered setup, the deviations from the purchased day-ahead energy are traded at the real-time market. We provide an analytic solution in the case that the retailer is not risk averse. Additionally, we formulate a stochastic programming model for optimal bidding under risk aversion. The price-responsiveness of the consumers is derived from a real-life dataset, and the modeling approach is non-parametric where non-linear relationships are allowed.

The analytic results show that, in the risk-unconstrained case, the optimal bid is a single price, meaning that there is no extra benefit from bidding a curve. On the other hand, the computational results from the risk-averse case show that a block-wise bidding curve successfully mitigates the risk in terms of profit volatility. Altogether, the proposed methodology allows the retailer to optimally bid in the day-ahead market, whether it is for expected-profit maximization (by leveraging arbitrage opportunities), or for the purpose of safely procuring energy.

Future work can follow several directions. From the viewpoint of the modeling of the retailer's trading problem, it could be of interest to compare our chance-constraint approach for risk control with alternative ones. For example, one could consider, instead of the chance constraint (1d), a weighted sum of the expected value of the retailer's profit and its Conditional Value at Risk (CVaR) in the objective function of problem (1), possibly establishing a link between the risk-related parameters that define both approaches. Likewise, it would be interesting to extend the retailer's trading problem (1) to allow for complex market bids. On a data-modeling level, one could explore how to extend the proposed method to generate scenarios when weather variables are considered. Finally, on a practical level, a possible avenue for future research would be to compare the results compiled in the case study of this paper with those that would be obtained for a control group of price-*irresponsive* loads. This way, we could properly assess how much the retailer benefits from supplying price-sensitive loads and accounting for their price-responsive nature.

References

[1] F. Rahimi and A. Ipakchi, "Demand response as a market resource under the smart grid paradigm," *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 82–88, 2010.

[2] N. O'Connell, P. Pinson, H. Madsen, and M. O'Malley, "Benefits and challenges of electrical demand response: A critical review," *Renewable and Sustainable Energy Reviews*, vol. 39, pp. 686–699, 2014.

[3] "California ISO Open Access Same-time Information System (OASIS)," Available: http://oasis.caiso.com.

[4] J. M. Morales, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, *Integrating Renewables in Electricity Markets: Operational Problems*. Springer Science & Business Media, 2013, vol. 205.

[5] H. Mohsenian-Rad, "Optimal demand bidding for time-shiftable loads," *IEEE Transactions on Power Systems*, vol. 30, no. 2, pp. 939–951, March 2015.

[6] M. Gonzalez Vaya and G. Andersson, "Optimal bidding strategy of a plug-in electric vehicle aggregator in day-ahead electricity markets under uncertainty," *IEEE Transactions on Power Systems*, vol. 30, no. 5, pp. 2375–2385, Sept 2015.

[7] S. Li, W. Zhang, J. Lian, and K. Kalsi, "Market-based coordination of thermostatically controlled loads–Part I: A mechanism design formulation," *IEEE Transactions on Power Systems*, vol. PP, no. 99, pp. 1–9, 2015.

[8] J. W. Taylor, L. M. De Menezes, and P. E. McSharry, "A comparison of univariate methods for forecasting electricity demand up to a day ahead," *International Journal of Forecasting*, vol. 22, no. 1, pp. 1–16, 2006.

[9] J. Hosking, R. Natarajan, S. Ghosh, S. Subramanian, and X. Zhang, "Short-term forecasting of the daily load curve for residential electricity usage in the smart grid," *Applied Stochastic Models in Business and Industry*, vol. 29, no. 6, pp. 604–620, 2013.

[10] K. Zare, A. J. Conejo, M. Carrión, and M. P. Moghaddam, "Multi-market energy procurement for a large consumer using a risk-aversion procedure," *Electric Power Systems Research*, vol. 80, no. 1, pp. 63–70, 2010.

[11] J. Saez-Gallego, J. M. Morales, M. Zugno, and H. Madsen, "A data-driven bidding model for a cluster of price-responsive consumers of electricity," *IEEE Transactions on Power Systems*, 2016.

[12] S.-E. Fleten and E. Pettersen, "Constructing bidding curves for a price-taking retailer in the Norwegian electricity market," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 701–708, 2005.

[13] Y. Liu and X. Guan, "Purchase allocation and demand bidding in electric power markets," *IEEE Transactions on Power Systems*, vol. 18, no. 1, pp. 106–112, 2003.

[14] A. B. Philpott and E. Pettersen, "Optimizing demand-side bids in day-ahead electricity markets," *IEEE Transactions on Power Systems*, vol. 21, no. 2, pp. 488–498, 2006.

[15] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, and B. Klöckl, "From probabilistic forecasts to statistical scenarios of short-term wind power production," *Wind Energy*, vol. 12, no. 1, pp. 51–62, 2009.

[16] J. M. Morales, R. Minguez, and A. J. Conejo, "A methodology to generate statistically dependent wind speed scenarios," *Applied Energy*, vol. 87, no. 3, pp. 843–855, 2010.

[17] X.-Y. Ma, Y.-Z. Sun, and H.-L. Fang, "Scenario generation of wind power based on statistical uncertainty and variability," *IEEE Transactions on Sustainable Energy*, vol. 4, no. 4, pp. 894–904, 2013.

[18] M. Kohansal and H. Mohsenian-Rad, "A closer look at demand bids in California ISO energy market," *IEEE Transactions on Power Systems*, 2016.

[19] H. Mohsenian-Rad. "Optimal bidding, scheduling, and deployment of battery systems in California day-ahead energy market." *IEEE Transactions on Power Systems*, vol.31, no. 1, pp. 442-453, 2016.

[20] J. Fortuny-Amat and B. McCarl, "A representation and economic interpretation of a two-level programming problem," *Journal of the Operational Research Society*, pp. 783–792, 1981.

[21] D. Hammerstrom, R. Ambrosio, J. Brous, T. Carlon, D. Chassin, J. DeSteese, R. Guttromson, G. Horst, O. Järvegren, R. Kajfasz *et al.*, "Pacific Northwest gridwise testbed demonstration projects. Part I. Olympic Peninsula project," October 2007.

[22] Burnham, Kenneth P., and David R. Anderson. "Multimodel inference understanding AIC and BIC in model selection." Sociological Methods & Research, vol. 33, no. 2, pp. 261-304, 2004.

[23] W. Hardle, *Applied nonparametric regression*, Cambridge, UK, 1990.

[24] J. Contreras, R. Espinola, F. J. Nogales, and A. J. Conejo, "ARIMA models to predict next-day electricity prices," *IEEE Transactions on Power Systems*, 18(3), 1014-1020. 2003

Appendix A
Proof of Theorem 1

We start by computing the expected profit, conditional on the day-ahead price (i.e., we treat $\Lambda^D = \lambda^D$ as a parameter). We disregard the first term in (2) since it is constant with respect to the decision variables u_b and X^D, and therefore, does not affect the solution. The expected profit (2) conditioned on the day-ahead price λ^D is thus given by

$$\mathbb{E}\left\{X^D\left(\Lambda^R-\lambda^D\right)|\Lambda^D=\lambda^D\right\}= \tag{7}$$

$$\sum_{b=1}^{B}\mathbb{I}\left(u_{t,b}\geq\lambda_t^D\right)C_b\left(\mathbb{E}\left\{\Lambda^R\right\}-\lambda^D\right). \tag{8}$$

Note that, since λ_t^D is given, X^D can be computed as $\sum_{b=1}^{B}\mathbb{I}\left(u_{t,b}\geq\lambda_t^D\right)C_b$. We distinguish three cases:

(a) When $\mathbb{E}\left\{\Lambda^R\right\}>\lambda^D$, the second term in (8) is positive, hence the expected profit is maximized when $u_b \geq u_{b+1}\geq\lambda^D,\ \forall b$. This implies that $u_B\geq\lambda^D$.

(b) When $\mathbb{E}\left\{\Lambda^R\right\}<\lambda^D$, the second term in (8) is negative, hence the profit is maximized when $u_{b+1}\leq u_b<\lambda^D,\ \forall b$. This implies that $u_1\leq\lambda^D$.

(c) When $\mathbb{E}\left\{\Lambda^R\right\}=\lambda^D$, any solution that satisfies $u_{b+1}\leq u_b$ is optimal.

Finally, we conclude that the expected value of the real-time price is an optimal price bid, since $u_b^*=\mathbb{E}\left\{\Lambda^R\right\}$ maximizes the retailer's expected income in the three cases above.■

Appendix B
Proof of Theorem 2

Analogously as in Appendix A, from Equation (8), the expected profit conditioned on $\Lambda_t^D=\lambda^D$ is proportional to

$$\sum_{b=1}^{B}\mathbb{I}\left(u_{t,b}\geq\lambda_t^D\right)C_b\Big(\mathbb{E}\left\{\Lambda^R|\Lambda^D=\lambda^D\right\}-\lambda^D\Big). \tag{9}$$

Next, recall that, from the basic properties of the expected value, $\mathbb{E}_X\left\{g(X)\right\}=\int_{-\infty}^{\infty}g(x)f_X(x)dx$. We compute the expected value of (9) with respect to Λ^D, which is equal to:

$$\int_{\Lambda^D}g(\lambda^D)f_{\Lambda^D}(\Lambda^D=\lambda^D)d\lambda^D \tag{10}$$

with $g(\lambda^D)$ equal to (9). Arranging terms, we obtain that (10) is equal to

$$\sum_{b=1}^{B}C_b\Bigg(\int_{-\infty}^{u_b}\left(\int_{-\infty}^{\infty}\lambda^R f_{\Lambda^R}(\lambda^R|\Lambda^D=\lambda^D)d\lambda^R\right)\times f_{\Lambda^D}(\lambda^D)d\lambda^D-\int_{-\infty}^{u_b}\lambda^D f_{\Lambda^D}(\lambda^D)d\lambda^D\Bigg). \tag{11}$$

Now we relax problem (2) by dropping constraint (1c). Then, the problem becomes decomposable by block, since (11) is a sum of B elements. For notational purposes, let us rename each of the B terms in the summation in (11) by $h_b(u_b)$. Note the functions $h_b(u_b)$ are continuous, since the integral of a continuous function is continuous. Then, for each block, the relaxed problem consists in maximizing the continuous function $h_b(u_b)$ subject to $\underline{\lambda}\leq u_b\leq\overline{\lambda}$. By the intermediate value theorem, we know that the maximum of each term in the summation will be achieved either at $u_b^* = \underline{\lambda}$, at $u_b^* = \overline{\lambda}$, or otherwise inside the interval $\left(\underline{\lambda}, \overline{\lambda}\right)$.

Considering the case when u_b^* is inside the interval, we proceed to find the u_b such that it maximizes $h_b(u_b)$. In order to achieve this, we calculate $\frac{d}{du_b} h_b(u_b) = 0$. Note that $\frac{d}{du} \int_{-\infty}^{u} \phi(x)dx = \phi(u)$. With this in mind, the derivative of $h_b(u_b)$ is equal to

$$\int_{-\infty}^{\infty} \lambda^R f_{\Lambda^R}(\lambda^R | \Lambda^D = u_b) d\lambda^R f_{\Lambda^D}(\lambda^D = u_b) \\ -u_b f_{\Lambda^D}(\lambda^D = u_b). \quad (12)$$

Assuming that $f_{\Lambda^D}(\lambda^D = u_b)$ is different than zero, and solving $\frac{d}{du_b} h_b(u_b) = 0$, we obtain the stationary point:

$$\left\{ u_b^* | u_b^* = \mathbb{E}\left\{ \Lambda^R | \Lambda^D = u_b^* \right\} \right\}. \quad (13)$$

Next we calculate the second derivative[2] of $\frac{d^2}{du_b^2} h_b(u_b^*)$. Its sign depends on the value of $(\frac{d}{du_b} \mathbb{E}\left\{ \Lambda^R | \Lambda^D = u_b^* \right\} - 1)$, which can be interpreted as the sensitivity of the expected real-time price to the day-ahead price at the stationary point. Depending on the sign of the second derivative, we distinguish three cases:

(a) When $\frac{d}{du_b} \mathbb{E}\left\{ \Lambda^R | \Lambda^D = u_b^* \right\} < 1$, u_b^* is a local maximum. From a practical point of view, it means that at day-ahead price $\lambda^D = u_b^*$, any marginal increase of this price will imply a comparatively lower marginal increase in the expected real-time price, hence, it becomes not profitable to buy energy from the day-ahead market at price levels greater than u_b^*.
(b) When $\frac{d}{du_b} \mathbb{E}\left\{ \Lambda^R | \Lambda_t^D = u_b^* \right\} > 1$, u_b^* is a local minimum.
(c) When $\frac{d}{du_b} \mathbb{E}\left\{ \Lambda^R | \Lambda^D = u_b^* \right\} = 1$, the solution u_b^* is an inflection point that delivers an expected profit equal to zero.

After having identified the possible candidates u_b that might maximize $h_b(u_b)$, it is easy to see that at least one global optimum to problem (2) satisfies that all u_b are all equal to each other, i.e., $u_b = u^*, \forall b$. This is so because functions h_b are all identical for all blocks, hence, the solution u_b^* that yields the highest expected profit for one block b will also deliver the highest expected profit for the remaining blocks.

Finally, we should point out that this global solution to the relaxed problem (2)—without constraint (1c)— naturally satisfies constraint (1c), hence, it must also be a global solution to the original problem (2).■

[2]The calculation of $\frac{d^2}{du_b^2} h_b(u_b^*)$, where u_b^* is given by (13), is available upon request.