

Robust self-scheduling of a price-maker energy storage facility in the New York electricity market

A. Barbry, M.F. Anjos,
E. Delage

G-2018-33

Mai 2018

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée: A. Barbry, M.F. Anjos, E. Delage (Mai 2018). Robust self-scheduling of a price-maker energy storage facility in the New York electricity market, document de travail, Les Cahiers du GERAD G-2018-33, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2018-33>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: A. Barbry, M.F. Anjos, E. Delage (May 2018). Robust self-scheduling of a price-maker energy storage facility in the New York electricity market, Working paper, Les Cahiers du GERAD G-2018-33, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2018-33>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2018
– Bibliothèque et Archives Canada, 2018

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2018
– Library and Archives Canada, 2018

Robust self-scheduling of a price-maker energy storage facility in the New York electricity market

Adrien Barbry^a

Miguel F. Anjos^b

Erick Delage^c

^a GERAD & Department of Energy Engineering, Polytechnique Montréal (Québec) Canada, H3C 3A7

^b GERAD & Department of Mathematics and Industrial Engineering, Polytechnique Montréal (Québec) Canada, H3C 3A7

^c GERAD & HEC Montréal, Montréal (Québec), Canada, H3T 2A7

adrien.barbry@polymtl.ca

miguel-f.anjos@polymtl.ca

erick.delage@hec.ca

Mai 2018

Les Cahiers du GERAD

G-2018-33

Copyright © 2018 GERAD, Barbry, Anjos, Delage

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract: Recent progress in energy storage have contributed to create large-scale storage facilities and to decrease their costs. This may bring economic opportunities for storage operators, especially via energy arbitrage. However, storage operation in the market could have significant impact on electricity prices. This work aims at evaluating jointly the potential operating profit for a price-maker storage facility and its impact on the electricity prices in New-York state. Based on historical data, lower and upper bounds on the supply curve of the market are constructed. These bounds are used as an input for the robust self-scheduling problem of a price-maker storage facility.

Keywords: Energy storage, electricity market, arbitrage, quantile regression, robust optimization

1 Introduction

Over the last five years, great progress have been achieved in the field of energy storage. Among the different technologies of energy storage, these progress have been especially significant in the field of batteries. A few years ago, provided their limited power and energy capacity, batteries were mainly considered as a mean to support renewables, damping the variability of wind and PV systems [1]. The recent deployment of large-scale batteries, highlighted by the 70 MW system brought online in California in late 2016 [16], brings new potential applications for batteries, which may now offer other services to the electricity grid.

At the same time, the deregulation of the electricity sector in many regions has led to the emergence of wholesale electricity markets and thus developed economic opportunities for energy storage [19]. The price volatility on these markets is indeed a potential source of profit for energy storage facilities, which can buy (and store) electricity during low demand (and low prices) periods and sell it back to the grid during high demand (and high prices) periods. This type of opportunity in the market is referred to as an energy arbitrage. Other applications exist and may be profitable for energy storage, especially ancillary services, and operating reserve [15]. Energy storage facilities may even participate simultaneously to several of these services to increase their profitability.

In this paper, we focus on energy arbitrage which is one of the best understood and more interesting applications in volatile markets. Such opportunities are especially present in New-York's day-ahead electricity market, which will be the focus of our study. In particular, Figure 1 shows the evolution of the hourly price during year 2016. Over the span of this year, the prices ranged between \$0.01/MWh and \$93.23/MWh. One can also observe a total of 1000 hours where a price lower than \$8.85/MWh was reached, and 1000 hours where prices were greater than \$32.55/MWh. The daily difference between off-peak price and peak-price is illustrated by Figure 2: this highlights the daily frequency of opportunities for energy arbitrage.

As a result of these opportunities, the energy storage sector is likely to attract investment in the years to come. However, investing in such large-scale facilities requires to evaluate accurately the potential benefits associated to energy arbitrage. One possibility to assess the profits of a large-scale energy storage facility is to perform backtesting to determine what would have been an optimal operating plan during the past years and compute the associated profit. The optimization of the strategy on electricity markets consist in formulating the optimal hourly bids for each day. In practice, in energy markets, the bids are classified into two types: self-schedule bids and economic bids [11]. A self-schedule bid does not include a price component. It indicates that the participant wants to buy (sell) electricity regardless of the price. An economic bid does include a price component. It indicates that the participant wants to buy (sell) electricity provided that the cleared market price is less (more) than the submitted price bid. In the case of energy storage, the self-scheduling formulation is generally preferred [17, 19, 23], since the hourly bids are interdependent. For instance, the storage operator needs to be certain that his purchase bids have been accepted if he plans to sell this electricity then.




Figure 1: Hourly electricity prices in NY market in 2016

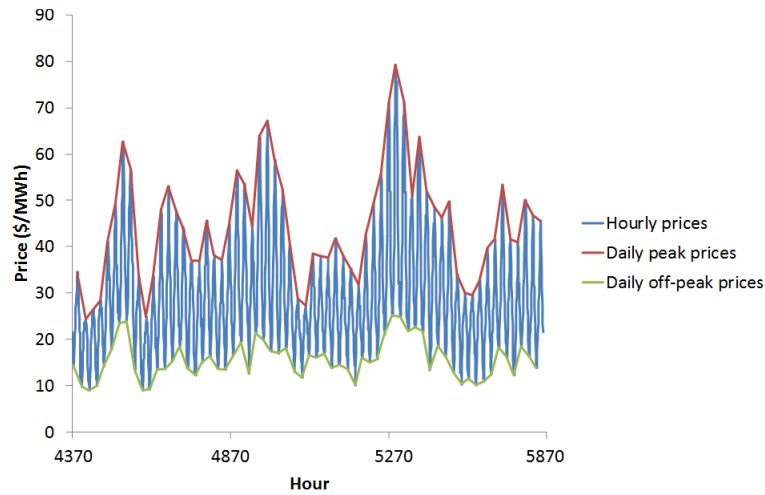


Figure 2: Off-peak and peak prices in July-August 2016

Many studies have been conducted to assess the profitability of energy storage facilities doing energy arbitrage in different electricity markets. Different storage technologies are considered in these studies, but storage facilities are generally characterized by three main features regardless of the technology: the power capacity (in MW), the energy capacity (typically given in MWh), and an efficiency illustrating the losses incurred during both charging and discharging operations. The energy capacity of a storage device may also be seen as a number of hours of full power output. These studies may be divided in two categories depending on the power capacity of the energy storage facility.

The first category assumes that the energy storage facilities are “price-taker”, which means that their operation does not affect the market price. This is generally the case when the storage power capacity is low compared to the total demand or supply in the market, so that the demand or supply is not affected by the storage operation. [19] and [26] estimate respectively the value of small energy storage devices in two jurisdictions of the US: PJM and New-York state. [19] examines the effects of different parameters on the value of storage (efficiency and energy capacity) and establishes that this value may range from \$60/kW-year to \$120/kW-year depending on these parameters. The impact of the prices of fuel (natural gas and coal) on the value of arbitrage is also highlighted: hourly on-peak prices are often set by natural gas or coal generation, therefore, increases in the price of commodities lead to increases in the value of storage. [26] underlines the economic opportunity in New York state for energy storage installations, and compares the economics of different technologies (sodium sulfur batteries and flywheel). The resulting revenues are comparable with those in PJM. In [18], Sioshansi et al. explore the value drivers behind energy arbitrage, for two different technologies (pumped-hydro storage, and compressed air storage). Especially, it is shown that, due to its hybrid nature, compressed air storage is more sensitive to gas price fluctuations. Finally, [7] propose a model to optimize the schedule of energy storage devices participating in both energy and reserve markets in different regions in US. The combination of energy arbitrage and operating reserve increase the value of energy storage in some markets.

This first set of studies provide a picture of the arbitrage value of energy storage in different markets. However, these studies assume that the storage operation does not affect the market price, which is a valid assumption provided that the storage capacity is sufficiently small. Our study focuses on larger-scale facilities, which may affect the demand and supply on the market when they operate. Charging during low-demand periods and discharging during peak periods will reduce the price gap, and therefore the arbitrage value. In this case, the energy storage facility is said to be “price-maker”. To reach our objective of assessing accurately the potential profits of a large-scale energy storage, it is crucial to find a formulation that takes into account the impact of storage operation on the price.

A few studies have addressed the self-scheduling of a large-scale energy storage facility [17, 19, 21, 23]. The general idea is to describe, at each time step, the market price p_t as a function of the quantity q_t submitted

by the energy storage on the market, i.e. $p_t = f(q_t)$, where q_t can either be positive (when discharging) or negative (when charging). t refers to the time slicing of the market, according to the frequency with which the bids are formulated. In the case of the day-ahead market, bids have to be formulated on a hourly basis. For the sake of simplicity, it is convenient to separate q_t in two positive variables: $q_t = P_t^d - P_t^c$ where P_t^d is the charging power and P_t^c is the discharging power at time step t . Based on these variables, the energy level of the storage E_t at each time step t can thus be computed. Hence, the self-scheduling problem may be formulated as follows:

$$\max_{\{P_t\}_{t=1}^T, \{E_t\}_{t=1}^{T+1}} \sum_{t=1}^T \left[P_t^d f(P_t^d) - P_t^c f(-P_t^c) \right] - \sum_{t=1}^T C(P_t^d + P_t^c) \quad (1a)$$

$$\text{subject to} \quad 0 \leq P_t^d \leq P_{max}^d \quad \forall t \in \llbracket 1, T \rrbracket \quad (1b)$$

$$0 \leq P_t^c \leq P_{max}^c \quad \forall t \in \llbracket 1, T \rrbracket \quad (1c)$$

$$E_{t+1} = E_t + \eta P_t^c - \frac{1}{\eta} P_t^d \quad \forall t \in \llbracket 1, T \rrbracket \quad (1d)$$

$$0 \leq E_t \leq E_{max} \quad \forall t \in \llbracket 1, T+1 \rrbracket \quad (1e)$$

$$E_1 = E_{T+1} = E_{cyc}, \quad (1f)$$

where $\llbracket 1, T+1 \rrbracket$ denotes all the integers between 1 and T .

The objective function (1a) describes the profit of the storage operator, which consists of the revenues from selling electricity to the market, minus the costs from buying electricity to the market, minus the operating costs of the facility. These operating costs are described in the second part of the objective function: C corresponds to the variable cost due to operation and degradation during the hours of charging and discharging. These revenues and costs are summed for each hour of the horizon, since the day-ahead market requires hourly bids. Let's note that, for each hour t , the price should be expressed as $f(P_t^d - P_t^c)$. But given the operating cost, it would be suboptimal to charge and discharge at the same time. Hence, the price can be expressed as $f(P_t^d)$ during the discharging hours and $f(P_t^c)$ during the charging hours. Constraints (1b), (1c) and (1e) describe the limits of storage in terms of power capacity and energy capacity. Constraint (1d) computes the state of charge of the storage at each period of time. The losses during both charging and discharging are taken into account by means of an efficiency η . The initial and final state of charge are specified in (1f).

Two main methods have been proposed to model the function f describing the impact of storage operation on the price. In both cases, it is necessary to understand the price formation process. For each hour, suppliers and consumers submit bids (which are composed of a quantity, and a price) to the market. After collecting and sorting all these bids, a supply curve and a demand curve may be constructed. The price is then given by the intersection between the two curves.

The first method is the most general, and uses the residual demand curve, which is defined as the market demand curve minus the quantity supplied by other producers. It provides a direct relation between the quantity submitted by the energy storage and the resulting market price. [23] and [21] approximate the residual demand curve of the Iberian market by a sigmoid function, and solve the corresponding non-linear self-scheduling problem. This formulation is particularly relevant when the demand on the market is elastic.

In practice, it is often the case that electricity demand can be assumed to be inelastic. It is therefore sufficient then to model the effect of storage operation on price through a supply curve $\pi(d)$, where d is the demand. Hence, the effect of the storage unit can be taken into account through $f(q) = \pi(d - q)$. Since the storage is self-scheduling, its operation is indeed equivalent to an increase (when charging) or a decrease (when discharging) in the demand. [19] and [17] exploit respectively the supply curve of the Alberta and PJM electricity markets to formulate the self-scheduling problem of a price-maker energy storage. In [19], a linear supply curve is constructed for each month based on historical data of prices and quantities. In [17], the supply curve is not approximated, actual supply curves from the Alberta market are used for years 2010 to 2014. For each hour, based on the supply curve and the demand, generation price quota curves (GPQC) and demand price quota curves (DPQC) are constructed to model respectively the price impacts of discharging

and charging. This stepwise approximation of the supply curve around the value of the demand allows to develop a mixed-integer linear formulation of the self-scheduling problem.

This paper addresses the economic assessment of energy arbitrage opportunities for a large-scale energy storage operator in the day ahead market of New York state. We will assume that at the moment of submitting his bids, while the operator of such a facility has an accurate idea of the hourly electricity demand, he is unaware of the market clearing price and in particular the exact effect of his bid on this price. This represents realistic conditions of operations given that such operators would usually be unaware of the bids that will be submitted by other participants, or even of their conditions of operations (e.g. cost of resources, periods of maintenance, etc.). We will instead assume that the operator employs historical observations of electricity demand, market price, available hourly wind power production to construct an uncertainty model for the potential supply curves, which consists of a nominal supply curve, a maximal supply curve, and a minimal supply curve. This uncertainty model will be employed by a robust formulation of the self-scheduling problem (1a) that will account for the level of aversion the operator has with respect to the possibility of daily losses. It is worth emphasizing that this is in sharp contrast with the approach presented in [17] and [19], who both assume that the supply curve for every hour of the day is exactly known in advance, or equivalently that the operator is insensitive to estimation errors. Furthermore, our approach will model the supply curve as a piecewise linear function which captures better the increasing marginal impact of supply on prices during high demand periods compared to the piecewise constant model employed in [19].

Overall, the contributions of this paper can be summarized as follows:

- We present for the first time a method that can be used to characterize the market price uncertainty that a participant is confronted to when submitting a self-scheduled bid in a day-ahead market. Our approach will be based on performing least squares and quantile regression on historical observations of total demand, market prices, and wind power contributions.
- We present for the first time a decision model that employs robust optimization to model the risk aversion of a energy storage operator. In particular, the model will control using a budget Γ under which magnitude of perturbation of a nominal daily profit curve is the operator comfortable with the possibility of a financial loss.
- We show that the robust bidding strategy obtained from this model using $\Gamma = 2$ allows to reduce the risk of a financial loss (from 3.01% to 1.09% with respect to the nominal strategy), while maintaining the expected profit at a reasonable level (10.8% inferior to that obtained with the nominal strategy).

The remainder of the paper is organized as follows. In section 2, the modelling of the supply curve, and of its variability is described. In section 3, the robust self-scheduling problem of a price-maker storage, and its reformulation are developed. In section 4, the developed model is applied and the robust strategy of the storage operator, as well as its impact on the market prices are explored.

2 Modelling the supply curve in New-York state day-ahead market

The methodology used in this paper, which is based on the construction of the supply curve, requires a good understanding of the organization of New-York state electricity markets, as well as an advanced study of the data extracted from this market.

2.1 New York electricity markets

In the state of New York, electricity is traded in a number of competitive electricity markets, all of which are administered by the regional transmission organization called New York Independent System Operator (NYISO). NYISO is also responsible for operating the state's bulk electricity grid, and for conducting long-term planning for the state's electric power system. The electricity grid serves about 20 million people and has historically been required to supply peaks of demand as large as 32 GW in 2015 (see [14]). In comparison, the total power capacity from sources within the state currently reaches 39 GW, with half of the capacity originating from dual fuel power plants (facilities capable of using natural gas as well as other fossil fuels).

The other half of the total capacity is mainly due to nuclear (14%), hydro (11%), and gas (10%) power plants. NYISO has also serious ambitions regarding the development of wind and solar power facilities. In particular, it introduced the “reforming the energy vision” in 2013 to pave the way towards a clean energy system and proposes revisions to the plan on a yearly basis. An important consequence of this initiative was to dramatically increase the wind power capacity from 48 MW in 2005 to 1746 MW in 2015. This has given rise to new opportunities for energy storage facilities given that the productivity of such renewable sources is highly variable.

Among all the markets operated by NYISO, this article focuses on the energy day-ahead market, which accounts for over 94% of energy exchanges [13]. In this market, energy suppliers and consumers submit economic bids for each hour of the following day. While the price curves for supply and demand are the key factors determining the market prices, the transmission of electricity also plays a noticeable role. Indeed, bottlenecks can occur on the electricity grid if large volumes need to be transmitted to meet demand in a particular zone. Thus, NYISO employs a nodal pricing scheme, which give rise to local marginal prices (LMP) for each of the 11 zones of New York state. These LMP’s are the result of three contributions: the marginal cost of energy (which is uniform over the state), the cost of losses in transmission lines, and a cost related to congestion in the zone considered. Our study will focus on the main contributor to market prices, namely the marginal cost of energy. The reasons for doing so are two-fold. First, our study aims at evaluating the general economic opportunities for doing arbitrage in New York state. We do not address the issue of determining the optimal location for the storage facility. Hence, the most consistent price to take into account is the marginal cost of energy, which is the same statewide. Secondly, to model the price-maker effect of energy storage, we will use a relation between the load (or supply) and the price via the supply curve. Yet, the only component of the price which is directly related to the load is the marginal cost of energy. As the load increases, energy sources with increasingly high marginal costs of production have to be dispatched to meet the demand. On the other hand, the two other contributions of the price are not directly related to the load, but rather to local transmission constraints.

As portrayed by Figure 3a, our data set consists of a list of historical $\{(p_i, d_i)\}_{i=1}^N$ pairs, ranging from January 1st, 2016 to December 31st, 2016, and describing on an hourly basis the market price and corresponding electricity demand that was observed on the energy day-ahead market supervised by NYISO. Based on this figure, one can observe that the supply curve describing the relationship between the demand and the price is subject to high variability. There are indeed many reasons why bids submitted by suppliers might vary from day to day (or even hour to hour):

- marginal costs incurred by each supplier in the market fluctuate depending on price of commodities such as natural gas, oil, uranium, etc.;
- the production capacity of renewable resources are sensitive to meteorological conditions (wind, rainfall, sunshine);
- each power plant may become unavailable at times because of power outages, maintenance, etc.

Given that many of the causes of variability are unobservable at the time of submitting a bid on this day-ahead market, it is reasonable to conclude that predicting exactly where the intersection between the supply curve and the inelastic demand curve will occur for any given hour of the day is an impossible task. Hence, under such conditions one should employ a representation that accounts for variability in the supply curve when searching for an optimal bidding strategy.

2.2 Constructing a nominal supply curve

We first attempt to identify a nominal representation of the supply curve by employing the least squares method to perform a regression, following the idea proposed in [19]. Specifically, under the assumption that the supply curve has a parametric form $\hat{\pi}(d; \delta)$, one can identify the best fit for $\delta \in \mathbb{R}^m$ by solving the following optimization problem :

$$\min_{\delta} (1/N) \sum_{i=1}^N (p_i - \hat{\pi}(d_i; \delta_i))^2 .$$

Figure 3 presents the nominal curves obtained when $\hat{\pi}(d; \delta)$ is chosen to be an affine function (a.k.a. linear regression) and a piecewise linear function with breakpoints at 25.558 and 28.098. Both of these regressions were performed using the software R version 3.2.0 with the “Segmented” package (available online through [12]). This package allows one to determine jointly the optimal breakpoints and slopes of a piecewise linear function, given that the number of breakpoints is pre-specified. One can also obtain the R^2 statistic of the fitted function which captures the amount of data variability that can be explained by the fitted model. The fact that this statistic increases from 0.5641 to 0.5923 when employing the piecewise linear function seem to confirm that the latter function provides a better fit. We can also expect that the piecewise linear model provides a more accurate description of how the marginal market price can be affected by the magnitude of the demand. This is indeed a key element in the context of the price-maker formulation such as in Problem (1a) given that it defines the impact that the storage facility will have on the market price.

In order to improve the accuracy of our nominal model, we also attempted to model the influence of wind variability on the supply curve. Indeed, given that wind power capacity represents approximately 10% of the average demand and can cause significant changes in the supply curve, and given that accurate predictors of this production can typically be used at the time when bids are submitted to the day-ahead market, it becomes relevant to perform a regression both on d and w of the market price in order to obtain a supply curve. In particular, we obtained from NYISO publicly available data information about the hourly wind power production $\{(w_i)\}_{i=1}^N$ for each period of our data set. Since one can usually assume that wind energy has a negligible marginal cost [25], it is reasonable to believe that it always comes first in the merit-order when market price is selected, which means that all the wind power available is systematically consumed. Hence, the objective to meet the total demand at the lowest cost is equivalent to the objective to meet the “net demand” (the total demand minus the wind power production) at the lowest cost, i.e. after removing the demand that will necessarily be satisfied by wind power. In other words, one might postulate that $\hat{\pi}(d, w) := \hat{\pi}(d - w; \delta)$. For this reason, we perform the same regressions as before but on the modified data set $\{(p_i, n_i)\}_{i=1}^N$ where each $n_i := d_i - w_i$. This gave rise to the linear and piecewise linear regressions presented in Figure 4. With the proposed modification that accounts for available wind power, the regressions produced an R^2 statistics of 0.6139 and 0.6485 respectively which seems to support our choice.

The conclusions of this fitting of a nominal supply curve support the use of the following function to model the impact of storage unit charge and discharge:

$$\hat{f}(P_t^d) := \hat{\pi}(d - P_t^d - w; \delta^*) = \hat{\pi}_w(n_t - P_t^d) \quad \hat{f}(-P_t^c) := \hat{\pi}(d + P_t^c - w; \delta^*) = \hat{\pi}_w(n_t + P_t^c),$$

where

$$\hat{\pi}_w(y) = \begin{cases} 2.086y - 17.354 & \text{if } y \leq 25.558 \\ 4.249y - 72.636 & \text{if } 25.558 < y \leq 28.098 \\ 6.705y - 141.45 & \text{if } 28.098 < y \end{cases}.$$

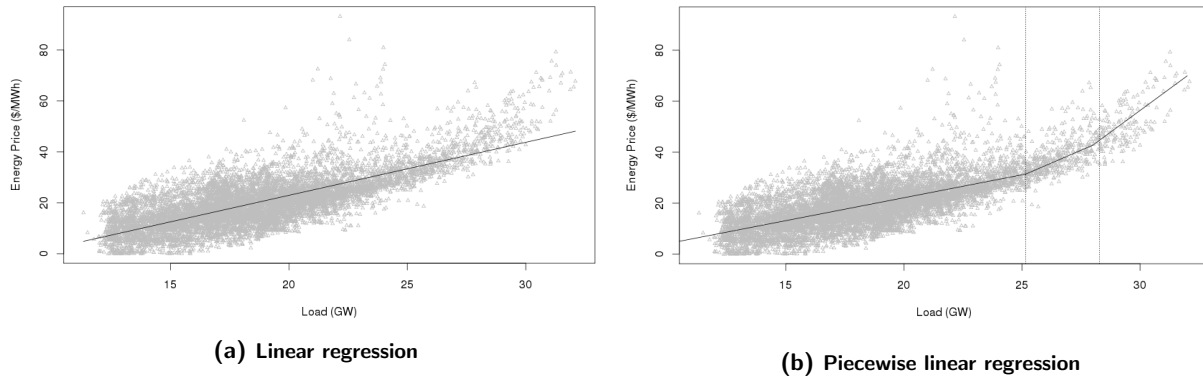


Figure 3: Best fitted models for the nominal supply curve based on historical data $\{(p_i, d_i)\}_{i=1}^N$. (a) presents the calibrated affine function $\hat{\pi}(d; \delta^*)$ with an R^2 of 0.5641. (b) presents the calibrated piecewise linear function $\hat{\pi}(d; \delta^*)$ with an R^2 of 0.5923

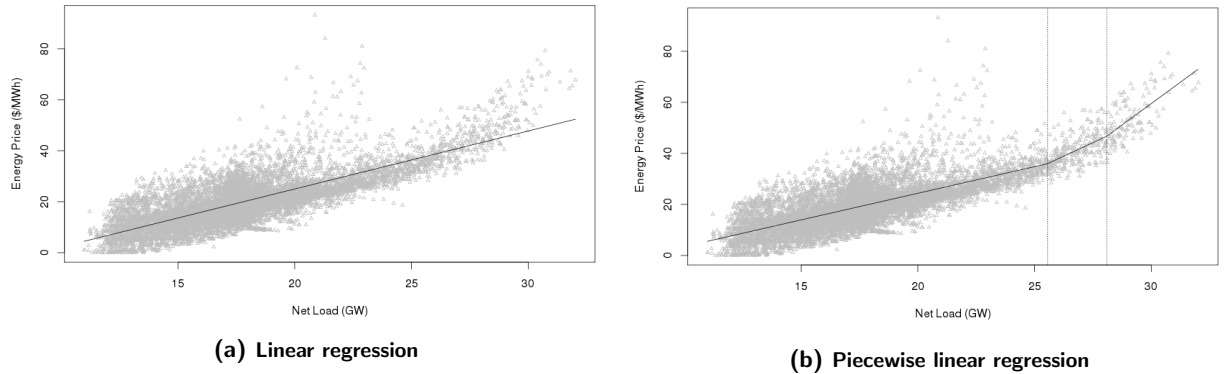


Figure 4: Best fitted models for the nominal supply curve based on historical data $\{(p_i, d_i, w_i)\}_{i=1}^N$. **(a)** presents the calibrated affine function $\hat{\pi}(d - w; \delta^*)$ with an R^2 of 0.6139. **(b)** presents the calibrated piecewise linear function $\hat{\pi}(d - w; \delta^*)$ with an R^2 of 0.6485. Note that the x -axis presents the net demand $n = d - w$ which was used in the regression instead of d

2.3 Constructing upper and lower bounds for the supply curve

We now turn ourselves toward the characterization of the variability of the supply curve and the effect of this variability on the cash flows that will be produced when scheduling a charge or discharge of the storage unit. In fact, we will exploit the historical data set $\{(p_i, d_i, w_i)\}_{i=1}^N$ to calibrate two bounding functions $\pi_w^+(n)$ and $\pi_w^-(n)$ so that they return for a given net load n , a confidence interval $[\pi_w^-(n), \pi_w^+(n)]$ for the realized market price. This can be done using quantile regression (as introduced in [10]).

Quantile regression is similar in spirit to the method of least squares. One first needs to identify a parametric form for $\pi_w^+(n)$ and $\pi_w^-(n)$, which we will refer to as $\pi_w(n; \delta_-)$ and $\pi_w(n; \delta_+)$. Given a confidence level η , which we choose to be $\eta := 10\%$, we will fit the δ_+ and δ_- parameters to the data set $\{(p_i, n_i)\}_{i=1}^N$ but this time using a regression function that aims at capturing the $\eta/2$ -th and $1 - \eta/2$ -th percentile respectively. Specifically, the optimization models take the form:

$$\delta_-^* = \arg \min_{\delta_-} (1/N) \sum_{i=1}^N \max \left((1 - \frac{\eta}{2})(\pi_w^-(n_i; \delta_-) - p_i), \frac{\eta}{2}(p_i - \pi_w^-(n_i; \delta_-)) \right)$$

$$\delta_+^* = \arg \min_{\delta_+} (1/N) \sum_{i=1}^N \max \left(\frac{\eta}{2}(\pi_w^+(n_i; \delta_+) - p_i), (1 - \frac{\eta}{2})(p_i - \pi_w^+(n_i; \delta_+)) \right).$$

Intuitively, the first optimization model will penalize more severely over-evaluations than under-evaluations of the price in order to return an under-estimator while the second model attempts the opposite. The connection to the notion of quantile estimation emerges when one assumes that, conditionally on observing n_i , the $\eta/2$ -th percentile of the market price can be computed using a member of the parametric family $\pi_w^-(n_i; \delta_-)$. In this case, as N converges to infinity, then δ_- can be shown to converge in probability to the true value, and similarly in the case of δ_+ . In contrast, the method of least squares offers a similar type of convergence but towards the conditional expected value of the market price. We refer interested readers to the Koenker's book [9] for a thorough presentation of this regression scheme.

In our implementation, we model the lower and upper bounds of the supply curve with piecewise linear functions. For the sake of consistency with the nominal supply curve determined in Subsection 2.2, the same breakpoints are used. This modeling decision also has the advantage of reducing the number of binary variables involved in the mixed integer quadratic program that is proposed in Section 3.3 to identify robust self-scheduled bids. It was however necessary to include an additional breakpoint at 12.817 GW in order to prevent the lower bounding function to return some negative prices. Negative prices are known not to occur in the New-York market because of the way the price selection mechanism is designed. This is however not the case in all electricity market given that negative prices do emerge temporarily in some markets because of producers that are unwilling or unable to interrupt suddenly their productions.

For completeness, we present below the resulting linear program that needs to be solved in order to obtain the calibrated parameters for the lower bounding function:

$$\begin{aligned}
& \min_{\delta, t, y} && (1/N) \sum_{i=1}^N t_i \\
& \text{subject to} && t_i \geq (1 - \frac{\eta}{2})(y_i - p_i), \forall i = 1, \dots, N \\
& && t_i \geq \frac{\eta}{2}(p_i - y_i), \forall i = 1, \dots, N \\
& && y_i = \sum_{j=1}^3 \delta_j (n_i - \gamma_j) \mathbf{1}\{n_i \geq \gamma_j\} \\
& && \delta \geq 0,
\end{aligned}$$

where $\gamma_1 = 12.817$, $\gamma_2 = 25.558$, and $\gamma_3 = 28.098$ are the three breakpoints at which the piecewise linear function changes slope, while each δ_j captures by how much the slope increases from one piece to the other. Finally, $\mathbf{1}\{y \geq 0\}$ is the indicator function that returns one if the condition is satisfied and zero otherwise.

Based on the result of our calibration, we will employ in the remainder of the paper the following calibrated curves to capture how low and how large the market price might be when submitting a bid of P_t^d or P_t^c :

$$\begin{aligned}
f^-(P_t^d) &:= \pi_w^-(d_t - w_t - P_t^d) = \pi_w^-(n_t - P_t^d) & f^+(P_t^d) &:= \pi_w^+(d_t - w_t - P_t^d) = \pi_w^+(n_t - P_t^d) \\
f^-(P_t^c) &:= \pi_w^-(d_t - w_t + P_t^c) = \pi_w^-(n_t + P_t^c) & f^+(P_t^c) &:= \pi_w^+(d_t - w_t + P_t^c) = \pi_w^+(n_t + P_t^c),
\end{aligned}$$

where

$$\pi_w^-(y) = \begin{cases} 0 & \text{if } y \leq \gamma_1 \\ 2.269y - 29.081 & \text{if } \gamma_1 < y \leq \gamma_2 \\ 3.508y - 60.767 & \text{if } \gamma_2 < y \leq \gamma_3 \\ 6.248y - 137.764 & \text{if } \gamma_3 < y \end{cases} \quad \text{and} \quad \pi_w^+(y) = \begin{cases} 2.272y - 9.023 & \text{if } y \leq \gamma_2 \\ 3.320y - 35.820 & \text{if } \gamma_2 < y \leq \gamma_3 \\ 7.884y - 164.069 & \text{if } \gamma_3 < y \end{cases}.$$

The three curves (nominal, lower bound, and upper bound) are presented in Figure 5. We observe that the lower and upper bounds allow to encapsulate most of the observations. However, for a certain range of demand (approximately from 20 GW to 24 GW), abnormally high values of the price occur. Explaining the origin of these outliers is beyond the scope of this work, however their existence should be kept in mind.

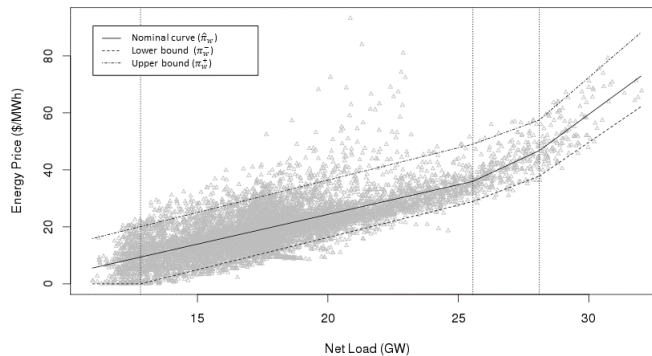


Figure 5: Lower, nominal, and upper supply curves excluding wind power

3 Robust formulation of the self-scheduling problem

In this section, we propose a robust optimization model in order to model a storage facility operator that is risk averse regarding the uncertainty he has about actual market price when submitting self-scheduled bids to a day-ahead market. We will initially review some background on the general methodology before

focusing on the choices we made in its application. Next, Section 3.2 will discuss how the approach presented in [6] can be used to robustify problem (1a)–(1f) in a way that immunizes the operator against potential daily losses. We then present in Section 3.3 how this robust problem can be reformulated as a mixed-integer convex quadratic problem.

3.1 Background on robust optimization

Robust optimization is a technique for optimization under uncertainty, which has received an increasing amount of interest in the past ten years. Contrary to other approaches that handle uncertainty, such as stochastic programming, it removes the need of identifying a probabilistic model that quantifies the likelihood of every possible future outcomes, replacing it with the characterization of a so-called uncertainty set. In principle, the robust optimization paradigm seeks solutions that remain feasible under any potential outcomes that fall within the prescribed uncertainty set. Its first applications to mathematical programming dates from [22] where the authors proposed that each uncertain parameter be circumscribed to its respective interval. This approach was quickly considered overly conservative as it allowed worst-case scenarios where each parameter takes on an extreme value simultaneously. This issue was addressed in [2], where Ben-tal and Nemirovski propose the use of ellipsoidal uncertainty sets, which does not allow for such events to be considered. Even more recently, the authors of [5] introduced a polyhedral set known as the budgeted uncertainty which allows one to control the level of conservatism through the use of a scaling parameter Γ which defines how many of the uncertain parameters are allowed to reach an extreme value. These important works contributed significantly to the popularization of the method. Overall, one might consider the following advantages that a robust optimization framework typically has over stochastic programming:

- For many classes of optimization problems, the robust optimization formulation is computationally tractable (see [3]) while a stochastic programming approach might be confronted to the challenge of performing high-dimensional integration.
- The non-probabilistic approach used in robust optimization allows the decision-maker to immunize against uncertainty without having to define a distribution for the uncertain parameters.

The latter advantage is especially practical in the case of data-driven problems, where there is no particular reason to represent the random vector with a distribution of a specific form, such as the normal distribution. In the case of stochastic programming, it is necessary to identify and calibrate a joint distribution for the vector of uncertain parameters. This distribution defines both the marginal likelihood of each parameter taken separately and the specifics of how each of them is correlated to others. When the random vector is large and the observations rather limited, then it can easily be the case that there are many distribution models that could explain the observations equally well thus making this choice rather arbitrary. This difficulty has given rise to what might be thought of as the “Optimizer’s curse” (see [20]) given that the solution that is identified by a stochastic program can easily over-exploit the selected distribution model resulting in an optimistic view of future performance which can lead to great post-decision disappointment.

For these reasons, the robust optimization approach has been applied in many different domains including power systems operations. In particular, from the market operator perspective, more and more sources of uncertainty have to be taken into account in the unit commitment problem. Given the increasing penetration of variable energy sources (wind, solar), and the recent development of price-responsive demand, solving this problem has become more challenging. In [4], Bertsimas et al. propose a two-stage adaptive robust optimization model for the security constrained unit commitment problem in the presence of nodal net injection uncertainty. In [27], a polytopic uncertainty set is constructed to capture wind uncertainty, and is then integrated in the robust formulation of the unit commitment problem. Finally, [8] proposes a robust optimization approach to provide a robust unit commitment schedule for the thermal generators in the day-ahead market that minimizes the total cost under the worst wind power output scenario. We also refer the reader to [24], where a robust model is developed to optimize the bidding strategy of a wind farm owner. The authors consider uncertainty about both future electricity price and wind production. To the best of our knowledge, there is no prior work on applying a robust optimization framework to the self-scheduling problem of a storage facility operator.

3.2 The robust optimization model

When confronted to historical observation of market prices such as those studied in Section 2, it is easy to see how a storage facility operator might express some concerns regarding the implementation of a self-scheduling bid strategy that does not account for price uncertainty. In particular, since supply curves are usually monotonic, it is often the case that such a “nominal strategy” would recommend to charge the battery during the lowest demand hour, and sell this electricity back when the demand is at its highest level. As seen in Figure 6, when price uncertainty is large, doing so exposes the operator to the risk that the realized market price for the period with low demand (i.e. a scheduled charge) be higher than during the period where a discharged was planned, hence leading to a net financial loss. This motivates the use of a robust optimization approach that will allow the storage facility operator to control his exposure to net financial losses.

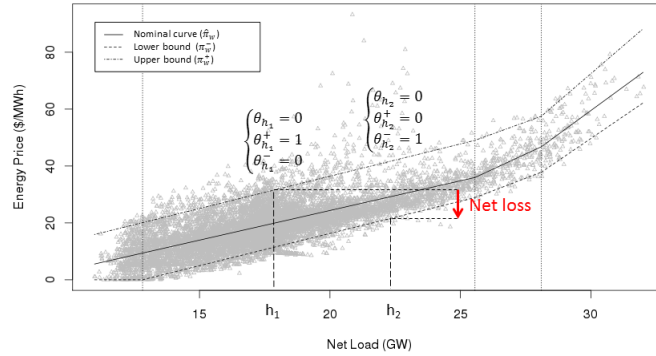


Figure 6: Situation where the nominal strategy exposes the operator to a net financial loss. The values h_1 and h_2 represent two periods with respectively low and high demand thus motivating a charge at period 1 followed by a discharge at period 2 when considering the nominal supply curve $\hat{\pi}_w$. When implementing this strategy, the operator is exposed to the risk that the realized market price coincide with the upper bound π_w^+ for period 1 and lower bound π_w^- for period 2 leading to a net loss

In what follows, we derive a robust optimization model based on the paradigm presented in [6] which can directly exploit the description of uncertainty that was presented in Section 2.3, defining a nominal, lower and upper bound for $f(\cdot)$. Specifically, Problem (1) is modified by adding a constraint that rejects a self-scheduling strategy if it has the potential of leading to a net loss when the nominal supply curve suffers a certain level Γ of perturbation. This gives rise to the following robust self-scheduling problem:

$$\max_{\{P_t^c, P_t^d, E_t\}_{t=1}^T} \sum_{t=1}^T [P_t^d \hat{f}(P_t^d) - P_t^c \hat{f}(-P_t^c)] - \sum_{t=1}^T C(P_t^d + P_t^c) \quad (2a)$$

$$\text{subject to} \quad \sum_{t=1}^T [P_t^d f_t(P_t^d) - P_t^c f_t(-P_t^c)] - \sum_{t=1}^T C(P_t^d + P_t^c) \geq 0, \forall (f_1, f_2, \dots, f_T) \in \mathcal{F}(\Gamma) \quad (2b)$$

(1b) – (1f),

where $\mathcal{F}(\Gamma)$ captures all supply curves that can be obtained by a Γ perturbation of \hat{f} , mathematically speaking

$$\mathcal{F}(\Gamma) := \left\{ (f_1, \dots, f_T) \left| \begin{array}{l} \exists \theta \in \mathbb{R}_+^T, \theta^- \in \mathbb{R}_+^T, \theta^+ \in \mathbb{R}_+^T, \\ \theta_t + \theta_t^+ + \theta_t^- = 1, \forall t \in \llbracket 1, T \rrbracket \\ \sum_t \theta_t^- + \theta_t^+ \leq \Gamma \\ f_t(\cdot) = \theta_t \hat{f}(\cdot) + \theta_t^- f^-(\cdot) + \theta_t^+ f^+(\cdot), \forall t \in \llbracket 1, T \rrbracket \end{array} \right. \right\}.$$

In a language similar to the one used by the authors of [5], one can interpret Γ as the maximum number of time periods during which the supply curve is allowed to reach either of the two supply curve bounds π_w^- or π_w^+ that were identified using the historical data. One might also recognize that in the construction of $\mathcal{F}(\Gamma)$, we model for each time period t a triplet $(\theta_t, \theta_t^-, \theta_t^+)$ that will let the market price at period t take on any convex combination of $f^-(\cdot)$, $\hat{f}(\cdot)$, and $f^+(\cdot)$. Furthermore, when $\Gamma = 0$, Problem (2) reduces to the nominal Problem (1) since in this case all $(\theta_t, \theta_t^+, \theta_t^-) = (1, 0, 0)$ leading to $\mathcal{F}(\Gamma) = \{\hat{f}(\cdot)\}$. Alternatively,

when $\Gamma = T$, constraint (2b) reduces to

$$\sum_{t:P_t^c=0} P_t^d f_t^-(P_t^d) - \sum_{t:P_t^d=0} P_t^c f_t^+(-P_t^c) - \sum_{t=1}^T C(P_t^d + P_t^c) \geq 0,$$

which effectively assumes that the market price always end up being the most unfavourable with respect to the self-scheduling strategy.

It is also possible to interpret the robust constraint (2b) as an approximation of the following chance constraint:

$$\mathbb{P}\left(\sum_{t=1}^T \left[P_t^d \tilde{f}_t(P_t^d) - P_t^c \tilde{f}_t(-P_t^c) \right] - \sum_{t=1}^T C(P_t^d + P_t^c) \geq 0\right) \geq 1 - \varepsilon,$$

where $\tilde{f}(\cdot)$ is the random mapping that is assumed to have produced the historical price observations, and $\varepsilon \in [0, 1]$ characterizes the amount of probability with which we are comfortable that the constraint might not be respected. Based on the definition of $\mathcal{F}(\Gamma)$, it is possible to evaluate the probability that the historical observation of market price be a member of our uncertainty set :

$$\mathbb{P}(\tilde{f} \in \mathcal{F}(\Gamma)) \approx (1/N) \sum_{i=1}^N \mathbf{1}\{\exists f \in \mathcal{F}(\Gamma), p_i = f(0)\},$$

where we count what is the proportion of historical observations for which the observed price could be a result of evaluating one of the functions in $\mathcal{F}(\Gamma)$ at zero (given that the contribution of the battery facility was null historically). Figure 7 presents the estimated level of protection depending on the size of Γ . This approach can provide the decision-maker an idea of the value of Γ he has to use depending on the level of protection needed. However, it leads to an overly conservative choice of Γ . We follow a more empirical approach, which consists in experimenting different values of Γ . This will be described in Section 4.




Figure 7: Chance-constraint oriented calibration of Γ

3.3 Mixed-integer convex quadratic reformulation

The robust model introduced in Section 3.1 cannot be solved using the shelf solvers given that it is formulated as a semi-infinite problem. To address this issue, one needs to identify a tractable reformulation for constraint (2b). To do so, we first rewrite the constraint using the following equivalent form:

$$\psi(P^d, P^c) \geq \sum_{t=1}^T C(P_t^d + P_t^c) \quad (3)$$

where

$$\psi(P^d, P^c) := \min_{(f_1, \dots, f_T) \in \mathcal{F}(\Gamma)} \sum_{t=1}^T P_t^d f_t(P_t^d) - P_t^c f_t(P_t^c).$$

Based on the definition of $\mathcal{F}(\Gamma)$, one can readily obtain a linear programming representation of $\psi(P^d, P^c)$, namely:

$$\psi(P^d, P^c) = \min_{\theta, \theta^+, \theta^-} \sum_{t=1}^T \theta_t (P_t^d \hat{f}(P_t^d) - P_t^c \hat{f}(-P_t^c)) + \theta_t^- (P_t^d f^-(P_t^d) - P_t^c f^-(P_t^c)) + \theta_t^+ (P_t^d f^+(P_t^d) - P_t^c f^+(-P_t^c)) \quad (4a)$$

$$\text{subject to} \quad \theta_t + \theta_t^+ + \theta_t^- = 1, \forall t \in \llbracket 1, T \rrbracket \quad (4b)$$

$$\sum_{t=1}^T \theta_t^- + \theta_t^+ \leq \Gamma \quad (4c)$$

$$\theta \geq 0, \theta^+ \geq 0, \theta^- \geq 0 \quad (4d)$$

Moreover, by linear program duality, one can show that the dual problem of problem (4) takes the form:

$$\psi(P^d, P^c) = \max_{s, \lambda} \sum_{t=1}^T \lambda_t - \Gamma s \quad (5a)$$

$$\text{subject to} \quad \lambda_t \leq P_t^d \hat{f}(P_t^d) - P_t^c \hat{f}(-P_t^c), \forall t = \llbracket 1, T \rrbracket \quad (5b)$$

$$\lambda_t - s \leq P_t^d f^-(P_t^d) - P_t^c f^-(P_t^c), \forall t = \llbracket 1, T \rrbracket \quad (5c)$$

$$\lambda_t - s \leq P_t^d f^+(P_t^d) - P_t^c f^+(-P_t^c), \forall t = \llbracket 1, T \rrbracket \quad (5d)$$

where $\lambda \in \mathbb{R}^T$ and s are the dual variables respectively associated to constraints (4b) and (4c). Indeed, the fact that $\psi(P^d, P^c)$ is equal to the optimum of this dual problem is guaranteed by strong duality which applies since Problem (4) is both feasible and bounded for all P^d and P^c and for all $\Gamma \geq 0$.

Given that we have obtained a representation of $\psi(P^d, P^c)$ as the optimal value of a maximization problem, one can establish that constraint (3) is equivalent to verifying whether there exists an assignment (λ, s) that satisfy constraints (5b)–(5d) and that also satisfies

$$\sum_{t=1}^T \lambda_t - \Gamma s \geq \sum_{t=1}^T C(P_t^d + P_t^c). \quad (6)$$

We can conclude that our robust optimization (2) is equivalent to:

$$\begin{aligned} & \max_{P^d, P^c, E, \lambda, s} \sum_{t=1}^T \left[P_t^d \hat{f}(P_t^d) - P_t^c \hat{f}(-P_t^c) \right] - \sum_{t=1}^T C(P_t^d + P_t^c) \quad (7a) \\ & \text{subject to} \quad (6) \\ & \quad (1b) - (1f) \\ & \quad (5b) - (5d). \end{aligned}$$

This new formulation allows to overcome the semi-infinite nature of the original robust model. However, the nature of the resulting problem depends on the nature of the three functions $\hat{f}(\cdot)$, $f^-(\cdot)$, and $f^+(\cdot)$ modelling the effect of the self-scheduled bid on market price under the nominal, upper and lower supply curve. In the case that the supply curves are piecewise linear, as estimated in Section 2, one can show that the resulting problem can be reformulated as a mixed integer convex quadratic problem provided that the slope of each piece of these supply curves are non-negative. For completeness, we present below the mixed integer convex quadratic program that is equivalent to Problem (2). This problem can be solved using a solver such as CPLEX. In particular, we used version 12.7 to produce the results presented in Section 4.

$$\max_{P^d, P^c, E, \lambda, s, \nu, \psi} \sum_{t=1}^T \psi_t - \sum_{t=1}^T C(P_t^d + P_t^c) \quad (8a)$$

$$\text{subject to } \psi_t \leq \sum_{j=1}^4 -\hat{a}_j (y_{tj}^d)^2 + y_{tj}^c{}^2 + (n_t \hat{a}_j - \hat{b}_j) (y_{tj}^d + y_{tj}^c) + n_t \hat{b}_j (\nu_j^d + \nu_j^c), \forall t \in \llbracket 1, T \rrbracket \quad (8b)$$

$$\sum_{t=1}^T \lambda_t - \Gamma s \geq \sum_{t=1}^T C(P_t^d + P_t^c) \quad (8c)$$

$$\lambda_t \leq \psi_t, \forall t \in \llbracket 1, T \rrbracket \quad (8d)$$

$$\lambda_t - s \leq \sum_{j=1}^4 -a_j^- (y_{tj}^d)^2 + y_{tj}^c{}^2 + (n_t a_j^- - b_j^-) (y_{tj}^d + y_{tj}^c) + n_t b_j^- (\nu_j^d + \nu_j^c), \forall t \in \llbracket 1, T \rrbracket \quad (8e)$$

$$\lambda_t - s \leq \sum_{j=1}^4 -a_j^+ (y_{tj}^d)^2 + y_{tj}^c{}^2 + (n_t a_j^+ - b_j^+) (y_{tj}^d + y_{tj}^c) + n_t b_j^+ (\nu_j^d + \nu_j^c), \forall t \in \llbracket 1, T \rrbracket \quad (8f)$$

$$d_t - w_t - P_t^d = \sum_{j=1}^4 y_{tj}^d, \forall t \in \llbracket 1, T \rrbracket \quad (8g)$$

$$d_t - w_t + P_t^c = \sum_{j=1}^4 y_{tj}^c, \forall t \in \llbracket 1, T \rrbracket \quad (8h)$$

$$\gamma_{j-1} \nu_{tj}^d \leq y_{tj}^d \leq \gamma_j \nu_{tj}^d, \forall j = 1, \dots, 4, \forall t \in \llbracket 1, T \rrbracket \quad (8i)$$

$$\gamma_{j-1} \nu_{tj}^c \leq y_{tj}^c \leq \gamma_j \nu_{tj}^c, \forall j = 1, \dots, 4, \forall t \in \llbracket 1, T \rrbracket \quad (8j)$$

$$\sum_{j=1}^4 \nu_{tj}^d = 1, \forall t \in \llbracket 1, T \rrbracket, \sum_{j=1}^4 \nu_{tj}^c = 1, \forall t \in \llbracket 1, T \rrbracket \quad (8k)$$

$$\nu^d \in \{0, 1\}^{T \times 4}, \nu^c \in \{0, 1\}^{T \times 4} \quad (8l)$$

$$(1b) - (1f),$$

where $\psi \in \mathbb{R}^T$, while the pairs of parameters (\hat{a}_j, \hat{b}_j) , (a_j^-, b_j^-) , and (a_j^+, b_j^+) respectively refer to the slopes and offsets of the j -th linear pieces of $\hat{\pi}_w(\cdot)$, $\pi_w^-(\cdot)$, and $\pi_w^+(\cdot)$. Finally, γ_1, γ_2 , and γ_3 are the locations of the three breakpoints, while $\gamma_0 = 0$ and $\gamma_4 = \max_t d_t + P_{max}^c$.

4 Results and discussion

In this section, we use the model developed in Section 3 for different purposes:

- analyze the bidding strategies obtained according to the chosen level of risk-aversion (namely the value of Γ);
- evaluate the performance of robust strategies, both in terms of expected profit, and protection against risk of financial loss. Especially, this evaluation will provide the decision-maker a mean to determine the optimal value of Γ .
- assess the potential operating profit for a price-maker storage facility participating in the New-York electricity market;
- investigate the impact of the implementation of energy storage on the market prices.

We study the case of a storage facility with 100 MW discharging and charging power and 300 MWh energy capacity (i.e. 3 hours of full power output), which is a bit more than the size of the largest battery project brought online so far [16]. For both charging and discharging operations, we consider a 90% efficiency, and a \$1/MWh variable cost due to operation and degradation [24]. We choose a scheduling horizon of one day, namely $T = 24$, which is the minimum horizon that can be considered, since bids have to be transmitted

for every hour of the subsequent day on the day-ahead market. Though longer-term horizons may allow to consider both intra- and inter-day arbitrage, it appeared that, in the case of NYISO day-ahead market, most of the arbitrage value was coming from intra-day peak price differential. This is especially the case when the energy capacity considered for the storage facility is limited: in our case, only 3 hours are necessary for a full charge. This can be done daily during off-peak hours to take advantage of the daily peak hours.

Besides, this shorter horizon allows to assume that the perfect forecast of future net demand during each one-day period is available. This assumption is important here because the idea is to focus on the uncertainty related to the supply curve. Considering simultaneously the uncertainty in the net demand and in the supply curve would increase dramatically the complexity of the model.

4.1 Robust self-scheduling strategy

In this subsection, to demonstrate the importance of the robust formulation, we analyze the bidding strategy according to the level of conservatism chosen by the storage operator.

To this end, it is necessary to choose a certain day of the year. The bidding strategy indeed depends on the expected price for each hour of the day, and these price forecasts are computed based on the hourly load, according to the supply curve. Thus, the bidding strategy varies with the load profile, which is different every day.

To provide a better comparison, two types of days can be distinguished according to their load profile. They lead to two different types of strategies, as we will observe further. We will see that for both types of day, the nominal strategy consists in purchasing electricity to charge the storage during off-peak hours, before selling electricity and discharging the storage during peak hours. However, this nominal strategy has to evolve when the storage operator seeks to be immuned against the risk of loss. This is where the strategies for the two types of days start to be different:

- The first type of day is the most common, an example of this kind is presented in 8a: it corresponds to the case when the price during off-peak hours in the most adverse situation (namely, the upper curve, since the storage wants to purchase electricity during off-peak hours) is greater to the price during peak hours in the most adverse situation (namely, the lower curve, since the storage wants to sell electricity during peak hours). For this type of day, if the most adverse situation occur (namely $(\theta_t, \theta_t^+, \theta_t^-) = (0, 1, 0)$ during the purchasing hours and $(\theta_t, \theta_t^+, \theta_t^-) = (0, 0, 1)$ during the selling hours), the storage may incur losses if the nominal strategy described previously is implemented.
- The second type of day offers larger arbitrage opportunities, with a greater peak price differential. An example is presented in 8b. In this case, the off-peak price in the most adverse situation is lower than the peak price in the most adverse situation. Thus, the storage operator is guaranteed to make profit when implementing the nominal strategy, even if the most adverse situation occurs.

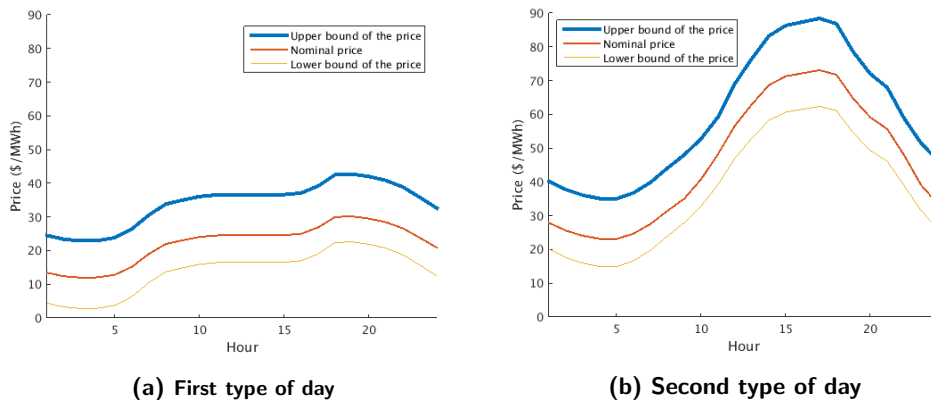


Figure 8: Price profile for two different types of day

We may now analyze the bidding strategies observed according to the level of conservatism (specified by the value of Γ). We start with the first type of day.

- When $\Gamma = 0$, the robust model is equivalent to the deterministic model. In this case, the optimal strategy (referred to as the nominal strategy), presented in Figure 9a, consists in charging during lowest price hours and discharge during highest price hours, and all these operations are conducted at maximum power capacity. The charging and discharging operations are concentrated during the smallest number of hours, to take profit of the most beneficial hours.
- When Γ increases, the strategy is unchanged until $\Gamma = 3$. For small budgets of uncertainty, the nominal strategy still ensures profits, even in the most adverse situations. Thus, our model allows to maintain the nominal strategy.
- For $\Gamma = 4$, the operator has to modify his strategy, to avoid being exposed to losses: the operating profit associated to the nominal strategy would be negative in the most adverse situation. The overall operational trend of the storage facility is still to charge when prices are low and discharge when prices are high. However, the robust strategy suggests spreading the purchases and the sales over a greater number of hours, even if these hours are less favorable in terms of nominal profit (see Figure 9b). This strategy avoids being exposed to maximal deviations in the supply curve during each of the operating hours. The value of Γ can indeed be interpreted in terms of a number of hours when the supply curve may deviate from its nominal position. When $\Gamma = 5$ (see Figure 9c), this trend is amplified (the purchases and sales are respectively distributed over 7 hours and 4 hours, when they were distributed over 5 hours and 3 hours with $\Gamma = 4$).
- For greater values of Γ , the uncertainty is such that no strategy can guarantee positive operating profit. The optimal strategy thus consists in not operating (see Figure 9d).

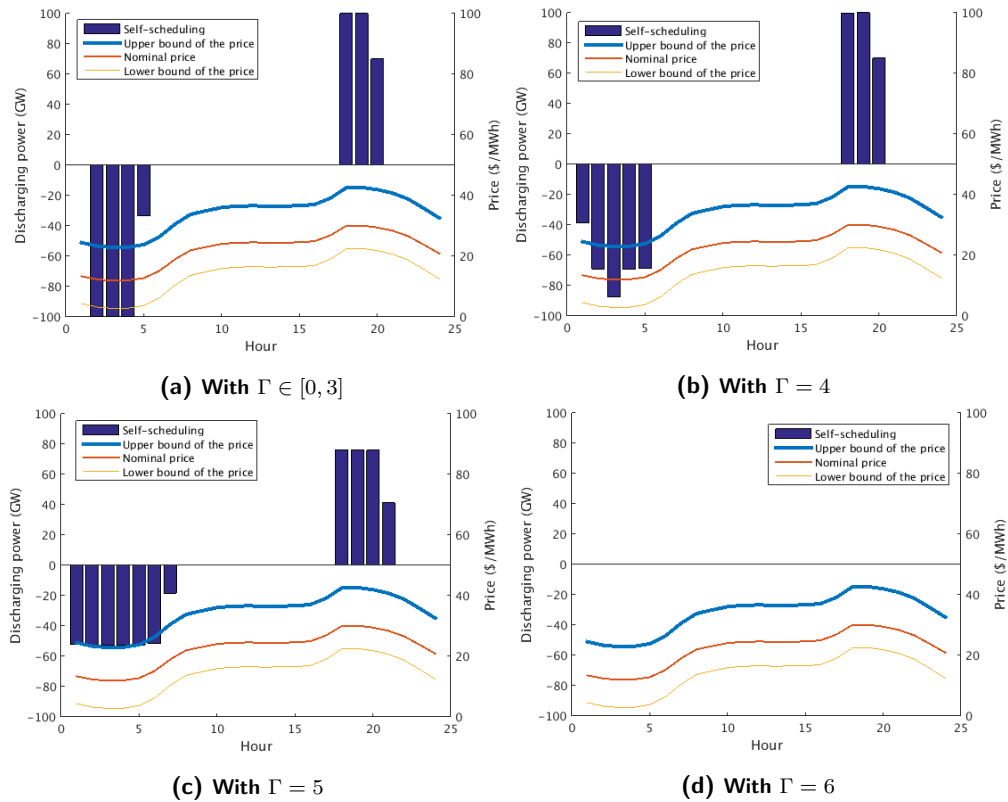


Figure 9: Self-scheduling strategy for the first type of day with different level of protection (Γ)

With the second type of load profile, the price differential between off-peak and peak hours is such that the nominal strategy guarantees a positive operating profit, even in the most adverse situation. Thus, the bidding strategy is the same regardless the value of Γ . This strategy is presented in Figure 10.

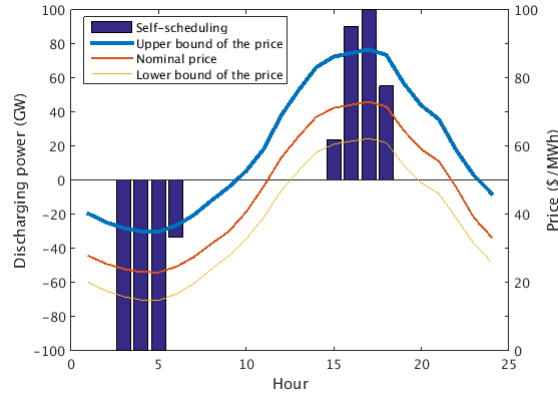


Figure 10: Self-scheduling strategy for the second type of day

From this analysis, it emerges that, from a certain level of risk-aversion, the robust model provides different bidding strategies, which allow to reduce the risk. But this reduction of the risk is done at the expense of the nominal profit: these conservative strategies are indeed suboptimal in the nominal situation where the supply curve does not deviate from its expected position (namely $(\theta, \theta^+, \theta^-) = (1, 0, 0)$). The reduction in the nominal profit, due to the risk aversion of the operator is quantified in Figure 11. To this end, we look at the profit that would have been obtained in the nominal situation with the strategies presented in Figures 9a-9d. For values of Γ ranging from 0 to 3, we observed previously that the bidding strategy was unchanged. Hence, the nominal profit is also unchanged. For $\Gamma = 4$ and $\Gamma = 5$, the wider distribution of the purchases and the sales causes a slight reduction in the nominal profit. From $\Gamma = 5.25$, the nominal profit is zero, since the optimal strategy for this level of uncertainty is to not operate.

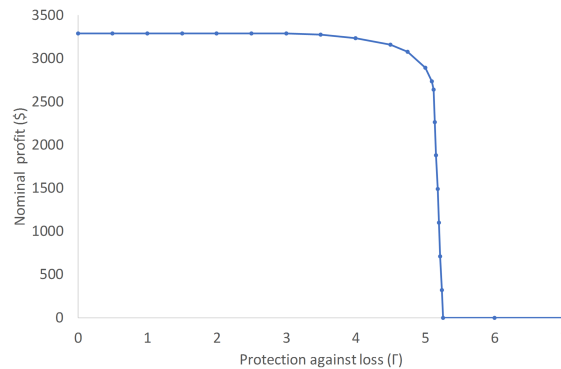


Figure 11: Nominal profit for different level of risk-aversion

Figure 11 provides, for an arbitrary day of the first type, an idea of the effect of robustness on the theoretical profit obtained in the nominal situation. Although this figure corresponds to an arbitrary day, the layout of the curve would be the same for any day of the first type: the first effect of robustness (for low values of Γ) is indeed to suggest to distribute the operations, resulting in a slight decrease of the nominal profit. For greater values of Γ the most conservative strategy is suggested (namely, no operation): it leads to a zero nominal profit. Yet, the tipping value of Γ is differing for depending of the considered day. As for the storage operator, he needs to decide a certain value of Γ that he will use for every day. Hence it is critical to assess the effect of the choice Γ over an entire year. This can be done by confronting the robust strategies obtained to real instances of the uncertain variables, to evaluate the performance of the robust model over a larger sample of realizations for $(\theta, \theta^+, \theta^-)$.

4.2 Experimental results

The data sample of year 2016 is used to confront the robust strategies to different instances of the uncertain variables $(\theta, \theta^+, \theta^-)$. It follows the idea introduced by Bertsimas in [5], and consists in experimenting different values of Γ , to compare the performance of the model according to the budget. It provides thus the decision-maker a tool to compare, when increasing Γ , the corresponding loss in the expected profit (which is referred to as the price of robustness) with the related increase in robustness. This empirical approach is different from the chance-constraint approximation introduced in Section 3.3, and will be privileged.

We understand that the choice of Γ corresponds to a tradeoff between the expected profit and the risk. The risk is here evaluated via two different metrics: the probability to incur losses, and the “worst-case” profit (represented by the 2nd percentile of the profits obtained over the year).

The results are presented in Figures 12 to 14, where the 90% confidence interval is represented.

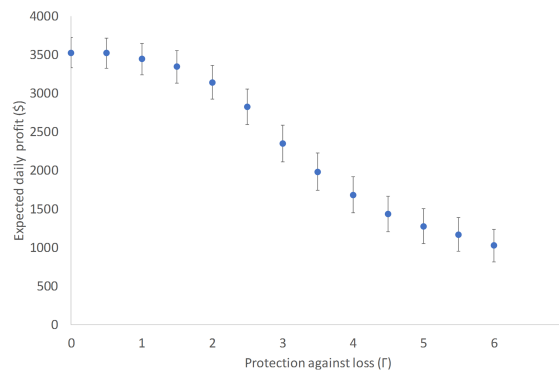


Figure 12: Experimental expected profit for different values of Γ




Figure 13: Experimental loss probability for different values of Γ

Figure 12 shows that the expected profit is reduced when Γ increases: this can be referred to as the “price of robustness”. However, this robustness allows to reduce the risk, both in terms of loss probability (see Figure 13) and in terms of “worst-case” profit (see Figure 14). For values of Γ greater than 4, the loss probability is even brought to zero: the robust model allows to identify the days when operating is too risky.

These figures highlight the necessity to find a compromise between expected profit, and risk. The objective of the decision-maker is to find a value of Γ for which the expected profit is not curtailed too much with respect to the deterministic model, while allowing to reduce the risk. Figures 15 and 16 illustrate this tradeoff, representing simultaneously the experimental value achieved for both objective (expected profit and risk management). It appears that, in the case of $\Gamma = 2$, the risk is significantly reduced, both in terms of loss probability (1.09% loss probability compared with 3.01% with the nominal strategy) and “worst-case” profit (-83.94\$ versus -1237.54\$), while the expected profit remains reasonably high (3143\$ compared with 3525\$ using the nominal strategy).

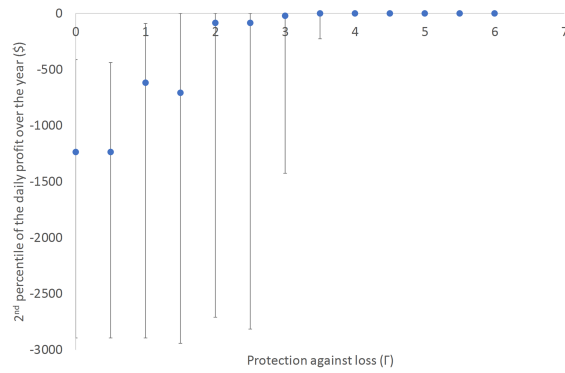


Figure 14: Experimental 2nd percentile of the profit for different values of Γ




Figure 15: Trade-off between the expected profit and the loss probability

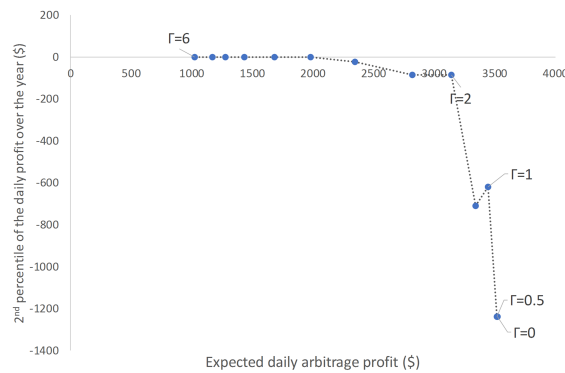


Figure 16: Trade-off between the expected profit and the worst-case profit (2nd percentile)

Therefore, $\Gamma = 2$ would be a relevant choice for the storage operator. Figure 17 and Figure 18 present the results that would have been obtained over the year 2016 for a storage facility operating using $\Gamma = 2$. Figure 17 shows the weekly profits obtained: we notice that the highest profits are reached during the summer, which is, in the case of New-York state, the period when demand peak occurs. 49% of the profits are realized between the 1st of June and the 31st of August. Besides, the general daily price profile observed in Figures 8a and 8b leads the storage facility to charge during the night (generally between 12 a.m. and 7 a.m.) and to discharge during the evening (mainly between 4 p.m. and 10 p.m.).




Figure 17: Weekly profit of the storage for year 2016




Figure 18: Hours of charging and discharging for year 2016

4.3 Impact of the storage operation on the market

The impact of the storage operation on the price was incorporated into the self-scheduling problem to avoid over-estimating the potential arbitrage profits for a large storage facility. This formulation allows to investigate the overall impact of the storage implementation on the prices. This may especially be interesting from the perspective of the market operator.

Though the size of the energy storage facility considered in the study is the biggest existing size for a battery system, it is still marginal with respect to the size of the market (the charging and discharging power considered is equivalent to 0.3% of the maximal load in New-York state). It might be of interest to study the price impacts of the storage for greater sizes. The highest sizes considered in the following study are not realistic for a single energy storage facility, but this might be seen as an approximation of the implementation of several storage facilities on the market, assuming that these facilities are operated simultaneously by a same player.

Tables 1 and 2 present respectively the impact of storage on the prices during charging and discharging hours, for different sizes of storage. The storage facilities are assumed to be scheduled using $\Gamma = 2$.

As expected, we notice that the storage operation reduces the price differential between off-peak hours (which are the discharging hours presented in Table 1) and peak hours (which are the charging hours presented in Table 2). During discharging operations, the energy storage facility causes a significant decrease in the average price, ranging from 0.67% to 4.54% depending on the size of the storage. During charging operations, the relative increase in the price varies from 1.30% to 11.25%.

We may also notice that the number of hours of operation, both in charge and in discharge, increases with the size of storage. This is due to the integration of the price-maker effect in the self-scheduling. When the storage size increases, it is sometimes better to distribute the purchases and sales over more hours to avoid reducing too much the peak price differential.

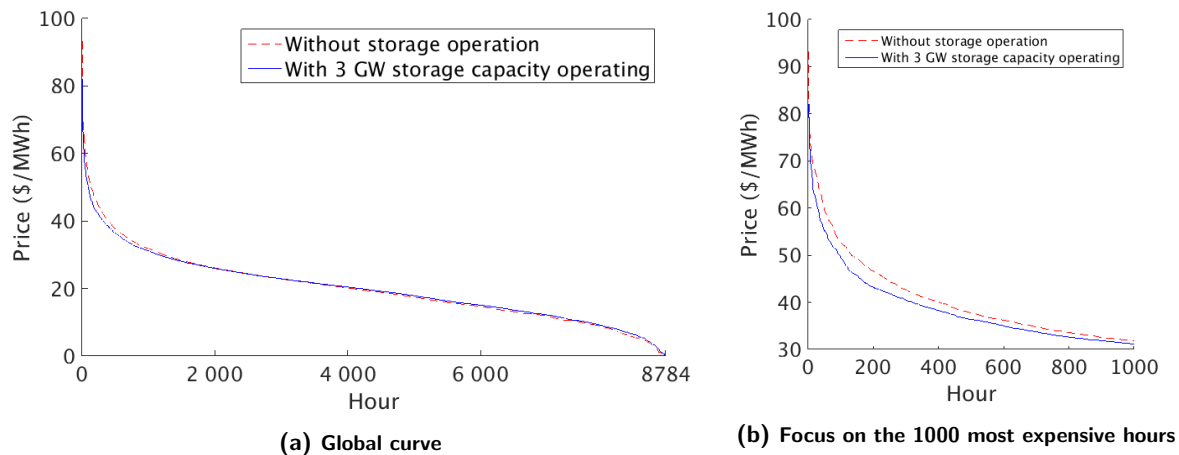
Table 1: Price comparison without and with storage during discharging hours for different size of storage

Size (% of the market)	Size (GW)	No. hours discharging	Price average (\$/MWh)		Price Change	
			Without	With	(\$/MWh)	%
0.3	0.1	1086	28.79	28.59	-0.19	-0.67
1	0.3	1458	27.80	27.38	-0.42	-1.52
2	0.6	1885	27.17	26.56	-0.61	-2.25
5	1.5	2403	26.97	26.02	-0.95	-3.53
10	3	2626	27.27	26.03	-1.24	-4.54

Table 2: Price comparison without and with storage during charging hours for different size of storage

Size (% of the market)	Size (GW)	No. hours charging	Price average (\$/MWh)		Price Change	
			Without	With	(\$/MWh)	%
0.3	0.1	1313	12.33	12.49	0.16	1.30
1	0.3	1465	12.36	12.77	0.42	3.40
2	0.6	1698	12.56	13.24	0.68	5.40
5	1.5	1956	13.19	14.33	1.15	8.71
10	3	2035	13.49	15.01	1.52	11.25

Figure 19 presents the price duration curves without storage and with the largest storage for the year 2016. For the sake of clarity, the duration curves obtained with other sizes for the storage facility are not represented. Especially, for lower sizes, the resulting curve is closer to the curve original curve (without storage operation). Figure 19 allows to measure the overall impact of storage on the prices for all the hours of the year. We notice that the presence of storage slightly reduces the volatility of the prices, lowering the highest prices and increasing the lowest prices. A focus on the 1000 first hours of the price duration curve (see Figure 19b) provides a more accurate overview of the impact during the most expensive hours of the year. For instance, the presence of storage allows to lower the maximal price by 12% (from 93.23\$/MWh to 82.12\$/MWh).

**Figure 19: Price duration curve over the year 2016 with and without storage**

The market operator may also be how this impact on the price is affected by the risk aversion of the storage. The results presented in Tables 1 and 2 assume that the storage operator will choose $\Gamma = 2$, but the choice may be different depending on the operator.

Tables 3 and 4 present respectively the impact of storage operation on the price during discharging and charging hours, for the maximal size of storage. We notice that the number of operating hours decreases when the risk-aversion increases. This result may be surprising, since we observed that risk aversion led the operator to distribute his purchases and sales over a larger time window. However, the increasing level of

risk-aversion concurrently reduce the number of days when the storage is operating, to avoid incur losses. The decrease in the number of operating hours is explained by this second effect.

Tables 3 and 4 both show that the impact of storage on the price is limited when the risk aversion of the storage increases. The increase during the charging hours drops from 13.36% for $\Gamma = 0$ to 7.62% for $\Gamma = 5$, while the decrease during the discharging hours drops from 7.35% for $\Gamma = 0$ to 3.51% for $\Gamma = 5$. This is due to the conservative strategies, which implies lower charging and discharging power, because of the distribution of purchases and sales.

Table 3: Price comparison without and with storage during discharging hours for different values of Γ

Gamma	No. hours discharging	Price average (\$/MWh)		Price Change	
		Without	With	(\$/MWh)	%
0	3562	25.42	23.55	-1.87	-7.35
1	3383	25.69	24.06	-1.64	-6.37
2	2626	27.27	26.04	-1.24	-4.54
3	1538	31.26	29.97	-1.29	-4.13
4	928	36.23	34.84	-1.38	-3.82
5	630	39.55	38.16	-1.39	-3.51

Table 4: Price comparison without and with storage during charging hours for different values of Γ

Gamma	No. hours charging	Price average (\$/MWh)		Price Change	
		Without	With	(\$/MWh)	%
0	2377	13.26	15.03	1.77	13.36
1	2302	13.27	15.03	1.76	13.26
2	2035	13.49	15.01	1.52	11.25
3	1427	14.55	15.92	1.37	9.45
4	938	16.42	17.8	1.38	8.39
5	666	17.6	18.94	1.34	7.62

5 Conclusion

This paper proposes an original approach to determine the optimal bidding strategy of a large energy storage facility doing arbitrage in a context of price uncertainty. In view of the size considered for the storage, it is critical to integrate the impacts of its operations on market prices: this is done using the supply curve of the market. Unlike previous works, this supply curve (resulting from the competitors strategies) is not assumed to be known perfectly. The storage operator actually needs to forecast the supply curve, which may fluctuate hour by hour. To address this issue, we investigated the drivers of the supply curve, which is mainly the result of the ranking of energy sources in the increasing order of their marginal costs. Based on this analysis, the wind power was excluded from the representation of the supply curve, hence reducing its variability. The remaining variability was then modeled by means of quantile regression, providing bounds that encapsulates the actual supply curve.

We proposed a robust model to deal with the uncertainty relating to this bounding. This robust model allows the storage operator to be immunized against the risk of loss, for a certain set of instances of the uncertain variables. The decision-maker may decide the level of conservatism he uses, choosing to consider a more or less wide set of feasible values for the uncertain variables. We observed that a broad variety of strategies could be implemented, depending on the risk-aversion of the operator. The nominal strategy (i.e. when the decision-maker does not take the uncertainty into consideration) consists in benefiting from the lowest (respectively highest) demand hours to charge (respectively discharge) the battery with high power output. The general trend is then to distribute the operations of storage on a greater range of time steps when the level of risk-aversion increases, to avoid being impacted by high perturbation at each hour of operation.

Implementing these strategies over a whole year, their efficiency in terms of profits and protection against loss was subsequently assessed, based on historical hourly net load and prices. It was found that an increase

in risk protection could be achieved, but at the expense of a reduction in the profit. Our study highlighted that certain levels of conservatism provided a good tradeoff between these two objectives. Especially, for $\Gamma = 2$, it was found that the loss probability could be lowered from 3% to 1%, reducing the expected profit by only 10% with respect to the nominal strategy.

Finally, the perspective of the market operator was considered, trying to assess the overall impacts of the implementation of energy storage on the market prices, depending on the size and on the strategy chosen by the storage operator. It was found that larger energy storage facilities could provoke a great increase in off-peak prices (up to 13.36% for zero risk-aversion operators) and a great decrease in peak-prices (up to 7.35%).

In the results presented, the strategy of the storage was optimized using a forecast of the supply curve, to take into account the imperfect information available for the storage operator. However, the actual supply curve should have been employed in a second step to compute the exact market price at each time step taking into account storage operation. Since this information was not available, we considered that the actual supply curve could be approximated using the actual hourly net load and price data to estimate the value of the uncertain variable, and hence the position of the supply curve between the two bounds. For a better modeling of the market impact of storage, it would be interesting to incorporate the storage strategies designed in this study in the market clearing process.

In this study, the issue relating to the location of the storage facility was not addressed. To assess the actual profits of the storage facility, transition costs should also be incorporated in the self-scheduling model. The representative supply curves and the developed robust model could be used to formulate an optimization framework to find the optimal location for the storage facility in future works.

Finally, we have to keep in mind that the hourly net demand for the 24 hours ahead was assumed to be known perfectly in this study. It could be interesting to assess the effect of forecast errors on the profits obtained. A further step could consist in incorporating the net demand uncertainty in the self-scheduling model. This would also allow to consider a further scheduling horizon.

References

- [1] Marc Beaudin, Hamidreza Zareipour, Anthony Schellenberglabe, and William Rosehart. Energy storage for mitigating the variability of renewable electricity sources: An updated review. *Energy for Sustainable Development*, 14(4):302–314, 2010.
- [2] Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of operations research*, 23(4):769–805, 1998.
- [3] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.
- [4] Dimitris Bertsimas, Eugene Litvinov, Xu Andy Sun, Jinye Zhao, and Tongxin Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems*, 28(1):52–63, 2013.
- [5] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.
- [6] Erick Delage, Luca Gianoli, and Brunilde Sanso. A practicable robust counterpart formulation for decomposable functions : A network congestion case study. *Les cahiers du GERAD*, (G-2015-27), 2015.
- [7] Easan Drury, Paul Denholm, and Ramteen Sioshansi. The value of compressed air energy storage in energy and reserve markets. *Energy*, 36(8):4959–4973, 2011.
- [8] Ruiwei Jiang, Jianhui Wang, and Yongpei Guan. Robust unit commitment with wind power and pumped storage hydro. *IEEE Transactions on Power Systems*, 27(2):800–810, 2012.
- [9] Roger Koenker. *Quantile regression*. Number 38. Cambridge university press, 2005.
- [10] Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica* 46(1):33–50. *Mathematical Reviews (MathSciNet)*, 1978.
- [11] Hamed Mohsenian-Rad. Coordinated price-maker operation of large energy storage units in nodal energy markets. *IEEE Transactions on Power Systems*, 31(1):786–797, 2016.
- [12] Vito MR Muggeo. Segmented: an r package to fit regression models with broken-line relationships. *R news*, 8(1):20–25, 2008.

- [13] NYISO. Nyiso markets. 2015.
- [14] NYISO. Power trends 2015. 2016.
- [15] US Department of Energy. Grid energy storage. 2013.
- [16] California Independent System Operator. 2016-2017 transmission plan. March 2017.
- [17] Soroush Shafiee, Payam Zamani-Dehkordi, Hamidreza Zareipour, and Andrew M Knight. Economic assessment of a price-maker energy storage facility in the alberta electricity market. *Energy*, 111:537–547, 2016.
- [18] Ramteen Sioshansi, Paul Denholm, and Thomas Jenkin. A comparative analysis of the value of pure and hybrid electricity storage. *Energy Economics*, 33(1):56–66, 2011.
- [19] Ramteen Sioshansi, Paul Denholm, Thomas Jenkin, and Jurgen Weiss. Estimating the value of electricity storage in pjm: Arbitrage and some welfare effects. *Energy economics*, 31(2):269–277, 2009.
- [20] James E Smith and Robert L Winkler. The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006.
- [21] Jorge AM Sousa, Fábio Teixeira, and Sérgio Faias. Impact of a price-maker pumped storage hydro unit on the integration of wind energy in power systems. *Energy*, 69:3–11, 2014.
- [22] A. L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21(5):1154–1157, 1973.
- [23] Fábio Teixeira, Jorge de Sousa, and Sérgio Faias. How market power affects the behavior of a pumped storage hydro unit in the day-ahead electricity market? In *European Energy Market (EEM), 2012 9th International Conference on the*, pages 1–6. IEEE, 2012.
- [24] Anupam A Thatte, Le Xie, Daniel E Viassolo, and Sunita Singh. Risk measure based robust bidding strategy for arbitrage using a wind farm and energy storage. *IEEE Transactions on Smart Grid*, 4(4):2191–2199, 2013.
- [25] Paul Twomey and Karsten Neuhoff. Wind power and market power in competitive markets. *Energy Policy*, 38(7):3198–3210, 2010.
- [26] Rahul Walawalkar, Jay Apt, and Rick Mancini. Economics of electric energy storage for energy arbitrage and regulation in new york. *Energy Policy*, 35(4):2558–2568, 2007.
- [27] Long Zhao and Bo Zeng. Robust unit commitment problem with demand response and wind energy. In *Power and Energy Society General Meeting, 2012 IEEE*, pages 1–8. IEEE, 2012.