

# Energy Storage Arbitrage in Two-settlement Markets: A Transformer-Based Approach

Saud Alghumayjan<sup>\*</sup>, Jiajun Han<sup>†</sup>, Ningkun Zheng<sup>\*</sup>, Ming Yi<sup>‡</sup>, Bolun Xu<sup>\*</sup>

<sup>\*</sup> Earth and Environmental Engineering, <sup>†</sup> Electrical Engineering, <sup>‡</sup> Data Science Institute

Columbia University

New York, NY 10027, United States

{saa2244, jh4316, nz2343, my2826, bx2177}@columbia.edu

**Abstract**—This paper presents an integrated model for bidding energy storage in day-ahead and real-time markets to maximize profits. We show that in integrated two-stage bidding, the real-time bids are independent of day-ahead settlements, while the day-ahead bids should be based on predicted real-time prices. We utilize a transformer-based model for real-time price prediction, which captures complex dynamical patterns of real-time prices, and use the result for day-ahead bidding design. For real-time bidding, we utilize a long short-term memory-dynamic programming hybrid real-time bidding model. We train and test our model with historical data from New York State, and our results showed that the integrated system achieved promising results of almost a 20% increase in profit compared to only bidding in real-time markets, and at the same time reducing the risk in terms of the number of days with negative profits.

**Index Terms**—Electricity markets, Energy storage, Machine learning

## I. INTRODUCTION

### A. Motivation

Energy storage participants are increasingly pivotal in electricity markets. In 2018, the Federal Energy Regulatory Commission (FERC) issued Order 841, mandating the inclusion of storage in all electricity markets [1]. Consequently, all system operators across the United States now permit storage to place both charge and discharge bids. Such a bidding system allows storage entities to craft their offers based on market prices. As storage capacity grows, the emphasis has transitioned from ancillary service markets, characterized by limited market capacity [2], to wholesale energy markets [3], [4].

United States' electricity markets follow a two-stage settlement design: a day-ahead market (DAM) and a real-time market (RTM). The DAM, covering the subsequent day's 24 trading intervals, requires bids to be placed by a specified closing time for each hour. The RTM is responsible for balancing the supply and demand deviations from day-ahead schedules. Existing literature offers multiple models for energy storage arbitrage but mostly focuses on either DAM or RTM. The most common method is model predictive control (MPC) [5], but its

effectiveness diminishes for long look-ahead windows, such as 24 hours [6]. Stochastic dynamic programming (SDP) [7] and reinforcement learning (RL) [8] are emerging approaches but often encounter computation difficulties.

RTM is significantly more attractive for storage participants to arbitrage as the prices are more volatile than DAM [9]. On the other hand, recent studies [10] have shown while RTM provides higher profits, storage participating in both DAM and RTM provides the best social welfare outcome. Hence, system operators like the California Independent System Operator (CAISO) have mandated storage must bid both into day-ahead and real-time markets [11]. However, the newest bidding data from CAISO [12] shows most participants submit unreasonably high bids day-ahead to purposefully avoid day-ahead settlements, like due to lacking good joint bidding solutions.

### B. Contribution

While it is lucrative for energy storage entities to engage in both DAM and RTM, the inherent challenges stem from multi-stage decision-making amidst volatile and unpredictable electricity prices. Notwithstanding the existing research [13], our study proposes a novel and practical approach integrating machine learning to bid energy storage into both day-ahead and real-time markets. Our salient contributions are:

- We propose a novel energy storage arbitrage in two-settlement markets framework that combines a transformer-based price prediction model for day-ahead bidding and a long short-term memory (LSTM)-dynamic programming hybrid real-time bidding.
- We build a transformer-based model for predicting electricity real-time prices and train it over the New York Independent System Operator (NYISO) market.
- We evaluated our proposed framework using actual price data obtained from four NYISO zones. The findings revealed that the model demonstrated a notable improvement, yielding up to a 29% increase in profit when compared to solely participating in the RTM.

Notably, to our knowledge, this will be the first study in the literature to demonstrate the benefit of systematic day-ahead and real-time bidding using real-world price data. Results from this paper will empower storage participants to systematically bid into two-settlement markets, generating more profits and social welfare.

The authors would like to acknowledge supports from Columbia Data Science Institute, Red River Clean Energy Inc., and King Abdulaziz City for Science and Technology (KACST).

### C. Related Work

Electricity price forecasting has been investigated heavily in the past years [14], [15]. Electricity is still hard to store and demand must be instantaneously balanced with demand in power systems, making the price of electricity have unique dynamics as it has special shapes for each time level and even sudden peaks. Many models have been developed previously to solve this problem; each has its own advantages and disadvantages [16]. These models vary from statistical models [17], economic equilibrium models [18], agent-based models [19], to experimental models [20]. Generally, there are three main types of price forecasting, short-term, mid-term, and long-term, focusing on forecasting hours ahead, weeks to months ahead, and years ahead. The horizon changes based on the needed application. This study focuses on the short-term horizon as it strongly influences the day-ahead and real-time markets operations.

Artificial intelligence (AI) models have risen as a leading solution, overcoming many limitations inherent in alternative approaches. Models such as multi-layered perceptron (MLP) have been shown to provide satisfactory performance and adaptability [21]. Furthermore, LSTM models have emerged as particularly suitable for price forecasting, outperforming hybrid methodologies like convolutional neural networks LSTM (CNN-LSTM) [22]. A comprehensive review of AI and statistical models in electricity price forecasting, encompassing datasets from diverse regions, is provided by [23].

Energy storage arbitrage in electricity wholesale markets has experienced rapid growth in recent years [24]. Storage entities in wholesale electricity markets can participate in arbitrage by charging during periods of low prices and discharging during periods of high prices, thereby maximizing their profits. To evaluate potential profits, various models have been introduced in the literature, including price taker and strategic-behavior models [25]. Price taker models operate under the assumption that their actions do not influence market operations, hence, they take the same price offered by the market. This simplifies the optimization problem by disregarding power balance constraints. Several approaches have been proposed to address this challenge, with Model Predictive Control (MPC) models often leading the literature due to their modeling flexibility [5], [6]. Additionally, SDP [7] and RL [8] models has been proposed frequently.

Two-settlement markets offer various participation options, including participation in DAM, RTM, and virtual bidding (VB). DAM participants are required to submit their bids one day in advance, necessitating accurate price forecasting for 24 hours-ahead. [26] proposes an MPC model using forecasted prices. On the other hand, RTM participants submit bids for 5-minute intervals during operational hours, facing highly volatile prices. [27] introduces a deep reinforcement learning (DRL) model aimed at discovering an optimal policy to maximize arbitrage profits in RTM. [7] propose an analytical SDP approach to address real-time price arbitrage with efficient computation. Additionally, [28] outline a highly transferable

price arbitrage model that combines model-based dynamic programming with neural networks.

VB option offers market participants the flexibility to place bids without the need for physical assets, enabling them to sell in the DAM and buy the same amount in the RTM, or vice versa. This option aims to improve market efficiency by minimizing the disparity between DAM and RTM prices [29]. Additionally, [30] discusses a virtual bidding framework utilizing a gradient boosting tree-based algorithm that takes price sensitivity into account. Market participants have the opportunity to participate in both DAM and RTM simultaneously, aiming to enhance their overall arbitrage profits. However, navigating multi-stage decision-making poses challenges, given the uncertainties in electricity prices. The proposed framework by [13] outlines an optimal bidding strategy for energy storage arbitrage across DAM and RTM, albeit without factoring in price uncertainty. Furthermore, [31] have introduced an SDP model for storage arbitrage in DAM and RTM using conventional statistical methods to model the uncertainties.

The remainder of the paper is organized as follows: Section II introduces the methodology. Section III describes the dataset, and Section IV presents the case study. Section V concludes the paper.

## II. ENERGY STORAGE ARBITRAGE FRAMEWORK




Fig. 1. Pipeline of the proposed solution.

Our profit-maximizing framework consists of a DAM bidding stage and sequential RTM bidding stages on a daily basis. The DAM stage uses the predicted hourly average real-time price as the charge and discharge bid for the same hour in DAM. The RTM stage designs real-time bids sequentially for each hour to maximize real-time arbitrage profits. Fig. 1 shows our full framework. In this framework, the DAM bids are based on predicted RTM prices, but the RTM bids are independent of DAM settlements. In the remainder of this section, we will present the theoretical analysis for the presented framework, our price prediction model, and the real-time arbitrage model.

### A. Problem Statement

We assume the energy storage is a price-taker and submits economic bids to both DAM and RTM. The storage does not model the impact of its action over price formations and hence does not seek to exercise market power. Each bid, encompassing both a charging (purchase) and discharging (sale) bid, is submitted for every market-clearing hour. Each day comprises 48 bids: 24 for the DAM submitted collectively before the operating day starts and 24 for real-time, where bids are placed hourly with clearances at five-minute intervals. We thus formulate the two-settlement arbitrage problem as follows:

$$\begin{aligned} \max_{B \in \mathcal{B}} \mathbb{E}_{\lambda, \pi \in \mathcal{M}} \left\{ \sum_{t \in \mathcal{T}} \lambda_t (p_t^D - b_t^D) - \sum_{s \in \mathcal{S}} c p_s^R \right. \\ \left. + \sum_{s \in \mathcal{S}} \pi_s (p_s^R - b_s^R - p_{\tau|s \in \tau}^D / S - b_{\tau|s \in \tau}^D / S) \right\} \quad (1a) \\ \{p_t^D, b_t^D, p_s^R, b_s^R\} \in B \quad \forall t \in \mathcal{T}, s \in \mathcal{S} \quad (1b) \end{aligned}$$

in which

- $t$  is the index of the hour at which the day-ahead market clears, the set  $\mathcal{T}$  represents all hours from one day.
- $s$  is the index of real-time market clearing intervals, usually per five minutes, the set  $\mathcal{S}$  represents all real-time market intervals in one day.
- $S$  is the number of real-time market clearing intervals in one day-ahead market interval. In our case study, we consider hourly day-ahead market clearing and 5-minute real-time market clearing intervals, so  $S$  will be 12.
- $\lambda_t$  is the day-ahead market clearing price.
- $\pi_s$  is the real-time market clearing price.
- $p_t^D$  and  $b_t^D$  are the storage's discharging and charging energy in day-ahead clearing results, respectively;
- $p_s^R$  and  $b_s^R$  are the storage's *physical* discharging and charging energy in real-time, respectively;
- $c$  is the cost to discharge; note that we assume the physical cost, such as degradation, only occurs during the discharge stage. This is to align with conventional generation cost models.
- $B$  is the bidding policy while  $\mathcal{B}$  represents the set of all feasible bidding policies.

The objective function maximizes the arbitrage profit in DAM and RTM under the expectation that the day-ahead price and real-time price are a stochastic process to the market clearing model ( $\lambda, \pi \in \mathcal{M}$ ). The first term represents DAM revenue, the second term is the physical discharge cost, and the third term represents real-time re-settlement based on day-ahead commitments. Note the subscript  $\tau|s \in \tau$  represents the hour corresponding to the current RTM interval.

In this objective, if we set  $p_t^D$  and  $b_t^D$  to zero, then the first term becomes zero and formulation becomes a real-time arbitrage objective. On the other hand, if we set  $p_s^R = p_{\tau|s \in \tau}^D / S$  and  $b_t^D = b_{\tau|s \in \tau}^D / S$ , then it indicates the storage's physical charge and discharge power fully follow the DAM clearing results. Hence, the third term becomes zero and the storage is only participating in the DAM. (1b) requires all storage

actions must be the result of a bidding policy  $B$ , while  $B$  must come from the set of all feasible bidding policy  $\mathcal{B}$ , which are policies that satisfy market clearing procedures and the following physical constraints ( $\forall t \in \mathcal{T}, s \in \mathcal{S}$ ):

$$0 \leq p_t^D, b_t^D \leq P \quad (1c)$$

$$0 \leq p_t^R, b_t^R \leq P/S \quad (1d)$$

$$e_t^D - e_{t-1}^D = -p_t^D / \eta + b_t^D \eta \quad (1e)$$

$$e_s^R - e_{s-1}^R = -p_s^R / \eta + b_s^R \eta \quad (1f)$$

$$0 \leq e_t^D \leq E \quad (1g)$$

$$0 \leq e_s^R \leq E \quad (1h)$$

$$p_s^R = 0 \text{ if } \pi_t < 0 \quad (1i)$$

where  $P$  is the storage power rating,  $\eta$  is the efficiency,  $e_t^D$  is the state-of-charge (SoC) constraint in DAM clearing, while  $e_s^R$  is the physical SoC constraint. (1c) and (1d) model the power rating constraint in day-ahead and real-time operations, (1e), (1f), (1g) and (1h) model the SoC evolution constraints in the same manner. (1i) ensures the storage doesn't discharge when the price is negative, a sufficient condition to avoid simultaneous charging and discharging [28]. Note that this constraint is only necessary in RTM as day-ahead prices are seldom negative.

### B. Decomposition Reformulation

Using (1) to directly find the optimal arbitrage policy can be extremely difficult as the electricity prices follow highly stochastic multi-stage processes and the bidding requirements are also non-trivial to model. Instead, we temporally relax the bidding constraint (1b) and treat the powers as direct decision variables. The goal thus is to establish intuition in the market participation and use it to design bids. We reformulate (1a) by partially merging the first and the third term to decouple day-ahead and real-time decision variables and stochastic processes, the resulting formulation is:

$$\begin{aligned} \max_{p_t^D, b_t^D} \mathbb{E}_{\lambda} \left\{ \sum_{t \in \mathcal{T}} \lambda_t (p_t^D - b_t^D) - \bar{\pi}_t (p_t^D - b_t^D) \right\} \\ + \max_{p_s^R, b_s^R} \mathbb{E}_{\pi} \left\{ \sum_{s \in \mathcal{S}} \pi_s (p_s^R - b_s^R) - \sum_{s \in \mathcal{S}} c p_s^R \right\} \quad (2) \end{aligned}$$

where  $\bar{\pi}_t$  is the expected hourly average real-time price. This reformulation is based on the following approximation to decouple price expectation from power expectation.

**Proposition 1.** *Real-time price expectation under price-taker assumptions.* The DAM clearing results are statistically independent of real-time price realizations. This allows us to convert the whole-day summation of RTM intervals ( $s \in \mathcal{S}$ ) into summation of hours ( $t \in \mathcal{T}$ ) and summation of intervals within a given hour ( $s \in t$ ), hence

$$\begin{aligned} \mathbb{E}_{\pi} \left\{ \sum_{s \in \mathcal{S}} \pi_s p_{\tau|s \in \tau}^D / S \right\} &= \sum_{t \in \mathcal{T}} \mathbb{E}_{\pi} \left\{ \sum_{s \in t} \pi_s p_{\tau|s \in \tau}^D / S \right\} \\ &= \sum_{t \in \mathcal{T}} (\mathbb{E}_{\pi} \sum_{s \in t} \pi_s) p_t^D = \sum_{t \in \mathcal{T}} \bar{\pi}_t p_t^D \end{aligned}$$

and similarly to  $b_t^D$ .

*Proof.* This proposition is based on the price-taker assumption that the storage does not assume its action would impact the market price and will not try to exercise market power, then it is trivial that  $p_\tau^D$  will not change the distribution of price  $\pi_s$ . This is also a common practice in bidding designs [32], [33] when assuming participants are price takers in the market  $\square$

Based on the previous proposition, we state our second proposition

**Proposition 2.** *Two-stage reformulation.* (1a) and (2) are equivalent formulations under the price-taker assumption.

*Proof.* Proposition 1 allows us to decouple the price expectation from the action, and further allows us to decouple the stochastic model of  $\lambda$  and  $\pi$ . Thus, we only need to justify the separation of the maximization objective from (1a) to (2), which is intuitively given there is no coupling physical constraint between DAM and RTM decisions. Thus, (1a) and (2) are equivalent.  $\square$

By observing the format of the reformulated objective (2): the first maximization term (associated with constraints (1c), (1e), and (1g)) corresponds to the day-ahead bidding problem with the bidding cost equal to the expected average real-time price over the corresponding hour; the second maximization term (associated with constraints (1d), (1f), (1h) and (1i)) is a real-time arbitrage model without considering any day-ahead results. Thus, this reformulation justifies our bidding model as shown in Figure 1, which also satisfies market bidding requirements as in the original formulation (1).

### C. Electricity Price Forecasting

According to our decomposition reformulation result, the key to designing day-ahead bids is to predict the real-time price of the same hour. Thus in this subsection, we will present a machine learning model for predicting real-time prices.

In the first stage of our framework, we forecast the real-time price, which serves as a critical input for our day-ahead arbitrage model. The foundation of our forecasting model is built upon the patch time series transformer (PatchTST), a model primarily developed as outlined by [34], [35]. Transformers [36], a recent class of deep learning models initially designed for applications in natural language processing, have consistently demonstrated remarkable performance across various domains. Among these, the PatchTST [34] stands out as a prominent model for time-series forecasting, consistently surpassing other transformers. Fig. 2 shows an overview of our real-time price forecasting model. It only uses the encoder from the traditional transform and has two main features: channel-independence, split the multivariate time series into univariate time series that share the same weights, and each channel may have a different attention map; patching, decompose the time series into subsequences, capture local information, and reduce model complexity.

We follow a similar formulation to [34], but with the input to be a multivariate time-series with a look-back window  $L$

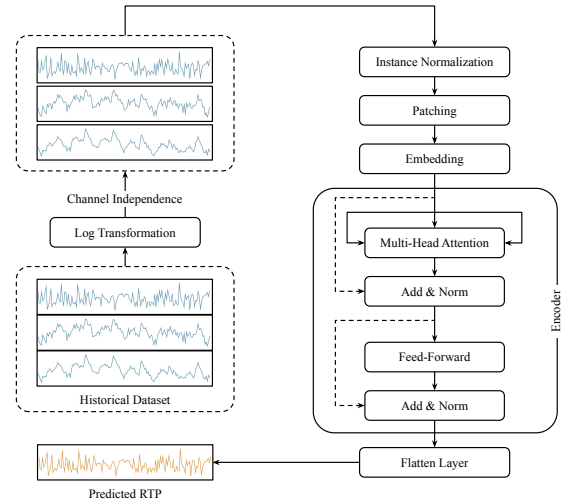


Fig. 2. Our real-time price forecasting model overview.

and the output a single univariate time-series with a forecasting horizon of  $T$  as described by the following equations:

$$\mathbf{x} = \begin{bmatrix} x_{1,1} & \cdots & x_{1,L} \\ \vdots & \ddots & \vdots \\ x_{M,1} & \cdots & x_{M,L} \end{bmatrix}, \quad \hat{\mathbf{x}} = \{\hat{x}_{1,1}, \dots, \hat{x}_{1,T}\} \quad (3)$$

The process begins by breaking down the multivariate time-series input into several distinct univariate time-series during a phase called channel-independence. Each of these individual univariate time series is then separately fed into the transformer backbone.

The transformer backbone primarily comprises three key components: instance normalization, patching, and the traditional transformer encoder. Every univariate time series will first be normalized with zero mean and unit standard deviation and then divided into a sequence of patches. This patching not only minimizes the number of input tokens required for the encoder but also serves to reduce computational complexity. Moreover, it helps the model to effectively capture extended historical sequence, enhancing its predictions.

The primary purpose of the transformer encoder is to transform input patches into a different representation. This transformation involves several steps. Initially, a linear projection of the input is obtained. Subsequently, scaled dot-product attention is applied, as expressed by the equation:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right) V \quad (4)$$

Here, the matrices  $Q$ ,  $K$ , and  $V$  correspond to the query, key, and value matrices, respectively, and have been transformed by each head in the multi-head attention mechanism. Following that, the resulting attention values are batch-normalized and passed through a feed-forward neural network before reaching the flattening layer, ultimately yielding the predictions as shown in (3).

Finally, the loss function that our model minimizes during training is the mean square error (MSE), which we employ to guide the model’s optimization process. This loss function is mathematically defined as follows:

$$\text{MSE} = \frac{1}{T} \sum_{t=1}^T (x_t - \hat{x}_t)^2 \quad (5)$$

#### D. Real-Time Energy Storage Arbitrage Model

Our storage arbitrage model for real-time bidding is based on [28], which solves the real-time arbitrage problem (second part of (2)) following a non-anticipatory bidding policy. The model predicts the opportunity value of the state of charge (SOC) and then maximizes the storage arbitrage profit. The model combines model-based dynamic programming with neural networks. Unlike MPC, which is very sensitive to price volatility over a 24-hour look-ahead horizon, this method focuses on predicting the value function within a single time period. Furthermore, [28] is highly transferable, which makes it suitable for different markets and participation scenarios. It has been evaluated using NYISO price data; it achieved a profit ratio ranging from 70% ~ 90% when compared to the perfect forecast. These features make this method ideal for our system. Due to page limitations, we will not present the full formulation in this paper as this model has already been published.

### III. DATA DESCRIPTION

This work utilizes a dataset spanning a five-year period, ranging from 2017 to 2021. This dataset encompasses multiple price zones within the NYISO market, including NYC, LONGIL, NORTH, and WEST:

- Hourly Day-Ahead Price (DAP): This variable represents the smallest available resolution for day-ahead LMP. It has been used in all parts of the framework, and it has a direct influence on both day-ahead and real-time markets.
- 5-minute Real-Time Price (RTP): This is the smallest resolution in NYISO. These granular data points help in finding the best solution by the real-time arbitrage model.
- Hourly Real-Time Price (RTP): This resolution was added to match the resolution of the day-ahead prices when used for the day-ahead arbitrage model.
- Hourly demand forecasting: This variable provides insights into the predicted energy demand on an hourly basis. It’s critical to understand how real-time pricing is influenced by changes in energy demand.

Our framework is constructed around these variables, driven by their direct impact on DAM and RTM, as well as their availability. These variables encapsulate pricing dynamics from both DAM and RTM. Fig. 3 shows the density distribution of the prices in the testing set. We can see that real-time prices have more outliers than the day-ahead prices. Also, the log-transformed prices gave a better distribution, hence expected to improve the forecasting task. Table I shows statistics about the historical price data for the four NYISO zones, it is mainly the price mean ( $\mu$ ), the price standard deviation ( $\sigma$ ), and the number of negative real-time prices.




Fig. 3. Density distribution of the day-ahead and real-time prices for NYC Zone price during 2021.

TABLE I  
NYISO PRICE ZONES’S STATISTICS

Zone	$\mu$ DAP	$\sigma$ DAP	$\mu$ RTP	$\sigma$ RTP	Negative RTP
NYC	42.54	18.94	42.49	26.89	1
LONGIL	54.11	31.40	54.85	63.72	42
NORTH	22.65	16.07	23.51	30.27	937
WEST	31.03	19.47	30.89	27.08	44

### IV. EXPERIMENT SETUP AND RESULTS

The study was performed using data collected from NYISO over the span of five years (2017-2021), four years for training and validation and one year for testing. Real-time prices, day-ahead prices, and demand forecasting were chosen as features due to their direct influence on the real-time price. A look-back window of two weeks and the prediction horizon of 24 hour-ahead were considered. The data were normalized before being fed into the model and the price features were log-transformed to reduce the impact of the extreme spikes on the model predictions. We consider battery storage with a 2-hour storage duration, 1 MWh storage capacity, 90% one-way efficiency, and \$10/MWh fixed discharge cost that covers degradation cost and other operational costs.

All experiments are run on a high performance computing cluster with Intel Xeon Platinum 8640Y 2Ghz CPU, 512GB Memory, and 2xNVIDIA A40 GPU. We used Tensorflow 2.15.0, CVXPY 1.4.2, PyTorch 2.2.0, and CUDA 12.1. All codes are available on Github<sup>1</sup>.

#### A. Electricity Prediction Results

In this subsection, we present the results of the electricity real-time price forecasting model. We benchmark the model with the DLinear model [37], which was shown to beat most time-series transformers. The DLinear is a special linear model designed for time-series problems; it utilizes the decomposition layer of the Autoformer model [38], where it splits the input data into multiple components that represent the trend and seasonality in the input and then predicts the output based on these components. Table II shows a comparison of our model (the transformer) and the baseline model (the DLinear) in terms of the prediction accuracy on the testing set using

<sup>1</sup><https://github.com/Alghumayjan/Energy-Storage-Arbitrage-in-Two-Settlement-Markets>

MSE for the NYC price zone. We can see that the our model beats the baseline model. Also, we can see that the model with log transformation achieved lower errors in general.

TABLE II  
TESTING SET PREDICTION ACCURACY

Model	Ours	log-Ours	Baseline	log-Baseline
MSE	0.458	<b>0.429</b>	0.506	0.457

We have considered different metrics to evaluate the models. Based on [15], one of the most suitable metrics for our case is the relative mean absolute error (rMAE) [39], which is the ratio of the MAE resulting from the original model and some naive model. (6) shows how the rMAE is calculated where the variable  $x$  represents the actual output values, while  $\hat{x}$  represents the forecasted output values. The naive model is described as  $\hat{x}_{d,h}^{\text{naive}} = x_{d-7,h}$  where in our case, we consider the 7-day lag model, where the forecast is the same hour one week ago. Table III shows the results based on the MAE and rMAE metrics after rescaling the predictions to the actual levels. These results align with the results from MSE, where our model with log transformation has the lowest error.

$$\text{rMAE} = \frac{\text{MAE}}{\text{MAE}_{\text{naive}}} = \frac{\sum_{d=1}^{365} \sum_{h=1}^{24} |x_{d,h} - \hat{x}_{d,h}|}{\sum_{d=1}^{365} \sum_{h=1}^{24} |x_{d,h} - \hat{x}_{d,h}^{\text{naive}}|} \quad (6)$$

TABLE III  
RESCALED TESTING SET PREDICTION ACCURACY

Model	Ours	log-Ours	Baseline	log-Baseline
MAE	11.01	<b>10.62</b>	11.95	10.97
rMAE	0.62	<b>0.60</b>	0.67	0.62

### B. Combined Profit Results

We now present the results of the arbitrage simulation for the entire proposed framework described in Section II. In this study, we investigate multiple settings and participating options. It mainly participates in DAM only, RTM only, VB, or both DAM and RTM. We consider these options under the perfect forecast real-time price case and the forecasted real-time price case:

- **DA+RT-PF, DA-PF, VB-PF**: participates in both DAM and RTM, only in the DAM, or only in VB, respectively, using perfect forecast real-time price.
- **DA+RT-F, RT-F, VB-F**: participates in both DAM and RTM, only in the DAM, or only in VB, respectively, using forecasted RT price via the developed models.

Table IV presents a summary model's performance for all the models in terms of prediction accuracy and total profit for the testing year (2021) across the four NYISO price zones considered. The profit presented is for the combined market participation option (DA+RT-F). We also report the improved

profit margin (IPM), which indicates the percentage of the increased profit we got by participating in the DAM and RTM over participating in the RTM only.

The log-ours model has the highest (Bold) or second-highest (underlined) profit across all zones. In addition, predictor errors do not strictly correlate with profits. For example, the model with the lowest error in the LONGIL zone is the baseline, but the model with the highest profit is ours. Moreover, the LONGIL zone has the highest prediction errors and the lowest IPM in all the models, which can be explained by Table I; LONGIL has a very high standard deviation in the real-time price, making it highly volatile and harder to predict. Furthermore, note that the baseline models suffer a lot in the NORTH zone, as they have IPM lower than 4% in comparison to our models, which have IPM higher than 14%; this suggests that the baseline models is too sensitive to negative prices, while our models are more robust. In general, the results show that our models is more profitable and reliable.

Fig. 4 shows the cumulative profit for all participating options considered in the study. We chose the log-ours model in all cases for its overall best performance. The proposed joint bidding approach provides the highest profit among all cases when assuming realistic price predictions. On the other hand, VB with perfect forecasts and DA+RT-PF case, which also includes VB, significantly surpasses all other approaches as it can fully take advantage of real-time price spikes. However, it is unrealistic to expect such high yields as price spikes are extremely difficult to predict [40], [41], and the VB profit becomes the lowest when using realistic price forecasts.

### C. Sensitivity Analysis

To assess the robustness of our bidding strategy against errors in real-time price forecasting, we demonstrate how the profitability of our proposed approach degrades with the performance of the real-time price forecaster. We have trained multiple versions of the real-time price forecaster, each with hyperparameters sampled from a uniform distribution. Subsequently, we conducted arbitrage simulations using the resulting forecasted prices. The left plot in Fig. 5 illustrates the cumulative profit of our approach compared to the baseline. Meanwhile, the right plot shows the errors of the predictions against the total profit. The figure shows that lower prediction errors correspond to higher profits. Furthermore, our model consistently exhibits superior profitability and reduced errors in general.

### D. Risk Analysis

By analyzing the cumulative profit curves and daily profit distributions, we observe that when predicted prices are utilized, many instances result in negative profits, posing a potential threat to energy storage owners. Ideally, we aim for results that closely resemble the scenarios with perfect forecasts. Table V provides insights into the frequency of daily negative profit occurrences for the three market participation options. Our analysis reveals that participating in both markets, while yielding on average 20% higher profits than RTM-only, does

TABLE IV  
MODEL PERFORMANCE SUMMARY

Price Zone	RTP mean	log-Ours				log-Baseline			
		MAE (\$/MWh)	rMAE	Profit (\$)	IPM	MAE (\$/MWh)	rMAE	Profit (\$)	IPM
NYC	42.74	10.62	0.60	<b>13875</b>	29%	10.97	0.62	12580	18%
LONGIL	55.05	27.08	0.76	<b>33849</b>	8%	29.40	0.82	32786	5%
NORTH	23.53	13.67	0.68	<b>20246</b>	18%	17.87	0.89	17318	1%
WEST	31.03	12.00	0.68	<b>16836</b>	26%	13.59	0.77	15289	14%

Price Zone	RTP mean	Ours				Baseline			
		MAE (\$/MWh)	rMAE	Profit (\$)	IPM	MAE (\$/MWh)	rMAE	Profit (\$)	IPM
NYC	42.74	11.01	0.62	<u>13802</u>	29%	11.95	0.67	12564	18%
LONGIL	55.05	26.06	0.73	<b>33995</b>	9%	25.95	0.73	33727	8%
NORTH	23.53	13.41	0.66	<u>19490</u>	14%	14.84	0.74	17856	4%
WEST	31.03	11.70	0.66	<u>16785</u>	26%	12.19	0.69	15345	15%




Fig. 4. The results of the system for all NYISO four price zones during 2021. Note that the dashed lines are participating options using perfect real-time price forecast, and the solid lines are participating options using forecasted real-time price generated by the log-Transformer model.

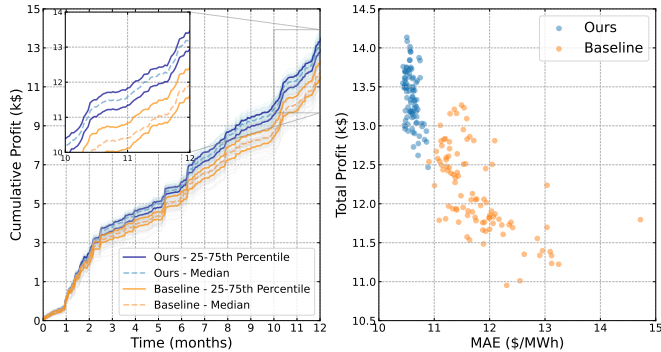


Fig. 5. Left: The cumulative profit of the trained models with randomly sampled hyperparameters. Right: the error of the models and their profits.

TABLE V  
NUMBER OF NEGATIVE PROFIT DAYS OF VARIOUS PARTICIPATION

Zone	VB	RT	RT+DA	IPM
NYC	104	90	59	29%
LONGIL	153	38	39	8%
NORTH	111	28	27	18%
WEST	115	39	33	26%

not increase the risk of profit loss. In particular, in NYC, joint participation reduced negative profit days by 34% compared to RTM-only. This indicates that combined market participation mitigates risk and enhances overall profitability or, at worst, maintains higher profitability levels while managing risk. VB, on the other hand, provides the highest number of negative profit days. This result reiterates that while VB is promising with perfect foresight, it is very risky in practice.

## V. CONCLUSION

We introduced an integrated model for optimizing energy storage bidding in two-settlement electricity markets. Combining a transformer-based model for day-ahead bidding and an LSTM-dynamic programming hybrid model for real-time bidding, we have demonstrated the potential to significantly enhance profit margins in two-settlement electricity markets. Our study addressed the problem of effectively managing energy storage bids in volatile real-time prices. Our experimentation with actual price data from the NYISO market shows that our proposed framework outperforms the sole real-time bidding approach. Specifically, we observed a substantial increase in profit, with improvements of up to 29% when compared to strategies focused solely on the real-time market. Our risk analysis over negative profit days also shows that

the proposed joint bidding algorithm provides less or similar risk than real-time market participation. In summary, using our proposed algorithm, bidding storage in day-ahead and real-time markets provides higher profits than real-time markets without added risk.

For the next steps, we plan to improve our forecasting model with an enhanced customized loss function to reflect the loss that arose from the decision-making process. In addition, we want to implement risk control measurements to reduce the risk that comes from price uncertainties.

## REFERENCES

- [1] Federal Energy Regulatory Commission, "Electric storage participation in markets operated by regional transmission organizations and independent system operators," *Order No. 841, 162 FERC*, vol. 61, p. 127, 2018.
- [2] R. Sioshansi, P. Denholm, and T. Jenkin, "Market and policy barriers to deployment of energy storage," *Economics of Energy & Environmental Policy*, vol. 1, no. 2, pp. 47–64, 2012.
- [3] O. Williams and R. Green, "Electricity storage and market power," *Energy policy*, vol. 164, p. 112872, 2022.
- [4] US Energy Information Association, "Form eia-860 detailed data with previous form data (eia-860a/860b)," 2022.
- [5] M. Arnold and G. Andersson, "Model predictive control of energy storage including uncertain forecasts," in *Power systems computation conference (PSCC), Stockholm, Sweden*, vol. 23. Citeseer, 2011, pp. 24–29.
- [6] K. Abdulla, J. De Hoog, V. Muenzel, F. Suits, K. Steer, A. Wirth, and S. Halgamuge, "Optimal operation of energy storage systems considering forecasts and battery degradation," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 2086–2096, 2016.
- [7] N. Zheng, J. Jaworski, and B. Xu, "Arbitraging variable efficiency energy storage using analytical stochastic dynamic programming," *IEEE Transactions on Power Systems*, vol. 37, no. 6, pp. 4785–4795, 2022.
- [8] H. Wang and B. Zhang, "Energy storage arbitrage in real-time markets via reinforcement learning," in *2018 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2018, pp. 1–5.
- [9] I. Milstein and A. Tishler, "Can price volatility enhance market power? the case of renewable technologies in competitive electricity markets," *Resource and Energy Economics*, vol. 41, pp. 70–90, 2015.
- [10] X. Qin, B. Xu, I. Lestas, Y. Guo, and H. Sun, "The role of electricity market design for energy storage in cost-efficient decarbonization," *Joule*, 2023.
- [11] R. H. Byrne, T. A. Nguyen, D. A. Copp, R. J. Concepcion, B. R. Chalamala, and I. Gyuk, "Opportunities for energy storage in caiso: Day-ahead and real-time market arbitrage," in *2018 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*. IEEE, 2018, pp. 63–68.
- [12] California Independent System Operator (CAISO), "Special report on battery storage," 2023.
- [13] 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, 2015.
- [14] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International journal of forecasting*, vol. 30, no. 4, pp. 1030–1081, 2014.
- [15] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, "Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark," *Applied Energy*, vol. 293, p. 116983, 2021.
- [16] P. L. Skantze and M. Ilic, *Valuation, hedging and speculation in competitive electricity markets: a fundamental approach*. Springer Science & Business Media, 2001, vol. 643.
- [17] S. Deng, *Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes*. University of California Energy Institute Berkeley, 2000.
- [18] B. F. Hobbs, C. B. Metzler, and J.-S. Pang, "Strategic gaming analysis for electric power systems: An mpec approach," *IEEE transactions on power systems*, vol. 15, no. 2, pp. 638–645, 2000.
- [19] P. Visudhiphan and M. D. Ilic, "Dynamic games-based modeling of electricity markets," in *IEEE Power Engineering Society. 1999 Winter Meeting (Cat. No. 99CH36233)*, vol. 1. IEEE, 1999, pp. 274–281.
- [20] M. J. Denton, S. J. Rassenti, V. L. Smith, and S. R. Backerman, "Market power in a deregulated electrical industry," *Decision Support Systems*, vol. 30, no. 3, pp. 357–381, 2001.
- [21] I. P. Panapakidis and A. S. Dagoumas, "Day-ahead electricity price forecasting via the application of artificial neural network based models," *Applied Energy*, vol. 172, pp. 132–151, 2016.
- [22] M. Lehna, F. Scheller, and H. Herwartz, "Forecasting day-ahead electricity prices: A comparison of time series and neural network models taking external regressors into account," *Energy Economics*, vol. 106, p. 105742, 2022.
- [23] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, "Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark," *Applied Energy*, vol. 293, p. 116983, 2021.
- [24] G. McGrath and O. Comstock, "Battery systems on the us power grid are increasingly used to respond to price," *Today in Energy*, 2022.
- [25] R. Sioshansi, P. Denholm, J. Arteaga, S. Awara, S. Bhattacharjee, A. Botterud, W. Cole, A. Cortes, A. De Queiroz, J. DeCarolis et al., "Energy-storage modeling: State-of-the-art and future research directions," *IEEE transactions on power systems*, vol. 37, no. 2, pp. 860–875, 2021.
- [26] M. U. Hashmi, D. Deka, A. Bušić, L. Pereira, and S. Backhaus, "Arbitrage with power factor correction using energy storage," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2693–2703, 2020.
- [27] J. Cao, D. Harrold, Z. Fan, T. Morstyn, D. Healey, and K. Li, "Deep reinforcement learning-based energy storage arbitrage with accurate lithium-ion battery degradation model," *IEEE Transactions on Smart Grid*, vol. 11, no. 5, pp. 4513–4521, 2020.
- [28] Y. Baker, N. Zheng, and B. Xu, "Transferable energy storage bidder," *IEEE Transactions on Power Systems*, 2023.
- [29] R. Li, A. J. Svoboda, and S. S. Oren, "Efficiency impact of convergence bidding in the california electricity market," *Journal of Regulatory Economics*, vol. 48, pp. 245–284, 2015.
- [30] Y. Li, N. Yu, and W. Wang, "Machine learning-driven virtual bidding with electricity market efficiency analysis," *IEEE Transactions on Power Systems*, vol. 37, no. 1, pp. 354–364, 2021.
- [31] D. Krishnamurthy, C. Uckun, Z. Zhou, P. R. Thimmapuram, and A. Botterud, "Energy storage arbitrage under day-ahead and real-time price uncertainty," *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 84–93, 2017.
- [32] E. Y. Bitar, R. Rajagopal, P. P. Khargonekar, K. Poolla, and P. Varaiya, "Bringing wind energy to market," *IEEE Transactions on Power Systems*, vol. 27, no. 3, pp. 1225–1235, 2012.
- [33] B. Xu, Y. Shi, D. S. Kirschen, and B. Zhang, "Optimal battery participation in frequency regulation markets," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 6715–6725, 2018.
- [34] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," *arXiv preprint arXiv:2211.14730*, 2022.
- [35] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 12, 2021, pp. 11 106–11 115.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [37] A. Zeng, M. Chen, L. Zhang, and Q. Xu, "Are transformers effective for time series forecasting?" in *Proceedings of the AAAI conference on artificial intelligence*, vol. 37, no. 9, 2023, pp. 11 121–11 128.
- [38] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," *Advances in Neural Information Processing Systems*, vol. 34, pp. 22 419–22 430, 2021.
- [39] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International journal of forecasting*, vol. 22, no. 4, pp. 679–688, 2006.
- [40] J. H. Zhao, Z. Y. Dong, X. Li, and K. P. Wong, "A framework for electricity price spike analysis with advanced data mining methods," *IEEE Transactions on Power Systems*, vol. 22, no. 1, pp. 376–385, 2007.
- [41] S. Baltaoglu, L. Tong, and Q. Zhao, "Algorithmic bidding for virtual trading in electricity markets," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 535–543, 2018.