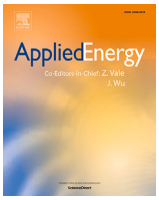
Applied Energy 389 (2025) 125754



Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apen





Optimal intraday power trading for single-price balancing markets: An adaptive risk-averse strategy using mixture models

Robin Bruneel*, Mathijs Schuurmans, Panagiotis Patrinos

Department of Electrical Engineering (ESAT-STADIUS), KU Leuven, Kasteelpark Arenberg 10, Leuven, 3001, Vlaams-Brabant, Belgium

HIGHLIGHTS

- A novel probabilistic mixture imbalance price model to effectively capture the bimodality in the future price distribution.
- A trading strategy that accounts for the uncertainties inherent in the model by using a coherent risk measure.
- The risk tuning parameters are adaptively optimized based on historical performance.
- The impact of the trading position on the system imbalance price is explicitly accounted for.
- The trading strategy is applied to the Belgian energy market.

ARTICLE INFO

Keywords:
Balancing markets
Electricity trading
Mixture models
Risk management
Stochastic optimization

ABSTRACT

Efficient markets are characterized by profit-driven participants continuously refining their positions towards the latest insights. Margins for profit generation are generally small, shaping a difficult landscape for automated trading strategies. This paper introduces a novel intraday power trading strategy tailored for single-price balancing markets. The strategy relies on a strategically devised mixture model to forecast future system imbalance prices and is formulated as a stochastic optimization problem with decision-dependent distributions to address two primary challenges: (i) the impact of trading positions on the system imbalance price and (ii) the uncertainty inherent in the model. The first challenge is tackled by adjusting the model to account for price changes after taking a position. For the second challenge, a coherent risk measure is added to the cost function to take additional uncertainties into account. This paper introduces a methodology to select the tuning parameter of this risk measure adaptively by continuously quantifying the performance of the strategy on a window of recently observed data. The strategy is validated with a simulation on the Belgian electricity market using real-time market data. The adaptive tuning approach leads to higher absolute profits, while also reducing the number of trades.

1. Introduction

The rising share of renewable energy sources is driving rapid evolution in electricity markets. Their inherent volatility and unpredictability, largely due to weather, make it more challenging for transmission system operators (TSOs) to balance consumption with production [1]. Rather than solely relying on the activation of traditional energy reserves, TSOs are increasingly engaging *balance responsible parties* (BRPs) to actively aid in balancing the grid. To facilitate this effort, TSOs utilize imbalance settlement schemes where BRPs are rewarded for maintaining the balance and penalized for creating an imbalance. In this context, accurately predicting future system imbalances becomes a crucial decision-making tool. By anticipating future energy shortages or surpluses in the local balancing market, electricity may be strategically purchased on the intraday market and used to balance the system for profit. This paper delineates an automated trading strategy based on this principle. The strategy comprises two main steps: (i) modelling future system imbalance prices, and (ii) converting these predictions into profitable positions on the intraday electricity market. The intraday market is used to purchase electricity shortly before it closes; therefore, we do not focus on adjusting our position throughout the day or modelling the intraday price itself. This section provides a brief introduction to electricity markets, discusses related work, and outlines the structure and key contributions of this paper.

* Corresponding author.
Email address: robin.bruneel@kuleuven.be (R. Bruneel).

https://doi.org/10.1016/j.apenergy.2025.125754
Received 21 January 2025; Received in revised form 26 February 2025; Accepted 16 March 2025
Available online 25 March 2025
0306-2619/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Nomenclature

Abbreviations

MDP	marginal decremental price
MIP	marginal incremental price
BRP	balance responsible party
TSO	transmission system operator
SI	system imbalance
aFRR	automatic frequency restoration reserve
mFRR	manual frequency restoration reserve
CVaR	conditional value at risk
EVaR	entropic value at risk
pdf	probability density function

Symbols

p_t^{SI}	System imbalance price [€/MWh]
p_t^{MDP}	Downregulaton price [€/MWh]
p_t^{MIP}	Upregulation price [€/MWh]
s_t	System imbalance volume [MW]
f_t^{SI}	pdf of the system imbalance price
f_t^{MDP}	pdf of the downregulation price
f_t^{MIP}	pdf of the upregulation price
π_t	Mixture weight model
n_q	Number of quantile models in the regulation price model
V_{aFRR}	Discretized range of aFRR reserve volumes
V_{mFRR}	Discretized range of mFRR reserve volumes
β	Market reactivity parameter
u	Trade position [MW]
q_t	Intraday price [€/MW]
K_{MDP}	Downregulation price sensitivity [€/MW]
K_{MIP}	Upregulation price sensitivity [€/MW]
α	Risk tuning parameter
$u^*_{\rho_\alpha,t}$	Optimal trade position w.r.t. risk measure ρ_α
N	Window size of historical trades

1.1. The electricity market

Modern electricity markets consist of several components. The largest volumes are traded on the day-ahead market, where transactions occur 1 day before delivery. Nowadays, this market is typically complemented by an intraday market which enables balance responsible parties (BRPs) to adjust their positions in the hours leading up to delivery. On these intraday markets, electricity may be traded through auctions (pay-as-cleared) or via continuous trading (pay-as-bid) [2]. Despite intraday trading efforts, perfectly balancing the actual production and consumption is impossible. The combined imbalance of all BRPs connected to the local grid is referred to as the *system imbalance volume* and to compensate for this imbalance, TSOs activate energy reserves. The costs associated with the activation of these reserves are reflected in a *system imbalance price*, which is a price used to settle energy imbalances among BRPs. Those who help balance the grid are rewarded and those who create the imbalance are penalized. The TSO is responsible for determining this system imbalance price and volume, as well as managing the settlement of energy imbalances.

Various methodologies for calculating this system imbalance price exist. The European Union Agency for the Cooperation of Energy Regulators (ACER) advocates for single-price imbalance settlement schemes where the price for having a positive or negative energy imbalance is the same [3]. This incentivizes market participants to actively aid in balancing the grid. In contrast, dual-price settlement schemes impose additional costs for imbalances making the rewards for contributing smaller than the penalties. This encourages participants to solely focus on minimizing their own imbalance rather than supporting overall grid stability.

In this paper, we specifically focus on the Belgian electricity market, which features a single-price imbalance settlement scheme and is characterized by a high share of renewable energy sources. In the Belgian electricity market, a single system imbalance price p_t^{SI} is settled every quarter hour (denoted by time index t) based on the average system imbalance volume s_t during that interval:

$$p_t^{\text{SI}}(s_t) = \begin{cases} p_t^{\text{MDP}} & \text{if} \quad s_t \geq 0, \\ p_t^{\text{MIP}} & \text{if} \quad s_t < 0. \end{cases} \tag{1}$$

Here, p_t^{MDP} represents the downregulation price or the price associated with the activated downward reserves in case of an energy surplus ($s_t \geq 0$) and p_t^{MIP} the upregulation price or the price associated with the activated upward reserves when there's an energy shortage ($s_t < 0$) [4]. Typically, the downregulation price, p_t^{MDP}, is low or even negative, as it reflects the need to eliminate an energy surplus. In contrast, the upregulation price, p_t^{MIP}, is high because it indicates the need to boost production. As visualized in Fig. 1, this formulation leads to a system imbalance price that continuously alternates between the two regulation prices. The Belgian electricity market is linked to the European intraday exchange where energy is continuously and freely traded on the interconnectors with neighbouring countries, as long as sufficient border capacity is available. The cross-border intraday market closes exactly 1 hour before delivery.

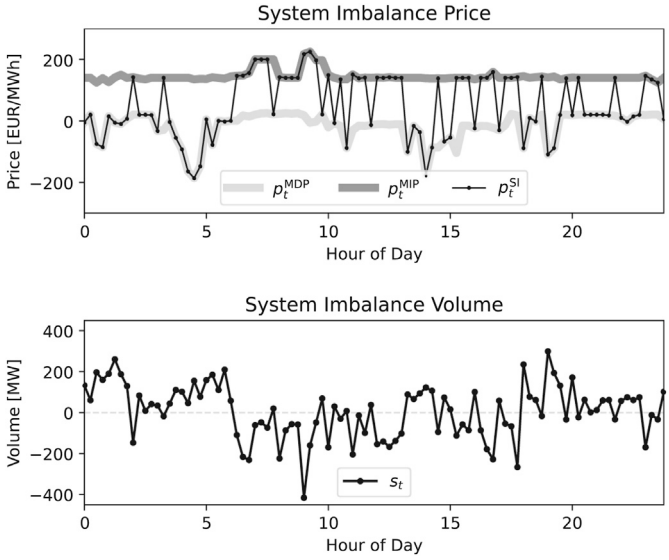


Fig. 1. System imbalance price and volume in Belgium on May 26 2023. Depending on the sign of the system imbalance volume s_t, the system imbalance price p_t^{SI} continuously alternates between the upregulation and downregulation price. This data has been obtained from the Belgian TSO Elia.

1.2. Related work

1.2.1. Modelling future system imbalance prices

The main difficulty in modelling the Belgian system imbalance price resides in its inherent alternating behavior. The prices exhibit a bimodal distribution, which traditional predictive models often fail to capture. In the literature, two methodologies to predict system imbalance prices can be distinguished: (i) those that model the balancing state implicitly, and (ii) those that model it explicitly. Implicitly modelling the balancing

state involves employing a single model to predict balancing prices, and in doing so, assumes that the typical alternating behavior is automatically discovered by the model [5–7]. Conversely, explicitly modelling the balancing state entails the use of multiple models to predict balancing prices. Typically, one model predicts the probability of arriving in a future balancing state, while separate models compute the prices for each balancing state, thereby reducing the overall model complexity.

A popular approach to explicitly model the balancing state is regime-switching Markov models [8–11]. Future balancing states are modelled as a Markov chain, and a distinct price is forecasted for each balancing state. Despite its popularity, a straightforward implementation of regime-switching Markov models in this study reveals two notable issues: (i) static state transition probability matrices are used, independent of other external inputs which might be relevant, and (ii) this method is not suited for longer prediction horizons since the state probabilities converge to a stationary distribution.

Instead of using Markov chains, Koch proposes a logistic regression model directly predicting the probability of a future balancing state [12]. This alternative approach mitigates the need to model all intermediate balancing states, rendering it particularly useful for longer prediction horizons. Unlike the static state transition probability matrices, the logistic regression model can easily incorporate other relevant external inputs.

1.2.2. Risk-averse trading strategy

Since the introduction of intraday trading, studies have demonstrated that trading contributes to the balance of the grid [13–15]. There are several ways to trade on the intraday market; however, in this study, we solely operate as a price-taker on the continuous intraday market just before the cross-border market closes. Given that the distribution of future prices was exactly known, one could simply select a trading strategy that maximizes the expected future return. In practice, however, these distributions are estimated with models that are fit to a finite amount of training data, and are therefore subject to some amount of error. When applied naively, this typically leads to larger trading losses than predicted. A popular approach for taking this model distributional uncertainty – also known as ambiguity – into account, is to replace the expectation with a *coherent risk measure* [16]. Based on the choice of risk measure, larger realizations of the cost will be penalized more, resulting in a more risk-sensitive behavior. Furthermore, coherent risk measures can be interpreted as robust counterparts to the expectation. Instead of minimizing the cost over the predicted distribution, the addition of the risk measure can be viewed as a minimization over a specific set of distributions, known as the *ambiguity set* [16, Thm. 6.4]. As a result, this approach robustifies the trading strategy against misestimations of the predicted cost distribution.

Risk-averse optimization is a decision making tool that has been used in power systems before [17–20]. Typically, the tuning parameters are selected in an offline fashion. Bottiau et al. pointed out that using a machine learning model to dynamically adjust the tuning parameters of the risk measure resulted in better performance when applied to very-short-term energy markets [21]. With online learning, the machine learning model is able to adaptively learn the model mismatch. In recent years, model-free alternatives relying entirely on machine learning methods have gained significant popularity, as they relieve the need for precise modelling of often complex market behaviors. In particular, deep reinforcement learning has emerged as a powerful technique for optimizing bidding strategies and asset management in the electricity market [22]. The main disadvantage of these techniques, however, is that their behavior remains largely unpredictable and difficult to interpret. Instead, we propose a hybrid model, where certain aspects of the market are directly encoded into the model, while others are learned in a data-driven manner, leading to a more efficient use of the available data.

1.3. Contribution and structure

In this paper, an automated intraday electricity trading strategy for single-price balancing markets is proposed. The contributions of this paper are summarized as follows:

- A novel probabilistic mixture model to predict future price distributions in single-price balancing markets (cf. Section 2).
- A risk-averse trading strategy to reduce the number of trades in uncertain scenarios (cf. Section 3). The tuning parameters are optimized in an online fashion based on a window of recently observed data, and the strategy explicitly accounts for the influence of the position on the imbalance price.

In Section 4, the model is trained on data from the Belgian electricity market. After benchmarking against several models, the trading strategy is validated through a simulation using real-time data. Finally, in Section 5, the results are discussed.

Notation

Let $\mathbb{R}$ ($\mathbb{R}_+$) denote the (nonnegative) reals. We denote the n-dimensional probability simplex as $\Delta_n = \{p \in \mathbb{R}^n_+ \mid \sum_{i=1}^n p_i = 1\}$. We introduce the shorthand $[n] = \{1, \dots, n\}$ to denote the n smallest positive natural numbers. We denote the positive part of x by $[x]_+ = \max\{x, 0\}$.

2. Mixture imbalance price model

To properly quantify the trading risks and profits, we are interested in predicting the probability density function (pdf) of the future system imbalance price. As mentioned in Section 1.1, this price continuously alternates between the two regulation prices, causing the future price to be bimodally distributed. We explicitly take this bimodality into account by modelling the pdf of the imbalance price as a mixture distribution [23]. In this context, the pdf $f : \Xi \to \mathbb{R}_+$ is said to be a mixture of a finite set of component densities $f_i(\cdot\,; \theta_i)$ with parameters θ_i, $i \in [d]$. A mixture distribution can be written in the form

$$f(\xi) = \sum_{i=1}^{d} \gamma_i f_i(\xi; \theta_i) \tag{2}$$

where $\gamma \in \Delta_d$ is a vector of nonnegative mixture weights that sum to one. A major advantage of mixture distributions is their ability to represent complex distributions as a combination of simpler component distributions. This approach is particularly effective for modelling multimodal behavior, such as the alternating dynamics of the imbalance price. A popular choice for the component densities is the normal distribution $\mathcal{N}(\mu_i, \sigma_i)$ with $\theta_i = (\mu_i, \sigma_i)$. Mixture models are often used when samples are observed from a mixture of distributions without knowing from which component distribution they originate. The components are considered *latent variables* and, in that case, specialized algorithms, such as Expectation-Maximization are used to jointly estimate the component weights γ_i and the parameters θ_i of the distributions, $i \in [d]$ [23]. Fortunately, in our application, this difficulty does not arise, as it will be straightforward to infer the active component for each observed data point.

Indeed, in the context of this work, mixture distributions are employed to incorporate prior knowledge of the system imbalance price settlement scheme into the model. As defined in Eq. (1), the system imbalance price p_t^{SI} is determined by the sign of the system imbalance volume s_t. Specifically, when the imbalance volume is positive, the system imbalance price equals the downregulation price p_t^{MDP}, whereas if it is negative, it matches the upregulation price p_t^{MIP}. To simplify the modelling task, we model the system imbalance price as a mixture of the two regulation price distributions. Here, the sign of the system imbalance volume acts as an exogenous variable indicating which density component is active. Using Bayes' theorem, the probability density function

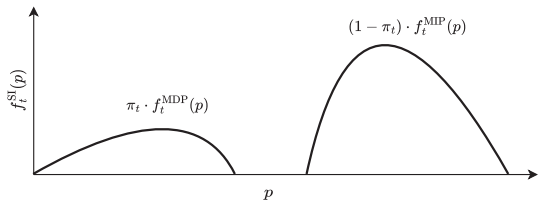


Fig. 2. A sketch of the pdf of the future system imbalance prices f_t^{SI}. Using a mixture model, we relieve the model from having to predict this bimodal distribution directly. Instead, we predict the regulation price distributions f_t^{MDP} and f_t^{MIP} separately, and combine these together with the mixture weight π_t.

of the future system imbalance price $f_t^{\text{SI}}(p)$ at time t is expressed as the following mixture distribution:

$$f_t^{\text{SI}}(p) = \pi_t \cdot f_t^{\text{MDP}}(p) + (1-\pi_t)\cdot f_t^{\text{MIP}}(p) \tag{3}$$

with

$$\begin{aligned} \pi_t &= \mathbb{P}(s_t > 0) \\ f_t^{\text{MDP}}(p) &= f_t^{\text{SI}}(p \mid s_t > 0) \\ f_t^{\text{MIP}}(p) &= f_t^{\text{SI}}(p \mid s_t \le 0) \end{aligned} \tag{4}$$

where $s_t \in S_t = \mathbb{R}$ is a random variable representing the future system imbalance volume, and $f_t^{\text{SI}}(\cdot \mid E)$ denotes the conditional pdf of the price p_t at time t, given event E. Thus, π_t is the probability to obtain a positive system imbalance volume, f_t^{MDP} is the conditional probability density function of the system imbalance price given a positive system imbalance volume, and f_t^{MIP} the conditional probability density function of the system imbalance price given a negative system imbalance volume. To model future system imbalance price distributions, we may choose to model these components separately to reduce the overall complexity of the problem. A visual representation of this mixture distribution is given in Fig. 2. It is important to note that the mixture weight and distributions are time-dependent (e.g. higher probability on energy shortage and more expensive prices during peak hours), yet, only one system imbalance price p_t is observed at each timestep t. That is, we only observe a single sample from the distribution $f_t^{\text{SI}}(p)$, making it practically impossible to estimate the probability density function directly. To address this issue, we assume that the time dependence is introduced through measurable inputs $\boldsymbol{x}_t \in \mathbb{R}^{n_x}$. With this assumption, we can parametrize the mixture distribution as time-invariant functions with parameters $\boldsymbol{\theta}_\pi$, $\boldsymbol{\theta}_{\text{MDP}}$ and $\boldsymbol{\theta}_{\text{MIP}}$:

$$\begin{aligned} \pi_t &= \pi(\boldsymbol{x}_t, \boldsymbol{\theta}_\pi) \\ f_t^{\text{MDP}}(p) &= f^{\text{MDP}}(p; \boldsymbol{x}_t, \boldsymbol{\theta}_{\text{MDP}}) \\ f_t^{\text{MIP}}(p) &= f^{\text{MIP}}(p; \boldsymbol{x}_t, \boldsymbol{\theta}_{\text{MIP}}). \end{aligned}$$

These time-invariant models can be estimated using statistical methods. In the upcoming sections, we elaborate on the choice of these models. The mixture weight is predicted with a logistic regression model, while the regulation price distributions are predicted with non-parametric probabilistic models.

2.1. Mixture weight model

The mixture weight in Eq. (3), denoted by $\pi_t = \mathbb{P}(s_t > 0)$, represents the probability of having a positive system imbalance at timestep t. Various approaches have been proposed for the problem of estimating the probability of such a binary event [24]. One of these approaches is logistic regression, an elementary method used to predict the probability of an event. Logistic regression has already been proven successful in forecasting future balancing states in the German electricity market, motivating its application in this study [12]. Formally, the logistic regression model is a sigmoid function applied to a linear model:

$$\pi_t = \pi(\boldsymbol{x}_t, \boldsymbol{\theta}_\pi) = \sigma(\boldsymbol{w}^\top \boldsymbol{x}_t + w_0) = \frac{\exp(\boldsymbol{w}^\top \boldsymbol{x}_t + w_0)}{1 + \exp(\boldsymbol{w}^\top \boldsymbol{x}_t + w_0)} \tag{5}$$

where the weights $\boldsymbol{\theta}_\pi = (w_0, \boldsymbol{w}) \in \mathbb{R}^{n_x+1}$ are optimized with maximum-likelihood estimation. The simplicity of this model renders it robust to overfitting. Theoretically, this model also requires the log odds of an event to be a linear combination of the input variables [25]. Despite this condition usually not being satisfied, logistic regression still yields meaningful results in practice.

2.2. Regulation price model

The regulation price models should output the distribution of the regulation prices. Models for probability distributions may either be parametric or non-parametric. Parametric probabilistic models are trained to output the parameters of a predefined family of distributions (e.g., the mean and variance of a normal distribution) and are particularly useful if the distribution type is known beforehand. When this is not the case, non-parametric probabilistic models are typically preferable. These models predict multiple quantiles of an output variable to build a discrete approximation of the cumulative density function (cdf), which fully characterizes the distribution of a scalar random variable. Since (i) there is no obvious choice for a parametric family of distributions that describes the complex nature of regulation prices, and (ii) the random variable is scalar, so a fine discretization of its support is computationally feasible, non-parametric models are used in this paper. Before introducing the price model, we briefly explain the energy reserves, a relevant input signal to determine the regulation price distribution.

2.2.1. Energy reserves

To compensate for energy shortages or surpluses, TSOs activate reserves. For example, in Belgium, two types of reserves can be activated: automatic reserves (aFRR) and manual reserves (mFRR). The aFRR reserves are activated almost instantly, but are limited in volume. The mFRR reserves, while slower to activate, are larger in volume. Depending on the size of the imbalance, a combination of these reserves may be activated. In Belgium, the costs of activating a certain amount of these reserves in the future are known beforehand and can be used as a mapping from reserve volume to price. We propose to predict the regulation price by internally modelling the activated volumes of each reserve type based on input features $\boldsymbol{z}_t$. Using the aforementioned price mapping, the corresponding reserve price is then calculated.

2.2.2. Regulation price model

For each reserve type $r \in R = \{\text{aFRR}, \text{mFRR}\}$, we discretize the range of reserve volumes $v \in V_r = \left\{v_r^{(1)}, \dots, v_r^{(n_r)}\right\}$ and use the available data to estimate a discrete probability distribution

$$\boldsymbol{w}(\boldsymbol{z}_t, \boldsymbol{\theta}) = \left(w_{r,v}(\boldsymbol{z}_t, \boldsymbol{\theta})\right)_{v \in V_r, r \in R} \in \Delta_{n_{\text{aFRR}}+n_{\text{mFRR}}}, \tag{6}$$

as a function of the input features $\boldsymbol{z}_t$ and model parameters $\boldsymbol{\theta}$. Here, $w_{r,v}(\boldsymbol{z}_t, \boldsymbol{\theta})$ represents the probability that the imbalance price is equal to the price of a given volume $v \in V_r$ of reserve r, given the input $\boldsymbol{z}_t$. Using the price mappings $o_{t,r} : V_r \to \mathbb{R}$, $r \in R$ at time t, we then predict the expected price for the activation of the reserves as

$$\hat{g}_{\boldsymbol{\theta}}(\boldsymbol{z}_t, \boldsymbol{o}_t) = \langle \boldsymbol{w}(\boldsymbol{z}_t, \boldsymbol{\theta}), \boldsymbol{o}_t \rangle, \tag{7}$$

where $\boldsymbol{o}_t$ is a vector containing the (known) prices for activating a volume of $v \in V_r$ of reserve r at time t:

$$\boldsymbol{o}_t = \left(o_{t,r}(v)\right)_{v \in V_r, r \in R} \in \mathbb{R}^{n_{\text{aFRR}}+n_{\text{mFRR}}}$$

The probability vector $\boldsymbol{w}_t$ in Eq. (6) is modelled with a linear model followed by a softmax activation function. This approach, also known as

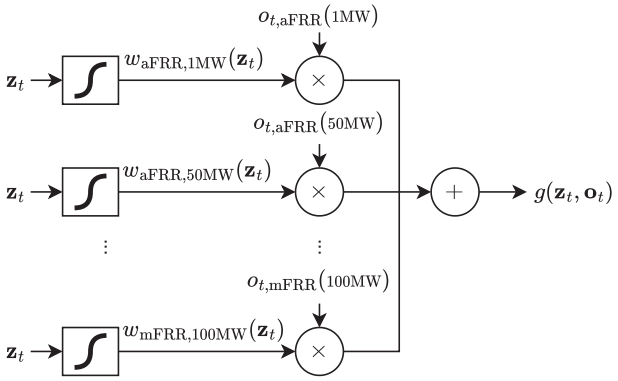


Fig. 3. A visual representation of the price model. The softmax operation ensures that $\boldsymbol{w}(\boldsymbol{z}_t)$ represents a valid probability distribution, with nonnegative elements that sum to one. This probability vector is then multiplied with the order book prices $\boldsymbol{o}_t$ to obtain a prediction of the regulation price $g(\boldsymbol{z}_t, \boldsymbol{o}_t)$.

multinomial logistic regression, is widely used for predicting the probabilities of multiple outcomes [26]. The structure of the model is visually depicted in Fig. 3.

We use the model structure from Eq. (7) to obtain an estimate of the cdf of both the upregulation and the downregulation price. To this end, we first discretize the interval $[0, 1]$ into n_q equally spaced values $\{\tau_1, \dots, \tau_{n_q}\}$. For both regimes $m \in \{\text{MDP}, \text{MIP}\}$, and for each $i = 1, \dots, n_q$, our goal is to find the model parameters $\boldsymbol{\theta}_{m,i}$, such that $g_i^m(\boldsymbol{z}_t, \boldsymbol{o}_t) = \hat{g}_{\boldsymbol{\theta}_{m,i}}(\boldsymbol{z}_t, \boldsymbol{o}_t)$ approximates the τ_i'th quantile of the up- or downregulation price (depending or m). We achieve this using a standard method known as quantile regression, which for a given quantile level $\tau \in [0, 1]$ is equivalent to a classical regression problem from $\boldsymbol{z}_t$ to p_t with a particular loss function, known as the *quantile loss* or *pinball loss* [27], i.e.,

$$L_\tau(e) = \begin{cases} \tau e & \text{if } e \geq 0 \\ (\tau - 1)e & \text{if } e < 0. \end{cases} \tag{8}$$

By minimizing this loss function, the optimal parameters are found. Occasionally, the estimated cdfs may not be monotone increasing with τ. In such cases, we enforce this constraint by reordering the quantiles. Since the quantiles are evenly spaced, the corresponding probability distribution can be represented by a sum of discrete point masses, each with a probability of $1/n_q$:

$$\begin{aligned} f_t^{\text{MDP}}(p) &= \frac{1}{n_q} \sum_{i=1}^{n_q} \delta\left(p - g_i^{\text{MDP}}(\boldsymbol{z}_t, \boldsymbol{o}_t)\right) \\ f_t^{\text{MIP}}(p) &= \frac{1}{n_q} \sum_{i=1}^{n_q} \delta\left(p - g_i^{\text{MIP}}(\boldsymbol{z}_t, \boldsymbol{o}_t)\right) \end{aligned} \tag{9}$$

Here, the function δ represents the Dirac distribution, which places unit probability mass at the origin and is 0 elsewhere [28, Ex. 4.7].

3. Trading strategy

Electricity markets, much like stock markets, enable traders to take both long and short positions. In this paper, a long position involves purchasing energy on the intraday market and settling the resulting imbalance on the balancing market, yielding a profit if the expected difference between the imbalance price and the intraday price is positive. Conversely, a short position is profitable when this price difference is negative. In this study, we consistently act as price-takers on the continuous intraday market. By examining the order books, we can identify the prices at which we can instantly buy or sell a specific amount of energy. In this section, we begin by formally defining how trade positions impact the imbalance price. Thereafter, we introduce the risk-averse trading strategy and present a methodology for adaptively tuning the risk parameters.

3.1. Impact on the balancing market

We assume that intraday trading positions are fully settled on the balancing market. That is, taking an intraday long position results in additional energy being produced or imported into Belgium, which naturally influences both the system imbalance volume and price. In this section, we present a simple model to take this effect into account.

For a long position of size u, we assume that the system imbalance volume will increase with a volume βu, where $\beta \in [0, 1]$ is a tuning parameter that models the market reactivity. In an open-loop system, where BRPs do not react to the actions of others, β would be equal to one. However, in reality, the activations of one party would lead to fewer activations in the same direction by others ($\beta < 1$). Accurately estimating β is highly challenging, as it requires placing numerous trades as a trader and detecting changes in the behavior of the system imbalance volume. For simplicity, we therefore consider it a tuning parameter of the method. To ensure meaningful simulations in our case study, we consider the most pessimistic scenario where our traded volume has the strongest negative impact on the market price by setting $\beta = 1$. Indeed, replacing the imbalance volume s_t by $s_t + \beta u$, the system imbalance price according to Eq. (1) becomes

$$p_t^{\text{SI}}(s_t + \beta u) = \begin{cases} p_t^{\text{MDP}}(s_t + \beta u) & \text{if } \quad s_t + \beta u \geq 0 \\ p_t^{\text{MIP}}(s_t + \beta u) & \text{if } \quad s_t + \beta u < 0 \end{cases} \tag{10}$$

If u is sufficiently large, it will not only impact the imbalance volumes, but the regulation prices as well. Based on empirical observations (cf. Fig. 4), we propose to model the impact of our position on the regulation prices using a piecewise-linear approximation

$$p_t^{\text{SI}}(s_t + \beta u) = \begin{cases} p_t^{\text{MDP}}(s_t) - K_{\text{MDP}}\beta u & \text{if } \quad s_t + \beta u \geq 0 \\ p_t^{\text{MIP}}(s_t) - K_{\text{MIP}}\beta u & \text{if } \quad s_t + \beta u < 0 \end{cases}. \tag{11}$$

Here, the constants K_{MDP} and K_{MIP} represent the sensitivity of the regulation price to the system imbalance volume. The larger these constants, the lower the system imbalance price will be and therefore also the profit of the trading strategy. We estimate these sensitivities by fitting the slope

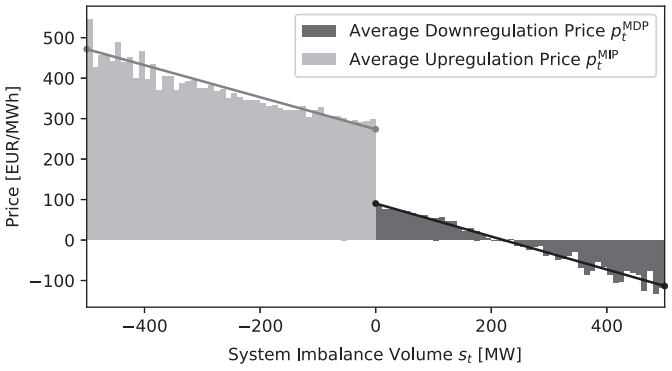


Fig. 4. The average regulation price for different system imbalance volumes on the Belgian balancing market between May 2021 and May 2023. The sensitivity of the regulation prices to the system imbalance volume is approximated by fitting two linear slopes. This seems like a reasonable approximation.

between the system imbalance volume and regulation price with a least squares regression:

$$K_{\text{MDP}} = \frac{\sum_{t\in T_{\text{MDP}}}(s_t - \bar{s}_t^{\text{MDP}})(p_t^{\text{MDP}} - \bar{p}_t^{\text{MDP}})}{\sum_{t\in T_{\text{MDP}}}(s_t - \bar{s}_t^{\text{MDP}})^2}$$
$$K_{\text{MIP}} = \frac{\sum_{t\in T_{\text{MIP}}}(s_t - \bar{s}_t^{\text{MIP}})(p_t^{\text{MIP}} - \bar{p}_t^{\text{MIP}})}{\sum_{t\in T_{\text{MIP}}}(s_t - \bar{s}_t^{\text{MIP}})^2} \tag{12}$$

where T_{MDP} and T_{MIP} are the sets of timesteps where the system imbalance volume is positive and negative, respectively. The terms $\bar{s}_t^x$ and $\bar{p}_t^x$ are the mean system imbalance volume and price over T_x. Fig. 4 visualizes the price sensitivity of the regulation prices to the system imbalance volume on the training set of the case study, where we obtain $K_{\text{MDP}} = 0.40$ €/MW and $K_{\text{MIP}} = 0.41$ €/MW. This implies that, on average, an increase of 1 MW in the system imbalance volume leads to a decrease of €0.41 in the upregulation price.

The modified imbalance price from Eq. (11) is used to build a position-aware model $f_{t,u}^{\text{SI}}$ that predicts the imbalance price distribution after placing a trade of size u. The modified mixture model now has to predict the probability that the system imbalance exceeds $-\beta u$:

$$\pi_t(u) = \mathbb{P}(s_t + \beta u > 0) = \mathbb{P}(s_t > -\beta u) \tag{13}$$

In this paper, the input u is integrated as an additional feature in the logistic regression model from Eq. (5). To achieve this, the training set is augmented with artificial trade positions u which are sampled from a uniform distribution over the interval $U = [0, u_{\max}]$. The model is then trained to predict the probability that the system imbalance volume exceeds $-\beta u$. For the regulation price models, we simply train the quantile models using the observed regulation prices $p_t^{\text{MDP}}(s_t)$, $p_t^{\text{MIP}}(s_t)$ and then subtract $K\beta u$ from each predicted quantile. The following expression represents the position-adjusted pdf of the system imbalance price:

$$f_{t,u}^{\text{SI}}(p) = \pi_t(u)\cdot f_{t,u}^{\text{MDP}}(p) + (1-\pi_t(u))\cdot f_{t,u}^{\text{MIP}}(p) \tag{14}$$

with

$$\begin{aligned}\pi_t(u) &= \mathbb{P}(s_t > -\beta u) = \pi\big((\boldsymbol{x}_t, u), \theta_\pi\big)\\ f_{t,u}^{\text{MDP}}(p) &= f_t^{\text{MDP}}(p + K_{\text{MDP}}\beta u)\\ f_{t,u}^{\text{MIP}}(p) &= f_t^{\text{MIP}}(p + K_{\text{MIP}}\beta u)\end{aligned} \tag{15}$$

This formulation may also be used for a short position by substituting u with a negative value and training a similar model.

3.2. Risk-averse trading strategy

The ultimate goal of an energy trader is to find the trade position in the electricity market that maximizes the expected profit, or equivalently, the position that minimizes the expected loss. To simplify the notation of the risk measures, we continue with minimization problems. The loss $Z_{p,t}(u)$ for a long strategy with position size $u \in U \subset \mathbb{R}_+$ and delivery time t may be expressed as:

$$Z_{p,t}(u) = \big(q_t(u) - p\big)u, \quad p \in P_t$$

where the random variable $p \in P_t$ denotes the future system imbalance price and $q_t(u)$ denotes the current (known) intraday price[1] to take a position of size u. The optimal position u_t^* that minimizes the expected loss then equals

$$u_t^* = \operatorname*{argmin}_{u\in U} \mathbb{E}_p[Z_{p,t}(u)], \tag{16}$$

where the expectation is taken with respect to the random variable p. While $q_t(u)$ is known exactly from the order book, the true distribution of the future imbalance price p is not. Using the estimated pdf $f_{t,u}^{\text{SI}}$, from Eq. (14), we may replace the expectation in Eq. (16) by an approximation

$$u_t^* \approx \operatorname*{argmin}_{u\in U} \mathbb{E}_{p\sim f_{t,u}^{\text{SI}}}[Z_{p,t}(u)].$$

However, with this methodology, the strategy ignores the fact that the price distribution was replaced by a prediction. This may lead to unexpected returns of the strategy when the predicted prices deviate a lot from the true distribution. Instead, we introduce risk-awareness by replacing the expectation with a coherent risk measure ρ_α:

$$u_{\rho_\alpha,t}^* = \operatorname*{argmin}_{u\in U} \rho_{\alpha,\, p\sim f_{t,u}^{\text{SI}}}[Z_{p,t}(u)]. \tag{17}$$

The tuning parameter α controls the level of risk aversion, allowing the trader to balance the desire for higher profits with the need to mitigate risk. An interesting property of coherent risk measures is that they can be viewed as a worst-case expected value over a specific set of distributions, known as the *ambiguity set* [16, Thm. 6.4]. That is, instead of minimizing the expected loss over a single price distribution f_t^{SI}, we consider a set of similar distributions $\mathcal{A}_{t,u}(\alpha)$ to account for potential inaccuracies in the cost:

$$u_{\rho_\alpha,t}^* = \operatorname*{argmin}_{u\in U} \max_{f\in\mathcal{A}_{t,u}(\alpha)} \mathbb{E}_{p\sim f}[Z_{p,t}(u)]$$

By taking the worst-case realization of the expected loss in this ambiguity set, the trading strategy is made robust against all realizations within this set. The size of this ambiguity set varies with the choice of α. A larger ambiguity set leads to a more conservative trading strategy that is less sensitive to potential modelling errors, while a smaller ambiguity set leads to a more aggressive trading strategy being more vulnerable to these errors. In this paper, we discuss two commonly used coherent risk measures, namely conditional value at risk and entropic value at risk [29,30].

The conditional value-at-risk (CVaR), also known as the expected shortfall, equals the expected loss in the worst α cases. CVaR may be expressed as a convex optimization problem:

$$\text{CVaR}_{1-\alpha}(Z) = \begin{cases}\operatorname{esssup} Z & \text{if } \alpha = 0\\ \inf_s\left\{s + \frac{1}{\alpha}\mathbb{E}[Z-s]_+\right\} & \text{otherwise}\end{cases} \tag{18}$$

where $\alpha \in [0,1]$. For CVaR, this problem can be solved analytically to obtain the average loss in the worst α cases. By considering only the worst cases, a higher weight is given to the losses making the trading strategy more risk sensitive. This risk measure converges to the expectation (risk-seeking) when α approaches 1. Conversely, when α approaches 0, this risk measure converges to the maximum value of Z or the largest loss (risk-averse).

The entropic value-at-risk (EVaR) is defined as the following convex optimization problem:

$$\text{EVaR}_{1-\alpha}(Z) = \begin{cases}\operatorname{esssup} Z & \text{if } \alpha = 0\\ \inf_{s>0}\left\{\frac{1}{s}\ln\left(\frac{\mathbb{E}[e^{sZ}]}{\alpha}\right)\right\} & \text{otherwise}\end{cases} \tag{19}$$

Similar to CVaR, EVaR also interpolates between the expectation and the maximum value as α decreases from 1 to 0. An interesting property of coherent risk measures is their positive homogeneity and translation invariance [30], i.e.,

$$\rho_\alpha[aZ + b] = a\,\rho_\alpha[Z] + b, \quad \forall a \in \mathbb{R}_+, \forall b \in \mathbb{R}.$$

Using this identity, Eq. (17) can be reformulated as

$$\begin{aligned}u_{\rho_\alpha,t}^* &= \operatorname*{argmin}_{u\in U} \rho_{\alpha,\, p\sim f_{t,u}^{\text{SI}}}[Z_{p,t}(u)]\\ &= \operatorname*{argmin}_{u\in U} \rho_{\alpha,\, p\sim f_{t,u}^{\text{SI}}}[(q_t(u) - p)u]\\ &= \operatorname*{argmin}_{u\in U} \underbrace{\Big(q_t(u) + \rho_{\alpha,\, p\sim f_{t,u}^{\text{SI}}}[-p]\Big)}_{=\varphi(u)} u\,.\end{aligned} \tag{20}$$

[1] The intraday price may arise from filling multiple orders at different price levels. In that case, the intraday energy cost $q_t(u)u$ is a piecewise-linear and monotone-increasing function of the position size u, as it is derived from the integration of piecewise-constant, increasing prices.

The reference distribution $f^{\mathrm{SI}}_{t,u}$ depends on the decision variable u, making this problem an instance of stochastic optimization with decision-dependent distributions. This class of problems is generally difficult to solve even for risk measures as simple as the expectation [31–33]. Indeed, without any assumptions on the risk measure ρ, the cost function φ may be non-smooth and non-convex. However, there are several cases of practical interest where the problem can be solved more efficiently:

Finite action space. Typically, the action space U is discretized into multiples of the *smallest tradable unit*, making this set finite and countable, i.e., $U = \{u_1, \dots, u_n\}$. In this case, Eq. (20) can be solved exactly by evaluating the cost in each u_i and selecting the smallest value.

Continuous cost. Alternatively, if φ is continuous and U is a closed interval, problem (20) can be solved with a line-search procedure.

Specific risk measures. For particular choices of the risk measure ρ, even more efficient methods to find the solution may exist. For instance, when the expected value is chosen as a risk measure, the optimization problem (20) can be shown to be both smooth and convex under typical conditions (cf. Appendix A), which allows the application of more efficient solution methods, such as Newton's method.

In the case study (cf. Section 4), the action space U is finite and can be effortlessly enumerated to identify the optimal trade position.

3.3. Adaptive risk parameters

Risk measures introduce a tradeoff between maximizing potential profits and minimizing associated risks. For CVaR and EVaR, this tradeoff is controlled by the tuning parameter α. In very risk-averse scenarios (e.g., safety in autonomous driving [34]), a low, constant α is chosen to have high robustness. However, in trading, this would lead to very conservative behavior where the strategy does not take any positions at all. Instead, we may afford to take larger risks at times when we are more confident about our predictions. To this end, we present a methodology for tuning the risk parameter α based on a window of historical trades. Concretely, we choose the next α_{t+1} as the value that would have resulted in the smallest absolute loss (in hindsight) over the last N trades:

$$\alpha_{t+1} = \operatorname*{argmin}_{\alpha\in[0,1]} \frac{1}{N}\sum_{i=0}^{N-1}\left(q_{t-i}\left(u^*_{\rho_\alpha,t-i}\right) - p^{\mathrm{SI}}_{t-i}\right)u^*_{\rho_\alpha,t-i} \tag{21}$$

Here, p^{SI}_t equals the actual system imbalance price at time t, $q_t(u)$ the intraday long price and $u^*_{\rho_\alpha,t}$ the risk-averse action that would have been chosen at time t according to Eq. (17) with risk measure ρ_α. Using this methodology, the trading strategy will converge to a more risk-averse choice for α when it has consistently generated losses over the past N time steps. Conversely, a more risk-seeking α will be chosen when consistent profits were made. This cost function, however, is completely non-convex in α as depicted in Fig. 5. Evaluating the cost in a single α requires the computation of N optimal trade positions, making it computationally expensive. However, since the optimization problem is solved on a rolling window of historical trades, previously computed solutions can be reused to significantly speed up subsequent evaluations. By discretizing the domain into a fine-grained grid and computing the optimal trade position at each grid point at every timestep, a sufficiently accurate solution can be obtained while also minimizing computational overhead. Although we have chosen to minimize absolute losses in this paper, this methodology can easily be adapted to other performance metrics.

4. Case study

In this section, the predictions generated by the previously developed mixture model are used to simulate trading positions on the Belgian electricity market. The strategy is designed to take quarter-hourly positions on the Belgian intraday electricity market 5 minutes before the cross-border intraday market closes (which is exactly 1 hour before delivery). To simulate these positions, real-time order books are used from the Nordpool intraday exchange. The smallest allowed trade unit on the intraday exchange is 0.1 MW and a maximum volume of 5 MW is a realistic volume for a small energy trader in Belgium. Hence, the action space is set to $U = \{0\,\mathrm{MW}, 0.1\,\mathrm{MW}, \dots, 5\,\mathrm{MW}\}$. The market influence, as outlined in Section 3.1, is used to calculate the modified imbalance prices, which are then employed for both training the model and for the profit calculation in the simulation. A 2-year training set spanning from May 1, 2021 to May 1, 2023 is selected to provide the model with sufficient data, while also ensuring that it remains up to date with recent market conditions. Testing is conducted from May 1, 2023 to May 1, 2024. To mitigate seasonal biases, both the training and test sets are chosen as multiples of a year. With renewable energy production increasing annually, market conditions evolve, making larger training sets less representative of current trends. The inputs to train this model are obtained from either the website of the Belgian TSO Elia or the intraday exchange Nordpool. In this section, we begin by training the statistical model, which is then benchmarked against several models. Finally, the results of the adaptive trading strategy are presented.

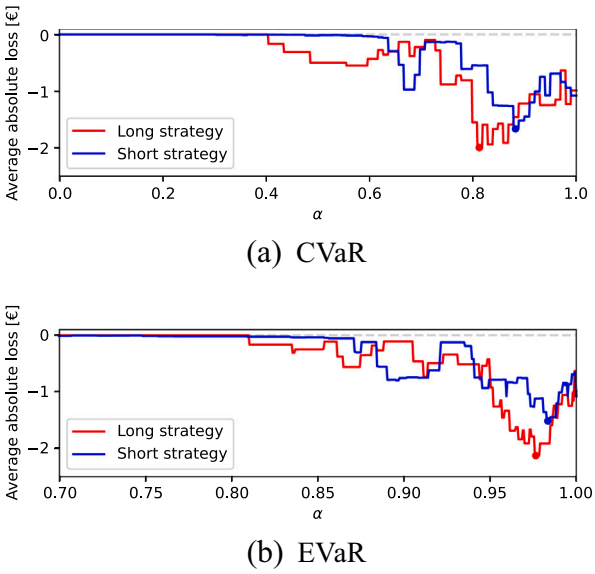


Fig. 5. The average historical absolute loss of the strategy over the past $N = 500$ trading periods for different values of α (cost function of Eq. 21) on the test set of the case study on 2023-06-26 18:30. Negative values indicate profits. For both strategies, the naive approach ($\alpha = 1$) results in suboptimal returns. The adaptive trading strategy chooses α_{t+1} as the minimizer of this function.

4.1. Statistical model

In Fig. 6, the imbalance prices of the test and training sets are visualized. There are no clear price trends, making the training set a good representation of the test set. To model the Belgian imbalance price, we employ the mixture model as proposed in Section 2. For the mixture weight model, the inputs are the same as those proposed by Koch's imbalance model [12]. The latter was built for the German electricity market which is very similar to the Belgian one. These inputs include historical system imbalance volumes from the last hour (until 1 hour before delivery), the quarter hour of the day (one-hot encoded into 96 binary variables), the difference between the intraday and dayahead forecast of the national solar, wind and load, the deviation of the intraday forecast

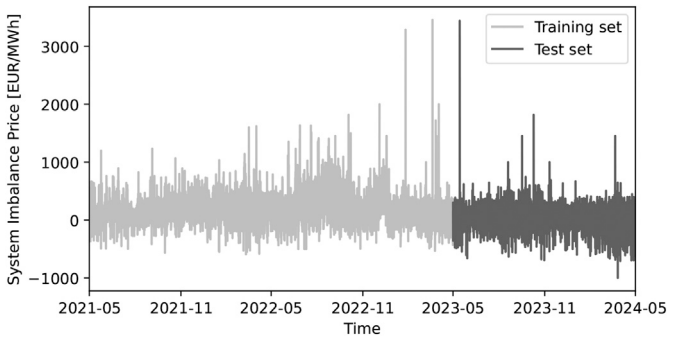


Fig. 6. The imbalance prices for both the training and test sets show no clear trends, indicating that the test set serves as a representative sample of the training data. This dataset includes several significant price spikes, primarily caused by nuclear power outages. These spikes were intentionally kept in the dataset to ensure that the model accounts for these extreme price events.

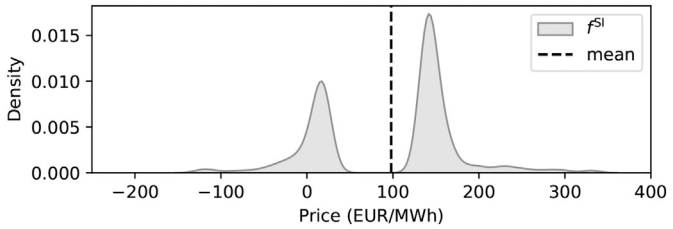


Fig. 7. The predicted system imbalance price probability density function f_t^{SI} on 2023-06-26 18:30 (visualized using kernel density estimation). Because of a higher probability of an energy shortage, the density of the upregulation prices is increased by the mixture weight.

of solar, wind and load from the hourly mean and the difference between the intraday and dayahead price.

For the regulation price model, we have to choose the inputs $\boldsymbol{z}_t$ and fixed sets of reserve volumes V_{aFRR} and V_{mFRR}. In our experiments, the fixed sets of reserve volumes are

$$V_{\text{aFRR}} = \{1\text{ MW}, 50\text{ MW}, 100\text{ MW}, 150\text{ MW}, 200\text{ MW}\}$$
$$V_{\text{mFRR}} = \{1\text{ MW}, 100\text{ MW}, 200\text{ MW}, 300\text{ MW}, 500\text{ MW}, 700\text{ MW}\}$$

These volumes are chosen because Elia typically reserves 150 MW of aFRR capacity and 700 MW of mFRR capacity for the next day [35]. Hence, this discretization serves as a realistic approximation of the actual range of reserve volumes. We choose the output of the regular mixture weight model as the input feature $(\boldsymbol{z}_t)_{t\in\mathbb{N}}$ of the regulation price model (cf. Section 2.2.2). This value summarizes the market and is strongly related to the amount of reserves that will be activated. For instance, if the probability on an energy shortage is very large, then there is also a high probability that the shortage itself will be large and that many reserves have to be activated. To prevent output data leaking from the mixture weight model to the regulation price model, we apply k-fold crossvalidation to prepare the input for the price model. For both regulation prices, we train 100 different models to predict $n_q = 100$ evenly distributed quantiles of that regulation price. The authors emphasize the importance of selecting a sufficiently large number of quantile models to minimize discretization errors and accurately represent extreme price events in the predicted distribution. A total of 100 quantiles was chosen as an optimal balance between computational efficiency and accuracy. Increasing the number of quantiles beyond this threshold does not yield any differences in performance. An example of the forecasted regulation price distributions weighted with the mixture weight is visualized in Fig. 7.

4.2. Probabilistic model benchmark

To evaluate the performance of the previously proposed price model, it is benchmarked against several existing alternatives. For completeness, we compare it against both implicit and explicit models (cf. Section 1.2.1). Implicit models are designed to predict the system imbalance price directly and in doing so, assume that the model implicitly learns to predict the balancing state. On the other hand, explicit models predict the regulation price distributions separately and then combine them by weighting each distribution with the predicted probability of its corresponding balancing state.

Explicit benchmark models. As mentioned in the introduction, another popular approach for explicitly modelling the imbalance price is regime-switching Markov models (RSMMs) [8–11]. Markov chains are commonly used to model stochastic processes in which future states depend on the current state. RSMMs extend this concept by assigning a different predictive model to each Markov state. By weighting the outputs of these predictive models according to their predicted state probabilities, RSMMs are able to forecast prices in the future. In this study, the Markov states correspond to the two balancing states (i.e., $s_t \geq 0$ and $s_t < 0$) and the price models are chosen to be the same regulation price models as those proposed by our model (i.e., f_t^{MDP} and f_t^{MIP}). The key difference between our proposed model and RSMMs lies in how future states are predicted. Our mixture weight model directly estimates the probability to be in a future balancing state, whereas an RSMM predicts each intermediate state sequentially using Markov-chain transition matrices. Since trading positions must be taken at least 1 hour before delivery, RSMMs are used to predict the balancing state probabilities five quarter-hour intervals in the future. In the benchmark comparison, we discuss two different variants:

1. *Static RSMM:* This RSMM employs a static transition matrix, which is determined by computing the empirical probabilities of transitioning from one balancing state to another in the next timestep based on historical training data.
2. *Dynamic RSMM:* In reality, the transition probabilities between balancing states are not constant and can be influenced by external factors. To better capture these dynamics, we introduce a more flexible approach where the transition probabilities are conditioned on exogenous input variables (production and load derived variables, quarter of the day as one-hot input, and difference between intraday and dayahead price). To estimate these transition probabilities, we employ two logistic regression models, a method that has been proven successful in previous studies [8].

Implicit benchmark models. We also present three different non-parametric implicit models: (1) linear quantile regression, (2) artificial neural networks (ANN) with one hidden layer, and (3) random forest quantile regression. Here, all inputs (previous system imbalance volumes, production and load derived variables, quarter of the day as one-hot input, difference between intraday and dayahead price and the order book prices of the fixed set of reserve volumes $\mathbf{o}_t$) are fed into the model at once.

1. *Linear Model:* The linear quantile regression model is the simplest quantile regression model. The output is a linear combination of the input variables. Quantiles are predicted by minimizing the quantile loss as proposed in Eq. (8). By training separate models for the different quantiles of the system imbalance price, an estimation of its probability density function is obtained.
2. *ANN:* Neural network quantile regression models extend traditional linear quantile regression models by replacing the internal

Table 1
Losses of different models on the test set. The mixture model outperforms the other models on every metric. The explicit models achieve lower MAE and predict denser distributions. However, these metrics do not necessarily translate into better RSME and CRPS scores.

Model	RMSE	MAE	Std	CRPS
Our Mixture Model	**139.18**	**96.09**	**108.45**	**67.33**
Static RSMM	144.25	100.03	114.00	70.23
Dynamic RSMM	141.29	98.75	110.56	68.45
Linear	139.36	101.75	132.45	70.24
ANN	139.32	104.77	130.36	69.54
Random Forest	142.45	102.24	148.77	71.66

model with a neural network. Unlike the linear counterpart, neural network quantile regression can handle complex, nonlinear relationships. Different configurations were tested and a model with one hidden layer and three neurons yielded the best performance. The neurons use the $tanh$ activation function.

3. *Random Forest:* Random forest is a popular ensemble learning method for regression and classification. By combining a multitude of decision trees, random forests are able to handle nonlinearity and complex interactions between variables. Meinshausen generalized random forests to not only forecast the mean of an output variable, but also forecast conditional quantiles [36]. This generalization is also used in this paper. The hyperparameters were chosen through grid search.

Typically, probabilistic forecasts are evaluated based on their *reliability* and *sharpness* [37]. Reliable models predict quantiles that are statistically correct with their observations, while the sharpness refers to the concentration of the predictive distributions. More concentrated distributions signify a higher certainty of the model. In this study, the sharpness is assessed by calculating the standard deviation (Std) of the predicted distribution. To jointly measure sharpness and reliability, we use the *continuous ranked probability score (CRPS)* [38], a widely used metric that quantifies the error between the predicted cdf and the true cdf. A lower CRPS equals a better probabilistic forecast. We also calculate the root mean square error (RMSE) between the expected value (or mean) of the forecasted distribution and the true observed system imbalance price and the mean absolute error (MAE) between the predicted median of the forecasted distribution and the imbalance price.

The performance metrics on the unseen test set are listed in Table 1. A snapshot of the forecasted distributions at one timestep is visualized in Fig. 8 and the probabilistic forecasts during a whole day are depicted in Fig. 9. Our proposed mixture model outperforms the other models on every metric. In this benchmark, the implicit models struggle to capture the bimodal nature of the imbalance price, leading to wider predicted distributions and relatively higher CRPS scores. Despite this, their mean predictions are quite accurate, as reflected in their relatively low RMSE. Unsurprisingly, the dynamic RSMM outperforms the simpler static RSMM. However, both RSMM variants still fall short of our proposed mixture model. Their Markov-based structure limits their ability to accurately forecast balancing states over longer horizons. Having to model all intermediate states increases the risk of accumulating prediction errors.

4.3. Risk-averse trading strategy

In this section, the adaptive risk-averse trading strategy is simulated on the Belgian electricity market. To ensure reliability, the order books from the continuous intraday market and inputs for the price model were recorded at the designated trading times, which is always 1 hour and 5 minutes before delivery. As mentioned in Section 3.3, we compute the optimal trade positions with an enumeration over the 51 volumes in U. To calibrate the risk tuning parameter α_{t+1} as outlined in Eq. (21), a window of $N = 500$ previous trades is utilized. The optimal parameter value is chosen by evaluating the expected profits for 200 distinct values of α. We found that choosing a finer grained grid did not lead to better results. On an Apple Macbook Pro M1, computing the optimal α_t for a single timestep takes 36ms for CVaR and 40ms for EVaR. This strategy is allowed to take both long and short positions on the electricity market and tracks a separate value for α_t for both position types. In the simulation, we consider the worst-case scenario where the market impact factor β from Section 3.1 is assumed to be 1.

The performance of the trading strategy on the test set is detailed in Table 2 and Fig. 10. The adaptive trading strategies consistently outperform their static counterparts. Moreover, the impact of misestimating the market influence β is analyzed by generating imbalance prices with a market impact different from the one assumed by the trading strategy. The findings, summarized in Table 3, highlight that the adaptive trading strategy is relatively robust to misestimations of the market impact.

5. Discussion

The proposed trading strategy is designed to be applied in a real-world setting. The integration of feedback in the adaptive strategy ensures that it remains robust to changing market conditions, which is a crucial element for successful trading in financial markets. When comparing the results in Table 2, the difference in profits between the adaptive EVaR and CVaR trading strategies is relatively small. This outcome aligns with our expectations, as the adaptive risk parameter selection seeks to filter out the same historically unprofitable trades and thereby approximates the same risk ambiguity set.

Analyzing the profits in Fig. 10, it clearly appears more challenging to achieve consistent profitability during the winter months compared to the summer months. This difference may be attributed to the higher levels of renewable energy production in the summer months, making the system imbalance price more predictable with the inputs considered in this study.

On May 22, 2024, the Belgian TSO Elia introduced a deadband mechanism to the imbalance price. Under this framework, the imbalance price is set to a predetermined mid-price when the imbalance volume lies within a specified deadband. To incorporate this deadband into the strategy, adaptations to the model are required to include the new balancing state. Furthermore, in November 2024, the Belgian TSO will transition to PICASSO, a platform for sharing reserves among TSOs [39]. Under this framework, the prices of shared reserves are only disclosed after delivery, in contrast to the local bids of which prices are known in advance. As a consequence, a more sophisticated model will be required to also predict future shared reserve prices.

In theory, this strategy may be extended to dual-price markets that include ex-post trading. In these markets, energy contracts are traded after delivery, enabling participants to settle energy imbalances at more favorable prices. Implementing this extension would require a model capable of predicting ex-post market prices to take advantage of this system. A notable example of an ex-post market is the Energy Trading Platform Amsterdam (ETPA), which operates within the Dutch energy market [40].

In addition to the simulation, this exact strategy has recently been implemented by a Belgian BRP and already shows promising results.

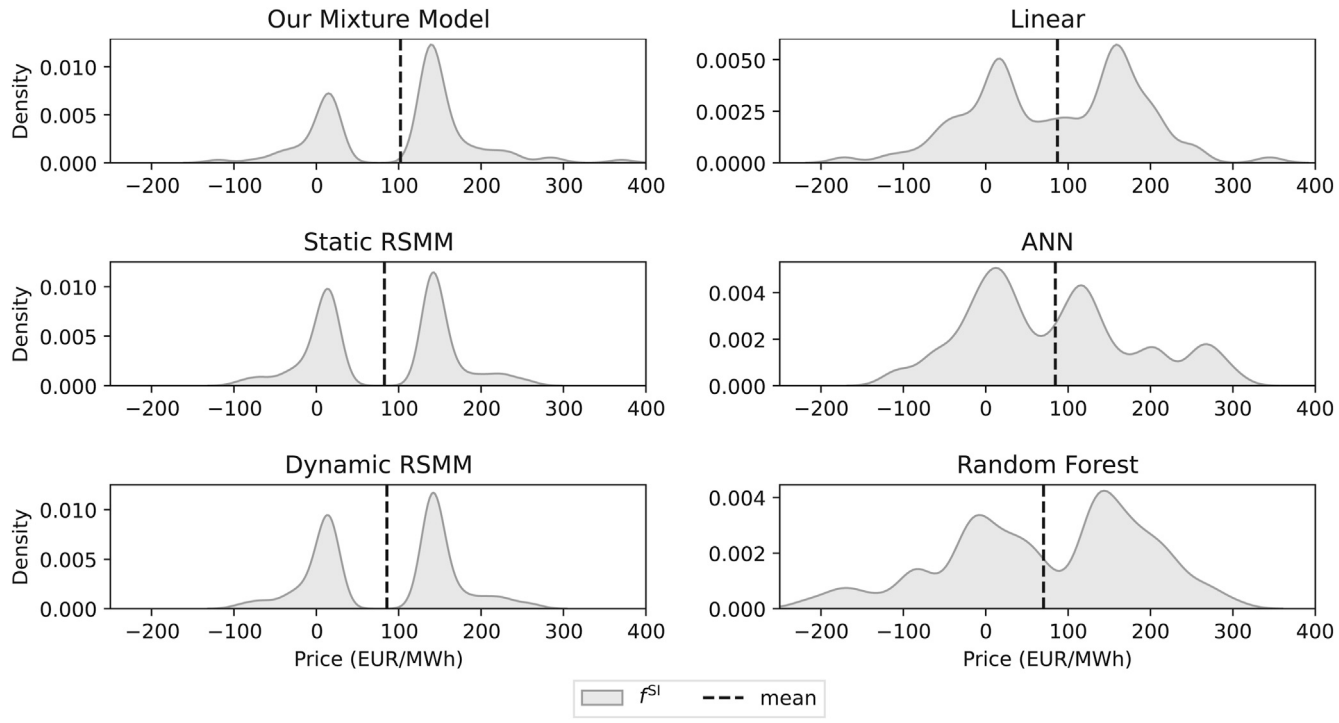


Fig. 8. Predicted system imbalance price probability density function of the different models on 2023-05-26 at 18:30 (visualized using kernel density estimation). The explicit models (left) make a clear distinction between the two regulation prices, while the implicit models (right) predict a less pronounced bimodal distribution.

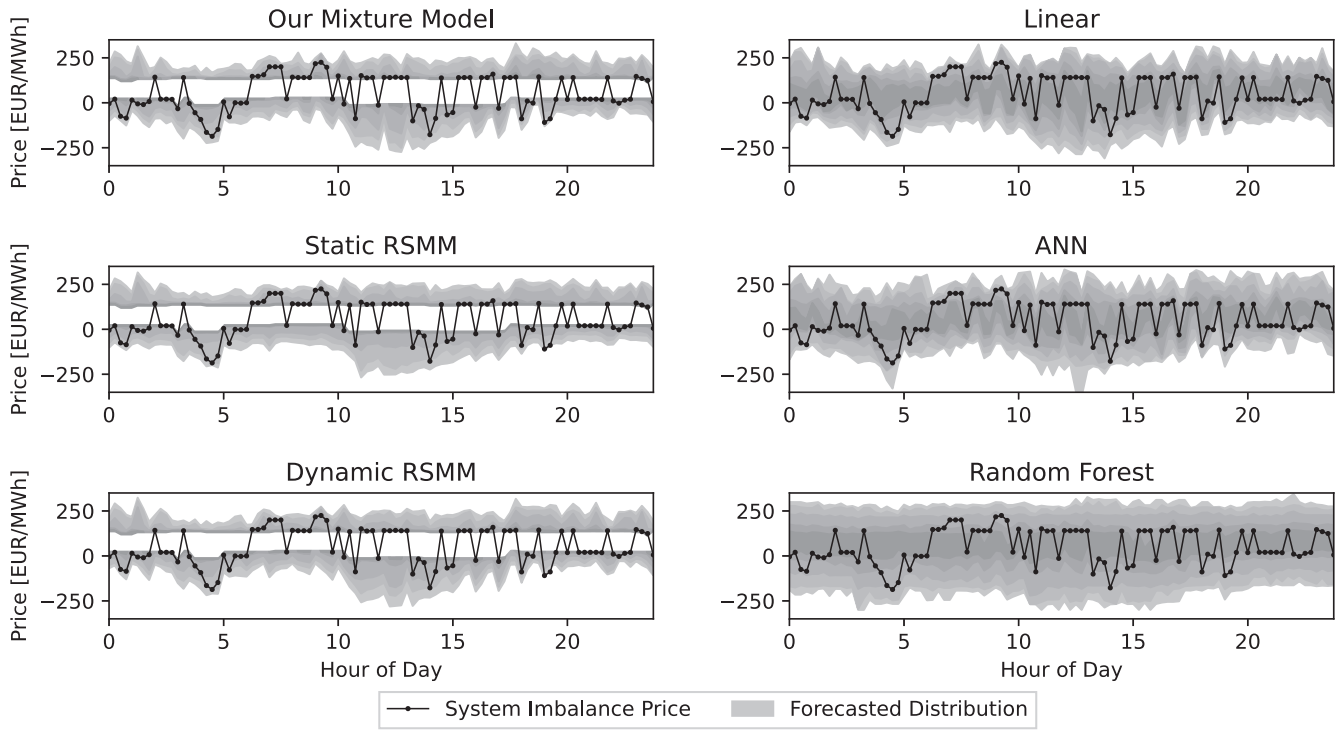


Fig. 9. Forecasted distributions on May 26, 2023. The alternating behavior of the system imbalance price (black line) is clearly visible. The explicit models (left) predict denser distributions compared to the implicit models (right). Furthermore, the structural design of the explicit models enables them to account for the gap between the regulation prices.

Table 2
Performance of the trading strategy on the test set using different risk measures. The adaptive strategies achieve higher absolute profits with fewer trades. As the risk parameter α decreases, the profit per trade increases, while the total number of trades decreases. The adaptive trading strategy automatically searches for the most optimal configuration.

Measure	α	Profit [€]	Trades [MWh]	Profit per trade [€/MWh]
Expectation		101k	28,595	3.48
CVaR	0.95	139k	22,457	6.16
	0.9	154k	17,166	8.96
	0.8	146k	10,002	14.55
	Adaptive	**171k**	14,402	11.87
EVaR	0.995	149k	22,105	6.72
	0.98	157k	16,328	9.60
	0.95	148k	9617	15.39
	Adaptive	**172k**	14,235	12.08

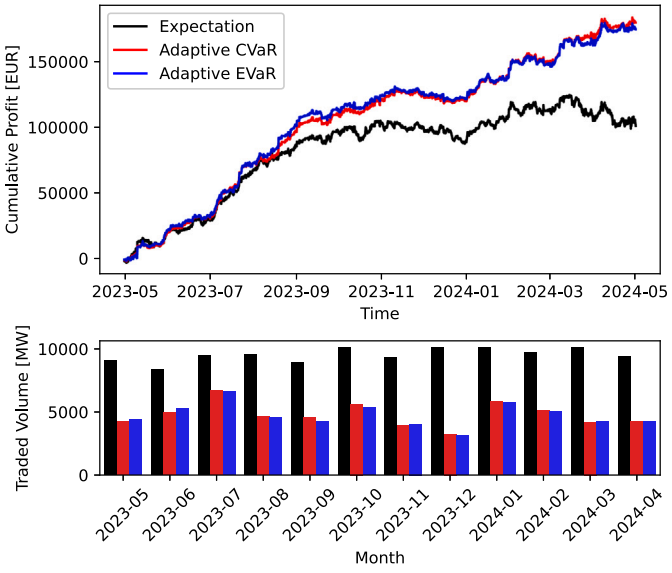


Fig. 10. The accumulated profits of the adaptive trading strategies. The adaptive strategies only require a fraction of the trades and generate higher profits.

Table 3
Influence of the market impact on the profit of the adaptive CVaR strategy. Here, β_{est} represents the estimated market impact (used by the strategy), while β denotes the actual market impact (used in the profit calculation). As expected, profit decreases as the actual market impact β increases. However, the differences in profit for a fixed β and varying β_{est} remain relatively small.

	$\beta = 0$	$\beta = 0.2$	$\beta = 0.4$	$\beta = 0.6$	$\beta = 0.8$	$\beta = 1$
$\beta_{\text{est}} = 0$	249k	228k	212k	193k	174k	161k
$\beta_{\text{est}} = 0.2$	**260k**	**238k**	216k	197k	180k	167k
$\beta_{\text{est}} = 0.4$	256k	237k	**218k**	**199k**	181k	170k
$\beta_{\text{est}} = 0.6$	251k	232k	215k	198k	178k	165k
$\beta_{\text{est}} = 0.8$	246k	230k	215k	194k	**182k**	168k
$\beta_{\text{est}} = 1$	244k	229k	215k	193k	178k	**171k**

6. Conclusion

Predicting future system imbalance prices is challenging due to their inherently alternating behavior. In this study, we employed a probabilistic mixture model to effectively capture the underlying distribution of the price, thereby reducing the overall complexity and enhancing robustness against overfitting. Compared to several benchmark models, the proposed mixture model outperforms its counterparts across every metric.

Using this model, we designed a strategy for automated trading on the intraday electricity market. This strategy explicitly accounts for (i) the impact of the trade positions on imbalance price and (ii) uncertainties inherent in the prediction model by including a coherent risk measure in the cost function. The tuning parameter of this risk measure is adaptively optimized using a window of recently observed data. After validating this strategy on the Belgian electricity market, the adaptive trading strategy resulted in larger absolute profits with only a fraction of the trades. This outcome highlights the importance of risk awareness in trading strategies to ensure robustness and profitability in real-world applications.

CRediT authorship contribution statement

Robin Bruneel: Writing – original draft, visualization, validation, software, methodology, investigation, funding acquisition, formal analysis, data curation, and conceptualization. **Mathijs Schuurmans:** Writing – review & editing, methodology, and conceptualization. **Panagiotis Patrinos:** Supervision, methodology, and conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Robin Bruneel reports that financial support was provided by Amplifino. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by Flanders Innovation & Entrepreneurship (VLAIO) under grant code “Baekeland Mandaat HBC.2022.0695”, Research Foundation Flanders (FWO) research projects G081222N, G033822N, and G0A0920N, and Research Council KUL grant C14/24/103. We gratefully acknowledge Smart@Energy and Amplifino for providing access to the Nordpool intraday energy market data.

Appendix A. Efficient optimization of Eq. (20) for the expected value

Before discussing more efficient optimization techniques, we begin by proving smoothness and convexity of the cost function φ when the risk measure equals the expected value. These are two necessary conditions to use gradient-based optimization methods.

Proposition. *Suppose $\rho = \mathbb{E}$, and the following conditions hold:*

(i) The regulation price sensitivities are equal: $K_{\text{MDP}} = K_{\text{MIP}} = K$

(ii) The intraday order book within U consists of a single bid such that the intraday energy cost is a linear function of u:

$$q_t(u)u = au + b,\ a \in \mathbb{R}_+,\ b \in \mathbb{R} \text{ for all } u \in U.$$

(iii) The maximum trade position is sufficiently small:

$$u_{\max} \leq \frac{20\ K n_q}{\beta w_u^2 \sum_{i=1}^{n_q} \left(g_{i,t}^{\text{MIP}} - g_{i,t}^{\text{MDP}}\right)}$$

Then, optimization problem (20) is smooth and convex.

Proof.

Combining condition (i) with Eq. (14), we can write

$$\mathbb{E}_{p\sim f_{t,u}^{\mathrm{SI}}}[-p] = K\beta u + \pi_t(u)\underbrace{\frac{1}{n_q}\sum_{i=1}^{n_q} -g_{i,t}^{\mathrm{MDP}}}_{c_t^{\mathrm{MDP}}} + (1-\pi_t(u))\underbrace{\frac{1}{n_q}\sum_{i=1}^{n_q} -g_{i,t}^{\mathrm{MIP}}}_{c_t^{\mathrm{MIP}}},$$

and thus, the cost φ in (20) is given as

$$\varphi(u) = \left(K\beta u + (c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}})\pi_t(u) + c_t^{\mathrm{MIP}} + a\right)u + b.$$

This function is twice differentiable:

$$\begin{aligned}\frac{\mathrm{d}\varphi(u)}{\mathrm{d}u} &= 2\ K\beta u + (c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}})\left(\pi_t(u) + \frac{\mathrm{d}\pi_t(u)}{\mathrm{d}u}u\right) + c_t^{\mathrm{MIP}} + a\\ \frac{\mathrm{d}^2\varphi(u)}{\mathrm{d}u^2} &= 2\ K\beta + (c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}})\left(2\frac{\mathrm{d}\pi_t(u)}{\mathrm{d}u} + \frac{\mathrm{d}^2\pi_t(u)}{\mathrm{d}u^2}u\right)\end{aligned} \tag{A.1}$$

where the derivative of the mixture weight $\pi_t(u)$, being a logistic regression model, is given by:

$$\begin{aligned}\frac{\mathrm{d}\pi_t(u)}{\mathrm{d}u} &= \pi_t(u)(1-\pi_t(u))\beta w_u\\ \frac{\mathrm{d}^2\pi_t(u)}{\mathrm{d}u^2} &= \pi_t(u)(1-\pi_t(u))(1-2\pi_t(u))\beta^2 w_u^2\end{aligned} \tag{A.2}$$

Here, w_u represents the weight of u in the logistic regression model. From the training set, we obtain $w_u = 0.007$. The derivatives of the cost function φ are continuous and bounded within U, which implies that the cost function is smooth. For convexity, it now suffices to demonstrate that the second-order derivative is positive within $U \subset \mathbb{R}_+$:

$$\begin{aligned}0 &\le \frac{\mathrm{d}^2\varphi(u)}{\mathrm{d}u^2}\\ 0 &\le 2\ K\beta + (c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}})\left(2\frac{\mathrm{d}\pi_t(u)}{\mathrm{d}u} + \frac{\mathrm{d}^2\pi_t(u)}{\mathrm{d}u^2}u\right)\end{aligned} \tag{A.3}$$

To solve this inequality, we derive a worst-case bound:

- The difference between predicted quantiles is always positive. The largest observed difference in the training set is: $\max_t\left(c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}}\right) = 700$
- The sigmoid function, as a function of u, always has a positive slope. Thus, the worst-case scenario for this inequality is: $\min_u \frac{\mathrm{d}\pi_t(u)}{\mathrm{d}u} = 0$
- The minimum second-order derivative of the sigmoid function with respect to u is approximately: $\min_u \frac{\mathrm{d}^2\pi_t(u)}{\mathrm{d}u^2} \approx -0.1\beta^2 w_u^2$

Under these worst-case conditions, inequality (A.3) simplifies to:

$$\begin{aligned}0 &\le 2\ K\beta - 0.1\beta^2 w_u^2\ u\left(c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}}\right)\\ u &\le \frac{20\ K}{\beta w_u^2\left(c_t^{\mathrm{MDP}} - c_t^{\mathrm{MIP}}\right)}\\ u &\le 233\ \mathrm{MW}\end{aligned} \tag{A.4}$$

As a result, the trade position would need to exceed at least 233 MW for the cost function to become concave. This is an unrealistically large volume for a small energy trader in Belgium, where typical trade sizes are approximately 5 MW. Eq. (A.4) is a sufficient condition for the cost function to be convex over its domain. □

The first assumption in the proposition is strictly not necessary, but holds in practice and significantly simplifies the computations in the proof. The second assumption ensures smoothness, as the intraday energy cost $q_t(u)u$ is a piecewise linear function when multiple order book bids have to be filled. In that specific case, the action space U can be subdivided into several segments, within each of which the intraday energy cost remains linear. By applying the proposition to each individual segment, the combined optimization problem becomes both piecewise smooth and piecewise convex. Concavity and non-smoothness may only occur at the (known) hinge points, allowing the problem to be solved efficiently nonetheless.

Our empirical findings indicate that the assumptions of the proposition are easily satisfied in practice. The smoothness of the method enables the use of gradient-based optimization techniques or, in this case, where a closed-form expression for the second-order derivative is available, Newton's method for efficiently solving the optimization problem. The convexity guarantees that any stationary point found is also a global minimizer.

Data availability

The authors do not have permission to share data.

References

[1] Goodarzi S, Perera HN, Bunn D. The impact of renewable energy forecast errors on imbalance volumes and electricity spot prices. Energy Policy 2019;134:110827. doi:https://doi.org/10.1016/j.enpol.2019.06.035. https://www.sciencedirect.com/science/article/pii/S0301421519304057.

[2] EPEX SPOT SE. Basics of the power market; 2024. https://www.epexspot.com/en/basicspowermarket [Accessed 15.11.2024].

[3] Agency for the Cooperation of Energy Regulators (ACER). ACER's monitoring shows broad implementation of the electricity imbalance settlement harmonisation methodology across the EU; 2024 July. https://www.acer.europa.eu/news/acers-monitoring-shows-broad-implementation-electricity-imbalance-settlement-harmonisation [Accessed 25.11.2024].

[4] Baetens J, Laveyne J, Van Eetvelde G, Vandevelde L. Imbalance pricing methodology in Belgium: implications for industrial consumers. In: 2020 17th International conference on the European Energy Market (EEM); 2020. p. 1–6. doi:https://doi.org/10.1109/EEM49802.2020.9221978.

[5] Narajewski M. Probabilistic forecasting of German electricity imbalance prices. Energies 2022;15:4976. doi:https://doi.org/10.3390/en15144976.

[6] Lima LM, Damien P, Bunn DW. Bayesian predictive distributions for imbalance prices with time-varying factor impacts. IEEE Trans Power Syst 2023;38(1):349–57. doi:https://doi.org/10.1109/TPWRS.2022.3165149.

[7] Makri E, Koskinas I, Tsolakis AC, Ioannidis D, Tzovaras D. Short term net imbalance volume forecasting through machine and deep learning: a UK case study. In: Maglogiannis I, Macintyre J, Iliadis L, editors. Artificial intelligence applications and innovations. AIAI 2021 IFIP WG 12.5 international workshops. Cham: Springer International Publishing; 2021. p. 377–89.

[8] Bunn DW, Inekwe JN, MacGeehan D. Analysis of the fundamental predictability of prices in the British balancing market. IEEE Trans Power Syst 2021;36(2):1309–16. doi:https://doi.org/10.1109/TPWRS.2020.3015871.

[9] Klæboe G, Eriksrud AL, Fleten S-E. Benchmarking time series based forecasting models for electricity balancing market prices. Energy Syst 2015;6(1):43–61. doi:https://doi.org/10.1007/s12667-013-0103-3.

[10] Olsson M, Soder L. Modeling real-time balancing power market prices using combined SARIMA and Markov processes. IEEE Trans Power Syst 2008;23(2):443–50. doi:https://doi.org/10.1109/TPWRS.2008.920046.

[11] Dumas J, Boukas I, de Villena MM, Mathieu S, Cornélusse B. Probabilistic forecasting of imbalance prices in the Belgian context. In: 2019 16th International conference on the European Energy Market (EEM); 2019. p. 1–7. doi:https://doi.org/10.1109/EEM.2019.8916375.

[12] Koch C. Intraday imbalance optimization: incentives and impact of strategic intraday bidding behavior. Energy Syst 2022;13(2):409–35. doi:https://doi.org/10.1007/s12667-021-00445-9.

[13] Koch C, Hirth L. Short-term electricity trading for system balancing: an empirical analysis of the role of intraday trading in balancing Germany's electricity system. Renew Sustain Energy Rev 2019;113:109275. doi:https://doi.org/10.1016/j.rser.2019.109275. https://www.sciencedirect.com/science/article/pii/S1364032119304836.

[14] Poplavskaya K, Lago J, de Vries L. Effect of market design on strategic bidding behavior: model-based analysis of European electricity balancing markets. Appl Energy 2020;270:115130. doi:https://doi.org/10.1016/j.apenergy.2020.115130. https://www.sciencedirect.com/science/article/pii/S0306261920306425.

[15] Lago J, Poplavskaya K, Suryanarayana G, De Schutter B. A market framework for grid balancing support through imbalances trading. Renew Sustain Energy Rev 2021;137:110467. doi:https://doi.org/10.1016/j.rser.2020.110467. https://www.sciencedirect.com/science/article/pii/S136403212030753X.

[16] Shapiro A, Dentcheva D, Ruszczynski A. Lectures on stochastic programming: modeling and theory. 3rd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics; 2021. doi:https://doi.org/10.1137/1.9781611976595.

[17] Bruninx K, Pandžić H, Le Cadre H, Delarue E. On the interaction between aggregators, electricity markets and residential demand response providers. IEEE Trans Power Syst 2020;35(2):840–53. doi:https://doi.org/10.1109/TPWRS.2019.2943670.

[18] Rashidizadeh-Kermani H, Vahedipour-Dahraie M, Shafie-Khah M, Siano P. A regret-based stochastic bi-level framework for scheduling of DR aggregator under uncertainties. IEEE Trans Smart Grid 2020;11(4):3171–84. doi:https://doi.org/10.1109/TSG.2020.2968963.
[19] Ghavidel S, Ghadi MJ, Azizivahed A, Aghaei J, Li L, Zhang J. Risk-constrained bidding strategy for a joint operation of wind power and caes aggregators. IEEE Trans Sustain Energy 2020;11(1):457–66. doi:https://doi.org/10.1109/TSTE.2019.2895332.
[20] Toubeau J-F, De Gréve Z, Vallée F. Medium-term multimarket optimization for virtual power plants: a stochastic-based decision environment. IEEE Trans Power Syst 2018;33(2):1399–410. doi:https://doi.org/10.1109/TPWRS.2017.2718246.
[21] Bottieau J, Bruninx K, Sanjab A, De Greve Z, Vallee F, Toubeau J-F. Automatic risk adjustment for profit maximization in renewable dominated short-term electricity markets. Int Trans Electr Energy Syst 2021 12; 31. doi:https://doi.org/10.1002/2050-7038.13152.
[22] Hua H, Chen X, Gan L, Sun J, Dong N, Liu D, et al. Demand-side joint electricity and carbon trading mechanism. IEEE Trans Ind Cyber-Phys Syst 2024;2:14–25. doi:https://doi.org/10.1109/TICPS.2023.3335328.
[23] McLachlan G, Peel D. Finite mixture models. New York: Wiley; 2000. Wiley series in probability and statistics. Applied probability and statistics section.
[24] Niculescu-Mizil A, Caruana R. Predicting good probabilities with supervised learning. 2005. p. 625–32. doi:https://doi.org/10.1145/1102351.1102430.
[25] Cucchiara A. Applied logistic regression. Technometrics 1992;34(3):358–9.
[26] Böhning D. Multinomial logistic regression algorithm. Ann Inst Stat Math 1992;44(1):197–200. doi:https://doi.org/10.1007/BF00048682.
[27] Koenker R, Hallock KF. Quantile regression. J Econ Perspect 2001;15(4):143–56. doi:https://doi.org/10.1257/jep.15.4.143.
[28] Schilling RL. Measures, integrals and martingales. Cambridge: Cambridge University Press; 2005.
[29] Rockafellar RT, Uryasev S, et al. Optimization of conditional value-at-risk. J Risk 2000;2:21–42.
[30] Ahmadi-Javid A. Entropic value-at-risk: a new coherent risk measure. J Optim Theory Appl 2012;155(3):1105–23. doi:https://doi.org/10.1007/s10957-011-9968-2.
[31] Drusvyatskiy D, Xiao L. Stochastic optimization with decision-dependent distributions. Math Oper Res 2023;48(2):954–98. doi:https://doi.org/10.1287/moor.2022.1287.
[32] Mendler-Dünner C, Perdomo J, Zrnic T, Hardt M. Stochastic optimization for performative prediction. In Larochelle H, Ranzato M, Hadsell R, Balcan M, Lin H, editors. Advances in neural information processing systems, vol. 33. Curran Associates, Inc; 2020. p. 4929–39. https://proceedings.neurips.cc/paper_files/paper/2020/file/33e75ff09dd601bbe69f351039152189-Paper.pdf.
[33] Perdomo J, Zrnic T, Mendler-Dünner C, Hardt M. Performative prediction, In: III HD, Singh A, editors. Proceedings of the 37th international conference on machine learning, vol. 119. Proceedings of Machine Learning Research, PMLR; 2020. p. 7599–609. https://proceedings.mlr.press/v119/perdomo20a.html.
[34] Schuurmans M, Katriniok A, Meissen C, Tseng HE, Patrinos P. Safe, learning-based MPC for highway driving under lane-change uncertainty: a distributionally robust approach. Artif Intell 2023;320:103920. doi:https://doi.org/10.1016/j.artint.2023.103920.
[35] Elia Group. Volume needs; 2024. https://www.elia.be/en/grid-data/balancing/capacity-volumes-needs [Accessed 18.06.2024].
[36] Meinshausen N. Quantile regression forests. J Mach Learn Res 2006;7(35):983–99.
[37] Gneiting T, Katzfuss M. Probabilistic forecasting. Annu Rev Stat Appl 2014;1(1):125–51. doi:https://doi.org/10.1146/annurev-statistics-062713-085831.
[38] Hersbach H. Decomposition of the continuous ranked probability score for ensemble prediction systems. Weather Forecast (2000);15(5):559–70. doi:https://doi.org/10.1175/1520-0434(2000)015{<}0559:DOTCRP{>}2.0.CO;2.
[39] ENTSO-E. PICASSO: the platform for the international coordination of automated frequency restoration and stable system operation; 2017.
[40] Energy Trading Platform Amsterdam (ETPA). ETPA trading platform; 2024. https://www.etpa.nl/ [Accessed 17.09.2024].