

Value-oriented forecast reconciliation for renewables in electricity markets

Honglin Wen^{a,b,*}, Pierre Pinson^{b,c,d,e,*}

^a*Department of Electrical Engineering, Shanghai Jiao Tong University, China,*

^b*Dyson School of Design Engineering, Imperial College London, United Kingdom,*

^c*Halfspace, Denmark,*

^d*Department of Technology, Management and Economics, Technical University of Denmark, Denmark,*

^e*CoRE, Aarhus University, Denmark,*

Abstract

Forecast reconciliation is considered an effective method to achieve coherence (within a forecast hierarchy) and to improve forecast quality. However, the value of reconciled forecasts in downstream decision-making tasks has been mostly overlooked. In a multi-agent setup with heterogeneous loss functions, this oversight may lead to unfair outcomes, hence resulting in conflicts during the reconciliation process. To address this, we propose a value-oriented forecast reconciliation approach that focuses on the forecast value for all individual agents. Fairness is ensured through the use of a Nash bargaining framework. Specifically, we model this problem as a cooperative bargaining game, where each agent aims to optimize their own gain while contributing to the overall reconciliation process. We then present a primal-dual algorithm for parameter estimation based on empirical risk minimization. From an application perspective, we consider an aggregated wind energy trading problem, where profits are distributed using a weighted allocation rule. We demonstrate the effectiveness of our approach through several numerical experiments, showing that it consistently results in increased profits for all agents involved.

Keywords: Forecasting; Reconciliation; Decision-making; Nash bargaining; Machine learning

1. Introduction

The geographically distributed nature of renewable energy sources, such as wind and solar, inherently creates a hierarchical structure for power generation. It is acknowledged that forecast accuracy can be improved by leveraging the spatial-temporal dependencies inherent in geographically distributed data (Tastu et al., 2013). However, various agents at different levels of aggregation

*Corresponding authors

Email addresses: linlin00@sjtu.edu.cn (Honglin Wen), p.pinson@imperial.ac.uk (Pierre Pinson)

in power systems and electricity markets, for instance, the portfolio and individual power generation assets, often independently generate their own forecasts. This can lead to *incoherence*, i.e., the forecast directly produced at the portfolio level is not necessarily equal to the sum of the forecasts for individual assets. Given that forecasts are frequently used in decision-making, incoherence across aggregation levels may eventually lead to dilemmas among decision-makers. For that, *forecast reconciliation* is applied, which is a post-forecasting process that involves transforming independently generated incoherent base forecasts into coherent forecasts that satisfy the constraints imposed by the hierarchical structures (Athanasopoulos et al., 2024).

Single-level reconciliation approaches, such as bottom-up, top-down, and middle-out approaches, have been in use for decades, dating back at least to Edwards and Orcutt (1969) and Zellner (1969). They require forecasts at a chosen level, which are then combined to generate forecasts for the remaining levels in the hierarchy. Hyndman et al. (2011) introduced a combination approach that leverages forecasts across the entire hierarchy. This was later extended by Van Erven and Cugliari (2015), Hyndman et al. (2016), and Wickramasuriya et al. (2019) through employing various projection matrices (Panagiotelis et al., 2021). In parallel, Ben Taieb and Koo (2019) developed an end-to-end reconciliation approach, which learns the projection matrices via machine learning. Forecast reconciliation in the energy sector has been employed since Van Erven and Cugliari (2015), and most often brings the added benefit of enhanced forecast accuracy (Zhang and Dong, 2018; Jeon et al., 2019).

While previous studies have substantially advanced forecast reconciliation through various methods and techniques, nearly all of them focus on leveraging hierarchical structures to improve forecast accuracy. Nonetheless, as Murphy (1993) suggests, in addition to evaluating the quality (accuracy) of forecasts, which is the correspondence between forecasts and observed outcomes, it is also important to assess the consistency, or the alignment between forecasts and judgments, and the value, which refers to the benefits they provide to users, in terms of profits or costs. Evidence suggests that more accurate forecasts do not always lead to higher value in downstream decision-making processes (Kourentzes, 2013; Barrow and Kourentzes, 2016). Although increasing emphasis is placed on developing forecast products that maximize value in decision-making, particularly within the energy industry (Carriere and Kariniotakis, 2019; Stratigakos et al., 2022; Zhang et al., 2025) and inventory management (Kourentzes et al., 2020; Theodorou et al., 2025),

especially after the theoretical analysis by [Bertsimas and Kallus \(2020\)](#) and [Elmachtoub and Grigas \(2022\)](#), value in forecast reconciliation remains largely under-investigated, with only a few notable exceptions. For instance, [Sanguri et al. \(2022\)](#) assessed the effectiveness of coherent forecasts in enhancing the value of port operations and empirically demonstrated that coherent forecasts improved decision-making at ports. In parallel, [Stratigakos et al. \(2024\)](#) proposed a method for combining multiple probabilistic forecasts using a linear opinion pool, with a particular emphasis on minimizing the expected costs associated with subsequent decision-making. However, all of these studies evaluate the value in forecast reconciliation from a centralized perspective, neglecting the fact that the agents involved may have heterogeneous loss (or cost) functions.

When agents have heterogeneous cost functions, the primary challenge lies in achieving consensus on the reconciliation approach before actually reconciling their forecasts. Since the commonly used reconciliation approaches aim to optimize an overall loss function, for instance, the mean squared error function ([Van Erven and Cugliari, 2015](#); [Hyndman et al., 2016](#)), some agents involved may even find that reconciled forecasts lead to an increase in their own costs, even though overall costs are reduced after reconciliation. As a result, the agents may perceive the reconciliation process as unfair, leading to disagreements on the reconciliation procedure. Indeed, studies like [Van Erven and Cugliari \(2015\)](#) and [Sprangers et al. \(2024\)](#) included weighting factors for forecasts at each level during the reconciliation process, either on an ad hoc basis or based on hierarchical levels. Nevertheless, the emphasis is primarily placed on the overall accuracy of the forecasts, thereby leaving the effects of this weighting on subsequent decision-making results unexplored. Therefore, how to account for the value of forecasts at each level and ensure fairness in the reconciliation process remains an unresolved issue in the forecast reconciliation literature. This paper addresses this gap by introducing a *value-oriented reconciliation* approach, which we contrast with the widely studied *quality-oriented reconciliation* approach.

To motivate our proposal and demonstrate its applicability, we focus on the well-established problem of trading wind energy in electricity markets, where a portfolio manager (PM) trades on behalf of multiple wind power producers (WPPs). In this setup, a natural two-level hierarchy emerges between the PM and the WPPs, which implies that the PM's offer must represent the collective output of all WPPs. As the forecasts issued by the PM and WPPs may be incoherent, they must be reconciled to ensure coherence before being used in the PM's trading strategies. After

trading in the markets, the PM allocates profits to each WPP by using pre-defined allocation rules.

Our reconciliation framework is inspired by Pennings and Van Dalen (2017), where the reconciled forecasts are expressed as the product of a structural matrix (defined by the hierarchy) and latent variables. The latent variables are derived from the base forecasts via learnable combination functions. Market data, including historical prices, are incorporated as context features for the combination functions, enabling the reconciliation to be adaptive to contextual information. The requirement for consensus on the reconciliation approach aligns with the group decision-making process, in which agents collaboratively select a combination function from the collection of mutually accepted functions. Particularly, we assume that WPPs exhibit *individual rationality*, indicating that each producer’s profit must be at least equal to their independent offering. Therefore, the forecast reconciliation process must ensure that each producer’s profit is either maintained or improved after reconciliation. To address this, we model the problem as a cooperative bargaining game (Nash, 1953), where each agent seeks to maximize their own gain by choosing from the available combination functions. The Nash bargaining solution is derived by maximizing the product of excess profits (known as the Nash product), which measures the difference between the post-reconciliation profits and the pre-reconciliation profits.

Given that the PM regularly engages in market trading, we apply the principle of fairness over time (Lodi et al., 2024) to determine the combination function in the reconciliation process, rather than relying on a single trade instance. Consequently, we calculate the average profit for each agent within the Nash bargaining framework. The combination function is derived through constrained optimization, with the objective of maximizing the product of average excess profits while ensuring that each agent’s average excess profit remains positive. It turns out to be a large-scale optimization problem, thus we frame it as a constrained learning problem by parameterizing the combination function (Donini et al., 2018). Specifically, we use the Lagrangian optimization approach outlined by Chamon and Ribeiro (2020) for parameter estimation. The parameters of the combination function model are estimated using the stochastic gradient descent algorithm, while the dual variables are updated through the dual ascent algorithm. Finally, following Muñoz et al. (2023), we validate the proposed method with case studies in an hourly forward market coupled with a dual-price settlement balancing market, using market data from the Danish TSO and wind power data from the Wind Toolkit. In summary, this work makes the following contributions:

- (1) We propose a value-oriented forecast reconciliation approach, with a particular emphasis on the value of forecasts for each agent involved. The approach ensures fairness through a Nash bargaining framework. It complements the existing forecast reconciliation literature by emphasizing forecast value rather than solely focusing on forecast quality.
- (2) We demonstrate the effectiveness of the proposed method in the context of aggregated wind energy trading and illustrate how consensus can be achieved when dealing with forecasts from multiple sources. This extends the existing literature on aggregated wind energy trading by introducing the concept of consensus on information, in addition to the traditional focus on consensus in profit allocations.
- (3) We prove in Corollary 3.2 that, by employing weighted proportional allocation rules, the agreement set of parameters in the Nash bargaining problem—within the context of aggregated wind energy trading—is guaranteed to be non-empty.
- (4) We propose an algorithm for parameter estimation within the Nash bargaining framework that alternates between maximizing over the dual variables and minimizing over the parameters of the reconciliation function.

The structure of the paper is as follows. Section 2 introduces the core concepts of value-oriented forecasting and decision-making, with a particular emphasis on wind energy trading, and discusses the challenges associated with aggregated wind energy trading. Section 3 formulates the problem and develops a value-oriented forecast reconciliation approach based on Nash bargaining, framing it as a learning problem. Section 4 presents two illustrative cases to explore the behavior of the proposed approach. Section 5 empirically validates the approach by constructing an hourly forward market using real market data. Finally, Section 6 concludes the paper and suggests directions for future research.

Notations: Vectors are denoted by bold lowercase letters (e.g., \mathbf{x}), while matrices are represented by bold uppercase letters (e.g., \mathbf{A}). Sets are denoted using calligraphic fonts, such as \mathcal{T} , with the cardinality of a set expressed as $|\mathcal{T}|$. The subscript F represents the forward market (e.g., π^F), while the superscripts UP and DW denote up- and down-regulation (e.g., π^{UP} and π^{DW}).

2. Preliminaries

Our setup closely mirrors that of [Guerrero-Mestre et al. \(2015\)](#). Specifically, we consider two trading floors: a forward market and a balancing market with a dual-price settlement mechanism for imbalances. As noted by [Pinson et al. \(2007\)](#), the trading problem for a risk-neutral, price-taker WPP is analogous to a newsvendor problem. Ultimately, it involves offering on a specific quantile of the predictive distribution for wind power generation. We assume that both the PM and WPPs are risk-neutral and act as price-takers, enabling us to model their interaction as a cooperative newsvendor game ([Montrucchio et al., 2012](#)). In what follows, we first outline the fundamentals of value-oriented forecasting and decision-making, with a specific focus on wind energy trading. We then present renewable energy offering in the forward market within the context of dual-price imbalance settlement. Finally, we describe the challenges associated with aggregated wind energy trading.

2.1. Value-oriented Forecasting and Decision Making

To begin, we outline the setups of wind energy offering. Let $y_{t,i}$ represent the wind energy generated by WPP i at time t , which is a realization of the random variable $Y_{t,i}$. In a forward market, WPP i must declare the energy amount $y_{t+k,i}^F$ it will generate before the market's closure time, t . Here, the subscript k indicates the time interval between trading and actual power delivery. In the balancing market, the WPP is required to balance the discrepancy between the actual generation, $y_{t+k,i}$, and the declared amount, $y_{t+k,i}^F$. Throughout this paper, we assume that WPPs act as price takers and are risk-neutral. This assumption is commonly observed in the electricity market literature ([Morales et al., 2013](#)), implying that the trading activities of a single WPP do not influence the clearing price. The prices for the forward market, upward and downward regulation are denoted by π_{t+k}^F , π_{t+k}^{UP} , and π_{t+k}^{DW} , respectively. It is established that $\pi_{t+k}^{\text{UP}} \geq \pi_{t+k}^F$ and $\pi_{t+k}^{\text{DW}} \leq \pi_{t+k}^F$, with at most one of these values differing from π_{t+k}^F .

The profit of a WPP at the trading hour $t+k$ is then given by the function $\rho_{t+k}(y_{t+k,i}^F, y_{t+k,i})$, which is expressed as

$$\rho_{t+k}(y_{t+k,i}^F, y_{t+k,i}) = \pi_{t+k}^F y_{t+k,i}^F - \pi_{t+k}^{\text{UP}} [y_{t+k,i}^F - y_{t+k,i}]^+ + \pi_{t+k}^{\text{DW}} [y_{t+k,i} - y_{t+k,i}^F]^+, \quad (1)$$

where $[x]^+ = \max(x, 0), \forall x \in \mathbb{R}$. When the power generation falls short of the offer, i.e., $y_{t+k,i} <$

$y_{t+k,i}^F$, the WPP has to purchase up-regulation at the price of π_{t+k}^{UP} to balance the discrepancy $y_{t+k,i}^F - y_{t+k,i}$. Conversely, if the power generation exceeds the offer, i.e., $y_{t+k,i}^F > y_{t+k,i}$, the WPP is required to purchase down-regulation at the price of π_{t+k}^{DW} . The penalties associated with overproduction and underproduction are calculated as follows: for overproduction, $\psi_{t+k}^+ = \pi_{t+k}^{\text{F}} - \pi_{t+k}^{\text{DW}}$, and for underproduction, $\psi_{t+k}^- = \pi_{t+k}^{\text{UP}} - \pi_{t+k}^{\text{F}}$. For mathematical convenience, we recast (1) as

$$\rho_{t+k}(y_{t+k,i}^F, y_{t+k,i}) = \pi_{t+k}^{\text{F}} y_{t+k,i} - \left[\psi_{t+k}^+ [y_{t+k,i} - y_{t+k,i}^F]^+ + \psi_{t+k}^- [y_{t+k,i}^F - y_{t+k,i}]^+ \right], \quad (2)$$

given that

$$y_{t+k,i} - y_{t+k,i}^F = [y_{t+k,i} - y_{t+k,i}^F]^+ - [y_{t+k,i}^F - y_{t+k,i}]^+.$$

The first term on the RHS of (2) represents the profit derived from the oracle forecast (i.e., an ideal forecast), and consequently, it is outside the producer's influence. The subsequent term on the RHS of (2) is identified as the *imbalance cost*, which is a non-negative quantity represented by a function $c_{t+k}(y_{t+k,i}^F, y_{t+k,i})$ as follows:

$$c_{t+k}(y_{t+k,i}^F, y_{t+k,i}) = \psi_{t+k}^+ [y_{t+k,i} - y_{t+k,i}^F]^+ + \psi_{t+k}^- [y_{t+k,i}^F - y_{t+k,i}]^+. \quad (3)$$

And we note that ψ_{t+k}^+ and ψ_{t+k}^- are non-negative and bounded.

It is important to note that the optimal offer can be determined by minimizing the expected imbalance cost, provided that the distribution of $Y_{t+k,i}$ is known. It is expressed as

$$y_{t+k,i}^{\text{F}*} = \arg \min_{y_{t+k,i}^{\text{F}} \in [0, u_i]} \mathbb{E}_{Y_{t+k,i}} [c_{t+k}(y_{t+k,i}^{\text{F}}, y_{t+k,i})], \quad (4)$$

where u_i represents the capacity of WPP i . This addresses a specific instance of the newsvendor problem (Choi, 2012), which leads to the identification of a particular quantile within the wind power generation distribution:

$$y_{t+k,i}^{\text{F}*} = F_{y_{t+k,i}}^{-1}(\alpha_{t+k}), \quad (5)$$

where $F_{y_{t+k,i}}$ represents the cumulative distribution function (CDF) of $Y_{t+k,i}$ and α_{t+k} signifies the

nominal level, expressed as:

$$\alpha_{t+k} = \frac{\psi_{t+k}^+}{\psi_{t+k}^+ + \psi_{t+k}^-}.$$

Thus, the expense associated with the imbalance is represented by $c_{t+k}(y_{t+k,i}^{\text{F}^*}, y_{t+k,i})$, and the expected cost is obtained as $\mathbb{E}_{Y_{t+k,i}}[c_{t+k}(y_{t+k,i}^{\text{F}^*}, y_{t+k,i})]$.

However, both the real generation $y_{t+k,i}$ and its distribution $F_{y_{t+k,i}}$ are unknown at the time of offering t , thus WPPs must rely on forecasts to determine their offers. It is common to use the contextual information available at time t such as weather and lagged values, denoted as $\mathbf{x}_{t,i}$, to predict the future value $Y_{t+k,i}$ using a specific model f_i . Denoting the point forecast as $\hat{y}_{t+k,i|t}$, forecasting can be mathematically expressed as

$$\hat{y}_{t+k,i|t} = f_i(\mathbf{x}_{t,i}; \phi_i), \quad (6)$$

where ϕ_i represents the model parameters. According to the literature on electricity markets, two main point forecasting products are identified for trading purposes: mean and quantiles. They are developed by employing various loss functions to train the model (6), as elaborated below.

Quality-oriented forecasting: The widely employed point forecasting seeks to estimate the expected value of $Y_{t+k,i}$. For that, the parameters of the model are estimated by minimizing the expectation of the mean squared error function $l(\hat{y}_{t+k,i|t}, y_{t+k,i})$, expressed as

$$\min_{\phi_i} \mathbb{E} [l(\hat{y}_{t+k,i|t}, y_{t+k,i})]. \quad (7)$$

We refer to this approach as quality-oriented forecasting since the loss function is directly related to the quality (or accuracy) of the forecasts.

Value-oriented forecasting: As various studies indicate that more precise forecasts do not necessarily enhance value in subsequent decision-making (Theodorou et al., 2025), we can create forecasting models that directly target decision-making value. Then, the parameters of the forecast model are estimated by minimizing the expected cost, expressed as

$$\min_{\phi_i} \mathbb{E} [c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i})], \quad (8)$$

where $c_{t+k}(\cdot)$ is defined in (3). Particularly for the wind power offering problem, [Bremnes \(2004\)](#) and [Pinson et al. \(2007\)](#) suggest that, for a specified nominal quantile level, it is allowed to directly estimate the optimal offers using quantile regression. Specifically, they employ the long-term mean values of ψ_{t+k}^+ and ψ_{t+k}^- during the model training process. It is worth mentioning that the loss function is not limited to the pinball loss; it can also be a more generalized form as explored in studies by [Kourentzes et al. \(2020\)](#); [Stratigakos et al. \(2022\)](#); [Zhang et al. \(2025\)](#), which address issues related to optimization or simulation. Since the loss function aligns with the value of downstream decisions, we refer to this approach as value-oriented forecasting.

2.2. Aggregated Wind Energy Trading

Leveraging the complementary nature of geographically diverse renewable energy sources to hedge against price and quantity uncertainty in electricity markets has attracted significant interest ([Bitar, 2011](#)). Consider a scenario where there are m WPPs, managed by a PM responsible for trading in the market on behalf of all the WPPs. Let $Y_{t+k,\text{sum}}$ denote the aggregate sum of the random variables $Y_{t+k,i}$, expressed as $Y_{t+k,\text{sum}} = \sum_i^m Y_{t+k,i}$. The observed value is indicated by $y_{t+k,\text{sum}} = \sum_i^m y_{t+k,i}$. Meanwhile, $y_{t+k,\text{sum}}^F$ stands for the trading offer by the PM, and $c_{t+k}(y_{t+k,\text{sum}}^F, y_{t+k,\text{sum}})$ denotes the total imbalance cost associated with the offer $y_{t+k,\text{sum}}^F$. When the generation assets are owned by the same company, the portfolio manager (PM) aims to minimize the total expected imbalance cost. In this case, the problem reduces to a classical newsvendor formulation, and the optimal offer for the aggregated portfolio is given by

$$y_{t+k,\text{sum}}^{F*} = F_{y_{t+k,\text{sum}}}^{-1}(\alpha_{t+k}), \quad (9)$$

where $F_{y_{t+k,\text{sum}}}$ denotes the CDF of the aggregated random variable $Y_{t+k,\text{sum}}$. The imbalance cost associated with the offer $y_{t+k,\text{sum}}^{F*}$ is then given by $c_{t+k}(y_{t+k,\text{sum}}^{F*}, y_{t+k,\text{sum}})$. Accordingly, the expected total imbalance cost is computed as $\mathbb{E}_{Y_{t+k,\text{sum}}} [c_{t+k}(y_{t+k,\text{sum}}^{F*}, y_{t+k,\text{sum}})]$.

The cost-saving benefits of aggregated trading—originally demonstrated by [Eppen \(1979\)](#) and further supported by [Hartman et al. \(2000\)](#)—are outlined below.

Theorem 2.1. [[Eppen \(1979\)](#); [Hartman et al. \(2000\)](#)] In a multi-location newsvendor system, the total optimal expected cost of the centralized system is lower than that of the decentralized system.

Using Theorem 2.1, it is evident that

$$\mathbb{E}_{Y_{t+k,\text{sum}}} [c_{t+k}(y_{t+k,\text{sum}}^{\text{F}^*}, y_{t+k,\text{sum}})] \leq \sum_i^m \mathbb{E}_{Y_{t+k,i}} [c_{t+k}(y_{t+k,i}^{\text{F}^*}, y_{t+k,i})], \quad (10)$$

which implies that offering as an aggregation can always lead to a reduction in costs for participants. Equality in (10) holds only when the variables $Y_{t+k,1}, \dots, Y_{t+k,m}$ are perfectly positively correlated.

In fact, Bitar (2011) and Baeyens et al. (2013) have demonstrated that WPPs can consistently increase their expected profits through aggregated trading. As discussed above, in the centralized context, it is optimal to submit $y_{t+k,\text{sum}}^{\text{F}^*}$ as the trading offer, as the primary objective is to maximize total profits. In contrast, in a multi-agent setting where the generation assets are owned by different companies, each agent seeks to maximize its own profit. As a result, the aggregate of independently determined offers often diverges from the PM’s optimal offer. This discrepancy necessitates determining both the total trading offer submitted to the market and the corresponding profit allocation for each individual agent. Applying simple cost allocation rules—such as evenly splitting the total cost among WPPs—may lead to dissatisfaction among participants due to perceived unfairness. To address this issue, various cooperative game-theoretic approaches have been employed to allocate costs among participating agents, including proportional sharing, the Shapley value, and the nucleolus. For instance, Freire et al. (2014); Nguyen and Le (2018) developed Benders decomposition and constraint-generation algorithms to solve mixed-integer linear programming formulations of the nucleolus problem, which aims to identify allocation rules within the core. Based on the selected allocation rule, the PM determines the market offer and subsequently distributes profits (or, equivalently, costs) to the WPPs. The operational framework implemented by the PM is illustrated in Figure 1.

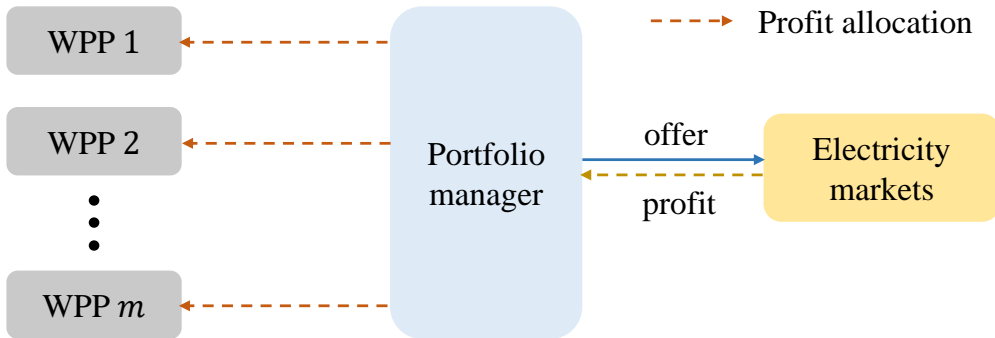


Figure 1: wind energy trading as an aggregation.

Although existing studies have investigated various allocation rules and computational techniques, they typically assume that the PM and WPPs have access to identical information and that WPPs unanimously agree on the submitted offers. In practice, however, diverse perspectives on future energy generation can lead to differing opinions between the PM and the WPPs regarding the appropriate offers. Specifically, each WPP typically uses a forecast model to estimate its own independent offer, as expressed in (6). Similarly, the PM employs a forecasting model f with parameters ϕ to predict the aggregated offer based on the feature vector \mathbf{x}_t . This model is expressed as

$$y_{t+k,\text{sum}}^{\text{F}} = \hat{y}_{t+k,\text{sum}|t} = f(\mathbf{x}_t; \phi). \quad (11)$$

As the forecasts $\{\hat{y}_{t+k,\text{sum}|t}, \hat{y}_{t+k,1|t}, \dots, \hat{y}_{t+k,m|t}\}$ are generated independently by the PM and WPPs, they are likely to be incoherent. Thus, WPPs may disagree with the offer proposed by the PM, and furthermore, the forecasts cannot be directly used to determine each WPP's share in the submitted offer. We address this challenge by developing a value-oriented forecast reconciliation approach in Section 3.

3. Methodology

In this section, we first formulate the forecast reconciliation problem within the context of aggregated wind energy trading. We then propose a value-oriented forecast reconciliation approach based on the Nash bargaining framework, followed by the presentation of a primal-dual algorithm for parameter estimation.

3.1. Problem Formulation

For ease of reference, the following concise notations are introduced: $[y_{t+k,\text{sum}}, y_{t+k,1}, \dots, y_{t+k,m}]^{\top}$ is represented by \mathbf{y}_{t+k} , while $[\hat{y}_{t+k,\text{sum}|t}, \hat{y}_{t+k,1|t}, \dots, \hat{y}_{t+k,m|t}]^{\top}$ is expressed as $\hat{\mathbf{y}}_{t+k|t}$, and referred to as the *base forecast*. As shown in Figure 1, the interaction between WPPs and the PM inherently establishes a two-level hierarchical structure. In the context of forecast reconciliation (Athanasopoulos et al., 2024), forecasts at the leaf nodes of the hierarchy are referred to as *bottom-level forecasts*, while forecasts at other nodes are called *aggregated forecasts*. In our specific WPP-PM hierarchy, we can decompose the vector \mathbf{y}_t into m bottom-level values $\mathbf{b}_t \in \mathbb{R}^m$ and an aggregated value $y_{t,\text{sum}}$ such that $\mathbf{y}_t = [y_{t,\text{sum}}, \mathbf{b}_t^{\top}]^{\top}$. By defining a structural matrix $\mathbf{S} \in \{0, 1\}^{(m+1) \times m}$, we

obtain

$$\mathbf{y}_t = \mathbf{S}\mathbf{b}_t \Leftrightarrow \begin{bmatrix} y_{t,\text{sum}} \\ \mathbf{b}_t \end{bmatrix} = \begin{bmatrix} \mathbf{s}_{\text{sum}} \\ \mathbf{I}_m \end{bmatrix} \mathbf{b}_t, \quad \forall t, \quad (12)$$

where $\mathbf{s}_{\text{sum}} = [1, 1, \dots, 1]$ is an aggregation vector and \mathbf{I}_m is the $m \times m$ identity matrix. Equivalently, (12) can be rewritten as

$$\mathbf{B}\mathbf{y}_t = \mathbf{0}, \quad (13)$$

where $\mathbf{B} = [1, -\mathbf{s}_{\text{sum}}]$. Notably, the hierarchy does not have to be confined to just two levels, as it can accommodate $n - m > 1$ aggregated series, meaning that $\mathbf{y}_t \in \mathbb{R}^n$. Correspondingly, the aggregation matrix is represented by $\mathbf{S}_{\text{sum}} \in \{0, 1\}^{(n-m) \times m}$, and \mathbf{B} defined as $[\mathbf{I}_{n-m} | -\mathbf{S}_{\text{sum}}]$. With the hierarchical structure, the coherent forecast is formally defined as follows.

Definition 3.1 (Coherent forecast). A forecast $\hat{\mathbf{y}}_{t+k|t}$ is said to be coherent if $\mathbf{B}\hat{\mathbf{y}}_{t+k|t} = \mathbf{0}$.

As previously stated, the base forecast, $\hat{\mathbf{y}}_{t+k|t}$ is independently generated by WPPs and PM, and is likely not coherent. Therefore, it is essential to reconcile the forecasts before using them in market trading. Let $\tilde{\mathbf{y}}_{t+k|t} = [\tilde{y}_{t+k,\text{sum}|t}, \tilde{y}_{t+k,1|t}, \dots, \tilde{y}_{t+k,m|t}]^\top$ denote the reconciled forecasts, where

$$\tilde{y}_{t+k,\text{sum}|t} = \sum_{i=1}^m \tilde{y}_{t+k,i|t}. \quad (14)$$

The reconciliation process can then be conceptually expressed as

$$\tilde{\mathbf{y}}_{t+k|t} = f_{\text{RE}}(\hat{\mathbf{y}}_{t+k|t}), \quad (15)$$

where f_{RE} represents the reconciliation function. In other words, we transform incoherent base forecast $\hat{\mathbf{y}}_{t+k|t}$ into coherent forecast $\tilde{\mathbf{y}}_{t+k|t}$ via the reconciliation function. This process can be viewed as a special approach to forecast combination, integrating both direct predictions for each series and indirect forecasts derived from the structure (Hollyman et al., 2021).

After reconciliation, the PM then places offers in the market using the reconciled aggregate forecast, $\tilde{y}_{t+k,\text{sum}|t}$. In parallel, the reconciled forecast for each WPP, $\tilde{y}_{t+k,i|t}$, is not submitted to the market; it serves instead to determine cost (or profit) allocation and is thus termed a pseudo-offer. In contrast to the independent offering outlined in Section 2, where all WPPs utilize an identical cost function, namely $c_{t+k}(\hat{y}_{t+k,i}, y_{t+k,i})$, in this context, the WPPs encounter distinct

cost functions. This variation arises because the WPPs are allocated as portions of the total balancing cost. Concretely, the allocated cost for WPP i is calculated based on the reconciled forecasts $\tilde{\mathbf{y}}_{t+k|t}$ and realizations \mathbf{y}_{t+k} , which is represented as $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$. The detailed definition of $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ is postponed to Section 3.2. Correspondingly, the profit for WPP i in the context of aggregated wind energy trading, is represented as $\rho_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$, which can be expressed as

$$\rho_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) = \pi_{t+k}^{\text{F}} y_{t+k,i} - c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}). \quad (16)$$

It is observed that, after reconciliation, the balancing cost of the aggregated forecast is always lower than or equal to the sum of costs of independent pseudo-offers, as described below.

Proposition 3.1. Let $\{\tilde{y}_{t+k,1|t}, \dots, \tilde{y}_{t+k,m|t}\}$ be a set of m independent pseudo-offers for the trading hour $t+k$. Then, it holds that,

$$c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}) \leq \sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}),$$

where $\tilde{y}_{t+k,\text{sum}|t} = \sum_{i=1}^m \tilde{y}_{t+k,i|t}$.

Proof. See the appendix. □

Since the pseudo-offers $\tilde{y}_{t+k,i|t}$ are not actually submitted to the markets, the above proposition illustrates the potential cost implications of using reconciled forecasts as offers in the markets for any realization of the actual observations. Specifically, in the case of bottom-up reconciliation, each reconciled forecast $\tilde{y}_{t+k,i|t}$ coincides with the corresponding base forecast $\hat{y}_{t+k,i|t}$. This implies that the balancing cost of an aggregated offering can be lower than or equal to the sum of the costs incurred under independent offerings. Based on this, we state the following corollary.

Corollary 3.1. After bottom-up reconciliation, the balancing cost of an aggregated offering is less than or equal to the sum of the costs incurred under independent offerings, for any realization of the actual observations.

Notably, coherence is often effectively ensured using existing forecast reconciliation methods. However, as noted by Athanasopoulos and Kourentzes (2023), heterogeneous objectives sometimes arise across different levels of the hierarchy, with specific decisions tied to individual nodes. Despite

this, most existing quality-oriented reconciliation approaches adopt a single loss function—typically the mean squared error—within the reconciliation process (Athanasopoulos et al., 2024). In contexts where agents possess heterogeneous cost functions, reconciliation naturally becomes a multi-objective optimization problem, highlighting the inherent challenge of reaching consensus on a suitable reconciliation function. Before presenting our proposed method, we outline several key assumptions, deferring detailed methodology to Section 3.2.

Assumption 3.1. The group of WPPs is compelled to form a grand coalition. If they fail to reach a consensus on the reconciliation process, each WPP will act independently.

In this context, we exclude sub-coalitions, as the forecasts for potential sub-coalitions are not available. Each WPP will decide whether to enter into an agency contract with the PM. Even if some WPPs choose to opt out of the contract, the PM can still manage the portfolio consisting of the remaining WPPs.

Assumption 3.2. Each WPP receives no less than the expected profit they would achieve through independent trading.

This assumption is valid because WPPs are primarily focused on the individual benefits they derive from the coalition, rather than on the collective benefits of the coalition. As both α_{t+k} and $Y_{t,i}$ are time-varying, we introduce the following assumptions in line with Besbes et al. (2015).

Assumption 3.3. We define the total variation of the sequence $\{\alpha_{t+k}\}$ as V_α , given by

$$V_\alpha = \sum_{t=1}^T |\alpha_{t+k} - \alpha_{t+k-1}|. \quad (17)$$

We assume that $V_\alpha = o(T)$, implying that α_{t+k} varies slowly over time.

This assumption is practically reasonable, as electricity prices typically do not exhibit frequent abrupt changes.

Assumption 3.4. For each distribution sequence $\{Y_{t,i}\}$ and $\{Y_{t,\text{sum}}\}$, we define the total variation distance between consecutive distributions as $V_{Y,i}$, given by

$$V_{Y,i} = \sum_{t=1}^T \frac{1}{2} \int |p_{t+k,i}(y) - p_{t+k-1,i}(y)| dy. \quad (18)$$

We assume that $V_{Y,i} = o(T)$, implying that $Y_{t,i}$ evolves slowly over time.

This assumption permits the time series to be non-stationary, provided that the underlying process exhibits only slow drift over time. Given that wind power time series are often quasi-stationary in practice, this assumption is also well justified.

3.2. Value-oriented Forecast Reconciliation

Our reconciliation model is inspired by [Pennings and Van Dalen \(2017\)](#). Specifically, the reconciliation model f_{RE} is implemented using a structural matrix \mathbf{S} combined with a learnable combination function g . Additionally, we incorporate contextual features, such as market data, represented as $\mathbf{z}_t \in \mathbb{R}^d$, throughout the reconciliation process. As a result, both the base forecasts $\hat{\mathbf{y}}_{t+k|t}$ and the contextual features \mathbf{z}_t serve as inputs to the model f_{RE} . The reconciliation procedure is characterized by

$$\tilde{\mathbf{y}}_{t+k|t} = \mathbf{S}\mathbf{h}_{t+k}, \quad (19)$$

$$\mathbf{h}_{t+k} = g(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta}) \quad (20)$$

where \mathbf{h}_{t+k} denotes the latent variable and $\boldsymbol{\theta}$ symbolizes the parameters within the combination function g . In accordance with [\(12\)](#), \mathbf{h}_{t+k} can also be interpreted as the reconciled bottom-level forecasts. This approach generalizes the widely used reconciliation technique based on linear combinations of base forecasts ([Hollyman et al., 2021](#)), by allowing the input vector \mathbf{z}_t to incorporate information beyond the original time series. Specifically, the learning-based reconciliation process leverages both the time series and external information to produce reconciled forecasts at the bottom level, which are then aggregated to generate forecasts at higher levels. This strategy aligns with the framework proposed by [Bertani et al. \(2025\)](#).

To estimate the parameters $\boldsymbol{\theta}$, machine learning techniques are commonly employed. This involves solving the following optimization problem, with a specific loss function denoted as $l(\cdot, \cdot)$:

$$\min_{\boldsymbol{\theta}} \mathbb{E} [l(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})]. \quad (21)$$

For example, [Ben Taieb and Koo \(2019\)](#) use the widely adopted mean squared error as the loss function, aiming to minimize the total root mean square error. In contrast, here each WPP i has

the cost function $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ and seeks to minimize its own cost. Therefore, for a group of WPPs, they must collectively decide on the parameters $\boldsymbol{\theta}$ within the reconciliation model. In such a multi-agent system, the parameter estimation procedure, following the classical utilitarian principle, entails minimizing the total cost incurred by all WPPs, formally expressed as:

$$\min_{\boldsymbol{\theta}} \sum_{i=1}^m \mathbb{E} \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \right].$$

The parameters obtained from this estimation are referred to as the utilitarian estimate. However, while each WPP aims to minimize its own costs, the sum of costs is neutral with respect to potential inequalities in the distribution of costs among the involved agents. Although weighting factors can be incorporated into the reconciliation process to account for forecasts at each hierarchical level (Van Erven and Cugliari, 2015), there is a lack of well-defined methodologies for determining these weights in the context of decision-making. In contrast, incorporating fairness considerations into the parameter estimation process offers a viable and rigorous alternative (Bertsimas et al., 2011). To achieve this, we use the cooperative bargaining framework to distribute costs fairly among the agents involved.

3.2.1. Nash Bargaining based Reconciliation

In a bargaining game of m agents, the agreement set \mathcal{A} is a set of all agreed parameters $\boldsymbol{\theta}$, whereas the disagreement point \mathbf{d} is the default decision when the group of agents cannot reach a consensus. Therefore, we redefine the loss function for each WPP i as

$$l_i(\boldsymbol{\theta}) = \begin{cases} \mathbb{E} \left[c_{t+k,i}^{(\text{AG})}(f_{\text{RE}}(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta}), \mathbf{y}_{t+k}) \right], & \text{if } \boldsymbol{\theta} \in \mathcal{A}, \\ \mathbb{E} [c_{t+k,i}(\hat{y}_{t+k,i}, y_{t+k,i})], & \text{if } \boldsymbol{\theta} = \mathbf{d}. \end{cases} \quad (22)$$

This implies that when the participating agents reach a consensus, they reconcile their forecasts and engage in collective market trading, thereby sharing the associated costs. Under Assumption 3.1, sub-coalitions are not permitted; thus, in the absence of consensus, each agent trades independently. The disagreement point \mathbf{d} can be assigned an unattainable value, such as an all-zero vector, since the reconciliation of forecasts is no longer required. Specifically, we assume that within the

agreement set, there are points that present Pareto improvements¹ over the disagreement point. If no such points exist, individual rationality will be violated, thereby rendering the design of a reconciliation function meaningless. This assumption is intuitively plausible, as Theorem 2.1 demonstrates how aggregate trading can reduce overall balancing costs. We will show that it holds within the framework of our allocation rule under a few mild conditions.

Particularly, the Nash bargaining framework proposes four axioms, namely Pareto optimality, symmetry, independence of irrelevant alternatives, and invariance to affine transformations (Nash, 1953), which are described as follows:

- (1) Pareto optimality: the solution $\boldsymbol{\theta}^{\text{NB}}$ is Pareto optimal in \mathcal{A} , that is, there does not exist $\boldsymbol{\theta} \in \mathcal{A}$ such that $l_i(\boldsymbol{\theta}) < l_i(\boldsymbol{\theta}^{\text{NB}})$, $\forall i = 1, 2, \dots, m$, and $\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\text{NB}}$.
- (2) Symmetry: the solution $\boldsymbol{\theta}^{\text{NB}}$ is invariant in the order of agents.
- (3) Independence of irrelevant alternatives: if we expand the feasible set \mathcal{A} and the solution stays in \mathcal{A} , then the solution $\boldsymbol{\theta}^{\text{NB}}$ stays the same.
- (4) Invariance to affine transformations: if affine transformations are applied to the costs, the fair allocation in the affine-transformed system equates to the affine transformation of the fair allocation in the original system.

According to Nash (1953), a unique solution exists for the bargaining game. In the context of the cost sharing game defined in this study, the Nash bargaining solution can be derived as follows:

$$\boldsymbol{\theta}^{\text{NB}} = \arg \min_{\boldsymbol{\theta} \in \mathcal{A}} \sum_i^m -\log [l_i(\mathbf{d}) - l_i(\boldsymbol{\theta})] \quad (23a)$$

$$\text{s.t. } l_i(\mathbf{d}) \geq l_i(\boldsymbol{\theta}), \quad i = 1, 2, \dots, m. \quad (23b)$$

The constraint in (23) ensures that the allocated cost for each agent does not exceed the cost incurred under independent offering, as stipulated by Assumption 3.2.

¹Consider a set of objective functions l_1, l_2, \dots, l_m to be minimized with respect to a parameter vector $\boldsymbol{\theta}$. We say that $\boldsymbol{\theta}$ is a Pareto improvement over $\boldsymbol{\theta}' \in \mathcal{A}$, if $l_i(\boldsymbol{\theta}) \leq l_i(\boldsymbol{\theta}')$, $\forall i = 1, 2, \dots, m$, and the inequality is strict for at least one index i .

In fact, neither $\mathbb{E} \left[c_{t+k,i}^{(\text{AG})}(f_{\text{RE}}(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta}), \mathbf{y}_{t+k}) \right]$ nor $\mathbb{E} [c_{t+k}(\hat{y}_{t+k,i}, y_{t+k,i})]$ is available. Following [Ban and Rudin \(2019\)](#), we employ the sample average approximation method to approximate the expectation components in (23) and apply the empirical risk minimization algorithm for parameter estimation. As a result, we convert problem (23) into a constrained learning problem. With a set of historical data $\{\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t, \mathbf{y}_{t+k}\}_{t=1,2,\dots,T}$, we can rewrite (23) as

$$\min_{\boldsymbol{\theta}} \sum_i^m -\log \left[\frac{1}{T} \sum_{t=1}^T \left[c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) - c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \right] \right], \quad (24a)$$

$$\text{s.t.} \quad \frac{1}{T} \sum_{t=1}^T \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) - c_{t+k}(\hat{y}_{t+k,i}, y_{t+k,i}) \right] \leq 0, \quad \forall i = 1, \dots, m, \quad (24b)$$

$$\tilde{\mathbf{y}}_{t+k|t} = \mathbf{S}g(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta}). \quad (24c)$$

Unlike [Ban and Rudin \(2019\)](#), which employs linear models for data-driven decisions, we leverage neural networks to model the combination function g , as described in (19). In the following subsections, we describe the allocation rule used and defer the details of model training to Section 3.3.

3.2.2. Weighted Proportional Allocation Mechanism

We design the cost allocation rule $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ for each WPP within the framework of value-oriented forecast reconciliation. Specifically, we employ the weighted proportional allocation rule introduced by [Anupindi et al. \(2001\)](#). This rule consists of two components: one associated with the bottom-level reconciled forecast $\tilde{y}_{t+k,i|t}$ and the other with the aggregate-level reconciled forecast $\tilde{y}_{t+k,\text{sum}|t}$. The cost $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ is mathematically expressed as

$$c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) = (1 - w) c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) + w \gamma_i c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}). \quad (25)$$

where $0 \leq w \leq 1$, $\gamma_i \geq 0$, and $\sum_i^m \gamma_i = 1$. The term $c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})$ represents the cost associated with the WPP i if it independently submits the pseudo-offer $\tilde{y}_{t+k,i|t}$ in the markets. In contrast, $c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})$ represents the total cost for the aggregation when the PM submits the offer $\tilde{y}_{t+k,\text{sum}|t}$ in the markets.

This allocation rule implies that each WPP reduces a fraction of its own cost associated with the pseudo-offer while simultaneously bearing an equivalent fraction of its share of the overall system cost, $\gamma_i c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})$. When $w = 0$, the PM retains the entire profit from the

aggregate trading and distributes none to the WPPs. Consequently, each WPP acts independently using the reconciled forecast $\tilde{y}_{t+k,i|t}$, as if aggregation does not exist. In contrast, when $w = 1$, the cost incurred by each WPP is represented by $\gamma_i c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})$, which denotes a portion of the system-wide cost. For $w \in (0, 1)$, the weighted proportional allocation rule reflects allocation schemes that fall between these extreme cases.

Specifically, γ_i can be defined according to various criteria, as outlined by [Hartman and Dror \(1996\)](#). For example, one definition uses the proportion of actual power generation², given by:

$$\gamma_i^{\text{ge}} = \frac{y_{t+k,i}}{y_{t+k,\text{sum}}}.$$

Another definition considers γ_i as the proportion of cost associated with pseudo offers, expressed as:

$$\gamma_i^{\text{pc}} = \frac{c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})}{\sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})}.$$

It is observed that these settings of γ_i ensure the allocated cost for WPP i is less than its corresponding cost associated with pseudo offers, as stated below.

Proposition 3.2. Designate γ_i in the cost allocation formula (25) as γ_i^{pc} . Thus, we obtain the following inequality:

$$c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \leq c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}),$$

which ensures that the cost allocation for WPP i is less than or equal to the corresponding cost associated with its pseudo offers.

Proof. See the appendix. □

Proposition 3.3. Define γ_i in the cost allocation (25) as γ_i^{ge} . If the following condition holds:

$$\frac{c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})}{y_{t+k,\text{sum}}} \leq \frac{c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})}{y_{t+k,i}}, \quad \forall i = 1, 2, \dots, m, \quad (26)$$

²When $y_{t+k,\text{sum}}$ is equal to 0, we set γ_i^{ge} as $1/m$.

then it follows that

$$c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \leq c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}).$$

The proof is straightforward and thus omitted here. In practice, the condition (26) is typically satisfied, as demonstrated through simulations and case studies. If the cost $c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})$ for the pseudo offer $\tilde{y}_{t+k,i|t}$ is lower than that of the independent offer $c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i})$, then the aforementioned propositions suggest that the cost experienced by WPP i under the proposed method, represented by $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$, is less than $c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i})$. Nevertheless, it implies a reduction in the cost of independent offers post-reconciliation, which is a strong requirement.

We note that Propositions 3.2 and 3.3 hold for any parameters θ . Particularly, in the context of the BU reconciliation approach, the bottom-level forecasts remain unchanged after reconciliation, i.e., $\tilde{y}_{t+k,i|t} = \hat{y}_{t+k,i|t}$. Consequently, it follows that the shared allocation cost $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ is less than or equal to the cost of independent offering $c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i})$. It is evident that,

$$\mathbb{E} \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \right] \leq \mathbb{E} \left[c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) \right].$$

Based on this result, we derive the following corollary.

Corollary 3.2. By employing the weighted proportional allocation rule, the agreement set of parameters θ in problem (23) is guaranteed to be non-empty.

Given that both $c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})$ and $c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})$ are convex functions, introducing γ_i^{pc} disrupts the convexity of the cost function $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$, as the ratio of convex functions is generally not convex. In contrast, γ_i^{ge} is a constant and thus preserves the convexity of $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$. For this reason, we primarily adopt γ_i^{ge} throughout this work.

3.2.3. Profit Allocation within the Aggregation

To further analyze the profit allocations, we formally define the extra profit from the forecast reconciliation as follows.

Definition 3.2 (Extra profit). Let $\tilde{y}_{t+k,\text{sum}|t}$ be the offer by the PM for the trading hour $t+k$,

and $\hat{y}_{t+k,1|t}, \dots, \hat{y}_{t+k,m|t}$ be a set of m independent offers. Then, the extra profit is defined as

$$R = \sum_{i=1}^m c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) - c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}),$$

where $c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})$ is the cost for the aggregate offer $\tilde{y}_{t+k,\text{sum}|t}$ post-reconciliation, and $c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i})$ is the cost for the independent offer using the base forecast $\hat{y}_{t+k,i|t}$.

This definition calculates the difference between the total costs of independent offers without reconciliation (i.e., disagreement point) and the cost of the aggregate offer by the PM post-reconciliation (i.e., agreement point). In parallel, we calculate the total cost assigned to all WPPs as

$$\sum_{i=1}^m c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) = (1-w) \sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) + w c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}).$$

It is noteworthy that the total cost of WPPs is higher than the actual balancing cost incurred by the PM, but remains lower than the sum of the cost of the pseudo offers, i.e.,

$$c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}) \leq \sum_{i=1}^m c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \leq \sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}), \quad (27)$$

where the inequality results from Proposition 3.1. Therefore, it can be seen that the designed cost allocation rule is not efficient, unless $w = 1$, as we have

$$\sum_{i=1}^m c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) - \sum_{i=1}^m c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \leq R. \quad (28)$$

In other words, the total extra profit allocated to WPPs is lower than the extra profit generated by aggregated trading. Specifically, the remaining extra profit is assigned to the PM as the payoff for their coordination efforts, denoted as R^{pm} , i.e.,

$$R^{\text{pm}} = (1-w) \left[\sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) - c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}) \right]. \quad (29)$$

3.3. Model Training via Empirical Risk Minimization

Indeed, problem (24) can be viewed as a machine learning problem with constraints. As demonstrated by Chamon and Ribeiro (2020), learning with constraints is, in theory, no more difficult

than unconstrained learning, and empirical risk minimization remains an appropriate approach for such problems. Specifically, they proposed a primal–dual gradient algorithm for constrained learning. Building on this, we adapt their framework and derive the empirical Lagrangian function for problem (24), i.e.

$$\begin{aligned}
L(\boldsymbol{\theta}, \boldsymbol{\mu}) = & - \sum_i^m \log \left[\frac{1}{T} \sum_{t=1}^T \left[c_{t+k}(\hat{y}_{t+k,i|t}, \mathbf{y}_{t+k,i}) - c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \right] \right] \\
& + \sum_i^m \mu_i \left[\frac{1}{T} \sum_{t=1}^T \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) - c_{t+k}(\hat{y}_{t+k,i|t}, \mathbf{y}_{t+k,i}) \right] \right]^+, \tag{30}
\end{aligned}$$

where $\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_m]^\top \in \mathbb{R}^m$ collects the dual variables μ_i associated with constraints. The empirical dual problem is then written as

$$\max_{\boldsymbol{\mu}} \min_{\boldsymbol{\theta}} L(\boldsymbol{\theta}, \boldsymbol{\mu}). \tag{31}$$

Obviously, the outer maximization is a convex optimization problem, and its gradient can be easily obtained by evaluating the minimizer of $L(\boldsymbol{\theta}, \boldsymbol{\mu})$. For instance, given the $\boldsymbol{\theta}$, the gradient of μ_i is calculated as

$$\nabla_{\mu_i} = \left[\frac{1}{T} \sum_{t=1}^T \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) - c_{t+k}(\hat{y}_{t+k,i|t}, \mathbf{y}_{t+k,i}) \right] \right]^+,$$

where $\tilde{\mathbf{y}}_{t+k|t} = \mathbf{S}g(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta})$. Hence, the main challenge of problem (31) lies in the inner minimization. However, abundant evidence has shown that the gradient descent algorithm can learn good parameters for neural networks (Shalev-Shwartz and Ben-David, 2014). Subsequently, we can iteratively minimize the Lagrangian with respect to $\boldsymbol{\theta}$ while keeping $\boldsymbol{\mu}$ constant, followed by adjusting the dual variables using the obtained minimizer.

We now analyze the convexity of the Lagrangian function. The convexity of the second component of $L(\boldsymbol{\theta}, \boldsymbol{\mu})$ is straightforward, as $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ is convex in $\tilde{\mathbf{y}}_{t+k|t}$. For the first component of $L(\boldsymbol{\theta}, \boldsymbol{\mu})$, we note that $-\log(\cdot)$ is convex and decreasing, while $-c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ is concave in $\tilde{\mathbf{y}}_{t+k|t}$. Since the combination of a convex, non-increasing function with a concave function preserves convexity (Boyd and Vandenberghe, 2004), it follows that the first term in $L(\boldsymbol{\theta}, \boldsymbol{\mu})$ is also convex. Therefore, the Lagrangian function $L(\boldsymbol{\theta}, \boldsymbol{\mu})$ is convex in $\tilde{\mathbf{y}}_{t+k|t}$, enabling the use of the

Algorithm 1 Primal-dual based iterative learning algorithm

Input: Optimizer and Lagrangian step sizes, $\lambda, (\nu_1, \nu_2, \dots, \nu_m)$

Initialize variables: $\mu_i^0 \leftarrow 1, \forall i = 1, 2, \dots, m$

for epoch $\eta = 1, 2, \dots$ **do**

 Randomly sample a batch of samples $\{\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t, \mathbf{y}_{t+k}\}_{t \in \mathcal{T}_b}$

$\tilde{\mathbf{y}}_{t+k|t} = \mathbf{S}g(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t; \boldsymbol{\theta}^{\eta-1}), \forall t \in \mathcal{T}_b$

$\boldsymbol{\theta}^\eta \leftarrow \boldsymbol{\theta}^{\eta-1} - \lambda \nabla_{\boldsymbol{\theta}} L^{(b)}(\boldsymbol{\theta}, \boldsymbol{\mu}^{\eta-1})$

$\mu_i^\eta \leftarrow \mu_i^{\eta-1} + \nu_i \left[\frac{1}{|\mathcal{T}_b|} \sum_{t \in \mathcal{T}_b} \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) - c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) \right] \right]^+$

end for

stochastic gradient descent algorithm to update $\boldsymbol{\theta}$ (Shalev-Shwartz and Ben-David, 2014). The explicit form of $L^{(b)}(\boldsymbol{\theta}, \boldsymbol{\mu})$ is deferred to the Appendix. Letting $\boldsymbol{\mu}^\eta$ and $\boldsymbol{\theta}^\eta$ denote the parameter estimates at the η iteration, the iterative learning procedure is summarized in Algorithm 1.

Remark 3.1. Note that the loss function is non-stationary. However, under Assumptions 3.3 and 3.4, Algorithm 1 is guaranteed to converge sublinearly, as established by Theorem 1 by Besbes et al. (2015).

4. Simulations

In this section, we present several illustrative examples to validate the proposed approach using data from the Wind Toolkit³. Specifically, we have collected hourly datasets from four nearby wind farms, with capacities of 1.7496 MW, 2.9646 MW, 3.3777 MW, and 2.5272 MW, covering the period from 2007 to 2008. In particular, we follow the framework of Muñoz et al. (2023), which examines an hourly online forward market followed by a balancing market with a dual-price settlement mechanism. The lead time, k , for both the base forecast and forecast reconciliation process is set to 1. The dataset is split into two segments: the first 80% is used for model training, while the remaining 20% is reserved for out-of-sample forecasting validation. In the following, we will introduce the setup and present the simulation results.

4.1. Experimental Setups

In our simulation studies, we investigate a simplified scenario where π_{t+k}^{F} is fixed at 25 €/MWh, ψ_{t+k}^+ is set to 12 €/MWh, and ψ_{t+k}^- is maintained at 4 €/MWh. This simplified setup is used to

³<https://www.nrel.gov/grid/wind-toolkit.html>

gain insight into the behavior of the proposed method, with a more practical case study deferred to Section 5. In this context, the optimal offering strategy is to use the 0.75 quantile of the wind power distributions. Specifically, this represents the independent optimal offering strategy for each WPP and the PM. Therefore, optimal offers can be constructed using quantiles, which represent a specific type of point forecast (Gneiting, 2011). Moreover, as point forecasts derived from mean regression are still widely used in practice, we consider both means and quantiles as base forecasts in our framework. While WPPs may act strategically by explicitly communicating quantiles, we assume that the PM remains agnostic to the specific form of the base forecasts. A description of the cases is provided below.

Case 1: Forecasts at both the aggregate and leaf levels are communicated as means, derived from mean regression models.

Case 2: Base forecasts for both the aggregate and leaf levels are communicated as quantiles, derived from quantile regression models.⁴

In this study, we evaluate various offering strategies, which include both independent and aggregate offering using reconciliation. The offering strategies under consideration are described as follows: **(1) Independent offering:** Each WPP submits offers based on their individual forecasts (i.e., the base forecasts). Thus, neither reconciliation nor the involvement of a PM is considered. We also include 3 aggregate offering strategies, in which the reconciled aggregate forecasts are used by the PM for offering. Concretely, the considered reconciliation approaches are described as follows. **(2) Bottom-up reconciliation:** The base forecasts at the leaf level are retained unchanged and then aggregated to generate the overall forecast. **(3) Quality-oriented reconciliation:** The base forecasts are reconciled within a learning-based model (19), where the loss function is defined as the mean squared error. **(4) Value-oriented reconciliation:** The base forecasts are reconciled using the learning-based model (19), with the training procedure outlined in Section 3. In addition, we establish linear counterparts for both quality-oriented and value-oriented reconciliation in Case 1, where the quality-oriented model reduces to the framework proposed by Hyndman et al. (2011).

It is important to note that all of the previously discussed aggregate offering strategies employ the proportional allocation rule and set γ_i as γ_i^{ge} , as outlined in Section 3.2. For both quality-

⁴Bottom-up and value-oriented reconciliation methods are behaviorally plausible but decision-theoretically sub-optimal in this case, which is the tension we intend to highlight.

oriented and value-oriented reconciliation models, we include the lagged wind power generation at both the leaf and aggregate levels as contextual features. To evaluate the performance of the offering strategies across a set of samples $\{(\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t, \mathbf{y}_{t+k})\}_{t \in \mathcal{T}_{\text{test}}}$, we compute the average profit (AP) for each WPP over the test set. Specifically, for independent offering, using ρ_{t+k} defined by (2), the average profit for WPP i is given by

$$\text{AP}_i = \frac{1}{|\mathcal{T}_{\text{test}}|} \sum_{t \in \mathcal{T}_{\text{test}}} \rho_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}).$$

In contrast, for portfolio-based aggregate offering, using $\rho_{t+k,i}^{(\text{AG})}$ from (16), the average profit for WPP i is given by

$$\text{AP}_i = \frac{1}{|\mathcal{T}_{\text{test}}|} \sum_{t \in \mathcal{T}_{\text{test}}} \rho_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}).$$

4.2. Case 1

We investigate the range of weights in the allocation rule, which varies between 0 and 1, and present the average profit over independent offers for each WPP across various models in Figure 2. Overall, each WPP's profit under independent offerings is invariant to w , whereas profits under reconciliation-based strategies increase as w rises. We have confirmed that the unit balancing cost condition (26) holds in this simulation. As expected by Proposition 3.3, the function $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ decreases monotonically with increasing w , which explains the observed monotonic increase in profits with respect to w . Specifically, within the frameworks of both BU reconciliation and quality-oriented reconciliation, profits are directly proportional to w . Notably, the quality-oriented reconciliation training is independent of w , and similarly, BU reconciliation does not require training. Thus, the balancing cost terms in (25) remain unaffected by changes in w , leading to a linear relationship.

Additionally, nearly every reconciliation-based model results in higher profits for all WPPs, except in the case of quality-focused forecast reconciliation when w equals 0 or 0.1. Given that the reconciled forecasts at the bottom level remain unchanged in the BU reconciliation framework, it can be inferred that the profits exceed those from independent offers when condition (26) is satisfied. Comparing the profits from BU reconciliation to those from quality-oriented reconciliation reveals that quality-based reconciliation generally yields higher profits, except for WPP 4. This outcome is

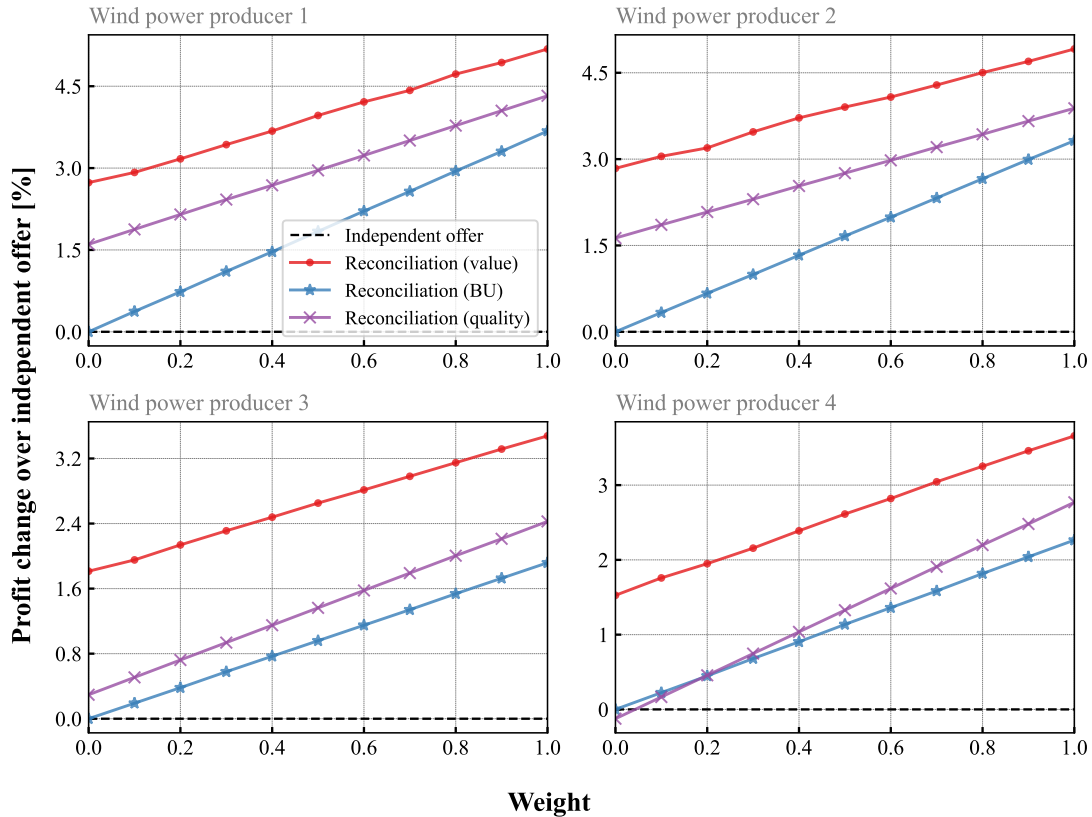


Figure 2: Average profit change over independent offers for each wind power producer across different models in Case 1

attributed to the ability of quality-oriented reconciliation to reduce the overall mean squared errors more effectively than BU reconciliation, as supported by existing research (Van Erven and Cugliari, 2015). However, the profits for WPP 4 suggest that quality-oriented reconciliation might lead to an unfair allocation of costs. Specifically, the value-oriented forecast reconciliation generates greater profits compared to other models because its loss function in the reconciliation process aligns with balancing costs. Evidently, the gains from aggregated trading differ across WPPs. In particular, WPPs 1 and 2 realize larger percentage profit increases relative to their independent offerings. This is consistent with the Nash product objective, which equalizes and safeguards excess profits over the independent baseline but does not enforce identical percentage gains. Because independent-offer profits differ across WPPs, even identical absolute excess profits translate into different percentage changes.

We present the results under both learning-based and linear reconciliation methods in Figure 3.

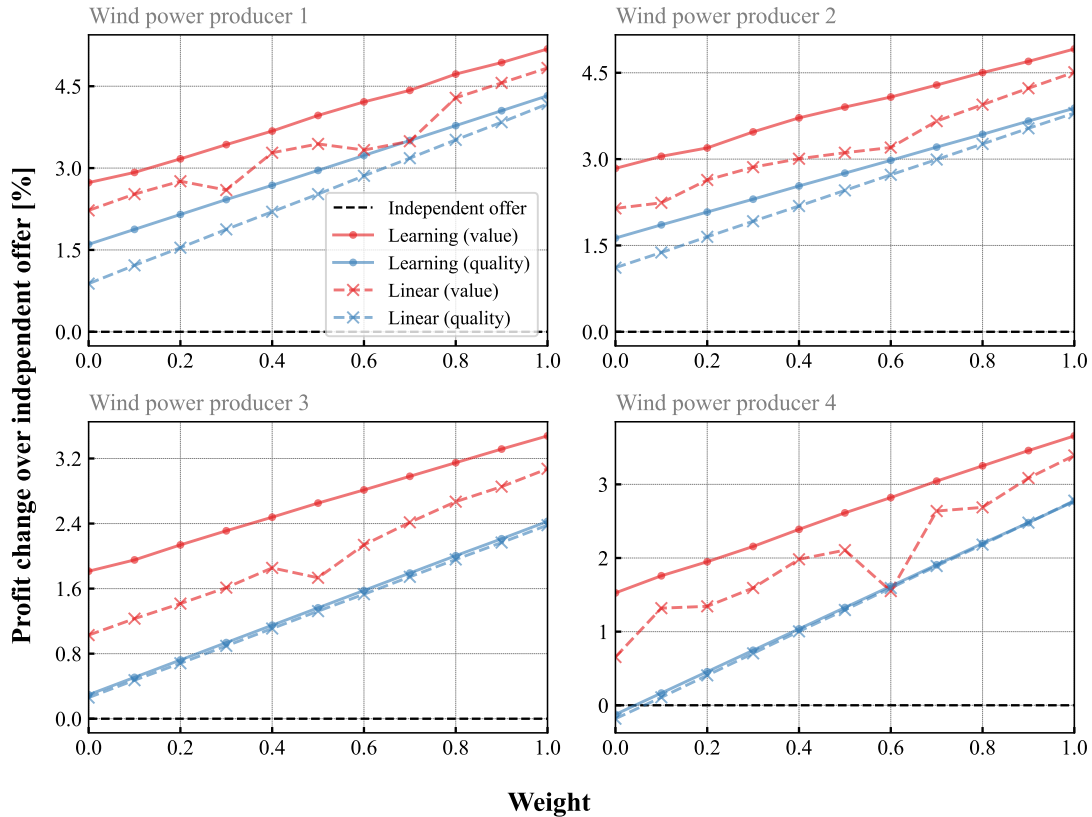


Figure 3: Average profit change over independent offers for each wind power producer under learning-based and linear reconciliation in Case 1

Overall, value-oriented reconciliation yields higher profits for each WPP compared to quality-oriented reconciliation, supporting the claim that value-oriented reconciliation is agnostic to the choice of reconciliation model. Notably, the performance of quality-oriented reconciliation is similar across both learning-based and linear implementations, highlighting the effectiveness of the optimal reconciliation approach proposed by Hyndman et al. (2011). In contrast, for value-oriented reconciliation, the learning-based model outperforms its linear counterpart in both performance and stability. This is expected, as value-oriented reconciliation involves a more complex training procedure, and neural networks offer greater flexibility in function approximation.

4.3. Case 2

We then analyze the average profit of each WPP when base forecasts are communicated as quantiles. The analysis considers varying the weight w in the allocation rule from 0 to 1, with the results shown in Figure 4. Similar to Case 1, the average profits for each WPP increase as w

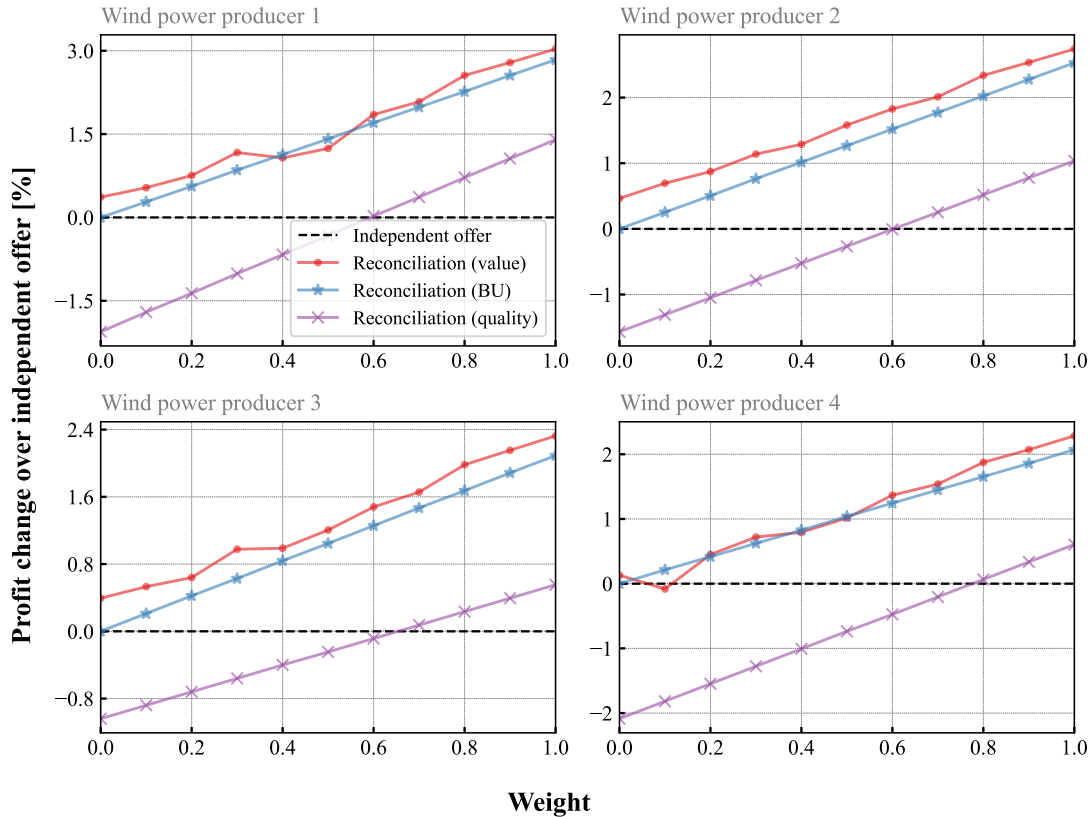


Figure 4: Average profit change over independent offers for each wind power producer across different models in Case 2

risers, although there are a few exceptions under the value-oriented forecast reconciliation approach. Provided that the unit balancing cost condition (26) is satisfied, the profits from BU reconciliation consistently exceed those from independent offering. For quality-oriented reconciliation, the results indicate that the resulting profits are lower than those achieved through strategic independent offering in most cases. This highlights a key limitation of quality-oriented reconciliation when applied in a decision-making context. For the proposed approach, WPP 4 is an exception when $w = 0.1$, primarily due to the challenges in model training at lower values of w . In fact, the base forecast profits in this case are substantial, suggesting that the agreement set during bargaining is minimal, which in turn complicates the training process. With more training samples, it is expected to attain estimated parameters within the agreement set. Nevertheless, the profits from the value-oriented forecast reconciliation approach tend to exceed those from BU reconciliation, owing to the information gain discussed in Section 3.2.2.

Table 1: RMSE of hierarchical forecasts in Cases 1 and 2 (MW)

Case	Method	Weight									
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
1	Reconciliation (value)	0.698	0.685	0.672	0.685	0.681	0.676	0.668	0.668	0.666	0.878
	Reconciliation (quality)	0.605	0.598	0.605	0.605	0.599	0.608	0.610	0.610	0.598	0.607
2	Reconciliation (value)	0.689	0.695	0.692	0.689	0.689	0.688	0.682	0.679	0.674	0.806
	Reconciliation (quality)	0.601	0.605	0.602	0.604	0.603	0.600	0.608	0.612	0.606	0.602

We compare the absolute average profits resulting from the proposed value-oriented forecast reconciliation approach in both scenarios: when the base forecasts are either mean forecasts or quantile forecasts. The results are presented in Figure 5, where the ‘m’ and ‘q’ in the legend indicate that the base forecasts are mean or quantile. The analysis shows that, when the base forecasts are mean forecasts, the profits generated by the proposed method exceed those of the alternative approach, even though profits under independent offering remain lower. Notably, when w is either 0.9 or 1.0, the profits obtained with the proposed method are almost identical in both scenarios. This finding reinforces the earlier discussion by suggesting that the agreement set becomes more restrictive when base forecasts are quantile forecasts, thereby complicating model training for smaller values of w . In contrast, this implies that wind power producers need only to provide their mean forecasts to the PM, which reduces their need to analyze the electricity markets. Furthermore, we report the RMSE of hierarchical forecasts for both cases in Tables 1. As observed, the forecast errors under the value-oriented reconciliation approach are consistently higher than those obtained from the quality-oriented reconciliation. When considered alongside the profit analysis, these results support our central claim: minimizing overall forecast errors does not necessarily lead to cost reductions for individual agents. Specifically, when the weight w is set to 1.0, the RMSE under the value-oriented reconciliation approach is relatively high, as the loss function exclusively prioritizes the aggregate level.

5. Real-world Case Study

Similar to Section 4, we also have an hourly forward market followed by a balancing market with a dual-price settlement mechanism. However, unlike the simulation studies, we obtain two years of hourly price data from the Danish TSO data portal⁵. The hourly penalties are calcu-

⁵<https://www.energidataservice.dk/>




Figure 5: Average profits for each wind power producer in cases where base forecasts are either mean or quantile forecasts

lated using the day-ahead spot and regulation prices, as specified in Section (2). We consider the four models described in Section 4. As noted in the previous section, the mean forecasts suffice as the base forecasts; therefore, we define the base forecasts as the mean forecasts. Specifically, the quality-oriented reconciliation model incorporates lagged wind power generation values as contextual features. While for the value-oriented reconciliation model, we include lagged penalties $\psi_{t-h}^+, \dots, \psi_t^+, \psi_{t-h}^-, \dots, \psi_t^-$ besides wind power generation values as contextual features. In particular, the weight w in the allocation rule is set to 0.9. For both quality-oriented and value-

Table 2: Average profits for each wind power producer under different models [€/hour]

Model	WPP 1	WPP 2	WPP 3	WPP 4
Independent offering	15.69	28.05	33.54	23.35
BU reconciliation	15.81	28.25	33.71	23.43
Quality-oriented reconciliation	15.83 ± 0.006	28.27 ± 0.012	33.71 ± 0.016	23.45 ± 0.011
Value-oriented reconciliation	15.92 ± 0.001	28.45 ± 0.004	33.95 ± 0.003	23.62 ± 0.002

oriented reconciliation, we repeat the experiments 10 times and report both the mean and standard deviation of the results.

We present the average profits for each wind power producer in Table 2. The aggregate offering models generate higher average profits compared to independent offering, aligning with the findings from simulation studies. Specifically, BU reconciliation and quality-oriented reconciliation exhibit similar performance, although the latter provides more accurate forecasts. As a result, the profit increase is primarily attributed to participation in aggregate offering. This supports our argument that quality-oriented reconciliation may not be the most effective approach for evaluating forecast accuracy. As expected, value-oriented reconciliation yields the highest average profits for each wind power producer, as it facilitates the development of a profitable trading strategy.

6. Concluding Remarks

Although value is considered a critical factor in evaluating forecasts, alongside quality (Murphy, 1993), it has been largely overlooked in the forecast reconciliation literature. Generally, forecast reconciliation can be viewed as an information pool, where a group of agents pools their information to make collective decisions. We emphasize the importance of integrating forecast value into the reconciliation process, demonstrating that quality-oriented approaches may fail to produce decisions that are fair and acceptable to all agents. This issue becomes particularly pronounced when agents with heterogeneous loss functions are involved, as quality-oriented reconciliation may lead to individual revenue losses, which in turn can cause potential conflicts. Consequently, we propose a value-oriented forecast reconciliation approach based on the Nash bargaining framework. Specifically, we frame the problem as a learning-based reconciliation problem, where parameters are estimated using a primal-dual algorithm.

We introduce the concept of consensus on information, distinguishing it from the well-established notion of consensus on allocations, which serves as a prerequisite for multi-agent cooperation. We demonstrate the effectiveness of the proposed value-oriented reconciliation approach through an aggregated wind energy trading problem, where profits are distributed among wind power producers according to a weighted proportional allocation rule. Furthermore, we prove that by employing this allocation rule, the agreement set of parameters within the Nash bargaining game is guaranteed to be non-empty. Given the estimated reconciliation and allocation mechanism, it is challenging

for individual WPPs to explicitly derive the functional form associated with their optimal offers. Instead, our framework aggregates the available information and reconciles the forecasts to determine an optimal offer for each WPP in a data-driven manner. Since the reconciled forecasts are directly used as decisions in the market, this approach may also be interpreted as a form of *decision reconciliation*. Specifically, we design simulations and case studies that consider an hourly forward market and a balancing market with a dual-price mechanism, utilizing market data from Denmark and wind power data from the Wind Toolkit. The results show that, under specific conditions, the proposed method ensures that the profits achieved through the aggregate offering approach consistently outperform those obtained through independent offering. This gain results from both information pooling and cooperative trading. Notably, the findings suggest that, in terms of subsequent decision-making value, quality-focused forecast reconciliation may not be the optimal choice.

However, we note that the allocation rule employed in this study does not lie within the core of the cooperative trading game. Furthermore, the proposed rule does not ensure that existing agents in the coalition remain unaffected by the inclusion of new members—an essential property referred to as the population monotonic allocation scheme (PMAS) (Sprumont, 1990). Future work should aim to develop allocation mechanisms that both lie within the core and support coalition expansion without disadvantaging any participant. Indeed, the proposed framework is not limited to energy trading; it is also applicable to various domains in supply chain management where cooperation among agents is essential, such as in retail operations. Although this study focuses primarily on point forecast reconciliation, the proposed framework is readily adaptable to broader group decision-making scenarios involving information pooling. Related algorithms may also be designed for reconciling probabilistic forecasts or aggregating multiple forecast sources. Another promising avenue for future research is the incorporation of fairness principles into value-oriented forecasting in more general and diverse application domains.

Appendix

Proof of Proposition 3.1

Proof. Using the definition of balancing cost,

$$\begin{aligned}
c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}}) &= \psi_{t+k}^+ [y_{t+k,\text{sum}} - \tilde{y}_{t+k,\text{sum}|t}]^+ + \psi_{t+k}^- [\tilde{y}_{t+k,\text{sum}|t} - y_{t+k,\text{sum}}]^+ \\
&\leq \psi_{t+k}^+ \sum_{i=1}^m [y_{t+k,i} - \tilde{y}_{t+k,i|t}]^+ + \psi_{t+k}^- \sum_{i=1}^m [\tilde{y}_{t+k,i|t} - y_{t+k,i}]^+ \\
&= \sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}).
\end{aligned}$$

where the inequality results from the sub-additivity property of $[x]^+$ for any $x \in \mathbb{R}$. \square

Proof of Proposition 3.2

Proof. Using the definition of $c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k})$ and γ_i^{pc} , we have

$$\begin{aligned}
c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) &= (1-w)c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) + w \frac{c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})c_{t+k}(\tilde{y}_{t+k,\text{sum}|t}, y_{t+k,\text{sum}})}{\sum_{i=1}^m c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i})} \\
&\leq (1-w)c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) + wc_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}) = c_{t+k}(\tilde{y}_{t+k,i|t}, y_{t+k,i}),
\end{aligned}$$

where the inequality results from Proposition 3.1. \square

Batch form of the Lagrangian function

Using a batch of samples $\{\hat{\mathbf{y}}_{t+k|t}, \mathbf{z}_t, \mathbf{y}_{t+k}\}_{t \in \mathcal{T}_b}$ where \mathcal{T}_b represents the set of indices, we write the Lagrangian of batch samples as

$$\begin{aligned}
L^{(b)}(\boldsymbol{\theta}, \boldsymbol{\mu}) &= - \sum_i^m \log \left[\frac{1}{|\mathcal{T}_b|} \sum_{t \in \mathcal{T}_b} \left[c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) - c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) \right] \right] \\
&\quad + \sum_i^m \mu_i \left[\frac{1}{|\mathcal{T}_b|} \sum_{t \in \mathcal{T}_b} \left[c_{t+k,i}^{(\text{AG})}(\tilde{\mathbf{y}}_{t+k|t}, \mathbf{y}_{t+k}) - c_{t+k}(\hat{y}_{t+k,i|t}, y_{t+k,i}) \right] \right]^+, \tag{32}
\end{aligned}$$

where $|\mathcal{T}_b|$ represents the cardinality of \mathcal{T}_b .

Acknowledgment

We are grateful for the detailed and constructive suggestions from three anonymous reviewers and the editor. Honglin Wen is funded by the National Science Foundation of China (52307119). Pierre Pinson acknowledges the support of UKRI through the Global NSF-UKRI Centre EPICS (Electric Power Innovation for a Carbon-free Society Centre – EPICS-UK; grant no. EP/Y025946/1).

References

- Anupindi, R., Bassok, Y., and Zemel, E. (2001). A general framework for the study of decentralized distribution systems. *Manufacturing & Service Operations Management*, 3(4):349–368.
- Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., and Panagiotelis, A. (2024). Forecast reconciliation: A review. *International Journal of Forecasting*, 40(2):430–456.
- Athanasopoulos, G. and Kourentzes, N. (2023). On the evaluation of hierarchical forecasts. *International Journal of Forecasting*, 39(4):1502–1511.
- Baeyens, E., Bitar, E. Y., Khargonekar, P. P., and Poolla, K. (2013). Coalitional aggregation of wind power. *IEEE Transactions on Power Systems*, 28(4):3774–3784.
- Ban, G.-Y. and Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108.
- Barrow, D. K. and Kourentzes, N. (2016). Distributions of forecasting errors of forecast combinations: implications for inventory management. *International Journal of Production Economics*, 177:24–33.
- Ben Taieb, S. and Koo, B. (2019). Regularized regression for hierarchical forecasting without unbiasedness conditions. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1337–1347.
- Bertani, N., Satopää, V., and Jensen, S. (2025). Joint bottom-up method for probabilistic forecasting of hierarchical time series. *Operations Research*, in press/available online.
- Bertsimas, D., Farias, V. F., and Trichakis, N. (2011). The price of fairness. *Operations research*, 59(1):17–31.
- Bertsimas, D. and Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044.
- Besbes, O., Gur, Y., and Zeevi, A. (2015). Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244.
- Bitar, E. Y. (2011). *Selling Random Energy*. PhD thesis, UC Berkeley.
- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- Bremnes, J. B. (2004). Probabilistic wind power forecasts using local quantile regression. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 7(1):47–54.
- Carriere, T. and Kariniotakis, G. (2019). An integrated approach for value-oriented energy forecasting and data-driven decision-making application to renewable energy trading. *IEEE transactions on smart grid*, 10(6):6933–6944.

- Chamon, L. and Ribeiro, A. (2020). Probably approximately correct constrained learning. *Advances in Neural Information Processing Systems*, 33:16722–16735.
- Choi, T.-M. (2012). *Handbook of Newsvendor problems: Models, extensions and applications*, volume 176. Springer Science & Business Media.
- Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S., and Pontil, M. (2018). Empirical risk minimization under fairness constraints. *Advances in neural information processing systems*, 31.
- Edwards, J. B. and Orcutt, G. H. (1969). Should aggregation prior to estimation be the rule? *The Review of Economics and Statistics*, pages 409–420.
- Elmachtoub, A. N. and Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1):9–26.
- Eppen, G. D. (1979). Note—effects of centralization on expected costs in a multi-location newsboy problem. *Management science*, 25(5):498–501.
- Freire, L., Street, A., Lima, D. A., and Barroso, L. A. (2014). A hybrid milp and benders decomposition approach to find the nucleolus quota allocation for a renewable energy portfolio. *IEEE Transactions on Power Systems*, 30(6):3265–3275.
- Gneiting, T. (2011). Quantiles as optimal point forecasts. *International Journal of forecasting*, 27(2):197–207.
- Guerrero-Mestre, V., de la Nieta, A. A. S., Contreras, J., and Catalao, J. P. (2015). Optimal bidding of a group of wind farms in day-ahead markets through an external agent. *IEEE Transactions on Power Systems*, 31(4):2688–2700.
- Hartman, B. C. and Dror, M. (1996). Cost allocation in continuous-review inventory models. *Naval Research Logistics (NRL)*, 43(4):549–561.
- Hartman, B. C., Dror, M., and Shaked, M. (2000). Cores of inventory centralization games. *Games and Economic Behavior*, 31(1):26–49.
- Hollyman, R., Petropoulos, F., and Tipping, M. E. (2021). Understanding forecast reconciliation. *European Journal of Operational Research*, 294(1):149–160.
- Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., and Shang, H. L. (2011). Optimal combination forecasts for hierarchical time series. *Computational statistics & data analysis*, 55(9):2579–2589.
- Hyndman, R. J., Lee, A. J., and Wang, E. (2016). Fast computation of reconciled forecasts for hierarchical and grouped time series. *Computational statistics & data analysis*, 97:16–32.
- Jeon, J., Panagiotelis, A., and Petropoulos, F. (2019). Probabilistic forecast reconciliation with applications to wind power and electric load. *European Journal of Operational Research*, 279(2):364–379.
- Kourentzes, N. (2013). Intermittent demand forecasts with neural networks. *International Journal of Production Economics*, 143(1):198–206.
- Kourentzes, N., Trapero, J. R., and Barrow, D. K. (2020). Optimising forecasting models for inventory planning. *International Journal of Production Economics*, 225:107597.
- Lodi, A., Sankaranarayanan, S., and Wang, G. (2024). A framework for fair decision-making over time with time-invariant utilities. *European Journal of Operational Research*, 319(2):456–467.
- Montrucchio, L., Norde, H., Özen, U., Scarsini, M., and Slikker, M. (2012). *Cooperative newsvendor games: A review*. Springer.

- Morales, J. M., Conejo, A. J., Madsen, H., Pinson, P., and Zugno, M. (2013). *Integrating renewables in electricity markets: operational problems*, volume 205. Springer Science & Business Media.
- Muñoz, M. A., Pinson, P., and Kazempour, J. (2023). Online decision making for trading wind energy. *Computational Management Science*, 20(1):33.
- Murphy, A. H. (1993). What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and forecasting*, 8(2):281–293.
- Nash, J. (1953). Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, pages 128–140.
- Nguyen, H. T. and Le, L. B. (2018). Sharing profit from joint offering of a group of wind power producers in day ahead markets. *IEEE Transactions on Sustainable Energy*, 9(4):1921–1934.
- Panagiotelis, A., Athanasopoulos, G., Gamakumara, P., and Hyndman, R. J. (2021). Forecast reconciliation: A geometric view with new insights on bias correction. *International Journal of Forecasting*, 37(1):343–359.
- Pennings, C. L. and Van Dalen, J. (2017). Integrated hierarchical forecasting. *European Journal of Operational Research*, 263(2):412–418.
- Pinson, P., Chevallier, C., and Kariniotakis, G. N. (2007). Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE transactions on Power Systems*, 22(3):1148–1156.
- Sanguri, K., Shankar, S., Punia, S., and Patra, S. (2022). Hierarchical container throughput forecasting: The value of coherent forecasts in the management of ports operations. *Computers & Industrial Engineering*, 173:108651.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- Sprangers, O., Wadman, W., Schelter, S., and de Rijke, M. (2024). Hierarchical forecasting at scale. *International Journal of Forecasting*, 40(4):1689–1700.
- Sprumont, Y. (1990). Population monotonic allocation schemes for cooperative games with transferable utility. *Games and Economic behavior*, 2(4):378–394.
- Stratigakos, A., Camal, S., Michiorri, A., and Kariniotakis, G. (2022). Prescriptive trees for integrated forecasting and optimization applied in trading of renewable energy. *IEEE Transactions on Power Systems*, 37(6):4696–4708.
- Stratigakos, A., Pineda, S., and Morales, J. M. (2024). Decision-focused linear pooling for probabilistic forecast combination. *International Journal of Forecasting*, in press/available online.
- Tastu, J., Pinson, P., Trombe, P.-J., and Madsen, H. (2013). Probabilistic forecasts of wind power generation accounting for geographically dispersed information. *IEEE Transactions on Smart Grid*, 5(1):480–489.
- Theodorou, E., Spiliotis, E., and Assimakopoulos, V. (2025). Forecast accuracy and inventory performance: Insights on their relationship from the m5 competition data. *European Journal of Operational Research*, in press/available online.
- Van Erven, T. and Cugliari, J. (2015). Game-theoretically optimal reconciliation of contemporaneous hierarchical time series forecasts. In *Modeling and stochastic learning for forecasting in high dimensions*, pages 297–317. Springer.
- Wickramasuriya, S. L., Athanasopoulos, G., and Hyndman, R. J. (2019). Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *Journal of the American Statistical Association*, 114(526):804–819.

- Zellner, A. (1969). On the aggregation problem: A new approach to a troublesome problem. In *Economic models, estimation and risk programming: Essays in honor of Gerhard Tintner*, pages 365–374. Springer.
- Zhang, Y. and Dong, J. (2018). Least squares-based optimal reconciliation method for hierarchical forecasts of wind power generation. *IEEE Transactions on Power Systems*, in press/available online.
- Zhang, Y., Jia, M., Bian, Y., Wen, H., and Shi, Y. (2025). Toward value-oriented renewable energy forecasting: An iterative learning approach. *IEEE Transactions on Smart Grid*, 16(2).