



HAL
open science

Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy

Akylas Stratigakos, Simon Camal, Andrea Michiorri, Georges Kariniotakis

► **To cite this version:**

Akylas Stratigakos, Simon Camal, Andrea Michiorri, Georges Kariniotakis. Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy. 2021. hal-03330017v1

HAL Id: hal-03330017

<https://hal.science/hal-03330017v1>

Preprint submitted on 31 Aug 2021 (v1), last revised 11 Jan 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy

Akylas Stratigakos, Simon Camal, Andrea Michiorri, and Georges Kariniotakis, *Senior Member, IEEE*

Abstract—Deriving decisions from data typically involves a sequential process with two components, forecasting and optimization. Forecasting models learn by minimizing a loss function that stands as a proxy for task-specific costs (e.g. trading, scheduling) without considering the downstream optimization problem, which in practice creates a performance bottleneck and obscures the impact of data on decisions. This work suggests leveraging the structure of the optimization component and directly learning a policy conditioned on explanatory data, effectively proposing a single data-driven module. For this purpose, we describe an algorithm to train ensembles of decision trees by directly minimizing task-specific costs, and prescribe decisions via a weighted Sample Average Approximation of the original problem. We then develop a generic framework to assess the impact of explanatory data on optimization efficacy. The proposed method is validated on the case of trading renewable energy in a day-ahead electricity market, where we design hybrid policies that balance optimal trading decisions and predictive accuracy. The empirical results demonstrate improved performance compared to solutions derived under the standard stochastic optimization framework. Further, we provide valuable insights on how explanatory data impact optimization performance and how this impact evolves under different market designs.

Index Terms—Data-driven optimization, decision trees, energy forecasting, energy trading, prescriptive analytics.

I. INTRODUCTION

ALMOST every aspect of power systems, including real-time operation, investment planning, and electricity markets, is subject to uncertainty stemming from multiple sources, such as stochastic renewable production and unknown market conditions. Facilitated by the influx of data associated with the transition towards the smart grid era, recent years have seen a burgeoning development of energy analytics tools that enable stakeholders to derive decisions in the presence of contextual information. Typically, data-driven decision-making comprises a sequential process with two components, (energy) forecasting and optimization. The first involves forecasting uncertain quantities, e.g. renewable production or market prices, at a future time interval and in a form that ranges from deterministic (point) forecasts to probabilistic forecasts, scenarios (trajectories) and uncertainty sets. Next, these forecasts are used as inputs in an optimization problem to derive optimal decisions (prescriptions).

In the context of statistical learning, forecasting models are trained by minimizing a surrogate loss function that serves

as a proxy for a task-specific cost. However, several issues arise with this approach, as moving from a prediction to a prescription might not be straightforward. Assessing the impact of forecasts on decision costs, i.e., *forecast value*, is considered to be one of the key challenges in energy forecasting in the coming years [1]. Further, directly optimizing towards forecast value rather than accuracy is identified as a high-leverage objective to employ machine learning as means of tackling climate change [2]. Studies on the economic impact of price forecasting errors [3] confirm that increased accuracy does not always translate into increased value, as the latter heavily depends on the specific task. In an early attempt [4] dating back to the late 1990s, the authors propose a metric to quantify operational cost savings for an autonomous system stemming from employing advanced forecasting tools instead of naive benchmarks. This is roughly analogous to comparing the value of stochastic optimization against the certainty equivalent solution [5], although the focus is placed on the forecasting rather than on the optimization component. A recently observed trend suggests moving beyond the simple statistical evaluation of prediction errors to assessing the quality of decisions obtained for different applications. For example, [6] proposes a multivariate probabilistic forecasting model and considers the economic benefits for an electricity retailer as a means of assessing the forecast value. In [7], a stochastic unit commitment problem is employed to evaluate a spatio-temporal scenario generation methodology. This trend highlights two pertinent issues that serve as our current motivation. First, it is pivotal for the forecasting model to be aware of the structure of the downstream optimization problem in order to maximize its value. Second, deploying multiple analytic tools in sequence increases the model chain complexity and obfuscates the impact of explanatory data on the efficacy of decisions.

This work addresses both of these issues by describing a single data-driven model for policy learning, allowing decisions to vary as a function of explanatory data. The direct mapping of data to decisions also enables us to assess the impact of the former to the efficacy of the optimization, while at the same time simplifying the overall model chain.

A. Literature Review

Classical stochastic optimization assumes that uncertain parameters have known distributions. In practice, however, only observational data are available. The problem of single-period

conditional stochastic optimization in the presence of explanatory data (features) has received a lot of attention recently, with proposed solutions working on the nexus of machine learning and operations research. Three general approaches are identified: i) modifying the optimization component, ii) modifying how the forecasting model is trained, iii) directly forecasting the decisions. For the first approach, [8] establishes a prescriptive analytics framework for data-driven decisions via a weighted Sample Average Approximation (SAA) [9] of the original problem, employing predictive machine learning algorithms. In [10] the residuals induced by the SAA solution are used to infer decision uncertainty. Integrated approaches merge the forecasting and optimization component, aiming either at learning a policy from data [11] or selecting from a set of discrete actions [12]. The second stream of literature considers training forecasting models under alternative loss functions in order to explicitly minimize downstream costs [13]. However, training a machine learning algorithm becomes challenging, as the alternative loss function might be non-convex and discontinuous. Gradient-based methods usually assume a smooth objective, as in an earlier work [14] that employs a specialized financial criterion as a loss function. Similarly, [15] describes an end-to-end learning approach for probabilistic forecasting, with applications for energy storage arbitrage and grid scheduling. Decision tree learning, on the other hand, can be used to directly minimize task-specific costs [16]. An alternative approach based on bilevel programming is presented in [17], where the lower problem computes the best decision given a forecast and the upper problem estimates the linear coefficients of a forecasting model that lead to minimum costs. Finally, the third approach suggests applying empirical risk minimization (ERM) to directly forecast the decisions, a form of policy selection. This is appropriate when the cost function can be employed as loss function in a learning algorithm. For example, [18] describes data-driven solutions to the newsvendor problem, which effectively result in a high-dimensional quantile regression. A significant drawback of this approach, however, is that it fails to ensure feasible solutions for out-of-sample observations.

The integrated method we propose follows the prescriptive analytics framework established in [8]. To that end, we develop a decision tree algorithm that learns by directly minimizing decision costs. Decision trees are also considered in a similar context in [11], although in this paper we propose an alternative approach to enhance tractability via a random split criterion. The final conditional prescriptions are derived by means of a weighted SAA of the original problem. Furthermore, we explicitly address the issue of individual feature importance based on optimization performance, which is not considered in the above-mentioned works, and propose two methods, based on the node-splitting criterion and on the permutation importance method.

From an application standpoint, the case of trading renewable energy in a day-ahead (DA) market is considered, which constitutes a well-researched subject. As forecasting is not our primary focus, we refer the reader to [19] and [20] regarding recent advances in renewable and price forecasting, respectively. Several works examine the problem of deriving

optimal energy offers for wind power plants (WPPs) based on probabilistic wind forecasts and regulation cost forecasts [21]–[23]. Extensions of this problem include considering adjustment markets and risk-aversion [24], jointly participating in energy and reserve capacity markets [25], and trading under a single price balancing mechanism [26]. In terms of applications of data-driven optimization, [27] extends the framework put forward in [18] and proposes linear decision rules to improve both the forecasting and trading performance of a WPP participating in a DA market. For a similar case study with photovoltaic (PV) plants, [28] proposes an ERM formulation, approximated with evolutionary algorithms. Finally, [29] suggests employing the trading strategy as an alternative loss function to forecast market quantities, prior to deriving trading decisions.

In summary, the main contributions of this work are:

- We develop a novel decision tree algorithm for policy learning, trained under an alternative loss function to minimize task-specific costs. The proposed methodology is generic and can be readily applied to other conditional stochastic optimization problems.
- We describe a generic framework for assessing feature importance in terms of prescriptiveness, i.e., impact on decision costs, and propose two approaches, the first based on the internal node-splitting criterion and the second on the permutation importance method.
- From an application standpoint, we design hybrid policies for trading renewable energy that trade off predictive accuracy against optimal trading decisions. By jointly considering uncertainty from various sources (i.e. energy and the market), we examine how feature importance evolves under different tasks and market designs and identify the key drivers of optimization performance.

The rest of the paper is organized as follows. Section II presents the mathematical background and the proposed methodology. Section III formulates the renewable trading problem. Results are presented in Section IV. Finally, we draw conclusions and provide directions for future research in Section V.

II. MATHEMATICAL BACKGROUND AND PROPOSED METHODOLOGY

A. Conditional Stochastic Optimization

We consider a single-period optimization problem with $y \in \mathcal{Y} \subseteq \mathbb{R}^{d_y}$ being a vector of uncertain parameters and $x \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ being vector of relevant features, following a joint probability distribution $(X, Y) \sim \mathbb{Q}$. We are interested in the conditional stochastic optimization problem

$$v = \min_{z \in \mathcal{Z}} \mathbb{E}_{\mathbb{Q}}[c(z; Y)|X = x] = \min_{z \in \mathcal{Z}} \mathbb{E}_{y \sim \mathbb{Q}_{\bar{x}}}[c(z; Y)], \quad (1)$$

where v is the objective value, $z \in \mathbb{R}^{d_z}$ is the decision vector, \mathcal{Z} is a convex set of feasible solutions, $c(\cdot)$ is a cost function and x is an observation of X . Instead of the true conditional distribution $\mathbb{Q}_{\bar{x}}$, we have access to a training data set $\{(y_1, x_1), \dots, (y_N, x_N)\}$ of N observations and aim at learning a *policy* \hat{z} that varies as a function of x .

The fundamental method of approximating (1) in the presence of observations of Y is via SAA [9]:

$$\hat{z}^{SAA} = \arg \min_{z \in \mathcal{Z}} \sum_{i=1}^N \frac{1}{N} c(z; y_i). \quad (2)$$

While the SAA enjoys several nice theoretical properties, such as consistency and asymptotic optimality, it fails to exploit the available information encoded in x . The standard approach for doing so is to employ a forecasting model $g: \mathcal{X} \times \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ that maps observations of x to y , prior to solving an optimization problem, which we denote Forecast-then-Optimize (FO). We differentiate between the case of deterministic forecasts

$$\hat{z}^{EV} = \arg \min_{z \in \mathcal{Z}} c(z; \mathbb{E}_{y \sim \mathbb{Q}_x} [Y | X = x]), \quad (3)$$

where EV stands for Expected Value, and the generalized case of probabilistic forecasts

$$\hat{z}^{FO} = \arg \min_{z \in \mathcal{Z}} \mathbb{E}_{y \sim \hat{\mathbb{Q}}_x} [c(z; Y)], \quad (4)$$

where $\hat{\mathbb{Q}}_x$ is the approximated conditional distribution of Y . In turn, (4) is solved with standard stochastic optimization techniques. While in most applications, this approach constitutes the current state of the art, note that an optimal solution to (4) is not an optimal policy, as various implicit assumptions apply (e.g. calibration of predictive density). Alternatively, one can directly forecast the decision via ERM [18]. Depending on (1), however, various solutions might overlap. As such a case is relevant to the applications presented in Section III, we omit a separate presentation for the ERM approach.

The framework proposed in [8] integrates predictive and prescriptive analytics by forming a weighted SAA of (1) to derive conditional decisions, termed *predictive prescriptions*. These prescriptions, which retain consistency and asymptotic optimality, are defined as:

$$\hat{z}(x) = \arg \min_{z \in \mathcal{Z}} \sum_{i=1}^N \omega_{N,i}(x) c(z; y_i), \quad (5)$$

where $\omega_{N,i}(x)$ denotes weights obtained from local learning algorithms, e.g. decision trees. In the original work [8] $\omega_{N,i}(x)$ are derived by training the various algorithms for prediction. This work proposes deriving $\omega_{N,i}(x)$ by directly minimizing decision costs, leading to *prescriptive trees*. Thus, we leverage the structure of (1) during learning, providing a more informed approximation of the decisions. This is further motivated by the fact that in Section III we consider uncertainty stemming from different sources (renewable production and market quantities), which in turn depend on a different set of features. Training a local learning algorithm to predict both of these, would inevitably lead to suboptimal performance. Instead, our approach allows the model to assess the relative impact of each uncertain parameter on the final decision costs and weight them accordingly during learning, while also exploiting possible cross-dependencies. Fig. 1 presents respective flow diagrams for the above-mentioned approaches.

B. Randomized Prescriptive Trees

Decision tree leaning is a widely popular machine learning algorithm, employed both for classification and regression tasks. The proposed method follows the classification and regression trees (CART) [30] approach, that recursively applies locally optimal binary splits, resulting in a set of L leaves, $R_{1:L}$. A node split separates the feature space at feature $j \in d_x$ and point s into two disjoint partitions: $R_1(j, s) = \{i \in [N] | x_{i,j} < s\}$ and $R_2(j, s) = \{i \in [N] | x_{i,j} \geq s\}$, where $[N]$ is shorthand notation for $\{1, 2, \dots, N\}$. Thus, observations that satisfy $x_{i,j} < s$ fall to the left of the node, while the rest fall to the right. For brevity of exposition we focus exclusively on quantitative features, although it is straightforward to also include categorical features. The locally optimal split at a root node is found by solving

$$\min_{j,s} \left[\min_{c_1} \sum_{x_i \in R_1(j,s)} \ell(y_i, c_1) + \min_{c_2} \sum_{x_i \in R_2(j,s)} \ell(y_i, c_2) \right], \quad (6)$$

where $\ell(\cdot)$ is the selected training loss function, e.g. the sum of squared errors for regression. In this case, node predictions are equal to the average observation in the new partition, as each split aims at minimizing the within-leaf variance. Decision trees are prone to overfitting, i.e., they suffer from high variance, which significantly hinders their predictive capacity. Randomization-based ensemble methods address this issue and lead to impressive predictive performance. Popular methods include bootstrap aggregation (bagging), Random Forests [31] and Extremely Randomized Trees (ExtraTrees) [32].

In this work, we employ the cost function of (1) as the split criterion, with the locally optimal split obtained from:

$$\min_{j,s} \left[\min_{z_1 \in \mathcal{Z}} \sum_{x_i \in R_1(j,s)} c(z_1; y_i) + \min_{z_2 \in \mathcal{Z}} \sum_{x_i \in R_2(j,s)} c(z_2; y_i) \right]. \quad (7)$$

The inner minimization problems correspond to the SAA solution of each partition, with \hat{z}_1, \hat{z}_2 being the estimated locally constant decisions. One approach to solve (7) is by ordering the observations in feature j , evaluating the solution of the inner problem at each candidate point, and selecting the split with the lowest cost. However, the discrete nature of (7) makes it intractable for all but the simplest problems. To that end, we employ a randomized split selection, following the ExtraTrees algorithm [32]. We refer to an ensemble of trees grown with random splits as *Randomized Prescriptive Trees (RPT)*.

For a single prescriptive tree, we start from the top with a full data set and recursively partition the feature space until no further improvements are possible or a stopping criterion is met. Typical stopping criteria include the maximum tree depth Δ^{max} , the minimum number of observations n_{min} that fall at each leaf, and a predefined threshold for cost reduction. At each node of each tree, we randomly select a subset of K features from x and for each feature randomly select a candidate split point within its range. Next, we estimate the aggregated cost of (7) for each candidate split and compare it with the cost at its root node, updating the tree structure

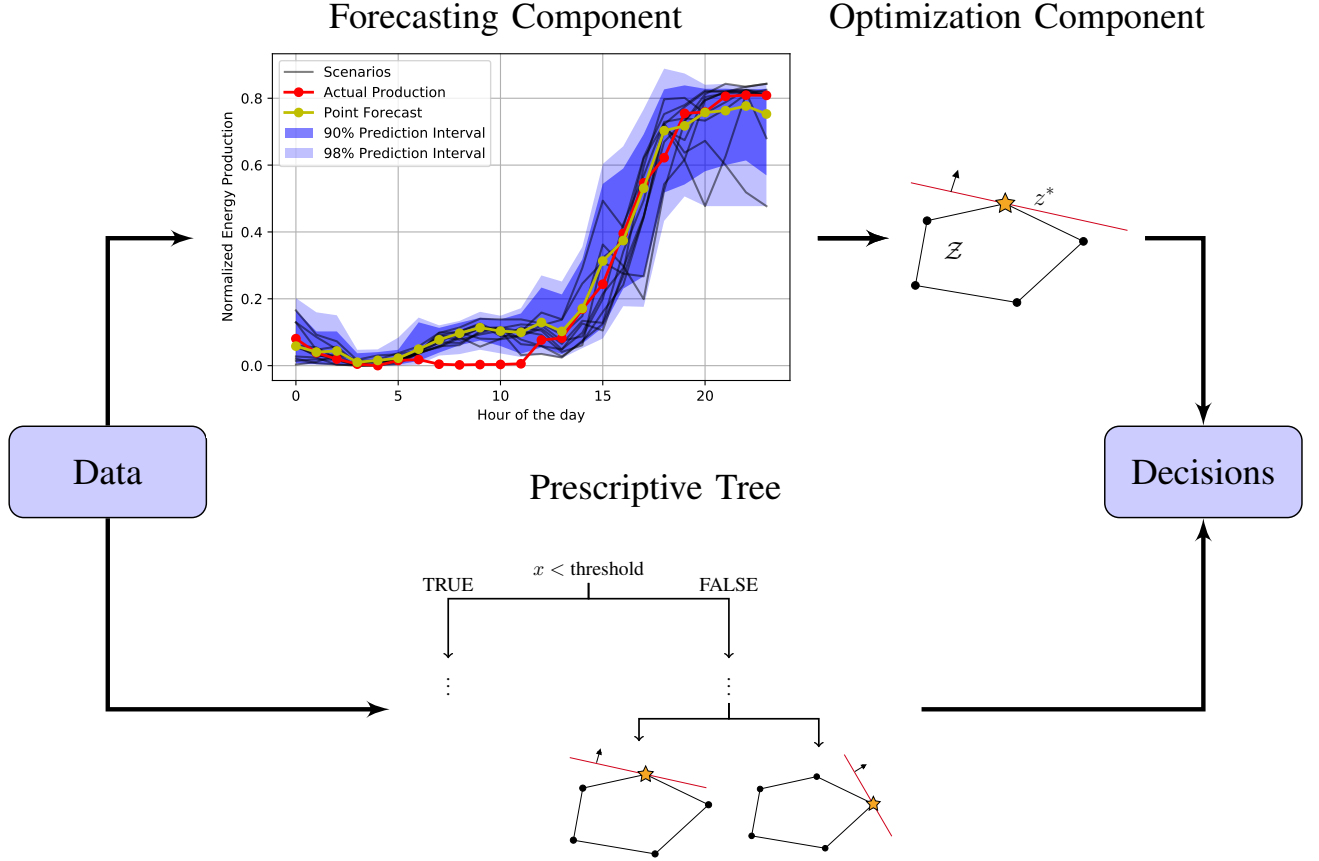


Fig. 1: Moving from data to decisions. On top, a standard model chain comprising of forecasting and optimization components is depicted, with different type of forecasts typically employed in practise being showcased. On bottom, a prescriptive decision tree that maps input data to decisions is shown. In both cases, the objective is uncertain and must be inferred from data.

accordingly. The subroutine of splitting a single tree node is detailed in Algorithm 1.

Algorithm 1 Node Split

Input: Training data (x_i, y_i) , current partition R_{init} , current depth Δ , hyperparameters $\{n_{min}, K, \Delta^{max}\}$

Output: Local split s

- 1: Find initial cost $v_{init} = \min_{z \in \mathcal{Z}} \sum_{x_i \in R_{init}} c(z; y_i)$ and initial prescription $\hat{z} = \arg \min_{z \in \mathcal{Z}} \sum_{x_i \in R_{init}} c(z; y_i)$
 - 2: **if** $\Delta < \Delta^{max}$ **then**
 - 3: Select K random features $\{x_1, x_2, \dots, x_K\}$
 - 4: Select random split point $s_\kappa \forall \kappa \in K$ from $U[x_\kappa^{min}, x_\kappa^{max}]$ such that $|R_1|, |R_2| \geq n_{min}$
 - 5: Solve (7) $\forall (\kappa, s_\kappa)$
 - 6: Find tuple (j^*, s^*) such that $v(s^*) = \min_{\kappa=1, \dots, K} v(s_\kappa)$
 - 7: **if** $v(s^*) < v_{init}$ **then**
 - 8: $\Delta \leftarrow \Delta + 1$
 - 9: **return** tuple (j^*, s^*) {Create split}
 - 10: **else**
 - 11: **return** nothing {Node becomes leaf with output \hat{z} }
 - 12: **end if**
 - 13: **end if**
-

For a single prescriptive tree, the corresponding weights $\omega_{N,i}(x)$ for a new query x are obtained as:

$$\omega_{N,i}(x) = \frac{\mathbb{I}[x_i \in R(x)]}{|R(x)|}, \quad (8)$$

with $R(x)$ is the leaf that x falls into, $|R(x)|$ the leaf cardinality and $\mathbb{I}[x_i \in R(x)]$ an indicator function that checks whether x_i falls into $R(x)$. For an ensemble of B trees the weights are obtained from:

$$\omega_{N,i}(x) = \frac{1}{B} \sum_{b=1}^B \frac{\mathbb{I}[x_i \in R^b(x)]}{|R^b(x)|}. \quad (9)$$

From Algorithm 1 we observe that a single tree is fully compiled, i.e., each leaf outputs a prescription, thus providing a direct mapping of input data to decisions, while also ensuring feasibility. For an ensemble of B trees, an additional step is required to approximate the solution from (5), which is detailed in Algorithm 2.

The main computational cost of Algorithm 1 occurs during the evaluation of candidate splits. The motivating factor behind the random split criterion lies in the expected reduction in computation time, as only a small number of splits are evaluated at each node. Computational experiments between the ExtraTrees and the Random Forest algorithm [32] suggest

Algorithm 2 Generation of Predictive Prescription

Input: New query x
Output: Prescription \hat{z}^{RPT}

- 1: Initialize $\omega_{N,i}(x) = 0$
 - 2: **for** tree $b = 1, \dots, B$ **do**
 - 3: $\omega_{N,i}(x) = \omega_{N,i}(x) + \frac{1}{B} \frac{\mathbb{I}[x_i \in R^b(x)]}{|R^b(x)|}$
 - 4: **end for**
 - 5: **return** $\hat{z}^{RPT}(x) = \arg \min_{z \in \mathcal{Z}} \sum_{i=1}^N \omega_{N,i}(x) c(z; y_i)$
-

an average reduction in training time by a factor of 3 for $K = \sqrt{d_x}$, which can rise up to a factor of 10 for wider data sets (larger d_x). Regarding the ensemble size B , the generalization error is expected to monotonically decrease as B increases, thus the computation time is the main consideration for its selection. Note that the task of training an ensemble is trivially parallelizable. Similarly, the rest of the hyperparameters K, n_{min} represent an inherent trade-off between model capacity and computation costs (single trees are maximally grown, thus Δ^{max} is set at infinity). The number of candidate splits K controls how strong individual splits are, with the number of split evaluations growing linearly with K . Regarding n_{min} , larger values result in shallower trees (and reduced computations), with higher bias and lower variance. Section IV presents detailed results on how hyperparameters B, K and n_{min} affect model performance under different loss functions.

C. An Illustrative Example

To illustrate the proposed method, we study a toy newsvendor problem. Consider an uncertain demand generated from $Y = 10 + 10\mathbb{I}[X > 0.5] + 10\mathbb{I}[X > 0.8] + \epsilon$, with X being a single feature following a uniform distribution $U[0, 1]$ and ϵ random noise following a normal distribution $N(0, 2)$. The cost function is set as $c(z; Y) = 2(Y - z)^- + 10(z - Y)^+$. We sample 1000 observations and train a single prescriptive tree with $\Delta^{max} = 2$. For the purpose of illustration, instead of random splits, we considered evaluating splits at 100 equally spaced quantiles of x . Fig. 2 presents a data scatterplot alongside the derived decisions and split thresholds, while Fig. 3 presents the tree structure, both of which highlight the fact that the tree does a good job of approximating the optimal decision.

D. Measuring the Prescriptiveness

Explainability is pivotal to disseminating the results to industry stakeholders and enabling large scale adoption of analytics tools in real-life applications. Here, we are interested in a quantitative assessment of the impact of various features on the efficacy of decisions (*prescriptiveness*). This is especially important in cases where obtaining explanatory data incurs in and of itself additional costs, e.g. acquiring weather forecasts for multiple locations. The proposed approach adapts the Mean Decrease Impurity (MDI) in a prescriptive analytics context. Provided a scoring rule, the MDI measures the total decrease in node impurity (dissimilarity) weighted by the probability




Fig. 2: Data scatterplot and derived decisions for the toy problem.




Fig. 3: Tree structure for the toy problem.

of reaching a specific node averaged over the ensemble [33]. Considering a prescriptive tree node R_0 partitioned at (j, s) into R_1, R_2 , the decrease in aggregated cost is measured as:

$$\Delta v(j, s) = v(R_0) - v(R_1) - v(R_2). \quad (10)$$

For an ensemble of B trees, the importance of a feature x_j in terms of prescriptiveness, $Imp(x_j)$, is measured as the weighted cost decrease over all trees in the ensemble:

$$Imp(x_j) = \frac{1}{B} \sum_{b=1}^B \sum_{x_j \in R_{1:L}} p(b) \Delta v(j, s), \quad (11)$$

with $p(b) = \frac{|R_l^b|}{N}$ being the proportion of observations reaching node R_l in tree b . The MDI is estimated internally during training, and is thus obtained without additional computational costs.

We also consider measuring prescriptiveness by adapting the permutation importance technique proposed in [31]. First, we estimate aggregated costs with respect to the selected objective function over a hold-out set, which defines the base score. Next, we iterate over all the features, permute (re-shuffle) each one, and derive new prescriptions, repeating the process a number of times. The permutation importance is then defined as the *expected increase* in decision costs compared to the base score. Preliminary analysis indicated that this approach leads to a significant increase in computational costs, as prescriptions need to be re-optimized at each query. Therefore, we omit it from the experimental results, but note

that it presents an attractive alternative for the case of a single prescriptive tree.

III. RENEWABLE TRADING

A. Problem Formulation

This section applies the proposed methodology to the case of trading renewable energy. Throughout, we assume a renewable producer participating as a price-taker in a DA market subject to imbalance costs. Prior to market closing, the producer submits an energy offer E_t^c for each clearing period t of the DA market. As temporal constraints do not apply, subscript t is dropped from the formulation. During real-time operation, the system operator activates balancing reserves to maintain the demand-supply equilibrium and hence stabilize the system frequency. The system assumes two states, namely *short*, i.e., demand exceeds supply and upward regulation is required, and *long*, i.e., supply exceeds demand and downward regulation is required. Based on real-time production, the producer buys back (sells) the amount of energy shortage (surplus) in order to balance its individual position. In the following, we present problem formulations that apply to different balancing market designs.

Let E^W denote the stochastic renewable production, π^{DA} the clearing price of the DA market and $\pi^{(\uparrow/\downarrow)}$ the marginal cost of activating upward/downward regulation services. Under the assumption of individual rationality, shortage of supply leads to increased real-time marginal costs. Thus, we assume that if the system is short, it holds that $\pi^\uparrow \geq \pi^{DA}$ and $\pi^\downarrow = \pi^{DA}$; while if the system is long, then $\pi^\downarrow \leq \pi^{DA}$ and $\pi^\uparrow = \pi^{DA}$. Let us further define $\lambda^\uparrow = \max(0, \pi^\uparrow - \pi^{DA})$ and $\lambda^\downarrow = \max(0, \pi^{DA} - \pi^\downarrow)$ as the respective upward and downward unit regulation costs. Evidently, it holds that $\lambda^\uparrow \cdot \lambda^\downarrow = 0$, i.e., only one of them (at most) assumes a value greater than zero for a given settlement period.

Under a *single-price* balancing mechanism, the revenue for each settlement period is defined as:

$$\begin{aligned} \rho^{\text{single}} &= \pi^{DA} E^c + \pi^\uparrow (E^W - E^c) + \pi^\downarrow (E^W - E^c) \\ &= \pi^{DA} E^W - \underbrace{[-\lambda^\uparrow (E^W - E^c) + \lambda^\downarrow (E^W - E^c)]}_{\text{imbalance cost}}. \end{aligned} \quad (12)$$

Here, $(E^W, \lambda^\uparrow, \lambda^\downarrow)$ define the uncertain problem parameters. Since revenue is affine with respect to the contracted energy, it is trivial to derive the optimal energy offer analytically as:

$$E^{c*} = \begin{cases} E^{\min}, & \text{if } -\hat{\lambda}^\uparrow + \hat{\lambda}^\downarrow \leq 0 \\ E^{\max}, & \text{if } -\hat{\lambda}^\uparrow + \hat{\lambda}^\downarrow > 0, \end{cases} \quad (13)$$

where $\hat{\lambda}^{\uparrow/\downarrow}$ denotes expected (forecast) values, see [22, Section II] for proof. We can interpret (13) as follows: the optimal offer equals zero if the system is expected to be short and the nominal capacity if the system is expected to be long. The case of zero costs is merged with the former without loss of generality. In practice, however, following this policy incurs great risks and could constitute market abuse, which motivates designing a policy that does not lead to excessive imbalances.

Alternatively, if the balancing market operates under a *dual-price* balancing mechanism, (12) is modified to impose non-arbitrage between the DA and the balancing market. The single period revenue is now defined as:

$$\rho^{\text{dual}} = \pi^{DA} E^W - \underbrace{[-\lambda^\uparrow (E^W - E^c)^- + \lambda^\downarrow (E^W - E^c)^+]}_{\text{imbalance cost}}, \quad (14)$$

where $(u)^- = \min\{u, 0\}$ and $(u)^+ = \max\{u, 0\}$. Hence, the term defining the imbalance cost is non-negative, which in turn means that no additional revenue can be attained in the balancing market. This contrasts the single-price market design, where deviations that help restore the system frequency result in negative imbalance costs, i.e., additional revenue. Provided probabilistic forecasts for all the uncertain parameters, the optimal offer is derived analytically as:

$$E^{c*} = \hat{F}^{-1}\left(\frac{\hat{\lambda}^\downarrow}{\hat{\lambda}^\downarrow + \hat{\lambda}^\uparrow}\right), \quad (15)$$

where \hat{F}^{-1} is the predicted inverse cumulative distribution function (c.d.f.) of E^W . The above solution holds without assuming independence between energy production and unit regulation costs, see [22, Section III] for details. Thus the optimal offering strategy requires probabilistic forecasts of energy production and deterministic forecasts of unit regulation costs.

For both cases, we propose a hybrid policy that balances profit maximization and energy forecasting accuracy. Specifically, the problem is formulated as:

$$\begin{aligned} \min_{E^c} \quad & \mathbb{E} \left[(1-k)(-\rho^{\text{single/dual}}) + k \|E^W - E^c\|_2^2 \right] \\ \text{s.t.} \quad & E^{\min} \leq E^c \leq E^{\max}. \end{aligned} \quad (16)$$

The objective function of problem (16) comprises two terms. The first term minimizes the trading costs (negative revenue), while the second term minimizes the energy forecast error. We can interpret this objective as adding a regularization term that penalizes excessive deviations from the expected energy production. Design parameter k controls the trade-off between the two objectives and defines different tasks. Specifically, for $k = 0$ we retrieve a purely prescriptive task, while for $k = 1$ we obtain a purely predictive task.

B. Energy and Price Forecasting

As discussed, the standard FO approach requires forecasting the uncertain quantities prior to solving a stochastic optimization problem. The particular instance described above requires estimating the conditional distribution of E^W and the conditional expectations of $\lambda^{\uparrow/\downarrow}$. To that end, the producer employs feature vectors x^E and x^{market} that include information relevant to the uncertain parameters, e.g. weather predictions, historical energy production, and historical market prices, among others. Subsequently, offering decisions are derived either numerically (sampling scenarios) or analytically. Note that since we consider single-period problems, scenario generation boils down to simply sampling from the conditional distribution of E^W .

As our purpose is not to provide a comprehensive analysis of forecasting models, we employ established benchmarks. For probabilistic energy forecasting, we select the Quantile Regression Forests model proposed in [25]. Contrary to energy forecasting, the literature on forecasting unit regulation costs is relatively scarce. A standard practice is to partition the problem into three prediction tasks, namely upward and downward regulation cost and direction prediction. Finally, individual components are combined according to the requirements of the specific policy by the law of total expectation. Specifically, the individual components are:

$$\hat{\phi} = \mathbb{P}(\lambda^\uparrow > 0), \quad (17)$$

$$\hat{\lambda}^\uparrow = \hat{\phi} \mathbb{E}[\lambda^\uparrow | \lambda^\uparrow > 0], \quad (18)$$

$$\hat{\lambda}^\downarrow = (1 - \hat{\phi}) \mathbb{E}[\lambda^\downarrow | \lambda^\downarrow > 0], \quad (19)$$

where $\hat{\phi}$ is the estimated probability of the system being short. Therefore, the prediction for the upward unit regulation cost λ^\uparrow equals the expectation of a regression model trained conditionally on the system being short, weighted by probability $\hat{\phi}$. Following [34], we apply exponential smoothing to model the individual components.

IV. EVALUATION AND RESULTS

A. Experimental Setting

For the computational experiments, we consider an aggregation of 3 WPPs and 1 PV plant, with a total capacity of 49MW (16% PV share), respectively located in northern and southern France. The results presented do not explicitly consider the effect of DA prices or possible control actions, e.g. curtailing production when DA prices are negative. Energy-related features x^E include weather forecasts for each plant location, namely wind speed, wind direction, temperature, cloud coverage and solar radiation, issued at 12am on day D-1 spanning a horizon of 24-48 hours ahead. Due to the close proximity of the WPPs, respective weather forecasts are aggregated via averaging. We employ market data for the French electricity market for the same period, downloaded from [35]. Market data comprises historical lags for DA prices (one day and one week prior), historical lags for system imbalance volumes (two days prior), and DA forecasts for available thermal generation, electricity demand, and renewable generation at transmission level. The DA forecasts from the system operator are processed to generate a net load series, defined as the expected electricity demand minus the expected renewable production, and a system margin series, defined as the ratio of net load to available thermal generation. In addition, we include categorical variables to model the calendar effect, namely day of the week and hour of the day. A total of $d_x = 20$ features are employed. Both DA and balancing markets are cleared at a half-hour settlement period, thus we derive 48 offers per day. Models are trained on one year of historical data spanning 2019 and evaluated on the first 4 months of 2020.

Our goal is to showcase the ability of the proposed approach to provide informed decisions under different objectives using

a single data-driven model, without the need to deploy multiple forecasting models. The following approaches are compared:

- FO^k : The standard FO approach, in which uncertain parameters are estimated via probabilistic forecasting prior to solving the respective stochastic optimization problem for trading on a single and dual-price market, respectively.
- RPT^k : Predictive prescriptions with weights derived from the proposed RPT algorithm.

Optimization problems are solved either analytically, when applicable, or numerically. As mentioned, for $k = 0$ both approaches result in the theoretical optimal offer for a risk-neutral participant, while for $k = 1$ the task coincides with forecasting energy production.

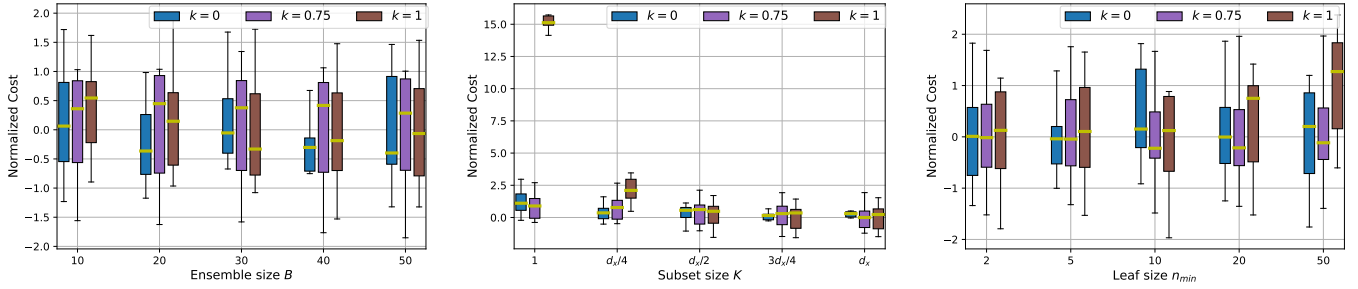
B. Hyperparameter Tuning

This subsection examines the performance of the RPT method with respect to hyperparameters $\{B, K, n_{min}\}$ in a controlled setting. We consider the problem of trading in a single-price market for values of $k = \{0, 0.75, 1\}$, randomly select a subset of 2000 training and 200 validation samples and estimate the (normalized) expected cost with respect to the different objective functions. The process is repeated 10 times and the expected costs for the different objectives are standardized as to have zero mean and unit variance.

Fig. 4a plots the standardized cost as a function of the ensemble size B for the different values of k . The performance appears to be insensitive to the size of the ensemble, with similar results across the different tasks. Next, we examine the effect of subset size K , which controls the model capacity. For $K = 1$ node splits are completely random (requiring minimum computations), while for $K = d_x$ all features are evaluated at each split. From Fig. 4b a significant discrepancy across tasks is evident. Specifically, the selection of K has a notable effect on the performance of the predictive task, with a significant decrease for lower values of K , while this effect is less pronounced for the other tasks. This discrepancy across tasks could be attributed to the fact that the prescriptive tasks are relatively more demanding than simple energy production forecasting. This fact is further analyzed in the following subsection where we estimate relative performance indices. Further, a non-negligible decrease in variance as K increases is also evident for the prescriptive task. Finally, we examine the impact of the minimum leaf size n_{min} . Generally, smaller values of n_{min} result in lower bias, while larger values provide a smoothing effect. Fig. 4c indicates a decrease in performance for values of leaf size greater than 10, with the effect being more pronounced for the predictive task. For the rest of this section, the results presented are estimated with hyperparameters $K = d_x/2$, $B = 50$, and $n_{min} = 5$.

C. Trading Results

We examine the performance for the different market designs in terms of revenue attained, trading risk and relative optimization performance. The conditional value at risk at 5% ($CVaR_{5\%}$) is used as a proxy for trading risk, defined



(a) Ensemble size B ($K = d_x/2, n_{min} = 5$). (b) Number of splits K ($B = 25, n_{min} = 5$). (c) Leaf size n_{min} ($B = 25, K = d_x/2$).

Fig. 4: Effect of hyperparameters B, K , and n_{min} in performance.

TABLE I: Results For Single-Price Market

Trading Results	$k = 0$		$k = 0.25$		$k = 0.5$		$k = 0.75$		$k = 1$	
	FO	RPT	FO	RPT	FO	RPT	FO	RPT	FO	RPT
Gross Revenue (10^3 EUR)	1172	1266	1165	1248	1156	1235	1155	1231	1176	1178
Expected Revenue (EUR/hour)	201.77	217.91	200.64	214.84	199.09	212.69	198.90	211.88	202.51	202.79
$CVaR_{5\%}$ (EUR)	-467.61	-363.08	-465.22	-353.42	-449.37	-315.95	-331.90	-174.87	-97.13	-99.48
Coefficient of prescriptiveness P	0.03	0.17	-0.01	0.12	-0.05	0.09	-0.08	0.11	0.86	0.86

TABLE II: Results For Dual-Price Market

Trading Results	$k = 0$		$k = 0.25$		$k = 0.5$		$k = 0.75$		$k = 1$	
	FO	RPT	FO	RPT	FO	RPT	FO	RPT	FO	RPT
Gross Revenue (10^3 EUR)	1136	1137	1136	1141	1136	1137	1137	1141	1138	1139
Expected Revenue (EUR/hour)	195.51	195.84	195.56	196.42	195.65	195.85	195.82	196.49	196.01	196.11
$CVaR_{5\%}$ (EUR)	-91.79	-102.96	-91.80	-91.94	-91.44	-104.19	-91.22	-94.53	-100.05	-97.67
Coefficient of prescriptiveness P	0.66	0.70	0.67	0.70	0.68	0.71	0.71	0.74	0.86	0.87

as the expected return over the worst 5% of cases. Revenue and risk are absolute metrics of performance and thus cannot be employed to compare relative performance across different use cases. The coefficient of prescriptiveness P [8] is a unitless metric, analogous to the R^2 for regression, that measures relative optimization performance against the naive (SAA) and perfect-foresight solution. Specifically, for each i in $\{FO, RPT\}$ and each value of k , the coefficient P is estimated as:

$$P_{i,k} = 1 - \frac{\widehat{v}_k^i - \widehat{v}_k^*}{\widehat{v}_k^{SAA} - \widehat{v}_k^*}, \quad (20)$$

where $\widehat{v}_k^i, \widehat{v}_k^{SAA}, \widehat{v}_k^*$ are the expected costs over the test set under the i , SAA and perfect-foresight method. Note that the SAA approach results in a constant energy offer derived from empirical distributions. The coefficient P is bounded above by one, while negative values indicate a failure to outperform the SAA.

First, we examine results for a single-price balancing market. Table I presents aggregated results for $k = \{0, 0.25, 0.5, 0.75, 1\}$. The RPT approach leads to a total revenue increase of 8.02%, 7.12%, 6.83%, 6.58% and 0.17%, for each respective value of k . As expected, a larger k leads to more conservative offers, as the minimization of the imbalance volume is weighted more heavily in the objective. Fig. 5 further showcases the improved risk-reward trade-off of the RPT, as it sets the efficient frontier, attaining higher revenue

for a given level of risk and vice versa. The 'elbow' of the graph present at $k = 0.75$ indicates a good trade-off point for the selection of k . The above observations regarding trading performance are further validated by examining the coefficient of prescriptiveness P . Interestingly, the FO approach fails to outperform the naive SAA solution for lower values of k (P is negative). The fact that both approaches converge to similar performance for $k = 1$, effectively means that the proposed method achieves similar performance with the Random Forest algorithm for standard regression. In addition, we observe significantly lower values for the coefficient P for $k = \{0, 0.25, 0.5, 0.75\}$ compared to the regression task, highlighting the fact that prescriptive tasks are relatively more demanding than simple energy forecasting, as the relative distance from the perfect-foresight solution is greater. This could be attributed to the fact that prescriptive tasks require an estimation of unit regulation costs in the DA horizon, which in practice is known to be difficult. Nonetheless, our results provide a quantitative estimation of this empirical knowledge, which we believe to be of use to both researchers and other stakeholders.

Table II presents aggregated results for a dual-price balancing mechanism. Contrary to the single-price case, results are rather insensitive to the selection of parameter k . This could be attributed to the non-arbitrage condition imposed by the market design. This notion is further supported by examining




Fig. 5: Risk versus reward plot for trading in a single-price market.

the evolution of P , which for lower values of k is significantly greater than the single-price market. Regarding the examined approaches, a small performance increase is evident for the RPT method for a lower k . Specifically, the gross revenue increase for the RPT approach ranges between 0.09% and 0.44%, which, however, might be negligible in terms of monetary value. Nonetheless, the single data-driven model proposed achieves similar or better performance than the benchmark, without requiring to develop multiple forecasting models.

D. Assessing Feature Importance

In this subsection, we investigate how different features affect optimization performance, as measured by the MDI method. Feature importances are normalized to add up to one. A subset of the most important features is plotted in Fig. 6. Considering a single-price market, we observe that for lower values of k , variables pertaining to estimating unit regulation costs assume greater weight. This can be attributed to the RPT placing more weight on the prescriptive term in the objective function. Specifically, the aggregation of the expected system margin, expected temperature at the WPP site, and lag observations for system imbalance volume assume approximately 65% of the total feature importance for $k = \{0, 0.25, 0.5, 0.75\}$. Note that the WPPs are located in close proximity to large metropolitan areas and interconnections with neighbouring countries, thus the expected temperature effectively serves as a proxy for electricity demand. As k increases, the importance of features related to energy forecasting gradually increases, with the expected wind speed at the WPP site reaching approximately 80% of the total feature importance for $k = 1$.

Under a dual-price mechanism, we observe significantly fewer variations in feature importance across the different values of k , resembling the results presented in Table II. Specifically, the expected wind speed at the WPP location is consistently the most important variable throughout, with its importance that ranging from 65% to 78%, followed by the expected system margin. Previous works on similar use cases mention that energy forecasting is relatively more important than price forecasting [27]. The results presented in Table II

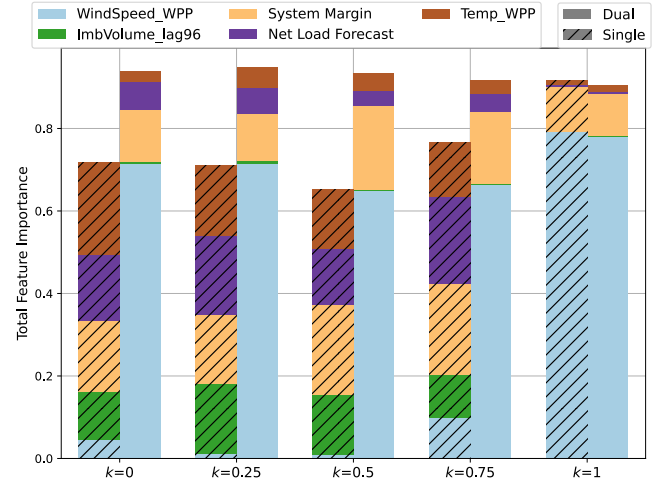


Fig. 6: Normalized feature importance based on prescriptiveness.

and Fig. 6 provide quantitative evidence for these assertions by jointly considering the two sources of uncertainty in the problem formulation and measuring the impact of different features.

Finally, comparing results across the different market designs, enables us to conclude that forecasting regulation costs is relatively more important if a single-price balancing mechanism is in place, while energy forecasting should be the primary focus for participants in dual-price markets.

V. CONCLUSIONS

This work presented a novel integrated forecasting and optimization approach to approximate conditional stochastic optimization problems employing weighted SAA. We proposed a decision tree learning algorithm that minimizes task-specific costs for generic convex problems during learning, employing a random split criterion to speed up computations. Further, we provided a framework to measure feature importance in terms of impact on optimization efficacy under different objective functions.

We validated the proposed approach in the case of a renewable producer participating in competitive electricity markets and proposed hybrid policies that balance optimal decisions and energy forecasting accuracy. In a detailed empirical analysis, we compared the proposed approach with the stochastic optimization benchmark in terms of monetary gains and relative optimization performance. Overall, our proposed solution achieved consistently better or similar performance against the current state of the art, with relative improvements being more pronounced for a single-price market. Additionally, we examined feature importance under different objectives and across different market designs, demonstrating the capability of the proposed solution to measure the impact of data on decision-making, and provided valuable insights on trading under different regulatory frameworks. Future work will examine multi-period problems, which are increasingly complex as they require accounting for the spatio-temporal correlations

of uncertain parameters, and enhancing the tractability of the proposed algorithm.

ACKNOWLEDGMENT

The authors wish to thank ENGIE GREEN for contributing historical production data.

REFERENCES

- [1] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, "Energy forecasting: A review and outlook," *IEEE Open Access Journal of Power and Energy*, 2020.
- [2] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown *et al.*, "Tackling climate change with machine learning," *arXiv preprint arXiv:1906.05433*, 2019.
- [3] H. Zareipour, C. A. Canizares, and K. Bhattacharya, "Economic impact of electricity market price forecasting errors: A demand-side analysis," *IEEE Transactions on Power Systems*, vol. 25, no. 1, pp. 254–262, 2009.
- [4] G. Kariniotakis, M. Matos, and V. Miranda, "Assessment of the benefits from advanced load & wind power forecasting in autonomous power systems," in *EWEC 1999*, 1999, pp. 391–394.
- [5] J. R. Birge and F. Louveaux, *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- [6] J.-F. Toubeau, J. Bottieau, F. Vallée, and Z. De Grève, "Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 1203–1215, 2018.
- [7] S. Camal, F. Teng, A. Michiorri, G. Kariniotakis, and L. Badesa, "Scenario generation of aggregated wind, photovoltaics and small hydro production for power systems applications," *Applied Energy*, vol. 242, pp. 1396–1406, 2019.
- [8] D. Bertsimas and N. Kallus, "From predictive to prescriptive analytics," *Management Science*, vol. 66, no. 3, pp. 1025–1044, 2020.
- [9] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- [10] R. Kannan, G. Bayraktar, and J. R. Luedtke, "Data-driven sample average approximation with covariate information," *Optimization Online*. URL: http://www.optimization-online.org/DB_HTML/2020/07/7932.html, 2020.
- [11] N. Kallus and X. Mao, "Stochastic optimization forests," *arXiv preprint arXiv:2008.07473*, 2020.
- [12] D. Bertsimas, J. Dunn, and N. Mundru, "Optimal prescriptive trees," *INFORMS Journal on Optimization*, vol. 1, no. 2, pp. 164–183, 2019.
- [13] A. N. Elmachtoub and P. Grigas, "Smart "predict, then optimize"," *Management Science*, 2021.
- [14] Y. Bengio, "Using a financial training criterion rather than a prediction criterion," *International journal of neural systems*, vol. 8, no. 04, pp. 433–443, 1997.
- [15] P. L. Donti, B. Amos, and J. Z. Kolter, "Task-based end-to-end model learning in stochastic optimization," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5490–5500.
- [16] A. N. Elmachtoub, J. C. N. Liang, and R. McNellis, "Decision trees for decision-making under the predict-then-optimize framework," *arXiv preprint arXiv:2003.00360*, 2020.
- [17] M. A. Muñoz, S. Pineda, and J. M. Morales, "A bilevel framework for decision-making under uncertainty with contextual information," *arXiv preprint arXiv:2008.01500*, 2020.
- [18] G.-Y. Ban and C. Rudin, "The big data newsvendor: Practical insights from machine learning," *Operations Research*, vol. 67, no. 1, pp. 90–108, 2019.
- [19] G. Kariniotakis, *Renewable Energy Forecasting: From Models to Applications*. Woodhead Publishing, 2017.
- [20] J. Nowotarski and R. Weron, "Recent advances in electricity price forecasting: A review of probabilistic forecasting," *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1548–1568, 2018.
- [21] P. Pinson, C. Chevallier, and G. N. Kariniotakis, "Trading wind generation from short-term probabilistic forecasts of wind power," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1148–1156, 2007.
- [22] C. J. Dent, J. W. Bialek, and B. F. Hobbs, "Opportunity cost bidding by wind generators in forward markets: Analytical results," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1600–1608, 2011.
- [23] E. Y. Bitar, R. Rajagopal, P. P. Khargonekar, K. Poolla, and P. Varaiya, "Bringing wind energy to market," *IEEE Transactions on Power Systems*, vol. 27, no. 3, pp. 1225–1235, 2012.
- [24] J. M. Morales, A. J. Conejo, and J. Pérez-Ruiz, "Short-term trading for a wind power producer," *IEEE Transactions on Power Systems*, vol. 25, no. 1, pp. 554–564, 2010.
- [25] S. Camal, A. Michiorri, and G. Kariniotakis, "Optimal offer of automatic frequency restoration reserve from a combined pv/wind virtual power plant," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 6155–6170, 2018.
- [26] J. Browell, "Risk constrained trading strategies for stochastic generation with a single-price balancing market," *Energies*, vol. 11, no. 6, p. 1345, 2018.
- [27] M. Munoz, J. M. Morales, and S. Pineda, "Feature-driven improvement of renewable energy forecasting and trading," *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3753–3763, 2020.
- [28] T. Carriere and G. Kariniotakis, "An integrated approach for value-oriented energy forecasting and data-driven decision-making application to renewable energy trading," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6933–6944, 2019.
- [29] A. Stratigakos, A. Michiorri, and G. Kariniotakis, "A Value-Oriented Price Forecasting Approach to Optimize Trading of Renewable Generation," in *2021 IEEE Madrid PowerTech*. Madrid, Spain: IEEE, Jun. 2021. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03208575>
- [30] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and regression trees*. CRC press, 1984.
- [31] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [32] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine learning*, vol. 63, no. 1, pp. 3–42, 2006.
- [33] G. Louppe, L. Wehenkel, A. Suter, and P. Geurts, "Understanding variable importances in forests of randomized trees," *Advances in neural information processing systems* 26, 2013.
- [34] T. Jónsson, P. Pinson, H. A. Nielsen, and H. Madsen, "Exponential smoothing approaches for prediction in real-time electricity markets," *Energies*, vol. 7, no. 6, pp. 3710–3732, 2014.
- [35] ENTSO-E. Transparency platform. [Online]. Available: <https://transparency.entsoe.eu/>