A bilevel framework for decision-making under uncertainty with contextual information☆

M.A. Muñoz, S. Pineda*, J.M. Morales

OASYS Group, University of Malaga, Malaga, Spain

ARTICLE INFO

ABSTRACT

In this paper, we propose a novel approach for data-driven decision-making under uncertainty in the presence of contextual information. Given a finite collection of observations of the uncertain parameters and potential explanatory variables (i.e., the contextual information), our approach fits a parametric model to those data that is specifically tailored to maximizing the decision value, while accounting for possible feasibility constraints. From a mathematical point of view, our framework translates into a bilevel program, for which we provide both a fast regularization procedure and a big-M-based reformulation that can be solved using off-the-shelf optimization solvers. We showcase the benefits of moving from the traditional scheme for model estimation (based on statistical quality metrics) to decision-guided prediction using three different practical problems. We also compare our approach with existing ones in a realistic case study that considers a strategic power producer that participates in the Iberian electricity market. Finally, we use these numerical simulations to analyze the conditions (in terms of the firm's cost structure and production capacity) under which our approach proves to be more advantageous to the producer.

1. Introduction

In the last couple of decades, the field of decision-making under uncertainty has regained momentum, spurred by the new opportunities that the Digital Age has brought to modern economies. As a result, this field has been prolific in the design and development of new tools capable of exploiting the vast amount of information that human societies currently generate, compile and record, mainly in the form of *data*.

From among all the exciting advances that have been achieved in the realm of decision making under uncertainty in recent years, we highlight the so-called *data-driven optimization under uncertainty*, which endows the decision maker with a powerful and versatile mathematical framework to hedge her decisions against both the intrinsic risk of an uncertain world and the limited and incomplete knowledge of the random phenomena that can be retrieved from a finite set of observations or data.

Data-driven optimization under uncertainty has been applied to a broad range of contexts and problems, for instance, inventory management [1,2], nurse staffing [3], portfolio optimization [4–6], shipment planning [4], network flow [5], power dispatch [7], and vehicle routing [8], just to name a few. For a recent survey on the topic and its applications, we refer the reader to [9,10].

In this paper, we first compare the proposed methodology with existing ones using two classical conditional stochastic optimization problems, namely, the newsvendor problem [3,11,12] and the product placement problem [1]. Additionally, we consider the problem of a strategic firm that has to decide the generation quantity that maximizes its expected profit while facing the uncertainty related to market conditions. This problem has a long tradition in the Economics and Management Science literature (see, for instance, [13–15]). In particular, we take *electricity* as the homogeneous good to be produced and thus, we place ourselves in the context of electricity markets, where this problem has received a great deal of attention since the deregularization of the power sector [16,17]. Most existing models address this problem by forecasting, as accurately as possible, the electricity market behavior. Then, such forecasts are used to compute the decision that maximizes the producer's profit. Here we present a novel and alternative data-driven procedure that considers the problem structure and leverage available auxiliary data to enhance market participation and increase profits. Our approach is formulated as a bilevel program that, under convexity assumptions, can be efficiently solved using commercially available optimization solvers. We demonstrate the superior performance of the proposed approach on a realistic case study that uses data from the Iberian electricity market.

In short, our contributions are threefold, namely:

- From a methodological point of view, we propose a novel data-driven framework for conditional stochastic optimization, whereby the parameters that are input to the decision-making problem are formulated as a function of some *covariates* or *features*. This function is, in turn, estimated factoring in its impact on the decision value. Finally, by way of this function, we construct a deterministic (single-scenario) surrogate optimization model that delivers decisions that are good in terms of the original conditional stochastic program. In Section 2, we introduce and mathematically formalize our proposal along with alternative state-of-the-art approaches available in the technical literature. Our approach is formulated as a bilevel optimization problem that can be reformulated as a single-level optimization problem and solved using off-the-shelf optimization solvers as discussed in Section 3.
- From a theoretical perspective, we compare our approach with existing ones in Section 4 for three different applications, namely, the newsvendor problem, the product placement problem, and the strategic producer problem.
- From a more practical point of view, Section 5 provides simulation results for the strategic producer problem using both an illustrative example and a realistic case study based on the Iberian electricity market. The numerical experiments show that our proposal can significantly increase the competitive edge of the strategic producer depending on her cost structure and the market demand elasticity.

We conclude the paper with a brief compilation of the most relevant observations in Section 6.

2. Mathematical framework and related work

In decision making we often model the uncertainty as a random vector of parameters ($y \in \mathcal{Y} \subseteq \mathbb{R}^m$) governed by a real unknown distribution Y and, typically, some relevant contextual information ($x \in \mathcal{X} \subseteq \mathbb{R}^p$) $\sim X$ is available before the decision is to be made. Following this scenario, the decision maker is interested in solving the *conditional* stochastic optimization problem

$$\min_{z \in Z} \quad \mathbb{E}[f(z; Y)|X = x] \tag{1}$$

where $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is a known function in the decision $z \in \mathbb{R}^n$, and $Z \subseteq \mathbb{R}^n$ is a nonempty, compact set known with certainty (i.e., independent of Y), to which the decision z must belong. In practice, neither the *joint* distribution of X and Y, nor the conditional distribution of Y given $X = x$ are known and therefore, problem (1) cannot be solved. On top of that, even if the true distribution were known and the decision z were fixed, problem (1) would typically require to compute the expectation of a function of a continuous random vector (i.e., a multivariate integral), which is, in itself, a hard task in general. Instead, the only information that the decision maker typically has is a sample $S = \{(y_i, x_i), \forall i \in \mathcal{N}\}$ where $y_i \in \mathbb{R}^m$ is a particular outcome of the uncertainty Y recorded under the context $x_i \in \mathcal{X}$, and $\mathcal{N}$ denotes the set of available samples.

Against this background, problem (1) is alternatively replaced with a *surrogate* optimization problem, in the hope that the solution to the latter is good enough for the former. In this line, different approaches have been proposed to construct such a surrogate optimization problem. For instance, the traditional *modus operandi* follows the rule "first predict, then optimize," which results in the following *surrogate problem* to approximate the solution to (1):

$$\min_{z \in Z} \quad f(z; \hat{y}) \tag{2}$$

where $\hat{y}$ denotes an estimate of the outcome of the uncertainty Y under the contextual information $x \in \mathcal{X} \subseteq \mathbb{R}^p$. The surrogate problem (2) is attractive for several reasons. First and foremost, it is

much simpler and faster to solve than (1). Actually, it is a *deterministic* optimization problem that, as opposed to (1), only requires evaluating the cost function $f(z; \cdot)$ at the single value or scenario $\hat{y}$. Furthermore, problem (2) seems intuitive and natural, especially when $\hat{y}$ represents "the most likely value" for Y given $X = x$. Indeed, the single scenario $\hat{y}$ is often chosen as an estimate of the expected value of the uncertainty Y conditional on $X = x$, that is, $\hat{y} \approx \mathbb{E}[Y|X = x]$, where, logically, the approximation is built from the available sample $S = \{(y_i, x_i), \forall i \in \mathcal{N}\}$. In the realm of forecasting, the estimate $\hat{y}$ is usually referred to as a *point prediction*.

In order to build the estimate $\hat{y} \approx \mathbb{E}[Y|X = x]$, a function $g^{FO} : \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}^m$ is normally chosen from a w-parameterized family G^{FO}, with $w \in \mathbb{R}^q$, to construct the forecasting model $\hat{y} = g^{FO}(x; w)$. The goodness of a certain parameter vector w is quantified in terms of a loss function $l^{FO}(y, \hat{y}) : \mathcal{Y} \times \mathbb{R}^m \to \mathbb{R}$ that measures the accuracy of the estimate. Then, given the sample $S = \{(y_i, x_i), \forall i \in \mathcal{N}\}$, the choice of w is driven by the minimization of the *in-sample* loss, as expressed below:

$$w^{FO} \in \arg\min_{w \in \mathbb{R}^q} \sum_{i \in \mathcal{N}} l^{FO}(y_i, g^{FO}(x_i; w)) \tag{3}$$

In this framework, the optimal decision z^{FO} under the context $X = x$ is thus obtained by solving the following deterministic problem:

$$z^{FO}(x) \in \arg\min_{z \in Z} \quad f(z; g^{FO}(x; w^{FO})) \tag{4}$$

We refer to this approach, which relies on a *good* forecast of the uncertainty Y (in particular, an estimate of $\mathbb{E}[Y|X = x]$), as FO (short for FOrecasting). Even though this approach is intuitive and may perform relatively well in many situations, it is fundamentally flawed for the following two basic reasons. First, since $\hat{y} \approx \mathbb{E}[Y|X = x]$ in FO, the surrogate problem (2) works as a proxy of the problem

$$\min_{z \in Z} \quad f(z; \mathbb{E}[Y|X = x]) \tag{5}$$

which, in general, is not equivalent to (1). Second, even in those cases where these two problems are indeed equivalent, the loss function l^{FO} that is typically used to compute w^{FO} (for example, the squared error) is solely intended to get a statistically good estimate of $\mathbb{E}[Y|X = x]$ and does not account for the nominal objective f or the constraints that the decision z must satisfy. For instance, approach (3) and (4) is unable to capture that overestimating $\mathbb{E}[Y|X = x]$ might worsen the objective function f much more than underestimating it.

In view of these design flaws, a number of works have proposed to replace the problem-agnostic l^{FO} that is generally used in (3) with a problem-aware loss function $l^{SP}(y, \hat{y}) = f(\dot{z}(\hat{y}); y)$ where $l^{SP} : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}$ and $\dot{z} : \mathcal{Y} \to Z$ defined as $\dot{z}(y) = \arg\min_{z \in Z} f(z; y)$. Therefore, function l^{SP} evaluates the loss of optimality associated with the decision $\dot{z}(\hat{y})$ that is prescribed by the surrogate decision-making problem (2) for the single value $\hat{y}$. Accordingly, the optimal parameter vector w^{SP} is obtained as the one that minimizes the *in-sample* optimality loss, that is:

$$w^{SP} \in \arg\min_{w \in \mathbb{R}^q} \sum_{i \in \mathcal{N}} f(\dot{z}(g^{SP}(x_i; w)); y_i) \tag{6}$$

where the function $g^{SP} : \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}^m$ is chosen from a family of functions G^{SP}. We use the acronym SP, which stands for "Smart Predict", to refer to this setup. Solving problem (6) using descent optimization methods requires to compute the gradient of the loss function $l^{SP}(y, \hat{y})$ with respect to w. This may not be feasible, since it involves the differentiation of the discontinuous function $\dot{z}(y)$ [18]. To overcome this difficulty, a great deal of research has been devoted to finding methods to approximate the gradient of (6) for

particular instances. The work developed in [19], for example, describes a procedure to solve (6) under the following three conditions: (i) f is quadratic, (ii) the uncertainty is only present in the coefficients of the linear terms of f, and (iii) no constraints are imposed on the decision z, which means $Z = \mathbb{R}^n$. Some years later, the authors [2] proposed a heuristic gradient-based procedure to solve (6) for strongly convex problems with deterministic equality constraints and inequality chance constraints. Almost concurrently, reference [5] discusses the difficulties of solving (6) in the case of linear problems, since such a formulation may lead to an uninformative loss function. To overcome this issue, they successfully develop a convex surrogate that allows to efficiently train $g^{SP}(x_i; w)$ in the linear case. More recently, the authors in [20] suggest a similar approach as in [2] to combinatorial problems with a regularized linear objective function.

In summary, the four references above propose ad-hoc gradient methods for specific instances of (6). However, the technical literature lacks, to the best of our knowledge, a general-purpose procedure to solve such a problem using available optimization solvers. To fill this gap, we propose the following bilevel program [21] as a generic mathematical formulation of (6):

$$w^{BL} \in \arg\min_{w \in \mathbb{R}^q; \hat{z}_i} \sum_{i \in \mathcal{N}} f(\hat{z}_i; y_i) \tag{7a}$$

$$\text{s.t.} \quad \hat{z}_i \in \underset{z \in Z}{\arg\min} \quad f\left(z; g^{BL}(x_i; w)\right), \quad \forall i \in \mathcal{N} \tag{7b}$$

where $g^{BL}: \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}^m$ is selected similarly to g^{FO} and g^{SP}. Problem (7) is formulated as a bilevel optimization model commonly used to mathematically characterize non-cooperative and sequential Stackelberg games in which the *leader* makes her decisions anticipating the reaction of the *follower* [22]. In this sense, the upper-level problem determines the optimal parameter vector w anticipating the decision provided by each lower-level problem (7b) if the value $\hat{y}_i$ is given by $g^{BL}(x_i; w)$. We denote this approach based on bilevel programming as BL (acronym of BiLevel). In Section 3, we discuss the assumptions that problem (1) must satisfy so that problem (7) can be reformulated as a single-level optimization problem to be solved using off-the-shelf optimization solvers. Although solving the bilevel problem (7) may be computationally expensive, this is a task that can be performed offline. Once w^{BL} is determined, the optimal decision z^{BL} under context $X = x$ is computed by solving the following problem:

$$z^{BL}(x) \in \arg\min_{z \in Z} \quad f(z; g^{BL}(x; w^{BL})) \tag{8}$$

The bilevel program (7a)–(7b) computes the value for the parameter vector w that maximizes the *in-sample* performance of the surrogate decision-making model (8). For this estimation to be of use, it must be guaranteed that under two contexts x_i, x_i', such that $x_i = x_i'$, it holds $\hat{z}_i = \hat{z}_i'$, i.e., under equal contexts, equal decisions. This is a condition that is reminiscent of the notion of *non-anticipativity* in Stochastic Programming. Importantly, this condition is automatically satisfied if the solution to the lower-level problem (7b) is unique for any value of w. Otherwise, the bilevel program (7a)–(7b) would choose the $\hat{z}_i$ from the optimal solution set of (7b) that minimizes the upper-level objective function (7a) *given* —i.e., by anticipating— the uncertainty outcome y_i. This is so because the bilevel program (7a)–(7b), as we have formulated it, delivers the optimistic Stackelberg solution. For instance, let us assume that there exists a value $\tilde{w}$ such that $f(z; g^{BL}(x_i; \tilde{w})) = \vartheta$ for all $i \in \mathcal{N}$, where ϑ is a constant. In this case, the lower-levels (7b) boil down to feasibility problems imposing that $z \in Z$ and therefore, $\hat{z}_i$ can violate non-anticipativity and adapt to realization y_i for all $i \in \mathcal{N}$. More importantly, using $\tilde{w}$ in (8) would lead to degenerate and highly suboptimal decisions under any context $X = x$. This issue is reported in [5] for linear objective functions,

where authors propose a convex surrogate function of l^{SP} to train meaningful instances of model $g^{SP}(\cdot; w^{SP})$. Similarly, we propose in Section 4.3 a modified lower-level surrogate model for the strategic producer problem in order to ensure non-anticipativity for any parameter vector w.

Next, we discuss other surrogate decision-making models different from (2), which have also been recently proposed to approximate the solution of (1). For this purpose, notice first that problem (1) can be equivalently recast as

$$\min_{z \in Z} \quad \mathbb{E}[f(z; Y)|X = x] = \min_{z \in Z} \quad \mathbb{E}_{\mathcal{Q}_{|x}}[f(z; Y)] \tag{9}$$

where $\mathcal{Q}_{|x}$ represents the conditional probability distribution of Y given $X = x$. Thus, a second family of surrogate decision-making models can be introduced with the following general form:

$$\min_{z \in Z} \quad \mathbb{E}_{\widehat{\mathcal{Q}}_{|x}}[f(z; Y)] \tag{10}$$

where $\widehat{\mathcal{Q}}_{|x}$ is an approximation of the unknown probability measure $\mathcal{Q}_{|x}$ that is constructed from the available sample $S = \{(y_i, x_i), \forall i \in \mathcal{N}\}$. For the surrogate problem (10) to be computationally tractable, the proxy $\widehat{\mathcal{Q}}_{|x}$ is often built as a discrete probability distribution supported on a finite number of points, more specifically, on the y-locations of the sample, i.e., $\{y_i, \forall i \in \mathcal{N}\}$. This way, the solution to (10) under context $X = x$, which we denote as $z^{ML}(x)$, can be generically expressed as:

$$z^{ML}(x) \in \arg\min_{z \in Z} \quad \sum_{i \in \mathcal{N}} g^{ML}(x, x_i; w) f(z; y_i) \tag{11}$$

with $\{g^{ML}(x, x_i; w), \forall i \in \mathcal{N}\}$ being the probability masses that the specific proxy $\widehat{\mathcal{Q}}_{|x}$ that is used places on $\{y_i, \forall i \in \mathcal{N}\}$. These masses or weights are determined as a function $g^{ML}: \mathcal{X} \times \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}$ of the historical contextual information x_i, the current context x, and some parameters w.

In essence, this scheme adapts the Sample Average Approximation (a well-known data-driven solution strategy in Stochastic Programming [23,24]) to the case of *conditional* stochastic programs. It was first formalized in [1] and, since then, has been subject to a number of improvements (e.g., regularization procedures for bias-variance reduction [25]; robustification [26]; and algorithmic upgrades [27]) and extensions, e.g., to a dynamic decision-making setting [4]. Recently, the work in [11] introduces a bilevel formulation to optimally tune the parameters w that determine the weights $g^{ML}(x, x_i; w)$. Using our notation, the method proposed in [11] can be formulated as follows:

$$w^{ML} \in \arg\min_{w \in \mathbb{R}^q; \hat{z}_i} \sum_{i \in \mathcal{N}} f(\hat{z}_i; y_i) \tag{12a}$$

$$\text{s.t.} \quad \hat{z}_i \in \arg\min_{z \in Z} \sum_{i' \in \mathcal{N}: i' \neq i} g^{ML}(x_i, x_{i'}; w) f(z; y_{i'}), \quad \forall i \in \mathcal{N} \tag{12b}$$

where the function $g^{ML}: \mathcal{X} \times \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}$ used to compute the weights can be chosen from a catalog of several classical machine learning algorithms G^{ML} such as k-nearest neighbors, Nadaraya–Watson kernel regression or Random Forest. The author of [11] resorts to tailor-made approximations and greedy algorithms for each machine learning technique that is used to construct function g^{ML}, but do not provide a general-purpose solution strategy valid for any function g^{ML}. This approach, which is based on machine learning techniques, is called ML (stands for Machine Learning). After solving (12), the optimal decision $z^{ML}(x)$ under context $X = x$ is obtained by solving (11) with $w = w^{ML}$.

The surrogate problems (2) and (10) are, by design, different, in part because they are the result of distinct frameworks to address the conditional stochastic program (1). The surrogate problem (2) is based on the assumption that it is possible to find a good decision z in terms of the conditional expected cost

$\mathbb{E}[f(z;Y)|X=x]$ by optimizing that decision for a single scenario $\hat{y}$ of the uncertainty Y. Naturally, all the complexity of this approach lies in how to infer, from the data sample S, the single scenario $\hat{y}$ that unlocks the best decision z. This inference process makes use of global methods that consider all data points in the sample to obtain more robust decision mappings. In contrast, all the difficulty of the surrogate problem (10) rests on how to retrieve a good approximation of the true conditional distribution $\mathcal{Q}_{|x}$ from the sample S. Such an approximation is performed using local machine learning methods that only employ data close to the given context x and consequently, a large amount of data is required to avoid overfitting. In more practical terms, embedding local machine learning methods into the estimation problem (12) makes this problem computationally intractable in most cases. Besides, the surrogate problem (2) is computationally less demanding than (10), because the latter requires evaluating the cost function $f(z;\cdot)$ for multiple values of the uncertainty Y.

Finally, there is a third class of surrogate decision-making models that arises from the idea of using the sample S to directly learn the optimal decision z as a function of the context x, this way bypassing the need for constructing the estimate $\hat{y}$ or the proxy distribution $\widehat{\mathcal{Q}}_{|x}$. Following this logic, we seek a decision rule or mapping $g^{\text{DR}} : \mathcal{X} \times \mathbb{R}^q \to \mathbb{R}^n$ from a family G^{DR} so that $\arg\min_{z\in Z} \mathbb{E}[f(z;Y)|X=x] \approx \hat{z} = g^{\text{DR}}(x;w)$. Particularizing for the empirical distribution of the data, this approach renders:

$$w^{\text{DR}} \in \arg\min_{w\in\mathbb{R}^q} \sum_{i\in\mathcal{N}} f(g^{\text{DR}}(x_i;w);y_i) \tag{13a}$$

$$\text{s.t.} \quad g^{\text{DR}}(x_i;w) \in Z, \quad \forall i \in \mathcal{N} \tag{13b}$$

One clear advantage of directly learning the optimal decision policy is that, after solving (13), the decision z^{DR} to be implemented under context $X = x$ is efficiently computed as follows:

$$z^{\text{DR}}(x) = g^{\text{DR}}(x;w^{\text{DR}}) \tag{14}$$

Actually, the mapping (14) constitutes the surrogate decision-making model itself. This method, which aims at determining an optimal decision rule, is denoted as DR (acronym of Decision Rule). Nevertheless, feasibility issues may arise as this approach does not necessarily guarantee that the resulting z^{DR} obtained through (14) belongs to Z for any plausible context x. The authors of [3] propose and investigate this approach for the popular newsvendor problem, for which they consider a linear decision rule. Their newsvendor formulation does not involve any constraint and therefore, decisions yielded by (14) are always valid. However, the use of this approach is questionable for many other practical applications in which decisions must satisfy a set of constraints.

In summary, the contributions of the proposed bilevel model (7) with respect to the other approaches presented in this section are:

- Unlike the traditional approach (3), ours provides estimations of y by leveraging information about the optimization problem to be solved.
- Unlike the existing "predict-then-optimize" methodology (6), our approach is formulated as a generic bilevel optimization that, under convexity assumptions, is reformulated as a single-level optimization problem that can be solved using off-the-shelf optimization software.
- Unlike approach (12), ours makes uses of global estimation methods that use all available data to infer the point forecast of the uncertainty that unlocks the best decision. Therefore, our approach is less prone to overfitting, especially for small data samples. In addition, formulation (12) is more difficult to solve than (7).
- Unlike approach (13), ours guarantees the feasibility of the resulting optimal decision under any context.

3. Solution strategy

In this section, we elaborate on how to solve the general-purpose bilevel program (7) we propose to compute the best single scenario $\hat{y}$ to be fed into the surrogate problem (2). To do so, we particularize the generic formulation (1) as follows:

$$\min_{z,s(Y)} \mathbb{E}[f_0(z,s(Y);Y)|X=x] \tag{15a}$$

$$\text{s.t.} \ f_j(z,s(Y);Y) \le 0, \quad \forall j \in J \tag{15b}$$

$$h_k(z,s(Y);Y) = 0, \quad \forall k \in K \tag{15c}$$

where z constitutes the vector of *here-and-now* variables independent of the uncertainty, $s(Y)$ represents the *wait-and-see* decisions, and constraints (15b), (15c) must be satisfied for $\mathcal{Q}_{|x}$-almost all y (i.e., with probability one). We also assume that f_0, f_j are convex functions with respect to all variables, h_k are affine functions, and function g^{BL} is continuous in the parameter vector w.

Our method solves the following bilevel optimization problem:

$$w^{\text{BL}} \in \arg\min_{w\in\mathbb{R}^q;\hat{z}_i} \ \sum_{i\in\mathcal{N}} f_0(\hat{z}_i,\hat{s}_i;y_i) \tag{16a}$$

$$\text{s.t.} \ \ f_j(\hat{z}_i,\hat{s}_i;y_i) \le 0, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{16b}$$

$$h_k(\hat{z}_i,\hat{s}_i;y_i) = 0, \quad \forall k \in K, \forall i \in \mathcal{N} \tag{16c}$$

$$\hat{z}_i \in \{\arg\min_{z,s} \ \ f_0(z,s;g^{\text{BL}}(x_i;w)) \tag{16d}$$

$$\text{s.t.} \ \ f_j(z,s;g^{\text{BL}}(x_i;w)) \le 0, \quad \forall j \in J \tag{16e}$$

$$h_k(z,s;g^{\text{BL}}(x_i;w)) = 0, \quad \forall k \in K\}, \forall i \in \mathcal{N} \tag{16f}$$

On the assumption that the lower-level minimization problems (16d)–(16f) satisfy some constraint qualification, the classical strategy to solve (16) is to replace each lower level (16d)–(16f) with its equivalent Karush–Kuhn–Tucker (KKT) conditions [28], that is,

$$w^{\text{BL}} \in \arg\min_{w\in\mathbb{R}^q;\hat{z}_i,\lambda_{ji}} \ \sum_{i\in\mathcal{N}} f_0(\hat{z}_i,\hat{s}_i;y_i) \tag{17a}$$

$$\text{s.t.} \ \ f_j(\hat{z}_i,\hat{s}_i;y_i) \le 0, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{17b}$$

$$h_k(\hat{z}_i,\hat{s}_i;y_i) = 0, \quad \forall k \in K, \forall i \in \mathcal{N} \tag{17c}$$

$$\nabla f_0(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) + \sum_{j\in J}\lambda_{ji}\nabla f_j(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) +$$
$$+ \sum_{k\in K}\upsilon_{ki}\nabla h_k(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) = 0, \quad \forall i \in \mathcal{N} \tag{17d}$$

$$f_j(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) \le 0, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{17e}$$

$$h_k(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) = 0, \quad \forall k \in K, \forall i \in \mathcal{N} \tag{17f}$$

$$\lambda_{ji} \ge 0, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{17g}$$

$$\lambda_{ji}f_j(\hat{z}_i,s_i;g^{\text{BL}}(x_i;w)) = 0, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{17h}$$

where $\lambda_{ji}, \upsilon_{ki} \in \mathbb{R}$ are, respectively, the Lagrange multipliers related to constraints (16e) and (16f) for each lower-level problem; (17a)–(17c) are, in that order, the objective function and constraints of the upper-level problem, and constrains (17d)–(17g), (17h), are, respectively, the stationarity, primal feasibility, dual feasibility and slackness conditions of the lower-level problems. As discussed in [29], problem (17) violates the Mangasarian–Fromovitz constraint qualification at every feasible point and therefore, interior-point methods fails to find even a local optimal solution to this problem. To overcome this issue, a regularization approach was first introduced in [30] and further investigated in [31]. This method replaces all complementarity constraints (17h) with inequality (18c) below:

$$w^{\mathrm{BL}} \in \arg \min_{w \in \mathbb{R}^q; \hat{z}_i, \lambda_{ji}} \sum_{i \in \mathcal{N}} f_0(\hat{z}_i, \hat{s}_i; y_i) \tag{18a}$$

$$\text{s.t.} \quad (17b) - (17g) \tag{18b}$$

$$-\sum_{j \in J} \lambda_{ji} f_j(\hat{z}_i, s_i; g^{\mathrm{BL}}(x_i; w)) \leq \epsilon, \quad \forall i \in \mathcal{N} \tag{18c}$$

where ϵ is a small non-negative scalar that allows to reformulate (17) as the parametrized nonlinear optimization problem (18), which typically satisfies a constraint qualification and can be then efficiently solved by standard non-linear optimization solvers. Scholtes [30] prove that, as ϵ tends to 0, the solution of (18) tends to a *local* optimal solution of (17). In the remaining of the manuscript, we will refer to this approach as BL-R.

Alternatively, the complementarity slackness conditions can be linearized according to Fortuny–Amat [32] as follows:

$$w^{\mathrm{BL}} \in \arg \min_{w \in \mathbb{R}^q; \hat{z}_i, \lambda_{ji}, u_{ji}} \sum_{i \in \mathcal{N}} f_0(\hat{z}_i, \hat{s}_i; y_i) \tag{19a}$$

$$\text{s.t.} \quad (17b) - (17g) \tag{19b}$$

$$\lambda_{ji} \leq u_{ji} M^D, \quad \forall j \in J, \quad \forall i \in \mathcal{N} \tag{19c}$$

$$f_j(\hat{z}_i, s_i; g^{\mathrm{BL}}(x_i; w)) \geq (u_{ji} - 1) M^P, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{19d}$$

$$u_{ji} \in \{0, 1\}, \quad \forall j \in J, \forall i \in \mathcal{N} \tag{19e}$$

where u_{ji} are binary variables, and $M^P, M^D \in \mathbb{R}^+$ are large enough constants whose values can be determined as proposed in [33]. The resulting model (19) is a single-level mixed-integer non-linear problem. We denote this method as BL-M.

Solving the bilevel problem (7) using either BL-R or BL-M is valid for a conditional stochastic problem that satisfies the conditions described in this section. Nonetheless, the complexity of solving the regularized non-linear problem (18) or the mixed-integer non-linear program (19) highly depends on functions $f_0, f_j, h_k, g^{\mathrm{BL}}$. In some cases (see, for instance, the particular applications discussed in Section 4), problem (19) can be reformulated as a mixed-integer linear/quadratic optimization problem that can be solved to global optimality using standard optimization solvers. In the general case, problems (18) and (19) can also be solved using off-the-shelf optimization solvers, but global optimality may not be guaranteed. Notwithstanding this, local optimal solutions of the proposed bilevel formulation (7) may still lead to optimal decisions that are significantly better than those computed by FO or DR.

4. Applications

In Section 2, we introduce a common mathematical framework to present five different approaches for contextual decision-making under uncertainty, namely, the *predict-then-optimize* strategies FO, SP, and BL; method ML, which relies on a proxy of the true conditional distribution that is built using machine-learning techniques, and the decision-rule approach DR. Unfortunately, in the technical literature, methods SP and ML have only been applied to conditional stochastic optimization problems with a specific structure and they both lack a solution strategy for more general conditional stochastic programs. For this reason, in this section, we limit ourselves to comparing approaches FO, BL, and DR on various contextual decision-making problems under uncertainty, each of which illustrates a certain relevant aspect of our proposal. Section 4.1 compares these methodologies using the newsvendor problem, a well-known stochastic programming problem with simple recourse. The proposed methodology is also applied in Section 4.2 to the product placement problem, a two-stage stochastic programming problem with full recourse. Finally, Section 4.3 presents a strategic producer problem formulated as a one-stage stochastic programming in which the uncertainty only affects the objective function.

4.1. Newsvendor problem

We start with the popular newsvendor problem in the spirit of Ban and Rudin [3], a work that elicited renewed interest [11,12] in the solution to the conditional stochastic program (1). In the newsvendor problem, the goal of the decision maker is to find the optimal ordering quantity for a product with unknown random demand Y. In turn, this (positive) demand may be influenced by a random vector of features X representing, for instance, product information, weather conditions, customer profiles, etc. The decision maker has, therefore, a collection of observations $\{(x_i, y_i), \forall i \in \mathcal{N}\}$, which s/he would like to exploit to make an informed ordering quantity z under the context $X = x$. Let d and r, with $r > d > 0$, be the cost and revenue of manufacturing and selling one product unit, respectively. This problem can be formulated as the following conditional stochastic program:

$$\min_{z \in \mathbb{R}} \quad \mathbb{E}[dz - r \min(z, Y)|X = x] \tag{20}$$

Approaches FO and BL both follow a "predict-then-optimize" strategy, whereby the ordering quantity is obtained as the solution to the following surrogate decision-making model:

$$\min_{z \in \mathbb{R}} \quad dz - r \min(z, \hat{y}) \tag{21}$$

We can use an auxiliary variable s to get rid of the inner minimization and write instead:

$$\min_{z, s} \quad dz - rs \tag{22a}$$

$$\text{s.t.} \quad s \leq z \tag{22b}$$

$$s \leq \hat{y} \tag{22c}$$

whose solution is trivial, namely, $z^* = s^* = \hat{y}$.

FO and BL differ in the particular single value or scenario $\hat{y}$ that each of them uses. In the case of FO, $\hat{y}$ is an estimate of $\mathbb{E}[Y|X = x]$. Consequently, it becomes apparent that, for the newsvendor problem, approach FO is fundamentally inconsistent, because it is well-known that the solution to (20) corresponds to the quantile $\frac{r-d}{r}$ of the demand distribution Y conditional on $X = x$. Naturally, this quantile is generally different from $\mathbb{E}[Y|X = x]$.

Now, if we take $\hat{y} = g^{BL}(x; w) = w^T x$ in our approach, the optimal vector of linear coefficients w^{BL} is computed as follows:

$$w^{BL} \in \arg\min_{w \in \mathbb{R}^p} \quad \sum_{i \in \mathcal{N}} d\hat{z}_i - r \min(\hat{z}_i, y_i) \tag{23a}$$

$$\text{s.t.} \quad \hat{z}_i \in \{\arg\min_{z_i, s_i} \quad dz_i - rs_i \tag{23b}$$

$$\text{s.t.} \quad s_i \leq z_i \tag{23c}$$

$$s_i \leq w^T x_i\}, \forall i \in \mathcal{N} \tag{23d}$$

which, based on our previous argument, boils down to:

$$w^{BL} \in \arg\min_{w \in \mathbb{R}^p} \quad \sum_{i \in \mathcal{N}} d\hat{z}_i - r \min(\hat{z}_i, y_i) \tag{24a}$$

$$\text{s.t.} \quad \hat{z}_i = w^T x_i, \quad \forall i \in \mathcal{N} \tag{24b}$$

Therefore, our approach coincides exactly with that proposed in [3], which, in turn, is given by problem (13) in Section 2 when $g^{DR}(x; w) = w^T x$. This equivalence is far from being general though, as we will see with the other applications below.

4.2. Product placement problem

Given a graph $\mathcal{G} = (\mathcal{B}, \mathcal{A})$ with node-arc matrix A, in the product placement problem, the goal is to decide the amount $z_b \in \mathbb{R}^+$ of a certain product to be placed in each node $b \in \mathcal{B}$ of the grid [1]. After this decision is made, the demand for the product at each node y_b is realized, and the inventories of product throughout the network are shipped across the arcs $\mathcal{A}$ so as to satisfy the actually observed nodal demands. As in the newsvendor problem, these demands may be affected by some exogenous factors X that may be also random, but that are disclosed before the product placement decision is to be made. Let $h \in \mathbb{R}^{|\mathcal{B}|}$ and $g \in \mathbb{R}^{|\mathcal{A}|}$ be the cost of initially placing products in the nodes of the network and the cost of shipping products through the edges of the graph, respectively. The product placement problem under uncertain demand, but with contextual information, can be formulated as follows:

$$\min_{z \geq 0} \quad \mathbb{E}[c(z; Y)|X = x] \tag{25}$$

where

$$c(z; y) = h^T z + \min_{f \geq 0, p \geq 0} \quad g^T f + r^T p \tag{26a}$$

$$\text{s.t.} \quad Af \leq z - y + p \tag{26b}$$

In problem (26), we have included a variable vector $p \in \mathbb{R}^{|\mathcal{B}|}_{\geq 0}$ to allow for unsatisfied demand, with the associated penalty cost $r \in \mathbb{R}^{|\mathcal{B}|}$. Furthermore, the decision vector $f \in \mathbb{R}^{|\mathcal{A}|}$ represents the amount of product shipped across the arcs of the network. The cost function (26a) takes the form of a two-stage linear cost, with the integration of a recourse problem. More importantly, unlike in the newsvendor problem, the recourse is given by a full-fledged (linear) minimization problem. The surrogate decision-making model associated with the *predict-then-optimize* strategies FO and BL is as follows:

$$\min_{z \geq 0, f \geq 0, p \geq 0} \quad h^T z + g^T f + r^T p \tag{27a}$$

$$\text{s.t.} \quad Af \leq z - \hat{y} + p \tag{27b}$$

To ease the exposition and the notation that follows, we make the additional assumption that $r > h > 0$, where the inequality holds component-wise. In this case, variable vector p in (27) is zero at the optimum and the surrogate model can be simplified to:

$$\min_{z \geq 0, f \geq 0} \quad h^T z + g^T f \tag{28a}$$

$$\text{s.t.} \quad Af \leq z - \hat{y} \tag{28b}$$

As previously discussed, problem (28) is a deterministic mathematical program whereby the decision z is solely optimized for the point prediction of demand $\hat{y}$. While the traditional FO approach sets such a prediction to $\mathbb{E}[Y|X = x]$, the rationale behind the approach BL is to compute a W-parameterized function such that the surrogate problem (28) delivers the decision z that minimizes the in-sample cost, that is:

$$W^{BL} \in \arg\min_{W \in \mathbb{R}^{|\mathcal{B}| \times p}} \quad \sum_{i \in \mathcal{N}} h^T \hat{z}_i + g^T \hat{f}_i + r^T \hat{p}_i \tag{29a}$$

$$\text{s.t.} \quad A\hat{f}_i \leq \hat{z}_i - y_i + \hat{p}_i, \quad \forall i \in \mathcal{N} \tag{29b}$$

$$\hat{f}_i, \hat{p}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{29c}$$

$$\hat{z}_i \in \{\arg\min_{z_i \geq 0, f_i \geq 0} \quad h^T z_i + g^T f_i \tag{29d}$$

$$\text{s.t.} \quad Af_i \leq z_i - Wx_i\}, \quad \forall i \in \mathcal{N} \tag{29e}$$

where we have taken $\hat{y} = g^{BL}(x; W) = Wx$ with $W \in \mathbb{R}^{|\mathcal{B}| \times p}$. As discussed in Section 2, the lower-level problem (29d) and (29e) must have a unique solution. This can be guaranteed if, for example, all the shipping routes that can be taken to satisfy each demand in the graph entail a different cost. If this condition is not satisfied, the degeneracy of the lower-level problem can be eliminated by using classical results from the linear programming literature as described in [34]. As stated in Section 3, the solution to (29) can be addressed by replacing the lower-level linear program (29d) and (29e) with its KKT optimality conditions:

$$W^{BL} \in \arg\min_{W \in \mathbb{R}^{|\mathcal{B}| \times p}} \quad \sum_{i \in \mathcal{N}} h^T \hat{z}_i + g^T \hat{f}_i + r^T \hat{p}_i \tag{30a}$$

$$\text{s.t.} \quad A\hat{f}_i \leq \hat{z}_i - y_i + \hat{p}_i, \quad \forall i \in \mathcal{N} \tag{30b}$$

$$\hat{f}_i, \hat{p}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{30c}$$

$$0 \leq (h - \alpha_i) \perp \hat{z}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{30d}$$

$$0 \leq (g + A^T \alpha_i) \perp f_i \geq 0, \quad \forall i \in \mathcal{N} \tag{30e}$$

$$0 \leq (\hat{z}_i - Af_i - Wx_i) \perp \alpha_i \geq 0, \quad \forall i \in \mathcal{N} \tag{30f}$$

where $\alpha_i \in \mathbb{R}^{|\mathcal{B}|}$ is the vector of Lagrange multipliers associated with constraint (29e). Thus, problem (30) can be solved by regularizing the complementary slackness conditions or by using their Fortuny-Amat big-M reformulation. In the latter case, we arrive to a MIP problem that can be solved using commercial optimization software such as CPLEX or GUROBI.

Finally, if we also take a linear decision mapping $z(x) = g^{DR}(x; W) = Wx$ where $W \in \mathbb{R}^{|\mathcal{B}| \times p}$, the DR approach solves the following minimization problem to compute the optimal matrix of

linear coefficients W^{DR}:

$$W^{\text{DR}} \in \arg \min_{W \in \mathbb{R}^{|\mathcal{B}| \times p}} \quad \sum_{i \in \mathcal{N}} h^T \hat{z}_i + g^T \hat{f}_i + r^T \hat{p}_i \tag{31a}$$

$$\text{s.t.} \quad A\hat{f}_i \leq \hat{z}_i - y_i + \hat{p}_i, \quad \forall i \in \mathcal{N} \tag{31b}$$

$$\hat{f}_i, \hat{p}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{31c}$$

$$\hat{z}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{31d}$$

$$\hat{z}_i = W x_i, \quad \forall i \in \mathcal{N} \tag{31e}$$

It is apparent that the estimation problems (30) and (31), which BL and DR solve, respectively, are structurally different and so are W^{BL} and W^{DR} in general. For instance, think of a graph for which $\min\{g_\ell\}_{\ell \in \mathcal{A}} > \max\{h_b\}_{b \in \mathcal{B}}$. This represents a network where it is always cheaper to satisfy the nodal demand y_b, $b \in \mathcal{B}$, through the amount z_b of product that is initially placed at the demand location, that is, a graph where product shipping would be uneconomical if the nodal demands were certainly known in advance. Indeed, take the $\ell - th$ row of $g + A^T \alpha_i$ in Eq. (30e) for any $i \in \mathcal{N}$, that is, $g_\ell + \alpha_{o(\ell),i} - \alpha_{e(\ell),i}$, where $o(\ell)$ and $e(\ell)$ denote the origin and end nodes of arc ℓ, respectively. We have that $\inf\{g_\ell + \alpha_{o(\ell),i} - \alpha_{e(\ell),i} : \alpha_{o(\ell),i} \in [0, h_{o(\ell)}], \alpha_{e(\ell),i} \in [0, h_{e(\ell)}]\} = g_\ell - h_{e(\ell)} > 0$. Hence, $f_\ell = 0, \forall \ell \in \mathcal{A}$ and the system of inequalities (30d)–(30f) boils down to:

$$0 \leq (h - \alpha_i) \perp \hat{z}_i \geq 0, \quad \forall i \in \mathcal{N} \tag{32a}$$

$$0 \leq (\hat{z}_i - W x_i) \perp \alpha_i \geq 0, \quad \forall i \in \mathcal{N} \tag{32b}$$

which, unlike (31d) and (31e), allows for feasible solutions in the form $\hat{z}_{b,i} = 0$ with $w_b^T x_i < 0$ (and $\alpha_{i,b} = 0$), where w_b is the $b - th$ row of matrix W. Furthermore, recasting (31e) as $\hat{z}_i - W x_i = 0$ and setting $\alpha_i = h, \quad \forall i \in \mathcal{N}$, it is trivial to see that any feasible point of DR is also feasible for BL. Since the feasible region of (31) is contained in the feasible region of (30), but the opposite is not true, the optimum of (30) is in general lower than that of (31).

4.3. Strategic producer problem

Here we apply our decision-making framework to the problem of a strategic producer partaking in a forward market [16]. This strategic player must decide the produced quantity $q \in \mathbb{R}$ that maximizes her profits while facing some uncertainty on market conditions. Let $c(q) : \mathbb{R} \to \mathbb{R}^+$ denote the generation cost function whose parameters are assumed to be known with certainty. Let $p(q; Y) : \mathbb{R} \times \mathbb{R}^m \to \mathbb{R}$ represent the inverse demand function expressing the impact of the generation quantity q on the good's price. For some goods such as electricity, the inverse demand function varies depending on the season of the year, the day of the week, or the hour of the day. Besides, this function is also uncertain when producers must make their generation decisions q, since it may depend, for example, on weather conditions. If Q represents the known feasible region of variable q according to technical or economic constraints, the strategic producer must solve the following conditional stochastic optimization problem:

$$\min_{q \in Q} \quad \mathbb{E}[c(q) - p(q; Y)q | X = x] \tag{33}$$

As it is customary, we assume that the price and the demand are linearly related as $p(q; \alpha, \beta) = \alpha - \beta q$ where $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}^+$ are unknown parameters. Similarly, we assume that the production

cost is computed through a quadratic cost function $c(q) = c_2 q^2 + c_1 q$ where $c_1, c_2 > 0$ are known parameters related, respectively, to proportional production costs (such as fuel cost) and the increase of marginal costs due to technological factors (such as efficiency loss) [35]. In order to ease the notation, we define $\alpha' = \alpha - c_1$ and $\beta' = \beta + c_2$. Finally, we consider that the production quantity q is bounded by known capacity limits, i.e., $\underline{q} \leq q \leq \overline{q}$ with $\underline{q}, \overline{q} \in \mathbb{R}^+$. Thus, problem (33) can be reformulated as:

$$\min_{\underline{q} \leq q \leq \overline{q}} \quad \mathbb{E}[\beta' q^2 - \alpha' q | X = x] \tag{34}$$

Since the quantity decision q is independent of the outcome of the uncertainty (β', α'), the above can be further simplified to:

$$\min_{\underline{q} \leq q \leq \overline{q}} \quad \mathbb{E}[\beta' | X = x] q^2 - \mathbb{E}[\alpha' | X = x] q \tag{35}$$

Therefore, the optimal solution q^* is driven by the conditional expected values of α' and β'. To be more precise, since $\beta' > 0$, q^* could be equivalently computed as follows:

$$q^*(x) \in \arg \min_{\underline{q} \leq q \leq \overline{q}} \quad q^2 - \frac{\mathbb{E}[\alpha' | x]}{\mathbb{E}[\beta' | x]} q \Rightarrow q^*(x) \in \left\{ \underline{q}, \frac{\mathbb{E}[\alpha' | x]}{2\mathbb{E}[\beta' | x]}, \overline{q} \right\} \tag{36}$$

Unfortunately, $\mathbb{E}[\alpha' | x]$ and $\mathbb{E}[\beta' | x]$ are both unknown and therefore, they need to be estimated somehow. As explained further in Section 5.2.1, the producer has available a set of historical observations $S = \left\{ (\alpha_i', \beta_i', x_i), \forall i \in \mathcal{N} \right\}$ with $\alpha_i' \in \mathbb{R}$, $\beta_i' \in \mathbb{R}^+$ and $x_i \in \mathbb{R}^p$ in order to accomplish such a task. At this point, it should be underlined that the strategic producer problem (33) is of a distinctly different nature from that of the newsvendor problem (20) and the product placement problem (25). Indeed, the conditional stochastic program (33) *has no recourse* and the uncertain parameters appear only in its objective function. Consequently, solving (33) is apparently as "simple" as estimating the two conditional expectations $\mathbb{E}[\alpha' | x]$ and $\mathbb{E}[\beta' | x]$. Our claim, however, is that the way the producer draws decisions from a *finite data sample* (all we usually have in practice) may have a significant impact on the actual expected performance of the producer's strategy. Actually, the best estimates of $\mathbb{E}[\alpha' | x]$ and $\mathbb{E}[\beta' | x]$ from a *statistical* sense do not necessarily result in the best offer q.

According to the *predict-then-optimize* strategies, the surrogate model of this problem is formulated as follows:

$$\min_{\underline{q} \leq q \leq \overline{q}} \quad \hat{\beta}' q^2 - \hat{\alpha}' q \tag{37}$$

As explained in Section 2, the traditional approach aims at learning the uncertain parameters α_i', β_i' as a function of the available information x_i. If we assume the family of linear functions, that is, $\hat{\alpha}_i' = w_\alpha^T x_i$, $\hat{\beta}_i' = w_\beta^T x_i$ with $w_\alpha, w_\beta \in \mathbb{R}^p$, and we choose the squared error as the loss function l^{FO}, then the standard implementation of (3) is:

$$w_\alpha^{\text{FO}} \in \arg \min_{w_\alpha \in \mathbb{R}^p} \sum_{i \in \mathcal{N}} (\alpha_i' - w_\alpha^T x_i)^2 \tag{38a}$$

$$w_\beta^{\text{FO}} \in \arg \min_{w_\beta \in \mathbb{R}^p} \sum_{i \in \mathcal{N}} (\beta_i' - w_\beta^T x_i)^2 \tag{38b}$$

The optimal quantity under context $X = x$ is the solution to the following optimization problem:

$$q^{\text{FO}}(x) \in \arg \min_{\underline{q} \leq q \leq \overline{q}} \quad (w_\beta^{\text{FO}})^T x q^2 - (w_\alpha^{\text{FO}})^T x q$$

$$\Rightarrow q^{\text{FO}}(x) \in \left\{ \underline{q}, \frac{(w_\alpha^{\text{FO}})^T x}{2(w_\beta^{\text{FO}})^T x}, \overline{q} \right\} \tag{39}$$

Alternatively, w_α and w_β can be determined following the proposed approach by solving the following bilevel formulation:

$$w_\alpha^{\mathrm{BL}}, w_\beta^{\mathrm{BL}} \in \arg\min_{w_\alpha, w_\beta \in \mathbb{R}^p} \quad \sum_{i \in \mathcal{N}} \beta_i' \hat{q}_i^2 - \alpha_i' \hat{q}_i \tag{40a}$$

$$\text{s.t.} \quad \hat{q}_i \in \arg\min_{\underline{q} \le q_i \le \overline{q}} w_\beta^T x_i q_i^2 - w_\alpha^T x_i q_i, \quad \forall i \in \mathcal{N} \tag{40b}$$

For this particular application, the bilevel optimization problem rendered by the proposed approach has a significant drawback, because the global optimal solution of (40) is $w_\alpha = w_\beta = 0$. Consequently, the lower-level problem (40b) can be replaced by the feasibility condition $\underline{q} \le \hat{q}_i \le \overline{q}$, and the optimal values of $\hat{q}_i$ are determined as if uncertain parameters α' and β' were known in advance, which violates non-anticipativity. While this solution does lead to the minimum value of objective function (40a), it is useless to determine the optimal decisions for any context $X = x$. This degenerate solution of the proposed approach occurs because all coefficients of the objective function (37) are uncertain. Interestingly, this shortcoming does not affect the newsvendor and product placement problems, because the uncertainty only affects the feasible region in those applications.

In this paper, we propose to ensure non-anticipativity by formulating a bilevel optimization problem that considers the following modified surrogate model:

$$\min_{\underline{q} \le q \le \overline{q}} \quad q^2 - \gamma q \tag{41}$$

where $\gamma = \frac{\alpha'}{\beta'}$. For known values of α' and β', the optimal solution of (37) and (41) coincide. However, surrogate model (41) is simpler since it only includes one uncertain parameter instead of two. Assuming a linear relationship between the new uncertain parameter γ and the contextual information, the proposed methodology yields the following bilevel problem:

$$w_\gamma^{\mathrm{BL}} \in \arg\min_{w_\gamma \in \mathbb{R}^p} \quad \sum_{i \in \mathcal{N}} \beta_i' \hat{q}_i^2 - \alpha_i' \hat{q}_i \tag{42a}$$

$$\text{s.t.} \quad \hat{q}_i \in \arg\min_{\underline{q} \le q_i \le \overline{q}} q_i^2 - w_\gamma^T x_i q_i, \quad \forall i \in \mathcal{N} \tag{42b}$$

Formulation (42) has the following advantages with respect to (40): (i) it includes fewer parameters and therefore, it is less prone to overfitting, (ii) it ensures non-anticipativity for any parameter vector w_γ, and (iii) under certain conditions, it is able to retrieve the true model that relates random variable γ and the context X and the optimal solution to (34) as the sample size $|\mathcal{N}|$ grows to infinity, as shown in Proposition 1 in Appendix A. By replacing the lower-level problem with its KKT conditions, we obtain the following single-level problem:

$$w_\gamma^{\mathrm{BL}} \in \arg\min_{w_\gamma \in \mathbb{R}^p} \quad \sum_{i \in \mathcal{N}} \beta_i' \hat{q}_i^2 - \alpha_i' \hat{q}_i \tag{43a}$$

$$\text{s.t.} \quad 2\hat{q}_i - w_\gamma^T x_i - \underline{\lambda}_i + \overline{\lambda}_i = 0, \quad \forall i \in \mathcal{N} \tag{43b}$$

$$0 \le (\hat{q}_i - \underline{q}) \perp \underline{\lambda}_i \ge 0, \quad \forall i \in \mathcal{N} \tag{43c}$$

$$0 \le (\overline{q} - \hat{q}_i) \perp \overline{\lambda}_i \ge 0, \quad \forall i \in \mathcal{N} \tag{43d}$$

where $\underline{\lambda}_i, \overline{\lambda}_i$ are the dual variables corresponding to the capacity limit constraints. Notice that if complementarity conditions (43c) and (43d) are reformulated using the Fortuny–Amat approach, problem (43) can be solved to global optimality as a quadratic mixed-integer program using off-the-shelf optimization

Table 1
Data sample S for the illustrative example.

i	x_i	α_i'	β_i'	γ_i
1	2	2	10	0.20
2	4	17	10	1.70
3	8	8	3	2.67
4	9	16	6	2.67

software. According to this procedure, optimal decisions under context $X = x$ are made by solving:

$$q^{\mathrm{BL}}(x) \in \arg\min_{\underline{q} \le q \le \overline{q}} \quad q^2 - (w_\gamma^{\mathrm{BL}})^T x q \Rightarrow q^{\mathrm{BL}}(x) \in \left\{ \underline{q}, \frac{(w_\gamma^{\mathrm{BL}})^T x}{2}, \overline{q} \right\} \tag{44}$$

Finally, we can directly learn the optimal production as a function of the known information as proposed in [3]. Assuming the linear mapping $\hat{q}_i = w_q^T x_i$ with $w_q \in \mathbb{R}^p$, problem (13) for this particular application is formulated as:

$$w_q^{\mathrm{DR}} \in \arg\min_{w_q \in \mathbb{R}^p} \sum_{i \in \mathcal{N}} \beta_i' (w_q^T x_i)^2 - \alpha_i' w_q^T x_i \tag{45a}$$

$$\text{s.t.} \quad \underline{q} \le w_q^T x_i \le \overline{q} \quad \forall i \in \mathcal{N} \tag{45b}$$

Formulation (45) is a convex quadratic optimization problem and can be then solved using commercial software such as CPLEX. In line with (14), the optimal quantity under context $X = x$ is directly computed as:

$$q^{\mathrm{DR}}(x) = (w_q^{\mathrm{DR}})^T x \tag{46}$$

Although not true in general, approaches (43) and (45) may lead to the same solution under specific conditions. For instance, if the produced quantity q is not limited by minimum/maximum bounds, then constraint (43b) boils down to $\hat{q}_i = w_\gamma^T x_i/2$. Consequently, the solutions of (43) and (45) satisfy that $w_q^{\mathrm{DR}} = w_\gamma^{\mathrm{BL}}/2$ and therefore, $q^{\mathrm{BL}}(x) = q^{\mathrm{DR}}(x)$ for any context $X = x$. As we show in the following section, the decisions q^{BL} delivered by our approach are significantly more profitable than q^{DR} in the *constrained* case.

5. Numerical simulations

As an additional contribution, we assess and compare the performance of the proposed approach for the strategic producer problem using numerical simulations. In Section 5.1 we illustrate the advantages of BL with respect to FO and DR using a small example with a reduced data sample. Additionally, Section 5.2 presents the numerical results of a realistic case study that uses real data from the Iberian Electricity Market and the Spanish Transmission System Operator [36,37].

5.1. Illustrative example

This section aims at gaining insight into the performance of the proposed approach with a small example of the strategic producer problem. For the sake of simplicity, we only consider four realizations of the uncertain parameters α_i', β_i' and a single feature $x_i \in [0, 10]$, whose values are shown in Table 1. Approach FO predicts the uncertain parameters using linear functions in the form $\hat{\alpha}_i = w_{\alpha,0}^{\mathrm{FO}} + w_{\alpha,1}^{\mathrm{FO}} x_i$ and $\hat{\beta}_i = w_{\beta,0}^{\mathrm{FO}} + w_{\beta,1}^{\mathrm{FO}} x_i$; approach BL assumes that $\hat{\gamma}_i = w_{\gamma,0}^{\mathrm{BL}} + w_{\gamma,1}^{\mathrm{BL}} x_i$; and approach DR considers $\hat{q}_i = w_{q,0}^{\mathrm{DR}} + w_{q,1}^{\mathrm{DR}} x_i$. These three approaches are compared with a benchmark method (BN) that assumes perfect knowledge of the uncertain parameters α', β' and, consequently, yields the best possible
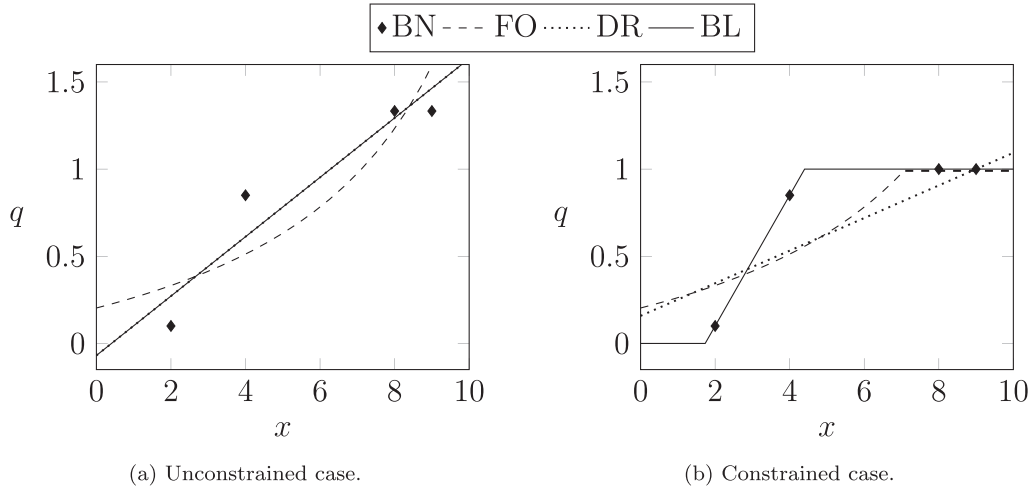
Fig. 1. Decision quantity q versus feature x for the illustrative example.

Table 2
Optimal offer and income for the unconstrained illustrative example (in-sample results). Parameter vectors w are: $w^{FO}_{\alpha,0} = 5.000$, $w^{FO}_{\alpha,1} = 1.000$, $w^{FO}_{\beta,0} = 12.298$, $w^{FO}_{\beta,1} = -0.878$, $w^{BL}_{\gamma,0} = -0.138$, $w^{BL}_{\gamma,1} = 0.341$, $w^{DR}_{q,0} = -0.069$, $w^{DR}_{q,1} = 0.170$.

	q_1	q_2	q_3	q_4	I(€)	RI(%)
BN	0.10	0.85	1.33	1.33	23.33	100.0
FO	0.33	0.51	1.23	1.59	21.21	91.0
DR	0.27	0.61	1.29	1.46	22.36	95.9
BL	0.27	0.61	1.29	1.46	22.36	95.9

Table 3
Optimal offer and income for the constrained illustrative example (in-sample results). Parameter vectors w are: $w^{FO}_{\alpha,0} = 5.000$, $w^{FO}_{\alpha,1} = 1.000$, $w^{FO}_{\beta,0} = 12.298$, $w^{FO}_{\beta,1} = -0.878$, $w^{BL}_{\gamma,0} = -1.300$, $w^{BL}_{\gamma,1} = 0.750$, $w^{DR}_{q,0} = 0.158$, $w^{DR}_{q,1} = 0.094$.

	q_1	q_2	q_3	q_4	I(€)	RI(%)
BN	0.10	0.85	1.00	1.00	22.33	100.0
FO	0.33	0.51	1.00	1.00	20.65	92.5
DR	0.35	0.53	0.91	1.00	20.50	91.8
BL	0.10	0.85	1.00	1.00	22.33	100.0

offer for each time period. Obviously, this method cannot be implemented in practice and, accordingly, is just used here for comparison purposes. Given the reduced size of this example, methods BL-R and BL-M provide the same results and are thus jointly referred to as BL.

First, we deal with the *unconstrained case*, that is, the case in which the capacity constraints are disregarded. Table 2 shows the in-sample results obtained from methods BN, FO, DR, and BL, namely, the optimal production quantity for each time period q_i, the absolute income (I), and the relative income with respect to the benchmark (RI). Notice that the income for each time period can be computed as $-\beta'_i q_i^2 + \alpha'_i q_i$. As discussed in Section 4.3, in connection with the unconstrained case, coefficients w^{DR} are equal to $w^{BL}/2$ and the decisions and incomes obtained by DR and BL are the same as a result. It is also interesting that the income of these two methods is 5% higher than that of FO. To explain this, we refer to Fig. 1a, which depicts the optimal production quantities given by the different methods as a function of the context $x \in [0, 10]$, namely,

$$q^{FO}(x) = \frac{w^{FO}_{\alpha,0} + w^{FO}_{\alpha,1} x}{2(w^{FO}_{\beta,0} + w^{FO}_{\beta,1} x)} \quad q^{BL}(x) = \frac{w^{BL}_{\gamma,0} + w^{BL}_{\gamma,1} x}{2}$$

$$q^{DR}(x) = w^{DR}_{q,0} + w^{DR}_{q,1} x \tag{47}$$

This figure shows that methods BL and DR can return decisions much closer to the benchmark ones than method FO for the four data points in the sample. Therefore, even for unconstrained optimization problems, the proposed methodology may outperform the classical "first-predict-then-optimize" approach, which is purely based on reducing the error of forecasting the uncertain parameters, simply because minimizing this error is not necessarily aligned with maximizing the decision value.

Now we consider the *constrained case*, that is, we bring the capacity constraints back into this small example. In particular, the

minimum and maximum outputs of the strategic producer are set to 0 and 1, respectively. Similarly to Table 2, the in-sample results obtained in the capacity-constrained case are collated in Table 3, where we can see that the optimal quantity q_i reaches its maximum value for some time periods and methods FO, DR and BL all provide different results. Methods FO and DR achieve an income 7.5% and 8.2% lower than the benchmark. This poor in-sample performance is better understood by means of Fig. 1b, which similarly to Fig. 1a, represents the optimal quantities as a function of the context x for the constrained case according to (39), (44) and (46). First, since method FO is unaware of the feasibility region of the original conditional stochastic problem, it provides the same prediction of the uncertain parameters α, β in the unconstrained and constrained cases. However, using these forecasts in the surrogate model (37) enforces $q = 1$ for $x \geq 7.1$ in the constrained case. As observed, reducing the forecast error of α, β does not lead to the maximization of the decision value in the constrained case either. Second, method DR must ensure feasible solutions for all samples, a condition that also leads to quite poor approximations of the optimal quantities for most values of the context x. Furthermore, this approach would return infeasible solutions $q > 1$ for $x > 9$ as shown in Fig. 1b. On the contrary, the proposed approach BL can find a linear relation between γ and x to be used in the surrogate model (41) that results in decisions q that perfectly match those provided by the benchmark for the four data points and therefore, this method achieves the highest possible income in sample.

In summary, this small example sheds light on the reasons why the proposed methodology outperforms existing ones for both unconstrained and constrained optimization problems under uncertainty: Our approach provides forecasts of the uncertain parameters that take into account the objective function and feasible region of the decision maker. Such enhanced forecasts translate into
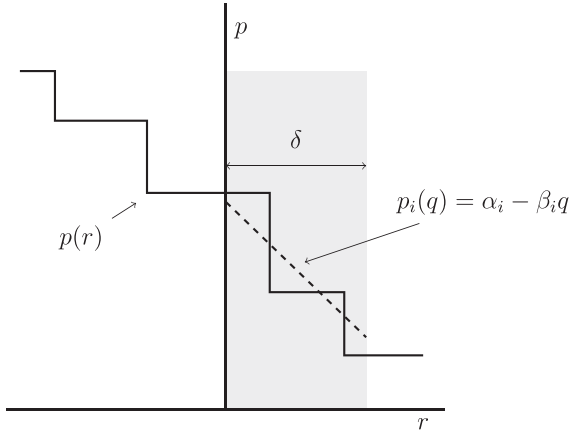
Fig. 2. Inverse residual demand curve $p(r)$ (solid) and fitted inverse demand function $p_i(q)$ (dashed) in the interval $[0, \delta]$. The intercept and slope of the fitted line are α_i and $-\beta_i$, respectively.

decisions that are much closer to those obtained in the ideal perfect information instance.

5.2. Case study

In this section, we compare the proposed approach with existing ones using realistic data from the Iberian electricity market, as described in detail in Section 5.2.1. Sections 5.2.2–5.2.4 investigate how the type of generation portfolio, the quadratic cost term c_2, and the residual demand elasticity impact the performance of the proposed methodology, respectively. These three sections only include the global optimal solutions given by method BL-M. Finally, Section 5.2.5 provides computational solution times for all the approaches and discusses the differences between BL-R and BL-M in that respect.

5.2.1. Experimental setup

In order to test our proposal, we consider a realistic case study based on actual data from the Iberian electricity market. We construct a data set of the form $\{(x_i, \alpha_i, \beta_i), \forall i \in \mathcal{N}\}$ from which we derive the rest of the parameters required for our simulations as explained in Section 4.3. We gather raw market data from the forward market OMIE [36] to compute parameters α_i, β_i of the inverse demand function. Furthermore, we collect wind and solar power forecasts of the aggregated production of Spain to be used as a vector of contextual information x_i. The wind and solar forecasts, originally published by the Spanish TSO, are downloaded from the ENTSO-e Transparency Platform [37].

Historical raw hourly block-wise bids and offers submitted by buyers and sellers to the Iberian day-ahead energy market are processed to obtain parameters α_i, β_i as follows. For each hour of the year, we have access to the set of bids and offers defined as $\{(q_b, p_b), \forall b \in B\}$, $\{(q_o, p_o), \forall o \in O\}$, respectively, where $q_{b/o}$ is the amount of energy to buy/sell at price $p_{b/o}$. Thus, the residual demand r to be potentially covered by a new producer entering the market for each possible price p is defined as $r := \sum_{b \in B: p_b \geq p} q_b - \sum_{o \in O: p_o \leq p} q_o$, that is, the aggregated demand minus the aggregated production. The step-wise function relating the residual demand r and the electricity price p is plotted in Fig. 2 for illustrative purposes.

Now consider that a new strategic producer enters the market with an offer to sell quantity q at offer price 0. If we assume that the remaining bids and offers stay unaltered, the market price would decrease following the right-hand part of the step-wise function depicted in Fig. 2. Therefore, a strategic producer

aiming at maximizing her profit is interested in modeling the dependence between her offered quantity q and the market price p in the shaded area, with parameter δ being a constant sufficiently larger than the producer's maximum generation capacity. In connection with Section 4.3, we approximate said dependency using a linear function such that $p_i(q) = \alpha_i - \beta_i q$ as illustrated in Fig. 2 and therefore, the values of α_i, β_i for each hour are obtained by determining the linear function that best approximates the blocks shaded in gray.

We collect data from November 2018 to October 2019 in order to build a data set of 8600 h (almost one year), which is divided into 43 bins of 200 consecutive samples. Each bin is randomly split into training and test sets with a ratio of 80%/20%, respectively. This process is repeated five times for each bin. Therefore, each approach is solved for 215 different training sets of 160 samples, and the obtained solutions are evaluated using the corresponding 215 test sets of 40 samples each. The out-of-sample results provided in Sections 5.2.2–5.2.5 are obtained by averaging the outcomes over these 215 test sets. We select a value of δ equal to 5 GW in order to encompass enough bids and offers to obtain accurate approximations of $p_i(q)$ throughout the whole data set. We determine the optimal parameters w through problems (38), (43), and (45), which we denote FO, BL and DR, respectively. More specifically, we name BL-M the Fortuny–Amat big-M reformulation of model (43) and BL-R the regularized counterpart, which are a particularization of (17) and (18), respectively.

Each bin is executed in parallel with the following resources: 1 CPUs Intel E5-2670 @ 2.6 GHz and 1 Gb of RAM. Each instance of model BL-M is solved using the MIQP solver CPLEX [38] for a maximum time of 20 min or a relative gap of 10^{-8}. On the other hand, BL-R is executed using the NLP solver CONOPT [39] without time limit.

5.2.2. Impact of the generation portfolio

As previously stated, the main advantage of our approach is that it yields forecast values for the uncertain parameters that are tailored to the optimization problem by which the strategic power producer determines her optimal market sale. However, such an advantage may translate into higher or lower incomes depending on the firm's generation portfolio. In this section, therefore, we evaluate the performance of the various approaches for three generic power plants characterized by different linear costs (c_1) and capacities ($\bar{q}$).

Table 4 provides the values of c_1, c_2 and $\bar{q}$ for these three generic units. For simplicity, the minimum output q of all units is assumed equal to 0 and the value of c_2 is set to 0.005 € /MWh² [35]. The base unit can represent a nuclear power station and is characterized by low fuel cost and high capacity. The medium unit can be, for example, a carbon-based power station with a lower capacity and higher fuel costs. Finally, peak units, such as combined cycle power plants, typically have the highest fuel cost and a smaller generation capacity. Table 4 also includes the percentage of time periods in which $q^{BN} = 0$, $0 < q^{BN} < \bar{q}$ and $q^{BN} = \bar{q}$ denoted as $\mathcal{N}_{q^{BN}=0}$, $\mathcal{N}_{0<q^{BN}<\bar{q}}$, and $\mathcal{N}_{q^{BN}=\bar{q}}$, respectively, where q^{BN} represents the optimal quantity that the strategic firm would place into the market under the *true* inverse demand function (that is, the solution given by the benchmark approach). It is observed that the base unit generates at maximum capacity for most times periods and is only shut down in 8% of the cases. The medium generating unit is idle 32% of the time (if prices are too low) and is at maximum capacity during the 39% of the time periods. Finally, the peak unit is not dispatched most of the time since electricity prices are usually below its marginal production cost.

Table 5 provides the out-of-sample results computed by averaging over the 215 test sets of 40 samples each described in Section 5.2.1. These results include the absolute income for the

Table 4
Generation technology data.

	c_1(€ /MWh)	c_2(€ /MWh^2)	$\bar{q}$(MW)	$\mathcal{N}_{q^{BN}=0}$(%)	$\mathcal{N}_{0<q^{BN}<\bar{q}}$(%)	$\mathcal{N}_{q^{BN}=\bar{q}}$(%)
Base	10	0.005	1000	8	16	76
Medium	35	0.005	500	32	29	39
Peak	50	0.005	250	79	12	9

Table 5
Impact of generation technology (Out-of-sample results).

	I^{BN}(M€)	RI^{FO}(%)	RI^{DR}(%)	RI^{BL-M}(%)	INFES^{DR}(%)
Base	176.7	96.1	94.6	96.3	4.9
Medium	20.9	77.4	62.5	80.0	1.7
Peak	1.2	44.1	18.9	58.7	0.1

Table 6
Income distribution for the peak generating unit (Out-of-sample results).

	BN	FO	DR	BL-M
$\mathcal{N}_{I>0}$(%)	20.6	9.0	6.6	10.1
$\mathcal{N}_{I<0}$(%)	0.0	3.7	3.0	3.4
$\mathcal{N}_{I=0}$(%)	79.4	87.3	90.4	86.4
I^+(M€)	1.23	0.73	0.37	0.87
I^-(M€)	0.00	-0.19	-0.14	-0.15

Table 7
Operating regime of a medium generating unit ($c_1 = 35$€ /MWh, $\bar{q} = 500$MW).

c_2(€ /MWh^2)	$\mathcal{N}_{q^{BN}=0}$(%)	$\mathcal{N}_{0<q^{BN}<\bar{q}}$(%)	$\mathcal{N}_{q^{BN}=\bar{q}}$(%)
0.01	32	43	25
0.005	32	29	39
0.001	32	15	53

Table 8
Impact of parameter c_2 on a medium generating unit (Out-of-sample results).

c_2(€ /MWh^2)	I^{BN}(M€)	RI^{FO}(%)	RI^{DR}(%)	RI^{BL-M}(%)	INFES^{DR}(%)
0.01	16.3	73.9	60.0	76.4	1.1
0.005	20.9	77.4	62.5	80.0	1.7
0.001	25.5	81.2	65.6	83.0	1.0

benchmark approach (I^{BN}) for the considered time horizon, the relative income (RI) for methods FO, DR and BL-M, and the percentage of time periods for which method DR provides infeasible solutions (INFES^{DR}). A first obvious observation is that, as expected, the absolute income is higher for base units and lower for peak units. A second, probably more interesting remark relates to the impact of the uncertainty about the inverse demand function on the market revenues accrued by each generating technology. Since the base unit is at full capacity most of the time, the uncertainty pertaining to the residual demand does not affect revenues that much, and the three methods obtain relative incomes above 94%. On the contrary, the participation of the medium and peak units highly depends on market conditions and therefore, this very same uncertainty remarkably deteriorates market revenues, with the eventual result that the maximum relative incomes amount to 80% and 59%, respectively, for the method featuring the best performance (which is BL-M).

On a different front, the DR approach produces infeasible offers in a considerable number of time periods, whereas FO and BL-M are guaranteed to provide feasible production quantities in all cases. The percentage of periods for which method DR results in an infeasible q is higher for the base unit because the medium and peak units are idle more frequently. For this particular application, making DR decisions feasible can be easily achieved by computing $\min(\max(\hat{q}_i, \underline{q}), \bar{q})$. However, this post-processing step to guarantee feasibility can be much more challenging in applications with general convex feasible sets. It is also apparent that the DR approach provides the lowest RI for the three cases considered and therefore, this method is not even recommended for decision-making models where the decision vector is simply bounded component-wise.

Finally, we notice that, for the three generation technologies, the proposed method BL-M always provides higher incomes than the FO approach. However, relative income improvements vary widely for each case. For the base unit, the relative income of BL-M is only 0.2% higher than that of FO. This is understandable since this power plant is at full capacity most of the time and thus, the impact of the uncertainty is comparatively minor, as we mentioned before. For the peak unit, in contrast, the relative income of BL-M is 14.6% higher than that of FO. Note that, unlike for base units, making small errors in the forecasts of the market conditions can be catastrophic for peak units, because such deviations may mean the difference between producing nothing or producing at maximum capacity. The ability of BL-M to reduce the forecast error when consequences are worse, together with the lower absolute incomes of peak units, explains this high difference in percentage. The gain of BL-M with respect to FO for the medium unit has an intermediate value of 2.6%.

To conclude this section, Table 6 includes, for the peak generating unit, the percentage of periods with a positive income, with

a negative income and with an income equal to zero, denoted as $\mathcal{N}_{I>0}$, $\mathcal{N}_{I<0}$ and $\mathcal{N}_{I=0}$, in that order. The total sum of positive and negative incomes is also provided in the last two rows, represented by the symbols I^+ and I^-, respectively. Interestingly, BL-M achieves the highest percentage of periods with a positive income and succeeds in providing the highest value of I^+.

5.2.3. Impact of parameter c_2

While parameter c_1 basically depends on the cost of the fuel used by each unit, the interpretation of c_2 is not as straightforward. Indeed, this parameter measures the decrease in the plant marginal cost as production increases and is connected to technological aspects of the plant's economy of scale, like the way the plant efficiency varies for different operating points. For this reason, in this section, we investigate the impact of c_2 on the performance of the proposed method. Notice that, if $\underline{q} = 0$ MW, then the unit marginal costs range from c_1 to $c_1 + c_2\bar{q}$. In a similar way, as Table 4 does, Table 7 shows the operating regime of a medium generating unit with $c_1 = 35$€ /MWh, $\bar{q} = 500$ MW and different values of c_2. As expected, a decrease in c_2 entails a reduction in the marginal production cost of the plant and, as a result, the amount of electricity the strategic firm places into the market increases.

Table 8 provides the same results as Table 5, but for different values of c_2 and the medium generating unit only. Naturally, reducing the plant marginal costs increases both the absolute income for the benchmark approach and also the relative income achieved by all methods. Nevertheless, BL-M proves to be between 1.8% and 2.6% more profitable to the producer than the traditional FO approach for the values of c_2 considered.

Table 9

Impact of residual demand elasticity (Out-of-sample results).

	$I^{BN}(M€)$	$RI^{FO}(\%)$	$RI^{DR}(\%)$	$RI^{BL-M}(\%)$	$INFES^{DR}(\%)$
Normal	20.9	77.4	62.5	80.0	1.7
Low-elast	18.6	74.7	60.0	77.1	1.7

Table 10

Comparison of BL-M and BL-R (Out-of-sample results).

	$RI^{BL-M}(\%)$	$RI^{BL-R}(\%)$
Base	96.3	96.3
Medium	80.0	79.2
Peak	58.7	58.4

Table 11

Average computing time.

	FO(s)	DR(s)	BL-R(s)	BL-M(s)
Base	0.24	0.65	3.90	197.77
Medium	0.35	1.06	6.80	149.89
Peak	0.26	0.78	4.62	22.68

light the higher computational burden required by BL-M to ensure global optimality. On the other hand, the computing times of BL-R are very affordable, especially considering the competitive edge that this method gives to the strategic firm.

5.2.4. Impact of the residual demand elasticity

So far we have centered our study on the cost structure of the generation portfolio owned by the strategic firm. Here, on the contrary, we focus on the elasticity of the market residual demand. Roughly speaking, this elasticity is inversely proportional to parameter β of the inverse demand function. Bearing this in mind, we compare the next two market situations, namely, the "Normal" and the "Low-elast" instances. The former corresponds to the values of β in the original data set, while the latter is obtained by multiplying these β-values by two.

Table 9 shows the incomes provided by each of the considered methods relative to those of the benchmark. The numbers correspond to the medium power plant of Table 4. The overall effect of increasing the residual demand elasticity (lower β-values) is analogous to that of decreasing parameter c_2, i.e., the involvement of the strategic producer in the market augments, thus leading to higher revenues. Results in Table 9 show that the proposed BL approach outperforms FO and DR for different values of the residual demand elasticity, improving the competitive edge of the strategic producer in more than 2% with respect to FO in terms of relative income.

5.2.5. Computational results

In Sections 5.2.2–5.2.4 we have only included results from BL-M, and not from BL-R, because the former variant of the bilevel framework we propose guarantees global optimality for the strategic producer problem for appropriate values of large constants M^P, M^D. However, solving model BL-M can be computationally very expensive. Alternatively, local optimal solutions of the proposed bilevel model (43) can be efficiently found by way of the particularization of the regularization approach (18) that we named BL-R.

Next, we first compare the solutions given by methods BL-M and BL-R. In order to solve model BL-R, we iteratively shrink the regularization parameter ϵ taking values from the discrete set $\{10^6, 10^4, 10^2, 1, 10^{-1}, 10^{-2}, 0\}$. In each iteration, we initialize the model with the solution provided by the previous problem. It is also worth mentioning that method BL-M is warm-started with the solution delivered by BL-R.

Results in Table 10 are intended to compare the relative incomes of BL-M and BL-R for each generating unit whose data is collated in Table 4. Although method BL-R logically yields lower incomes, the differences with respect to BL-M are below 0.8%. This means that if model (40) does not satisfy the conditions to be reformulated as a MIQP or the computational resources are limited, then a good solution (i.e., a solution with a small loss of optimality) can be efficiently computed by solving the regularized NLP version of our approach.

Finally, we compare the computational burden of methods FO, DR, BL-M, and BL-R. The average simulation time invested in solving problems (38), (43) and (45) for the three generation technologies are indicated in Table 11, where the maximum solution time has been limited to 20 min for all methods. These results high-

6. Conclusions

In this paper, we have addressed the problem of data-driven decision-making under uncertainty in the presence of contextual information. More precisely, our ultimate purpose has been to construct a parametric model to predict, based on some covariate information, the uncertain parameters that are input to the optimization model by which the decision is made. To this end, we have proposed a bilevel framework whereby such a parametric model is estimated taking into account the impact of its outputs on the feasibility and value of the decision. Under convexity assumptions, we have provided two single-level reformulations of the bilevel program, namely, a non-linear regularized optimization problem and a mixed-integer non-linear reformulation based on the use of large enough constants. When compared to alternative approaches available in the technical literature, ours features two major advantages: it guarantees feasibility in constrained decision-making problems, and its solution can be directly tackled using off-the-shelf optimization solvers under convexity assumptions.

We have theoretically compared our approach with existing ones for three different applications, namely, the newsvendor problem, the product placement problem, and the strategic producer problem. Additionally, we have evaluated the performance of our approach and its practical relevance through a realistic case study of a strategic producer that participates in the Iberian electricity market. Specifically, numerical results show that our framework not only significantly increases the revenue streams of the firm in general, but also proves to be critical to generation portfolios mainly consisting of peak power units. Indeed, the market revenues of a strategic peak generation portfolio are specially sensitive to the uncertainty in the inverse demand function. Therefore, in this case, the strategic firm may put at risk the bulk of its market incomes, by being left out of the market or trading in deficit. Our approach, however, is, by construction, aware of that sensitivity and thus, is able to retain most of the profit the firm would make under a perfectly predictable inverse demand function.

Potential extensions of this work would include the use of more advanced techniques in the resolution of our bilevel framework such as those employed in more general MPCC problems. Likewise, the generalization of our approach to multi-stage decision-making problems under uncertainty requires further analysis.

CRediT authorship contribution statement

M.A. Muñoz: Data curation, Software, Methodology, Investigation, Formal analysis, Validation, Writing – original draft. **S. Pineda:** Conceptualization, Methodology, Investigation, Formal analysis, Writing – original draft, Supervision. **J.M. Morales:** Conceptualization, Methodology, Investigation, Formal analysis, Writing – original draft, Supervision, Funding acquisition.

Acknowledgments

Appendix A. Asymptotic consistency

Proposition 1. *Let* $S = \left\{ (\alpha'_i, \beta'_i, x_i), \forall i \in \mathcal{N} \right\}$ *be an i.i.d sample of size N and suppose that there exists a linear relationship between α' and $\beta' > 0$ given by $\frac{\alpha'}{\beta'} = a^T x + \xi$, with ξ being a zero-mean noise independent of x, α' and β', and that the expectations $\mathbb{E}[\alpha']$, $\mathbb{E}[\beta']$ and $\mathbb{E}[\alpha'x]$ are all finite. Then, it almost surely holds in the limit $N \to \infty$ that the optimizer of the problem*

$$\min_{w_\gamma \in \mathcal{W}; \hat{q}_i} \quad \frac{1}{N} \sum_{i \in \mathcal{N}} \beta'_i \hat{q}_i^2 - \alpha'_i \hat{q}_i \tag{A.1a}$$

$$\text{s.t.} \quad \hat{q}_i \in \arg\min_{\underline{q} \le q_i \le \overline{q}} q_i^2 - w_\gamma^T x_i q_i, \quad \forall i \in \mathcal{N} \tag{A.1b}$$

with $\mathcal{W} \subset \mathbb{R}^p$ being a compact set containing a, is attained at $w_\gamma = a$.

Proof. First, notice that $\frac{\alpha'}{\beta'} = a^T x + \xi$ implies that $\frac{\mathbb{E}[\alpha'|x]}{\mathbb{E}[\beta'|x]} = a^T x$, since $\alpha' = \beta' a^T x + \beta' \xi$, and thus, $\mathbb{E}[\alpha'|x] = a^T x \mathbb{E}[\beta'|x]$ given the independent nature of the noise ξ.

The *true* expectation problem associated with the sample average approximation (A.1) is given by:

$$\min_{w_\gamma \in \mathcal{W}; \hat{q}(x)} \quad \int_{\mathcal{X} \times \mathbb{R}^+ \times \mathbb{R}} \left(\beta' \hat{q}^2(x) - \alpha' \hat{q}(x) \right) Q(dx, d\beta', d\alpha') \tag{A.2a}$$

$$\text{s.t.} \quad \hat{q}(x) \in \arg\min_{\underline{q} \le q \le \overline{q}} q^2 - w_\gamma^T x q, \quad \forall x \in \mathcal{X} \tag{A.2b}$$

where Q is the joint probability law governing the random parameters β' and α' and the feature vector X.

We first show that a is the unique solution to problem (A.2). To this end, we note that the lower-level problem (A.2b) renders the following decision mapping for almost all $x \in \mathcal{X}$:

$$\hat{q}(x) = \max\left(\underline{q}, \min\left(\frac{w_\gamma^T x}{2}, \overline{q} \right) \right)$$

which is a continuous function in w_γ.

Now let Q_X be the probability measure of X. Consider the following optimization problem, which is a relaxation of (A.2).

$$\min_{q(x) \in [\underline{q}, \overline{q}], \forall x \in \mathcal{X}} \quad \int_{\mathcal{X} \times \mathbb{R}^+ \times \mathbb{R}} \left(\beta' q^2(x) - \alpha' q(x) \right) Q(dx, d\beta', d\alpha') =$$

$$\min_{q(x) \in [\underline{q}, \overline{q}], \forall x \in \mathcal{X}} \quad \int_{\mathcal{X}} \left(q^2(x) \mathbb{E}[\beta'|x] - q(x) \mathbb{E}[\alpha'|x] \right) Q_X(dx) =$$

$$\int_{\mathcal{X}} \left(\min_{q(x) \in [\underline{q}, \overline{q}]} \quad q^2(x) \mathbb{E}[\beta'|x] - q(x) \mathbb{E}[\alpha'|x] \right) Q_X(dx)$$

The inner pointwise minimum results in the following optimal decision rule:

$$q(x) = \max\left(\underline{q}, \min\left(\frac{\mathbb{E}[\alpha'|x]}{2\mathbb{E}[\beta'|x]}, \overline{q} \right) \right) = \max\left(\underline{q}, \min\left(\frac{a^T x}{2}, \overline{q} \right) \right)$$

for almost all $x \in \mathcal{X}$.

Therefore, since $w_\gamma = a$ is feasible in the true expectation problem (A.2), then it is also an optimal solution to this problem. Furthermore, this solution is unique, if there exists a subset of $\mathcal{X}$ with measure greater than zero such that $\underline{q} < \frac{\mathbb{E}[\alpha'|x]}{2\mathbb{E}[\beta'|x]} < \overline{q}$.

In addition, note that all the samples in S are i.i.d. and that $\beta' q^2(x) - \alpha' q(x)$ is dominated by the function $\max\left(\beta'\overline{q}^2 - \alpha'\overline{q}, \beta'\underline{q}^2 - \alpha'\underline{q}, \frac{\alpha'^2}{4\beta'} \right)$, which is integrable because the expectations $\mathbb{E}[\alpha']$, $\mathbb{E}[\beta']$ and $\mathbb{E}[\alpha'x]$ are all finite. Indeed, since $\frac{\alpha'}{\beta'} = a^T x + \xi$ by assumption, we have that $\mathbb{E}[\frac{\alpha'^2}{4\beta'}] = \frac{1}{4}\mathbb{E}[\alpha'\frac{\alpha'}{\beta'}] = \frac{1}{4}\mathbb{E}[\alpha'(a^T x + \xi)] = \frac{a^T}{4}\mathbb{E}[\alpha'x]$.

Therefore, by invoking Theorems 5.3 and 7.48 in [40], we have that the minimizer of the sample average approximation problem (A.1) converges to a almost surely as the sample size N grows to infinity. □

References

[1] Bertsimas D, Kallus N. From predictive to prescriptive analytics. Manag Sci 2020;66(3):1025–44.

[2] Donti P, Amos B, Kolter JZ. Task-based end-to-end model learning in stochastic optimization. Adv Neural Inf Process Syst 2017:5484–94.

[3] Ban G-Y, Rudin C. The big data newsvendor: practical insights from machine learning. Oper Res 2019;67(1):90–108.

[4] Bertsimas et al.(2019)Bertsimas, McCord and Sturt Bertsimas D., McCord C., Sturt B.. Dynamic optimization with side information. arXiv preprint arXiv:1907073072019.

[5] Elmachtoub AN, Grigas P. Smart "predict, then optimize". Manag Sci 2021. doi:10.1287/mnsc.2020.3922.

[6] Esfahani PM, Kuhn D. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. Math Program 2018;171(1–2):115–66.

[7] Li B, Jiang R, Mathieu JL. Ambiguous risk constraints with moment and unimodality information. Math Program 2019;173(1–2):151–92.

[8] Hao Z, He L, Hu Z, Jiang J. Robust vehicle pre-allocation with uncertain covariates. Prod Oper Manag 2020;29(4):955–72.

[9] Keith AJ, Ahner DK. A survey of decision making and optimization under uncertainty. Ann Oper Res 2019:1–35.

[10] Bakker H, Dunke F, Nickel S. A structuring review on multi-stage optimization under uncertainty: aligning concepts from theory and practice. Omega 2020;96:102080.

[11] Mundru N. Predictive and prescriptive methods in operations research and machine learning: an optimization approach. Massachusetts Institute of Technology; 2019. Ph.D. thesis.

[12] Bertsimas D, Gupta V, Kallus N. Robust sample average approximation. Math Program 2018;171(1):217–82.

[13] Vives X. Duopoly information equilibrium: Cournot and Bertrand. J Econ Theory 1984;34(1):71–94.

[14] Wu J, Zhai X, Huang Z. Incentives for information sharing in duopoly with capacity constraints. Omega 2008;36(6):963–75.

[15] Bimpikis K, Ehsani S, Ilkılıç R. Cournot competition in networked markets. Manag Sci 2019;65(6):2467–81.

[16] Allaz B, Vila J-L. Cournot competition, forward markets and efficiency. J Econ Theory 1993;59(1):1–16.

[17] Ruiz C, Conejo A, García-Bertrand R. Some analytical results pertaining to Cournot models for short-term electricity markets. Electr Power Syst Res 2008;78(10):1672–8.

[18] Mandi et al.(2019)Mandi, Demirović, Stuckey, Guns et al. Mandi J., Demirović E., Stuckey P., Guns T., et al. Smart predict-and-optimize for hard combinatorial optimization problems. arXiv preprint arXiv:1911100922019.

[19] hao Kao Y, Roy BV, Yan X. Directed regression. In: Proceedings of the advances in neural information processing systems 22; 2009. p. 889–97.

[20] Wilder B, Dilkina B, Tambe M. Melding the data-decisions pipeline: decision–focused learning for combinatorial optimization. In: Proceedings of the AAAI conference on artificial intelligence, 33; 2019. p. 1658–65.

[21] Dempe S. Bilevel optimization: reformulation and first optimality conditions. Singapore: Springer Singapore; 2017. p. 1–20. ISBN 978-981-10-4774-9.

[22] Casorrán C, Fortz B, Labbé M, Ordóñez F. A study of general and security Stackelberg game formulations. Eur J Oper Res 2019;278(3):855–68.

[23] Kleywegt AJ, Shapiro A, Homem-de Mello T. The sample average approximation method for stochastic discrete optimization. SIAM J Optim 2002;12(2):479–502.

[24] Ruszczynski A., Shapiro A.. Stochastic programming (handbooks in operations research and management science). 2003.

[25] Ho C.P., Hanasusanto G.A.. On data-driven prescriptive analytics with side information: a regularized Nadaraya–Watson approach. 2019. URL: http://www.optimization-online.org/DB_FILE/2019/01/7043.pdf.

[26] Bertsimas D, Van Parys B. Bootstrap robust prescriptive analytics. Math Program 2017;2021:1–40.

[27] Diao S, Sen S. Distribution-free algorithms for learning enabled optimization with non-parametric estimation. Manag Sci 2020;66(3):1025–44.

[28] Boyd S, Boyd SP, Vandenberghe L. Convex optimization. Cambridge University Press; 2004.

[29] Scheel H, Scholtes S. Mathematical programs with complementarity constraints: stationarity, optimality, and sensitivity. Math Oper Res 2000;25(1):1–22.

[30] Scholtes S. Convergence properties of a regularization scheme for mathematical programs with complementarity constraints. SIAM J Optim 2001;11(4):918–36.

[31] Ralph D, Wright SJ. Some properties of regularization and penalization schemes for MPECs. Optim Methods Softw 2004;19(5):527–56.

[32] Fortuny-Amat J, McCarl B. A representation and economic interpretation of a two-level programming problem. J Oper Res Soc 1981;32(9):783–92.

[33] Pineda S, Bylling H, Morales J. Efficiently solving linear bilevel programming problems using off-the-shelf optimization software. Optim Eng 2018;19(1):187–211.

[34] Garcia J.D., Street A., de Mello T.H., Muñoz F.D.. Application-driven learning via joint prediction and optimization of demand and reserves requirement. 2021.

[35] Djurović MŽ, Milačić A, Kršulja M. A simplified model of quadratic cost function for thermal generators. In: Proceedings of the 23rd international DAAAM symposium, 23; 2012. p. 25–8.

[36] OMIE offer and demand aggregated curves. 2020. Available (online): https://www.omie.es/. Last accessed on February 6, 2020.

[37] ENTSO-E. Transparency platform. 2020. Available (online): https://transparency.entsoe.eu/. Last accessed on February 6, 2020.

[38] CPLEX optimizer. IBM; 2020. Available (online): https://www.ibm.com/analytics/cplex-optimizer.

[39] CONOPT optimizer. ARKI Consulting Development A/S; 2020. Available (online): http://www.conopt.com/.

[40] Shapiro A, Dentcheva D, Ruszczynski A. Lectures on stochastic programming: modeling and theory. SIAM; 2021.