This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TSG.2019.2891747

Learning from Past Bids to Participate Strategically in Day-Ahead Electricity Markets*

Ruidi Chen,† Ioannis Ch. Paschalidis,‡ *Fellow, IEEE,* Michael C. Caramanis,† *Senior Member, IEEE,* and Panagiotis Andrianesis†

***Abstract*—We consider the process of bidding by electricity suppliers in a day-ahead market context, where each supplier bids a linear non-decreasing function of her generating capacity with the goal of maximizing her individual profit given other competing suppliers' bids. Based on the submitted bids, the market operator schedules suppliers to meet demand during each hour and determines hourly market clearing prices. Eventually, this game-theoretic process reaches a Nash equilibrium when no supplier is motivated to modify her bid. However, solving the individual profit maximization problem requires information of rivals' bids, which are typically not available. To address this issue, we develop an inverse optimization approach for estimating rivals' production cost functions given historical market clearing prices and production levels. We then use these functions to bid strategically and compute Nash equilibrium bids. We present numerical experiments illustrating our methodology, showing good agreement between bids based on the estimated production cost functions with the bids based on the true cost functions. We discuss an extension of our approach that takes into account network congestion resulting in location-dependent prices.**

***Index Terms*—Day-ahead market, Equilibrium bids, Learning, Inverse equilibrium, Inverse optimization.**

I. Introduction

In the past several decades that followed the seminal work on spot market pricing [1], the electricity industry has evolved from vertical integrated regulated monopolies to competitive supply and demand market participants with equal access to a regulated transmission and distribution network. Nevertheless, due to special features of the power industry, including a limited number of producers (electricity suppliers), large capital investments that introduce barriers to entry, and congestion caused by occasionally binding transmission constraints, the electricity market is characterized by oligopolistic conditions [2]. Under perfect competition, suppliers would bid their marginal costs, a necessary condition for social welfare and efficiency maximization. In an imperfect oligopolistic energy market setting, however, suppliers can exploit market manipulation opportunities to increase their profits by bidding above their marginal cost.

The investigation of such behavior, referred to as *strategic bidding*, is of dual interest. First, to market participants (mainly suppliers), who are interested in devising optimal bidding strategies that would allow them to "outsmart" competitors and realize profits exceeding those that a perfectly competitive market would allow. Second, to market regulators, who are interested in identifying market power abuses and developing policies to increase efficiency and social welfare.

There is an immense amount of literature on strategic bidding in the context of electricity markets (see e.g., [3] for a related non-exhaustive literature review) particularly when one takes into account the specific market rules that apply. Currently, U.S. markets involve multi-part bids for energy and commitment costs, as well as several types of ancillary services, resulting in location-dependent hourly and real-time (5 min) prices. In the European day-ahead market coupling problem [4], even more complex bids are allowed. Regardless of the underlying framework and market rules, an optimal bidding strategy aims to answer the same question: how to bid in order to maximize profits.

From a game-theoretic point of view, the approaches for equilibrium analysis of the strategic bidding problem can be further classified as Bertrand models, Cournot and Stackelberg models, and *Supply Function Equilibrium (SFE)* models. In the latter approach, instead of setting their price bids (Bertrand) or quantities (Cournot), suppliers bid their supply functions that link prices with quantities. The SFE literature originates from the seminal work of Klemperer and Meyer [5], and, since its first application in electricity markets [6], it has been extensively studied [7], [8] — an overview is presented in [9], in both stylized examples and actual electricity markets (see, e.g., [10], [11], [12] for analytical and numerical results, and [13] for an empirical analysis).

One of the main criticisms regarding game-theoretic approaches is the unrealistic assumption that the payoff functions of all participants are publicly available. Most related works deal with this issue by assuming some type of uncertainty. An early work [14] proposes a recursive dynamic programming approach for determining the optimal bid price for each block of generation, in which each supplier models the uncertainty about rival bid prices by a probability distribution. In [15], the developed bidding scheme maximizes the hourly profit assuming all other producers' bids are represented by a multivariate normal distribution whose parameters are estimated from historical data. In [16], each supplier assumes types (based on the cost structure) of other suppliers and their joint probability distribution, based on the published information on fuel contracts, availability of transmission lines, and operating parameters. In [17], a decomposition-based

* Research partially supported by the NSF under grants DMS-1664644, CNS-1645681, CCF-1527292, AitF-9500307423, and IIS-1237022, by the ARO under grant W911NF-12-1-0390, by the ONR under grant N00014-16-1-2832, and by the Sloan Foundation under grant G-2017-9723.

† Division of Systems Engineering, Boston University, Boston, MA 02215, e-mail: {rchen15, mcaraman, panosa}@bu.edu.

‡ Department of Electrical and Computer Eng., Division of Systems Eng., and Dept. of Biomedical Eng., Boston University, 8 St. Mary's St., Boston, MA 02215, e-mail: yannisp@bu.edu, http://sites.bu.edu/paschalidis/.

particle swarm optimization method is proposed to solve the expected profit maximization problem with the market clearing price modeled as an uncertain, exogenous variable. In [18], a decentralized Nash equilibrium learning strategy is presented in a Bertrand competition framework to solve the economic dispatch problem. Recently, in [19], a Bayesian inference approach is proposed to reveal the aggregate supply curve in a day-ahead electricity market. In [20], a non-cooperative game with incomplete information among demand response aggregators is considered under different market conditions, where a Bayesian approach is used to estimate the unknown information such as the types of competitors. In [21], a multi-period market equilibrium problem is considered to study the strategic behavior of energy storage systems, where the optimality conditions of all participants' profit maximization problems are collected and solved together.

In this paper, we consider SFE-based equilibrium strategies for suppliers in the context of a day-ahead electricity market. We address the aforementioned criticism by estimating payoff functions using an *inverse optimization* approach combined with the theory of variational inequalities [22]. Inverse optimization seeks to recover input data to optimization problems from optimal solutions; it was first introduced in [23] but recently revisited in new settings [24], [22], [25]. Interestingly, inverse optimization has not been extensively used in the context of electricity markets. In [26], inverse optimization is used to identify the bids of marginal suppliers in a multi-period network-constrained electricity pool. In [27], it is employed to address the market-bidding problem of a cluster of price-responsive consumers of electricity. In [28], it is used to determine market structure from commodity and transportation prices; the methodology is applied to data from the MISO electricity market. Recently, [29] used inverse optimization to estimate how loads respond to demand response price signals.

A preliminary conference version of this paper has appeared in [30]. That paper focused on comparisons between inverse optimization-based strategic bidding versus earlier approaches [15]. The present paper uses a different, more realistic parametrization of the unknown cost function with respect to observable market variables, has an extensive numerical validation of the proposed approach, establishes new rigorous results on algorithm termination, and offers extensions to location-dependent prices.

To the best of our knowledge, we are the first to leverage inverse optimization for estimating cost parameters in payoff functions and obtaining equilibrium bids in the context of electricity markets. Our method is driven purely by data, in the sense that only the observed samples are utilized for inference and estimation, without relying on any distributional assumptions on the observed data. We note that any hypothesis on the data generating pattern could be questionable, due to the lack of supporting evidence on such assumption and the noisy nature of the data. By contrast, data-driven approaches receive input from the observed samples and are self-adjusted in the estimation process as more samples are available.

The main idea of this paper is to learn from past bids of electricity suppliers that bid strategically in a day-ahead market context. We develop an inverse optimization approach for estimating suppliers' cost functions, based on historical bidding data. We propose an algorithm that randomly searches for cost function parameters with good out-of-sample performance, among multiple possible values that are compatible with past data. Our proposed framework is validated with extensive numerical experimentation, and is extended to accommodate location-dependent prices.

The remainder of the paper is organized as follows. In Section II, we present the general market framework. In Section III, we formulate the strategic bidding problem (referred to as the "forward" problem), and in Section IV we present the inverse optimization framework as it applies to our day-ahead market setting. In Section V, we discuss the specific algorithm we use to estimate competitors' cost parameters, based on which equilibrium bids are obtained, and establish its convergence properties. In Section VI we illustrate our approach with numerical examples. In Section VII, we discuss the extension of our approach to location-dependent prices. We conclude and provide directions for further research in Section VIII.

II. Market Framework

We consider a day-ahead electricity market setting, which is composed of N electricity suppliers, and a market operator instantiating a *Power Exchange (PX)*. Each supplier submits a bid curve (supply curve) that describes the relationship between energy price and production quantity, for each of the 24 hours of the next day. After receiving the bidding functions from all suppliers, the market operator clears the market by balancing aggregate supply and demand; the output is the hourly market clearing price and the supplier specific dispatch schedules. Assuming no inter-temporal coupling in the PX setting, the auctions for different hours are performed separately and independently. This allows us to consider the bidding strategy for a specific hour and omit the time index in our analysis.

In actual power markets, the bidding functions are piecewise-constant curves, reflecting the constant bid price (marginal as-bid cost) for each block of electricity generation. These piecewise-constant curves correspond to piecewise-linear functions for the total as-bid costs of the suppliers, which approximate a quadratic cost function of typical generators. Piecewise-linear functions are used in practice to allow reliance on available commercial optimization solvers (for solving large-scale security-constrained unit commitment and economic dispatch problems typically formulated as mixed integer linear programming problems). In this paper, we assume the same affine bid curve as commonly used in the SFE literature; this assumption not only facilitates our analysis, but also corresponds to a quadratic approximation of generator cost functions.

Assume that supplier i submits a linear non-decreasing bid function to the market operator, $\alpha_i + \beta_i P_i$, $i = 1, \ldots, N$, that denotes the marginal as-bid cost of power at production level P_i, and α_i, β_i are the bidding coefficients to be determined under the optimal bidding strategy. After receiving these linear bidding functions, the market operator derives the clearing price and the generator dispatch schedule as follows:

$$\begin{aligned}
&R = \alpha_i + \beta_i P_i,\ i \in \mathcal{I}, \qquad (1)\\
&Q = Q^{\text{forecast}} - \textstyle\sum_{i\in\underline{\mathcal{I}}} P_i^{\min} - \sum_{i\in\overline{\mathcal{I}}} P_i^{\max},\\
&Q = \textstyle\sum_{i\in\mathcal{I}} P_i,\\
&\overline{\mathcal{I}} = \{i:\ R \geq \alpha_i + \beta_i P_i^{\max}\},\ \underline{\mathcal{I}} = \{i:\ R \leq \alpha_i + \beta_i P_i^{\min}\},\\
&\mathcal{I} = \{1,\ldots,N\} \setminus \{\overline{\mathcal{I}} \cup \underline{\mathcal{I}}\},
\end{aligned}$$

where R is the market clearing price, Q^{forecast} is the demand forecast (or as-bid load as is also the case in actual electricity markets) that is publicly announced by the market operator, $P_i^{\min}, P_i^{\max}$ are the minimum and maximum generation levels, respectively, of supplier i, $\overline{\mathcal{I}}$ is the set of suppliers producing at $P_i^{\max}$, $\underline{\mathcal{I}}$ is the set of suppliers producing at $P_i^{\min}$, $\mathcal{I}$ the set of marginal suppliers, and Q the effective demand met by marginal suppliers.

Since for $i \in \mathcal{I}$ the capacity constraints are not binding, for a given Q^{forecast} (hence, Q) the solution to (1) becomes

$$\begin{aligned}
&R(\boldsymbol{\alpha},\boldsymbol{\beta};Q) = \frac{Q + \sum_{i\in\mathcal{I}} \frac{\alpha_i}{\beta_i}}{\sum_{i\in\mathcal{I}} \frac{1}{\beta_i}}, \qquad (2)\\
&P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q) = \frac{R(\boldsymbol{\alpha},\boldsymbol{\beta};Q) - \alpha_i}{\beta_i},\ i \in \mathcal{I},\\
&P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q) = P_i^{\min},\ i \in \underline{\mathcal{I}},\ P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q) = P_i^{\max},\ i \in \overline{\mathcal{I}},
\end{aligned}$$

where we write $R(\boldsymbol{\alpha},\boldsymbol{\beta};Q), P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q)$ to explicitly express the dependency on $\boldsymbol{\alpha} \triangleq (\alpha_1,\ldots,\alpha_N)$ and $\boldsymbol{\beta} \triangleq (\beta_1,\ldots,\beta_N)$ for a given demand forecast resulting in effective demand Q.

III. Forward Problem Formulation

The forward problem deals with the individual profit maximization problem, in which supplier i determines her bidding curve (α_i, β_i) to maximize her profit $\phi_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q)$ defined as

$$\phi_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q) = R(\boldsymbol{\alpha},\boldsymbol{\beta};Q)P_i(\boldsymbol{\alpha},\boldsymbol{\beta},Q) - C_i(P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q)),$$

where $C_i(P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q))$ is the production cost at generation level $P_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q)$. The problem is formulated as follows:

$$\begin{aligned}
\max_{\alpha_i,\beta_i} \quad & \phi_i(\boldsymbol{\alpha},\boldsymbol{\beta};Q)\\
\text{s.t.} \quad & 0 \leq \alpha_i \leq \bar{\alpha},\\
& \beta_i > 0,
\end{aligned} \qquad (3)$$

where $\bar{\alpha}$ is an upper bound on α_i, and is related to the price cap in electricity markets. [1]

Note that the form of the profit function is generic (defined as revenues minus costs) and, in general, each supplier can have her own cost function. For the case of electricity generators, a common assumption is a quadratic cost function:

$$C_i(P_i) = c_{i0} + c_{i1}P_i + c_{i2}P_i^2, \qquad (4)$$

which implies a marginal cost equal to $c_{i1} + 2c_{i2}P_i$. Since the intercept of the quadratic cost function (parameter c_{i0}) is constant (this practically refers to the so-called "no-load" cost of a generator), its value will not affect the profit maximization problem of the supplier and, without loss of generality, it can be set to zero.

A direct comparison of the marginal cost $(c_{i1} + 2c_{i2}P_i)$ and the linear bid function $(\alpha_i + \beta_i P_i)$ indicates that the cost parameters c_{i1} and $2c_{i2}$ correspond to the bidding curve parameters α_i and β_i, respectively. Hence, truthful bidding would result in $\alpha_i = c_{i1}$ and $\beta_i = 2c_{i2}$.

In this paper, we assume that suppliers game only with parameter α_i, and that β_i is small and equal to $2c_{i2}$, representing a publicly known, technology-specific efficiency decline associated with approaching generating capacity. This assumption corresponds to the "bid-α" game in [31], implying that β_i is known to other suppliers for all intents and purposes. We elaborate on this assumption next.

The technology and capacity of individual generating plants is public information that provides useful partial information about their cost functions. Nevertheless, their fuel and variable maintenance cost, and their exact heat rate (efficiency), reflected primarily in parameter c_{i1}, is proprietary and not known with sufficient accuracy to competitors so as to allow them to bid optimally. On the other hand, as also pointed out in [32], the marginal cost functions of individual suppliers usually have very shallow slopes, and thus β_i is relatively small (about two orders of magnitude lower than α_i); furthermore, if both α_i and β_i can be chosen, the existence of a unique equilibrium is rare. Hence, one can argue that the small value of the slope of the marginal cost $(2c_{i2})$ is more or less known, and that the suppliers reflect this cost in parameter β_i, [2] as in [32], thus avoiding bidding non credible high slopes. Still, in our results section, we mainly explore cases in which we allow errors in the estimates of past bids or the knowledge of parameters (including parameter β_i). Furthermore, we note that if we consider the framework from the perspective of a regulator, the technology-specific data (e.g., heat-rate curves) that mainly drive the slope of the marginal cost are declared by the participants; as such, the slope is relatively easily calculated.

For the purposes of this paper, we assume that c_{i1} consists of two cost components: (i) one non-fuel cost component that reflects operational and maintenance variable costs (e.g., labor, parts, consumables, lubricants, chemicals, consumption from power station supplies, etc.), and (ii) a fuel-cost component (essentially depending on the heat rate and the fuel price). As such, c_{i1} is defined as

$$c_{i1} = \theta_{i1} + \theta_{i2}\xi, \qquad (5)$$

where ξ is a variable reflecting the publicly known fuel price, and θ_{i1}, θ_{i2} are the unknown cost coefficients. This decomposition is in line with the declared characteristics of the generation units, which comprise the heat rate curve and operational (other than fuel) and maintenance variable costs. We note that the framework can support even more detailed

[1] Strictly speaking, the price cap imposes a bound on the marginal cost at the maximum capacity, $\alpha_i + \beta_i P_i^{\max} \leq \bar{\alpha}$, implying that the upper bound is different for each supplier i, i.e., $\bar{\alpha}_i = \bar{\alpha} - \beta_i P_i^{\max}$. However, in practice, the price cap is high enough, and we can assume without loss of generality a common upper bound on α_i.

[2] The parameter β_i reflects (i) the smoothing/regularization of the bid conforming to the monotonically increasing market rule requirement (marginal costs are physically not strictly monotonic), and (ii) the advantage of and desire for achieving unique price-directed marginal generator schedules. In the experiments, the values of β_i are on the order of 0.1.

decompositions (e.g., consider separately a carbon price for emissions). Also, we note that the unknown cost coefficients can be interpreted in various ways. For instance, assuming a publicly known fuel price ξ, coefficient θ_{i2} may contain the combined effect of the heat rate curve and potential discounts that suppliers may have secured; such contract information is not available to either regulators or competitors.

Following the above assumptions, i.e., setting $c_{i0} = 0$ and $c_{i2} = (1/2)\beta_i$ in the quadratic cost function (4), and using (5), the profit function for supplier i can be rewritten as

$$\begin{aligned}\phi_i(\boldsymbol{\alpha}, \boldsymbol{\beta}; Q, \xi) =& [R(\boldsymbol{\alpha}, \boldsymbol{\beta}; Q) - (\theta_{i1} + \theta_{i2}\xi)]P_i(\boldsymbol{\alpha}, \boldsymbol{\beta}; Q) \\ & - \frac{1}{2}\beta_i[P_i(\boldsymbol{\alpha}, \boldsymbol{\beta}; Q)]^2. \end{aligned} \tag{6}$$

Note that for a given Q and ξ, the profit of supplier i is determined by the actions $\boldsymbol{\alpha}$ of all players and her own cost parameter $\boldsymbol{\theta}_i \triangleq (\theta_{i1}, \theta_{i2})$. We therefore write $\phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi)$ to emphasize this dependency ($\boldsymbol{\beta}$ is removed since it is constant and known).

Since all suppliers choose their bids by solving the profit maximization problem (3), we can construct a SFE model describing the game among all profit-maximizing suppliers. By definition, a specific $\boldsymbol{\alpha}$ is a Nash equilibrium if no single supplier can increase her profit by unilaterally changing her own bid. We know (see [30]) that there exists a unique Nash equilibrium $\boldsymbol{\alpha}^* = (\alpha_1^*, \ldots, \alpha_N^*)$ in this SFE model, since $\phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi)$ is strictly concave in α_i, i.e., its second partial derivative is strictly negative.

Next, we compute both the first and the second partial derivatives of the profit function with respect to α_i. From (6), the first derivative is

$$\begin{aligned}&\nabla_i \phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi) = \\ &\quad \frac{1}{\beta_i}[\hat{\beta}_i Q_i + \alpha_i(\hat{\beta}_i^2 - 1) - (\hat{\beta}_i - 1)(\theta_{i1} + \theta_{i2}\xi)], \end{aligned} \tag{7}$$

where $0 < \hat{\beta}_i = (1/\beta_i)/\sum_{l\in\mathcal{I}}(1/\beta_l) < 1$, and $Q_i = (Q + \sum_{k\in\mathcal{I}, k\neq i} \alpha_k/\beta_k)/\sum_{l\in\mathcal{I}}(1/\beta_l)$. We observe that the first derivative is linear in α_i and also linear in θ_{i1} and θ_{i2}. From (7), the second derivative is

$$\nabla^2_{ii}\phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi) = \frac{1}{\beta_i}(\hat{\beta}_i - 1)(\hat{\beta}_i + 1). \tag{8}$$

From (8), it is easy to verify that $\nabla^2_{ii}\phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi) < 0$, which implies strict concavity since $\beta_i > 0$ and $0 < \hat{\beta}_i < 1$.

IV. Inverse Problem Formulation

The inverse problem seeks to estimate rivals' cost parameters. This knowledge is required for estimating the objective of the profit maximization problem (3). The main theoretical foundation is attributed to [22], where the authors estimate the utility functions of the players in a Nash game from the observed equilibrium.

In the context of this paper, we are given (or we can obtain/estimate) M past equilibrium bids (observations) $\boldsymbol{\alpha}^j = (\alpha_i^j;\ i \in \mathcal{I}^j)$, $j = 1, \ldots, M$, where $\mathcal{I}^j$ is the set of marginal suppliers for observation j and is defined similar to $\mathcal{I}$ in (2). The $\boldsymbol{\alpha}^j$ are realized under different residual demand levels Q^j and fuel prices ξ^j, and we are interested in estimating $\boldsymbol{\theta}_i$ of supplier $i = 1, \ldots, N$. Without loss of generality, we assume that for each supplier i there are sufficient observations j at which supplier i was marginal (i.e., $i \in \mathcal{I}^j$), so that there is enough information to estimate $\boldsymbol{\theta}_i$. In case this is not true for some suppliers, we can *a priori* remove them from the set $\{1, \ldots, N\}$ (and appropriately adjust (1)). For such suppliers, we simply do not have enough information to estimate their cost parameters. It should be noted, however, that these suppliers will generally correspond to base loaded units that do not compete in the market. Given the quadratic cost function representation, the resulting linear supply curve is associated with a broadly construed notion of marginality that will render non-base-loaded units marginal during some hours. As long as each unit is marginal in some of the observations – not all units need to be marginal in all observations – the proposed framework is broadly applicable, and there is no loss of generality from the exclusion of base loaded units.

The estimates for the cost parameters are obtained by applying [22, Theorem 3], which is derived through duality, and leads to the following optimization problem:

$$\begin{aligned}\min_{\substack{\mathbf{y}, \boldsymbol{\epsilon} \\ \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_N}} \quad & \|\boldsymbol{\epsilon}\|_\infty \\ \text{s.t.} \quad & y_i^j \geq 0,\ i \in \mathcal{I}^j;\ j = 1, \ldots, M, \\ & y_i^j \geq \nabla_i \phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}^j, Q^j, \xi^j),\ \forall i \in \mathcal{I}^j, j, \\ & \sum_{i\in\mathcal{I}^j} \Big(\bar{\alpha} y_i^j - \alpha_i^j \nabla_i \phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}^j, Q^j, \xi^j)\Big) \leq \epsilon_j, \forall j, \\ & \nabla_i \phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}^{k_i}, Q^{k_i}, \xi^{k_i}) = \phi_i^{\text{norm}},\ \forall i, \end{aligned} \tag{9}$$

where $\mathbf{y} = (y_i^j)_{j=1,\ldots,M}^{i\in\mathcal{I}^j}$ is the decision variable (introduced as a dual variable in [22, Theorem 2]); $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_M)$, $\|\boldsymbol{\epsilon}\|_\infty = \max_j |\epsilon_j|$ is the infinity norm, and the last constraint is used for normalization purposes and will be discussed below. We note that the variables in $\nabla_i \phi_i$ are $\boldsymbol{\theta}_i$, and that $\nabla_i \phi_i$ is linear in $\boldsymbol{\theta}_i$, where $\boldsymbol{\alpha}^j, Q^j$, and ξ^j are parameters of the optimization problem. From (7), we have

$$\begin{aligned}\nabla_i \phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}^j, Q^j, \xi^j) = \frac{\hat{\beta}_i}{\beta_i} \frac{Q^j + \sum_{m\in\mathcal{I}^j, m\neq i} \frac{\alpha_m^j}{\beta_m}}{\sum_{l\in\mathcal{I}^j} \frac{1}{\beta_l}} + \\ \frac{\alpha_i^j}{\beta_i}(\hat{\beta}_i^2 - 1) + \theta_{i1}\frac{1 - \hat{\beta}_i}{\beta_i} + \theta_{i2}\frac{1 - \hat{\beta}_i}{\beta_i}\xi^j. \end{aligned} \tag{10}$$

Interestingly, we can reformulate the optimization problem (9) as a *Linear Programming (LP)* problem, which can be solved very efficiently. Specifically, instead of the infinite norm objective, we can introduce constraints that impose an upper bound to each $|\epsilon_j|$ and then minimize this upper bound.

The last constraint in (9) is a normalization constraint, which is equivalent to [22, Eqs. (39d) and (39e)]. The right hand side (rhs) of the constraint, ϕ_i^{norm}, is some estimate of the partial derivative at a specific point, which is evaluated at an observation k_i (potentially different for each i). In [22], for illustration purposes, the rhs estimate is obtained using the actual values of $\boldsymbol{\theta}_i$ at a median observation, considering some lower bounds for the bidding coefficients. We further elaborate on the implementation of this constraint in Section V.

We note that this inverse optimization technique still applies even when more constraints are imposed in the forward problem setup or when the bid function is changed, as long as R and P_i have closed-form expressions w.r.t. the bidding coefficients.

The quality of the computed equilibrium strategies depends on the explanatory value of the estimated cost parameters. Indeed, good estimators should explain future equilibria as well as the equilibria used to estimate them. The following result, which is a restatement of [22, Theorem 6], ensures the quality of the estimated cost functions under mild conditions. To simplify the notation, we assume that all suppliers are marginal at all past observations j; otherwise, proper adjustments to the statement of the theorem can be made.

Theorem IV.1. *Suppose that $\boldsymbol{\alpha}^j,\ j = 1,\dots,M$ are i.i.d. realizations of a random variable $\tilde{\boldsymbol{\alpha}}$, and $\tilde{\boldsymbol{\alpha}} \in \{\boldsymbol{\alpha} : 0 \le \alpha_i \le \bar{\alpha},\ \forall\ i\}$ almost surely. Then, for any $0 < \delta < 1$, with probability at least $1 - \eta$ w.r.t. the sampling,*

$$\mathbb{P}\big(\tilde{\boldsymbol{\alpha}} \text{ is a } z\text{-approximate equilibrium for the game with payoffs defined through } \hat{\boldsymbol{\theta}}_1,\dots,\hat{\boldsymbol{\theta}}_N\big) \ge 1-\delta, \tag{11}$$

where $\eta = \sum_{i=0}^{2N}\binom{M}{i}\delta^i(1-\delta)^{M-i}$; z is the optimal value of problem (9); and $\hat{\boldsymbol{\theta}}_1,\dots,\hat{\boldsymbol{\theta}}_N$ are the optimal solutions to (9).

Roughly speaking, the z-approximate equilibrium describes the situation where each supplier does not necessarily play her best action given what others are doing, playing instead a strategy that is no worse than z relative to the best response. For the definition of z-approximate equilibrium, we refer the interested reader to [22, Section 2.2].

There are two probability measures in the statement of Theorem IV.1. One is related to the new data $\tilde{\boldsymbol{\alpha}}$, while the other is related to the samples $\boldsymbol{\alpha}^1,\dots,\boldsymbol{\alpha}^M$. The probability in (11) is taken w.r.t. the new data $\tilde{\boldsymbol{\alpha}}$. For a fixed set of samples, (11) holds with probability at least $1-\eta$ w.r.t. the measure of samples. Theorem IV.1 essentially states that given typical samples, the probability that the estimated cost functions explain well a new future equilibrium is bounded below. It guarantees the accuracy of the estimated cost parameters under mild conditions.

V. Algorithmic Implementation

In this section, we present the algorithmic implementation for estimating the rivals' profit functions (or cost parameters $\boldsymbol{\theta}_i$), which can then be used to obtain equilibrium bids. We use historical data from which we can derive the past bids. Suppose we are aware of the market clearing price and the dispatch schedules of all suppliers. [3] Using this information, the past bids $\boldsymbol{\alpha}^j = (\alpha_i^j;\ i \in \mathcal{I}^j),\ j = 1,\dots,M$, can be computed via the market-clearing condition in (1), where $\boldsymbol{\beta}$ is constant and known. As before, we assume that for each supplier i there are sufficient observations j at which supplier i was marginal ($i \in \mathcal{I}^j$).

[3] Such information is publicly available in some European power markets, or it can be assumed to be discoverable at a later point in time by market participants. It is certainly available to regulators even in pool-based markets, and, to some extent, it can be estimated by entities with market knowledge; as we will discuss later, errors in the estimates can be viewed as noise in the data.

It is worth mentioning that (9) might give multiple optimal solutions. Our goal is to recover the true cost parameters from this set. Although there might be multiple cost function estimates that can explain the observed equilibria well, only true costs are expected to have good out-of-sample performance. The following Algorithm is thus proposed to identify the true cost functions, based on which equilibrium bids could be computed via an iterative best response process. We refer to this algorithm as "random search," since it searches randomly in the set of optimal solutions until the one that performs well on a validation dataset is found. The out-of-sample performance is measured by the average discrepancy, d, between the computed and true bids on the validation dataset. The variable k serves as a counter of the iterations (random searches); the algorithm terminates when the discrepancy d is smaller than the tolerance level (τ) or the maximum number of iterations is reached ($MaxIter$). In the latter case, we select the iteration with the best out-of-sample performance (smallest d).

In what follows, we present the main steps of the Algorithm.

1: **Input:** N suppliers, with constant and known bidding slopes $\beta_i,\ i = 1,\dots,N$; M past bids (observations), and for each bid $j = 1,\dots,M$, the market-clearing price R^j, residual demand Q^j, fuel price ξ^j, dispatch schedules P_i^j, upper bound for bids $\bar{\alpha}$; percentage of training samples p; tolerance level τ; maximum number of iterations $MaxIter$.
2: **Initialize:** $d = \infty$, $k = 0$.
3: **while** $d \ge \tau$ and $k < MaxIter$ **do**
4: $k \leftarrow k+1$. Randomly choose $M_t = \lfloor Mp \rfloor$ samples from all past bids (observations) to constitute the training dataset (as a percentage p of the entire set), and use the remaining bids ($M_v = M - M_t$) as the validation dataset.
5: Obtain $\hat{\boldsymbol{\theta}}_i,\ i = 1,\dots,N$, by solving problem (9) using the training dataset.
6: Compute equilibrium strategies (solving (3) via an iterative best response process) $\hat{\boldsymbol{\alpha}}_{\text{val}}^j,\ j = 1,\dots,M_v$, on the validation dataset using $\hat{\boldsymbol{\theta}}_i,\ i = 1,\dots,N$.
7: Evaluate the discrepancy between computed and true bids on the validation dataset as

$$d = \frac{\sum_{j=1}^{M_v}\|\boldsymbol{\alpha}_{\text{val}}^j - \hat{\boldsymbol{\alpha}}_{\text{val}}^j\|_1/N}{M_v}, \tag{12}$$

where $\boldsymbol{\alpha}_{\text{val}}^j$ is the j-th true bid (obtained form the historical data) on the validation dataset, and $\|\cdot\|_1$ is the ℓ_1 norm operator defined as the sum of the absolute elements of the argument.
8: **end while**
9: Compute equilibrium bids using $\hat{\boldsymbol{\theta}}_i,\ i = 1,\dots,N$ for given Q and ξ.

The iterative best response process mentioned in Step 6 for computing equilibrium bids (which also applies to Step 9) is a standard fixed point iteration process. Each supplier solves problem (3) assuming all other suppliers are fixed in their previous bids (in fact, problem (3) in our case can be solved even analytically). Then the bids are updated and the process is repeated until an equilibrium is reached, i.e., no supplier can gain by unilaterally changing her bid. In practice, this process terminates in a few iterations since the profit functions are strictly concave.

We also note that the algorithm is amenable to parallelization, as essentially, given adequate resources, all iterations (Steps 3 to 8) could be run in parallel. Our next result establishes that the algorithm requires more than T iterations with a probability that diminishes exponentially with T. Equivalently, we can select a large enough maximum number of iterations, $MaxIter$, so that the algorithm will terminate before $MaxIter$ is reached with a desirable large probability. The result further establishes that the convergence rate of the algorithm improves as we increase the size M_t of the training set. The proof of the result is included in the Appendix. We numerically explore in Section VI the out-of-sample performance of the cost estimators obtained through this algorithm.

Theorem V.1. *Assume that for some $\gamma > 0$, $(1 - \hat{\beta}_i^2)/\beta_i \geq \gamma, \forall i = 1, \ldots, N$, and the conditions of Thm. IV.1 hold. Assuming that all the past bids are at most $\bar{\epsilon}$ away from the equilibrium, and for a threshold $\tau = \sqrt{\bar{\epsilon}/(N\gamma)}$, it follows:*

1) *for any $T \geq 1$, the probability that the algorithm terminates after T iterations is no more than η^T;*
2) *for any $0 < \epsilon < 1$, when $T \geq (\log \epsilon)/(\log \eta)$, the probability that the algorithm terminates after T iterations is no more than ϵ,*

where η is defined in the statement of Thm. IV.1. Moreover, as we increase the training sample size M_t, the number of iterations that are needed for termination decreases when M_t is large enough.

Another issue mentioned in the previous section is the implementation of the normalization constraint, i.e., the last constraint in (9). For the purposes of this paper, unlike [22], we do not use the true costs (since they are indeed unknown); instead, we set the rhs (estimate of the partial derivative) to zero for the median observation of the training dataset.

Finally, we note that the algorithm can handle cases in which "noise" is present in the data, e.g., in the past bids (observations). This is perhaps the most interesting — and not trivial, application which we also explore in Section VI.

VI. Numerical Illustration

In this section, we use synthetic input data to test the validity of our approach. We first describe the experimental setup.

We consider a setup with $N = 2, 3, 4, 5$, and 10 suppliers. For each case, we assume that the true cost parameters θ_{i1} and θ_{i2}, as well as c_{i2} of supplier i are equispaced in the intervals $[7, 5]$, $[0.7, 0.9]$, and $[0.05, 0.07]$, respectively. [4]

We generate $M = 200$ past observations, in which demand Q and fuel price ξ are randomly selected within the intervals $[50, 100]$, and $[10, 30]$, respectively. For each demand and fuel price realization, i.e., for each observation j (among the 200), using the true cost estimators, we generate equilibrium bids $\boldsymbol{\alpha}^j$, assuming that all suppliers are marginal at all observations. The upper bound $\bar{\alpha}$ is set to 200. [5] The training and validation datasets are assumed to be of equal size, $M_t = M_v = 100$, using $p = 0.5$. For evaluating the out-of-sample performance we generate a test dataset with 100 additional demand and fuel price values, randomly selected within the aforementioned intervals.

The algorithm was implemented using Matlab R2017a and Gurobi 7.5.1 (for solving the LP problem (9)), without any parallelization, and the computational experiments were run on an Intel i7 5500U, at 2.4GHz, with 8 GB RAM.

In what follows, we consider two setups: a "clean" setup without noise in the data (Subsection VI-A), and a setup with noise (Subsection VI-B). We evaluate the out-of-sample performance for the noisy data case (in Subsection VI-C), and we perform sensitivity analysis with respect to key parameters (in Subsection VI-D). Lastly, we present an interesting comparison of our approach with a method introduced in [15] (Subsection VI-E).

A. "Clean" Setup

As a measure of error for the cost estimators of $\boldsymbol{\theta}_i$, we use the *Mean Absolute Percentage Error (MAPE)*, defined as $(100/2N) \sum_{i=1}^{N} \sum_{l=1}^{2} |(\theta_{il} - \hat{\theta}_{il})/\theta_{il}|$. Notably, we expect to see that cost parameters with low discrepancy values would exhibit low MAPE values. In all cases, the algorithm managed to exactly reveal the true costs within 1 or 2 iterations for $N = 2, 3, 4, 5$ and 86 iterations for $N = 10$. We plot the results for $N = 10$, in Fig. 1. Each circle corresponds to the discrepancy and MAPE values on the validation dataset for a certain partition of the samples into training/validation datasets

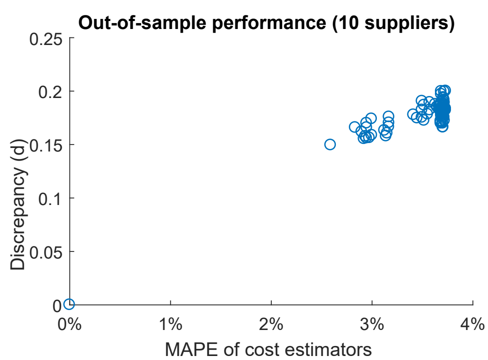

Fig. 1. Validation dataset performance ($N = 10$ suppliers, "clean" setup).

In Fig. 1 we see that in general, the lower the discrepancy (the better the performance in the validation set), the lower the MAPE of the cost estimators (the better the estimates of the true costs).

B. Setup with Noise

We introduce noise in past bid observations, by generating them within 1% of the optimal value obtained by the iterative best response process. Alternatively, one may think of this

[4]For instance, for the case $N = 3$, we have for supplier 1, $\theta_{11} = 7$, $\theta_{12} = 0.7$, $c_{12} = 0.05$, for supplier 2, $\theta_{21} = 6$, $\theta_{22} = 0.8$, $c_{22} = 0.06$, and for supplier 3, $\theta_{31} = 5$, $\theta_{32} = 0.9$, $c_{32} = 0.07$.

[5]Note that, in practice, we would estimate these bids, using the market outcomes.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TSG.2019.2891747

noise as an error in estimating past bids from the market results. For instance, noisy data may be due to errors in estimates of the market outcome when not all data required (e.g., the exact schedules) are available, or due to e slopes of the marginal cost function (reflected in pa Further, noise may account for inaccurate estimat marginal or extra-marginal suppliers that are ren the observations. In a more loose interpretation of a equilibria, one may also think of this noise as obs which suppliers do not play exactly their equilibri

For practical purposes, we set $MaxIter =$ 1 a tight tolerance level $\tau = 10^{-3}$. In all cases t limit is reached first; computational times rangec 10 min. We note that by selecting such a tight tol in absolute figures (with an average bid α of arc algorithm would terminate when reaching discrepa than 0.005%), and with 1% noise present in th almost certain that the algorithm will terminate the iteration limit. Hence, it is highly unlikely tha tolerance limits (which are achieved in the clean be also achievable in the case of noisy data. N we keep both termination conditions in the fra the sake of completeness in case of noise-free da unnecessary iterations. We also note that, by setting high iteration limit, we enhance the confidence in (see also Thm. V.1); we elaborate further on the this limit in Subsection VI-D.

For all cases ($N = 2, \ldots, 5, 10$), the best achieved discrepancy at the validation dataset ranges from 0.111 to 0.154 in absolute figures — an average bid α of around 20 implies that the discrepancy is less than 1%. The results indicate reasonably good cost estimators with MAPE ranging from 0.86% to 3.44% (for the aforementioned best achieved discrepancies). In Fig. 2, we show the performance in the validation dataset for $N = 10$ suppliers — compare with Fig. 1 for the "clean" setup. Fig. 2 also verifies the expected behavior when noise is present in the data, i.e., good performance in the validation dataset is associated with good cost estimators.

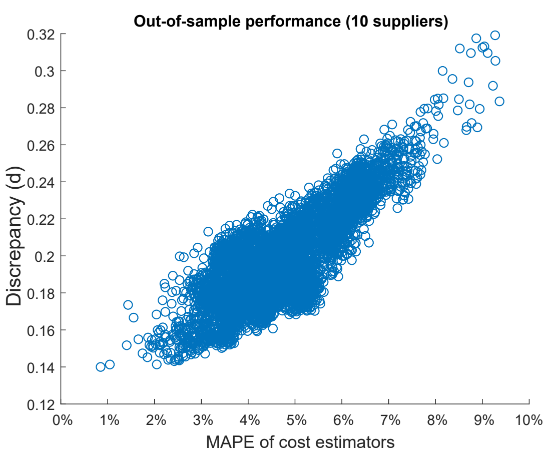

Fig. 2. Validation dataset performance ($N = 10$ suppliers, setup with noise).

In Fig. 3, we plot the true costs θ_{i1} and θ_{i2} and their estimates for $N = 5$ and $N = 10$ suppliers. We illustrate the best estimate (the one that corresponds to the lowest discrepancy calculated at the validation dataset), as well as the average cost estimators and their standard deviation (σ) over the 10,000 iterations.

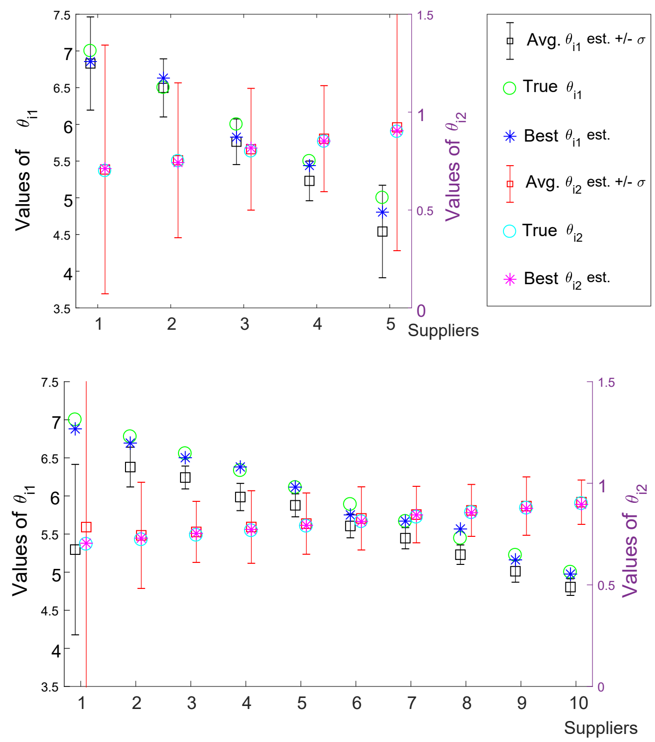

Fig. 3. True costs and estimates (best and average $\pm\sigma$) for $N = 5$ (upper figure) and $N = 10$ (lower figure) suppliers. Values of θ_{i1} are shown on the left axis; values of θ_{i2} are shown on the right axis.

The results indicate that the best estimates range from −3.9% (−2.3%) to 2% (2.2%) for θ_{i1}, and from −1% (−0.4%) to 2% (1.3%) for θ_{i2} for the case of $N = 5$ ($N = 10$) suppliers. The best estimates, which the Algorithm is designed to obtain, are reasonably close to the true cost parameters, and hence, they are expected to exhibit good out-of-sample performance. In fact, even the average estimates over the 10,000 iterations are not too far from the true cost parameters.

Next, we evaluate the out-of-sample performance of the best cost estimators using the test dataset (100 observations, different from the 200 observations that were used for the training and validation datasets).

C. Out-of-Sample Performance

For each observation (i.e., value of Q and ξ) of the test dataset, we compare the equilibrium bids using the estimates of $\hat{\boldsymbol{\theta}}_i$ with bids derived using the true costs $\boldsymbol{\theta}_i$, and we calculate the discrepancy (d) — using (12) with $M_v = 1$. We summarize the results (average discrepancy — in absolute figures, and its standard deviation) in Table I. In Fig. 4, we illustrate the values of the discrepancy for each observation of the test dataset for the case of $N = 5$ suppliers.

The results in Table I and Fig. 4 indicate a satisfactory out-of-sample performance. In fact, the average discrepancy is lower than the best achieved over the validation set. The reason is that the validation dataset contains noise, whereas the test dataset is a "clean," free of noise setup.

TABLE I
OUT-OF-SAMPLE PERFORMANCE OF BEST C

Suppliers	Avg. Discrepancy (d)
2	0.086
3	0.047
4	0.052
5	0.063
10	0.104

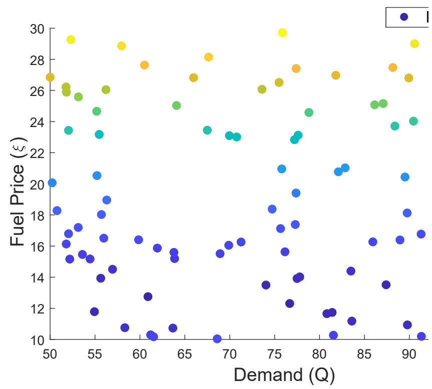

Fig. 4. Discrepancy of bids using the best cost estima
($N = 5$ suppliers).

We then take a closer look at the equilibrium bids and profits. We consider 3 instances of demand and fuel price: (a) $Q = 45$, $\xi = 8$, (b) $Q = 75$, $\xi = 20$, and (c) $Q = 110$, $\xi = 35$. Note that instance (b) refers to the mean demand and mean fuel price of the past bids, whereas instances (a) and (c) contain values that are outside the intervals used for generating the past bids, i.e., that fall outside the range of previously observed values. For each instance, we list the discrepancy values (in absolute figures) in Table II.

TABLE II
DISCREPANCY OF EQUILIBRIUM BIDS (ESTIMATES VS TRUE COSTS)

Suppliers	Instance (a) $Q = 45, \xi = 8$	Instance (b) $Q = 75, \xi = 20$	Instance (c) $Q = 110, \xi = 35$
2	0.154	0.037	0.278
3	0.048	0.041	0.114
4	0.104	0.035	0.111
5	0.062	0.059	0.140
10	0.064	0.100	0.162

We show the equilibrium bids ($\boldsymbol{\alpha}_i$) for the three instances, and for N = 2, 3, 4, 5 and 10 suppliers, in Fig. 5.

From Fig. 5 we see that the difference between the equilibrium bids using the estimates compared to the ones using the true costs is very small. For example, for the case $N = 5$ suppliers, the difference ranges from -0.135 to 0.047 for instance (a), from -0.074 to 0.133 for instance (b), and from -0.076 to 0.351 for instance (c). The differences reported as percentages range from -1.1% to 0.4%, from -0.3% to 0.6%, and from -0.2% to 1%, for instances (a), (b), and (c), respectively.

Lastly, we provide in Table III the total profits for the three instances, calculated using (6); the profits using the estimated costs are shown first and the profits using the true costs follow in parenthesis for comparison. Apart from

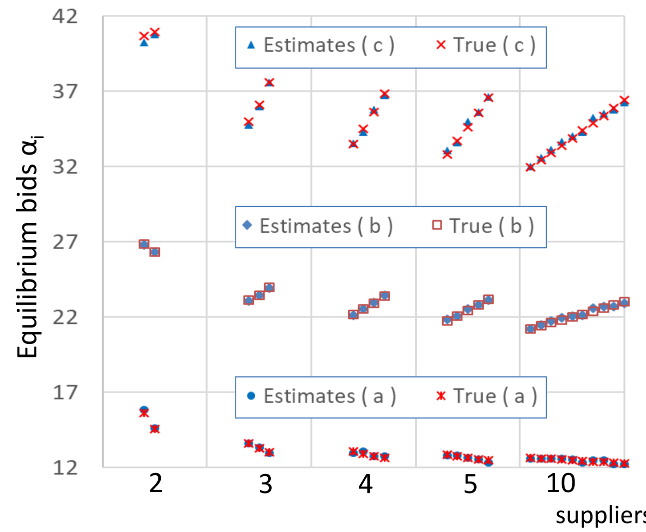

Fig. 5. Equilibrium bids using true costs and estimates for instances (a), (b) and (c) and N = 2, 3, 4, 5, and 10 suppliers (setup with noise).

the expected behavior of total profits decreasing when the number of supplier increases, and total profits increasing with increasing demand, the results also show that the differences using the estimated costs are very small (in fact the relative differences are within 5%).

TABLE III
TOTAL PROFITS USING ESTIMATES (VS TRUE COSTS)

Suppliers	Instance (a) $Q = 45, \xi = 8$	Instance (b) $Q = 75, \xi = 20$	Instance (c) $Q = 110, \xi = 35$
2	188.8 (181.7)	516.0 (518.7)	1122.5 (1151.0)
3	81.3 (80.4)	230.6 (233.7)	525.7 (537.5)
4	51.5 (50.3)	150.8 (150.0)	358.4 (361.1)
5	34.5 (36.4)	113.5 (111.7)	294.9 (283.6)
10	15.1 (15.3)	61.4 (58.3)	206.0 (195.8)

D. Sensitivity Analysis

In this subsection, we perform sensitivity analysis with respect to the level of noise, the number of available observations, and the number of iterations. As a base case, we consider the case for $N = 5$ suppliers.

1) Noise Level: As already mentioned, introducing noise in the past bids can be thought of as introducing errors in obtaining the past bids from the available or estimated market data. In the previous subsections, we assumed a noise level of 1%; in this subsection, we explore higher noise levels (2%, 3%, 4%, 5%, and 10%), and we present the results for the performance on the validation dataset (discrepancy vs. MAPE) in Fig. 6.

The best discrepancy achieved ranges from 0.116 (noise level 1%) to 1.204 (noise level 10%). MAPE values range from 1.75% to 10.46%. The results verify the expected behavior, i.e., discrepancy and MAPE increase with the noise level. They also verify that good performance in the validation dataset (low discrepancies) is associated with good estimators (low MAPE), under all noise levels. Table IV presents the out-of-sample performance under various noise levels. The results

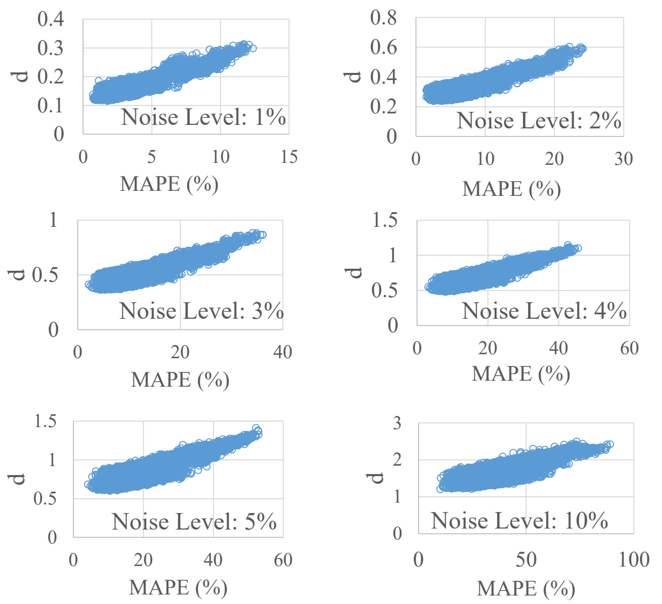

Fig. 6. Validation dataset performance, $N = 5$ suppliers, various noise levels.

are in good agreement with the ones presented in Table I. They also indicate low average discrepancies, which increase with the noise level. Note that at high noise levels (see e.g., 10%) the standard deviation becomes low, indicating that the discrepancy is mostly affected by the noise.

TABLE IV
OUT-OF-SAMPLE PERFORMANCE OF BEST COST ESTIMATORS, $N = 5$ SUPPLIERS, VARIOUS NOISE LEVELS

Noise Level	Avg. Discrepancy (d)	Std of d
1%	0.063	0.0220
2%	0.152	0.0373
3%	0.219	0.0802
4%	0.303	0.0895
5%	0.380	0.1210
10%	0.614	0.0181

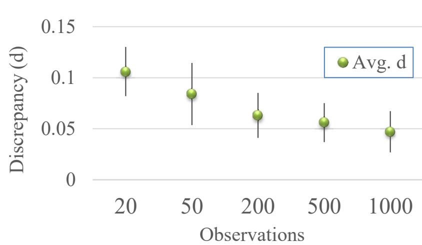

Fig. 7. Out-of-sample performance of best cost estimators (average discrepancy $\pm\sigma$), $N = 5$ suppliers, noise level 1%, various numbers of past observations.

Not surprisingly, and as predicted by Thm. IV.1, Fig. 7 shows that the out-of-sample performance improves with the number of available past observations. From Thm. IV.1 we deduce that as the number of past observations M increases, the probability that our estimates $\hat{\boldsymbol{\theta}}_1, \ldots, \hat{\boldsymbol{\theta}}_N$ yield an equilibrium that is close to the true one increases. This is due to the fact that, as shown in the Appendix (cf. Eqs. (14) and (15)), η decreases with the number of past observations. As a result, increasing the training sample size could lead to a small discrepancy between computed and true bids, and thus an improved out-of-sample performance, which is consistent with Fig. 7. Interestingly, even with 20 available observations, the performance discrepancy is reasonable. We also checked the average discrepancies (out-of-sample performance) for the various noise levels with only 20 observations and the results are also good; average discrepancies range from 0.107 to 0.852 (increasing with the level of noise) with standard deviations that range from 0.0241 to 0.0644.

3) Number of Iterations: Last but not least, we elaborate on the selection of the maximum number of iterations, i.e., the termination condition. Thm. V.1 yielded a rigorous result that provides guidance on how to select the iteration limit. Here, we numerically verify that a small number of iterations is sufficient to provide some good enough estimators.

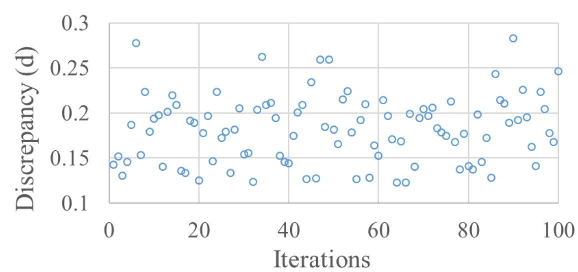

Fig. 8. Discrepancy values calculated in the validation dataset (Step 7 of the Algorithm) at each iteration (shown for the first 100 iterations), $N = 5$ suppliers, noise level 1%.

In Fig. 8, we show the values of the discrepancy calculated in the validation dataset at each iteration k, for the case of $N = 5$ suppliers with noise level 1%. For ease of exposition, we plotted the first 100 iterations of the algorithm (out of the 10,000). The results indicate that good enough estimators (with low discrepancy) can be obtained early in the process. In fact, in the first 100 iterations, we observed 7 instances with discrepancies that are less than 10% higher than the best achieved (which in our tests was 0.116). We elaborate on this indication in the following figure.

Fig. 9 shows the maximum average discrepancy ($\pm\sigma$) computed over the top x% of the iterations ranked in ascending order of discrepancy computed over the same test dataset. For instance, for the top 1%, i.e., the top 100 iterations out of 10,000, the worst achieved out-of-sample performance was an average $d = 0.100$ with $\sigma = 0.0167$.

Lastly, we investigate the relationship between discrepancy values and the total profits. In Fig. 10, we show the values of the discrepancy and the total profits calculated in the validation dataset at each iteration, for the case of $N = 5$ suppliers, $M = 200$ past observations with noise level 1%, and $MaxIter = 10,000$ iterations. It can be seen that the larger the discrepancy is, the smaller the total profits, which validates the use of the

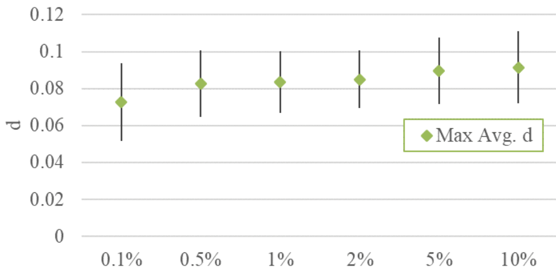

Fig. 9. Out-of-sample performance of cost estimators (maximum average discrepancy $\pm\sigma$), $N = 5$ suppliers, noise level 1%, top x% (x-axis) of iterations (ranked in ascending order of discrepancy).

discrepancy as a performance metric in identifying the best cost estimators.

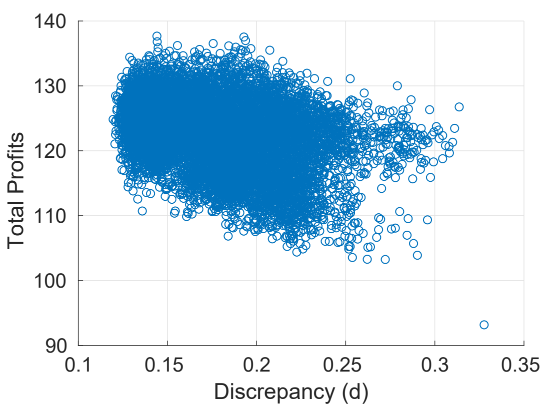

Fig. 10. Total profits vs. discrepancy values calculated in the validation dataset, $N = 5$ suppliers, noise level 1%, 200 observations, 10,000 iterations.

E. Comparison with the Method in [15]

In this subsection, we compare the performance of our approach with the method introduced in [15], which assumes that the bidding coefficients of any rival supplier follow a bivariate normal distribution whose mean and covariance could be inferred from the past observed bids. The profit maximization problem for a single supplier i is then reformulated to involve only α_i and β_i, with the market clearing price R and the production level P_i evaluated at the mean bids of the rival suppliers.

We note that [15] does not estimate rivals' cost functions; it calculates the bids using as input the mean and covariance of rivals' bidding coefficients. In our setup, since we assume that the bidding slope β is known, only the mean and variance of α need to be inferred from the historical data. For the needs of our comparisons, we use the noisy setup with 1% noise level.

We evaluate the performance of [15] on the test dataset (that contains 100 equilibrium bids). The average discrepancy between the solution obtained using [15] and the equilibrium bids on the test dataset is shown in Table V (the differences with the values presented in Table I that are obtained using our approach are shown in parentheses). The results show that [15] leads to significantly higher discrepancies (as well as higher standard deviations), indicating that our approach obtains a more accurate prediction of the bids.

TABLE V
AVERAGE DISCREPANCY OF BIDS ON THE TEST DATASET USING [15] (DIFFERENCES WITH TABLE I)

Suppliers	Avg. Discrepancy (d)	Std. of d
2	1.406 (+1.32)	0.7891 (+0.73)
3	1.026 (+0.98)	0.5709 (+0.55)
4	0.818 (+0.77)	0.4551 (+0.44)
5	0.681 (+0.62)	0.3793 (+0.36)
10	0.372 (+0.27)	0.2071 (+0.19)

We then take a closer look at the bids and profits by considering the 3 instances illustrated in Section VI-C. In Table VI we list the discrepancy of bids obtained using [15] (in parentheses we show the differences with the values listed in Table II). The results show that for instance (b), which represents instances that have been constantly observed in the past, [15] achieves a comparable performance with our approach, and occasionally even lower discrepancies. However, the performance of [15] is much worse for instances (a) and (c), which contain values that were outside the range of past observations. The reason is that [15] implicitly assumes that rivals' bidding behavior is similar to what has been observed in the past, which results in a large bias when a new, unseen scenario occurs. By contrast, through estimating rivals' cost functions from the past bids, our approach acquires more information, and exhibits a stronger out-of-sample inference capability that guarantees a low estimation bias for every possible scenario. Interestingly but not surprisingly, [15] achieves better results when the number of suppliers is large, in which case the bids are very close (due to competition), see e.g., Fig. 5.

TABLE VI
DISCREPANCY OF BIDS USING [15] (DIFFERENCES WITH TABLE II)

Suppliers	Instance (a) $Q = 45, \xi = 8$	Instance (b) $Q = 75, \xi = 20$	Instance (c) $Q = 110, \xi = 35$
2	3.773 (+3.62)	0.048 (+ 0.01)	4.776 (+4.50)
3	2.487 (+2.44)	0.050 (+0.01)	3.211 (+3.10)
4	1.923 (+1.82)	0.043 (+0.01)	2.497 (+2.39)
5	1.581 (+1.52)	0.037 (-0.02)	2.057 (+1.92)
10	0.849 (+0.79)	0.021 (-0.08)	1.108 (+0.95)

In Table VII we list the total profits for the bids derived using [15], and we show in parentheses the differences with the values in Table III obtained using the cost estimates. Recall that the profits obtained using the cost estimates are very close to the profits obtained using the true costs (within +/- 5%). For instance (b), the two methods obtain similar profits. For instance (c) which represents higher demands and fuel prices, our approach achieves much higher profits. The opposite is true for instance (a), which represents lower demands and fuel prices. This is an interesting result, which is explained by the fact that the bids generated using [15] are biased by the mean past bids. Hence, they tend to be values of past observations, i.e., they are inflated for the instances of lower demands and fuel prices, and they are reduced for higher demands and fuel prices. More specifically, the values of the equilibrium bids are about 25% higher in instance (a), and about 12% lower in instance (c) compared to the ones shown in Fig. 5. Of course, in instance (a) where the bids (and hence, profits) are inflated, the conditions are ripe for a new supplier to come in, underbid, and capture significant market share.

TABLE VII
TOTAL PROFITS OF ESTIMATED BIDS USING [15] (DIFFERENCES WITH TABLE III USING ESTIMATES)

Suppliers	Instance (a) $Q = 45, \xi = 8$	Instance (b) $Q = 75, \xi = 20$	Instance (c) $Q = 110, \xi = 35$
2	350.2 (+161.40)	515.3 (-0.70)	636.0 (-486.50)
3	192.4 (+111.10)	230.0 (-0.60)	186.2 (-339.50)
4	137.1 (+85.60)	146.7 (-4.10)	87.0 (-271.40)
5	107.7 (+73.20)	108.9 (-4.60)	57.4 (-237.50)
10	53.7 (+38.60)	56.8 (-4.60)	73.7 (-132.30)

In conclusion, our approach possesses a stronger out-of-sample inference capability attributed to the estimation of the cost functions. The method proposed in [15] ignores the interaction among suppliers, assumes a normal distribution and uses only the mean values of the past bids to infer rivals' behavior, which accounts for its unsatisfactory performance in new, unseen scenarios.

VII. EXTENSION TO LOCATION-DEPENDENT PRICES

So far, we assumed competition among suppliers in uncongested networks. Indeed, several day-ahead electricity markets clear ignoring network congestion. In the instance of such day-ahead market rules, the system operator adjusts the generation dispatch to observe line flow capacity constraints and ensure secure and reliable operation.

In U.S. markets, however, the transmission system representation has been part of the standard market design for many years, with resulting "Locational Marginal Prices" (LMPs) representing the marginal cost at each node of the transmission system. In practice, without entering in a detailed analysis of how LMPs are formed, we note that "price islands" may characterize clearing prices, differing only slightly to reflect varying loss factors.

Our method appears to assume away the fact that network connected markets result in location dependent clearing price differentials driven by (i) small effects of location-specific line loss contributions, but also, (ii) significant contributions during network congestion events. It can capture and address, however, significant congestion-caused differentials by detecting market-splitting occurrences that result in "price islands" with essentially homogeneous prices within each island. Although this limits the number of relevant observations when price islanding occurs, it utilizes the unusually high or low price events associated with congestion.

Our approach applies to price islands where congestion is a result of generators outside of the relevant island, which do not set the price and, hence, are not part of our analysis. In this context, the market clearing price in each island s is $R_s = \alpha_i + \beta_i P_i$, $i \in \mathcal{I}_s$, where $\mathcal{I}_s$ is the set of marginal suppliers in island s. The residual demand in each island is $Q_s = \sum_{i \in \mathcal{I}_s} P_i$. Similarly to (2), the solution for island s, in terms of price and quantity, is

$$R_s(\boldsymbol{\alpha}_s, \boldsymbol{\beta}_s; Q_s) = \frac{Q_s + \sum_{i \in \mathcal{I}_s} \frac{\alpha_i}{\beta_i}}{\sum_{i \in I_s} \frac{1}{\beta_i}},$$

$$P_i(\boldsymbol{\alpha}_s, \boldsymbol{\beta}_s; Q_s) = \frac{R_s(\boldsymbol{\alpha}_s, \boldsymbol{\beta}_s; Q_s) - \alpha_i}{\beta_i}, \ i \in \mathcal{I}_s,$$

where $\boldsymbol{\alpha}_s, \boldsymbol{\beta}_s$ refer to suppliers in island s.

In the forward problem, the profit of the suppliers in each island is therefore straightforwardly defined. Considering the inverse problem, we note that the past bids may contain both congested and uncongested instances. The set of marginal suppliers may be different in each instance, and furthermore the islands in the congested instances may be different. But in either case, our approach can handle these different sets, since for each observation j, we can have different price islands $s \in \mathcal{S}_j$, where $\mathcal{S}_j$ is the set of price islands for the j-th observation, and for each price island of that observation we have a set of marginal suppliers denoted by $\mathcal{I}_s^j$. Hence, the inverse optimization problem in (9) is applied for each price island s for observations j within this island, and for the respective set of marginal suppliers for the specific island. Essentially, the uncongested case represents one single island (and can still be described by the above notation).

VIII. CONCLUSIONS AND FURTHER RESEARCH

In this paper, we proposed an inverse optimization method to estimate electricity suppliers' cost functions in the day-ahead electricity market based on historical bidding data. The problem of computing optimal bidding strategies can be seen as an equilibrium computation problem given the estimated payoff functions. We applied a "random search" algorithm to estimate the cost function parameters of electricity generators; specifically, the parameters that are proportional to their generation output. The algorithm essentially seeks cost function parameters (among multiple possible values compatible with the past data) which have good out-of-sample generalization performance. We established strong, exponential-type probabilistic convergence guarantees for this algorithm. Extensive numerical experimentation verifies that one can recover accurate estimates of the cost function parameters, which, in turn, allows generators to bid with knowledge of how competitors would respond. Even though we considered a simple setting involving no congestion or transmission network effects, we discussed an extension of the methodology to location-dependent prices.

Regarding future research directions, it would be of interest to develop non-parametric approaches that do not require to assume a specific parametric form for the cost functions. Finally, in addition to estimating competing generators' cost functions, our methodology is particularly useful in estimating the underlying cost functions of market participants who bid synthetic or virtual generators corresponding to contracts with either physical generation owners or a portfolio of demand-response-capable consumers.

REFERENCES

[1] F. C. Schweppe, M. C. Caramanis, R. D. Tabors, and R. E. Bohn, *Spot pricing of electricity*. Kluwer Academic Publishers, 1988.

[2] A. K. David and F. Wen, "Strategic bidding in competitive electricity markets: a literature survey," in *Power Engineering Society Summer Meeting*, vol. 4. IEEE, 2000, pp. 2168–2173.

[3] G. Li, J. Shi, and X. Qu, "Modeling methods for GenCo bidding strategy optimization in the liberalized electricity spot market–a state-of-the-art review," *Energy*, vol. 36, no. 8, pp. 4686–4700, 2011.

[4] "Euphemia public description. PCR market coupling algorithm," [Online]. https://www.nordpoolspot.com, 2016.

[5] P. D. Klemperer and M. A. Meyer, "Supply function equilibria in oligopoly under uncertainty," *Econometrica: Journal of the Econometric Society*, pp. 1243–1277, 1989.

[6] R. J. Green and D. M. Newbery, "Competition in the British electricity spot market," *Journal of Political Economy*, vol. 100, no. 5, pp. 929–953, 1992.

[7] R. Baldick, R. Grant, and E. Kahn, "Theory and application of linear supply function equilibrium in electricity markets," *Journal of regulatory economics*, vol. 25, no. 2, pp. 143–167, 2004.

[8] A. Rudkevich, "On the supply function equilibrium and its applications in electricity markets," *Decision Support Systems*, vol. 40, no. 3-4, pp. 409–425, 2005.

[9] P. Holmberg and D. Newbery, "The supply function equilibrium and its policy implications for wholesale electricity auctions," *Utilities Policy*, vol. 18, no. 4, pp. 209–226, 2010.

[10] E. J. Anderson and A. B. Philpott, "Using supply functions for offering generation into an electricity market," *Operations Research*, vol. 50, no. 3, pp. 477–489, 2002.

[11] E. J. Anderson and X. Hu, "Finding supply function equilibria with asymmetric firms," *Operations Research*, vol. 56, no. 3, pp. 697–711, 2008.

[12] P. Holmberg, "Numerical calculation of an asymmetric supply function equilibrium with capacity constraints," *European Journal of Operational Research*, vol. 199, no. 1, pp. 285–295, 2009.

[13] R. Sioshansi and S. Oren, "How good are supply function equilibrium models: an empirical analysis of the ERCOT balancing market," *Journal of Regulatory Economics*, vol. 31, no. 1, pp. 1–35, 2007.

[14] A. K. David, "Competitive bidding in electricity supply," in *IEE Proceedings C (Generation, Transmission and Distribution)*, vol. 140, no. 5. IET, 1993, pp. 421–426.

[15] F. Wen and A. David, "Strategic bidding for electricity supply in a day-ahead energy market," *Electric Power Systems Research*, vol. 59, no. 3, pp. 197–206, 2001.

[16] T. Li and M. Shahidehpour, "Strategic bidding of transmission-constrained GENCOs with incomplete information," *IEEE Transactions on power Systems*, vol. 20, no. 1, pp. 437–447, 2005.

[17] A. D. Yucekaya, J. Valenzuela, and G. Dozier, "Strategic bidding in electricity markets using particle swarm optimization," *Electric Power Systems Research*, vol. 79, no. 2, pp. 335–345, 2009.

[18] A. Cherukuri and J. Cortés, "Decentralized Nash equilibrium learning by strategic generators for economic dispatch," in *American Control Conference (ACC)*, 2016, pp. 1082–1087.

[19] L. Mitridati and P. Pinson, "A bayesian inference approach to unveil supply curves in electricity markets," *IEEE Transactions on Power Systems*, 2017.

[20] M. Motalleb and R. Ghorbani, "Non-cooperative game-theoretic model of demand response aggregator competition for selling stored energy in storage devices," *Applied Energy*, vol. 202, pp. 581–596, 2017.

[21] P. Zou, Q. Chen, Q. Xia, G. He, C. Kang, and A. J. Conejo, "Pool equilibria including strategic storage," *Applied energy*, vol. 177, pp. 260–270, 2016.

[22] D. Bertsimas, V. Gupta, and I. C. Paschalidis, "Data-driven estimation in equilibrium using inverse optimization," *Mathematical Programming, Series A*, vol. 153, no. 2, pp. 595–633, 2015.

[23] R. Ahuja and J. Orlin, "Inverse Optimization," *Operations Research*, vol. 49, no. 5, pp. 771–783, 2001.

[24] D. Bertsimas, V. Gupta, and I. C. Paschalidis, "Inverse optimization: A new perspective on the Black-Litterman model," *Operations Research*, vol. 60, no. 6, pp. 1389–1403, December 2012.

[25] Q. Zhao, A. Stettner, E. Reznik, I. C. Paschalidis, and D. Segrè, "Mapping the landscape of metabolic goals of a cell," *Genome Biology*, vol. 17, no. 1, p. 109, 2016.

[26] C. Ruiz, A. J. Conejo, and D. J. Bertsimas, "Revealing rival marginal offer prices via inverse optimization," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 3056–3064, 2013.

[27] J. Saez-Gallego, J. M. Morales, M. Zugno, and H. Madsen, "A data-driven bidding model for a cluster of price-responsive consumers of electricity," *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 5001–5011, 2016.

[28] J. R. Birge, A. Hortaçsu, and J. M. Pavlin, "Inverse optimization for the recovery of market structure from market outcomes: An application to the MISO electricity market," *Operations Research*, vol. 65, no. 4, pp. 837–855, 2017.

[29] T. Lu, Z. Wang, J. Wang, Q. Ai, and C. Wang, "A data-driven Stackelberg market strategy for demand response-enabled distribution systems," *IEEE Transactions on Smart Grid*, 2018.

[30] R. Chen, M. Caramanis, and I. C. Paschalidis, "Strategic equilibrium bidding for electricity suppliers in a day-ahead market using inverse optimization," in *Proceedings of the 56th IEEE Conference on Decision and Control*, Melbourne, Australia, December 12–16 2017, pp. 220–225.

[31] X. Hu and D. Ralph, "Using EPECs to model bilevel games in restructured electricity markets with locational prices," *Operations Research*, vol. 55, no. 5, pp. 809–827, 2007.

[32] B. F. Hobbs, C. B. Metzler, and J.-S. Pang, "Strategic gaming analysis for electric power systems: An MPEC approach," *IEEE Transactions on Power Systems*, vol. 15, no. 2, pp. 638–645, 2000.

Appendix
Proof of Thm. V.1

Proof: Assume that the optimal solutions to the inverse problem at iteration k are $\hat{\boldsymbol{\theta}}_1^k, \ldots, \hat{\boldsymbol{\theta}}_N^k$, and the optimal value is z_k. We will first show that the function $f^k(\boldsymbol{\alpha}) \triangleq \big(-\nabla_1\phi_1(\hat{\boldsymbol{\theta}}_1^k; \boldsymbol{\alpha}, Q, \xi), \ldots, -\nabla_N\phi_N(\hat{\boldsymbol{\theta}}_N^k; \boldsymbol{\alpha}, Q, \xi)\big)$ is strongly monotone. For simplicity we suppress the dependence of f on $\boldsymbol{\theta}_i, Q$ and ξ. By definition, a function $f^k(\boldsymbol{\alpha})$ is *strongly monotone* if $\exists\gamma > 0$ such that

$$\big(f^k(\boldsymbol{\alpha}_1) - f^k(\boldsymbol{\alpha}_2)\big)'(\boldsymbol{\alpha}_1 - \boldsymbol{\alpha}_2) \geq \gamma\|\boldsymbol{\alpha}_1 - \boldsymbol{\alpha}_2\|_2^2, \forall \boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2.$$

Plugging in the formula for $\nabla_i\phi_i(\boldsymbol{\theta}_i; \boldsymbol{\alpha}, Q, \xi)$, we have:

$$f^k(\boldsymbol{\alpha}_1) - f^k(\boldsymbol{\alpha}_2) = \Big(\frac{1-\hat{\beta}_1^2}{\beta_1}(\alpha_{1,1} - \alpha_{2,1}), \ldots, \frac{1-\hat{\beta}_N^2}{\beta_N}(\alpha_{1,N} - \alpha_{2,N})\Big),$$

where $\alpha_{1,i}, \alpha_{2,i}$ are the i-th elements of $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$, respectively. Using $(1-\hat{\beta}_i^2)/\beta_i \geq \gamma$, $\forall i$, it follows

$$\big(f^k(\boldsymbol{\alpha}_1) - f^k(\boldsymbol{\alpha}_2)\big)'(\boldsymbol{\alpha}_1 - \boldsymbol{\alpha}_2) \geq \gamma\|\boldsymbol{\alpha}_1 - \boldsymbol{\alpha}_2\|_2^2.$$

With the strongly monotone function $f^k(\boldsymbol{\alpha})$, we can use [22, Thm. 8], which shows that for any $0 < \delta < 1$, with probability at least $1 - \eta$ with respect to the sampling,

$$\|\boldsymbol{\alpha}_{\text{val}}^j - \hat{\boldsymbol{\alpha}}_{\text{val}}^j\|_2 \leq \sqrt{z_k/\gamma}, \quad \forall j = 1, \ldots, M_v,$$

where z_k is the optimal value of the inverse optimization problem (9) at iteration k. Using the norm inequality

$$\|\boldsymbol{\alpha}_{\text{val}}^j - \hat{\boldsymbol{\alpha}}_{\text{val}}^j\|_1 \leq \sqrt{N}\|\boldsymbol{\alpha}_{\text{val}}^j - \hat{\boldsymbol{\alpha}}_{\text{val}}^j\|_2,$$

we obtain that at iteration k, the discrepancy satisfies:

$$d \leq \sqrt{z_k/(N\gamma)} \leq \sqrt{\bar{\epsilon}/(N\gamma)},$$

which yields that

$$\mathbb{P}(d \leq \sqrt{\bar{\epsilon}/(N\gamma)}) \geq 1 - \eta. \tag{13}$$

Since the iterations are independent from each other, setting $p \triangleq \mathbb{P}(d \leq \sqrt{\bar{\epsilon}/(N\gamma)})$, and using (13), we have:

$$\mathbb{P}\Big(\text{Algorithm terminates after } T \text{ iterations}\Big) = (1-p)^T \leq \eta^T.$$

Therefore, for any small $0 < \epsilon < 1$, if the probability that the algorithm terminates after T iterations is below ϵ, we need $T \geq (\log\epsilon)/(\log\eta)$.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TSG.2019.2891747

We next show that as the training sample size M_t increases, the number of iterations that are needed decreases for a large enough M_t. First note that

$$\begin{aligned}\eta &= \sum_{i=0}^{2N} \binom{M_t}{i} \delta^i (1-\delta)^{M_t-i} \\ &= \sum_{i=0}^{2N} \frac{(M_t-i+1)\ldots M_t}{i!} \delta^i (1-\delta)^{M_t-i} \\ &\le \sum_{i=0}^{2N} \frac{(M_t)^i}{i!} \delta^i (1-\delta)^{M_t-i}.\end{aligned} \tag{14}$$

Define $h_i(M_t) \triangleq \frac{(M_t)^i}{i!} \delta^i (1-\delta)^{M_t-i}$, and take its derivative:

$$\nabla h_i(M_t) = \frac{(M_t)^{i-1}}{(i-1)!} \delta^i (1-\delta)^{M_t-i} \Big(1 + \frac{M_t}{i} \log(1-\delta)\Big). \tag{15}$$

We see that for a large enough M_t, $\nabla h_i(M_t) < 0, \forall i$, since $\log(1-\delta) < 0$. Therefore, as M_t increases, η decreases, and the number of iterations, i.e., $(\log \epsilon)/(\log \eta)$, decreases as well. ■

Ruidi Chen received the B.S. degree in statistics from Fudan University, Shanghai, China, in 2011. From 2014 to 2015, she worked as a visiting scholar at George Mason University. She is currently a Ph.D. candidate in systems engineering at Boston University, Boston, MA, USA. Her research interests lie in the areas of statistical machine learning and optimization with applications to health care. Ms. Chen was a recipient of several awards, including the 2015 College of Engineering Dean's Fellowship Award at Boston University, the 2016-2018 Boston University–Brigham and Women's Hospital (BU-BWH) Imaging Fellowship Award, a 2017 IFIP WG 7.3 Performance Conference Travel Award, and a 2017 IEEE Conference on Decision and Control Travel Award.

Ioannis Ch. Paschalidis (M'96–SM'06–F'14) received the Diploma in ECE from the National Technical University of Athens, Athens, Greece, in 1991, and the M.S. and Ph.D. degrees, both in EECS, from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1993 and 1996, respectively.

In September 1996 he joined Boston University where he has been ever since. He is a Professor and Data Science Fellow at Boston University with appointments in the Department of Electrical and Computer Engineering, the Division of Systems Engineering, and the Department of Biomedical Engineering. He is the Director of the Center for Information and Systems Engineering (CISE). He has held visiting appointments with MIT and Columbia University, New York, NY, USA. His current research interests lie in the fields of systems and control, networks, applied probability, optimization, operations research, computational biology, and medical informatics. He has more than 180 refereed publications in these and related areas.

Dr. Paschalidis is a recipient of the NSF CAREER award (2000), several best paper and best algorithmic performance awards, and a 2014 IBM/IEEE Smarter Planet Challenge Award. He was an invited participant at the 2002 Frontiers of Engineering Symposium, organized by the U.S. National Academy of Engineering and the 2014 U.S. National Academies Keck Futures Initiative (NAFKI) Conference. He is the founding Editor-in-Chief of the IEEE Transactions on Control of Network Systems.

Michael Caramanis received the B.S. in Chemical Eng. from Stanford University, Palo Alto, CA, USA, in 1971 and the M.S. and Ph.D. degrees in Engineering from Harvard University, Cambridge, MA, USA, in 1972 and 1976, respectively. Since 1982 he is at Boston University where he is Professor of Systems and Mechanical Eng. He chaired the Greek Regulatory Authority for Energy and the International Energy Charter's Investment Group (2014-2008), was personally involved in power market implementations in England (1989-90) and Italy (2000-03), and his written work has influenced Power Market design in the U.S. and Europe. His current application domain focus is Marginal Costing and Dynamic Pricing on smart Power grids, grid topology control for congestion mitigation, and the extension of power markets to include distribution connected loads, generation, and resources. He is coauthor of "Spot Pricing of Electricity," Kluwer, 1987, and has more than 100 refereed publications. His disciplinary background is in Mathematical Economics, Optimization, and Stochastic Dynamic Decision Making.

Panagiotis Andrianesis is a graduate of the Hellenic Army Academy, Greece, also holding a B.Sc. degree in economics (2004) from the National and Kapodistrian University of Athens, Athens, Greece, and a Diploma degree in electrical and computer engineering (2010) from the National Technical University of Athens, Athens, Greece. He received his M.Sc. degree in production management (2011) and his Ph.D. degree (2016) in the area of design and analysis of electricity market mechanisms from the University of Thessaly, Greece.

He is currently a Postdoctoral Associate in the Division of Systems Engineering, at Boston University, also affiliated with the Information and Data Science Research Group, Electrical and Computer Engineering. He has been a Research Associate of the Production Management Laboratory in the Department of Mechanical Engineering, at the University of Thessaly, and a Consultant and Research Associate of ECCO International Inc. His research interests include power system economics, electricity markets, operations research, optimization, and applied mathematics.

Dr. Andrianesis is a Member of INFORMS, the Mathematical Optimization Society (MOS), and the Production and Operations Management Society (POMS). He is the recipient of the 2010 IEEE Antennas and Propagation Society Pre-Doctoral Research Award.