Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 50-1 (2017) 6761–6767

Learning to Buy (and Sell) Demand Response ⋆

Kia Khezeli * **Weixuan Lin** * **Eilyan Bitar** *

* *School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA. (e-mails:[kk839,wl476,eyb5]@cornell.edu).*

Abstract: We adopt the perspective of an aggregator, which seeks to coordinate its *purchase* of demand reductions from a fixed group of residential electricity customers, with its *sale* of the aggregate demand reduction in a two-settlement wholesale energy market. The aggregator procures reductions in demand by offering its customers a uniform price for reductions in consumption relative to their predetermined baselines. Prior to its realization of the aggregate demand reduction, the aggregator must also determine how much energy to sell into the two-settlement energy market. In the day-ahead market, the aggregator commits to a forward contract, which calls for the delivery of energy in the real-time market. The underlying aggregate demand curve, which relates the aggregate demand reduction to the aggregator's offered price, is assumed to be affine and subject to unobservable, random shocks. Assuming that both the parameters of the demand curve and the distribution of the random shocks are initially unknown to the aggregator, we investigate the extent to which the aggregator might dynamically adapt its DR prices and forward contracts to maximize its expected profit over a window of T days. Specifically, we design a data-driven pricing and contract offering policy that resolves the aggregator's need to learn the unknown demand model with its desire to maximize its cumulative expected profit over time. The proposed pricing policy is proven to exhibit a *regret* over T days that is at most $O(\sqrt{T})$.

Keywords: Demand response, dynamic pricing, online learning, electricity markets.

1. INTRODUCTION

The large scale utilization of demand response (DR) resources has the potential to substantially improve the reliability and efficiency of electric power systems. Accordingly, several state and federal mandates have been established to facilitate the integration of demand response resources into wholesale electricity markets. For example, FERC Order 719 mandates that Independent System Operators (ISOs) permit the direct sale of DR services into wholesale electricity markets (FERC, 2008). As individual residential customers often posses insufficient capacity [1] to participate in such markets directly, there emerges the need for an intermediary, or *aggregator*, with the ability to coordinate the demand response of large numbers of residential customers for direct sale into the wholesale electricity market.

In this paper, we adopt the perspective of an aggregator, which seeks to coordinate its *purchase* of an aggregate demand reduction from a fixed group of residential electricity customers, with its *sale* of the aggregate demand reduction into a two-settlement wholesale energy market. [2] Formally, this amounts to a two-sided optimization problem, which requires the aggregator to balance the cost it incurs in procuring the demand reduction from customers against the revenue it derives from its sale into the wholesale energy market. We develop the problem more formally in what follows.

We consider the setting in which the aggregator purchases demand reductions from customers using a non-discriminatory, price-based mechanism. That is to say, each participating customer is payed for her reduction in electricity demand according to a uniform per-unit energy price determined by the aggregator. Pricing mechanisms of this form fall within the more general category of DR programs that rely on peak time rebates (PTR) as incentives for demand reduction. Prior to its realization of the aggregate demand reduction, the aggregator must also determine how much energy to sell into the two-settlement energy market. In the day-ahead (DA) market, the aggregator commits to a forward contract, which calls for delivery of energy in the real-time (RT) market. If the realized reduction in demand exceeds (falls short of) the forward contract, then the difference is sold (bought) in the RT market. In order to maximize its profit, the aggregator must, therefore, co-optimize the DR price it offers its customers with the forward contract that it commits to in the wholesale energy market.

⋆ This work was supported in part by NSF grant ECCS-1351621, NSF grant CNS-1239178, NSF grant IIP- 1632124, US DoE under the CERTS initiative, and the Simons Institute for the Theory of Computing.

[1] For example, the Proxy Demand Resource (PDR) program currently being operated by the California ISO has a minimum curtailment capacity requirement of 100 kW (Wolak et al., 2009).

[2] We note that a measurable reduction in demand is equivalent to an increase in supply.

2405-8963 © 2017, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2017.08.1193

There are a myriad of challenges that the aggregator faces in the deployment of such programs. The most basic challenge is the prediction of how customers will adjust their aggregate demand in response to different DR prices, i.e., the aggregate demand curve. If the offered price is too low, consumers may be unwilling to curtail their demand; if the offered price is too high, the aggregator pays too much and gets more reduction than is needed. As the aggregator is initially ignorant to customers' aggregate demand curve, the aggregator must attempt to learn a model of customer behavior over time through repeated observations of demand reductions in response to its offered DR prices. Simultaneously, the aggregator must jointly adjust its DR prices and forward contract offerings in such a manner as to facilitate profit maximization over time. As we will later show, such tasks are intimately related, and give rise to a trade-off between the need to *learn* (explore) and *earn* (exploit).

Contribution and Related Work: We study the setting in which the aggregator is faced with an aggregate demand curve that is affine in price, and subject to unobservable, additive random shocks. We assume that both the parameters of the demand curve and the probability distribution of the random shocks are fixed, and *initially unknown* to the aggregator. Faced with such ignorance, we explore the extent to which the aggregator might dynamically adapt its posted DR prices and offered contracts to maximize its expected profit over a time frame of T days. Specifically, we design a causal pricing and contract offering policy that resolves the aggregator's need to learn the unknown demand model with its desire to maximize its cumulative expected profit over time. The proposed pricing policy is proven to exhibit *regret* (relative to an oracle) over T days that is at most $O(\sqrt{T})$. In addition, the proposed policy generates a sequence of posted DR prices and forward contracts that converge to the oracle optimal DR price and forward contract in the mean square sense.

The literature – as it relates to the problem of co-optimizing an aggregator's decisions in both the retail and wholesale electricity markets – is sparse. Campaigne and Oren (2015) consider a market model that is perhaps closest in nature to the one considered in this paper. They adopt a mechanism design approach to eliciting demand response, where customers are rationed and remunerated according to their reported types. A related line of literature includes (Chao, 2012) and (Crampes and Léautier, 2015). In this paper, we take a posted price approach to the procurement of demand response. This is in sharp contrast to the mechanism design approach, as it gives rise to the need to learn customers' types (i.e., demand function) over time.

Organization: The remainder of the paper is organized as follows. In Section 2, we formulate the aggregator's profit maximization problem. In Section 3, we propose a recursive estimation scheme to learn the unknown demand model. In Section 4, we propose a joint pricing and contract offering policy for the aggregator, and provide a theoretical analysis of the regret incurred by the policy. In Section 5, we illustrate the performance of our proposed policy with a numerical example. All mathematical proofs are omitted in this version of the paper due to space constraints.

2. MODEL

We adopt the perspective of an aggregator who seeks to purchase demand reductions from a fixed group of N customers for sale into a two-settlement energy market. The market is assumed to repeat over multiple time periods (e.g., days) indexed by $t = 1, 2, \ldots$. The actions taken by the aggregator and their timing are specified in the following subsections.

2.1 Two-Settlement Market Model

At the beginning of each time period t, the aggregator commits to a forward contract for energy in the day-ahead (DA) market in the amount of Q_t (kWh). The forward contract is remunerated at the *DA energy price*. The forward contract calls for delivery in the real-time (RT) market. If the energy delivered by the aggregator (i.e., demand reduction) falls short of the forward contract, the aggregator must purchase the shortfall in the RT market at the *shortage price*. If the energy delivered exceeds the forward contract, the aggregator must sell the excess supply in the RT market at the *overage price*. We denote the wholesale energy prices (\$/kWh) by

- π, DA energy price,
- π_-, RT shortage price,
- π_+, RT overage price.

Although we assume throughout the paper that the wholesale energy prices are fixed and known across time, all results stated in this paper can be generalized to accommodate the more general setting in which the wholesale energy prices exhibit known variation with time. We also assume that the wholesale energy prices satisfy $\pi > 0$ and $\pi_+ < \pi < \pi_-$. Such assumption serves to facilitate clarity of exposition and analysis in the sequel, as it preserves concavity of the aggregator's expected profit function (2).

2.2 Demand Response Model

In order to meet its forward contract commitment Q_t, the aggregator must elicit an aggregate reduction in demand from its customers. It does so by broadcasting a uniform DR price $p_t \geq 0$, to which each customer i responds with a reduction in demand in the amount of D_{it} (kWh), thereby entitling each customer i to receive a payment of $p_t D_{it}$. Implicit in this model is the assumption that each customer's reduction in demand is measured against a predetermined baseline.

We model the response of each customer i to the posted price p_t at time t according to the *affine* function

$$D_{it} = a_i p_t + b_i + \varepsilon_{it}, \quad \text{for} \quad i = 1, \ldots, N,$$

where $a_i \in \mathbb{R}$ and $b_i \in \mathbb{R}$ are the demand model parameters, and ε_{it} is an unobservable demand shock, which we model as a zero-mean random variable. We assume that both the model parameters a_i and b_i, and the distribution function of the demand shock are initially