

A Two-Stage Mechanism for Demand Response Markets

Bharadwaj Satchidanandan, Mardavij Roozbehani, and Munther A. Dahleh
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology

Abstract—Demand response involves system operators using incentives to modulate electricity consumption during peak hours or when faced with an incidental supply shortage. However, system operators typically have imperfect information about their customers' baselines, that is, their consumption had the incentive been absent. The standard approach to estimate the reduction in a customer's electricity consumption then is to estimate their counterfactual baseline. However, this approach is not robust to estimation errors or strategic exploitation by the customers and can potentially lead to overpayments to customers who do not reduce their consumption and underpayments to those who do. Moreover, optimal power consumption reductions of the customers depend on the costs that they incur for curtailing consumption, which in general are private knowledge of the customers, and which they could strategically misreport in an effort to improve their own respective utilities even if it deteriorates the overall system cost. The two-stage mechanism proposed in this paper circumvents the aforementioned issues. In the day-ahead market, the participating loads are required to submit only a probabilistic description of their next-day consumption and costs to the system operator for day-ahead planning. It is only in real-time, if and when called upon for demand response, that the loads are required to report their baselines and costs. They receive credits for reductions below their reported baselines. The mechanism for calculating the credits guarantees incentive compatibility of truthful reporting of the probability distribution in the day-ahead market and truthful reporting of the baseline and cost in real-time.

Index Terms—Demand response, Mechanism design, Incentive compatibility, Stochastic baseline.

I. INTRODUCTION

Demand Response (DR) refers to a variety of mechanisms that aim to actively engage otherwise passive consumers, with the aim to modulate their electricity consumption during peak hours. It has been estimated that in some systems, over 20% of energy costs are driven by just 2% of peak hours. Thus, targeted demand reductions during particular periods can provide substantial value.

The most common form of demand response involves voluntary load reduction by participating consumers in exchange for monetary compensation. These programs are often referred to as incentive-based DR, and share a common challenge: How to determine the reduction in a customer's consumption during a demand response period? This involves counterfactual estimation of consumer baselines and is prone to estimation errors and strategic exploitation. Errors in baseline estimation can lead to overpayments to customers who do not actually reduce their consumption, and underpayments to those who do. Moreover, optimal consumption reductions

of the customers depend on the costs that they incur for curtailing consumption, which in general are private knowledge of the customers and which they could strategically misreport in an effort to improve their own respective utilities even if it results in an increase in the overall system cost.

In this paper, we present a mechanism that circumvents the issues that arise from baseline estimation errors and private costs by incentivizing demand response providers to truthfully self-report their baselines and costs. The mechanism is based on the mechanism developed in [1] for two-stage repeated stochastic games. In the day-ahead market, participating consumers may not know exactly their baselines and costs the following day. Rather, they may only know a probabilistic description thereof. Hence, the mechanism allows for the demand response providers to submit only a *probabilistic description* of their next-day consumption and costs to the system operator for day-ahead planning. It is only in real time, if and when called upon for demand response, that the mechanism requires the customers to report their baselines and costs. Based on these reports, optimal demand reductions are computed for all customers, and they receive credits for reductions below their reported baselines. Our mechanism for calculating the credits ensures ex post incentive compatibility of truthful bidding of the probability distribution in the day-ahead market and truthful reporting of the baseline and cost in real time. While the mechanism itself is adapted from [1], the main contribution of the paper lies in modeling the interaction between demand response providers and the system operator as a two-stage repeated stochastic game, thereby setting the stage for the application of the results developed in [1] to design an efficient and incentive-compatible demand response market.

Related work: Prior works that are the closest to our paper are [2]–[5]. While there are many differences between the formulations in [2]–[5] and this paper, we highlight three major ones here. First, the aforementioned papers restrict attention to piecewise linear valuations for demand response providers and leave the case of general concave valuations as an open problem. On the other hand, the results of this paper not only apply to concave valuations, but also to arbitrary valuation functions. Secondly and more importantly, all of these papers assume the baselines and costs of demand response providers to be deterministic so that they are perfectly known in the day-ahead market. However, an important characteristic of real-world loads is that their baselines and costs are in general random and not perfectly known in the day-ahead

market. Motivated by this, we develop a mechanism that requires the demand response providers to only report a probabilistic description of their baselines and costs in the day-ahead market, and report their actual realizations only in real time. The mechanism incentivizes demand response providers to truthfully report the probability distribution *and* the realization of these quantities in the day-ahead market and in real time respectively. Finally, the mechanisms developed in [2], [3] implement truth-telling only as a Nash equilibrium whereas the mechanism presented in this paper implements truth-telling in a stronger notion of equilibrium, namely Dominant Strategy Non-Bankrupting equilibrium. References [4], [5] implement truth-telling in dominant strategies but under the assumptions of deterministic piecewise linear costs and deterministic baselines. Since our formulation does not assume any specific form for the cost functions, nor does it assume the baselines and costs to be deterministic, it is considerably more general than the formulations in [2]–[5]. **Notation:** Given a sequence $\{x(1), x(2), \dots\}$, we denote by x^l the sequence $\{x(1), \dots, x(l)\}$ and by x^∞ the entire sequence. Given a vector \mathbf{x} , we denote by x_i its i th component and by \mathbf{x}_{-i} the vector of all other components. The paper is rather notation-heavy, and the Arxiv version of the paper [6] contains a table of all symbols for quick reference.

II. PROBLEM FORMULATION

Fig. 1 illustrates the overall setup, the individual components of which are elaborated upon in the ensuing subsections. The interaction between the Independent System Operator (ISO) and DR providers occurs over two stages, namely, the day-ahead market and the real-time market. Optimal operation of the grid involves the ISO solving the Economic Dispatch (ED) problem in the day-ahead market. In a power system containing DR providers, the ED problem must take into account the costs incurred by the DR providers in addition to the costs incurred by the generators and the reserves. In Sections II-A thru II-D, we introduce the quantities that are necessary to formulate the ED problem with DR providers. In Section II-E, we formulate the ED problem with DR providers. The solution to the problem, as we will see, is a function of certain probability distributions that are privately known to the DR providers (denoted by θ_i s in Fig. 1). Hence, a DR provider could potentially misreport its probability distribution to cause the ISO to take a day-ahead market decision that improves its own utility even if it deteriorates the overall system cost. Optimal grid operation also requires the ISO to take certain recourse actions in real time as a function of certain random variables that realize privately to each DR provider in real time (denoted by δ_i s in Fig. 1). The DR providers could misreport these quantities too in real time in an effort to improve their own respective utilities. These issues motivate the mechanism design problem which we formulate in Section II-I.

Consider a power system consisting of an ISO, a generator that should be scheduled in the day-ahead market, a fast-acting reserve from which the ISO can purchase energy in the

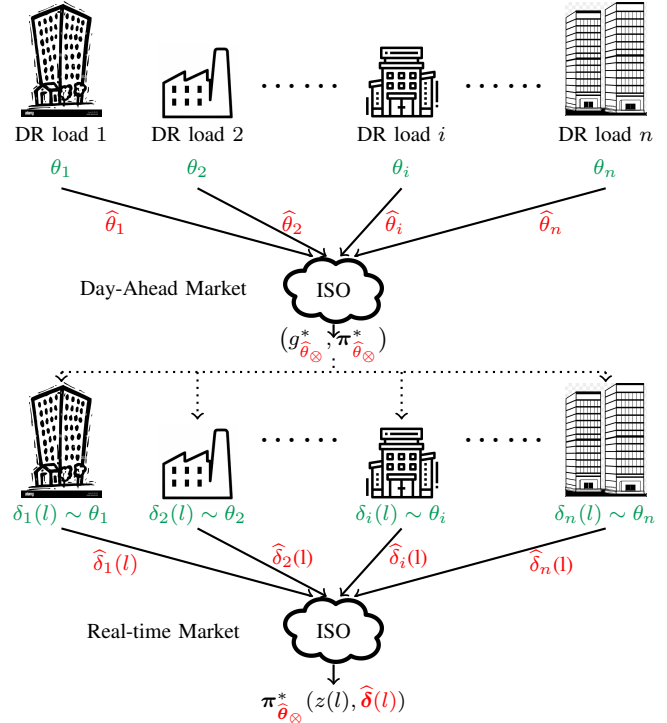


Fig. 1. The DR loads observe their type distributions $\{\theta_i\}_{i=1}^n$ in the day-ahead market. They report $\{\hat{\theta}_i\}_{i=1}^n$ to the ISO. The ISO takes day-ahead market decision $(\mathbf{g}_{\hat{\theta}}^*, \pi_{\hat{\theta}}^*)$ based on the reported distributions, which the DR loads observe. In real time on any day l , the types $\{\delta_i(l)\}_{i=1}^n$ of the DR loads realize. They report $\{\hat{\delta}_i(l)\}_{i=1}^n$ as their types. The ISO computes real-time decisions on day l based on the reported types on day l .

spot market to balance real-time demand-supply mismatches, a renewable energy source such as solar or wind, n demand response-providing loads which we refer to henceforth as the DR loads, and inelastic loads.

A. Demand of inelastic loads and DR providers

Denote by $z(l)$ the net demand of inelastic loads on day $l \in \mathbb{N}$ (at an arbitrarily fixed time interval), by which we mean the difference between the sum demand of all inelastic loads and the renewable energy that is produced on day l (at the arbitrarily fixed time interval). A quintessential feature of renewable energy sources is that the energy that they produce on any day l is a random variable in the day-ahead market. Consequently, $z(l)$ is a random variable in the day-ahead market and we denote by θ_z its probability distribution. We also assume that $\{z(1), z(2), \dots\}$ is an Independent and Identically Distributed (IID) sequence.

Denote by $d_i(l)$ the *baseline* of DR load i on day l , $i \in \{1, \dots, n\}$ and $l \in \mathbb{N}$. By *baseline*, we mean the energy that DR load i consumes on day l if it does not adjust its energy consumption for demand response.

B. Types

Real-world loads exhibit two important characteristics: (i) Their baselines could be different on different days, and

(ii) their appetite for reducing energy consumption could be different on different days. We model these two factors by the notion of a *type*. Hence, the baseline of a DR load and its appetite for reducing energy consumption being different on different days is tantamount to its type being different on different days. Section II-C provides a precise operational meaning for the type of a DR load. We denote by Δ the type space of the DR loads.

We model the time variation of a DR load i 's type by modeling it as a random variable that is drawn independently on each day l from a probability distribution θ_i . Hence, denoting by $\delta_i(l) \in \Delta$ the type of DR load i on day l , the sequence $\{\delta(1), \delta(2), \dots\}$ is an IID sequence of random variables with $\delta(1) \sim \theta_1 \times \dots \times \theta_n =: \theta_\otimes$.

Both the type distribution θ_i and the type realization $\{\delta_i(1), \delta_i(2), \dots\}$ of DR load i are private knowledge of the DR load. Moreover, these quantities are revealed to the DR load at different times. Specifically, in the day-ahead market, the DR load does not know precisely its baseline the next day, nor does it know precisely its appetite for reducing consumption the next day. Rather, it only knows the probability distribution θ_i according to which its type the next day is drawn. The DR load observes its type $\delta_i(l)$ only on day l .

Note that we have assumed the type distribution of each DR load to remain the same on all days. This is only for ease of exposition and to simplify to some extent the already complicated notation. With minor modifications, the results of this paper will continue to hold even in the case when the type distributions are different on different days but the number of different distributions in L days is $o(L)$.

C. Cost functions

Providing demand response involves the DR loads reducing their power consumption from their baseline to a lower quantity during times of power supply shortage. Power supply shortages could occur, for example, due to low renewable energy generation. Power consumption reductions come at a certain discomfort or productivity loss to the DR loads which we capture by means of a cost function. The cost incurred by a DR load on any given day is a function of both the consumption reduction and its type on that particular day. Specifically, we denote by $c_i(x, \delta_i(l)) \in \mathbb{R}$ the cost incurred by DR load i on day l for consuming x units of energy below its baseline.

Generators that have low ramp rates must have their power dispatch scheduled well in advance of the time of power delivery, and this is typically done in the day-ahead market. We denote by $c_g : \mathbb{R}_+ \rightarrow \mathbb{R}$ the production function of the generator so that $c_g(x)$ specifies the cost that the generator incurs for producing x units of energy.

We assume the availability of a fast-acting reserve generator that can produce or consume energy in real time to balance real-time demand-supply mismatches. We denote by $c_r : \mathbb{R} \rightarrow \mathbb{R}$ the production function of the reserve.

D. Demand response policy

On each day l , the ISO observes the net demand $z(l)$ and each DR load i observes its type $\delta_i(l)$. The type realizations are private to the DR loads but the ISO must know them to determine the optimal consumption reduction of each DR load on day l . Hence, the ISO requests the DR loads to report their types, but for strategic reasons that will become clear in Section II-I, they may not bid their types truthfully. Hence, we denote by $\widehat{\delta}_i(l)$ the type reported by DR load i on day l . A *demand response policy* $\pi : \mathbb{R} \times \Delta^n \rightarrow \mathbb{R}^n$ is a rule by which the ISO determines the amount by which each DR load should reduce its energy consumption from its baseline as a function of $z(l)$ and the reported types $\widehat{\delta}(l)$.

E. The day-ahead market

There are two decisions that the ISO should compute in the day-ahead market: (i) the power dispatch of the generator, and (ii) the demand response policy of the DR loads. These quantities must be computed only based on the information that is available in the day-ahead market, namely, the type distributions $\theta_1, \dots, \theta_n$ of the DR loads (which are private knowledge) and the distribution θ_z (which is public knowledge). The ISO requests the DR loads to report their type distributions in the day-ahead market. However, for strategic reasons that we describe in Section II-I, the DR loads may not bid their type distributions truthfully and so we denote by $\widehat{\theta}_i$ the type distribution reported by DR load i .

Suppose for a moment that the DR loads bid their type distributions truthfully so that $\widehat{\theta}_\otimes := \widehat{\theta}_1 \times \dots \times \widehat{\theta}_n = \theta_\otimes$, and that they bid their type realizations truthfully in real time so that $\widehat{\delta}(l) = \delta(l)$ for all l . How should the ISO compute the power dispatch of the generator and the demand response policy of the DR loads in this case?

To answer this, suppose that the energy dispatch of the generator is chosen to be g and that the demand response policy is chosen to be π . Then, the demand-supply mismatch that occurs in real time on any day l is

$$z(l) - g + \sum_{i=1}^n [\widehat{d}_i(l) - \pi_i(z(l), \widehat{\delta}(l))] =: g_r(l)$$

where $\pi_i(z(l), \widehat{\delta}(l))$ denotes the i th component of $\pi(z(l), \widehat{\delta}(l))$ and $\widehat{d}_i(l)$ denotes the baseline reported by DR load i on day l (which, recall, is specified by $\widehat{\delta}_i(l)$). The ISO must purchase $g_r(l)$ units of energy from the reserve at cost $c_r(g_r(l))$. Note that $g_r(l)$ is a random variable in the day-ahead market whose distribution depends on θ_\otimes , which we have assumed for the moment is known to the ISO. Consequently, the expected social cost — defined as the total cost incurred by the generator, the reserve, and the DR loads — on day l is

$$W(g, \pi, \widehat{\theta}_\otimes) = \mathbb{E}_{(z(l), \widehat{\delta}(l)) \sim \theta_z \times \widehat{\theta}_\otimes} \left[c_g(g) + c_r(g_r(l)) + \sum_{i=1}^n c_i(\pi_i(z(l), \widehat{\delta}(l)), \widehat{\delta}_i(l)) \right]. \quad (1)$$

The goal of the ISO in the day-ahead market is to minimize the total expected social cost and therefore determines the power dispatch $g_{\hat{\theta}_{\otimes}}^*$ and demand response policy $\pi_{\hat{\theta}_{\otimes}}^*$ as

$$(g_{\hat{\theta}_{\otimes}}^*, \pi_{\hat{\theta}_{\otimes}}^*) = \underset{g, \pi}{\operatorname{argmin}} W(g, \pi, \hat{\theta}_{\otimes}). \quad (2)$$

We denote by $W^*(\hat{\theta}_{\otimes})$ the optimal value of the above objective.

The problem of course is that the reported type distributions $(\hat{\theta}_1, \dots, \hat{\theta}_n)$ may not be equal to the true type distributions $(\theta_1, \dots, \theta_n)$, and so the decision $(g_{\hat{\theta}_{\otimes}}^*, \pi_{\hat{\theta}_{\otimes}}^*)$ which the ISO computes may not be equal to the optimal decision $(g_{\theta_{\otimes}}^*, \pi_{\theta_{\otimes}}^*)$. We describe in Section III the mechanism by which the ISO can elicit the type distributions truthfully.

It is worth noting that the optimal energy dispatch and the optimal demand response policy are functions of only the distributions θ_{\otimes} and θ_z , which we assume remain the same on all days. Hence, it suffices for the DR loads to report their type distributions just once and for the ISO to compute the optimal energy dispatch and demand response policy just once, namely, in the day-ahead market before day 1. It can reuse these decisions on all days without any loss of optimality. As mentioned before, with minor modifications, the results of this paper will continue to hold even if the distributions are different on different days but the number of different distributions in L days is $o(L)$.

F. The real-time market

After the day-ahead decisions are made, the net demand $z(l)$ of the inelastic loads and the type profile $\delta(l)$ of the DR loads realize on day l . The ISO requests the DR loads to report their type realizations and we denote by $\hat{\delta}_i(l)$ the type reported by DR load i on day l , which may or may not be equal to its true type $\delta_i(l)$. The ISO then computes $\pi_{\hat{\theta}_{\otimes}}^*(z(l), \hat{\delta}(l))$, informs each DR load the amount by which it should reduce its consumption, and purchases the residual mismatch from the spot market. Each DR load i is contractually obligated to set its energy consumption $y_i(l)$ as

$$y_i(l) = \hat{d}_i(l) - \pi_{\hat{\theta}_{\otimes}, i}^*(z(l), \hat{\delta}(l)),$$

where $\pi_{\hat{\theta}_{\otimes}, i}^*(z(l), \hat{\delta}(l))$ denotes the i th component of $\pi_{\hat{\theta}_{\otimes}}^*(z(l), \hat{\delta}(l))$. Note that the ISO can check for each DR load i on each day l whether its actual energy consumption satisfies the above equality, and declare it to be non-compliant with the DR program if it does not.

Note also that the ISO conducting the above check does not imply that a DR load i truthfully reduces its energy consumption by $\pi_{\hat{\theta}_{\otimes}, i}^*(z(l), \hat{\delta}(l))$ since it could misreport its baseline. For example, it could inflate its baseline by reporting $\hat{d}_i(l) > d_i(l)$ to give the impression of curtailing consumption without actually doing so. The mechanism presented in Section III incentivizes the DR loads to truthfully report their baselines, thereby obligating them to reduce their real-time consumption by the ISO-specified amounts.

G. Bidding strategies of DR loads

As described in Section II-E, the DR loads only know their type distributions in the day-ahead market which they have to report to the ISO. We denote by $\sigma_i : \Theta \rightarrow \Theta$ the day-ahead bidding strategy of DR load i so that $\hat{\theta}_i = \sigma_i(\theta_i)$. Here, Θ denotes the set of probability distributions on Δ .

In real time, DR loads bid their type realizations. We allow for the type bid $\hat{\delta}_i(l)$ of DR load i to be constructed based on all information that is available to it until day l , and in accordance with any arbitrary, randomized, history-dependent policy. Hence, a real-time bidding policy ϕ of DR load i specifies $\mathbb{P}_{\phi}(\hat{\delta}_i(l) | \delta_i^l, (\pi_{\hat{\theta}_{\otimes}, i}^*)^{l-1}, g_{\hat{\theta}_{\otimes}}^*)$ for each $l \in \mathbb{N}$. That is, it specifies a probability distribution over the type space Δ according to which $\hat{\delta}_i(l)$ is chosen as a function of all observations available to DR load i until day l . We denote by Φ_i the set of all real-time bidding policies.

Observe that the real-time bidding policy is a rule which specifies how a DR load should construct its type bid on any given day. While the output of the rule on any given day is a random variable which depends on the realization of the DR loads' types, *there is nothing random about the rule itself*. A DR load, without any loss of generality, can choose the rule in the day-ahead market corresponding to day 1, and as a function of θ_i — the only information that is available to it at that time. This observation leads to the notion of a *real-time bidding strategy*. A real-time bidding strategy of DR load i is a function that maps its type distribution θ_i to a real-time bidding policy in Φ_i . We denote by $\mu_i : \Theta \rightarrow \Phi_i$ the real-time bidding strategy of DR load i so that $\mu_i(\theta_i)$ is its real-time bidding policy.

Note that both the day-ahead bidding strategy σ_i and the real-time bidding strategy μ_i are functions on the set Θ , and a DR load i without any loss of generality can choose these functions “offline,” i.e., even before it observes θ_i .

We refer to the combination $S_i := (\sigma_i, \mu_i)$ as the *strategy* of DR load i and denote by \mathcal{S}_i the set of strategies available to DR load i . Note that once all DR loads fix their strategies, a functional relationship is established between $\hat{\delta}^{\infty}$ and δ^{∞} , and all random variables become well defined.

Definition 1. A strategy (σ, μ) of DR load i , $i \in \{1, \dots, n\}$, is *truthful* if

- 1) $\sigma(\theta) = \theta$ for all $\theta \in \Theta$, and
- 2) there exists $\mathcal{L} \subset \mathbb{N}$ with $\sum_{k=1}^L \mathbf{1}_{\{k \in \mathcal{L}\}} = o(L)$ such that for all $l \notin \mathcal{L}$,

$$\mathbb{P}_{\mu}(\hat{\delta}_i(l) | \delta_i^l, (\pi_{\hat{\theta}_{\otimes}, i}^*)^{l-1}, g_{\hat{\theta}_{\otimes}}^*) = \mathbf{1}_{\{\hat{\delta}_i(l) = \delta_i(l)\}}.$$

In words, a truthful strategy reports the type distribution truthfully and reports the type realization truthfully “almost all days.” We denote by \mathcal{T}_i the set of truthful strategies available to DR load i .

H. Payments and Utilities

The ISO defines for every $(i, l) \in \{1, \dots, n\} \times \mathbb{N}$ a payment rule $p_{i,l} : \theta_{\otimes} \times \Delta^{n \times l} \rightarrow \mathbb{R}$ that determines the payment

that DR load i receives on day l for setting its consumption on that day to be equal to $\widehat{d}_i(l) - \pi_{\widehat{\theta}_\otimes, i}^*(z(l), \widehat{\delta}(l))$. The payment is determined as a function of all information that is available to the ISO until day l , namely, $\widehat{\theta}_\otimes$ and $\widehat{\delta}$.

The utility that DR load i accrues on day l is defined as

$$u_{i,l}(S_i, \mathbf{S}_{-i}, \theta_\otimes, \delta^\infty) := p_{i,l} - c_i(\pi_{\widehat{\theta}_\otimes, i}^*(z(l), \widehat{\delta}(l)), \delta_i(l))$$

and its average utility is defined as

$$u_i^\infty(S_i, \mathbf{S}_{-i}, \theta_\otimes, \delta^\infty) := \liminf_{L \rightarrow \infty} \frac{1}{L} \sum_{l=1}^L u_{i,l}(S_i, \mathbf{S}_{-i}, \theta_\otimes, \delta^\infty). \quad (3)$$

Note that the above utility is a function not only of the strategy S_i that DR load i employs, but also of the strategies \mathbf{S}_{-i} that the other DR loads employ. We will return to this point in the next subsection. Before that, it is necessary to recall the notion of a *non-bankrupting strategy* that was recently introduced in [1].

Definition 2. A strategy S_i of DR load i , $i \in \{1, \dots, n\}$, is *non-bankrupting* if for all $(\mathbf{S}_{-i}, \theta_\otimes)$,

$$u_i^\infty(S_i, \mathbf{S}_{-i}, \theta_\otimes, \delta^\infty) > -\infty \quad (4)$$

almost surely.

A strategy profile $\mathbf{S} = (S_1, \dots, S_n)$ is *non-bankrupting* if S_i is non-bankrupting for all $i \in \{1, \dots, n\}$.

To elaborate, note that when the cost function of a DR load is bounded, the only way for its average utility to be $-\infty$ is for its average payment to be $-\infty$. Hence, the above definition essentially states that a DR load's strategy is non-bankrupting if it is guaranteed to not expend an infinite amount by employing that strategy, regardless of what strategies the other DR loads employ and what the type distributions are.

I. The Mechanism Design Problem

There are three problems that the ISO faces in operating the grid optimally. The first is that the optimal day-ahead decision $(g_{\widehat{\theta}_\otimes}^*, \pi_{\widehat{\theta}_\otimes}^*)$ is a function of the type distributions $\theta_1, \dots, \theta_n$ which may not be reported truthfully. Secondly, even if the ISO were to somehow compute the optimal day-ahead decisions, the optimal real-time curtailment $\pi_{\widehat{\theta}_\otimes}^*(z(l), \delta(l))$ on any day l is a function of the type realizations $\delta_1(l), \dots, \delta_n(l)$ which may not be reported truthfully. Finally, the ISO cannot verify if a DR load i faithfully reduces its consumption by $\pi_{\widehat{\theta}_\otimes, i}^*(z(l), \widehat{\delta}(l))$ since it can only observe its actual consumption $y_i(l)$ and not its counterfactual baseline $d_i(l)$.

All of these problems disappear if each DR load i employs a truthful strategy. However, note from (3) that the utility of a DR load i is a function not only of the strategy S_i that it employs, but also of the strategies \mathbf{S}_{-i} that the other DR loads employ. Consequently, a DR load i may not employ a truthful strategy if there exists $(\mathbf{S}_{-i}, \theta_\otimes)$ such

that with some non-zero probability, DR load i accrues a larger average utility by employing a non-truthful strategy than by employing a truthful one. This brings us to the mechanism design problem. We wish to design the payment rule $\{p_{i,l} : (i, l) \in \{1, \dots, n\} \times \mathbb{N}\}$ such that each DR load i 's utility (3) is almost surely maximized by choosing $S_i \in \mathcal{T}_i$ regardless of what non-bankrupting strategy profile \mathbf{S}_{-i} the other DR loads employ, and regardless of what the distribution θ_\otimes is. The next section presents such a payment rule.

III. AN EFFICIENT AND INCENTIVE-COMPATIBLE MECHANISM FOR DEMAND RESPONSE MARKETS

For any $(i, l) \in \{1, \dots, n\} \times \mathbb{N}$, the payment function $p_{i,l}$ consists of two components: (i) a first-stage payment $p_{i,l}^1$ that is determined based on only the type distributions reported in the day-ahead market and (ii) a second-stage settlement $p_{i,l}^2$ that is determined at the end of day l based on the history of type realizations reported until day l . We describe these payment rules next, which are adapted from the payment rule developed in [1] for two-stage repeated stochastic games.

A. First-stage payment

For every DR load i , $i \in \{1, \dots, n\}$, the first stage payment that it receives on any day l is a Vickrey-Clarke-Groves (VCG) payment defined as

$$p_{i,l}^1(\widehat{\theta}_\otimes) := W^*(\widehat{\theta}_{\otimes, -i}) - \left[W^*(\widehat{\theta}_\otimes) - \mathbb{E}_{(z, \widehat{\delta}) \sim \theta_z \times \widehat{\theta}_\otimes} [c_i(\pi_{\widehat{\theta}_\otimes, i}^*(z, \widehat{\delta}), \widehat{\delta}_i)] \right]$$

where $W^*(\widehat{\theta}_{\otimes, -i})$ denotes the optimal social cost that would be attained if DR load i were absent.

B. Second-stage settlement

One of the primary functions of the second-stage settlement is to penalize DR loads for type bids whose empirical distributions are not consistent with the type distribution that they bid in the day-ahead market. Towards this, define for each $i \in \{1, \dots, n\}$, $\nu \in \Delta$ and $L \in \mathbb{N}$, the empirical deviation

$$f_{i,\nu}(L) := \left[\frac{1}{L} \sum_{l=1}^L \mathbf{1}_{\{\widehat{\delta}_i(l)=\nu\}} \right] - \widehat{\theta}_i(\nu) \quad (5)$$

where $\widehat{\theta}_i(\nu)$ denotes the probability that a random variable distributed according to $\widehat{\theta}_i$ takes the value ν . It is easy to see that if DR load i employs a truthful strategy, then $f_{i,\nu}(L) \rightarrow 0$ almost surely for all $\nu \in \Delta$.

Similarly, for every $i \in \{1, \dots, n\}$, $\nu \in \Delta$, $\eta \in \Delta^{n-1}$, and $L \in \mathbb{N}$, define

$$h_{i,\nu,\eta}(L) := \left[\frac{1}{L} \sum_{l=1}^L \mathbf{1}_{\{\widehat{\delta}_i(l)=\nu, \widehat{\delta}_{-i}(l)=\eta\}} \right] - \left[\widehat{\theta}_i(\nu) \right] \left[\frac{1}{L} \sum_{l=1}^L \mathbf{1}_{\{\widehat{\delta}_{-i}(l)=\eta\}} \right] \quad (6)$$

and note that $h_{i,\nu,\eta}(L) \rightarrow 0$ almost surely if the DR loads employ a truthful strategy.

The second stage settlement rule checks on each day L if $f_{i,\nu}(L)$ or $h_{i,\nu,\eta}(L)$ exceeds a certain threshold $r(L)$ for some (ν, η) and imposes penalty $J_p(L)$ if one of them does. Towards this, define

$$E_i(l) := \left\{ \sup_{\nu} |f_{i,\nu}(l)| \geq r(l) \cup \sup_{(\nu, \eta)} |h_{i,\nu,\eta}(l)| \geq r(l) \right\}.$$

The second stage settlement that DR load i , $i \in \{1, \dots, n\}$, receives on any day l , $l \in \mathbb{N}$, is defined as

$$p_{i,l}^2(\widehat{\theta}_{\otimes}, \widehat{\delta}^l) := \left[c_i(\pi_{\widehat{\theta}_{\otimes}, i}^*(z(l), \widehat{\delta}(l)), \widehat{\delta}_i(l)) - \mathbb{E}_{(z, \widehat{\delta}) \sim \theta_{\otimes} \times \widehat{\theta}_{\otimes}} [c_i(\pi_{\widehat{\theta}_{\otimes}, i}^*(z, \widehat{\delta}), \widehat{\delta}_i)] \right] - J_p(l) \mathbf{1}_{E_i(l)}. \quad (7)$$

How should the threshold sequence $\{r\}$ and the penalty sequence $\{J_p\}$ be chosen? The sequence $\{r\}$ should be chosen to balance two competing objectives. On the one hand, $r(l)$ must tend to 0 as $l \rightarrow \infty$ since otherwise, the set of type bids that fall within the threshold will be “large,” thereby violating incentive compatibility. However, if the sequence shrinks too quickly, then even truthful type bids may fall outside the threshold often, thereby incurring large penalties which would violate individual rationality. To balance the objectives, $\{r\}$ must be chosen such that (i) $\lim_{l \rightarrow \infty} r(l) = 0$, and (ii) for some $\gamma > 0$,

$$r(l) \geq \sqrt{\frac{\ln 2l^{1+\gamma}}{2l}}.$$

The penalty sequence must be chosen to satisfy

$$\lim_{l \rightarrow \infty} \frac{J_p(l)}{l} = \infty.$$

See [1, Section III] for further intuition for these conditions.

The total payment received by DR load i on day l is

$$p_{i,l}(\widehat{\theta}_{\otimes}, \widehat{\delta}^l) = p_{i,l}^1(\widehat{\theta}_{\otimes}) + p_{i,l}^2(\widehat{\theta}_{\otimes}, \widehat{\delta}^l). \quad (8)$$

We now state the main result of the paper.

Theorem 1. *Consider the mechanism defined by the decision rule (2) and the payment rule (8).*

- 1) *For every $i \in \{1, \dots, n\}$, $S_i \in \mathcal{S}_i$, $T_i \in \mathcal{T}_i$, non-bankrupting strategy profile \mathbf{S}_{-i} , and θ_{\otimes} ,*

$$u_i^{\infty}(T_i, \mathbf{S}_{-i}, \theta_{\otimes}, \delta^{\infty}) \geq u_i^{\infty}(S_i, \mathbf{S}_{-i}, \theta_{\otimes}, \delta^{\infty}) \quad (9)$$

almost surely.

That is, every DR load accrues a larger utility by employing a truthful strategy than by employing any other strategy, regardless of what non-bankrupting strategies the other DR loads employ.

- 2) *Suppose that $W^*(\theta_{\otimes, -i}) - W^*(\theta_{\otimes}) \geq 0$ for all θ_{\otimes} . For every $i \in \{1, \dots, n\}$, $T_i \in \mathcal{T}_i$, $\mathbf{S}_{-i} \in \mathcal{S}_{-i}$, and θ_{\otimes} ,*

$$u_i^{\infty}(T_i, \mathbf{S}_{-i}, \theta_{\otimes}, \delta^{\infty}) \geq 0 \quad (10)$$

almost surely.

That is, every DR load accrues a nonzero utility by employing a truthful strategy regardless of what strategies the other DR loads employ.

- 3) *If for all $i \in \{1, \dots, n\}$, $S_i \in \mathcal{T}_i$,*

$$\lim_{L \rightarrow \infty} \frac{1}{L} \sum_{l=1}^L \left[c_g(g_{\widehat{\theta}_{\otimes}}^*) + c_r(g_r(l)) + \sum_{j=1}^n c_j(\pi_{\widehat{\theta}_{\otimes}, j}^*(z(l), \widehat{\delta}(l)), \delta_j(l)) \right] = W^*(\theta_{\otimes}) \quad (11)$$

almost surely.

That is, if all DR loads employ a truthful strategy, then the average social cost that is incurred is almost surely equal to its optimal value.

Proof. The proof mimics that of [1, Theorem 1] and is omitted in the interest of space. \square

Certain numerical results illustrate the cost benefits afforded by the proposed mechanism as compared to the posted price mechanism — a popular alternative that has been employed in certain demand response trials in Europe. In the interest of space, they have been moved to [6].

IV. CONCLUSION

A key difficulty in demand response schemes is to determine if a load indeed reduces its power consumption when called upon for demand response and if so by how much. This is challenging because the system operator can only observe the actual consumption of a load and not its counterfactual baseline. Additionally, determining the optimal curtailment levels of the loads is also challenging as they depend on the loads' privately-known costs which the loads may misreport. We have designed a two-stage mechanism to address these issues. The mechanism guarantees ex post incentive compatibility of truthful bidding in both the day-ahead market and in real time, ex post individual rationality, and attains the optimal social cost.

REFERENCES

- [1] B. Satchidanandan and M. A. Dahleh, “Incentive compatibility in two-stage repeated stochastic games,” arXiv preprint arXiv:2203.10206 [econ.TH], March 2022.
- [2] B. Xia, K.-Y. Lee, S. Shakkottai, and D. Kalathil, “A market for retail electric provider based demand response,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 8429–8434.
- [3] D. Muthirayan, D. Kalathil, K. Poolla, and P. Varaiya, “Mechanism design for self-reporting baselines in demand response,” in *2016 American Control Conference (ACC)*, 2016, pp. 1446–1451.
- [4] —, “Baseline estimation and scheduling for demand response,” in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 4857–4862.
- [5] —, “Mechanism design for demand response programs,” *IEEE Transactions on Smart Grid*, vol. 11, no. 1, pp. 61–73, 2020.
- [6] B. Satchidanandan, M. Roozbehani, and M. A. Dahleh, “A two-stage mechanism for demand response markets,” arXiv preprint arXiv:2205.12236 [eess.SY], May 2022.