

Firming renewable power with demand response: an end-to-end aggregator business model

Clay Campaigne & Shmuel S. Oren

Journal of Regulatory Economics

ISSN 0922-680X

Volume 50

Number 1

J Regul Econ (2016) 50:1-37

DOI 10.1007/s11149-016-9301-y



Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media New York. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Firming renewable power with demand response: an end-to-end aggregator business model

Clay Campaigne¹ · Shmuel S. Oren¹

Published online: 18 May 2016

© Springer Science+Business Media New York 2016

Abstract Environmental concerns have spurred greater reliance on variable renewable energy resources (VERs) in electric generation. Under current incentive schemes, the uncertainty and intermittency of these resources impose costs on the grid, which are typically socialized across the whole system, rather than born by their creators. We consider an institutional framework in which VERs face market imbalance prices, giving them an incentive to produce higher-value energy subject to less adverse uncertainty. In this setting, we consider an “aggregator” that owns the production rights to a VER’s output, and also signs contracts with a population of demand response (DR) participants for the right to curtail them in real time, according to a contractually specified probability distribution. The aggregator bids a day ahead offer into the wholesale market, and is able to offset imbalances between the cleared day-ahead bid and the realized VER production by curtailing DR participants’ consumption according to the signed contracts. We consider the optimization of the aggregator’s end-to-end problem: designing the menu of DR service contracts using contract theory, bidding into the wholesale market, and dispatching DR consistently with the contractual agreements. We do this in a setting in which wholesale market prices, VER output, and participant demand are all stochastic, and possibly correlated.

Keywords Electricity markets · Demand response · Aggregator · Business model · Renewables integration · Market design · Screening mechanisms

JEL Classification D11 · D45 · D47 · Q41 · Q42

✉ Shmuel S. Oren
oren@ieor.berkeley.edu

¹ Department of Industrial Engineering and Operations Research, University of California, Berkeley, 4141 Etcheverry Hall, Berkeley, CA 94720, USA

1 Introduction

1.1 Background and motivation

Environmental concerns regarding global warming and the adverse health effects of emissions produced by fossil fuel generation have led to a greater reliance on renewable sources of generation, such as solar and wind, which are inherently variable and uncertain. This trend is accompanied by increased proliferation of distributed resources, storage, and smart grid technologies for metering and control, which facilitate demand response and greater observability of the grid. As a result, the electric power industry faces new challenges in planning and operation of the power system that require new institutional and regulatory frameworks, along with appropriate market mechanisms to achieve productive and allocative efficiencies. While the conventional approach to mitigating adverse uncertainty and variability on the supply and demand sides has been increased reliance on reserves and flexible generation units, this approach is expensive, and will undermine the economic and environmental goals of renewables integration. Mobilizing demand side flexibility enabled by smart metering and other smart grid technologies to mitigate the uncertainty and variability of renewable resources is a sustainable solution for addressing the operational challenges posed by massive integration of renewables.

Alternative approaches to integrate renewable resources into the power grid and facilitate demand response have been proposed and experimented with by policy makers around the world, and have been the subject of numerous academic studies in the economics and power system literature. From an economics perspective, the gold standard approach to achieving production and allocative efficiency is a centralized market where all renewable resources and conventional resources are pooled together with demand side resources, responding to real time marginal prices set through a market clearing mechanism. However, while such an approach may serve as a useful benchmark, it is impractical, as it would require the system operator to collect information and co-optimize the dispatch of a vast number of resources including conventional generation, renewables and participating demand side resources (PDR). The computational and institutional barriers to such a centralized approach calls for more pragmatic second-best alternatives with more manageable scope. Recent regulatory initiatives such as “Reforming the Energy Vision” (REV) initiated by the New York Public Service Commission (PCS) promote a more decentralized approach as a way to facilitate the integration of decentralized renewable resources and demand response ([MDPT Working Group 2015](#)). Likewise, the concept of aggregators that can pool demand side resources and act as intermediaries, offering load reduction into the wholesale market, has been popularized by the emergence of commercial entities such as EnerNOC. The scope of such aggregation can be expanded to include behind-the-meter resources and distributed renewable resources.

In this paper we propose and analyze an aggregator business model that assembles a portfolio of variable energy resources (VER) such as wind, and of flexible demand response (DR), with the purpose of producing a firm and controllable bundled energy resource that can be offered into the ISO wholesale day ahead market. We presume that the aggregator is in a position to acquire detailed information and enter into

contractual arrangements that enables it to mobilize the DR flexibility so as to offset the VER uncertainty and variability. Such a bundled resource will relieve the ISO from having to procure additional reserves or other ancillary service products for the purpose of mitigating renewables intermittency.

Our premise in this paper is that future regulatory reforms will provide incentives to VER to firm up their output and induce loads to surrender their flexibility. On the VER side, such incentives will be enabled when subsidies to renewables such as feed in tariffs will be replaced by nondiscriminatory market mechanisms. Under such a mechanism, uncertain resources bear the cost they impute on the system, whereas flexible resources are rewarded for the flexibility. Furthermore, VER will have to schedule their forecasted production and be subject to deviation settlements in the real time market like other resources, whereas firmed up VER will be eligible for capacity payments through resource adequacy mechanisms. On the demand side, ex ante contractual agreements with an aggregator that compensate the customer for forgone consumption and “information rents” should provide incentives for load to reveal and trade their flexibility.

The two principal forms of demand response are direct load control, wherein the aggregator physically constrains participants’ consumption during scarcity events, and price-based control, wherein the participants face real-time prices that reflect current system conditions.¹ Direct load control has been studied in theory (Chao 1983) and implemented in practice, particularly in contexts such as air conditioner cycling (RLW Analytics 2007). It has the advantage in terms of system reliability, because the response is more predictable; as well as with respect to billing simplicity and predictability, because the customer does not face state-dependent prices. On the other hand, price-based control provides customers with more flexibility (Braithwait et al. 2006). According to standard microeconomic models, the most economically efficient form of control is real-time pricing, because it ensures that customers consume exactly when their marginal benefit is greater than the instantaneous marginal cost of power production (Borenstein 2005; Caramanis et al. 1983; Holland and Mansur 2006).² If the consumer’s demand curve for power were constant over time, then a direct load control contract linked to spot prices would result in the same consumption decisions as real time pricing (Chao and Wilson 1987).

Restructured electricity markets are premised on treating electricity at the wholesale level as a homogeneous commodity that is produced and traded based on fluctuating price signals. We argue, however, that at the retail level electricity can be offered as a quality differentiated service with predetermined prices and uncertain availability (quantity control). Such uncertainty is realized through direct load control, or customer response to a load control signal subject to a noncompliance penalty. The above perspective, which has been articulated by Oren (2013), is the underlying paradigm explored in this paper and we will not attempt to contrast it with a real time pricing approach, which as we concede, represents the “economic gold standard.” Specifically,

¹ Caramanis et al. (1983) categorized direct load control as “price/quantity transactions,” and price-based control as “price only transactions”.

² The standard analysis ignores intertemporal interactions; but see, for example, Tsitsiklis and Xu (2015) for an extension to pricing for contribution to system-wide ramping cost.

we consider a profit-maximizing aggregator contracting ex ante with DR participants for the right to send a curtailment signal with a specified probability (or, more generally, in specified states of the world, as reflected by a publicly observable index). The curtailment signal effectively raises the participant's price for calling energy from a particular capacity increment from its original retail rate R , to an exogenously determined "penalty price," $H > R$. That is, the capacity increment is an option, and curtailment raises the strike price. We assume that demand response load pays a regulated retail rate, and has no other venue for participating in wholesale markets. The case where $H = \infty$ can be interpreted as direct load control. This generalizes plans like PG&E's SmartRate plan, which raises the customer's tariff for 15 days a year or less. In our generalization, different slices of the household's consumption capacity have different probabilities of facing curtailment/penalty rate signal. Combined with a model of stochastic valuations for service, this approach models two kinds of imperfect or fractional DR yield: DR that fails to materialize because the customer would not have consumed in the first place (the ex post valuation of consumption is less than R), and DR that fails to materialize because the customer's ex post valuation is higher than the penalty price, H . In either case we assume that the valuations are constant throughout the time interval, and each valuation is for the energy from an infinitesimal capacity slice, so we do not consider the possibility of partial exercise of a capacity increment within a period. However, by "stacking" these increments, the model generalizes to horizontal load slices that can be fractionally utilized, at a constant level during the period. Less-than-infinite penalties may be a happy medium between the intrusiveness of a hard constraint, and the complexity of a real-time price.

Our proposed business model is based on a "fuse-control paradigm" (Margellos and Oren 2015) where the aggregator manages the service quality for the aggregate consumption by imposing a capacity constraint, or by signaling a capacity threshold above which the penalty will be imposed, and leaves the decision of allocating the available power to devices behind the meter to the household. This is a less intrusive alternative to direct curtailment of individual devices, such as air-conditioner cycling programs, for instance. Furthermore, delegating the behind-the-meter allocation allows the customer to reflect intertemporal variations in preferences for different electricity uses, and capture the effect of behind-the-meter variable resources such as solar panels, local storage devices, and deferrable energy uses such as electric vehicle charging, HVAC etc. In our model the aggregator is assumed to submit price-contingent hourly offers into the ISO day ahead market and dispatch curtailment signals to its contracted load based on the awarded quantities in the day ahead market, the realized renewable output, and the deviation settlement prices. Our end-to-end approach seeks to co-optimize the contract design on the demand side with the aggregator's bidding strategy in the ISO day ahead market and the DR deployment strategy.

1.2 Related work

Motivated by the same concern about the subsidization of VERs' contribution to reserve costs, Bitar et al. (2011, 2012) consider several stylized market models for renewable power. They use a newsvendor-type model to quantify the effect of imbal-

ance charges on the offer behavior and profit of a renewable producer, and to quantify, for example, the value of forecast improvement in this policy environment. Their second model is a market for reliability-differentiated power, originally studied by [Tan and Varaiya \(1991, 1993\)](#). In this model, the producer owns a stochastic power resource, and sells its entire production in advance without using reserves, by offering contracts with imperfect service reliability. Our model can be seen as a synthesis and generalization of these two models.

We cast the problem of designing an optimal menu of variable-reliability demand response contracts as a variation on the classic monopsony screening problem from contract theory. Our approach to embedding this screening problem in a wholesale electricity market follows the literature on priority service, particularly [Chao and Wilson \(1987\)](#) and [Chao \(2012\)](#). However, that literature has focused on perfect competition or regulated social welfare maximization,³ and abstracts away from the scheduling and recourse decisions of individual producers. Because we are interested in new business models that manage imbalance, we update the priority service approach in a profit maximization setting, where imbalance cost is reflected by imbalance prices. We also consider preliminary extensions of our analysis to competitive settings.

Another point of contrast with [Chao \(2012\)](#) is that our stochastic demand model disaggregates the aggregate demand curve along the quantity axis, and then adds post-contracting noise to *valuations*, in a manner similar to [Courty and Li \(2000\)](#)'s sequential screening model. However, in contrast to most screening environments, including that of [Courty and Li \(2000\)](#), our producer's contracting problem is embedded in a newsvendor-like problem, with asymmetric linear prices for positive and negative imbalance. As a consequence, the aggregator's benefit is not linear (i.e. is not an expectation) over a type distribution. The aggregator-*cum*-producer co-optimizes its demand response menu with a day ahead offer quantity, with the demand response providing recourse in case of real-time imbalance.

Recently, [Crampes and Léautier \(2015\)](#) have used contract theory to study the welfare effects of allowing demand response participation in adjustment markets, when DR participants have private information about their utility from consumption. In their setting, vertically integrated producers contract as profit-maximizing monopoly⁴ retailers with consumers in the first stage, where consumers "buy their baseline" on which adjustment is settled, and producers incur the obligation to produce the contracted amount. Then, in the second stage, all producers experience an identical supply shock (capacity failure), and both producers and consumers can participate in a competitive adjustment market. They employ a stylized, two-type model with asymmetric information to show that there exist cases in which allowing consumers to participate in a competitive adjustment market reduces social welfare, by creating sufficiently large distortions in first-stage retail contracting.

There are two major points of contrast between our model and that of [Crampes and Léautier \(2015\)](#) worth mentioning. [Crampes and Léautier \(2015\)](#) treat retail contracting as monopolistic, and view the adjustment market as competitive. In contrast, we take

³ But see [Wilson \(1993\)](#), who treats profit maximization but in a slightly different setting from ours.

⁴ However, [Crampes and Léautier](#) argue that the qualitative insights carry over to imperfectly competitive settings.

both retail and wholesale prices as exogenous, and we consider monopsony contracting in the adjustment market, with a preliminary extension to Cournot oligopsony. This reflects our focus on the medium-term future, in which the aggregation market has few participants, and is small in toto relative to wholesale markets. We view retail rates as administratively determined, in a manner that is exogenous to consumers' and aggregators' decision-making. This is because we are interested in the normative business decisions of aggregators.

The second point of contrast with [Crampes and Léautier \(2015\)](#) is purely a modeling choice. [Crampes and Léautier \(2015\)](#) consider a two-type demand model, to give the clearest demonstration of a distortion effect. We model a continuous market demand curve, comprising a continuum of types. This provides a more detailed, less stylized account of how an aggregator should optimize a production offer and DR dispatch policy with knowledge of market statistics, renewable output, demand conditions, etc. While a two type demand model may suffice to illustrate welfare implications, our modeling choice is motivated by a market design perspective, addressing the operational question of "how to" construct and utilize a DR contract menu.

1.3 Introduction to the model

We consider the profit maximization problem of an aggregator. This aggregator has two sources from which it produces energy: a VER ("wind") with known probability distribution over production quantities, and a population of DR participants, with whom it signs contracts *ex ante* (say, at the beginning of the season) giving the aggregator the right to curtail them with specified probabilities. The market system operator treats reductions in participants' consumption, induced by curtailment, as the aggregator's production. The DR participants have private information regarding their valuation for service. For simplicity we assume that the aggregator acts as a monopsonist purchaser of rights to curtail increments of their capacity with specified probabilities. The monopsony assumption is obviously questionable from an institutional perspective unless there are regulatory barriers to entry for aggregators. Our main motivation for this assumption is to focus on the contracting details. We will discuss how this assumption can be relaxed somewhat allowing for a symmetric Cournot oligopsony model of aggregators competing by offering exclusive contracts to DR load, which is used to firm up their VER supply that they offer into the wholesale market. The exclusivity assumption can be justified on technological grounds, since implementing a curtailment policy either through direct load control or penalty signal may require specialized aggregator-owned equipment. We analyze the aggregator's problem as a "screening problem" ([Börger 2010](#)) in which the aggregator's benefit function reflects its participation in the wholesale electricity market, as we describe presently.

The aggregator bundles the VER and DR production for sale into a wholesale electricity market by choosing an energy offer quantity q into the day-ahead (DA) market, contingent on DA information. If the DA offer is made contingent only on the price p , then this price-contingent offer policy can be interpreted as a supply offer curve. In the day-ahead, the aggregator receives revenue $p q$. In the real time dispatch (RT) stage, it learns the wind outcome s , the prices a and b for positive and negative deviations respectively from the DA commitment quantity q , and chooses a set of

DR participants to curtail. Ex post, this results in a net (“aggregated” or “bundled”) production quantity, $s + DR$. The aggregator then pays $b(q - s - DR)^+ - a(q - s - DR)^-$, $a < p < b$, to settle the difference between the ex post production and the DA commitment. The joint probability distribution over all information is known in advance. One might say that from the ex ante perspective, the aggregator’s problem is a probability distribution over newsvendor problems, and in each newsvendor problem, after the initial quantity choice, the aggregator can take recourse actions in response to observed “demand” (here we mean negative wind), by dispatching DR. The DR cost is nonlinear, determined by the economics of the screening problem. In general the probability distribution over DR actions may be constrained to conform to contracts negotiated ex ante, but this advance commitment has no *economic* effect: assuming the aggregator faces no statistical or computational limitations, DR can be treated purely as scenario-dependent recourse (see Sect. 2.4.1), whose optimal distribution it can foresee and therefore commit to from the ex ante stage.

The population of demand response participants is modeled as a continuum of “increments” of capacity—i.e., potential consumption. Ignoring stochasticity of valuations, each increment is a differential, dx , on the quantity axis of the population aggregate demand curve.⁵ Each infinitesimal increment has private information, indexed by its type $\tau \in [\underline{\tau}, \bar{\tau}] \subset \mathbb{R}$, parameterizing the distribution over its “ex post valuation” θ for power at the time of consumption. The types of increments are distributed according to a measure with associated distribution function G and density g , with convex compact support $[\underline{\tau}, \bar{\tau}] \subset \mathbb{R}$. The measure of a set of increments under G represents the total potential consumption capacity of that set, in MW.

Before laying out the microeconomic model of how DR is produced and how much it costs, we can informally write the aggregator’s problem, from the ex ante perspective, as⁶:

$$\begin{aligned} & \max_{q, DR, T} J_{EA}(q, DR, T) \\ & = \max_{q, DR, T} \mathbb{E} \left[\underbrace{p q}_{\substack{\uparrow \\ \text{day-ahead revenue}}} + a \overbrace{(DR + s - q)^+}^{\text{overproduction}} - b \overbrace{(q - DR - s)^+}^{\text{shortfall}} \right] - \underbrace{T}_{\substack{\uparrow \\ \text{payment to DR}}} . \end{aligned} \tag{1}$$

Here q , DR , and T are policy variables. The DR dispatch is determined in real time, although in accordance with a policy determined ex ante, and the corresponding payment T is made ex ante. The exogenous random variables are⁷:

⁵ Our model implies that the aggregator offers a menu of quality-differentiated service options (Mussa and Rosen 1978) to each increment, independently of contracting with the rest. A consumer chooses a menu item for each of its increments, and its total utility and transfer are Lebesgue integrals over the corresponding utility and transfers for each increment. This treatment rules out quantity-based nonlinear price discrimination.

⁶ The EA subscript indicates that this objective is an expectation from the ex ante perspective.

⁷ We use the same symbols for random variables and their realized values.

- p : day ahead (“DA”) price $\in [p, \bar{p}]$
- a : overproduction payment rate $\in [a, \bar{a}]$, realized in RT
- b : shortfall penalty rate $\in [b, \bar{b}]$, realized in RT
- s : VER (“wind”) realization, $\in [s, \bar{s}]$, realized in real time (RT),
- $\{\theta_\tau : \tau \in [\underline{\tau}, \bar{\tau}]\}$: DR participants’ valuations, realized ex post.

We generally assume that $0 < a < p < b$. Allowing a penalty for overproduction, i.e. $a < 0$, reflecting the frequent occurrence of negative real time prices, would involve minor complications.⁸

The last item is a continuum of random variables: a *process*, although indexed by the type of the DR participant, rather than by time.⁹ It does not show up in the informal objective, but we will explain how it affects the DR quantity DR , and payment T .

1.4 Information and decision structure

In our most general analysis, random variables are realized, and decisions taken, in four temporal stages. As mentioned above, the contracting decision that is made at the ex ante stage is an expectation over scenario-contingent decisions which can be treated, for analytical purposes, as if they are postponed until real time, after all uncertainties affecting aggregator decisions are realized. At each subsequent stage, information from the previous stage is retained, new payoff-relevant random variables are realized, and forecasts of the random variables in future stages may be updated. We denote the tuple of random variables at each time-stage, drawn from the set of possible events, as “ $\omega \in \Omega$,” with a subscript denoting that time-stage, and we denote the non-payoff-relevant component as ξ with the same subscript.

⁸ In this case (let us still assume $0 < p < b$), we would allow the curtailment of the VER if technologically feasible, represented by replacing the wind outcome s with a policy variable that is a truncation of s , $\tilde{s} \leq s$. Whenever $a < 0$, it is clear that the aggregator would set $DR + \tilde{s} \leq q$. In cases of overproduction, the aggregator would reduce DR until $DR \leq q - s$ if possible, maintaining $\tilde{s} = s$. But if this is not sufficient, because $q - s < 0$, then at the optimum, $DR = 0$ and $\tilde{s} = q < s$. This is because DR is nonnegative and costly to produce (the aggregator dispatches DR by purchasing consumption options), and wind has zero operating cost. (We ignore the possibility of renewable production subsidies, as discussed in the introduction.) The general first-order conditions introduced up through Sect. 3.2 continue to hold with the minor technical adjustments, since $a < 0$ only makes the aggregator’s objective “more concave.” We exclude cases where $a < 0$ in the analytical example of Sect. 3.3 for simplicity, since the solution presented there in Table 1 already comprises a profusion of cases.

⁹ We assume that this process is jointly measurable with respect to the product measure induced by g and the probability measure over θ_τ . Therefore the process admits an essentially unique decomposition into an idiosyncratic component and an aggregate component (Al-Najjar 1995). We are not particularly interested in technical issues regarding measurability. Instead we simply posit, as suggested in Judd (1985), that the idiosyncratic noise obeys an exact strong law of large numbers. In a similar spirit, we make whatever assumptions necessary to license the application of Fubini’s theorem to exchange the order of integration with respect to g and expectation over the process θ_τ , which should not be demanding, since each valuation takes a value in a compact interval $[\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$. Further, the increments’ decision that is contingent on this information process involves no strategic interaction, so most of the potential technical complications do not arise.

0. Ex ante (EA) stage. The aggregator learns all probability distributions. The aggregator offers the same menu of contracts to each member of the population of DR participants, and the DR participants select their preferred plans. The aggregator makes the aggregate payment T from Eq. (1).¹⁰
1. Day-ahead (DA) stage. The aggregator learns the DA price p : $\omega_{DA} = (p, \xi_{DA}) \in \Omega_{DA}$, and chooses its offer quantity, $q(\omega_{DA})$. The function $\omega_{DA} \mapsto q(\omega_{DA})$ can be interpreted as a supply function offered in the ISO DA market, if it is p -measurable.¹¹
2. Real time dispatch (RT) stage. The aggregator learns the imbalance prices (a, b) and wind outcome s : $\omega_{RT} = \omega_{DA} \times (a, b, s, \xi_{RT}) \in \Omega_{RT}$, and chooses the set of DR increments to send curtailment signals to: $\{\tau : k(\tau, \omega_{RT}) = 1\}$.¹² A general curtailment function is denoted as $k : [\underline{\tau}, \bar{\tau}] \times \Omega_{RT} \rightarrow [0, 1]$, where the value $k(\tau, \omega_{RT})$ is the ex ante probability that type τ is curtailed in RT event ω_{RT} . In Assumption 2 below, we restrict attention to curtailment rules of the form $k(\tau, \omega) = \mathbb{1}_{\{\tau \leq \hat{\tau}(\omega_{RT})\}}$.
3. Ex post (EP) stage. The participants' valuations are realized: $\omega_{EP} = \omega_{RT} \times \{\theta_\tau : \tau \in [0, N]\}$; this determines the realized quantity of demand response. We denote the latter random variable (or its realization in event ω_{EP}) as $DR(\omega_{EP}, \hat{\tau}(\omega_{RT}))$.

The aggregator's primary policy variables are thus

$$q : \omega_{DA} \mapsto q(\omega_{DA}) \in \mathbb{R}$$

$$k(\cdot, \cdot) : (\tau, \omega_{RT}) \mapsto k(\tau, \omega_{RT}) \in [0, 1].$$

We alternate between $\hat{\tau}$ and k notation for the curtailment policy as convenient. The payment T is a decision, but the screening analysis lets us express the optimal T given a curtailment policy k as a functional of that policy.

This is a rather general description, which is suitable to our analysis of the DR contract design component of the aggregator's problem. However, we only solve the aggregator's whole problem (the "end-to-end problem"), which embeds the DR contracting into a wholesale offer problem, in special cases. In these special cases, some of the information is realized at earlier stages than in the general case, or is never stochastic; that is, certain advance forecasts are assumed to be perfect.

We occasionally omit the DA, RT and EP subscripts on random outcomes when the referent is clear from the context.

1.5 Organization of remaining sections

The remainder of the paper proceeds as follows: In Sect. 2, we characterize the class of merit order curtailment policies and their corresponding contracts in our setting.

¹⁰ This is a slight simplification: the collection of any penalties from increments, for violating curtailment signals, is also netted out from T ex post.

¹¹ In our example in Sect. 3.3, the DA forecasts of imbalance prices are assumed to be perfect, so they are effectively revealed in this stage as well.

¹² The product notation " \times " denotes concatenation of ordered tuples.

This determines the *cost* of implementing a curtailment policy. In Sect. 3, we analyze the aggregator's end-to-end problem, which embeds DR contracting and dispatch into a newsvendor-style wholesale market. In Sect. 3.1, we present the general model of the aggregator's benefits from demand response. In Sects. 3.2 and 3.3, we consider two special cases of the end-to-end problem that we can solve to successive degrees of explicitness. These give us insight into the structure of the aggregator's end-to-end problem. Finally, we conclude and outline extensions and future directions.

2 Demand response: utility model, production, and payment

In this section, we analyze the DR contracting process. In Sect. 2.1, we lay out the key economic assumptions that make our problem tractable. In Sect. 2.2, we introduce direct mechanisms and specify our capacity increments' utility function in a direct mechanism. In Sect. 2.3 we invoke the Revelation Principle, the foundation of mechanism design. Section 2.4.1, we note that contracting for merit order curtailment policies can also be implemented "in real-time," at least if we treat the economic model literally, which gives us economic insight into the set of implementable contracts and policies.

2.1 A one-dimensional type space

The key feature of our analysis of the DR contractual screening problem is that we make two assumptions that are jointly sufficient to render the increments' type space "one-dimensional."¹³

The first assumption is that conditional on any real-time outcome ω_{RT} , "higher types" have a higher distribution over ex post valuations, in the sense of first-order stochastic dominance (FOSD). At the ex post stage, each capacity increment will consume if its realized valuation for consumption is sufficiently high. Since there is a continuum of increments, each one is infinitesimal. Therefore we can assume that no two increments have the same type, and that a distinct ex post valuation random variable is associated with each type¹⁴:

Definition 2.1 (*Ex post valuation*) The "ex post valuation" is the dollar value that an increment of type τ derives from consumption in ex post event ω_{EP} . It is a random variable, with value

$$\theta(\tau, \omega_{EP}).$$

The cdf for the valuation $\theta(\tau, \omega_{EP})$, conditional on the type and the information publicly available in real time, is

¹³ See Börgers (2010, Chap. 5.6).

¹⁴ This could easily be generalized, but assuming a very high density over a short interval of types should be a reasonable approximation to a point mass on a single type, and this setup allows us to distinguish increments anonymously—i.e. only by type—which is convenient.

$$\Pr\{\theta(\tau, \omega_{EP}) \leq \theta|\omega_{RT}\} \triangleq F(\theta|\tau, \omega_{RT}).$$

We denote the conditional pdf $\partial/\partial\theta F(\theta|\tau, \omega_{RT})$ as $f(\theta|\tau, \omega_{RT})$.

The set of distributions over ex post valuations obeys a monotonicity and smoothness condition:

Assumption 1 (*First-order stochastic dominance—FOSD*) The distribution over ex post valuations is ordered by, and differentiable with respect to, ex ante type:

1. $F(\theta|\tau, \omega_{RT}) < F(\theta|\tau', \omega_{RT}) \forall \tau > \tau', \forall \theta, \omega_{RT}$;
2. $\partial F(\theta|\tau, \omega_{RT})/\partial\tau < 0$; and
3. $\exists M \in \mathbb{R}_+$ s.t. $|\partial F(\theta|\tau, \omega_{RT})/\partial\tau| < M$ for a.e. $(\omega, \tau) \in \Omega_{RT} \times [\underline{\tau}, \bar{\tau}]$ under the product measure.

In fact, this condition can be weakened so that there is a different FOSD ordering of the type space for each real time outcome, but the corresponding informational requirements for the aggregator may seem unrealistically onerous.

Our second assumption is a restriction on the set of DR curtailment policies:

Assumption 2 (*Merit order curtailment policy*) We restrict attention to DR curtailment policies with a “Merit Order,” or cutoff, form: $k(\tau, \omega_{RT}) = \mathbb{1}_{\{\tau \leq \hat{\tau}(\omega_{RT})\}}$. That is, in each real-time RT outcome ω_{RT} , the aggregator sends the curtailment signal to all increments with ex ante type less than some event-specific cutoff type, $\tau = \hat{\tau}(\omega_{RT})$.

(We explain in Sect. 2.3 below how the curtailment policy can take the increments’ types as an argument, despite that the types are the increments’ private information.) The choice of DR dispatch policy $\hat{\tau}(\cdot)$ determines the quantity of demand response, which is a random variable, whose value is realized ex post: $DR(\omega_{EP}; \hat{\tau}(\omega_{RT}))$ (see Definition 3.1).

The combination of Assumptions 1 and 2 ensures that the type space $[\underline{\tau}, \bar{\tau}]$ is “one-dimensional” in a key economic sense.¹⁵ Consider any pair of types $\tau_2 > \tau_1$ and any merit order allocation rule $k(\tau, \omega_{RT}) = \mathbb{1}_{\{\tau \leq \hat{\tau}(\omega_{RT})\}}$. The marginal utility for type τ_2 being switched from allocation $k(\tau_1, \cdot)$ to allocation $k(\tau_2, \cdot)$ is greater than the marginal utility for τ_1 undergoing the same switch. So higher types value “higher allocations” (that is, being curtailed less) more than do lower types, a fact which allows the aggregator to “separate” the types. This essentially reduces DR contracting problem to a textbook single-stage screening problem by separability across events $\omega_{RT} \in \Omega_{RT}$.

In a standard sequential screening problem (Courty and Li 2000), a *single* agent’s ex post valuation is realized conditional on its ex ante type. Our problem is putatively dynamic, but because we exogenously specify how the increment makes the final consumption decision, it becomes effectively static, except that the increments’ first-stage utility function reflects the information dynamics.¹⁶ In addition to the dynamic

¹⁵ See Börgers (2010, Chap. 5.6).

¹⁶ This “stiction” holds in a more specific sense than the more general result that sequential mechanisms can be reduced to a particular kind of static mechanism (Krähmer and Strausz 2015). This reflects our concrete interest in demand response as embedded in our end-to-end problem.

aspect, in our problem, our aggregator simultaneously derives benefit from a whole *population* of DR participants who have stochastic and possibly correlated valuations, rather than drawing a single agent from a common distribution. But, as reflects our application area, DR participants are “too small” to affect each other, so they are not strategically relevant to each other. These facts reduce our problem to a variation on a textbook screening problem, resulting in an expression for the optimal payment for a given curtailment policy, in Proposition 2.10 (“revenue equivalence”). This makes it straightforward to embed the contracting decision in the wholesale market problem.

For technical reasons, we also assume that the distribution $F(\theta|\tau, \omega_{RT})$ has constant support for all τ, ω_{RT} .

Finally, we make a very plausible simplifying assumption, that the highest type, $\bar{\tau}$, is “very high”¹⁷:

Assumption 3 For any optimal curtailment policy $\hat{\tau}^*(\cdot), \bar{\tau} > \hat{\tau}^*(\omega_{RT}), \forall \omega_{RT} \in \Omega_{RT}$.

2.2 Direct mechanisms

We assume that ex ante, DR increments are anonymous: the only distinguishing feature of each one is its type. Therefore, the result of contracting will be that, one way or another, each capacity increment is assigned a curtailment status, contingent only on its type, and the real-time dispatch outcome ω_{RT} ; this curtailment status function is its “allocation,” which we denote as $k(\tau, \omega_{RT})$. The contracting outcome will also involve a payment from the aggregator to the increment. Because both the aggregator and the increments are risk-neutral and we assume that the aggregator is capable of commitment, it makes no difference whether the payment is made ex post, or whether the same payment is made ex ante as an expectation over the ex post payments.

There might be many mechanisms by which such a contracting outcome could come about, but the Revelation Principle shows that any equilibrium contracting outcome can be achieved by a “direct revelation mechanism.”¹⁸

Definition 2.2 A direct mechanism in our setting is a pair of functions

$$[\underline{\tau}, \bar{\tau}] \ni \tau \mapsto k(\tau, \cdot) \in [0, 1]^{\Omega_{RT}}$$

(where $[0, 1]^{\Omega_{RT}}$ is the set of functions having domain Ω_{RT} and codomain $[0, 1]$) and

$$[\underline{\tau}, \bar{\tau}] \ni \tau \mapsto t(\tau) \in \mathbb{R}.$$

Here $k(\tilde{\tau}, \cdot)$ is the real-time dispatch curtailment function, assigned ex ante to an increment reporting type $\tilde{\tau}$, and $t(\tilde{\tau})$ is the ex ante or expected gross payment made to an increment reporting $\tilde{\tau}$.

¹⁷ We just mean that there are some units of demand whose valuation for power is high enough that they would not accept a reasonable payment curtailment in any contingency. For example, hospitals with life support units may have some quantity level at which their demand is, for all intents and purposes, perfectly inelastic.

¹⁸ Here we closely follow the development of Börgers (2010, Chap. 2).

The interpretation is that the increment of type τ gives a “report” of its type, $\tilde{\tau} \in [\underline{\tau}, \bar{\tau}]$, and the aggregator commits to effect the corresponding allocation function and payment, $(k(\tilde{\tau}, \cdot), t(\tilde{\tau}))$. We call $t(\tilde{\tau})$ the “gross ex ante payment” to an increment reporting type $\tilde{\tau}$. It is “gross” because the aggregator may also collect a penalty from the increment ex post if it violates the curtailment signal, and we will net the latter quantity out of the aggregate payment.

2.2.1 The increments' utility function in a direct mechanism

Here we display the increments' utility function in a direct mechanism in a manner that reflects our demand response setting. In the following subsection, we show that any contracting equilibrium that allocates a curtailment function and an expected payment can be implemented by a direct mechanism—in particular, a “direct revelation mechanism.”

A capacity increment of infinitesimal magnitude dx , with ex post valuation θ , pays the variable charge Rdx or Hdx (depending on its curtailment status) and enjoys utility θdx from consuming the quantity dx MWh. Henceforth, we normalize the an increment quantities produces, pays or enjoys—DR, tariffs or expected utility—by dividing them by dx ; resulting in the corresponding quantity “per unit mass” of increments.

A priori, in its “outside option”, each increment has a retail service contract that permits it to consume if it pays the retail rate R \$/MWh. This service contract is an option, in the financial sense, for physical delivery of a commodity at the point of service. Before the institution of a DR policy, this option only has value to its owner, since the commodity cannot be transferred and thus enjoyed by anyone else. But a DR policy establishes (in our model), by administrative fiat, that if the aggregator prevents an increment from exercising its service option when it otherwise would have, then the resulting reduction in consumption is treated as *production* of that same quantity by the aggregator.

An increment's expected utility from holding its original service option is its “option value.” This is its consumption utility net of the retail price, provided that it is positive. Through contracting, the aggregator purchases the right to send a curtailment signal to each increment in every event where the reported type $\tilde{\tau} \leq \hat{\tau}(\omega_{RT})$. The curtailment signal penalizes consumption by raising the effective price (exercise price) of the increment's service option from the retail rate, R , to an exogenously determined penalty rate, $H > R$. This reduces both the option value and the quantity of consumption. Generally, we assume that the aggregator collects the difference or “penalty fee” $H - R$ if the increment consumes despite receiving the signal, but we are particularly interested in the special case of direct load control, which is modeled by $H = \infty$. In that case a penalty fee is never collected, because the increment always complies with the curtailment signal.

We quantify the increment's utility when under contract as net of the outside option value:

Definition 2.3 (*Net option value from curtailment*) Denote an increment's change in ex post value given curtailment as

$$L(\theta) \triangleq ((\theta - H)^+ - (\theta - R)^+) \leq 0. \tag{2}$$

The lost option value for type τ , from perspective of real time event ω_{RT} conditional on curtailment, is

$$z(\tau, \omega_{RT}) \triangleq \int_{\Theta} L(\theta) f(\theta|\tau, \omega_{RT}) d\theta \leq 0. \tag{3}$$

The net option value of an increment of type τ , reporting type $\tilde{\tau}$, given curtailment policy $k(\cdot, \cdot)$, is

$$U(\tilde{\tau}|\tau) \triangleq \int_{\Omega_{RT}} z(\tau, \omega) k(\tilde{\tau}, \omega) dP_{RT}(\omega). \tag{4}$$

We assume that increments have quasi-linear utility, so that it affords the following decomposition:

Definition 2.4 (*Direct mechanism representation of the increments' expected utility function*) An increment that reports its type as $\hat{\tau}$, given that it has true type τ , enjoys expected net utility

$$u(\tilde{\tau}|\tau) \triangleq U(\tilde{\tau}|\tau) + t(\tilde{\tau}). \tag{5}$$

2.3 The revelation principle

The Revelation Principle establishes that without loss of generality, we can restrict attention to “direct revelation mechanisms,” in which the increment reports its *true* type to the aggregator’s direct mechanism, and in which the incentive compatibility constraint is satisfied.

We quote Proposition 2.1 of [Börger \(2010\)](#) with minor substitutions, the proof of which can be found there:

Proposition 2.5 (Revelation principle) *For every mechanism Γ and every optimal increment strategy σ in Γ , there is a direct mechanism Γ' and an optimal buyer strategy σ' in Γ' such that*

1. *The strategy σ' satisfies:*

$$\sigma'(\tau) = \tau \text{ for every } \tau \in [\underline{\tau}, \bar{\tau}],$$

i.e. σ' prescribes telling the truth;

2. *For every $\tau \in [\underline{\tau}, \bar{\tau}]$, the curtailment allocation $k(\tau, \cdot)$ and the payment $t(\tau)$ equal the allocation function and the expected payment that result in Γ if the buyer plays her optimal strategy σ .*

Definition 2.6 (*Incentive compatibility*) A direct mechanism in our problem is incentive compatible if truth-telling is optimal for every type; that is if:

$$\tau \in \arg \max_{\tilde{\tau} \in [\underline{\tau}, \bar{\tau}]} \{U(\tilde{\tau}|\tau) + t(\tilde{\tau})\}.$$

The previous result and definition allow us to simplify our problem, by restricting attention to direct revelation mechanisms which satisfy the (IC) constraint.¹⁹ An increment participating in an incentive compatible mechanism derives its *equilibrium utility*:

Definition 2.7 (*Equilibrium utility*) For a given direct revelation mechanism, an increment of type τ enjoys net expected utility

$$u(\tau) \triangleq u(\tau|\tau) = U(\tau|\tau) + t(\tau)$$

in the contracting equilibrium, i.e., when it truthfully reports its type in a direct mechanism.

Another crucial constraint results from our assumption that the increments participate in the contracting scheme voluntarily:

Assumption 4 (*Individual rationality*)

$$u(\tau) \geq 0 \quad \forall \tau \in [\underline{\tau}, \bar{\tau}].$$

This is to say that contracting cannot leave the increment worse off than in its outside option, which we normalize to zero. (Remember that $U(\tilde{\tau}|\tau)$ is the *change* in consumption utility, and penalties, resulting from contracting.) Further, we assume that all increments participate, without loss of generality: for suppose an increment found it better not to participate, so that it is never curtailed and generates no DR for the aggregator, and receives no payment from the aggregator. This could be equivalently represented by $k(\tau, \cdot) \equiv 0$, and $t(\tau) = 0$. In our model, this allocation makes it so the increment makes no contribution to the aggregator's objective, and by definition of the outside option, the increment's net utility is also zero. So non-participation can be modeled as the degenerate form of participation just mentioned, and without loss of generality we can enforce the assumption that all increments participate and achieve nonnegative net utility.

2.4 The optimal payment T needed to effect a demand response policy

We are now ready to derive an expression for the aggregate payment T as a function of the curtailment policy.

¹⁹ We should note, however, that the Revelation Principle assumes away problems of multiple equilibria. However, this is not a problem for us, because we have already restricted the curtailment policy set in Assumption 2.

Lemma 2.8 (Continuity and differentiability of equilibrium consumption utility with respect to true type) *The increment's utility function, as a function of its true type while holding its report constant, is differentiable, Lipschitz continuous, and thus absolutely continuous, for any given report $\tilde{\tau}$. Therefore there exists an integrable function $b(\tau)$ such that $|\frac{\partial}{\partial \tau} U(\tilde{\tau}|\tau)| \leq b(\tau)$, and $U(\tilde{\tau}|\cdot)$ is uniformly continuous, $\forall \tilde{\tau}$.*

Proof See Appendix. □

At this point it would be standard to state a necessary condition that incentive compatibility places on the curtailment allocation $k(\cdot, \cdot)$. However, since we are restricting attention to the class of cutoff policies $\{\hat{\tau}(\cdot)\}$, we will skip this, and show in Proposition 2.12 that all cutoff policies can be implemented, given the appropriate payment.

Lemma 2.9 *Incentive compatibility implies that the equilibrium net expected utility $u(\tau)$ is decreasing in ex ante type, that $u(\tau)$ is absolutely continuous, and that*

$$u(\tau) = u(\bar{\tau}) - \int_{\tau}^{\bar{\tau}} \frac{\partial}{\partial s} U(\tilde{\tau}|s) \Big|_{\tilde{\tau}=s} ds \tag{6}$$

$$= u(\bar{\tau}) + \int_{\tau}^{\bar{\tau}} \int_{\Omega} \int_{\Theta} \underbrace{L'(\theta)}_{\leq 0} \underbrace{\frac{\partial F(\theta|s, \omega)}{\partial s}}_{\leq 0} \underbrace{k(s, \omega)}_{\geq 0} dP(\omega) ds \tag{7}$$

$$= u(\bar{\tau}) - \int_{\tau}^{\bar{\tau}} \int_{\Omega} \int_{\Theta} L(\theta) \frac{\partial f(\theta|s, \omega)}{\partial s} k(s, \omega) dP(\omega) ds . \tag{8}$$

Proof The first line follows from the envelope theorem: [Milgrom and Segal \(2002\)](#), Theorem 2, the conditions of which we have established in Lemma 2.8. Since the integrand in line (7) is nonnegative, and τ is the lower limit of integration, $u(\tau)$ is nonincreasing. □

Proposition 2.10 (Necessary condition on the payment for incentive compatibility/revenue equivalance) *Under a merit-order curtailment policy, incentive compatibility requires that the ex ante gross payment to an increment of type τ be*

$$t(\tau) = u(\bar{\tau}) - \int_{\Omega_{RT}} z(\hat{\tau}(\omega), \omega) k(\tau, \omega) dP_{RT}(\omega). \tag{9}$$

Proof See Appendix. □

The next proposition shows that when the aggregator optimizes its profit, $u(\bar{\tau})$ is zero. So Eq. (9) can be interpreted as saying that, for each real time state, every curtailed increment is paid the reservation utility of the highest curtailed increment in that state.

Proposition 2.11 (Individual rationality “binds at the top”) *Given incentive compatibility, individual rationality holds only if $u(\bar{\tau}) \geq 0$. Maximization of the aggregator's profit implies that $u(\bar{\tau}) = 0$.*

Proof Proposition 2.9 implies that $u(\tau)$ is nonincreasing. Since the aggregator's payment to each increment includes the constant term $u(\bar{\tau}) \geq 0$, the aggregator maximizes profit by reducing $t(\bar{\tau})$ until $u(\bar{\tau}) = 0$. \square

Proposition 2.12 (Sufficient condition for incentive compatibility) *Any merit-order curtailment policy is rendered incentive compatible, when the corresponding payment is as in Proposition 2.10.*

Proof See Appendix. \square

This proposition shows that the DR contracting problem is separable over Ω_{RT} . In the next section we consider the implications of this result.

2.4.1 State-contingent interpretation of the payment form

The formula in Proposition 2.10 for the payment (setting $u(\bar{\tau}) = 0$), has an intuitive interpretation, corresponding to what we call the “state-contingent representation.” Proposition 2.12 shows that under First Order Stochastic Dominance (Assumption 1), merit order curtailment guarantees the satisfaction of any linking constraints across random events ω_{RT} that are implied by incentive compatibility, so that those linking constraints can be discarded. The aggregator's DR contracting problem is therefore separable across realized RT states ω_{RT} , and we can think of the aggregator as committing to separate DR purchases contingent on each ω_{RT} . In the merit order setting, the curtailment policy for each fixed real time outcome is binary and decreasing over the type argument: i.e., it is constant over all curtailed types. This is because incentive compatibility forces the aggregator to pay each curtailed increment in a particular state the reservation value of the highest curtailed type in that state, since any increment could impersonate that type.²⁰

The formula for the payment in Proposition 2.10 is in fact the ex ante expectation of the payments that would be made if curtailment contracting were performed in this state-contingent manner.

2.5 Expressing the net payment as a linear functional of the curtailment policy

In order to solve the end-to-end problem, we need expressions of a certain form for the *aggregate* payment, i.e. integrated over the DR population.

Definition 2.13 (*Payment*) The aggregate net payment is the integral of the type-specific payment over the increment population, net of penalty receipts:

$$T \triangleq \int_{\underline{\tau}}^{\bar{\tau}} t(\tau)g(\tau)d\tau - \underbrace{(H - R) \int_{\Omega_{EP}} \int_{\underline{\tau}}^{\bar{\tau}} \mathbb{1}\{\tau \leq \hat{\tau}(\omega_{RT})\} \mathbb{1}\{\theta(\tau, \omega_{EP}) > H\}g(\tau)d\tau}_{\text{aggregate expected penalty fee receipts}}$$

²⁰ See the discussion at the end of Börgers (2010, Chap. 2.2).

Proposition 2.14 (Expressing the aggregate payment as an expectation of a linear functional of the increment population) *Defining the virtual net payment to type τ in event ω_{RT} as*

$$\psi(\tau, \omega_{RT}) \triangleq - \left(z(\tau, \omega_{RT}) + \frac{G(\tau)}{g(\tau)} \frac{\partial}{\partial \tau} z(\tau, \omega_{RT}) \right) - \underbrace{(H - R) \Pr\{\theta(\tau, \omega_{EP}) \geq H | \omega_{RT}\}}_{\text{expected penalty fee from curtailing type } \tau}$$

The aggregate payment can be expressed as

$$T = \int_{\Omega_{RT}} \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \psi(\tau, \omega_{RT}) g(\tau) d\tau dP(\omega_{RT}). \tag{10}$$

Proof See Appendix. ²¹ □

We call $\psi(\tau, \omega_{RT})$ the “virtual net payment” to type τ for curtailment in event ω_{RT} . In addition to penalty fee receipts, the marginal aggregate payment from raising the marginally curtailed type to τ in event ω_{RT} has two components: the aggregator must pay the marginal segment of the type population its lost utility $z(\tau, \omega_{RT})g(\tau)$; and it must also raise the payment it makes to infra-marginal types to that same level, incurring a payment $G(\tau) \frac{\partial}{\partial \tau} z(\tau, \omega_{RT})$, because the infra-marginal types, the measure of which is $G(\tau)$, can impersonate the marginal type. This second term corresponds to the information rent in the mechanism design literature. The expression $\psi(\tau, \omega_{RT})$ attributes both of these components to the marginal type, so that we obtain the marginal change in the *aggregate* payment from recruiting type τ . The same economic insight arises when we interpret the product rule in the first order condition for the elementary monopsony pricing problem, as we discuss in the next section. This allows us to express the first order conditions for the end-to-end problem.

2.6 State-contingent monopsony procurement of DR, and competitive extension

As we have just seen that the contracting problem is separable across RT scenarios, we will focus on single scenario ω_{RT} , and suppress the notation for it. Examining the formula for $\psi(\tau)$ in Proposition 2.14, we note that since first order stochastic dominance guarantees that $\partial/\partial \tau z(\tau) \leq 0$, and since $\frac{G(\tau)}{g(\tau)} > 0$, the information rent makes the virtual payment greater than the marginal benefited increment’s lost utility. Assume that the aggregator has a marginal benefit from procuring DR, equal to society’s benefit, which is nonincreasing in the quantity curtailed, and thus nonincreasing in the marginal type. Then we have the standard monopsony distortion: the aggregator will purchase the quantity that sets its marginal *expenditure*, $\psi(\hat{\tau})$, equal to its marginal

²¹ It turns out that this formula holds for general $k(\cdot, \cdot)$ satisfying incentive compatibility (i.e. we would integrate the above expression over the whole population, multiplying the integrand by k), not just for the cutoff form $\hat{\tau}$, but this is not an immediate concern of ours here, so we leave this result unproved.

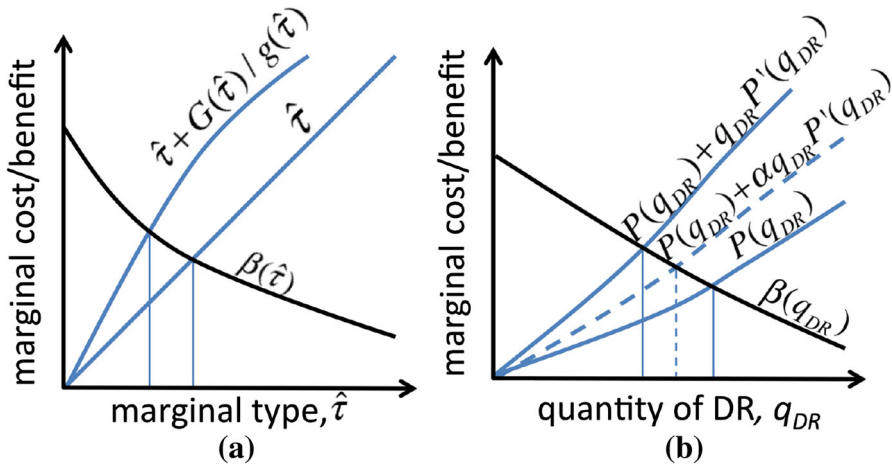


Fig. 1 The monopsony and oligopsony purchase problem

benefit, rather than the quantity that sets the marginal social cost (i.e. the marginal increment's lost consumption utility) equal to its marginal benefit.

In keeping with the direct mechanism representation of contracting, we mostly focus on the marginal increment's type as the decision variable in this paper. But to develop economic intuition and suggest an extension to Cournot competition between aggregators, we consider a special case, and then a change of variables that maps the marginal type to the corresponding DR procurement quantity. First, for the purpose of illustration, we parameterize the type variable (and its distribution), so that $\tau = -z(\tau)$, the (positive) lost option value in the RT scenario under consideration. Then $\frac{\partial}{\partial \tau} z(\tau) = -1$, and $\psi(\tau) = \tau + \frac{G(\tau)}{g(\tau)}$.²² Also, for simplicity, assume that the DR yield is perfect: dispatching a unit of increments produces a unit of DR. The aggregator's marginal analysis on the type domain is portrayed in Fig. 1a, where the benefit as a function of marginal curtailed type is $\beta(\hat{\tau})$.

Next we display the corresponding curves, but parameterize them as functions of the quantity of DR procured, rather than the marginal curtailed type. So the x -axis is $q_{DR} = G(\hat{\tau})$. We denote $G^{-1}(\cdot)$ as $P(\cdot)$, so that $\hat{\tau} = P(q_{DR})$: this is the state-contingent price that the aggregator must offer to procure the quantity q_{DR} of DR. The virtual payment, expressed as a function of DR quantity q_{DR} , is $\hat{\tau} + \frac{G(\hat{\tau})}{g(\hat{\tau})} = P(q_{DR}) + q_{DR} P'(q_{DR})$. (This is because $P'(q) = \frac{d}{dq} G^{-1}(q) = \frac{1}{g(G^{-1}(q))} = \frac{1}{g(\hat{\tau})}$.) Of course this is the marginal expenditure as a function of quantity, i.e., $\frac{d}{dq} (qP(q))$. So we see (Fig. 1b) that this instance of our problem is the same as the elementary monopsony pricing problem. It is the aggregator's effect on the price of DR, in the second term, that causes it to purchase less than the efficient level.

The effect of competition among aggregators in DR procurement can be captured as in classical oligopoly or oligopsony models by scaling the second term, i.e. the information rent as a function of q_{DR} , by $0 < \alpha < 1$. We illustrate this with the

²² We assume that ψ is monotonically increasing in this representation, as we also do below when we solve the end-to-end problem.

dashed line in Fig. 1b. In the case of a symmetric state-contingent Cournot oligosony, it can be shown that $\alpha = 1/n$, where n is the number of competitors (Varian 1992, Chap. 16). As n increases, the distortion approaches zero, and the equilibrium procurement approaches the efficient level. In order to combine state-contingent procurement contracts offered by a competitive aggregator into a single ex-ante contract offer that specifies expected payment versus curtailment probability, we may need to assume that the competition for contracting takes place up-front while contracts are exclusive. Such an assumption can be justified by the need for specialized aggregator-owned technology for managing load curtailments.

In most of our subsequent presentation of the end-to-end model, we maintain the DR monopsony assumption, making informal remarks about the possible effects of competition among aggregators in Sect. 3.3 and the conclusion.

3 Analyzing the end-to-end problem

In the previous section, we derived useful expressions for the *cost* of recruiting DR. To maximize its profit, the aggregator balances costs against benefits. Having determined the cost of DR as a function of the policy, we now analyze the aggregator’s problem as a two-stage problem: first, we characterize the optimal dispatch of DR, conditional on an arbitrary DA offer policy q . Then, holding the DR policy at its optimal setting as a function of q , we optimize q .

3.1 Benefit from DR dispatch

Recall the informal sketch of the objective from Eq. (1):

$$\begin{aligned} & \max_{q, DR, T} J(q, DR, T) \\ &= \max_{q, DR, T} \mathbb{E} \left[\underbrace{p q}_{\substack{\uparrow \\ \text{day-ahead revenue}}} + \underbrace{a(DR + s - q)^+}_{\text{overproduction}} - \underbrace{b(q - DR - s)^+}_{\text{shortfall}} \right] - \underbrace{T}_{\substack{\uparrow \\ \text{payment to DR}}} \end{aligned}$$

We now give the definition of the DR production quantity:

Definition 3.1 (*DR quantity, $DR(\hat{\tau}(\cdot), \omega_{EP})$*) The quantity of DR in ex post event ω_{EP} is the measure under g of increments that forego consumption as a result of receiving the curtailment signal:

$$DR(\hat{\tau}(\omega_{RT}); \omega_{EP}) \triangleq \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \mathbb{1}\{\tau : R \leq \theta(\tau, \omega_{EP}) \leq H\} g(\tau) d\tau$$

We assume that the valuation process allows the application of Fubini’s theorem to the DR process:

Assumption 5 (*Fubini property of ex post DR process*) We assume that for all events $\omega_{RT} \in \Omega_{RT}$, and all subsets $\mathcal{B} \subset [\underline{\tau}, \bar{\tau}]$,

$$\begin{aligned} & \mathbb{E}_{EP} \left[\int_{\mathcal{B}} \mathbb{1}\{\theta(\tau, \omega_{EP}) \in [R, H]\} g(\tau) d\tau \mid \omega_{RT} \right] \\ &= \int_{\mathcal{B}} \Pr_{EP}\{\theta(\tau, \omega_{EP}) \in [R, H]\} | \omega_{RT} \} g(\tau) d\tau \\ &\triangleq \int_{\mathcal{B}} y(\tau, \omega_{RT}) g(\tau) d\tau. \end{aligned}$$

(This equals $DR(\hat{\tau})$ when $\mathcal{B} = [\underline{\tau}, \hat{\tau}]$.) Here we define $y(\tau, \omega_{RT})$ as the “expected DR yield” of type τ conditional on event ω_{RT} . Also, recall that the DA offer policy is a function of the day-ahead information: $q : \Omega_{DA} \rightarrow \mathbb{R}$. We will analyze the aggregator’s end-to-end problem in two special cases.

3.2 Example 1: Purely idiosyncratic valuation shocks

In this section, we assume that conditional on the real-time outcome, uncertainty regarding the ex post valuation process is i.i.d. noise.

Assumption 6 The valuation process decomposes as

$$\theta(\tau, \omega_{EP}) = m(\tau, \omega_{RT}) + \epsilon(\tau, \omega_{EP})$$

where $m(\tau, \omega_{RT})$ is a deterministic function, $\epsilon(\tau, \omega_{EP})$ is i.i.d. conditional on ω_{RT} over τ , and $\mathbb{E}[\epsilon(\tau, \omega_{EP})] = \mathbb{E}[\epsilon(\tau, \omega_{EP}) | \omega_{RT}] = 0, \forall \omega_{RT}, \omega_{EP}$. Each $\epsilon(\tau, \cdot)$ has conditional cdf $\Phi(\epsilon; \omega_{RT})$, and conditional pdf $\varphi(\epsilon; \omega_{RT})$.

Remark 3.2 Assumptions 1 (FOSD) and 6 (idiosyncratic noise) jointly imply that $m(\cdot, \omega_{RT})$ is increasing, for each ω_{RT} . The common conditional cdf of valuation shocks can be written as $\Pr(\theta \leq z | \tau, \omega_{RT}) = \Phi(z - m(\tau, \omega) | \omega_{RT})$. This determines the virtual payment function $\psi(\tau, \omega_{RT})$, the formula for which we omit. Further, our assumption of constant support for $\{\theta_\tau\}_\tau$ implies that each θ_τ has full support on \mathbb{R} , for every ω_{RT} .

We make the following assertion without proof:

Remark 3.3 Under Assumption 6, the DR output is almost surely deterministic, conditional on $\hat{\tau}(\omega_{RT})$:

$$\begin{aligned} DR(\hat{\tau}(\omega_{RT}); \omega_{EP}) &= \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \mathbb{1}\{\theta(\tau, \omega_{EP}) \in [R, H]\}g(\tau)d\tau \\ &\stackrel{a.s.}{=} \mathbb{E}_{EP} \left[\int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \mathbb{1}\{\theta(\tau, \omega_{EP}) \in [R, H]\}g(\tau)d\tau \middle| \omega_{RT} \right] \\ &= \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \Pr_{EP}\{\theta(\tau, \omega_{EP}) \in [R, H]\}g(\tau)d\tau \\ &= \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} y(\tau, \omega_{RT})g(\tau)d\tau. \end{aligned}$$

See Al-Najjar (1995) for further discussion on this issue. Our purpose in making Assumption 6 was simply to license the above result, so the reader may take this consequence as the operative assumption instead.

Having derived expressions for the DR production quantity and the payment contribution in each RT outcome ω_{RT} , we can now plug them into the stylized objective of Eq. (1). Note that the real-time contribution to the ex ante component is the integrand from Eq. (10).

Definition 3.4 (*Aggregator’s objective, and its constituent parts*) The real-time contribution to the aggregator’s ex ante objective in event ω_{RT} is

$$\begin{aligned} J_{RT}(q, \hat{\tau}(\omega_{RT}); \omega_{RT}) &\triangleq pq + a \left(\int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} y(\tau, \omega_{RT})g(\tau)d\tau + s - q \right)^+ \\ &\quad - b \left(q - \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} y(\tau, \omega_{RT})g(\tau)d\tau - s \right)^+ \\ &\quad - \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \psi(\tau, \omega_{RT})g(\tau)d\tau; \end{aligned} \tag{11}$$

The day-ahead contribution is the expectation over the real time contribution:

$$J_{DA}(q, \hat{\tau}(\cdot); \omega_{DA}) \triangleq \mathbb{E}_{\omega_{RT}|\omega_{DA}} [J_{RT}(q, \hat{\tau}(\omega_{RT}); \omega_{RT})|\omega_{DA}]; \tag{12}$$

And the aggregator’s objective, full stop, is the expectation over the DA contribution:

$$J_{EA}(q(\cdot), \hat{\tau}(\cdot)) \triangleq \mathbb{E}_{\omega_{DA}} [J_{DA}(q(\omega_{DA}), \hat{\tau}(\cdot); \omega_{DA})]. \tag{13}$$

Approaching the problem as a two-stage decision problem with recourse, we first derive a first order necessary condition for optimizing Eq. (11) with respect to the optimal DR dispatch $\hat{\tau}(\omega_{RT})$ given a DA offer q . We denote this optimal recourse

policy as $\hat{\tau}^*(\omega_{RT}|q)$. Since we will optimize q below, we also occasionally drop the “ q ”, or “conditioned on q ” argument for brevity.

3.2.1 First order necessary conditions for $\hat{\tau}^*(\omega_{RT}|q)$ in Example 1

To make the first order conditions and consequent results easier to read, we define the following quantities:

Definition 3.5 Under Assumption 6 (idiosyncratic noise), we define the marginal type that must be curtailed to cancel out a nonnegative RT imbalance, given wind realization s , as

$$DR^{-1}(q - s; \omega_{RT}) \triangleq \begin{cases} -\infty & \text{if } q < s \\ \tau & \text{s.t. } DR(\tau, \omega_{RT}) = q - s, \text{ if } DR(\bar{\tau}, \omega_{RT}) \geq q - s \geq 0 \\ \infty & \text{if } q - s \geq DR(\bar{\tau}, \omega_{RT}). \end{cases}$$

Definition 3.6 The marginal cost of DR from type τ (i.e. the dollar amount the aggregator must pay to curtail increments of type τ , per unit of resulting DR yield), is $c(\tau; \omega_{RT}) \triangleq \psi(\tau, \omega_{RT})/y(\tau, \omega_{RT})$.

We will occasionally suppress ω_{RT} arguments for DR , c , etc. for conciseness.

Proposition 3.7 (Optimal DR curtailment policy) *The first order necessary condition requires that if $\hat{\tau}^*(\omega_{RT}) > \underline{\tau}$, the following condition holds (here $\partial(\cdot)$ denotes the subgradient mapping):*

$$\psi(\hat{\tau}^*)g(\hat{\tau}^*) \in \partial \left(a \left(\int y(\tau; \omega_{RT}) \mathbb{1}_{\tau \leq \hat{\tau}^*} g(\tau) d\tau + s - q \right)^+ - b \left(q - s - \int y(\tau, \omega_{RT}) \mathbb{1}_{\tau \leq \hat{\tau}^*} g(\tau) d\tau \right)^+ \right)$$

which implies

$$c(\hat{\tau}^*) \in \{a\} \mathbb{1}_{DR(\hat{\tau}^*)+s > q} + \{b\} \mathbb{1}_{DR(\hat{\tau}^*)+s < q} + [a, b] \mathbb{1}_{DR(\hat{\tau}^*)+s = q}.$$

If $c(\cdot)$ is monotonically increasing, this requires that (suppressing the ω_{RT} argument of DR^{-1})

$$c(\hat{\tau}^*(\omega_{RT})) = \begin{cases} c(DR^{-1}(q - s)) & \text{if } c(DR^{-1}(q - s)) \in (a, b), \\ a & \text{if } c(DR^{-1}(q - s)) < a, \text{ and} \\ b & \text{if } c(DR^{-1}(q - s)) > b. \end{cases} \quad (14)$$

Proof Subdifferentiation of the previous formula for the aggregator’s objective with respect to $\hat{\tau}^*$. □

Assuming that c , the marginal cost of DR, is monotonic, in the optimal DR recourse policy $\hat{\tau}^*(\omega_{RT}|q)$, the aggregator exactly meets its commitment if the marginal cost of DR needed to do so is strictly between the two imbalance prices, a and b .²³ Otherwise, if the marginal cost of DR needed to do so would be greater than b , the aggregator curtails all increments up to the “upper critical valuation,” $c^{-1}(b)$. If the marginal cost of curtailing up to the point of zero imbalance is below a , then the aggregator produces *more* than it offered day ahead, curtailing all increments up to the “lower critical type,” $c^{-1}(a)$. Increments with type below the minimum lower critical type, $\tau < \inf_{\omega_{RT}} c^{-1}(a(\omega_{RT}); \omega_{RT})$ are curtailed in all wind outcomes; increments with valuation above the maximum upper critical type are not curtailed in any wind outcome. These features of the real time curtailment decision are depicted in the upper left panel of Fig. 2, although that figure is further specialized to Example 2, introduced below. That panel is essentially the same as Fig. 1b, but with a piecewise constant benefit function, reflecting our aggregator’s newsvendor-style revenue function.

Decomposing events by imbalance status, the first order conditions for monotonic c admit the following formulation,²⁴ which is also helpful for the optimization of q .

$$\hat{\tau}^*(\omega_{RT}|q) = \begin{cases} DR^{-1}(q - s) & \text{if } DR(\hat{\tau}^*(\omega_{RTD})) = q - s, \\ c^{-1}(b) & \text{if } DR(\hat{\tau}^*(\omega_{RTD})) < q - s, \text{ and} \\ c^{-1}(a) & \text{if } DR(\hat{\tau}^*(\omega_{RTD})) > q - s. \end{cases} \quad (15)$$

3.2.2 The optimal day-ahead offer policy, q^*

To find the optimal q^* , we first plug the expression for the optimal DR dispatch in Eq. (15) into the aggregator’s objective, conditional on ω_{DA} :

$$\begin{aligned} J_{DA}(q, \hat{\tau}^*(\cdot|q); \omega_{DA}) &= \sup_{\hat{\tau}(\cdot)} J_{DA}(q, \hat{\tau}(\cdot); \omega_{DA}) \\ &= pq + \mathbb{E}_{\omega_{RT}|\omega_{DA}} \left[a \left(\int_{\underline{\tau}}^{\hat{\tau}^*(\omega_{RT}|q)} y(\tau)g(\tau)d\tau + s - q \right)^+ \right] \\ &\quad - \mathbb{E}_{\omega_{RT}|\omega_{DA}} \left[b \left(q - s - \int_{\underline{\tau}}^{\hat{\tau}^*(\omega_{RT}|q)} y(\tau)g(\tau)d\tau \right)^+ \right] \\ &\quad - \mathbb{E}_{\omega_{RT}|\omega_{DA}} \int_{\underline{\tau}}^{\hat{\tau}^*(\omega_{RT}|q)} \psi(\tau; \omega_{RT})g(\tau)d\tau. \end{aligned} \quad (16)$$

We obtain the first order condition by differentiation:

²³ We will argue in a followup work that the cost c should typically be monotonic. Another setting perhaps worth considering is for c “single-troughed,” i.e. quasi-convex, so that $-c$ is single-peaked. This might obtain if it is prohibitively expensive to obtain DR from types low valuation types, because their yield is too low. We may consider the optimality conditions for this case elsewhere.

²⁴ If c is weakly increasing, then the equality is replaced with \in , and c^{-1} is interpreted as the set-valued preimage function.

Proposition 3.8 (First order necessary condition for optimal day-ahead offer) *The optimal day-ahead offer, q^* , conditional on ω_{DA} , satisfies the following condition:*

$$\begin{aligned}
 p &= \mathbb{E}[a|DR(\hat{\tau}^*) + s > q^*] \Pr\{DR(\hat{\tau}^*) + s > q^*\} \\
 &\quad + \mathbb{E}[b|DR(\hat{\tau}^*) + s < q^*] \Pr\{DR(\hat{\tau}^*) + s < q^*\} \\
 &\quad + \mathbb{E}[c(\hat{\tau}^*)|DR(\hat{\tau}^*) + s = q] \Pr\{DR(\hat{\tau}^*) + s = q^*\}.
 \end{aligned}$$

Proof See Appendix. □

This is not an explicit solution; we present this condition because of its clear economic interpretation: The aggregator increases its DA offer quantity until the marginal change in expected real-time expenditures (imbalance prices plus payments to DR) rises to meet the marginal DA revenue. If the RT imbalance prices a and b are known day-ahead, then Fig. 2 and the accompanying discussion in Sect. 3.3.1 apply to this case. That is, we show in that section how the aggregator’s problem can be viewed as an elaboration of the elementary monopsony pricing problem.

3.3 Example 2: Parameterized uniform distributions

Now we consider a simple concrete example, where renewable power output and increment types are distributed uniformly, valuation noise is degenerate (i.e. nonexistent), the imbalance prices are known day-ahead, the lowest valuation is equal to the retail rate, and the aggregator employs direct load control. In this case, we can derive formulas for the aggregator’s optimal policy and its relevant features, in order to display the solution graphically, and obtain quantitative sensitivity results. The main objects in the model are

$$\begin{aligned}
 g(\tau) &= \frac{d}{d\tau} G(\tau) = N \mathbb{1}_{\tau \in [\underline{\tau}, \bar{\tau}]} \\
 s &\sim \text{Uniform}[0, \bar{s}] \\
 \theta_\tau &\equiv \tau \text{ (degenerate distribution at } \tau) \\
 R &= \underline{\tau} \\
 H &= \infty.
 \end{aligned} \tag{17}$$

With no valuation noise, this model does not satisfy our Assumption 1 (FOSD). But we only needed Assumption 1 in order to prove Lemma 2.8 (monotonicity and differentiability of the equilibrium utility with respect to true type), which we can directly verify. In our current setting,

$$U(\tilde{\tau}|\tau) = - \int_{\Omega_{RT}} \tau k(\tilde{\tau}, \omega_{RT}) dP(\omega_{RT}) = -\tau \Pr\{\tilde{\tau} \leq \hat{\tau}(\omega_{RT})\}, \tag{18}$$

$$\therefore \frac{\partial}{\partial \tau} U(\tilde{\tau}|\tau) = - \Pr\{\tilde{\tau} \leq \hat{\tau}(\omega_{RT})\} \leq 0. \tag{19}$$

The remaining lemmas and propositions of Sect. 2 therefore follow. Before obtaining formulas for the payment, etc., we can simplify the model by reparameterization. First,

by a change of variables on τ , we normalize the valuations to express them as *net* of the retail rate, and set $\underline{\tau} = R = 0$.²⁵ This specification results in the following model quantities: $\forall \tau \in [\underline{\tau}, \bar{\tau}]$,

$$\begin{aligned} z(\tau) &= -\tau \\ \frac{\partial}{\partial \tau} z(\tau) &= -1 \\ G(\tau)/g(\tau) &= \tau \\ \psi(\tau) &= 2\tau \\ y(\tau) &= 1 \\ c(\tau) &= 2\tau. \end{aligned} \tag{20}$$

Instead of expressing the problem in terms of N , the DR population size, and \bar{s} , the VER nameplate capacity, we write the problem in terms of the parameter $v = N/(\bar{s}(\bar{\tau} - \underline{\tau}))$: the density of increment-valuations per dollar, per MW nameplate capacity. Now power (and energy) quantities are expressed as a fraction of the nameplate capacity: the DA offer quantity now takes the form $q \leftarrow q/\bar{s}$, the fraction of VER nameplate capacity bid day ahead; DR quantities are in the same units; the random variable s is reparameterized as $s \leftarrow s/\bar{s}$, now the wind realization's cdf value. Correspondingly, the aggregator's profit is now denominated in dollars per unit VER nameplate capacity, per hour.

Following the same two-stage solution method as before, we first consider the optimization of the real-time demand response dispatch. In this setting,

$$DR(\hat{\tau}) = v\hat{\tau}. \tag{21}$$

The aggregator's "real time objective" is:

$$\begin{aligned} J_{RT}(q, \hat{\tau}(\omega_{RT}); \omega_{RT}) &= p q + a(s + v\hat{\tau}(\omega_{RT}) - q)^+ \\ &\quad - b(q - s - v\hat{\tau}(\omega_{RT}))^+ - T(\omega_{RT}), \end{aligned} \tag{22}$$

where the last term the net ex post payment in event ω_{RT} , i.e. $\int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \psi(\tau, \omega_{RT})g(\tau)d\tau$, from Proposition 2.14:

$$T(\omega_{RT}) = v \int_0^{\hat{\tau}(\omega_{RT})} 2\tau d\tau \tag{23}$$

$$= v\hat{\tau}(\omega_{RT})^2. \tag{24}$$

The first order conditions for $\hat{\tau}(\omega_{RT}|q)$ give us that in the optimal recourse curtailment policy, Eq. (15) takes the form

²⁵ The derivation of model parameters from elasticity estimates must be done before the change of variables, because an elasticity is a ratio involving the retail price, which we are eliminating from the problem.

$$\hat{\tau}^*(\omega_{RT}|q) = \begin{cases} a/2 & \text{if } (q - s)/v < a/2 \Leftrightarrow s > q - va/2 \\ (q - s)/v & \text{if } a/2 < (q - s)/v < b/2 \\ & \Leftrightarrow q - vb/2 < s < q - va/2 \\ b/2 & \text{if } (q - s)/v > b/2 \Leftrightarrow s < q - vb/2. \end{cases} \quad (25)$$

That is, the lower and upper critical valuations mentioned above are $c^{-1}(a) = a/2$ and $c^{-1}(b) = b/2$ respectively. Referring back to Eq. (14), which applies here as well, we see that the optimal marginal DR cost is set equal to the imbalance cost that obtains given the resulting DR quantity. A marginal analysis of this optimization is depicted in the upper left panel of Fig. 2. Henceforth we assume that $\hat{\tau}$ is held at its optimal recourse value given q , and proceed to analyze the “first stage” problem, $\max_q J_{DA}(q, \hat{\tau}^*(\cdot|q); \omega_{DA})$.

The assumptions made for our symmetric oligopsony extension, in Sect. 2.6, hold in this case. So with n aggregators, then the distortion term becomes $\frac{G(\tau)}{ng(\tau)}$, and $c(\tau) = \frac{n+1}{n}\tau$ instead of 2τ , deflecting the marginal expenditure curve downward as the number of competitors increases.

For a given policy and $\omega_{DA} = (p, a, b)$, there are three possible types of event with respect to wind: shortfall, zero imbalance, and overproduction. To calculate the expected value of the objective over the real-time information s , we define the wind levels that constitute breakpoints between these regimes:

$$d \triangleq \min(1, \max(0, q - vb/2)) \quad (26)$$

$$e \triangleq \min(1, \max(0, q - va/2)) \quad (27)$$

Since s is now normalized to refer to the wind realization’s cdf value, d is the probability of shortfall, and the level e is one minus the probability of overproduction.²⁶ Plugging in the optimal $\hat{\tau}^*(\omega_{RT})$ and taking the expectation over s , we get

$$\begin{aligned} \max_q J_{DA}(q, \hat{\tau}^*(\cdot|q); \omega_{DA}) = \max_q & p q - \int_0^d \overbrace{v(b/2)^2 + b(q - vb/2 - s)}^{\text{payment}} ds \\ & - \int_d^e \overbrace{v\left(\frac{q-s}{v}\right)^2}^{\text{payment}} ds - \int_e^1 \overbrace{v(a/2)^2 - a(s + va/2 - q)}^{\text{payment}} ds. \end{aligned} \quad (28)$$

no imbalance overproduction

This “first-stage” (DA-stage) objective is concave (it is a day ahead expectation with respect to s of a concave real time benefit function), and piecewise polynomial in q , with breakpoints where d and e hit zero or one. The combinations of possible sets of events based on parameter values generate many cases.

²⁶ For example, d is the probability that wind is less than $q - vb/2$, which would imply that the wind quantity plus the maximum economical amount of DR be less than the offer quantity q .

Proposition 3.9 (Solution to the end-to-end problem, Example 2) *The solution to the aggregator's end-to-end problem in Example 2 is presented in Table 1.*

By “solution,” we mean that we exhibit the optimal policies q and $\hat{\tau}$, as a function of ω_{DA} and ω_{RT} respectively. If the aggregator is to publish a menu of contract choices ex ante, it would do this by taking expectations of the curtailment status, and payment, over Ω_{DA} , according to average market statistics. We provide an example of this in Sect. 3.3.2.

3.3.1 Graphical marginal analysis of the optimization of DR and q in Example 2

We can obtain more intuition regarding the optimization of the aggregator's policy by considering the graphical depiction of the marginal analysis of the aggregator's DR quantity decision in Fig. 2. Here we treat the DR procurement decision as a quantity decision, as described in Sect. 2.6. First we consider a single wind and price outcome, in the upper-left panel. With respect to DR, the aggregator is a monopsonist, i.e., the sole buyer in a commodity market with many sellers. In each real-time realization ω_{RT} , it faces the same marginal purchase cost curve (the marginal cost of DR, Definition 3.6) and it makes the optimal real-time curtailment decision by ensuring that its marginal cost of DR is between the imbalance prices that it faces on the margin. The marginal benefit from purchasing the q_{DR} th MW of DR is denoted by $\beta(q_{DR})$, a piecewise constant decreasing function depicted in black. The marginal cost from purchasing the q_{DR} th MW of DR is equal to the virtual payment to the marginally curtailed increment, which is increasing: $mc(q_{DR}) = 2q_{DR}/v$. The optimal curtailment quantity is the point on the quantity axis where the marginal benefit curve crosses the marginal cost curve. Projecting the intersection of the two curves onto the y axis, we obtain the optimal DR recourse cost associated with that real time outcome, which we depict with a circle on the y axis.

Next we step back to the DA optimization of q . At the DA stage, the aggregator foresees that a random RT outcome will realize, at which point it will take a DR recourse action in the manner just described. In the current example, the imbalance prices a and b are known day-ahead, and the only random variable at the DA stage is the wind, s . The distribution over wind outcomes (here with pdf h), together with the aggregator's choice of q , induces a distribution over breakpoints in the marginal benefit curve, $h(q - s)$, which we depict under the x axis of the last three panels of Fig. 2. Adjusting q slides the pdf of breakpoints along the x axis. We depict the distribution over marginal benefit curves induced by a choice of q as a regularly spaced finite sample from it (curves in light gray). Assuming optimal RT recourse, a choice of q induces a distribution over recourse costs, which we depict vertically on the y axis. This distribution has a density component corresponding to the virtual payment to the marginally curtailed type when there is no imbalance, as well as two point-masses, at a and b , corresponding to scenarios where the aggregator pays imbalance prices, depicted as circles with area proportional to their probability. The aggregator's optimal DA offer, q^* , sets the expected recourse cost equal to the DA revenue. That is, imagining that gravity is pulling the recourse cost distribution to the right, the optimal q^* balances that distribution on the DA price, p .

Table 1 $J_{DA}(q)$ and the optimal offer quantity q^* in Example 2

| Region $[q_i, \bar{q}_i]$ | $[q_i, \bar{q}_i] =$ | for $q \in [q_i, \bar{q}_i], J_{DA}(q) =$ | $q_i < q^* < \bar{q}_i$ if | if $q_i < q^* < \bar{q}_i$, then $q^* =$ |
|---|---------------------------|--|---|---|
| Case 1: $v(b-a)/2 \leq 1$ | | | | |
| 1a | $0 < q < va/2$ | $(p-a)q + (a/2)(va/2 + 1)$ | $p < a$ | 0 |
| 1b | $va/2 < q < vb/2$ | $(p-a - \frac{a^2v}{4})q + \frac{a}{2}q^2 - \frac{1}{3v}q^3 + \frac{a}{2} + \frac{a^3v^2}{24} + \frac{a^2v}{4}$ | $a \leq p \leq a + v(\frac{a-b}{2})^2$ | $va/2 + \sqrt{v}\sqrt{p-a}$ |
| 1c | $vb/2 < q < 1 + va/2$ | $(p-a - \frac{a^2v}{4} + \frac{b^2v}{4})q + (\frac{a-b}{2})q^2 + \frac{a^3v^2}{24} + \frac{a}{2} - \frac{b^3v^2}{24} + \frac{a^2v}{4}$ | $a + v(\frac{a-b}{2})^2 \leq p \leq b - v(\frac{a-b}{2})^2$ | $(p-a)/(b-a) + v(a + b)/4$ |
| 1d | $1 + va/2 < q < 1 + vb/2$ | $(p + \frac{b^2v}{4} + \frac{1}{v})q - (\frac{b}{2} + \frac{1}{v})q^2 + \frac{1}{3v}q^3 - \frac{1}{24}b^3v^2 - \frac{1}{3v}$ | $b - v(\frac{a-b}{2})^2 < p < b$ | $1 + vb/2 - \sqrt{v}\sqrt{b-p}$ |
| 1e | $1 + vb/2 < q$ | $(p-b)q + \frac{b}{2} + \frac{b^2v}{4}$ | $b < p$ | ∞ |
| Case 2: $v(b-a)/2 > 1, \Leftrightarrow 1 + va/2 \leq vb/2$ | | | | |
| 2a | $0 < q < va/2$ | $(p-a)q + (a/2)(va/2 + 1)$ | $p < a$ | 0 |
| 2b | $va/2 < q < 1 + va/2$ | $(p-a - \frac{a^2v}{4})q + \frac{a}{2}q^2 - \frac{1}{3v}q^3 + \frac{a}{2} + \frac{a^3v^2}{24} + \frac{a^2v}{4}$ | $a \leq p \leq a + 1/v$ | $va/2 + \sqrt{v}\sqrt{p-a}$ |
| 2c | $1 + va/2 < q < vb/2$ | $(p + \frac{1}{v})q - \frac{q^2}{v} - \frac{1}{3v}$ | $a + 1/v \leq p \leq b - 1/v$ | $(1 + vp)/2$ |
| 2d | $vb/2 < q < 1 + vb/2$ | $(p + \frac{b^2v}{4} + \frac{1}{v})q - (\frac{b}{2} + \frac{1}{v})q^2 + \frac{1}{3v}q^3 - \frac{1}{24}b^3v^2 - \frac{1}{3v}$ | $b - 1/v < p < b$ | $1 + vb/2 - \sqrt{v}\sqrt{b-p}$ |
| 2e | $1 + vb/2 < q$ | $(p-b)q + \frac{b}{2} + \frac{b^2v}{4}$ | $b < p$ | ∞ |

Differences from Case 1 highlighted in bold

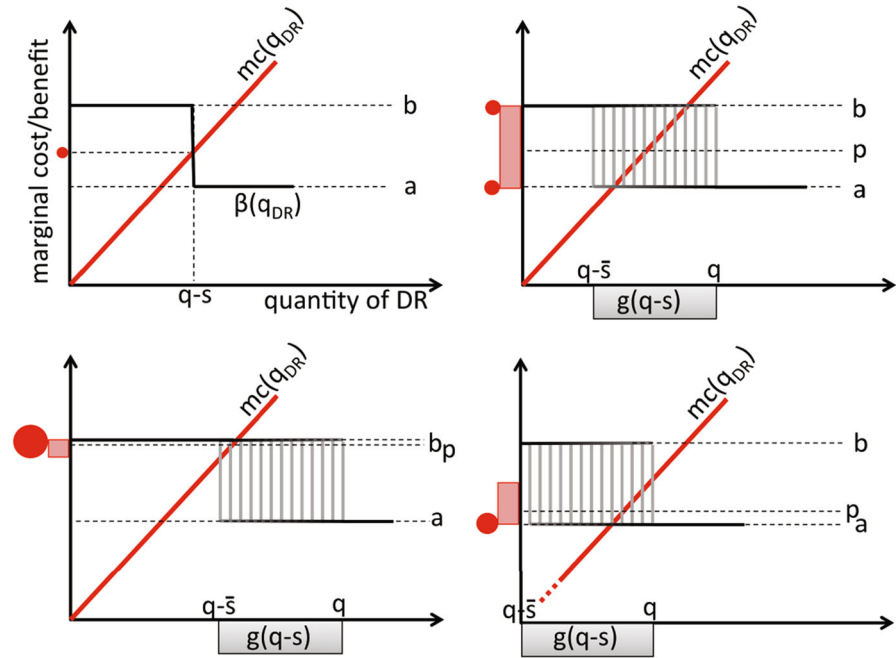


Fig. 2 Monopsony DR purchase decision and choice of DA offer q

The result of optimizing q contingent on DA information can be represented as a supply curve. However, in Example 2, since we assume that the imbalance prices a and b are revealed simultaneously with p , the policy mapping DA information to the DA offer quantity q is actually a “supply surface.” We display several representative slices of this surface in Fig. 3. In this figure, we consider three cases where the imbalance prices are equal to the day ahead price, plus a premium ϵ . As ϵ increases in this figure, the shortfall penalty and the overproduction rate are increased. Both of these effects move in the same direction, encouraging the producer to bid a smaller fraction of nameplate. (A higher overproduction payment encourages the producer to bid less, because it reduces the producer’s opportunity cost in scenarios where its supply exceeds its bid.) In future research we will discuss how one can solve the end-to-end problem numerically, via simulation, in a general model. The output of such an optimization can be offered as a supply curve in an ISO auction market.

In the symmetric Cournot oligopsony setting, the DR cost curve is deflected downward. (We assume that the aggregators have equal fractional shares of the DR market; to maintain comparability, we also need to assume that the aggregator owns a fractional share of a common wind resource, which is represented by a scaling of the wind distribution.) This should result in a lower marginal cost of curtailment, a higher optimal day-ahead offer, and of course, lower profits. We intend to explore the oligopsony setting in more depth in future work.

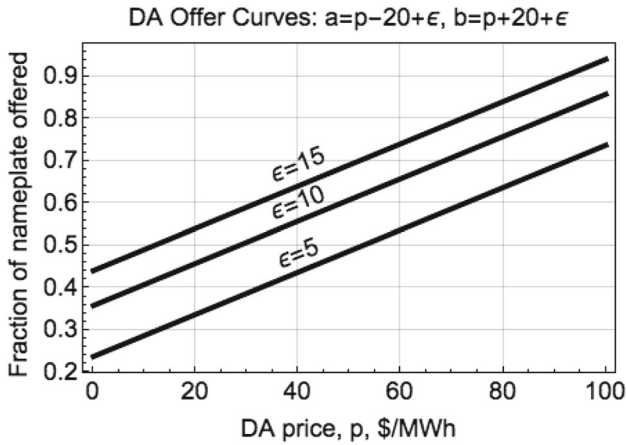


Fig. 3 Supply curves

3.3.2 Graphical depiction of the contract menu

Figure 2 characterizes the aggregator’s optimal DA action, q^* , and RT recourse, $\hat{\tau}^*$. Stepping back to the ex ante stage, we consider the ex ante contract menu that would implement $\hat{\tau}^*(\cdot)$. This requires some assumption of market statistics over Ω_{DA} . We consider the example:

$$\begin{aligned}
 p &\sim \text{Uniform}[10, 100] \\
 a &= (1 - \delta)p \\
 b &= (1 + \delta)p \\
 \delta &\sim \text{Uniform}[0.1, 0.9].
 \end{aligned}
 \tag{29}$$

We also assume that $\nu = 1/100$, which is derived from the following parameter choices under linear demand²⁷:

$$\begin{aligned}
 \bar{s} &= 100 \text{ MW} \\
 R &= \$30/\text{MW} \text{ (generation component of the retail price)} \\
 N &= D(R) = 100 \text{ MW} \text{ (aggregate demand at } R = \$30) \\
 \eta(R) &= 0.3 \text{ (elasticity at } R = \$30).
 \end{aligned}
 \tag{30}$$

In Fig. 4, we plot the probability of curtailment, as well as the ex ante payment, as a function of type.²⁸ We also display as a dotted line what the payment *would* be, if the aggregator maintained the same curtailment allocation, but were able to perfectly price discriminate (paying each increment its reservation price for curtailment, rather than the reservation price of the highest curtailed increment). The shaded region between

²⁷ Note that $\nu = N/(\bar{s}(\bar{\tau} - \underline{\tau})) = g(R)/\bar{s}$. The elasticity at the retail rate is $\eta(R) \triangleq g(R) \frac{R}{N}$. The parameter chosen yield $\nu = 1/100$.

²⁸ We did this by analytically solving the aggregator’s problem pointwise over Ω_{DA} , and then taking expectations of the quantities with respect to our market statistics.

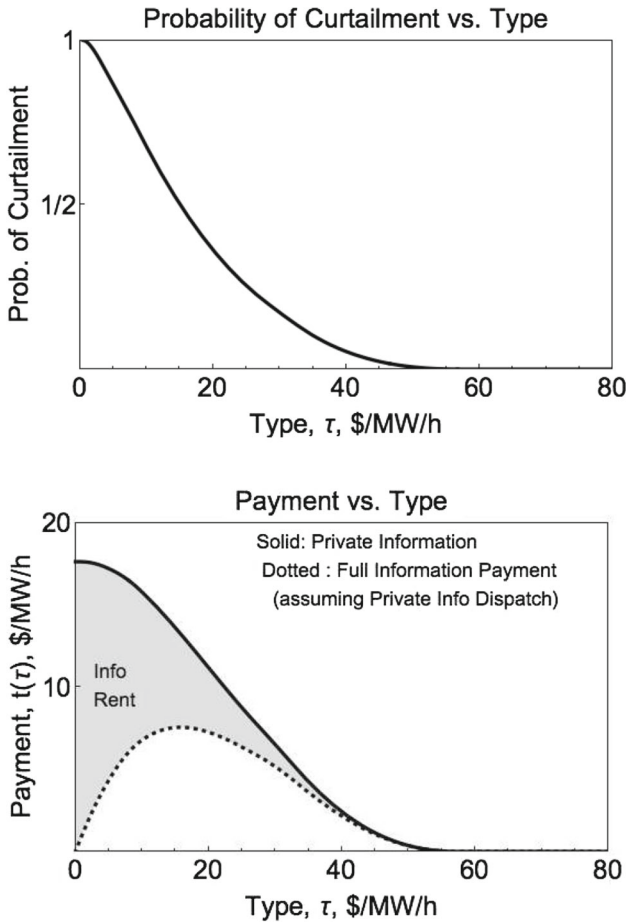


Fig. 4 DR contract as a function of type

those two lines is the information rent: the surplus payment that the increments are able to extract by virtue of their private information.

In Fig. 5, we eliminate the type parameter from the contract menu, by plotting the payment to each increment type against the probability that that type is curtailed. This gives a more realistic depiction of a menu that the aggregator might offer in practice. Here we see that the payment rises sub-linearly in probability. If we imagine only a single tuple of prices (p, a, b) , then the probability of curtailment is piecewise linear: there is no curtailment for valuations above $b/2$, valuations below $a/2$ are always curtailed, and between those two levels, the probability varies linearly in type, as determined by the uniform distributions over wind and type (see Eq. 15). However, while the probability of curtailment falls piecewise linearly in type, the corresponding valuation for service rises linearly. If each increment were paid its reservation utility (i.e. no private information) the payment would be a concave quadratic function of type, and the lowest type would accept curtailment for no payment. With private information,

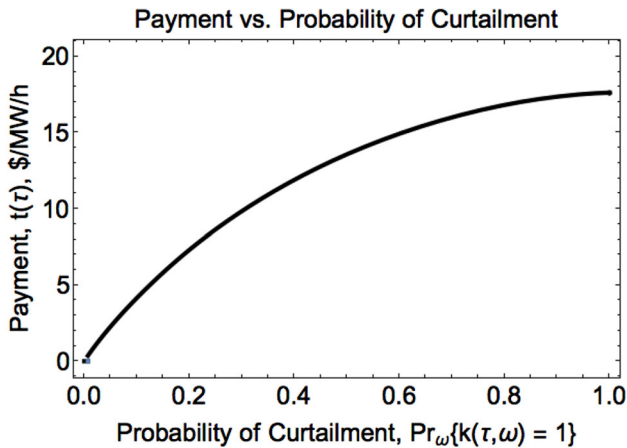


Fig. 5 Payment to DR as a function of probability of curtailment

the payment curve has a similar quadratic form, but it is decreasing in type, and the payments are inflated. (Incentive compatibility, combined with merit order curtailment, forces the payment to be nonincreasing in type. Merit order curtailment implies that service plans designed for higher types have a lower probability of curtailment. There is no way to compel a low-type increment to accept a package intended for it, with high probability of curtailment and low payment, if an alternative package is available for higher types, which has lower probability of curtailment and a higher payment.) When the market prices are stochastic, the payment as a function of type is an average, or ex ante expectation, over such decreasing piecewise quadratic curves.

4 Discussion and conclusion

We mentioned in Sect. 2 that the aggregator purchases less than the efficient level of demand response. This is because the aggregator is a monopsony purchaser of DR contracts. Under current institutional norms, customers effectively have option rights to consumption in the quantity of their physical fuse. But going forward, particularly with the rise of distributed generation, we anticipate that demand charges will become more prevalent. A demand charge requires the consumer to “buy the baseline,” or the capacity rights, which they would then re-sell to the aggregator. If the increments purchase less firm capacity up-front, then more flexibility will be available for real time adjustment and recourse in general. In order to investigate this, a model must incorporate an antecedent stage in which the DR participants purchase the baseline, as in, for example, [Crampes and Léautier \(2015\)](#).

The incorporation of a demand charge would typically remove lower-type increments from the distribution (or demand curve) faced by the aggregator. The effect of this can be ambiguous for the aggregator’s profit. On the one hand, increments with low but positive net ex post valuation can provide cheap DR for the aggregator to purchase and deploy in the wholesale market; removing these from the market would

reduce aggregator profit. On the other hand, increments with ex post valuation likely to be *below* the retail rate must be paid despite that they provide no DR, raising the total cost curve for DR in each scenario. In our Example 2, if we add increments with valuation below the retail rate, the aggregator would always have to pay them to dispatch a positive quantity of DR, but it would not get demand response from them. (This can be considered a form of adverse selection.) The result is that in each RT scenario, the aggregator would purchase either the same amount of DR as if there were no such increments (because of marginal cost calculations), or the aggregator would switch to purchasing zero DR in that scenario. We may explore the effects of demand charges more systematically in later work. Generally, we anticipate that a properly set demand charge can increase social welfare, but that a welfare-improving demand charge has an ambiguous effect on the aggregator's profits, depending on specific conditions.

Acknowledgments The authors thank Felipe Castro for his help with an earlier version of this project. They would also like to thank the anonymous referees, for their careful review and insightful comments. This work was supported by the Department of Energy, through a grant administered by the Center for Electric Reliability Technology Solutions (CERTS).

Appendix: Proofs

Proof of Lemma 2.8 Assumption 1 ensures the conditions sufficient for this claim. Note that $L(\cdot)$ is absolutely continuous, and differentiable except at two points. Further, recall that the support of $F(\cdot|\tau, \omega)$ is constant. Therefore we can apply integration by parts, and the antiderivative term drops out:

$$U(\tilde{\tau}|\tau) = \int_{\Omega_{RT}} \int_{\Theta} -L'(\theta)F(\theta|\tau, \omega)d\theta k(\tilde{\tau}, \omega)dP_{RT}(\omega).$$

The partial derivative in question is therefore

$$\frac{\partial}{\partial \tau} U(\tilde{\tau}|\tau) = \int \int -L'(\theta) \frac{\partial F(\theta|\tau, \omega)}{\partial \tau} k(\tilde{\tau}, \omega)dP_{RT}(\omega) \leq 0.$$

Each term in the integrand is uniformly bounded, which guarantees that this partial derivative is uniformly bounded, so that the function $U(\tilde{\tau}|\cdot)$ is Lipschitz continuous, and thus absolutely continuous.

Proof of Proposition 2.10 By rearrangement, the payment is the equilibrium utility minus the net option value:

$$t(\tau) = u(\tau) - U(\tau). \tag{31}$$

First we simplify the expression for $u(\tau)$ from line (8). Under merit order curtailment, $k(\tau, \omega) = \mathbb{1}_{\{\tau \leq \hat{\tau}(\omega)\}}$. Exchanging the order of integration so that we integrate first with respect to τ , we get that:

$$u(\tau) = u(\bar{\tau}) - \int_{\Omega_{RT}} \int_{\Theta} L(\theta)(f(\theta|\hat{\tau}(\omega), \omega) - f(\theta|\tau, \omega))k(\tau, \omega)dP_{RT}(\omega). \quad (32)$$

Subtracting from this the net option value $U(\tau)$, we arrive at the desired result.

Proof of Proposition 2.12 Consider the difference in utility between the case where an increment of ex ante type τ reports truthfully, versus mis-reporting as $\tilde{\tau}$.

$$\begin{aligned} & u(\tau) - (U(\tilde{\tau}|\tau) + t(\tilde{\tau})) \\ &= \int_{\Omega} \int_{\Theta} L(\theta) f(\theta|\tau, \omega) d\theta (k(\tau, \omega) - k(\tilde{\tau}, \omega)) dP(\omega) + t(\tau) - t(\tilde{\tau}) \\ &= \int_{\Omega} \int_{\Theta} L(\theta) (f(\theta|\tau, \omega) - f(\theta|\hat{\tau}(\omega), \omega)) d\theta (k(\tau, \omega) - k(\tilde{\tau}, \omega)) dP(\omega) \\ &= \int_{\Omega} \int_{\Theta} \underbrace{-L'(\theta)}_{\geq 0} (F(\theta|\tau, \omega) - F(\theta|\hat{\tau}(\omega), \omega)) d\theta (k(\tau, \omega) - k(\tilde{\tau}, \omega)) dP(\omega). \end{aligned}$$

Suppose $\tau > \tilde{\tau}$. Merit order curtailment implies that $k(\tau, \omega) - k(\tilde{\tau}, \omega) = -\mathbb{1}_{\{\tilde{\tau} \leq \hat{\tau}(\omega) < \tau\}} \leq 0$. On this set, FOSD implies that $F(\theta|\tau, \omega) \leq F(\theta|\hat{\tau}(\omega), \omega)$. This implies that the integral's value is nonnegative. Similarly, if $\tau < \tilde{\tau}$, then $k(\tau, \omega) - k(\tilde{\tau}, \omega) = \mathbb{1}_{\{\tau \leq \hat{\tau}(\omega) < \tilde{\tau}\}} \geq 0$. On this set, $F(\theta|\tau, \omega) \geq F(\theta|\hat{\tau}(\omega), \omega)$. Again, the integral's value is nonnegative. \square

Proof of Proposition 2.14 The substance of the Proposition is that (ignoring the penalty receipts term)

$$\begin{aligned} & \int_{\underline{\tau}}^{\bar{\tau}} t(\tau)g(\tau)d\tau \\ &= - \int_{\Omega_{RT}} \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} \left(\frac{G(\tau)}{g(\tau)} \frac{\partial}{\partial \tau} z(\tau, \omega_{RT}) + z(\tau, \omega_{RT}) \right) g(\tau) d\tau dP(\omega_{RT}). \quad (33) \end{aligned}$$

Integrating Eq. (9) under merit order curtailment, we get that

$$\int_{\underline{\tau}}^{\bar{\tau}} t(\tau)g(\tau)d\tau = \int_{\Omega_{RT}} \int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} z(\hat{\tau}(\omega_{RT}))g(\tau)d\tau dP(\omega_{RT}) \quad (34)$$

$$= \int_{\Omega_{RT}} z(\hat{\tau}(\omega_{RT}))G(\hat{\tau}(\omega_{RT}))d\tau dP(\omega_{RT}). \quad (35)$$

To obtain the desired formula, we first apply the fundamental rule of calculus to express the integrand as $\int_{\underline{\tau}}^{\hat{\tau}(\omega_{RT})} z(\tau)G(\tau)d\tau$. in (35) differentiate the integrand in (35) using the product rule, factor out a “ $g(\tau)$,” and re-integrate.

Proof of Proposition 3.8 (FONC for q^)* In order to differentiate this objective with respect to q , we decompose RT outcomes into three sets: overproduction = $\{DR + s >$

q }, shortfall = $\{DR + s < q\}$, and no imbalance = $\{DR + s = q\}$.²⁹ From expression (15), we see that

$$\begin{aligned} \frac{\partial}{\partial q} \hat{\tau}^*(\omega_{RT}|q) &= \frac{1}{DR'(DR^{-1}(q - s))} \mathbb{1}_{\{DR(\hat{\tau}^*)=q-s\}} \\ &= \frac{1}{DR'(\hat{\tau}^*)} \mathbb{1}_{\{DR(\hat{\tau}^*)=q-s\}} = \frac{1}{y(\hat{\tau}^*)g(\hat{\tau}^*)} \mathbb{1}_{\{DR(\hat{\tau}^*)=q-s\}}. \end{aligned}$$

Here we ignore certain “edge” events, which we assume have probability zero: abusing notation,

$$\begin{aligned} s_{\text{inf}} &= \inf\{s : DR(\hat{\tau}^*(s)) + s = q\} \\ s_{\text{sup}} &= \sup\{s : DR(\hat{\tau}^*(s)) + s = q\}. \end{aligned}$$

Also, note that the contributions to the expected derivative from the overproduction and shortfall terms are both zero when $DR + s = q$, because by assumption, the shortfall quantity is constantly zero on this set. When overproduction and underproduction are strict, then $\frac{\partial}{\partial q} \hat{\tau}^*(\omega_{RT}|q) = 0$.

This gives us that

$$\begin{aligned} \frac{\partial}{\partial q} J(q|\omega_{DA}) &= p - \mathbb{E}[a|DR(\hat{\tau}^*) + s > q] \Pr\{DR(\hat{\tau}^*) + s > q\} \\ &\quad - \mathbb{E}[b|DR(\hat{\tau}^*) + s < q] \Pr\{DR(\hat{\tau}^*) + s < q\} \\ &\quad - \mathbb{E}[c(\hat{\tau}^*)|DR(\hat{\tau}^*) + s = q] \Pr\{DR^* + s = q\}. \end{aligned}$$

The first order condition follows. □

References

- Al-Najjar, N. I. (1995). Decomposition and characterization of risk with a continuum of random variables. *Econometrica*, 63(5), 1195–1224. doi:[10.1111/1468-0262.00060](https://doi.org/10.1111/1468-0262.00060).
- Bitar, E., Poolla, K., Khargonekar, P., Rajagopal, R., Varaiya, P., & Wu, F. (2011). Selling random wind. In *Proceedings of the Annual Hawaii International Conference on System Sciences*, pp 1931–1937. doi:[10.1109/HICSS.2012.523](https://doi.org/10.1109/HICSS.2012.523).
- Bitar, E. Y., Rajagopal, R., Khargonekar, P. P., Poolla, K., & Varaiya, P. (2012). Bringing wind energy to market. *IEEE Transactions on Power Systems*, 27(3), 1225–1235. doi:[10.1109/TPWRS.2012.2183395](https://doi.org/10.1109/TPWRS.2012.2183395).
- Borenstein, S. (2005). The long-run efficiency of real time electricity pricing. *The Energy Journal*, 26, 93–116. doi:[10.5547/ISSN0195-6574-EJ-Vol26-No3-5](https://doi.org/10.5547/ISSN0195-6574-EJ-Vol26-No3-5).
- Börgers, T. (2010). *An introduction to the theory of mechanism design*. Oxford: Oxford University Press. <http://aida.econ.yale.edu/~dirkb/teach/521b-08-09/reading/2008-mechanismdesign.pdf>.
- Braithwait, S., Hansen, D., & Kirsch, L. (2006). *Incentives and rate designs for efficiency and demand response*. Lawrence Berkeley National Laboratory Collaborative Paper 60132 (March). <http://smartgrid.lbl.gov/sites/all/files/60132.pdf>.

²⁹ Despite the fact that the DR policy is held at its optimal values, we cannot apply an envelope theorem, because there is a positive probability that the value function is not differentiable with respect to $\hat{\tau}$ at $\hat{\tau} = \hat{\tau}^*$ (in all ω_{DA} events where there is zero imbalance at the optimum, the right derivative with respect to $\hat{\tau}$ has an a coefficient, and the left derivative has a b coefficient).

- Caramanis, M., Schweppe, F. D., & Tabors, R. D. (1983). *Spot pricing and its relation to other load management methods*. Technical report. MIT Energy Laboratory Report MIT-EL 83-001
- Chao, H. P. (1983). Peak load pricing and capacity planning with demand and supply uncertainty. *Bell Journal of Economics*, 14(1), 179–190.
- Chao, H. P. (2012). Competitive electricity markets with consumer subscription service in a smart grid. *Journal of Regulatory Economics*, 41, 155–180. doi:10.1007/s11149-011-9179-7.
- Chao, H. P., & Wilson, R. (1987). Priority service: Pricing, investment, and market organization. *The American Economic Review*, 77(5), 899–916. <http://www.jstor.org/stable/1810216>.
- Courty, P., & Li, H. (2000). Sequential screening. *Review of Economic Studies*, 67(4), 697–717. doi:10.1111/1467-937X.00150.
- Crampes, C., & Léautier, T. O. (2015). Demand response in adjustment markets for electricity. *Journal of Regulatory Economics*, 48(2), 1–33. doi:10.1007/s11149-015-9284-0.
- Holland, S. P., & Mansur, E. T. (2006). The short-run effects of time-varying prices in competitive electricity markets. *The Energy Journal*, 27, 127–155.
- Judd, K. L. (1985). The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory*, 35(1), 19–25. doi:10.1016/0022-0531(85)90059-6.
- Krähmer, D., & Strausz, R. (2015). Sequential versus Static Screening: An equivalence result, 1–13.
- Margellos, K., & Oren, S. (2015). Capacity controlled demand side management: A stochastic pricing analysis. *IEEE Transactions on Power Systems* 99, 1–12. doi:10.1109/TPWRS.2015.2406813. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7061537>.
- MDPT Working Group. (2015). Report of the market design and platform technology working group. Technical report. http://newyorkrevworkinggroups.com/wp-content/uploads/MDPT-Report_150817_Final.pdf.
- Milgrom, P., & Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2), 583–601. doi:10.1111/1468-0262.00296.
- Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 317, 301–317.
- Oren, S. S. (2013). A historical perspective and business model for load response aggregation based on priority service. In *2013 46th Hawaii international conference on system sciences (HICSS)*, pp 2206–2214. doi:10.1109/HICSS.2013.39.
- RLW Analytics. (2007). *Deemed savings estimates for legacy air conditioning and water heating direct load control programs in PJM region*. Technical report. Middletown: RLW Analytics.
- Tan, C. W., & Varaiya, P. (1991). *A model for pricing interruptible electric power service*. Technical report. Berkeley: EECS Department University of California. <http://www.eecs.berkeley.edu/Pubs/TechRpts/1991/1695.html>.
- Tan, C. W., & Varaiya, P. (1993). Interruptible electric power service contracts. *Journal of Economic Dynamics and Control*, 17, 495–517.
- Tsitsiklis, J. N., & Xu, Y. (2015). Pricing of fluctuations in electricity markets. *European Journal of Operational Research*, 246(1), 199–208. doi:10.1016/j.ejor.2015.04.020. <http://www.sciencedirect.com/science/article/pii/S0377221715003136>. arXiv:1207.2848v1.
- Varian, H. R. (1992). *Microeconomic analysis* (2nd ed.). New York: WW Norton.
- Wilson, R. B. (1993). *Nonlinear pricing*. New York: Oxford University Press.