BESS Wholesale market arrangements for battery energy storage systems

Code amendment consultation paper

Submission by Electric Power Optimization Centre

The University of Auckland

This document was prepared by

Professor Andy Philpott

https://www.epoc.org.nz

June 30, 2026

Executive Summary

1. This document is a response to the call for submissions on the Electricity Authority Consultation paper "BESS Wholesale market arrangements for battery energy storage systems" published on May 19, 2026.
2. The market arrangements that we wish to comment on in this submission are those that pertain to:
    a. bids and offers of BESS energy;
    b. fixing gate closure at 1 hour;
    c. state-of-charge constraints.
3. We take the position advocated by NewPower that BESS bids and offers should be linked to their state of charge. This means that in each period a BESS would offer a (possibly) different supply function of bids and offers for each potential storage level in their battery. We call such offers "Agent Decision Rules" or ADRs. ADRs would obviate the need for state-of-charge constraints and would provide the flexibility in offers that is constrained by a one-hour gate closure with conventional single supply functions.
4. An ADR is designed to represent an opportunity cost of discharging. This has obvious connections to the concept of marginal water value for a (possibly pumped) hydro storage reservoir. An ADR offer can then be represented as a family of marginal value curves, or as a single total expected future value of storage function. This makes the implementation in a dispatch system no more complicated for the system operator than requiring a single curve defining the expected future value of storage.
5. A number of submitters have argued for reduced gate closure to give more flexibility for BESS to respond to five-minute price signals. We contend that an ADR offered for every five-minute period in the next hour provides more optionality for BESS owners than a set of twelve time-stamped offer curves. A one-hour gate closure gives the system operator some advance warning of possible system state while providing optionality to the BESS.
6. An ADR passes control of the state of charge of a BESS from the system operator to the BESS owner who will take a position on possible future price outcomes and offer accordingly. This passes price forecasting from the system operator to each BESS owner who will make offers that put "their money where their mouth is".
7. Since ADR offers can co-exist alongside conventional offers that may be preferred by some BESS owners, they can be introduced incrementally to SPD as an optional form of offer and bid.
8. We recommend that the Electricity Authority investigate the possibility of implementing ADRs in SPD as an option for BESS market participants.

Agent decision rules for dispatching BESS

We propose that in each period a BESS would offer a (possibly) different supply function of bids and offers for each potential storage level in their battery. We call such offers "Agent Decision Rules" or ADRs and discuss them in a recent EPOC paper "Electricity dispatch and pricing using agent decision rules"[1].

An ADR for a BESS owner could be represented as a family of supply functions, one for each possible state of charge. The system operator when dispatching the BESS in a five-minute interval would choose the appropriate supply function for the observed state of charge at the beginning of this interval.

An example can serve to illustrate this. Suppose a BESS owner has a 40MWh battery and can charge and discharge at a maximum rate of 180MW. Suppose that they have a proprietary model of wholesale electricity prices over the next 24 hours that can be used to generate a probability distribution of price for each five-minute period. For the purposes of the example suppose the random price in any period is independent of other periods. Then the BESS owner can solve a stochastic dynamic program to maximize its expected revenue over the day. This defines a Bellman function for each period defining expected future revenue for each level of storage. A piecewise linear example is shown in Figure 1. The slope of the Bellman function is shown in Figure 2.

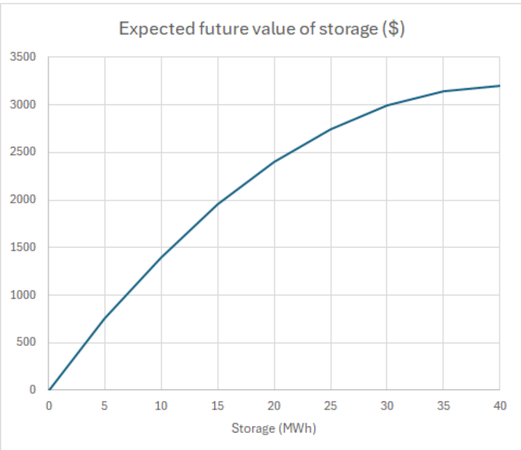

Figure 1: Example Bellman function

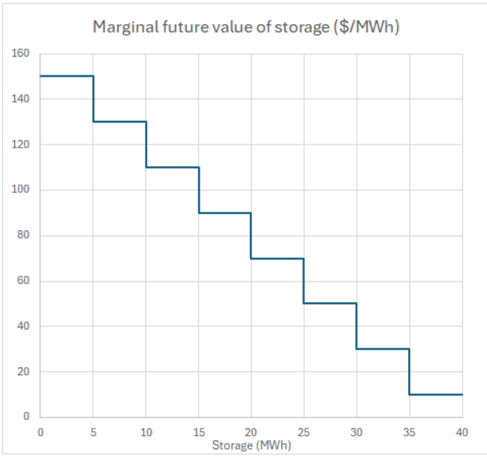

Figure 2: Marginal future value

Suppose the battery has 25 MWh of charge in it. Then it offers to generate at its marginal cost of discharge. This is shown as the orange step function in Figure 3 (that needs to be reversed to be an increasing offer curve). The maximum it can discharge in 5 minutes is 180*5/60=15MWh.

The BESS operator will also bid to charge when the energy price is below the marginal value of battery storage. This is shown as the green step function in Figure 3. The maximum charge the BESS can absorb in 5 minutes is 180*5/60=15MWh.

The offer tranches and bid tranches are shown in Table 1.

[1]A.B. Philpott, M.C. Ferris and J.P. Mays, Electricity dispatch and pricing using agent decision rules, https://www.epoc.org.nz/papers/ADRPaperForEPOC.pdf.

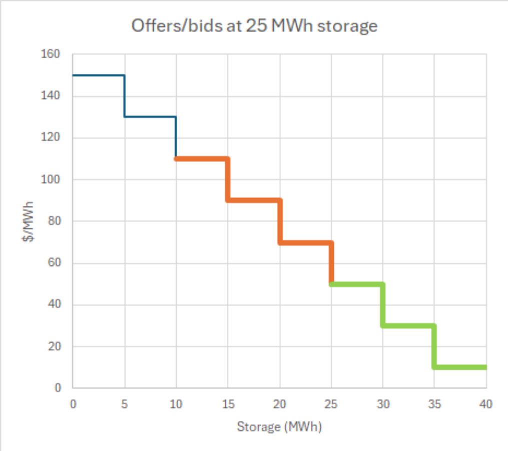

Sell tranche (MWh)	Buy tranche (MWh)	Price ($/MWh)
	5	10
	5	30
	5	50
5		70
5		90
5		110

Figure 3

Table 1: Sell and buy tranches

The offer-bid curve in Table 1 is for a state of charge of 25MWh. If the initial state of charge were 15MWh then the offers in Table 1 become those shown in Table 2 that reflect the higher marginal value placed on storage at the lower level of charge.

Sell tranche (MWh)	Buy tranche (MWh)	Price ($/MWh)
	5	50
	5	70
	5	90
5		110
5		130
5		150

Table 2: offer-bid curve with lower storage

This example shows that the single curve in Figure 1, or its derivative in Figure 3, can be used to supply all the offers and bids needed to define an ADR.

State-of-charge constraints

With standard offers, BESS operators must determine offers in advance that do not adjust to the BESS state of charge. This means that an offer to sell might be dispatched more than the available BESS storage. In this case the BESS discharge must be curtailed by the system operator by imposing a state-of-charge constraint in real time. In contrast, an ADR offer will constrain the maximum dispatch to the available charge, and maximum energy purchase to the BESS headroom. These constraints are imposed on the ADR offers by the BESS trader, not the system operator. This gives each BESS participant more control over their desired dispatch.

Gate closure

Real time prices are computed every five minutes. BESS owners earn revenue by arbitraging these prices. This is difficult to do effectively using a set of offers that must be set at gate closure one hour before dispatch. An ADR provides more flexibility for the BESS to profit from price variations. To see this observe that a classical offer made an hour ahead for each five-minute period is a specific example of a feasible ADR, so an optimal ADR will be no worse in terms of expected revenue, and is likely to be better. If the price process estimated by the BESS owner is the same as the price process that emerges in equilibrium, and the BESS owner constructs an optimal ADR using this price process then they will make the same return as they would if optimal offers could be submitted in real time.

In practice, of course, each BESS owner will make their own assumptions regarding future prices, and the equilibrium price outcome will integrate the views of different market participants into system prices. So the BESS dispatch (based on an ADR computed using the BESS operator's forecast prices) is likely to produce a lower revenue than an optimal ADR computed using these system prices.

Implementation

The implementation of ADRs into dispatch software SPD is straightforward. ADRs could be offered by each BESS as a family of supply/demand curves. Alternatively, each BESS could supply a suitably defined piecewise linear Bellman function (e.g. using cutting planes) that could be summed over BESS participants and (with a sign change) used as cost-to-go function in each five-minute dispatch. Neither approach precludes market participants offering in a single supply/demand curve for each period. For example, some participants might view the gains from offering ADRs as not worth the effort required to compute them. If they do generate more revenue for BESS participants then this will incentivize more entry of these into the market, enhancing its efficiency.

Electricity Dispatch and Pricing using Agent Decision Rules

Andy Philpott* Michael Ferris† Jacob Mays‡

October 21, 2025

Abstract

Models for computing dispatch and prices in wholesale electricity market pools are typically deterministic multiperiod mathematical programs that are solved in a rolling horizon fashion. In convex settings with perfect foresight these optimization problems yield dispatch outcomes and locational marginal prices that solve a competitive equilibrium problem. The use of these models in practice is challenging for several reasons, particularly in the context of increased uncertainty resulting from growing investment in renewable energy. Deterministic models can miscalibrate the value of holding energy in storage or positioning the system to meet future ramping constraints, leading to inefficient dispatch decisions. Pricing outcomes from the models are dependent on the point forecasts used as inputs, leading to inefficient remuneration and uplift payments that compensate participants for the fact that the system operator forecasts the future incorrectly. To address these challenges, researchers and practitioners have proposed a variety of model enhancements (e.g., the use of longer lookahead periods or the implementation of stochastic programming models using scenario trees) that increase the computational and informational demands placed on the system operator. We present a class of new economic dispatch models that instead attempt to overcome these drawbacks through the use of agent decision rules. Forecasting future outcomes or scenarios passes from the system operator to market participants, who implicitly make state-dependent offers of energy through these decision rules. We show how storage and ramping can be priced correctly in convex markets and illustrate the advantages of the approach through simple examples.

*Electric Power Optimization Centre, University of Auckland, *a.philpott@auckland.ac.nz*

†Department of Computer Sciences and Wisconsin Institute for Discovery, University of Wisconsin-Madison, *ferris@cs.wisc.edu*

‡School of Civil and Environmental Engineering, Cornell University, *jacobmays@cornell.edu*

1. Introduction

Wholesale electricity markets across the world are confronting challenges brought by a rapid transition away from traditional technologies toward solar, wind, storage, and distributed energy resources. As increasing numbers of batteries and flexible loads shift consumption "behind the meter", market operators face greater uncertainty in net demand. This challenge has become particularly acute in jurisdictions with high renewable penetration, such as Australia (Nelson et al. 2025).

Growth in renewable supply has motivated a great deal of research investigating improved algorithmic approaches for managing uncertainty and for coordinating storage and distributed resources, e.g., through stochastic or robust optimization. To date, however, real-world systems have stopped short of an explicit treatment of uncertainty in algorithms used for unit commitment, economic dispatch, and price formation. An alternative to purely operator-managed uncertainty is to enable batteries and other flexible devices (often through aggregators) to participate directly in wholesale dispatch and pricing mechanisms by submitting supply and demand bids. Such participation can increase market efficiency by coordinating these bids with other dispatchable services, such as peaking plant. To improve the participation of load-shifting and ramping services, market operators have implemented a variety of design adaptations, such as new market participation models for batteries and new ancillary service products for ramping.

In this paper we consider alternative ways to incorporate stochastic elements in short-term electricity market design and propose a change in the format of bids and offers supplied by market participants. In current practice, market operators rely on a series of deterministic lookahead models solved in a rolling horizon fashion, taking bids and offers from market participants as an input into the models. This approach leads to three potential issues. First, the use of a deterministic formulation may lead to suboptimal decisions and inefficient prices within the lookahead horizon. Second, the use of a finite lookahead horizon can lead to myopic decisions that fail to prepare the system for operations beyond that horizon. Third, the parameterization of lookahead models requires the system operator to make decisions (e.g., regarding demand forecasts) that can meaningfully affect prices, with unclear consequences for efficiency in both short-run operations and long-run investment. With these three issues in mind, the goal of this paper is to shift auctions away from the conceptual framework of model predictive control toward that of dynamic programming. In place of classical price–quantity pairs, our proposal aims to enable market participants to submit offers that amount to Agent Decision Rules (ADRs) that can be adapted to any scenario that arises. In a dynamic setting, decisions depend not just on the profits earned in a given interval, but also on a value function reflecting the future benefit of being in a different state at the end of that interval. The classical supply functions used in current formulations for commitment, dispatch, and market clearing do not include a direct way of expressing this value function. Effectively, the current format embeds an assumption that future benefits will be unrelated to the state transition that results from current decisions. While such an assumption may have been reasonable in the past, when intertemporal constraints were less of a concern for operators, it is increasingly suspect given issues with ramping constraints and battery state-of-charge limits.

Rather than relying primarily on the system operator, the ADR approach to uncertainty management depends on market participants to develop their own view of future uncertainty and incorporate it in their bids and offers. As such, the proposal reflects a continuation of debates about the split of responsibilities between market participants and the market operator that have been ongoing since the introduction of competition. Some markets, such as in New Zealand and most of Australia, rely on self-commitment of thermal generators, while others, including most of the U.S., rely on central commitment. Some markets, like in Alberta, rely purely on participant offers, while others, like much of South America, rely on costs estimated by the system operator. Still others, as in the U.S., rely on a complicated mix of these, with participant offers replaced by cost-based offers in cases of significant market power. In other words, while sharing common theoretical underpinnings, competitive markets in different jurisdictions have evolved in very different directions to accommodate their specific technological and regulatory contexts. In the U.S. context, growth of wind and solar has led many to argue for a more centralized approach. The California Independent System Operator (CAISO), for example, describes its operational issues as stemming from "the challenges of having a limited optimization horizon," indicating a desire to extend its modeled lookahead horizon further into the future if computationally feasible (Department of Market Monitoring 2024). Our analysis instead suggests a clean division of responsibility at the frequency of market clearing (e.g., 5 minutes in U.S. markets), with the market operator solving a single-period economic dispatch model to generate prices that balance projected supply and demand in each interval but employing control mechanisms to ensure more precise balancing within the interval. At the same time, the analysis provides insight into the conditions under which a more operator-driven approach to managing uncertainty across market-clearing intervals may be required.

In an idealized setting with a fully specified scenario tree, a socially optimal schedule for dispatch under uncertainty can be determined through stochastic programming. This observation has motivated many studies examining the use of stochastic programming in dispatch and market clearing, primarily examining simpler two-stage models (see, e.g., Pritchard et al. (2010), Zavala et al. (2017), Cory-Wright et al. (2018), Zakeri et al. (2019)). Rolling horizon models implement the dispatch from the current period and re-optimize the model with a new scenario tree starting in the next period. The prices generated by these models support the schedules that would be chosen by profit-maximizing market participants, as long as agents are risk neutral and agree on the probabilities attached to each scenario. With assumptions enabling complete trading in risk, this result can be extended to situations with risk aversion (Ferris and Philpott 2022). While moving from current deterministic formulations to stochastic programming could improve the management of uncertainty within the lookahead horizon, it would not address the issues of myopia and conflicting beliefs noted above. Given that the implementation of stochastic programming would require simplifications from the full scenario tree, a key question is how the market operator would choose scenarios for use in the model. The construction of scenarios could have meaningful effects on the prices ultimately formed, leading to divergence between the schedule determined by the market operator and the ones preferred by individual agents (Mays 2024).

Issues connected to dispatch under uncertainty have led to ongoing evolution in the participation models used by batteries in U.S. markets. Given its early deployment of significant battery capacity, discussions of new participation models are most active in CAISO, which employs a deterministic lookahead economic dispatch model in real-time market clearing that extends two hours into the future. In CAISO, gate closure occurs 75 minutes in advance of each operating hour and offers must be constant through an operating hour. From a dynamic programming standpoint, efficiency-maximizing offers in later periods of the lookahead horizon should depend on decisions made earlier in the horizon. Typically, it can be expected that the residual value of energy stored in a battery will increase as the battery gets closer to being empty. The current rules introduce inefficiency into the dispatch in two ways. First, due to gate closure, the battery operator does not know what its state of charge will be at the beginning of the operating hour and so is not able to match its offer to its estimate of residual value. Second, if the battery offers a set of constant price–quantity pairs for the hour, the market clearing engine can select the cheapest segment of the curve in each 5-minute interval rather than moving up the residual value curve as the battery is discharged. In response to these issues, CAISO is contemplating changes that would allow state-of-charge-dependent battery offers but would introduce non-convexity to the otherwise convex economic dispatch (Zheng et al. 2023, Chen and Tong 2023). In principle, the ADR-based approach advanced in this paper would enable offers to depend on state of charge and avoid issues with gate closure without introducing non-convexity.

In addition to changes to participation models, several authors have proposed to resolve issues with mismatched incentives by modifying price formation in a way that mitigates the potential for market participant losses. Here the issue is that energy prices in the first period of each lookahead model solution are binding for settlement, whereas the remaining prices are only provisional. The lookahead model may hold energy in a battery or pre-position a ramp-constrained generator in anticipation of higher future prices, even if the market participant does not believe that such high prices will arise. More generally, the sequence of prices generated in a rolling horizon fashion with lookahead models relying on forecasts often turn out to be different ex-post from the prices that would be obtained from solving a perfect foresight model with the observed values of the parameters. This results in a so-called *lost opportunity cost* faced by market participants who would have acted differently from their dispatched quantities if they had known the prices in advance. U.S. markets use side payments to encourage compliance with operator instructions. These payments are large and growing, amounting to 7 percent of battery revenues in CAISO in 2023 (Department of Market Monitoring 2024). In an effort to limit these payments, mechanisms to ensure consistency between rolling horizon and perfect foresight models in the deterministic setting were proposed by Hogan (2016) and studied by Hua et al. (2019). Real-time price consistency in a stochastic setting is addressed by Cho and Papavasiliou (2023), who propose a pricing model that minimizes expected ex-post lost opportunity cost, a measure of the regret experienced by market participants when they view their historical dispatch in the realized sequence of prices. A challenge in the analysis of these alternative pricing models is that adjustments to price formation can lead to different opportunity costs for resources, leading to different participant offers and inefficient commitment and dispatch solutions. As in Eldridge et al. (2023a,b), we adopt a different approach to Cho and Papavasiliou (2023) by placing the focus more on ex-ante outcomes. When generators and battery owners face future uncertainty they take positions that risk losses. Some of these losses result from being dispatched in advance of a realized random price under which they would have preferred to be dispatched differently. In our dispatch model, we propose that the generators should factor this possibility into their ADRs and not be compensated with an uplift payment should they experience some ex-post losses.

The specification of state-dependent future cost functions has a natural interpretation using dynamic programming, where market participants construct offers that fully encode decision rules applicable to any potential state of the system. These policies form the basis of Lagrangian relaxation techniques for solving deterministic economic dispatch problems that have a long history dating back to Muckstadt and Koenig (1977). Over the past twenty years Lagrangian relaxation models have been superseded by mixed integer programming formulations that generally yield better solutions (Hobbs 2001, Li and Shahidehpour 2005). In recent years, the increase in renewable energy and battery storage has resulted in a renewal of interest in Lagrangian relaxation for solving stochastic economic dispatch problems (Brown and Smith 2025), resulting in price-directed decision rules for optimizing the generating decisions of plants in any observed state of the system. Our approach is similar, but constructs ADRs that provide state-dependent energy offers to a system operator. The ADRs will involve a short-run marginal cost and a future cost function that enables the system operator to dispatch resources and generate prices by solving a single-period economic dispatch problem.

Ideally, decision rules from dynamic programming solutions will approximate socially optimal policies. We show that a dispatch that maximizes expected social welfare can be duplicated by ADRs when all agents share the same information and beliefs as the social planner, and the costs and constraints of the social planner satisfy some separability conditions. While it is difficult to demonstrate convergence to a socially optimal equilibrium more generally, we give some simple examples showing how dispatching based on imperfect decision rules can achieve results that approximate the social optimum.

Our contributions can be summarized as follows.

1. We propose a new form of energy offer for market participants, an ADR, that encapsulates their view of future market conditions.

2. We show how an optimal dispatch can be computed by solving a sequence of single-period problems without requiring lookahead. We show that prices from this process approximate the correct prices for competitive equilibrium.

3. We define the concept of ADR partial equilibrium, a situation in which agent's beliefs of future prices give rise to ADRs that together result in the conjectured prices, and give an example of such an equilibrium.

4. We show how an ADR partial equilibrium in a setup with a complete market for contingent contracts will result in agents assuming the same probability distributions about future events.

5. We show how a dispatch that maximizes expected social welfare can be duplicated by agent decision rules when all agents share the same information, and the costs and constraints of the social planner satisfy some separability conditions.

6. We describe an approach for separating the decisions of the system operator (who should ensure a reliable supply of power) from market participants (who seek to benefit financially from their foresight into future market conditions).

The paper is laid out as follows. In Section 2, we formulate a simple deterministic example of an electricity dispatch model to establish an optimization framework and notation for the rest of the paper. Section 3 then defines ADRs and the notion of an ADR partial equilibrium in this setting. Section 4 establishes a correspondence between social optimization and partial equilibrium when agents share the same probability beliefs. To illustrate the use of ADRs in a practical setting, we discuss two numerical examples in Section 5. Section 6 gives a selection of optimization problems in electricity markets that are amenable to modeling using ADRs, and Section 7 concludes.

2. Multiperiod economic dispatch

In this section we define a simplified version of a convex multiperiod economic dispatch problem for an electricity system, where we ignore many of the complex constraints that are a feature of these models in practice. Our purpose here is to to fix notation and to provide a framework and definition of Agent Decision Rules in the convex setting. To do this consider the following social optimization problem.

$$\begin{aligned} \text{SOP:}\ \min\ & \textstyle\sum_{t=1}^{T}(c^t(x(t)) + L^\top z(t)) \\ \text{s.t.}\ & Ax(t) + z(t) \geq d^t, \\ & z(t) \in [0, d^t], \quad x(t) \in \mathcal{X}^t, \quad t = 1, 2, \ldots, T. \end{aligned}$$

Here, decision variables $x(t)$ incur cost $c^t(x(t))$ and must lie in the feasible set $\mathcal{X}^t = \mathcal{X}(x(t-1))$. The total output in period t is $Ax(t)$ where A is a real matrix. This must meet demand in period t denoted by d^t, where $z(t)$ represents lost load, penalized at a shortage cost denoted as L. To allow for load at different locations we penalize $L^\top z(t)$ where L is a vector of (possibly) different lost load costs. We assume throughout that SOP is a convex optimization problem.

The model SOP is quite general.

Example 1. If $x(t) = (q(t), f(t))$ consists of generation $q(t)$ and transmission flows $f(t)$, where the kth row of $Ax(t)$ would equal the net supply in node k of a transmission grid, giving

$$q_k(t) + \sum_l (f_{lk}(t) - f_{kl}(t)) + z_k(t) \geq d_k^t.$$

The constraints $x(t) \in \mathcal{X}(x(t-1))$ would then be

$$(q(t), f(t)) \in \mathcal{Q}(q(t-1)) \times \mathcal{F}$$

where

$$\mathcal{Q}(\bar{q}) = \{q : 0 \leq q \leq q^{\max}, q - \bar{q} \leq \rho, \bar{q} - q \leq \sigma\} \tag{1}$$

represents generation capacities ($q^{\max}$) and ramping constraints on generation, and $\mathcal{F}$ denotes the set of feasible transmission flows meeting thermal limits and loop-flow constraints in a DC-Load flow model. The cost function $c^t(q(t), f(t))$ would measure a variable cost of generation and zero cost for transmission. □

Example 2. One could set $x(t) = (q(t), y(t), u(t), v(t))$ in a single location model, where $q(t)$ is generation, $y(t)$ is storage, $v(t)$ is storage charging rate and $u(t)$ is storage discharging rate. Then $Ax(t)$ would give

$$\textstyle\sum_{i\in\mathcal{G}} q_i(t) + \sum_{i\in\mathcal{B}} u_i(t) - \sum_{i\in\mathcal{B}} v_i(t) + z(t) \geq d^t,$$

where $\mathcal{G}$ is a set of generators, $\mathcal{B}$ is a set of batteries,

$$x(t) = (q(t), y(t), u(t), v(t)) \in \mathcal{Q}(q(t-1)) \times \mathcal{Y}(y(t-1)),$$

$\mathcal{Q}$ is defined by (1) and

$$\mathcal{Y}(\bar{y}) = \{(y, u, v) : 0 \leq y \leq E, 0 \leq u \leq r, 0 \leq v \leq s, y = \bar{y} - u + \eta v\}.$$

Here r and s are bounds on rates of discharge and charge respectively, and the diagonal matrix η measures the round-trip efficiency of the batteries. As before we would set $c^t(q(t), y(t), u(t), v(t))$ to be the variable cost of generation, where battery cycling costs could be imposed with a cost on $u(t)$. $\square$

3. Agent decision rules

Moving away from the deterministic social optimization problem, we now consider a competitive market consisting of a set $\mathcal{I}$ of agents. Each agent $i \in \mathcal{I}$ seeks to choose actions to maximize its expected surplus over periods $t = 0, 1, 2, \ldots$, where agent i operates in probability space $(\Omega, \mathcal{F}, \mathbb{P}_i)$ endowed with a filtration $\{\mathcal{F}_t\}_{t=0}^{T}$. We assume that the sample space and filtration is fixed and known by all agents, but they may have different probability measures. Demand d^t is a random variable measurable with respect to $\mathcal{F}_t$.

To simplify the analysis we will henceforth use a finite sample space, so the probability space for each agent can be interpreted as a single scenario tree $\mathcal{T}$ with root node 0, and demand d^n defined at each node $n \in \mathcal{T}$, but with (possibly) different probability measures. Here we use notation $n-$ to denote the predecessor node of n and $n+$ to denote the set of immediate successors of n. We assume a hazard-decision information structure, in which each agent $i \in \mathcal{I}$ in node n observes their own incoming state $x_i(n-)$ and the states of other agents $x_{-i}(n-)$ (in node 0, the initial incoming state $x(0-)$ is given), and takes action $x_i(n)$ given information revealed in node n. The action is chosen to optimize their current and future payoff given current prices $\pi(n)$ and their probability distribution of future prices. The outgoing state in node n resulting from this action is then denoted $x_i(n)$ (by an overload of notation). The action $x_i(n)$ of each agent i in node n is constrained to lie in a convex set $\mathcal{X}_i(x_i(n-))$. Note that this set depends only on the incoming state of agent i, as is the case in Example 1 and Example 2.

To align the agent problems with the social optimization problem SOP we make the following assumption.

Assumption 1. *For $n \in \mathcal{T}$, $c^n(x) = \sum_i c_i^n(x_i)$, and $\mathcal{X}(x(n-)) = \prod_i \mathcal{X}_i(x_i(n-))$*

In practice this is not restrictive as long as ownership structures are simple. Hydro stations operated by different agents but linked in a cascade would be an exception.

For each $n \in \mathcal{T}$, agent $i \in \mathcal{I}$ constructs a concave function $V_i^n(\bar{x}_i)$ that estimates their future expected surplus from node n onwards given their incoming state $\bar{x}_i$. This requires an agent probability measure $\mathbb{P}_i$ on the tree, where $\mathbb{P}_i(0) = 1$, and for $m \in n+$ we write $\mathbb{P}_i(m \mid n) = \mathbb{P}_i(m)/\mathbb{P}_i(n)$. (If n is a leaf node of $\mathcal{T}$ then we set $\mathbb{P}_i(m \mid n) = 0$.) The future value function $V_i^n(\cdot)$ is defined by the recursion

$$V_i^n(\bar{x}_i) = \max_{x_i \in \mathcal{X}_i(\bar{x}_i)} \{\pi_i(n)^\top A_{\cdot i} x_i - c_i^n(x_i) + \sum_{m \in n+} \mathbb{P}_i(m \mid n) V_i^m(x_i)\}. \tag{2}$$

Here $A_{\cdot i}$ denotes the columns of A that correspond to agent i. Since we make no assumption that all agents share the same beliefs about the probability distribution of demand and future prices, π_i has subscript i. The expectation is conditioned on the history of demand, and is computed using agent i's conditional probability $\mathbb{P}_i(m \mid n)$. Note that V_i^m is a function of the outgoing state x_i in node n which will be the incoming state in node $m \in n+$. The function $R_i^n(\cdot) = \sum_{m \in n+} \mathbb{P}_i(m \mid n) V_i^m(\cdot)$ is called the Agent Decision Rule (ADR) for agent i.

Market agents can communicate their views of the future to the system operator using an ADR that will describe how to dispatch their plant in each period (say an hour) over the next day and could capture the expected future revenues via a value function. A typical ADR will be a function of observable parameters in the electricity system. For example these could be the time of day, the air temperature, the state of charge of a battery, the previous hour's dispatch, and the previous hour's electricity price. In node n, each agent i supplies the function c_i^n, the participant constraints $\mathcal{X}_i^n = \mathcal{X}_i(x_i(n-))$ and ADR $R_i^n(\cdot)$ to the system operator who solves

$$\begin{aligned} \mathrm{DP}(n, x(n-)):\ \min\ & \textstyle\sum_i c_i^n(x_i) + L^\top z - \sum_i R_i^n(x_i) \\ \text{s.t. } & \textstyle\sum_i A_{\cdot i} x_i + z \geq d^n, \quad [\pi(n)], \\ & z \in [0, d^n], x_i \in \mathcal{X}_i^n(x_i(n-)) \end{aligned}$$

to yield a dispatch $x(n)$ and prices $\pi(n)$. Demand then pays $\pi(n)^\top d^n$, and each agent i is paid $\pi(n)^\top A_{\cdot i} x_i(n)$. Observe that the system output Ax is now expressed as the sum of agent outputs $\sum_i A_{\cdot i} x_i$. Also note that this is a single-stage deterministic dispatch problem. The system operator does not forecast any values for future demand, but relies on the ADRs to ensure that they are in a position (either by ramping up generation or charging their battery) to meet a future demand peak. Nevertheless, the system operator is responsible for reliability of operations that are manifested via a variety of different operational and security constraints. Such constraints, and aspects of frequency and voltage control within the dispatch interval, should be incorporated directly into the dispatch problem. Determining the right form of these constraints could involve additional (offline) analyses or machine learning models that inform the composition and structure of these constraints.

In this setting, we assume that $R_i^n(\cdot)$ is provided by agent i in "real time" in node n after the history of n (including the demand in node n) has been observed. An alternative "long-lived" ADR could be provided at $n=0$; to make this replicate the real-time ADR for every node n the agent would need to provide an adapted history-dependent ADR sequence for each node $n \in \mathcal{T}$. This sequence will depend of course on what information is observed by each agent in each node. In this paper we take the view that all agents have complete information on other agents incoming states (and the history of these).

It is clear that the prices that are paid after market clearing might not be the same as those assumed by the agents when deciding their ADRs. This motivates the following definition.

Definition 1. Given a scenario tree $\mathcal{T}$ with demand $\{d^n\}_{n\in\mathcal{T}}$, an *ADR partial equilibrium* is a set of prices $\{\pi^*(n)\}_{n\in\mathcal{T}}$, agent actions $\{x^*(n)\}_{n\in\mathcal{T}}$, and ADRs $\{R_i^n(\cdot)\}_{n\in\mathcal{T}}, i \in \mathcal{I}$, such that:

1. *for each* $n \in \mathcal{T}$, $x^*(n)$ *solves DP*$(n, x^*(n-))$ *with shadow prices* $\pi^*(n)$;
2. *for each* $n \in \mathcal{T}$ *and* $i \in \mathcal{I}$, $x_i^*(n)$ *solves*

$$\max_{x_i \in \mathcal{X}_i(x^*(n-))} \{\pi^*(n)^\top A_{\cdot i} x_i - c_i^n(x_i) + R_i^n(x_i)\};$$

3. *for each* $n \in \mathcal{T}$ *that is not a leaf node and* $i \in \mathcal{I}$,

$$R_i^n(x_i^*(n)) = \sum_{m\in n+} \mathbb{P}_i(m \mid n)\{\pi^*(m)^\top A_{\cdot i} x_i^*(m) - c_i^n(x_i^*(m)) + R_i^m(x_i^*(m))\}.$$

An ADR partial equilibrium occurs when the stochastic process of ADRs, dispatch and prices assumed by each agent i yields the same stochastic process of prices and dispatch in market clearing. This is a form of *rational expectations equilibrium* (see Lucas Jr and Prescott (1971)) but also has similarities to *recursive competitive equilibrium* (see Prescott and Mehra (2005), Aiyagari (1994)) and to *mean-field games* (see Lasry and Lions (2007), Gomes et al. (2010)). Unlike these papers, we work with a finite number of competing agents. We make the assumption that the agents do not behave strategically with respect to other agents' actions in the same state of the world (as might happen, e.g., in a Nash equilibrium where agents seek to influence the market price).

The following two-period example illustrates an ADR partial equilibrium.

Example 3. Consider a scenario tree with three nodes and three agents: batteries A and B with capacity 1 MWh and equal efficiency $\eta = 0.8$, and a thermal plant C with unlimited capacity and no ramping constraints. Suppose the marginal cost of plant C with output q is

$$\tilde{c}(q) = \begin{cases} 1, & 0 \le q \le 2, \\ 2q-3, & 2 < q \le 3, \\ 3, & q > 3. \end{cases}$$

At time zero in node $n=0$, demand is $d(0)=3.5$ and each battery holds 1 MWh of charge at the start of the period. At time 1, two demand outcomes are possible, defined by successor nodes $n=1$ with $d(1)=2$ and $n=2$ with $d(2)=4$. Agent A assigns $\mathbb{P}(1)=0.6$, $\mathbb{P}(2)=0.4$, and agent B assigns $\mathbb{P}(1)=0.4$, $\mathbb{P}(2)=0.6$.

The ADR partial equilibrium is defined by charge and discharge actions $v_i(n)$, $u_i(n)$ of A and B, and generation actions $q(n)$ of C, prices $\pi(n)$, and ADRs for A, B and C. The actions and prices are as follows.

$$\begin{array}{lll}
v_A(0) = 0, & v_A(1) = 0, & v_A(2) = 0; \\
v_B(0) = 0, & v_B(1) = 0, & v_B(2) = 0; \\
u_A(0) = 1, & u_A(1) = 0, & u_A(2) = 0; \\
u_B(0) = 0, & u_B(1) = 1, & u_B(2) = 1; \\
q(0) = 2.5, & q(1) = 1, & q(2) = 3; \\
\pi(0) = 2.0, & \pi(1) = 1, & \pi(2) = 3;
\end{array}$$

The ADR for A in node 0 is the expected future value of storage y_A at prices 1 and 3 is $1.8y_A$, for $y_A \in [0,1]$. Similarly the ADR for B in period 0 is $2.2y_B$, for $y_B \in [0,1]$. In nodes 1 and 2 the ADR of each battery is assumed to be identically 0. The ADR for C in all nodes is also 0.

We now verify that this solution is an ADR partial equilibrium, by checking each of the three conditions in the definition. First,

$$\begin{array}{lll}
\text{DP}(0)\text{: min} & \int_0^q \tilde{c}(z)dz - 1.8(1-u_A) - 2.2(1-u_B) & \\
& q + u_A + u_B \geq 3.5 & [\pi] \\
& u_A, u_B \in [0,1], q \geq 0 &
\end{array}$$

has optimal solution defined by thermal dispatch $q(0)=2.5$, discharge $u_A(0)=1$, $u_B(0)=0$, price $\pi=2.0$, so the solution satisfies condition (1). For node $n=1$ we have $\pi(1)=1$, and marginal cost of generation at net demand of 1 is \$1/MWh. For node $n=2$ we have that $\pi(n)=3$, and the marginal cost of generation at net demand of 3 is \$3/MWh. So thermal dispatch satisfies condition (2). Each battery satisfies condition (2) also: $u_A(1)=u_A(2)=0$ is the only feasible solution for A in the second stage, and the action $u_B(1)=u_B(2)=1$ maximizes πu_B plus a zero future reward in each node. Finally at time 0, for $y_A, y_B \in [0,1]$, $R_A^0(y_A)=1.8y_A$ and $R_B^0(y_B)=2.2y_B$ are easily shown to be expected future rewards of batteries given prices and agent probabilities, so the solution satisfies condition (3). $\square$

It is not clear for a scenario tree $\mathcal{T}$ with demand $\{d^n\}_{n\in\mathcal{T}}$ that an ADR partial equilibrium will always exist. If there are additional instruments that enable agents to speculate on future energy prices then a necessary condition for the existence of an ADR partial equilibrium is that the agents share a common probability distribution. This statement is a form of no-arbitrage condition that is made precise as follows.

Definition 2. An *Arrow-Debreu security* indexed on node $m \in \mathcal{T} \setminus \{0\}$ is a financial instrument acquired in predecessor node $m-$ that pays \$1 in node m.

Assumption 2. *In each node $n \in \mathcal{T}$ apart from leaf nodes, there is a complete market for Arrow-Debreu securities indexed on each successor of n.*

Definition 3. A *complete-market ADR partial equilibrium* is an ADR partial equilibrium together with a price $\mu(m)$ and a set of trades $W_i(m)$ for each $m \in \mathcal{T} \setminus \{0\}$, so that

$$0 \leq \sum_{i \in \mathcal{I}} W_i(m) \perp \mu(m) \geq 0, \quad m \in \mathcal{T} \setminus \{0\},$$

and in each node n each agent i maximizes

$$\pi^*(n)^\top A_{\cdot i} x_i - c_i^n(x_i) - \sum_{m \in n+} \mu(m) W_i(m) + \sum_{m \in n+} \mathbb{P}_i(m \mid n) W_i(m) + R_i^n(x_i).$$

As we show in Example 3, it is not necessary for agents to share the same probability measure in an ADR partial equilibrium. However Assumption 2 makes this a property of every complete-market ADR partial equilibrium.

Theorem 1. *Suppose Assumption 2 holds, and a set of prices $\{\pi^*(n)\}_{n \in \mathcal{T}}$, agent actions $\{x^*(n)\}_{n \in \mathcal{T}}$, and ADRs $\{R_i^n(\cdot)\}_{n \in \mathcal{T}}$ exists forming a complete-market ADR partial equilibrium (with prices $\mu(m)$ and trades $W_i(m)$, $m \in \mathcal{T} \setminus \{0\}$). Then $\mathbb{P}_i$ is the same for each agent $i \in \mathcal{I}$.*

Proof. Consider an arbitrary node $n \in \mathcal{T}$ that is not a leaf node. Suppose $\mu(m)$ is the price for an Arrow-Debreu security indexed on $m \in n+$. Suppose for some agent i $\mathbb{P}_i(m) > \mu(m)\mathbb{P}_i(n)$. Then agent i could make an infinite expected reward from buying infinitely many Arrow-Debreu securities indexed on m. Thus $\mathbb{P}_i(m) \leq \mu(m)\mathbb{P}_i(n)$. A similar argument shows that $\mathbb{P}_i(m) \geq \mu(m)\mathbb{P}_i(n)$. Since $\mathbb{P}_i(0) = 1$, this shows that $\mathbb{P}_i(m) = \prod_{l \in P(m)} \mu(l)$ where $P(m)$ is the set comprising m and the path of nodes back to node 0, where we set $\mu(0) = 1$. $\square$

Theorem 1 assumes that all agents are risk neutral. It has a natural extension to agents endowed with coherent risk measures, where the common probability distribution emerges from trading Arrow-Debreu securities to yield a social coherent risk measure defining this distribution (see Ferris and Philpott (2022)).

Theorem 1 states that agents sharing the same probability distribution is a necessary condition for a complete-market ADR partial equilibrium. In the next section we show that a common probability distribution is a sufficient condition for existence of an ADR partial equilibrium, without Assumption 1, but subject to an assumption on the separability of the subdifferential of a social Bellman function.

4. Welfare theorems for ADRs

This section demonstrates the existence of an ADR equilibrium that solves the social optimization problem. The idea underlying the construction is quite simple. We assume that agents have access to all parameters defining the social optimization problem (including other agents' costs and capacities) and share the same probability distribution as the system operator. Agents then solve the social optimization optimization problem, and construct their ADRs by decomposing the social Bellman function. This decomposition requires some separability assumptions on the problem data and constraint qualifications on the optimization problems to be solved, and its construction is rather technical. The reader could skip this section without losing the thread of the paper.

To simplify the argument, we focus first on a deterministic setting, and then show how the construction extends naturally to a scenario tree. In the deterministic setting, recall the multiperiod economic dispatch problem over T periods, defined as follows:

$$\begin{aligned} \text{SOP:}\ \min\ & \textstyle\sum_{t=1}^{T}(c^t(x(t)) + L^\top z(t)) \\ \text{s.t.}\ & Ax(t) + z(t) \geq d^t, \quad t = 1,2,\ldots,T, \\ & z(t) \in [0, d^t], \quad x(t) \in \mathcal{X}^t, \end{aligned}$$

where $\mathcal{X}^t = \mathcal{X}(x(t-1))$ provides the linkage between periods. SOP can be solved by dynamic programming. The stage problem at time t is

$$\begin{aligned} \text{SSP}(t)\text{:}\ \min\ & c^t(x) + L^\top z + C^{t+1}(x) \\ \text{s.t.}\ & Ax + z \geq d^t, \quad [\pi(t)], \\ & z \in [0, d^t], \quad x \in \mathcal{X}^t. \end{aligned}$$

Here $\pi(t)$ are the optimal Lagrange multipliers on the demand constraint, and $C^{t+1}(x)$ denotes an optimal future cost incurred from the end of stage t if the action in stage t is the vector x.

Assumption 3 **(CQ)**. *We assume throughout this section that c^t and C^t are convex functions of x for every t with $C^{T+1}(x) = 0$, that $\operatorname{dom} C^t$ is the whole space, and $\operatorname{ri}\operatorname{dom} c^t \cap \operatorname{ri}\mathcal{X}^t \neq \emptyset$.*

Suppose that a social planner, instead of solving SOP, solves a sequence of T problems of the form SSP(t). If $C^{t+1}(x)$ is the true Bellman function then this solution will recover the optimal solution of SOP. In the stage problem the system marginal cost is $\pi(t)$, so under uniform pricing, load should pay this price at the margin and generators should be paid at this price.

We would like to apply a Lagrangian argument to express the solution to SSP(t) as an ADR partial equilibrium. In general, however, $C^{t+1}(x)$ will not be separable into a sum of functions $\sum_{i\in\mathcal{I}} C_i^{t+1}(x_i)$, which appears to make such an argument impossible.

However, to see when this is possible, consider a Lagrangian for SSP(t):

$$\mathcal{L}(x,z,\pi) = c^t(x) + L^\top z + C^{t+1}(x) + \pi^\top(d^t - Ax - z).$$

Theorem 2. *Suppose $f(x) := c^t(x) - \pi^\top Ax + C^{t+1}(x)$ and Assumption 3 holds. Then (x,z) solves SSP(t) if and only if there is some π such that*

$$0 \le Ax + z - d^t \perp \pi \ge 0, \tag{3}$$

$$0 \in L - \pi + \mathcal{N}_{[0,d^t]}(z), \tag{4}$$

and

$$0 \in \partial f(x) + \mathcal{N}_{\mathcal{X}^t}(x) \tag{5}$$

hold, where $\partial f(x)$ is the subdifferential of f at x, and $\mathcal{N}_{\mathcal{X}^t}(x)$ is the normal cone to $\mathcal{X}^t$ at x. In this case, x minimizes f over $x \in \mathcal{X}^t$.

Proof. We just apply the saddlepoint optimality conditions. When is (x,z,π) a saddle point of the Lagrangian $\mathcal{L}(x,z,\pi)$? Clearly π maximizes $L(x,z,\pi)$ over $\pi \ge 0$ is equivalent to (3). For the minimization of the Lagrangian, observe that we can minimize over z separately, giving (4). This means the remaining optimization seeks to minimize f over $x \in \mathcal{X}^t$. Assumption 3 guarantees that x is optimal for this problem if and only if (5) holds. But these conditions are equivalent to (x,z) solving SSP(t). $\square$

We wish to decompose (5) by agent, and in this deterministic case we need the following separability assumption:

Assumption 4. *$c^t(x) = \sum_i c_i^t(x_i)$ and $\mathcal{X} = \prod_i \mathcal{X}_i$.*

Since $\partial \sum_i c_i^t(x_i) = \prod_i \partial c_i^t(x_i)$ and $\mathcal{N}_\mathcal{X}(x) = \prod_i \mathcal{N}_{\mathcal{X}_i}(x_i)$, (5) is true if there exists a subgradient $g(x)$ of C^{t+1} at x with

$$0 \in \prod_i \partial c_i^t(x_i) - \pi^\top A + g(x) + \prod_i \mathcal{N}_{\mathcal{X}_i}(x_i).$$

Consider a fixed $\hat{x} \in \mathcal{X}$ and let $F(x)$ be any convex function of x. For each $i \in \mathcal{I}$ we define a univariate *conditional* function $F_{i,\hat{x}_{-i}}(x_i)$ by replacing all components j of x except the ith by $\hat{x}_j$. Thus (using $\hat{x}_{-i}$ notation) we have $F_{i,\hat{x}_{-i}}(x_i) = F(x_i, \hat{x}_{-i})$.

Now consider the conditional functions $C^{t+1}_{i,\hat{x}_{-i}}(x_i)$, $i \in \mathcal{I}$, defined from the (social) optimal future cost function $C^{t+1}(x)$. For $\hat{x} \in \mathcal{X}^t$ and $\pi \ge 0$ we define the agent stage problem:

$$\begin{aligned} \mathrm{P}_i^t(\pi, \hat{x}_{-i}) \colon \ \max\ & (\pi^\top A)_i x_i - c_i^t(x_i) - C^{t+1}_{i,\hat{x}_{-i}}(x_i) \\ \text{s.t. } & x_i \in \mathcal{X}_i^t. \end{aligned}$$

The next result is used to show that if (x^*, z^*) solves SSP(t) then there exists π such that each component x_i^* will solve $\mathrm{P}_i^t(\pi, x_{-i}^*)$.

Proposition 1. *Let $\pi \ge 0$ be given and Assumption 4 hold. If x^* minimizes $f(x) := \sum_{i \in \mathcal{I}} c_i^t(x_i) - (\pi^\top A)_i x_i + C^{t+1}(x)$ over $x \in \mathcal{X}^t$, then x_i^* solves $P_i^t(\pi, x_{-i}^*)$.*

Proof. Suppose x_i^* does not maximize $(\pi^\top A)_i x_i - c_i^t(x_i) - C^{t+1}_{i,x^*_{-i}}(x_i)$. Then there is some $\tilde{x}_i \in \mathcal{X}_i^t$ with

$$c_i^t(\tilde{x}_i) - \left(\pi^\top A\right)_i \tilde{x}_i + C^{t+1}_{i,x^*_{-i}}(\tilde{x}_i) < c_i^t(x_i^*) - \left(\pi^\top A\right)_i x_i^* + C^{t+1}_{i,x^*_{-i}}(x_i^*).$$

It follows that

$$\begin{aligned}&\sum_{j\in\mathcal{I}\setminus\{i\}} (c_j^t(x_j^*) - \left(\pi^\top A\right)_i x_j^*) + c_i^t(\tilde{x}_i) - \left(\pi^\top A\right)_i \tilde{x}_i + C^{t+1}(\tilde{x}_i, x^*_{-i}) \\ &< \sum_{j\in\mathcal{I}\setminus\{i\}} (c_j^t(x_j^*) - \left(\pi^\top A\right)_i x_j^*) + c_i^t(x_i^*) - \left(\pi^\top A\right)_i x_i^* + C^{t+1}(x^*)\end{aligned}$$

and $(\tilde{x}_i, x^*_{-i}) \in \mathcal{X}^t$ which violates the optimality of x^*. $\square$

We would like to establish a converse result that states that optimal solutions of P_i^t will also solve SSP(t) as long as the complementary slackness conditions (3) and (4) hold. This requires a connection between the subdifferential of each $C_{i,x^*_{-i}}$ and the subdifferential of C. We have the following result.

LEMMA 1. *Let $\hat{x} \in \mathcal{X}$. If $g(\hat{x}) = [g_1(\hat{x}) \dots g_n(\hat{x})]^\top \in \partial C(\hat{x})$ then for each $i \in \mathcal{I}$, $g_i(\hat{x}) \in \partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$, so $\partial C(\hat{x}) \subset \prod_i \partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$.*

Proof. We have for any x

$$\begin{aligned} C(x) &\geq C(\hat{x}) + g(\hat{x})^\top (x - \hat{x}) \\ &= C(\hat{x}) + \sum_{i\in\mathcal{I}} g_i(\hat{x})(x_i - \hat{x}_i)\end{aligned}$$

Thus

$$\begin{aligned} C_{i,\hat{x}_{-i}}(x_i) &= C(x_i, \hat{x}_{-i}) \\ &\geq C(\hat{x}) + g_i(\hat{x})(x_i - \hat{x}_i)\end{aligned}$$

which gives the result. $\square$

If C is differentiable everywhere then for each i, $C_{i,\hat{x}_{-i}}$ is also differentiable. This means that $\partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$ is a singleton $\{g_i(\hat{x})\}$ that defines the unique subgradient that comprises $\partial C(\hat{x})$, so $\partial C(\hat{x}) = \prod_i \partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$. Note also that if C is separable $C(x) = \sum_i f_i(x_i)$, then it also follows that $\partial C(\hat{x}) = \prod_i \partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$. Equality does not hold in general, since it is easy to find nonsmooth convex functions C where $g_i(\hat{x}) \in \partial C_{i,\hat{x}_{-i}}(\hat{x}_i)$ but $g(\hat{x}) \notin \partial C(\hat{x})$. See for example $C(x,y) = |x-y| + |x+y-2|$.

We can now establish the converse to Proposition 1.

PROPOSITION 2. *Let $\pi \geq 0$, $\hat{x} \in \mathcal{X}^t$ be given and suppose $\partial C^{t+1}(x) = \prod_{i\in\mathcal{I}} \partial C^{t+1}_{i,\hat{x}_{-i}}(x_i)$ and Assumptions 3 and 4 hold. If for each $i \in \mathcal{I}$, $\hat{x}_i$ solves $P_i^t(\pi, \hat{x}_{-i})$ then $\hat{x}$ minimizes $\sum_{i\in\mathcal{I}}(c_i^t(x_i) - (\pi^\top A)_i x_i) + C^{t+1}(x)$ over $x \in \mathcal{X}^t$.*

Proof. If $\hat{x}_i$ maximizes $(\pi^\top A)_i x_i - c_i^t(x_i) - C_{i,\hat{x}_{-i}}^{t+1}(x_i)$ over $x_i \in \mathcal{X}_i^t$ then $\hat{x}_i$ minimizes $c_i^t(x_i) - (\pi^\top A)_i x_i + C_{i,\hat{x}_{-i}}^{t+1}(x_i)$ over $x_i \in \mathcal{X}_i^t$ so under Assumption 3

$$0 \in \partial c_i^t(\hat{x}_i) - \left(\pi^\top A\right)_i + \partial C_{i,\hat{x}_{-i}}^{t+1}(\hat{x}_i) + \mathcal{N}_{\mathcal{X}_i^t}(\hat{x}_i)$$

which under Assumption 4 gives

$$0 \in \prod_i \partial c_i^t(\hat{x}_i) - A^\top \pi + \partial C^{t+1}(\hat{x}) + \prod_i \mathcal{N}_{\mathcal{X}_i^t}(\hat{x}_i)$$

implying the optimality of $\hat{x}$. □

Propositions 1 and 2 can now be combined to give the following welfare therorems.

Theorem 3. *Suppose (x^*, z^*) solves SSP. If Assumptions 3 and 4 hold then there exists π such that x_i^* solves $P_i^t(\pi, x_{-i}^*)$, $i \in \mathcal{I}$.*

Proof. By Theorem 2 there exists π such that (3), (4) and (5) hold. For this π, (5) and Proposition 1 then implies that x_i^* solves $\mathrm{P}_i^t(\pi, x_{-i}^*)$. □

Theorem 4. *Suppose $(\hat{x}, \hat{z}, \pi)$ satisfy (3) and (4), Assumptions 3 and 4 hold and $\partial C^{t+1}(x) = \prod_{i\in\mathcal{I}} \partial C_{i,\hat{x}_{-i}}^{t+1}(x_i)$. Then if $\hat{x}_i$ solves $P_i^t(\pi, \hat{x}_{-i})$, $i \in \mathcal{I}$ it follows that $(\hat{x}, \hat{z})$ solves SSP(t).*

Proof. If $\hat{x}_i$ solves $\mathrm{P}_i(\pi, \hat{x}_{-i})$ then Proposition 2 implies (5) whereby the result follows from Theorem 2. □

We can extend these theorems in a straightforward way to a setting in a scenario tree $\mathcal{T}$ with demand $\{d^n\}_{n\in\mathcal{T}}$ and probability measure $\mathbb{P}$. This gives a stochastic formulation of SOP:

$$\begin{aligned}
\text{SSOP:}\ \min\ & \textstyle\sum_{n\in\mathcal{T}} \mathbb{P}(n)(c^n(x(n)) + L^\top z(n)) \\
\text{s.t.}\ & \textstyle\sum_i A_{\cdot i} x_i(n) + z(n) \geq d^n, \quad [\pi(n)], \quad n \in \mathcal{T}, \\
& z(n) \in [0, d^n], \quad n \in \mathcal{T}, \\
& x_i(n) \in \mathcal{X}_i(x(n-)), \quad n \in \mathcal{T} \setminus \{0\}.
\end{aligned}$$

The stochastic optimal dispatch problem can then be solved (in principle at least) in a recursive fashion. The social planning stage problem at node n is now

$$\begin{aligned}
\text{SSP}(n, \bar{x}):\ C^n(\bar{x}) = \min\ & \textstyle\sum_i c_i^n(x_i) + L^\top z + \sum_{m\in n+} \mathbb{P}(m \mid n) C^m(x) \\
\text{s.t.}\ & \textstyle\sum_i A_{\cdot i} x_i + z \geq d^n, \quad [\pi(n)], \\
& z \in [0, d^n], \quad x_i \in \mathcal{X}_i^n := \mathcal{X}_i(\bar{x}).
\end{aligned}$$

Here $\mathbb{P}(m \mid n)$ is assumed to be 0 when n is a leaf node of $\mathcal{T}$ and $C^n(x)$ denotes the minimum expected cost from the optimal policy implemented in node n and all its successors when the incoming state is x,

Definition 4. $C^m(x)$ is *gradient separable* at $\hat{x}(n)$ if for every x,

$$\partial C^m(x) = \prod_{i\in\mathcal{I}} \partial C^m_{i,\hat{x}_{-i}}(x_i)$$

where $C^m_{i,\hat{x}_{-i}}(x_i) = C^m(x_i, \hat{x}_{-i}(n))$, $i \in \mathcal{I}$.

The deterministic analysis can be applied to $\mathrm{SSP}(n, \bar{x})$, but we assume now that every agent shares the same probability distribution as the social planner.

Assumption 5. *For every $i \in \mathcal{I}$, $\mathbb{P}_i = \mathbb{P}$.*

For $\hat{x} \in \mathcal{X}$ and $\pi \geq 0$ we define the agent stage problem

$$\begin{array}{rl}\mathrm{P}^n_i(\pi(n), \hat{x}_{-i})\colon & \max\ \pi(n)^\top A_{\cdot i} x_i - c^n_i(x_i) + S^n_i(x_i) \\ & \text{s.t. } x_i \in \mathcal{X}_i,\end{array}$$

where $S^n_i(x_i) = -\sum_{m\in n+} \mathbb{P}(m \mid n) C^m_{i,\hat{x}_{-i}}(x_i)$. We also write complementary slackness conditions

$$0 \leq Ax(n) + z(n) - d^n \perp \pi(n) \geq 0, \quad n \in \mathcal{T}, \tag{6}$$

and

$$0 \in L - \pi(n) + \mathcal{N}_{[0,d^n]}(z(n)), \quad n \in \mathcal{T}. \tag{7}$$

The following assumption extends Assumption 3 to every node in the scenario tree.

Assumption 6. *For every $n \in \mathcal{T}$, c^n and C^n are convex functions of x, $\mathrm{dom}\, C^n$ is the whole space, Assumption 1 holds, and $\mathrm{ri}\,\mathrm{dom}\, c^n \cap \mathrm{ri}\, \mathcal{X}^n \neq \emptyset$.*

The same arguments as used in the proofs of Theorem 3 and Theorem 4 yield the following welfare theorems.

Theorem 5. *Suppose $(x^*(n), z^*(n))_{n\in\mathcal{T}}$ solves SSOP, Assumption 5 and Assumption 6 hold, and we let $S^n_i(x_i) = -\sum_{m\in n+} \mathbb{P}(m \mid n) C^m_{i,x^*_{-i}}(x_i)$, $i \in \mathcal{I}$. Then there exist $\pi(n)$ such that $x^*_i(n)$ solves $P^n_i(\pi(n), x^*_{-i})$, $i \in \mathcal{I}$.*

Theorem 6. *Suppose $(\hat{x}(n), \hat{z}(n), \pi(n))_{n\in\mathcal{T}}$ satisfy (6) and (7), Assumption 5 and Assumption 6 hold, and for every $n \in \mathcal{T}$, $C^n(x)$ is gradient separable at $\hat{x}(n-)$. Then if $\hat{x}_i$ solves $P_i(\pi(n), \hat{x}_{-i})$, $i \in \mathcal{I}$ it follows that $(\hat{x}, \hat{z})$ solves SSOP.*

Theorem 5 allows us to establish the existence of an ADR partial equilibrium. In this equilibrium all agents assume the same probability distribution as the social planner and choose ADRs based on the expected future cost functions of the social planner.

Theorem 7. *Suppose Assumption 5 and Assumption 6 hold and $(x^*(n), z^*(n), C^n)_{n\in\mathcal{T}}$ solves SSOP, with each Bellman function C^n gradient separable at $x^*(n-)$, $n\in\mathcal{T}$. If $S_i^n(x_i) = -\sum_{m\in n+}\mathbb{P}(m\mid n)C^m_{i,x^*_{-i}}(x_i)$, $i\in\mathcal{I}$, then there exists $\pi^*(n)$ so that $(x^*(n), \pi^*(n), S_i^n(\cdot))_{n\in\mathcal{T}}$ is an ADR partial equilibrium.*

Proof. By Theorem 5 there exists $\pi^*(n)$ satisfying (6) and (7) such that $x_i^*(n)$ solves $\mathrm{P}_i^n(\pi^*, x^*_{-i})$, $i\in\mathcal{I}$, where $S_i^n(x_i) = -\sum_{m\in n+}\mathbb{P}(m\mid n)C^m_{i,\hat{x}_{-i}}(x_i)$. We show that $\{\pi^*(n), x^*(n), S_i^n(\cdot)\}_{n\in\mathcal{T}}$ is an ADR partial equilibrium, where $S_i^n(\cdot)$ takes the place of $R_i^n(\cdot)$ in the definition (since $\mathbb{P}$ replaces $\mathbb{P}_i$). The second property of the definition is immediate by Theorem 5. The optimality conditions for $\mathrm{P}_i^n(\pi, x^*_{-i})$, $i\in\mathcal{I}$ along with (6) and (7) show that $x^*(n)$ solves $\mathrm{DP}(n, x^*(n-))$ with Bellman function $\sum_i S_i^n$ and prices $\pi^*(n)$. Finally since $(x^*(n), z^*(n)), C^n)_{n\in\mathcal{T}}$ solves SSOP, and C^n is gradient separable at $x^*(n-)$, and satisfies the system dynamic programming recursion at every $n\in\mathcal{T}$, the definition of $S_i^n(x_i)$ ensures that it satisfies the third condition in the definition of an ADR equilibrium. $\square$

5. Illustrative example

In this section we study a particular dispatch problem at a single node with a single generator and one battery operator over a 24-hour period. The problem SOP becomes the example problem

$$\begin{aligned}
\text{EP: } \min\ & \sum_{t=1}^{24}\left(c^t(q(t)) + Lz(t)\right) \\
\text{s.t. } & q(t) + u(t) - v(t) + z(t) \geq d^t, \\
& q(0) = q^0, \quad q(t)\in\mathcal{Q}(q(t-1)), \\
& y_i(0) = y^0, \quad (y_i(t), u(t), v(t))\in\mathcal{Y}(y(t-1)), \\
& z(t)\in[0, d^t], \quad t = 1, 2, \ldots, 24,
\end{aligned}$$

where

$$\mathcal{Q}(\bar{q}) = \left\{q \mid 0\leq q\leq q^{\max}, q-\bar{q}\leq\rho\right\}, \tag{8}$$

$$\mathcal{Y}(\bar{y}) = \left\{(y,u,v)\mid 0\leq y\leq E, 0\leq u\leq r, 0\leq v\leq s, y = \bar{y} - u + \eta v\right\}. \tag{9}$$

Here $q(t)$ denotes generation dispatched in period t, and $y(t)$ is storage of energy in the battery at the end of period t. These variables have initial values q^0 and y^0 at the start of the day. The dispatch $q(t)$ is constrained by a ramp-up limit ρ and capacity $q^{\max}$ and incurs a cost of $c^t(q(t))$. Battery storage is increased by charging using variable v and decreased by discharging an amount u. Round trip losses are modeled using the factor η, which multiplies v. Charging and discharging rates are limited by the parameters s and r respectively, and the battery has a maximum charge E. The total amount generated should meet demand d^t. Any shortfalls $z(t)$ are penalized at a value of lost load L.

In the example, the generator has an increasing marginal cost with ten steps defined by Table 1. The values chosen for the other parameters are given in Table 2. Meeting the duck-curve shaped

energy tranche (MWh)	5.0	5.0	5.0	5.0	5.0	5.0	10.0	10.0	10.0	10.0
marginal cost ($/MWh)	10.0	20.0	30.0	40.0	50.0	70.0	90.0	110.0	150.0	200.0

Table 1 **Marginal cost of generator**

$q^{\max} = 70.0$	$E = 30.0$	$\eta = 1.0$
$r = 15.0$	$s = 15.0$	$\rho = 10.0$
$L = 1000.0$	$q^0 = 35.0$	$y^0 = 0.0$

Table 2 **Parameter values for example**

demand in our model will require the generator to ramp up in periods 14 to 18 and the battery to discharge in periods 19 through 21.

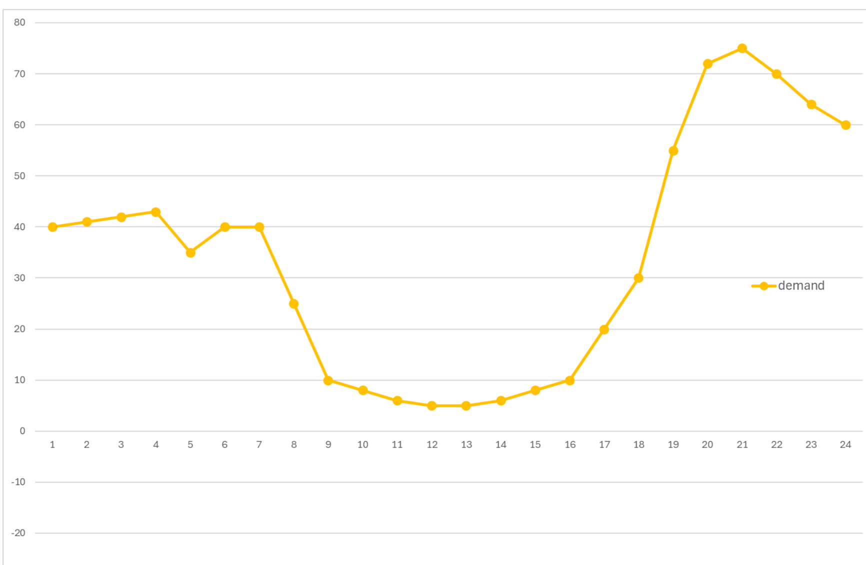

Figure 1 **Values of d^t for $t = 1, 2, \ldots, 24$.**

We first solve the perfect foresight model EP where the demand is known ex-ante for all $t = 1, 2, \ldots, 24$. Given perfect foresight, the model can be solved as a deterministic linear program covering all 24 hours of operations. The optimal solution to EP has cost $48,470, with optimal dispatch and battery charge shown in Figure 2.

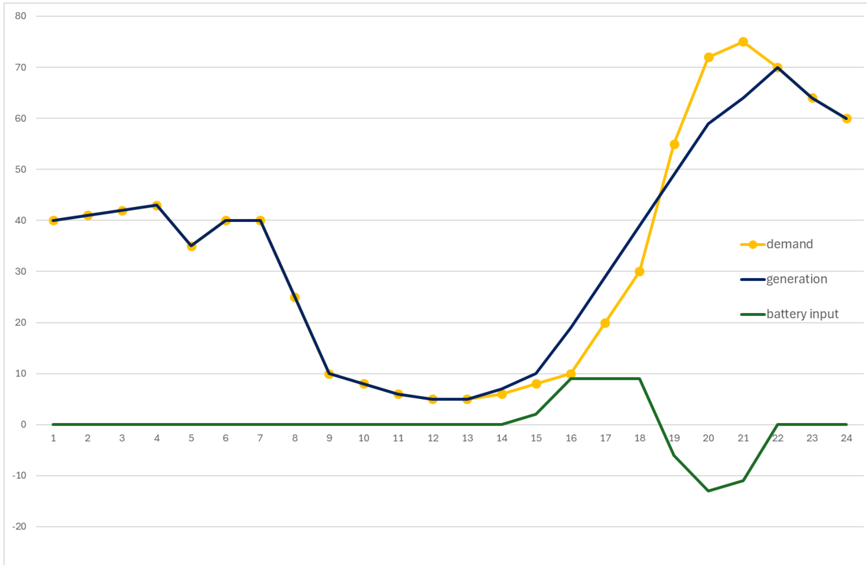

Figure 2 **Solution of DP showing generation x superimposed on demand, and battery net input $v-u$ for $t=1,2,\ldots,24$.**

The problem EP can alternatively be considered as a discrete-time optimal control problem. It can be solved using dynamic programming using the recursion:

$$\begin{aligned} F^{t-1}(q(t-1), y(t-1)) = \min\ & c^t(q) + Lz + F^t(q,y) \\ \text{s.t.}\ & a + u - v + z \geq d^t, \\ & q \in \mathcal{Q}(q(t-1)), \\ & (y,u,v) \in \mathcal{Y}(y(t-1)), \\ & z \in [0, d^t], \end{aligned}$$

where $F^{24}(q,y) = 0$ and $q(0) = q^0$, $y(0) = y^0$.

This yields future cost functions F^t that can be used to solve the sequence of optimization problems DP(1,$q(0)$,$y(0)$), DP(2,$q(1)$,$y(1)$), ..., DP(24,$q(23)$,$y(23)$) where $q(0)$ and $y(0)$ are given and F^t is substituted directly into DP. (Here we use the notation F to denote a *deterministic* future cost function.) These problems involve no explicit lookahead forecasts, but when solved in sequence they replicate the socially optimal dispatch obtained by directly solving EP.

In practice, the system operator does not have perfect foresight. The system operator could estimate the parameters of a stochastic process of future demand (and other parameters) and use these to approximate an *expected* future cost function denoted $C^t(q,y)$. This is the approach followed by the system operator to evaluate the expected future value of stored water in the Brazilian electricity system (Diniz et al. 2018). The value is based on a centrally determined stochastic model of inflows, and does not explicitly incorporate differing views of market participants.

Experiment 1

The example above can be extended to accommodate uncertainty in demand by adding independent equally likely noise terms chosen from $\{-4,-2,0,2,4\}$ at each t to the demand in Figure 1.

We then solve the system optimization problem using stochastic dynamic programming (SDP) and extract expected future cost functions $C^t(q,y)$ at each stage. The SDP solution has optimal expected cost of $C^1(q^0,y^0) = \$52,377$. This is the socially optimal solution for a perfectly competitive equilibrium where every agent shares the same probability distribution of future demand.

We can investigate alternative agent policies by simulation. In all simulations we use 10,000 sample paths with common random numbers. Recall that the perfect foresight solution gives a cost of \$48,470 in the deterministic case. Simulating the policy defined by the deterministic future cost functions F^t with the noisy demand defined above yields an estimated expected cost of \$54,255 with standard error \$57. By contrast, simulating the socially optimal solution computed using SDP gives an expected cost of \$52,406 with standard error \$50.

As discussed above, we propose that market participants provide ADRs defined by their own future reward functions. These could be provided in real time at the start of each period, or be long-lived over many future trading periods. Here we investigate the latter option in the example problem, where at the start of the day the generator agent provides functions $S_g^t(q)$, $t = 1, 2, \ldots, 24$ and the battery operator provides functions $S_b^t(y)$, $t = 1, 2, \ldots, 24$, and then the system operator solves DP(1,$q(0)$,$y(0)$), DP(2,$q(1)$,$y(1)$), ..., DP(24,$q(23)$,$y(23)$) using

$$C^t(q,y) = -S_g^t(q) - S_b^t(y).$$

First, suppose that each agent solves the deterministic system problem with expected demand and computes the generation $\hat{q}(t)$ and finishing storage $\hat{y}(t)$ for each interval t as well as the deterministic future cost function $F^t(q,y)$. The agents then compute univariate value functions $F_g^t(q) = F^t(q,\hat{y}(t))$ and $F_b^t(y) = F^t(\hat{q}(t),y)$ respectively, with each assuming that the other agent follows the socially optimal dispatch from the deterministic case. The system operator then solves DP(1,$q(0)$,$y(0)$), DP(2,$q(1)$,$y(1)$), ..., DP(24,$q(23)$,$y(23)$) using

$$F^t(q,y) = F_g^t(q) + F_b^t(y).$$

Using the data in the above example, this approach gave ADRs for each agent which yielded a social cost of \$48,470 when simulated in the perfect foresight scenario, the same as the deterministic social optimum as predicted by Theorem 4. Simulating these ADRs with noise added to demand yielded an estimated expected social cost of \$55,430 (standard error \$76). This turns out to be worse than the \$54,255 obtained by simulating with the function $F^t(q,y)$ rather than the separable ADRs that approximate $F^t(q,y)$. In other words, when the system operator and the agents all use a deterministic simplification to compute future cost functions, there is an advantage to the centralized solution. Observe that this need not be the case in general, as $F^t(q,y)$ is itself an approximation of the true social cost function $C^t(\hat{q}(t),y)$.

Second, suppose that the agents are more sophisticated and solve the stochastic social planning problem to yield $C^t(q,y)$. Each agent then computes respective univariate value functions $C_g^t(q) = C^t(q,\hat{y}(t))$, and $C_b^t(y) = C^t(\hat{q}(t),y)$, using $\hat{q}(t)$ and finishing storage $\hat{y}(t)$ from the deterministic problem solved with expected demand. In the example, this approach gave ADRs for each agent which yielded an estimated social cost of \$52,688 (standard error \$55) when simulated with 10,000 sample paths. This compares well with the social optimum policy that yields \$52,406. With that said, these ADRs have higher cost since they are constructed to approximate $C^t(q,y)$ only at those values visited in the deterministic solution.

In order to recover the example's socially optimal solution using ADRs, we would need to supply a value function for each agent at each stage that varies with the noise outcome in that stage, and history of dispatch up to that point. At each stage t, when presented with noise outcome ω, the agent g would provide the system operator with $S_g^{t+1}(q) = -C^{t+1}(q,\hat{y}(t,\omega))$, and the agent b would provide the system operator with $S_b^{t+1}(y) = -C^{t+1}(\hat{q}(t,\omega),y)$, where $(\hat{q}(t,\omega),\hat{y}(t,\omega))$ is the state at the end of period t in scenario ω when the social optimal policy is applied. The ADR in this case is provided in real time at the beginning of each trading period, assuming knowledge of the demand in that period. If each participant's ADR for all periods must be supplied to the system operator at the beginning of the day, then to recover system optimality a univariate function $S_g^t(q)$ would need to be specified at each t for all possible $\hat{y}(t)$ values, along with a univariate function $S_b^t(y)$ specified at each t for all possible $\hat{q}(t)$ values.

The expected cost resulting from the four simulated policies is summarized in Table 3.

		Dispatch Input	
		Separable ADRs	Social Cost-to-Go
Cost Function Estimate	Deterministic	\$55,430	\$54,255
	Stochastic	\$52,688	\$52,406

Table 3 Expected cost under simulated policies

Experiment 2

We repeated the above experiment with different noise terms added to the demand outcomes. The noise terms are now chosen from {-4, -2, 0, 8, 16} with probabilities {0.2, 0.2, 0.5, 0.05, 0.05} which gives the same mean demand but higher variance. The optimal policy now has an expected cost of \$57,438. Simulation with 10000 sample paths gives an estimated expected cost of this policy equal to \$57,506 (standard error 102)

The solution for this example in the deterministic case is identical to that of Experiment 1, yielding a social cost of \$48,470. Simulating the deterministic dynamic programming policy computed assuming deterministic demand with 10000 (high variance) sample paths gives an estimated expected cost equal to \$60,984 (standard error 112).

Simulating the (separated) ADRs derived from the deterministic dynamic programming policy with 10000 (high variance) sample paths gives an estimated expected cost equal to \$63,895 (standard error 126).

Finally we simulate the ADRs derived from the stochastic Bellman functions evaluated at $\hat{q}(t)$ and $\hat{y}(t)$ values obtained by the perfect foresight problem (solved with expected demand). Simulating these ADRs with 10000 (high variance) sample paths gives an estimated expected cost equal to \$59,409 (standard error 105). As such, the rank ordering of the four tested policies is the same in example 2 as it was in example 1. However, the higher variance in demand as compared with example 1 has increased the difference in cost between the ADR policy and the socially optimal policy.

6. Further examples

The ADRs we used in the example deal with battery storage and ramping generation when these technologies are operated by different agents. A straightforward extension would consider some agents having a mix of technologies and future value functions that depend on the states of each. In this section we discuss other settings in which ADRs might play a role.

6.1. Classical bids as agent decision rules

Classical offer curves (i.e., supply functions) are a special case of an ADR, but typically do not add any more state-dependent information than is currently available in a conventional dispatch mechanism. In this limited case, the rule consists of data pairs (c_i, q_i) describing the supply function. The ADR is specified by the immediate cost data c_i and the parameterized set $\mathcal{X}_i$ defined in (1). In this case both ρ_i and σ_i are infinite (no explicit ramping constraints) and capacities are given by $q^{\max}$. To construct a more nuanced ADR, a generator might have a forecast of future electricity prices as a function of current observations and define offer curves that depend on these forecasts.

An ADR defined for period t cannot depend arbitrarily on the observed price $\pi(t)$ in period t. To illustrate this, consider a simple form of ADR defined by a supply function offered to a single node convex dispatch model without ramping constraints. Such a function will yield a dispatch of plant with marginal costs below the computed system marginal price, without specifying this price explicitly in the supply curve. The form of dependence of dispatch on the observed price $\pi(t)$ in period t is restricted by the convexity of the dispatch problem. To be clear, an ADR that dispatched 10 units if $\pi(t) \in [0, 50]$ and 5 units if $\pi(t) \in [50, 100]$ would not be acceptable in our framework.

Demand response (also known as peak shaving) refers to demand that is decreased when prices are high. When this is "behind the meter" it must be treated as a variation in net demand by the dispatcher who must use a forecast. On the other hand it can be offered as a demand-side bid to a dispatch model and dispatched by the system operator. This is modelled in dispatch problems using a nonincreasing inverse demand curve that specifies the price in period t when the quantity demanded is x. Demand curves for each consumer can be submitted to the system operator and summed to give a system demand curve for use in a conventional dispatch model (see above).

Demand response can make use of ADRs when electricity is used to make a product that is stored for later sale. In this case, the demand curve offered to the market might depend on the state of storage of the product. When storage is low the demand curve will buy more at each price to replenish stock, and when it is high the demand curve will buy less.

6.2. ADRs for pump-storage hydro plants

Pump-storage hydro plants release water from a reservoir through turbines in peak periods when prices are high, and then refill the reservoir by pumping water uphill when prices are low. These facilities can be optimized using the same ADRs as batteries, possibly over a longer time scale.

6.3. ADRs for flexible demand

Instead of just reducing load, some industrial loads (or even data centers) can shift demand from peak periods to off-peak periods. This flexible demand can be offered to wholesale markets using ADRs. If the industry has a battery or some other mechanism to store energy then the ADR takes a similar form to those for battery storage. This extends to settings where products that use electricity in production can be stored for later sale.

The use of electricity for a particular task can be deferred from a high price period until later when prices are lower. If the task has to be completed by the end of the time horizon then the shifting of load can be optimized using an ADR. The variable $x(t) = (y(t), v(t))$ where the state $y_j(t)$ is the proportion of task j that has been completed, and $v_j(t)$ denotes the electricity consumed by task j in period t where the task requires η_j units of energy. We have

$$y_j(t) = y_j(t-1) + v_j(t)/\eta_j,$$

and the ADR uses a future cost function $-W_j(y_j)$ that is zero when $y_j = 1$.

6.4. ADRs for hydroelectric generators

Although not represented explicitly in our formulation, a decision rule could be defined for hydroelectric generators who release water from reservoirs to generate electricity, and replenish the reservoir contents with (stochastic) inflows. Hydroelectric generators price the release of water by estimating the *expected marginal water value*, which represents the expected opportunity cost of releasing water now rather than in the future. This cost can be viewed as the derivative of a function

$W_j^t(y_j)$ of the same form used to express the future value for batteries. This enables hydroelectric generators to offer decision rules to the system operator as if they were a battery operator.

There is a range of models that can be used to dispatch hydroelectric generators. An isolated reservoir with no inter-temporal constraints can compute an expected marginal water value through dynamic programming. This can be used by the system operator to dispatch the hydro plant efficiently. When hydroelectric stations are located at different points on a river network the marginal water value will vary with time and location. Indeed these values will be Lagrange multipliers to flow balance constraints in a complicated multiperiod optimization problem. In principle, the system operator can use these as a guide to dispatch each station in the river system.

An alternative model cedes control of the hydroelectric river system to the electricity system operator who solves a lookahead problem. The river constraints are incorporated into the dispatch model (like transmission constraints) which is optimized by the system operator accounting for all water released from storage over the time horizon using an end-of-horizon future value function $W_j^T(y_j)$. This hydro-enhanced dispatch model requires a forecast of demand to inform some form of lookahead in the dispatch model.

6.5. ADRs for reserve

Electricity generators and batteries can assign part of their capacity for reserve, and be paid a price for this. This can be incorporated into an ADR dispatch model. The exact form of this model depends on how reserve is defined. We outline a simple model where reserve is spare generation capacity made available in each period by generators to deal with contingencies in that period only. (The model for reserve being offered by batteries is similar.)

Suppose the amount of reserve required in period t is $d^{r,t}$, and at the start of period t the generation levels are $\bar{q}$ and the battery charge levels are $\bar{y}$. Suppose that generator $i \in \mathcal{G}$ is dispatched q_i^r of reserve at cost g_i. The security-constrained dispatch model allows available generation to be split into immediate demand satisfaction and reserve requirements.

$$
\begin{aligned}
\text{SCDP}(t,\bar{x})\colon \min \; & \sum_{i\in\mathcal{G}} (c_i(q_i) + g_i(q_i^r)) + L^\top z - \sum_i R_i^n(x_i), \\
\text{s.t.} \; & \sum_{i\in\mathcal{G}} q_i^r = d^{r,t}, \\
& \sum_{i\in\mathcal{G}} q_i + \sum_{j\in\mathcal{J}} u_j - \sum_{j\in\mathcal{J}} v_j + z \ge d^t, \\
& (q_i, q_i^r) \in \tilde{\mathcal{Q}}_i(\bar{q}_i), \quad i \in \mathcal{G} \\
& (y_j, u_j, v_j) \in \tilde{\mathcal{Y}}_j(\bar{y}_j), \quad j \in \mathcal{J} \\
& z(t) \in [0, d^t],
\end{aligned}
$$

where

$$
\tilde{\mathcal{Q}}_i(\bar{q}_i) = \{(q, q^r) \mid 0 \le q + q^r \le q_i^{\max}, q + q^r - \bar{q}_i \le \rho_i, \bar{q}_i - q - q^r \le \sigma_i\},
$$

$$\tilde{\mathcal{Y}}_j(\bar{y}_j) = \{(y,u,v) \mid 0 \le y \le E_j, 0 \le u \le r_j, 0 \le v \le s_j, y = \bar{y}_j - u + \eta_j v\}.$$

The definition of $\tilde{\mathcal{Q}}_i(\bar{q}_i)$ can include extra constraints on q^r that depend on each generator's plant. Some care is needed in defining the expected future cost.

6.6. ADRs for frequency regulation

Batteries can assign part of their capacity for frequency regulation, and be paid a price for this. Suppose in period t that the total amount of battery capacity required for frequency regulation is F^t. At time t each battery operator j offers some capacity k_j MW at a price of φ_j dollars per MW. The amount of frequency regulation they are dispatched is f_j, which requires them to allocate some of their battery storage to this task. Regulating the frequency involves charging and discharging which consumes energy because of round-trip losses. Suppose this energy is $\psi_j f_j$.

The frequency regulating dispatch model is as follows.

$$\begin{aligned}
\text{FRDP}(t,\bar{x})\colon \min \; & \sum_{i\in\mathcal{G}} c_i(q_i) + \sum_{j\in\mathcal{J}} \varphi_j(t) f_j + L^\top z + C^t(x), \\
\text{s.t.} \; & \sum_{j\in\mathcal{J}} f_j \ge F^t, \\
& \sum_{i\in\mathcal{G}} q_i + \sum_{j\in\mathcal{J}} u_j - \sum_{j\in\mathcal{J}} v_j + z \ge d^t, \\
& q_i \in \tilde{\mathcal{Q}}_i(\bar{q}_i), \quad i \in \mathcal{G}, \\
& (y_j, u_j, v_j, f_j) \in \tilde{\mathcal{Y}}_j(\bar{y}_j), \quad j \in \mathcal{J}, \\
& z(t) \in [0, d(t)],
\end{aligned}$$

where

$$\begin{aligned}
\tilde{\mathcal{Q}}_i(\bar{q}_i) &= \{q \mid 0 \le q \le q_i^{\max},\ q - \bar{q}_i \le \rho_i,\ \bar{q}_i - q \le \sigma_i\}, \\
\tilde{\mathcal{Y}}_j(\bar{y}_j) &= \{(y,u,v,f) \mid 0 \le y \le E_j, \quad 0 \le u \le r_j, \quad 0 \le v \le s_j, \\
& \qquad 0 \le f \le k_j, \quad y = \bar{y}_j - u - \psi_j f + \eta_j v\}.
\end{aligned}$$

7. Conclusions

In this paper we have described a new electricity dispatch and pricing model based on agent decision rules (ADRs). We have demonstrated how ADRs can be used in storage, ramping, reserve and frequency regulation. This model has the advantage of dealing with uncertainty in future net demand for electricity without requiring the system operator to make forecasts or estimate probability distributions. The individual views of the future taken by market participants are incorporated into their ADRs and aggregated by the system operator in making the current period's dispatch.

In practice a market participant could devise their ADR to account for their attitude to risk and possible trades in derivative contracts. As long as the ADR gives convex future cost functions it can be easily handled in the formulation DP. Although our analysis in Section 4 focuses on the expected efficiency of ADR dispatch in a risk-neutral setting, a similar analysis could be performed when agents are risk-averse and endowed with coherent risk measures and markets for risk are complete. As shown by Ferris and Philpott (2022), a risked competitive equilibrium in a scenario tree is equivalent to a risk-averse solution to a system optimization, using a coherent risk measure derived from those of the agents. This gives similar results to Theorems 5 and 6, as long as the system risked future cost function can be separated into the sum of agent functions.

The welfare results we derive assume that ADRs can be derived and submitted to the system operator in real time at the beginning of a dispatch interval (typically five minutes) when the incoming states of all agents are known. Gate-closure conditions that preclude changes in offer in a given period (typically more than an hour) before dispatch make this impossible in practice. Long-lived ADR offers that are contingent on all possible dispatch histories during the gate closure period could be shown to yield the socially optimal solution in theory but are not realistic in any practical setting. Nevertheless a single ADR that is contingent on the state of a battery or ramping plant can persist over several consecutive dispatch periods to guide the system operator, who may either solve a sequence of one-period problems with the ADR in each, or in case a reliable forecast is available, solve a deterministic problem over T periods with the ADR guiding the outgoing state in period T. We see experimentation to quantify any efficiency gains in such long-lived ADRs as a fruitful area for further research.

Our analysis has not dwelt on how agents should generate their ADRs. The example we have presented has a social planning problem with two state variables, making it amenable to solution by dynamic programming. In most applications these future cost functions will have higher state dimension and not be separable by agent at every stage. An approximate dynamic programming method such as SDDP (Pereira and Pinto 1991) might then be required to generate suitable ADRs. Even so, some battery operators will find the effort required to compute an optimal ADR too much. However, there is nothing in our proposed dispatch process precluding them from using heuristic rules to specify their ADR, e.g., a collection of buy and sell prices that are parameterized by their state of charge.

Throughout this paper we have assumed a convex dispatch process. In many electricity markets the dispatch involves the start-up and shut-down of generating units with minimum operating levels and minimum up and down times. These are modeled using binary variables in multi-period mixed-integer programs. Deriving suitable prices from these models remains a challenge. Furthermore constructing ADRs for such problems is not straightforward, although extensions of SDDP to incorporate binary variables (Zou et al. 2019, Philpott et al. 2020) can be used to construct approximate future cost functions to use as a guide for deriving good ADRs.

Acknowledgments

This research was performed while the authors were participating in the Architecture of Green Energy Systems Program hosted by the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-1929348). Andy Philpott acknowledges support from UOCX2117 MBIE Catalyst Fund New Zealand German Platform for Green Hydrogen Integration (HINT).

References

Aiyagari S (1994) Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109(3):659–684.

Brown D, Smith J (2025) Unit commitment without commitment: A dynamic programming approach for managing an integrated energy system under uncertainty. *Operations Research* 73(4):1744–1766.

Chen C, Tong L (2023) Convexifying market clearing of SoC-dependent bids from merchant storage participants. *IEEE Transactions on Power Systems* 38(3):2955–2957, URL `http://dx.doi.org/10.1109/TPWRS.2023.3242470`.

Cho J, Papavasiliou A (2023) Pricing under uncertainty in multi-interval real-time markets. *Operations Research* 71(6):1928–1942, URL `http://dx.doi.org/10.1287/opre.2022.2314`.

Cory-Wright R, Philpott A, Zakeri G (2018) Payment mechanisms for electricity markets with uncertain supply. *Operations Research Letters* 46(1):116–121, ISSN 0167-6377, URL `http://dx.doi.org/https://doi.org/10.1016/j.orl.2017.11.017`.

Department of Market Monitoring (2024) 2023 special report on battery storage. Technical report, California ISO.

Diniz A, Costa F, Maceira M, dos Santos T, dos Santos L, Cabral R (2018) Short/mid-term hydrothermal dispatch and spot pricing for large-scale systems-the case of Brazil. *2018 Power Systems Computation Conference (PSCC)*, 1–7 (IEEE).

Eldridge B, Knueven B, Mays J (2023a) Rethinking the price formation problempart 1: Participant incentives under uncertainty. *IEEE Transactions on Energy Markets, Policy and Regulation* 1(4):480–489, URL `http://dx.doi.org/10.1109/TEMPR.2023.3315956`.

Eldridge B, Knueven B, Mays J (2023b) Rethinking the price formation problempart 2: Rewarding flexibility and managing price risk. *IEEE Transactions on Energy Markets, Policy and Regulation* 1(4):490–498, URL `http://dx.doi.org/10.1109/TEMPR.2023.3315953`.

Ferris M, Philpott A (2022) Dynamic risked equilibrium. *Operations Research* 70(3):1933–1952.

Gomes D, Mohr J, Souza R (2010) Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées* 93(3):308–328.

Hobbs B (2001) *The next generation of electric power unit commitment models*, volume 36 (Springer Science & Business Media).

Hogan W (2016) Electricity market design. *Workshop on Optimization and Equilibrium in Energy Economics, Institute for Pure and Applied Mathematics (IPAM), University of Southern California, Los Angeles, January, 2016.*

Hua B, Schiro D, Zheng T, Baldick R, Litvinov E (2019) Pricing in multi-interval real-time markets. *IEEE Transactions on Power systems* 34(4):2696–2705.

Lasry JM, Lions PL (2007) Mean field games. *Japanese journal of mathematics* 2(1):229–260.

Li T, Shahidehpour M (2005) Price-based unit commitment: A case of Lagrangian relaxation versus mixed integer programming. *IEEE transactions on power systems* 20(4):2015–2025.

Lucas Jr R, Prescott E (1971) Investment under uncertainty. *Econometrica: Journal of the Econometric Society* 659–681.

Mays J (2024) Sequential pricing of electricity. *Energy Economics* 137:107790, ISSN 0140-9883, URL `http://dx.doi.org/https://doi.org/10.1016/j.eneco.2024.107790`.

Muckstadt J, Koenig S (1977) An application of Lagrangian relaxation to scheduling in power-generation systems. *Operations Research* 25(3):387–403.

Nelson T, Conboy P, Hancock A, Hirschhorn P (2025) National Electricity Market wholesale market settings review, https://www.dcceew.gov.au/energy/markets/nem-wms-review. Technical report, Australian Government: Department of Climate Change, Energy, Environment and Water.

Pereira M, Pinto L (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52(1):359–375.

Philpott A, Wahid F, Bonnans J (2020) MIDAS: A mixed integer dynamic approximation scheme. *Mathematical Programming* 181(1):19–50.

Prescott E, Mehra R (2005) Recursive competitive equilibrium: The case of homogeneous households. *Theory Of Valuation*, 357–371 (World Scientific).

Pritchard G, Zakeri G, Philpott A (2010) A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations Research* 58(4-part-2):1210–1219, URL `http://dx.doi.org/10.1287/opre.1090.0800`.

Zakeri G, Pritchard G, Bjorndal M, Bjorndal E (2019) Pricing wind: A revenue adequate, cost recovering uniform price auction for electricity markets with intermittent generation. *INFORMS Journal on Optimization* 1(1):35–48, URL `http://dx.doi.org/10.1287/ijoo.2018.0002`.

Zavala V, Kim K, Anitescu M, Birge J (2017) A stochastic electricity market clearing formulation with consistent pricing properties. *Operations Research* 65(3):557–576, URL `http://dx.doi.org/10.1287/opre.2016.1576`.

Zheng N, Qin X, Wu D, Murtaugh G, Xu B (2023) Energy storage state-of-charge market model. *IEEE Transactions on Energy Markets, Policy and Regulation* 1(1):11–22, URL `http://dx.doi.org/10.1109/TEMPR.2023.3238135`.

Zou J, Ahmed S, Sun X (2019) Stochastic dual dynamic integer programming. *Mathematical Programming* 175:461–502.