Submitted to
manuscript (Please, provide the manuscript number!)

Electricity Dispatch and Pricing using Agent Decision Rules

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

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

Jacob Mays
Cornell University, jacobmays@cornell.edu

Models for computing economic 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. Growing investment in renewable energy has increased the uncertainty in net demand to be met by dispatchable generation. To accommodate this, stochastic programming models formulated using scenario trees have been proposed for economic dispatch. The use of these models in practice is challenging for several reasons. Market participants need to agree on the scenarios used for uncertain parameters in the model, and realizations of these parameters will be different from those in any modelled scenario. When updated in a rolling horizon fashion, stochastic models can misprice the option value of energy storage and the value of changing current dispatch to meet future ramping constraints. This leads to uplift payments that compensate participants for the fact that the system operator forecasts the future incorrectly. We present a class of new economic dispatch models that attempt to overcome these drawbacks, based on agent decision rules (ADRs). 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 new model through simple examples.

Key words: electricity markets, storage, batteries, energy prices

History: This paper was first submitted on August 26, 2024.

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. 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. In parallel with these research efforts, market operators have implemented a variety of market design adaptations, such as new market participation models for batteries and new ancillary service products for ramping. 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. A major roadblock to the adoption of advanced decision support tools is a lack of a shared understanding of their effect on price formation: implementation of new models depends not only on their ability to support efficient operations, but also the ability of stakeholders to forecast uncertain values and interpret the resulting prices.

To enable better uncertainty management, this paper considers alternative ways to incorporate stochastic elements in short-term electricity market design and proposes a change in the format of bids and offers supplied by market participants. In place of classical price–quantity pairs, the 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 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 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 duration limits.

Rather than relying primarily on the system operator, the ADR approach to uncertainty management depends on market participants to develop their own view on 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 our analysis, a clean division of responsibility occurs 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 enables 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 and competing beliefs (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. The scenario tree for (net) demand depends on many factors such as solar insolation and wind strength that are difficult to model accurately. Any model for determining the demand scenarios will be subject to some litigation by market participants if realized demand values fall outside the range of the modelled scenarios. 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 the California Independent System Operator (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, the issue is that 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 stored energy 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 last 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 2023), 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 when they are separable by agent. While it is difficult to demonstrate convergence to a socially optimal equilibrium more generally, we show examples of how dispatching based on decision rules can achieve results that approximate the social optimum. In principle, this approach allows the system operator to avoid generating its own forecasts or scenarios of supply and demand and prevents the mismatch in incentives described above. Uncertainty about future energy prices is instead encoded in participant ADRs. At the same time, we expect that many system operators will prefer to play a more active role in managing variability and uncertainty, going beyond a single-period dispatch and employing lookahead models and ancillary service products that act across market clearing intervals. CAISO, for example, describes its operational issues as stemming from "the challenges of having a limited optimization horizon," indicating a desire to extend the lookahead horizon further into the future if computationally feasible (Department of Market Monitoring 2024). In this vein, we discuss how the ADR approach can be modified depending on the desired split of responsibility between the market operator and market participants. Such a modified approach could enable system operators to better manage near-term issues within the lookahead horizon without creating issues outside of the lookahead horizon.

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, and give conditions for this approximation to be exact.

3. 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 the next section we demonstrate the key ideas of the paper by formulating an ADR dispatch model and solving a simple deterministic example of the model with one battery and one generator. For this model we compute an optimal decision rule for dispatch and show how this replicates the optimal solution to the perfect foresight problem. This is compared with approximately optimal ADRs that use imperfect information about the state variables. In Section 3, we present a version of the dispatch model that works with random demand, and illustrate this with a stochastic version of our previous example. This is solved over a 24-hour time horizon using stochastic dual dynamic programming, and the socially optimal dispatch is compared with that obtained from a sequence of single dispatches using ADRs for the generator and the battery. Section 4 presents some theoretical results that give conditions under which ADRs and the dispatch model we propose will yield a stochastic competitive equilibrium. In Section 5 we discuss extensions of ADRs and the ADR dispatch model to include transmission and reserve. The paper concludes in Section 6.

2. Agent decision rules and dispatch

A dispatch model is one that could be solved at the beginning of a planning horizon $1, \ldots, T$ to determine the dispatch at each time t. If the entire sequence of demands are known, then a clairvoyant system operator would solve the following perfect foresight formulation:

$$\begin{aligned}
\text{EP:}\min\ & \sum_{t=1}^{T}\left(\sum_{i\in\mathcal{G}} c_i(t)x_i(t) + Lz(t)\right) \\
\text{s.t.}\ & \sum_{i\in\mathcal{G}} x_i(t) + \sum_{j\in\mathcal{J}} u_j(t) - \sum_{j\in\mathcal{J}} v_j(t) + z(t) = d(t) + w(t), \\
& x_i(0) = x^0, \quad x_i(t) \in \mathcal{X}_i(x(t-1)), \quad i \in \mathcal{G}, \\
& y_j(0) = y^0, \quad (y_j(t), u_j(t), v_j(t)) \in \mathcal{Y}_j(y(t-1)), \quad j \in \mathcal{J}, \\
& w(t) \geq 0, z(t) \in [0, d(t)], \quad t = 1, 2, \ldots, T,
\end{aligned}$$

where

$$\mathcal{X}_i(\bar{x}) = \{x \mid 0 \leq x \leq q_i, x - \bar{x}_i \leq \rho_i, \bar{x}_i - x \leq \sigma_i\}, \tag{1}$$

$$\mathcal{Y}_j(\bar{y}) = \{(y, u, v) \mid 0 \leq y \leq E_j, 0 \leq u \leq r_j, 0 \leq v \leq s_j, y = \bar{y}_j - u + \eta_j v\}. \tag{2}$$

Here $x_i(t)$ denotes generation dispatched to generator i in period t, and $y_j(t)$ is storage of energy in device j at the end of period t. Both of these variables have initial values at the start of the day. The dispatch $x_i(t)$ is constrained by a ramp-up limit ρ_i and a ramp-down limit σ_i. Dispatch in period t is limited by capacity q_i and incurs a marginal cost of $c_i(t)$. Storage in device j is increased by charging using variable v_j and decreased by discharging an amount u_j. Round trip losses are modelled using the factor η_j which multiplies v_j. Degradation in battery performance from cycling can be modelled by imposing a cost on u_j, but we choose to ignore this in our simple model. Charging and discharging rates for device j are limited by the parameters s_j and r_j respectively. Each storage device j has a maximum charge E_j. The total amount generated should meet demand $d(t)$. Any shortfalls $z(t)$ are penalized at a value of lost load L. In general we also allow free disposal of energy using variable $w(t)$. Note that the technical constraints on generators and batteries are encapsulated in the definitions of $\mathcal{X}_i$ and $\mathcal{Y}_j$ given in (1) and (2) that could be generalized if necessary to capture additional restrictions.

This model must use a forecast of demand. The dispatch will be optimal for the planning horizon if the forecast is correct. Such perfect foresight is impossible in practice, especially if $d(t)$ is demand net of renewable energy. Nevertheless the perfect foresight model is convenient as a benchmark for assessing other types of model.

The problem EP is a convex programming model of economic dispatch. Observe that the model does not have transmission constraints, and has no discrete decisions arising, for example, from unit commitment switching costs and constraints. Transmission constraints are omitted for simplicity of exposition, and are discussed briefly in Section 5. In contrast, unit commitment variables make EP a nonconvex model, a setting that lies outside our modelling framework, which relies heavily on being able to compute energy prices that yield a competitive equilibrium. Some comments on this are also provided in Section 5.

In our approach, 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. Given the parameters, the system operator can determine a dispatch by solving a deterministic optimal dispatch problem in the current trading period incorporating these ADRs. Since this dispatch happens after all the ADR parameters have been observed, there is no need for a scenario tree to model the variation in net demand. The view of the future is encapsulated in the ADR. In our examples, ADRs will consist of a triple that includes an immediate cost, participant constraints, and a future cost function. However, generalizations are possible, provided the ADRs are implementable within a convex dispatch problem.

Although the ADR is enacted in real time, it is computed (offline) and supplied to the system operator in advance. One might imagine such rules being supplied to the system operator at 11:30 pm to apply for all 24 hours of the following day. This does not mean the ADRs are fixed - a different ADR could be defined for each hour - but they need not be re-defined during the day of operation after market outcomes have been observed in previous periods.

2.1. 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_i. 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.

2.2. Agent decision rules for batteries

A battery is any storage device that can be charged and discharged to shift energy between periods. This encompasses grid-scale lithium ion batteries, pumped storage facilities, and aggregated demand shifting devices such as plug-in electric vehicles. An ADR for a battery determines its charge or discharge decision. For example, a battery operator could set an upper threshold electricity price above which it would discharge at the maximum rate and a lower threshold price below which it would charge at the maximum rate. These threshold prices could be a function of the current state of charge y_j of the battery.

Another approach is to provide the system operator with a function $W_j^t(y)$ that specifies the expected future revenue that the battery operator can earn from the end of the period t till the end of T if its charge at the end of t is y. This can then be included in a $s$ingle-period dispatch problem.

2.3. Agent decision rules for ramping plant

An ADR for a ramping plant determines its dispatch. Ramping constraints impose a bound ρ on the increase in dispatch in consecutive periods. If generation is needed to meet a demand peak, then it must be ramped up in preceding periods to enable it to generate at its maximum during the peak. A generator might have a model of future electricity prices as a function of current observations and devise a decision rule based on this model.

An alternative approach is to provide the system operator with a function $V_i^t(x)$ that specifies the expected future revenue that generator can earn from the end of the period t till the end of T if its generation level at the end of t is x.

2.4. An ADR dispatch model

Provided with an ADR by generators and storage agents, the system operator can solve a single period dispatch problem in each period t using the observed values of parameters that map any ADR to an action. We will assume that each agent defines these rules for period t using value functions V_i^t and W_j^t. In period t, given $\bar{x} = x(t-1)$ and $\bar{y} = y(t-1)$, the system operator then solves

$$\begin{aligned}
\text{DP}(t,\bar{x},\bar{y})\colon \min \; & \sum_{i\in\mathcal{G}} c_i(t)x_i + Lz + C^t(x,y) \\
\text{s.t.} \; & \sum_{i\in\mathcal{G}} x_i + \sum_{j\in\mathcal{J}} u_j - \sum_{j\in\mathcal{J}} v_j + z = d(t) + w, \\
& x_i \in \mathcal{X}_i(\bar{x}), i \in \mathcal{G} \\
& (y_j, u_j, v_j) \in \mathcal{Y}_j(\bar{y}), j \in \mathcal{J} \\
& w \geq 0, z \in [0, d(t)],
\end{aligned}$$

where $\mathcal{X}_i$, $\mathcal{Y}_j$ are defined in (1) and (2) and

$$C^t(x,y) = -\sum_{i\in\mathcal{G}} V_i^t(x_i) - \sum_{j\in\mathcal{J}} W_j^t(y_j). \tag{3}$$

Observe 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.

2.5. A simple example

Our example simplifies exposition using a single generator and a single battery, and $T = 24$. We drop the subscripts i and j in (1) and (2) since these are superfluous in this setting.

The values chosen for the scalar parameters are given in Table 1, where σ is chosen as ∞ to remove the ramp-down constraint. We set $c(t) = 7.0$, and plot values for $d(t)$ in Figure 1. Meeting the duck-curve shaped demand in our model will require the generator and battery to be in a position to respectively ramp up and discharge in periods 18 through 22.

$q = 70.0$	$E = 8.0$	$\eta = 0.8$
$r = 10.0$	$s = 10.0$	$\rho = 10.0$
$L = 35.0$	$x^0 = 35.0$	$y^0 = 4.0$

Table 1 **Parameter values for example**

We first solve the perfect foresight model EP where the demand is known ex-ante for all $t = 1, 2, \dots, 24$. The optimal solution to EP has cost 6062, with optimal dispatch and battery charge shown in Figure 2.

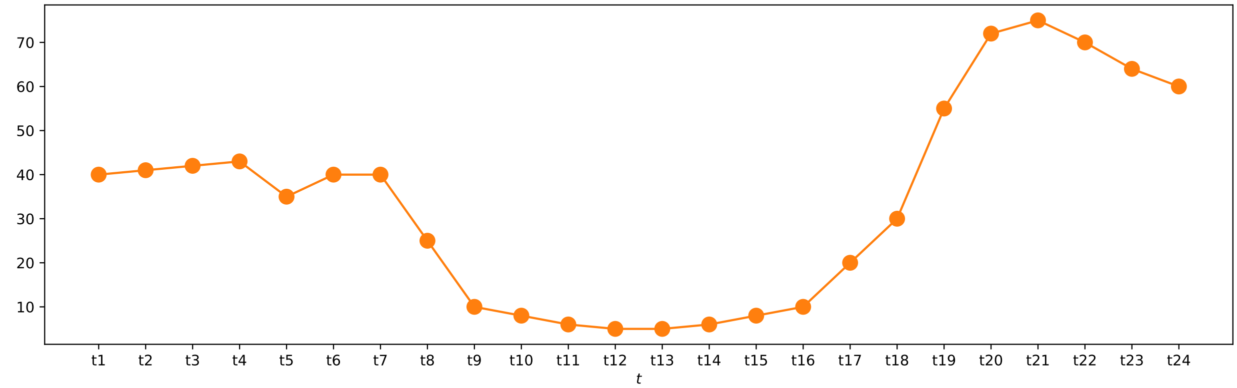

Figure 1 **Values of $d(t)$ for $t = 1, 2, \dots, 24$.**

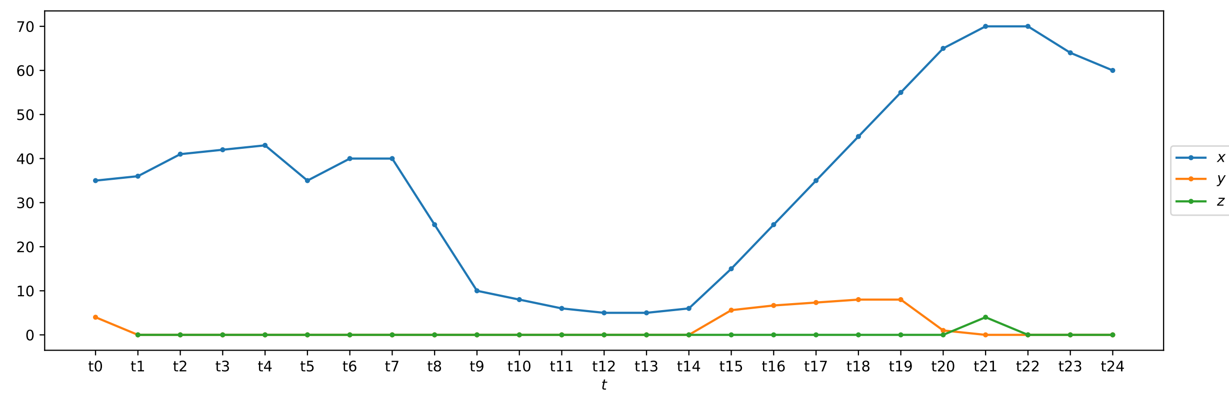

Figure 2 **Solution of DP showing generation x, battery charge y and lost load z for $t = 1, 2, \dots, 24$.**

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

$$
\begin{aligned}
C^{t-1}(x(t-1), y(t-1)) = \min \;& c(t)x + Lz + C^t(x, y) \\
\text{s.t.} \;& x + u - v + z = d(t) + w, \\
& x \in \mathcal{X}(x(t-1)), \\
& (y, u, v) \in \mathcal{Y}(y(t-1)), \\
& w \geq 0, z \in [0, d(t)],
\end{aligned}
$$

where $C^{24}(x,y)=0$ and $x(0)=x^0$, $y(0)=y^0$. Since each stage problem is a linear program, the future cost functions are convex polyhedral functions that can be represented epigraphically as

$$C^t(x,y) = \min\ \theta$$
$$\text{s.t.}\quad \theta \geq \alpha_k(t)+\beta_k(t)x+\gamma_k(t)y, \quad k\in\mathcal{K}(t).$$

The values of $\alpha_k(t)$, $\beta_k(t)$ and $\gamma_k(t)$ can be computed for $k\in\mathcal{K}(t)$, $t=1,2,\ldots,24$, using stochastic dual dynamic programming (SDDP) (Dowson and Kapelevich 2021). Given these values the system operator can solve the sequence of optimization problems DP(1,$x(0)$,$y(0)$), DP(2,$x(1)$,$y(1)$), …, DP(24,$x(23)$,$y(23)$) where $x(0)$ and $y(0)$ are given and the definition of C^t via θ is substituted directly into DP. These problems involve no explicit lookahead forecasts and 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, e.g., in SDDP, to approximate an expected future cost function $C^t(x,y)$. 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.

Our proposal is to estimate the functions $C^t(x,y)$ by soliciting information from market participants. In other words, at the start of the day the generator agent would provide functions $C_g^t(x,y)$, $t=1,2,\ldots,24$ and the battery operator would provide functions $C_b^t(x,y)$, $t=1,2,\ldots,24$, and then the system operator would solve DP(1,$x(0)$,$y(0)$), DP(2,$x(1)$,$y(1)$), …, DP(24,$x(23)$,$y(23)$) using

$$C^t(x,y)=C_g^t(x,y)+C_b^t(x,y).$$

The future cost functions C_g^t and C_b^t are estimated by each agent. In this example they will depend on the values of state variables (x,y) that are realized at the end of period t. Here the dependence of C_g^t on storage levels y, and the dependence of C_b^t on dispatch levels x might be hard for each agent to define. A simpler dispatch model uses a separable form of C^t as in (3). If the true function $C^t(x,y)$ (derived from SDDP) is not separable then a sequence of dispatches using (3) can be socially suboptimal as illustrated by experiments with the above example. Specifically, when this separable approach is used with functions trained using SDDP, the optimal solution from solving DP(1,$x(0)$,$y(0)$), DP(2,$x(1)$,$y(1)$), ..., DP(24,$x(23)$,$y(23)$) has a higher objective value of 6174, with the solution plotted in Figure 3. This shows that the separable ADR begins charging the battery later and incurs more lost load than the optimal ADR.

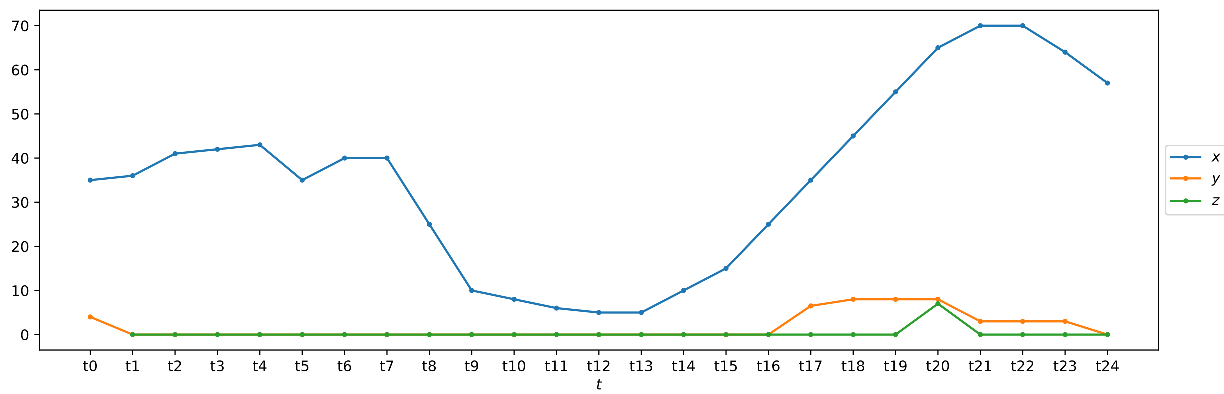

Figure 3 **DP solution with separable C showing generation x, battery charge y and lost load z for $t = 1,2,\ldots,24$**

3. Stochastic demand and decision rule dispatch

In this section we show how the example of the previous section can be extended to accommodate uncertainty in demand. The demand process we assume adds 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 SDDP and extract future cost functions $C^t(x,y)$ defined by the cutting planes at each stage. At each t we define a separable ADR using the approximation

$$\hat{C}^t(x,y) = \frac{1}{2}C^t(x,\tilde{y}(t)) + \frac{1}{2}C^t(\tilde{x}(t),y) \tag{4}$$

where $\tilde{x}(t)$ and $\tilde{y}(t)$ are chosen to be close to the values we would expect in an optimal solution.

Each ADR is tested by simulation. N realizations of demand are created by random sampling and dispatch problems DP(1,$x(0)$,$y(0)$), DP(2,$x(1)$,$y(1)$), ..., DP(24,$x(23)$,$y(23)$) are solved sequentially for each demand scenario. The cost of the ADR applied to each scenario is recorded and averaged. We choose $N = 1000$.

We test three ADRs using these simulations. The *deterministic* ADR uses the separable future cost functions generated from the original deterministic demand. The *optimal* ADR uses (non-separable) future cost functions $C^t(x,y)$ generated from a solution to SDDP. The *separable* ADR uses (separable) approximate future cost functions $\hat{C}^t(x,y)$ generated from $C^t(x,y)$ using (4).

For the data given above, the estimate of the expected cost of the deterministic ADR (with two standard error confidence interval) was 6226.37 ± 10.09. The estimated expected cost of the optimal ADR is 6109.51 ± 10.85. The estimated expected cost of the separable ADR is 6208.27 ± 9.17. As in the deterministic case the separable ADR does not achieve the optimal expected cost.

In Figure 4 we plot some price realizations from simulations of the system optimal policy. This shows that the optimal price process is not necessarily stagewise independent. In more detailed practical settings, the price process undoubtedly becomes even more complicated and hard to represent using a simple stochastic model. This makes agent decisions difficult to optimize using Lagrangian relaxation as agents must optimize stochastic dynamic programs with complicated price processes. The ADR approach generates system prices as a consequence of the dispatch and therefore has no need to represent these as exogenous parameters in agent optimization models.

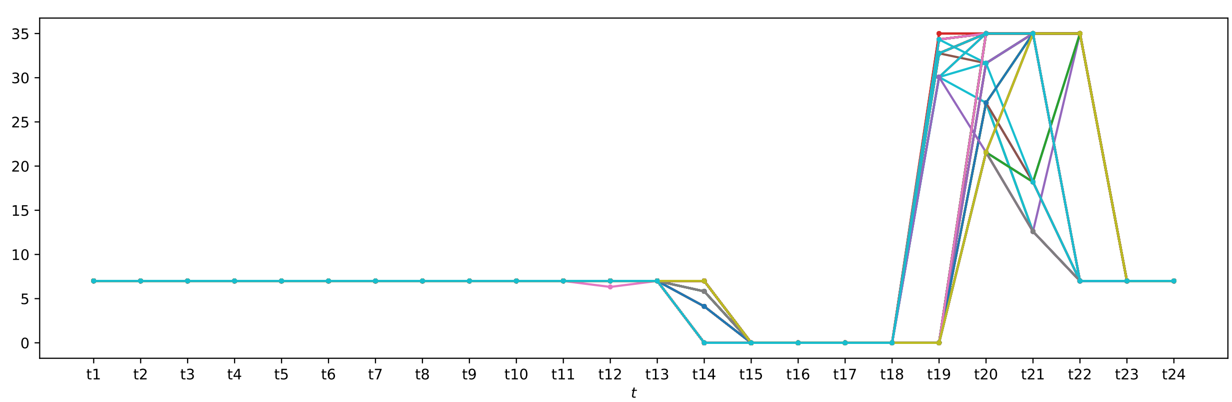

Figure 4 System marginal prices from 100 simulations of optimal stochastic policy.

Although our goal in defining ADR dispatch is to try and avoid lookahead and issues related to forecasting or scenarios, it is possible to get closer to the optimal solution by adding some lookahead to DP. To illustrate this, we extend the dispatch model DP by adding a lookahead window of τ steps using the notation $\tau(t) = \{t, \ldots, t+\tau\}$:

$$
\begin{aligned}
\text{LA}((\tau(t), x(t-1), y(t-1))) \colon \min \; & \sum_{k\in\tau(t)} \sum_{i\in\mathcal{G}} c_i(k)x_i(k) + Lz(k) + C^{t+\tau}(x(t+\tau), y(t+\tau)) \\
\text{s.t.} \; & \sum_{i\in\mathcal{G}} x_i(k) + \sum_{j\in\mathcal{J}} u_j(k) - \sum_{j\in\mathcal{J}} v_j(k) + z(k) = d(k) + w(k), \\
& x_i(k) \in \mathcal{X}_i(x(k-1)), i\in\mathcal{G} \\
& (y_j(k), u_j(k), v_j(k)) \in \mathcal{Y}_j(y(k-1)), j\in\mathcal{J} \\
& w(k) \geq 0, z(k) \in [0, d(k)], k \in \tau(t).
\end{aligned}
$$

For this to work, the system operator has to estimate the demand $d(k)$ for $k\in\tau(t)$, and makes the assumption that such demand is realized in each of the periods k. The model still generates the dispatch $x(t)$ and $y(t)$, and only uses the decision rule at the end of the window. In our example, with $\tau = 1$, if we use the deterministic demand of Figure 1 for demand forecasts in the window, then the deterministic (separable) ADR gives objective 6062 (equal to the optimal value for the perfect foresight model). For the random data, the same forward simulation process with 1000 runs leads to a confidence interval of 6153.00 $\pm$ 32.50 for the deterministic (separable) ADR with one-step lookahead, which gives an improvement in expected cost from the previous value 6226.37 without lookahead, but short of the optimal (SDDP) ADR that has estimated expected cost 6109.51. Note that this rolling horizon approach is more effective at capturing the future, but the dispatch action in each period is not necessarily optimal for the future cost function at the end of the period, since it is optimized over the whole window using the future cost function at the end of the window.

4. Stochastic programming and equilibrium

We now consider a more standard stochastic version of DP and use this to discuss settings where we can prove that ADR dispatch minimizes expected system costs. To do this we model random variables using a scenario tree with node set $\mathcal{N}$ as described in Ferris and Philpott (2022). Each node $n\in\mathcal{N}$ corresponds to a state of the world with an associated demand $d(n)$ and probability $P(n)$. The root of the tree is $n = 1$, the leaf nodes of the tree are denoted $\mathcal{L}$, and n_- is the parent of node n. The set of children nodes of n are denoted by n_+, and the subtree of descendants rooted at n is denoted $\mathcal{T}(n)$. Each node n corresponds to a time interval $t(n)$, where $t(n_-)=t(n)-1$. We will assume that each node n at the same time interval has the same number of children nodes.

Suppose we have generators $i\in\mathcal{G}$ producing quantities $x_i(n)$ at marginal cost $c_i(n)$ for $n\in\mathcal{N}$. Battery operators $j\in\mathcal{J}$ buy $v_j(n)$ or sell $u_j(n)$ units of energy in each node n. A social planner minimizes the expected cost of meeting demand over all the states of the world represented by the scenario tree, by solving the following problem:

$$
\begin{aligned}
\text{SDP: }\min\ & \sum_{n\in\mathcal{N}} P(n)\left(\sum_{i\in\mathcal{G}} c_i(n)x_i(n) + Lz(n)\right) - \sum_{n\in\mathcal{L}} P(n)\bar{V}^n(x(n),y(n)) \\
\text{s.t. } & \sum_{i\in\mathcal{G}} x_i(n) + \sum_{j\in\mathcal{J}} u_j(n) - \sum_{j\in\mathcal{J}} v_j(n) + z(n) = d(n) + w(n), \quad [P(n)\pi(n)], \quad n\in\mathcal{N}, \\
& x_i(1)=x_0, \quad x_i(n)\in\mathcal{X}_i(x(n_-)), \quad i\in\mathcal{G}, n\in\mathcal{N}\setminus\{1\}, \\
& y_j(1)=y_0, \quad (y_j(n),u_j(n),v_j(n))\in\mathcal{Y}_j(y(n_-)), \quad j\in\mathcal{J}, n\in\mathcal{N}\setminus\{1\}, \\
& w(n)\geq 0, z(n)\in[0,d(n)], \quad n\in\mathcal{N}.
\end{aligned}
$$

To construct an equilibrium, it is necessary to assume that the total expected future social benefit function $\bar{V}^n$ in state $n\in\mathcal{L}$ is separable by agent, so

$$
\bar{V}^n(x(n),y(n)) = \sum_{i\in\mathcal{G}} V_i^n(x_i(n)) + \sum_{j\in\mathcal{J}} W_j^n(y_j(n)).
$$

$V_i^n(x)$ for $n\in\mathcal{L}$ represents the future benefit earned by generator i at the end of the decision horizon when their dispatch in node n is x. Similarly $W_j^n(y)$ for $n\in\mathcal{L}$ represents the future benefit to be earned by device owner j at the end of the decision horizon when they have stored energy y. Each V_i^n and W_j^n is assumed to be a concave function. The separability assumption enables us to decompose SDP by price into agent optimization problems. Since SDP is a convex problem it can be solved by minimizing its Lagrangian

$$
\begin{aligned}
\mathcal{L} = & \sum_{n\in\mathcal{N}} P(n)\sum_{j\in\mathcal{J}} c_i(n)x_i(n) - \sum_{n\in\mathcal{L}} P(n)\left(\sum_{i\in\mathcal{G}} V_i(x_i(n)) + \sum_{j\in\mathcal{J}} W_j(y_j(n))\right) \\
& + \sum_{n\in\mathcal{N}} P(n)\sum_{j\in\mathcal{J}} Lz(n) \\
& + \sum_{n\in\mathcal{N}} P(n)\pi(n)(d(n) - \sum_{i\in\mathcal{G}} x_i(n) - \sum_{j\in\mathcal{J}} u_j(n) + \sum_{j\in\mathcal{J}} v_j(n) - z(n) + w(n))
\end{aligned}
$$

subject to simple bounds on the variables. The shadow prices $P(n)\pi(n)$ obtained in each node n at the optimal solution to SDP give marginal prices that can be used to decouple the problem by agent. Each agent solves a profit maximization problem at these prices to give a dynamic market equilibrium as defined in Ferris and Philpott (2022).

For this particular application the equilibrium is defined as follows. We first define a generator optimization problem that maximizes their profit at given prices.

$$\begin{aligned}\text{GP}(i): \quad \max \sum_{n\in\mathcal{N}} P(n)(\pi(n)-c_i(n))x_i(n)) + \sum_{n\in\mathcal{L}} P(n)V_i(x_i(n))\\ \text{s.t. } x_i(1)=x_0, \quad x_i(n)\in\mathcal{X}_i(x(n_-)), \quad i\in\mathcal{G}, n\in\mathcal{N}\setminus\{1\}.\end{aligned}$$

The consumer chooses to shed load $z(n)$ to solve

$$\begin{aligned}\text{CO}(n)\text{: } \max\ (\pi(n)-L)z(n)\\ \text{s.t. } 0\le z(n)\le d(n),\ n\in\mathcal{N}.\end{aligned}$$

The optimization problem for storage agent j is

$$\begin{aligned}\text{BP}(j)\text{: } \max \sum_{n\in\mathcal{N}} P(n)\pi(n)(u_j(n)-v_j(n)) + \sum_{n\in\mathcal{L}} P(n)W_j(y_j(n))\\ \text{s.t. } y_j(1)=y_0, \quad (y_j(n),u_j(n),v_j(n))\in\mathcal{Y}_j(y(n_-)), \quad j\in\mathcal{J}, n\in\mathcal{N}\setminus\{1\}.\end{aligned}$$

Definition 1. A *multistage stochastic equilibrium* is a stochastic process of prices $\{\pi(n), n\in\mathcal{N}\}$ and a corresponding collection of actions $\{u(n),v(n),x(n),y(n),z(n),n\in\mathcal{N}\}$ with the properties

1. $x_i(n)$, $n\in\mathcal{N}$ solves the problem GP(i), for $i\in\mathcal{G}$,
2. $z(n)$ solves the problem CO(n), for $n\in\mathcal{N}$,
3. $(u_j(n),v_j(n),y_j(n))$, $n\in\mathcal{N}$ solves the problem BP(j), for $j\in\mathcal{J}$, and
4. $0\le\pi(n) \quad \perp \quad \sum_{i\in\mathcal{G}} x_i(n) + \sum_{j\in\mathcal{J}} u_j(n) - \sum_{j\in\mathcal{J}} v_j(n) + z(n) - d(n) \ge 0, \quad n\in\mathcal{N}.$

In a multistage stochastic equilibrium, the system clearing agent announces a set of prices $\{\pi(n), n\in\mathcal{N}\}$, and each agent chooses a sequence of actions adapted to the filtration defined by the scenario tree that maximizes their profit with these prices as viewed in the root node of the tree.

The following welfare theorems are easily demonstrated (see Ferris and Philpott (2022)).

THEOREM 1. *Suppose* $\{\pi(n), n\in\mathcal{N}\}$, *and* $\{u(n), v(n), x(n), y(n), z(n), n\in\mathcal{N}\}$ *form a multistage stochastic equilibrium. Then* $\{u(n), v(n), x(n), y(n), z(n), w(n),\ n\in\mathcal{N}\}$ *is a solution to SDP, where*

$$w(n)=\sum_{i\in\mathcal{G}} x_i(n)+\sum_{j\in\mathcal{J}} u_j(n)-\sum_{j\in\mathcal{J}} v_j(n)+z(n)-d(n).$$

THEOREM 2. *Suppose* $\{u(n), v(n), x(n), y(n), z(n), w(n), n\in\mathcal{N}\}$ *is a solution to SDP, with Lagrange multipliers* $\{P(n)\pi(n), n\in\mathcal{N}\}$ *for the demand constraints. Then* $\{\pi(n), n\in\mathcal{N}\}$, *and* $\{u(n), v(n), x(n), y(n), n\in\mathcal{N}\}$ *is a multistage stochastic equilibrium.*

A special case of multistage equilibrium occurs when $\mathcal{N}$ is a singleton. This is easily seen to correspond to an instance of the ADR dispatch problem. We call this a *stage equilibrium*. It is characterized by prices at which each agent maximizes their single-period revenue minus cost plus an expected future reward defined by their ADR.

We would like to show that an ADR dispatch can be socially optimal in expectation if the decision rules for the agents are chosen to maximize current and future expected profit. To do this we will assume that the system operator constructs a scenario tree $\mathcal{N}$, by approximating some known random process, e.g., by sampling. The random process will be assumed to represent a ground-truth model of the future shared by the system operator and all the agents. (This assumption is reasonable in the risk-neutral setting that we adopt. On the other hand, if agents are risk averse with coherent risk measures, the system operator would optimize SDP using a probability distribution that emerges in equilibrium from risk trading as discussed in Ferris and Philpott (2022)).

At this point it is helpful to simplify the discussion by making the following assumption (stagewise independence).

ASSUMPTION 1. *Each node* $n\in\mathcal{N}$ *with the same value* $t(n)$ *has an identical subtree* $\mathcal{T}(n)$ *(with corresponding values of* d, c_i, *and* P*).*

Assumption 1 would hold if SDP was obtained by sampling a stochastic optimal control problem where outcomes $c_i(n)$ and $d(n)$ come from deterministic forecasts, each with a random noise term that is independent from the noise at other times. In this case SDP can be solved (approximately) by SDDP. Some exceptions to Assumption 1 are possible by extending the state space of the dynamic program, but we do not consider these here to make the presentation as simple as possible.

Although the stochastic process of prices $\{\pi(n), n \in \mathcal{N}\}$ that forms a dynamic stochastic equilibrium is adapted to the filtration defined by $\mathcal{N}$, it does not necessarily share any independence properties of the data processes that define $d(n)$ or $c_i(n)$ (see Barty et al. (2010), Brown and Smith (2023)). This makes the agents' stochastic optimization problems that define the equilibrium more difficult to solve. Without explicitly solving them, we show that the solutions to the agents' problems when viewed as decision rules will maximize expected social welfare.

To do this we need to introduce a more restrictive separability assumption on expected future cost functions. Suppose SDP is solved using SDDP applied to $\mathcal{N}$. Under Assumption 1 this yields an approximate future cost function $C^t(x, y)$ at each stage t, that is the same for all nodes $n \in \mathcal{N}$ with $t(n) = t$. The assumption we need is as follows:

Assumption 2. *For every $t = 1, 2, \ldots, T-1$, the expected future cost function can be written*

$$C^t(x, y) = -\sum_{i \in \mathcal{G}} V_i^t(x_i) - \sum_{j \in \mathcal{J}} W_j^t(y_j).$$

As shown in our numerical examples, expected future cost functions are typically not separable functions at every t even if they are separable at the end of the horizon. On the other hand, Assumption 2 holds when the devices defining state variables are identical. When these variables represent m storage devices beginning with equal storage, and the only random variable is demand, then under a socially optimum policy each device will follow the same charging and discharging strategy. This means at all times t, $y_1(t) = y_2(t) = \ldots = y_m(t) = y(t)$, so the future cost of storage $C^t(y_1, y_2, \ldots, y_m)$ at any time t will therefore be some function $\tilde{C}^t$ of the (equal) charge in each device. Since the devices are identical, this future cost will be shared equally, and so device j incurs future cost $C_j^t(y_j) = \frac{1}{m}\tilde{C}^t(y_j)$, giving

$$C^t(y_1, y_2, \ldots, y_m) = \sum_{j \in \mathcal{J}} C_j^t(y_j).$$

Now consider a system operator who is provided with the functions $V_i^t(x_i)$, $i \in \mathcal{G}$ and $W_j^t(y_j)$, $j \in \mathcal{J}$ for each $t = 1, 2, \ldots, T$. The system operator can compute actions that solve SDP by solving the following single-stage dispatch problem for each $n \in \mathcal{N}$, starting with the root node.

$$\begin{aligned}
\text{DP}(n)\colon \min \ & \sum_{i\in\mathcal{G}} (c_i(n)x_i + Lz)) - \sum_{i\in\mathcal{G}} V_i^{t(n)}(n)(x_i) - \sum_{j\in\mathcal{J}} W_j^{t(n)}(n)(y_j) \\
\text{s.t.} \ & \sum_{i\in\mathcal{G}} x_i + \sum_{j\in\mathcal{J}} u_j - \sum_{j\in\mathcal{J}} v_j + z = d(n) + w, \qquad [\pi(n)] \\
& x_i \in \mathcal{X}_i(x(n_-)), i \in \mathcal{G}, \\
& (y_j, u_j, v_j) \in \mathcal{Y}_j(y(n_-)), j \in \mathcal{J}, \\
& w \geq 0, \quad z \in [0, d(n)].
\end{aligned}$$

Here we assume that the values $x(n_-)$ and $y(n_-)$ are interpreted as initial values when n is the root node, and DP(n_-) is solved before DP(n). The dual variables $\pi(n)$ for DP(n) when scaled by $P(n)$ can be shown to be the optimal Lagrange multipliers for SDP, and so the actions $\{u(n), v(n), x(n), y(n), z(n), n \in \mathcal{N}\}$ that optimize DP(n) solve SDP and by Theorem 2 form a stochastic dynamic equilibrium with prices chosen to be $\pi(n)$. Each generator agent i in this equilibrium chooses $x_i(n)$ to solve

$$\begin{aligned}
\text{GP}(i,n)\colon \quad & \max \ (\pi(n) - c_i(n))x_i(n)) + V_i^{t(n)}(x_i(n)) \\
\text{s.t.} \quad & x_i(n) \in \mathcal{X}_i(x(n_-)), \quad i \in \mathcal{G}.
\end{aligned}$$

and each storage agent j solves

$$\begin{aligned}
\text{BP}(j,n)\colon \quad & \max \ \pi(n)(u_j(n) - v_j(n)) + W_j^{t(n)}(y_j(n)) \\
\text{s.t.} \quad & (y_j(n), u_j(n), v_j(n)) \in \mathcal{Y}_j(y(n_-)).
\end{aligned}$$

This shows that the functions W_i^t and V_j^t can be interpreted as future value functions for each agent when they optimize at equilibrium prices.

We can summarize this discussion in the following theorem.

THEOREM 3. *Suppose Assumption 1 holds and the system operator computes an optimal solution to SDP in a sampled scenario tree $\mathcal{N}$ using SDDP. Assume that the future cost functions at each stage are separable by agent, i.e. Assumption 2 holds. Suppose generators and storage devices use* *optimal future value functions* $V_i^t(x_i)$*,* $i \in \mathcal{G}$ *and* $W_j^t(y_j)$*,* $j \in \mathcal{J}$ *to define their ADRs and we simulate sequences of ADR dispatches in every scenario in the tree. Then*

1. *each optimal ADR dispatch and corresponding prices on each demand constraint gives a competitive stage equilibrium;*
2. *the ADR dispatches and prices give a stochastic equilibrium;*
3. *the ADR dispatches solve SDP.*

It is important to recognize that the scenario tree in this theorem is merely a device to enable us to model the equilibrium price process. Even though SDP in this sampled tree may be impossible to solve, in principle we can apply the above argument with an arbitrarily large tree. Since sample average approximation applied to multistage finite horizon problems can be shown to converge almost surely with increasing sample size (Shapiro 2006) this indicates asymptotic optimality for ADRs computed using SDDP. It is also worth noting that ADRs obtained from SDDP can be applied in a dispatch model and simulated using any stochastic process of demand to investigate their average performance.

5. Further extensions

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.

5.1. ADRs for demand response

Demand response (also known as peak shaving) refers to demand that is decreased when prices are high. 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 curve. This is currently possible in conventional dispatch models, and fits in our model (see Subsection 2.1).

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.

5.2. ADRs for flexible demand

Instead of just reducing load, some industrial loads (or even datacenters) 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 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 state $y_j(t)$ here 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 $y_j = 1$.

5.3. 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.

5.4. 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 of the same form as in the dispatch model LA.

5.5. Transmission constraints

The examples in this paper focus on the dispatch modelled at a single location. This model can be extended to a DC load flow model with locational marginal prices and ADRs defining future costs as functions of state variables. Suppose f_{ab} denotes transmission in the directed line $(a,b) \in \mathcal{A}$ that connects buses $a,b \in \mathcal{B}$. The set $\mathcal{G}(a)$ contains indices of generators at bus a and $\mathcal{J}(a)$ contains indices of storage devices at bus a. Each line has reactance X_{ab} and thermal limits $\bar{f}_{ab}$ on its flow. The voltage angles θ_a and θ_b determine the flow. We amend the constraints of DP as follows:

$$\sum_{i\in\mathcal{G}(a)} x_i + \sum_{j\in\mathcal{J}(a)} u_j - \sum_{j\in\mathcal{J}(a)} \eta_j v_j - d^a(t) + z^a - w^a = \sum_{\{b|(a,b)\in\mathcal{A}\}} f_{ab} - \sum_{\{b|(b,a)\in\mathcal{A}\}} f_{ba}, \quad a \in \mathcal{B},$$

$$(\theta_a - \theta_b) = X_{ab} f_{ab}, \quad (a,b) \in \mathcal{A},$$

$$-\bar{f}_{ab} \le f_{ab} \le \bar{f}_{ab}, \quad (a,b) \in \mathcal{A}.$$

The future cost function in DP now takes the form

$$-\sum_{a\in\mathcal{B}} \Big(\sum_{i\in\mathcal{G}(a)} V_i^t(x_i) + \sum_{j\in\mathcal{J}(a)} W_j^t(y_j) \Big).$$

5.6. 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 $\bar{x}$ and the battery charge levels are $\bar{y}$. Suppose that generator $i \in \mathcal{G}$ is dispatched x_i^r of reserve at marginal cost g_i. The security-constrained dispatch model allows available generation to be split into immediate demand satisfaction and reserve requirements.

$$
\begin{aligned}
\text{SCDP}(t,\bar{x},\bar{y})\colon \min \; & \sum_{i\in\mathcal{G}} (c_i(t)x_i + g_i(t)x_i^r) + Lz + C^t(x,y), \\
\text{s.t.} \; & \sum_{i\in\mathcal{G}} x_i^r = d^r(t), \\
& \sum_{i\in\mathcal{G}} x_i + \sum_{j\in\mathcal{J}} u_j - \sum_{j\in\mathcal{J}} v_j + z = d(t) + w, \\
& (x_i, x_i^r) \in \tilde{\mathcal{X}}_i(\bar{x}), \quad i \in \mathcal{G} \\
& (y_j, u_j, v_j) \in \tilde{\mathcal{Y}}_j(\bar{y}), \quad j \in \mathcal{J} \\
& w(t) \geq 0, \quad z(t) \in [0, d(t)],
\end{aligned}
$$

where

$$
\tilde{\mathcal{X}}_i(\bar{x}) = \{(x, x^r) \mid 0 \leq x + x^r \leq q_i, x + x^r - \bar{x}_i \leq \rho_i, \bar{x}_i - x \leq \sigma_i\},
$$

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

The definition of $\tilde{\mathcal{X}}_i(\bar{x})$ can include extra constraints on x^r that depend on each generator's plant. Some care is needed in defining the expected future cost. The problem SCDP can be solved for each realization $\xi(\omega)$ of net demand $d(t)$ and $c(t)$, yielding optimal cost $\tilde{C}^t(\bar{x},\bar{y},\xi(\omega))$ assuming that reserve is not called in scenario ω. If reserve is called (say with probability $P_t(\omega)$) then the optimal cost in scenario ω is $\tilde{C}^t(\bar{x}+x^r(\omega),\bar{y},\xi(\omega))$. This means

$$
C^t(\bar{x},\bar{y}) = \mathbb{E}_{\xi(\omega)}[(1-P_t(\omega))\tilde{C}^t(\bar{x},\bar{y},\xi(\omega)) + P_t(\omega)\tilde{C}^t(\bar{x}+x^r(\omega),\bar{y},\xi(\omega))].
$$

5.7. 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.

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

where

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

5.8. Integer variables

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 modelled 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.

6. 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 for this. 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 a similar result to Theorem 3, as long as the system risked future cost function can be separated into the sum of agent functions.

Our analysis has not dwelt on how agents should generate their ADRs. Some battery operators will find the computation of an optimal ADR too complicated. There is nothing in our proposed dispatch process precluding them from using heuristic rules to specify their ADR, or using a linear decreasing function that specifies a constant marginal value of discharge. The example we have presented uses the sets of cutting planes defining an optimal policy computed by SDDP. In most applications these future cost functions will not be separable by agent at every stage. Some heuristic is then required to generate suitable ADRs. This paper is the first of what we hope will be several on this market design. The limited experiments reported here show that it has promise. Some more benchmarking is needed on realistic dispatch systems.

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

Barty K, Carpentier P, Girardeau P (2010) Decomposition of large-scale stochastic optimal control problems. *RAIRO-Operations Research* 44(3):167–183.

Brown D, Smith J (2023) Unit commitment without commitment: A dynamic framework for managing an integrated energy system under uncertainty. *Available at SSRN* .

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 AL, Costa F, Maceira T MEand dos Santos, 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).

Dowson O, Kapelevich L (2021) SDDP.jl: a Julia package for stochastic dual dynamic programming. *INFORMS Journal on Computing* 33(1):27–33.

Eldridge B, Knueven B, Mays J (2023a) Rethinking the price formation problem–part 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 problem–part 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.

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*, volume 13.

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.

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.

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.

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

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`.

Shapiro A (2006) On complexity of multistage stochastic programs. *Operations Research Letters* 34(1):1–8.

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.