Sequential Pricing of Electricity

Jacob Mays*

School of Civil and Environmental Engineering, Cornell University, Ithaca, NY 14850

October 10, 2023

Abstract

This paper investigates the design and analysis of price formation in wholesale electricity markets given variability, uncertainty, non-convexity, and intertemporal operating constraints. The paper's primary goal is to develop a framework to assess the many resource participation models, reserve product definitions, and enhanced pricing methods that have arisen in U.S. systems, especially in the context of growing contributions from wind, solar, and storage. Departing from the static models typically used for electricity auctions based on thermal resources, the paper situates price formation within the sequential decision problem faced by system operators. This more complete description of the problem has several implications for price formation. Since prices are derived from operational models, algorithmic choices in the design of policies for the sequential decision problem influence the prices ultimately formed. In numerical tests, policy variants with comparable operational performance (within 3% in terms of total cost) lead to substantial differences in prices and resource remuneration. Storage is particularly affected, earning revenues ranging from 68% to 116% of the amount suggested as economically efficient by a benchmark approximated through stochastic programming.

Keywords: Electricity markets, price formation, reserves, storage

1 Introduction

Wholesale electricity markets operate under the premise that power can be treated as a commodity, with no differentiation between potential suppliers of the product at a given location for a given delivery period. Markets necessarily make many compromises relative to the full complexity of electricity systems, e.g., by discretizing time periods and using linear approximations to the power flow equations, leading to divergent opinions on the validity of this premise. This paper concerns one aspect of the debate, namely, how to generate sequences of prices reflecting the variability, uncertainty, non-convexity, and intertemporal constraints that govern power generation and storage. Representing one side, Tesfatsion (2023) calls for a fundamental reconsideration of pricing, arguing that intertemporal constraints render the concept of spot pricing to be conceptually flawed. Representing the other side, Biggar and Hesamzadeh (2022) questions if efforts to incorporate intertemporal constraints in price formation are even necessary. Operating between these poles, debates about the treatment of variability, uncertainty, non-convexity, and intertemporal constraints have manifested as discussion about and in some cases implementation of new storage participation models,

*jacobmays@cornell.edu

reserve products, and price formation policies in several markets across the U.S. and worldwide (Electric Power Research Institute, 2016; Sun et al., 2021). The effects of such alternative specifications are material. An analysis performed by the Independent Market Monitor of the Electric Reliability Council of Texas (ERCOT) after introduction of a new Contingency Reserve Service, for example, estimated that it had increased charges by $8.5B over a three-month period (Potomac Economics, 2023b).

This paper develops a framework to assess the effect of such proposals on market outcomes, investigating how choices made by wholesale market operators regarding algorithms for commitment, dispatch, and market clearing can affect incentives for operation and investment. As such, the paper connects both to the higher-level debate on the concept of spot pricing of electricity and to several more detailed price formation reforms currently contemplated in the academic literature and in practice. The paper's conceptual goal is a shift from the static picture of the merit-order curve in thermal-dominant markets to a dynamic understanding of price formation. Electricity prices are often colloquially described as the cost to serve an additional unit of load for a given period. In a static, convex economic dispatch model, prices that maximize efficiency both in short-run operations and long-run investment can be calculated as the dual variables corresponding to power balance constraints equating supply and demand. With no intertemporal operating constraints, dual values are typically determined by the fuel cost of thermal resources. In a dynamic model, an additional unit of load in a given period not only entails a direct cost in the present period, but also places the system in a slightly different state entering the subsequent operating period (e.g., with more or less energy stored in batteries). Dynamic models have been long been understood as necessary to the design and analysis of markets in regions with significant reservoir hydropower (see, e.g., Pereira and Pinto (1991) as well as more recent reviews in Steeger et al. (2014) and Aasgård et al. (2019)). In other regions, including the U.S. systems that serve as the primary motivation of this paper, markets have evolved in a context where storage was negligible. This state of affairs is set to change rapidly over the coming decade, as models typically find that decarbonized electricity systems will feature substantial quantities of storage (Jenkins et al., 2018; Williams et al., 2021; Frazier et al., 2021). With an estimated 680 GW of storage included in U.S. interconnection queues at the end of 2022 (Lawrence Berkeley National Laboratory, 2023), this transition appears to be well under way, necessitating a shift toward dynamic models.

The analysis begins in Section 2 by generalizing the classical result on efficiency-maximizing prices to a multi-stage stochastic setting including storage. While this theoretical property makes such prices an important benchmark, the arbitrarily large model from which they are derived cannot be used in actual operations. Thus, our interest in this paper is investigating the degree to which different implementable models for operations can lead to a stochastic process of spot prices that mimics the theoretical benchmark. To assess this question, Section 3 casts the price formation problem as one of sequential decisions under uncertainty (Powell, 2022). A central point of the paper is to emphasize a clear separation between the steps of modeling the decision problem, designing policies for it, and deriving prices consistent with the decisions made under the chosen policy. The typical approach in the analysis of impacts from non-convexity and intertemporal constraints on price formation in U.S. markets omits a full description of the sequential decision problem, instead focusing on a forward (e.g., day-ahead) market covering many time periods and neglecting the back-propagation of expectations for eventual real-time prices to those markets. This paper argues for the centrality of real-time markets, which are cleared sequentially with a single binding interval. Since profit-maximizing market participants will exploit differences between day-ahead and expected real-time prices, a model of price formation in the real-time market is a necessary precondition for describing price formation in day-ahead or other forward markets. Modeling as a sequential decision problem enables extensions from the analysis of Mays (2021b) and Eldridge et al. (2023a,b), which demonstrate on limited examples that price formation methods with attractive features in deterministic contexts do not necessarily perform well when uncertainty is considered. Without reviewing the many participation models, reserve products, and price formation proposals that have been put forward in the literature, several examples in Section 3.4 highlight how adopting a more complete dynamic model of the problem is required to completely specify how prices are formed and describe their economic properties.

After re-casting the price formation problem in this way, Section 4 formulates four different models for market clearing that represent four different policies for the sequential decision problem. Through simple numerical examples, we describe how the choice of policy and its parameterization can affect the prices that emerge. The first policy, most appropriate for a decentralized market, is a single-period deterministic economic dispatch in which the effects of intertemporal constraints and opportunity costs are incorporated into the bids and offers of market participants. Since it in principle allows the system to take maximal advantage of diffuse and continuously updating information, Biggar and Hesamzadeh (2022) argues that this approach may be more effective than attempting to directly model multiple periods in market clearing. However, since researchers typically do not have direct access to market participant views, this approach is of limited use for the analysis of markets. The second is a deterministic lookahead model with net load biasing similar to that employed by U.S. systems for reliability unit commitment processes and/or for market clearing. Relative to the literature on multi-period market clearing, this paper's focus in this context is on the challenge posed by the need to calibrate bias parameters when using this approach (cf. Mays (2021b)). The third is a deterministic lookahead model with reserve tuning, in which parameters associated with financially binding reserve products are calibrated rather than non-binding energy forecasts. While a primary focus of the literature has been how to choose reserve *quantities*, this paper's focus is on how the third approach differs from the second in its effect on *prices* for energy and reserves. Lastly, the fourth is a stochastic programming lookahead model in which uncertainty is represented through sampling. In contrast to prior literature on stochastic market clearing, the stochastic program does not represent a day-ahead market or other forward process (Pritchard et al., 2010; Zavala et al., 2017; Cory-Wright et al., 2018; Zakeri et al., 2019); instead, as in Cho and Papavasiliou (2023), it is a multi-period model in which only the first period is binding, making it a natural and readily implementable extension from the market-clearing approaches used by the California Independent System Operator (CAISO) and New York Independent System Operator (NYISO). While these four approaches are neither exhaustive nor mutually exclusive, the overarching goal with these descriptions is to demonstrate the value of the modeling framework in evaluating different potential policies.

To illustrate the effect of different algorithmic choices, Section 5 shows results from a series of simulations covering 20 days of operations on a system similar to ERCOT in 2018 with additional storage. The numerical experiments confirm the potential for significant revenue impacts due to alternative models and parameterizations, with the total amount charged to load varying by nearly \$1.5B. Based on their connection to the idealized prices established in Section 2, we take the prices computed under the fourth, stochastic programming policy as a reference point. While the tested operating policies result in similar total cost (within 3% of the reference), fleetwide revenues vary over a much wider range, from 84.4 to 119.6% of the reference. The practical consequences of the pricing impacts of different policies are possibly most significant with respect to storage. At present, it is not clear how well U.S. markets are equipped to efficiently incorporate this new class of resource (Karaduman, 2021). In market contexts storage relies on intertemporal arbitrage, charging when prices are low and discharging when prices are high. While price volatility should be sufficient to support efficient operation of and investment in storage under idealized assumptions (Korpås and Botterud, 2020; Schmalensee, 2022), in practice systems suppress real-time price volatility through various mechanisms (Mays, 2021a). In debates about resource adequacy, it is well understood that suppressing the level of prices leads to a missing money problem and necessitates the introduction of supplemental revenue streams to achieve an efficient level of capacity (Joskow, 2008; Joskow and Tirole, 2007; Cramton et al., 2013). Similarly, a failure to allow full-strength price volatility implies a need for supplemental revenues to achieve an efficient level of flexibility. While in theory, the "first-best" approach to inducing flexibility is restoring price volatility, system operators in the U.S. have preferred "second-best" strategies, introducing new rules and instruments into the market. Thus, a key question is the degree to which alternative policies for price formation affect incentives for investment in storage. The numerical tests in Section 5 confirm that storage revenues could vary even more than that for other resources, ranging from 68.5 to 116.4% of that suggested by the reference.

The shift from static to dynamic models has several implications for the design and analysis of price formation in electricity markets. The paper closes with a discussion of policy implications and future research directions. While the experiments conducted in the paper are not sufficient to conclusively support a specific operating policy, a broad point is that the first and fourth policies described above, i.e., single-period auctions with participant-submitted opportunity costs and multi-period stochastic programs, have the strongest connection to the idealized prices that would support an efficient system in the long run. Along these lines, the primary policy recommendation is to question many recent reforms that have added complexity to market designs through new reserve products or price formation methods. These changes are typically implemented in hopes that they will improve short-run efficiency with no consideration of long-run consequences. Instead of explicitly defining reserve products, the algorithmic approaches in the first and fourth policies convey the cost of uncertainty and value of flexibility through volatility in energy prices. These approaches avoid the need for added complexity, stakeholder processes, and regular updates to ensure that reserve product specifications keep pace with an evolving resource mix. Given the theoretical disadvantages, market operators must make a strong empirical case for new market products. Accordingly, a second policy recommendation is to advocate for high-fidelity simulation tools enabling system operators to compare operational performance and pricing outcomes with alternative algorithmic choices and reserve product specifications, putting them in position to credibly demonstrate the value of new reserve products and parameterize them efficiently.

2 Benchmark prices for long-run efficiency

This section conveys a result on the prices that would support a competitive equilibrium at the socially optimal long-run resource mix assuming convex cost functions, perfect competition and complete markets in risk (Ferris and Philpott, 2022). The formulation of the problem will make clear that due to computational and informational limits, forming these prices in real-world applications is unlikely. Nevertheless, because of their connection to the long-run social optimum, they provide an important benchmark against which to assess the quality of the implementable pricing policies discussed in subsequent sections. The formulation and proof is a generalization of similar results on long-run competitive equilibrium including storage in simpler settings (Korpås and Botterud, 2020; Schmalensee, 2022).

2.1 Capacity expansion with operational uncertainty

We first define the joint capacity expansion and operational problem faced by a social planner in a risk neutral setting. The primary difference between this formulation and the standard capacity expansion model used in the economics literature is the explicit inclusion of uncertainty in the operational stage. Consider a scenario tree with nodes $n \in \mathcal{N}$. At the root node $n = 0$ we make investments in generation and storage resources, with which we operate the system through nodes $\mathcal{N}^{OP} = \mathcal{N} \setminus \{0\}$. We denote the unique predecessor of any node $n \neq 0$ as $n-$ and the set of immediate successors of node n as $n+$. The depth $\tau(n)$ of node $n \in \mathcal{N}^{OP}$ along a path is interpreted as a time index $t \in \mathcal{T}$. For convenience we will assume annualized investment costs and hourly time steps, define $\tau(0) = 0$, and define $\mathcal{T} = \{1, \ldots, 8760\}$, such that the tree represents a year of operations. We use ϕ_n to represent the probability that the path taken through the scenario tree will include node n, with $\sum_{n \in \mathcal{N}:\tau(n)=t} \phi_n = 1 \; \forall t \in \mathcal{T}$. Here we define notation for a relatively simple version of the capacity expansion problem. The important result of the proof extends to more complicated formulations as long as convexity is maintained. We complicate the formulation further, including by the introduction of non-convexity, in the next section.

While most of the simplifications are unimportant for the theoretical result, we draw attention to the absence of reserve products. Since the stochastic model endogenously determines the need to prepare for future intervals, products intending to address uncertainty across intervals do not need to be specified. At the same time, products to address uncertainty within a given pricing interval or to compensate for features not included in the model (e.g., voltage support) would still be necessary.

2.1.1 Notation

Sets:

- $g \in \mathcal{G}$: generators
- $b \in \mathcal{B}$: storage resources
- $l \in \mathcal{L}$: demand bids

Parameters:

- C_g^{INV}: investment cost for generator g (\$/MW-yr)
- C_b^{INV}: investment cost for storage resource b (\$/MW-yr)
- C_g^{OP}: marginal cost for operating generator g (\$/MWh)
- E_b: duration of storage resource b (hrs)
- V_l^L: value of load l (\$/MWh)
- A_{gn}: availability of generator g in node n, as a fraction of installed capacity
- D_{ln}^+: quantity of demand bid l in node n (MW)

Decision Variables:

- y_g: capacity of generator g to install in node 0 (MW)
- y_b: capacity of storage resource b to install in node 0 (MW)
- p_{gn}: power output of generator g in node n (MW)
- k_{bn}^c: power charged by storage resource b in node n (MW)
- k_{bn}^d: power discharged by storage resource b in node n (MW)

- j_{bn}: state of charge for storage resource b in node n (MWh)
- d_{ln}: amount of demand l cleared in node n (MW)

2.1.2 Formulation

The social planner's problem is stated as

(SOC)

$$\max_{y,p,k^c,k^d,j,d} \quad -\sum_{g\in\mathcal{G}} C_g^{INV} y_g - \sum_{b\in\mathcal{B}} C_b^{INV} y_b + \sum_{n\in\mathcal{N}^{OP}} \sum_{l\in\mathcal{L}} \phi_n V_l^L d_{ln} - \sum_{n\in\mathcal{N}^{OP}} \sum_{g\in\mathcal{G}} \phi_n C_g^{OP} p_{gn} \tag{1a}$$

subject to

$$\sum_{l\in\mathcal{L}} d_{ln} - \sum_{g\in\mathcal{G}} p_{gn} - \sum_{b\in\mathcal{B}} (k_{bn}^d - k_{bn}^c) = 0 \qquad \forall n \in \mathcal{N}^{OP} \qquad (\phi_n \lambda_n) \tag{1b}$$

$$d_{ln} \le D_{ln}^+ \qquad \forall l \in \mathcal{L}, n \in \mathcal{N}^{OP} \qquad (\phi_n \alpha_{ln}^D) \tag{1c}$$

$$p_{gn} - A_{gn} y_g \le 0 \qquad \forall g \in \mathcal{G}, n \in \mathcal{N}^{OP} \qquad (\phi_n \alpha_{gn}^{P+}) \tag{1d}$$

$$k_{bn}^c - y_b \le 0 \qquad \forall b \in \mathcal{B}, n \in \mathcal{N}^{OP} \qquad (\phi_n \alpha_{bn}^{K-}) \tag{1e}$$

$$k_{bn}^d - y_b \le 0 \qquad \forall b \in \mathcal{B}, n \in \mathcal{N}^{OP} \qquad (\phi_n \alpha_{bn}^{K+}) \tag{1f}$$

$$j_{bn} - j_{b,n-} + k_{bn}^d - k_{bn}^c = 0 \qquad \forall b \in \mathcal{B}, n \in \mathcal{N}^{OP} \qquad (\phi_n \beta_{bn}) \tag{1g}$$

$$j_{bn} - E_b y_b \le 0 \qquad \forall b \in \mathcal{B}, n \in \mathcal{N}^{OP} \qquad (\phi_n \gamma_{bn}) \tag{1h}$$

$$j_{b0} = 0 \qquad \forall b \in \mathcal{B} \qquad (\phi_0 \gamma_{b0}) \tag{1i}$$

$$d_{ln} \ge 0 \qquad \forall l \in \mathcal{L}, n \in \mathcal{N}^{OP} \tag{1j}$$

$$p_{gn} \ge 0 \qquad \forall g \in \mathcal{G}, n \in \mathcal{N}^{OP} \tag{1k}$$

$$j_{bn}, k_{bn}^d, k_{bn}^c \ge 0 \qquad \forall b \in \mathcal{B}, n \in \mathcal{N}^{OP}. \tag{1l}$$

The social planner maximizes the value of load served across all nodes $n \in \mathcal{N}^{OP}$, subtracting the upfront cost of investment in node $n = 0$ and the operating cost in all other nodes. Dual variables are scaled by ϕ_n in order to produce unscaled prices (e.g., such that λ_n corresponds to the normal interpretation of the cost to serve an additional unit of load in node n). Constraint (1b) enforces power balance. With the objective, constraint (1c) models load curtailment. Constraint (1d) represents the maximum output of generators, which depends on their node-dependent availability and the decision to install capacity in node 0. Constraints (1e) and (1f) represent maximum charging and discharging from storage resources. Constraint (1g) enforces consistency in the state of charge for storage resources and constraint (1h) encodes the maximum state of charge, with constraint (1i) establishing an initial condition.

2.2 Market participants

With the same notation, we can formulate models for each generation and storage resource to maximize their operating profit given a fixed capacity. Without explicitly modeling financial trades, the assumptions of complete markets in risk allow us to make the simplifying assumption that all participants are risk-neutral with respect to the same socially determined probabilities ϕ_n (Philpott et al., 2016; Gérard et al., 2018). With capacity fixed at y_g and λ_n representing the price of power in node n, generator $g \in \mathcal{G}$ solves

$$(GEN)_g \quad \max_{p} \quad \sum_{n \in \mathcal{N}^{OP}} \phi_n p_{gn}(\lambda_n - C_g^{OP}) \tag{2a}$$

$$\text{subject to} \quad p_{gn} \leq A_{gn} y_g \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \alpha_{gn}^{P+}) \tag{2b}$$

$$p_{gn} \geq 0 \quad \forall n \in \mathcal{N}^{OP}. \tag{2c}$$

Similarly, with capacity given at y_b, storage resource $b \in \mathcal{B}$ solves

$$(STOR)_b \quad \max_{k^c, k^d, j} \quad \sum_{n \in \mathcal{N}^{OP}} \phi_n \lambda_n (k^d - k^c) \tag{3a}$$

$$\text{subject to} \quad k_{bn}^c \leq y_b \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \alpha_{bn}^{K-}) \tag{3b}$$

$$k_{bn}^d \leq y_b \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \alpha_{bn}^{K+}) \tag{3c}$$

$$j_{bn} - j_{b,n-} + k_{bn}^d - k_{bn}^c = 0 \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \beta_{bn}) \tag{3d}$$

$$j_{bn} \leq E_b y_b \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \gamma_{bn}) \tag{3e}$$

$$j_{b0} = 0 \tag{3f}$$

$$j_{bn}, k_{bn}^d, k_{bn}^c \geq 0 \quad \forall n \in \mathcal{N}^{OP}. \tag{3g}$$

Lastly, load $l \in \mathcal{L}$ solves the problem

$$(LOAD)_l \quad \max_{d} \quad \sum_{n \in \mathcal{N}^{OP}} \phi_n d_{gn}(V_l^L - \lambda_n) \tag{4a}$$

$$\text{subject to} \quad d_{ln} \leq D_{ln}^+ \quad \forall n \in \mathcal{N}^{OP} \quad (\phi_n \alpha_{ln}^D) \tag{4b}$$

$$d_{ln} \geq 0 \quad \forall n \in \mathcal{N}^{OP}. \tag{4c}$$

2.3 Long-run equilibrium

Using these problem definitions, and retaining the assumptions of perfect competition and complete markets in risk, we define a long-run competitive equilibrium for the market as follows.

Definition 1. A long-run competitive equilibrium is a set of prices $(\lambda_n)_{n\in\mathcal{N}^{OP}}$, installed generation capacities $(y_g)_{g\in\mathcal{G}}$, installed storage capacities $(y_b)_{b\in\mathcal{B}}$, generator schedules $(p_{gn})_{g\in\mathcal{G},n\in\mathcal{N}^{OP}}$, storage schedules $(k^c_{bn}, k^d_{bn}, j_{bn})_{b\in\mathcal{B},n\in\mathcal{N}^{OP}}$, and load consumption quantities $(d_{ln})_{l\in\mathcal{L},n\in\mathcal{N}^{OP}}$ such that:

1. the generator schedules solve $(GEN)_g$ $\forall g \in \mathcal{G}$, the storage schedules solve $(STOR)_b$ $\forall b \in \mathcal{B}$, and the load consumption quantities solve $(LOAD)_l$ $\forall l \in \mathcal{L}$ at the given prices and installed capacities;
2. the market clears in every operating node, i.e., $\sum_{l\in\mathcal{L}} d_{ln} - \sum_{g\in\mathcal{G}} p_{gn} - \sum_{b\in\mathcal{B}} (k^d_{bn} - k^c_{cn}) = 0 \quad \forall n \in \mathcal{N}^{OP}$; and
3. all generators and storage resources have zero expected profit at the assumed annualized investment cost, i.e.,

$$0 \le y_g \perp C^{INV}_g y_g - \sum_{n\in\mathcal{N}^{OP}} \phi_n p_{gn}(\lambda_n - C^{OP}_g) \ge 0 \quad \forall g \in \mathcal{G}$$

and

$$0 \le y_b \perp C^{INV}_b y_b - \sum_{n\in\mathcal{N}^{OP}} \phi_n \lambda_n (k^d - k^c) \ge 0 \quad \forall b \in \mathcal{B}.$$

The following theorem shows that a solution to model (SOC) satisfies these conditions, with the proof provided in Appendix A.

Theorem 1. Let $(y^*, p^*, k^{c*}, k^{d*}, j^*, d^*)$ be an optimal solution to (SOC) and let $(\lambda^*, \alpha^*, \beta^*, \gamma^*)$ be an optimal solution to its dual. Then the prices $(\lambda^*_n)_{n\in\mathcal{N}^{OP}}$, installed generation capacities $(y^*_g)_{g\in\mathcal{G}}$, installed storage capacities $(y^*_b)_{b\in\mathcal{B}}$ generator schedules $(p^*_{gn})_{g\in\mathcal{G},n\in\mathcal{N}^{OP}}$, storage schedules $(k^{c*}_{bn}, k^{d*}_{bn}, j^*_{bn})_{b\in\mathcal{B},n\in\mathcal{N}^{OP}}$, and load consumption quantities $(d^*_{ln})_{l\in\mathcal{L},n\in\mathcal{N}^{OP}}$ represent a long-run competitive equilibrium.

The relationship between competitive equilibrium and social optimum is an important justification for the regulatory decision to establish competitive markets in electricity. Accordingly, this paper takes prices derived in this theoretically efficient way to be an appropriate benchmark for evaluating alternative price formation proposals. However, a clear drawback of model (1) is that, despite simplifications in the technical constraints, it is impossible to solve as formulated in any practical setting. In the next section, we turn to implementable models for the design and analysis of price formation.

3 Modeling the sequential price formation problem

We now turn to a more realistic description of how system operators solve the decision problem faced at each node $n \in N^{OP}$ in model (1). Taking investment decisions as fixed and following Powell (2022), a canonical model for the operational problem over time periods $\mathcal{T} = \{1, \dots, T\}$ can be stated as

$$\max_{\pi\in\Pi} \mathbb{E}\left[\sum_{t=1}^{T} C_t(S_t, X^\pi_t(S_t))\right]. \tag{5}$$

Here C_t represents benefits achieved at time t, which depends on the state of the system S_t and decision x_t, which is determined by some policy π in a set of potential policies Π, with $X_t^\pi(S_t)$ mapping states to actions. In the electricity context, S_t may include the current status of generators, the state of charge for batteries, and beliefs regarding future load, component failures, and the availability of wind and solar. Decisions in x_t can then include commitment status and dispatch levels for each resource in the system. Further, with W_{t+1} representing the new information available between time t and $t+1$, the transition function $S_{t+1} = S^M(S_t, x_t, W_{t+1})$ describes how the system state may evolve into the next interval depending on actions taken and the realization of random variables. The decision problem is to identify an eligible policy π that will maximize net benefits accruing over the time period under consideration. Noting that system operators may prefer to use an averse risk measure, for the time being Eq. (5) assumes risk neutrality.

3.1 Base policy: deterministic lookahead unit commitment

In electricity systems, the most common policy approach involves the solution of a series of deterministic lookahead unit commitment problems (e.g., every 15 minutes, with lookahead periods ranging from several hours to a full day ahead) along with more frequent deterministic economic dispatch problems with a shorter lookahead horizon (e.g., a single-period model every 5 minutes). As a starting point for the discussion and since it is commonly used in the analysis of power systems, we formulate a deterministic lookahead unit commitment model to be solved sequentially for each hour-long period t. Because of their relevance to debates about price formation, the formulation includes several technical constraints beyond those in model (1). To simplify notation relative to the state S_t, the model includes H hours of history relevant to current decisions along with F future hours for which the system operator wishes to make provisional decisions. Section 4 explores modifications to this base policy.

3.1.1 Notation

Retaining the definition of the sets $\mathcal{G}$, $\mathcal{B}$, and $\mathcal{L}$ as well as parameters C_g^{OP} and V_l^L from above, we alter the indices of most other parameters and variables to reflect the sequential decision framework, replace the investment decision variables with parameters, and define additional notation to describe additional technical constraints.

Sets:

- $t' \in \mathcal{T}_t^{H,F}$: time periods included in model t
- $t' \in \mathcal{T}_t^{H} = \{t-H, t-H+1, \ldots, t-1\}$: prior time periods included in model t
- $t' \in \mathcal{T}_t^{F} = \{t+1, \ldots, t+F\}$: lookahead time periods included in model t

Known Parameters:

- C_g^{NL}: no-load cost for generator g (\$/period)
- C_g^{SU}: start-up cost for generator g (\$)
- UR_g^+, DR_g^+: maximum up and down per-period ramping capability of generator g (MW/h)
- UT_g, DT_g: minimum up and down time for generator g
- J_b: energy capacity of storage resource b (MWh)
- K_b: power capacity of storage resource b (MW)
- η_b^c: charging efficiency of storage resource b (p.u.)
- η_b^d: discharging efficiency of storage resource b (p.u.)
- V^R: penalty cost for reserve shortfall (\$/MWh)
- $R_{tt'}^+$: reserve requirement in time period t' in model t (MW)

Estimated Parameters:

- $\tilde{P}_{gtt'}^+, \tilde{P}_{gtt'}^-$: maximum and minimum availability for generator g in time period t', as estimated at time t (MW)
- $\tilde{D}_{ltt'}^+$: quantity of demand bid l in time period t', as estimated at time t (MW)

Decision Variables:

- $\tilde{u}_{gtt'}$: (binary) commitment status of generator g in time period t', provisional at time t
- $\tilde{v}_{gtt'}$: (binary) start-up decision for generator g time period t', provisional at time t
- $\tilde{q}_{gtt'}$: (binary) shut-down decision for generator g in time period t', provisional at time t
- $\tilde{p}_{gtt'}$: power output of generator g in time period t', provisional at time t (MW)
- $\tilde{r}_{gtt'}$, $\tilde{r}_{btt'}$: reserve provided by generator g or storage resource b in time period t', provisional at time t (MW)
- $\tilde{k}_{btt'}^c$: power charged by storage resource b in time period t', provisional at time t (MW)
- $\tilde{k}_{btt'}^d$: power discharged by storage resource b in time period t', provisional at time t (MW)
- $\tilde{j}_{btt'}$: state of charge for storage resource b in time period t', provisional at time t (MWh)
- $\tilde{d}_{ltt'}$: amount of demand l cleared in time period t', provisional at time t (MW)
- $\tilde{z}_{tt'}$: total quantity of reserves provided by all resources in time period t, provisional at time t (MW)

A tilde over a parameter indicates that it is estimated based on the state S_t in which the system is currently positioned. For $t' \in \mathcal{T}_t^H$ as well as $t' = t$, we assume that these values are known with certainty. Given this assumption for $t' = t$, this paper focuses on variability and uncertainty across pricing intervals, rather than within them. A tilde over a decision variable indicates that it is a provisional decision made at time t for time t'. We also define a separate set of parameters without tildes to indicate the past decisions made at time $t' \in \mathcal{T}_t^H$ for period t', i.e., the actual binding decisions made under the chosen policy in the period t' model. For example, the parameter $u_{gt'}$ indicates the actual commitment status of generator g in time period t', as determined at time t'. We note that although they are irrelevant as primal decision variables, we define notation for the decisions for time periods preceding time t because some price formation methods in the literature and in practice incorporate start-up and no-load costs arising from prior decisions. We use the notation $\tilde{x}_t$ to represent the whole set of provisional decisions made in the period t model, i.e., $\tilde{x}_t = (\tilde{u}_t, \tilde{v}_t, \tilde{q}_t, \tilde{p}_t, \tilde{r}_t, \tilde{k}^c_t, \tilde{k}^d_t, \tilde{j}_t, \tilde{d}_t, \tilde{z}_t)$, and x_t from above to represent the entire set of binding decisions made in the period t model for time period t.

It is worth highlighting certain assumptions implicit in the notation, some of which could have important consequences for a complete analysis of pricing policies. First, we assume that certain parameters relevant to generator cost functions and technical constraints are "known." In reality, these parameters are provided by the generators themselves, and as such could reflect strategic decisions that are a function of the price formation policy implemented by market operators (Duggan and Sioshansi, 2019). Further, we assume that the penalty cost for reserve shortfall V^R and the nominal reserve requirements $R^+_{tt'}$ are determined in advance rather than calculated dynamically. We note, however, that the presence of two indices on $R^+_{tt'}$ implies that the model may use a smaller reserve requirement as the relevant operating hour approaches to reflect smaller forecast uncertainty. We return to the price formation consequences of reserve tuning later in the paper.

3.1.2 Formulation

A deterministic lookahead unit commitment policy for the period t problem can be stated as

$(DLAC)$

$X_t^{DLAC}(S_t|H,F) \in$

$$\arg\max_{\tilde{x}_t} \quad \sum_{l\in\mathcal{L}} \sum_{t'\in\mathcal{T}_t^{H,F}} V_l^L \tilde{d}_{ltt'} + \sum_{t'\in\mathcal{T}_t^{H,F}} V^R \tilde{z}_{tt'} - \sum_{g\in\mathcal{G}} \sum_{t'\in\mathcal{T}_t^{H,F}} C_g^{OP} \tilde{p}_{gtt'} - \sum_{g\in\mathcal{G}} \sum_{t'\in\mathcal{T}_t^{H,F}} \left(C_g^{NL} \tilde{u}_{gtt'} + C_g^{SU} \tilde{v}_{gtt'}\right) \tag{6a}$$

subject to

$$\sum_{l\in\mathcal{L}} \tilde{d}_{ltt'} - \sum_{g\in\mathcal{G}} \tilde{p}_{gtt'} - \sum_{b\in\mathcal{B}} \left(\tilde{k}^d_{btt'} - \tilde{k}^c_{btt'}\right) = 0 \qquad \forall t' \in \mathcal{T}_t^{H,F} \qquad (\tilde{\lambda}_{tt'}) \tag{6b}$$

$$\tilde{z}_{tt'} - \sum_{g\in\mathcal{G}} \tilde{r}_{gtt'} - \sum_{b\in\mathcal{B}} \tilde{r}_{btt'} = 0 \qquad \forall t' \in \mathcal{T}_t^{H,F} \qquad (\tilde{\mu}_{tt'}) \tag{6c}$$

$$\tilde{d}_{ltt'} \leq \tilde{D}^+_{ltt'} \qquad \forall l \in \mathcal{L}, t' \in \mathcal{T}_t^{H,F} \qquad (\tilde{\alpha}^D_{ltt'}) \tag{6d}$$

$$\tilde{z}_{tt'} \leq R^+_{tt'} \qquad \forall t' \in \mathcal{T}_t^{H,F} \qquad (\tilde{\alpha}^R_{tt'}) \tag{6e}$$

$$\tilde{P}^{-}_{gtt'}\tilde{u}_{gtt'} \le \tilde{p}_{gtt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\alpha}^{P-}_{gtt'}) \quad (6f)$$

$$\tilde{p}_{gtt'} + \tilde{r}_{gtt'} \le \tilde{P}^{+}_{gtt'}\tilde{u}_{gtt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\alpha}^{P+}_{gtt'}) \quad (6g)$$

$$\tilde{k}^{c}_{btt'} \le K_b \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\alpha}^{K-}_{btt'}) \quad (6h)$$

$$\tilde{k}^{d}_{btt'} + \tilde{r}_{btt'} \le K_b \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\alpha}^{K+}_{btt'}) \quad (6i)$$

$$\tilde{j}_{btt'} - \tilde{j}_{bt,t'-1} + \frac{1}{\eta^d_b}\tilde{k}^{d}_{btt'} - \eta^c_b\tilde{k}^{c}_{btt'} = 0 \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\beta}_{btt'}) \quad (6j)$$

$$\tilde{j}_{btt'} \le J_b \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\gamma}_{btt'}) \quad (6k)$$

$$\tilde{p}_{gtt'} - \tilde{p}_{gt,t'-1} \le UR^{+}_{g} + \tilde{P}^{+}_{gtt'}\tilde{v}_{gtt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\delta}^{+}_{gtt'}) \quad (6l)$$

$$\tilde{p}_{gt,t'-1} - \tilde{p}_{gt,t'} \le DR^{+}_{g} + \tilde{P}^{+}_{gtt'}\tilde{q}_{gtt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\delta}^{-}_{gtt'}) \quad (6m)$$

$$\tilde{v}_{gtt'} + \tilde{u}_{gt,t'-1} - \tilde{u}_{gtt'} - \tilde{q}_{gtt'} = 0 \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\zeta}_{gtt'}) \quad (6n)$$

$$\sum_{t''=\max(t-H,t'-UT_g+1)}^{t'} \tilde{v}_{gtt''} - \tilde{u}_{gtt'} \le 0 \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\zeta}^{+}_{gtt'}) \quad (6o)$$

$$1 - \sum_{t''=\max(t-H,t-DT_g+1)}^{t'} \tilde{q}_{gtt''} + \tilde{u}_{gtt'} \le 0 \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \qquad (\tilde{\zeta}^{-}_{gtt'}) \quad (6p)$$

$$\tilde{d}_{ltt'} \ge 0 \qquad \forall l \in \mathcal{L}, t' \in \mathcal{T}^{H,F}_t \quad (6q)$$

$$\tilde{z}_{tt'} \ge 0 \qquad \forall t' \in \mathcal{T}^{H,F}_t \quad (6r)$$

$$\tilde{p}_{gtt'}, \tilde{r}_{gtt'} \ge 0 \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \quad (6s)$$

$$\tilde{j}_{btt'}, \tilde{k}^{d}_{btt'}, \tilde{k}^{c}_{btt'} \ge 0 \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H,F}_t \quad (6t)$$

$$\tilde{u}_{gtt'}, \tilde{v}_{gtt'}, \tilde{q}_{gtt'} \in \{0,1\} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H,F}_t \quad (6u)$$

$$\tilde{u}_{gtt'}, \tilde{v}_{gtt'}, \tilde{q}_{gtt'} = u_{gt'}, v_{gt'}, q_{gt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H}_t \quad (6v)$$

$$\tilde{p}_{gtt'}, \tilde{r}_{gtt'} = p_{gt'}, r_{gt'} \qquad \forall g \in \mathcal{G}, t' \in \mathcal{T}^{H}_t \quad (6w)$$

$$\tilde{k}^{c}_{btt'}, \tilde{k}^{d}_{btt'}, \tilde{j}_{btt'} = k^{c}_{bt'}, k^{d}_{bt'}, j_{bt'} \qquad \forall b \in \mathcal{B}, t' \in \mathcal{T}^{H}_t \quad (6x)$$

$$\tilde{d}_{ltt'} = d_{lt'} \qquad \forall l \in \mathcal{L}, t' \in \mathcal{T}^{H}_t \quad (6y)$$

$$\tilde{z}_{tt'} = z_{t'} \qquad \forall l \in \mathcal{L}, t' \in \mathcal{T}^{H}_t. \quad (6z)$$

The formulation omits several features of real-world unit commitment problems (e.g., transmission constraints, notification times for generators, and battery degradation), but is intended to have enough fidelity to include several important issues affecting price formation. The objective function maximizes projected surplus over the decision horizon, adding the value of served load and reserves minus the total cost of serving that load. Provision of reserves is assumed to have no direct cost. We model a single reserve product, assumed to be in the upwards direction only; since the numerical examples operate at an hourly time step, we are not able to capture constraints arising at the five-minute or sub-five-minute levels important in real-world markets. Costs are modeled with the traditional three-part structure used in U.S. markets. We note that this formulations does not include either opportunity costs or residual value for storage, leaving these modifications for future sections.

The formulation preserves the constraints from model (1), incorporating changes in notation consistent with the omission of investment variables and the switch to the sequential decisions framework. Constraint (6b) enforces power balance and constraint (6c) calculates total reserves. Constraint (6d) models load curtailment while constraint (6e) models a reserve shortfall. Constraints (6f) and (6g) represent the minimum and maximum output of generators, which depends on their commitment status u_{gt}, and constraints (6h) and (6i) represent maximum charging and discharging from storage resources. Constraint (6j) encodes logic for the state of charge for storage resources and constraint (6k) encodes the maximum state of charge. Where the time index $t - H - 1$ appears, constraints can either be dropped or variables can be replaced with parameters describing initial conditions. Constraints (6l) and (6m) describe ramping limits for generators. Constraints (6n), (6o), and (6p) ensure consistency between start-up, shut-down, and commitment variables given generator minimum run times. Inclusion of these constraints reinforces the need for prior time periods $t' \in \mathcal{T}_t^H$ to fully describe the state of the system. Constraints (6q)–(6t) enforce non-negativity, constraint (6u) introduces non-convexity through the binary commitment variables, and the "non-revisionist" constraints (6v)–(6z) ensure that decisions from prior time periods ($\mathcal{T}_t^H$) cannot be altered in the current decision epoch. Greek letters to the right of each constraint define dual variables to be used in relaxations or other modifications of the mixed-integer program. The most important of these are $\tilde{\lambda}_{tt'}$ and $\tilde{\mu}_{tt'}$, which give rise to binding prices for energy and reserves in period t, $\lambda_t = \tilde{\lambda}_{tt}$ and $\mu_t = \tilde{\mu}_{tt}$, as well as advisory prices for energy and reserves in $t' \in \mathcal{T}_t^F$. We return to the topic of day-ahead markets, in which prices for multiple time periods are binding, after establishing the real-time prices around which the day-ahead markets are based.

3.2 Defining real-time pricing policies

Having defined the sequential decision problem faced by electricity system operators, as well as a base policy describing one possible approach to the problem, the remainder of the section draws out several implications for the design and analysis of price formation. The first contrast with the prior literature is in fully specifying what is required to define a policy for real-time price formation. Broadly, the literature on price formation seeks a set of prices that will support the solution identified for an instance of model ($DLAC$), with complications arising due both to uncertainty and non-convexity. Typically, this is accomplished by finding dual values corresponding to the system-wide constraints in a model that is closely related, but not identical, to model ($DLAC$). For example, the fixed binaries approach in O'Neill et al. (2005) that could be considered "traditional" locational marginal pricing (LMP) first solves the mixed-integer program in model ($DLAC$), then fixes all binary variables to their optimal values, then resolves the resulting linear program to compute dual values as prices, while the convex hull pricing of Gribik et al. (2007) instead relaxes the feasible region to its convex hull. By casting the problem in a sequential decision lens, however, it can be seen that fully describing a pricing policy requires at least three distinct steps. Of these, only the third has received major attention in the literature (exceptions include Singhal and Ela (2020), which looks at step 1, and Frew et al. (2021), which looks at step 2). The remainder of this subsection delineates these three steps.

3.2.1 Step 1: specifying the model for operations

Most analyses of price formation take the form of model ($DLAC$), or some variant, as a given (see, e.g., the review in Liberopoulos and Andrianesis (2016)). The several design choices required in the construction of model ($DLAC$) suggest why omitting this step may cause issues. Consider the lookahead horizon F. A key source of concern for California is the evening ramp, during which the system sees growing demand combined with falling production from solar. At mid-day in California, operators may already wish to begin positioning the system for the evening ramp. Necessarily, changing the form of the model (e.g., by using a larger number of hours F in the lookahead horizon) will change the prices calculated by the model. If the lookahead horizon does not include this evening ramp, then potential price impacts from the ramp will not be directly reflected in prices calculated in the model.

3.2.2 Step 2: specifying a parameterization

Model ($DLAC$) includes parameters for demand and generator availability that are estimated using information available in state S_t. Thus, a full specification of an operating policy requires a description of how these forward-looking estimates are chosen given the uncertainty in the system. It is widely understood in the context of stochastic optimization that using the expected value of uncertain parameters in a deterministic model would be unwise; model ($DLAC$) therefore entails a judicious selection of values in the parameters $\tilde{P}^{+}_{gtt'}$, $\tilde{P}^{-}_{gtt'}$, and $\tilde{D}^{+}_{ltt'}$. A standard practice in U.S. systems is to bias the amount of load upwards or bias the availability of wind and solar downwards to induce a better solution to the underlying stochastic problem (Mays, 2021b). Decisions on which parameters to bias, and the amount to do so, will necessarily affect the dual values computed when solving the model.

3.2.3 Step 3: specifying a pricing policy

Only after clarifying the choice of model and its parameterization can a pricing policy be fully defined. In general, proposals in the literature demonstrate the theoretical or empirical properties of a given policy having assumed the form of the model and its parameterization. Using the notation above, most studies assess the properties of the prices $\tilde{\lambda}_{tt'}$ $\forall t' \in \mathcal{T}_t^{H,F}$ in supporting the provisionally optimal decisions $\tilde{x}_t$ over the lookahead horizon. For example, the analysis in Gribik et al. (2007) shows that convex hull prices minimize a definition of lost opportunity costs for market participants, a property that leads Chao (2019) to claim that "convex-hull pricing is generally considered an ideal solution" to the non-convex pricing problem. With a more complete description of the sequential decision problem, however, it can be seen that many different sets of convex hull prices could be calculated from different models under different parameterizations. Among these, it is not clear which if any could be considered ideal. In a sequential decision context, defining lost opportunity costs in terms of a single instance of model $(DLAC)$ is incomplete. We do not need a set of prices $\tilde{\lambda}_{tt'}$ $\forall t' \in \mathcal{T}_t^{H,F}$ supporting the provisionally optimal decisions $\tilde{x}$ found by model $(DLAC)$, but rather a pricing policy that supports the operational policy sought in Eq. (5) and, ultimately, realized prices λ_t consistent with the binding decisions x_t determined under that policy and the realization of random variables in the system. Given that prices calculated will depend on decisions in steps 1 and 2, it is not possible to interpret theoretical results on the properties of prices derived starting only from step 3.

3.3 Defining day-ahead pricing policies

The discussion to this point has focused on real-time prices generated in a sequential manner, with the period t model giving rise to binding prices and quantities only in period t. It might be objected that day-ahead markets, on which analysis is often performed, typically cover a much larger set of time periods, simultaneously determining prices and quantities for a 24-hour period. In the context of the modeling framework above, day-ahead markets constitute special instances of model $(DLAC)$ occurring once per day, rather than every hour.

The major counterpoint to this potential objection is that for consistency in the economic analysis, a day-ahead market should converge on prices that equal the expected value of those that will eventually emerge in the real-time market in each hour. Using the notation above, if a day-ahead market executed at time t creates binding financial positions for some future time period $t' \in \mathcal{T}_t^F$, then a no-arbitrage condition entails that

$$\tilde{\lambda}_{tt'} = \mathbb{E}_{W_{t+1},\ldots,W_{t'}}[\lambda_{t'}]. \tag{7}$$

Without an explanation for why this no-arbitrage condition is not met, or at least approximated, a failure to satisfy Eq. (7) implies that the behavior of market participants has not been fully characterized. In markets that allow financial participants, virtual bids and offers will enter the day-ahead market; in markets that do not allow financial participants, physical participants will alter their bids and offers. An implication is that employing an alternative pricing policy in only the day-ahead market would have little if any impact. For example, the Independent Market Monitor for the Midcontinent ISO (MISO) reports that in 2022, the ELMP model resulted in real-time prices that were $1.45/MWh higher than they would have been under traditional locational marginal pricing (LMP), but "as expected, ELMP had almost no effect in the day-ahead market because the supply is far more flexible and includes virtual transactions" (Potomac Economics, 2023a). In other words, the formation of prices in the day-ahead markets is downstream of the policy chosen for price formation in real time. As a result, analyses that focus on price formation policies in day-ahead markets without describing how they are instantiated in real-time markets are incomplete.

3.4 Contrasts with prior literature

The goal of this subsection is to more specifically contrast the view of price formation advanced in this paper with several strands of prior literature, describing the implications for design and analysis of electricity markets along five dimensions.

3.4.1 Price formation policies and long-run efficiency

As highlighted in the analysis of Section 2, prices in wholesale electricity markets are intended to support efficiency not only in short-run operations but also in long-run investments. In the broader view, the problem in Eq. (5) is conditioned on the set of resources available in the system, and for consistency in the analysis a pricing policy must result in sufficient revenue to support an investment equilibrium including those resources. While some studies have investigated the long-run implications of alternative pricing policies (Herrero et al., 2015; Mays et al., 2021; Guo et al., 2022; Byers and Hug, 2022), the primary focus of the literature on price formation has been on short-run effects.

3.4.2 Arbitrage and production cost modeling

In addition to creating difficulties in the design of pricing policies, the no-arbitrage condition in Eq. (7) presents a challenge for analysis of market outcomes using typical production cost models, which do not include virtual bidders. A standard approach in such analyses is to mimic the two-settlement system of U.S. day-ahead and real-time markets, conducting first a day-ahead market covering the entire operating day and then a series of real-time markets with updated information. When simulating one year of operations, such analyses typically report only a single realized time series of real-time prices, making it difficult to assess conformance with Eq. (7). However, it is not uncommon for such analyses to result in persistently lower prices in day-ahead markets than in real-time markets, suggesting under-commitment of resources in the day-ahead time frame. One possible reason for this common result is the use of expected demand in day-ahead market simulations. Given the "hockey stick" shape of real-time electricity prices as a function of demand, prices calculated under the assumption of expected demand will tend to underestimate expected real-time prices (Mays, 2021b).

3.4.3 Participation models and dispatch authority

Using a model similar to model ($DLAC$), Jiang and Sioshansi (2023) investigate the question of whether market operators can be granted the authority to dispatch storage resources given that the decisions they make could affect prices. Based on duality arguments, the analysis suggests that prices will be equivalent regardless who retains authority over dispatch. The analysis, however, does not reflect how enabling independent storage operations would necessitate changes to model ($DLAC$) to include storage opportunity costs in the objective function, and in this sense does not fully consider step 1 (Section 3.2.1). Further, by not making explicit the parameterization of step 2 (Section 3.2.2), the analysis does not address how different estimates of future demand and generator availability will affect the prices calculated. Given that storage owners and system operators have different knowledge, beliefs, incentives, and levels of risk aversion, it is unlikely that the two groups would employ the same estimates for uncertain parameters and therefore unlikely that the two participation models would lead to equivalent prices. Accordingly, expanding the analysis to include these factors could affect the policy recommendation to grant dispatch authority to system operators.

3.4.4 Ramp constraints and multi-interval pricing

Several recent papers put forward proposals for price formation in multi-interval market clearing models in order to address ramping constraints that link consecutive periods (Hua et al., 2019; Chen et al., 2021; Guo et al., 2021; Cho and Papavasiliou, 2023). One of the animating concerns behind these proposals, summarized in Biggar and Hesamzadeh (2022) as the "time-inconsistency problem," is that consecutive instances of model ($DLAC$) will not necessarily result in the same prices in the presence of binding ramping constraints. In our notation, even if we modify H and F such that the period t model and period $t+1$ models cover the same time horizon, moving to period $t+1$ changes the definition of the set of prior time periods $\mathcal{T}_t^H$, expanding the number of constraints in Eqs. (6v)–(6z). These added constraints lead to different calculated dual values, even if all parameters in the models are identical. The analysis in Biggar and Hesamzadeh (2022) argues that time inconsistency is a special case of linear cost and utility functions, while the review in Hogan (2022) suggests this case is sufficiently relevant to warrant implementation of pricing logic addressing it. Viewed in a sequential decision context, however, the problem disappears: given the arrival of new information W_{t+1}, there is no reason to believe that the estimated parameters $\tilde{P}^+_{gtt'}$, $\tilde{P}^-_{gtt'}$, and $\tilde{D}^+_{ltt'}$ could be the same over two consecutive periods. Along the lines of Biggar and Hesamzadeh (2022), applying special logic to address this perceived problem could impede the incorporation of new information in prices.

Because it also includes uncertainty and a rolling horizon simulation, the modeling framework in this paper is closest to that of Cho and Papavasiliou (2023). However, it is worth mentioning the important difference that the analysis in Cho and Papavasiliou (2023) focuses on results related to a stochastic and convex version of model ($DLAC$), rather than the underlying sequential decision problem in Eq. (5). This difference has especially interesting implications when storage or non-convexity, which are not present in Cho and Papavasiliou (2023), are added to the model. In particular, because storage resources bid based on opportunity costs that are calculated based on the distribution of futures prices, different pricing policies lead to different storage offers and thus different operational decisions. We return to this discussion in Section 4.2.

3.4.5 Avoidable fixed costs and non-convex pricing

With the stated aim of "allow[ing] the commitment costs of fast-start resources to be reflected in prices" (Federal Energy Regulatory Commission, 2019), all of the FERC-jurisdictional markets except CAISO have implemented new pricing mechanisms under the umbrella of Fast-Start Pricing or Extended LMP (ELMP). Following earlier analyses of market clearing in general settings with non-convexities (Scarf, 1990), many proposals have emerged in the literature attempting to resolve the issues that arise due to the lack of uniform market-clearing prices in the context of electricity markets (see the review in Liberopoulos and Andrianesis (2016)). A prominent example is the convex hull pricing policy sometimes considered an "ideal" in deterministic settings. As discussed above, proposals in the literature typically bypass steps 1 and 2 when specifying a price formation policy. Viewing the problem in a sequential decision context thus gives a new criticism of convex hull pricing, and by extension many similar proposals. While it is sometimes suggested that the major barrier to implementing convex hull prices is computational (Chao, 2019; Knueven et al., 2021), a more fundamental issue is convex hull pricing has been defined on an incomplete model of the pricing problem. This incompleteness has important consequences. In order for commitment costs to be reflected in prices, proposals include some relaxation of the binary variables in model ($DLAC$). However, start-up and no-load costs will not be reflected in real-time prices unless a) the set of prior time periods T_t^H includes the relevant start-up decision and b) the constraints on historical consistency, Eqs. (6v)–(6z), are also relaxed in some way, such that the resulting feasible region is no longer the convex hull of the operational problem. In other words, models must (implicitly or explicitly) be allowed to adjust decisions retroactively in order for the cost of those decisions to participate directly in setting prices. Further, real-time market clearing models in practice may exclude historical decisions altogether (i.e., $H = 0$), necessitating that market operators take a different approach altogether (e.g., forming a different model in which costs from decisions in prior intervals are allocated to the present one (Wang et al., 2016)). Examples in Mays (2021b) and Eldridge et al. (2023b) demonstrate how relaxing prior decisions in real-time price formation can degrade efficiency and create poor incentives for the commitment of resources, contrary to the intentions of convex hull pricing.

4 Policies for sequential market clearing

Having described a model of the problem, this section considers the effect that different commitment and dispatch policies may have on pricing outcomes. The modifications to model ($DLAC$) considered here are not mutually exclusive, and system operators already use a combination of the first three approaches. We analyze the variants separately in order to gain intuition on the different price formation effects of each approach, using toy examples to explore the relationships between them. Throughout the examples, which assume convexity, we as a default use the simplest possible pricing policy in "step 3" described above, taking the dual variable $\lambda_t = \tilde{\lambda}_{tt}$ from Eq. (6b) to be the price in period t. In other words, the focus is on the impacts of steps 1 and 2 rather than the step 3 that has received the most attention in the literature. We first construct three toy examples, deriving the optimal dispatch and benchmark prices for each. We then consider four implementable operating policies (i.e., variants of model ($DLAC$)) and pricing policies (derived from the chosen variant of model ($DLAC$)), discussing their ability to identify the optimal dispatch decisions and approximate the benchmark prices. To distinguish between the prices generated from different operating policies, we attach a superscript to λ_t indicating the policy in use.

4.1 Toy Examples

Consider an example system with a single load $l \in \mathcal{L}$ valued at $V_l^L = \$10,000$/MWh at a known demand of $D_{lt} = 10$ MW for all periods $t \in \mathcal{T}$. The problem extends for 3 hour-long periods, i.e., $\mathcal{T} = \{1, 2, 3\}$. There are two generators, one wind and one gas, as well as one storage resource. Wind is the lone source of uncertainty. Let wind availability in MW be certain at $P^+_{wind,1} = 16$ in the first hour, distributed as $P^+_{wind,2} \sim \mathcal{U}(5, 10)$ in the second hour, and distributed as $P^+_{wind,3} \sim \mathcal{U}(0, 20)$ in the third hour, with output in the second and third hours assumed to be independent. At the beginning of period 2 we learn the actual availability of wind in period 2, but no additional information about its availability in period 3. The gas plant has a capacity of $P^+_{gas,t} = 9$ MW in each time period $t \in \mathcal{T}$, has no binary constraints or minimum operating level, and operates at a cost of $c^{OP}_{gas} = \$100$/MWh. We initially consider a case without ramping constraints, then add a limit of $UR^+_{gas} = DR^+_{gas} = 4$ MW/h for two other cases. The battery is able to store $J_b = 5$ MWh of energy and we enter the first hour with $j_{b0} = 0$ MWh stored. Storage can charge and discharge at a maximum power of $K_b = 5$ MW with perfect efficiency, no cost for degradation, and no residual value to stored energy at the end of the third hour.

The problem is simple enough that we can solve it by inspection. In period 1, since the wind exceeds demand plus the size of the storage, we serve all the load with wind, fill up the storage, and curtail the extra 1 MW of wind. The period 3 problem is also straightforward by virtue of being deterministic. We use wind and storage interchangeably up to their availability, since stored energy has no residual value at the end of the problem. After exhausting the wind and storage, we deploy gas until either its ramping or capacity constraint binds. Lastly, we curtail any residual demand that cannot be served. The period 2 problem is the most interesting, requiring us to weigh the additional cost incurred in period 2 when operating the gas plant against the lower cost expected by entering period 3 with more energy in storage and the gas plant able to ramp to a higher level. In expected value terms, the optimal decision can be found by equating the certain marginal cost of $c_{gas}^{OP} = \$100$/MWh with the expected marginal savings in period 3.

4.1.1 Example 1: no ramping constraint

First, as **Example 1**, suppose there were no ramping constraint, such that the state S_t is determined entirely by the amount of stored energy entering period t, i.e., $j_{b,t-1}$. Wind is dispatched first up to its availability. Since $P_{wind,2}^{+} \sim \mathcal{U}(5, 10)$, $j_{b1} = 5$, and $D_{l2} = 10$, it is always possible to meet demand with wind and storage, but the question is whether it would be superior to reserve more energy in storage and instead use more gas. Depending on the realization of wind in period 3, the marginal recourse action in period 3 may also be using gas, in which case having additional energy in storage allows us to save \$100/MWh, or it may be curtailing load, in which case it allows us to save \$10,000/MWh. In expected value terms the operator is willing to use storage down to a level of 0.89 MWh, at which the probability of curtailing load in period 3 (which occurs when $P_{wind,3}^{+} < 0.11$) is $.11/20 = 0.0055$ and the probability of gas being on the margin (which occurs when $0.11 \leq P_{wind,3}^{+} < 9.11$) is 0.45. Leaving $j_{b2} = .89$ MWh of energy in storage equates the expected period 3 marginal savings of $0.0055 \cdot \$10000/\text{MWh} + 0.45 \cdot \$100/\text{MWh}$ with the certain marginal cost of \$100/MWh incurred by deploying additional gas in period 2. This efficient trade-off holds regardless of the amount of wind available in period 2, so the optimal dispatch will be first to dispatch wind up to its full availability, then storage down to a level of 0.89 MWh unless all demand is met first, then gas as needed to meet any remaining demand.

4.1.2 Examples 2 and 3: with ramping constraint

Next, consider the case with the ramping constraint. In this case, the state S_t includes the state of charge for the battery as well as the output of the gas plant in the previous period, i.e., $p_{g,t-1}$. Since this added complexity makes it harder to write down a solution for all potential values of $P_{wind,2}^{+}$, we consider the situations of $P_{wind,2}^{+} = 10$ (**Example 2**) and $P_{wind,2}^{+} = 5$ (**Example 3**). In Example 2, we are able to serve load entirely with wind, but we prefer to use some gas in order to enable it to ramp higher in period 3. In instances where there is insufficient ramping capability, having used extra gas in period 2 leads to savings of

\$10,000 - \$100 = \$9,900. Otherwise, there are no savings from having pre-positioned gas. Thus the optimal dispatch reduces the probability of shortfall in period 3 to 100/9900 = 0.0101, which occurs when wind availability $P^+_{wind,3} < 0.0101 \cdot 20 = 0.202$. This implies gas dispatch of 0.798 MW in period 2.

With the ramping constraint and $P^+_{wind,2} = 5$, we are unable to meet period 2 load with wind but it is always optimal to meet period 2 load in full. Since demand $D_{l2} = 10$, discharge of storage will mirror the dispatch of gas such that they combine to the 5 MW required to meet demand after accounting for wind. Since stored energy $j_{b1} = 5$, this observation implies that the amount left in storage entering period 3, j_{b2}, will be precisely equal to the amount of gas dispatched in period 2, $p_{gas,2}$. As a consequence, we can write the potential savings in period 3 as a function of $p_{gas,2}$, the amount of gas dispatched in period 2. If it turns out that there is excess wind available, i.e., $P^+_{wind,3} > D_{l3} - p_{gas,2}$, then having additional energy in storage and ramping capability will not affect period 3 costs. If gas is on the margin and within its ramping capability, i.e., $D_{l3} - p_{gas,2} > P^+_{wind,3} > D_{l3} - 2 \cdot p_{gas,2} - UR^+_{gas}$, then having additional energy in storage saves \$100/MWh but having greater gas availability does not make a difference. Lastly, if wind is very low and the gas ramping constraint is binding, i.e., $P^+_{wind,3} < D_{l3} - 2 \cdot p_{gas,2} - UR^+_{gas}$, then having additional energy in storage saves \$10,000/MWh and having additional ramp capability saves \$9,900/MWh, for a total of \$19,900/MWh. Setting the expected period 3 savings equal to the certain period 2 cost of \$100/MWh leads to an optimal gas dispatch of $p_{gas,2} = 2.967$ MW. This dispatch implies discharge from storage of $k^d_{b2} = 2.033$ MW and curtailment of demand in period 3 whenever $P^+_{wind,3} < 0.065$ MW.

4.1.3 Benchmark prices

In setting up these examples, we intentionally determined optimal decisions from first principles rather than by solving an optimization model. The problem is simple enough that, however, that with small modifications we could have used the idealized stochastic program from model (SOC) with arbitrarily many operating nodes. More specifically, the examples can be considered as sub-trees starting from nodes with depth $\tau(n) = 1$ and adding the generator ramping constraint. In order to construct examples that make the pricing results clearer, the investment stage of node 0 is omitted; as a consequence the resource mix may not be optimal for these specific sub-trees. Nevertheless, the prices in these sub-trees still contribute to overall remuneration, so it is still possible to meaningfully discuss the different long-run incentives provided by different pricing policies. The optimal dispatch and benchmark prices for period 2 in the examples are summarized in Table 1.

It can be observed that the benchmark price in many instances corresponds to the residual value of storage, rather than the operating cost of wind or gas. To discuss further how these prices arise, we turn to the first implementable policy.

Table 1: Optimal Dispatch and Benchmark Prices in Period 2

	Wind Avail. $P^+_{wind,2}$	Wind Disp. $p^*_{wind,2}$ (MW)	Gas Disp. $p^*_{gas,2}$ (MW)	Storage Disc. k^{d*}_{b2} (MW)	Price λ^*_2 (\$/MWh)
	$6 < P^+_{wind,2} \leq 10$	$P^+_{wind,2}$	0	$10 - P^+_{wind,2}$	$25 + 5 \cdot (P^+_{wind,2} - 6)$
Ex. 1	$5.89 < P^+_{wind,2} \leq 6$	$P^+_{wind,2}$	0	$10 - P^+_{wind,2}$	$45 + 500 \cdot (P^+_{wind,2} - 5.89)$
	$5 \leq P^+_{wind,2} \leq 5.89$	$P^+_{wind,2}$	$5.89 - P^+_{wind,2}$	4.11	100
Ex. 2	10	9.202	0.798	0	0
Ex. 3	5	5	2.967	2.033	67.58

4.2 Opportunity cost bidding

The first approach modifies the objective function in model $(DLAC)$ to include bids and offers from operators of battery storage. In the simplest setting, this implementation can take the form of a single-period model (Biggar and Hesamzadeh, 2022). In this case, views on the future evolution of prices must be incorporated in the behavior of market participants. With $V^C_{btt'}$ representing the bid price at which storage resource b will charge in period t' and $V^D_{btt'}$ representing the offer price at which it will discharge, with the additional t index indicating that these values can be different in each sequential model, we add terms to the objective function in Eq. (6a) to formulate the opportunity cost bidding policy as follows:

$$
\begin{aligned}
&(DLAC - OC)\\
&X_t^{DLAC-OC}(S_t|H,F) \in\\
&\arg \max_{\tilde{u},\tilde{v},\tilde{q},\tilde{p},\tilde{r},\tilde{k}^c,\tilde{k}^d,\tilde{j},\tilde{d},\tilde{z}} \quad \sum_{l\in\mathcal{L}} \sum_{t'\in\mathcal{T}_t^{H,F}} V_l^L \tilde{d}_{ltt'} + \sum_{t'\in\mathcal{T}_t^{H,F}} V^R \tilde{z}_{tt'} + \sum_{b\in\mathcal{B}} \sum_{t'\in\mathcal{T}_t^{H,F}} \left(V^C_{btt'} \tilde{k}^c_{btt'} - V^D_{btt'} \tilde{k}^d_{btt'} \right)\\
&\qquad - \sum_{g\in\mathcal{G}} \sum_{t'\in\mathcal{T}_t^{H,F}} C_g^{OP} \tilde{p}_{gtt'} - \sum_{g\in\mathcal{G}} \sum_{t'\in\mathcal{T}_t^{H,F}} \left(C_g^{NL} \tilde{u}_{gtt'} + C_g^{SU} \tilde{v}_{gtt'} \right) \qquad (8)\\
&\text{subject to} \qquad \text{(6b)–(6z)}.
\end{aligned}
$$

We note that while real-world markets allow submission of piecewise linear bid and offer segments, the formulation here simplifies to a single price for storage charging and discharging (as it does for load bids and generator offers).

In Example 1 described above, consider the situation of the battery in period 2 and assume we employ a single-period model for dispatch (i.e., $H = F = 0$). The battery will submit offers based on its view of prices in period 3. In period 3, we formulate a deterministic, single-period version of model (8) without binary variables in which the battery offers at \$0/MWh. In the simplest pricing policy, the price is computed as a dual value from the power balance constraint in Eq. (6b). Since wind is uniformly distributed as $P^+_{wind,3} \sim \mathcal{U}(0, 20)$, it is straightforward to calculate the distribution of prices that arise in period 3 as a function of stored energy j_{b2}. If $P^+_{wind,3} + j_{b2} \geq D_{l3}$, then wind and storage can satisfy all the load and $\lambda_3 = \$0$/MWh. If $D_{l3} - P^+_{gas,3} < P^+_{wind,3} + j_{b2} < D_{l3}$, then the gas plant must be engaged and $\lambda_3 = \$100$/MWh. Lastly, if $P^+_{wind,3} + j_{b2} < D_{l3} - P^+_{gas,3}$ then load must be curtailed and $\lambda_3 = \$10,000$/MWh. Thus, if we were to enter period 3 with 0 MWh in storage, there would be a 50% chance that the price would be \$0/MWh, a 45% chance it would be \$100/MWh, and a 5% chance it would be \$10,000/MWh. If instead we enter period 3 with 1 MWh in storage, there would still be a 45% chance of a price of \$100/MWh, but there would be no chance of curtailing load. If on the other hand we enter period 2 with 5 MWh in storage, the probability of engaging gas would drop to 25% and there would be still be no chance of curtailing load.

These probabilities imply prices for period 2 determined through the offers of storage. Assuming perfectly competitive, risk neutral offers, the marginal offer from storage in period 2 equals the expected price in period 3, conditional on the amount of residual stored energy. Thus, storage enters period 2 with $j_{b1} = 5$ MWh of energy and submits an offer curve that grows linearly from \$25/MWh to \$45/MWh over the range [0,4], then grows linearly from \$45/MWh to \$545/MWh over the range [4,5]. These offers reflect residual values of stored energy consistent with the conditional expectation of prices in period 3 if storage is dispatched up to a given point in period 2. This offer curve implies that in a competitive auction for period 2, storage will clear before gas, which offers at \$100/MWh, up to a level of $4 + (100 - 45)/(545 - 45) = 4.11$ MW, i.e., precisely the efficient amount determined in Section 4.1. Correspondingly, the price is set by storage whenever wind availability $P^+_{wind,2} < 5.89$, and the resulting prices match the benchmark shown in Table 1.

This single-period model is perhaps the most general in terms of real-world markets, as it does not require any special assumptions about the market operator's ability to make decisions centrally. In more decentralized settings (e.g., the Western Energy Imbalance Market or the Southeast Energy Exchange Market in the U.S.), participants submit storage offers based on their private view of opportunity costs, based on their own projected needs as well as their estimates of future price distributions.

Example 1 shows how, in certain cases and under certain assumptions, the optimal dispatch and prices in a multi-period problem can be achieved using only a single-period model with appropriate definition of opportunity costs. Despite this promise, under more realistic assumptions a lack of coordination can lead to production inefficiency Philpott et al. (2010). At least three potential issues arise in determining opportunity costs in the examples, to which we devote the remainder of this subsection.

4.2.1 Pre-positioning and market power

In Example 2, since wind is curtailed in period 2, marginal consumption would be met by reducing the curtailment of wind, implying an efficient price of \$0/MWh. Despite a price of \$0/MWh, gas is dispatched at 0.798 MW. In an analysis similar to that performed for Example 1, we would find that the distribution of prices in period 3 implied by the gas unit's ramping constraint make it willing to incur a loss in period 2 in order to position itself for higher potential profits in period 3. In other words, the efficient dispatch could be found if it submitted an offer below \$0/MWh for its first 0.798 MW of production in period 2, leading it to be dispatched ahead of wind in the single-period model. In doing so, it would incur an operating loss of \$79.80 in period 2. In this example, the unit would make up this loss in the 1.01% of cases when the ramping constraint is binding in period 3 and prices are set by curtailed demand at \$10,000/MWh.

While the separation between the \$100/MWh cost of the gas plant and the \$10,000/MWh value of lost load heightens the contrast in the toy example, the more general question is whether system operators can trust the process by which operators of technically constrained plants (gas in the example) would pre-position equipment on the basis of future price expectations, especially when doing so would entail an upfront cost. Recalling the assumption of complete markets in risk from the idealized model in Section 2, the viability of relying on such pre-positioning could be enhanced if generators had the ability to "lock in" the higher prices expected in future intervals.

An additional related question is whether effective market power mitigation is possible in such cases: in the example, the gas plant could increase the frequency of scarcity pricing in period 3 by submitting an offer of \$100/MWh in period 2 for its first 0.798 MW of production instead of submitting at a negative price. Since such an offer would be easily defensible as consistent with its operating costs, it is unclear how a market power screen would identify it.

4.2.2 State-dependent offers and gate closure

In Example 3, the optimal solution is to dispatch storage at 2.033 MW and gas at 2.967 MW, based on which the expected value of prices in period 3 can be calculated as \$67.58/MWh. This expected price implies that to achieve this dispatch in period 2, the marginal price from storage at this point of its offer curve, and as such the clearing price, must be \$67.58/MWh. As in Example 2, this price entails that the gas plant incurs a loss in period 2 and requires that it nevertheless submit its first 2.967 MW of production at a price below \$67.58/MWh. An additional complexity of Example 3 relative to the previous examples, however, is that the distribution of future prices depends on both components of the state variable. In order for the battery operator to construct its offer curve it must assess the residual value of storage. Since prices in period 3 depend on how high the gas plant is able to ramp, the residual value is conditional on the dispatch of the gas plant in period 2, which the storage operator does not know when submitting its offers for period 2. Heightening the challenge, gate closure in some real-world systems occurs well in advance of operations (e.g., 75 minutes in CAISO).

4.2.3 Market design vs. market analysis

The updated objective function in Eq. (8) makes use of the parameters $V^C_{btt'}$ and $V^D_{btt'}$, which represent bid and offer prices for storage resources. Since researchers do not typically have access to participant bids and offers, the need to derive these values limits the usefulness of this policy variant for the analysis of markets. In the context of design, however, market operators can simply request that these values be provided by market participants prior to each auction. For that reason, the notation we use omits the tilde that we use for estimated load and generator availability. However, there are several contexts in which system operators may need to independently assess appropriate values to use. First, as mentioned in Section 3.4.3, some have suggested that system operators be authorized to dispatch storage themselves. Second, as a market power mitigation practice, many operators have the ability to substitute participant offers with lower "cost-based" offers. In the case of storage resources, developing sound cost-based offers requires that market monitors take a view on participant opportunity costs (Muñoz et al., 2018; Lavin et al., 2021). We note in this context that these opportunity costs are unrelated to the cost of charging the resource, which is zero in the examples. Third, many systems run non-market processes (e.g., reliability unit commitment models) that may have longer lookahead horizons or other aspects that differ from the market clearing models. Since participant bids and offers may only be available for time periods relevant to market clearing, system operators may need their own estimate on these values for other time periods included in these models.

4.3 Deterministic lookahead with net load biasing

The second implementable policy we consider, one that is common in both the design and the analysis of markets, uses the deterministic model ($DLAC$) directly, but biases the values of $\tilde{P}^+_{gtt'}$ for maximum available power, $\tilde{P}^-_{gtt'}$ for minimum available power, and $\tilde{D}^+_{ltt'}$ for demand to induce a higher quality solution to the underlying stochastic problem. Removing the tildes that indicate the estimate to be used in the model, let P^+_{gtt}, $P^-_{gtt'}$, and $D^+_{ltt'}$ be random variables representing the corresponding uncertain parameters. Based on the conditional distribution assessed given information in state S_t, let $\hat{P}^+_{gtt',\theta}$ be an estimate of the θ quantile of the maximum power available from generator g in time t', as estimated at time t, with $\hat{P}^-_{gtt',\theta}$ and $\tilde{D}^+_{ltt',\theta}$ defined in an analogous way. Since this approach to modeling uncertainty is best suited to variable renewables, let $\mathcal{G}^V \subseteq \mathcal{G}$ represent a subset of generators whose uncertainty we wish to address in the model. To simplify the notation and guarantee feasibility, and because these inverter-based resources are typically curtailable, we assume $P^-_{gtt} = 0$ and thus $\hat{P}^-_{gtt',\theta} = 0 \; \forall g \in \mathcal{G}^V$.

To define the deterministic lookahead with net load biasing ($DLAC - NLB - \theta$) policy, we use the following values for the estimated parameters in model ($DLAC$):

$$\tilde{P}^{-}_{gtt'} = 0 \qquad \forall g \in \mathcal{G}^V, \forall t' \in \mathcal{T}^F_t;$$
$$\tilde{P}^{+}_{gtt'} = \hat{P}^{+}_{gtt',\theta} \qquad \forall g \in \mathcal{G}^V, \forall t' \in \mathcal{T}^F_t; \text{ and}$$
$$\tilde{D}^{+}_{ltt'} = \hat{D}^{+}_{ltt',1-\theta} \quad \forall l \in \mathcal{L}, \forall t' \in \mathcal{T}^F_t.$$

The remaining estimated parameters, i.e., those for generators $g \in \mathcal{G} \setminus \mathcal{G}^V$, are set equal to the nominal physical parameters of the generators (with adjustments for, e.g., planned maintenance outages). If the set of potential policies Π is limited to this form, the problem in Eq. (5) becomes a search for the best possible value of θ. In general, we expect system operators to use $0 < \theta < 0.5$, such that we have a pessimistic estimate of the availability of uncertain generators and assume load will be higher than its median value. Generalizations of this policy are immediately apparent: for example, operators may wish to choose different values of $\theta_{ltt'}$ and $\theta_{gtt'}$ for different loads and generators in the system or that change over time. For the purposes of the discussion, however, we restrict to a single value.

4.3.1 The inevitability of biasing

Implementing a "bias-free" model of this form is not in general a viable operating strategy. From a modeling standpoint we could have inserted expected values instead of quantile estimates for generator availability and demand. However, this policy would likely perform poorly in practice. The intuition is similar to that of the newsvendor problem, with results along those lines shown for two-stage examples in Zavala et al. (2017) and Zakeri et al. (2019). Since the value of lost load is typically two to three orders of magnitude greater than the cost of serving load, it is optimal to prepare for much higher than expected load. Without biasing, deterministic lookahead models like ($DLAC$) will tend to under-commit thermal units and over-dispatch storage in early time intervals.

In Example 1, suppose we solve a period 2 model that includes period 3, i.e., $F = 1$, and used the expected availability of wind in period 3, i.e., $\tilde{P}^{+}_{wind,2,3} = 10$. Anticipating sufficient wind in period 3 to meet all demand, this model would see no reason to avoid discharging storage to zero in period 2. If we instead use a value of $\tilde{P}^{+}_{wind,2,3} = 0.11$ (i.e., we use $\theta = 0.11/20$) in the ($DLAC - NLB - \theta$) policy), then the model will leave at least 0.89 MWh in storage, such that all demand is provisionally served in the lookahead model. By biasing the load in this way, system operators can induce the deterministic model to produce an optimal solution to the stochastic problem in this instance, replicating the period 2 dispatch shown in Table 1. To be sure, we should not expect such an approach to recover an optimal dispatch in all cases. Nevertheless, the wide use of net load biasing in practice suggests it can improve the quality of solutions.

4.3.2 Biasing and prices

While in this example operators can reproduce the optimal dispatch, they cannot in most cases reproduce the benchmark price. Table 2 shows the optimal bias term θ^* to use in each example in period 2 in addition to the price $\lambda_t^{NLB}(\theta^*)$ calculated by model $(DLAC - NLB - \theta)$ with that optimal θ.

Table 2: Dispatch and Prices in Period 2 with Net Load Biasing

	Wind Avail. $P^+_{wind,2}$	Optimal bias param. (θ^*)	Price $\lambda_2^{NLB}(\theta^*)$ (\$/MWh)
	$6 < P^+_{wind,2} \leq 10$		0
Ex. 1	$5.89 < P^+_{wind,2} \leq 6$	0.11/20	0
	$5 \leq P^+_{wind,2} \leq 5.89$		100
Ex. 2	10	0.202/20	0
Ex. 3	5	0.065/20	100

In Table 1, in the range $5.89 < P^+_{wind,2} < 10$ prices follow the opportunity cost of storage. In model $(DLAC - NLB - \theta)$, the price is \$0 in the same range, since the model believes that exactly 0.89 MWh and no more will be useful in useful in period 3. As such, the model calculates an implicit opportunity cost of \$0 for discharging storage down to that level. In other cases the policy can lead to a higher price. In Example 3, assuming wind availability of $\tilde{P}^+_{wind,2,3} = 0.065$ restores the optimal dispatch, but leads to a price of \$100/MWh, above the benchmark of \$67.58/MWh. Moreover, increasing the amount of bias in a future time period t' can affect prices in the binding period t in either direction. For example, if higher projected load in future intervals causes the model to charge or keep more energy in storage it may cause prices to rise, whereas if it causes ramping constraints to bind it may cause prices to fall.

Because a "neutral" approach is suboptimal and potentially even infeasible operationally, biasing is inevitable under this policy. The primary point we wish to highlight in this regard is that, as a direct consequence, operator choices about the degree of biasing necessarily impact price formation.

4.3.3 Multi-period pricing policies

In Table 1 we calculated prices in such a way that gas operates at a loss in period 2 in Examples 2 and 3. Losses occur due to the presence of a ramping constraint between periods 2 and 3 that sometimes binds under the optimal dispatch. Table 2 shows that this potential for loss persists with net load biasing. However, the same constraint leads to the potential for higher prices in period 3 under which the gas plant recoups its losses. Shortfalls occur $.202/20 = 1.01\%$ of the time, leading to a benchmark price of \$10,000/MWh. Since it produces 4.798 MW in such cases, expected profit for the plant in period 3 is \$479.80, well above its certain loss of \$79.80 in period 2. This result holds generally: in a stochastic competitive equilibrium assuming convexity, units have the potential for short-term losses but will always be profitable in expectation ((Kazempour et al., 2018; Mays, 2021b)). However, since outcomes in period 3 are non-binding, following the optimal dispatch leads to a situation where the plant takes a guaranteed loss without any guarantee it will recoup its losses in future intervals.

To address this potential issue, proposals have emerged for multi-interval pricing policies that would adjust prices to limit the potential for short-losses (e.g., by lifting the price in period 2 to $100/MWh). A key argument against these proposals is that they depart from the benchmark prices that support an efficient resource mix in the long run, in particular weakening incentives for flexibility (Mays, 2021a; Wang et al., 2023). A separate argument relevant to the examples is that the proposals have an unclear impact on opportunity costs and thus operational efficiency. In order to not result in excess charges to load, the higher price in period 2 must be accompanied by a lower price in period 3. Changing the distribution of prices that emerges in period 3, however, would lead to different offers from storage (and all resources with inter-temporal constraints) in period 2, and thus lead to different commitment and dispatch decisions.

4.3.4 Multi-period settlement

Another strategy to resolve the potential for loss in the first, binding interval is to make all intervals in the lookahead model financially binding Zhao et al. (2020). In Example 2, the gas plant's loss in period 2 would be accompanied by a financial contract for period 3 that guaranteed positive operating profit overall, regardless of the price that emerged in period 3. Model $(DLAC - NLB - \theta)$ gives a provisional price $\tilde{\lambda}_{23}$ = $200/MWh, since an additional unit of load in period 3 requires an additional unit of output from gas in both periods 2 and 3. Thus the deterministic model would have the gas plant sell forward 4.798 MWh in period 3 at $200/MWh. Beyond the complications it presents for settlement, a major challenge with such proposals is the same issue raised in Section 3.3. There is no mechanism in model $(DLAC - NLB - \theta)$ to ensure that $\tilde{\lambda}_{tt'} = \mathbb{E}_{W_{t+1},\ldots,W_{t'}}[\lambda_{t'}]$. Without such a mechanism, settling on future intervals could present arbitrage opportunities or compel market participants to take forward financial positions that did not accord with their beliefs about future prices.

4.4 Reserve tuning

A third strategy system operators use is to adjust the quantities of reserves $R^+_{tt'}$ procured in the model or the penalty cost V^R for a reserve shortfall. As in the previous strategy, tuning these quantities can enable better solutions to the stochastic problem while maintaining a deterministic modeling framework. Relative to net load biasing, however, there are significantly more degrees of freedom in this approach: operators can define operating reserve demand curves with multiple steps and can even define new reserve products. For example, separate reserve products may be defined for upwards and downwards adjustments, whereas operators have to choose a single direction (typically upwards) when biasing load. The different implementations can also have different pricing impacts. For example, biasing load upwards in period $t' = t + 1$ could lead to a binding ramping constraint and a lower price in period t, whereas introducing a ramping product in period t instead increases demand across all products and could therefore lead to a higher price in period t.

With these additional degrees of freedom, there is a larger space of potential policies that may be pursued under the umbrella of reserve tuning. We note the difference between reserve products addressing uncertainty within a pricing interval and those addressing uncertainty across pricing intervals, with the modeling framework in this paper focused on the latter. Appendix B provides a reserve formulation that holds extra capacity in future intervals, with effects similar to that of net load biasing, as well as a formulation that specifies a value of residual energy at the end of the lookahead horizon, conceptually similar to opportunity cost bidding. Here we define reserves in period t that reflect uncertainty only in period $t+1$. To define an upward reserve tuning $(DLAC - RT - \theta)$ policy, we define the random variable $M_{tt'} = \sum_{l \in \mathcal{L}} D^+_{ltt'} - \sum_{g \in \mathcal{G}^V} P^+_{gtt'}$ representing load net of variable generation at the system level in time t' and make the substitution

$$R^+_{tt'} = \overline{R}^+_{t'} + \max\{0, \hat{M}_{t,t'+1,1-\theta} - \overline{M}_{t,t'+1}\} \quad \forall t' \in \{t, ..., t + F - 1\},$$

where $\overline{R}^+_{t'}$ represents a minimum quantity of reserves that must be procured in period t' (e.g., due to other sources of uncertainty). As before, $\hat{M}_{t,t'+1,1-\theta}$ denotes an estimate of the $1 - \theta$ quantile of net load and $\overline{M}_{t,t'+1}$ denotes its expected value. Expected values are used for the parameters $\tilde{P}^+_{gtt'}$ and $\tilde{D}^+_{ltt'}$, and operators are presumed to choose a value of θ such that there is an upwards shift in the reserves procured. This specification is accompanied by a constraint restricting the supply of reserves from individual resources based on their physical capabilities. For example, specifying $\tilde{r}_{gtt'} \leq UR^+_g$ would limit the reserves supplied by a generator to its upward ramping capability.

We note several differences between the $(DLAC - NLB - \theta)$ and $(DLAC - RT - \theta)$ policies. First, while we assumed in Section 4.3 that net load biasing would not result in any binding forward financial positions, the $(DLAC - RT - \theta)$ policy directly affects the quantity of reserves procured in the binding interval. Second, as defined here the reserve adjustment is made at a system level, summing over all generators and loads rather than handling each individually. Reserve products are typically defined zonally or system-wide, while energy must be balance at a nodal level; while we have not written network constraints here, this difference can affect congestion and losses. Third, reserve shortfall penalties determined by V^R will typically be below the value of load represented by V^L_l, thus generating different prices in the event of a shortfall. Fourth, the set of generators qualified to provide reserves is typically smaller than the set able to provide energy.

4.4.1 Calibration and dispatch

As was the case with net load biasing, the problem is then to determine which value of θ to use. Beyond this narrowly defined family, the problem could be expanded to test the performance of different reserve product definitions. Since this paper's focus is price formation rather than identifying optimal operating policies, we consider the simpler question of identifying specifications that would enable a single-period version of the $(DLAC - RT - \theta)$ policy executed in period 2 to result in dispatch solutions that are optimal for the underlying stochastic problems in the examples. In Example 1, the challenge is to provoke the model to dispatch gas instead of discharging storage below 0.89 MWh. This can be accomplished by relaxing constraint (6g) to allow the gas plant to sell reserves up to its full capacity of 9 MW despite also selling energy in the period. Then, setting a reserve requirement $R^+_{2,2} = 9.89$ will lead to storage selling 0.89 MW of reserves and gas selling 9 MW. In Example 2, the system operator can set a requirement of $R^+_{2,2} = 9.798$, allow the gas plant to sell up to $p_{gas,2} + UR^+_{gas}$, and disallow the wind plant from selling reserves. Then storage will sell 5 MW and gas 4.798 MW to satisfy the demand for reserves. Using the same limit in Example 3 while setting a requirement $R^+_{2,2} = 9.934$ leads to storage selling 2.967 MW and gas selling 6.967 MW. In all of the examples, the reserve requirements correspond to the same optimal bias parameters θ determined for the net load biasing policy.

4.4.2 Prices, products, and revenues

As with the $(DLAC - NLB - \theta)$ policy, the $(DLAC - RT - \theta)$ policy tuned in this way can identify an optimal dispatch on each of the example problems. However, the two policies have different effects on prices and revenues. Table 3 shows the cleared reserve quantities as well as prices for energy reserves under the $(DLAC - RT - \theta)$ assuming the same optimal θ values as in Table 2.

Table 3: Prices in Period 2 with Reserve Tuning

	Wind Avail. $P^+_{wind,2}$	Gas Reserve $r^*_{gas,2}$ (MW)	Storage Reserve r^*_{b2}	Energy Price $\lambda^{RT}_2(\theta^*)$ (\$/MWh)	Reserve Price $\mu^{RT}_2(\theta^*)$ (\$/MWh)
	$6 < P^+_{wind,2} \leq 10$			0	0
Ex. 1	$5.89 < P^+_{wind,2} \leq 6$	9.00	0.89	0	0
	$5 \leq P^+_{wind,2} \leq 5.89$			100	100
Ex. 2	10	4.798	5.00	0	100
Ex. 3	5	6.967	2.967	50	50

In Example 1, whenever $P^+_{wind,2} > 5.89$ storage is not fully allocated across energy and reserves, so an increase in demand for either does not lead to increased cost. In Example 2, an increase in demand for energy is met by curtailed wind, while an increase in demand for reserves must be met by gas, entailing that it is dispatched to a higher level at cost of \$100/MWh. In Example 3, an increase in demand for either product is met half by storage and half by gas, leading to prices of \$50/MWh for both energy and reserves. As such, the resulting energy prices depart from both the idealized benchmark in Table 1 and the net load biasing policy in Table 2. An additional difference is the introduction of a new product for reserves with a price of its own. In some instances, the new product can significantly change the compensation going to resources in the system: in Example 2, whereas the benchmark price $\lambda_2^* = \$0$/MWh, the $(DLAC - RT - \theta)$ policy gives rise to a payment of \$479.80 to the gas plant and \$500.00 to the storage resource.

4.4.3 Challenges in defining reserve products

In Section 2 we identified benchmark prices based on their ability to support a socially optimal resource mix in the long run. In this benchmark we did not need to define any reserve products aimed at addressing uncertainty across pricing intervals, since the need for additional flexibility is calculated endogenously by the stochastic model. In this vein, it is worth highlighting the different economic interpretations that may be attached to payments for reserves. In some cases (e.g., black start and reactive power), payments are for services that are explicitly omitted from real-time auctions. For active power services deployed on timescales faster than market clearing (e.g., fast frequency response and regulation), payments can reflect unpriced fluctuations in marginal cost that occur within the pricing interval. By contrast, for services deployed on timescales longer than market clearing, reserve products reflect algorithmic choices about how to commit and dispatch the power system in ways that cannot replicate the stochastic ideal. A separate possible economic justification for reserves relates to externalities due to the possibility of a system collapse (Joskow and Tirole, 2007). However, the general assumption is that the possibility of a cascading failure is unaffected on the margin by the quantity of this slower class of reserves available.

Different algorithmic choices about reserves will result in different prices, different revenue for market participants, and therefore a different long-run competitive equilibrium. The reserve products currently in use in U.S. markets arose from engineering practice prior to the introduction of competitive markets and were initially implemented in a context when settlements occurred on an hourly basis, longer than the ten- to thirty-minute timescale on which contingency reserves are defined. Since settlement intervals have since been reduced to five minutes, the economic justification for these reserve products must be reestablished. The examples show that after the introduction of reserves, total payments to generators can be either higher or lower in individual instances. To avoid an overall transfer from loads to generators relative to the benchmark without reserve products, however, it must be the case that higher payments for intertemporal reserve products lead to lower payments for energy in the long run.

One challenge in this context is the allocation of cost associated with the reserve product. Whereas energy purchases can straightforwardly be charged to the entity that consumed the energy, the presence of reserve payments necessitates a methodology to recover costs from market participants. If the additional cost in Example 2 were allocated to the load, as is the default in U.S. markets, it could justifiably ask why it should pay for the service instead of the wind plant that is the source of uncertainty. The load could similarly ask why the gas plant should be rewarded for its inflexibility, which exacerbates the effect of the uncertainty. Given potential disagreements along these lines, the need for a cost allocation methodology may be seen as a disadvantage of reserve products relative to the more straightforward outcomes that would result with more reliance on energy markets.

A second challenge pertains to consistency in product definitions across markets and across time. Whereas the benchmark with endogenous reserves is formulated in a general way and can apply to all systems, reserve products are defined and parameterized in the context of Eq. (5). The policy identified by system operators for their sequential decision problem can depend on the resource mix and other characteristics of the particular system. This calibration implies that if they are near-optimal, product definitions will not be consistent across systems and across time. This lack of consistency presents challenges for both the design and analysis of electricity markets. In terms of design, one problem is that introducing a new product is a several-year effort involving stakeholder processes, regulatory reviews, and software development, raising the possibility that the product actually implemented will no longer reflect the needs of the system. Additionally, a lack of consistency implies less ability to transfer knowledge across different markets and less ability to hedge exposure to future reserve prices. In terms of analysis, it implies that projections of future market outcomes should in principle include assumptions on the reserve market changes that could accompany shifts in the resource mix.

4.5 Stochastic programming

Given the drawbacks in the three policies discussed above, here we describe how to compute prices in a stochastic program that may be seen as a better approximation of the benchmark in Section 2. Among the many potential alternative policies that might be employed to better handle uncertainty, we choose a stochastic program because it is a natural extension from current practice, with prices in the binding interval calculated in the same way they that occurs in a deterministic program. In the period t model we solve a stochastic program using scenarios $\omega \in \tilde{\Omega}_t$ generated based on the state S_t, with scenario ω assigned probability $\tilde{\rho}_\omega$. In order to write the stochastic program in a concise way, we add the subscript ω to all the primal and dual decision variables as well as estimated parameters from model $(DLAC)$ to create scenario-specific copies. Let $\tilde{x}_{t\omega} = (\tilde{u}_{t\omega}, \tilde{v}_{t\omega}, \tilde{q}_{t\omega}, \tilde{p}_{t\omega}, \tilde{r}_{t\omega}, \tilde{k}^c_{t\omega}, \tilde{k}^d_{t\omega}, \tilde{j}_{t\omega}, \tilde{d}_{t\omega}, \tilde{z}_{t\omega})$ represent a full set of scenario-specific provisional decisions made at time t for scenario ω and let $\mathcal{X}_{t\omega}$ represent the feasible region of model $(DLAC)$ for period t with the scenario-specific values of the estimated parameters inserted and all constraints multiplied by $\tilde{\rho}_\omega$. With $\tilde{x}_t$ enlarged to include all the scenarios, we can write a two-stage stochastic lookahead unit commitment ($SLAC$) policy for the period t problem as follows:

$$
\begin{aligned}
&(SLAC) \\
&X_t^{SLAC}(S_t|H,F) \in \\
&\arg\max_{x_t,\tilde{x}_t} \quad \sum_{\omega\in\tilde{\Omega}_t}\sum_{l\in\mathcal{L}}\sum_{t'\in\mathcal{T}_t^{H,F}} \tilde{\rho}_\omega V_l^L \tilde{d}_{ltt'\omega} + \sum_{\omega\in\tilde{\Omega}_t}\sum_{t'\in\mathcal{T}_t^{H,F}} \tilde{\rho}_\omega V^R \tilde{z}_{tt'\omega} \\
&\qquad - \sum_{\omega\in\tilde{\Omega}_t}\sum_{g\in\mathcal{G}}\sum_{t'\in\mathcal{T}_t^{H,F}} \tilde{\rho}_\omega C_g^{OP} \tilde{p}_{gtt'\omega} \\
&\qquad - \sum_{\omega\in\tilde{\Omega}_t}\sum_{g\in\mathcal{G}}\sum_{t'\in\mathcal{T}_t^{H,F}} \tilde{\rho}_\omega \left(C_g^{NL}\tilde{u}_{gtt'\omega} + C_g^{SU}\tilde{v}_{gtt'\omega}\right) && \text{(9a)} \\
&\text{subject to} \quad \tilde{x}_{t\omega} \in \mathcal{X}_{t\omega} \qquad \forall \omega \in \tilde{\Omega}_t && \text{(9b)} \\
&\qquad \tilde{x}_{t\omega} - x_t = 0 \qquad \forall \omega \in \tilde{\Omega}_t. && \text{(9c)}
\end{aligned}
$$

Structurally, the first stage in this model corresponds to decisions in the binding interval t while the second stage covers the entire lookahead period $\mathcal{T}_t^F$. While we defined multiple copies of the decision variables for the binding interval t for notational convenience, the nonanticipativity constraint in Eq. (9c) guarantees that the model returns a single decision in this period. It should also be noted that while the model returns scenario-specific prices $\tilde{\lambda}_{tt\omega}$ for the binding interval, we can set $\lambda_t^{SLAC} = \sum_{\omega\in\tilde{\Omega}_t} \tilde{\rho}_\omega \tilde{\lambda}_{tt\omega}$. This sum can be shown to be an optimal dual value in an equivalent formulation of the stochastic program that used a single copy of the period t decisions. We note that model ($SLAC$) is a simplification of the full scenario tree in that we have a smaller set of scenarios, the decision horizon is truncated, and all nonanticipativity constraints that would otherwise appear for $t' \in \mathcal{T}_t^F$ have been relaxed. In period 2 of the toy examples, the second and third of these simplifications are not relevant, making model ($SLAC$) enabling the model to converge on the idealized benchmark of model (SOC) as the number of included scenarios grows.

4.5.1 Opportunity costs and prices

The motivation behind moving from a deterministic variants of the lookahead unit commitment to the stochastic variant is to generate a better approximation λ_t^{SLAC} to the benchmark price λ_n^* that would be produced at the corresponding node of the full scenario tree in model (SOC). Unlike in Section 4.2, however, the opportunity is calculated endogenously by the centralized model rather than being submitted by decentralized market participants. Because of these simplifications from the full scenario tree, we cannot make guarantees regarding the similarity of the prices. In the toy examples, however, we are able both to recover the optimal dispatch and form the benchmark prices.

In creating these estimates, the model provides an approximation of the marginal opportunity costs of storage resources. Defining dual variable $\tilde{\nu}^{K+}_{b\omega}$ for the constraints within Eq. (9c) that $\tilde{k}^d_{btt\omega} - k^d_{bt} = 0 \ \forall b \in \mathcal{B},\ \forall \omega \in \tilde{\Omega}_t$, we can write the following dual constraint corresponding to the primal variable copies for discharge from a storage resource b in the binding interval, i.e., $\tilde{k}^d_{btt\omega}$:

$$\tilde{\rho}_\omega(-\tilde{\lambda}_{tt\omega} + \tilde{\alpha}^{K+}_{btt\omega} + \frac{1}{\eta^d_b}\tilde{\beta}_{btt\omega} + \tilde{\nu}^{K+}_{b\omega}) \geq 0. \tag{10}$$

We can also write the dual constraint corresponding to the primal variable k^d_{bt} as follows:

$$-\sum_{\omega \in \tilde{\Omega}_t} \tilde{\rho}_\omega \tilde{\nu}^{K+}_{b\omega} = 0. \tag{11}$$

Suppose in some period t that a storage resource b is discharging above 0 but below its maximum ability, i.e., $0 < k^d_{bt} < K_b$. In other words, it is the marginal resource on the system. Due to the nonanticipativity constraint, these inequalities hold for each of the scenario copies $\tilde{k}^d_{btt\omega}$. By complementary slackness we know that each of the constraints in Eq. (10) holds with equality and that $\tilde{\alpha}^{K+}_{btt\omega} = 0 \ \forall \omega \in \tilde{\Omega}_t$. Summing Eq. (10) over all scenarios thus gives

$$\sum_{\omega \in \tilde{\Omega}_t} \tilde{\rho}_\omega(-\tilde{\lambda}_{tt\omega} + \frac{1}{\eta^d_b}\tilde{\beta}_{btt\omega} + \tilde{\nu}^{K+}_{b\omega}) = 0. \tag{12}$$

Substituting Eq. (11) and using the previously defined $\lambda^{SLAC}_t = \sum_{\omega \in \tilde{\Omega}_t} \tilde{\rho}_\omega \tilde{\lambda}_{tt\omega}$ gives

$$\lambda^{SLAC}_t = \sum_{\omega \in \tilde{\Omega}_t} \tilde{\rho}_\omega \frac{1}{\eta^d_b}\tilde{\beta}_{btt\omega}. \tag{13}$$

The dual values $\tilde{\beta}_{btt\omega}$ correspond to the state of charge consistency constraints in the scenario copies of constraint (6j) from model $(DLAC)$. Accordingly, these dual variables can be interpreted as the value of entering period $t+1$ with an additional unit of stored energy. The price λ^{SLAC}_t generated when a storage resource is on the margin is thus a weighted average opportunity cost adjusted by discharge efficiency η^d_b.

4.5.2 Market design vs. market analysis revisited

Section 4.2.3 raised the distinction between market design and market analysis, arguing that the opportunity cost policy $(DLAC - OC)$ could be appropriate for design but presents difficulties for analysis due to the lack of direct access to participant opportunity costs. With an endogenous estimate of opportunity costs across potential future scenarios, the $(SLAC)$ policy may be superior to the $(DLAC - NLB - \theta)$ and $(DLAC - RT - \theta)$ policies in this regard. As an additional advantage, model $(SLAC)$ creates many scenario-specific estimates of the price in future intervals. Whereas deterministic models struggle to enforce the no-arbitrage rule in Eq. (7), an estimate given by $\hat{\lambda}_{tt'}^{SLAC} = \sum_{\omega \in \tilde{\Omega}_t} \tilde{\rho}_\omega \tilde{\lambda}_{tt'\omega}^{SLAC}$ may lead to lower error.

From a design perspective, the question is whether using model $(SLAC)$ would be appropriate in actual market clearing. One objection raised in the use of stochastic models in the generation of forward prices is that market participants will not agree on the scenarios or other assumptions that act as inputs into the market clearing model. Model $(SLAC)$ avoids this criticism in that the only the spot price it produces is binding, just as in present markets. While the price in the binding interval does depend on choices regarding the scenarios, prices under net load biasing and reserve tuning are perhaps even more influenced by operator choices.

In the toy examples, either model $(DLAC - OC)$ or model $(SLAC)$ could in principle replicate the benchmark prices of model (SOC). In practice, the comparative performance of the two policies may be dependent on the market structure and system characteristics. In competitive systems with ample flexibility, the $(DLAC - OC)$ policy could offer a simple way to get high-quality operational decisions and prices. In systems where market power mitigation or a lack of flexibility present significant challenges, the more centralized approach of model $(SLAC)$ may lead to better performance.

5 Evaluating price formation policies

To demonstrate the effects that competing operating policies may have for price formation, this section describes simulations using the $(DLAC-NLB-\theta)$, $(DLAC-RT-\theta)$, and $(SLAC)$ policies on an ERCOT-like test system. As previously mentioned, simulating the $(DLAC-OC)$ policy would require an additional step to estimate opportunity costs and construct participant offer curves. The primary intent with these tests is not to calibrate and compare optimal versions of each policy, but to validate the observations in Section 4 regarding the impact of model choice and parameterization on price outcomes. Thus, the results include tests at multiple levels of θ reflecting, in an informal way, different levels of risk aversion among operators.

Thermal generator costs and technical constraints are taken from the FERC unit commitment test system Krall et al. (2012), with a subset of 456 generators selected to approximate the capacity mix in ERCOT. The simulations assume 5,000 MW of 4-hr batteries with charge and discharge efficiencies $\eta_b^c = \eta_b^d = 0.93$, for a round-trip efficiency of 86%. Scenarios for wind, solar, and load are generated based on forecasted marginal distributions for each hour described in Bryce et al. (2023) and covering the 48-hour period of July 18–19, 2018, selected due to price spikes that occurred on those days. Scenarios are constructed separately for wind, solar, and load by sampling from a standard uniform distribution in time $t = 1$, then in each subsequent hour sampling from a triangular distribution centered on the current observation and extending

0.1 in either direction, truncated (if applicable) to 0.0 and 1.0. These series are then mapped to the marginal distributions from Bryce et al. (2023). Without performing a full statistical analysis (cf. Carmona and Yang (2022)), the intention is to generate plausible sample paths given strong intertemporal correlation. The experiments cover 10 sample paths of wind, solar, and load, with no outages among the batteries or thermal plants, with each policy variant tested on the same set of realizations.

For all of the tested policies and in all time periods $t \in \mathcal{T} = \{1, \ldots, 48\}$, the entire set of time periods is included in the model, i.e., $\mathcal{T}_t^{H,F} = \mathcal{T}$. For the $(DLAC - NLB - \theta)$ and $(DLAC - RT - \theta)$ policies, 100 scenarios for future intervals $t' \in \mathcal{T}_t^F$ are generated at each time step t using the sampling procedure described above; that is, scenarios are sampled from the correct distribution. The relevant quantile of each random variable is estimated by ranking the 100 outcomes and using the value in position $100 \cdot \theta$. In the $(SLAC)$ policy, each instance uses 20 scenarios generated in the same way. Results are reported for the entire 48-hour period, with no hours discarded; by observation, any end-of-horizon effects are limited since results are driven by high-demand hours in the afternoon of each day. AMPL code and data for the simulations will be available in an online appendix.

It can be understood that, while the experiments do seek some realism in pricing outcomes through the incorporation of a large set of generators with detailed cost information, the tests cover a relatively small number of options from the universe of potential operating policies. Additionally, all of the experiments relax binary variables, resulting in linear programs for which prices can be straightforwardly calculated. As described in Section 3.4.5 the analysis has significant implications for interpretation of policies adopted to address non-convexity; numerical investigation of this topic is left for future work. The broader point of the examples is to establish the need for the proposed framework in evaluating price formation: even with the limited scope of tests included here, the models produce a wide range of price outcomes.

5.1 Operational performance

Table 4 reports the total production cost incurred when operating the system under each of the tested policies, as well as the cost relative to the $(SLAC)$ policy. As an additional comparison, the table reports results using a policy $(DLAC - AVG)$ that uses the expected value of demand and renewable availability in all future time periods. A mild surprise is that the best performance among the tested policies is achieved through a small amount of net load biasing under $(DLAC - NLB - 45)$, outperforming $(SLAC)$ by 0.3%. The surprise is only mild, however; enhancements to $(SLAC)$ (e.g., with more careful scenario selection, as in Feng et al. (2015)) could improve its performance. While each of the policies could be further refined, overall cost is within a relatively small window, ranging from 99.7 to 103.1% of the total calculated under $(SLAC)$. Moreover, perfect foresight models for the ten instances compute a cost of \$803.2M, indicating relatively little room for improvement beyond these simple policies. The relative ease with which high-quality policies can be identified may relate to the arbitrarily chosen addition of 5 GW of highly flexible battery storage to the system.

More important for purposes of this paper are the pricing outcomes. As with the toy examples in Section 4, prices vary much more widely than cost. Table 4 also shows the total charges for energy and reserves to buyers of electricity under each of the policies. Taking the prices from $(SLAC)$ to be the best available approximation to the idealized benchmark of (SOC) and noting that the peak summer days included in the simulation were intentionally chosen to include the potential for extreme price spikes, these charges range from \$3,556M to \$5,040M, from 84.4 to 119.6% of those calculated under $(SLAC)$.

Table 4: Outcome Comparison for Policy Variants

Policy	Cost (\$M)	Relative Cost (%)	Total Charges (\$M)	Prediction Bias (\$/MWh)
$(DLAC-NLB-50)$	833.5	103.1	4,076	-53.3
$(DLAC-NLB-45)$	806.1	99.7	4,099	-51.8
$(DLAC-NLB-40)$	806.3	99.7	3,993	-45.1
$(DLAC-NLB-35)$	806.6	99.7	3,633	-30.4
$(DLAC-NLB-30)$	818.8	101.2	3,556	-20.7
$(DLAC-NLB-25)$	813.5	100.6	4,229	-21.5
$(DLAC-NLB-20)$	825.9	102.1	4,989	-10.2
$(DLAC-NLB-15)$	814.3	100.7	4,325	75.8
$(DLAC-NLB-10)$	815.9	100.9	5,040	214.9
$(DLAC-NLB-5)$	819.9	101.4	5,034	406.5
$(DLAC-RT-60)$	816.3	100.9	4,071	-55.4
$(DLAC-RT-55)$	822.7	101.7	4,072	-55.4
$(DLAC-RT-50)$	829.4	102.5	4,061	-53.5
$(DLAC-RT-45)$	831.5	102.8	4,079	-53.7
$(DLAC-RT-40)$	832.2	102.9	4,092	-52.0
$(DLAC-AVG)$	812.9	100.9	4,053	-54.2
$(SLAC)$	808.8	100.0	4,214	-11.4

5.2 Price consistency

Along with the cost and price outcomes, Table 4 shows a metric for the bias in forward price predictions calculated under each pricing policy. This metric is calculated for each iteration of the simulation as an average of the average difference between provisional prices and the actual realized price in each time period for which any predictions are made, i.e.,

$$\frac{1}{47} \sum_{t' \in \{2,\ldots,48\}} \sum_{t \in \{1,\ldots,t'-1\}} (\tilde{\lambda}_{tt'} - \lambda_{t'})/(t'-1), \tag{14}$$

and then averaged across the ten iterations. Except in the policies with high net load biasing, provisional prices are on average below the prices actually observed. As might be expected (cf. Zavala et al. (2017)), the ($SLAC$) policy performs relatively well in terms of price consistency. However, perhaps related to the naive sampling procedure (cf. Mak et al. (1999)) and the simplification to a two-stage program (cf. Brigatto et al. (2017)), a negative bias remains.

As previously mentioned in Sections 3.3 and 3.4.2, a lack of price consistency presents a challenge for analysis. In particular, while it is common to base economic analysis on "day-ahead" prices calculated from lookahead models covering several intervals, a safer approach is to report simulation results (e.g., total revenues) using only spot prices.

5.3 Long-run consequences

Prices in electricity markets are intended to support efficiency not only in short-run operations, but also in long-run investment. It is thus important to consider the implications of different pricing policies on the total remuneration for different types of resources. One challenge in interpreting the numerical results is that we do not have access to the prices that would be computed by model (SOC), which are optimal in the sense of supporting an efficient long-run resource mix. Despite the fact that the ($SLAC$) policy as implemented is outperformed in short-run operations, its resemblance to the idealized stochastic program makes it a defensible point of reference for the distribution of revenues across different resource types. As previously discussed in Table 4, the choice of operating policy can have a significant impact on the total revenue across all resources. Table 5 expands on this total, showing the revenue available to each resource type in the system under each policy. In general, resource-specific revenues track the total charges. Storage is a notable exception, with revenues deviating significantly from that seen by the market as a whole (e.g., 68.5% vs. 84.4% under ($DLAC - NLB - 30$) and 116.4% vs. 97.1% under ($DLAC - RT - 40$). While the numerical examples do not precisely replicate the rules of ERCOT or any other specific market, they indicate the potential for substantial misallocation of investment due to inefficiencies in spot price formation.

In energy-only markets, such as ERCOT, the total revenue available dictates the level of investment and therefore directly affects reliability outcomes. In systems with supplementary capacity payments, changes in the revenue available through sale of energy and reserves are in principle balanced by a compensatory shift in the revenue available through sale of capacity. However, since capacity payments are typically intended to be non-discriminatory, they would be unlikely to affect distributional shifts embedded in the pricing of energy and reserves. In principle, additional capacity-like products (e.g., the flexible capacity requirement in CAISO) could restore near-efficient distributional outcomes in cases where price formation policies for energy and reserves result in under-compensation of particular resource types.

Table 5: Total Revenues Relative to (SLAC) Policy (%)

Policy	Steam	Combined	Gas	Wind	Solar	Storage	All
$(DLAC-NLB-50)$	96.6	96.4	96.1	99.1	96.7	81.7	96.7
$(DLAC-NLB-45)$	97.1	96.9	96.7	99.2	97.4	102.1	97.3
$(DLAC-NLB-40)$	94.7	94.5	93.9	98.2	94.8	97.1	94.7
$(DLAC-NLB-35)$	86.6	86.0	84.7	94.6	85.2	80.7	86.2
$(DLAC-NLB-30)$	84.9	84.2	82.7	93.7	83.3	68.5	84.4
$(DLAC-NLB-25)$	100.5	100.4	100.4	100.0	101.9	95.3	100.4
$(DLAC-NLB-20)$	118.4	118.7	120.3	104.9	120.7	114.0	118.4
$(DLAC-NLB-15)$	103.1	102.9	103.1	98.0	101.9	98.3	102.6
$(DLAC-NLB-10)$	120.1	120.3	122.1	102.8	119.9	115.3	119.6
$(DLAC-NLB-5)$	120.0	120.2	122.1	102.2	119.8	114.8	119.5
$(DLAC-RT-60)$	95.7	95.5	95.1	99.1	96.0	114.4	96.6
$(DLAC-RT-55)$	95.6	95.3	94.9	99.0	95.8	115.3	96.6
$(DLAC-RT-50)$	95.1	94.8	94.5	98.8	95.3	114.7	96.4
$(DLAC-RT-45)$	95.5	95.2	94.8	99.0	95.7	115.5	96.8
$(DLAC-RT-40)$	95.8	95.4	95.1	99.0	95.9	116.4	97.1

Note: Thermal generators in the data set are aggregated based on prime mover rather than fuel

6 Future directions

The goal of this paper is to promote a shift in the discussion of price formation in wholesale electricity markets from a static to a dynamic modeling framework. While the design and analysis of systems with significant reservoir hydropower have long relied on dynamic models, most other systems have come to rely on simpler static models that have nevertheless been useful in contexts with limited variability, uncertainty, and intertemporal constraints. The entry of large quantities of renewables and battery storage has increased the salience of all of these factors, necessitating a richer modeling framework.

The shift has important consequences for several lines of research in electricity markets. The paper demonstrates that many price formation methods proposed in the literature to address non-convexity or intertemporal constraints are not fully specified in a dynamic context. With the addition of storage, a particular challenge with enhanced price formation proposals is to ensure that the opportunity costs arising from future price distributions are consistent with efficient operational decisions in the present. The paper also invites some skepticism regarding the implementation of new reserve products intending to address variability, uncertainty, and intertemporal constraints. While economic analysis of electricity markets sometimes treats reserves as a strict requirement, this paper describes them as reflections of the algorithmic choices made by system operators. Additionally, the paper highlights that standard production cost simulations will result in questionable pricing outcomes in cases with significant uncertainty. In numerical tests, provisional forward prices calculated in deterministic lookahead models tended to underestimate the actual prices generated. Perhaps most importantly, the paper highlights that algorithmic choices can have a significant impact on the total revenue available to resources in the market. A particular concern in this regard is the degree to which current practices can efficiently incentivize investment in storage and other resources providing flexibility. In recent years, the trend in U.S. markets has been toward introduction of new reserve products corresponding to new constraints in the deterministic models used by system operators. The effect of these products on long-term investment is not well understood. Several challenges attend the introduction and calibration of reserve products, including cost allocation, determining which resources qualify to provide the service, and ensuring that the product definition keeps pace with the needs of a changing resource mix.

While the simulations conducted in this paper intend to demonstrate the value of adopting a dynamic framework, they leave many avenues for future exploration. First, the simulations cover only four policy families out of the universe of potential options, and furthermore do not exhaust even those four. Second, while the paper uses prices from a stochastic program as a point of reference because of its superficial resemblance to the idealized benchmark prices that would arise in a model with a fully specified scenario tree, further examination would be required to firmly establish the relationship. Third, since different algorithmic choices lead to differences not only in expected prices but also in price volatility, future work could investigate the implications of different policies with risk averse market participants and incomplete risk trading. Fourth, the simulations limit participation to a set of well-defined resources on the bulk power system, but extensions could explore the consequences of alternative price formation and reserve product configurations for distribution-level and other non-traditional resources. Fifth, while the results highlight the potential for choices regarding spot price formation to lead to suboptimal investment incentives, more detailed simulations would be required to assess the degree to which current rules in any given market lead to addressable inefficiencies.

Acknowledgments

The author thanks Brent Eldridge, Jennie Jorgenson, Fredrich Kahrl, Bernard Knueven, Andrew Mills, and David Morton for many helpful comments and suggestions. This work was funded by the U.S. Department of Energy's Solar Energy Technologies Office, under Contract No. DE-AC02-05CH11231. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

References

Aasgård, E. K., S.-E. Fleten, M. Kaut, K. Midthun, and G. A. Perez-Valdes (2019, Aug). Hydropower bidding in a multi-market setting. *Energy Systems 10*(3), 543–565.

Biggar, D. R. and M. R. Hesamzadeh (2022). Do we need to implement multi-interval real-time markets? *The Energy Journal 43*(2), 111–131.

Brigatto, A., A. Street, and D. M. Valladão (2017). Assessing the cost of time-inconsistent operation policies in hydrothermal power systems. *IEEE Transactions on Power Systems 32*(6), 4541–4550.

Bryce, R., G. Buster, K. Doubleday, C. Feng, R. Ring-Jarvi, M. Rossol, F. Zhang, and B.-M. Hodge (2023, May). Solar PV, wind generation, and load forecasting dataset for ERCOT 2018: Performance-based energy resource feedback, optimization, and risk management (P.E.R.F.O.R.M.). Available at `https://www.nrel.gov/docs/fy23osti/79498.pdf`.

Byers, C. and G. Hug (2022). Long-run optimal pricing in electricity markets with non-convex costs. *European Journal of Operational Research*.

Carmona, R. and X. Yang (2022). Joint stochastic model for electric load, solar and wind power at asset level and monte carlo scenario generation.

Chao, H. (2019). Incentives for efficient pricing mechanism in markets with non-convexities. *Journal of Regulatory Economics 56*(1), 33–58.

Chen, C., Y. Guo, and L. Tong (2021). Pricing multi-interval dispatch under uncertainty part ii: Generalization and performance. *IEEE Transactions on Power Systems 36*(5), 3878–3886.

Cho, J. and A. Papavasiliou (2023). Pricing under uncertainty in multi-interval real-time markets. *Operations Research 0*(0), 0–0.

Cory-Wright, R., A. Philpott, and G. Zakeri (2018). Payment mechanisms for electricity markets with uncertain supply. *Operations Research Letters 46*(1), 116–121.

Cramton, P., A. Ockenfels, and S. Stoft (2013). Capacity market fundamentals. *Economics of Energy & Environmental Policy 2*(2), 27–46.

Duggan, J. E. and R. Sioshansi (2019). Another step towards equilibrium offers in unit commitment auctions with nonconvex costs: Multi-firm oligopolies. *The Energy Journal 40*(6), 249–282.

Eldridge, B., B. Knueven, and J. Mays (2023a). Rethinking the price formation problem–part 1: Participant incentives under uncertainty. *IEEE Transactions on Energy Markets, Policy and Regulation*, 1–10.

Eldridge, B., B. Knueven, and J. Mays (2023b). Rethinking the price formation problem–part 2: Rewarding flexibility and managing price risk. *IEEE Transactions on Energy Markets, Policy and Regulation*, 1–9.

Electric Power Research Institute (2016). Wholesale electricity market design initiatives in the United States: Survey and research needs. 3002009273.

Federal Energy Regulatory Commission (2019, April). Docket no. EL18-34-000. Available at `https://www.ferc.gov/whats-new/comm-meet/2019/041819/E-3.pdf`.

Feng, Y., I. Rios, S. Ryan, K. Spürkel, J. Watson, R. Wets, and D. Woodruff (2015). Toward scalable stochastic unit commitment. part 1: load scenario generation. *Energy Systems 6*(3), 309–329.

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

Frazier, A. W., W. Cole, P. Denholm, S. Machen, N. Gates, and N. Blair. (2021, May). Storage futures study: Economic potential of diurnal storage in the U.S. power sector. Available at `https://www.nrel.gov/docs/fy21osti/77449.pdf`.

Frew, B., G. Brinkman, P. Denholm, V. Narwade, G. Stephen, A. Bloom, and J. Lau (2021). Impact of operating reserve rules on electricity prices with high penetrations of renewable energy. *Energy Policy 156*, 112443.

Gérard, H., V. Leclère, and A. Philpott (2018). On risk averse competitive equilibrium. *Operations Research Letters 46*(1), 19–26.

Gribik, P., W. Hogan, and S. Pope (2007, December). Market-clearing electricity prices and energy uplift. Available at `https://www.hks.harvard.edu/fs/whogan/Gribik_Hogan_Pope_Price_Uplift_123107.pdf`.

Guo, C., M. Bodur, and D. J. Papageorgiou (2022). Generation expansion planning with revenue adequacy constraints. *Computers & Operations Research 142*, 105736.

Guo, Y., C. Chen, and L. Tong (2021). Pricing multi-interval dispatch under uncertainty part i: Dispatch-following incentives. *IEEE Transactions on Power Systems 36*(5), 3865–3877.

Herrero, I., P. Rodilla, and C. Batlle (2015). Electricity market-clearing prices and investment incentives: The role of pricing rules. *Energy Economics 47*, 42–51.

Hogan, W. W. (2022). Electricity market design and zero-marginal cost generation. *Current Sustainable/Renewable Energy Reports 9*(1), 15–26.

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

Jenkins, J. D., M. Luke, and S. Thernstrom (2018). Getting to zero carbon emissions in the electric power sector. *Joule 2*(12), 2498–2510.

Jiang, Y. and R. Sioshansi (2023). What duality theory tells us about giving market operators the authority to dispatch energy storage. *The Energy Journal 44*(3).

Joskow, P. (2008). Capacity payments in imperfect electricity markets: Need and design. *Utilities Policy 16*(3), 159–170.

Joskow, P. and J. Tirole (2007). Reliability and competitive electricity markets. *The RAND Journal of Economics 38*(1), 60–84.

Karaduman, Ö. (2021). Economics of grid-scale energy storage in wholesale electricity markets. Available at `https://ceepr.mit.edu/wp-content/uploads/2021/09/2021-005.pdf`.

Kazempour, J., P. Pinson, and B. Hobbs (2018). A stochastic market design with revenue adequacy and cost recovery by scenario: Benefits and costs. *IEEE Transactions on Power Systems 33*(4), 3531–3545.

Knueven, B., J. Ostrowski, A. Castillo, and J.-P. Watson (2021). A computationally efficient algorithm for computing convex hull prices. *Computers & Industrial Engineering*, 107806.

Korpås, M. and A. Botterud (2020, March). Optimality conditions and cost recovery in electricity markets with variable renewable energy and energy storage. Available at `http://ceepr.mit.edu/files/papers/2020-005.pdf`.

Krall, E., M. Higgins, and R. O'Neill (2012, July). RTO unit commitment test system. Technical report.

Lavin, L., N. Zheng, and J. Apt (2021). Market power challenges and solutions for electric power storage resources. Available at `https://www.cmu.edu/ceic/research-publications/ceic_21_02-esr-policy.pdf`.

Lawrence Berkeley National Laboratory (2023, April). Queued up: Characteristics of power plants seeking transmission interconnection. Available at `https://emp.lbl.gov/queues`.

Liberopoulos, G. and P. Andrianesis (2016). Critical review of pricing schemes in markets with non-convex costs. *Operations Research 64*(1), 17–31.

Mak, W.-K., D. P. Morton, and R. Wood (1999). Monte carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters 24*(1), 47–56.

Mays, J. (2021a). Missing incentives for flexibility in wholesale electricity markets. *Energy Policy 149*, 112010.

Mays, J. (2021b). Quasi-stochastic electricity markets. *INFORMS Journal on Optimization 3*(4), 350–372.

Mays, J., D. Morton, and R. O'Neill (2021). Investment effects of pricing schemes for non-convex markets. *European Journal of Operational Research 289*(2), 712–726.

Muñoz, F. D., S. Wogrin, S. S. Oren, and B. F. Hobbs (2018). Economic inefficiencies of cost-based electricity market designs. *The Energy Journal Volume 39*(3), 51–68.

O'Neill, R., P. Sotkiewicz, B. Hobbs, M. Rothkopf, and W. Stewart (2005). Efficient market-clearing prices in markets with nonconvexities. *European Journal of Operational Research 164*(1), 269–285.

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

Philpott, A., M. Ferris, and R. Wets (2016). Equilibrium, uncertainty and risk in hydro-thermal electricity systems. *Mathematical Programming 157*(2), 483–513.

Philpott, A., Z. Guan, J. Khazaei, and G. Zakeri (2010). Production inefficiency of electricity markets with hydro generation. *Utilities Policy 18*(4), 174–185. Designing Electricity Auctions.

Potomac Economics (2023a). 2022 state of the market report for the MISO electricity markets. Available at `https://www.potomaceconomics.com/wp-content/uploads/2023/06/2022-MISO-SOM_Report_Body-Final.pdf`.

Potomac Economics (2023b). IMM concerns with the AS methodology and recommended improvements. Available at `https://www.ercot.com/files/docs/2023/09/15/imm-as-methodology-for-wmwg-091523.pdf`.

Powell, W. B. (2022). *Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions*. John Wiley & Sons.

Pritchard, G., G. Zakeri, and A. Philpott (2010). A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations Research 58*(4-part-2), 1210–1219.

Scarf, H. E. (1990). Mathematical programming and economic theory. *Operations Research 38*(3), 377–385.

Schmalensee, R. (2022). Competitive energy storage and the duck curve. *The Energy Journal 43*(2).

Singhal, N. G. and E. G. Ela (2020). Pricing impacts of state of charge management options for electric storage resources. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pp. 1–6.

Steeger, G., L. A. Barroso, and S. Rebennack (2014). Optimal bidding strategies for hydro-electric producers: A literature survey. *IEEE Transactions on Power Systems 29*(4), 1758–1766.

Sun, Y., B. Frew, T. Levin, R. Hytowitz, J. Kwon, A. Mills, Q. Xu, M. Heidarifar, N. Singhal, P. de Mello, E. Ela, A. Botterud, Z. Zhou, B. Hobbs, and C. Crespo Montanes (2021, May). Research priorities and opportunities in United States competitive wholesale electricity markets. Available at `https://www.nrel.gov/docs/fy14osti/61765.pdf`.

Tesfatsion, L. (2023, May). Locational marginal pricing: When and why not? Available at `https://www2.econ.iastate.edu/tesfatsi/LMPWhenAndWhyNot.LTesfatsion.pdf`.

Wang, C., P. Luh, P. Gribik, T. Peng, and L. Zhang (2016). Commitment cost allocation of fast-start units for approximate extended locational marginal prices. *IEEE Transactions on Power Systems 31*(6), 4176–4184.

Wang, S., B. Sun, L. Ni, S. H. Low, and D. H. Tsang (2023). Financial incentives for energy resource investments in stochastic economic dispatch. *IEEE Transactions on Sustainable Energy*, 1–15.

Williams, J. H., R. A. Jones, B. Haley, G. Kwok, J. Hargreaves, J. Farbes, and M. S. Torn (2021). Carbon-neutral pathways for the united states. *AGU Advances 2*(1), e2020AV000284. e2020AV000284 2020AV000284.

Zakeri, G., G. Pritchard, M. Bjorndal, and E. Bjorndal (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.

Zavala, V., K. Kim, M. Anitescu, and J. Birge (2017). A stochastic electricity market clearing formulation with consistent pricing properties. *Operations Research 65*(3), 557–576.

Zhao, J., T. Zheng, and E. Litvinov (2020). A multi-period market design for markets with intertemporal constraints. *IEEE Transactions on Power Systems 35*(4), 3015–3025.

A Proof of Theorem 1

Proof of Theorem 1. The dual to (SOC) can be stated as follows:

$$(DSOC)$$

$$\min_{\lambda,\alpha,\beta,\gamma} \quad \sum_{n\in\mathcal{N}^{OP}}\sum_{l\in\mathcal{L}} \phi_n D_{ln}^{+}\alpha_{ln}^{D} \tag{15a}$$

$$\text{subject to} \quad -\sum_{n\in\mathcal{N}^{OP}} \phi_n A_{gn}\alpha_{gn}^{P+} \geq -C_g^{INV} \quad \forall g\in\mathcal{G} \quad (y_g) \tag{15b}$$

$$-\sum_{n\in\mathcal{N}^{OP}} \phi_n(\alpha_{bn}^{K-}+\alpha_{bn}^{K+}+E_b\gamma_{bn}) \geq -C_b^{INV} \quad \forall b\in\mathcal{B} \quad (y_b) \tag{15c}$$

$$\phi_n(\lambda_n+\alpha_{ln}^{D}) \geq \phi_n V_l^L \quad \forall l\in\mathcal{L}, n\in\mathcal{N}^{OP} \quad (d_{ln}) \tag{15d}$$

$$\phi_n(-\lambda_n+\alpha_{gn}^{P+}) \geq -\phi_n C_g^{OP} \quad \forall g\in\mathcal{G}, n\in\mathcal{N}^{OP} \quad (p_{gn}) \tag{15e}$$

$$\phi_n(\lambda_n+\alpha_{bn}^{K-}-\beta_{bn}) \geq 0 \quad \forall b\in\mathcal{B}, n\in\mathcal{N}^{OP} \quad (k_{bn}^{c}) \tag{15f}$$

$$\phi_n(-\lambda_n+\alpha_{bn}^{K+}+\beta_{bn}) \geq 0 \quad \forall b\in\mathcal{B}, n\in\mathcal{N}^{OP} \quad (k_{bn}^{d}) \tag{15g}$$

$$\phi_n(\beta_{bn}-\sum_{m\in n+}\beta_{bm}+\gamma_{bn}) \geq 0 \quad \forall b\in\mathcal{B}, n\in\mathcal{N}^{OP} \quad (j_{bn}) \tag{15h}$$

$$\phi_0(-\sum_{m\in 0+}\beta_{bm}+\gamma_{b0}) \geq 0 \quad \forall b\in\mathcal{B} \quad (j_{b0}) \tag{15i}$$

$$\alpha_{ln}^{D} \geq 0 \quad \forall l\in\mathcal{L}, n\in\mathcal{N}^{OP} \tag{15j}$$

$$\alpha_{gn}^{P+} \geq 0 \quad \forall g\in\mathcal{G}, n\in\mathcal{N}^{OP} \tag{15k}$$

$$\alpha_{bn}^{K-}, \alpha_{bn}^{K+}, \gamma_{bn} \geq 0 \quad \forall b\in\mathcal{B}, n\in\mathcal{N}^{OP}. \tag{15l}$$

The dual to $(GEN)_g$ can be stated as follows:

$$(DGEN)_g \quad \min_{\alpha_g^{P+}} \quad \sum_{n\in\mathcal{N}^{OP}} \phi_n A_{gn} y_g \alpha_{gn}^{P+} \tag{16a}$$

$$\text{subject to} \quad \phi_n\alpha_{gn}^{P+} \geq \phi_n(\lambda_n - C_g^{OP}) \quad \forall n\in\mathcal{N}^{OP} \quad (p_{gn}) \tag{16b}$$

$$\alpha_{gn}^{P+} \geq 0 \quad \forall n\in\mathcal{N}^{OP}. \tag{16c}$$

The dual to $(STOR)_b$ can be stated as follows:

$$(DSTOR)_b \quad \min_{\alpha^{K-},\alpha^{K+},\beta,\gamma} \quad \sum_{n\in\mathcal{N}^{OP}} \phi_n y_b(\alpha_{bn}^{K-}+\alpha_{bn}^{K+}+E_b\gamma_{bn}) \tag{17a}$$

$$\text{subject to} \quad \phi_n(\alpha_{bn}^{K-}-\beta_{bn}) \geq -\phi_n\lambda_n \quad \forall n\in\mathcal{N}^{OP} \quad (k_{bn}^{c}) \tag{17b}$$

$$\phi_n(\alpha^{K+}_{bn} + \beta_{bn}) \geq \phi_n \lambda_n \qquad \forall n \in \mathcal{N}^{OP} \qquad (k^d_{bn}) \tag{17c}$$

$$\phi_n(\beta_{bn} - \sum_{m \in n+} \beta_{bm} + \gamma_{bn}) \geq 0 \qquad \forall n \in \mathcal{N}^{OP} \qquad (j_{bn}) \tag{17d}$$

$$\phi_0(- \sum_{m \in 0+} \beta_{bm} + \gamma_{b0}) \geq 0 \qquad (j_{b0}) \tag{17e}$$

$$\alpha^{K-}_{bn}, \alpha^{K+}_{bn}, \gamma_{bn} \geq 0 \qquad \forall n \in \mathcal{N}^{OP}. \tag{17f}$$

The dual to $(LOAD)_l$ can be stated as follows:

$$(DLOAD)_l \qquad \min_{\alpha^D} \qquad \sum_{n \in \mathcal{N}^{OP}} \phi_n D^+_{ln} \alpha^D_{ln} \tag{18a}$$

$$\text{subject to} \qquad \phi_n \alpha^D_{ln} \geq \phi_n (V^L_l - \lambda_n) \qquad \forall n \in \mathcal{N}^{OP} \qquad (d_{ln}) \tag{18b}$$

$$\alpha^D_{ln} \geq 0 \qquad \forall n \in \mathcal{N}^{OP}. \tag{18c}$$

Constraints included in the primal and dual problems of the market participants are a subset of those included in the social problem (SOC). Given optimal solutions $(y^*, p^*, k^{c*}, k^{d*}, j^*, d^*)$ to (SOC) and $(\lambda^*, \alpha^*, \beta^*, \gamma^*)$ to $(DSOC)$, the complementary slackness conditions hold and are a superset of the conditions for the market participant problems. Thus, the solution satisfies the first, profit-maximizing condition of Definition 1. Next, since $(y^*, p^*, k^{c*}, k^{d*}, j^*, d^*)$ is a feasible solution to (SOC), it also satisfies the second, market-clearing condition of Definition 1.

For the third, zero-profit condition, from (SOC) and $(DSOC)$ we can form the complementary slackness conditions

$$0 \leq y_g \perp C^{INV}_g - \sum_{n \in \mathcal{N}^{OP}} \phi_n A_{gn} \alpha^{P+}_{gn} \geq 0 \quad \forall g \in \mathcal{G}$$

and

$$0 \leq y_b \perp C^{INV}_b - \sum_{n \in \mathcal{N}^{OP}} \phi_n (\alpha^{K-}_{bn} + \alpha^{K+}_{bn} + E_b \gamma_{bn}) \geq 0 \quad \forall b \in \mathcal{B}.$$

If $y_g = 0$ or $y_b = 0$, then the relevant condition for generator g or storage resource b in Definition 1 is automatically satisfied. If $y_g > 0$, then the right-hand condition is satisfied with equality. Multiplying yields the expression $C^{INV}_g y_g - \sum_{n \in \mathcal{N}^{OP}} \phi_n A_{gn} y_g \alpha^{P+}_{gn} = 0$. By strong duality, $\sum_{n \in \mathcal{N}^{OP}} \phi_n A_{gn} y_g \alpha^{P+}_{gn} = \sum_{n \in \mathcal{N}^{OP}} \phi_n p_{gn} (\lambda_n - C^{OP}_g)$. Similarly, if $y_b > 0$, then multiplying the right-hand condition by y_b yields the expression $C^{INV}_b y_b - \sum_{n \in \mathcal{N}^{OP}} \phi_n y_b (\alpha^{K-}_{bn} + \alpha^{K+}_{bn} + E_b \gamma_{bn}) = 0$. By strong duality, $\sum_{n \in \mathcal{N}^{OP}} \phi_n y_b (\alpha^{K-}_{bn} + \alpha^{K+}_{bn} + E_b \gamma_{bn}) = \sum_{n \in \mathcal{N}^{OP}} \phi_n \lambda_n (k^d - k^c)$. Thus, substituting in both cases reconstructs the third condition in Definition 1 and establishes zero expected profit in equilibrium. Thus, the prices $(\lambda^*_n)_{n \in \mathcal{N}^{OP}}$, installed generation capacities $(y^*_g)_{g \in \mathcal{G}}$, installed storage capacities $(y^*_b)_{b \in \mathcal{B}}$ generator schedules $(p^*_{gn})_{g \in \mathcal{G}, n \in \mathcal{N}^{OP}}$, storage schedules $(k_{bn}^{c*}, k_{bn}^{d*}, j_{bn}^{*})_{b\in\mathcal{B},n\in\mathcal{N}^{OP}}$, and load consumption quantities $(d_{ln}^{*})_{l\in\mathcal{L},n\in\mathcal{N}^{OP}}$ constitute a long-run competitive equilibrium. □

B Alternate reserve formulations

Continuing the discussion in Section 4.4, a second way to think about reserves is to increase the capacity available in the future based on uncertainty in the interval itself (rather than in subsequent intervals). Since the supply of energy and reserves is linked through the constraint in Eq. (6g), doing so is conceptually similar to the net load biasing described in Section 4.3. Using the notation from above, we can define the future reserve tuning $(DLAC - FRT - \theta)$ policy with the substitutions

$$R_{tt'}^{+} = \max\{0, \sum_{g\in\mathcal{G}^V}(\overline{P}_{gtt'}^{+} - \hat{P}_{gtt',\theta}^{+}) + \sum_{l\in\mathcal{L}}(\hat{D}_{ltt',1-\theta}^{+} - \overline{D}_{ltt'})\} \quad \forall t' \in \mathcal{T}_t^F.$$

Whereas the $(DLAC - RT - \theta)$ policy reflected uncertainty in the forecast for the subsequent interval, this policy focuses on the uncertainty in forecasts for interval itself. To achieve the optimal dispatch, this definition must be accompanied by a constraint limiting the quantity of capacity reserves able to be supplied by resources. In Example 2, the optimal dispatch in period 2 can be achieved by setting $R_{23}^{+} = 9.798$ and adding the constraints $\tilde{r}_{gas,2,3} \leq p_{gas,2} + UR_{gas}^{+}$ and $\tilde{r}_{b2,3} \leq j_{b2}$. Then, even though the deterministic model will expect to be able to serve all load in period 3 with wind at its expected availability of $\tilde{P}_{wind,2,3}^{+} = 10$, it will dispatch gas in period 2 in order to provide sufficient reserves to satisfy the constraint.

A third way to define reserves is not through the parameters $R_{tt'}^{+}$ and V^R, but instead by defining a residual value for energy in storage. We can formulate the residual value policy as follows:

$$\begin{aligned}
&(DLAC - RV)\\
&X_t^{DLAC-RV}(S_t|H,F) \in\\
&\arg\max_{\tilde{u},\tilde{v},\tilde{q},\tilde{p},\tilde{r},\tilde{k}^c,\tilde{k}^d,\tilde{j},\tilde{d},\tilde{z}} \quad \sum_{l\in\mathcal{L}}\sum_{t'\in\mathcal{T}_t^{H,F}} V_l^L \tilde{d}_{ltt'} + \sum_{t'\in\mathcal{T}_t^{H,F}} V^R \tilde{z}_{tt'} + \sum_{b\in\mathcal{B}} V_{t,t+F}^J \tilde{j}_{bt,t+F}\\
&\qquad - \sum_{g\in\mathcal{G}}\sum_{t'\in\mathcal{T}_t^{H,F}} C_g^{OP}\tilde{p}_{gtt'} - \sum_{g\in\mathcal{G}}\sum_{t'\in\mathcal{T}_t^{H,F}} \left(C_g^{NL}\tilde{u}_{gtt'} + C_g^{SU}\tilde{v}_{gtt'}\right) \qquad (19)\\
&\text{subject to} \qquad (6b)\text{–}(6z),
\end{aligned}$$

with the new parameter $V_{t,t+F}^J$ employed to estimate the value of stored energy at the end of the lookahead horizon. As before, we could substitute this value with a piecewise linear function without changing the analysis but increasing the degrees of freedom for operators. In a single-period model, the demand curve for stored energy could be defined in a manner similar to the offer values in model $(DLAC-OC)$. If the value of this function crossed below $V^J_{2,2} = \$100$/MWh at the quantity $j_{b2,2} = 0.89$, the model would then switch from discharging storage to using gas at this point. Similarly, in a two-period model with $\tilde{P}^+_{wind,2,3} = 10$, setting a residual value $V^J_{2,3} = \$100$/MWh at the quantity $j_{b2,3} = 0.89$ (even though the residual value is by definition \$0 in period 3) would induce an optimal dispatch in period 2. In either case, the model could also produce the optimal price λ^*_2. As such, modeling energy reserves in this way does not necessarily require introduction of a reserve product to produce appropriate incentives. However, achieving this price would be contingent on decisions made by system operators or market participants in establishing the value curve.