Electricity Dispatch and Pricing using Agent Decision Rules

Andy Philpott[1] Michael C. Ferris[2] Jacob Mays[3]

SIAM Optimization Meeting, Edinburgh

June 3, 2026

[1]University of Auckland
[2]University of Wisconsin-Madison
[3]Cornell University

Outline

- Setting: renewables, storage, and uncertainty in wholesale markets
- Current practice: deterministic lookahead and its drawbacks
- Proposal: Agent Decision Rules (ADRs)
- Theory: ADR partial equilibrium and existence
- Example: 24-period dispatch with one battery + one generator
- Extensions and conclusions

Our setting

- Renewable energy (wind and solar) growing in scale.
- Grid-connected storage increasing rapidly.
- Stochastic multiperiod dispatch and pricing being proposed:
  - Pricing rules for minimizing uplift payments.
  - Pricing the option value of storage.
  - Consistency of prices across multiperiod solutions.
  - Consensus on system operator's scenarios?
- **Our proposal:** return to single-period dispatch but use decision rules defined by a dynamic programming policy.

uplift payments for operational costs that are not recovered through the standard market-clearing energy price, such as start-up, no-load, and minimum-generation costs.

Why this matters: the duck curve deepens

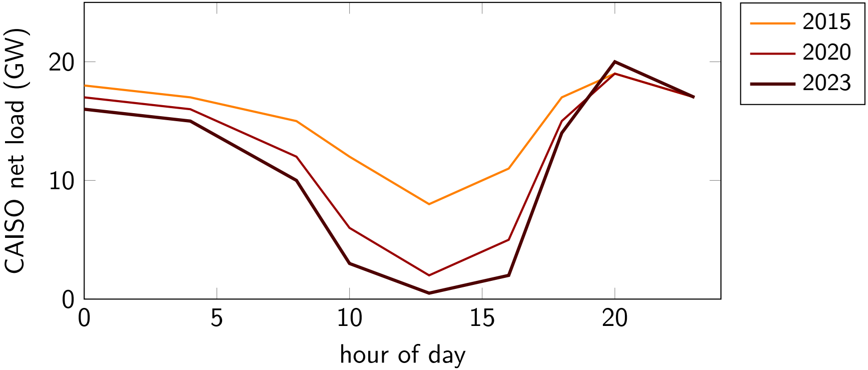

Lowest net-load day each spring (March–May), 2015–2023. Source: CAISO/EIA.

Storage is being deployed at scale

California (CAISO)

- First state to deploy 10 GW of battery storage (April 2024).
- Operating capacity 16.9 GW of battery storage (July 2025).
- Battery revenues large enough that uplift payments reached 7% of revenues in 2023.

New Zealand

- Rotohiko BESS (35 MW, Dec 2023).
- Ruakākā BESS (100 MW, May 2025).
- Contact Glenbrook–Ohurua (100 MW, Mar 2026).

New participation models for storage are being actively debated.

Three issues with deterministic lookahead

Market operators run rolling-horizon deterministic models, with bids/offers as inputs.

1. Inefficiency within the lookahead horizon (uncertainty is ignored).
2. Myopia: a finite lookahead can fail to prepare the system for what happens beyond it.
3. Forecast leakage: the operator's demand forecast affects prices, with unclear consequences for short-run efficiency and long-run investment.

Goal of this work:
Shift auctions from model predictive control toward dynamic programming.
Replace price–quantity offers with Agent Decision Rules that encode each participant's view of the future.

Self dispatch versus central dispatch

Self dispatch

- Battery forecasts prices and self-optimizes its storage revenue.
- System operator forecasts exogenous battery operation as part of net demand.

Central dispatch

- Battery provides a supply/demand curve.
- System operator co-optimizes single-period dispatch and endogenous battery operation.

But a static curve cannot express that *future* value of energy in storage depends on *current* state of charge.

Battery offer curve depends on state of charge

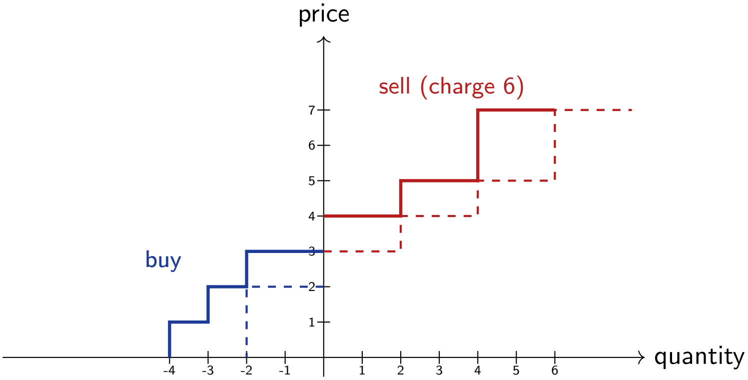

Solid: state of charge $= 6$. Dashed: state of charge $= 8$. Capacity: $= 10$

We need a way to communicate this state dependence to the operator.

Multiperiod economic dispatch (SOP)

A generic convex multiperiod social-optimization problem:

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

- $x(n)$: decision (generation, charge, discharge, storage, flows).
- $z(n)$: lost load, penalized at value of lost load L.
- $\mathcal{X}_i(x_i(n-))$: ramping, capacity, storage dynamics.
- Demand d^n is random; in current practice the operator forecasts.

Stage problem and the missing piece

The single-stage problem with an expected future cost function:

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

- This is the basis of a dynamic-programming algorithm.
- Who supplies $C^n(x)$? Today: the system operator (via forecasts/scenarios).
- **Our proposal:** let the agents supply it.

Agent Decision Rules (ADRs)

Each agent i in scenario node n supplies the operator with:

- immediate cost c_i^n, feasible set $\mathcal{X}_i(\bar{x}_i)$,
- an ADR $R_i^n(\cdot)$ that encodes their expected future surplus.

The operator solves the *single-period* problem

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

to yield a dispatch $x(n)$ and prices $\pi(n)$. Demand pays $\pi(n)^\top d^n$; each agent is paid $\pi(n)^\top A_{.i} x_i(n)$.

No operator forecast required — and no scenario tree imposed on participants.

How an agent constructs an ADR

Agent i uses its own probability $\mathbb{P}_i$ to recurse a future-value function:

$$V_i^n(\bar{x}_i) = \max_{x_i \in \mathcal{X}_i(\bar{x}_i)} \left\{ \pi_i(n)^\top A_{\cdot i} x_i - c_i^n(x_i) + \mathbb{E}_{\mathbb{P}_i(m|n)}[V_i^m(x_i)] \right\}.$$

The ADR for agent i at node n is

$$R_i^n(\cdot) \; = \; \mathbb{E}_{\mathbb{P}_i(m|n)}[V_i^m(\cdot)]\,.$$

- Concave V, so DP$(n, \cdot)$ remains a convex single-period problem.
- Naturally depends on observable state: time of day, state of charge, previous dispatch, previous price . . .
- Can be supplied *long-lived* (start of day) or *real time* (each interval).

ADR partial equilibrium

Definition. Given a scenario tree $\mathcal{T}$ with demand $\{d^n\}$, an ADR partial equilibrium are prices $\{\pi^*(n)\}$, actions $\{x^*(n)\}$, ADRs $\{R_i^n(\cdot)\}$ s.t.

1. for each n, $x^*(n)$ solves $\mathrm{DP}(n, x^*(n-))$ with shadow prices $\pi^*(n)$;
2. for each n, $x_i^*(n)$ solves $\max_{x_i \in \mathcal{X}_i(x^*(n-))} \{\pi^*(n)^\top A_{\cdot i} x_i - c_i^n(x_i) + R_i^n(x_i)\}$;
3. $R_i^{n-}(x_i^*(n-)) = \mathbb{E}_{\mathbb{P}_i(n|n-)}\left[\pi^*(n)^\top A_{\cdot i} x_i^*(n) - c_i^n(x_i^*(n)) + R_i^n(x_i^*(n))\right]$.

A form of rational-expectations equilibrium: the stochastic process of prices each agent *assumes* matches what *clears* in the market.

A two-period illustration

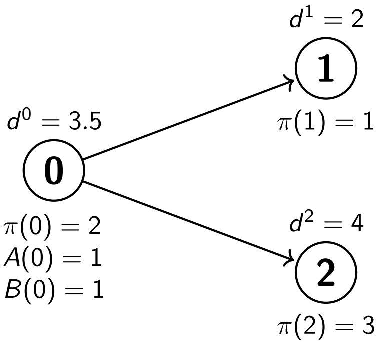

- Three nodes
- Three agents: (A,B,C)
- batteries A, B (cap 1 MWh, $\eta = 0.8$)
- thermal C with marginal cost

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

- A believes $\mathbb{P}_A = (0.6, 0.4)$; B believes $\mathbb{P}_B = (0.4, 0.6)$.

Equilibrium: A discharges 1 at $t = 0$; B discharges 1 at $t = 1$.

ADRs at $t = 0$: $R_A^0(y_A) = 1.8\, y_A (= 0.6 * 1 * y_A + 0.4 * 3 * y_A)$, $R_B^0(y_B) = 2.2\, y_B$, $R_C^n(y_C) = 0$.

Agents need not share beliefs to form an ADR partial equilibrium.

When does an ADR equilibrium exist?

Theorem (no-arbitrage). If a complete-market ADR partial equilibrium exists (Arrow–Debreu securities indexed on each successor node), then $\mathbb{P}_i$ is the same for every agent $i \in \mathcal{I}$.

Theorem (sufficient conditions). If all agents share the system operator's $\mathbb{P}$, and the social Bellman function is gradient separable at $x^*(n-)$, then there exists π^* such that (x^*, π^*, S_i^n) is an ADR partial equilibrium, where

$$S_i^n(x_i) = -\mathbb{E}_{\mathbb{P}(m|n)}\left[C_{i,\hat{x}_{-i}}^m(x_i)\right]$$

is the agent's slice of the social future-cost function.

Common beliefs $\Rightarrow$ ADRs can deliver the social optimum.

Illustrative example: one battery, one generator

24-period dispatch problem EP : $\min \sum_{t=1}^{24}(c^t(q(t)) + Lz(t))$ with

$$\mathcal{Q}(\bar{q}) = \{q \mid 0 \leq q \leq q^{\max},\ q - \bar{q} \leq \rho\},$$
$$\mathcal{Y}(\bar{y}) = \{(y, u, v) \mid 0 \leq y \leq E,\ 0 \leq u \leq r,\ 0 \leq v \leq s,\ y = \bar{y} - u + \eta v\}.$$

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

Marginal generator cost has 10 increasing tranches ($10 to $200 per MWh).

The peak in the early evening forces ramping and discharge.

Demand profile and deterministic solution

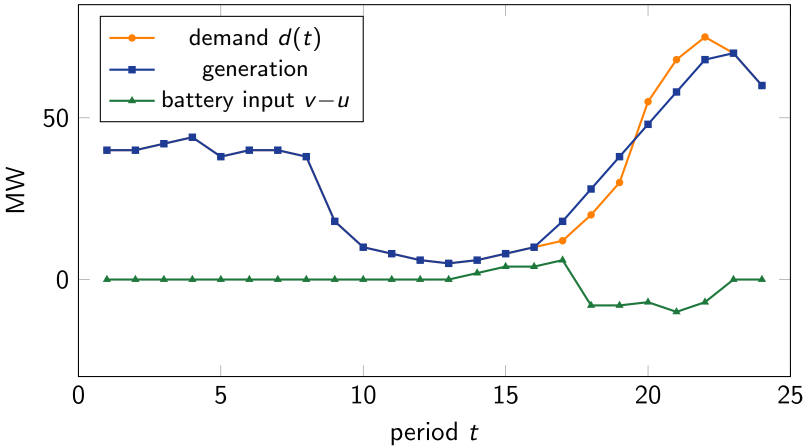

Optimal cost (perfect foresight): $48,470.

Adding uncertainty: stochastic dynamic programming

Add stagewise-independent noise to demand, $\varepsilon \in \{-4, -2, 0, 2, 4\}$, equiprobable.

- Solve full SDP to obtain expected future cost $C^t(q, y)$.
- SDP optimal expected cost: $52,377 (social optimum).

Four candidate policies for the operator:

	Dispatch input	
Cost function estimate	Separable ADRs	Social cost-to-go
Deterministic	$55,430	$54,255
Stochastic	$52,688	$52,406

Stochastic ADRs $\approx$ social optimum ($52,688 vs $52,406).

Why ADRs work here

- Each agent computes its own future-value function from a perfect-foresight solution (or SDP).
- Battery: $S_b^t(y) = -C^t(\hat{q}(t), y)$.
- Generator: $S_g^t(q) = -C^t(q, \hat{y}(t))$.
- Operator solves a single-period convex problem each interval.

Higher-variance experiment (noise from $\{-4, -2, 0, 8, 16\}$):

- Optimal expected cost rises to \$57,438.
- Deterministic dispatch costs jump to \$60,984; stochastic ADRs deliver \$59,409.
- Gap to social optimum widens with variance, but *ranking is preserved*.

What goes back to the system operator

- Generator agents: marginal costs and ABFs $G_a^t(x_a)$.
- Battery agent a: convex ABF $B_a^t(y_a)$.
- Operator solves single-stage ADR(t):

$$\begin{aligned} \min\ & \textstyle\sum_{a\in\mathcal{G}} c_a x_a(t) + Lz(t) + \sum_{a\in\mathcal{G}} G_a^t(x_a) + \sum_{a\in\mathcal{B}} B_a^t(y_a) \\ \text{s.t. } & \textstyle\sum_{a\in\mathcal{G}} x_a + \sum_{a\in\mathcal{B}}(u_a - v_a) + z \geq d^t, \quad [\pi(t)] \\ & x_a \in \mathcal{X}_a(x_a(t-1)),\ (y_a, u_a, v_a) \in \mathcal{Y}_a(y_a(t-1)),\ z \in [0, d^t]. \end{aligned}$$

Generator i paid $\pi(t)x_i(t)$; battery j paid $\pi(t)(u_j(t) - v_j(t))$.

Convex, single-period, deterministic. Prices give budget balance and revenue adequacy.

Extensions

ADRs are not just for batteries:

- Classical supply functions are a special case (limited state dependence).
- Demand response via state-dependent demand-side bids.
- Pump-storage hydro — same form as batteries, longer time scale.
- Flexible demand (industrial loads, data centers) — shift consumption to off-peak.
- Hydroelectric generators — expected marginal water value plays the role of B^t.
- Reserve and frequency regulation — co-optimized within the single-period model.

Open: unit commitment with binaries; extensions of SDDP (e.g. MIDAS) for approximate future-cost functions.

Contributions

1. A new form of offer (an ADR) that encapsulates a participant's view of future market conditions.
2. Computation of an optimal dispatch through a sequence of single-period problems, with prices that approximate competitive equilibrium (exact under shared beliefs).
3. A clean separation of responsibility: the operator delivers reliable supply; participants benefit financially from their foresight.

Conclusions

- In a perfectly competitive, convex, complete market where all agents share the same probability distribution, ADRs implement the welfare-maximizing dispatch.
- Agents with different beliefs "put their money where their mouth is" — and equilibrium prices reveal the implied common measure.
- The operator no longer pushes its own forecast through prices; aggregation of participants' views does that work.
- Changes to single-period dispatch are modest; the offer format is richer.