ERCOT Residential Battery Dispatch: A 2022 to 2025 Backtest

Deep Hayer
Energy | ERCOT | Optimization

May 2026

Abstract. We backtest a Powerwall-class residential battery (13.5 kWh, 5 kW, 90% round trip efficiency) on four years of ERCOT day-ahead settlement point prices across five load zones. A perfect-foresight linear program earns \$474 to \$583 per battery per year on average across zones, peaking at \$971 in Austin during the 2023 heat dome. A naive sort-based heuristic earns 93% to 97% as much. The implication is uncomfortable but important: on the day-ahead market, where tomorrow's prices are known when you dispatch, optimization has very little value over a simple rule. The interesting forecasting and optimization problem lives in the 5-minute real-time market and in ancillary services. We close with what this means for distributed battery operators like Base Power.

ERCOT and the Residential Battery Thesis

ERCOT is the most actively traded electricity market in the United States and operates with one of the highest renewable penetrations in the world. The combination of high wind and solar shares, no formal capacity market, and recurring scarcity events (Winter Storm Uri 2021, Heat Dome 2023, multiple near-misses since) generates the kind of price volatility that, in theory, makes distributed batteries economically interesting.

Base Power is the most visible Texas-based player on the residential side. The pitch: install batteries in customer homes, aggregate them into a virtual power plant, monetize across energy arbitrage, ancillary service awards, and capacity availability payments. This paper isolates the simplest of those revenue streams, energy arbitrage on the day-ahead market, and asks what a single Powerwall earns if you dispatch it optimally and how much the optimal beats a homeowner who follows a simple rule.

Data and Methodology

Day-ahead market settlement point prices for 2022 through 2025 are pulled from ERCOT MIS via the open-source `gridstatus` library. We keep five load zones (LZ_AEN, LZ_NORTH, LZ_SOUTH, LZ_HOUSTON, LZ_WEST) and four hubs. Prices are hourly, in dollars per MWh, in US Central time.

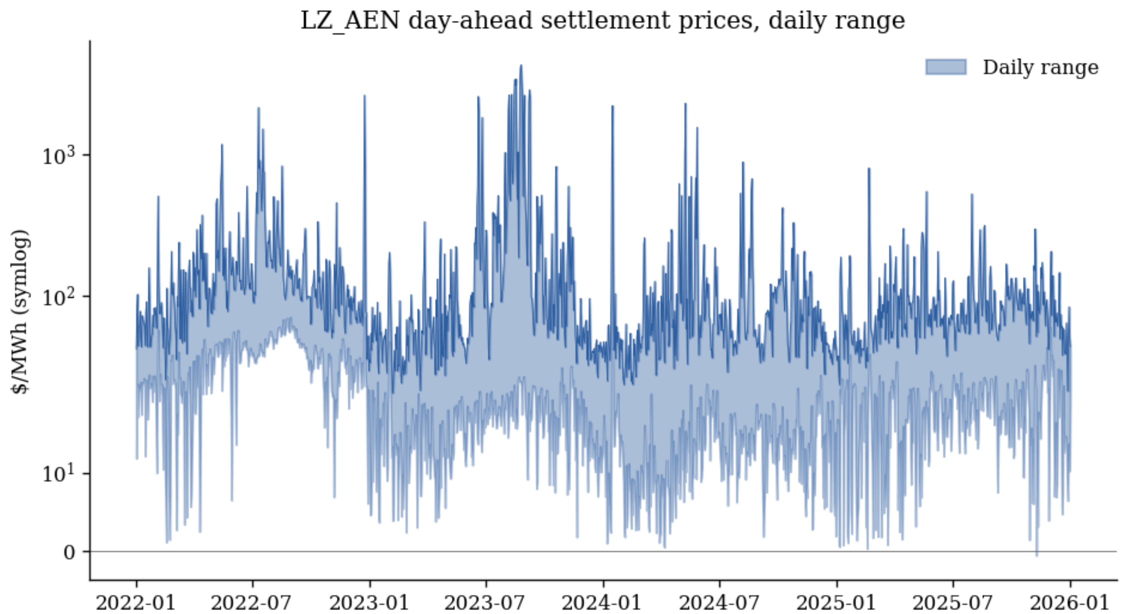

Figure 1: LZ_AEN daily price range, 2022 to 2025. Scarcity events visible as upward spikes; oversupply events as negative-price floors.

The battery is a Powerwall-class unit: 13.5 kWh usable capacity, 5 kW peak power, 95% charge and discharge efficiency for a round-trip efficiency of about 90%. At 5 kW power, a full charge from empty takes roughly three hours, which sets the natural number of daily cycles.

Two strategies

The first strategy is perfect-foresight linear programming. For each day, solve:

$$\max \sum_t p_t(d_t - c_t) \text{ subject to}$$

$$s_{t+1} = s_t + \eta_c c_t - \frac{d_t}{\eta_d},$$

$$0 \le c_t, d_t \le P, 0 \le s_t \le C, s_0 = s_T = 0.$$

where p_t is the hourly price, c_t and d_t are charge and discharge flows in kWh per hour, s_t is state of charge, η_c and η_d are efficiencies, P is max power, C is capacity. Solved with `scipy.optimize.linprog` (HiGHS backend).

The second strategy is a naive heuristic. Given a day's known prices, charge during the three cheapest hours at full power and discharge during the three most expensive hours at the power level that returns exactly the round-trip-efficiency-adjusted energy. No optimization, no forecasting beyond sorting tomorrow's published prices.

The gap between the two is the "value of forecasting" on this market.

Results

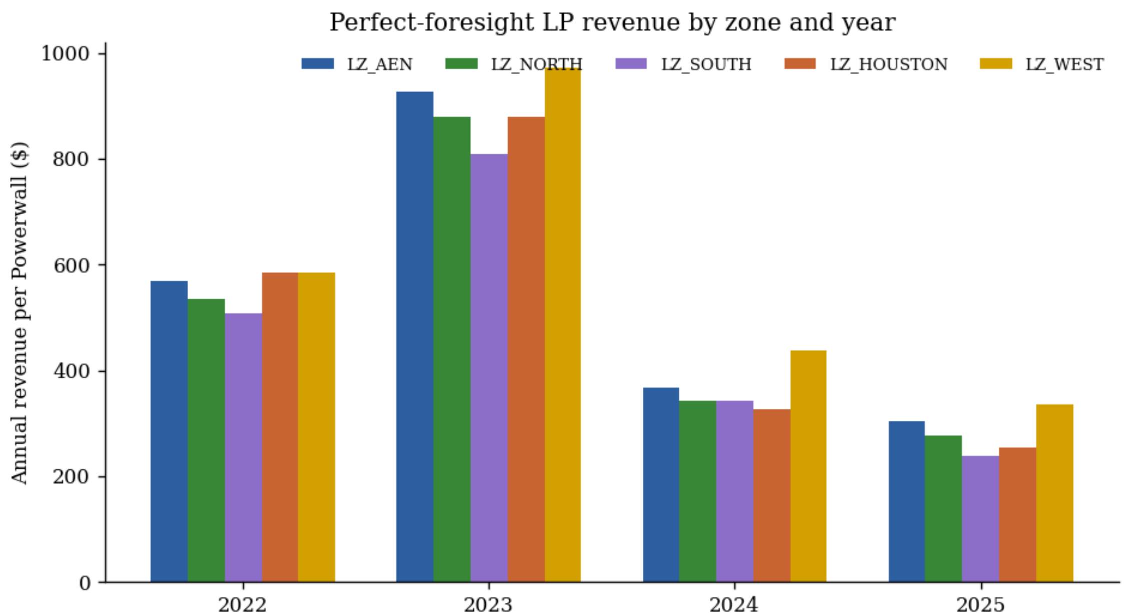

Figure 2: Annual revenue per Powerwall by load zone, perfect-foresight LP. 2023 was the standout year across every zone, driven by the summer heat dome.

Across four years and five zones the LP earns $474 to $583 per battery per year on average. Annual variation is large: 2023 dominates every other year by a factor of two to three, driven by recurring $1,000+ per MWh prints during the summer heat dome. 2024 and 2025 are about a third of 2023 in revenue terms, reflecting more normal weather and stronger reserve margins.

Across zones, LZ_WEST tops the list slightly. West Texas concentrates wind generation; high wind output combined with congestion to load centers produces the highest hour-to-hour volatility in the system, which is exactly what an arbitrage battery wants. LZ_AEN sits a close second.

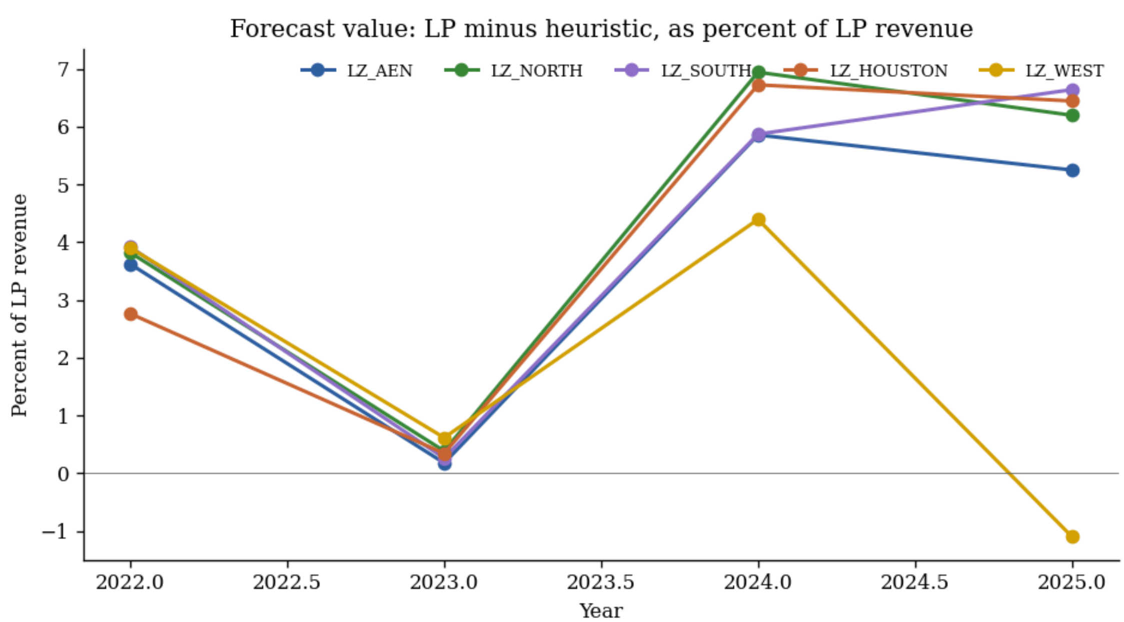

Figure 3: LP minus heuristic, as percent of LP revenue, by zone and year. The "forecast value" on the day-ahead market is consistently 0 to 7 percent.

The interesting result is in figure 3. The naive heuristic captures 93 to 97 percent of the LP's revenue in every zone, every year. In 2023, when prices were highest, the gap narrowed to under 1 percent because the extreme price spikes made the right hours obvious. In 2024 and 2025 the gap widened slightly, but never above 7 percent.

On the day-ahead market, where dispatch is decided after tomorrow's prices are already known and published, optimization buys you very little. The problem of "which hours should I charge and discharge tomorrow" is nearly-solved by sorting tomorrow's price vector.

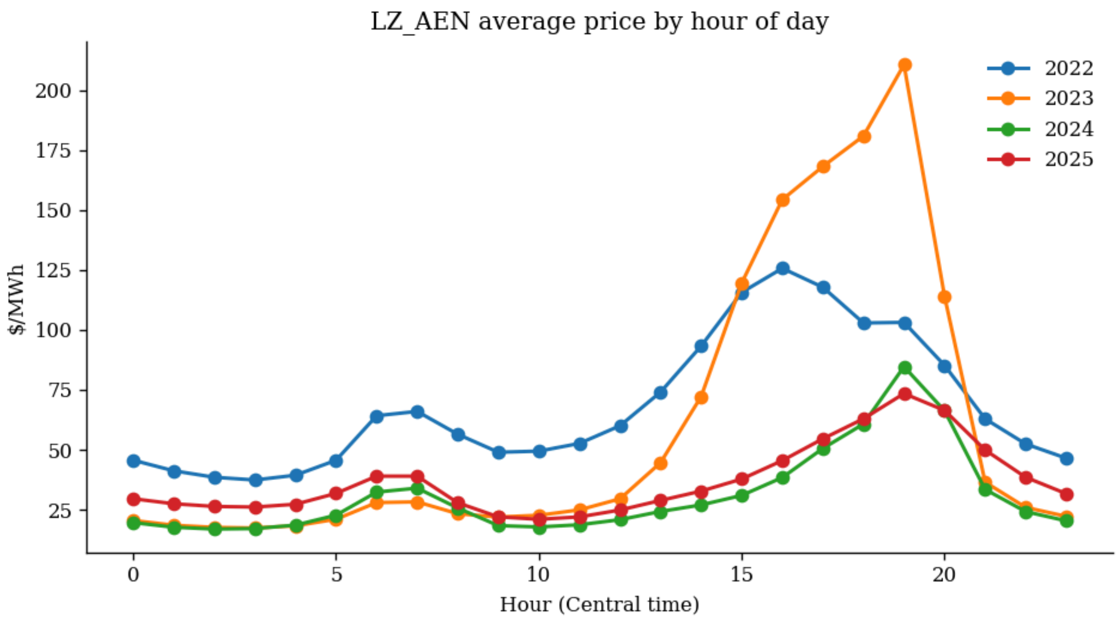

Figure 4: Average LZ_AEN price by hour of day, by year. The morning ramp and evening peak structure is consistent, but the magnitude of the daily spread varies dramatically year to year.

The average intraday profile is stable in shape: morning ramp into the business day, dip mid-day as solar floods the system, evening peak as solar falls off and load stays high, overnight trough. The magnitude of the daily spread varies enormously, and the daily spread is what generates arbitrage revenue.

What This Means

The honest conclusion is that day-ahead energy arbitrage on a single residential battery is a marginal business. At $500 per year per Powerwall and an installed system cost of roughly $10,000 to $11,500, simple payback from arbitrage alone is 20-plus years. That is below the battery's warrantied life. Energy arbitrage is the floor of the residential battery thesis, not the ceiling.

The ceiling lives in three places:

- **Real-time market dispatch.** The 5-minute real-time settlement prices are much more volatile than day-ahead. Forecasting matters there because the price is not known when you dispatch. A capable price model and a closed-loop controller capture meaningful spread that this paper does not attempt to estimate.
- **Ancillary services.** ERCOT pays separately for ECRS, RRS, and Non-Spin reserves. Awarded batteries earn capacity payments whether or not they deploy. As of 2025, AS revenue is the dominant share of large BESS revenue in ERCOT.

- **Aggregation.** A single Powerwall is too small to clear most market thresholds. The interesting unit is a fleet of thousands of homes treated as a single dispatchable resource. That is exactly the structure Base Power operates.

For a residential battery operator like Base Power, the implication is that the technical edge is not "I run a slightly better LP than the homeowner." The technical edge is in aggregation, real-time dispatch under price uncertainty, and AS market participation. This paper measures the floor; the actual business sits well above it.

Limitations and Next Steps

Three honest caveats:

- Day-ahead only. The richer arbitrage value lives in real-time. The paper understates revenue.
- No degradation modeling. We assume zero capacity fade across four years. Real Powerwall capacity drops 1 to 2 percent per year under typical cycling.
- No ancillary services. The 2024-2025 ERCOT data shows AS as a larger revenue line than arbitrage for utility-scale BESS. A residential aggregator can capture a share of that. Not modeled here.

A v2 should add (1) real-time 5-minute price backtest with a rolling-forecast controller, (2) ECRS and RRS award modeling, (3) capacity fade and round-trip efficiency drift.

Appendix: Reproducibility

Companion essay at deephayer.com. Run `python -m data.ercot_pull && python -m src.backtest && python -m src.plots && cd paper && typst compile ercoiq.typ`. The interactive Streamlit dashboard is `streamlit run app.py`.